

A Pose-based Feature Fusion and Classification Framework for the Early Prediction of Cerebral Palsy in Infants

Kevin D. McCay , Pengpeng Hu , Hubert P. H. Shum , Wai Lok Woo ,
 Claire Marcroft, Nicholas D. Embleton , Adrian Munteanu  and Edmond S. L. Ho 

Abstract—The early diagnosis of cerebral palsy is an area which has recently seen significant multi-disciplinary research. Diagnostic tools such as the General Movements Assessment (GMA), have produced some very promising results. However, the prospect of automating these processes may improve accessibility of the assessment and also enhance the understanding of movement development of infants.

Previous works have established the viability of using pose-based features extracted from RGB video sequences to undertake classification of infant body movements based upon the GMA. In this paper, we propose a series of new and improved features, and a feature fusion pipeline for this classification task. We also introduce the RVI-38 dataset, a series of videos captured as part of routine clinical care. By utilising this challenging dataset we establish the robustness of several motion features for classification, subsequently informing the design of our proposed feature fusion framework based upon the GMA. We evaluate our proposed framework's classification performance using both the RVI-38 dataset and the publicly available MINI-RGBD dataset. We also implement several other methods from the literature for direct comparison using these two independent datasets.

Our experimental results and feature analysis show that our proposed pose-based method performs well across both datasets. The proposed features afford us the opportunity to include finer detail than previous methods, and further model GMA specific body movements. These new features also allow us to take advantage of additional body-part specific information as a means of improving the overall classification performance, whilst retaining GMA relevant, interpretable, and shareable features.

Index Terms—Cerebral Palsy, Early Diagnosis, Explainable AI, General Movements Assessment, Machine Learning, Motion Analysis, Skeletal Pose Estimation

I. INTRODUCTION

CEREBRAL palsy (CP) is the collective term given to a group of lifelong neurological conditions caused by non-progressive damage to the brain [5], occurring before, during, or shortly after birth [49]. CP typically affects movement, muscle tone, posture and co-ordination, but can also cause difficulties with swallowing, speech-articulation, hearing, vision, and can impact upon an individual's ability to learn new skills [17].

CP is the most prevalent physical disability found in children, with 2.11 diagnoses per 1000 live births [41]. There is also an increased prevalence of CP in infants born prematurely, with 32.4 diagnoses per 1000 infants born very preterm (28-32 weeks gestation), and

Manuscript received October ??, 2020; revised ?? ??, 2020; accepted ?? ??, 2021. Date of publication ?? ?, 2021; date of current version August 16, 2020. (Corresponding author: Edmond S. L. Ho, e-mail: e.ho@northumbria.ac.uk)

Kevin D. McCay, Edmond S. L. Ho and Wai Lok Woo are with the Department of Computer and Information Sciences, Northumbria University, Newcastle upon Tyne, UK

Pengpeng Hu and Andrian Munteanu are with the Department of Electronics and Informatics, Vrije Universiteit Brussel, Belgium

Hubert P. H. Shum is with the Department of Computer Science, Durham University, UK

Claire Marcroft and Nicholas D. Embleton are with the Newcastle upon Tyne Hospitals Trust, Royal Victoria Infirmary, Newcastle upon Tyne, UK.

70.6 diagnoses per 1000 infants born extremely preterm (<28 weeks gestation) [47]. As the degree of prematurity increases so does the likelihood of severe disability, with one-in three of those surviving at 22 weeks gestation being severely disabled [14], [28]. As such, the early diagnosis of CP is an ongoing area of multidisciplinary research, as it has the potential to allow for early intervention clinical care. Early interventions look to optimise the neuroplasticity of the developing infant brain, thereby inhibiting the impact of impairment and subsequently reducing the likelihood of the child fully developing associative conditions [24]. However, early diagnosis can be difficult and time consuming [31].

Manual examinations, such as the General Movements Assessment (GMA) [16], have been developed to identify the emerging signs of CP. Studies suggest that the GMA compares favourably with other methods, such as cranial ultrasound and neurological examination, by producing more consistent and reliable individual results [6], with the GMA being identified as the leading method of predicting later CP [26]. The GMA is typically carried out by evaluating the quality of body movements at a specific window in an infant's development (Section II-A); however, the application of these assessments can be challenging, chiefly due to the availability of appropriately skilled clinicians [31].

There is increasing motivation to utilise technology to aid with clinical decision making, helping with logistical constraints, increasing the predictive accuracy and targeting early intervention [39]. As such, several works [1], [33]–[35], [40], [53] have proposed automated solutions to help address the challenges faced in the diagnosis of CP. The proposed methods typically make use of machine learning frameworks to automatically assess infants based upon the movement patterns associated with the GMA. These methods suggest that a machine learning framework could substantiate the decision making process, allowing for intuitive, quantitative, cost-effective, evidence-based evaluation [29], and also provide a means of fully remote diagnostic assessment.

However, it is also clear that the interpretability of the model must be considered, particularly when using machine-learning based approaches in the medical domain. Machine learning models are often seen as 'black boxes', in which the underlying structures can be difficult to comprehend. Consequently, clarity is required as to why a system arrives at a decision, and as such the mechanisms behind classification frameworks have to be transparent, understandable and explainable [21]. We therefore suggest that by using pose data in combination with histogram representations, we retain human interpretability throughout the classification pipeline. As such, we propose a feature-extraction, feature-fusion, and classification framework, which extracts several new GMA relevant features using pose data generated from standard 2D RGB video. By using this approach we aim to ensure that the extracted features are not only closely mapped to the assessment criteria specified in the GMA, but are also able to be understood by clinicians as a means of retaining predictive accountability and explainability. This approach also provides opportunities for remote data analysis without the need for specialist equipment, making wider clinical diagnosis feasible.

To evaluate our system and to inform the design of our proposed features, we subsequently collected a new video dataset (RVI-38), which reflects the complexity associated with video data gathered in real-world clinical settings for use in the GMA. By utilising the challenging RVI-38 dataset, we have developed and evaluated combinations of several different supplementary pose-based features. In doing so, we have produced a feature-set capable of generalising well across datasets of varying size, quality, and duration, whilst simultaneously dealing with superfluous information present in shot. We compare our method with several methods from the literature using both the RVI-38 dataset and the MINI-RGBD [20] dataset to evaluate the comparative robustness of our features. The source code and dataset used in this paper is available at (<https://github.com/emmondslho/Pose-basedCerebralPalsyPrediction>).

The contributions of this work can therefore be summarised as:

- A series of new and enhanced pose-based interpretable features, extracted from 2D video sequences, and based upon specific criteria set out in the GMA.
- Experimental re-implementation, comparative evaluation, and discussion of several prominent previous methods, along with our proposed method, are undertaken using shared datasets, for unbiased assessment and the generation of a new benchmark.
- The introduction of a challenging new video dataset, composed of real-world patient data. This GMA specific dataset informs our proposed feature design and comparative analysis. Additionally, given the difficulty in acquiring data in this sensitive area, we also make the extracted pose dataset and associated GMA labels available to the community
- Analysis of a new automated feature extraction, fusion, and classification pipeline, for the prediction of cerebral palsy based upon the GMA. We also make this framework available to the community to further encourage research in this field.

II. RELATED WORKS

In this section, we discuss the GMA as well as relevant studies which have attempted to automate this diagnostic process using computer vision based approaches. We provide an overview of the associated feature extraction processes, with a view to including several of the discussed methods as comparative evaluation baselines.

A. General Movements Assessment

The GMA is a non-invasive and non-intrusive physical examination of infants for the early detection of neurological anomalies associated with CP. General movements (GMs) are body movements which are spontaneous and variable, and form part of a distinct pattern of movements, called Fidgety Movements (FMs), which can be observed from early fetal life through to around 20 weeks post-term [16].

In a typically developing infant, FMs wax and wane in intensity, speed, frequency, amplitude, and range of motion, with notable fluctuations to the rotation, orientation, and displacement around the limb axes [42]. Conversely, abnormal GMs are identified by the absence of FMs, with a lack of duration, variability, and complexity throughout the movement sequence [48]. FMs are seen as a reliable indicator of brain dysfunction and as such, the presence of these abnormal GM patterns is seen as a strong predictor that the infant will receive a confirmed diagnosis of CP in later life [15].

The ability to provide an early diagnosis of CP is considered key, as early intervention care can contribute towards improved motor-function ability due to the application of neuro-developmental treatments, which optimise the “brain plasticity” found in developing infants [37]. As such, the earlier a confirmed diagnosis can be made, the better the outlook for the infant’s long term development [58].

However, whilst diagnostic tools such as the GMA have been proven to provide accurate prediction of CP, the training required to become a suitably proficient assessor is significant.

Motivated by this, an initial study by Adde et al. [1], proposed a system in which computer vision techniques were employed to assess the feasibility of undertaking automated GMA. In this study, a representation of the associated motion was created by generating a “motion image”, which quantified the amount of infant movement present in each video through frame differencing. The generated motion images were then used to generate features for binary classification of normal or abnormal infant movements. However the use of difference images is subject to several limitations, such as sensitivity to camera movement, issues with self-occlusion, strict background pre-processing, and a lack of information about the speed and direction of objects moving in frame [43], [53]. As such, several studies have subsequently attempted to further evaluate the viability of automated GMA through more advanced computer vision-based techniques, such as optical flow.

B. Automated Optical Flow-based Methods

In optical flow based methods, the motion is represented by a displacement vector field of the pixels between consecutive frames of an image sequence, rather than simply calculating the gross quantity of motion present in the frame and localising this centrally as previously proposed [53]. This allows for more detailed tracking of infant body-parts and the associated movements, and subsequently deeper analysis of the correlated motion.

In [53], optical flow data was extracted from video sequences before wavelet frequency analysis was implemented in order to calculate the time-dependant trajectory signals found in the data. However, using this optical flow method presented issues with tracking larger movements, and as such it was suggested that video captured at a higher frame rate would be required for future analysis. In order to address this problem several works [23], [40], [44], [45] implemented large displacement optical flow (LDOF) [8], to track infant body movements and extract features for classification. The LDOF method extends optical flow to better deal with large displacements of foreground objects and camera movement, making it more suitable for detailed infant motion analysis. This enhancement allowed for the improvement of existing motion quantity features [23], the development of velocity-based features [40], [44], and the introduction of frequency-based features [45] for infant motion analysis and binary classification. Through comprehensive analysis of these movement-based features, it was subsequently determined that dynamic features are typically more predictive than the previous statistical features [45]. However, similar challenges are present, such as issues with occlusion, drift and noise [54], as well as susceptibility to unrelated movements (such as equipment, parents or clinicians in shot), and sensitivity to illumination changes [43]. Additionally, the low generalizability and interpretability of extracted features make translation to clinical practice less likely [43], [45]. Finally, the motion data does not easily assess body parts separately, typically relying upon whole body movement analysis [10], [46]. By incorporating part-based assessment, the potential to provide further analytic cues for CP prediction and the identification of CP sub-types is enhanced [43]. As such, alternative methods have been proposed to better model human shape and motion, such as pose estimation [9], [22].

C. Pose Estimation-based Methods

Pose estimation is the task of using computer vision and machine learning techniques, to detect human figures in images and video,

and determine their pose by estimating the spatial locations of key body joints. In this particular use case, it is suggested that pose estimation is more robust to several factors negatively affecting the discussed methods [11], such as illumination changes, camera movement, changes in resolution, inconsistencies in infant size, and larger movements between frames. Since pose-estimation provides localised joint estimation it is also better able to deal with external influences, as well as provide important motion information based upon individual body-part movements, requiring comparatively minimal manual tuning [10], [36].

Research into the feasibility of using pose-based assessment for GMA is ongoing, with several papers contributing advances in this area. In [10], [18], [27], [30] and [36], pose estimation and domain adaptation methods for video sequences of infants are proposed, with each suggesting that their methods are a step towards automated pose-based GMA. However, each typically assesses the effectiveness of the extracted pose rather than the effect this might have upon final classification. In contrast, [33], [34], [38], [51], [56], examine the viability of using pose-based features for the prediction of CP using a feature extraction and classification pipeline.

Whilst reasonable results have been reported in each of the previously discussed related methods, evaluation of the robustness of the proposed features has not been demonstrated across multiple datasets. Indeed, given the sensitive nature of the data required for evaluation of each of the discussed methods, it is difficult to assess their viability without full re-implementation and assessment on shared datasets. As such, we not only propose a wide range of new GMA relevant features capturing orientation, displacement, and frequency information (Section III-E), and a feature fusion pipeline (Section III-G), but also a full re-implementation of several methods for comparative testing across multiple datasets, including videos captured in a real-world clinical setting (Section IV).

III. METHODOLOGY

In this section we provide details of the study design, data collection and the composition of the two datasets used for evaluation. We also discuss the proposed framework, data pre-processing, feature extraction, and the classification techniques used.

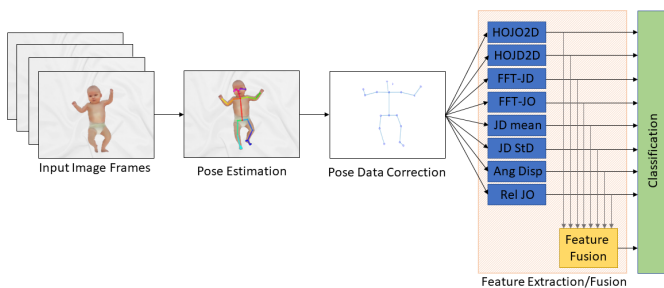


Fig. 1: Overview of the pose estimation (Section III-D), feature extraction (Section III-E) and classification (Section III-F) framework.

A. Study Design and Data Collection

For this collaborative project, a retrospective cohort study design was implemented for data collection to produce the RVI-38 dataset (detailed in Section III-B.1). Ethical approval was obtained from the host organisation (Ref: 9865), the Research Ethics Committee (REC), the Health Research Authority (HRA), and Health and Care Research Wales (HCRW) (Ref: 19/LO/0606, IRAS project ID: 252317). Parental consent for the use of video recordings was obtained by the clinical staff associated with the project prior to implementation. The

study population included infants who had a clinical GMA, with a video recording at 3–5 months post-term, as part of their routine follow-up care. The MINI-RGBD dataset (detailed in Section III-B.2) is an open source dataset, and all terms specified in the author’s license agreement were met.

B. Datasets

In order to comprehensively evaluate our feature extraction, fusion and classification pipeline, we made use of two separate datasets, details of which are discussed in this section. These datasets were also used to evaluate the selected baseline methods for comparison, as discussed in Section IV.

1) *RVI-38 Dataset*: An important part of this study is that the framework has to have the ability to generalise well across different datasets, particularly when processing real-world video data. To reflect this, the challenging new RVI-38 dataset was collected to inform the design of our framework, and for evaluative analysis. This dataset is composed of real patient video data gathered as part of routine clinical care at the Royal Victoria Infirmary (RVI) in Newcastle upon Tyne, and reflects the genuine intra-class variance and subsequent complexity present in the real-world clinical setting. The RVI-38 dataset consists of 38 videos, of 38 different infants aged between three and five months post-term. The videos were recorded using a handheld Sony DSC-RX100 Advanced Compact Premium Camera with a resolution of 1920x1080 @ 25 FPS. The duration of each video varied between a minimum of 40 seconds and a maximum of 5 minutes, with an average duration of 3 minutes and 36 seconds. The footage was captured from above, in a top-down orientation, with the infant lying in a supine position per the GMA guidelines. However, unlike many of the related works (e.g [1], [2], [10], [23], [40]), all video recordings were used as part of our evaluation, with no prior screening for inconsistencies such as poor lighting, camera movement, external factors, significant shadows, or additional selective pre-processing. Whilst this produced a more challenging dataset, it also represents a real-world evaluation of footage captured in a clinical setting. By including this challenging footage we hope to demonstrate that our system is more capable of being robust to variations in data capture, making it more suitable to clinical implementation. The videos were classified by two experienced assessors, using the GMA, into one of two categories; 1) FM+ where the infant demonstrates normal movements indicative of typical development; and 2) FM- where the infant demonstrates abnormal movement patterns that may be of concern to clinicians. This resulted in 32 videos being annotated as FM+ and 6 videos being annotated as FM-.

2) *MINI-RGBD Dataset*: One of the challenges facing researchers attempting to automate the GMA is the availability of suitable data. Given that the video data required for the GMA is of a sensitive nature, baseline datasets are not currently publicly available. Additionally, since human pose estimation frameworks are almost exclusively trained and tested using images of adults, a dataset consisting of images of infants for research purposes can understandably be difficult to obtain. To address this problem the Moving INfants In RGB-D (MINI-RGBD) dataset was generated and made publicly available [20]. This dataset maps real-world 3D infant movements, captured in a clinical setting, to virtual 3D models of infants. Photo-realistic videos of the 3D infant models were produced using computer graphics rendering, allowing for the generation of anonymised, and subsequently shareable footage, which retains the real-world movement characteristics required for the GMA. This dataset consists of 12 top down videos of infants lying in a supine position, each 40 seconds in duration. The videos were once again

analysed by two experienced GM assessors and similarly labelled, resulting in 8 videos being annotated as FM+ and 4 videos being annotated as FM-.

C. Proposed Framework Overview

In this paper, we make use of traditional machine learning classification algorithms for practicality and interpretability (Fig 1). We use raw video as input for the pose estimation framework to compute joint positions, which are then corrected to remove outliers and inconsistencies, as discussed in Section III-D. This corrected data is then used to generate features based upon the GMA for further analysis. Details of the features used and the extraction processes are discussed in Section III-E. These individually extracted features are then fed into a classification framework for evaluation, as discussed in Section III-F. Finally, we fuse the features together to create a more robust representation and evaluate this against both the individual features and the selected baselines, as discussed in Section IV.

D. Pose Estimation and Data Pre-processing

In order to extract meaningful features for subsequent analysis and classification, pose estimation methods are utilised in this paper. Specifically, we make use of the OpenPose framework [9] to extract joint positions from 2D RGB video data. The extracted joint positions form a skeletal pose representation consisting of 25 predefined joints, as shown in Figure 2. As such, each frame of each video is represented by 25 sets of 2D (x and y) coordinates and an associated confidence score for the prediction.

In this work, we make use of all the extracted joints with the exception of the facial landmarks (joints 16 to 19), and the feet (joints 20 to 25), as it was determined that these joints were less reliable than the other body landmarks acquired through OpenPose due to self occlusion errors, and were found to play a less important role in the final GMA-based classification results.

1) **Automated Data Correction:** To ensure consistency throughout the pipeline, the exported OpenPose data is pre-processed prior to feature extraction. This pre-processing involves remapping anomalous joint positions caused by self occlusion or inaccuracies in the OpenPose joint prediction process. The first stage is a qualitative evaluation of the extracted OpenPose data, to check that predicted joint positions and the associated confidence scores correctly align, and are consistent with the input video. We then use the predicted confidence scores to calculate a confidence threshold. The confidence threshold is calculated by taking the average confidence score, per joint, across each video sequence, and subtracting 5% from this. This means that we are able to remove joint positions with a lower confidence score than the confidence threshold, removing outliers on a frame by frame basis, for each joint, in each video sequence. As such, we compute the confidence threshold value t_i for joint i by

$$t_i = \left(\frac{1}{n} \sum_{j=1}^n c_{i,j} \right) \times 95\% \quad (1)$$

where n is the total number of frames (or postures), $c_{i,j}$ is the confidence score of joint i at frame j returned by OpenPose.

For joints identified as outliers, we calculate a revised joint position based upon the coordinates of the joint in the nearest neighbouring frames with confidence scores higher than the confidence threshold. To calculate the revised joint position we use modified Akima interpolation [4] as proposed in [34].

We then apply a moving-average filter between frames on a per joint basis to reduce jitter present in the sequence. With the filter, smoother, more reliable movements are generated for motion analysis. Empirically, we found that using a filter calculated over a 5 frame sliding window provides the best results.




Fig. 2: The 25 joint OpenPose [9] output skeleton, with associated joint reference numbers overlaid on an example input RGB image from the MINI-RGBD dataset [20].

2) **Automated Data Normalisation:** To ensure that the orientation and displacement is comparable between videos, we rotate and normalise the landmark coordinates within each frame. Using joint 9 as the root, we amend all joint coordinates so that the root is fixed at 0,0 whilst the relative distance of each of the joints is unchanged. We then calculate the rotation θ_{align}^f required at frame f to align the spinal column (i.e. the central line between joints 2 and 9) with the $y_{axis} = (0, 1)$ by

$$\theta_{align}^f = dir \times \arccos \frac{(fp_2^f - fp_9^f) \cdot y_{axis}}{\|(fp_2^f - fp_9^f)\| \|y_{axis}\|} \quad (2)$$

where fp_2^f and fp_9^f are the filtered 2D coordinates of joint 2 and 9, respectively, and $dir = sign(fp_2^f - fp_9^f) \times y_{axis}$ is used to determine the direction of the rotation (i.e. clockwise or counter-clockwise). Finally, the normalized position p of each joint can be computed by

$$p_i^f = \begin{bmatrix} \cos(\theta_{align}^f) & -\sin(\theta_{align}^f) \\ \sin(\theta_{align}^f) & \cos(\theta_{align}^f) \end{bmatrix} (fp_i^f - fp_9^f)^T \quad (3)$$

where $i \in [1, 15]$.

E. Feature Extraction

Motivated by the encouraging results obtained in our pilot study [33], we suggest that pose-based histogram features can effectively represent the motion and distribution of postures over time related to the GMA. Using the corrected and normalised pose data, we therefore propose several new pose-based features for the analysis of infant body movements and subsequent prediction of CP based upon distinct features from the GMA. Given that GM assessors typically look for specific movement patterns, we attempt to model these patterns through a set of orientation-based, displacement-based and frequency-based features. Specifically we aim to model the movements associated with the assessment criteria set out in the GMA checklist [3], and the passive movement assessment section of the Optimality Score neurological examination [19]. In this section we discuss details of the proposed new features, their relevance to the GMA, and the methods used for feature fusion.

1) **Angular Displacement (HOAD2D):** Angular Displacement represents the change in angular orientation across a specified time interval for each body part in the video. This histogram-based feature captures the distribution of the angular displacement between a predefined regular offset interval. As such, the smoothness of the body part movements can be represented. For example, a smooth

movement should be characterised by a histogram which has only a small number of bins with high values. This feature is therefore designed to help identify spasmodic, abrupt, and sporadic movements of short duration. The orientation of each joint is the 2D vector pointing from the parent joint to the child joint:

$$o_i^f = p_i^f - p_j^f \quad (4)$$

where o_i^f is the orientation (2D vector) of joint i at frame f , p_i^f and p_j^f are the 2D coordinates of joint i and j , respectively, with joint j being the parent of joint i . This method of joint orientation calculation was also used to extract the HOJO2D feature in our pilot study [33], however, here, we further compute the angular displacement θ_i^f of joint i at frame f by:

$$\theta_i^f = \arccos \frac{o_i^f \cdot o_i^{f-\Delta t}}{\|o_i^f\| \|o_i^{f-\Delta t}\|} \quad (5)$$

where Δt is a predefined regular offset interval of 10 frames, the value of which we determined empirically from our experiments. Equation 5 represents the cosine similarity of o_i^f and $o_i^{f-\Delta t}$. By applying this method, the direction (i.e. clockwise or counter-clockwise) of the angular displacement is discarded such that the feature solely focuses on the magnitude of the orientation change.

Having computed the angular displacements of every joint in the whole pose sequence, the HOAD2D for each joint is computed by quantizing the displacements into a finite number of bins. Therefore, the number of bins and the size (i.e. range of values) of each bin will significantly affect the discriminative power of the feature. We observed that most of the angular displacements are very small, while the maximum theoretical displacement is 180° . As a result, we propose using non-uniform bin sizes to better represent the distribution for the angular displacements:

$$bs_i = \frac{180^\circ}{2^{n-i}} \quad (6)$$

where bs_i is the bin size for the i -th bin and n is the total number of bins for the histogram feature. For HOAD2D, we empirically found that $n = 16$ yields the best results.

2) Relative Joint Orientation (HORJO2D): To analyse the coordination and synchronisation of different body part movements, it is important to extract features from different joints simultaneously. Inspired by [50], we propose representing the distribution of the relative orientation of the joints using a histogram-based feature. Here, the pairwise relative joint orientation is computed in a similar manner as in Equation 4:

$$o_{i \rightarrow j}^f = p_j^f - p_i^f \quad (7)$$

, although the two joints are not necessarily physically connected. In order to capture the synchronisation of different parts of the body, we compute the relative orientation for all pairs of joints.

Since the relative joint orientation has a potential range of 0° to 359° , a uniform bin size is used. In doing so, we empirically found that $n = 16$ produced the best performance. Once the individual joint histograms have been extracted, we combine these to form histogram representations for each limb prior to concatenation for classification. HORJO2D intuitively represents the body synchronisation, as such, a histogram which has only a small number of bins with high values means that the joints are moving in the same direction together.

3) Relative Joint Angular Displacement (HORJAD2D): To further capture the change in body part movement synchronisation over time, the angular displacement of the relative joint orientation is also extracted as a histogram feature. This feature is crafted to evaluate the relationship between body parts, such that whole body coordination,

dystonia, and ataxic movements can be assessed. As with extracting the HORJO2D feature, the pairwise relative joint orientation (RJO) is computed and similarly combined. We further compare the RJOs before and after the predefined frame offset interval, and angular displacement is calculated using the cosine similarity of the two RJO vectors similar to the calculation of HOAD2D, as in Equation 5.

Again, most of the angular displacements computed have a small values. As a result, we again implement a non-uniform bin size (Equation 6) to increase the discriminative power of the HORJAD2D feature. From our experiments, we empirically found that the best results were obtained when $n = 8$.

4) Fast Fourier Transform of Joint Displacement (FFT-JD):

Whilst the aforementioned histogram features represent the distribution of different kinds of spatial features at a coarse level, information pertaining to temporal ordering is discarded. Inspired by previous work analysing body movements in the frequency domain [45], we propose the FFT-JD feature. This feature contains the magnitude of each of the frequency components extracted from the motion such that the variability of the motion can be better assessed. By using the Fast Fourier Transform (FFT) we convert the extracted joint displacement signal $D_i = [\|p_i^2\|, \|p_i^3\|, \dots, \|p_i^m\|]$ of joint i from a motion with m frames to a representation in the frequency domain, allowing us to model the complexity, fluidity, and variety of the movements, whilst highlighting any repetitive, athetoid, tremulous, or myoclonic characteristics. Additionally it is reported that analyzing human motion in the frequency domain is more robust to noisy data [55], and as such helps with the task of assessing some of the smaller, more detailed movements associated with the GMA.

We extract the FFT-JD by applying FFT to the vector D_i :

$$Y_i^k = \frac{1}{m} \sum_{f=0}^{m-1} D_i e^{-\frac{j2\pi kf}{m}} \quad (8)$$

where Y_i^k is a vector which contains the magnitude of the frequency component at index k for joint i , $e^{-\frac{j2\pi kf}{m}}$ is a primitive m^{th} root of 1.

Having computed the frequency component Y_i from D_i , Y_i is partitioned into 16 bins with non-uniform bin sizes:

$$bs_{FFT-JD,b} = \begin{cases} F \frac{b^2}{n^2}, & \text{if } b = 1, \\ F \frac{b^2}{n^2} - \sum_{k=1}^{b-1} bs_{FFT-JD,k}, & \text{if } 2 \leq b < n, \\ F - \sum_{k=1}^{b-1} bs_{FFT-JD,k}, & \text{if } b = n. \end{cases} \quad (9)$$

where $bs_{FFT-JD,b}$ is the size of the b -th bin and F is the number of frequency components obtained from D_i using FFT. The last bin (i.e. $bs_{FFT-JD,n}$) will occupy the remaining space.

5) Fast Fourier Transform of Joint Orientation (FFT-JO): Similar to the FFT-JD feature we once again make use of FFT to model repetitive movements by looking into the frequency components. This feature provides information relating to the rigidity, directional variation, and range of movement associated with the infant's posture. In this case we model the repetition and frequency of similar postures from a joint orientation sequence $O_i = [o_i^1, o_i^2, \dots, o_i^m]$ for joint i using FFT as in Equation 8. The histogram-based FFT-JO is computed using the same method described in Equation 9, where the frequency components are computed using O_i instead. In this case, the bins for the lower frequency components will be smaller and can more effectively represent the more relevant low frequency components, as opposed to the high frequency motion signals which potentially contain noise.

6) Histograms of Joint Orientation (HOJO2D) and Histograms of Joint Displacement (HOJD2D): In this work we improve upon our previously reported method [33], by extracting individual HOJO2D and HOJD2D joint histograms and concatenating these

to form limb-based representations. This method means that we are able to incorporate a greater range of motion detail than the previous method, which extracts an individual per-limb histogram representation grouping several joints together.

F. Classification

Once we have extracted features, we then use Z-score to standardise the feature data h , to ensure it is on the same scale prior to further analysis by $z = \frac{h-\mu}{\sigma}$, where μ and σ are the mean and standard deviation of all samples in the training set, and z contains the normalized features. In our implementation we use Z-score as this allows our system to retain the shape properties of the original data set, with our initial classification results showing improvements using this method over min-max normalization. We then feed the features into a classification framework to obtain an overall prediction on the prevalence of CP based upon the annotations provided by the GMA assessors (Section III-B). Our classification framework consists of several machine learning algorithms, namely Logistic Regression (LR), Support Vector Machine (SVM), Decision Tree (Tree), Linear Discriminant Analysis (LDA), Ensemble of classification models (Ens), and k-Nearest Neighbour where $k = 1$ and $k = 3$. This approach allows us to generate a suitable interpretation of the strength of the features assessed, and the performance of each classifier in this task. The metrics used to evaluate the features and associated classification performance are discussed in Section IV-B.

G. Feature Fusion

In addition to evaluating the classification performance of each of the individual features, we also fuse our selected features together for further analysis. This fusion process concatenates two separate feature sets into *pose-based* features, and *velocity-based* features which are used for classification. The pose-based features represent the angular feature information extracted from the pose data, as such these representations are indicative of the overall quality of the infant posture and the predominant directions of movement. The pose-based features consist of a concatenation of HOJO2D, HOAD2D, HORJO2D, and FFT-JO. The velocity-based features represent the displacement of the joints over predefined time intervals, and as such model the speed, fluidity, coordination, and complexity of the infant movements. The velocity-based features consist of a concatenation of HOJD2D, HORJAD2D, and FFT-JD. Lastly, we fuse the pose-based feature set with the velocity-based feature set for classification. By concatenating the features through early fusion, it is expected that improved classification performance will be achieved, provided the classifier is capable of handling the higher dimensionality input data.

IV. EVALUATION

In this section we provide details of our evaluation methods. In Section IV-A we discuss the baseline methods used for comparison with our proposed method. In Section IV-B we discuss the metrics used for our comparative evaluation. In Section IV-C we discuss the experimental settings used and the rationale behind each selected test. In Section IV-D we discuss our classification results and in Section IV-E we examine the hyperparameter optimisation. Finally, in Section IV-F we include an analysis of the proposed features.

A. Baseline Methods

In order to assess the effectiveness and robustness of our system, we reimplement several video-based methods from the literature to serve as baselines for comparison, including Centroid of Motion and

Quantity of Motion [1], [2]; Absolute Motion Distance, Relative Frequency, and Magnitude of Wavelet Coefficients [53]; and Frequency Analysis [45]. We also compare our results with those reported in [56] and [57], as well as conducting our experiments using the source code provided by the authors of [38] and [51].

1) *Centroid of Motion*: The centroid of motion is the spatial centre point of the motion image which highlights the pixels with detected changes (i.e. body movement), and in our case, represents the centre point of the movements of the infant. As discussed in [2], the mean and standard deviation of centroid of motion in the x-and y-directions (CX_m , CX_{SD} , CY_m and CY_{SD}) are calculated and exported as features for classification.

2) *Quantity of Motion*: The quantity of motion is also calculated through the generation of a motion image [2]. It is the sum of all pixels with positive values from the motion image, divided by the total number of pixels contained within the image. The standard deviation (Q_{SD}) and mean (Q_m) of the quantity of motion are therefore calculated and used for classification.

3) *Cerebral Palsy Predictor*: We also include the Cerebral Palsy Predictor (CPP) feature set as discussed in [2]. This is the concatenated combination of the centroid of motion standard deviation (C_{SD}), the quantity of motion mean (Q_m), and the quantity of motion standard deviation (Q_{SD}).

4) *Absolute Motion Distance*: The Absolute Motion Distance (AMD), Relative Frequency (RF), and Magnitude of Wavelet Coefficients (MWC) methods are all based upon optical flow information, which is used for motion-based tracking. We followed the technical details presented in [53] in our implementation. The AMD is proposed as a holistic measure of activity, as it captures the absolute values of the optical flow velocities and stores them in histogram format. Since the bin size is not specified in [53], we experiment with bin sizes of 8, 16, 32, 64 and 128 in our implementation. We reported the best results obtained using the bin size of 8 in this paper.

5) *Relative Frequency*: The relative frequency (RF) of the signal represents the occurring frequencies found in the movement patterns, which are converted into a histogram representation for classification. In our implementation, we followed [53] and again used bin sizes of 8, 16, 32, 64 and 128 since this is not specified. From our experiments we empirically found a bin size of 8 returns the best results.

6) *Magnitude of Wavelet Coefficients*: The Magnitude of Wavelet Coefficients (MWC) power spectrum is used to demonstrate the variety of the observed movement at different resolution levels, providing insight into the complexity of the movement. We followed [53] to compute the associated histogram features.

7) *Frequency Analysis*: Given that normal FMs are defined as an ongoing and variable stream of movements, [45] suggest that these motions can be better studied in frequency domain. As such, Fast Fourier Transform (FFT) was used to obtain the frequency components of the motion. We followed [45] to extract the mean and standard deviation values of the Fourier coefficients in horizontal (FFT_{x_m} and $FFT_{x_{SD}}$) and vertical directions (FFT_{y_m} and $FFT_{y_{SD}}$) as features for classification, using 100 bins with non-uniform sizes as specified in the literature.

B. Evaluation Metrics

In this paper we make use of several evaluation metrics to assess the performance of each feature and the associated classifier. In our evaluation, true positive (TP) is a measure of the cases in which impaired infants are correctly classified as impaired, true negative (TN) represents unimpaired infants correctly classified as unimpaired, false positive (FP) represents unimpaired infants incorrectly classified as impaired, and false negative (FN) represents impaired infants

incorrectly classified as unimpaired. Based upon these metrics, the sensitivity (SE) is defined as the percentage of correctly identified positive classifications amongst the positive population of the dataset, the specificity (SP) is the percentage of correctly identified negative classifications amongst the negative population of the dataset, and the accuracy (AC) is defined as the percentage of all correctly classified instances. We also calculate the precision (PR), recall (RE), and F1 Score (F1). PR represents the percentage of correctly identified positive cases from all positive predictions, RE measures the correctly identified positive cases from all of the actual positive cases, and F1 is the harmonic mean of PR and RE and typically provides one of the most valuable measures of performance [13]. In addition to this, we also calculate the Matthews Correlation Coefficient (MCC) [32], which is regarded as a reliable measure where there are significant differences between the classes sizes, and as such it is widely used as a consistent performance metric on imbalanced datasets [7], [12].

C. Experimental Settings

To evaluate the generality of the classifiers using both our proposed features and the baseline approaches, a leave-one-subject-out cross validation method is used. Splitting the data in this manner ensures that the classifiers will be evaluated using unseen data, as suggested in previous works [33], [34], [56]. We further evaluate the performance of each type of feature on different classifiers (as explained in Section III-F). We report the best result for each method along with the associated classifier.

D. Experimental Results

The classification results on the MINI-RGBD and RVI-38 datasets are presented in Table I and Table II, respectively. From our evaluation, we observe that our individual features are typically performing at a similar level to that of the best features proposed in the related works. However, we note that the fusion of our proposed features is providing state-of-the-art performance across both of the datasets used in our experiments, as discussed the following subsections.

1) *Classification Performance on Individual Features*: From Table I, it can be seen that our proposed feature FFT-JO achieved a 100.00% classification accuracy on the MINI-RGBD dataset. Only one of the 20 baseline methods (BPB) evaluated achieved this perfect classification result in our tests, highlighting the remarkable performance of this new feature. Encouraging results are also obtained using our other frequency-based feature FFT-JD with 91.67% classification accuracy, 85.71% F1 score, and 81.65% MCC. This performance is higher than all of the 20 baselines in the experiments, with the exception of AMD (F1:88.89%, MCC:83.67%), RF (F1:88.89%, MCC:83.67%), FFT-Y_m (F1:85.71%, MCC:81.65%), MCI (F1:88.89%, MCC:83.67%), and STAM (F1:88.89%, MCC:83.67%). Similarly, our newly proposed HORJO2D feature achieved the same performance of 91.67% classification accuracy, 85.71% F1 score, and 81.65% MCC. However, the other relative orientation-based feature HORJAD2D is not performing as well on this dataset with 83.33% classification accuracy, an F1 Score of 75% and MCC of 62.50%. Although this performance still outperforms most of the baselines, the noticeably lower specificity (87.50%) results in a lower overall classification performance for this feature. For the angular displacement based feature HOAD2D, an average performance is obtained on this dataset with an F1 score of 66.67%, matching or outperforming 8 of 20 baselines.

For the RVI-38 dataset, we note a general drop in performance due to the challenging nature of the dataset, as shown in Table II. This is particularly noticeable in the baseline methods where we see

Feature	Class.	AC	SE	SP	F1	MCC
CX _m [2]	Ens	83.33	50.00	100.00	66.67	63.25
CX _{SD} [2]	Ens	83.33	75.00	87.50	75.00	62.50
CY _m [2]	LDA	33.33	100.00	0.00	50.00	0.00
CY _{SD} [2]	k=3	75.00	75.00	75.00	66.67	47.81
Q _m [2]	Ens	58.33	50.00	62.50	44.44	11.95
Q _{SD} [2]	k=3	75.00	75.00	75.00	66.67	47.81
CPP [1]	Tree	66.67	75.00	62.50	60.00	35.36
AMD [53]	LDA	91.67	100.00	87.50	88.89	83.67
MWC [53]	LDA	83.33	75.00	87.50	75.00	62.50
RF [53]	LDA	91.67	100.00	87.50	88.89	83.67
FFT _{x_m} [45]	Tree	83.33	75.00	87.50	75.00	62.50
FFT _{x_{SD}} [45]	Ens	58.33	50.00	62.50	44.44	11.95
FFT _{y_m} [45]	Tree	91.67	75.00	100.00	85.71	81.65
FFT _{y_{SD}} [45]	Tree	75.00	75.00	75.00	66.67	47.81
FFT _m [45]	Tree	83.33	75.00	87.50	75.00	62.50
FFT _{SD} [45]	LR	75.00	75.00	62.50	72.73	59.76
MCI [56]	n/a	91.67	100.00	87.50	88.89	83.67
CA [57]	DNN	91.67	-	-	-	-
BPB [51]	DNN	100.00	100.00	100.00	100.00	100.00
STAM [38]	DNN	91.67	100.00	87.50	88.89	83.67
HOJO2D	Ens	91.67	75.00	100.00	85.71	81.65
HOJD2D	Ens	83.33	75.00	87.50	75.00	62.50
FFT-JO	Ens	100.00	100.00	100.00	100.00	100.00
FFT-JD	LR	91.67	75.00	100.00	85.71	81.65
HOAD2D	LR	66.67	100.00	50.00	66.67	50.00
HORJO2D	LDA	91.67	75.00	100.00	85.71	81.65
HORJAD2D	Ens	83.33	75.00	87.50	75.00	62.50
Pose	Ens	100.00	100.00	100.00	100.00	100.00
Velocity	Ens	91.67	100.00	87.50	88.89	84.32
Pose & Vel.	Ens	100.00	100.00	100.00	100.00	100.00

TABLE I: Classification results using the MINI-RGBD dataset [20].

a significant drop for each baseline, with the exception of MWC (F1:75%, MCC:62.50%). This drop is most likely associated with the challenging nature of the captured data and the full frame analysis of these methods. We are seeing that methods which are able to deal with external influences better, such as the pose-based methods, are generally producing more accurate results. This is also reflected in the results produced using our proposed individual features. In this setting we note that the HOAD2D feature is performing particularly well, representing the strongest individual feature on this dataset, recording the highest F1 Score (83.33%) and MCC (80.21), along with the joint highest accuracy (94.74%) and sensitivity (83.33%). The reworked HOJO2D (F1:72.73%, MCC:68.54) and HOJD2D (F1:80.00%, MCC:79.21%) again perform well, showing the robustness of these improved features. The HORJO2D feature is also performing well, with an accuracy of 92.11%, F1 score of 76.92%, and MCC of 72.51%. We note that FFT-JD is once again performing well, with an accuracy of 92.11%, an F1 score of 72.73%, and MCC of 68.54%. We also observe that whilst FFT-JO and HORJA2D achieve a reasonable performance on the RVI-38 dataset, with 84.21% and 86.84% classification accuracy respectively, the F1 and MCC scores are lower than our other proposed features. However, whilst the scores for these features are not class leading, they are still higher than those achieved by 16 of the 18 baseline methods evaluated in this setting.

2) *Classification Performance on Feature Fusion*: We observe that on the MINI-RGB dataset the pose-based fusion is extracting the strongest feature representation and retaining the perfect classification performance provided by the FFT-JO individual feature. Our evaluation also suggests that whilst the pose-based fused features are generally outperforming the velocity-based features, fusing both of these feature sets further improves performance on both the MINI-RGB dataset (F1: 100%, MCC 100.00%) and the RVI-38 dataset (F1: 90.91%, MCC: 89.89%). We note that on the RVI-38 dataset,

Feature	Class.	AC	SE	SP	F1	MCC
CX _m [2]	LR	50.00	83.33	43.75	34.48	20.20
CX _{SD} [2]	Ens	68.42	33.33	75.00	25.00	6.90
CY _m [2]	k=3	84.21	50.00	90.63	50.00	40.63
CY _{SD} [2]	LR	63.16	66.67	65.63	36.36	21.54
Q _m [2]	LR	52.63	50.00	53.13	25.00	2.28
Q _{SD} [2]	k=1	86.84	50.00	93.75	54.44	47.19
CPP [1]	Ens	84.21	50.00	90.63	50.00	40.63
AMD [53]	LDA	83.33	50.00	100.00	66.67	63.25
MWC [53]	Tree	83.33	75.00	87.50	75.00	62.50
RF [53]	LDA	84.21	66.67	87.50	57.14	48.45
FFT-X _m [45]	Ens	84.21	50.00	90.63	50.00	40.63
FFT-X _{SD} [45]	LR	63.16	66.67	62.50	36.36	21.54
FFT-Y _m [45]	k=1	81.58	33.33	90.63	36.36	25.84
FFT-Y _{SD} [45]	LDA	55.26	50.00	56.25	26.09	4.58
FFT _m [45]	Tree	84.21	50.00	90.63	50.00	40.63
FFT _{SD} [45]	LR	42.11	66.67	37.50	26.67	3.15
BPB [51]	DNN	84.21	33.33	93.75	40.00	32.18
STAM [38]	DNN	81.58	33.33	90.63	36.36	25.85
HOJO2D	Ens	92.11	66.67	96.88	72.73	68.54
HOJD2D	Ens	94.74	66.67	100.00	80.00	79.21
FFT-JO	LDA	84.21	83.33	84.38	62.50	56.07
FFT-JD	Ens	92.11	66.67	96.88	72.73	68.54
HOAD2D	Ens	94.74	83.33	96.88	83.33	80.21
HORJO2D	Tree	92.11	83.33	93.75	76.92	72.51
HORJAD2D	LR	86.84	66.67	90.63	61.54	53.89
Pose	Ens	94.74	83.33	96.88	83.33	80.21
Velocity	Ens	94.74	66.67	100.00	80.00	79.21
Pose & Vel.	Ens	97.37	83.33	100.00	90.91	89.89

TABLE II: Classification results using the RVI-38 dataset.

the strengths from each feature set combine to provide this improved overall classification performance, with the higher sensitivity found in the pose-based features (83.33%) and the higher specificity found in the velocity-based features (100.00%) directly translating to the concatenated fusion of these feature sets. This observation aligns well with our feature design, which looks to incorporate the combined positional, directional, postural, and transitory information specified in the GMA guidelines. We also note that on the fused features we are seeing a consistently high performance using the Ensemble classifier, with the best results obtained on both datasets for all fused feature sets using this classification method. The evaluation metrics also highlight the robustness of the proposed feature fusion method, given that only one positive sample video was misclassified across both datasets.

E. Hyperparameter Optimisation

To refine the framework performance, we further investigate hyperparameter optimisation using Bayesian Optimisation [52]. Using this method, we evaluate the results of optimisation in an informed manner, by tuning the learning rate, the number of learning cycles, the minimum observations per leaf, and the maximum number of branch nodes, to minimize the cross-validation loss of the classifier.

We present several plot representations of our Bayesian Optimisation in the supplementary information. In the plots, the number of function evaluations relates to the iteration number of the objective function, the min objective is the minimum value that the objective function has reached up to the current iteration, and the estimated minimum objectives are the mean values of the posterior distribution of the Gaussian process model of the objective function [25]. We also map the hyper-parameter variables to the classification performance metrics to determine the optimal hyper-parameters. In this setting, we found the optimal hyper-parameters to be: 0.1045800 learning rate, 11 learning cycles, 1 minimum observation per leaf, and 32 split branch nodes, providing an objective function of 0.026316 and an accuracy of 100% on the MINI-RGBD dataset, and 97.37% on the RVI-38 dataset, per our reported results.

F. Feature Analysis

To further evaluate the discriminative power of the newly proposed features, chi-square tests are used for testing if the predictor variables (i.e. the multi-dimensional features proposed in this work) and the response variable (i.e. the label of each video) are related. In particular, such tests have been widely used for feature selection and are thus able to reflect the quality of the features we propose. We further conducted the chi-square tests on both the MINI-RGBD and the RVI-38 datasets to highlight the differences between these two datasets. Specifically, the p-value for each predictor variable is calculated and the median values are reported in Table III. Here, we consider the predictor variables as *significant predictor (sp)* if $p < 0.05$. Since the features used in the experiments are mostly multi-dimensional, the dimensionality (dim.), number of *sp* (# *sp*) and percentage of *sp* (% *sp*) are also reported in Table III. We also include the top performing baselines from our experiments (i.e. AMD, MWC and RF proposed in [53]) in this analysis to further highlight the effectiveness of the proposed features.

From Table III, it can be seen that HOJD2D and HORJO2D achieve the highest % *sp* on MINI-RGBD and RVI-38, respectively. On MINI-RGBD, our proposed features are having the same or lower median p-value when compared with AMD, MWC and RF. On RVI-38, HORJO2D and HORJAD2D achieved significantly lower median p-values than other top performers. Also, it can be seen that AMD performed better than some of our proposed features in terms of the median p-values. However, the results from the 2 datasets also indicate the robustness of our proposed features since more consistent results are obtained by using our features.

Feature	Dim.	MINI-RGBD			RVI-38		
		Median	# <i>sp</i>	% <i>sp</i>	Median	# <i>sp</i>	% <i>sp</i>
AMD [53]	8	0.1020	0	0%	0.1280	2	25%
MWC [53]	18	0.6456	0	0%	0.8931	6	33.33%
RF [53]	16	0.1020	4	25.00%	0.2388	6	37.50%
HOJO2D	64	0.1020	18	28.13%	0.1370	10	31.25%
HOJD2D	128	0.0860	52	40.63%	0.1849	45	35.16%
FFT-JO	48	0.1020	8	16.67%	0.1849	18	32.14%
FFT-JD	64	0.1020	23	35.94%	0.1849	18	28.13%
HOAD2D	64	0.0847	17	26.56%	0.1849	22	34.38%
HORJO2D	32	0.1020	5	15.63%	0.050	14	43.75%
HORJAD2D	32	0.1020	9	28.13%	0.071	8	25%
Pose & Vel.	792	0.1020	228	28.78%	0.1849	260	32.83%

TABLE III: The p-values of the features computed from chi-square tests on the MINI-RGBD and RVI-38 datasets.

To better visualize the distribution of the p-values, boxplots of the p-values of different features on MINI-RGBD and RVI-38 are shown in Figure 3a and 3b, respectively. In particular, the maximum, minimum, first quartile, third quartile and median (red line) values are illustrated in the figures. It can be seen that the majority (from the first to third quartile) of the predictor variables in our proposed features are having a small range with low p-values. This indicates the majority of our proposed features are of higher importance and quality when compared with the AMD, MWC and RF.

V. DISCUSSION

As discussed in Section II-C, we suggest that pose-estimation based approaches provide several advantages over the previously proposed methods in data acquisition and analysis. Our pose and velocity-based method is simpler to understand, retains understandable information, and has less parameters to tune than the related methods, making it more accessible in a clinical setting. We also suggest that, due to the relative assessment of joint motion, our framework is better able to

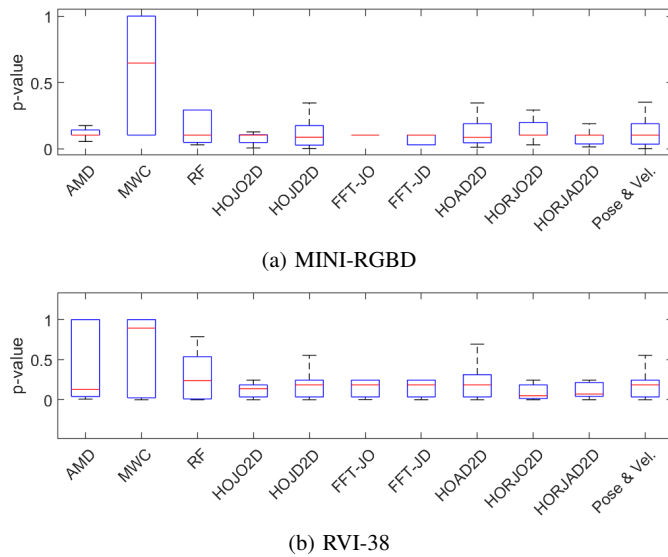


Fig. 3: Boxplots of the p-values of different features on each dataset.

deal with camera movement, changes in resolution, variable infant sizes, and larger motion changes between frames.

However, whilst the results we have achieved in this paper are encouraging, there are several key areas which we would like to address in future works. Our method is heavily dependant upon the quality of the extracted joint positions and as such the pose estimation method. In this paper we made use of OpenPose [9], however this framework is trained using adult data and as such the extracted pose has the potential to have some inconsistencies. Although the results demonstrated here suggest that these inconsistencies are largely dealt with by our qualitative assessment and automated pre-processing techniques, we would like to enhance the pose-estimation by integrating domain adaptation to make the framework more specific to infant body dimensions and posture. We would also like to further explore the temporal aspect of the GMA by extracting spatio-temporal features for further analysis.

Our results on the RVI-38 dataset in particular represent a particularly robust performance given the difficulty of the associated dataset, with only one misclassified video. In this case the misclassified video was one of the positive samples which, in practice presents a greater issue than a misclassified negative sample. In future works we will look to methods by which we can improve the sensitivity of the classification performance, perhaps through data augmentation to help deal with the moderate class imbalance found in the datasets.

VI. CONCLUSION

In this paper we have proposed and evaluated several new interpretable motion features relating directly to the criteria associated with the GMA checklist [3] and optimality score [19]. Additionally, we have fused these features together to produce a more robust representation of infant body movement for classification. We compared these features with several other methods from the literature by re-implementing them for assessment using shared datasets. The datasets used in the study consist of the publicly available MINI-RGBD dataset [20], and the RVI-38 dataset, a challenging new video dataset gathered as part of routine clinical care. We find that our proposed fused features achieve state-of-the-art performance across both datasets whilst retaining clinical interpretability. Additionally, we also suggest that by utilising pose-based features, we make the likelihood of collaborative working within the healthcare domain

more viable, due to the inherently anonymised and unidentifiable patient data. As such, we make the pose data, and our feature extraction and classification code available to the community to further improve related research.

In future works we hope to explore improving the interpretability of CP prediction models by allowing clinicians to form part of the feedback loop. We will also explore methods of extracting more detailed information about the movement characteristics, through temporal analysis and improved annotation.

ACKNOWLEDGMENT

This work was supported in part by the Royal Society (Ref: IES\R2\181024 and IES\R1\191147). The authors would like to thank the extended team at the RVI, and in particular, Patricia Dulson for her work annotating and transferring the video data for analysis.

REFERENCES

- [1] Lars Adde, Jorunn L. Helbostad, Alexander Refsum Jensenius, Gunnar Taraldsen, and Ragnhild Støen. Using computer-based video analysis in the study of fidgety movements. *Early Human Development*, 2009.
- [2] Lars Adde, Hong Yang, Rannei Sæther, Alexander Refsum Jensenius, Espen Ihlen, Jia yan Cao, and Ragnhild Støen. Characteristics of general movements in preterm infants assessed by computer-based video analysis. *Physiotherapy Theory and Practice*, 34(4):286–292, 2018.
- [3] Carolina Yuri Panvequio Aizawa, Christa Einspieler, Fernanda Francisco Genovesi, Silvia Maria Ibbidi, and Renata Hydee Hasue. The general movement checklist: A guide to the assessment of general movements during preterm and term age. *Jornal de Pediatria*, 2020.
- [4] Hiroshi Akima. A new method of interpolation and smooth curve fitting based on local procedures. *J. ACM*, 17(4):589–602, October 1970.
- [5] Martin Bax, Murray Goldstein, Peter Rosenbaum, Alan Leviton, Nigel Paneth, Bernard Dan, Bo Jacobsson, and Diane Damiano. Proposed definition and classification of cerebral palsy, april 2005. *Developmental Medicine and Child Neurology*, 47(8):571–576, 2005.
- [6] Margot Bosanquet, Lisa Copeland, Robert Ware, and Roslyn Boyd. A systematic review of tests to predict cerebral palsy in young children. *Developmental Medicine & Child Neurology*, 55(5):418–426, 2013.
- [7] Sabri Boughorbel, Fethi Jarray, and Mohammed El-Anbari. Optimal classifier for imbalanced data using matthews correlation coefficient metric. *PLOS ONE*, 12(6):1–17, 06 2017.
- [8] Thomas Brox, Christoph Bregler, and Jitendra Malik. Large displacement optical flow. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 41–48, 2009.
- [9] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2019.
- [10] Claire Chambers, Nidhi Seethapathi, Rachit Saluja, Helen Loeb, Samuel R. Pierce, Daniel K. Bogen, Laura Prosser, Michelle J. Johnson, and Konrad P. Kording. Computer vision to automatically assess infant neuromotor risk. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 28(11):2431–2442, nov 2020.
- [11] Yucheng Chen, Yingli Tian, and Mingyi He. Monocular human pose estimation: A survey of deep learning-based methods. *Computer Vision and Image Understanding*, 192(December 2019):102897, 2020.
- [12] Davide Chicco and Giuseppe Jurman. The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genomics*, 21(1):1–13, 2020.
- [13] Ryan Cunningham, María B. Sánchez, Penelope B. Butler, Matthew J. Southgate, and Ian D. Loram. Fully automated image-based estimation of postural point-features in children with cerebral palsy using deep learning. *Royal Society Open Science*, 6(11):191011, 2019.
- [14] ES Draper, ID Gallimore, JJ Kurinczuk, PW Smith, T Boby, LK Smith, and BN Manktelow. *MBRRACE-UK – Perinatal Mortality Surveillance Report 2017*. The Infant Mortality and Morbidity Studies, 2018.
- [15] Christa Einspieler and Heinz F. R. Prechtl. *Prechtl's method on the qualitative assessment of general movements in preterm, term, and young infants*. Mac Keith Press, 2004.
- [16] Christa Einspieler, Heinz F.R. Prechtl, Fabrizio Ferrari, Giovanni Cioni, and Arend F. Bos. The qualitative assessment of general movements in preterm, term and young infants. *Early Human Development*, 50(1):47 – 60, 1997. Spontaneous Motor Activity as a Diagnostic Tool Functional Assessment of the Young Nervous System.

- [17] National Institute for Health and Care Excellence. NICE seeks to improve diagnosis and treatment of cerebral palsy, Jan 2017.
- [18] Daniel Groos, Heri Ramampiaro, and Espen Af. EfficientPose : Scalable single-person pose estimation. *Applied Intelligence*, 2021.
- [19] Leena Haataja, Eugenio Mercuri, Rivka Regev, Frances Cowan, Mary Rutherford, Victor Dubowitz, and Lilly Dubowitz. Optimality score for the neurologic examination of the infant at 12 and 18 months of age. *The Journal of Pediatrics*, 135(2):153–161, 1999.
- [20] Nikolas Hesse, Christoph Bodensteiner, Michael Arens, Ulrich G. Hoffmann, Raphael Weinberger, and A. Sebastian Schroeder. Computer vision for medical infant motion analysis: State of the art and RGB-D data set. In *Computer Vision - ECCV 2018 Workshops*. Springer International Publishing, 2018.
- [21] Andreas Holzinger, Chris Biemann, Constantinos S. Pattichis, and Douglas B. Kell. What do we need to build explainable ai systems for the medical domain?, 2017.
- [22] Ying Huang, Hubert P. H. Shum, Edmond S. L. Ho, and Nauman Aslam. High-speed multi-person pose estimation with deep feature transfer. *Computer Vision and Image Understanding*, 197-198:103010, 2020.
- [23] Espen Ihlen, Ragnhild Støen, Lynn Boswell, Raye-Ann Deregnier, Toril Fjærtøft, Deborah Gaebler-Spira, and Lars Adde. Machine learning of infant spontaneous movements for the early prediction of cerebral palsy: A multi-site cohort study. *Journal of Clinical Medicine*, 9, 12 2019.
- [24] Fatima Yousif Ismail, Ali Fatemi, and Michael V. Johnston. Cerebral plasticity: Windows of opportunity in the developing brain. *European Journal of Paediatric Neurology*, 21(1):23–48, 2017. *Advances in Neuromodulation in Children*.
- [25] Anezka Kazikova, Michal Pluhacek, and Roman Senkerik. How does the number of objective function evaluations impact our understanding of metaheuristics behavior? *IEEE Access*, 9:44032–44048, 2021.
- [26] Amanda K.L. Kwong, Tara L. Fitzgerald, Lex W. Doyle, Jeanie L.Y. Cheong, and Alicia J. Spittle. Predictive validity of spontaneous early infant movement for later cerebral palsy: a systematic review. *Developmental Medicine and Child Neurology*, 60(5):480–489, 2018.
- [27] Min Li, Fan Wei, Yu Li, Sicong Zhang, and Guanghua Xu. Three-dimensional pose estimation of infants using data from a kinect sensor with low training cost. *IEEE Sensors Journal*, 21(5):6904–6913, 2021.
- [28] Helen MacTier, Sarah Elizabeth Bates, Tracey Johnston, Caroline Lee-Davey, Neil Marlow, Kate Mulley, Lucy K. Smith, Meekai To, and Dominic Wilkinson. Perinatal management of extreme preterm birth before 27 weeks of gestation: A framework for practice. *Arch. of Disease in Childhood: Fetal and Neonatal Edition*, 105(3):F232–F239, 2020.
- [29] Nathalie Maitre. Skepticism, cerebral palsy, and the General Movements Assessment. *Developmental Medicine and Child Neurology*, 2018.
- [30] Viviana Marchi, Anna Hakala, Andrew Knight, Federica D’Acunto, Maria Luisa Scattoni, Andrea Guzzetta, and Sampsa Vanhatalo. Automated pose estimation captures key aspects of general movements at eight to 17 weeks from conventional videos. *Acta Paediatrica*, 108(10):1817–1824, 2019.
- [31] Claire Marcroft, Aftab Khan, Nicholas D. Embleton, Michael Trenell, and Thomas Plötz. Movement recognition technology as a method of assessing spontaneous general movements in high risk infants. *Frontiers in Neurology*, 6(JAN):284, 2015.
- [32] B.W. Matthews. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA) - Protein Structure*, 405(2):442–451, 1975.
- [33] K. D. McCay, E. S. L. Ho, C. Marcroft, and N. D. Embleton. Establishing pose based features using histograms for the detection of abnormal infant movements. In *EMBC 2019*, pages 5469–5472, July 2019.
- [34] K. D. McCay, E. S. L. Ho, H. P. H. Shum, G. Fehringer, C. Marcroft, and N. D. Embleton. Abnormal infant movements classification with deep learning on pose-based features. *IEEE Access*, 2020.
- [35] Kevin D. McCay, Edmond S. L. Ho, Dimitrios Sakkos, Wai Lok Woo, Claire Marcroft, Patricia Dulson, and Nicholas D. Embleton. Towards explainable abnormal infant movements identification: A body-part based prediction and visualisation framework, 2021.
- [36] Sara Moccia, Lucia Migliorelli, Virgilio Carnielli, and Emanuele Frontoni. Preterm Infants’ Pose Estimation With Spatio-Temporal Features. *IEEE transactions on bio-medical engineering*, 67(8):2370–2380, 2020.
- [37] Catherine Morgan, Iona Novak, Russell C. Dale, Andrea Guzzetta, and Nadia Badawi. Single blind randomised controlled trial of GAME (Goals - Activity - Motor Enrichment) in infants at high risk of cerebral palsy. *Research in Developmental Disabilities*, 55:256–267, 2016.
- [38] Binh Nguyen-Thai, Vuong Le, Catherine Morgan, Nadia Badawi, Truyen Tran, and Svetha Venkatesh. A spatio-temporal attention-based model for infant movement assessment from videos. *IEEE Journal of Biomedical and Health Informatics*, 25(10):3911–3920, 2021.
- [39] Joy Olsen, Peter Marschik, and Alicia Spittle. Do fidgety general movements predict cerebral palsy and cognitive outcome in clinical follow-up of very preterm infants?. 2018.
- [40] Silvia Orlandi, Kamini Raghuram, Corinna R. Smith, David Mansueto, Paige Church, Vibhuti Shah, Maureen Luther, and Tom Chau. Detection of Atypical and Typical Infant Movements using Computer-based Video Analysis. in *EMBC*, 2018:3598–3601, 2018.
- [41] Maryam Oskoui, Franzina Coutinho, Jonathan Dykeman, Nathalie Jette, and Tamara Pringsheim. An update on the prevalence of cerebral palsy: A systematic review and meta-analysis. *Developmental medicine and child neurology*, 55, 01 2013.
- [42] Colleen Peyton and Christa Einspieler. General Movements: A Behavioral Biomarker of Later Motor and Cognitive Dysfunction in NICU Graduates. *Pediatric Annals*, 47(4):e159–e164, 2018.
- [43] Kamini Raghuram, Silvia Orlandi, Paige Church, Tom Chau, Elizabeth Uleryk, Petros Pechlivanoglou, and Vibhuti Shah. Automated movement recognition to predict motor impairment in high-risk infants: a systematic review of diagnostic test accuracy and meta-analysis. *Developmental Medicine and Child Neurology*, 63(6):637–648, 2021.
- [44] Kamini Raghuram, Silvia Orlandi, Vibhuti Shah, Tom Chau, Maureen Luther, Rudaina Banihani, and Paige Church. Automated movement analysis to predict motor impairment in preterm infants: a retrospective study. *Journal of Perinatology*, 39(10):1362–1369, 2019.
- [45] H. Rahmati, H. Martens, O. M. Aamo, Ø. Stavadahl, R. Støen, and L. Adde. Frequency analysis and feature reduction method for prediction of cerebral palsy in young infants. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 24(11):1225–1234, Nov 2016.
- [46] Simon Reich, Dajie Zhang, Tomas Kulvicus, Sven Böltje, Karin Nielsen-Saines, Florian B. Pokorny, Robert Peharz, Luise Poustka, Florentin Wörgötter, Christa Einspieler, and Peter Marschik. Novel AI approach to classify infant motor functions. *Scientific Reports*, 11(1):1–13, 2021.
- [47] Susan M Reid, Elaine Meehan, Sarah McIntyre, Shona Goldsmith, Nadia Badawi, Dinah S Reddihough, and AusACPDM. Temporal trends in cerebral palsy by impairment severity and birth gestation. *Developmental Medicine & Child Neurology*, 58(S2):25–35, 2016.
- [48] Eileen Ricci, Christa Einspieler, and Alexa K. Craig. Feasibility of Using the General Movements Assessment of Infants in the United States. *Physical and Occupational Therapy in Pediatrics*, 2018.
- [49] Peter Rosenbaum, Nigel Paneth, Alan Leviton, Murray Goldstein, Martin Bax, Diane Damiano, Bernard Dan, and Bo Jacobsson. A report: The definition and classification of cerebral palsy april 2006. *Developmental medicine and child neurology. Supplement*, 109:8–14, 03 2007.
- [50] W. Rueangsirarak, J. Zhang, N. Aslam, E. S. L. Ho, and H. P. H. Shum. Automatic musculoskeletal and neurological disorder diagnosis with relative joint displacement from human gait. *IEEE Trans. on Neural Systems and Rehabilitation Engineering*, 26(12):2387–2396, Dec 2018.
- [51] Dimitrios Sakkos, Kevin D. McCay, Claire Marcroft, Nicholas D. Embleton, Samiran Chattopadhyay, and Edmond S. L. Ho. Identification of abnormal movements in infants: A deep neural network for body part-based prediction of cerebral palsy. *IEEE Access*, 2021.
- [52] Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms, 2012.
- [53] Annette Stahl, Christian Schellewald, Øyvind Stavadahl, Ole Morten Aamo, Lars Adde, and Harald Kirkerød. An optical flow-based method to predict infantile cerebral palsy. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(4):605–614, 2012.
- [54] Pavan Turaga, Rama Chellappa, and Ashok Veeraraghavan. Advances in video-based human activity analysis: Challenges and approaches. In Marvin V. Zelkowitz, editor, *Advances in Computers*, volume 80 of *Advances in Computers*, pages 237–290. Elsevier, 2010.
- [55] Jiang Wang, Zicheng Liu, Ying Wu, and Junsong Yuan. Learning actionlet ensemble for 3d human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(5):914–927, 2014.
- [56] Q. Wu, G. Xu, F. Wei, L. Chen, and S. Zhang. Rgb-d videos-based early prediction of infant cerebral palsy via general movements complexity. *IEEE Access*, 9:42314–42324, 2021.
- [57] Manli Zhu, Qianhui Men, Edmond S. L. Ho, Howard Leung, and Hubert P. H. Shum. Interpreting Deep Learning based Cerebral Palsy Prediction with Channel Attention. pages 1–4, 2021.
- [58] Dragan Zlatanović, Hristina Colovic, Zivkovic Vesna, Mirjana Kocic, Anita Stanković, Jelena Vučić, Nikola Bojovic, Maja Raičević, Mladjan Golubovic, Ljubomir Dinic, and Tamara Stanković. The importance of the prechtl method for ultra-early prediction of neurological abnormalities in newborns and infants. *AMM*, pages 111–115, 09 2019.