# Infant Movement Detection via Eigenvalue-Entropy Based Subspace Method

**Leonardo Yuto Suzuki Camelo[1], Bernardo Bentes Gatto[2], Ayrles Mendonça[3], Rafael Giusti[1], Eulanda Miranda dos Santos[1]**

[1]Federal University of Amazonas, Institute of Computing, Manaus, AM, Brazil

[2]MTI, Tokyo, Japan

[3]Federal University of Amazonas, Faculty of Physical Education and Physiotherapy, Manaus, AM, Brazil

{leonardo.yuto, rgiusti, emsantos}@icomp.ufam.edu.br, gatto_b@mti.co.jp, ayrles@ufam.edu.br

***Abstract.*** *The early identification of anomalous movements in infants is crucial for intervening in potential neuromotor development disorders. The clinical method General Movement Assessment (GMA) is devoted to this identification task. However, since GMA is intensive and requires experts, new machine learning-based approaches and keypoints extracted from videos have emerged. However, challenges such as the underrepresentation of infants with writhing movements (WM)—general movements presented by infants in their first weeks of life; the scarcity of public datasets; and the fact that only video segments showing infants performing movements must be analyzed, are limitations to identify anomalous movements in infants automatically. This work introduces a method which uses spatial distance features extracted from skeletal data and employs subspace method based on the statistical analysis of the eigenvalue-entropy to enhance the detection of infants movements in video data, especially video from infants exhibiting WMs. The proposed method applies a subspace approach as an initial step to filter infant movements for further detection and subsequent classification, aiming to improve the detection and understanding of these critical early indicators. The results show that the proposed method is able to detect subtle nuances in infant movements more effectively than the baseline method, making it a promising tool for automatic developmental monitoring.*

## 1. Introduction

Early and accurate detection of abnormal movements in infants is crucial for identifying potential neuromotor disorders. Motor development in infancy is a significant indicator of neurological health and overall well-being. Hence, the development of effective methods to analyze and classify infant movements is vital [Silva et al. 2021].

The GMA [Einspieler et al. 1997], introduced by Prechtl as a clinical method for analyzing infants' general movements (GMs), focuses on quantifying the variability and smoothness of a baby's spontaneous body movements. These GMs are part of the spontaneous movement repertoire, which, depending on the infant's age, is divided into writhing movements (WMs), occurring from 8 to 12 weeks of age, and fidgety movements (FMs), occurring from 12 to 20 weeks [Irshad et al. 2020]. However, GMA is a costly and challenging process, requiring an expert professional due to the irregular occurrence of movements and demanding prolonged observation of the subject [Gong et al. 2022, Irshad et al. 2020].

In light of these facts, recent advancements in artificial intelligence have led to numerous studies employing machine learning techniques as alternatives to GMA [Groos et al. 2022, Gao et al. 2023]. The primary methodology of these works involves using keypoint representations extracted from videos of infants for subsequent classification. The keypoint representations are especially used due to their invariance to appearance and background information that may be present videos and for better modeling of infant movements, serving as input for extracting motion representation information [Nguyen-Thai et al. 2021, Gao et al. 2023].

However, several challenges remain in the field. Most studies focus on term infants and primarily work with FMs [Gong et al. 2022], which may result in delayed diagnoses. There is also a significant use of high-complexity methods for classifying normal and abnormal movements, such as Graph Convolutional Networks (GCN) [Groos et al. 2022] and Transformers [Gao et al. 2023]. Moreover, there is a limited number of publicly available datasets for classifying and recognizing infant movements [Marschik et al. 2022]. Another observed issue in the collected database is the presence of video segments where the infant shows no movement, which can lead to redundant analysis both for the classifier and for the GMA expert, by analysing data that does not contain relevant information.

Recognizing whether the infant is moving or not is crucial for the classification process. Video recordings in accordance with the Prechtl GMA are typically lengthy, and depending on the age of the infant, the recording length may vary from 10 minutes up to 1 hour. Those recordings are usually analysed by an expert and shorter sequences of 2 to 5 minutes of the awake and moving infant are used for the assessment [Einspieler et al. 1997, Ferrari et al. 2004, Aizawa et al. 2021]. However, infants may spend considerable time immobile, for both the lengthy and shorter videos, making the process time consuming by rendering segments less useful for analysis.

Notably, much of the existing literature related to automating the movement recognition task overlooks the movement detection phase, opting instead for short, pre-segmented videos that concentrate solely on the movement classification stage. On the other hand, an ideal automated system should encompass a stage earlier to the movement classification, ensuring a focused analysis on periods of movements, so as to enhance efficiency. Figure 1 illustrates a conceptual diagram of such an ideal automated system.

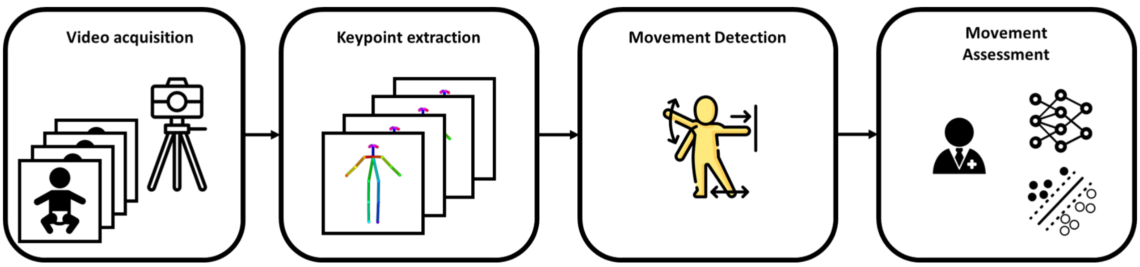


**Figure 1. Conceptual diagram of an ideal infant movement analysis and classification system.**

With these challenges in mind, in this paper, we focus on the movement detection stage. The goal is to detect and differentiate when the infant is moving, aiming for an early and more accurate identification of potential irregularities in neuromotor development. Our proposed method employs a subspace-based method, which requires low computational resources, and can generalize data distribution even for few samples, as the dataset used in this work. Additionally, subspace-based methods are statistically robust to noise in input data and are effective in compressing data into a low-dimensional format, offering a promising approach for primary movement detection [Gatto et al. 2021].

The main contributions of this paper are summarized as follows:

1. We introduce a discriminative method based on subspace representation, which uses a distance matrix constructed between the joints in each frame to capture the potential coordinates relationship between nonadjacent joints and represent the overall motion patterns;
2. We propose a movement detection method based on the entropy of the eigenvalues generated by the motion of the infant, demonstrating a more discriminative capability for this task.
3. We evaluate the performance of our method compared to a traditional image-based method using optical flow.

## 2. Related Work

In this section, we review recent works related to machine learning algorithms used to detect abnormal motor patterns, as a mean to assess infants at risk of developing neuromotor impairments, with a focus on approaches that work directly on visual signals.

In the study by [Groos et al. 2022], a deep learning method was applied to analyze FMs in infants at high risk of perinatal brain injury. These infants included preterm babies with very low birth weights and those with various neurological conditions. The videos were recorded in accordance with Prechtl's GMA standards, showing awake infants in a supine position for an average of 5 minutes with a rate of 30 frames per second. The proposed method, employing the EfficientPose model for skeleton tracking and an ensemble of GCNs, showed promising results in identifying infants with and without cerebral palsy (CP), outperforming traditional machine learning algorithms that rely on manually selected movement features. However, the authors divided the original videos in 5 second windows, using the median CP score across all 5 second windows as final score. Therefore, a movement detection step was not performed.

FidgetyFind is proposed in [Morais et al. 2023]. It is an approach designed for the detection and assessment of FMs in infants, closely mimicking expert evaluations through the analysis of video segments by assigning "fidgety scores" to body joints based on multidirectional movement intensity. The authors did not specify the recordings length, but it is mentioned that they used overlapping windows of 50 frames and a stride of 20 frames as input. This method achieved superior performance in recognizing FMs without needing annotated labels, providing clear insights into the assessment process by identifying influential joint movements and times. Although it excels in interpretability and accuracy, skeleton detection inaccuracies and distinguishing non-FM multidirectional movements are pointed out by the authors as challenges to deal with.

In [Gao et al. 2023], the Motor Assessment Model (MAM) is introduced, which leverages 3D pose estimation and a multi-instance learning (MIL) framework combined with Transformers to automate GMA for early cerebral palsy (CP) screening. The videos were recorded at a standardized value of 25 frames per second, and the infant clips were cropped to 9.6 seconds, with a step size of 6 seconds, resulting in input windows of 240 frames length, with 150 frames step. MAM predicts infant joint coordinates and uses a distance representation to analyze motion patterns, incorporating a dedicated FMs reference branch and a specialized Closeness loss function to enhance interpretability and alignment with GMA standards. The approach outperformed existing methods by integrating multimodal data, including infants' basic characteristics, although video features were predominantly more important for the prediction accuracy. This quantitative and explainable model offers a new direction for video-based medical diagnostics, emphasizing the potential for automated GMA in early CP detection.

Unlike the previously described methods, which focused on FMs, [Palheta et al. 2023] propose an approach for classifying infant general movements dealing with videos of infants ranging in the writhing movements (WMs). Their method employed two 3D Convolutional Neural Networks (CNNs): one trained with RGB video frames and the other with keypoints. Three different fusion functions were compared to combine the two models. Experiments were conducted on a small proprietary dataset containing both normal and abnormal infant movements. The videos were segmented into 30 second sequences to standardize the video duration, which originally varied from 1:30 minutes to 3 minutes. This segmentation, performed at a frame rate of 30 fps, resulted in 900 frames per instance. Their results indicated that the combination of two CNNs slightly outperformed the individual channel-based models, highlighting the potential for improved classification through combined modalities.

With the analysis of the related works, we concluded that, although the recording length of the GMs of infant are conducted according to the Prechtl guideline [Ferrari et al. 2004, Aizawa et al. 2021], all the reviewed works in literature segment the videos in small sequences of infant movements to feed their methods, performing local analysis that may not reflect the global outcome of the infant GMs.

Similar to the mentioned approaches, our method uses video as input, employing body skeletal data by applying pose estimation algorithm as the main feature. The main difference, however, is that our method aims to facilitate the GMA process by extracting only the portions of the videos containing relevant movement information, being a step prior to the approaches listed here.

## 3. Proposed Method

In this section we describe the proposed movement detection method, and a description of the dataset employed in our experiments. We introduce the procedure for information acquisition and pre-processing. After that, we describe the method for extracting the distance feature and the subspaces for human keypoints. Finally, we present an approach to classify the movement features subspaces by analysing the entropy generated by the eigenvalues of the subspaces. Figure 2 illustrates a diagram summarizing the proposed method:

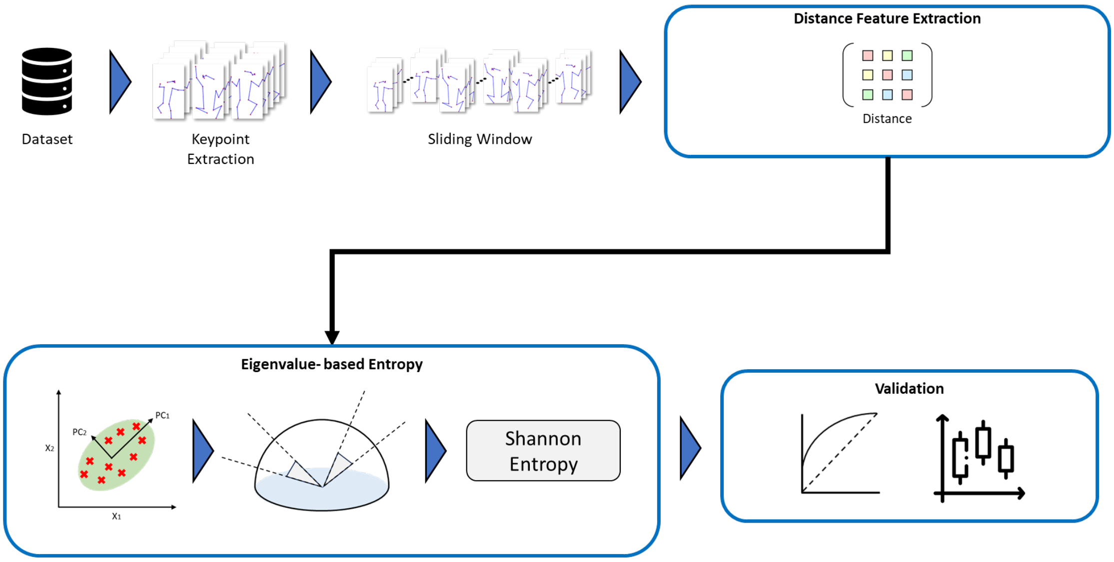


**Figure 2. The process of the proposed infant movement detection method.**

### 3.1. Dataset

The original dataset used in our experiments consists of 22 videos featuring preterm or low-weight infants, employed previously in [Palheta et al. 2023]. The videos range from 40 seconds to 3 minutes and were recorded at 30 frames per second. Figure 3 shows frames of two videos of the dataset, anonymized and with the keypoints overlapping.

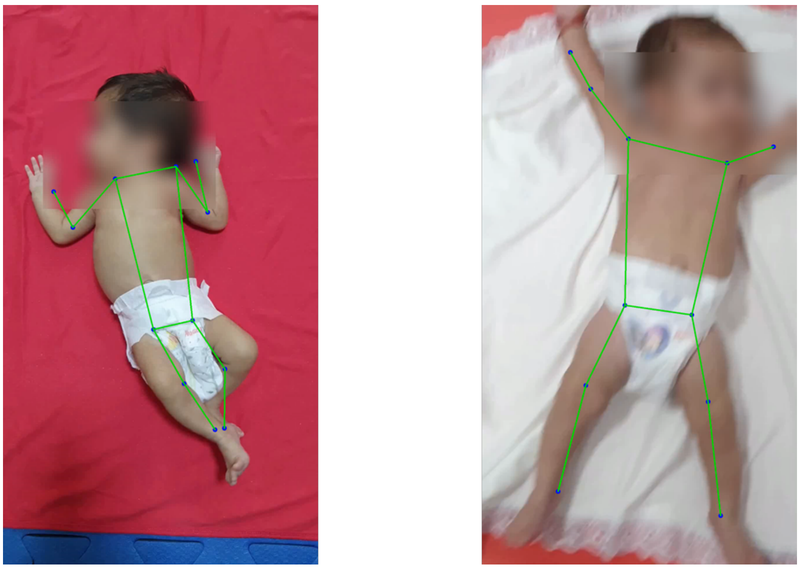

**Figure 3. Examples representing one frame of two instances of the dataset.**

Despite being a very small dataset, we still had to reduce the original dataset, since we had to remove videos presenting the following characteristics: showing the infant interacting with external individuals; where the infant is in continuous movement during the whole video; or videos with improper framing. Then, the final dataset was reduced to only 12 videos. These selected videos were manually labeled to indicate time intervals of movement or stillness.

### 3.2. Data Acquisition and Preprocessing

The raw videos were processed using the YOLO-pose model [Maji et al. 2022], an advanced deep learning pose estimation tool, chosen for its effectiveness in detecting two-dimensional coordinates of key body points, thus creating a skeletal representation of the infant. For each frame, the model outputs a set of 17 keypoints, each with $x$ and $y$ coordinates, providing a detailed basis for subsequent analysis. The selection of YOLO-pose, implemented by Ultralytics library [Jocher et al. 2023], is due to its accuracy, efficiency, and ease implementation, as demonstrated in various contexts and pretrained on the COCO dataset [Lin et al. 2015]. The keypoint sequence was segmented into 60-frame windows (2 seconds), each overlapping by 50%, for processing through feature extraction and classification to estimate movement in each window.

### 3.3. Keypoint Subspace Extraction

After the acquisition and preprocessing of keypoints, our methodology involves two stages: feature extraction and decomposition. Initially, the raw keypoints are transformed into a set of informative features. Subsequently, the extracted features are subjected to a decomposition process, to provide information of their variance.

In our work, this step is conducted based on the feature construction employed in [Gao et al. 2023], as follows. First, for each frame window, we construct pairwise Euclidean distance matrices between all body joints for each frame. These matrices encapsulate the spatial relationships between body joints, providing a rich representation of the body's posture and configuration at each moment in time. Given two points $P_1(x_1, y_1)$ and $P_2(x_2, y_2)$, the distance $d$ between them is calculated as $d(P_1, P_2) = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$. The result matrix $D$ is as shown below:

$$D = \begin{bmatrix} d_{1,1} & \cdots & d_{1,J} \\ \vdots & \ddots & \vdots \\ d_{J,1} & \cdots & d_{J,J} \end{bmatrix}$$

where $J$ is the number number of joints of the infant in each frame extracted by the pose estimation model, and $d_{ij}$ is the Euclidean distance between the $i$-th joint and the $j$-th joint. The input feature for each frame window can be represented by a $[T, J, J]$ tensor, where $T$ is the number of the frames.

By constructing pairwise Euclidean distance matrices for each frame window, this method captures complex spatial relationships between body joints, ensuring temporal consistency and allowing for precise, quantitative analysis. It facilitates detailed comparisons between time frames, enhancing the understanding of movement dynamics and patterns. For instance, by examining and contrasting the Euclidean distance matrices $D$ of different babies, clinicians gain a refined tool for assessing and comparing their movement dynamics, facilitating a deeper understanding of each infant's developmental progression. Then, each distance matrix is subsequently transformed into a feature vector. This transformation is done by extracting the upper triangle of each distance matrix and flattening it, converting complex joint configurations to obtain a more compact form.

To capture the temporal evolution of movements, we analyze the sequence of feature vectors derived from consecutive frames within the time window. This involves computing the autocorrelation matrix for each time window features, which is decomposed using Singular Value Decomposition to extract eigenvalues and eigenvectors that represent the variance in the data and the directions of maximal variance, respectively.

Finally, to quantify the complexity and uncertainty in the movement patterns, we compute the Shannon entropy of the normalized eigenvalue distribution. This metric provides a measure of the information content and diversity in the temporal dynamics of infant movements, offering a nuanced understanding of motor development, as detailed in the next subsection.

### 3.4. Eigenvalue-based Entropy

As described in the previous section, the eigenvalues represent the variance in the data. As a means to further extract discriminative properties of the movement features, we propose computing the Shannon entropy of the normalized eigenvalue distribution. Eigenvalue-based entropy have been extensively used in various fields, for instance in network topology [Sun et al. 2021] and data sampling [Huang et al. 2023].

Shannon's entropy provides a description of the information content and diversity in a message, and is computed as follows:

$$H = -\sum_{i=1}^{N} p_i \log p_i$$

where $N$ is the number of values a random variable can have, and $p_i$ is the probability of the random variable having the value of $i$.

Once equipped with the normalized eigenvalues decomposed by the features of each time window, the eigenvalue-based entropy is computed as follows:

$$H = -\sum_{i=1}^{N} \lambda_i^n \log \lambda_i^n$$

where $\lambda_i^n$ is the $i^{th}$ eigenvalue of the autocorrelation matrix, with the superscript $n$ indicating the normalized value for the eigenvalue.

The intuition behind the analysis of eigenvalue-based entropy in the context of infant movement detection lies in its ability to quantify the complexity of movement patterns. When an infant is in motion, the variance across the data points, reflected in the eigenvalues, tends to increase due to the dynamic and diverse nature of the movements. Higher variance in the data implies a broader spread of eigenvalues, indicating that the movement is not confined to a few predictable patterns but rather spans a more extensive range of motion. Shannon entropy, in this scenario, serves as a measure of this spread or diversity. A higher entropy value suggests a more complex and less predictable pattern of movements, as is typically the case when an infant is actively moving. Each eigenvalue contributes to the entropy calculation in proportion to its magnitude, with larger eigenvalues (indicating greater variance and, therefore, more significant movement) contributing more significantly to the entropy value.

In periods of rest or minimal movement, the variance in the data, and consequently the eigenvalues, would be lower, leading to a more concentrated eigenvalue distribution and lower entropy values. Conversely, active movement periods are characterized by higher variance and a more spread-out eigenvalue distribution, resulting in higher entropy. By analyzing the entropy of the eigenvalue distribution, we can infer the level of activity in the infant's movements. Higher entropy values indicate more complex, varied, and less predictable movements, suggesting active motion.

Finally, with the entropy value for each time window for each video, the classification of movement or stillness is done by using a fixed threshold for all videos, obtained by calculating the average optimal entropy threshold that maximized the subtraction between the true positive rate and false positive rate of the area under the ROC curve (ROC-AUC) score for each video.

### 3.5. Evaluation Metrics

In order to assess the performance of our method for detecting infant movements, labels were attributed to each time window based on the timestamps associated with the infant's movement. For each video, metrics of accuracy, precision, recall, f1-score and ROC-AUC were obtained. Table 1 illustrates each metric and their corresponding formulas:

| **Metric** | **Formula** |
|---|---|
| Accuracy | $\frac{\text{Number of Correct Predictions}}{\text{Total Number of Predictions}}$ |
| Precision | $\frac{\text{True Positives}}{\text{True Positives + False Positives}}$ |
| Recall | $\frac{\text{True Positives}}{\text{True Positives + False Negatives}}$ |
| F1-Score | $2 \times \frac{\text{Precision}\times\text{Recall}}{\text{Precision + Recall}}$ |
| ROC-AUC | $\int_{x=0}^{1} \text{ROC}(x)\, dx$ |

**Table 1. Evaluation Metrics and their Corresponding Formulas**

## 4. Experimental Results

In our experiments, we evaluated the proposed method on the infant movement dataset. We adopted a fixed sliding window size of 60 frames (equivalent to 2 seconds), with an overlap of 50%. In terms of data representation, we utilized all the eigenvalues generated in the decomposition process to retain the full variance present in the data.

To establish a benchmark for our study, we compared our method to a baseline utilizing optical flow, a classical technique for analyzing motion. Optical flow calculates the motion between two consecutive image frames based on the motion of an object or camera or both [Shah and Xuezhi 2021]. It quantifies the displacement of pixels from one frame to the next, facilitating the computation of flow magnitude for each image frame within the same time window. This average flow magnitude is then used for classification.

For both the proposed method and the optical flow approach, a fixed threshold was applied to differentiate between movement and stillness. The threshold for computing the optical flow was obtained by calculating the average optical magnitude that maximized the ROC curve score, as described in Subsection 3.4.

Figure 4 presents a box plot showing the performance of our method (in green) compared to the baseline (in orange) across five different metrics. The proposed method's performance is robust, surpassing the baseline across all metrics, suggesting a more effective detection of infant movements. However, a noticeable variance in recall for the proposed method indicates variability in the detection of true movements across different instances, which demands further investigation. Precision, on the other hand, is more stable, indicating that when the method predicts movement, it is likely correct. The proposed method's ROC-AUC scores are consistently above the no-discrimination level, highlighting its superior discriminative capability.

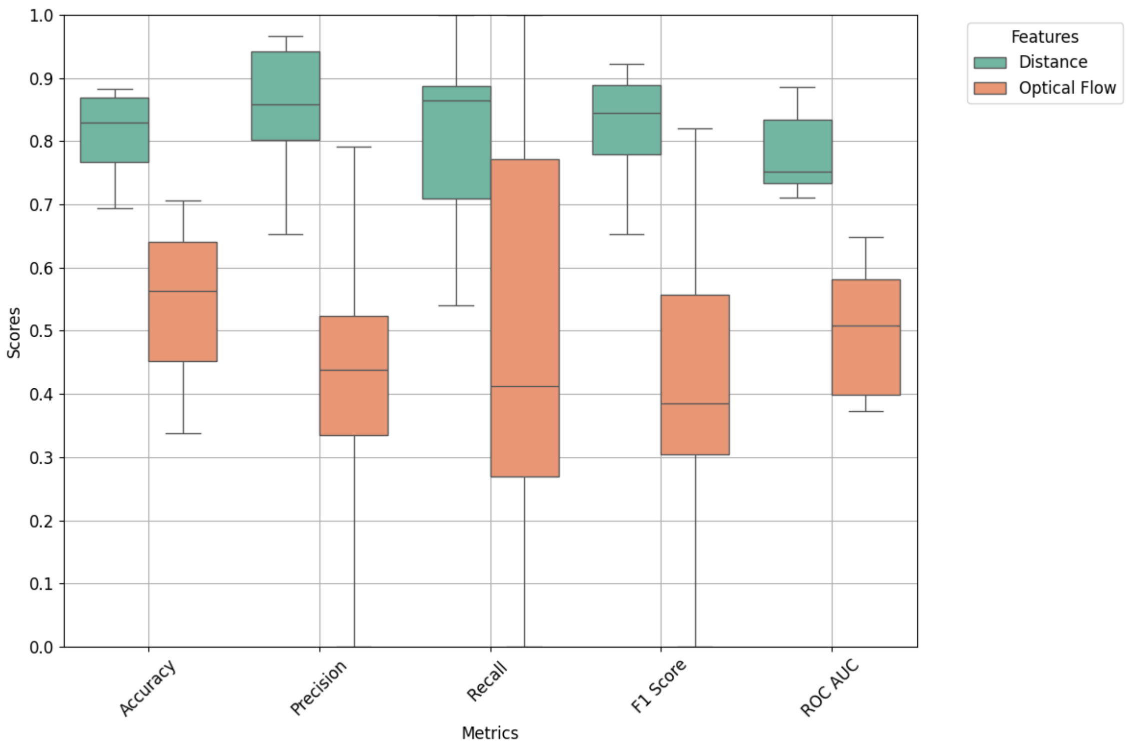


**Figure 4. Comparing the performance of our method (in green) and the performance of the baseline (in orange).**

Table 2 summarizes the average rates for each metric across all videos in the dataset. The upper values are related to the movement class, and the lower values denote the stillness class, with standard deviations provided in parentheses. The proposed method demonstrates high accuracy, precision, and recall for detecting movement, with lower but still significant scores for stillness detection. This difference suggests a potential area for improvement in recognizing stillness as precise as movements are recognized.

The results show the proposed method's efficacy in detecting infant movement using distance features, suggesting that spatial relationships between body joints provide a more nuanced depiction of motion than optical flow. Nevertheless, the considerable spread in recall and the decreased performance in stillness detection indicate the need for a refined approach that can more uniformly capture the diverse range of infant movements.

| Features | Accuracy | Precision | Recall | F1 Score | ROC AUC |
|---|---|---|---|---|---|
| Distance | **0.813 (0.062)** | **0.855 (0.092)**<br>**0.719 (0.182)** | **0.808 (0.144)**<br>**0.758 (0.164)** | **0.823 (0.094)**<br>**0.714 (0.124)** | **0.783 (0.064)** |
| Optical Flow | 0.545 (0.119) | 0.579 (0.303)<br>0.429 (0.209) | 0.520 (0.288)<br>0.483 (0.326) | 0.506 (0.251)<br>0.418 (0.231) | 0.501 (0.100) |

**Table 2. Performance metrics for Distance and Optical Flow features with standard deviations in parentheses.**

The high precision for the stillness class implies that the proposed method is accurate when it predicts stillness but could benefit from improved sensitivity to avoid missing instances of inactivity. The F1 Score, while high for the movement class, highlights a trade-off in detecting stillness, which could be addressed by balancing the dataset or adjusting the classification threshold. The superior ROC-AUC values for the proposed method validate its ability to distinguish between movement and stillness effectively. The use of a comprehensive eigenvalue spectrum to represent data ensures that no variance is disregarded, which is crucial in capturing the full complexity of infant movements. However, this approach may also include noise, contributing to the variability observed in recall scores.

In addition to yielding encouraging results, the use of eigenvalue of pairwise Euclidean distance matrices in our method enhances robustness to external noise and provides a detailed spatial analysis of body joint relationships, outperforming the sensitivity and pixel-level focus of optical flow on detecting infant movements. The critical analysis indicates that while the proposed method is promising, there is room for improvement in detecting stillness. Future enhancements could include employing class-weighting techniques, refining thresholding strategies, or exploring advanced feature engineering. Moreover, employing sophisticated classification algorithms or deep learning models could further optimize the detection accuracy and reliability across all metrics.

## 5. Conclusion and future directions

This study presented a methodology for analyzing infant movement using subspace method and Shannon entropy with keypoint distance features. The method demonstrated substantial promise, outperforming the baseline optical flow approach in terms of several standard evaluation metrics, including accuracy, precision, recall, F1 score, and ROC-AUC. The use of distance features to capture the spatial configuration of body keypoints over time provided an important approach that yielded a high degree of predictive capability for identifying infant movement.

However, the experimental results also revealed areas for improvement. While the system showed high accuracy and precision, the variability in recall and the lower performance in detecting stillness suggest that the method may benefit from further refinement to better capture the subtleties of infant movements, particularly during less active periods. These limitations may partially be due to constraints encountered in this study. The manual, non-expert labeling of movement and stillness intervals could have affected the precision of data annotation due to the intricate nature of infant movements. Additionally, inconsistencies in video captures, such as non-standardized environments and occasional camera movement, introduced challenges like frame shifts and off-center subjects.

Despite these challenges, our method provides a practical tool by transforming raw movement data into high-level features. Clinicians gain a direct pathway to analyse the intricacy and variability in infant movements, providing a robust framework for monitoring developmental milestones and identifying potential anomalies.

In future work, we aim to compensate the non-stationary camera by employing video stabilization algorithms, as we expect our solution to be used by non-experts and the infants' families. We also propose to integrate our methodology with a classifier for anomalous infant movements in an end-to-end system to further evaluate its efficacy and impact on the overall classification process.

## Acknowledgement

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES-PROEX) - Finance Code 001. This work was partially supported by Amazonas State Research Support Foundation - FAPEAM - through the POSGRAD 23-24 project. This research, carried out within the scope of the Samsung-UFAM Project for Education and Research (SUPER), according to Article 39 of Decree n°10.521/2020, was funded by Samsung Electronics of Amazonia Ltda., under the terms of Federal Law n°8.387/1991 through agreement 001/2020, signed with UFAM and FAEPI, Brazil.

## References

[Aizawa et al. 2021] Aizawa, C. Y. P., Einspieler, C., Genovesi, F. F., Ibidi, S. M., and Hasue, R. H. (2021). The general movement checklist: A guide to the assessment of general movements during preterm and term age. *Jornal de Pediatria*, 97:445–452.

[Einspieler et al. 1997] Einspieler, C., Prechtl, H. F., Ferrari, F., Cioni, G., and Bos, A. F. (1997). The qualitative assessment of general movements in preterm, term and young infants—review of the methodology. *Early human development*, 50(1):47–60.

[Ferrari et al. 2004] Ferrari, F., Einspieler, C., Prechtl, H., Bos, A., Cioni, G., et al. (2004). *Prechtl's method on the qualitative assessment of general movements in preterm, term and young infants*. Mac Keith Press.

[Gao et al. 2023] Gao, Q., Yao, S., Tian, Y., Zhang, C., Zhao, T., Wu, D., Yu, G., and Lu, H. (2023). Automating general movements assessment with quantitative deep learning to facilitate early screening of cerebral palsy. *Nature Communications*, 14(1):8294.

[Gatto et al. 2021] Gatto, B. B., Fukui, K., Júnior, W. S., and dos Santos, E. M. (2021). Advances in subspace learning and its applications. In *Anais Estendidos do XXXIV Conference on Graphics, Patterns and Images*, pages 35–41. SBC.

[Gong et al. 2022] Gong, X., Li, X., Ma, L., Tong, W., Shi, F., Hu, M., Zhang, X.-P., Yu, G., and Yang, C. (2022). Preterm infant general movements assessment via representation learning. *Displays*, 75:102308.

[Groos et al. 2022] Groos, D., Adde, L., Aubert, S., Boswell, L., De Regnier, R.-A., Fjørtoft, T., Gaebler-Spira, D., Haukeland, A., Loennecken, M., Msall, M., et al. (2022). Development and validation of a deep learning method to predict cerebral palsy from spontaneous movements in infants at high risk. *JAMA network open*, 5(7):e2221325–e2221325.

[Huang et al. 2023] Huang, J., Yoon, H., Wu, T., Candan, K. S., Pradhan, O., Wen, J., and O'Neill, Z. (2023). Eigen-entropy: A metric for multivariate sampling decisions. *Information sciences*, 619:84–97.

[Irshad et al. 2020] Irshad, M. T., Nisar, M. A., Gouverneur, P., Rapp, M., and Grzegorzek, M. (2020). Ai approaches towards prechtl's assessment of general movements: A systematic literature review. *Sensors*, 20(18):5321.

[Jocher et al. 2023] Jocher, G., Chaurasia, A., and Qiu, J. (2023). Ultralytics yolov8.

[Lin et al. 2015] Lin, T.-Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C. L., and Dollár, P. (2015). Microsoft coco: Common objects in context.

[Maji et al. 2022] Maji, D., Nagori, S., Mathew, M., and Poddar, D. (2022). Yolo-pose: Enhancing yolo for multi person pose estimation using object keypoint similarity loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2637–2646.

[Marschik et al. 2022] Marschik, P. B., Kulvicius, T., Flügge, S., Widmann, C., Nielsen-Saines, K., Schulte-Rüther, M., Hüning, B., Bölte, S., Poustka, L., Sigafoos, J., et al. (2022). Open video data sharing in developmental and behavioural science. *arXiv preprint arXiv:2207.11020*.

[Morais et al. 2023] Morais, R., Le, V., Morgan, C., Spittle, A., Badawi, N., Valentine, J., Hurrion, E. M., Dawson, P. A., Tran, T., and Venkatesh, S. (2023). Robust and interpretable general movement assessment using fidgety movement detection. *IEEE Journal of Biomedical and Health Informatics*.

[Nguyen-Thai et al. 2021] Nguyen-Thai, B., Le, V., Morgan, C., Badawi, N., Tran, T., and Venkatesh, S. (2021). A spatio-temporal attention-based model for infant movement assessment from videos. *IEEE journal of biomedical and health informatics*, 25(10):3911–3920.

[Palheta et al. 2023] Palheta, M., Santos, G., Mendonça, A., Gonçalves, P., Albuquerque, R., Souto, E., and Santos, E. (2023). Fusão de dados de vídeos rgb e pontos-chaves para classificação de movimentos gerais de bebês. In *Anais do XXIII Simpósio Brasileiro de Computação Aplicada à Saúde*, pages 384–394, Porto Alegre, RS, Brasil. SBC.

[Shah and Xuezhi 2021] Shah, S. T. H. and Xuezhi, X. (2021). Traditional and modern strategies for optical flow: an investigation. *SN Applied Sciences*, 3(3):289.

[Silva et al. 2021] Silva, N., Zhang, D., Kulvicius, T., Gail, A., Barreiros, C., Lindstaedt, S., Kraft, M., Bölte, S., Poustka, L., Nielsen-Saines, K., et al. (2021). The future of general movement assessment: The role of computer vision and machine learning–a scoping review. *Research in developmental disabilities*, 110:103854.

[Sun et al. 2021] Sun, Y., Zhao, H., Liang, J., and Ma, X. (2021). Eigenvalue-based entropy in directed complex networks. *Plos one*, 16(6):e0251993.