

Dissertation Project

Data Science for Health and Social Care (M. Sc.)

© Benjamin Gross

23.04.2026 10:45

This dissertation explores the application of AI-based posture recognition using Convolutional Neural Networks (CNNs) to detect and analyse habitual patterns of human movement and posture. The work focuses on transfer learning with TensorFlow Lite and Google MediaPipe for identifying head-neck-torso imbalances.

Table of contents

1	MSc in Data Science for Health and Social Care	3
1.1	Dissertation Title	3
2	1. Introduction	3
2.1	1.1 Background and rationale	3
2.2	1.2 Research question, aim, and objectives	4
2.3	1.3 Dissertation structure	4
3	2. Literature Review	4
4	3. Methodology	5
4.1	3.1 Study design and setting	5
4.2	3.2 Data source and characteristics	5
4.3	3.3 Variables, features, and preprocessing	6
4.4	3.4 Model development and transfer learning	6
4.5	3.5 Analysis and performance evaluation	7
5	4. Implementation	7
5.1	4.1 Experimental setup	7
5.2	4.2 Training pipeline	7
5.3	4.3 Inference pipeline and deployment artefacts	7
6	5. Results	7
6.1	5.1 Dataset and preprocessing outcomes	7
6.2	5.2 Model performance	7
6.3	5.3 Error analysis	8
7	6. Discussion	8
7.1	6.1 Interpretation of findings	8
7.2	6.2 Strengths and limitations	8
7.3	6.3 Clinical and practical implications	8

8	7. Ethics, Information Governance, and Data Management	8
9	8. Impact and Dissemination	9
10	9. Conclusion and Future Work	9
11	Appendix A. Figures	9
12	References	10

1 MSc in Data Science for Health and Social Care

Examination Number: B248593 First reviewer: Dr. Ahmar Shah Second reviewer: John Wilson

1.1 Dissertation Title

Detecting Movement Imbalances using a custom TensorFlow Lite Model utilizing Google MediaPipe framework.

2 1. Introduction

Healthcare is increasingly recognizing the importance of movement quality and postural balance in preventing injury, managing chronic conditions, and optimizing physical performance. However, traditional methods for assessing posture and movement often rely on subjective clinical observation or expensive motion capture systems, limiting their accessibility and consistency. Advances in artificial intelligence (AI) and computer vision offer promising avenues for developing objective, scalable tools to analyze human movement patterns. This dissertation explores the application of Convolutional Neural Networks (CNNs) within the Google MediaPipe framework to detect and analyze habitual patterns of imbalance in head-neck-torso alignment, with a focus on transfer learning to adapt pre-trained models for clinical relevance.

2.1 1.1 Background and rationale

Human posture and coordination are central to both health and performance. Impairments such as muscular tension, poor motor control, or postural imbalance contribute significantly to discomfort, injury, and reduced physical performance in daily and professional life [1]. Advanced motor skills rely on precise kinesthetic differentiation, including fine motor control, proprioceptive sensitivity, accurate force regulation, coordinated timing, and the economy of movement. These elements form the basis of effective movement in sports, rehabilitation, and everyday function [2]. Recent research highlights that chronic low back pain and other movement system impairments are closely associated with sustained postural deviations and maladaptive movement patterns that often develop gradually and may not be apparent during isolated clinical assessments [3], [4].

Traditional approaches to assessing posture and coordination rely on subjective evaluations, clinical observation, expensive motion analysis systems, or extensive practitioner expertise. While commercial software for posture assessment exists, these solutions often require professional oversight and remain limited in their ability to capture subtle or habitual patterns of imbalance consistently [5]. This highlights a clear gap between the qualitative nature of clinical assessments and the potential for more objective, data-driven methods enabled by modern machine learning techniques.

Recent advances in computer vision and artificial intelligence provide opportunities to bridge these limitations. Convolutional Neural Networks (CNNs) have demonstrated high accuracy in identifying patterns within image and video data and are increasingly applied in gait analysis, posture detection, and motion tracking [6]. Among available frameworks, the Google MediaPipe framework brings a lightweight approach that can be run in the browser or on mobile devices, utilizing a statistical 3D human shape modeling pipeline offered as a trainable and modular deep

learning framework. It can accurately represent full-body details to detect pose and posture [7], [8].

From my perspective as a physiotherapist with a professional focus on movement quality, these gaps are not purely academic but deeply practical. It required years of training and clinical experience to develop the ability to assess movement with precision. The possibility of using lightweight AI to complement clinical expertise, enhance rehabilitation outcomes, and support performance improvement is therefore highly compelling.

2.2 1.2 Research question, aim, and objectives

Research question

How can transfer learning on a pre-existing TensorFlow Lite model for image classification be used to detect and analyze posture and movement imbalances?

Aim

Finetuning and evaluation of a machine learning model for posture and movement quality assessment that detects and interprets imbalances in the head, neck and upper body parts to make it suitable for clinical context.

Objectives

1. To finetune an existing image classification model (Google MediaPipe) for detecting posture and movement imbalances in clinical contexts.
2. To evaluate the performance of the finetuned image classification model.
3. To detect and analyze indicators of upper body imbalances, including the alignment angles of the head, neck and torso, as well as asymmetries in muscle contours.

2.3 1.3 Dissertation structure

Lore Ypsum placeholder text: This section will briefly explain the purpose and content of each chapter in the final dissertation.

3 2. Literature Review

However, recent validation studies comparing MediaPipe with gold-standard motion capture systems reveal measurable discrepancies in landmark accuracy, particularly for subtle postural features, highlighting the need for domain-specific refinement for clinical use [9]. More broadly, current pose estimation models remain limited by sensitivity to noise and 2D–3D estimation gaps, reducing their reliability for detecting the subtle asymmetries critical to movement imbalance analysis [9], [10]. Fairness and representativeness in pose estimation systems are also important concerns: current Pose Estimation Models (PEMs) exhibit performance disparities across demographic groups, with reduced accuracy for underrepresented populations in training datasets [10]. Available datasets often focus on males and light-skinned subjects in age groups 19–50 and under-represent females, older adults, and people with darker skin tones, creating the possibility of systematic errors in movement analysis.

Consequently, while MediaPipe Pose serves as a robust foundation for pose detection, it does not inherently capture the fine-grained symmetry metrics, sagittal and coronal balance indicators, or compensatory movement patterns that are essential for clinical postural and movement analysis. AI-based markerless systems nonetheless represent an important step towards closing the gap

between theoretical research and practical implementation in remote clinical assessment [11], [12].

This project uses transfer learning [9] to adapt a pre-trained Google MediaPipe CNN model, originally trained on extensive ImageNet data, to a new Kaggle dataset with more relevant and higher-quality images for posture imbalance detection. By utilizing pre-learned feature representations, the model can achieve faster convergence and improved accuracy even with a comparatively smaller or domain-specific dataset.

Lore Ypsum placeholder text: This section will be expanded with a deeper synthesis of related work, critical comparison of methods, and identification of research gaps.

4 3. Methodology

In this work, we will apply a quantitative, observational design to explore the feasibility of using transfer learning on a pre-trained TensorFlow Lite model for posture and movement imbalance detection. The study will involve secondary data analysis of an existing image dataset curated for posture analysis, with a focus on evaluating the performance of a finetuned CNN model within the Google MediaPipe framework.

4.1 3.1 Study design and setting

This study is a secondary data analysis exploring postural and movement imbalance detection using artificial intelligence. The project applies a quantitative, observational design, analysing existing image data through supervised machine learning, specifically Convolutional Neural Networks (CNNs) implemented via TensorFlow Lite within the Google MediaPipe framework. The aim is to evaluate whether lightweight AI models can detect subtle patterns of imbalance in posture and coordination and to adapt them for potential clinical use through transfer learning on a posture-specific dataset.

The study will be conducted remotely using a secondary dataset. All data analysis will be performed on secured infrastructure provided by the University of Edinburgh. No new data collection or participant interaction is planned.

4.2 3.2 Data source and characteristics

This study will use secondary image data from the openly available “Posture Keypoints Detection – Photos & Labels” dataset on Kaggle. The dataset consists of 300 images (approximately 250 for training and 50 for validation) in static side view with image resolutions of up to around 1.6 MP. The images are accompanied by annotations to key points in the YOLO pose format, a standardized labeling structure that specifies the pixel coordinates of essential anatomical landmarks [12].

Unlike large-scale, general-purpose image datasets, this Kaggle dataset is specifically curated for posture analysis. It focuses on high-quality, side-view images suitable for assessing sagittal-plane alignment and upper-body posture. The relatively small dataset size and potential demographic limitations may constrain generalizability and fairness.

4.3 3.3 Variables, features, and preprocessing

Primary outcomes

- Head-neck-torso alignment: Angular relationships and relative positioning of the head, cervical spine, and torso.
- Postural coordination: Upper-body skeletal alignment patterns derived from keypoints and CNN feature maps.
- Movement imbalance indicators: Detection of asymmetries, misalignments, or inefficient posture as inferred from pose landmarks.

Exposures/independent variables

- Image category/context: Static posture images in standardized side view.
- Camera/view angle: Minor variation in pose estimation accuracy depending on exact lateral positioning.

Covariates/control features

- Image characteristics: Resolution, background, and lighting, which may influence model performance.
- Posture variation: Degree of anterior head translation, kyphosis, and other sagittal-plane deviations.

Extracted features

- Skeletal keypoints: Annotated landmarks in YOLO pose format and MediaPipe Pose outputs.
- Kinematic proxies: Angular relationships between key joints and segments.
- Postural imbalance markers: Anterior head translation, sagittal imbalance, and deviations from expected alignment patterns.

All selected images from the Kaggle dataset will undergo preprocessing prior to model training, including resizing, normalization, and verification of annotation integrity. Where appropriate, data augmentation (e.g., small rotations, cropping, brightness adjustments) will be applied to improve generalization while preserving clinically meaningful posture characteristics.

4.4 3.4 Model development and transfer learning

This study employs a transfer learning approach to develop a domain-specific TensorFlow Lite model suitable for clinical posture and movement analysis. Transfer learning enables the adaptation of pre-trained CNN architectures—such as those used within MediaPipe and trained on large-scale datasets like ImageNet—to the focused task of detecting head-neck-torso imbalances. In practice, base model weights are kept frozen or partially frozen, while classification and selected high-level layers are retrained using the Kaggle posture dataset.

Deploying the resulting model as a quantized TensorFlow Lite artefact supports real-time, on-device inference on mobile and edge devices, reducing latency and preserving user privacy by processing data locally.

4.5 3.5 Analysis and performance evaluation

The analysis will proceed in two phases:

1. **Descriptive analysis** of dataset composition, annotation completeness, and posture-related metrics.
2. **AI-based classification and pattern recognition** using transfer learning and supervised deep learning methods.

Model performance will be evaluated using accuracy as the primary metric and, where feasible, precision, recall, F1-score, and confusion matrices. Robustness and generalization will be explored through limited distribution-shift experiments (e.g., withheld images with slight background or lighting variation). All hyperparameters, data splits, and preprocessing steps will be documented to ensure reproducibility.

5 4. Implementation

Lore Ypsum placeholder text: This chapter will detail the final implementation of the model, including the training pipeline, inference pipeline, and deployment artefacts.

5.1 4.1 Experimental setup

Lore Ypsum placeholder text: This section will describe the final implementation environment, hardware/software configuration, and reproducible execution setup.

5.2 4.2 Training pipeline

Lore Ypsum placeholder text: This section will document the final end-to-end training workflow, including dataset split strategy, augmentation settings, and training schedules.

5.3 4.3 Inference pipeline and deployment artefacts

Lore Ypsum placeholder text: This section will present the final inference pipeline, exported TensorFlow Lite artefacts, and deployment considerations.

6 5. Results

6.1 5.1 Dataset and preprocessing outcomes

Lore Ypsum placeholder text: This section will report final dataset counts, exclusions, and preprocessing outcomes.

6.2 5.2 Model performance

Lore Ypsum placeholder text: This section will report validation metrics, confusion matrices, and benchmark comparisons.

6.3 5.3 Error analysis

Lore Ypsum placeholder text: This section will analyse misclassifications, edge cases, and expected failure modes.

7 6. Discussion

7.1 6.1 Interpretation of findings

Lore Ypsum placeholder text: This section will interpret the main findings in relation to posture science and AI-based movement analysis.

7.2 6.2 Strengths and limitations

Lore Ypsum placeholder text: This section will critically discuss methodological strengths, data limitations, and generalisability constraints.

7.3 6.3 Clinical and practical implications

Lore Ypsum placeholder text: This section will explain implications for physiotherapy, rehabilitation, and digital health applications.

8 7. Ethics, Information Governance, and Data Management

This study exclusively uses publicly available secondary image data and includes no participant recruitment, intervention, or direct interaction. The “Posture Keypoints Detection – Photos & Labels” dataset is made available on Kaggle under the Apache 2.0 License, and its use in this project follows the corresponding attribution and license requirements.

No attempt will be made to identify individuals from the dataset, and no personally identifiable information is included to the best of current knowledge. As no personal data are processed, the work is classified as secondary data research under University of Edinburgh governance guidance.

Data and derived outputs will be stored on approved, access-controlled systems. A complete record of preprocessing steps, model training pipelines, and derived features will be maintained. Code and workflows will be version-controlled and shared where licensing permits, aligned with FAIR principles.

9 8. Impact and Dissemination

The primary expected output is a proof-of-concept lightweight AI model, built with TensorFlow Lite and integrated with Google MediaPipe Pose, capable of detecting and analysing subtle postural and movement imbalances in upper-body alignment. Supporting outputs include annotated code, workflow documentation, and visual examples of pose-based analysis.

The intended impact is to advance accessible, low-cost, and non-invasive approaches for posture and movement assessment, with potential relevance for physiotherapy, sports science, education, and digital health.

Dissemination will prioritize practitioner- and public-facing channels, including open repository documentation, blog-style communication, and concise visual explanations.

10 9. Conclusion and Future Work

Lore Ypsum placeholder text: This chapter will summarise the final contributions, reflect on research objectives, and define priorities for future work.

11 Appendix A. Figures

Figure 1: DatabaseData.png




Fig A1: Smart Walker with the mobile recording setup; depth image overlaid with projected 2D skeleton; point-cloud overlaid with aligned 3D skeleton. **Source:** *Palermo et al. (2021)*

Figure 2: Timeline.png

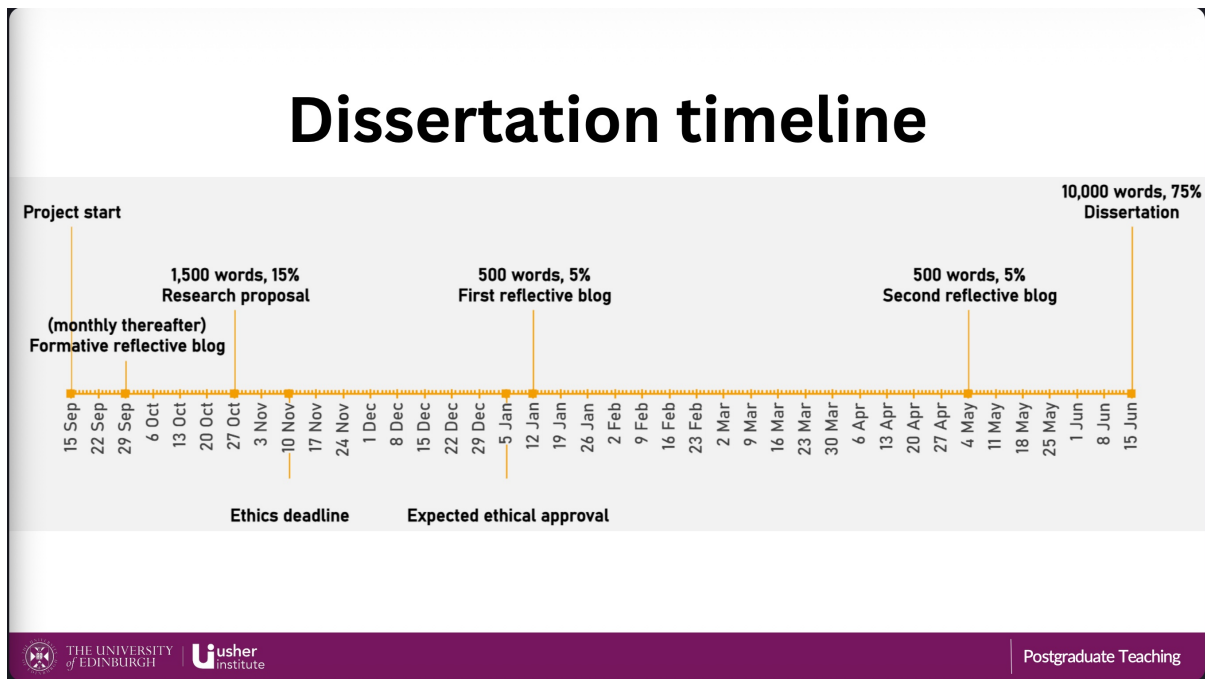


Fig A2: Project timeline outlining key milestones and deliverables. **Source:** *University of Edinburgh MSc in Data Science for Health and Social Care program*

12 References