Appendix: MoiréLens: Bringing Schlieren Imaging into Real-World Environments Using Moiré Patterns

Paper #229

1. CAMERA-TO-BACKGROUND GEOMETRIC CALIBRATION.

We show the examples of each step in camera-to-Background geometric calibration. Fig. 1(a) shows a high-contrast chessboard pattern covering the display is rendered, with known grid geometry in screen coordinates. Fig. 1(b) shows a rendered Moiré stimulus with a pre-warped sinusoidal grid using homography **H**.

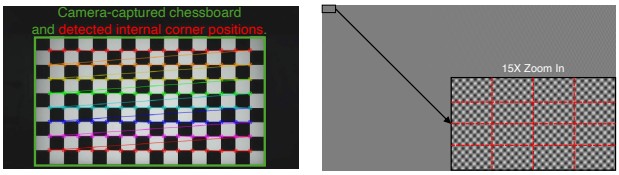

(a) Homography calculation using chessboard **(b) Example of Moiré stimulus after perspective warping**

Figure 1: Calibration and rendering procedure. (Red dashed lines in (b) are not rendered in practice but are used to illustrate perspective warping.)

2. ALGORITHM OF ESTIMATING THE NUMBER OF MOIRÉ CELLS IN THE ROI.

We design Algorithm 1 to estimates $m(f_b^i)$ by locating Moiré-cell centroids and measures their spatial density to infer the effective (possibly fractional) number of cells $\mathrm{m}(f_b^i)$ within the ROI. It enables fine-grained frequency tuning rather than discrete integer jumps in cell count.

Algorithm 1: Estimating the observed Moiré cell $m(f_b^i)$.

Input: Moiré stimulus rendered frequency f_b^i; Captured image frame I_{cam}; ROI width W.
Output: Estimated number of Moiré cells $m(f_b^i)$.

Compute the 2D DCT of the ROI in I_{cam} and identify the dominant spectral peak (f_x, f_y).
Convert I_{cam} to CIELAB (Lab*) and extract the a^* channel;
Binarize $|a^*|$ using an global threshold to obtain mask B;
Perform connected-component analysis on B to obtain component areas $\{A_i\}$ and centroids $\{(c_i^x, c_i^y)\}$;
Estimate the expected Moiré-cell area $A_{\text{ref}} \leftarrow \left(\frac{W}{f_x}\right)\left(\frac{W}{f_y}\right)$;
Identify the dominant histogram cluster of A_i centered around A_{ref} within a tolerance factor;
Select valid centroids $\mathcal{C} = \{(c_i^x, c_i^y) \mid A_i \text{ in cluster}\}$;
Compute nearest neighbor distance among $\mathcal{C}$ and estimate $m(f_b^i)$ from their mean spacing;
return $m(f_b^i)$;

3. DNN MODEL PARAMETERS.

Classification model. Table 1 summarizes the architecture and training hyperparameters of the lightweight 2D CNN used for binary classification. The model consists of three Conv–BN–ReLU–MaxPool blocks followed by a compact fully connected head with dropout regularization.

Regression model. Table 2 details the R(2+1)D-style regression model used to estimate oil temperature from short video windows. The architecture employs three spatiotemporal R(2+1)D blocks to extract plume dynamics, followed by a SiLU-activated regression head. The coord_y channel is a normalized vertical positional encoding that provides the network with an explicit notion of the gravity axis. Training uses Huber loss with AdamW, cosine learning-rate scheduling, and early stopping for robust and stable regression performance.

Table 1: Classification model archiecture and hyperparameters.

Component	Specification
Input	$1 \times 160 \times 160$ variance map
Backbone	3× Conv–BN–ReLU–MaxPool blocks:
	Block 1: $1 \rightarrow 32$
	Block 2: $32 \rightarrow 64$
	Block 3: $64 \rightarrow 128$
	Output spatial size: 20×20
Head	Flatten ($128 \times 20 \times 20 = 51{,}200$)
	FC(51,200 → 256)
	Dropout p=0.5
	FC(256 → 128)
	Dropout p=0.3
	FC(128 → 2 logits)
Loss	Cross-entropy
Optimizer	Adam (lr=10^{-3})
Batch size	32
Epochs	10

DATASET SCALE

We summarize the scale of our colleted datasets in Table 4.

4. MOIRÉLENS VS. GAS SENSOR

We compare MoiréLens with a conventional chemical gas sensor (MQ-5) using the (1.a) butane-leakage scenario to highlight MoiréLens 's advantages in responsiveness, range, and versatility. The comparison results are summarized in Table 3.

Sensing latency and range. The MQ-5 requires slow molecular diffusion to reach its minimum detectable concentration (200,PPM of LPG [4]): it responds in 2.55,s at 2,cm, 2.75,s at 5,cm, and fails entirely at 10,cm. In contrast, MoiréLens provides *near-instant* visualization of refractive-index disturbances with 3.01 ms (331.92 FPS). Using a Pixel 9 Pro camera, MoiréLens successfully visualizes the same butane flow at s = 215 cm, over 40× the sensing range and more than 125× faster detection.

Gas type generality. Conventional chemical sensors (e.g., MQ-5) are gas-specific and require separate calibration for each target analyte (e.g., LPG, methane, propane, butane). MoiréLens is a

Table 2: R(2+1)D-style regression model architecture and training hyperparameters.

Component	Specification
Input	Two channels: feature + coord_y
	Shape: $(2, T, 160, 160)$
Backbone	R(2+1)D-style network [1]
	Width: 64
	Stages: [2, 2, 2] blocks
	Temporal kernels: [9, 5, 3]
	Dilations: [1, 2, 4]
	Downsampling: [1, 1, 1]
Head	Flatten spatiotemporal features
	SiLU activation [2]
	Dropout p=0.1
	FC($d_{\text{feat}} \rightarrow 1$)
Loss	Huber loss [3]
Optimizer	AdamW (wd = 10^{-4})
Learning rate	5×10^{-4}, cosine decay + warmup
Training	30 epochs

Table 3: Comparison of Butane Leakage Detection Abilities.

Sensor	Distance	Latency	Gas Type	Spatial Information	Warmup Time
MQ- 5	✗ 10cm [a]	✗ 2.75 s[a]	◆ Specified	✗ No	✗ >48h [c]
MoiréLens	✓ 215cm	✓ 20ms	◆ Generalized	✓ Visible	✓ <1s

[a] 200 PPM; [b] at 5 cm sensing range; [c] >48h, after <1 month storage time.

Table 4: Overview of Collected Dataset Scale. Data from "Base powers" + "W/o WP" Blue marked are used for training and validation, while the remaining data are used for testing.

# of sessions		Base powers (W) 200	600	1000	1400	Unseen powers (W) 400	800	1200
Water	W/o WP	4	4	4	4	2	2	2
Oil	W/o WP	6	6	6	6	2	2	2
	WP	1	1	1	1	1	1	1
	Hood	1	1	1	1	1	1	1
	WP+Hood	1	1	1	1	1	1	1

general-purpose optical framework that responds to refractive-index variations. This cross-factor generality enables deployment in complex or dynamic environments (e.g., industrial facilities) where the nature of a disturbance may be unknown.

Spatial information. MoiréLens provides immediate spatial localization of leak sources and visualizes flow propagation in real time. The MQ-5 outputs only a scalar concentration value, offering no information about flow structure or direction.

Sensor warm-up time. MoiréLens becomes fully operational in under 1.4 s without target-size Moiré size search. The MQ-5 requires prolonged warm-up, often over 48 hours, if left unpowered for a month, to regain stable readings.

5. GASEOUS PLUME FLOW ANALYSIS

We evaluate MoiréLens's applicability to analyze gaseous plume dynamics by applying dense optical flow We feed the visualization results (100 FPS) directly into a dense optical flow algorithm, i.e.,

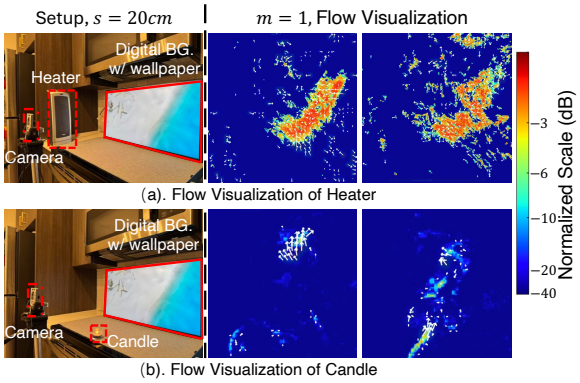

Figure 2: Gaseous Plume Flow Analysis

Gunnar Farneback's method [5]. The resulting vector fields, visualized as arrow overlays (Fig. 2), capture both the direction and magnitude of refractive-index motion.

(1.c) Heater Airflow (Fig. 2(a)): A 1500,W ceramic heater generates a strong, turbulent thermal plume. MoiréLens reveals clear refractive-index gradients and the associated rising-hot-air convection, while optical-flow vectors quantify plume directionality and turbulence intensity. It highlights MoiréLens 's ability to provide quantitative insight into HVAC-related airflow analysis.

(2.a) Candle Plume (Fig. 2(b)). MoiréLens also captures the faint, laminar plume from a candle flame, including its characteristic upward drift and oscillatory motion. Despite the small temperature gradients, the optical-flow field traces subtle convection patterns, highlighting sensitivity to weak thermal disturbances.

MoiréLens provides quantitative flow insights suitable for indoor HVAC monitoring and fine-grained plume analysis.

(2.a) Candle Plume Flow (Fig. 2(b)): MoiréLens also captures the faint, laminar plume generated by a candle flame, a challenging scenario due to the small temperature gradient and low refractive-index variation. The visualizations reveal the characteristic upward convection and oscillatory motion, while the optical-flow vectors trace the gradual ascent and diffusion of heated air. This sensitivity to subtle thermal anomalies indicates MoiréLens 's potential for detecting small leaks, drafts, or inefficient air circulation in indoor environments.

MoiréLens's capability to visualize gaseous plume flow makes it an ideal tool for indoor HVAC monitoring.

REFERENCES

[1] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A Closer Look at Spatiotemporal Convolutions for Action Recognition. In *Proceedings of the IEEE CVPR*, pages 6450–6459, 2018.

[2] Stefan Elfwing, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, November 2018.

[3] Gregory P. Meyer. An Alternative Probabilistic Interpretation of the Huber Loss. pages 5261–5269, 2021.

[4] Grove - Gas Sensor(MQ5) | Seeed Studio Wiki, January 2023.

[5] G. Farneback. Fast and accurate motion estimation using orientation tensors and parametric motion models. In *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, volume 1, pages 135–139 vol.1, September 2000. ISSN: 1051-4651.