

Minimalist Visual Inertial Odometry

Francesco Pasti, Jeremy Klotz, Nicola Bellotto, and Shree K. Nayar

Abstract—Visual-Inertial Odometry (VIO), which is critical to mobile robot navigation, uses cameras with a large number of pixels. Capturing and processing camera images requires significant resources. This work presents a minimalist approach to planar odometry, demonstrating that just four visual measurements and an IMU can provide robust motion estimation for differential-drive robots. Our key insight is that four downward-facing photodiodes that sense the world through optical Gabor masks produce signals that encode speed. Based on this, we jointly optimize the mask parameters alongside a Temporal Convolutional Network (TCN) using a physically-grounded simulator. The resulting model decodes speed from just the four measurements produced by the photodiodes. Pairing these estimates with the angular speed from an IMU yields a continuous planar trajectory. We validate our approach with a prototype sensor mounted on a differential drive robot. Across diverse indoor and outdoor terrains, our system closely tracks the reference ground truth without any real-world fine-tuning. Our work shows that minimalist sensing enables efficient and accurate planar odometry.

I. INTRODUCTION

Autonomous mobile robots often rely on Visual-Inertial Odometry (VIO) for robust navigation [1]. VIO fuses rich visual cues with measurements from an Inertial Measurement Unit (IMU) to compute the trajectory of the robot. These visual cues are extracted from camera images with a large number of pixels. Since the power consumption to sense and process images is roughly linear in the number of pixels, traditional VIO using high resolution image sensors can be unsuitable for resource-constrained platforms [2], [3].

Our work draws inspiration from minimalist vision [4], [5], which explores the lower bound of visual information needed to solve a vision task. We show that for differential drive robots, robust planar odometry can be achieved using just four ground-facing pixels, where each pixel is a photodiode with an optical mask (Fig. 1). We show that when the masks represent Gabor functions, i.e., sinusoidal waves modulated by a Gaussian envelope [6], the masks isolate a specific spatial frequency from the ground texture. As the robot moves, the sensor generates a temporal signal whose dominant frequency directly encodes the robot’s linear speed.

In any real-world setting, odometry has to contend with non-linear motion dynamics, unknown and varying ground textures, and hardware noise. Therefore, decoding the signals produced by our four Gabor pixels is a non-trivial problem. To achieve this, we develop an end-to-end differentiable

F. Pasti and N. Bellotto are with the Department of Information Engineering, University of Padua, Padua, Italy. F. Pasti, J. Klotz, and S. K. Nayar are with the Computer Science Department, Columbia University, New York, NY, USA. Corresponding author: francesco.pasti@dei.unipd.it

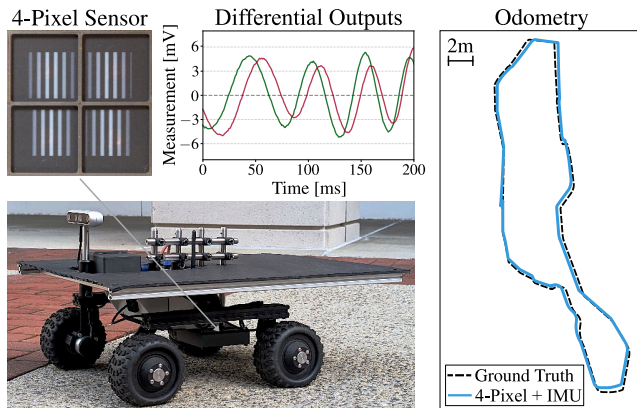


Fig. 1. The minimalist odometry system utilizes a custom sensor consisting of four masked photodiodes. The masks act as analog spatial filters that isolate specific spatial frequencies from the ground texture. As the differential drive robot moves, this optical filtering generates continuous temporal signals. We decode these signals to regress the robot’s instantaneous speed. Fusing this speed with the IMU’s gyroscope yaw rate yields the full planar trajectory.

framework that jointly optimizes the Gabor mask parameters alongside a Temporal Convolutional Network (TCN) decoder [7]. This system is trained using data produced by a physically grounded simulator that uses a diverse set of real-world textures and motion profiles.¹ The end result is a robust mapping of our pixel measurements to linear speed. This speed estimate is fused with the yaw rate from an IMU’s gyroscope to obtain the full planar trajectory of the robot.

We developed a hardware prototype of our sensor that we mounted on a differential drive robot. To validate our approach, we drove the robot for 87 minutes over 920 meters of diverse indoor and outdoor terrains. Even though our minimalist system has been optimized purely on simulated data, it achieves robust real-world odometry performance that closely tracks the reference ground truth computed via standard high-resolution VIO. Our method achieves a mean absolute trajectory error (ATE) of 0.34 meters and an average endpoint drift of 0.60%. In contrast, standard wheel encoder and IMU fusion yields a 0.74 meters ATE and a 1.55% drift. This performance is achieved despite the significant reduction in sensing resources of our 4-pixel sensor compared to high-resolution VIO, demonstrating that minimalist vision can be a viable solution for resource-constrained robot odometry.

¹4-Pixel sensor simulator: <https://github.com/pastifra/four-pixel-vio>.

II. RELATED WORKS

A. Odometry in Robotics

Robust mobile robot localization relies on the fusion of proprioceptive sensors like wheel encoders and IMUs with exteroceptive sensors like cameras and LiDARs [1]. Wheel encoders provide a reliable baseline for short-term localization but can suffer from drift due to wheel slippage over longer trajectories. Similarly, IMUs offer high-frequency and low-power proprioception, but require double integration of noisy acceleration readings which are known to cause a large drift in position estimation over time. While data-driven methods can constrain this drift for pedestrian dead-reckoning [8], [9], they fail to generalize to the smooth, non-oscillatory kinematics of wheeled robots [10].

To mitigate these errors, Visual-Inertial Odometry (VIO) is often used in robotics as a standard solution [11]. VIO fuses exteroceptive visual cues with IMU data to deal with rapid maneuvers and unreliable visual information [12]. However, conventional VIO systems require the processing of video streams containing thousands or millions of pixels per frame. This leads to computational and energy requirements that are prohibitive for resource-constrained platforms [3], especially to operate autonomously over extended period of time. In contrast, our minimalist odometry framework utilizes only 4 pixels and an IMU, significantly reducing sensing resources compared to standard VIO, while maintaining robust odometry performance.

B. Efficient Optical Flow Estimation

An alternative approach to motion estimation uses specialized optical flow sensors adapted for robotics [13], [14]. These systems utilize small pixel arrays (ranging from 18x18 to 30x30) and dedicated digital signal processors (DSPs) to compute spatial cross-correlation between images in high-framerate videos. While optimized for low latency, the need to digitize and process 2D images thousands of times per second limits their efficiency. Furthermore, these downward-facing sensors are highly sensitive to their standoff distance from the ground, which should be fixed with millimeter-scale accuracy to prevent large errors. This limits their robustness on uneven terrain, where chassis vibrations cause the standoff distance to vary continuously. Similarly, event cameras offer a relatively low-power alternative for motion estimation [15]. By outputting sparse, asynchronous events that encode per-pixel brightness changes, they drastically reduce data rates while maintaining high temporal resolution. Nevertheless, extracting motion from these event streams requires dense spatial sampling of the scene.

Rather than using small pixel arrays to compute optical flow digitally, our approach draws inspiration from a model of biological motion perception, which shows that spatio-temporal frequency filtering of light measurements can directly lead to motion estimation [16], [17]. We implement this principle in the optical domain. As the robot moves, our sensor optically convolves the scene texture with Gabor filters to produce a low-dimensional temporal signal. This

temporal signal encodes motion information in the frequency domain.

C. Minimalist Vision

The goal of minimalist vision is to directly sense the smallest number of task-relevant visual measurements without sensing a full image. For instance, the minimalist camera in [4] employs handcrafted optical masks placed over photodetectors to process scene information directly in the optical domain. Freeform pixels [5] extend this concept by modeling masked photodetectors as linear projections of the scene. Such physical masks can be included as initial layers of a neural network. An end-to-end optimization of this network for any given task leads to not just a trained inference network but also the design of the optical masks to be used with the network.

Our approach builds on this framework but with a critical distinction. We show that since motion estimation is inherently based on spatio-temporal frequency analysis [17], Gabor filters are ideally suited to isolate the motion-related frequency components. Rather than designing freeform pixels for speed estimation, which is not a well-constrained problem, we constrain the masks to be Gabor functions with learnable parameters. These parameters are left unknown so that they can be optimized for the complex real-world conditions (texture, lighting, mechanical) faced by the sensor. The outputs of our Gabor pixels are fed into a TCN decoder that estimates speed. The Gabor parameters and the TCN are jointly trained to obtain a robust speed estimator.

III. THE MINIMALIST ODOMETRY SYSTEM

A. Theoretical Intuition

Consider the simplified scenario illustrated in Fig. 2. A sensor, consisting of a single detector and an optical mask, is pointed down at a surface and moves over it along its longitudinal x -axis with a constant speed v .

While the surface texture is two-dimensional, let the optical mask have translational symmetry along the lateral y -axis. This reduces their optical interaction to a single dimension, allowing us to model the system strictly as one-dimensional along the direction of motion.

Let the texture of the surface along the direction of motion be the function $I(x)$ and the transmittance function of the mask, when projected onto the surface, be $M(x)$. If the sensor is positioned at an arbitrary displacement x , the detector integrates the light passing through the mask at that location. Therefore, the sensor output as a function of position is simply the spatial cross-correlation of the texture and the mask: $S_x(x) = (I \star M)(x)$.

If the sensor moves at a constant speed v , the detector's displacement relative to the surface is $x = vt$. Therefore, the continuous temporal output $s(t)$ of the sensor is simply a scaled version of the spatial cross-correlation $S_x(x)$, such that $s(t) = S_x(vt)$. This directly links the temporal and spatial domains via $x = vt$.

To understand how this temporal signal $s(t)$ encodes speed, we analyze it in the frequency domain. Let f denote




Fig. 2. **Theoretical intuition.** A detector integrates the light from a surface texture $I(x)$ passing through an optical mask $M(x)$. As the sensor moves at a constant speed v , it performs a continuous spatial cross-correlation in the optical domain. This physical process maps the scene's spatial frequencies into a temporal signal $s(t)$.

temporal frequency and \mathcal{F} denote the Fourier transform. Then, the transform of the sensor's temporal output is $\hat{S}(f) = \mathcal{F}\{S_x(vt)\}$. By applying the time-scaling property of the Fourier transform, we get:

$$\hat{S}(f) = \frac{1}{|v|} \hat{S}_x\left(\frac{f}{v}\right). \quad (1)$$

Now let ξ denote frequency in the spatial domain. Using the cross-correlation theorem, the Fourier transform of $S_x(x)$ can be expressed as $\hat{S}_x(\xi) = \hat{I}(\xi)\hat{M}^*(\xi)$, where $*$ denotes the complex conjugate, and \hat{I} and \hat{M} are the transforms of the texture and the mask, respectively. By substituting $\hat{S}_x(\xi)$ in Eq. 1, we get:

$$\hat{S}(f) = \frac{1}{|v|} \hat{I}\left(\frac{f}{v}\right) \hat{M}^*\left(\frac{f}{v}\right). \quad (2)$$

The above equation represents the core working principle of our sensor. The temporal frequencies (f) produced by the moving sensor are the spatial frequencies (ξ) of the texture, filtered by the mask, and scaled by the sensor speed v , such that $f = \xi v$.

Our challenge is to estimate the sensor speed v robustly, irrespective of the texture $I(x)$ of the surface. Consider an infinite cosine mask $M(x) = \cos(2\pi\xi_0x)$, where ξ_0 is the known frequency of the cosine. In the frequency domain, its spectrum $\hat{M}(\xi)$ consists of two symmetric impulses at $\pm\xi_0$. Although the spectrum of the texture $\hat{I}(\xi)$ is unknown, we can assume it to be broadband. Then, our ideal mask has the effect of passing through to the sensor a single spatial frequency ξ_0 from the texture. In Eq. 2, this makes the spectrum of the sensor output $\hat{S}(f)$ collapse to symmetric impulses at $\pm f_0$, where $f_0 = \xi_0 v$. Therefore, if we can detect f_0 from $s(t)$, the speed can be estimated as $v = f_0/\xi_0$.

While an infinite cosine mask provides a theoretical baseline, any physical mask must have a finite aperture. In the frequency domain, a cosine mask with an aperture is no longer a pair of impulses at $\pm\xi_0$, but rather a pair of broader functions, which makes the estimation of speed harder.

To mitigate this broadening of the impulses, we make the mask a Gabor function [6]:

$$G_{\cos}(x) = \alpha \exp\left(-\frac{x^2}{2\sigma^2}\right) \cos(2\pi\xi_0x), \quad (3)$$

where α is an amplitude scaling factor and σ^2 is the variance of a Gaussian envelope.² By modulating the mask with a Gaussian envelope, we smoothly taper the cosine to zero, avoiding its abrupt spatial truncation. So long as the Gaussian envelope is broad, the spectrum of the Gabor mask has narrow peaks at $\pm\xi_0$. Similar to the ideal cosine, the Gabor mask restricts the spectrum of $\hat{S}(f)$ in Eq. 2 to symmetric peaks at $\pm f_0$, where $f_0 = \xi_0 v$. The resulting temporal sensor output $s(t)$ can be approximated as an amplitude-modulated cosine whose frequency remains $f_0 = \xi_0 v$. This allows us to recover the speed v by finding the fundamental frequency f_0 of the sensor's output $s(t)$.

This brings us to a limitation of the above method. If the sensor reverses its direction of motion, the speed v becomes $-v$. This theoretically yields a negative frequency $f_0 = \xi_0(-v)$ in the sensor output $s(t)$. However, since $s(t)$ is a real-valued signal, it still contains symmetric peaks at $\pm f_0$. Therefore, we can only determine the magnitude of the speed (the fundamental frequency f_0), but not the direction of motion (the sign of v). We will now show that this directional ambiguity in the sensor speed can be resolved using more than one sensor.

B. Directional Ambiguity and Positive Masks

To resolve the above directional ambiguity, we introduce a second masked sensor, co-located and in quadrature with the first sensor. Let its mask be a sine Gabor function, $G_{\sin}(x)$, which is simply $G_{\cos}(x)$ in Eq. 3 phase shifted by 90° , which corresponds to a translation in space. In the frequency domain, the sine Gabor also has peaks at $\pm\xi_0$. However, the spatial shift between the masks of the two sensors naturally produces a 90° temporal phase shift between the outputs of the two sensors.

Let the outputs of the two sensors be $s_{\cos}(t)$ and $s_{\sin}(t)$. Since we know that both outputs are narrowband with peaks at $f_0 = \xi_0 v$, we can approximate them as:

$$\begin{aligned} s_{\cos}(t) &\approx A(t) \cos(2\pi\xi_0 vt + \theta(t)), \\ s_{\sin}(t) &\approx A(t) \sin(2\pi\xi_0 vt + \theta(t)). \end{aligned} \quad (4)$$

Here, $A(t)$ and $\theta(t)$ represent the time-varying amplitude and spatial phase of the texture $I(x)$ at ξ_0 . Crucially, the fundamental frequency of the two signal is the same $f_0 = \xi_0 v$.

Intuitively, the sign (direction) of the speed v is encoded in the relative phase between $s_{\cos}(t)$ and $s_{\sin}(t)$. Note that this relative phase is either $+90^\circ$ or -90° , one corresponding to forward motion and the other to backward motion of the sensors, while, as before, f_0 reveals the speed magnitude.

The above results generalize to dynamic motion as well. Even when the speed varies with time, the frequency and relative phase represent the speed and direction at that time.

In practice, there is an additional constraint we need to consider when building our sensor. While our Gabor filters have positive and negative values, the optical transmittance

²This is analogous to using window functions in digital signal processing for reducing spectral leakage in the Fourier transform of finite signals.

of a mask can only be positive. We address this problem by decomposing each Gabor mask function into a pair of strictly non-negative masks:

$$M^+ = \max(G(x), 0), \quad M^- = \max(-G(x), 0). \quad (5)$$

The difference between the outputs of two sensors with the above masks is equivalent to the output of a single sensor with a Gabor mask. Therefore, in order to obtain our two quadrature sensor outputs, we use four sensors with the masks M_{cos}^+ , M_{cos}^- , M_{sin}^+ , and M_{sin}^- . This is the basic design of our minimalist sensor.

C. Height Dependency

The above derivation of minimalist sensing for speed estimation is valid when the four masked detectors are co-located. In practice, however, they must be spatially offset with respect to each other, as shown in Fig. 3. Given the nominal height h_{nom} of the masks from the ground plane, the masks are positioned with respect to their detectors such that their projections onto the ground plane are perfectly aligned. In other words, the projected spatial frequencies of the four masks are the same (ξ_0), and the temporal phase shift between the quadrature signals is exactly 90° . This makes the system equivalent to the co-located one in Sec. III-B for the height h_{nom} .

In any real setting, however, due to the unevenness of the ground and the robot’s vibrations, the height of the sensor will vary ($h = h_{nom} + \Delta h$). This has the effect of changing the scales of the projections of the masks onto the ground plane, and hence their effective spatial frequency ($\xi_0 + \Delta\xi_0$). In addition, as the four detectors are offset with respect to each other, any deviation from the nominal height h_{nom} introduces parallax, and the four detectors no longer observe the same patch on the ground plane. While this corrupts the sensor outputs, it also provides a latent cue that can be exploited. Since the spatial displacement between the projected masks changes as a function of height, the resulting phase difference between s_{cos} and s_{sin} encodes information regarding deviations from the nominal height.

Since the offset sensors observe slightly different ground patches when $h \neq h_{nom}$, the height dependent phase shifts are confounded by the random undulations of the ground plane. The effects of these phase shifts on the speed estimate are hard to model analytically. However, we are able to exploit them using our learning framework for speed estimation (Sec. IV). By training model on simulated sequences with realistic textures and height variations, the model estimates speed robustly even in the presence of such perturbations.

D. Planar Kinematics and IMU Integration

Our goal is to recover the planar odometry of a differential-drive robot. We have shown that our minimalist sensor encodes the signed speed v along its direction of motion. We are therefore able to measure the robot’s speed v_x in the “forward” direction. We mount our sensor such that its center lies on the longitudinal axis of the robot and its




Fig. 3. **Hardware implementation of our 4-pixel speed sensor.** (a) Since light transmission is strictly positive, the Gabor filters (G_{cos} and G_{sin}) are split into their positive and negative components to obtain the masks M^+ and M^- . The resulting four masks are printed on film. (b) They are placed in front of a 2×2 grid of photodiodes. The distance between adjacent photodiodes is $d = 1.9$ cm.

mask stripes are perpendicular to the direction of motion (see Fig. 3(a)). Note that there are a few conditions for which the speed estimation from the detector outputs can be challenging. For instance, when the robot undergoes simultaneous rotation and translation, the speed of the ground plane is non-uniform within the sensor’s field of view. Also, if the robot experiences lateral slippage, the detector outputs are impacted by the variation in ground texture due to the slip. To deal with these effects, our model for forward speed estimation is trained using a simulator that includes these motion dynamics.

Note that full planar odometry for a differential-drive robot also requires the rotational component, i.e., the yaw rate w_z . To this end, we integrate an IMU whose gyroscope independently measures the yaw rate, ω_z . By fusing our optically estimated forward speed v_x with the IMU’s yaw rate ω_z , we obtain the full planar odometry.

E. Sensor Prototype

We designed our minimalist sensor based on the custom hardware architecture developed for freeform pixels [5]. The Gabor filters function, G_{cos} and G_{sin} , are split into their positive and negative components to obtain the four physical masks M_{cos}^+ , M_{cos}^- , M_{sin}^+ , and M_{sin}^- . These masks, each 16×16 mm² in size, are printed on transparent film as shown in Fig. 3(a). The masks are placed in front of four Hamamatsu S9119-01 photodiodes arranged in a 2×2 grid, as shown in Fig. 3(b). The distance between adjacent photodiodes is 1.9 cm, and the distance between the photodiodes and the masks is 11.4 mm. This results in a 70° field of view for each detector. We have positioned the masks to ensure that the four detectors observe the same area on the ground plane at a nominal height $h_{nom} = 6$ cm.

In our experiments, we used an external data acquisition (DAQ) system to digitize the four analog signals produced by the above sensor. Our sensor produces four analog detector outputs consuming just 2.5 mW. In comparison, the image sensor in a conventional camera consumes hundreds of milliwatts [2]. This translates to a reduction of two orders of magnitude in the power consumed by the sensor.

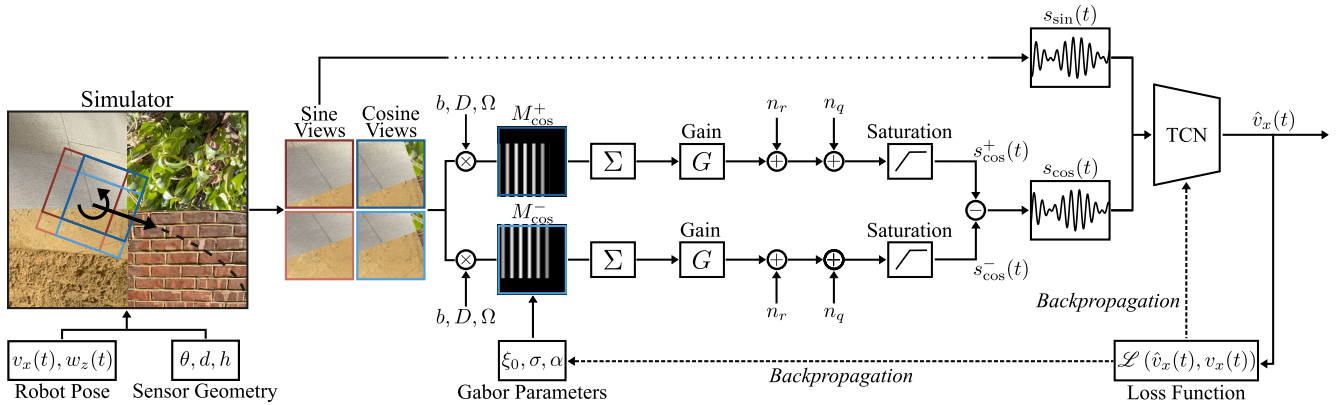


Fig. 4. **Overview of our learning framework.** The simulator generates the four detector views $I_k(x, y, t)$ of our sensor during kinematic motion. The effects of the finite area of the detector (b), the directional response of the detector (D) and the foreshortening effect (Ω) are applied to the detector views before they are modulated by the masks $M_k(x, y, t)$. These modulated views are integrated (Σ) to get four signals to which a gain (G), read noise (n_r), quantization noise (n_q) and saturation are applied. From these four detector outputs, $s_{\cos}(t)$ and $s_{\sin}(t)$ are computed and fed into a TCN decoder. The TCN takes a temporal window of the input signals to predict the linear speed $\hat{v}_x(t)$ of the sensor. Since the pipeline is fully differentiable, the loss \mathcal{L} is backpropagated to optimize not just the TCN parameters but also the Gabor parameters (ξ_0, σ, α).

IV. LEARNING FRAMEWORK FOR SPEED ESTIMATION

Finding a robust mapping from our minimalist sensor signals to speed is a non-trivial problem. The theoretical model described in Sec. III-A must cope with various factors including uneven ground, complex robot motions, and noise in the sensor outputs. Furthermore, although we have established our masks will be Gabor functions, we have not yet determined what their parameters (ξ_0, σ, α) should be. This is a challenging problem as it depends on the properties of the wide range of textures the sensors will encounter.

We address the above problems by taking a learning-based approach, where the Gabor parameters and the parameters of a network for speed estimation are jointly learned using simulated sensor data. To this end, we built a physically based simulator (Fig. 4) that generates sensor outputs for a wide range of robot motions and ground plane textures.

A. Minimalist Odometry Sensor Simulation

Fig. 4 shows the complete pipeline of our learning framework. We generate realistic training data for our speed estimator by using a variety of high-resolution surface textures and dynamic kinematic trajectories. To ensure our speed estimator can cope with diverse (indoor and outdoor) environments, our simulator uses real-world textures from the Matador dataset [18], comprising ~ 7200 high-quality images across 57 material categories. For realistic movements, we simulate our minimalist sensor across over 80 Km and 12 h from the TartanGround [19] dataset, encompassing linear speeds up to 5 m/s and angular speeds up to 1 rad/s. To avoid aliasing effects, we upsample the 100 Hz pose trajectories in [19] to 1000 Hz.

At each timestep t , the instantaneous pose of the robot defines the spatial coordinates of the four detectors, indexed by $k \in \{\cos^+, \cos^-, \sin^+, \sin^-\}$. The geometric configuration of the sensor (Sec. III-E) is parametrized by its inter-detector distance d and detector field of view θ . We randomly perturb the nominal sensor height ($h = h_{nom} + \Delta h$) to simulate the scale error discussed in Sec. III-C.

For each detector, we use resampling with bilinear interpolation to construct a 128×128 pixel image from the high-resolution texture in its field of view. The resulting images $I_k(x, y, t)$, corresponding to the four detectors, are converted to grayscale and their brightness values are normalized to $[0, 1]$. We also render the aligned Gabor masks $M_k(x, y)$ as 128×128 images, parametrized by the spatial frequency ξ_0 , sigma σ , and amplitude α of the Gabor function. Since our physical masks are printed on transparency film, we constrain their transmittance to lie within $[0, 1]$.

Using the image patches $I_k(x, y, t)$, we simulate the optical pipeline of our sensor. Following the physical sensor model described in [5], we first blur I_k with a kernel $b(x, y)$, which represents the finite size of the detector. Then, we multiply the resulting image with the directional response of the detector $D(x, y)$ and $\Omega(x, y)$, which accounts for the foreshortening effect. The resulting image is multiplied with the Gabor mask $M_k(x, y)$ and then integrated to obtain the final detector output. This entire process of simulating a detector output can be mathematically written as:

$$s_k(t) = \sum_{\mathbf{u}} \left[I_k(\mathbf{u}, t) * b(\mathbf{u}) \right] D(\mathbf{u}) \Omega(\mathbf{u}) M_k(\mathbf{u}), \quad (6)$$

where s_k is the detector output and $\mathbf{u} = (x, y)$.

Finally, we need to convert our sensor output $s_k(t)$ to a voltage. As in [5], we scale $s_k(t)$ by a calibrated hardware gain $G = 1.22 \times 10^{-4}$ and inject both Gaussian read noise ($\sigma = 175 \mu\text{V}$) and uniform quantization noise. The result is then clipped at 3.2 V to model the finite dynamic range of the detectors. To construct the final inputs for our speed decoder, we compute $s_{\cos}^+(t) - s_{\cos}^-(t)$ and $s_{\sin}^+(t) - s_{\sin}^-(t)$, obtaining $s_{\cos}(t)$ and $s_{\sin}(t)$, respectively.

B. Temporal Speed Decoder and End-to-End Optimization

To estimate the sensor's linear speed, we process a 1-second observation window with 1000 temporal samples of our simulated signals $s_{\cos}(t)$ and $s_{\sin}(t)$ (i.e. 2×1000 input matrix), using a TCN [7] with 184K parameters.. The TCN extracts multi-scale temporal features via stacked

TABLE I
INFLUENCE OF MASK DESIGN ON SPEED ACCURACY

Masks	RMSE (m/s) ↓	MAE (m/s) ↓
Freeform	0.147	0.101
Fixed Gabor	0.076	0.052
Learned Gabor	0.054	0.034

dilated convolutions, and passes the features to an attention pooling layer. Since the amplitudes ($A(t)$ in Eq. 4) of the signals $s_{cos}(t)$ and $s_{sin}(t)$ are dependent on the spectrum of the ground plane texture, the signal to noise ratio (SNR) of the signals can vary within our 1-second observation window. The attention pooling layer of the TCN mitigates this problem by allowing the network to suppress noisy portions of the window.

The pooled features feed into a residual fully-connected head, which outputs a linear speed prediction $\hat{v}_x(t)$ and a log-variance $\log(\hat{\sigma}^2)$ that quantifies prediction uncertainty.

We train the network using a gaussian negative log-likelihood loss, which optimizes both the speed prediction and its uncertainty. To make training stable in the presence of high accelerations, we define the ground-truth target speed $v_x(t)$ as the average speed of the sensor over the last 0.1 s of the observation window.

Since our simulator (Sec. IV-A) is fully differentiable, the loss gradients can be propagated from the TCN all the way back to the simulated sensor. This enables us to jointly optimize both the parameters of the TCN as well as the parameters of our Gabor masks (ξ_0, σ , and α). In other words, we are able to jointly design our sensor (hardware) and speed estimator (software) to achieve high speed accuracy.

V. EXPERIMENTAL EVALUATION

A. Learning the Masks

We conducted simulations to evaluate the impact of optimizing the masks (Gabor parameters) on speed estimation accuracy. We refer to the masks that were designed by jointly optimizing the Gabor parameters (ξ_0, σ, α) with the TCN parameters as *Learned Gabor*. We compared them with *Fixed Gabor* masks, where the Gabor parameters were empirically selected to be $\xi_0 = 6$, $\sigma = 1$, and $\alpha = 1$. We also compared them with *Freeform Pixels* [5], where the masks can assume an arbitrary form and are jointly optimized with the TCN.

In these experiments, we used an inter-detector distance of $d = 1.9$ cm and detector field of view of $\theta = 70^\circ$ (see Sec. IV-A). We set the nominal height to $h_{nom} = 20$ cm to prevent aliasing at the high speeds (± 5 m/s) present in the TartanGround dataset. We split the Matador textures and TartanGround trajectories into train (70%), validation (10%), and test (20%) sets. The three different mask designs described above and the TCN were trained for 100 epochs using Adam optimizer with a learning rate of 10^{-4} and a batch size of 32.

Table I summarizes the results in terms of root mean square error (RMSE) and mean absolute error (MAE). *Freeform Pixels* perform poorly; without the Gabor prior, the masks converge to blurred patterns, forcing the TCN to rely

TABLE II
SPEED RMSE (m/s) AS HEIGHT DEVIATES FROM $h_{nom} = 20$ CM

Train Height Range	Test Height Range			
	h_{nom}	$\pm 10\%$	$\pm 25\%$	$\pm 50\%$
h_{nom}	0.054	0.070	0.266	1.542
$h_{nom} \pm 10\%$	0.050	0.054	0.063	1.154
$h_{nom} \pm 25\%$	0.048	0.050	0.056	1.259
$h_{nom} \pm 50\%$	0.071	0.071	0.075	0.088

on the temporal lag between the displaced detectors' signals to regress speed. Using Gabor masks instead significantly improved performance. In particular, the joint optimization of *Learned Gabor* yielded a 29% reduction of RMSE and 35% reduction of MAE compared to the *Fixed Gabor* masks.

B. Robustness to Height Variations

As discussed in Sec. III-C, deviations from the nominal height alter the effective spatial frequency of the Gabor ($\xi_0 + \Delta \xi_0$), thereby impacting the frequency-to-speed mapping. However, height variations also change the phase difference between the detector outputs s_{cos} and s_{sin} , which proves to be a useful cue.

We evaluated via simulations the TCN's ability to exploit this phase difference to achieve robustness to height variations ($h_{nom} + \Delta h$). We used the same sensor configuration parameters as in Sec. V-A, but we trained and tested the *Learned Gabor* masks and the TCN with the random height variations uniformly sampled from within the ranges of $\pm 10\%$, $\pm 25\%$ and $\pm 50\%$, relative to $h_{nom} = 20$ cm.

As shown in Table II, the performances of the model trained at a fixed height h_{nom} degrade significantly when exposed to height variations. In contrast, we see that training with the uniform height randomization makes speed estimation resilient to them. A $\pm 25\%$ uniform height randomization during training yields robust performance for the same range of height variation during testing. In fact, it slightly improves the accuracy at the nominal height, yielding a speed RMSE of 0.048 m/s. This shows that exposing the model to these variations during training results in a speed estimator that implicitly learns the relationship between height, phase, frequency and speed.

However, when trained and tested on substantial height variations ($\pm 50\%$), extreme parallax between the views seen by the four detectors makes the phase cue unreliable and the overall performance drops. In our real-world experiments, we found that training with the $\pm 25\%$ uniform height randomization results in a system that is resilient to uneven terrains and vibrations encountered in both indoor and outdoor settings.

C. Robot with Minimalist Odometry Sensor

We evaluated our minimalist odometry sensor prototype (Fig. 3) using the LeoRover³ differential-drive robot shown in Fig. 5. The sensor is mounted at a height h_{nom} of 6 cm, and centered on the robot's longitudinal axis so

³<https://fictionlab.pl/leo-rover/>




Fig. 5. **Differential drive robot with minimalist odometry sensor.** Our 4-pixel sensor is mounted downward-facing at a nominal height of $h_{nom} = 6$ cm on a differential-drive robot. A shield is used to suppress strong specular reflections from the ground plane, and an LED lamp provides illumination in poorly lit indoor environments. An Intel RealSense D455 camera is used to capture the ground-truth trajectories. The IMU of this camera provides the angular speed which is fused with our sensor’s linear speed estimates to compute planar odometry.

as to measure the linear speed v_x of the robot. A shield is used to avoid strong specular reflections from light sources in the environment towards the sensor. To cope with low-light environments that can arise indoors, we used a low-powered LED source that is suspended from the chassis of the robot. This source is not used outdoors or in well-lit indoor environments.

We jointly trained the Gabor masks and the TCN with a nominal sensor height h_{nom} of 6 cm and a uniform height randomization of $\pm 25\%$. For this training, we scaled the speed in the TartanGround trajectories so that they match the maximum speed of the robot (0.4 m/s). This resulted in a trained model with a speed RMSE of 0.004 m/s and MAE of 0.003 m/s. The trained Gabor masks were printed on transparency film and installed on the prototype without further modifications.

An external DAQ samples the four detectors signals at 41.6 kHz. Then, we process the resulting signals with a 60 Hz notch filter to remove AC lighting fluctuations and a 450 Hz low-pass filter to denoise the signals before down-sampling them to 1 kHz, which is the temporal resolution the TCN is designed for.

D. Indoors and Outdoors Experiments

We have evaluated the performance of our minimalist odometry sensor by teleoperating the robot in diverse indoor and outdoor environments, totaling 920 meters of travel distance and 87 minutes of travel time over 11 indoor trajectories (618 m in 61 min) and 5 outdoor trajectories (306 m in 26 min). During teleoperation, we only controlled the direction of the robot while its speed was randomly varied.

We processed the 1 kHz sensor signals offline so that we could evaluate various update rates of the TCN speed predictions. Specifically, we applied a 1-second sliding window to the collected signals, using strides of 1, 10, and 33 milliseconds to update the speed estimates at 1 kHz, 100 Hz, and 30 Hz, respectively. We leveraged the network’s prediction uncertainty to drop low-confidence speed estimates, followed by a simple median filter to eliminate transient outliers.

TABLE III
MINIMALIST ODOMETRY PERFORMANCE

Method	Indoor		Outdoor	
	ATE (m)↓	Drift (%)↓	ATE (m)↓	Drift (%)↓
Encoders	9.75 ± 8.45	25.60 ± 30.00	14.09 ± 7.74	30.25 ± 26.13
Encoders + IMU	0.75 ± 0.47	1.62 ± 1.66	0.74 ± 0.19	1.37 ± 1.20
4-Pixel + IMU (1 kHz)	0.28 ± 0.11	0.60 ± 0.22	0.42 ± 0.08	0.62 ± 0.25
4-Pixel + IMU (100 Hz)	0.29 ± 0.12	0.65 ± 0.24	0.43 ± 0.07	0.72 ± 0.12
4-Pixel + IMU (30 Hz)	0.28 ± 0.11	0.65 ± 0.21	0.44 ± 0.08	0.85 ± 0.22

Finally, we fused this filtered linear speed with calibrated gyroscope readings from the IMU of an Intel Realsense D455 camera to compute planar odometry.

To obtain the reference odometry we processed visual, depth, and inertial data from the RealSense camera using the RTAB-Map library [20]. While this is not absolute ground truth, it provides a strong upper bound VIO reference. We compared our sensor with wheel encoders, which represent the standard proprioceptive baseline for differential-drive robots. In particular, we considered two specific configurations. The first one, *Encoders*, computes planar odometry by using the mean and differential measurements from the robot’s four wheel encoders to derive linear and angular speed. The second, *Encoders + IMU*, retains the encoder-based linear speed but, for angular speed, it integrates the same calibrated gyroscope readings used in our approach. This second configuration serves as a direct comparison with our sensor, as they both are being relied on for their linear speed estimates.

To quantify the performance of our system, we consider the mean absolute trajectory error (ATE) and the percentage of endpoint drift [1]. The ATE represents the global geometric consistency of the path by computing the root-mean-square difference between the position estimates and the ground-truth trajectories. On the other hand, the endpoint drift isolates the accumulated relative error over the total traveled distance.

The quantitative results, summarized in Table III, show that our minimalist odometry approach outperforms the baselines for both the indoor and outdoor trajectories. As expected, the performance of our approach degrades slightly outdoors, as uneven terrains and drastic lighting variations corrupt the sensor signals. Reducing the update rate of the TCN speed estimates (which corresponds to lower computational load) from 1 kHz to 100 Hz and 30 Hz, causes only a negligible drop in odometry performance.

Fig. 6 shows the trajectories computed for some of our indoor and outdoor experiments. Note that our system (4-Pixel + IMU) closely tracks the reference VIO ground truth for a wide variety of environments with diverse ground textures. Despite the fact that the sensor masks and the TCN were optimized entirely using simulated data, our sensor is able to generalize to arbitrary real-world environments.

VI. CONCLUSIONS

We presented a minimalist approach to planar odometry using a sensor comprising just four masked detectors. The-

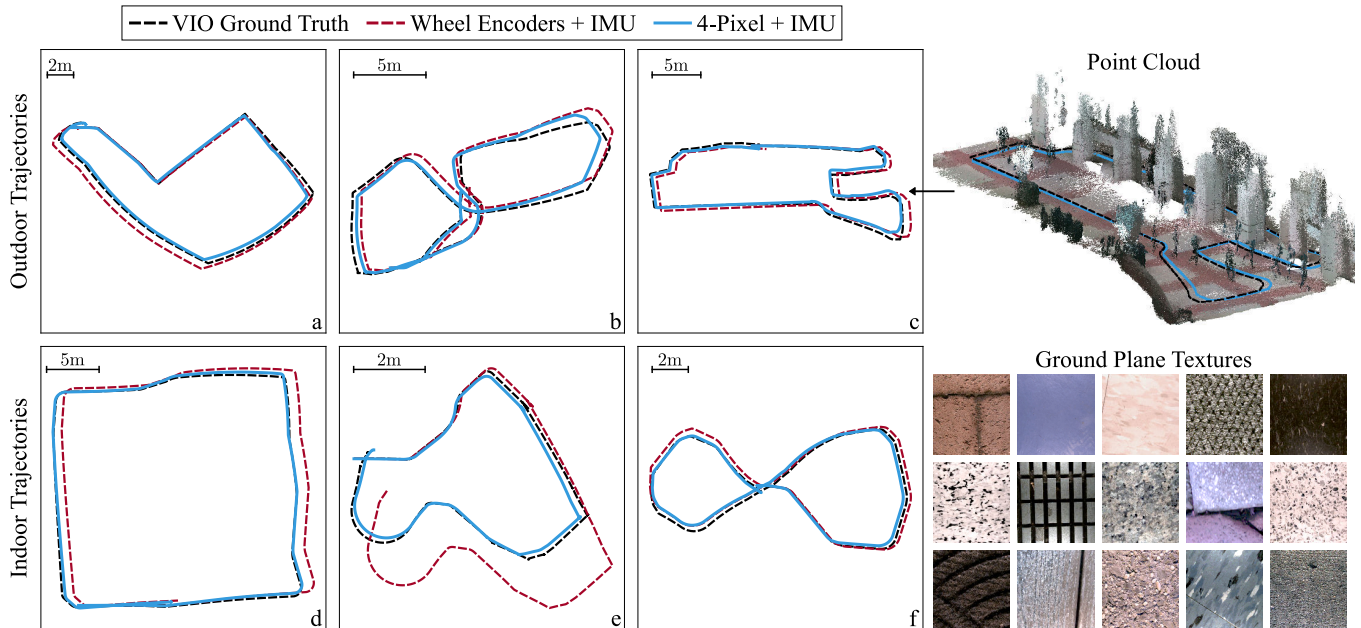


Fig. 6. **Indoor and outdoor experiments.** The trajectories (solid blue) computed using our minimalist odometry sensor paired with an IMU closely follows the reference VIO ground truth (dashed black) across a variety of outdoor (a-c) and indoor (d-f) environments. It outperforms the wheel encoder baseline (dashed red) in all cases. The point cloud corresponding to path 'c' is shown on the top-right. Some of the arbitrary real-world textures (tile, brick, marble, concrete, carpet, pebbles) encountered during the experiments are shown on the bottom-right.

oretical analysis showed Gabor masks were ideal for speed estimation. Through a custom simulator, we jointly optimized the parameters of the Gabor masks and a TCN network. We installed the optimized masks on a prototype sensor and used it with an IMU on a differential drive robot to compute planar odometry. Our minimalist approach achieves robust performance across diverse indoor and outdoor environments.

This work can be extended in several directions. First, we plan to develop custom embedded hardware for onboard implementation of an optimized version of our TCN. This will enable ultra-low-powered planar odometry. Second, we would like to explore the use of masks for odometry with more degrees of freedom. This will make our approach applicable to other platforms, such as drones. Finally, minimalist sensing can be exploited to solve other robotics tasks such as terrain classification and collision avoidance.

VII. ACKNOWLEDGMENTS

This work was supported in part by the Office of Naval Research award N00014-23-1-2096 and in part by the National Science Foundation (NSF) and Center for Smart Streetscapes (CS3) under NSF Cooperative Agreement No. EEC-2133516. Jeremy Klotz was supported by a National Defense Science and Engineering Graduate (NDSEG) Fellowship. The authors thank Matthew Beveridge for technical feedback and Dhruv Yalamanchi for helping with experiments.

REFERENCES

- [1] L. Carlone, A. Kim, T. Barfoot, D. Cremers, and F. Dellaert, *SLAM Handbook: From Localization and Mapping to Spatial Intelligence*. Cambridge University Press, 2025.
- [2] R. LiKamWa, B. Priyantha, M. Philipose, L. Zhong, and P. Bahl, "Energy characterization and optimization of image sensing toward continuous mobile vision," in *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*, 2013, pp. 69–82.
- [3] S. M. Neuman, B. Plancher, B. P. Duisterhof, S. Krishnan, C. Banbury, M. Mazumder, S. Prakash, J. Jabbour, A. Faust, G. C. de Croon *et al.*, "Tiny robot learning: Challenges and directions for machine learning in resource-constrained robots," in *2022 IEEE 4th international conference on artificial intelligence circuits and systems (AICAS)*. IEEE, 2022, pp. 296–299.
- [4] P. Pooj, M. Grossberg, P. N. Belhumeur, and S. K. Nayar, "The minimalist camera," in *BMVC*, 2018, p. 141.
- [5] J. Klotz and S. K. Nayar, "Minimalist vision with freeform pixels," in *European Conference on Computer Vision*. Springer, 2024, pp. 329–346.
- [6] D. Gabor, "Theory of communication. part I: The analysis of information," *Journal of the Institution of Electrical Engineers-part III: radio and communication engineering*, vol. 93, no. 26, pp. 429–441, 1946.
- [7] C. Lea, M. D. Flynn, R. Vidal, A. Reiter, and G. D. Hager, "Temporal convolutional networks for action segmentation and detection," in *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 156–165.
- [8] H. Yan, Q. Shan, and Y. Furukawa, "Ridi: Robust imu double integration," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 621–636.
- [9] W. Liu, D. Caruso, E. Ilg, J. Dong, A. I. Mourikis, K. Daniilidis, V. Kumar, and J. Engel, "Tlio: Tight learned inertial odometry," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5653–5660, 2020.
- [10] S. A. Mohamed, M.-H. Haghbayan, T. Westerlund, J. Heikkonen, H. Tenhunen, and J. Plosila, "A survey on odometry for autonomous navigation systems," *IEEE access*, vol. 7, pp. 97466–97486, 2019.
- [11] G. Huang, "Visual-inertial navigation: A concise review," in *2019 international conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 9572–9582.
- [12] S. Zhao, S. Zhou, Y. Zhang, J. Zhang, C. Wang, W. Wang, and S. Scherer, "Resilient odometry via hierarchical adaptation," *Science Robotics*, vol. 10, no. 109, p. eadv1818, 2025.
- [13] H. Chao, Y. Gu, and M. Napolitano, "A survey of optical flow techniques for robotics navigation applications," *Journal of Intelligent & Robotic Systems*, vol. 73, no. 1, pp. 361–372, 2014.
- [14] L. Bergantin, C. Coquet, J. Dumon, A. Negre, T. Raharijaona, N. Marchand, and F. Ruffier, "Indoor and outdoor in-flight odometry

- based solely on optic flows with oscillatory trajectories,” *International Journal of Micro Air Vehicles*, vol. 15, 2023.
- [15] E. Mueggler, G. Gallego, H. Rebecq, and D. Scaramuzza, “Continuous-time visual-inertial odometry for event cameras,” *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1425–1440, 2018.
- [16] M. F. Land and R. D. Fernald, “The evolution of eyes,” *Annual review of neuroscience*, vol. 15, no. 1, pp. 1–29, 1992.
- [17] E. H. Adelson and J. R. Bergen, “Spatiotemporal energy models for the perception of motion,” *Journal of the optical society of america A*, vol. 2, no. 2, pp. 284–299, 1985.
- [18] M. Beveridge and S. K. Nayar, “Hierarchical material recognition from local appearance,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 8165–8176.
- [19] M. Patel, F. Yang, Y. Qiu, C. Cadena, S. Scherer, M. Hutter, and W. Wang, “Tartanground: A large-scale dataset for ground robot perception and navigation,” in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025, pp. 20524–20531.
- [20] M. Labbé and F. Michaud, “Rtab-map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation,” *Journal of field robotics*, vol. 36, no. 2, pp. 416–446, 2019.