

Article

Not peer-reviewed version

Exploring Optical Flow Methods for Automated Fall Detection System

[Simeon Karpuzov](#) , [Stiliyan Kalitzin](#) , [Stefan Tabakov](#) , [Dobromir Tsolyov](#) , [Georgi Petkov](#) *

Posted Date: 26 February 2026

doi: 10.20944/preprints202602.1391.v1

Keywords: optical flow; fall detection; video analysis



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Exploring Optical Flow Methods for Automated Fall Detection System

Simeon Karpuzov ¹, Stiliyan Kalitzin ^{2,3}, Stefan Tabakov ¹, Dobromir Tsolyov ¹
and George Petkov ^{1,*}

¹ GATE Institute, Sofia University, 1164 Sofia, Bulgaria

² Stichting Epilepsie Instellingen Nederland (SEIN), Achterweg 5, 2103 SW Heemstede, The Netherlands

³ Image Sciences Institute, University Medical Center Utrecht, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands

* Correspondence: georgi.petkov@gate-ai.eu; Tel.: +359-888134817

Abstract

Falls pose severe risks to vulnerable populations, particularly the elderly and individuals with adverse neurological conditions, necessitating reliable and non-obstructive detection systems. While previous multi-modal approaches utilizing video and audio have demonstrated strong performance, they face significant limitations regarding sensitivity to environmental noise. This paper presents a robust, video-only fall detection framework that eliminates reliance on acoustic data to enhance universality. We conduct a comprehensive comparative analysis of five Optical Flow (OF) algorithms—Horn-Schunck, Lucas-Kanade (LK), LK-Derivative of Gaussian, Farneback, and the spectral method SOFIA—to determine the range of applicability of each technique for capturing fall dynamics. Beyond detection accuracy, we investigate computational efficiency of each configuration. This optimized, privacy-centric pipeline offers a scalable solution for continuous monitoring in home and clinical settings, addressing the critical need for immediate intervention following high-impact falls.

Keywords: optical flow; fall detection; video analysis

1. Introduction

Motivation

Falls represent a significant health risk for vulnerable populations, particularly the elderly and individuals with epilepsy. For these groups, a fall is often not merely a mechanical failure of balance but a potential indicator of an acute medical event, such as a seizure, or a precursor to a long-term injury. The consequences of such incidents are worsened when the individual is unable to get help due to injury or loss of consciousness. Consequently, the development of automated fall detection systems has become a critical area of research in biomedical engineering and ambient assisted living.

To address the limitations of wearable devices which are prone to an array of issues, remote sensing approaches using computer vision have emerged as a non-obstructive alternative. In our previous work [1], we proposed an automated analysis algorithm for the remote detection of high-impact falls. That system utilized a multi-modal approach, combining Optical Flow (OF) features from video with sound amplitude analysis. While the inclusion of audio features increased detector specificity, reliance on sound presents significant challenges in real-world deployment such as sensitivity to environmental noise and privacy concerns.

This article presents a substantial evolution of that pipeline, moving towards a robust, video-only fall detection system. By removing the dependency on audio, we aim to increase the universality of our solution. To compensate for the loss of audio data, we focus on maximizing the information extraction from the video stream. The specific choice of Optical Flow algorithm and machine learning

classifier plays a decisive role in the system's ability to distinguish falls from everyday activities. Therefore, this study systematically evaluates multiple Optical Flow techniques, ranging from classical global methods to modern robust estimators to identify the optimal configuration for real-time fall detection. We aim to investigate which OF method is most suitable for which use case. We also propose an extracted feature set to achieve a better detection performance.

While numerous Optical Flow methods exist in isolation, their performance varies drastically when deployed in live clinical environments. We proceed from the understanding that different methods behave differently depending on scene dynamics—such as the sudden lighting changes or occlusions common in hospital corridors. Therefore, we are not merely ranking algorithms; we are examining the OF methods to identify which domains allow for computationally cheaper estimators and which strictly require robust, complex solvers. This comparative analysis serves as the foundation for an adaptive logic into our existing hospital-deployed system, ensuring that the most efficient method is utilized without compromising patient safety.

Related Works

Camera-based fall detection has emerged as a prominent research direction within computer vision and digital healthcare. Numerous methodologies have been explored, and we refer the reader to existing surveys for a broader overview [2], [3]. In what follows, we focus specifically on approaches that incorporate Optical Flow (OF) analysis as a core component of their processing pipeline.

In [4], the authors propose a system aimed at improving classification accuracy in dynamic lighting conditions and optimizing pre-processing performance. The methodology employs an Enhanced Dynamic Optical Flow technique that leverages rank pooling to summarize temporal video data into a single image representation

The method shown in [5] uses a combination of standard Optical Flow, background subtraction [6] and Kalman filtering [7] to determine whether the recorded movement is a fall. The algorithm proceeds by organizing the isolated motion patterns into a feature vector. This structure aggregates the silhouette angle, the width-to-height ratio of the bounding rectangle, and the ratio derivative—a parameter specifically included to quantify the rate of shape deformation over time. Subsequently, this data is used as an input for a k-Nearest Neighbor (k-NN) algorithm [8]. The classifier processes these features to determine whether the observed motion dynamics correspond to a fall event.

Another contemporary method is shown in [9]. Optical Flow is calculated and then the extracted information is utilized in a distinct manner. Descriptors of points of interest are derived and subsequently serve as an input to a Convolutional Neural Network (CNN) [10], which executes the final classification. To increase robustness, the framework incorporates a two-stage rule-based motion detection module. This component identifies large, abrupt movements and applies predefined rules once the optical flow magnitude surpasses specific thresholds.

In [11], principal motion parameters are passed into a lightweight CNN that classifies falls. It works in real-time and avoids computationally expensive optical flow calculation. The videos are divided into portions of five seconds and the system analyzes each subsection for falls.

Original Contributions

In this work, we extend our original method by including tests and analysis on the following:

- **Comparative Evaluation of Optical Flow Methods:** We move beyond the exclusive use of Horn-Schunck (HS) [12] by implementing and testing four additional Optical Flow algorithms: Lucas-Kanade (LK) [13], Lucas-Kanade Derivative of Gaussian (LKDoG) [14], Farneback (FB) [15], and SOFIA (SF) [16]. This comparison allows us to determine which method provides the most reliable motion estimation for fall dynamics under varying conditions.
- **Expanded Feature Extraction:** To attempt compensation for the removal of the sound channel, we introduce new motion features derived from the Optical Flow field, moving beyond simple vertical velocity and acceleration to capture more subtle kinematic characteristics of falls.

- FPS and Computational Efficiency Analysis: Real-time processing is a prerequisite for emergency alerting systems. Thus, we provide a detailed analysis of the Frames Per Second (FPS) performance for the different OF methods. This ensures that the recommended solution is not only accurate but also viable for deployment on standard hardware.

Article Structure

The remainder of this article is organized in the following way: Section 2 details the methodology, outlining the modified detection pipeline that excludes sound analysis. We describe the working principles of the five Optical Flow methods selected for evaluation and the expanded feature set extracted from the flow fields. Section 3 presents the experimental results, comparing various performance metrics of each Optical Flow method. Finally, Section 4 discusses the implications of these results for real-world deployment, analyzes the trade-offs between accuracy and computational cost, and suggests directions for future research.

2. Materials and Methods

In our original work we combine audio and video features into a single feature vector that is then used as input for an SVM classifier. A general scheme of the system is presented in Figure 1 a). Here, we remove the audio branch entirely. We examine different OF calculation methods and add a new motion feature to our feature vector. The updated system can be viewed in Figure 1 b):

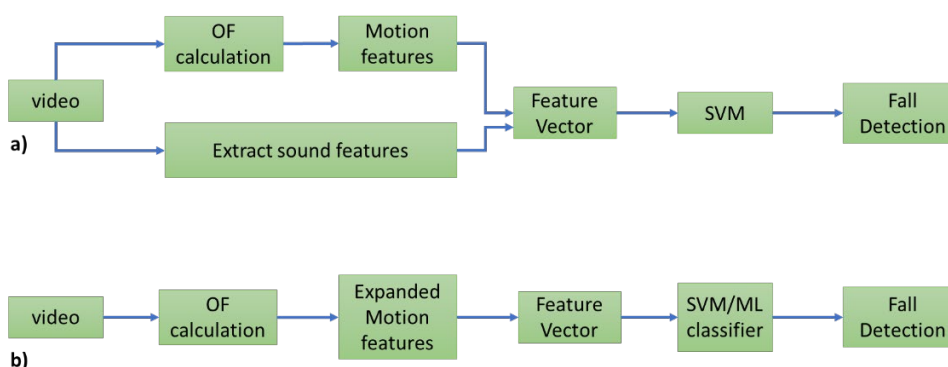


Figure 1. a): Outline of our original system. Both sound and motion features are combined into a feature vector, that is then used as input of an SVM classifier. b) Updated system pipeline in the current work. We no longer use audio for fall detection.

We begin this chapter by introducing different Optical Flow calculation techniques.

Optical Flow Methods

The Horn-Schunck [12] method is a global regularization method that addresses optical flow estimation as a variational minimization problem. The method minimizes a global energy functional:

$$J(h) = \iint \left[(I_x u + I_y v + I_t)^2 + \alpha^2 (|\nabla u|^2 + |\nabla v|^2) \right] dx dy, \quad (1)$$

where the first term enforces brightness constancy and the second term imposes smoothness regularization on the velocity field, with the parameter α controlling the trade-off between these competing objectives. The solution is obtained through Euler-Lagrange equations. The resulting global regularization produces dense flow fields with velocity estimates at every pixel, propagating motion information from high-gradient regions into textureless areas through the smoothness constraint. Horn-Schunck produces dense flow fields with velocity estimates at every pixel. Even in textureless regions where the brightness constraint provides no useful information, the global smoothness constraint propagates flow from surrounding areas, filling in the velocity field throughout the image. This makes the method particularly valuable for applications requiring

complete motion coverage across the entire image domain. This is the OF method we have used in our original work.

We implemented and tested the Farneback [15] OF algorithm to assess its suitability for fall-related motion analysis. Farneback provides a dense flow field with good sensitivity to broad motion patterns, estimating motion for every pixel, allowing it to capture these patterns across the entire image. This dense representation is particularly sensitive to large, abrupt movements typical of fall events. From the resulting flow fields, we extracted frame-wise vertical velocity profiles and computed the standard fall-related features used in prior work, (maximum vertical velocity, acceleration, deceleration).

Additionally, we also integrated the Lucas–Kanade [13] method as a lightweight alternative to Farneback’s dense estimation. Lucas–Kanade tracks small sets of coherent key points, making it faster, sparser, and less memory intensive. Using the vertical component of the tracked object, we derived the same fall-related features, enabling a direct comparison between the two methods within the unified evaluation framework.

A different variation of the standard LK method is the so-called Lucas-Kanade Derivative of Gaussians (LKDoG) [17]. Compared with the original Lukas-Kanade (LK) method, LKDoG differs primarily in how it computes derivatives. For spatial derivatives, LKDoG first applies Gaussian smoothing to the whole input image. Then, it computes the DoG filters, followed by additional Gaussian smoothing applied on the gradient fields. The method uses separate standard deviation parameters for image smoothing and gradient smoothing to control the characteristics of these filters. For temporal derivative, it applies a DoG filter for temporal filtering not on two consecutive frames, but on a window of multiple frames. Using DoG filters provides better noise reduction and smoother gradient estimates compared to simple finite difference approximations used in basic Lucas-Kanade implementations. The LKDoG method also incorporates temporal filtering across multiple frames rather than just two consecutive frames, making it more robust to insignificant temporal variations.

The Spectral Optical Flow Iterative Algorithm (SOFIA) [16] proposes a novel approach for optical flow reconstruction that integrates local structural information from multiple spectral components, such as color channels, to address the inverse problem. To tackle singularities, the method constructs a structural tensor using aggregated spatial-spectral gradient data, effectively enhancing the rank of the system. Furthermore, the algorithm introduces a spatial smoothing functionality applied directly to the structural tensor within local neighborhoods, which avoids the gradient cancellation effects often observed when smoothing images directly. Subsequently, an iterative multi-scale scheme is utilized where the vector field reconstructed at coarser scales deforms the source image to refine the input for finer scale calculations. Validations demonstrate that this approach significantly improves reconstruction accuracy and functional association compared to standard methods like Horn-Schunck, particularly when processing complex global translations.

In summary, a brief description of the methods is available in Table 1:

Table 1. Summary of different optical flow methods.

Method	Strengths	Weaknesses	Use cases
Lucas–Kanade (LK)	<ul style="list-style-type: none"> - Simple, efficient - Good for small, coherent displacements - Works well in sparse, feature-based tracking (corners, textures) 	<ul style="list-style-type: none"> - Assumes brightness constancy & small motion - Fails with large displacements - Sensitive to noise and texture less regions 	<ul style="list-style-type: none"> - Tracking a few good features (e.g., Shi–Tomasi corners) - Real-time applications with limited compute
Lucas–Kanade Derivative of Gaussians (LKDoG)	<ul style="list-style-type: none"> - More accurate gradient estimation than standard LK 	<ul style="list-style-type: none"> - Higher computational cost than basic LK - Still assumes small motion 	<ul style="list-style-type: none"> - Scientific/medical imaging

	<ul style="list-style-type: none"> - Improves robustness to noise - Better handling of textured patterns 	<ul style="list-style-type: none"> - Cases where accurate gradient estimation matters more than speed 	
Farneback (Polynomial Expansion)	<ul style="list-style-type: none"> - Dense optical flow - Handles larger displacements better than LK - Provides smooth, detailed flow fields 	<ul style="list-style-type: none"> - More computationally expensive - May over smooth motion boundaries - Can be less accurate with sudden motion 	<ul style="list-style-type: none"> - Dense motion analysis - Video stabilization, background/foreground segmentation - Gesture recognition
Horn-Schunck (HS)	<ul style="list-style-type: none"> - Produces dense, smooth flow - Global regularization makes it robust to noise - Well-established classical method 	<ul style="list-style-type: none"> - Over-smooths at motion boundaries - Sensitive to parameter tuning (smoothness λ) - Computationally heavier than LK 	<ul style="list-style-type: none"> - Global motion estimation - Medical imaging (e.g. tissue/organ motion) - Scenarios where smoothness is desired
SOFIA	<ul style="list-style-type: none"> - Robust to noise via spectral representation - Can handle complex, non-rigid deformations - Iterative refinement improves accuracy 	<ul style="list-style-type: none"> - Higher computational load - Less common as it is new - Requires careful parameter tuning 	<ul style="list-style-type: none"> - Biomedical signals (EEG, brain activity mapping) - Applications with periodic/oscillatory motion - Research where spectral methods outperform spatial-only methods

Data

In the current work we have used two publicly available datasets of people falling down on video. The first dataset is LE2I [18]. It contains 190 videos with framerate of 25 frames per second. Various camera positions are available, most often the camera is placed in one of the upper corners of the room. Frame resolution is 320 by 240 pixels. Only one moving person is present in the videos. Only one person is falling down. This is a very widespread dataset in the current topic.

The second dataset is UG [19]. It consists of 40 videos of five seconds each. Twenty are with falls, twenty are with other movement activities. There is a single moving person in front a homogenous background. Camera resolution is 1920 by 1080 pixels, with framerate of 30 frames per second. It is indoor and the brightness in the scene is constant. It is a useful out-of-distribution set to test detection and classification performance. It differs from LE2I with respect to the background of the scenes.

Both datasets are open access and available to everyone for testing, which was our primary reason for selection – availability. Both sets are labeled.

Expanded Motion Features

In our previous work we extract three fall descriptors from the mean vertical velocity curve – the maximum acceleration A_{max} , the maximum deceleration D_{max} and the maximum vertical velocity V_{max} . We choose the maximum in value (A, V, D) triplet, such that these values are sequential in time. The mean velocity curve may have several peaks for a given time interval and imposing a restriction related to the time of occurrence of the above features allows us to select the proper triplet that we have shown is related to the fall event.

Here we introduce another feature that can help detect a falling person. It is related to the width of the velocity curve where a (A, V, D) triplet is found. The width is related to the standard fall time

of the average adult – somewhere between 0.5-3.0 seconds. Because of this, we add the the full width at half maximum σ to our feature vector. All motion extracted features are presented in Figure 2.




Figure 2. All four motion features that we extract from the video - $(A_{max}, V_{max}, D_{max}, \sigma)$.

The width of the curve is helpful because it allows us to distinguish between events with different rates of movement. When a person sits down, σ will be lower. When a person falls we expect to see a wider mean velocity curve. After we have calculated σ , we incorporate it to the feature vector $(A_{max}, V_{max}, D_{max}, \sigma)$.

Evaluation Parameters

We split events into two categories: falls and all other types of movement. For evaluation, we use a set of statistical values, derived from the confusion matrix which is a two by two table that is useful when comparing ground-truth labels with model predictions. It has the following elements: True positives (TP)—number of times we have correctly predicted an event as a fall; False positives (FP)—number of times we have incorrectly predicted an event as fall; True negative (TN)—number of times we have correctly predicted an event as other movement; False positives (FN)—number of times we have incorrectly predicted an event as other movement. Using these values, we introduce the following metrics:

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \quad , \quad Sensitivity = \frac{TP}{TP+FN} \quad , \quad Specificity = \frac{TN}{TN+FP} \quad , \quad (2)$$

$$Precision = \frac{TP}{TP+FP} \quad , \quad F1 = 2 * \frac{Precision \cdot Sensitivity}{Precision+Sensitivity}$$

To comprehensively evaluate the system's performance, Accuracy is utilized to quantify the overall proportion of correctly predicted events. Complementary metrics include Sensitivity, which measures the model's capacity to correctly classify positive instances, and Specificity, which assesses the prediction reliability for negative cases. Precision is defined as the ratio of true positive classifications to the total predicted positives. Furthermore, the F1-score is computed to provide a balanced metric that simultaneously accounts for false positives and false negatives. Graphically, the discriminatory capability of the model is illustrated using Receiver Operating Characteristic (ROC) curves, which plot Sensitivity against the False Positive Rate (1-Specificity) across varying thresholds. The Area Under the Curve (AUC) is subsequently derived to quantify effectiveness; a value approaching 1 indicates superior performance, whereas a value near 0.5 suggests results equivalent to random guessing. Collectively, all these dimensionless metrics provide us with a robust assessment of the fall detection model.

3. Results

A Lenovo® ThinkPad (Lenovo, Hong Kong) with an Intel® Core i5 CPU (Intel, Santa Clara, CA, USA) and 16 Gb of RAM is used to process videos from our sets. Algorithm realization is carried out in MATLAB® R2023b environment.

Our first contribution is investigating how different optical flow methods affect the overall detection algorithm. We have tested the following five: Lucas-Kanade, Horn-Schunck, Lucas-

Kanade-Derivative-of-Gaussians, SOFIA and Farneback. Results for the LE2I dataset, presented as ROC curves and performance metrics are available in Figure 3:

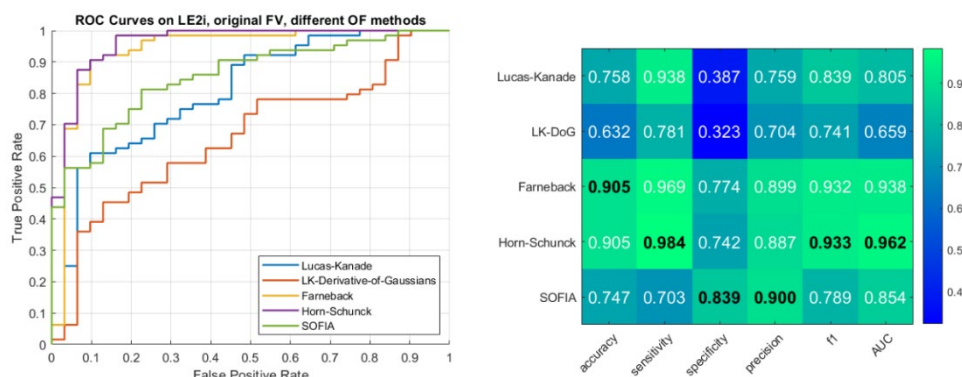


Figure 3. Left: ROCs for each of the tested optical flow methods. Right: Various performance metrics of our technique, based on different OF methods for velocity field calculation. The maximum value for each metric is given in bold and black. Stronger shades of green indicate a better result. All values are in the range from 0 to 1. Higher value indicates a better result. LE2I dataset with the original feature vector.

Here we can see that the three best performing methods in terms of maximum area below the curves are Horn-Schunck, Farneback and SOFIA. Detailed performance metrics - accuracy, sensitivity, specificity, precision, f1-score and AUC are also presented. These values confirm the findings from the ROC curves for the three best performing OF methods.

In addition, we present ROC curves and performance metrics for the second dataset - UG. Results are available in Figure 4:

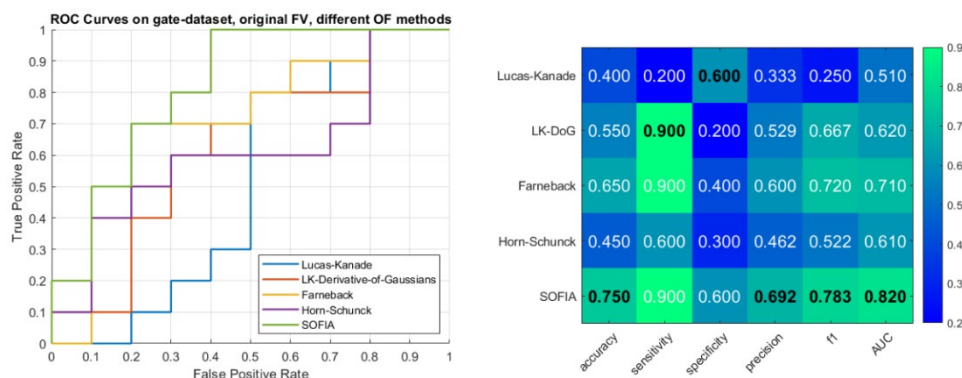


Figure 4. Left: ROCs for each of the tested optical flow methods. Right: Various performance metrics of our technique, based on different OF methods for velocity field calculation. The maximum value for each metric is given in bold and black. Stronger shades of green indicate a better result. All values are in the range from 0 to 1. Higher value indicates a better result. UG dataset with the original feature vector.

From these results it would seem that the best performing OF method is SOFIA. Adding a new feature to the feature vector – namely the width of the velocity curve where V_y is at a maximum, gives us the following results:

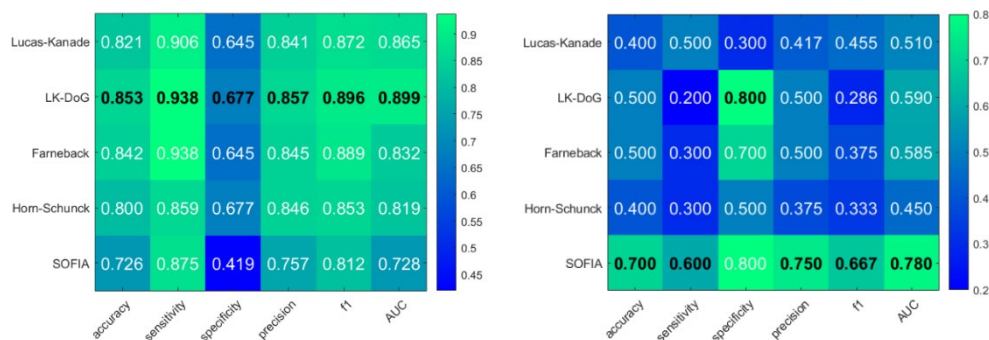


Figure 5. *Left:* Performance metrics for the new feature vector on LE2I data. *Right:* Performance metrics for the new feature vector on UG data. The maximum value for each metric is given in bold and black. Stronger shades of green indicate a better result. All values are in the range from 0 to 1. Higher value indicates a better result.

We can see that for LE2I LKDoG excels, while for the UG dataset SOFIA is still the best performing. We can also compare how the addition of the new feature affected performance. In Figure 6 we show the difference in performance metrics between the original and updated feature vector.

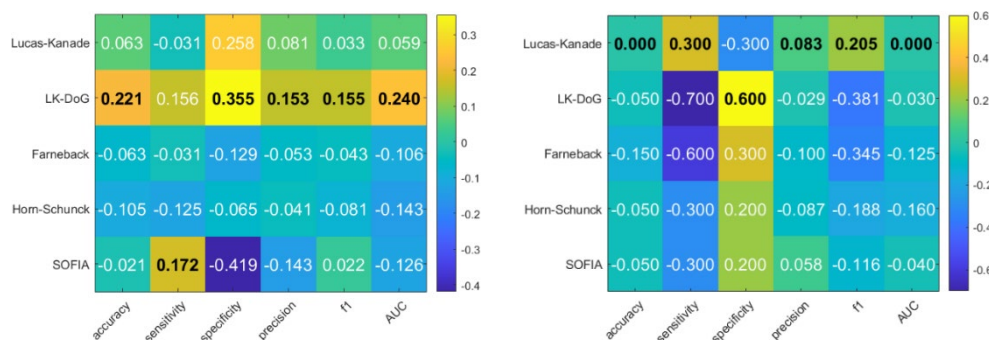


Figure 6. *Left:* Difference in performance between new feature vector and original feature vector for LE2I dataset. *Right:* Difference in performance between new feature vector and original feature vector for first occurrence for the UG dataset. The maximum value for each metric is given in bold and black. Stronger shades of yellow/red indicate a better result. All values are in the range from 0 to 1. Higher value indicates a better result.

Finally, we examine OF methods runtimes for a video with FPS = 25 and N = 150 total frames. Video frame size is 320 by 240 pixels². Results are presented in Figure 7:

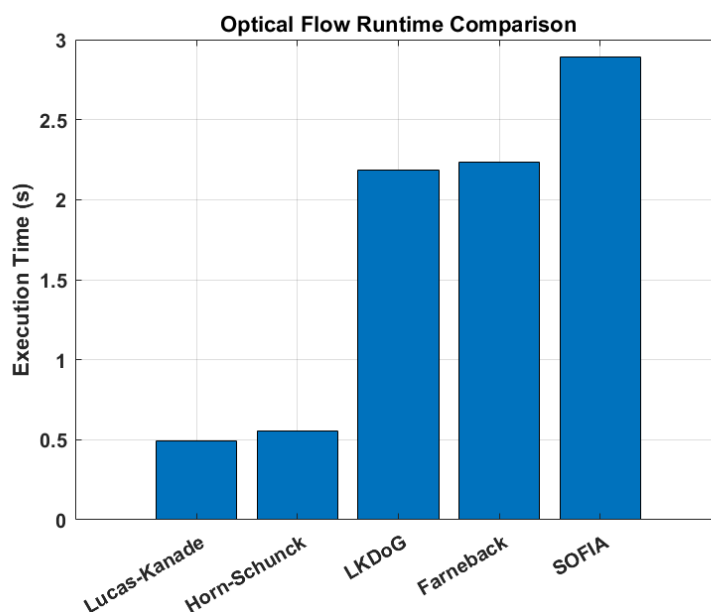


Figure 7. Computation times for each optical flow method. Lower value is better.

These results are expected and are related to the nature of the used algorithms – sparse versus dense. In the case of LKDoG, Farneback and SOFIA, the extra assumptions, conditions and information used by the methods, lead to a higher time-performance cost. Their reported calculation speed does not prevent their use in real-time calculations.

4. Discussion

In this study, we presented a comprehensive evaluation of a vision-based fall detection system that integrates variable Optical Flow (OF) algorithms with a standard ML classifier (namely the SVM). We tested the OF methods on two datasets and can summarize our findings in the following way:

On the LE2I dataset, Horn-Schunck and Farneback methods demonstrate the strongest performance with the original feature vector. Horn-Schunck achieves the highest overall metrics, with a sensitivity of 98.4%, an F1-score of 93.3% and an AUC of 0.962. Farneback shows the second-best performance across all categories. The introduction of the new feature vector significantly shifts the results for LE2I. The Lucas-Kanade Derivative of Gaussians method sees a dramatic improvement, becoming the top performer with increases in accuracy (+22.1%) and specificity (+35.5%). This suggests that LKDoG benefits most from the additional temporal shape information in simpler environments.

In the UG dataset, SOFIA emerges as the most robust method using the original feature vector, achieving the highest accuracy (75.0%), sensitivity (90.0%), and F1-score (78.3%). While Horn-Schunck struggles here with a low accuracy of 45.0%, SOFIA maintains high reliability, outperforming standard methods like Lucas-Kanade and Farneback. Unlike the LE2I dataset, adding the new feature vector to the UG dataset generally degrades performance for most methods. SOFIA's sensitivity drops by 30%, and Horn-Schunck sees decreases across multiple metrics. Only Lucas-Kanade shows a mild improvement in sensitivity (+30%) but suffers in specificity (-30%), indicating that the new feature may introduce noise or overfitting in this dataset.

The performance disparity on the UG dataset—where Horn-Schunck (HS) failed (45.0% accuracy) while SOFIA excelled (75.0%)—can be attributed to the handling of homogeneous backgrounds. In textureless regions, standard methods like HS face an ill-posed inverse problem (the Aperture Problem [20]) and often rely on global smoothness constraints. This artificial effect distorts the silhouette, degrading the geometric features used for classification. SOFIA mitigates this

singularity by employing Tikhonov regularization on the structural tensor in under-determined areas. This effectively stabilizes the flow in homogeneous regions, preserving the object's true boundaries and ensuring the feature vector remains discriminative even in sparse-texture environments.

One future direction of work is to determine whether the SOFIA algorithm can benefit from calculation on a graphics processing unit. We will explore full GPU acceleration (utilizing CUDA kernels) to offload the structural tensor computation and Gaussian smoothing steps.

Our findings show that Horn-Schunck excels in high-texture zones, while SOFIA dominates in homogeneous backgrounds. This suggests the potential for a hybrid framework. That would include developing a pre-processing module that evaluates how homogenous the background is. Based on this texture information, the system could dynamically switch between Horn-Schunck (for textured areas) and SOFIA (for textureless areas), which will allow us to maximize both accuracy and computational efficiency across diverse domestic environments.

While this study focuses on classical and spectral Optical Flow methods suitable for CPU-bound edge devices, we acknowledge the emergence of deep learning-based estimators like RAFT (Recurrent All-Pairs Field Transforms). While RAFT offers superior accuracy in detecting fine micro-movements, its heavy GPU requirement currently makes it cost-prohibitive for large-scale, privacy-preserving edge deployment in every hospital room. Future iterations of our system may explore lightweight distilled versions of such networks (e.g., FlowSeek or LiteFallNet) to balance this trade-off.

Finally, we aim to explore how changing the ML method for classification can increase performance.

Author Contributions: Conceptualization, S.K. (Stiliyan Kalitzin); methodology, S.K. (Stiliyan Kalitzin), Sm.K. (Simeon Karpuzov) and G.P.; software, Sm.K. (Simeon Karpuzov), G.P. and S.K. (Stiliyan Kalitzin); data, Sm.K. (Simeon Karpuzov); S.T. (Stefan Tabakov) and D.T. (Dobromir Tsolyov); verification, S.T. (Stefan Tabakov) and D.T. (Dobromir Tsolyov); writing—original draft preparation, Sm.K. (Simeon Karpuzov). All authors have read and agreed to the published version of the manuscript.

Funding: This research is part of the GATE project funded by the Horizon 2020 WIDESPREAD-2018–2020 TEAMING Phase 2 programme under grant agreement no. 857155, the programme “Research, Innovation and Digitalization for Smart Transformation” 2021–2027 (PRIDST) under grant agreement no. BG16RFPR002-1.014-0010-C01. Stiliyan Kalitzin is partially funded by “Anna Teding van Berkhout Stichting”, Program 35401, Remote Detection of Motor Paroxysms (REDEMP).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets, used in the study, are publicly available at: <https://doi.org/10.5281/zenodo.17777971>
<https://www.kaggle.com/datasets/tuyenldvn/falldataset-imvia>.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

OF	Optical Flow
LK	Lucas-Kanade
LKDoG	Lucas-Kanade Derivative of Gaussians
SF	SOFIA
FB	Farneback
HS	Horn-Schunck

References

1. E. E. Geertsema, G. H. Visser, M. A. Viergever, and S. N. Kalitzin, "Automated remote fall detection using impact features from video and audio," *J. Biomech.*, vol. 88, pp. 25–32, May 2019, doi: 10.1016/j.jbiomech.2019.03.007.
2. J. Gutiérrez, V. Rodríguez, and S. Martín, "Comprehensive Review of Vision-Based Fall Detection Systems," *Sensors* 2021, Vol. 21, Page 947, vol. 21, no. 3, p. 947, Feb. 2021, doi: 10.3390/S21030947.
3. A. Singh, S. U. Rehman, S. Yongchareon, and P. H. J. Chong, "Sensor Technologies for Fall Detection Systems: A Review," *IEEE Sens. J.*, vol. 20, no. 13, pp. 6889–6919, Jul. 2020, doi: 10.1109/JSEN.2020.2976554.
4. R. Tarik et al., "Deep Learning for Vision-Based Fall Detection System: Enhanced Optical Dynamic Flow", doi: 10.1111/coin.12428.
5. K. De Miguel, A. Brunete, M. Hernando, and E. Gamba, "Home camera-based fall detection system for the elderly," *Sensors (Switzerland)*, vol. 17, no. 12, Dec. 2017, doi: 10.3390/s17122864.
6. M. Piccardi, "Background subtraction techniques: A review," *Conf. Proc. IEEE Int. Conf. Syst. Man Cybern.*, vol. 4, pp. 3099–3104, 2004, doi: 10.1109/ICSMC.2004.1400815.
7. R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, Mar. 1960, doi: 10.1115/1.3662552.
8. T. M. Cover and P. E. Hart, "Nearest Neighbor Pattern Classification," *IEEE Trans. Inf. Theory*, vol. 13, no. 1, pp. 21–27, 1967, doi: 10.1109/TIT.1967.1053964.
9. Y. Z. Hsieh and Y. L. Jeng, "Development of Home Intelligent Fall Detection IoT System Based on Feedback Optical Flow Convolutional Neural Network," *IEEE Access*, vol. 6, pp. 6048–6057, Nov. 2017, doi: 10.1109/ACCESS.2017.2771389.
10. G. E. Hinton, "Computation by neural networks," *Nat. Neurosci.*, vol. 3, no. 11s, p. 1170, 2000, doi: 10.1038/81442;KWRD=BIOMEDICINE.
11. S. Karpuzov, S. Kalitzin, O. Georgieva, A. Trifonov, T. Stoyanov, and G. Petkov, "Automated Remote Detection of Falls Using Direct Reconstruction of Optical Flow Principal Motion Parameters," *Sensors* 2025, Vol. 25, Page 5678, vol. 25, no. 18, p. 5678, Sep. 2025, doi: 10.3390/S25185678.
12. B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artif. Intell.*, vol. 17, no. 1–3, pp. 185–203, Aug. 1981, doi: 10.1016/0004-3702(81)90024-2.
13. B. D. Lucas and T. Kanade, "An Iterative Image Registration Technique with an Application to Stereo Vision (IJCAI) An Iterative Image Registration Technique with an Application to Stereo Vision." [Online]. Available: <https://www.researchgate.net/publication/215458777>
14. R. K. Bhogal and V. Devendran, "Motion estimating optical flow for action recognition:(farneback, horn schunck, lucas kanade and lucas-kanade derivative of Gaussian)," in *2023 International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT)*, 2023, pp. 675–682.
15. G. Farneback, "Two-frame motion estimation based on polynomial expansion," in *Scandinavian conference on Image analysis*, 2003, pp. 363–370.
16. S. Kalitzin, E. Geertsema, and G. Petkov, "Scale-Iterative Optical Flow Reconstruction from Multi-Channel Image Sequences," *Frontiers in Artificial Intelligence and Applications*, vol. 310, pp. 302–314, 2018, doi: 10.3233/978-1-61499-929-4-302.
17. J. L. Barron, D. J. Fleet, S. S. Beauchemin, and T. A. Burkitt, "Performance of Optical Flow Techniques," *Springer Nature*, vol. 12, pp. 43–77, 1994, doi: <https://doi.org/10.1007/BF01420984>.
18. I. Charfi, J. Miteran, J. Dubois, M. Atri, and R. Tourki, "Optimized spatio-temporal descriptors for real-time fall detection: comparison of support vector machine and Adaboost-based classification," *J. Electron. Imaging*, vol. 22, no. 4, p. 041106, Jul. 2013, doi: 10.1117/1.jei.22.4.041106.
19. S. Karpuzov and G. Petkov, "Fall detection dataset UG (Underground GATE)," *Dataset*. [Online]. Available: <https://zenodo.org/records/17777971>
20. K. Nakayama and G. H. Silverman, "The aperture problem—I. Perception of nonrigidity and motion direction in translating sinusoidal lines," *Vision Res.*, vol. 28, no. 6, pp. 739–746, 1988, doi: [https://doi.org/10.1016/0042-6989\(88\)90052-1](https://doi.org/10.1016/0042-6989(88)90052-1).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.