Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira,
Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria
 Published by Research Publishing, Singapore.
doi: 10.3850/ESREL2026061419_esrel26-p26281-cd

Reliable Perception under Veiling Glare: Synthetic Glare Generation and Realism Validation for Automated Vehicles

Kazi Kowshin Raihana
German Aerospace Center (DLR), Germany. E-mail: kazi.raihana@dlr.de

Shrijal Pradhan
German Aerospace Center (DLR), Germany.. E-mail: shrijal.pradhan@dlr.de

Redge Melroy Castelino
German Aerospace Center (DLR), Germany. E-mail: redge.castelino@dlr.de

Arne Lamm
German Aerospace Center (DLR), Germany. E-mail: arne.lamm@dlr.de

Veiling glare caused by sunlight critically compromises camera-based perception in Automated Vehicles (AVs) by reducing image contrast and visibility. This phenomenon, particularly prominent during daytime driving, degrades object and lane detection performance and undermines safety-critical AV decision-making. Although several studies have addressed glare detection and mitigation, the limited availability and diversity of real-world glare datasets hinder the development and validation of automotive perception systems under such challenging conditions. This paper introduces a synthetic veiling glare injection approach for camera-based AV datasets, enabling controlled synthesis of glare effects by parameterizing glare shape, size, position, color, number, intensity, and transparency to realistically emulate sun-induced veiling glare. Two multi-metric evaluations are presented to assess the perceptual, photometric, and physical consistency of the injected glare through comparative analyses: (i) synthetic glare versus glare-free data, and (ii) synthetic glare versus naturally captured glare from the same dataset. The comparative results demonstrate that the proposed injection process effectively reproduces the dominant characteristics of real-world veiling glare with mild to moderate severity, showing strong agreement in metric behavior. Consequently, the proposed injection and evaluation framework provides a validated, glare-aware foundation for safety assessment in automated driving systems.

Keywords: synthetic glare, veiling glare, glare injection, SSIM, PSNR, Contrast difference, spatial frequency power difference.

1. Introduction

Automated Driving Systems (ADS) represent a transformative technology which promises numerous benefits to society, such as enhanced road safety, improved traffic efficiency and improved mobility access Shafiee et al. (2021). However, the deployment of fully ADS are constrained by challenges to assure safe and reliable perception given real world environment conditions that compromise system performance Kumar and Muhammad (2023); Yuxiao Zhang et al. (2023); Fenglian Pan et al. (2024). Camera-based perception, which plays a central role in environmental understanding for modern AVs, is particularly vulnerable to adverse conditions that degrade visual sensor capabilities Kumar and Muhammad (2023); Yuxiao Zhang et al. (2023); Yuan et al. (2026).

Among the diverse environmental factors affecting camera perception reliability, glare from intense sunlight, headlights, or surface reflections presents a critical safety risk Chen et al. (2021); Lucie Yahiaoui et al. (2020). When camera-based perception systems encounter glare, the resulting degradation manifests as pixel saturation, reduced contrast, and obscured visual information, leading to diminished object detection reliability and potentially catastrophic failures in downstream perception and planning modules.

Existing perception benchmarks and datasets, such as Esfahani and Wang (2021); Chen et al. (2021), fail to adequately represent the prevalence and variability of glare conditions. This creates a critical gap between algorithm performance in controlled testing and real-world conditions.

The scarcity of diverse, labeled glare data poses a fundamental barrier to developing and validating robust perception systems capable of handling these challenging edge case conditions Gray et al. (2023). Real-world glare datasets exist, such as the GLARE traffic sign dataset containing 2,157 images with sun glare. However, they remain limited in terms of glare parameter variation (intensity, size, position, spectral characteristics) and scenario diversity.Zhu and Lee (2025)

Synthetic data generation has emerged as a powerful paradigm for addressing data scarcity challenges in ADS. Synthetic data enables controlled generation of diverse scenarios including rare edge cases and safety-critical conditions.

This synthetic data is generated on top of an ADS dataset to account for domain adoption for the generated glare. Pandey et al. (2011) was selected as the dataset for the application due to available camera positions and camera type e.g. Point Grey Ladybug3 omnidirectional camera system which is generally used in ADS. The Ford dataset contains camera based data on the urban road that matches the type of scenario that are of interest when testing glare-based effects in ADS.

This paper addresses the critical gap between the need for glare-robust perception systems and the scarcity of suitable training and validation data. We present a novel synthetic glare injection pipeline designed to generate realistic glare artifacts with controllable parameters including position, shape, size and numbers, intensity, and transparency

Furthermore, in order to validate the synthetic glare, we analyse multiple image similarity and distance metrics such as Structural Similarity Index Metric (SSIM), Peak Signal-to-Noise Ratio (PSNR), Contrast Difference (CD) and Spatial Frequency Power Difference (PD) to perform a comprehensive quantitative along with a qualitative comparisons against real-world glare data.

The related previous works are discussed in Section 2. The process of veiling glare generation is presented in Section 3. The qualitative and quantitative analysis between synthetic and real images are evaluated in Section 5. The paper is concluded in Section 6.

2. Related Work

Veiling glare appears as a uniform haze with slight color tint, reducing contrast across the image. McCann and Rizzi (2007) show that sun-induced veiling glare under clear-sky conditions often exhibits a bluish component due to atmospheric scattering. Raskar et al. (2008) explained that veiling glare often arises when intense sunlight causes diffuse scattering within camera optics and lens barrels. This leads to reduced visibility of low-contrast scene elements. Such type of glares are common in urban driving scenarios during sunrise and sunset. Hullin et al. (2011) demonstrated physically accurate glare simulation using ray-tracing and path-tracing methods, which are computationally expensive. This motivates simpler, computationally efficient alternatives based on additive composition of static glare textures aligned with the light source. Talvala et al. (2007) proposed a deconvolution-based glare removal approach and modeled Glare Spread Function (GSF), which is commonly cited for understanding glare formation. However, their assumption of a spatially constant GSF neglects variations in glare shape and color across the image. Shoshin and Shvets (2021) generated synthetic veiling glare by superimposing low-frequency luminance masks onto glare-free images for training and evaluation. Their method employed predefined glare masks and reported SSIM values of $\sim$0.8–0.9 and PSNR of $\sim$25–30 dB, indicating effective glare removal. However, the approach lacks controlled glare parametrization and physical validation against real-world measurments. For safety and training evaluation, metrics that quantify veiling glare–induced image quality distortion, as well as photometric and physical characteristics, are essential in automotive imaging.

Prior works report indicative metric ranges for glare-free and glare-affected images in automotive imaging. Existing works by Keatkongsang

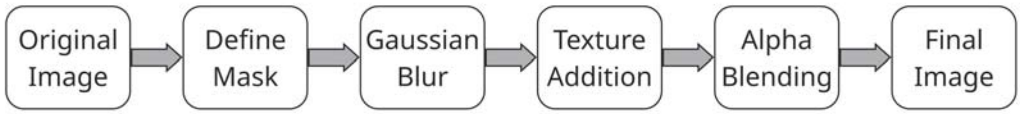

Fig. 1.: Synthetic Glare Injection workflow.

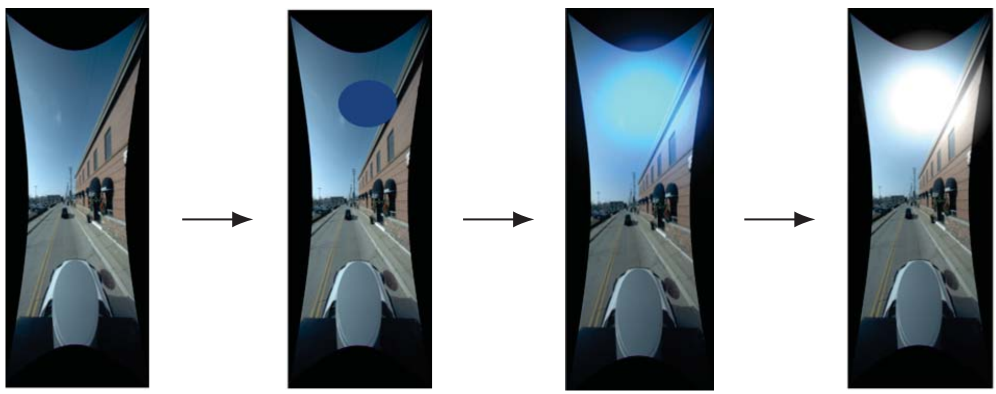

(a) (b) (c) (d)

Fig. 2.: Veiling glare injection step-by-step with arrows indicating the transformation sequence from the original image to the final blended glare result.

et al. (2024); Wang et al. (2004); Imatest LLC (2026b); Sara et al. (2019) reported SSIM of close to 1.0 for reference or undistorted images and SSIM reduced often below about 0.7 for strong distortions w.r.t. scene and method. PSNR values are commonly reported to range from 30–50 dB or higher for high-quality images and decrease to ∼20–30 dB under strong distortions such as glare National Instruments (2024). Moreover, according to Imatest LLC (2015, 2026a), ISO 16505–based Imatest optical methods show contrast reductions of 40% or more under glare. Wolf (2014) observed that stable spectral frequency power corresponds to undistorted images, while glare causes a reduction or alteration.

3. Methodology

This work adopts and extends the approach of Shoshin and Shvets (2021) by superimposing smoothed low-frequency luminance onto glare-free images to simulate real-world veiling glare. Following the concept of Talvala et al. (2007), the injection process models key veiling glare characteristics: large-scale diffused scattering modeled via blurred masks, controlled intensity, spectral color tinting, and partial background visibility through alpha blending. The Sun is assumed as the glare source for evaluation. The proposed injection process extends these concepts by including systematic and controlled parameters for veiling glare injection that is explained in this chapter.

3.1. *Glare Injection Procedure*

The processing stages are summarized in Figure 1. For visual demonstration of the intermediate results, the stages in Figure 2 are additionally processed with alpha blending, the final step in the glare injection process.

3.1.1. *Mask Generation*

A parameterized glare mask matching the image dimensions is defined, allowing control over the following properties:

Glare position: The upper half of the image with X and Y coordinate defined by pixel numbers stating logical position of the blue sky.

Glare shape: Circular or elliptical shape representing the sun.

Glare Size: Dimensions defined by pixel numbers, proportional to the image aspect ratio, generated using the skimage draw function Van Der Walt et al. (2014).

Number of glares: As sun is the source of the veiling glare, the number is set to 1.

Color selection: As evident from the observed Ford data sample in Figure 3a, the Ladybug3 camera sensor perceives lighting from sun with a bluish tint, a uniform base color of blue is a suitable starting point for injecting the veiling glare.

The originally glare-free image shown in Figure 2a is taken and a mask is generated with controlled parameters i.e. glare position= (x:250,y:350), shape= ellipse, size= (122,140), number of glares= 1, Base color= blue. The generated glare in the initial phase is visible on the image as the result of Step1 in Figure 2b.

3.1.2. *Optical blur using Gaussian blur*

Gaussian blur is applied to the initial glare mask to produce a soft, diffused appearance. A controlled blur parameter σ for the Gaussian filter was applied utilizing Skimage library. Larger σ values produce softer glare edges. In Figure 2c, a Gaussian blur with σ = 150 was applied, due to which the solid colored glared area achieves blue colored hazy glow with soft edges.

3.1.3. *Texture addition to Mask*

This step introduces subtle texture variations to simulate a uniform haze around the light source. To achieve such effect this injection process applies texture to the flat glowing spot achieved in Figure 2c. This paper applies the concept of utilizing an array with the size of the image height and width along with random integers from discrete uniform distribution within the interval of the pixel-wise base color (blue) to mathematically add the smooth and subtle variation of uniform haze at the defined position of veiling glare. A full-image array with discrete uniform random intensity variations in the glare mask region is used to introduce low-amplitude texture within the glare region as illustrated in Figure 2d.

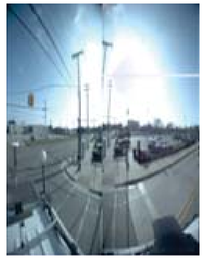

(a)

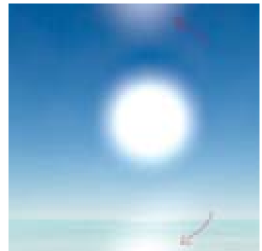

(b)

Fig. 3.: Original glare from Ford Campus Vision and LiDAR dataset (left) and Synthetic glare from Delavennat (2021) (right).

3.1.4. *Alpha blending*

Alpha blending is used to seamlessly overlay the glare mask onto the original image by combining the pixel values calculating the weighted sum and gamma correction. The OpenCV library by Bradski (2000) is used to control mask opacity via the α parameter. Finally, the resulting image after texture addition and the original image are blended using Alpha blending. As mentioned earlier, for the purpose of better visualization, each step in Figure 2 is illustrated after applying alpha blending with α= 0.2.

4. Evaluation metrics

This section describes the evaluation of realism of injected glare through quantitative metrics and a preliminary qualitative filtering step to assess similarity to real-world glare.

4.1. *Qualitative Evaluation*

Qualitative evaluation was conducted by visually comparing injected glare samples to two references: (i) Original images with real-world glare from the Ford Campus Vision and LiDAR dataset, Figure 3a, and (ii) physically-based simulated glare from , Figure 3b.

The evaluation focused on low-frequency luminance haze, bluish tint (specific to the Ladybug3 sensor), logical sun positions, and contrast degradation with partially visible background typical of veiling glare. This served as a filtering step before quantitative evaluation where samples such as Figure 2d were retained while samples with illogical sun glare position or missing bluish tint were discarded.

4.2. *Quantitative Evaluation*

Quantitative evaluation analyzed the perceptual, photometric, and physical realism of injected veiling glare using a multi-metric comparative analysis. Two similarity metrics (SSIM, PSNR) and two difference metrics (CD, PD) were selected as each of these metrics capture the exhibited key aspects of real-world veiling glare in Figure 4.

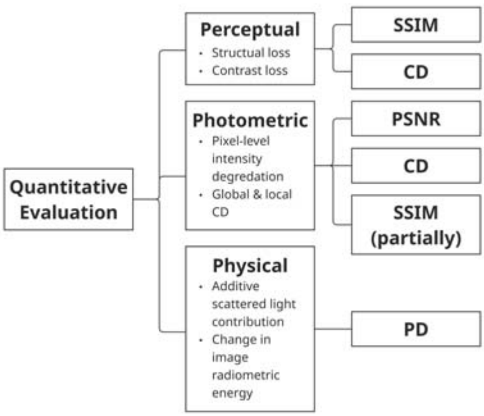

Fig. 4.: Overview of Quantitative Evaluation Metrics.

4.2.1. *Structural Similarity Index (SSIM)*

SSIM measures the structural similarity of two images by jointly evaluating luminance, contrast, and structure. It ranges from 0, indication total dissimilarity, to 1, indicating identical structure. Veiling glare flattens contrast while preserving scene structure, reducing SSIM scores.

4.2.2. *Peak Signal-to-Noise Ratio (PSNR)*

PSNR quantifies pixel-wise similarity between two images by comparing maximum possible signal power and the Mean Squared Error (MSE) computed across pixels Chandler and Hemami (2007). PSNR is sensitive to global brightness shifts caused by veiling glare. Combined with SSIM, it effectively evaluates structural preservation in real-world glare conditions.

4.2.3. *Contrast Difference (CD)*

CD captures local intensity variations, highlighting glare-induced brightening or flattening that global metrics may miss Jiang and Fairchild (2021). CD of 0 indicates identical contrast; higher values indicate greater local mismatch. Veiling glare reduces contrast non-uniformly, lowering this metric value relative to glare-free images.

4.2.4. *Spatial Frequency Power Difference (PD)*

PD measures the frequency domain energy difference between two images. It is calculated by performing 2D Fourier transform of each grayscale image to frequency domain and then squared magnitude of the Fourier coefficient spectrum is computed to connect spatial domain energy to summed spectral energy described by Young et al. (2004). The absolute difference between the mean spectral power for each image representing two images is taken as the value of PD. It mainly captures the total radiometric energy change between two images that can indicate additive low frequency energy contribution by veiling glare according to the concept of Oppenheim and Lim (1980).

A summary of empirical ranges for (i) glare-free image comparisons and (ii) glare-free vs. injected images is shown in Table 1 which is vital for guiding interpretation in subsequent analyses results discussed in Chapter 5.

5. Results

This paper performed multi-metric assessment in the form of the following two comparative analyses showing ranges of observed metric values across 30 samples in order to capture the variability span of glare severity, and verify whether injected glare effects reside within the expected real-world glare magnitude indicated by literature ranges in Table 1.

5.1. *Analysis 1: synthetic vs. glare-free*

This analysis evaluates whether injected veiling glare produces the expected perceptual, photometric, and physical degradation patterns observed

Metrics	Glare free images comparison	Glare free vs. Glare-Affected
SSIM	Very high similarity ∼0.9 to 1.0	Under severe glare can be ∼0.3 to 0.7
PSNR (db)	Less distortion, ∼30-50 dB	Noticeable distortion caused by Severe glare, decreased to ∼20–30 dB
CD	Glare free scenes, High contrast difference	Depending on glare and scene complexity, reduced by about 40% percent or more;
PD	scene texture matched, table power spectrum matching	Visible drop or change in frequency domain components

Table 1.: Indicative literature ranges of image quality metrics for glare-free and glare-induced conditions.

in real-world glare affected images. We evaluate 30 injected glare samples (such as Figure 2d) by comparing them with their corresponding glare-free images (such as Figure 2a) using SSIM, PSNR, CD, and PD, as detailed in Table 2. SSIM values ranged from 0.363–0.758, indicating perceptual and structural degradation from mild (0.75) to strong (0.38) veiling glare while preserving realistic structural retention. PSNR values (28.3–30.6 dB) indicated mild pixel-level deviations, with stronger haze near the light source (28.3 dB) and weaker effects away from it (30.6 dB), producing mild contrast reduction. CD values (0.0–0.222) confirmed local Michelson contrast reduction of 20–25% for samples with upper bound CD (0.222), while the lower bound of CD (0) reflects the non-uniform contrast reduction characteristic of veiling glare. PD values (16,000–170,000) showed strictly positive and substantial frequency-energy redistribution due to blur, confirming the additive scattered-light property of veiling glare. The injected veiling glare samples thus demonstrate high realism matching the perceptual, photometric, and physical effects of mild-to-moderate severity real-world glare, remaining distinctly separate from glare-free images.

Metrics	Analysis-1	Analysis-2
SSIM	0.363–0.758	0.498-0.619
PSNR (dB)	28.3–30.6	28.8-29.5
CD	0.0–0.222	0.02-0.2
PD	16000-170000	51000-180000

Table 2.: Metric value ranges across 30 samples for synthetic vs. glare-free and synthetic vs. real-world comparisons.

5.2. *Analysis 2: synthetic vs. real*

In this analysis, images with injected glare were compared to real-world glare from the Ford Campus Vision and LiDAR dataset to quantify similarity in perceptual, photometric, and physical behavior. Since veiling glare occurs randomly—especially during daytime—this analysis compares 30 injected samples (such as Figure 2d) to 30 real-world glare samples (such as Figure 3a), where image pairs depict different driving scenarios but both contain veiling glare, enabling multi-metric pattern assessment regardless of background differences, as shown in Table 2. SSIM (0.498–0.619) proved comparable to Analysis 1, with the upper bound (0.619) reflecting injected glare's similarity to real-world contrast degradation despite differing scene structures. PSNR (28.8–29.5 dB) exhibited a narrower range (0.7 dB) than injected vs. glare-free comparisons, indicating consistent pixel-level distortions across scenes. CD (0.02–0.2) overlapped the glare-free range, with lower values (0.02) reflecting localized glare and higher values (0.2) indicating realistic contrast distortion, demonstrating that local contrast reduction does not imply global contrast change. PD (51,000–180,000) showed a large, strictly positive shift overlapping Analysis 1, indicating slightly stronger frequency-domain changes in real-world glare.

Despite varying scenarios, this multi-metric evaluation demonstrates that injected veiling glare exhibits perceptual, photometric, and physical characteristics similar to mild-to-moderate real-world glare, ignoring scene-specific artifacts. Real-world samples included stronger glare instances, but injected glare closely matches the reported metric ranges, confirming high realism.

6. Conclusion

This paper presents a controlled, parameterized synthetic veiling glare injection method, tunable via position, shape, size, number, intensity, and transparency, to support automotive safety and reliability validation through scenario and data generation tailored to daytime urban driving The accompanying multi-metric quantitative analysis validates the perceptual, photometric, and physical realism of the injected glare across spatial, pixel, contrast, and frequency domains. By assessing SSIM, PSNR, CD, and PD metric sensitivities in both matching and divergent scenarios, comparisons against glare-free originals and real-world glare samples confirm that the injected effects replicate mild-to-moderate veiling glare severity. Among these metrics, PD excelled in glare-specific validation by capturing low-frequency energy redistribution tied to veiling glare's physical properties. CD effectively highlighted non-uniform local contrast flattening, while SSIM and PSNR collaboratively quantified structural and pixel-level degradation

The proposed glare injection method effectively captures sun-induced veiling glare but remains limited to daytime urban scenarios from the Ford Campus Vision and LiDAR dataset. Future work could extend this approach to veiling glare from diverse light sources (e.g., headlights, reflections) across varied urban driving conditions.

Additionally, the current framework addresses veiling glare exclusively. Future developments could adapt the method to other glare types, such as lens flare, sensor flare and stray light glare.

References

Bradski, G. (2000). The OpenCV Library. *Dr. Dobb's Journal of Software Tools*.

Chandler, D. M. and S. S. Hemami (2007, September). Vsnr: A wavelet-based visual signal-to-noise ratio for natural images. *IEEE Transactions on Image Processing 16*(9), 2284–2298.

Chen, Y., F. Liu, and K. Pei (2021). Self-supervised sun glare detection cnn for self-aware autonomous driving. In *NeurIPS 2021 Workshop on Machine Learning for Autonomous Driving (ML4AD)*.

Delavennat, J. (2021, December). Physically-based real-time glare. Master's thesis, Linköping University, Norrköping, Sweden. Supervisor: Mark E. Dieckmann; Examiner: Jonas Unger.

Esfahani, M. A. and H. Wang (2021). Robust glare detection: Review, analysis, and dataset release. *CoRR abs/2110.06006*.

Fenglian Pan, Yinwei Zhang, Jian Liu, Larry Head, Maria Elli, and Ignacio Alvarez (2024). Reliability modeling for perception systems in autonomous vehicles: A recursive event-triggering point process approach. *Transportation Research Part C: Emerging Technologies 169*, 104868.

Gray, N., M. Moraes, J. Bian, A. Wang, A. Tian, K. Wilson, Y. Huang, H. Xiong, and Z. Guo (2023). Glare: A dataset for traffic sign detection in sun glare.

Hullin, M., E. Eisemann, H.-P. Seidel, and S. Lee (2011, July). Physically-based real-time lens flare rendering. *ACM Trans. Graph. 30*(4), 108:1–108:10.

Imatest LLC (2015). Iso 16505 — camera monitor systems. Accessed: 2026-01-.

Imatest LLC (2026a). Image quality factors (key performance indicators). Current documentation on image quality factors used by Imatest. Accessed: 2026-01-.

Imatest LLC (2026b). Ssim: Structural similarity index. Current documentation. Accessed: 2026-01-.

Jiang, F. and M. D. Fairchild (2021, September). Preliminary result on the direct assessment of perceptible simultaneous luminance dynamic range. *Journal of Imaging Science and Technology 65*, 1–13.

Keatkongsang, W., W. Maythamaluang, A. Maha, and O. Chitsobhuk (2024, April). Reflection removal and facial detection of individuals in vehicles. In *Proceedings of the 2024 10th International Conference on Computing and Data Engineering*, ICCDE '24, New York, NY, USA, pp. 35–41. Association for Computing Machinery.

Kumar, D. and N. Muhammad (2023). Object detection in adverse weather for autonomous driving through data merging and yolov8. *Sensors 23*(20).

Lucie Yahiaoui, Michal Uřičář, Arindam Das, and Senthil Yogamani (2020). Let the sunshine in: Sun glare detection on automotive surround-view cameras. *Electronic Imaging 32*(16), 80–1–80–1.

McCann, J. J. and A. Rizzi (2007). Veiling glare: the dynamic range limit of HDR images. In *Proc. SPIE Electronic Imaging: Human Vision and Electronic Imaging XII*, Volume 6492, pp. 649213. SPIE.

National Instruments (2024). Peak signal-to-noise ratio as an image quality metric. Updated Jun 28, 2024. Accessed: 2026-01-30.

Oppenheim, A. V. and J. S. Lim (1980). The importance of phase in signals. *Proceedings of the IEEE 69*, 529–541.

Pandey, G., J. R. McBride, and R. M. Eustice (2011, November). Ford campus vision and lidar data set. *The International Journal of Robotics Research 30*(13), 1543–1552.

Raskar, R., A. Agrawal, C. A. Wilson, and A. Veeraraghavan (2008, August). Glare aware photography: 4D ray sampling for reducing glare effects of camera lenses. *ACM Trans. Graph. 27*(3), 1–10.

Sara, U., M. Akter, and M. S. Uddin (2019, March). Image quality assessment through fsim, ssim, mse and psnr—a comparative study. *7*, 8–18.

Shafiee, M. J., A. Jeddi, A. Nazemi, P. Fieguth, and A. Wong (2021). Deep neural network perception models and robust autonomous driving systems: Practical solutions for mitigation and improvement. *IEEE Signal Processing Magazine 38*(1), 22–30.

Shoshin, A. and E. Shvets (2021, 07). Veiling glare removal: synthetic dataset generation, metrics and neural network architecture. *Computer Optics 45*, 615–626.

Talvala, E.-V., A. Adams, M. Horowitz, and M. Levoy (2007, 07). Veiling glare in high dynamic range imaging. *ACM Trans. Graph. 26*(3), 37–es.

Van Der Walt, S., J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu (2014, june). scikit-image: image processing in python. *PeerJ 2*, e453.

Wang, Z., A. Bovik, H. Sheikh, and E. Simoncelli (2004, April). Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing 13*(4), 600–612.

Wolf, D. C. (2014). *Modelling Image Quality for Automotive Display Technologies*. Ph. D. thesis, Abertay University. accessed: 2026-01-30.

Young, I., J. Gerbrands, L. Van Vliet, C.-d. Bibliotheek, D. Haag, Y. Theodore, G. Jacob, V. Vliet, and L. Jozef (2004). Fundamentals of image processing.

Yuan, Y., W. Dong, S. Yang, and T. Wu (2026). Awd-yolo enhancing autonomous driving perception reliability in adverse weather. *Scientific Reports 16*(1), 338.

Yuxiao Zhang, Alexander Carballo, Hanting Yang, and Kazuya Takeda (2023). Perception and sensing for autonomous vehicles under adverse weather conditions: A survey. *ISPRS Journal of Photogrammetry and Remote Sensing 196*, 146–177.

Zhu, J. and S. Lee (2025, October). Pbfg: A new physically-based dataset and removal of lens flares and glares. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5448–5457.