Improving a neural network to classify movement disorders.

Research Internship Computing Science
University of Groningen

Tanja de Vries

Daily supervisor: dr. I. Giotis
Second supervisor: prof. dr. M. Biehl

July 13, 2020

Abstract

The aim of this project was to improve an existing Long Short Term Memory Recurrent Neural Network on movement disorders. The network was implemented by project Next Move in Movement Disorders, a collaboration between ZiuZ Visual Intelligence and the University Medical Centre Groningen. The network classifies patients based on 3D video data of them performing certain tasks. In this report we perform several experiments to get more clarity on the problematic areas of the network. The focus was mostly on myoclonus and tremor and partially on dystonia. We did an elaborate search on the possible feature vectors and values of the parameters. With these experiments we were not able to improve the classifier. After that, we explored the data itself. We found that the movement in the task is too dominant compared to the small involuntary movements that manifest due to the movement disorder. In the end we were able to find satisfactory results for a task that did not involve movement, here the network could distinguish myoclonus from dystonia. However, it could not distinguish myoclonus from tremor. More research will be needed to implement a network that can classify the disorders correct independent from the task and disorders we choose. The report is finished with an extensive discussion on future work.

Contents

1 Introduction

This project took place within project Next Move in Movement Disorders (NEMO). In this chapter we first introduce project NEMO, then the relevance of this report is described, and lastly, the structure of this report is explained.

1.1 Next Move in Movement Disorders

Project NEMO is a collaboration between ZiuZ Visual Intelligence and the University Medical Centre Groningen (UMCG). The research is focused on making a classifier to distinguish between multiple movement disorders. The goal is to make an application that can be used by physicians, to have a correct classification in the first hospital without need for a second opinion. This way the patient can receive earlier and correct treatment and will have less hinder from the disorder.

Project NEMO focuses on multiple data modalities, including accelerometers and 3D video. The researchers hope to find only a selection is necessary to build a reliable classifier. Because the data modalities are very different and complex, each data modality has its own pipeline, and in the end they will be combined with ensemble learning.

As part of the NEMO project, research has been done on what information the physicians use when diagnosing a patient. This is done with help of questionnaires. Several physicians have told us what body parts they find most important for each task/disorder, and they have labeled the data. The results of these questionnaires are used to make an initial guess of the parameter choices needed in the classifier.

1.2 Contribution of this report

At the start of the internship the classifier for the video data did not give good results. Our contribution to project NEMO consists of doing experiments to find out why the classifier for video data was not working. In this project we only used the data from the 3D camera. Moreover, we mostly focused on the movement disorders tremor and myoclonus. The main difference between tremor and myoclonus is in the pattern. Tremor is a rhythmic movement and can thus be identified by a frequency. Myoclonus patients have nonrhythmic jerks. The experiments include focusing only on the body parts the physicians find most useful and varying the parameters of the classifier. As a result, we hoped to find an adaptation of the network that is working well.

1.3 Overview of report

Firstly, the theory and the setup that we started with are explained in Chapter 2. In Chapter 3, the experiments are introduced, a reasoning for the choice of experiments is given and the results are shown. Then, a discussion is and suggestions for future work are given in Chapter 4. Lastly, Chapter 5 provides some concluding remarks.

2 Theory

The first part of my project consisted of understanding the code that was already there. To understand the code, I first had to do research on the theory that was used as a base for the algorithms. This theory is documented in this chapter. Firstly, the data is explained, then the way features were built is described, then the classifier is explained, and lastly, an introduction to the results is given.

2.1 Data

As mentioned in the introduction, project NEMO focuses on multiple data streams. In this project we only consider the video data. Each patient is filmed while performing certain tasks that are commonly used by the physicians to diagnose the patients. In total they perform 24 tasks of which 12 are performed twice, both with the right and left arm. This gives a total of 36 recorded tasks per patient. The tasks can be divided into six categories. The categories are, posture, rhythmic movement, fine motor skills, coordination, standing and others. Respectively, these categories include tasks such as holding a finger in front of the nose, finger tapping, spiral drawing, moving a finger to the nose and back, walking, and drinking.

The NEMO project did two surveys among the physicians. In the first survey the physicians rated each task with regards to its overall diagnostic potential (extremely, very, moderately, somewhat, not at all). The second survey is focused on what body parts give most information for each task and disorder. The results of these surveys are used to make an initialization of the parameters of the classifier. However, it is possible that the classifier will use different information than the physicians. So, we have to keep in mind that other tasks and body parts could be more important for the classifier.

Due to the tasks being of different nature we first focused on classifying one task. In the chosen task, task 22, the patient moves his right hand multiple times from his right side to his nose and back. This task is chosen because good results were obtained with this task in previous NEMO experiments. Moreover, on average the physicians consider this task very useful and some physicians even classified this task as extremely useful.

We also minimised the disorders we were classifying. At first we focus on myoclonus and tremor patients. We had data of 8 myoclonus and 6 tremor patients. In the second survey we can find which body parts are important for task 22. For this task it is the same for both myoclonus and tremor. The physicians say the ipsilateral hand and wrist are most important, where ipsilateral means the hand performing the task. Next important are the ipsilateral elbow and shoulder.

The camera that is used is the Intel RealSense D415[2] 3D camera. It is placed at a distance of approximately 2.5 meters facing the patient. Then, in real-time, the raw frames are processed by the Nuitrack skeleton tracking software. The software produces the coordinates (x, y, z) of 20 joints of the human body. The Nuitrack skeleton provides us with a way to focus on the important pixels of the frame. Resulting in a feature vector of size $60 = (3 \times 20)$ reduced from a 640×680 RGB frame. Moreover, the legs are not important except for the three standing tasks. Therefore, the legs

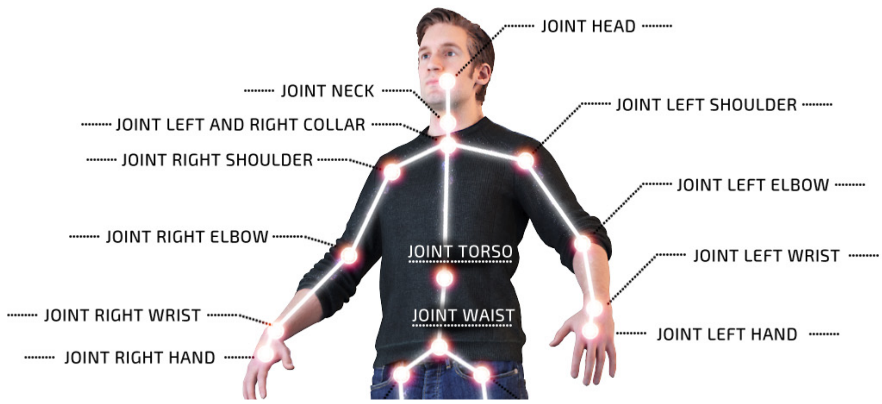

Figure 1: Joints of the human body without the legs, in Nuitrack skeleton tracking.[1]

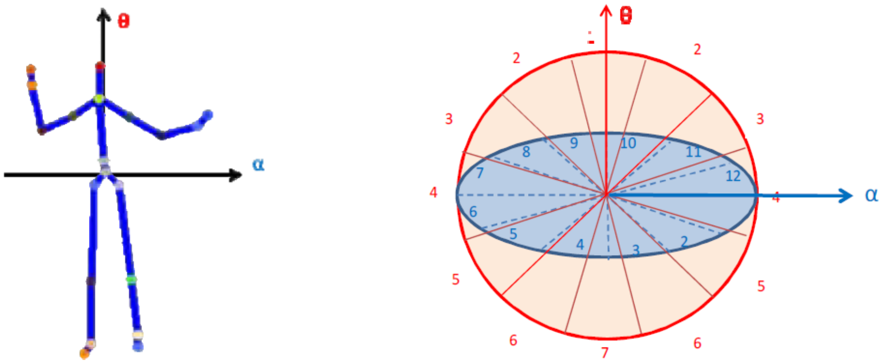

(a) Skeleton joints with α and θ axes[3]. (b) Histogram[3].

Figure 2: The α and θ axes with the corresponding histogram bins.

are removed from the skeleton. Hence, we use all body parts without the legs, thus starting at the waist, as is shown in Figure 1.

2.2 Features

We want the data to be viewpoint invariant, such that the precise position of the camera does not affect the quality of the data. To achieve this the method from the article "View Invariant Human Action Recognition Using Histograms of 3D Joints" by Lu Xia et al.[3] is partially used.

In this article the researchers try to find body postures from 3D coordinates of skeleton joints. To achieve this, the joints are converted to spherical coordinates. To make the system viewpoint invariant, the spherical coordinates' center is at the waist. The horizontal reference vector α goes through both the waist and the reference vector θ is perpendicular, which is shown in Figure 2a. The planes defined by the axes are divided into 12 pieces, of which an illustration is shown in Figure 2b. Note that in the article by LuXia et al. [3] the circle of θ is divided into 7 bins of which bins 2 to 6 repeat on both sides, that is, it only measures the height. We divide this circle in 12 bins without combining the bins. Hence, they use 84 bins in the article where we use 144.

To make the coordinates invariant to the viewpoint, the coordinates of the origin are subtracted from the coordinates of each joint. The angle θ of each joint with respect to the origin is then computed with respect to the θ-axis connecting the origin and reference projected on the vertical plane. The angle α is computed with respect to the horizontal axis with a constant depth (z-coordinate) intersecting the origin. In order to make the method scale-invariant the radius is not used.

For each joint we compute in which bin it is located, and then we use smoothing to reduce noise from the skeleton tracking. To make the method robust against noise from the skeleton tracking, the mean and standard deviation of each joint over the frames is computed. The Gaussian weight function based on the mean and standard deviation is then taken as a base value for the joint. Note that the Gaussian distribution is different for each person and task. The distribution of the right hand with respect to the waist and torso of one patient in task 22 is given in Figure 3. In Figure 3a the position of the right hand in one frame is shown. In Figure 3b the Gaussian distribution of the right hand is plotted for the entire video. In Figure 3c the combination of the two is given. This last histogram is computed for all joints and then for each frame the sum of the histograms of the joints is taken. The resulting histogram is the feature vector that goes into the classifier.

2.3 Classifier

The video is recorded at 30 frames per second and divided into blocks of 15 frames. We used a window offset of 5 frames, thus adjacent windows have an overlap of 10 frames. Note that only full blocks are taken into account, so a video of 34 frames does not have more blocks than the video of 30 frames.

Each patient is classified by the medical experts. Every window of a patient is given the same label. Note that the symptoms of the disorders are not always visible, so it may occur that the

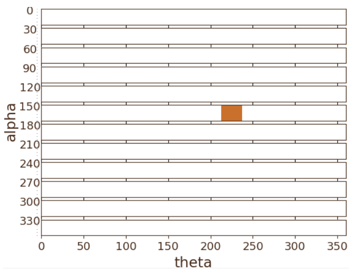

(a) The position of the right hand in one frame.

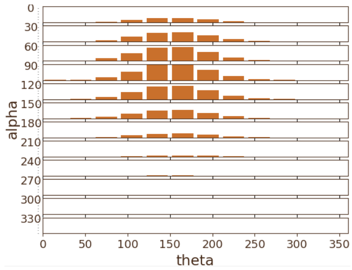

(b) Gaussian distribution based on the mean and variance of the right hand over all frames in a video.

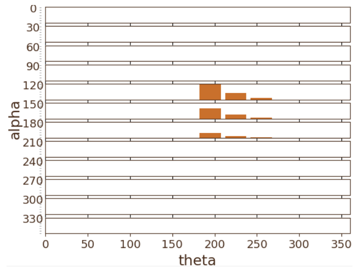

(c) Smoothed position of the right hand in the same frame as a).

Figure 3: Histogram of the right hand.

Truth \ Classified	Myoclonus	Tremor
Myoclonus	10	0
Tremor	9	1

(a) Classification table

Precision		Recall	
M	T	M	T
0,53	1,00	1,00	0,10

(b) Precision and recall.

Figure 4: Example to explain the precision and recall method.

disorder symptoms are not present in a certain window. It is possible that this gives some clutter in the data, but we do not know whether this has an influence and how big this influence is.

The input data is thus a time series of 15 frames where each frame is represented by a histogram. For time-series, a recurrent neural network (RNN) is the natural choice. An RNN has the ability to use context of the other frames for classification. A Long Short Term Memory network (LSTM) is a variant of the RNN that has the ability to learn long term dependencies as well as short term dependencies[4]. Therefore, we have chosen to use an LSTM network.

For each disorder we choose one patient that will be used to test the network. The network is trained on the other patients. To make sure the model is not biased, the number of block used for both disorders is equal. For each window of the test patients the classifier makes a decision on the disorder. For easy interpretation of the results, the test sets are also of equal size. A future plan is to combine these decisions to give one final classification per patient instead of per frame.

2.4 Results

The results are given as the precision and recall. Firstly, we explain these terms, than we give a clarifying example and lastly the real results are introduced.

The precision is a number between zero and one, it represents the percentage of how many of the windows classified as myoclonus are indeed myoclonus. Recall gives the percentage of myoclonus windows that were classified correctly. We assume that the disorder manifests itself in at least 50 percent of the windows, therefore we want to have a recall of at least 0.5. Furthermore, 2 physicians classifying the patients from the videos have an agreement around 70 percent. We want to have a classifier that performs at least as good as the physicians, so we aim for a precision of 0.75.

Suppose we have 20 windows, 10 of both disorders. Let us assume a classifier classified all but one windows as myoclonus and the last window is correctly classified as tremor, as shown in Table 4a. The values of the precision and recall are shown in Figure 4b. The precision for myoclonus (M) is 10/19. The recall for myoclonus is 10/10, we found all windows. The precision for tremor (T) is 1/1, all windows that we found were tremor. The recall for tremor is 1/10. From this we learn that a high number is not always good, we have to look at all 4 results. The average precision is here 0.765 and the average recall is 0.55, but the classifier clearly did not learn the difference between tremor and myoclonus. This shows that it is necessary for the precision of both tremor and myoclonus to be above 0.75, and the recall for both above 0.5. Therefore, in all figures in this report we mark the precision and recall passing these thresholds in green and yellow, respectively.

Setting	Value
Disorders	Myoclonus and tremor
Origin	Waist
Reference	Torso
Joints	All (without legs)
Number of orientations	12
Task	22
Window size	15
Window offset	5

Figure 5: Settings of the base case

		Mean				Standard deviation				Median			
Test patient		Precision		Recall		Precision		Recall		Precision		Recall	
M	T	M	T	M	T	M	T	M	T	M	T	M	T
24	14	0,20	0,52	0,05	1,00	0,40	0,03	0,11	0,00	0,00	0,50	0,00	1,00
49	42	0,84	0,79	0,99	0,72	0,21	0,40	0,02	0,39	1,00	1,00	1,00	1,00
26	30	0,50	0,00	1,00	0,00	0,00	0,00	0,00	0,00	0,50	0,00	1,00	0,00
31	35	0,46	0,00	0,89	0,00	0,08	0,00	0,23	0,00	0,50	0,00	1,00	0,00
36	45	0,80	0,83	0,70	1,00	0,40	0,21	0,40	0,00	1,00	1,00	1,00	1,00
32	17	0,21	0,00	0,36	0,00	0,19	0,00	0,35	0,00	0,27	0,00	0,38	0,00

Figure 6: Results base case. In green the precision higher than 0.75 and in yellow the recall higher than 0.5.

2.4.1 Base case results

In this chapter the parameters have been introduced. The values used as base are summarized in Figure 5. The results of the classifier trained using these settings are shown in Figure 6. The mean, standard variance and median are taken over 5 runs. The mean is the sum of the values divided by the number of values, the standard deviation measures how much the values deviate from the mean, and the median is the middle value separating the high and low values. If the standard deviation is low, the mean and median are close to each other. However, if the standard deviation is high, the median can give additional information. For example for test patients pair (36, 45) the mean recall for myoclonus is 0.7 and the corresponding standard deviation is 0.4. The median is 1, that tells us 3 out of 5 experiments have a recall of 1 and there are one or two experiments that resulted in a low recall. We see the pattern of the coloured indices are the same for the mean and the median. So the mean is robust enough over five experiments. From now on we will focus only on the mean.

The experiment is performed among different test pairs. As mentioned before, the classifier is trained on all patients except for the two test patients. We see the result depends a lot on the patients in the test set. It has good results for patients (49,42) and (36,45). However, if we look closer the network classifies (24, 14) both as tremor patients, and (26,30) and (31,35) both as myoclonus. Moreover, in case of (32,17) it classifies patient 17 as myoclonus instead of tremor. Looking at the mean and median, the classifier seems to divide 32 over both myoclonus and tremor.

Overall, it is possible that the classifier finds the tremor and myoclonus in some patients, but the results are not better then the results would be for guessing, so we do not know for sure. One thing that stands out is that, except for patient 32, it seems to give a consistent label for all windows of a patient throughout the runs.

3 Experiments

In this chapter, the experiments are explained and their results are shown. The results of each experiment have influenced the choice of the next experiments, therefore we have chosen to show the experiments and results combined in one chapter. Due to the amount of results, only the important results are shown in this chapter. The other results are shown in Appendix A.

Unless otherwise specified, the base settings from Figure 5 are used. As explained before, we use a green colour to mark the precision greater than or equal to 0.75, and a yellow background for the recall greater than or equal to 0.5. We have chosen to display the results for 3 pairs of patients that show the disorders clearly according to the experts. If we are not able to classify these pairs, we decide the classifier is not performing good enough. We have seen that the mean covers most information, so from now on only the mean is shown in the results, it is taken over 5 runs.

3.1 Moving the origin and reference.

As mentioned in Chapter 2, the origin is at the torso and the reference at the waist. However, as found in the surveys of the physicians, the focus is mostly on the arms and hands performing the task. Task 22 is performed with the right arm, therefore, the first experiment is to move the origin and reference to the right arm. We have chosen multiple combinations of origin and reference for this experiment. In Figure 7 the results are shown. The results of patients (24, 14) suggest that (right shoulder, right elbow) and (right elbow, right wrist) are performing best. However, if we look at the other patient, we see very different results. Hence, choosing the best performance of origin and reference depends on the chosen test patients. There is too much variance in these results to conclude which origin is the best. Therefore, we will vary the origin and reference during the other experiments. Because we expect the upper arm of the right arm to be most suitable for the origin and reference and the 4 top rows in Figure 7 all predict one patient correctly, we decided to keep these four settings for now.

3.2 Smoothing

While reading the article of the histogram of joints [3], we realised they used a lot of smoothing in the histogram. The histogram itself is a way of smoothing, all locations in a bin of 30 degrees are set to be equal. Moreover, as explained in Section 2.2, the Gaussian distribution is used to smooth over the bins. The goal of the report is to find the posture of the person. But, depending on the disorder, we want to ignore most of that posture and focus on the small movement. Therefore, we were wondering if we could get better results by changing the feature vector. On the other hand, there is noise in the data and the tracking of the joints that is presumably removed with the smoothing. Therefore, it is nontrivial to predict what method is the best. In this section we explore two alternative approaches. Firstly, we remove the smoothing within the histogram. Secondly, we remove the computation of the histogram altogether and use the second derivative of the coordinates with respect to time instead.

		Test patient 24 and 14				Test patient 49 and 42				Test patient 32 and 17			
		Precision		Recall		Precision		Recall		Precision		Recall	
Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Waist	Torso	0,20	0,52	0,05	1,00	0,84	0,79	0,99	0,72	0,21	0,00	0,36	0,00
Right elbow	Right shoulder	0,51	0,20	1,00	0,05	0,00	0,50	0,00	0,99	1,00	0,94	0,93	1,00
Right shoulder	Right elbow	0,79	0,92	0,91	0,68	0,53	0,42	0,15	0,69	0,52	0,80	1,00	0,09
Right elbow	Right wrist	0,96	0,93	0,92	0,96	0,00	0,50	0,00	1,00	0,50	0,49	0,72	0,29
Right wrist	Right elbow	0,77	0,60	0,48	0,79	0,24	0,25	0,26	0,26	0,18	0,00	0,24	0,00
Right shoulder	Left shoulder	0,31	0,14	0,52	0,14	0,10	0,50	0,00	0,98	0,51	0,40	1,00	0,03

Figure 7: Results for multiple origins and references.

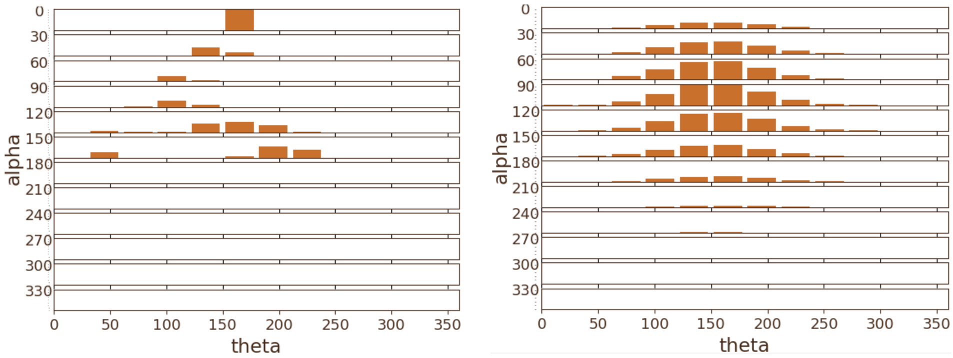

(a) Real distribution of the right hand. (b) Gaussian distribution of the right hand.

Figure 8: Comparison of the real distribution and the corresponding Gaussian distribution. The right hand with respect to the origin waist, and reference torso.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,20	0,52	0,05	1,00	0,84	0,79	0,99	0,72	0,21	0,00	0,36	0,00
	Right elbow	Right shoulder	0,51	0,20	1,00	0,05	0,00	0,50	0,00	0,99	1,00	0,94	0,93	1,00
	Right elbow	Right wrist	0,96	0,93	0,92	0,96	0,00	0,50	0,00	1,00	0,50	0,49	0,72	0,29
	Right shoulder	Right elbow	0,79	0,92	0,91	0,68	0,53	0,42	0,15	0,69	0,52	0,80	1,00	0,09
No smoothing	Waist	Torso	0,37	0,29	0,40	0,28	0,73	0,82	0,83	0,68	0,31	0,06	0,45	0,04
	Right elbow	Right shoulder	0,86	0,68	0,56	0,91	0,31	0,40	0,19	0,56	0,44	0,42	0,53	0,33
	Right elbow	Right wrist	0,68	0,66	0,62	0,70	0,46	0,48	0,24	0,71	0,30	0,30	0,30	0,30
	Right shoulder	Right elbow	0,56	0,71	0,84	0,35	0,63	0,67	0,68	0,61	0,45	0,39	0,61	0,25

Figure 9: Results for no smoothing in histogram.

3.2.1 No smoothing in histogram

The smoothing within the histogram is based on the Gaussian distribution. However, as shown in Figure 8, the Gaussian distribution is in our case not a close representation of the actual distribution. Furthermore, the smoothing might remove not only the noise but the symptoms of the disorder as well.

In the method that we propose here we remove all smoothing in the histogram. This means that we take the histogram from Figure 3a for each joint. As before, we take the sum over all joints and take this as the new feature vector. This is given as input to the same classifier as before. The results are shown in Figure 9. The values are less extreme. It seems good that we do not have so many low numbers. Although, it might be the case we are not guessing per patient but per window. It is interesting to see that both methods seem to agree somewhat on which settings get high values. Specifically patients (24, 14) for the (right elbow, right wrist) and patients (49, 42) for the (waist, torso). However, they do not agree on all values. From this data, it is difficult to say which setting is performing better. The original smoothing results in higher values. However, is also results in lower values. Thus, we do not think we have enough reason to choose one above the other.

We tried to find out which method could distinguish the data of the two disorders better on a low level. The idea was to take two synchronised videos and compute the distance between the frames. However, it was not possible to synchronise the frames exactly because different patients performed the task at a different speed. Moreover, it was difficult to find a metric that could compare the two methods in a meaningful way. We were not able to find a way to make a distinction between the within and between class distances.

24 and 14				49 and 42				32 and 17			
Precision		Recall		Precision		Recall		Precision		Recall	
M	T	M	T	M	T	M	T	M	T	M	T
0,74	0,73	0,72	0,74	0,65	0,55	0,32	0,83	0,49	0,41	0,76	0,21

Figure 10: Results for second derivative.

3.2.2 Second derivative

The histogram itself is a way of smoothing the data. Therefore, we tried a method without using the histogram.

As mentioned in the introduction, Chapter 1, project NEMO makes use of multiple data modalities. So far, the accelerometer seems to produce better results. The accelerometer measures acceleration, which is equal to the second derivative over time. We have the coordinates (x, y, z) per joint in each frame. This gives us enough data to estimate the second derivative over the frames (time). This is a method that does not use the histogram. Additionally, this gives us data that is independent of viewpoint. Therefore, we do not need an origin and reference in this experiment. The results are shown in Figure 10. If we look at the range of the values, the results are comparable to the results without smoothing. The results are again not satisfactory.

3.2.3 Conclusion

There is no method that stands out enough to conclude it is the best method to use. Therefore, we keep comparing the different methods in the rest of this report.

3.3 Varying parameters

Since we did not see the results that we hoped for we did an experiment on the values of the parameters. It is possible that there was a parameter that was not chosen optimally. We use the base settings from Figure 5 where the joints are now changed to only include the arms. In this section we vary the values of the number of orientations, the window size and the window offset. Due to the experiments being similar the results of the original method and the histogram without smoothing are combined in one figure. The results of the second derivative are given separately in Appendix A, Figure 19.

The approach we have is as follows. For each parameter we determine which value works best and we keep this value when we move towards the next parameter. This might not be the best method, however testing all combinations would take too long.

3.3.1 Body parts

Because the medical experts say they focus on the right arm for task 22, the other body parts might give more noise than information to the classifier. So, in this experiment we vary the joints that we give as input to the classifier. Note that the histogram is computed with respect to the origin and reference, but these are not included themselves. So, it is still possible to have the waist and torso as origin and reference even though we focus on the arms.

In Appendix A, Figure 20 the results are shown. If we say an arm is included, it means hand, wrist, elbow and shoulder are included. We see that the joints of the left arm are included have a big influence on the results. Since the left arm is not moving at all in this task, we did not expect these results. Therefore, it is unsure if the classifier makes conclusions based on the disorders or on something else.

The best result is seen for the right arm combined with the left hand. However, we can not explain this based on the theory. To keep things simple, we would like to look at whole arms. Thus, we either add both arms or only the right arm. If we compare these two, the result is better when including both arms. So, from now on, we include only the arms instead of all body parts.

3.3.2 Number of orientations

As explained in Chapter 2, the number of orientations determines the size of the histogram. A higher number of orientations gives smaller pieces, so more detail in the movement is stored. However, there might be more noise because there is more information, both good and bad, stored.

The results are shown in Appendix A, Figure 21. It immediately stands out that for 6 orientations patient pair (24, 14) has high precision and recall for all but one origin reference pair. However, the values for the other patients have not increased. For 18 orientations we see slightly higher results than for 12 orientations and for 24 orientations the values are lower. Dividing a sphere into 6 bins results in bins of 60 degrees. That means a movement of 50 degrees does not necessarily result in a shift of a bin. Therefore, we expect this setting not to be able to find a small involuntary movement. Combining all information, we choose 18 orientations to be the new setting. It has higher results than 12 orientations and matches better with the expectations based on theory. Note that for 18 orientation the setting (original smoothing, right elbow, right shoulder) predicts two out of three pairs of patients correctly.

3.3.3 Window size

The next parameter that we explore is the window size. The frame rate is 30 frames per second. Right now the window size is 15 frames, which is half a second. A larger frame gives contains more information, as a result it also contains more noise. Therefore the right window size is finding a balance between enough information and not too much noise. The results are shown in Appendix A, Figure 22. We do not see a lot of difference between the different window sizes, therefore, we keep the window size at 15.

3.3.4 Window offset

The last parameter that is explored is the window offset. The offset determines how many overlapping frames adjacent windows have. Right now the offset is set to five, which means we start a new window every five frames. This means adjacent windows overlap 10 frames. This is a way of data augmentation. We have more data to train the classifier on. However, if we use too many windows that are similar we risk to overfit on the training data. The results are shown in Appendix A, Figure 23. One thing that stands out are the results for window offset 20, they are as good if not better than offset 15. However, in both cases there is no overlap between the windows, there are just less windows with an offset of 20. The best performance seems to be with an offset of 10. Moreover, using an offset of 10 is simpler than an offset of 5, due to having less input. Therefore, we decide upon a window offset of 10.

3.3.5 Conclusion

The final combination of parameters is shown in Appendix A, Figure 11. Overall, the results seem to be slightly better, however it is still not satisfactory. Therefore, we conclude that the parameter values is not the main issue of the classifier. Since the results are not satisfactory either way, it is unsure if the better performance of the new settings will generalize to a classifier that is working.

3.4 Is it even possible with this data?

At this point we were wondering if the problem could be in the data itself. Maybe it consisted of too much noise, or the frame rate was too low to find the disorders.

3.4.1 Myoclonus - dystonia

The involuntary movement of a tremor patient is usually smaller than that of the other patients. Therefore, it is possible that tremor is more difficult to find in the data and the other disorders perform better. Dystonia is a disorder that shows up in posture and larger movements. Therefore, in this experiment we try to distinguish between myoclonus and dystonia. If this does produces satisfactory results, it might be an issue with the small tremors.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,86	0,56	0,60	0,79	0,79	0,98	0,98	0,65	0,41	0,00	0,76	0,00
	Right elbow	Right shoulder	0,76	0,98	0,99	0,65	0,00	0,50	0,00	0,99	1,00	0,98	0,98	0,99
	Right elbow	Right wrist	0,99	0,99	0,99	0,99	0,00	0,49	0,00	0,95	0,58	0,73	0,80	0,43
	Right shoulder	Right elbow	0,45	0,39	0,57	0,29	0,70	0,64	0,59	0,74	0,00	0,00	0,00	0,00
No smoothing	Waist	Torso	0,67	0,54	0,39	0,76	0,84	0,81	0,81	0,84	0,18	0,01	0,24	0,01
	Right elbow	Right shoulder	0,89	0,73	0,67	0,92	0,46	0,45	0,22	0,66	0,59	0,56	0,54	0,61
	Right elbow	Right wrist	0,87	0,72	0,64	0,90	0,42	0,43	0,31	0,55	0,42	0,45	0,41	0,47
	Right shoulder	Right elbow	0,36	0,38	0,35	0,40	0,49	0,49	0,50	0,49	0,09	0,00	0,10	0,00
2nd derivative	-	-	0,80	0,64	0,52	0,87	0,60	0,58	0,54	0,63	0,47	0,39	0,74	0,17

Figure 11: Results vary parameters, with parameters shown in Figure 12.

Setting	Value
Disorders	Myoclonus and tremor
Joints	Arms
Number of orientations	18
Task	22
Window size	15
Window offset	10

Figure 12: Final values parameters

The experiment is run for the base settings and the new parameters found in Section 3.3. Moreover, we have two options for the joints, the arms and the right arm. The results are shown in Appendix A, Figure 24.

The results where most patients are classified correctly are shown in Figure 13. We conclude that the results are not better than they were for tremor. Note that one of the medical experts told us this is not a good task to use for recognising dystonia. Therefore it would be interesting to see if the results are different if we use a better task. However, due to time constraints, we were not able to include that in this report.

3.4.2 Tasks that are holding a position

We are searching for the small movements, so maybe the larger movement of the task is overwhelming the feature vector and thus the classifier. Therefore, in this experiment we look at a task that has no movement itself. Task 24 consists of holding the right finger in front of the nose. The results are shown in Appendix A, Figure 25. We see the task works best in combination with dystonia.

The best result is obtained with myoclonus and dystonia, looking at the right arm, with 12 orientations and offset 5. This result is shown in Figure 14. The setting with original smoothing and the right shoulder as origin and right elbow as reference is the first setting that performs well for all three patient pairs. To see whether the classifier is completely working we look at more patients in Figure 15. Note that we only have four dystonia patients, thus they are repeated. Five out of eight pairs are found. So, this is clearly better than the original results in Figure 6. Moreover, patient 41 has a dystonia that is in this task mostly present in the fingers, since the fingers are not tracked it is very understandable that this patient was not classified correctly in this task.

3.4.3 Spectral characteristics

Since tremor is a rhythmic movement, it is possible to find the frequencies in the data corresponding to the tremor. In this section we compare the spectral characteristics of the data of the accelerometer and the video data. The accelerometer records at 150 fps, hence for the comparison we have scaled

			24 and 46				49 and 21				36 and 15			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	D	M	D	M	D	M	D	M	D	M	D
Original	Waist	Torso	0,44	0,27	0,76	0,06	0,39	0,11	0,62	0,07	0,50	0,00	1,00	0,00
	Right elbow	Right shoulder	0,68	0,63	0,54	0,73	0,55	0,59	0,78	0,36	0,66	0,87	0,92	0,52
	Right elbow	Right wrist	0,98	0,78	0,70	0,98	0,17	0,35	0,13	0,49	0,69	0,68	0,69	0,66
	Right shoulder	Right elbow	0,43	0,34	0,58	0,22	0,49	0,25	0,91	0,06	0,66	0,82	0,88	0,55
No smoothing	Waist	Torso	0,46	0,42	0,66	0,25	0,43	0,31	0,65	0,15	0,52	0,69	0,94	0,13
	Right elbow	Right shoulder	0,70	0,57	0,35	0,85	0,48	0,46	0,63	0,33	0,68	0,82	0,87	0,58
	Right elbow	Right wrist	0,66	0,55	0,33	0,84	0,43	0,43	0,40	0,45	0,68	0,86	0,91	0,57
	Right shoulder	Right elbow	0,30	0,41	0,20	0,55	0,45	0,40	0,57	0,29	0,76	0,95	0,96	0,70
2nd derivative	-	-	0,64	0,62	0,60	0,64	0,48	0,48	0,57	0,39	0,70	0,64	0,57	0,75

Figure 13: Results for myoclonus and dystonia, with right arm, 18 orientations, offset 10.

			24 and 46				49 and 21				36 and 15			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	D	M	D	M	D	M	D	M	D	M	D
Original	Waist	Torso	0,70	1,00	1,00	0,40	0,50	0,00	1,00	0,00	0,71	0,80	1,00	0,53
	Right elbow	Right shoulder	0,00	0,50	0,00	1,00	0,89	0,80	1,00	0,79	0,80	0,90	0,80	1,00
	Right elbow	Right wrist	0,00	0,31	0,00	0,61	0,00	0,50	0,00	1,00	1,00	1,00	1,00	1,00
	Right shoulder	Right elbow	0,80	1,00	1,00	0,61	1,00	1,00	1,00	1,00	0,80	0,90	0,80	1,00
No smoothing	Waist	Torso	0,50	1,00	1,00	0,02	0,88	1,00	1,00	0,80	0,59	1,00	1,00	0,29
	Right elbow	Right shoulder	0,00	0,23	0,00	0,29	0,33	0,75	0,60	0,32	0,52	0,61	0,69	0,37
	Right elbow	Right wrist	0,50	0,89	0,93	0,08	0,02	0,10	0,02	0,15	0,63	0,84	0,84	0,51
	Right shoulder	Right elbow	0,00	0,02	0,00	0,03	1,00	1,00	1,00	1,00	0,59	1,00	1,00	0,30
2nd derivative	-	-	0,27	0,29	0,27	0,35	0,47	0,42	0,75	0,19	0,23	0,21	0,34	0,23

Figure 14: Results for task 24, myoclonus and dystonia, rightarm, offset 5, 12 orientations.

		Mean				Standard deviation				Median			
Test patient		Precision		Recall		Precision		Recall		Precision		Recall	
M	D	M	D	M	D	M	D	M	D	M	D	M	D
24	46	0,80	1,00	1,00	0,61	0,24	0,00	0,00	0,48	1,00	1,00	1,00	1,00
49	21	0,76	0,90	0,80	0,94	0,39	0,20	0,40	0,11	1,00	1,00	1,00	1,00
36	15	1,00	1,00	1,00	1,00	0,00	0,00	0,00	0,00	1,00	1,00	1,00	1,00
26	41	0,00	0,50	0,00	1,00	0,00	0,00	0,00	0,00	0,00	0,50	0,00	1,00
31	46	0,70	1,00	1,00	0,41	0,24	0,00	0,00	0,48	0,50	1,00	1,00	0,02
32	21	1,00	1,00	1,00	1,00	0,00	0,00	0,00	0,00	1,00	1,00	1,00	1,00
40	15	1,00	1,00	1,00	1,00	0,00	0,00	0,00	0,00	1,00	1,00	1,00	1,00
38	41	0,00	0,31	0,00	0,61	0,00	0,24	0,00	0,48	0,00	0,50	0,00	1,00

Figure 15: Results for task 24, with origin right shoulder, reference right elbow.

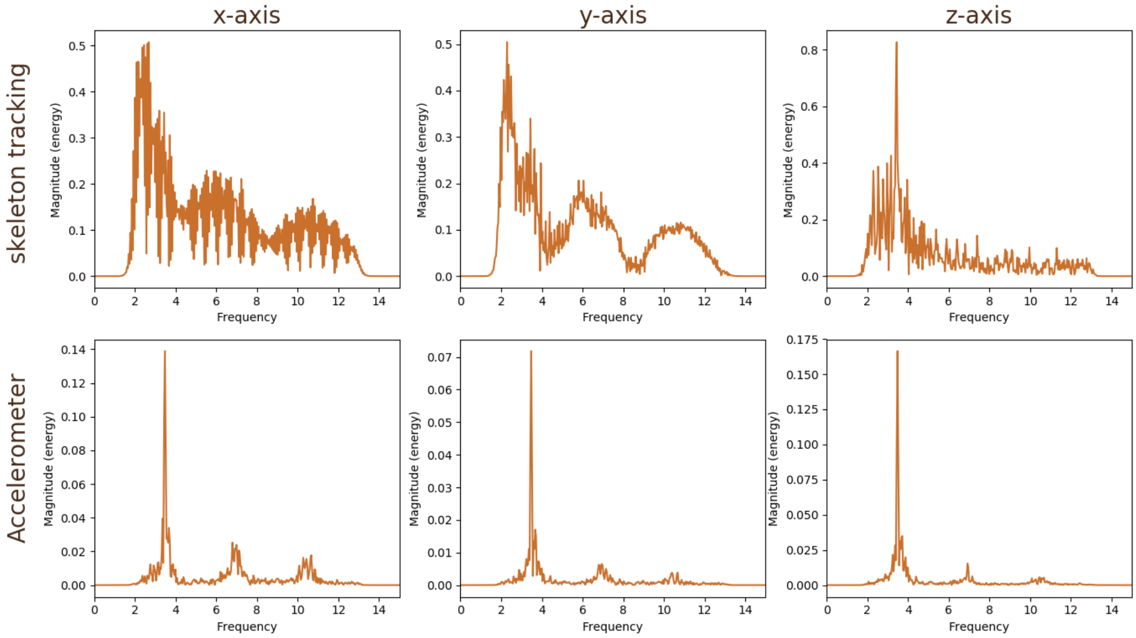

Figure 16: Spectral characteristics of patient 17 performing task 21.

the data of the accelerometer down to 30 fps to match the frame rate of the camera. This is done by only taking every first of five frames. A butterworth bandpass filter in range $[2, 13]$ is used to remove the frequency of a task, which is usually around 0.5. Moreover, we know that the frequencies of tremors are usually within this domain.

In Figure 16 the spectral characteristics corresponding to patient 17 performing task 21 are shown. In the top row the frequencies in the video data are shown. These are computed by taking the displacement of the relative position of the hand with respect to the elbow. The bottom row displays the data of the accelerometer. In the bottom row we see clear peaks around 3.5, 7 and 10.5, although more cluttered, these peaks are also visible in the video data. However, the accelerometer is often more noisy in other tasks, in those cases the video data does not show the frequencies as well. This is fore example the case in task 22, as is shown in the top and bottom row of Figure 17. Here, we do see the first peak in the video data, however the peak around 8 that can be seen in the accelerometer on the y-axis is not visible on the video data. Hence, we concluded that the frequencies are more difficult to find in the video data, however they still seem to be present.

We were wondering if the noise could be a result of the skeleton tracking. Therefore, we have tracked the hand and elbow manually in task 22. In Figure 17 the middle row contains the results of the manual tracking. We do not see a lot of difference in the x-axis and y-axis. However, the z-axis clearly contains more noise, this is presumably due to the way the z-coordinates are computed. Since the figures are very similar, we conclude the tracking is not necessarily the problem.

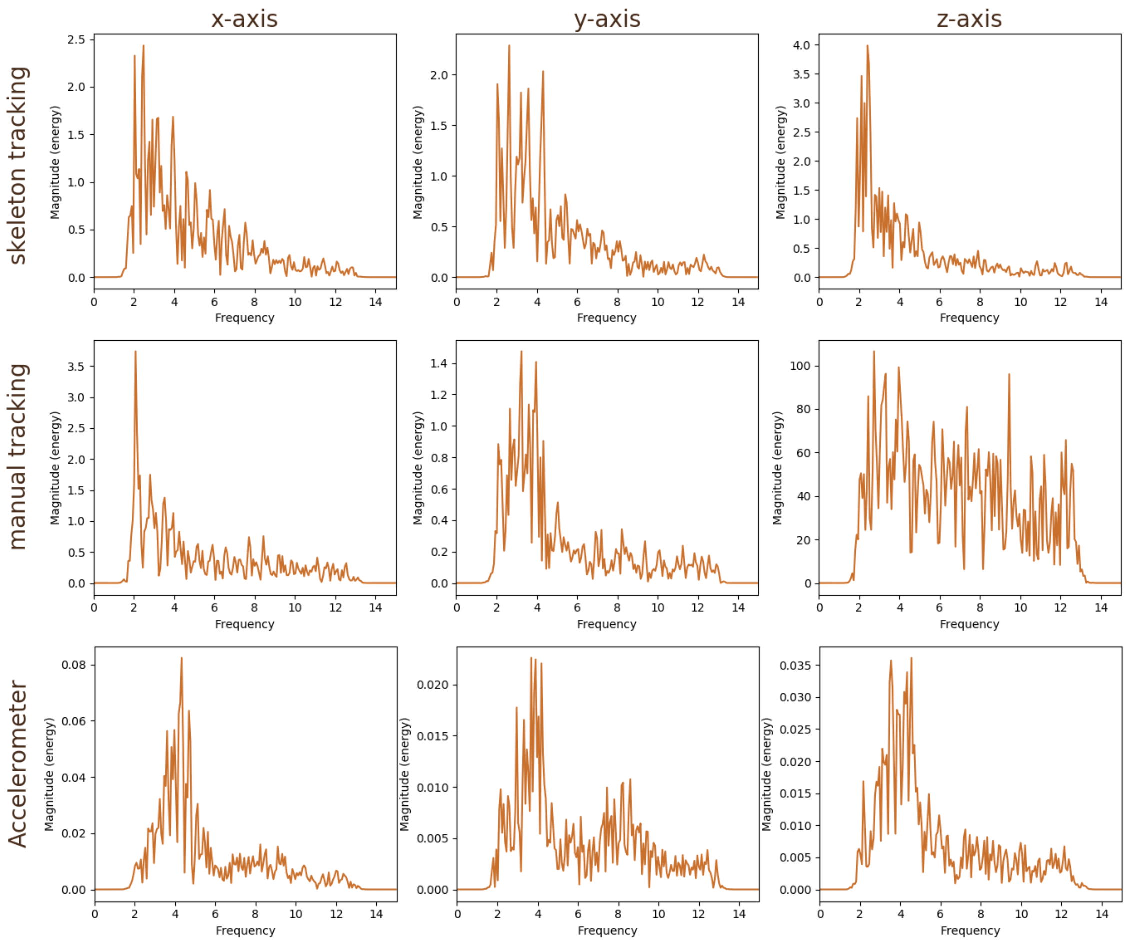

Figure 17: Spectral characteristics of patient 17 performing task 22.

4 Discussion

In this chapter a discussion on the most important problems that we encountered is given. The first problem is the quality of the training data. Secondly the dominant movement in the task itself. The last problem is the difficulty of capturing the small movements.

Problem 1: Quality of training data
The first problem with the training data is the labeling. Right now, we have one label per patient. However, the disorder manifests itself only in a selection of the tasks. Unfortunately, the selection of tasks seems to be different for each patient. It might be possible to make a selection of the tasks for each sub type of the disorders, but we do not have enough patients to test this hypothesis. So, we do not know beforehand what task can be best chosen for each patient. Moreover, within a task that shows the disorder, it is possible that the disorder does not manifests itself in all windows. Therefore, the training data that we have used includes windows where the patients looks healthy. In other words, the labeling method we have used adds noise to the training data.

A possible solution is to label the patients for each task, or selecting the best tasks per patient. For each task, we can use only the patients of which the disorder manifests itself in this task. It would be even better to label the patients per window. Unfortunately, this is a time consuming task. Moreover, it requires a medical expert to do so. If we have labels for each task it is an option to use the whole video instead of windows. However, we do not have enough patients to train a network in this way.

Another solution is to take the confidence of the classifier into account. It is possible that the classifier is able to learn to recognize the disorder, but has a low precision and recall due to the 'healthy' windows it has to classify from the test set. If a high confidence corresponds to a clearly visible disorder, we can maybe find a threshold and classify only the windows that exceed this threshold.

Alternatively, we could add a healthy class to the classifier. In this case we would train the classifier on healthy patients as well such that the classifier can learn the difference between the disorders, as well the healthy windows. Hopefully, the windows that do not clearly show the disorder will be classified as healthy, so the windows classified as a disorder contain less noise.

The second problem with the training data is the number of patients. We have divided the videos into windows to provide enough data to the classifier. Another approach would be to combine tasks. Since we have 36 tasks per patient, we could significantly increase the number of videos. However, the tasks are very different. Therefore, it might be interesting to combine the tasks in groups that are similar enough to feed to the same classifier. The simplest example are the tasks performed with both the right and left arm. These tasks are very similar and could even be mirrored to provide the classifier with simple data. Some of the patients show the disorder in a different way depending on which arm they use to perform the task. Hence, we expect it to give additional information if the two tasks are combined.

Problem 2: Movement in the task
We are trying to classify movement disorder. These manifest itself in small involuntary movements. However, the tasks the patients are performing consist of movement itself. We have found that the movement of the task is too dominant in the feature vector and thus in the classifier.

There are several solutions we can think of. One would be to ignore the tasks that contain movement, however this would result in a lot of data that we can not use. Another possible solution is to combine all tasks. This way the classifier cannot use the movements of the tasks and might be able to find the disorder. However, we expect we do not have enough data for the classifier to learn what is the common factor in the data. Alternatively, it could help to have a third class of healthy patients as discussed before. This way the task is the norm and anything deviating from the task might contain information. Lastly, it would be interesting to think of a way where we can subtract the data of a healthy patient, so that we only feed the other information to the classifier. However, this is a nontrivial task. Especially due to the fact that every patient performs the task following another path and pace.

Problem 3: Capture small movements

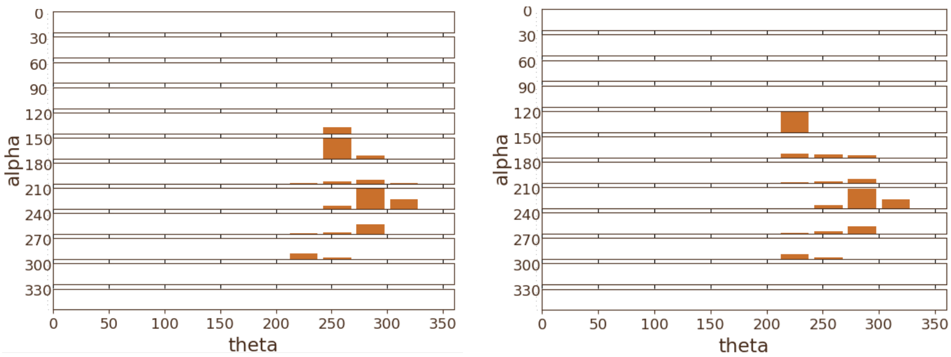

(a) Right wrist

(b) Right hand

Figure 18: Distribution of the location of the right hand and wrist. Origin right shoulder, reference right elbow

Even in the static tasks it is difficult to find the important movement. Some possible causes are the distance of the camera to the patient, the resolution of the camera and the skeleton tracking that can make some mistakes. The second option is easily solved by buying a better camera.

Moreover, we compute the location of every joint with respect to the same origin and reference. This has as a result that the small movements are not very present. If we compute the location of the hand with respect to the waist, as in Figure 18b, the movement that is shown is mostly determined by the movement of the wrist shown in Figure 18a. A possible solution is to change the computation of the location such that we take for every joint a specific origin and reference. For example, the position of the hand is computed with respect to the lower arm, and the position of the wrist is computed with respect to to the upper arm, etc.

Another option is to use motion magnification. Here we can amplify the motion of the disorder. Since tremor is a rhythmic movement we can amplify the frequencies that are known to relate to tremor. However, for the other disorders it is not trivial to see what can be used to amplify the important movements and not those of the task itself.

Alternatively, we could start over and construct a completely different feature vector. Ideally one that is invariant of task. One of the options is to zoom in on the hands and track a bounding box around it. Then run an algorithm on this box, for example with optical flow.

5 Conclusion

In this report we have done a lot of experiments to find out more about the classifier we were using. We explored different focus points by changing the origin and the body parts that were given as input. We have done research on the feature vector mostly focused on the smoothing that was present. After that we did a extensive parameter search. Unfortunately, none of these experiments gave us the results we were looking for.

In the last part of the experiments we decided to try to classify dystonia instead of tremor. Moreover, we changed the task we focused on for a new task in which the patient was asked to hold the finger still in front of the nose instead of moving it. Combining these two experiments, we were able to find a good result where five out of eight pairs of test patients were classified correctly. Based on this result we can conclude that the original smoothing with the original parameters, the right shoulder as origin, the right elbow as reference focused on only the right arm is the best setting. Moreover, we should use a static task. However, note that we have performed the experiments only for 2 out of 36 tasks. Hence, more research on the other tasks is needed to know if these results can be generalized to all tasks.

During the experiments we have found certain parts that could be improved. For example, the labels are per patient and the classification is per window. It would be better if these were both only on the patient or both per window. Moreover, we expect to have better results if a third class is added to the classifier consisting of healthy patients. This does not only neutralize the movement of the task, it also gives more training data.

6 Bibliography

[1] Nuitrack tracked body parts. https://community.nuitrack.com/t/joints-names-and-locations-scheme/685.

[2] Specifications realsense camera. https://www.intelrealsense.com/depth-camera-d415/.

[3] Lu Xia, Chia-Chih Chen, and Jake K Aggarwal. View invariant human action recognition using histograms of 3d joints. In *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 20–27. IEEE, 2012.

[4] Colah's blog. Understanding lstm. http://colah.github.io/posts/2015-08-Understanding-LSTMs/.

Appendix A Results experiments

				24 and 14				49 and 42				32 and 17			
				Precision		Recall		Precision		Recall		Precision		Recall	
Task	Joints	Orientations	Offset	M	T	M	T	M	T	M	T	M	T	M	T
22	All	12	5	0,74	0,73	0,72	0,74	0,65	0,55	0,32	0,83	0,49	0,41	0,76	0,21
		18	10	0,77	0,63	0,50	0,84	0,52	0,51	0,38	0,64	0,46	0,36	0,69	0,19
	Arms	12	5	0,72	0,67	0,65	0,74	0,71	0,59	0,43	0,83	0,48	0,36	0,82	0,12
		18	10	0,80	0,64	0,52	0,87	0,60	0,58	0,54	0,63	0,47	0,39	0,74	0,17
	Right arm	12	5	0,63	0,59	0,52	0,68	0,28	0,37	0,20	0,48	0,38	0,36	0,42	0,33
		18	10	0,54	0,51	0,42	0,61	0,39	0,39	0,38	0,39	0,48	0,48	0,49	0,47
24	All	12	5	0,81	0,84	0,84	0,80	0,70	0,59	0,46	0,80	0,49	0,00	0,97	0,00
		18	10	0,81	0,68	0,59	0,85	0,76	0,65	0,53	0,81	0,49	0,00	0,97	0,00
	Arms	12	5	0,81	0,78	0,76	0,83	0,78	0,63	0,49	0,86	0,50	0,00	0,98	0,00
		18	10	0,82	0,69	0,59	0,87	0,84	0,60	0,38	0,93	0,49	0,00	0,98	0,00
	Right arm	12	5	0,84	0,81	0,80	0,84	0,86	0,61	0,39	0,94	0,50	0,00	0,99	0,00
		18	10	0,86	0,76	0,71	0,89	0,90	0,63	0,44	0,95	0,49	0,00	0,98	0,00

(a) Results for myoclonus and tremor.

		24 and 46				49 and 21				36 and 15			
		Precision		Recall		Precision		Recall		Precision		Recall	
Joints	Task	M	D	M	D	M	D	M	D	M	D	M	D
All	22	0,59	0,75	0,87	0,40	0,51	0,53	0,71	0,32	0,65	0,63	0,62	0,65
	24	0,51	0,52	0,82	0,20	0,36	0,45	0,26	0,60	0,46	0,28	0,76	0,11
Arms	22	0,56	0,74	0,90	0,29	0,43	0,40	0,53	0,32	0,65	0,72	0,77	0,59
	24	0,52	0,56	0,76	0,29	0,45	0,47	0,55	0,35	0,46	0,17	0,81	0,04
Right arm	22	0,65	0,66	0,68	0,61	0,47	0,46	0,54	0,39	0,69	0,71	0,73	0,66
	24	0,27	0,29	0,27	0,35	0,47	0,42	0,75	0,19	0,23	0,21	0,34	0,23

(b) Results for myoclonus and dystonia.

Figure 19: Results for the second derivative.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,27	0,44	0,22	0,80	0,71	0,51	0,83	0,44	0,16	0,00	0,28	0,00
	Right elbow	Right shoulder	0,54	0,60	1,00	0,12	0,00	0,50	0,00	0,99	1,00	1,00	1,00	1,00
	Right elbow	Right wrist	1,00	1,00	1,00	1,00	0,00	0,50	0,00	0,99	0,64	0,71	0,72	0,61
	Right shoulder	Right elbow	0,51	0,57	0,48	0,59	0,59	0,56	0,64	0,52	0,60	1,00	1,00	0,34
No smoothing	Waist	Torso	0,35	0,22	0,45	0,16	0,77	0,84	0,85	0,73	0,36	0,08	0,57	0,02
	Right elbow	Right shoulder	0,79	0,67	0,59	0,82	0,67	0,55	0,38	0,76	0,48	0,44	0,63	0,30
	Right elbow	Right wrist	0,78	0,68	0,62	0,82	0,75	0,56	0,28	0,90	0,34	0,33	0,35	0,33
	Right shoulder	Right elbow	0,50	0,57	0,81	0,22	0,60	0,60	0,59	0,60	0,58	0,68	0,79	0,43

(a) Results for arms.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,40	0,51	0,03	1,00	0,91	1,00	1,00	0,86	0,49	0,00	0,97	0,00
	Right elbow	Right shoulder	0,49	0,50	0,41	0,58	0,38	0,35	0,43	0,30	0,18	0,04	0,22	0,04
	Right elbow	Right wrist	0,73	0,58	0,41	0,82	0,03	0,08	0,02	0,09	0,16	0,00	0,19	0,00
	Right shoulder	Right elbow	0,39	0,41	0,36	0,44	0,42	0,20	0,61	0,14	0,29	0,04	0,47	0,02
No smoothing	Waist	Torso	0,81	0,64	0,44	0,89	0,75	0,68	0,63	0,79	0,43	0,00	0,75	0,00
	Right elbow	Right shoulder	0,13	0,39	0,07	0,59	0,33	0,24	0,32	0,28	0,23	0,27	0,21	0,29
	Right elbow	Right wrist	0,73	0,59	0,41	0,85	0,20	0,25	0,18	0,28	0,31	0,34	0,29	0,37
	Right shoulder	Right elbow	0,22	0,41	0,11	0,63	0,39	0,33	0,47	0,27	0,19	0,26	0,17	0,30

(b) Results for right arm.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,40	0,64	0,31	1,00	0,85	0,99	0,99	0,79	0,05	0,00	0,05	0,00
	Right elbow	Right shoulder	0,92	1,00	1,00	0,90	0,78	0,51	0,06	0,99	0,86	0,79	0,75	0,87
	Right elbow	Right wrist	1,00	0,93	0,92	1,00	0,00	0,47	0,00	0,91	0,27	0,06	0,38	0,04
	Right shoulder	Right elbow	0,87	0,57	0,28	0,93	0,43	0,44	0,49	0,38	0,55	0,82	0,90	0,25
No smoothing	Waist	Torso	1,00	0,65	0,46	1,00	0,81	0,80	0,81	0,81	0,11	0,00	0,16	0,00
	Right elbow	Right shoulder	0,71	0,65	0,60	0,75	0,53	0,51	0,34	0,68	0,28	0,32	0,25	0,36
	Right elbow	Right wrist	0,78	0,70	0,65	0,81	0,49	0,49	0,25	0,73	0,32	0,37	0,27	0,43
	Right shoulder	Right elbow	0,55	0,54	0,40	0,68	0,25	0,34	0,19	0,42	0,60	0,58	0,53	0,65

(c) Results for right arm + left hand.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,72	0,81	0,65	0,86	0,54	0,20	1,00	0,12	0,50	0,00	1,00	0,00
	Right elbow	Right shoulder	0,52	0,51	0,96	0,11	0,00	0,50	0,00	1,00	0,90	1,00	1,00	0,87
	Right elbow	Right wrist	0,94	1,00	1,00	0,94	0,00	0,43	0,00	0,77	0,47	0,34	0,69	0,20
	Right shoulder	Right elbow	0,61	0,83	0,89	0,43	0,46	0,30	0,66	0,21	0,60	1,00	1,00	0,33
No smoothing	Waist	Torso	0,31	0,36	0,46	0,18	0,69	0,68	0,64	0,72	0,52	0,60	0,87	0,21
	Right elbow	Right shoulder	0,75	0,62	0,49	0,82	0,43	0,44	0,40	0,47	0,42	0,33	0,57	0,22
	Right elbow	Right wrist	0,73	0,62	0,51	0,81	0,69	0,56	0,31	0,86	0,44	0,42	0,53	0,33
	Right shoulder	Right elbow	0,54	0,75	0,92	0,23	0,52	0,51	0,49	0,54	0,62	0,78	0,85	0,46

(d) Results for arms - left hand.

Figure 20: Results for body parts.

Smoothing	Origin	Reference	24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
			M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	1,00	0,90	0,80	1,00	0,60	0,20	1,00	0,20	0,08	0,00	0,12	0,00
	Right elbow	Right shoulder	0,71	1,00	1,00	0,55	0,00	0,50	0,00	1,00	1,00	1,00	1,00	1,00
	Right elbow	Right wrist	0,99	1,00	1,00	0,99	0,00	0,50	0,00	1,00	0,48	0,59	0,85	0,14
	Right shoulder	Right elbow	0,24	0,23	0,27	0,21	0,60	0,40	1,00	0,20	0,00	0,10	0,00	0,19
No smoothing	Waist	Torso	0,96	0,77	0,68	0,97	0,75	0,73	0,69	0,73	0,21	0,00	0,28	0,00
	Right elbow	Right shoulder	0,88	0,70	0,61	0,92	0,60	0,53	0,34	0,76	0,54	0,55	0,64	0,46
	Right elbow	Right wrist	0,90	0,76	0,70	0,92	0,47	0,46	0,27	0,63	0,40	0,41	0,39	0,43
	Right shoulder	Right elbow	0,37	0,39	0,37	0,40	0,61	0,88	0,94	0,38	0,46	0,46	0,51	0,41

(a) Results for 6 orientations.

Smoothing	Origin	Reference	24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
			M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,27	0,44	0,22	0,80	0,71	0,51	0,83	0,44	0,16	0,00	0,28	0,00
	Right elbow	Right shoulder	0,54	0,60	1,00	0,12	0,00	0,50	0,00	0,99	1,00	1,00	1,00	1,00
	Right elbow	Right wrist	1,00	1,00	1,00	1,00	0,00	0,50	0,00	0,99	0,64	0,71	0,72	0,61
	Right shoulder	Right elbow	0,51	0,57	0,48	0,59	0,59	0,56	0,64	0,52	0,60	1,00	1,00	0,34
No smoothing	Waist	Torso	0,35	0,22	0,45	0,16	0,77	0,84	0,85	0,73	0,36	0,08	0,57	0,02
	Right elbow	Right shoulder	0,79	0,67	0,59	0,82	0,67	0,55	0,38	0,76	0,48	0,44	0,63	0,30
	Right elbow	Right wrist	0,78	0,68	0,62	0,82	0,75	0,56	0,28	0,90	0,34	0,33	0,35	0,33
	Right shoulder	Right elbow	0,50	0,57	0,81	0,22	0,60	0,60	0,59	0,60	0,58	0,68	0,79	0,43

(b) Results for 12 orientations.

Smoothing	Origin	Reference	24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
			M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,39	0,42	0,36	0,71	0,77	0,90	0,80	0,96	0,21	0,00	0,41	0,00
	Right elbow	Right shoulder	0,78	1,00	1,00	0,71	0,00	0,50	0,00	1,00	0,95	0,99	0,99	0,94
	Right elbow	Right wrist	1,00	1,00	1,00	1,00	0,00	0,42	0,00	0,75	0,56	0,64	0,74	0,43
	Right shoulder	Right elbow	0,40	0,17	0,56	0,13	0,59	0,62	0,65	0,55	0,00	0,01	0,00	0,01
No smoothing	Waist	Torso	0,48	0,46	0,37	0,57	0,84	0,81	0,81	0,83	0,19	0,07	0,22	0,06
	Right elbow	Right shoulder	0,86	0,71	0,63	0,90	0,49	0,49	0,29	0,69	0,46	0,45	0,54	0,37
	Right elbow	Right wrist	0,89	0,74	0,68	0,92	0,68	0,56	0,34	0,84	0,35	0,32	0,36	0,32
	Right shoulder	Right elbow	0,39	0,36	0,45	0,31	0,60	0,67	0,75	0,48	0,13	0,02	0,15	0,02

(c) Results for 18 orientations.

Smoothing	Origin	Reference	24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
			M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,00	0,50	0,00	1,00	0,98	0,88	0,78	0,98	0,21	0,00	0,42	0,00
	Right elbow	Right shoulder	0,53	0,78	1,00	0,11	0,00	0,50	0,00	1,00	0,74	0,61	0,48	0,83
	Right elbow	Right wrist	0,92	0,99	0,99	0,91	0,02	0,44	0,01	0,79	0,36	0,18	0,52	0,11
	Right shoulder	Right elbow	0,57	0,68	0,75	0,45	0,67	0,57	0,43	0,77	0,15	0,22	0,14	0,28
No smoothing	Waist	Torso	0,40	0,36	0,41	0,36	0,56	0,66	0,87	0,30	0,51	0,50	0,71	0,31
	Right elbow	Right shoulder	0,72	0,69	0,67	0,75	0,78	0,54	0,21	0,93	0,45	0,40	0,60	0,26
	Right elbow	Right wrist	0,81	0,75	0,72	0,83	0,82	0,57	0,31	0,93	0,44	0,41	0,47	0,38
	Right shoulder	Right elbow	0,41	0,34	0,49	0,28	0,52	0,52	0,47	0,56	0,36	0,32	0,38	0,31

(d) Results for 24 orientations.

Figure 21: Results for varying the number of orientations, arms.

Smoothing	Origin	Reference	24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
			M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,77	0,46	0,41	0,76	0,95	0,99	0,99	0,94	0,40	0,00	0,80	0,00
	Right elbow	Right shoulder	0,71	1,00	1,00	0,57	0,00	0,50	0,00	0,99	0,96	1,00	1,00	0,96
	Right elbow	Right wrist	1,00	0,99	0,99	1,00	0,00	0,40	0,00	0,67	0,58	0,71	0,85	0,37
	Right shoulder	Right elbow	0,38	0,06	0,59	0,03	0,57	0,56	0,52	0,61	0,00	0,03	0,00	0,03
No smoothing	Waist	Torso	0,46	0,43	0,37	0,51	0,81	0,71	0,64	0,85	0,14	0,06	0,16	0,05
	Right elbow	Right shoulder	0,80	0,74	0,72	0,82	0,48	0,48	0,26	0,70	0,47	0,47	0,58	0,37
	Right elbow	Right wrist	0,85	0,70	0,63	0,89	0,56	0,52	0,33	0,73	0,42	0,43	0,38	0,47
	Right shoulder	Right elbow	0,40	0,32	0,50	0,24	0,50	0,49	0,54	0,44	0,18	0,00	0,23	0,00

(a) Results for window size 10.

Smoothing	Origin	Reference	24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
			M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,39	0,42	0,36	0,71	0,77	0,90	0,80	0,96	0,21	0,00	0,41	0,00
	Right elbow	Right shoulder	0,78	1,00	1,00	0,71	0,00	0,50	0,00	1,00	0,95	0,99	0,99	0,94
	Right elbow	Right wrist	1,00	1,00	1,00	1,00	0,00	0,42	0,00	0,75	0,56	0,64	0,74	0,43
	Right shoulder	Right elbow	0,40	0,17	0,56	0,13	0,59	0,62	0,65	0,55	0,00	0,01	0,00	0,01
No smoothing	Waist	Torso	0,48	0,46	0,37	0,57	0,84	0,81	0,81	0,83	0,19	0,07	0,22	0,06
	Right elbow	Right shoulder	0,86	0,71	0,63	0,90	0,49	0,49	0,29	0,69	0,46	0,45	0,54	0,37
	Right elbow	Right wrist	0,89	0,74	0,68	0,92	0,68	0,56	0,34	0,84	0,35	0,32	0,36	0,32
	Right shoulder	Right elbow	0,39	0,36	0,45	0,31	0,60	0,67	0,75	0,48	0,13	0,02	0,15	0,02

(b) Results for window size 15.

Smoothing	Origin	Reference	24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
			M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,74	0,60	0,90	0,58	0,89	0,77	0,97	0,79	0,44	0,00	0,85	0,00
	Right elbow	Right shoulder	0,86	1,00	1,00	0,80	0,00	0,50	0,00	1,00	0,94	0,95	0,93	0,94
	Right elbow	Right wrist	1,00	1,00	1,00	1,00	0,00	0,43	0,00	0,76	0,54	0,64	0,70	0,41
	Right shoulder	Right elbow	0,35	0,00	0,55	0,00	0,68	0,70	0,73	0,58	0,00	0,03	0,00	0,03
No smoothing	Waist	Torso	0,78	0,63	0,51	0,84	0,66	0,71	0,77	0,59	0,13	0,12	0,15	0,13
	Right elbow	Right shoulder	0,88	0,73	0,67	0,91	0,44	0,48	0,26	0,67	0,60	0,58	0,56	0,61
	Right elbow	Right wrist	0,98	0,78	0,72	0,99	0,69	0,56	0,31	0,86	0,42	0,43	0,45	0,41
	Right shoulder	Right elbow	0,38	0,31	0,47	0,25	0,57	0,66	0,74	0,46	0,10	0,00	0,12	0,00

(c) Results for window size 30.

Smoothing	Origin	Reference	24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
			M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,43	0,45	0,38	0,75	0,93	0,90	0,88	0,91	0,30	0,00	0,60	0,00
	Right elbow	Right shoulder	0,88	1,00	1,00	0,86	0,20	0,51	0,04	1,00	1,00	1,00	1,00	1,00
	Right elbow	Right wrist	1,00	1,00	1,00	1,00	0,00	0,37	0,00	0,63	0,52	0,45	0,96	0,11
	Right shoulder	Right elbow	0,39	0,01	0,66	0,01	0,92	0,85	0,80	0,91	0,00	0,05	0,00	0,06
No smoothing	Waist	Torso	0,30	0,33	0,17	0,45	0,73	0,65	0,63	0,71	0,03	0,01	0,03	0,01
	Right elbow	Right shoulder	0,92	0,84	0,82	0,92	0,79	0,65	0,61	0,73	0,56	0,52	0,63	0,42
	Right elbow	Right wrist	0,97	0,79	0,72	0,97	0,38	0,54	0,24	0,86	0,50	0,45	0,40	0,55
	Right shoulder	Right elbow	0,35	0,12	0,53	0,11	0,59	0,63	0,73	0,45	0,08	0,00	0,09	0,00

(d) Results for window size 60.

Figure 22: Results for varying window size, 18 orientations.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,39	0,35	0,53	0,59	0,78	0,94	0,92	0,64	0,50	0,00	1,00	0,00
	Right elbow	Right shoulder	0,64	1,00	1,00	0,40	0,00	0,50	0,00	1,00	0,99	1,00	1,00	0,99
	Right elbow	Right wrist	0,99	1,00	1,00	0,99	0,00	0,38	0,00	0,61	0,59	0,67	0,79	0,44
	Right shoulder	Right elbow	0,42	0,02	0,72	0,01	0,70	0,58	0,50	0,74	0,00	0,00	0,00	0,00
No smoothing	Waist	Torso	0,52	0,45	0,44	0,53	0,91	0,80	0,76	0,93	0,09	0,04	0,10	0,04
	Right elbow	Right shoulder	0,85	0,86	0,85	0,85	0,54	0,50	0,32	0,70	0,31	0,32	0,31	0,33
	Right elbow	Right wrist	0,84	0,83	0,82	0,85	0,67	0,56	0,37	0,80	0,33	0,35	0,30	0,39
	Right shoulder	Right elbow	0,44	0,33	0,65	0,18	0,57	0,60	0,70	0,46	0,13	0,00	0,14	0,00

(a) Results for window offset 1.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,86	0,56	0,60	0,79	0,79	0,98	0,98	0,65	0,41	0,00	0,76	0,00
	Right elbow	Right shoulder	0,76	0,98	0,99	0,65	0,00	0,50	0,00	0,99	1,00	0,98	0,98	0,99
	Right elbow	Right wrist	0,99	0,99	0,99	0,99	0,00	0,49	0,00	0,95	0,58	0,73	0,80	0,43
	Right shoulder	Right elbow	0,45	0,39	0,57	0,29	0,70	0,64	0,59	0,74	0,00	0,00	0,00	0,00
No smoothing	Waist	Torso	0,67	0,54	0,39	0,76	0,84	0,81	0,81	0,84	0,18	0,01	0,24	0,01
	Right elbow	Right shoulder	0,89	0,73	0,67	0,92	0,46	0,45	0,22	0,66	0,59	0,56	0,54	0,61
	Right elbow	Right wrist	0,87	0,72	0,64	0,90	0,42	0,43	0,31	0,55	0,42	0,45	0,41	0,47
	Right shoulder	Right elbow	0,36	0,38	0,35	0,40	0,49	0,49	0,50	0,49	0,09	0,00	0,10	0,00

(b) Results for window offset 10.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,38	0,53	0,38	0,62	0,82	1,00	1,00	0,73	0,50	0,00	1,00	0,00
	Right elbow	Right shoulder	0,56	1,00	1,00	0,21	0,00	0,50	0,00	1,00	1,00	1,00	1,00	1,00
	Right elbow	Right wrist	1,00	1,00	1,00	1,00	0,00	0,46	0,00	0,84	0,64	0,74	0,80	0,55
	Right shoulder	Right elbow	0,40	0,11	0,63	0,08	0,69	0,71	0,71	0,67	0,00	0,10	0,00	0,11
No smoothing	Waist	Torso	0,42	0,45	0,26	0,62	0,76	0,73	0,74	0,73	0,27	0,11	0,37	0,08
	Right elbow	Right shoulder	0,84	0,85	0,86	0,83	0,47	0,49	0,26	0,70	0,53	0,51	0,51	0,53
	Right elbow	Right wrist	0,85	0,70	0,61	0,89	0,51	0,52	0,30	0,74	0,49	0,48	0,44	0,53
	Right shoulder	Right elbow	0,35	0,23	0,44	0,19	0,50	0,52	0,63	0,39	0,14	0,00	0,16	0,00

(c) Results for window offset 15.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,27	0,50	0,32	0,77	0,80	0,78	0,98	0,61	0,31	0,00	0,61	0,00
	Right elbow	Right shoulder	0,78	1,00	1,00	0,66	0,00	0,50	0,00	0,99	1,00	1,00	1,00	1,00
	Right elbow	Right wrist	0,96	1,00	1,00	0,95	0,03	0,46	0,01	0,84	0,67	0,84	0,79	0,53
	Right shoulder	Right elbow	0,52	0,42	0,65	0,37	0,61	0,61	0,60	0,62	0,00	0,00	0,00	0,00
No smoothing	Waist	Torso	0,60	0,53	0,35	0,74	0,84	0,76	0,73	0,86	0,35	0,03	0,55	0,02
	Right elbow	Right shoulder	0,85	0,67	0,56	0,89	0,60	0,51	0,35	0,69	0,47	0,49	0,55	0,41
	Right elbow	Right wrist	0,88	0,76	0,70	0,90	0,36	0,36	0,36	0,37	0,45	0,44	0,45	0,44
	Right shoulder	Right elbow	0,42	0,43	0,43	0,43	0,53	0,54	0,52	0,55	0,20	0,00	0,26	0,00

(d) Results for window offset 20.

Figure 23: Results for varying window offset, with window size 15, original smoothing.

			24 and 46				49 and 21				36 and 15			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	D	M	D	M	D	M	D	M	D	M	D
Original	Waist	Torso	0,50	0,00	1,00	0,00	0,50	0,00	1,00	0,00	0,28	0,00	0,52	0,00
	Right elbow	Right shoulder	0,05	0,22	0,04	0,27	0,44	0,00	0,80	0,00	0,51	0,20	1,00	0,04
	Right elbow	Right wrist	0,71	0,79	0,81	0,59	0,61	0,84	0,90	0,43	0,55	0,75	0,99	0,17
	Right shoulder	Right elbow	0,41	0,36	0,51	0,27	0,51	0,70	0,99	0,05	0,00	0,01	0,00	0,01
No smoothing	Waist	Torso	0,50	0,00	0,99	0,00	0,44	0,00	0,81	0,00	0,14	0,00	0,16	0,00
	Right elbow	Right shoulder	0,37	0,37	0,36	0,37	0,33	0,18	0,43	0,13	0,63	0,94	0,97	0,42
	Right elbow	Right wrist	0,61	0,58	0,50	0,68	0,59	0,66	0,78	0,45	0,53	0,54	0,64	0,42
	Right shoulder	Right elbow	0,43	0,39	0,53	0,30	0,49	0,07	0,95	0,01	0,02	0,00	0,02	0,00

(a) Results for arms, 18 orientations, offset 10.

			24 and 46				49 and 21				36 and 15			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	D	M	D	M	D	M	D	M	D	M	D
Original	Waist	Torso	0,50	0,00	1,00	0,00	0,57	0,20	1,00	0,16	0,20	0,00	0,33	0,00
	Right elbow	Right shoulder	0,04	0,21	0,03	0,26	0,46	0,00	0,88	0,00	0,50	0,00	1,00	0,00
	Right elbow	Right wrist	0,74	0,75	0,76	0,71	0,43	0,46	0,51	0,36	0,48	0,21	0,92	0,01
	Right shoulder	Right elbow	0,00	0,16	0,00	0,20	0,50	0,40	1,00	0,01	0,40	0,55	0,34	0,74
No smoothing	Waist	Torso	0,54	0,80	1,00	0,13	0,49	0,28	0,96	0,01	0,20	0,00	0,28	0,00
	Right elbow	Right shoulder	0,45	0,46	0,37	0,54	0,47	0,44	0,65	0,28	0,56	0,72	0,87	0,32
	Right elbow	Right wrist	0,61	0,56	0,46	0,69	0,43	0,41	0,51	0,34	0,62	0,75	0,81	0,51
	Right shoulder	Right elbow	0,02	0,31	0,01	0,46	0,45	0,10	0,81	0,02	0,38	0,35	0,41	0,33

(b) Results for arms, 12 orientations, offset 5.

			24 and 46				49 and 21				36 and 15			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	D	M	D	M	D	M	D	M	D	M	D
Original	Waist	Torso	0,44	0,27	0,76	0,06	0,39	0,11	0,62	0,07	0,50	0,00	1,00	0,00
	Right elbow	Right shoulder	0,68	0,63	0,54	0,73	0,55	0,59	0,78	0,36	0,66	0,87	0,92	0,52
	Right elbow	Right wrist	0,98	0,78	0,70	0,98	0,17	0,35	0,13	0,49	0,69	0,68	0,69	0,66
	Right shoulder	Right elbow	0,43	0,34	0,58	0,22	0,49	0,25	0,91	0,06	0,66	0,82	0,88	0,55
No smoothing	Waist	Torso	0,46	0,42	0,66	0,25	0,43	0,31	0,65	0,15	0,52	0,69	0,94	0,13
	Right elbow	Right shoulder	0,70	0,57	0,35	0,85	0,48	0,46	0,63	0,33	0,68	0,82	0,87	0,58
	Right elbow	Right wrist	0,66	0,55	0,33	0,84	0,43	0,43	0,40	0,45	0,68	0,86	0,91	0,57
	Right shoulder	Right elbow	0,30	0,41	0,20	0,55	0,45	0,40	0,57	0,29	0,76	0,95	0,96	0,70

(c) Results for right arm, 18 orientations, offset 10.

			24 and 46				49 and 21				36 and 15			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	D	M	D	M	D	M	D	M	D	M	D
Original	Waist	Torso	0,50	0,10	0,95	0,05	0,67	0,54	0,91	0,44	0,71	1,00	1,00	0,55
	Right elbow	Right shoulder	0,72	0,76	0,77	0,70	0,62	0,65	0,70	0,56	0,54	0,65	0,85	0,28
	Right elbow	Right wrist	0,71	0,84	0,86	0,63	0,04	0,17	0,04	0,21	0,05	0,18	0,05	0,21
	Right shoulder	Right elbow	0,50	0,48	0,52	0,47	0,58	0,59	0,99	0,24	0,58	0,84	0,92	0,34
No smoothing	Waist	Torso	0,61	0,78	0,88	0,43	0,50	0,49	0,62	0,38	0,45	0,33	0,71	0,15
	Right elbow	Right shoulder	0,83	0,68	0,57	0,88	0,54	0,55	0,63	0,46	0,65	0,81	0,88	0,52
	Right elbow	Right wrist	0,58	0,55	0,50	0,63	0,36	0,28	0,44	0,22	0,65	0,93	0,96	0,49
	Right shoulder	Right elbow	0,22	0,40	0,11	0,60	0,54	0,62	0,79	0,32	0,70	0,88	0,92	0,61

(d) Results for right arm, 12 orientations, offset 5.

Figure 24: Results for myoclonus versus dystonia.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,00	0,50	0,00	1,00	0,50	0,00	1,00	0,00	0,00	0,50	0,00	1,00
	Right elbow	Right shoulder	0,40	0,70	0,40	1,00	0,20	0,00	0,40	0,00	1,00	1,00	1,00	1,00
	Right elbow	Right wrist	0,00	0,20	0,00	0,40	0,47	0,71	0,60	0,58	0,80	0,90	0,80	1,00
	Right shoulder	Right elbow	0,00	0,00	0,00	0,00	0,00	0,43	0,00	0,78	0,20	0,10	0,40	0,20
No smoothing	Waist	Torso	1,00	0,76	0,68	1,00	0,53	1,00	1,00	0,11	0,40	0,51	0,02	1,00
	Right elbow	Right shoulder	0,80	0,53	0,10	1,00	0,31	0,13	0,41	0,10	0,65	0,40	0,58	0,52
	Right elbow	Right wrist	0,50	0,13	0,57	0,20	0,57	1,00	1,00	0,24	1,00	0,77	0,68	1,00
	Right shoulder	Right elbow	0,00	0,00	0,00	0,00	0,33	0,47	0,14	0,76	1,00	0,76	0,59	1,00

(a) Results for tremor, arms, 12 orientations, offset 5.

			24 and 14				49 and 42				32 and 17			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,50	0,00	1,00	0,00	0,57	0,40	1,00	0,16	0,00	0,50	0,00	1,00
	Right elbow	Right shoulder	0,90	0,80	1,00	0,80	0,51	0,60	1,00	0,03	0,80	0,90	0,80	1,00
	Right elbow	Right wrist	0,40	0,70	0,40	1,00	0,40	0,00	0,80	0,00	0,00	0,50	0,00	1,00
	Right shoulder	Right elbow	0,50	0,00	1,00	0,00	0,39	0,69	0,40	0,98	0,00	0,40	0,00	0,80
No smoothing	Waist	Torso	0,22	0,34	0,40	0,24	0,54	0,69	0,90	0,22	0,00	0,35	0,00	0,63
	Right elbow	Right shoulder	0,80	0,80	0,80	0,80	0,52	0,52	0,81	0,25	0,13	0,47	0,08	0,81
	Right elbow	Right wrist	1,00	1,00	1,00	1,00	0,30	0,39	0,48	0,14	0,00	0,10	0,00	0,20
	Right shoulder	Right elbow	0,50	0,00	1,00	0,00	0,39	0,45	0,19	0,66	0,01	0,26	0,00	0,36

(b) Results for tremor, rightarm, 12 orientations, offset 5.

			24 and 46				49 and 21				36 and 15			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	T	M	T	M	T	M	T	M	T	M	T
Original	Waist	Torso	0,00	0,20	0,00	0,40	0,00	0,50	0,00	1,00	0,50	0,00	1,00	0,00
	Right elbow	Right shoulder	0,60	0,20	1,00	0,20	0,00	0,10	0,00	0,20	0,00	0,07	0,00	0,12
	Right elbow	Right wrist	0,40	0,00	0,80	0,00	0,00	0,40	0,00	0,80	0,51	0,40	1,00	0,04
	Right shoulder	Right elbow	0,50	0,00	1,00	0,00	0,30	0,10	0,60	0,20	0,65	0,74	0,92	0,41
No smoothing	Waist	Torso	0,52	0,23	0,24	0,40	0,20	0,50	0,02	1,00	0,51	0,57	0,93	0,11
	Right elbow	Right shoulder	0,79	0,54	0,90	0,60	0,06	0,00	0,09	0,00	0,30	0,19	0,43	0,22
	Right elbow	Right wrist	0,59	0,23	0,40	0,40	0,70	0,50	0,46	0,70	0,39	0,41	0,54	0,19
	Right shoulder	Right elbow	0,50	0,00	1,00	0,00	0,27	0,09	0,48	0,18	0,45	0,44	0,43	0,47

(c) Results for tremor, rightarm, 18 orientations, offset 10.

			24 and 46				49 and 21				36 and 15			
			Precision		Recall		Precision		Recall		Precision		Recall	
Smoothing	Origin	Reference	M	D	M	D	M	D	M	D	M	D	M	D
Original	Waist	Torso	0,50	0,00	1,00	0,00	0,00	0,50	0,00	1,00	1,00	1,00	1,00	1,00
	Right elbow	Right shoulder	0,00	0,50	0,00	1,00	0,00	0,50	0,00	1,00	0,00	0,50	0,00	1,00
	Right elbow	Right wrist	0,00	0,14	0,00	0,25	0,70	0,70	0,80	0,80	1,00	1,00	1,00	1,00
	Right shoulder	Right elbow	0,94	1,00	1,00	0,91	1,00	1,00	1,00	1,00	0,00	0,50	0,00	1,00
No smoothing	Waist	Torso	0,50	1,00	1,00	0,02	1,00	0,72	0,58	1,00	0,72	0,86	0,81	0,57
	Right elbow	Right shoulder	0,09	0,21	0,18	0,24	0,07	0,34	0,09	0,49	0,00	0,24	0,00	0,32
	Right elbow	Right wrist	0,31	0,11	0,45	0,09	0,48	0,60	0,58	0,44	0,27	0,02	0,36	0,01
	Right shoulder	Right elbow	0,54	0,40	1,00	0,14	0,59	1,00	1,00	0,29	0,75	0,74	0,73	0,75

(d) Results for dystonia, rightarm, 18 orientations, offset 10.

Figure 25: Results for task 24.