



BRNO UNIVERSITY OF TECHNOLOGY

VYSOKÉ UČENÍ TECHNICKÉ V BRNĚ

FACULTY OF INFORMATION TECHNOLOGY

FAKULTA INFORMAČNÍCH TECHNOLOGIÍ

DEPARTMENT OF COMPUTER GRAPHICS AND MULTIMEDIA

ÚSTAV POČÍTAČOVÉ GRAFIKY A MULTIMÉDIÍ

**WEB APPLICATION FOR SETTING UP AN ERGONOMIC
POSITION ON A BICYCLE**

WEBOVÁ APLIKACE PRO NASTAVENÍ ERGONOMICKÉHO POSEDU NA JÍZDNÍM KOLE

MASTER'S THESIS

DIPLOMOVÁ PRÁCE

AUTHOR

AUTOR PRÁCE

Bc. JIŘÍ VLASÁK

SUPERVISOR

VEDOUCÍ PRÁCE

doc. Ing. MARTIN ČADÍK, Ph.D.

BRNO 2024

Master's Thesis Assignment



154460

Institut: Department of Computer Graphics and Multimedia (DCGM)
Student: **Vlasák Jiří, Bc.**
Programme: Information Technology and Artificial Intelligence
Specialization: Computer Vision
Title: **Web application for setting up an ergonomic position on a bicycle**
Category: Computer vision
Academic year: 2023/24

Assignment:

1. Familiarize yourself with the issue of computer support for sitting ergonomics on a bicycle.
2. Research existing methods and applications for setting up the correct posture.
3. Design and implement a video and/or camera based web application for custom bicycle seating.
Compare the approach based on the automatic detection of the rider's pose with the traditional marker-based motion capture method.
4. Experiment with the application, assess its properties in a user study, and/or further synthetic experiments. Discuss avenues for future development.
5. Present the achieved results in the form of a video, poster, article, etc.

Literature:

- Burt, P.: Bike Fit: Optimise Your Bike Position for High Performance and Injury Avoidance. A&C Black, 2014.
- Alberto Menache: Motion Capture Primer, In The Morgan Kaufmann Series in Computer Graphics, Understanding Motion Capture for Computer Animation (Second Edition), Morgan Kaufmann, 2011.
- Wade L, Needham L, McGuigan P, Bilzon J. Applications and limitations of current markerless motion capture methods for clinical gait biomechanics. PeerJ. 2022 Feb 25;10:e12995. doi: 10.7717/peerj.12995. PMID: 35237469; PMCID: PMC8884063.

Requirements for the semestral defence:
Points 1-3 of the assignment.

Detailed formal requirements can be found at <https://www.fit.vut.cz/study/theses/>

Supervisor: **Čadík Martin, doc. Ing., Ph.D.**
Head of Department: Černocký Jan, prof. Dr. Ing.
Beginning of work: 1.11.2023
Submission deadline: 17.5.2024
Approval date: 9.11.2023

Abstract

This work focuses on the development of a web application that helps users to set up an ergonomic position on a bicycle. The application uses the RTMPose pose estimation model to analyze a side view video of the user pedaling their bicycle, then computes various joint angles and provides the user suggestions for changing the bicycle saddle height, saddle setback, handlebar height and handlebar reach to achieve more ergonomic position. The application is implemented using the SvelteKit framework. The pose estimation model runs on the client-side using the TensorFlow.js library. To evaluate and improve the accuracy of the pose estimation model, a dataset of 2309 photos of the author pedaling his bicycle was created. After fine-tuning the model on this dataset, the model Normalized Mean Error significantly decreased from 2.54 to 1.51 (40% improvement). Experiments showed that that the application successfully detects angle changes caused by the saddle height changes and that the application can be used to set up the saddle height.

Abstrakt

Tato práce se zaměřuje na vývoj webové aplikace, která uživatelům pomůže nastavit ergonomickou polohu na jízdním kole. Aplikace využívá model odhadu pózy RTMPose k analýze videa uživatele šlapajícího na kole z profilu, poté vypočítá rozličné úhly kloubů a poskytne uživateli návrhy na změnu výšky sedla kola, předozadní polohu sedla, výšky řídítek a předozadní polohy řídítek pro dosažení ergonomičtějšího posedu. Aplikace je implementována pomocí frameworku SvelteKit. Model pro odhad pózy běží na straně klienta pomocí knihovny TensorFlow.js. Pro vyhodnocení a zlepšení přesnosti modelu odhadu pózy byla vytvořena datová sada 2309 fotografií autora šlapajícího na kole. Po fine-tunování modelu na tomto datovém souboru se normovaná průměrná chyba modelu výrazně snížila z 2,54 na 1,51 (zlepšení o 40 %). Experimenty ukázaly, že aplikace detekuje změny úhlu způsobené změnami výšky sedla a že ji lze použít k nastavení výšky sedla.

Keywords

pose estimation, bikefit, video processing, web application, RTMPose

Klíčová slova

odhad pózy, bikefit, zpracování videa, webová aplikace, RTMPose

Reference

VLASÁK, Jiří. *Web application for setting up an ergonomic position on a bicycle*. Brno, 2024. Master's thesis. Brno University of Technology, Faculty of Information Technology. Supervisor doc. Ing. Martin Čadík, Ph.D.

Rozšířený abstrakt

Tato práce se zaměřuje na vývoj webové aplikace, která uživatelům pomáhá nastavit ergonomickou polohu na jízdním kole. Aplikace využívá model odhadu pózy RTMPose k analýze videa uživatele šlapajícího na kole z profilu, poté vypočítá rozličné úhly kloubů a tyto úhly porovná s doporučenými hodnotami. Na základě tohoto porovnání aplikace následně poskytne uživateli návrhy na změnu výšky sedla kola, předozadní polohu sedla, výšky řídítek a předozadní polohy řídítek pro dosažení ergonomičtějšího posedu. Aplikace umožňuje nastavení posedu pro 3 typy kol (silniční, gravel a horské) a umožňuje uživateli vybrat si ze 3 typů posedů (komfort, výkon a kompromis mezi komfortem a výkonem).

První část práce se zabývá pozadím nastavování posedu na jízdním kole, možnými nastaveními kola a existujícími aplikacemi a metodami pro nastavení posedu na kole a jejich limitacemi. Další části se zabývají popisem metod sledování klíčových bodů těla ve videu. Nejdříve je představena jednoduchá metoda pro sledování nálepek a jsou diskutovány její nedostatky. Následně jsou představeny rozličné metody pro odhad pózy, které nálepky nepotřebují.

Pro vyhodnocení a zlepšení přesnosti modelů odhadu pózy byla vytvořena datová sada 2309 fotografií autora šlapajícího na kole. Anotace pro tuto datovou sadu byly získány pomocí sledování 8 barevných nálepek nalepených na straně těla otočené ke kameře. Nálepky byly sledovány pomocí jednoduchého algoritmu založeného na barevném prahování a trackování, který je v této práci také popsán. Zbylé anotace byly získány automaticky pomocí modelu odhadu pózy RTMPose.

Evaluace na této datové sadě ukázala, že modely RTMPose dosahují dobrých výsledků v odhadu pózy, ale chyba modelu je výrazně vyšší u klíčových bodů na noze (špička nohy a pata). Pro zlepšení přesnosti modelu byl model RTMPose fine-tunován na této datové sadě. Po fine-tunování modelu na tomto datovém souboru se normovaná průměrná chyba modelu výrazně snížila z 2,54 na 1,51 (zlepšení o 40 %). Fine-tunování proběhlo pomocí knihovny MMpose.

Aplikace je implementována jako single page aplikace pomocí frameworku SvelteKit. Model pro odhad pózy běží na straně klienta pomocí knihovny TensorFlow.js. Zpracování videa je rovněž prováděno na straně klienta pomocí WebCodecs API. Výsledné predikce modelu a vypočítané úhly kloubů jsou po zpracování zobrazeny uživateli nad videem za využití HTML canvas elementu. Aplikace kromě prezentování doporučených změn pro posed uživatele také zobrazuje pokročilejší statistiky a grafy jak se mění úhly kloubů pro jednotlivé cykly šlapání.

Experimenty ukázaly, že aplikace detekuje změny úhlu způsobené změnami výšky sedla a že ji lze použít k nastavení výšky sedla. Zároveň se ukázalo, že doporučení aplikace dobře korespondovaly s pocity autora při různých výškách sedla. Následné měření výkonu ukázaly, že aplikace zpracuje jeden snímek videa za asi 160-200 ms na starším stolním počítači s využitím WebGPU backendu a za asi 325-400 ms na mobilním telefonu s využitím WebGL backendu.

Web application for setting up an ergonomic position on a bicycle

Declaration

I hereby declare that this Master's thesis was prepared as an original work by the author under the supervision of Mr. doc. Ing. Martin Čadík, Ph.D and all the relevant information sources, literature, and other sources have been properly cited and included in the bibliography.

.....
Jiří Vlasák
May 13, 2024

Acknowledgements

I would like to express my gratitude to my supervisor, Mr. doc. Ing. Martin Čadík, Ph.D, for his guidance, valuable advice, and support throughout the work on this thesis.

Contents

1	Introduction	3
2	Bikefitting and Bikefitting Applications	4
2.1	Adjusting a Bicycle to the Rider	5
2.1.1	Saddle Height	5
2.1.2	Handlebar Height and Reach	6
2.1.3	Saddle Setback	7
2.2	Existing Solutions for Automated Bikefitting	8
2.2.1	Retül	8
2.2.2	MyVeloFit	9
2.2.3	Bike Fast Fit	11
3	Marker-based Tracking	13
3.1	Manual Localization of the Markers	13
3.2	Detection	14
3.3	Tracking	15
3.4	Advantages and Limitations	16
4	Pose Estimation Models	18
4.1	RTMPose	18
4.1.1	Architecture	19
4.1.2	Model Versions	20
4.2	HRNet	20
4.3	YOLOX-Pose	21
4.4	ViPNAS	22
4.5	Results On The COCO Keypoint Dataset	23
5	Pose Estimation Dataset for Bikefitting	25
5.1	Video Acquisition	25
5.2	Obtaining Ground Truth Annotations	25
5.3	Extracting Individual Frames	26
5.4	Removing Colored Markers From the Frames	26
5.5	Training and Validation Split	27
6	Fine-Tuning and Evaluation of Pose Estimation Models for Bikefitting	29
6.1	Evaluation Metrics	29
6.1.1	Percentage of Correct Keypoints (PCK)	29
6.1.2	Object Keypoint Similarity (OKS)	30

6.1.3	Mean Average Precision (mAP)	30
6.1.4	Normalized Mean Error (NME)	30
6.2	Initial Evaluation of Pre-Trained Models	31
6.2.1	6 Keypoint Evaluation	32
6.2.2	8 Keypoint Evaluation	32
6.3	Fine-Tuning On Custom Bikefit Dataset	35
6.3.1	Training Procedure	36
6.4	Evaluation Of Fine-Tuned Models	36
7	Bikefit Web Application	39
7.1	User Interface	40
7.1.1	Video Upload	40
7.1.2	Select Fit Goals	40
7.1.3	Examine Suggestions	41
7.2	Video Upload and Frames Extraction	42
7.2.1	Accessing the Video Frames	45
7.3	Pose Estimation	45
7.3.1	Model Conversion and Deployment	46
7.3.2	Model Download and Caching	46
7.3.3	Preprocessing	47
7.3.4	Postprocessing	47
7.4	Computing Joint Angles and Measurements	47
7.4.1	Determining the Pedal Position	48
7.5	Generating Suggestions	48
7.6	Video Player	49
8	Web Application Experiments	50
8.1	Testing Multiple Saddle Heights	50
8.2	Video Processing Benchmark	51
9	Conclusion	53
	Bibliography	54
A	Poster	58

Chapter 1

Introduction

Cycling enjoys immense popularity worldwide, attracting enthusiasts of all ages and skill levels, from daily commuters to professional athletes. However, the enjoyment of cycling hinge significantly on the proper fit of the bike to its rider. A poorly fitted bike can not only cause discomfort and pain but also increase the risk of long-term injuries and hinder performance potential.

Unfortunately, the setup of one’s bike is not a trivial task. Information on how to setup the bike are often limited and fail to account for the individual differences between riders and the dynamic nature of pedaling.

This led to the creation of a new profession called bikefitting. Bikefitters are experts who help cyclists set up their bikes. This was traditionally done in person and relied solely on the bikefitter’s experience and rudimentary tools such as plumb bobs, goniometers, etc. However, in recent years, to make it more accessible, bikefitting has been increasingly relying on technology, mainly motion capture systems. However, these in-person bikefitting sessions are still often not available in all areas and can be too costly for amateur cyclists.

To make bikefitting even more accessible, there has been a rise in bikefitting applications that aim to automate some of the bikefitting process. Despite the success of these applications, they are often limited to a single platform, provide only very basic information to the user, or are too costly.

This work aims to create a bikefitting application that does not suffer from these limitations. The application provides the user with detailed information, while being easy to use. It can be used to fit road, gravel or mountain bikes and allows the user to select their fit goal (performance, comfort, or both). The application uses a fine-tuned pose estimation model to estimate the rider’s pose from a side view video of the rider pedaling their bike. It then makes suggestions for adjusting the saddle height, saddle fore and aft position, handlebar height and handlebar reach. The application runs directly in the browser, which keeps the cost of the application low.

Chapter 2 explains how can a rider’s position on the bike be adjusted to make the rider more comfortable and perform better. It also compares several commonly used software systems for bikefitting. Chapter 3 describes how a marker-based tracking approach can be used to estimate the rider’s pose. Chapter 4 describes modern pose estimation models that can be used to estimate the rider’s pose. Chapter 5 describes a custom bikefitting dataset that is used to train and evaluate the pose estimation models. Chapter 6 describes how the pose estimation models are evaluated and fine-tuned on the custom dataset. Chapter 7 describes the bikefit application and chapter 8 assesses the application’s real-world performance. Chapter 9 concludes the work and proposes future work.

Chapter 2

Bikefitting and Bikefitting Applications

Bikefitting is a process of adjusting the bike to the rider. The goal of bikefitting is to make the rider as comfortable as possible on the bike, while also maximizing the rider's performance. This is done by adjusting the bike parts position, such as the height of the saddle, the fore and aft position of the saddle, the height of the handlebars or changing some of the bike parts, such as the stem, saddle or the handlebars. Main motivation behind bikefitting is to prevent injuries and pain, which can be over time caused by an incorrect bike setup.

While a rider can fit his bike himself, publicly available information on how to do it is often only composed of general guidelines and rules of thumb. Such guidelines include setting saddle height based on inseam length or so that when the heel is on the pedal at the bottom of the pedal stroke, the leg is straight. While they can be used as a starting point, they don't work for everyone.

Bikefitting is traditionally done in person by specialists, often with experience in physiotherapy, biomechanics, or sports science. These specialists are known as bikefitters. They can rely on their experience or use static measurements such as manually measuring angles at the knee at the bottom of the pedal stroke using a goniometer. However, measuring angles statically can be misleading, as the angles change dynamically during the pedal stroke and setting the leg in a static position can change the angles. This is in why recently bikefitting has been increasingly relying on motion capture technology, which provides the bikefitter dynamic measurements which are more accurate.

Additionally to bikefitters, technology has also been used to automate the bikefitting process altogether. This has led to the creation of bikefitting applications, which aim to provide the user with the information needed to set up the bike themselves. These applications are often based on pose estimation models, which predict the keypoints of the rider's body from the video. The keypoints are then used to compute the measurements used in bikefitting.

The following sections describe the most ways a bike can be adjusted to the rider and compares several commonly used systems and user applications for bikefitting.

2.1 Adjusting a Bicycle to the Rider

This section describes the easiest and most impactful ways a bike can be fitted to the rider. These changes can be done by the rider themselves and can have a significant impact on the rider's comfort and performance. The adjustments mentioned can also be suggested by a bikefitting application as they are easily measurable from a video. More advanced adjustments, which are often done by a bikefitter, are not described here.

2.1.1 Saddle Height

The saddle height is a fundamental measurement that significantly impacts a rider's comfort and pedaling efficiency. It measures how high the saddle is in relation to the bottom bracket. Figure 2.1 shows how is the saddle height measured.



Figure 2.1: Saddle height measures how high is the saddle compared to the bottom bracket.

It is argued to be the most important measurement in bikefitting and should be the first measurement to be adjusted [7]. Bad saddle height can even cause problems commonly associated with other measurements such as knee pain, lower back pain, neck pain, saddle sores, etc.

Optimal saddle height is nowadays determined by the rider's knee angle at the bottom of the pedal stroke. Knee angle is the angle between the hip and the ankle measured at the knee joint. Burt recommends a knee angle of 35-40 degrees for average riders and even up to 30 degrees for professional cyclists [7]. Figure 2.2 shows the knee angle at the bottom of the pedal stroke.

Higher saddle height can help to better recruitment of glutes and hamstrings, which can lead to more power output. However, it requires more flexibility and can lead to injury if the rider is not flexible enough. Similarly too low saddle height can increase the compressive forces on the knee and lead to pain and injury.

Proper saddle height is therefore a balance between power output and comfort. It is also important to note that the saddle height is not the only factor that affects the knee angle. The saddle fore and aft position and the cleat position also affect the knee angle.




Figure 2.2: Knee extension angle of 140-145 degrees, which is often referred to as 35-40 degrees (being the angle of deviation from straight leg). Is considered optimal for the average rider. Taken from [7].

2.1.2 Handlebar Height and Reach

Handlebar height measures how high the handlebars are in relation to the saddle. Handlebar reach measures how far the handlebars are in relation to the saddle. Figure 2.3 shows how are the handlebar height and reach measured.




Figure 2.3: Handlebar height and reach measure the vertical and horizontal distance to the handlebars from the bottom bracket.

Handlebar height influences mainly torso angle and to a lesser extent also shoulder angle. Torso angle is the angle between the shoulder and the level plane measured at the hip. It is also known as the back angle. Shoulder angle is the angle between the hip and wrist measured at the shoulder. Handlebar reach influences shoulder angle more and torso angle less.

Handlebar height can be adjusted by changing the number of spacers under the stem, by changing the stem orientation or by changing the stem itself. Handlebar reach can be

adjusted by changing the stem orientation, changing the stem or by fitting a handlebars with different shape.

While optimal handlebar height and reach are highly individual, there are some general guidelines. For example, a more upright position is more comfortable and easier to maintain and is therefore recommended for less experienced riders or for longer rides. A more aggressive position is more aerodynamic and is therefore recommended for faster riders. Burt recommends back angle of about 45 degrees for average riders and shoulder angle of about 90 degrees with the elbow slightly bent [7]. On the contrary, for faster riders, the back angle can be lowered down to 30 degrees with more open shoulder angle. For more upright riders, the back angle can be increased up to 55 degrees with more closed shoulder angle.

2.1.3 Saddle Setback

Saddle setback or saddle fore and aft position refers to the horizontal position of the saddle with respect to the bottom bracket. Figure 2.4 illustrates the saddle setback.



Figure 2.4: Saddle setback is the horizontal position of the saddle with respect to the bottom bracket.

Setback is most often measured at the 3 o'clock position of the crank. At this position, the rider's knee should be directly above the pedal spindle. Having the knee too far back makes it harder to generate power, while having the knee too far forward can lead to knee pain due to increased forces on the kneecap [7].

Saddle setback also affects the rider's balance on the bike. Having the saddle too far forward can lead to the rider's weight being too far forward, which can lead to hand pain because of too much weight on the handlebars. Having the saddle too far back can lead to the rider's weight being too far back, which can make the front wheel light and make the bike harder to control.

Setback also affects the rider's hip angle. Hip angle is the angle between the shoulder and the knee measured at the hip. Having the saddle more forward can lead to a more open hip angle, which can lead to more power output and more space between the rider's torso and legs at the top of the pedal stroke. This is why time trial bikes often have a more forward saddle position.

While saddle setback is an important measurement, it is often the last measurement to be adjusted. This is because the saddle height and the handlebar height and reach need to be set up first. Saddle height also changes the saddle setback, because the saddle seatpost is

not vertical, but angled. Meanwhile handlebar position also changes the weight distribution on the bike, so saddle setback needs to be adjusted last.

2.2 Existing Solutions for Automated Bikefitting

This section describes several commonly used software systems for bikefitting. These include both professional bike fitting systems, that are only available in bike shops and are used by trained bike fitters, and bike fitting applications, that are available to the general public and can be run on a mobile device or a computer.

2.2.1 Retül

Retül¹ is a professional bike fitting system employing 3D motion capture technologies. It is often available in bike shops, where a bike shop employee, who is trained by Retül, performs the bike fitting. Retül was the first company to use motion capture technology and it is still the most popular bike fitting system.

It uses LED markers placed on specific body points to track the rider's movements dynamically while cycling. The LED markers are tracked by multiple infrared cameras placed around the rider. The cameras surprisingly capture only 18 frames per second. Despite this, research [30] shows that the system is relatively reliable even when compared to 3D motion capture systems with higher frame rates.

Retül uses 16 markers placed on 8 joints on both sides of the rider's body. These markers are placed on the following locations: fifth metatarsal of the foot, heel, ankle, knee, hip, shoulder, elbow and wrist.

The markers are placed by the fitter on the rider's body. Accurate placement of the markers is crucial for the system to work properly. Even small deviations can lead to inaccurate results. Figure 2.5 shows the placement of the markers.

Retül's fitting process involves setting up the bike on a trainer equipped with the system. During the session, the rider performs various motions and pedal strokes while the Retül system captures real-time data on joint angles and movements.

Data Captured by Retül includes a wide range of joint angles and movements such as knee angles at top of the pedal stroke and bottom of the pedal stroke, hip angles throughout the pedal stroke, shoulder, elbow, and wrist positions in relation to handlebar reach and drop, as well as ankle and foot movement concerning cleat positioning and alignment. Figure 2.6 shows a small part of the captured data in Retül's software.

The fitter can then use normal ranges for these angles to make changes to the bike position. The normal ranges for these angles were constructed based on the data collected from thousands of cyclists. However, one common criticism of Retül is that these cyclists were not necessarily optimally fitted to their bikes. Therefore, the normal ranges may not be based on the optimal position for the rider.

Despite the fact that Retül is a very popular bike fitting system, it has some important disadvantages:

- Costly equipment and setup requirements, limiting accessibility to some individuals or smaller bike shops.
- The need for trained Retül bike fitters to interpret and implement fitting recommendations effectively.

¹<https://www.retul.com/>



Figure 2.5: Placement of the markers used by Retül. (Screenshot from instructional instructional video by Retül [29].)

Measurement Title	1	2	3	Units	Range
Ankle Angle at Bottom	107	15	122	deg	90 to 100
Maximum Knee Flexion	113	-11	102	deg	107 to 113
Maximum Knee Extension	36	-5	31	deg	32 to 42
Knee Forward of Spindle	-63	5	-58	mm	-35 to -5
Foot Rotation Mean	-11	3	-8	deg	N/A
Hip Vertical Travel	45	10	55	mm	N/A
Hip Lateral Travel	10	-1	9	mm	5 to 20
Cadence Avg	82	-4	78	rpm	N/A

Figure 2.6: Retül’s software showing a small part of the captured data. (Screenshot from instructional video by Retül [28].)

- Requires in-person fitting sessions. These sessions can be time-consuming and costly.

2.2.2 MyVeloFit

MyVeloFit² is a very popular bikefitting web application that uses a pose estimation model to predict the joint locations for a side view video of the user pedaling their bike on an indoor trainer. Based on location of these joints, joint angles are then computed. On the basis of these angles and their relation to average angles, suggestions are made for adjusting the position of the saddle and handlebars.

The fitting process starts with the rider filling out a questionnaire about their mobility. This is then used to adjust the recommended angle ranges. For example: if the user has lower shoulder mobility, recommended ranges for shoulder angle will be increased so the user is not stretched forward so much.

²<https://www.myvelofit.com/>

After creating the user profile, user can create a fit session for one specific bike. In the process, the user selects their fit goal (performance, comfort, or balanced) and the type of bike they are using (road, gravel, mtb, triathlon, hybrid, or stationary). This also changes the recommended angle ranges.

MyVeloFit predicts 6 joint locations for the camera facing side of the body. Most common pose estimation models predict similar keypoints. However, keypoints, which are needed to compute ankle angles and are commonly used to adjust the position of the saddle, such as the heel and the fifth metatarsal of the foot, are missing. The keypoints used by MyVeloFit are: ankle, knee, hip, shoulder, elbow and wrist. Figure 2.7 shows the predicted keypoints.


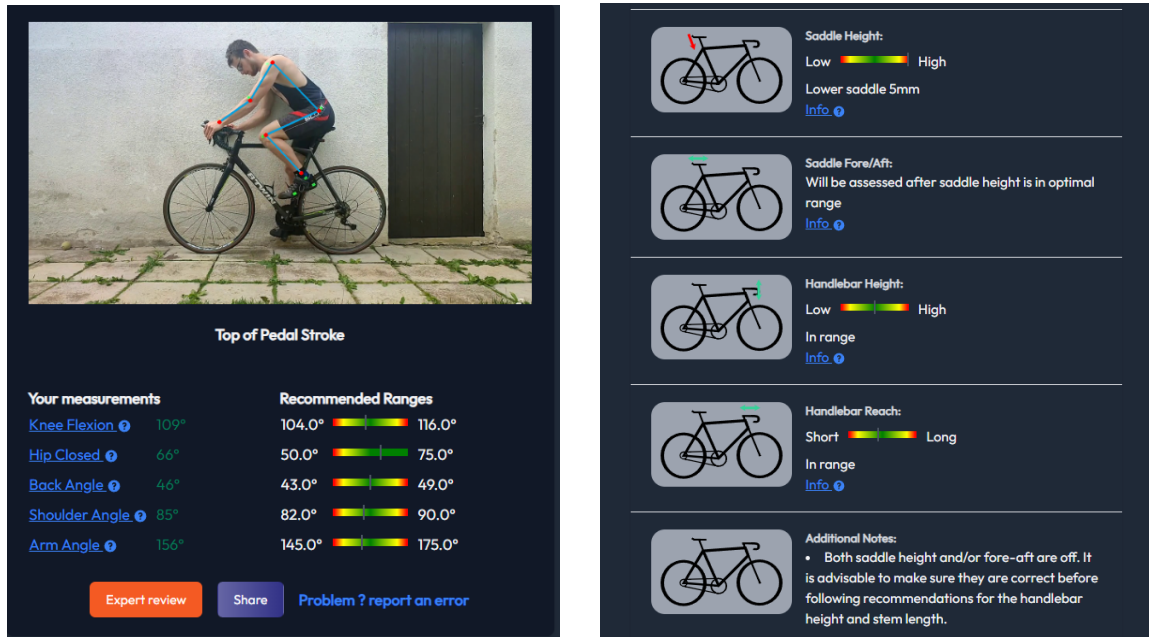


Figure 2.7: Side view image with predicted keypoints in MyVeloFit.

From the joint angles for every frame, some are selected for computing the joint angles at the top of the pedal stroke, front of the pedal stroke and bottom. Every position uses different angle ranges. Based on the angles computed for parts of the pedal stroke, MyVeloFit then makes suggestions for saddle height, saddle fore and aft, handlebar height and handlebar reach. Figure 2.8 shows the predicted joint angles at the top of the pedal stroke and the recommended changes to the bike position.

Overall, MyVeloFit is relatively easy to use and its joint predictions are fairly accurate. However, it has few disadvantages:

- Only the most basic keypoints are used and therefore it does not take into account the ankle angles and the foot position.
- Every video is converted to 30 FPS and cut down to 10 seconds.
- Video processing and keypoint predictions are slow (3-5 minutes).
- It is not cheap to get joint angles and recommended changes. It requires either a one time payment of 35 US dollars for access for 1 person and 1 bike for 2 weeks or 75 US dollars annually for unlimited number of bikes and people.
- More advanced features are missing, such as the ability to see how the angles change over time, visualizations of the angles, etc.



(a) Predicted joint angles at the top of the pedal stroke in MyVeloFit. (b) Recommended changes to the bike position based on the angles computed by MyVeloFit.

Figure 2.8: Predicted joint angles at the top of the pedal stroke and recommended changes to the bike position in MyVeloFit.

2.2.3 Bike Fast Fit

Bike Fast Fit³ are iOS mobile applications that use pose estimation model to predict the joint locations for a side view video of the user pedaling their bike on an indoor trainer. Compared to MyVeloFit, they use additional keypoints for the fifth metatarsal of the foot and the heel. There are 2 version of the app: Bike Fast Fit Elite and Bike Fast Fit EZ. The EZ version provides suggestions for saddle height and fore and aft position. The Elite version provides suggestions for saddle height, fore and aft position and handlebar position and additional features such as front view knee tracking. Figure 2.9 shows a screenshot of the Elite version of the app.

Both the EZ and Elite versions of the app seem to be popular, with the Elite version having a 4.6 star rating on the App Store from 269 ratings⁴. The EZ version has a 4.5 star rating from 91 ratings⁵.

The Elite version costs 19.99 US dollars and does not require a subscription. The EZ version can be used free for 7 days and then requires a subscription of 4.99 US dollars per year.

Despite the popularity of the apps, the developer stated that he is not interested in developing an Android version of the app. Another limitation is that the apps only capture 3.5 seconds of video, which translates to about 3-5 pedal strokes. This can limit the accuracy of the predictions.

³<https://www.bikefastfit.com/>

⁴<https://apps.apple.com/us/app/bike-fast-fit-elite/id1145619812>

⁵<https://apps.apple.com/us/app/bike-fast-fit-ez/id1578925255>

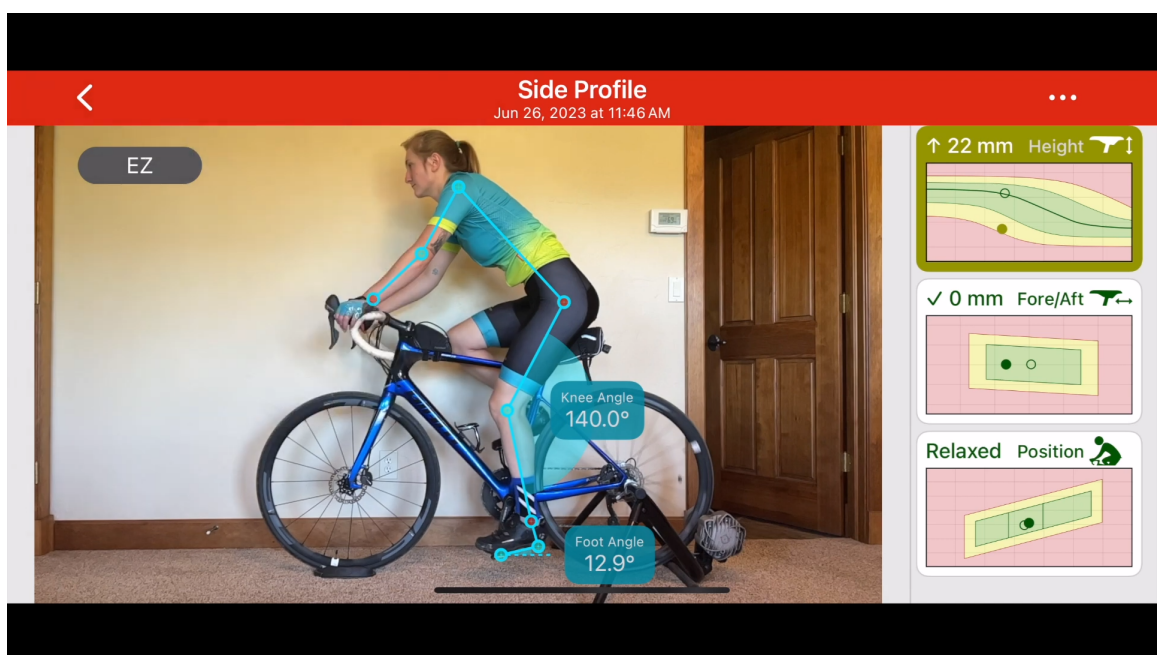


Figure 2.9: Side view image with predicted keypoints in BikeFast Fit Elite. (Screenshot from promotional video [5] by BikeFast Fit Elite.)

Chapter 3

Marker-based Tracking

To automatically fit a bike to a rider, a bikefitting system needs to estimate the rider's pose. One way to estimate the cyclist's pose is to use markers, which are placed on the parts of the body. These markers can then be accurately detected and tracked in the video. The position of the markers can then be used to estimate the keypoints of the rider's body.

This chapter describes an approach to estimating the keypoints of the rider's body using colored markers and color thresholding. The markers are first manually localized by the user in the first frame of the video. The color of the markers is then used to detect them in the following frames. The markers are then tracked using a method similar to SORT [4]. This method was implemented in a web application with the help of the OpenCV.js [34] library, which is a JavaScript binding for WebAssembly port of OpenCV [6].

Following sections describe the implementation of this approach and discusses its advantages and limitations and why it is not suitable for a user application.

3.1 Manual Localization of the Markers

At first, the user is asked to manually localize the markers by clicking on them in the first frame of the video. Figure 3.1 shows an example of the manual localization of the markers. The colors of these markers are saved and used for detection of the markers by thresholding in the following frames. The user is asked to first localize the marker on the foot, then the marker on the heel, then the marker on the ankle, etc. Figure 3.1 shows the manual localization process of the markers.

After clicking on a marker, multiple steps are performed:

1. The x,y coordinates of the marker are saved.
2. If there is already some marker localized and therefore some color thresholds defined and markers detected, the detected markers are searched if there is some marker close to the clicked marker. If there is, the location of the clicked marker is updated to the location of the detected marker. This is done so that user does not have to precisely click on the marker. This is especially useful when the markers are small and the user is using a mobile device with a small screen.
3. Based on the x,y coordinates of the marker, the color of the marker is extracted from the image.
4. The color of the marker is used to update the color thresholds used to detect the markers. More on this in section 3.2.

Mark the joints

The second joint should be the heel.




Figure 3.1: Example of the manual localization of the markers. The user is asked to click on the markers in the first frame of the video. The single localized foot marker is highlighted in green.

Figure 3.2 shows a closeup of the localization of the markers. The single localized foot marker is highlighted in green. The contours of the thresholded regions are highlighted in white. The centroids of the contours are highlighted by small green cross.

3.2 Detection

After obtaining the markers' locations and colors, the markers are detected in the following frames. The detection is performed in the following steps:

1. The frame is converted to the CIELAB color space.
2. The frame is thresholded using the minimum and maximum A and B channel values of the localized markers and some margin. The margin is used to account for the changes in lighting conditions. Empirically, the margin of 15 works well. Figure 3.3 shows an example of the thresholded image.
3. Contours are found in the thresholded image. The `cv.findContours()` method is used for this.
4. The centroids of the contours are calculated using the `cv.moments()` method.
5. The markers are then found in the detected centroids.



Figure 3.2: Closeup of the localization of the markers. The single localized foot marker is highlighted in green. The contours of the thresholded regions are highlighted in white. The centroids of the contours are highlighted by small green cross.

The CIELAB color space is used because it is designed to approximate human vision. The A and B channels are used because they represent the color of the marker. The L channel is not used because it represents the lightness of the color and is not very useful for detecting the markers due to the changes in the lighting conditions between different parts of the body. Experiments with using the H channel of the HSV color space were also performed, but there were some problems with the detection of the markers on the edges of the color spectrum ($H=0$ and $H=255$). The CIELAB color space does not have this problem, because it uses two channels for the color. Figure 3.3 shows an example of the thresholded image.

3.3 Tracking

The detected markers are then tracked using a simple method similar to SORT [4]. The tracking algorithm works as follows:

1. Every possible assignment of the detected markers and the markers in the previous frame is created. To limit the number of possible assignments, the assignments are created only if the distance between the markers is less than 50 pixels.
2. For every found assignment, the cost of the assignment is calculated as the sum of the distances between the markers in the assignment.
3. To make tracking more robust, the cost of the assignment is increased by the difference of the angles between the markers in the assignment. This is done to penalize the assignments where the angles between the markers change too much. This happens when the markers for two or more joints are switched.
4. Assignment with the most markers assigned and the lowest cost is selected.




Figure 3.3: Example of the image thresholded using the minimum and maximum A and B channel values of the localized markers and a margin of 15. The markers are then detected as the centroids of the contours in the thresholded image.

Experiments with using the Kalman filter were also performed, unfortunately the author was not able to get the Kalman filter to work properly.

3.4 Advantages and Limitations

The marker-based approach was mainly chosen due to its theoretically perfect accuracy, which is its biggest advantage. However, in practice there are several limitations to this approach:

1. Accuracy of the system heavily depends on the placement of the markers. If the markers are not placed correctly, the results are not accurate. Since the placement is done by the user, this is a real issue. Correct placement of the markers is not trivial due to possibly changing body positions, loose clothing, camera angle, etc.
2. The user needs to have the markers available. They don't have to be special video markers, but they need to be of a different color than the background and the clothing of the user and need to be glued to the user.
3. The biggest limitation is the reduced robustness of the system. The system works well if the color of the markers is not present elsewhere in the image. However, in other cases, the markers are detected in the wrong places and the system does not work. Figure 3.4 shows an example of the failed localization of the markers. The knee marker is not localized correctly, because the leaves in the background have similar color to the markers.

Due to these serious limitations, the marker-based approach was not used in the final version of the application. And instead, a pose estimation model was used to estimate the



Figure 3.4: Example of a failed localization of the markers. The leaves in the background have similar color to the markers and the knee marker is therefore not localized correctly. The location of the markers is show as the green dots.

rider's pose without the need for markers. The following chapter describes modern pose estimation models that were evaluated for this task.

Chapter 4

Pose Estimation Models

Pose estimation is a computer vision task that involves predicting the locations of keypoints of a person in an image or a video. The keypoints are usually the joints of the person, such as the ankle, knee, hip, shoulder, elbow, wrist, etc, but can also be other points of interest, such as the nose, eyes, ears, etc. The pose estimation models are trained on datasets that contain images or videos with annotated keypoints. The models are then evaluated on the same datasets.

Not all of the models predict the same keypoints. The models usually predict 17 keypoints, as defined in the COCO dataset [22]. These keypoints are composed of prominent joints and few head keypoints. As the need for more keypoints and better understanding of the human pose has increased, the COCO-WholeBody dataset [18] was created. This dataset contains 133 keypoints, which include the 17 COCO keypoints and additional keypoints such as the fingers, toes, heels, etc. However, the models often struggle to predict all of the keypoints accurately due to the low amount of data for some of the keypoints. Another newly created dataset is Halpe26 dataset [11], which contains 26 keypoints. This dataset is designed to be more lightweight than the COCO-WholeBody dataset, while still containing more keypoints than the COCO dataset. Figure 4.1 shows the sets of keypoints used by the MS-COCO dataset, the Halpe26 dataset and the COCO-WholeBody dataset.

There are multiple approaches to pose estimation. The two most common approaches are the top-down approach and the bottom-up approach. The top-down approach first detects the people in the image and their location and then predicts the keypoints for each person. The bottom-up approach first predicts the keypoints and then groups them into poses using hand written post processing algorithms. Another approach is the one-stage approach, which predicts the keypoints directly from the image without the need for separate person detecting or post processing step. This chapter describes three pose estimation models using the top-down approach (RTMPose, HRNet and ViPNAS), described in sections 4.1, 4.2 and 4.4 and one pose estimation model using the one-stage approach (YOLOX-Pose), described in section 4.3.

4.1 RTMPose

RTMPose [17] is a pose estimation model from the authors of the MMPose framework [10]. It is designed to bridge the gap between the excellent performance of 2D pose estimation on public benchmarks and its application in the industrial community, which still suffers from heavy models and high latency.

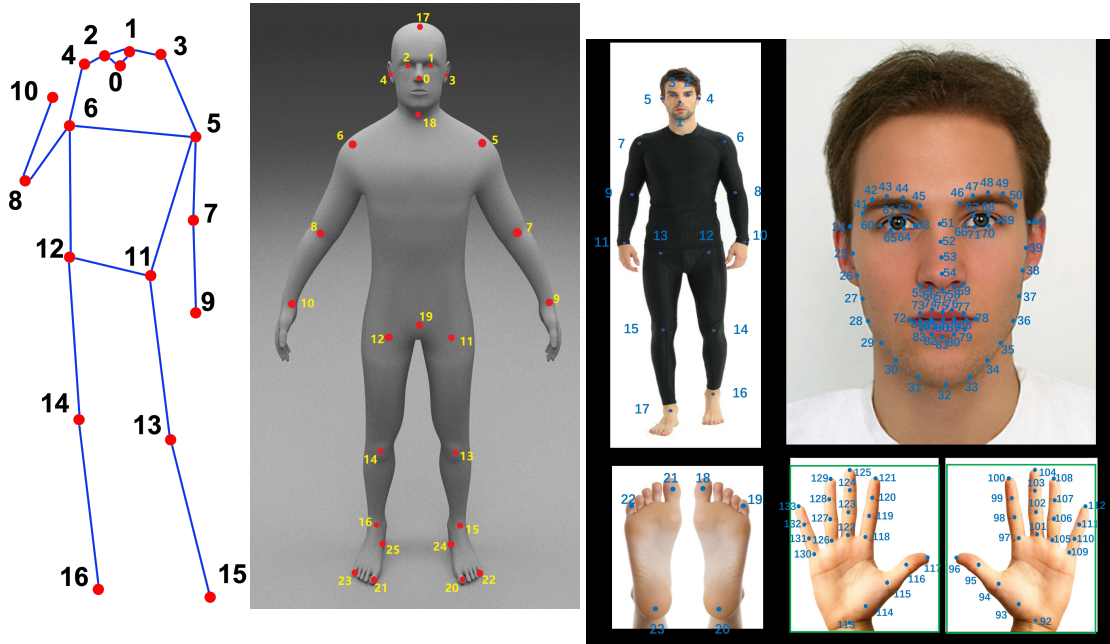


Figure 4.1: Sets of keypoints used by (from left to right) MS-COCO dataset, Halpe26 dataset and COCO-WholeBody dataset (images taken from [22] and [11]).

The RTMPose models are designed to be lightweight and fast. The authors claim, that the RTMPose-m achieves 75.8% AP on COCO with 90+ FPS on an Intel i7-11700 CPU and 430+ FPS on an NVIDIA GTX 1660 Ti GPU. The RTMPose-s model, achieving 72.2% AP on COCO was also tested on a mobile device with the Snapdragon 865 chip, running at 70+ FPS.

4.1.1 Architecture

The RTMPose models are based on the top-down pose estimation approach. They use a two-stage prediction, where the first stage is a person detector and the second stage is a pose estimator. The person detector is used to crop the image to the bounding box of the person. The pose estimator then predicts the keypoints for the cropped image. The authors claim that this approach is more accurate than the bottom-up approach, while still being faster in cases where the number of people in the image is lower than 6.

The CSPNeXt backbone is used in the RTMPose models. This backbone is primarily designed for object detection. Authors claim that backbones designed for image classification (such as ResNet, MobileNet, etc.), are not optimal for dense tasks such as pose estimation, object detection, semantic segmentation, etc. Some backbones using high-resolution feature maps or advanced transformer architectures achieve good results, but suffer from high computational cost, high latency or difficulties in deployment. The CSPNeXt backbone is designed to have a good balance of speed and accuracy. Figure 4.2 shows the architecture of the RTMPose models.

For prediction of the keypoints, the RTMPose models utilize the SimCC [20] algorithm. SimCC reformulates human pose estimation as two classification tasks for horizontal and vertical coordinates. To reduce quantization error, SimCC can use a larger number of bins for the classification tasks. This can lead to better accuracy, while still being faster than

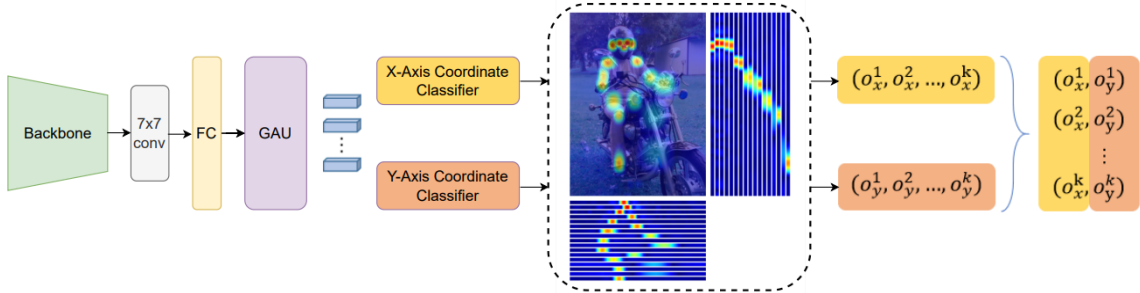


Figure 4.2: Architecture of the RTMPose models, which comprises of a CSPNeXt backbone, a convolutional layer, a fully-connected layer, and a Gated Attention Unit (GAU) designed to enhance K keypoint representations. Subsequently, the process of 2D pose estimation is treated as two separate classification tasks for the x-axis and y-axis coordinates. This involves predicting the horizontal and vertical positions of keypoints. Taken from [17].

post-processing methods commonly used with traditional heatmap based pose estimation models.

To better use spatial information, the RTMPose models use a Gated Attention Unit (GAU) [14] module. The GAU has faster speed, lower memory cost and better accuracy than the commonly used self-attention module, proposed in the Transformer architecture [37].

4.1.2 Model Versions

This work evaluates the following RTMPose models:

RTMPose Body8 models: These models are trained and evaluated on the Body8 dataset consisting of 8 pose estimation datasets (AI Challenger [38], MS-COCO [22], Crowd-Pose [19], MPII [2], sub-JHMDB [16], Halpe [11], PoseTrack18 [1] and OCHuman [40]). They predict 17 keypoints, as defined in the COCO dataset.

RTMPose Halpe26 models: These models are also trained and evaluated on the Body8 dataset. However, they predict 26 keypoints, as defined in the Halpe26 dataset.

RTMPose WholeBody models: These models are trained and evaluated on the COCO-WholeBody [18] and UBody [21] datasets. They predict 133 keypoints, as defined in the COCO-WholeBody dataset.

4.2 HRNet

The HRNet (High-Resolution Network) [33] pose estimation architecture is characterized by its emphasis on maintaining high-resolution representations throughout the network. Unlike conventional methods that downsample the input image early in the network, HRNet adopts a multi-resolution approach, preserving detailed information essential for accurate pose estimation.

A notable feature of the HRNet pose estimation architecture is its ability to capture both local and global context effectively. The network consists of parallel branches, each processing different resolutions of the input image, and these branches are interconnected,

facilitating the exchange of information between different scales. This design allows HRNet to address challenges associated with scale variations in human poses, ensuring robust performance in capturing both fine details and overall pose structure. Figure 4.3 shows the architecture of the HRNet models.

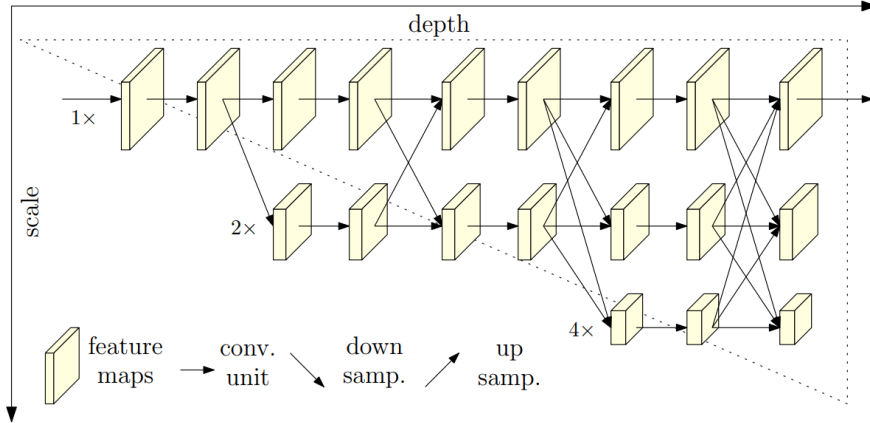


Figure 4.3: Architecture of the HRNet models. The HRNet models consists of three parallel branches, each processing different resolutions of the input image. These branches are interconnected, facilitating the exchange of information between different scales. Taken from [33].

This work evaluates the HRNet-W32 model, which uses heatmap to predict the keypoints. The HRNet-W32 model is trained on the COCO dataset and predicts 17 keypoints, as defined in the COCO dataset.

4.3 YOLOX-Pose

YOLOX-Pose is a one-stage human pose estimation model that distinguishes itself from traditional top-down and bottom-up methods. The authors claim that YOLOX-Pose is the first one-stage human pose estimation model that achieves comparable performance to the state-of-the-art two-stage methods.

The one-stage approach has some advantages over bottom-up and top-down approaches. Bottom-up approaches first detect keypoints and then group them into poses. The grouping step is not learned and needs to be hand-crafted. This can lead to inaccurate results, especially in cases where the keypoints from different people are close together.

Top-down approaches first detect people and then predict keypoints for each person. This approach is favored by SOTA models, but the processing time grows linearly with the number of people in the image. This can lead to slow processing times in cases where there are many people in the image. SOTA methods also often use heavy detectors, which further increases the processing time.

The original YOLO-pose [27] model uses the anchor-based YOLOv5 detector [36]. The YOLOX-Pose model, which is evaluated in this work, uses the YOLOX detector [12]. The YOLOX detector is based on the YOLOv5 detector, but uses anchor-free approach, which does not require the anchor boxes to be predefined. This allows the YOLOX detector to be more flexible. The YOLOX detector generally outperforms the YOLOv5 detector on the COCO dataset. Figure 4.4 shows the architecture of the YOLOX-Pose model.

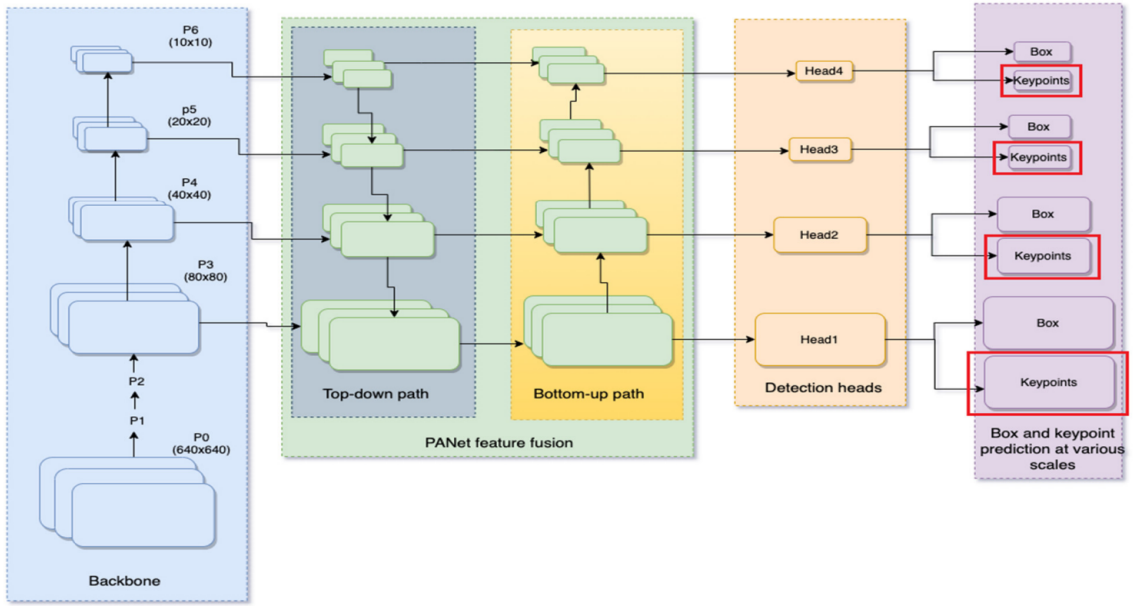


Figure 4.4: Architecture of the original YOLO-Pose model, which was based on the anchor-based YOLOv5 detector. The YOLOX-Pose model, which is evaluated in this work, is based on the anchorless YOLOX detector. Taken from [27].

The YOLOX-Pose model is trained on the COCO dataset and predicts 17 keypoints, as defined in the COCO dataset.

4.4 ViPNAS

The ViPNAS (Video Pose Estimation via Neural Architecture Search) [39] model focuses on video pose estimation, employing two networks: S-ViPNet and T-ViPNet.

The S-ViPNet is used on individual frames and keyframes in the video to predict the keypoints. The lightweight T-ViPNet network is used on non-keyframes to propagate the pose estimations. It extracts features from the current frame and also fuses heatmaps from the previous frame. The fused features are then used to predict the keypoints.

To find the optimal architectures for the S-ViPNet and T-ViPNet networks, the authors used neural architecture search (NAS). The search space for the S-ViPNet network is called the spatial-level search space. It consists of 5 dimensions (depth, width, kernel size, group and attention). The search space for the T-ViPNet network is called the temporal-level search space. It searches for the optimal fusion operation (addition, multiplication or concatenation) and the best stage of features to fuse.

This work evaluates the modified S-ViPNet model with the MobileNetV3 [13] backbone and SimCC [20] algorithm instead of predicting heatmaps. The model is trained on the COCO dataset and predicts 17 keypoints, as defined in the COCO dataset.

4.5 Results On The COCO Keypoint Dataset

To compare the potential of the methods, the results of the models on the COCO dataset are shown in table 4.1. All of the methods were trained only on the COCO dataset and predict the same keypoints. The results are taken from MMPose documentation [10].

The RTMPose models achieve strong results, with the RTMPose-l 256x192 model achieving the best results. The HRNet-W32 model also performs well but it is a much larger and slower model. However, the other lightweight models such as YOLOX-Pose and ViPNAS-MobileNetV3 perform significantly worse than even the medium and small RTMPose models. Overall the RTMPose models show very promising results compared to the other models.

Architecture	Input Size	AP	AP-50	AP-75	AR	AR-50
RTMPose-l	256x192	0.758	0.906	0.826	0.806	0.942
HRNet_w32	256x192	0.749	0.906	0.821	0.804	0.945
RTMPose-m	256x192	0.746	0.899	0.817	0.795	0.935
RTMPose-s	256x192	0.716	0.892	0.789	0.768	0.929
SimCC_ViPNAS-MBNetv3	256x192	0.695	0.883	0.772	0.755	0.927
Yoloxpose_m	640x640	0.695	0.899	0.766	0.733	0.926
RTMPose-t	256x192	0.682	0.883	0.759	0.736	0.920

Table 4.1: Results of the models on the COCO dataset. The results are sorted by the Average precision (AP). The AP describes the average precision over Object Keypoint Similarity (OKS) thresholds from 0.5 to 0.95 with a step size of 0.05. The AR describes the average recall over OKS thresholds from 0.5 to 0.95 with a step size of 0.05. The AP-50 and AR-50 are the same as AP and AR, but only for OKS threshold of 0.5. Similarly AR-75 is average recall with OKS threshold of 0.75. The results are taken from MMPose documentation [10].

Table 4.2 compares the performance of the models with their computational costs. The RTMPose models are the most computationally efficient. Especially the medium model, which achieves similar results to the HRNet-W32 model, but is much faster and smaller. The YOLOX-Pose model computational costs are higher, because it is a one-stage model, which predicts the keypoints directly from the image. The ViPNAS-MobileNetV3 model is also very small and fast but it achieves significantly worse results than the RTMPose models.

Architecture	Input Size	AP	FLOPS(G)	Params(M)
RTMPose-l	256x192	0.758	4.2	28.1
HRNet_w32	256x192	0.749	7.7	28.5
RTMPose-m	256x192	0.746	1.9	13.9
RTMPose-s	256x192	0.716	0.7	5.7
SimCC_ViPNAS-MBNetv3	256x192	0.695	0.8	8.6
Yoloxpose_m	640x640	0.695	5.8	29.3
RTMPose-t	256x192	0.682	0.4	3.5

Table 4.2: Results of the models on the COCO dataset and their computational cost. The results are sorted by the Average precision (AP). The AP describes the average precision over Object Keypoint Similarity (OKS) thresholds from 0.5 to 0.95 with a step size of 0.05. The FLOPS describe the number of floating point operations needed to process the image. The numbers in the table are in billions. The Params describe the number of parameters in the model. The numbers in the table are in millions. The computational cost numbers are obtained using the MMPose library.

Chapter 5

Pose Estimation Dataset for Bikefitting

To evaluate and improve the performance of pose estimation models in the setting of bikefitting, it is necessary to have a dataset that most closely resembles the real-world use case (videos of cyclists pedaling their bikes, usually in a side view). Since there is no such dataset, a new small dataset was created.

5.1 Video Acquisition

To construct the dataset, the author shot 23 videos of himself pedaling on an indoor trainer or leaning on a wall and pedaling a bike. The videos feature both a road bike and a mountain bike. To ensure variety, they were shot on multiple days in different lighting conditions and environments, with various camera angles, clothes and riding positions. The videos were shot in FullHD resolution at 30 or 60 FPS using a Redmi Note 10 Pro smartphone. Both the standard and ultra-wide camera lens were used. The videos are 8-35 seconds long. Figure 5.1 shows examples of some of the videos present in the dataset.

5.2 Obtaining Ground Truth Annotations

Before shooting the videos, colored markers were placed on the camera-facing side of the rider's body. These markers were placed on the following locations: fifth metatarsal of the foot, heel, ankle, knee, hip, shoulder, elbow and wrist. The marker placement was guided by the instructions from the Retül bike fitting system¹. These markers were then tracked using the methodology described in chapter 3. In contrast to manual annotation, tracking the placed markers allows for obtaining the ground truth locations of the keypoints for all of the frames in the videos. Another advantage is that feeling the keypoint locations in real world should be more accurate than annotating the keypoints in the video.

After obtaining annotations for the 8 markers facing the camera, the annotations for the remaining keypoints need to be obtained. To do this, the RTMPose-l 256x192 model was used to predict the keypoints for the videos. The model was trained on the Body8 dataset and predicts 26 keypoints. These keypoints are not needed for evaluation of the models for bikefitting, but are needed for the fine-tuning of the models. If the models

¹<https://youtu.be/kF0Zjtp04FE?si=JRLX71DNZ72oTaiC>



Figure 5.1: Examples of some the videos present in the dataset. The videos feature both a road bike and a mountain bike. The videos were shot on multiple days in different lighting conditions and environments, with various camera angles, clothes and riding positions.

are fine-tuned only on the 8 keypoints, the model will unlearn how to localize the other keypoints. Another reason for predicting the keypoints is to get the bounding box of the person. The bounding box is needed to crop the frames to the bounding box of the person. The bounding box is predicted by the RTMDet-nano model [25] from the MMDetection framework [8]. Figure 5.2 shows an example frame from the dataset with the annotations.

5.3 Extracting Individual Frames

After obtaining the annotations, the individual frames are extracted from the videos. This is done since the pose estimation models are trained on individual frames and not on videos.

To reduce the size of the dataset and remove duplicate frames, not all of the frames are extracted. The frames to extract are selected using the k-means clustering algorithm [26]. Then only the centroids of the clusters are used as the frames for the dataset. The clustering is ran on the locations of the 8 markers. The number of clusters is set to 1/3 of the number of frames.

5.4 Removing Colored Markers From the Frames

To not bias the training and evaluation of the models, the colored markers are removed from the frames. This is done using an inpainting algorithm. The inpainting algorithm takes the frames with the colored markers and a mask of the colored markers and outputs the frames with the colored markers removed.

The mask of the colored markers is created the same way as in the marker-based tracking method, described in chapter 3. It works by converting the frames to the CIELAB color



Figure 5.2: Example frame with annotations from the dataset. Ground truth locations of the 8 camera facing keypoints are obtained by tracking colored markers. The remaining keypoints are predicted by the RTMPose-l-256x192 model. The bounding box of the person is obtained by the RTMDet-nano model.

space and thresholding the A and B channels using the minimum and maximum A and B channel values of the colored markers.

The inpainting was done using the opencv library [6]. The library contains implementations of the Navier-Stokes inpainting algorithm [3], which is based on fluid dynamics, and the Fast Marching Method [35], which works by propagating the information from the known parts of the image to the unknown parts. After testing both algorithms, the Navier-Stokes inpainting algorithm was chosen, as it produced better results. Figure 5.3 shows an example of the inpainting algorithm. It works well in most cases, but struggles with non-uniform backgrounds and can inpain unwanted parts of the image if the mask is not accurate.

5.5 Training and Validation Split

The extracted frames are split into training and validation sets. Frames from the same video are not split between the training and validation sets. The training set contains **1833** frames extracted from **19** videos and the validation set contains **476** frames extracted from **4** videos.




Figure 5.3: Example of the removal of the colored markers. The red colored markers are removed from the original image (top) using the Navier-Stokes inpainting algorithm. Most of the markers are removed well, but the algorithm produces blurry results on non-uniform backgrounds (pants with logo) and sometimes inpaints unwanted parts of the image (face), if the mask is not accurate.

Chapter 6

Fine-Tuning and Evaluation of Pose Estimation Models for Bikefitting

This chapter evaluates the performance of the pose estimation models on side view pose estimation of cyclists. Based on these results the best performing model is selected and fine-tuned on the bikefitting dataset. The fine-tuned model is then once again evaluated on the bikefitting dataset and compared with the pre-trained models. This model is then used in the bikefit application. Models are evaluated and fine-tuned on the dataset described in section 5.

The evaluation metrics are described in section 6.1. Initial evaluation of the standard, pre-trained models is done in section 6.2. The fine-tuning of the models is described in section 6.3. The results of the fine-tuned models are shown in section 6.4.

All of the evaluations and training is done using the MMDetection framework [8] and the MMPose framework [10].

6.1 Evaluation Metrics

This section describes commonly used metrics for evaluating pose estimation models and their limitations. The Normalized Mean Error (NME) metric, described in section 6.1.4, is used in the evaluation of the models on the bikefitting dataset.

6.1.1 Percentage of Correct Keypoints (PCK)

Percentage of Correct Keypoints (PCK) is a simple metric that calculates the percentage of keypoints that are within a certain distance from the ground truth keypoints. The distance is normalized by the size of the bounding box of the person. It provides a simple and intuitive measure of the accuracy of the model.

However, for the bikefitting dataset, commonly used thresholds of 0.1 or 0.05 are too big and the models achieved close to 100% PCK on the validation set. With a lower threshold, the PCK would swing wildly if the distance between the keypoints and the ground truth keypoints is close to the threshold. This is why the PCK metric is not used in the evaluation of the models on the bikefitting dataset.

6.1.2 Object Keypoint Similarity (OKS)

The Object Keypoint Similarity (OKS) metric is a widely used metric for evaluating pose estimation models. It calculates the similarity between the predicted keypoints and the ground truth keypoints based on the distance between the keypoints and defined variance for each keypoint.

The equation for keypoint similarity is:

$$KS_i = \exp\left(\frac{-d_i^2}{2s^2\sigma_i^2}\right).$$

KS_i is the keypoint similarity for keypoint i .

d_i is the Euclidean distance between the ground truth and predicted keypoint i .

k_i is the variance for keypoint i .

s is the scale of the object, s^2 is the area of the bounding box of the person.

The final OKS is defined as sum of the keypoint similarities of visible keypoints divided by the number of visible keypoints.

It calculates a score between 0 and 1, where 1 means that the predicted keypoints are in the same position as the ground truth keypoints and 0 means that the predicted keypoints are far from the ground truth keypoints. However, it relies on defining the variance for each keypoint, which defines how much can each keypoint vary from the ground truth keypoint. Figure 6.1 shows the effect of the variance in the OKS metric for different keypoints. This setting of variance is dataset dependent and can be hard to define for a new dataset. This is why the OKS metric is not used in the evaluation of the models on the bikefitting dataset.

6.1.3 Mean Average Precision (mAP)

In assessing the performance of pose estimation models, the most widely used metric is the Mean Average Precision (mAP) which is calculated as the average of the precision-recall curve at different thresholds of Object Keypoint Similarity (OKS). This metric is often used in large scale datasets, such as COCO [22], where the keypoints of interest are not always visible and therefore are hard to accurately predict. Moreover, these datasets are labeled by multiple annotators, which can lead to inconsistencies in the annotations. This metric intuitively tells us how many keypoints are in some proximity to the ground truth keypoints. However, in the case of the bikefitting dataset, where only few well visible keypoints are of interest, the mAP metric may not be the most suitable for evaluation, since all of the keypoints should be in close proximity to the ground truth keypoints and we are more interested in the exact error of the keypoints.

6.1.4 Normalized Mean Error (NME)

As was discussed in the previous sections, the commonly used metrics for evaluating pose estimation models, such as PCK, OKS, and mAP, may not be suitable for the bikefitting dataset. This is why a simple Normalized Mean Error (NME) metric is used for evaluation of the models on the bikefitting dataset.

Normalized Mean Error (NME) is calculated by obtaining the Euclidean distance between the predicted keypoints and the ground truth keypoints, and then normalizing this




Figure 6.1: Different variances in the OKS metric for different keypoints. The variance for the eye and wrist keypoints are different, leading to different keypoint similarities (0.25 and 0.75), even though the Euclidean distance to ground truth is the same for both keypoints. Taken from [31].

distance with respect to the average bounding box dimension. The normalization is done to account for the scale of the video. It is then expressed as a percentage for better interpretability. The formula for NME is as follows:

$$\text{NME}_i(\%) = \frac{d_i}{\frac{s.w + s.h}{2}} \cdot 100,$$

where:

NME_i: is the Normalized Mean Error for the i -th keypoint, expressed as a percentage of the average bounding box dimension

d_i : is the Euclidean distance between the ground truth and predicted keypoint i

$s.w$: is the width of the bounding box of the person

$s.h$: is the height of the bounding box of the person

The NME for keypoint i is then averaged over all keypoints to obtain the NME for a single frame. The final NME is then calculated over all frames in the dataset. NME offers easy to interpret results, with a lower value indicating a more accurate prediction.

6.2 Initial Evaluation of Pre-Trained Models

This section evaluates the performance of the pre-trained pose estimation models on the bikefitting dataset. This is done, because there can be a significant difference in the per-

formance of the models on COCO dataset and the bikefitting dataset. The models are evaluated using the NME metric, described in section 6.1.

The models are evaluated using the MMPose framework [10], with the Inferencer API and the PyTorch backend. Topdown models (models that require inputs cropped to the bounding box of the person) use the RTMDet-nano model [25] from the MMDetection framework [8] for person detection.

As was described in chapter 4, some models predict more keypoints than others. Most of the models are trained to predict 17 keypoints, as defined in the COCO dataset. The RTMPose, which are also trained on other datasets, including Halpe26, predict 26 keypoints. The WholeBody models predict 133 keypoints. The keypoints that are not needed for the evaluation are ignored. Since no models predict the fifth metatarsal of the foot, which is a crucial point for bikefitting, the small toe landmark is used instead.

6.2.1 6 Keypoint Evaluation

At first the models are evaluated using only the 6 keypoints that are both in the COCO dataset and the bikefitting dataset (camera facing ankle, knee, hip, shoulder, elbow and wrist). The results are shown in table 6.1. The best results are achieved by the RTMPose-m 384x288 model, but most of the bigger RTMPose models outperform the other models. This is in contrast to the results on the COCO dataset, where HRNet-W32 model performed similarly to the bigger RTMPose models. The relative decrease in performance of the HRNet-W32 model is probably caused by the fact that it was trained only on the COCO dataset, while the RTMPose models were trained on multiple datasets and are more robust to different poses and camera angles.

The best model achieves a mean NME of **2.12%**. This means that the average distance between the predicted keypoints and the ground truth keypoints is **2.12%** of the average bounding box dimension.

The good results of the RTMPose models are not surprising, since they were trained on multiple datasets, while the other models were trained only on the COCO dataset. The RTMPose models trained on the Body8 dataset perform slightly better than those trained on the Body8-Halpe26 dataset, which were trained to predict 26 keypoints. However, the models trained on COCO-WholeBody dataset perform significantly worse, this is probably because the WholeBody dataset is much smaller than the Body8 and standard COCO datasets.

The other lightweight models (YOLOX-Pose and ViPNAS-MobileNetV3) perform significantly worse than the RTMPose models. The YOLOX-Pose model is the worst performing model, with a mean NME of **2.74%**. The ViPNAS-MobileNetV3 model performs slightly better, but still worse than even the tiny RTMPose model.

Overall, we can see few unsurprising results. Models that are larger, or trained on multiple datasets, perform better than the smaller models trained only on the COCO dataset. It is interesting though, that there is not a huge difference between the large and medium RTMPose models, but there is a clear advantage to RTMPose models with larger input size. This is probably because the larger input size allows the model to capture more details in the image.

6.2.2 8 Keypoint Evaluation

The models which also output heel and foot fingers keypoints are evaluated on all of the 8 keypoints present in the bikefitting dataset. Meaning the camera facing fifth foot metatarsal


Architecture	Input Size	Trained On	Keypoints	NME (%)
RTMPose-m	384x288	Body8	17	2.12
RTMPose-l	384x288	Body8	17	2.13
RTMPose-m	256x192	Body8	17	2.22
RTMPose-m	384x288	Body8-Halpe26	26	2.23
RTMPose-l	256x192	Body8	17	2.27
RTMPose-m	256x192	Body8-Halpe26	26	2.28
HRNet-w32	256x192	COCO	17	2.34
RTMPose-l	384x288	Body8-Halpe26	26	2.42
RTMPose-l	256x192	Body8-Halpe26	26	2.43
RTMPose-s	256x192	Body8	17	2.52
RTMPose-s	256x192	Body8-Halpe26	26	2.56
RTMPose-l	256x192	COCO-WholeBody	133	2.59
RTMPose-m	256x192	COCO-WholeBody	133	2.61
RTMPose-t	256x192	Body8-Halpe26	26	2.68
SimCC_ViPNAS-MBNetv3	256x192	COCO	17	2.70
RTMPose-t	256x192	Body8	17	2.70
Yoloxpose_m	640x640	COCO	17	2.74

Table 6.1: Evaluation results on the custom bike fitting dataset, using only 6 keypoints most important to bikefitting (camera facing ankle, knee, hip, shoulder, elbow and wrist). The results are sorted by the Normalized Mean Error, defined in section 6.1.4. It is an average distance between the predicted keypoints and the ground truth keypoints, normalized by the average bounding box dimension, expressed as a percentage. The larger RTMPose models trained on the Body8 datasets outperform the other models.


and heel are also evaluated. The results are shown in table 6.2. We can see, that the Normalized Mean Error (NME) for the foot and heel keypoints is much higher than for the other keypoints. This highlights the difficulty of predicting the foot and heel keypoints, as these keypoints are less represented in the training datasets. The models that output the WholeBody keypoints perform slightly worse than the models that only output Halpe26 keypoints, but the difference is smaller than when comparing only the 6 keypoints. What is surprising, is that the COCO-WholeBody models perform better than the Halpe26 on the foot and heel keypoints. This can be because the COCO-WholeBody dataset keypoints may match the foot and heel keypoints defined in the bikefitting better than the Halpe26 dataset. However, overall the best performing models are still the RTMPose models trained on the Body8 dataset, with the RTMPose-m 384x288 model achieving the best results with a mean NME of **2.49%**.

Figure 6.2 shows example predictions of the models on the bikefitting dataset. The RTMPose-m 384x288 Body8 model achieves the best results. The RTMPose-m 384x288 Body8-Halpe model also achieves good results, and also predicts the heel and foot keypoints. The RTMPose-l 256x192 COCO-WholeBody model performs much worse on the foot and heel keypoints. The accuracy of the SimCC_ViPNAS-MBNetv3 is also worse than the RTMPose models.

The RTMPose models show good results on the bikefitting dataset. However, the accuracy on the foot and heel keypoints is still not satisfactory even for the largest models trained on the Body8-Halpe26 dataset. The results are even worse for the models trained




(a) RTMPose-m 384x288 Body8



(b) RTMPose-m 384x288 Body8-Halpe



(c) RTMPose-l 256x192 COCO-WholeBody



(d) SimCC_ViPNAS-MBNetv3

Figure 6.2: Example predictions of the models on the bikefitting dataset. Models (a) and (d) predict standard 17 COCO keypoints, while model (b) predicts 26 Halpe keypoints and model (c) predicts 133 WholeBody keypoints. The RTMPose-m 384x288 Body8 model (a) achieves the best results.

Architecture	Input Size	Trained On	Keypoints	Foot (%)	Heel (%)	Mean (%)
RTMPose-m	384x288	Body8-Halpe26	26	3.45	3.40	2.49
RTMPose-m	256x192	Body8-Halpe26	26	3.29	3.62	2.54
RTMPose-l	256x192	Body8-Halpe26	26	3.56	3.21	2.64
RTMPose-l	384x288	Body8-Halpe26	26	3.48	3.58	2.67
RTMPose-l	256x192	COCO-WholeBody	133	2.55	3.45	2.68
RTMPose-s	256x192	Body8-Halpe26	26	3.14	3.52	2.73
RTMPose-m	256x192	COCO-WholeBody	133	3.25	3.15	2.74
RTMPose-t	256x192	Body8-Halpe26	26	4.21	3.20	2.91

Table 6.2: Evaluation results for the models with extended keypoint set on the custom bike fitting dataset, including the heel and foot landmarks. Results for the other keypoints can be found in table 4.1. Foot column shows the Normalized Mean Error (NME) for the foot landmark, Heel column shows the NME for the heel landmark and Mean column shows the average NME for all of the 8 keypoints present in the bikefitting dataset. The values are represented as a percentage of the average bounding box dimension for a given frame. Details of the NME metric are described in section 6.1.4.

on the COCO-WholeBody dataset. To improve the accuracy of the models on the foot and heel keypoints, the models need to be fine-tuned on the bikefitting dataset. The details of the fine-tuning are described in the next section.

6.3 Fine-Tuning On Custom Bikefit Dataset

For the fine-tuning, the RTMPose-m 256x192 Body8-Halpe26, RTMPose-l 256x192 Body8-Halpe26 and RTMPose-m 384x288 Body8-Halpe26 models are selected. These models are selected, because they achieved good results on the bikefitting dataset, predict the heel and foot keypoints and are still relatively fast. The models are fine-tuned on the bikefitting dataset using the MMPose framework [10].

To train a model using the MMPose framework, the model needs to have a custom config file for the dataset. The config file contains the paths to the dataset, the model architecture, the optimizer, the learning rate scheduler and other hyperparameters. The original config files for the RTMPose models were used as a starting point and modified to use the bikefitting dataset different hyperparameters and evaluation metrics.

Custom datasets in MMPose are easiest to define using the COCO format. The dataset is defined in a JSON file, which contains the paths to the images, the annotations and the categories. The annotations contain the bounding box of the person and the keypoints. The keypoints are defined as a list of x and y coordinates and a visibility flag. The visibility flag is set to 2 if the keypoint is not visible, 1 if the keypoint is occluded and 0 if the keypoint is visible. MMPose internally uses only annotated keypoints (visibility flag 2 or 1) for training and evaluation and discards keypoints with visibility flag 0. The training set has all of the keypoints defined as annotated (visibility flags 1 or 2), so the models are trained on all of the keypoints. This ensures that the models do not unlearn the other keypoints. The validation set has only the 8 keypoints annotated by colored markers defined as annotated, so the models are evaluated only on these keypoints.

6.3.1 Training Procedure

The RTMPose models originally used the AdamW [24] optimizer with base learning rate of $4e-3$, linear warmup learning rate scheduler for first 1000 iterations and cosine annealing learning rate scheduler. The models were in total trained for 700 epochs with training batch size of 512. These are tuned for a large scale dataset like COCO. For the bikefitting dataset, smaller learning rate and batch size are used. The models are trained with the AdamW optimizer with base learning rate of $5e-4$, no learning rate schedulerers were used as this slowed down the training. The batch size was reduced to 16, as this sped up the training significantly. The models were trained for 130 iterations, which corresponds to 1 epoch on the bikefitting dataset. This low number of iterations was chosen, so that the models do not overfit on the bikefitting dataset. Using larger batch sizes of 32 and 64 resulted in much slower training, where the models needed 5-10 epochs of training and still achieved worse results then when using batch size of 16. The models were fast enough to be trained on an single Intel i7-7500U CPU.

Figure 6.3 shows the loss of the models during the fine-tuning. The RTMPose large and medium models with 256x192 input size achieve basically the same loss throughout the training. The RTMPose-m 384x288 model has significantly lower loss than the other models throughout the whole training. Since its validation results are on par with the 256x192 models (as discussed in the next section), it is probably because different parameters are used to construct the ground truth heatmaps for the 384x288 model, which results in lower loss. Overall the loss of the models is decreasing throughout the training, at first rapidly and then from step 60 forward more slowly.

6.4 Evaluation Of Fine-Tuned Models

Table 6.3 shows the results for the models before and after fine-tuning on the bikefitting dataset. The models are evaluated using the NME metric, described in section 6.1. All of the fine-tuned models achieve significantly better results than the pre-trained models. Despite the differences in the performance of medium and large pre-trained models, all of the fine-tuned models achieve similar results. The best results are achieved by the RTMPose-m 384x288 model, which achieves a mean NME of **1.50%**, but the other fine-tuned models achieve very similar results. The fine-tuned models achieve a significant improvement in the accuracy of the foot and heel keypoints. The foot and heel keypoints are still the most difficult to predict, but the fine-tuned models achieve much better results than the pre-trained models. Because there is not a significant difference in the performance of the fine-tuned models, the RTMPose-m 256x192 model is selected for the bikefit application, as it is the smallest and fastest model that achieves good results.

Figure 6.4 shows example predictions of the RTMPose-m 256x192 model before and after fine-tuning. The fine-tuned model achieves much better results on the foot and heel keypoints. It also shifts the knee and shoulder keypoints more closely to the location which are used for example in the Retul bikefitting system. However, it also shifts the hip keypoint closer to the knee keypoint, which is not desirable. Similarly, the wrist keypoint is now no longer predicted on the center of rotation of the wrist joint, but more upwards. These changes are probably caused by misplacement of the colored markers, which is especially hard to avoid when muliple poses and camera angles are present in the dataset. Nevertheless, the fine-tuned model overall achieves better results than the pre-trained model and is used in the bikefit application.





Figure 6.3: Loss of the RTMPose models during fine-tuning on the bikefitting dataset. The models were trained for 130 iterations, which corresponds to 1 epoch on the bikefitting dataset. For optimization, AdamW optimizer with base learning rate of $5e-4$ and batch size of 16.

Architecture	Input Size	Trained On	Keypoints	Foot (%)	Heel (%)	Mean (%)
RTMPose-m	384x288	Bikefit dataset	26	1.25	2.11	1.50
RTMPose-m	256x192	Bikefit dataset	26	1.45	1.99	1.51
RTMPose-l	256x192	Bikefit dataset	26	1.26	1.94	1.52
RTMPose-m	384x288	Body8-Halpe26	26	3.45	3.40	2.49
RTMPose-m	256x192	Body8-Halpe26	26	3.29	3.62	2.54
RTMPose-l	256x192	Body8-Halpe26	26	3.56	3.21	2.64

Table 6.3: Evaluation results for the models before and after fine-tuning on the custom bike fitting dataset. The models are fine-tuned on the bikefitting dataset for 130 iterations. The results are sorted by the Normalized Mean Error, defined in section 6.1.4. It is an average distance between the predicted keypoints and the ground truth keypoints, normalized by the average bounding box dimension, expressed as a percentage. All of the fine-tuned models achieve significantly better results than the pre-trained models.



(a) RTMPose-m 256x192 Body8-Halpe26



(b) RTMPose-m 256x192 Body8-Halpe26 Fine-Tuned

Figure 6.4: Example predictions of the RTMPose-m 256x192 model before and after fine-tuning on the bikefitting dataset. The fine-tuned model achieves much better results on the foot and heel keypoints, but also undesirably shifts the hip and wrist keypoints, probably due to misplacement of the colored markers.

Chapter 7

Bikefit Web Application

This chapter describes the web application for bike fitting. Implementing the application as a web application has several advantages over implementing it as a mobile or desktop application. The web application can be used on any device with a web browser. This allows the application to be used on mobile devices, tablets, laptops and desktop computers. The user also does not need to manually download and install anything.

Most web applications are implemented using the client-server architecture with the video processing and other more complicated tasks (such as pose estimation) being performed on the server. This has several disadvantages. The server needs to be powerful enough to handle the load of multiple users. This can be quite expensive for heavy models and a lot of users. The user data also needs to be sent to the server, which means the user needs to be online when using the application and the users can be concerned about their privacy.

To avoid these disadvantages, all of the logic of the application (including video processing, keypoint estimation and generation of the suggestions) is implemented on the client side.

The application is implemented using the Svelte framework¹. Svelte is a component framework similar to React² or Vue³. It is used to create reactive user interfaces.

The application consists of several parts. The user can upload a video of the cyclist from the side view. The video is then processed frame by frame by the pose estimation model. The keypoint predictions are then used to compute the joint angles. The joint angles are used to generate the suggestions for the bike fit. The suggestions are displayed to the user. The application from the user perspective is more thoroughly described in section 7.1.

The following sections describe the implementation of the application in more detail. To process the video frame by frame, the application uses the WebCodecs API. The details of the implementation are described in section 7.2. To estimate the keypoint positions, the application uses the RTMPose-m Body8-Halpe26 pose estimation model, which was fine-tuned on the bikefitting dataset. The details of the deployment and implementation are described in section 7.3. From the keypoint positions, the application computes various joint angles and measurements. The details of the implementation are described in section 7.4. The application then uses the computed joint angles and measurements to generate the suggestions for the bike fit. The details of the implementation are described in section 7.5. Lastly, to play the video with the keypoint predictions overlaid on the video, a custom

¹<https://svelte.dev/>

²<https://react.dev/>

³<https://vuejs.org/>

canvas based video player is implemented. The details of the implementation are described in section 7.6.

7.1 User Interface

From the user's perspective, the application consists of three main steps: video upload, select fit goals and examine suggestions. The user is guided through these steps by a wizard-like interface. The following sections describe each of the steps in more detail.

7.1.1 Video Upload

The first step is the video upload. The user is prompted to upload a video of the cyclist from the side view. The user can also choose to use a demo video for testing the application. After the video is uploaded, the user can proceed to the next step. Figure 7.1 shows the first step of the application.

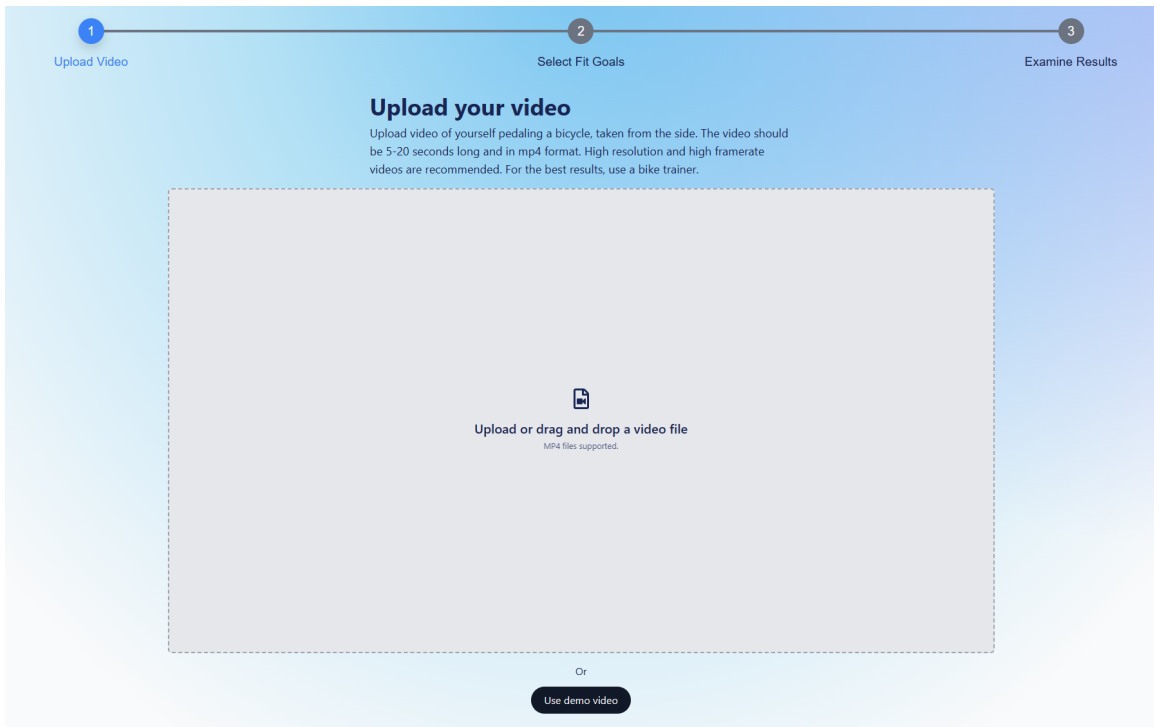


Figure 7.1: The first step of the bikefit application. The user is prompted to upload a video of the cyclist from the side view. The user can also choose to use a demo video for testing the application. After the video is uploaded, the user can proceed to the next step.

7.1.2 Select Fit Goals

On the second step, shown in figure 7.2, the user can review and play the uploaded video and is asked to select the type of the bike (road, gravel and mountain bike) and the fit goals (comfort, performance, balanced). These fit goals and the type of the bike are used to adjust the suggestions for the bike fit. After the fit goals are selected, the user can process

the video by the pose estimation model by clicking on the blue „Process Video“ button on the right.

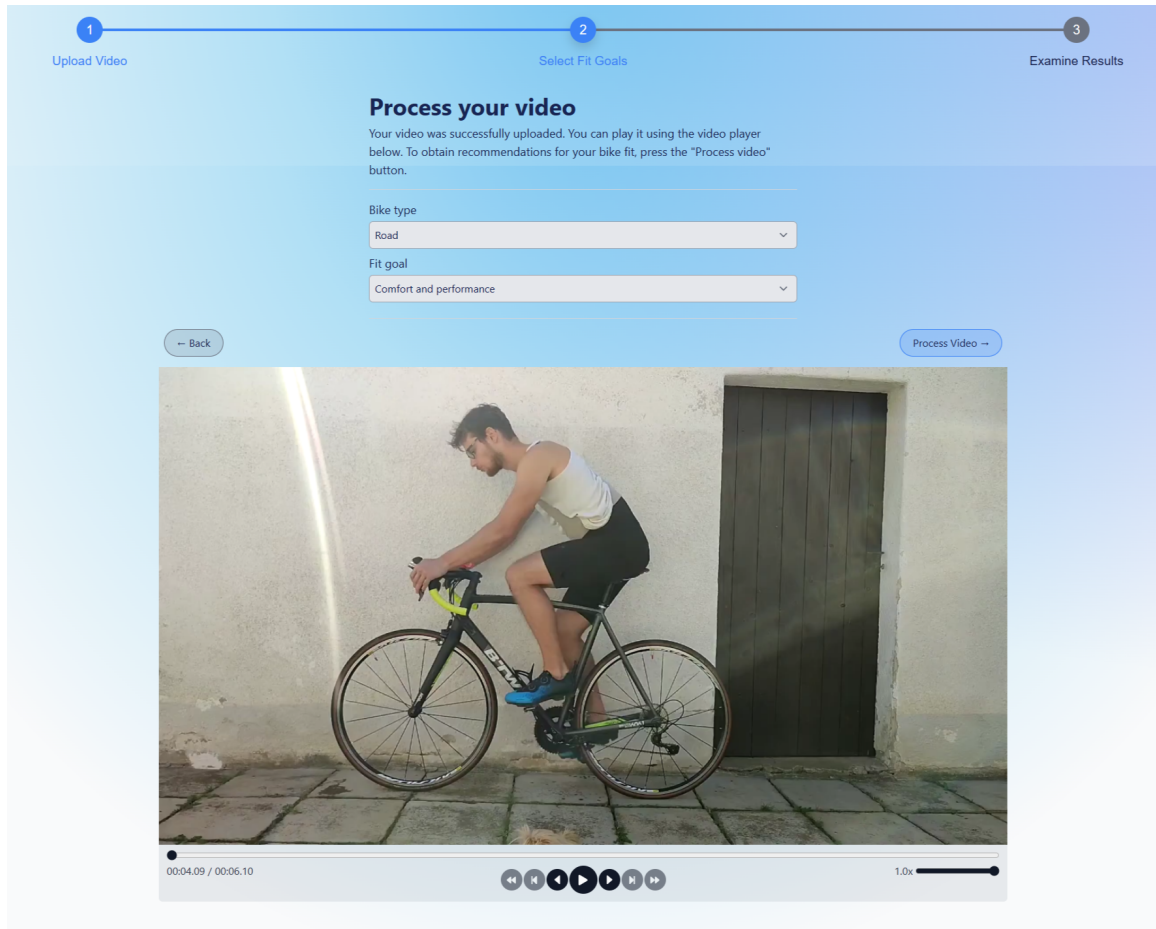


Figure 7.2: The second step of the bikefit application. The user can review and play the uploaded video and is asked to select the type of the bike (road, gravel and mountain bike) and the fit goals (comfort, performance, balanced). After the fit goals are selected, the user can process the video by the pose estimation model. The video speed can be adjusted by the slider on the right. The user can also seek forward and backward 1 frame by clicking the buttons on the left and right of the middle play button.

After the user clicks on the „Process Video“ button, the video is processed by the pose estimation model and the user is shown a progress circle in the middle of the screen. Figure 7.3 shows the progress circle shown to the user while the video is being processed.

7.1.3 Examine Suggestions

On the last step, after the video is processed, the user can examine the keypoint predictions from the pose estimation model, computed joint angles, suggested angle ranges and the suggestions for the bike fit. The user can also play the video with the keypoint predictions overlaid on the video. Figure 7.4 shows the first part of this step. The user is presented with a video player with the keypoint predictions and measured angles overlaid on the video.

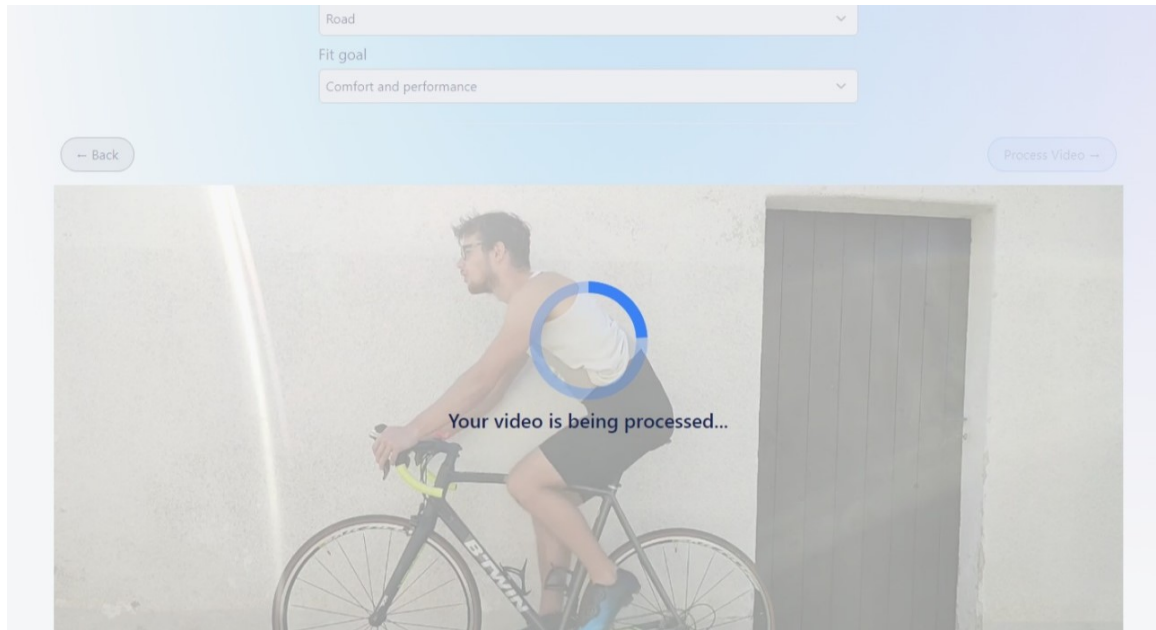


Figure 7.3: After the user selects the fit goals, the video is processed by the pose estimation model. The user can see the progress of the processing by the progress circle in the middle of the screen.

Figure 7.5 shows the second part of the last step of the application. The „Recommendations“ section shows the bike type and fit goal selected by the user. The user can also see the assessment and suggested changes to their saddle height, saddle setback, handlebar height and handlebar reach. For each measurement, the user sees a verbal assesment of a measurement (e.g. „Too high“, „Too far back“, „Withing range“, etc.) a graphical representation of the measurement by a colored slider input and a verbal suggestion for the change (e.g. „Lower by 1 cm“, „Move saddle forward by 1 cm“, etc.). The „Results“ section shows the recommended ranges and measured values for the angles and measurements in each of the pedal positions (bottom, top and front). The measurements are organized in a table-like structure, with columns for the measurement name, the measured value, a colored slider input indicating the measurement value and the recommended range. Additionally, the user can click on the visibility indicator on the left to hide or show the measurement in the video player. After clicking on a row in the table, the row expands to show more detailed information about the measurement. The user can see a chart with the angle values for each pedal stroke, with sections of the chart colored according to the recommended ranges (green for within range, yellow for close to end of the range and red for outside of the range). The user can also see the minimum, maximum, median and range of the measurement. In the future, additional information about the measurement and its effect on the bike fit could be added.

7.2 Video Upload and Frames Extraction

At first, the video needs to be uploaded to the application. The video is uploaded using the HTML5 file input. The file is then converted to a URL using the `URL.createObjectURL()` method, because the WebCodecs API requires a URL as input.

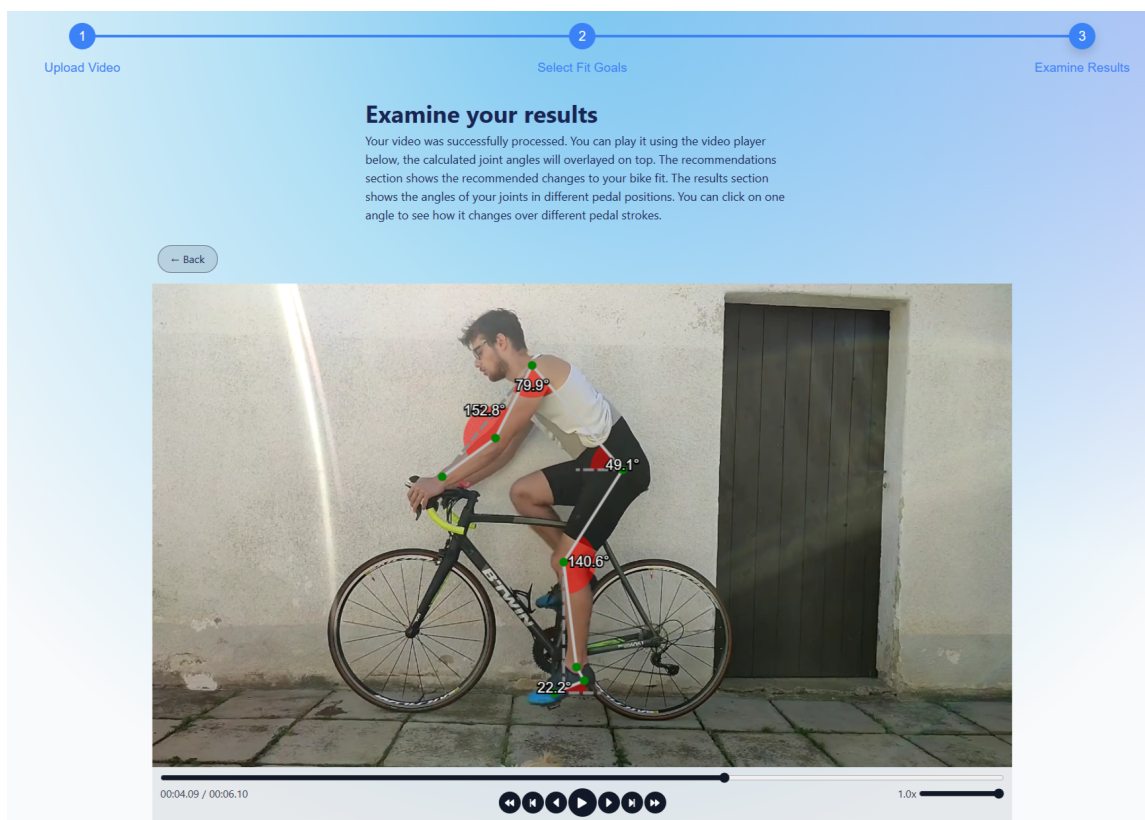


Figure 7.4: The first part of the last step of the bikefit application. The user can play the video with the keypoint predictions and measured angles overlaid on the video. For each angle, the user sees the current value and a sector of a circle connecting angle points. The video speed can be adjusted by the slider on the right. The user can also seek forward and backward 1 frame by clicking the buttons on the left and right of the middle play button. The buttons on the side can be used to seek to the next or previous pedal position and next or previous pedal stroke.

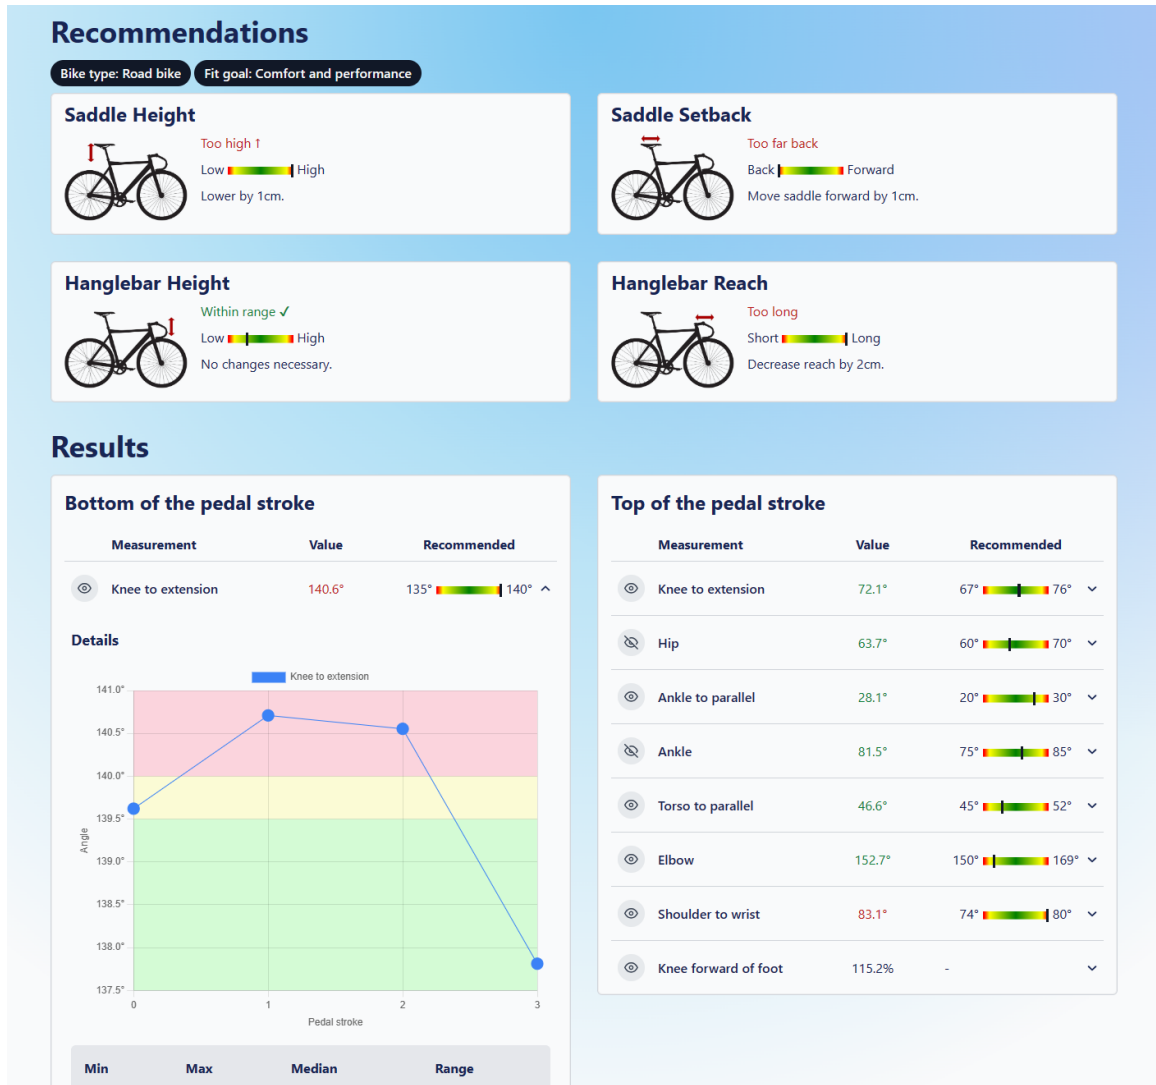


Figure 7.5: The second part of the last step of the bikefit application. In the „Recommendations“ section, the user can see the suggested changes to their saddle height, saddle setback, handlebar height and handlebar reach. In the „Results“ section, the user can see the recommended ranges and measured values for the angles and measurements in each of the pedal positions (bottom, top and front (not shown)). After clicking on a measurement, the user can see a chart with the angle values for each pedal stroke, with sections of the chart colored according to the recommended ranges.

7.2.1 Accessing the Video Frames

To accurately estimate the keypoints, the application needs to access all of the video frames and process them. This is not trivial to do in the browser. The browser does not allow the application to access the video frames directly.

To process the video frame by frame, the application uses the WebCodecs API⁴. The WebCodecs API is a low-level API for encoding and decoding audio and video. It gives developers access to the individual frames of the video. It is the only way to access the frames of the video in the browser. Simple and commonly used method is to simultaneously play the video and draw the video frames to a canvas element and then access the pixels of the canvas. However, this method does not work well if the processing is computationally expensive, since the video is played in real time, and callbacks to process the frames are called only at best effort to keep up with the video playback. This leads to skipping of the frames and not all of the frames are processed. The WebCodecs API allows the application to process the video at its own pace.

Because the API is relatively new, it is not yet supported by all browsers. Data from CanIUse.com⁵ show that the API is currently supported by all browsers based on Chromium (Google Chrome, Microsoft Edge, Opera, Brave, etc.). It is Partially supported by Safari and Safari on iOS (video only) and not supported by Firefox. Overall, the API is supported by 87.6% of the users, where 74.5% of the users have full support and 13.1% of the users have partial support.

Due to the low-level design of the WebCodecs API, it is not easy to use. Moreover, the API is quite new and not very well documented so there is not a lot of examples. The developer also needs to take care of demuxing the media containers. This work uses the `getVideoFrames.js` library⁶ from josephrocca. The library internally uses the WebCodecs API and `MP4Box.js`⁷ for demuxing mp4 files, but provides a simple interface for accessing the video frames. It defines three callback functions. The `onFrame` callback is called for each frame of the video. The `onConfig` callback is called when the video metadata, which contains the total number of frames and other useful information, is loaded. The `onFinish` callback is called when all of the frames are processed. Overall, WebCodecs API is a powerful tool for processing video in the browser, but it relies on third party demuxing libraries and is not yet supported by all browsers.

7.3 Pose Estimation

To get the locations of the 8 keypoints (foot, heel, ankle, knee, hip, shoulder, elbow, wrist) in the video, the RTMPose-m 256x192 Body8-Halpe26 model is used. The model was trained on the custom bikefitting dataset, as described in chapter 6.3. The model is run on each frame of the video and the keypoints are extracted.

Since JavaScript is single-threaded and the model is computationally expensive, the model is run in a Web Worker. Web Worker is a JavaScript feature that allows the application to run a script in a separate thread. This allows the application to run long tasks in the background without blocking the main thread.

⁴https://developer.mozilla.org/en-US/docs/Web/API/WebCodecs_API

⁵<https://caniuse.com/webcodecs>

⁶<https://github.com/josephrocca/getVideoFrames.js>

⁷<https://github.com/gpac/mp4box.js/>

The keypoints are then used to compute the joint angles and generate the suggestions for the bike fit.

7.3.1 Model Conversion and Deployment

The RTMPose models are originally stored in the PyTorch `.pt` format. To run the model in the browser, the model needs to be converted to either ONNX or TensorFlow.js format. ONNX model can be run in the browser using ONNX Runtime Web⁸. However, the ONNX Runtime Web only provides API for running the model, but does not provide tensor operations needed for preprocessing the input and postprocessing the output. Due to these reasons, TensorFlow.js is used instead.

To convert the model to TensorFlow.js format, the model is first converted to ONNX format using the MMDeploy [9] toolset. The MMDeploy toolset is a set of tools for converting models from the various OpenMMLab frameworks (MMPose, MMAAction, MMDetection, etc.) to ONNX and other formats and provides an easy way to deploy the RTMPose models.

The ONNX model is then converted to the TensorFlow SavedModel format using the unofficial `onnx2tf` [15] tool. The SavedModel format is a universal format for storing TensorFlow models. The SavedModel format is then converted to the TensorFlow.js format using the official TensorFlow.js converter⁹ provided by the TensorFlow.js team.

The TensorFlow.js converter also quantizes the model to float16 to reduce the size of the model and speed up the inference. Quantization is a process of converting the weights and activations of the model from float32 to float16. This reduces the size of the model by half and can speed up the inference without significant loss of accuracy. The quantized model is then used in the application.

TensorFlow.js provides several different backends for running the model. CPU backend runs the model on CPU, using JavaScript for the computation and is the slowest, but has support for all of the operations. WebGL backend runs the model on the GPU using WebGL shaders and is the faster, but introduces very noticeable latency when the model is first run, because the WebGL API is not designed for running deep learning models and the model needs to be compiled to WebGL graphical shaders. WebAssembly backend runs the model on the CPU using WebAssembly, with support for SIMD instructions and multi-threading and is the faster than the CPU backend, but slower than WebGL for larger models. WebGPU backend runs the model on the GPU and should be the fastest, because it is designed with deep learning in mind, but is not yet supported by all browsers. According to CanIUse.com¹⁰, WebGPU is only supported by Google Chrome and Chromium based browsers as well as Chrome on Android, which combines for 70.5% of the users. This application uses the WebGPU backend if it is available, otherwise it falls back to the WebGL backend, which is universally supported.

7.3.2 Model Download and Caching

To run the pose estimation model, it first needs to be downloaded to the user's device. The converted RTMPose model takes up 26.6 MB of space. The COCO-SSD object detection model requires a download of 18.1 MB. This amounts to a total of 44.7 MB of data that needs to be downloaded. The models are then cached and do not need to be downloaded again for the next session.

⁸<https://onnxruntime.ai/docs/tutorials/web/>

⁹<https://github.com/tensorflow/tfjs/tree/master/tfjs-converter>

¹⁰<https://caniuse.com/webgpu>

7.3.3 Preprocessing

The RTMPose model is a top-down pose estimation model, which means that the model requires the input image to be cropped around the person. To crop the image, the bounding box of the person is first detected using the COCO-SSD¹¹ model, available in the TensorFlow.js model zoo. It is a Single Shot Detector (SSD) [23] model trained on the COCO [22] dataset. To speed up processing the detector model is ran only once on the first frame of the video, because the person is assumed to be in the same position in all of the frames.

The bounding box of the person is then expanded by 2.5% on the left and right and by 20% on the bottom to account for changing legs position of the user due to pedaling. The bounding box is then used to crop the image. The cropped image is then resized to 256x192 pixels, which is the input size of the RTMPose model and normalized by the mean and standard deviation of the Body8 dataset (mean = [123.675, 116.28, 103.53], std = [58.395, 57.12, 57.375]).

The COCO-SSD is not the most accurate or the fastest model, but it is readily available in the TensorFlow.js model zoo and is good enough for using it on the first frame of the video.

7.3.4 Postprocessing

The RTMPose models output a SimCC [20] heatmap, which is a heatmap of the keypoints where the x and y axis are separated. This significantly reduces the number of parameters of the model and speeds up the inference. For each keypoint, the model outputs a list of scores for x and y axis. To find the location of the keypoint, the argmax of x and y axis is taken. The x and y coordinates are then divided by 2 because the models output the coordinates multiplied by 2 to increase accuracy. The x and y coordinates are then adjusted for the resizing and added to the bounding box coordinates to get the coordinates in the original image. Only 8 body keypoints on the camera facing side (foot, heel, ankle, knee, hip, shoulder, elbow, wrist) are used further, the other keypoints are discarded. The keypoints over all of the frames are then smoothed using a Savitzky-Golay filter [32] to reduce the noise in the predictions.

7.4 Computing Joint Angles and Measurements

After the location of the 8 keypoints is estimated, the joint angles and measurements for the current frame are computed. Joint angles are specified either by three keypoints if the angle is computed between two bones or by two keypoint if the angle is computed between one bone and the horizontal axis.

Other than joint angles, application also computes the knee forward of foot measurement, which is the distance between the knee and the vertical line going through the fifth metatarsal of the foot. Due to the absence of real world units in the video, the measurement is expressed as the percentage of the length from the fifth metatarsal to the heel. This measurement is used to determine saddle setback as its very similar to the widely used knee over pedal spindle (KOPS) measurement.

¹¹<https://github.com/tensorflow/tfjs-models/tree/master/coco-ssd>

7.4.1 Determining the Pedal Position

For measurements that barely change over the pedal stroke (mainly upper body joint angles), the average value over all of the frames can be used. But for measurements that change over the pedal stroke (lower body measurements such as knee angle, knee forward of foot, ankle angle), the measurements are compared to the recommended ranges in a specific pedal position (bottom, top, front). To do that, the current pedal position needs to be determined for every frame, and frames for each pedal position need to be determined.

To determine the pedal position for each frame, individual pedal strokes need to be detected. To do that, a circle is fitted to the locations of the foot keypoint over all of the frames using a least squares method. This circle is then used to determine the pedal position. For example, to determine the front pedal position, the angle of the foot keypoint on the circle is computed for each frame in a pedal stroke. The front pedal position is then determined as the frame with the angle closest to 90 degrees (horizontal forward). Similarly, the bottom pedal position is determined as the frame with the angle closest to 180 degrees (vertical down) and the top pedal position is determined as the frame with the angle closest to 0 degrees or 360 degrees (vertical up). After obtaining the list of frames for each pedal position, and list of frames for each pedal stroke, we can finally compute the measurements and joint angles for each pedal position. To do this, median value of the measurements for each pedal position is used. This is done to avoid the influence of the outliers in the measurements.

7.5 Generating Suggestions

After the joint angles and measurements are computed for each pedal position, the suggestions for the bike fit are generated. The suggestions are generated based on comparing some of the measurements in the pedal positions to the recommended ranges. The recommended ranges are based on the bike fit guidelines and are then adjusted based on the type of the bike and the fit goals selected by the user. For example, saddle height recommendations are based on knee angle at the bottom of the pedal stroke. If the knee angle is inside the recommended range, saddle height is considered correct. If the knee angle at the bottom of the pedal stroke is at the lower 10 percent of the recommended range, a suggestion to increase the saddle height slightly (by 1 cm) is generated. If the knee angle is lower than the recommended range, a suggestion to increase the saddle height significantly (by 2 cm) is generated. Suggestions are made similarly when the knee angle is at the upper 10 percent of the recommended range or higher than the recommended range.

In total, suggestions are generated for the saddle height, saddle setback, handlebar height and handlebar reach. As was mentioned before, saddle height is determined by the median knee angle at the bottom of the pedal stroke. Saddle setback is determined by the median knee forward of foot measurement at the front of the pedal stroke. Handlebar height is determined by the median shoulder angle at all of the frames and handlebar reach is determined by the median shoulder angle at all of the frames.

The recommended ranges that are used to generate the suggestions are shifted by the selected bike type and fit goal. For example, the recommended back angle will be increased for gravel and mountain bikes to recommend a more upright position. Similarly, the comfort fit goal also prioritize a more upright position. Conversely, the performance fit goal will prioritize more aerodynamic position with a lower back angle, higher shoulder angle and a

higher knee angle at the bottom of the pedal stroke, resulting in a slightly higher saddle height and overall more aggressive position.

Basing suggestions only on a single measurement is not ideal, because in reality the bike fit is a complex process and the measurements are interdependent. For example, increasing the saddle height will also increase the saddle setback, handlebar reach also influences back angle not only shoulder angle and so on. In the future, the suggestions could be generated using a more complex model, which would take into account all of the measurements and their interdependencies.

7.6 Video Player

To play the video with the predicted keypoints and measurements overlaid on the video, the application implements a custom video player based on the HTML canvas and video elements. The video is played on a hidden video element and the frames are drawn to a canvas element. This is done so that joint angles and measurements can be drawn on top of the video frames. The video player allows the user to play the video, seek to a specific time, seek to the next or previous pedal position and seek to the next or previous pedal stroke. The video player also allows the user to adjust the speed of the video.

However, there does not exist a standard way to get the index of the current frame of the video, which is needed to get the current joint angles and measurements and draw them on the video. The video element has a `currentTime` property, which returns the current time of the video in seconds, but there is no property that returns the current frame index, or even the video frames per second. Another reason why relying on the `currentTime` property is not ideal is that the `currentTime` property returns the time of the video but the video can be longer than the video track, also the video track can have variable frame rate.

This is why, to get the current frame index, the application saves the frame timestamps obtained when processing the video with the WebCodecs API. And then uses the `requestVideoFrameCallback`¹² to get the current frame index. The `requestVideoFrameCallback` is a new API that allows the application to accurately get the current video track time using the `metadata.mediaTime` property. The `metadata.mediaTime` property returns the current time of the video track in seconds. The application then finds the closest frame timestamp to the current video track time and uses the index of the frame timestamp as the current frame index. This way the application can accurately get the current frame index of the video.

¹²<https://developer.mozilla.org/en-US/docs/Web/API/HTMLVideoElement/requestVideoFrameCallback>

Chapter 8

Web Application Experiments

This chapter describes the experiments done to evaluate the bikefitting web application. The experiments are divided into two parts. Section 8.1 tests the application's ability to assess and make suggestions for saddle height. Section 8.2 measures how long the application takes to process one frame of the video on different devices.

The application was also confirmed to work on Google Chrome and Chromium based browsers on a desktop computer with dedicated GPU, on a laptop with integrated GPU, on a MacBook with M1 Chip and on an android smartphone. The application also ran on Safari on a MacBook. The application unfortunately did not run on an Iphone, when using Safari or Chrome. Due to lack of console logs, it is not clear what is the cause of the problem. The application also did not run on Firefox, because the WebCodecs API is not supported by Firefox.

8.1 Testing Multiple Saddle Heights

To test the typical use case of the application, the application is used to set up the saddle height of a cyclist. The author filmed 6 videos of himself cycling on a road bike in a stationary trainer with various saddle heights. The saddle heights ranged from 3.5 cm lower than the usual saddle height to 2.5 cm higher than the usual saddle height. The goal of this test is to see if the measured knee angle at the bottom of the pedal stroke, which is the main measurement used to determine the saddle height, correctly increases with the saddle height and if the application recommended angles correspond to the feelings of the author about the saddle heights.

The videos were filmed outside in daylight with a resolution of 1920x1080 pixels and a frame rate of 30 frames per second using a smartphone. The author has experience riding a road bike and is moderately fit, but is not a professional cyclist. Therefore the balanced fit goal is selected for the test. The type of the bike is set to road bike. This configuration corresponds to recommended knee angle at the bottom of the pedal stroke of 135° to 140°.

Table 8.1 shows the results of the test. The measured knee angle at the bottom of the pedal stroke, the recommended knee angle and the feelings of the author about the saddle height are shown for each video. The results show that the measured knee angle at the bottom of the pedal stroke increases with the saddle height, as expected. The measured knee angle is within the recommended range of 135° to 140° for the saddle heights of 0 cm -0.5 cm and -1.5 cm. This surprisingly well correspond to the feelings of the author about the saddle height. The extreme saddle heights of -3.5 cm and +2.5 cm are correctly

assessed as too low and too high, respectively. The saddle height of -1.5cm felt low, but much better than the previous height of -3.5 cm, and could be used for beginners or for riders with some flexibility issues. The saddle height of -0.5 cm felt lower than the author is used to, but again, better than the previous height of -1.5 cm and could definitely be used long term. The saddle height of 0 cm felt ok and the saddle height of +0.5 cm felt distinctly too high, it was a small change, but very noticeable and the author had trouble to pedal smoothly with this saddle height.

Saddle Height	Knee Angle	Rider Feelings	App Assesment
-3.5 cm	128.5°	Too low.	Too low. Raise saddle by 1 cm and try again.
-1.5 cm	135.2°	Low end of rideable window.	On lower end of range. Try raising by 5 mm and try again.
-0.5 cm	137.5°	Lower than used to, but ok.	Within range.
0 cm	138.8°	Ok.	Within range.
+0.5 cm	141.3°	Too high. Not comfortable.	Too high. Lower saddle by 1 cm and try again.
+2.5 cm	152.3°	Definetely too high.	Too high. Lower saddle by 1 cm and try again.

Table 8.1: Results of the saddle height setup test. The measured knee angle at the bottom of the pedal stroke, the recommended knee angle and the feelings of the author about the saddle height are shown for each video.

Figure 8.1 shows screenshots of the videos with saddle heights of -3.5 cm and +2.5 cm. We can see how the knee angle at the bottom of the pedal stroke changes with the saddle height. The keypoint predictions and angles are overlaid on the video.

The test showed, that even small changes in the saddle height have effect on the knee angle at the bottom of the pedal stroke measured by the application. It also showed that the recommended knee angles are well chosen and correspond to the feelings of the author about the saddle height. It also corresponds to bikefit literature, which states that having a slightly lower saddle height is better than having a slightly higher saddle height. The test also showed that the application can be used to set up the saddle height of a cyclist. In the future, similar tests could be done for other measurements and joint angles to further validate the application and adjust the recommended ranges.

8.2 Video Processing Benchmark

To test the speed of the pose estimation model inference and the overall processing pipeline, the time needed to process a video is measured. The benchmark is run on a desktop computer with an NVIDIA GeForce GTX 950 GPU, using Microsoft Edge browser and WebGPU backend and on a Xiaomi Redmi Note 10 Pro with Adreno 618 GPU, using Google Chrome browser and WebGL backend.



Figure 8.1: Screenshots of the videos with saddle heights of -3.5 cm (left) and +2.5 cm (right) with the keypoint predictions and angles overlaid on the video. Both screenshots show the knee angle at the bottom of the pedal stroke.

Table 8.2 shows the time needed to process a video frame by frame on both devices. The results show that the WebGPU backend on desktop is around 2 times faster than the WebGL backend on the smartphone. Another interesting observation is on WebGPU backend, the inference time of the model is only a small part of the total frame time, while on the WebGL backend, the inference time is the majority of the total frame time. This is probably due to the fact that the WebGPU backend is more optimized for deep learning models, but has slow synchronous transfer of the data between the CPU and GPU. Another reason for long frame time on the WebGPU compared to inference time is probably not optimal implementation of the application, which could be improved in the future.

Device	Backend	Inference Time	Model Time	Frame Time
Desktop	WebGPU	45-70 ms	85-120 ms	160-200 ms
Smartphone	WebGL	220-300 ms	250-320ms	325-400 ms

Table 8.2: Time needed to process a video frame by frame on a desktop computer with an NVIDIA GeForce GTX 950 GPU and on a Xiaomi Redmi Note 10 Pro with Adreno 618 GPU. Inference time is the time needed to run the model inference on a single frame. Model time is the time needed to run the model on a single frame, including the preprocessing and postprocessing. Frame time is the time needed to process a single frame, including the model time and time needed to decode the frame.

Chapter 9

Conclusion

This work focused on the development of a web application for bike fitting using pose estimation. The application was successfully implemented and deployed¹.

The first chapter of the work described the needed bikefitting background and existing bikefitting applications and systems and their limitations. Then the following chapters described a marker-based method for tracking the cyclist's body and marker-less pose estimation models. Next, a custom dataset for bikefitting featuring 2309 photos of the author pedaling on a stationary trainer was created. The dataset was used to evaluate the pose estimation models and fine-tune the RTMPose model, whose Normalized Mean Error on the bikefitting dataset was significantly reduced from 2.54 to 1.51. Next chapter describes the implementation of the web application and the deployment process of the pose estimation model. The last chapter describes the experiments done to evaluate the application.

The resulting application, which runs fully in the browser, is able to accurately estimate the keypoint positions of the cyclist from the side view video and compute the joint angles and measurements. The application then generates suggestions for the bike fit based on the computed joint angles and recommended ranges. The application also allows more advanced users to examine the keypoint predictions and joint angles in the video player and see how the angles change over multiple pedal strokes. The experiments showed that that the application detects angle changes caused by the saddle height changes and that the application can be used to set up the saddle height. Benchmarks also showed that an older desktop computer can process a frame in 160-200 ms and a modern smartphone can process a frame in 325-400 ms.

In the future, the model accuracy could be further improved by fine-tuning it on a larger dataset, perhaps obtained by annotating the videos uploaded by the users. Using a model which also predicts hand keypoints could be used to also generate suggestions based on the wrist angle. The user experience could be further improved by creating a landing page for the application, or by connecting the application to a database of users to store the videos and the results of the bike fit. Additionally, suggestions for fit changes could be generated by taking into account more measurements and their interdependencies.

¹<https://vlasakjiri.github.io/bikefit/>

Bibliography

- [1] ANDRILUKA, M., IQBAL, U., INSAFUTDINOV, E., PISHCHULIN, L., MILAN, A. et al. PoseTrack: A Benchmark for Human Pose Estimation and Tracking. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2018.
- [2] ANDRILUKA, M., PISHCHULIN, L., GEHLER, P. and SCHIELE, B. 2D Human Pose Estimation: New Benchmark and State of the Art Analysis. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2014.
- [3] BERTALMIO, M., BERTOZZI, A. and SAPIRO, G. Navier-stokes, fluid dynamics, and image and video inpainting. In: *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*. 2001, vol. 1, p. I-I. DOI: 10.1109/CVPR.2001.990497.
- [4] BEWLEY, A., GE, Z., OTT, L., RAMOS, F. and UPCROFT, B. Simple online and realtime tracking. In: *IEEE. 2016 IEEE international conference on image processing (ICIP)*. 2016, p. 3464–3468.
- [5] BIKE FAST FIT. *Bike Fast Fit Elite v2 - Basic Bike Fitting*. 2023. Available at: <https://www.youtube.com/watch?v=4xhFhvIM7R4>.
- [6] BRADSKI, G. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*. 2000.
- [7] BURT, P. *Bike Fit 2nd Edition: Optimise Your Bike Position for High Performance and Injury Avoidance*. Bloomsbury Publishing, 2022.
- [8] CHEN, K., WANG, J., PANG, J., CAO, Y., XIONG, Y. et al. MMDetection: Open MMLab Detection Toolbox and Benchmark. *ArXiv preprint arXiv:1906.07155*. 2019.
- [9] CONTRIBUTORS, M. *OpenMMLab's Model Deployment Toolbox*. [<https://github.com/open-mmlab/mmdploy>]. 2021.
- [10] CONTRIBUTORS, M. *OpenMMLab Pose Estimation Toolbox and Benchmark* [<https://github.com/open-mmlab/mmpose>]. 2020.
- [11] FANG, H.-S., LI, J., TANG, H., XU, C., ZHU, H. et al. AlphaPose: Whole-Body Regional Multi-Person Pose Estimation and Tracking in Real-Time. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2022.
- [12] GE, Z., LIU, S., WANG, F., LI, Z. and SUN, J. Yolox: Exceeding yolo series in 2021. *ArXiv preprint arXiv:2107.08430*. 2021.

- [13] HOWARD, A., SANDLER, M., CHU, G., CHEN, L.-C., CHEN, B. et al. Searching for mobilenetv3. In: *Proceedings of the IEEE/CVF international conference on computer vision*. 2019, p. 1314–1324.
- [14] HUA, W., DAI, Z., LIU, H. and LE, Q. Transformer quality in linear time. In: PMLR. *International Conference on Machine Learning*. 2022, p. 9099–9117.
- [15] HYODO, K. *Tools to convert ONNX files (NCHW) to TensorFlow format (NHWC)*. Available at: <https://github.com/PINT00309/onnx2tf>.
- [16] JHUANG, H., GALL, J., ZUFFI, S., SCHMID, C. and BLACK, M. J. Towards understanding action recognition. In: *International Conf. on Computer Vision (ICCV)*. December 2013, p. 3192–3199.
- [17] JIANG, T., LU, P., ZHANG, L., MA, N., HAN, R. et al. RTMPose: Real-Time Multi-Person Pose Estimation based on MMPose. *ArXiv preprint arXiv:2303.07399*. 2023.
- [18] JIN, S., XU, L., XU, J., WANG, C., LIU, W. et al. Whole-Body Human Pose Estimation in the Wild. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. 2020.
- [19] LI, J., WANG, C., ZHU, H., MAO, Y., FANG, H.-S. et al. CrowdPose: Efficient Crowded Scenes Pose Estimation and a New Benchmark. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2019.
- [20] LI, Y., YANG, S., LIU, P., ZHANG, S., WANG, Y. et al. Simcc: A simple coordinate classification perspective for human pose estimation. In: Springer. *European Conference on Computer Vision*. 2022, p. 89–106.
- [21] LIN, J., ZENG, A., WANG, H., ZHANG, L. and LI, Y. One-Stage 3D Whole-Body Mesh Recovery with Component Aware Transformer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2023, p. 21159–21168.
- [22] LIN, T.-Y., MAIRE, M., BELONGIE, S., HAYS, J., PERONA, P. et al. Microsoft COCO: Common Objects in Context. In: FLEET, D., PAJDLA, T., SCHIELE, B. and TUYTELAARS, T., ed. *Computer Vision – ECCV 2014*. Cham: Springer International Publishing, 2014, p. 740–755. ISBN 978-3-319-10602-1.
- [23] LIU, W., ANGUELOV, D., ERHAN, D., SZEGEDY, C., REED, S. et al. SSD: Single Shot MultiBox Detector. In: LEIBE, B., MATAS, J., SEBE, N. and WELLING, M., ed. *Computer Vision – ECCV 2016*. Cham: Springer International Publishing, 2016, p. 21–37. ISBN 978-3-319-46448-0.
- [24] LOSHCHILOV, I. and HUTTER, F. Decoupled weight decay regularization. *ArXiv preprint arXiv:1711.05101*. 2017.
- [25] LYU, C., ZHANG, W., HUANG, H., ZHOU, Y., WANG, Y. et al. Rtmddet: An empirical study of designing real-time object detectors. *ArXiv preprint arXiv:2212.07784*. 2022.

- [26] MACQUEEN, J. et al. Some methods for classification and analysis of multivariate observations. In: Oakland, CA, USA. *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*. 1967, vol. 1, no. 14, p. 281–297.
- [27] MAJI, D., NAGORI, S., MATHEW, M. and PODDAR, D. Yolo-pose: Enhancing yolo for multi person pose estimation using object keypoint similarity loss. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022, p. 2637–2646.
- [28] RETÜL TECHNOLOGY. *XY Fit Step 4: Saddle Fore-Aft*. 2021. Available at: <https://www.youtube.com/watch?v=hZgI87DUUbU>.
- [29] RETÜL TECHNOLOGY. *XY Fit Steps 1-8: Markers and Harness*. 2021. Available at: <https://www.youtube.com/watch?v=k-FsejpbkE>.
- [30] RIBEIRO BRANCO, G., DE MICHELIS MENDONÇA, L., ALVES RESENDE, R. and PIVETTA CARPES, F. Does the Retül System provide reliable kinematics information for cycling analysis? *Journal of Science and Cycling*. Dec. 2022, vol. 11, no. 3, p. 76–84. DOI: 10.28985/1322.jsc.15. Available at: <https://www.jsc-journal.com/index.php/JSC/article/view/759>.
- [31] RONCHI, M. R. and PERONA, P. Benchmarking and Error Diagnosis in Multi-instance Pose Estimation. In: October 2017. DOI: 10.1109/ICCV.2017.48.
- [32] SAVITZKY, A. and GOLAY, M. J. E. Smoothing and Differentiation of Data by Simplified Least Squares Procedures. *Analytical Chemistry*. 1964, vol. 36, no. 8, p. 1627–1639. DOI: 10.1021/ac60214a047. Available at: <https://doi.org/10.1021/ac60214a047>.
- [33] SUN, K., XIAO, B., LIU, D. and WANG, J. Deep High-Resolution Representation Learning for Human Pose Estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2019.
- [34] TAHERI, S., VEDIENBAUM, A., NICOLAU, A., HU, N. and HAGHIGHAT, M. R. Opencv.js: Computer vision processing for the open web platform. In: *Proceedings of the 9th ACM Multimedia Systems Conference*. 2018, p. 478–483.
- [35] TELEA, A. C. An Image Inpainting Technique Based on the Fast Marching Method. *Journal of Graphics Tools*. 2004, vol. 9, p. 23 – 34. Available at: <https://api.semanticscholar.org/CorpusID:5908881>.
- [36] ULTRALYTICS. *YOLOv5: A state-of-the-art real-time object detection system* [<https://docs.ultralytics.com>]. 2021. Accessed: 14.01.2024.
- [37] VASWANI, A., SHAZEER, N., PARMAR, N., USZKOREIT, J., JONES, L. et al. Attention is all you need. *Advances in neural information processing systems*. 2017, vol. 30.
- [38] WU, J., ZHENG, H., ZHAO, B., LI, Y., YAN, B. et al. Ai challenger: A large-scale dataset for going deeper in image understanding. *ArXiv preprint arXiv:1711.06475*. 2017.

- [39] XU, L., GUAN, Y., JIN, S., LIU, W., QIAN, C. et al. Vipnas: Efficient video pose estimation via neural architecture search. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, p. 16072–16081.
- [40] ZHANG, S.-H., LI, R., DONG, X., ROSIN, P., CAI, Z. et al. Pose2Seg: Detection Free Human Instance Segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2019.

Appendix A

Poster

Web application for setting up an ergonomic position on a bicycle

2024

Bc. Jiří Vlasák

Supervisor: doc. Ing. Martin Čadík, Ph.D.

1. Cycling video recording

- Video should be at least 5 seconds long
- Use an indoor trainer for best results



Figure 1: Example input video.

2. Video upload and fit goals selection

- Can be used with road, gravel and mountain bikes
- Supports comfort, performance and balanced fit goals

3. Pose estimation and angle calculations

- Uses state of the art RTMPose model
- Fine-tuned on custom bikefitting dataset
- Everything runs on front-end using tensorflow.js with WebGPU or WebGL backends




Figure 2: Video after processing, with measured joint angles drawn on top.



vlasakjiri.github.io/bikefit

Excel@FIT 2024


4. Angle analysis

- App tracks joint angles at different phases of the pedal stroke (top, bottom and front)
- Joint angles at individual pedal strokes are compared to recommended ranges

Bottom of the pedal stroke

Measurement	Value	Recommended
Knee to extension	141.7°	135° - 140°

Details




Min	Max	Median	Range
138.0°	144.0°	141.7°	6.0°

Figure 3: Angle analysis for knee to extension angle on the bottom of the pedal stroke.

5. Recommendations


- Based on measured angles, app recommends changes to saddle height, saddle setback, handlebar height and handlebar reach

Saddle Height



Lower by 1cm.

Handlebar Height



No changes necessary.

Figure 4: Recommendations for changes in saddle height and handlebar height.

