LiDAS: Lighting-driven Dynamic Active Sensing for Nighttime Perception

Supplementary Material

A. Downstream Tasks Studies

Combining detection and semantic segmentation. As shown in Sec. 5.4, each downstream task induces a distinct illumination policy. We report single-task results for detection and semantic segmentation in Tab. S2. When combined in LiDAS, the two guidance signals regularize the policy and yield further gains for both tasks, highlighting their complementarity. This joint training also produces a more general illumination pattern that remains informative for multiple downstream models.

Downstream model generalization. Beyond the models in Sec. 5.4, we evaluate LiDAS on additional downstream models unseen during training. In particular, Tab. S1 reports results for YOLO12L [21]. Consistent with the other models, LiDAS provides informative illumination for YOLO12L and yields consistent performance gains, further demonstrating the generality of the learned policy across multiple downstream models.

Depth estimation. We evaluate our method's ability to enhance depth estimation task in Sec. 5.4. We compare against LED [9], another HD-lighting-based approach specifically designed to improve depth estimation perception at night. To this end, we reproduce the LED pattern, a checkerboard with 0.25° cells, and apply it using our relighting operator (see Sec. 3.2) on our synthetic dataset. Table S3 shows that LED reduces RMSE by 1.61 m and improves the other metrics compared with low-beam illumination. However, LED requires fine-tuning the perception models to accommodate its structured illumination pattern. In contrast, LiDAS improves frozen downstream models without any retraining. It learns a more informative illumination pattern for frozen depth model, achieving a comparable 1.38 m RMSE reduction while using 40% less power. This demonstrates that our policy can also benefit other tasks when they are included during training. While depth and semantic tasks favor different light structures, LiDAS does not degrade depth performance when optimized for detection/segmentation: baseline accuracy is maintained (RMSE: $LiDAS^{[0.6]}$:8.35m, $LB^{[1]}$:8.79m, $LiDAS^{[1]}$:8.13m, $HB^{[1.8]}$:8.36m). Depending on the application, we choose which task maintains performances and which is improved.

Table S1. **Downstream model generalization.** We report YOLO12L [21] results, which was not used during training. In this zero-shot setting, LiDAS improves perception performance, showing its generality. Training downstream models are YOLO11L, YOLOv8L, YOLOv8L-Worldv2, Mask2Former (see Sec. 3)

Method	Power	Detection P ↑	R ↑	mAP^{50} ↑	mAP$^{50-90}$ ↑
		Yolo12L			
LiDAS (Ours)	0.6	**62.9**	38.9	44.3	28.0
Low Beam	1	53.2	31.4	34.7	19.7
LiDAS (Ours)	1	60.5	<u>40.9</u>	**45.0**	<u>28.4</u>
High Beam	1.8	<u>62.1</u>	36.5	42.1	25.9
LiDAS (Ours)	1.8	58.0	**41.3**	<u>44.4</u>	**28.7**

Table S2. **Impact of type of supervision (multi-task, single-task) on performance.** LiDAS trains with both detection and semantic segmentation as downstream tasks, showing improvements for each compared to task-specific training. Power = 1.

Training Task: Detection	Semantic S.	Detection: P ↑	R ↑	mAP^{50} ↑	mAP$^{50-90}$ ↑	Semantic Segmentation: mIoU ↑	mAcc ↑
✓	✗	60.6	40.5	46.3	29.2	-	-
✗	✓	-	-	-	-	72.0	85.1
✓	✓	**66.9**	**40.6**	**47.3**	**30.0**	**72.8**	**87.0**

Table S3. **Depth estimation task.** We evaluate LiDAS on depth estimation using DepthAnythingV2 [48] as the downstream model. We compare against LED [9], which requires fine-tuning the perception model, whereas our results use the pretrained DepthAnythingV2 in a zero-shot setting.

Method	Power	RMSE ↓	Abs Rel ↓	SiLog ↓	δ_1 ↑
With fine-tuning					
Low Beam	1	5.14	0.117	0.157	0.903
LED [9]	1	**3.53**	**0.066**	**0.103**	**0.950**
High Beam	1.8	4.79	0.107	0.148	0.909
Zero-shot					
LiDAS (Ours)	0.6	<u>7.41</u>	<u>0.178</u>	<u>0.225</u>	<u>0.745</u>
Low Beam	1	8.79	0.305	0.309	0.343
LED [9]	1	11.3	0.250	0.381	0.586
LiDAS (Ours)	1	7.68	<u>0.178</u>	0.233	0.732
High Beam	1.8	8.36	0.297	0.294	0.358
LiDAS (Ours)	1.8	**7.23**	**0.165**	**0.216**	**0.777**

B. Performance Over Distance

Fig. S1 shows that $LiDAS^{[1]}$ leads in the 20-60 m band, a safety-critical region for early Autonomous Emergency Braking (AEB) triggers. Methods are statistically similar at 0-20 m, where all are sufficiently bright and the main concern is self-glare. Beyond 70 m, performance converges again due to limited energy on small, distant targets. Contrary to intuition, HB does not dominate at long range: its light pattern does not illuminate lateral areas and thus primarily benefits centered targets. In contrast, LiDAS widens illumination around the horizon, improving accuracy for all far objects.

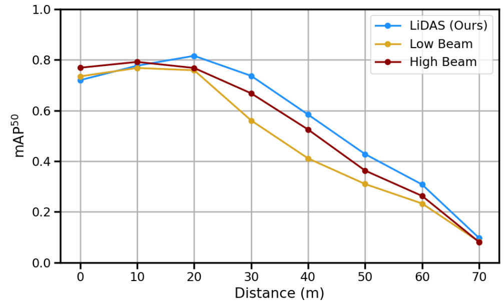

Figure S1. **Performance over distance.** LiDAS improves detection in the 20-60 m range, a safety-critical band for applications such as autonomous emergency braking.

C. Performance Over Time

We study the sequential refinement parameters (see Sec. 3.3) in Fig. S2. With K=1, the model sees only the initial random pattern M_0 during training and thus degrades quickly once exposed to its own illumination at test time. Increasing the unroll length N mitigates this drift: N=10 maintains performance longer, N=40 effectively stabilizes it, while N=100 brings no further benefit, yet substantially increases training time. Thus, making N=40 the best compute-accuracy trade-off. Varying K also matters as it sets the balance between the random start M_0 and LiDAS-generated light seen during training. We find K=5 best balances robustness to arbitrary initial light fields with stability under self-generated illumination. We note that LiDAS requires a brief warm-up to reach peak performance: 20-30 iterations suffice, corresponding to 1-2 s of initialization at vehicle startup, which is negligible in our active perception setting.

D. Evaluation on Public Datasets

Evaluating on existing datasets is challenging because we cannot physically alter the car's illumination as captured by the camera. We assume that night frames are LB-illuminated (M_{LB}), which is the expected setup in urban areas. Since captured images cannot gain photons, we only attenuate regions where our model requests less light than the LB. Accordingly, we apply our relighting operator to simulate new lighting under this darken-only constraint: Given an input I_{LB} and a desired illumination map $M \in [0, 1]$, we interpolate with a black image:

$$M' = 1 - \max(M_{\text{LB}} - M, 0), \quad \hat{I}' = I_{\text{LB}} \odot M'. \tag{6}$$

This protocol cannot reveal new details, it can only remove information. Table S5 shows that even under this unfavorable setting, *LiDAS* performances are comparable to the *Low beam* on the nuImages [4] dataset while reducing power by 30%. It demonstrates that it has learned to unlit only the uninformative regions of the scene as shown in Fig. S3. Following the same protocol, we report results on BDD100K [49] in Tab. S6 and NightCity [35] in Tab. S7. Results are consistent with performances observed on nuImages, reducing power consumption by 20% while maintaining performances. In an active-illumination system, the energy saved from dimmed areas would be reallocated to informative area of the scene, which could further improve perception.

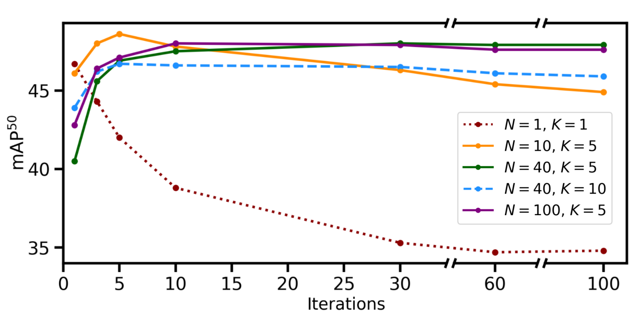

Figure S2. **Performance over time.** Increasing the step count N mitigates performance drift. Adjusting K balances exposure to the initial random pattern and the model's own light during training.

E. Public Baselines

As a bolt-on system, LiDAS is compatible with arbitrary perception architectures, we integrated the DTP [42] SOTA nighttime model as a downstream head on our closed-loop dateset. Results are reported in Tab. S4. LB and HB represent DTP's native performance under standard lighting. LiDAS consistently improves metrics for this nighttime-specialized model in a zero-shot setting, demonstrating it's compatibility with both DA/DG methods and models dedicated for nighttime perception.

F. Runtime Details

During training we used an RTX 4090 on which the inference takes 6.8 ms. For our real-world deployment, we used a laptop-GPU (RTX 2000 Ada) and achieved 32 FPS

Table S4. **Results using DTP in closed-loop.** LiDAS improves performances of the DTP nighttime-specialized model over the standard illumination.

Method	Power	Semantic Segmentation mIoU ↑	mAcc ↑
LiDAS (Ours)	0.6	19.0	19.3
Low Beam	1	18.3	18.5
LiDAS (Ours)	1	**23.2**	**23.9**
High Beam	1	18.0	18.3

(31.2 ms) at full and 97 FPS (10.3 ms) at half resolution. Our parallel pipeline, handling camera I/O, LiDAS inference, and headlight projection ($\leq$15 ms), maintains a total system latency of $\leq$46.2 ms, which has proven to be responsive enough. LiDAS runs in parallel with the perception stack, adding no latency to downstream models. While the size of the model was not our focus, standard compression techniques (distillation, pruning) can optimize it for ECUs. Successful real-world tests shows the system's feasibility.

G. Relighting Operator Limitations

The linear interpolation is a differentiable proxy designed for computational efficiency only during training. We acknowledge that it does not model complex physical effects (specular reflections, camera ISP non linearities...), still, its effectiveness is validated by LiDAS improvement on real-world perception in zero-shot transfer from synthetic data. It shows the operator provides a sufficient signal to learn a robust illumination policy.

H. Lighting Regulations

Our policy does not explicitly prevent glare toward other road users, so practical deployment will likely require integration with anti-glare systems [34]. LiDAS is compatible with these methods, which could mask sensitive regions in our output. We consider this a critical safety trade-off: with 77% of pedestrian crashes occurring at night[1], LiDAS prioritizes illumination for life-saving detection, in less-populated very low-light environments. Regulations remain a near-term barrier: many jurisdictions do not yet authorize fully dynamic HD headlight functions, and LiDAS may need adaptations (e.g., exclusion zones, intensity and update-rate limits) as rules evolve. Still, regulation is evolving toward dynamic HD lighting: active working groups are studying the use of light to assist ADAS, it will help define the balance between visibility and glare.

[1]NHTSA. Traffic Safety Facts: 2023 Data-Pedestrians. US Dept. of Transp. 2025. Report No. DOT HS 813 727.

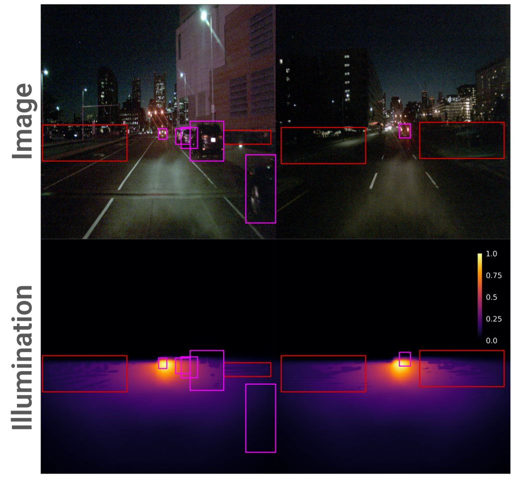

Figure S3. **Qualitative results on nuImages (night).** Red boxes indicate darkened areas. Purple boxes denote ground-truth objects. LiDAS selectively reduces illumination only in uninformative regions of the scene.

Table S5. **Results on nuImages (night).** LiDAS achieves performance comparable to Low Beam while using 30% less energy, indicating that it suppresses light only in uninformative regions.

Method	Power	Detection			
		P $\uparrow$	R $\uparrow$	mAP^{50} $\uparrow$	mAP$^{50-90}$ $\uparrow$
LiDAS (Ours)	0.7	61.1	23.7	41.0	23.8
Low Beam	1	62.1	24.8	43.0	26.3

Table S6. **Results on BDD100K (night).** LiDAS achieves performance comparable to Low Beam while using 20% less energy, indicating that it suppresses light only in uninformative regions.

Method	Power	Detection			
		P $\uparrow$	R $\uparrow$	mAP^{50} $\uparrow$	mAP$^{50-90}$ $\uparrow$
LiDAS (Ours)	0.8	51.1	23.2	36.1	24.0
Low Beam	1	50.0	24.2	36.1	24.6

Table S7. **Results on NightCity.** LiDAS achieves performance comparable to Low Beam while using 20% less energy, indicating that it suppresses light only in uninformative regions. We used SoMA as downstream model for this experiment, demonstrating its compatibility.

Method	Power	Semantic Segmentation	
		mIoU $\uparrow$	mAcc $\uparrow$
LiDAS (Ours)	0.8	51.4	74.5
Low Beam	1	51.5	74.7