Yunfan Lu | Research Statement

ylu066@connect.hkust-gz.edu.cn

Research Aim

Perception is a core bottleneck for intelligent systems. While biological vision delivers low power usage, high dynamic range, and microsecond temporal precision, modern machine vision still struggles under fast motion, extreme lighting, and limited bandwidth. This gap at the sensor and imaging-pipeline level constrains the upper limits of robotics, drones, and embodied intelligence. My long-term research aim is to narrow this gap by pursuing computational imaging systems co-designed with next-generation sensors, enabling machines to perceive with the robustness of biological vision.

Over the past several years, I have explored this direction through event-based sensing, using it as a testbed to rethink how dynamic range, spatial–temporal resolution, and color quality should be integrated into an imaging pipeline. This work has led to new algorithmic principles and system-level insights.

Looking forward, I plan to develop task-driven, end-to-end designs that jointly optimize sensor architecture and learning-based imaging algorithms. Just as biological visual systems evolve specialized structures for specific functions, future machine vision should adopt a similar co-design philosophy.

Research Experience

My research asks how event-based sensors can remove core limits of conventional cameras and support practical, robust imaging systems. Over the past years, my work has evolved along three connected dimensions: robust visibility under extreme conditions, continuous spatial–temporal reconstruction, and human-centric, system-level co-design.

High dynamic range and robust visibility under extreme conditions: I started from the observation that many real scenes are limited not by recognition models, but by the visibility of the input. Building on this idea, Evlight++ [9] constructs a large-scale real-world low-light dataset and an event-guided enhancement pipeline that stabilizes textures when RGB signals are close to sensor noise. In SEE [8], I formulate adaptive brightness adjustment as an event-guided regression problem, enabling flexible control over a broad light range. Most recently, we move beyond lighting to atmospheric degradation, using an event-guided diffusion framework for dehazing [7]. Across these works [7,8,9], events evolve from a high-dynamic-range sensor to a general signal for stabilizing visibility under hard conditions.

Exploiting temporal precision for spatial–temporal reconstruction: Event cameras offer microsecond temporal resolution, which I use to rethink how we reconstruct and correct videos. My CVPR 2023 work [10] introduces spatial–temporal implicit neural representations (INRs) driven by events, treating video as a continuous function of space and time. UniINR [6] then unifies rolling-shutter correction, deblurring, and interpolation within a single INR framework, showing that diverse motion-induced artifacts can be handled by one continuous representation. To further exploit temporal structure, Continuous Space–Time Video Super-Resolution [4] uses events as a motion-aware prior for dense high-frame-rate reconstruction. In parallel, I develop a self-supervised method for event-guided frame interpolation under rolling-shutter distortion [5], which removes the need for dense ground-truth labels and makes deployment more realistic. Together, these works [4,5,6,10] trace a clear line: from task-specific super-resolution [10] to unified correction of multiple artifacts [6], to continuous space–time reconstruction [4], and finally to self-supervised temporal modeling [5]. The common theme is to use events to parameterize continuous, motion-aware video representations rather than isolated frames.

Aligning with human visual perception: Practical systems must also satisfy human visual expectations in color and tone and must be compatible with real hardware. On the algorithm side, I first explore RAW domain demosaicing for event cameras [3]. I then lead the RGB-Event ISP project [2], which provides the first event-guided ISP dataset and benchmark for hybrid event sensors. These works [2,3] link event signals directly to human-perceived image quality.

To support future sensor–algorithm co-design, I develop a unified modeling, calibration, and simulation framework for hybrid event sensors [1]. This work integrates imaging noise, event generation, and calibration into one platform, allowing researchers to prototype new sensing and algorithm strategies before hardware is fixed.

Across this experience, my work moves from robust visibility, to continuous spatial–temporal modeling, and finally to human-centric and hardware-aware system design. This trajectory provides both the scientific basis and the practical tools for my future agenda on task-driven sensor–algorithm co-design and next-generation event-based perception systems.

Future Directions

My long-term goal is to build a framework that can automatically design both sensors and imaging algorithms for specific tasks, so that vision systems can be optimized directly for deployed scenarios such as VR/AR, robotics, and mobile devices. To reach this goal, I plan to advance the research agenda in three stages.

First, I will develop an automated end-to-end imaging model that can handle a wide range of degradations. Recent vision foundation models already show strong performance across multiple restoration tasks, but they still fall short of supporting real-world, device-level imaging. I aim to build a full computational pipeline that integrates physics-based modeling, event-guided priors, and continuous space–time representations to generate more reliable outputs across diverse conditions.

Second, I will incorporate sensor chip design into the imaging model. The core insight is to treat the sensor as the first layer of the imaging network. By creating differentiable sensor models—covering exposure control, threshold dynamics, readout timing, and noise behavior—I will enable joint optimization of sensor parameters and imaging algorithms. This will provide a principled way to design sensor architectures that are tailored to specific tasks, environments, and constraints.

Third, I will extend the framework to include optical design as part of the optimization loop. By embedding lens properties, optical distortions, and spectral responses into the differentiable pipeline, the system can co-optimize optics, sensor hardware, and algorithms together. This integrated view will allow imaging systems to be created from the ground up, customized for VR/AR, high-speed robotics, and other challenging domains.

Selected Works

1. Yunfan LU, Nico Messikommer, Xiaogang Xu, Liming Chen, Yuhan Chen, Nikola Zubic, Davide Scaramuzza, Hui Xiong; Hybrid Event Frame Sensors: Modeling, Calibration, and Simulation, Under Review
2. Yunfan LU, Yanlin Qian, Ziyang Rao, Junren Xiao, Liming Chen, Hui Xiong; RGB-Event ISP: The Dataset, Benchmark and Direction, ICLR 2025
3. Yunfan LU, Yijie Xu, Wenzong MA, Weiyu Guo, Hui Xiong; Event Camera Demosaicing via Swin Transformer and Pixel-focus Loss, CVPRW 2024
4. Yunfan LU, Zipeng Wang, Yusheng Wang, Hui Xiong; Continuous Space-Time Video Super-Resolution via Event Camera, IJCV 2025
5. Yunfan LU, Guoqiang Liang, YiranShen, Lin Wang; Self-supervised learning of event-guided video frame interpolation for rolling shutter frames, IEEE TVCG 2025
6. Yunfan LU, Guoqiang Liang, Yusheng Wang, Lin Wang, Hui Xiong; UniINR: Event-guided unified rolling shutter correction, deblurring, and interpolation, ECCV 2024
7. Ling Wang*, Yunfan LU*, Wenzong Ma, Huizai Yao, Pengteng Li, Hui Xiong; From Events to Clarity: The Event-Guided Diffusion Framework for Dehazing, Under Review * These authors contributed equally.
8. Yunfan LU, Xiaogang Xu, Hao LU, Yanlin Qian, Bin Yang, Junyi Li, Qianyi Cai, Weiyu Guo, Hui Xiong; SEE: See Everything Every Time - Adaptive Brightness Adjustment for Broad Light Range Images via Events, IJCV 2025
9. Kanghao Chen*, Guoqiang Liang*, Yunfan LU*, Hangyu Li, Lin Wang; Evlight++: Low-light Video Enhancement with an Event Camera: A Large-scale Real-world Dataset, Novel Method, and More, IEEE T-PAMI 2025 * These authors contributed equally.
10. Yunfan LU, Zipeng Wang, Minjie Liu, Hongjian Wang, Lin Wang; Learning Spatial-Temporal Implicit Neural Representations for Event-Guided Video Super-Resolution, CVPR 2023