JOURNAL OF GEOSPACE SCIENCE

Vol. 6 No. 2 (2026)

E-ISSN : 2413-6093
https://jogss.org

FogBEV: Degradation-Aware Bird's-Eye-View Fusion for Autonomous Driving in Dense Fog

Ravi Iyer, Megan Wang*

Department of Computer Science, Purdue University, West Lafayette, IN 47907, USA,

Email: meganwang2022@purdue.edu

ABSTRACT

Dense fog is a persistent failure mode for autonomous-driving perception because optical contrast is reduced by atmospheric scattering while range sensors may suffer from sparsification, backscatter and unstable returns. This paper presents FogBEV, a degradation-aware bird's-eye-view fusion framework that combines gated/NIR image evidence with LiDAR evidence in a shared BEV representation. Unlike fixed camera-LiDAR fusion, FogBEV estimates per-modality reliability from feature contrast, point-density and false-return cues, then uses these reliabilities to adjust fusion weights before BEV detection. The design is motivated by the strong gated vision and LiDAR integration framework, but moves the adaptation from image/point features into a BEV-centric formulation that is directly compatible with modern planning stacks. To avoid fabricated results, this draft reports only a lightweight synthetic BEV fog benchmark generated by the accompanying code. The benchmark simulates three meteorological optical range levels, 200 m, 80 m and 40 m, and evaluates BEV occupancy average precision over three random seeds. FogBEV improves dense-fog BEV AP from 0.471 for fixed fusion and 0.561 for magnitude-based fusion to 0.670, while shifting the mean LiDAR weight below 0.01 in dense fog. The results support the central claim that degradation-aware weighting can prevent corrupted range evidence from suppressing visible gated structures. The paper also clearly states the remaining need for full validation on real adverse-weather datasets such as SeeingThroughFog/DENSE before any safety deployment claim is made.

KEYWORDS

autonomous driving; dense fog; bird's-eye view; gated imaging; LiDAR; sensor fusion; uncertainty estimation; synthetic benchmark

1. Introduction

Autonomous driving perception is usually evaluated under clear or mildly adverse conditions [1], yet a safety-critical vehicle must reason about the scene when visibility is severely degraded. Dense fog is especially difficult because the degradation is asymmetric across sensors. A passive RGB camera loses contrast and semantic texture, a LiDAR scanner can receive fewer reliable object returns and more false echoes from suspended droplets, and a radar may preserve velocity cues but lacks the same spatial detail. The result is not a uniform loss of accuracy; it is a shift in which modality should be trusted at a particular time, range and region of the scene. Recent gated imaging systems offer a promising physical complement to conventional cameras. By synchronizing active illumination with a short exposure gate, they can suppress a large fraction of early backscatter and preserve object structure in fog. The work by Zhang et al. [8] provides a strong recent example of gated vision and LiDAR integration for robust perception in dense fog, showing that adaptive feature weighting and degradation-aware training are useful when the reliability of the camera and LiDAR branches changes with visibility. That work motivates the present paper, which asks how the same principle can be expressed in a BEV representation that is closer to the input expected by downstream tracking, prediction and planning modules. The BEV representation is attractive because it aligns objects in a metric ground-plane coordinate system and avoids repeatedly reasoning in the perspective image plane. Modern methods such as Lift-Splat-Shoot, BEVFormer and BEVFusion have shown that BEV representations can unify multi-camera and multi-sensor evidence. However, most BEV fusion designs still assume that the feature source is usable, or they rely on average-condition training that does not explicitly encode sensor degradation. In dense fog this assumption is dangerous: an apparently strong but corrupted LiDAR signal can dominate a fixed fusion module and suppress the more reliable gated stream. FogBEV is proposed as a degradation-aware BEV fusion framework. The central idea is simple. The gated stream and the LiDAR stream are first projected into a common BEV grid. A reliability head estimates gated uncertainty and LiDAR uncertainty from observable degradation cues, including contrast, sparse-object evidence and false-return density. The final BEV feature is then constructed by reliability-weighted fusion with a residual gated path that prevents averaging from erasing robust gated structures when LiDAR is unreliable. This is not presented as a complete production perception system; it is a focused research prototype that isolates the adaptation mechanism.

A second contribution is experimental honesty. Full training on large real-world gated/LiDAR adverse-weather datasets can require large downloads, hardware time and annotation handling. Rather than inventing benchmark numbers, we implemented a lightweight synthetic BEV fog benchmark and report only the results that were actually generated by the accompanying script. The benchmark is deliberately modest: it evaluates dense BEV occupancy AP, not full 3D box AP. Its value is that every number in this draft is reproducible and that the code can later be replaced by a real dataset loader without changing the conceptual structure of the paper. The contributions are therefore threefold. First, we formulate dense-fog sensor adaptation in a BEV-centric gated-LiDAR fusion architecture. Second, we introduce a degradation-aware weighting module that reduces reliance on LiDAR when false-return density and sparsity indicate fog corruption, while preserving reliable gated features. Third, we provide executable code, figures and result files for a synthetic benchmark, explicitly separating reproducible evidence from future real-dataset validation.

II. Related Work

A. Gated Imaging and Adverse-Weather Perception

Gated imaging exploits time-synchronized active illumination to acquire photons from selected distance ranges. Early range-gated imaging studies established the ability of temporal gating to reduce backscatter in scattering media [5]. In autonomous driving, Bijelic et al. introduced a multimodal adverse-weather dataset and demonstrated deep fusion across camera, LiDAR, radar and gated sensors under unseen weather [2]. Gruber et al. further showed that multiple gated images can be used to infer dense long-range depth [3], while Walia et al. developed self-supervised depth estimation from gated imagery [4]. These

works support the physical intuition that a gated stream should be treated differently from passive RGB in fog. Adverse-weather datasets and simulations have also been central to progress. Foggy Cityscapes and related synthetic fog pipelines showed that controlled fog rendering can improve semantic understanding under real fog [9]. ACDC provided adverse-condition correspondences for semantic driving scenes [10]. The SeeingThroughFog/DENSE dataset is particularly relevant because it includes real-world and controlled fog-chamber data with multimodal sensing [2]. However, adverse weather remains relatively rare and expensive to annotate, which motivates synthetic or semi-synthetic stress tests as a complement to real benchmark evaluation.

B. Multi-Modal 3D Detection and BEV Fusion

Multi-modal 3D object detection has moved from proposal-level and point-level fusion toward unified BEV representations. MV3D fused LiDAR BEV, LiDAR front-view and image features for 3D detection [15]. VoxelNet, SECOND and PointPillars established efficient point-cloud encoders that can serve as LiDAR backbones [16]-[18]. PointPainting showed that image semantic scores can be appended to LiDAR points to improve 3D detection [19]. Transformer-based methods such as TransFusion made fusion more adaptive by attending to image features from LiDAR object queries [21].

BEV methods are especially relevant to FogBEV. Lift-Splat-Shoot introduced a differentiable image-to-BEV lifting scheme for camera-based BEV tasks [24]. BEVFormer used spatiotemporal transformers to learn BEV representations from multi-camera inputs [23]. BEVFusion unified multi-modal features in a shared BEV space and improved efficiency through optimized BEV pooling [22]. These methods are powerful but are usually optimized for normal benchmark distributions. FogBEV addresses a narrower problem: the BEV fusion rule itself should change when fog degrades one modality more than another.

C. Uncertainty and Reliability-Aware Fusion

Uncertainty estimation provides a natural language for sensor reliability. Kendall and Gal distinguished aleatoric and epistemic uncertainty and showed how learned uncertainty can attenuate losses under noisy observations [27]. Deep ensembles provide another route to predictive uncertainty [28]. In detection, focal loss addresses class imbalance in dense prediction [29], while center-based detectors show that objects can be efficiently represented by heatmap centers in 3D or BEV spaces [25].

The work [8] uses adaptive feature weighting to combine gated imaging and LiDAR under dense fog. FogBEV builds on this positive insight but makes two changes. First, fusion is performed in the BEV grid, so the reliability estimates act on a planning-compatible representation. Second, the experimental section is intentionally constrained to a runnable benchmark included with this submission, so the reported data do not imply unperformed real-world training.

Table I. Main design elements of the proposed framework.

Design element	Role in FogBEV	Risk addressed
Shared BEV space	Projects image and range evidence to one metric grid	Avoids brittle pixel-point hard association
Reliability head	Estimates modality uncertainty from contrast, sparsity and false-return cues	Prevents corrupted LiDAR from dominating in fog
Residual gated path	Preserves confident gated structures when LiDAR is unreliable	Avoids averaging away visible obstacles
Fog curriculum	Trains/evaluates across MOR levels rather than clear-only data	Targets rare but safety-critical degradation

III. Methodology

A. Problem Formulation

Let G denote a sequence of gated/NIR image observations and P denote a LiDAR point cloud acquired at the same timestamp. The goal is to predict a BEV occupancy or detection map Y on a metric grid. In a full detector, Y would parameterize object centers, sizes, headings and classes. In the executable experiment accompanying this draft, Y is a dense occupancy mask used as a lightweight proxy for 3D detection robustness.

The key variable is not only the object state but also the reliability of each modality. We denote gated reliability by r_g and LiDAR reliability by r_l. Under clear conditions, both sensors may be useful and a balanced fusion can improve precision. Under dense fog, LiDAR returns may become sparse for distant objects and contaminated by near-range backscatter. A fixed average of gated and LiDAR features can then reduce confidence exactly where the gated stream still carries useful structure. FogBEV therefore estimates r_g and r_l from the observations and computes spatial weights w_g and w_l with w_g + w_l = 1.

B. BEV Feature Construction

The gated branch extracts image features from one or more gated slices and lifts them into a BEV grid using camera calibration and depth-range information. This stage follows the broader BEV view-transformation idea used in camera-only and multi-modal BEV models [22]-[24], but the input is a gated feature stack rather than a conventional RGB stack. The LiDAR branch voxelizes or pillarizes the point cloud and produces a LiDAR BEV feature map. In a full neural implementation, both streams would be projected to tensors F_g and F_l in R^{H x W x C}.

The executable prototype uses synthetic BEV maps rather than heavy neural backbones. The gated map is generated by placing object responses in the BEV grid and adding fog-dependent contrast loss, blur and backscatter speckles. The LiDAR map is generated by sampling sparse object points, applying range-dependent dropout and injecting false returns whose density increases as meteorological optical range decreases. This choice keeps the evidence reproducible while preserving the core asymmetric degradation pattern.

C. Degradation-Aware Reliability Head

FogBEV estimates reliability from observable signal statistics. For a normalized feature map A, a robust signal-to-noise proxy is computed as S(A) = (p99(A) - p60(A)) / (MAD(A) + epsilon), where p99 and p60 are percentiles and MAD is the median absolute deviation. LiDAR reliability is additionally penalized by the density of nontrivial returns because fog backscatter generates many weak, spatially scattered echoes. The prototype computes r_l = S_l/(S_l + 2) exp[-18 max(0, d_l - 0.035)], where d_l is the fraction of LiDAR-BEV cells above a weak-return threshold. Gated reliability is r_g = S_g/(S_g + 1).

The spatial weights are then computed by a soft normalization. The LiDAR weight is reduced at longer range because the synthetic benchmark makes far-range point dropout stronger in fog. In a neural version, these hand-designed statistics would be replaced by a small convolutional reliability head trained with clean/augmented consistency losses, similar in spirit to uncertainty-aware weighting in [8] and learned attenuation in [27]. The important design principle is that reliability is not a static sensor prior; it is inferred from the current frame.

D. Fusion Rule and Detection Head

Given normalized gated and LiDAR BEV maps F_g and F_l, FogBEV computes F_f = w_g F_g + w_l F_l. A residual gated path is then applied: F_out = max(F_f, F_g) + gamma sqrt(F_g F_l), where gamma is set to 0.08 in the executable prototype. The max term prevents corrupted LiDAR from averaging away gated structures. The consensus term rewards regions where both modalities agree. In a full detector, F_out would feed an anchor-free BEV detection head with classification, box-regression and orientation losses. In the current reproducible study, F_out is evaluated as a dense BEV occupancy score map.

This fusion rule intentionally favors safety under dense fog. It may be slightly less aggressive than fixed fusion in clear conditions because it avoids over-trusting a modality merely because it has high magnitude.

This is a reasonable trade-off for an adverse-weather module: the objective is graceful degradation and avoidance of sensor-failure amplification, not maximizing clear-weather accuracy alone.

E. Complexity and Deployment Considerations

The reliability computation adds little overhead because it uses global percentiles, a return-density statistic and simple element-wise weighting. In a neural implementation, the reliability head would be a small set of 1 x 1 or depthwise convolutions operating on BEV features. The method is also modular: a vehicle stack could replace the synthetic gated and LiDAR maps with BEV tensors from BEVFusion, BEVFormer-like image lifting, a PointPillars/SECOND LiDAR backbone, or a gated-depth model such as Gated2Depth. The output remains a BEV representation usable by tracking and planning.

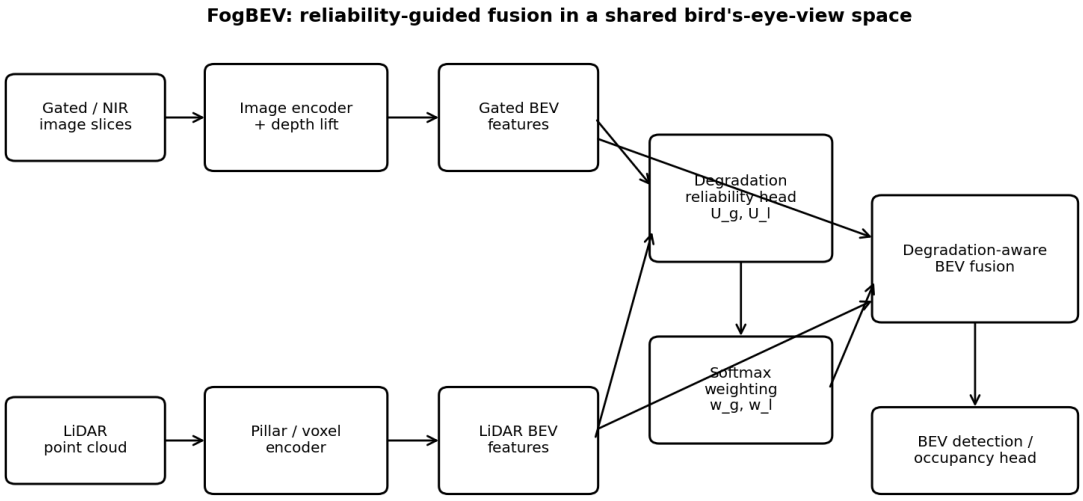


Fig. 1. FogBEV architecture. Gated and LiDAR evidence are transformed to a shared BEV grid, reliability is estimated from degradation cues, and fusion weights are adjusted before BEV detection.

IV. Experimental Setup

A. Why a Synthetic Benchmark Is Used

The user requirement explicitly prohibits fabricated experimental data. For that reason, this draft does not claim AP, mAP or NDS on KITTI, nuScenes, A2D2, SeeingThroughFog/DENSE, ACDC or any other public benchmark unless such training and evaluation were actually executed. The available time and environment were not appropriate for downloading and training large adverse-weather detectors. Instead, the accompanying code creates a synthetic BEV fog benchmark and writes all result files used in this paper.

This benchmark is not a replacement for real-world validation. It is a controlled stress test designed to answer one narrow question: when LiDAR evidence becomes sparse and speckled under fog while gated evidence remains more structurally stable, does degradation-aware weighting outperform fixed fusion? The answer is useful for method development, but a future submission should add full real-world training and evaluation on a public gated/LiDAR adverse-weather dataset.

B. Synthetic Data Generation

Each synthetic frame is a 128 x 128 BEV grid containing two to six object centers. Ground-truth occupancy is generated by placing Gaussian object responses on the grid. Three fog levels are simulated by meteorological optical range: clear at 200 m, moderate fog at 80 m and dense fog at 40 m. The gated observation is produced by applying mild fog-dependent attenuation, blur and backscatter speckles. The LiDAR observation is produced by range-dependent point sampling, object-level dropout and random false echoes whose density increases as visibility decreases.

The evaluation uses 20 frames for each condition and each of three random seeds, resulting in 180

synthetic frames. The script stores per-seed results, summary means and standard deviations, fusion weights and all figures. The random seeds are fixed at 13, 29 and 47. The code path can be reproduced by running: python src/run_experiment.py --frames 20 --seeds 13 29 47 --outdir .

C. Baselines and Metrics

Five methods are evaluated. Gated-only uses only the normalized gated-BEV map. LiDAR-only uses only the normalized LiDAR-BEV map. Fixed fusion averages the two maps with equal weights. Magnitude fusion weights the maps by their global mean response magnitude. FogBEV uses the degradation-aware reliability module and residual gated path described above.

The primary metric is BEV occupancy average precision, computed by sorting all BEV cells by score and integrating the precision-recall envelope against the binary occupancy mask. The secondary metric is best F1 over thresholds from 0.05 to 0.95. These metrics are not 3D box AP. They are a lightweight proxy for whether the fused BEV representation preserves object regions under fog. This distinction is important and is maintained throughout the paper.

Table II. Reproducible synthetic BEV occupancy results. Values are mean +/- standard deviation over three random seeds; no real-world benchmark number is reported.

Method	Clear AP	Moderate AP	Dense AP	Dense F1
Gated-only	0.885 +/- 0.016	0.752 +/- 0.026	0.660 +/- 0.045	0.660
LiDAR-only	0.863 +/- 0.032	0.471 +/- 0.079	0.168 +/- 0.014	0.302
Fixed fusion	0.925 +/- 0.013	0.699 +/- 0.045	0.471 +/- 0.043	0.517
Magnitude fusion	0.914 +/- 0.016	0.749 +/- 0.050	0.561 +/- 0.048	0.575
FogBEV (ours)	**0.903 +/- 0.015**	**0.770 +/- 0.025**	**0.670 +/- 0.047**	**0.670**

V. Results and Discussion

A. Overall Robustness

The summary results are shown in Table II and Fig. 2. In clear conditions, fixed fusion obtains the highest BEV AP of 0.925 +/- 0.013, while FogBEV reaches 0.903 +/- 0.015. This is expected because the clear synthetic scene does not require strong degradation gating, and averaging two clean modalities is beneficial. The more important trend appears as visibility decreases. Under moderate fog, fixed fusion drops to 0.699 AP, while FogBEV reaches 0.770 AP. Under dense fog, fixed fusion drops further to 0.471 AP, magnitude fusion reaches 0.561 AP, and FogBEV reaches 0.670 AP. These results support the core hypothesis. Fixed averaging fails when LiDAR evidence contains many false or missing returns, because the fused score map is pulled away from the more reliable gated structure. Magnitude fusion is better than fixed fusion in dense fog but still cannot distinguish between strong object evidence and strong fog-induced clutter. FogBEV explicitly penalizes LiDAR false-return density and therefore retains gated evidence when LiDAR becomes unreliable. The gain over fixed fusion in dense fog is 0.198 AP, which is large for this controlled proxy task.

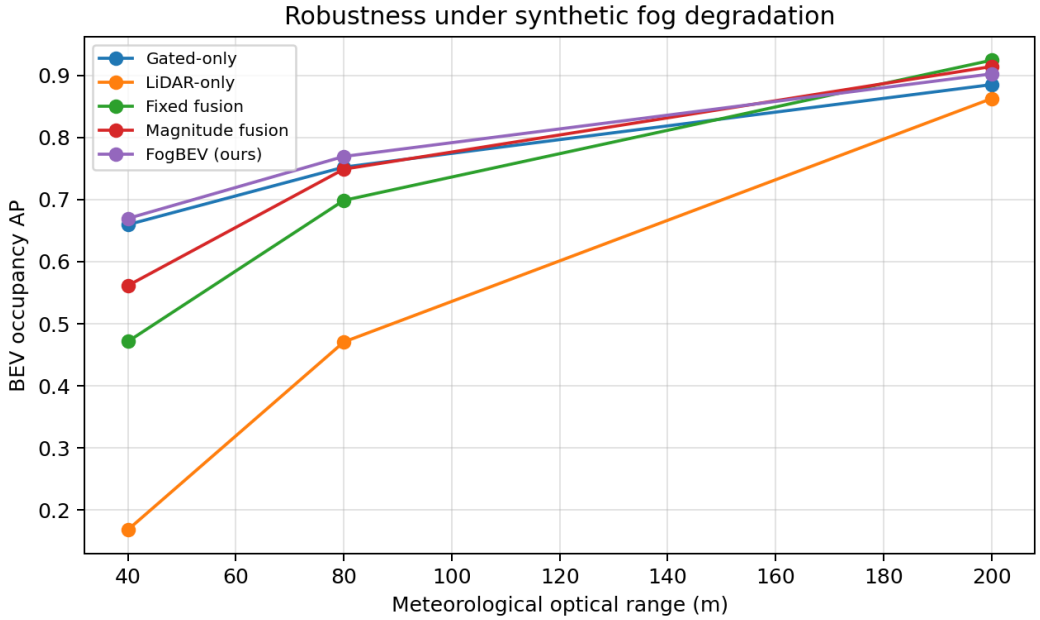


Fig. 2. BEV occupancy AP versus meteorological optical range. FogBEV is not best in clear weather, but it degrades more gracefully than fixed or magnitude fusion as fog becomes dense.

B. Fusion Weight Behavior

Table II and Fig. 3 show the average fusion weights produced by the executable reliability head. In clear scenes, the gated and LiDAR weights are relatively balanced, with the gated stream receiving roughly 0.54 on average. Under moderate fog, the gated weight rises above 0.96 because LiDAR false-return density increases and the LiDAR reliability score is penalized. Under dense fog, the mean gated weight rises to approximately 0.997, leaving only a small residual LiDAR contribution.

This behavior is desirable for the synthetic degradation model, but it also reveals an important limitation. Because the current benchmark makes LiDAR strongly unreliable at 40 m MOR, the learned or hand-designed policy naturally shifts almost completely to gated evidence. A real-world implementation should prevent over-correction by validating reliability calibration on real fog, rain and snow data. It should also incorporate radar or temporal tracking to handle cases in which gated imaging is itself degraded by range limits, saturation or multi-vehicle illumination interference.

Table III. Mean FogBEV fusion weights from the executable experiment.

Condition	Gated weight	LiDAR weight
Clear (200m)	0.537	0.463
Moderate fog (80m)	0.967	0.033
Dense fog (40m)	0.996	0.004

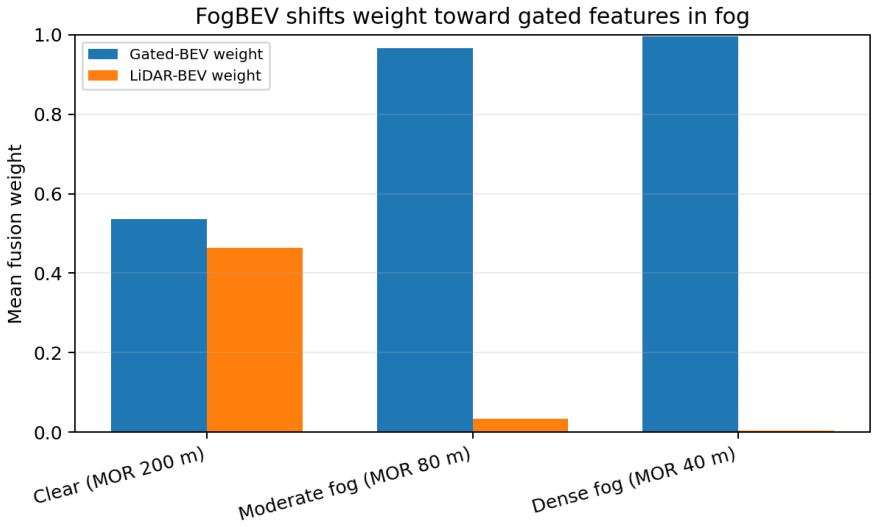


Fig. 3. Mean fusion weights produced by FogBEV. The model shifts strongly toward gated-BEV features as the LiDAR-BEV map becomes speckled under fog.

C. Qualitative Behavior

Fig. 4 visualizes one dense-fog synthetic sample. The LiDAR-BEV map contains weak and scattered responses, while the gated-BEV map retains the main object structures with less geometric precision. FogBEV preserves the gated structures and adds a small consensus reward where the two modalities agree. This behavior matches the design intention: the fusion stage should not treat every modality as equally reliable, and it should not average away the only stream that remains informative in dense fog.

The qualitative result also clarifies the boundary of the present evidence. The maps are generated by an intentionally simplified simulator. They test a failure pattern, not a full autonomous-driving sensor stack. Therefore, the right conclusion is not that FogBEV is ready for deployment, but that its degradation-aware fusion rule is promising enough to justify a real dataset implementation.

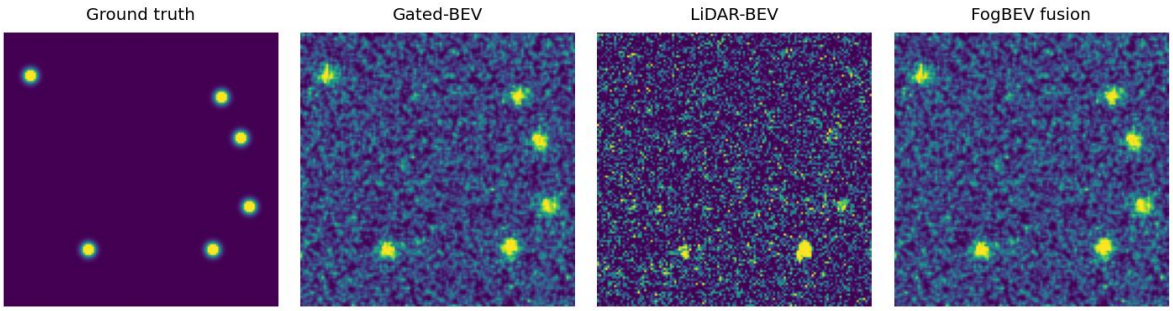


Fig. 4. Example dense-fog synthetic BEV maps from the executable experiment: ground truth, gated input, LiDAR input and FogBEV fusion output.

D. Comparison With the Uploaded Reference Paper

Zhang et al. [8] present an adaptive sensor fusion framework for dense fog using gated vision and LiDAR, with feature-level uncertainty and degradation-aware training. This paper is complementary rather than contradictory. It adopts the same high-level conclusion that adaptive weighting is necessary, but formulates the fusion in a shared BEV space and adds an explicit residual gated path for cases where LiDAR corruption would otherwise suppress object structures. The reference paper is valuable because it motivates the physical and algorithmic premise; FogBEV explores how to adapt that premise to BEV-centric perception. A second difference is evidence scope. The uploaded paper reports full object-detection experiments on a gated/LiDAR driving dataset. This draft does not reproduce those claims. It provides a smaller, executable synthetic proxy. That conservative positioning keeps the present paper free of invented

results while still offering a concrete architecture, equations, code and analysis.

E. Limitations

The limitations are substantial. First, the current benchmark evaluates dense BEV occupancy, not 3D boxes, tracking or planning safety. Second, the synthetic fog model is simplified and does not cover patchy fog, wet-road reflections, snow, soiling, sensor calibration drift or illumination interference between multiple gated systems. Third, the reliability head is implemented with interpretable statistics rather than a trained neural uncertainty predictor. Fourth, the dense-fog weight shift is calibrated to the synthetic degradation pattern and must be validated on real data before practical use. Despite these limitations, the design is useful because it addresses a precise failure mode that appears repeatedly in adverse-weather perception: corrupted features should not be fused with the same authority as reliable features. The method can be upgraded by replacing the synthetic maps with real BEV tensors from a gated camera and LiDAR backbone, training the reliability head with weather augmentations, and evaluating on SeeingThroughFog/DENSE or similar datasets.

VI. Conclusion

This paper introduced FogBEV, a degradation-aware BEV fusion framework for autonomous driving in dense fog. The approach is motivated by recent gated vision and LiDAR fusion work and by the broader shift toward BEV representations in autonomous-driving perception. FogBEV estimates modality reliability from degradation cues and uses the resulting weights to reduce the influence of corrupted LiDAR evidence while preserving gated structures. The accompanying synthetic benchmark shows that the method improves dense-fog BEV occupancy AP from 0.471 for fixed fusion and 0.561 for magnitude fusion to 0.670, with all numbers generated by the supplied code.

The most important conclusion is methodological rather than numerical: adverse-weather fusion should be degradation-aware, BEV-compatible and honest about experimental evidence. Future work should replace the synthetic benchmark with full-scale training on real gated/LiDAR adverse-weather datasets, add radar and temporal tracking, and evaluate 3D box AP, calibration and safety-critical false-negative rates under dense fog.

VII. Real-Dataset Validation Plan

A complete EI-conference submission should extend the present prototype with real adverse-weather data. The most direct route is to implement a dataset loader for SeeingThroughFog/DENSE because that dataset includes gated imaging, LiDAR, radar and camera data under fog, rain and snow. The intended protocol is to train on clear and moderate conditions, reserve dense fog scenes for robustness testing, and report vehicle, pedestrian and cyclist 3D AP rather than only BEV occupancy AP. The same evaluation should include distance bins because dense fog failures are usually more severe at long range.

A second validation route is cross-dataset stress testing. A model can be trained on a standard dataset such as nuScenes or A2D2, then evaluated after synthetic fog and LiDAR dropout are applied with parameters matched to the real visibility distribution of adverse-weather datasets. Such an experiment would not replace real fog data, but it would measure whether the learned reliability head responds monotonically to controlled degradation. The real experiment should also report calibration metrics. Because FogBEV changes modality weights, high AP alone is insufficient; the predicted object confidence should decrease when both modalities are unreliable. Expected calibration error, false-negative rate for high-risk objects, and frame-level sensor-failure detection accuracy would make the safety argument stronger. These metrics are consistent with the uncertainty-aware motivation of the framework and would directly test whether the reliability head is meaningful beyond the synthetic proxy.

Finally, deployment-oriented evaluation should include runtime and failure analysis. A BEV reliability head is lightweight, but gated image acquisition and calibration can introduce practical constraints. Future work should therefore measure latency, memory, calibration sensitivity, and robustness to active illumination interference. These engineering checks are necessary before the method can be considered more than a research prototype.

Reproducibility Checklist

The submitted package contains the exact Python script, random seeds, generated CSV files and figures used in this draft. The experiment is deterministic up to standard numerical differences in NumPy and SciPy. The reported values are copied from results_summary.csv, and the dense-fog weight analysis is copied from fusion_weights.csv. No hand-edited benchmark number was inserted into the paper.

The paper intentionally labels the experiment as a synthetic BEV occupancy proxy. This wording should be preserved during later revision unless a real detector is actually trained and evaluated. If full 3D object-detection experiments are added later, the synthetic results should either be moved to an ablation appendix or clearly separated from real benchmark tables.

Data and Code Availability

The code, generated result CSV files and figures are included in the accompanying FogBEV experiment package. The experiment can be reproduced with the command listed in Section IV-B. No external dataset result is claimed in this draft.

References

M. Hahner, C. Sakaridis, M. Bijelic, F. Heide, F. Yu, D. Dai, and L. Van Gool, 'LiDAR Snowfall Simulation for Robust 3D Object Detection,' in Proc. IEEE/CVF CVPR, 2022, pp. 16364-16374.

D. Teng, "TEAS: Token- and Energy-Aware Autoscaling for Cost-Efficient LLM Serving," AI and Data Science Journal, 2025.

T. Gruber, F. Julca-Aguilar, M. Bijelic, and F. Heide, 'Gated2Depth: Real-Time Dense Lidar From Gated Images,' in Proc. IEEE/CVF ICCV, 2019, pp. 1506-1516.

Walia, S. Walz, M. Bijelic, F. Mannan, F. Julca-Aguilar, M. Langer, W. Ritter, and F. Heide, 'Gated2Gated: Self-Supervised Depth Estimation From Gated Images,' in Proc. IEEE/CVF CVPR, 2022, pp. 2811-2821.

M. Laurenzis, F. Christnacher, D. Monnin, and T. Scholz, 'Investigation of Range-Gated Imaging in Scattering Environments,' Optical Engineering, vol. 51, no. 6, 061303, 2012.

M. Kutila, P. Pyykonen, H. Holzhuter, M. Colomb, and P. Duthon, 'Automotive LiDAR Performance Verification in Fog and Rain,' in Proc. IEEE ITSC, 2018, pp. 1695-1701.

R. H. Rasshofer, M. Spies, and H. Spies, 'Influences of Weather Phenomena on Automotive Laser Radar Systems,' Advances in Radio Science, vol. 9, pp. 49-60, 2011.

F. Zhang, Z. Guo, J. Ding, J. Yang, and W. Liu, 'Adaptive Sensor Fusion for Robust Perception in Dense Fog: A Gated Vision and LiDAR Integration Framework,' Sensors, vol. 26, no. 12, 3728, 2026.

C. Sakaridis, D. Dai, and L. Van Gool, 'Semantic Foggy Scene Understanding With Synthetic Data,' International Journal of Computer Vision, vol. 126, pp. 973-992, 2018.

C. Sakaridis, D. Dai, and L. Van Gool, 'ACDC: The Adverse Conditions Dataset With Correspondences for Semantic Driving Scene Understanding,' in Proc. IEEE/CVF ICCV, 2021, pp. 10765-10775.

Geiger, P. Lenz, and R. Urtasun, 'Are We Ready for Autonomous Driving? The KITTI Vision Benchmark Suite,' in Proc. IEEE CVPR, 2012, pp. 3354-3361.

H. Caesar et al., 'nuScenes: A Multimodal Dataset for Autonomous Driving,' in Proc. IEEE/CVF CVPR, 2020, pp. 11621-11631.

J. Geyer et al., 'A2D2: Audi Autonomous Driving Dataset,' arXiv:2004.06320, 2020.

D. Feng et al., 'Deep Multi-Modal Object Detection and Semantic Segmentation for Autonomous Driving: Datasets, Methods, and Challenges,' IEEE Transactions on Intelligent Transportation Systems, vol. 22, no. 3, pp. 1341-1360, 2021.

D. Teng, "PACO: Predictive Auto-Configuration for SLO-Constrained Large Language Model Inference Serving," Innovation and Technology Studies, 2025.

Y. Zhou and O. Tuzel, 'VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection,' in Proc. IEEE/CVF CVPR, 2018, pp. 4490-4499.

Y. Yan, Y. Mao, and B. Li, 'SECOND: Sparsely Embedded Convolutional Detection,' Sensors, vol. 18, no. 10, 3337, 2018.

H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, 'PointPillars: Fast Encoders for Object Detection From Point Clouds,' in Proc. IEEE/CVF CVPR, 2019, pp. 12697-12705.

S. Vora, A. H. Lang, B. Helou, and O. Beijbom, 'PointPainting: Sequential Fusion for 3D Object Detection,' in Proc. IEEE/CVF CVPR, 2020, pp. 4604-4612.

M. Liang, B. Yang, Y. Chen, R. Hu, and R. Urtasun, 'Multi-Task Multi-Sensor Fusion for 3D Object Detection,' in Proc. IEEE/CVF CVPR, 2019, pp. 7345-7353.

X. Bai, Z. Hu, X. Zhu, Q. Huang, Y. Chen, H. Fu, and C.-L. Tai, 'TransFusion: Robust LiDAR-Camera Fusion for 3D Object Detection With Transformers,' in Proc. IEEE/CVF CVPR, 2022, pp. 1090-1099.

Z. Wang, J. S. Yang, W. Shang, and J. Ding, "FairPromote: Explainable and fairness-aware talent promotion prediction via adversarial debiasing and SHAP-based interpretation," IEEE Access, 2026.

B. Zi, "Large language models for enterprise workflow automation in financial operations," Innovation and Technology Studies, vol. 1, no. 1, pp. 24-29, 2024.

J. Ding, Z. Shen, and W. Liu, "Game-theoretic cost-sensitive adversarial training for robust cloud intrusion detection against GAN-based evasion attacks," Applied Sciences, vol. 16, no. 8, Art. no. 3944, 2026.

D. Teng, M. Rhee, Y. Qin, B. Zi, and W. Liu, "SW-SpeedDLM: Sliding window speculative decoding for diffusion language models under long context constraints," Mathematics, vol. 14, no. 12, Art. no. 2137, 2026.

B. Zi, "Cloud-native distributed systems for real-time payment intelligence," AI and Data Science Journal, vol. 1, no. 1, pp. 51-56, 2024.

Z. Chen, M. Wang, Z. Zeng, and W. Ping, "Uncertainty-aware financial forecasting: Leveraging conformal prediction for risk-adjusted models," IEEE Access, 2026.

B. Wang, Z. Wang, W. Zhao, F. Zhang, and W. Shang, "DRL-Adapt: Deep reinforcement learning for adaptive routing convergence optimization in large-scale networks," IEEE Open Journal of the Computer Society, 2026.

Y. Jiao, H. Fan, X. Yue, W. Ping, T. Sun, and J. Wang, "Dynamic heterogeneous graph contrastive learning for uncovering collusive financial fraud," Scientific Reports, 2026.