

***Using a V-Shaped Learning Algorithm to
Predict Stability Control on Curved Objects***

Bellamie Persad

S4355474

Bachelors Thesis in Artificial Intelligence

Radboud University Nijmegen

Supervisor:

Dr. ing. L.P.J. Selen

Radboud University Nijmegen

Artificial Intelligence

January 2017

Abstract

In this study, research was done on the model for motor learning as described by Franklin et al. (2008) to see whether the model was able to predict stability on curved objects. The results from the model were promising. However, the model was based on a task that involved performing a forward reaching movement in different force fields, such as a divergent field, which is based on external forces. In our research we looked if this model was able to predict stability in a situation where the instability came from task geometry and signal-dependent noise. In the simulation, three different object with varying curvatures were simulated. The simulation started inside the object, 1cm to be precise, and the goal for the simulation was to remain in that position. When inside the object a force of 10N was exerted onto the subject to push it out of the object. The subject had to therefore, produce a net force in the muscles to counteract this force. Because the simulation was dealing with curvatures the direction of the force was dependent upon the location of the hand. Our results show that when the curvature increases the stiffness also increases. The direction of stiffness ellipse turned towards the direction of the instability (as was concluded in research by Selen et al. (2009)). However, this rotation was mirrored to the rotation from the results by Selen et al. (2009).

Table of Contents

Abstract	2
Introduction	4
Previous research	4
Materials and Methods	6
V-Shaped learning algorithm.....	6
Muscles.....	7
Task used in Selen et al. (2009)	9
Simulation.....	10
Results	13
Hand paths	13
Stiffness Ellipses	13
Muscle activation	15
Discussion	17
References	20

Introduction

Robots are becoming more imperative in daily life. The image of robots will no longer be just present in an industry setting. One important role for robots could be to assist in health care. This setting requires a very different set of skills than a welding robot in a car manufactory. In healthcare, the robot will be interacting with humans, and therefore safety is a big and important issue that needs to be looked at (Ikeura & Inooka, 1995). Having a robot that knows how to use the right amount of force at the right time is therefore incredibly relevant.

Humans are capable of performing a wide range of different complex tasks, a lot of which require precise stiffness control (Burdet et al., 2001, 2014). By doing research on muscle stiffness, we want to learn how humans control their muscles and see if we can eventually apply this ability to adapt to the external and internal environment onto robots. By using algorithms that are based on human behavior they might become more intuitive to interact with and therefore safer (Burdet et al., 2013).

Another aspect of this research is for instance that by precisely understanding how humans control their own muscles we might also be able to improve prosthetic devices (Hogan, 1984). The hope is to eventually be able to control prosthetics like a human would control their own hand.

Previous work

Stiffness control is an important aspect for humans. In 2001, Burdet et al. showed that humans are able to stabilize their hand in an unstable environment by increasing the stiffness of that arm (Burdet et al., 2001). Given that a lot of tasks are essentially unstable (for instance using a screwdriver to place a screw), this would mean that without stiffness control most tasks would not be able to be done. But what is the strategy the brain uses that results in increased stiffness?

Previous standard learning functions on how motor learning works, used linear schemes as shown in **figure 1** (Kawato et al., 1987; Emken et al, 2007; Franklin et al, 2008). These schemes work as follows: when a muscle in a previous trial is shorter than it should be, the activation of that muscle will be lower in the following trial. When a muscle in a previous trial is longer than it should be, the activation of that muscle will be higher in the following trial. The problem with these learning schemes is that there is no possibility of increasing the stiffness. This requires co-contraction of the opposing muscles, which is not covered in these learning schemes. Therefore, these schemes are only useful in environments that are predictable. Given that real world situations are anything but predictable, a new learning scheme had to be thought of.

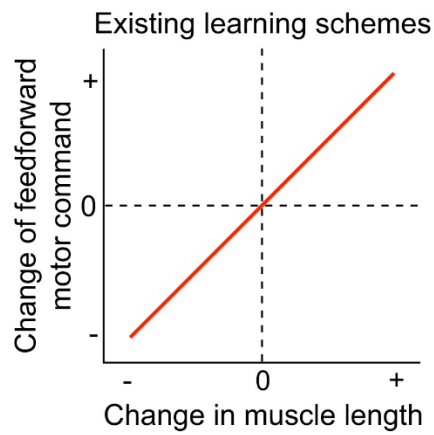


Figure 1 Existing learning schemes on motor control. Adapted from Franklin et al. (2008). Existing learning schemes work with linear models. When a muscle is longer than it needs to be, the activity of the muscle in the next trial will be higher. When a muscle is shorter than it needs to be, the activity of the muscle in the next trial will be lower.

Franklin et al. (2009) thought of a learning scheme that is able to work with unstable environments. They tested this algorithm in two specific tasks, and claim that this algorithm is generalizable. The aim of our study is to investigate if this model is indeed general, by applying the algorithm to a different task.

Materials and Methods

V-Shaped learning algorithm

In 2008, Franklin et al. wrote a paper describing a new non-linear model on motor learning. Instead of a linear function in the learning scheme, they used a V-Shaped function. The great advantage of this function is that it allows for an increase or decrease in co-activation of the muscles, and therefore regulate the joint impedance.

The model by Franklin et al. (2008) is based on three principles. First, if there is an unexpected muscle lengthening (positive error) in the previous trial, it means that in the next trial the feedforward muscle activity needs to be increased. Second, if there is an unexpected muscle shortening (negative error) in the previous trial, then in the next trial the feedforward muscle activity also needs to be increased. And third, if the error is below a certain threshold, then the feedforward muscle activation needs to be reduced. **Figure 2** provides an example of how this would work with the biceps and the triceps. As visible in the figure, there is a desired joint angle of the hand and the actual joint angle. In this example, there are two changes necessary. First, observing the actual joint angle it can be seen that the muscle length of the biceps is longer than it should have been (unexpected muscle lengthening). This means that we are dealing with principle one, and therefore it will result in an increase of muscle activity in the next trial. Second, the muscle length of the triceps is shorter than it should have been (unexpected muscle shortening). Here the second principle applies, because we are dealing with a negative error. Therefore, the muscle activity in the next trial should also be increased (but to a lesser extent compared to the biceps).

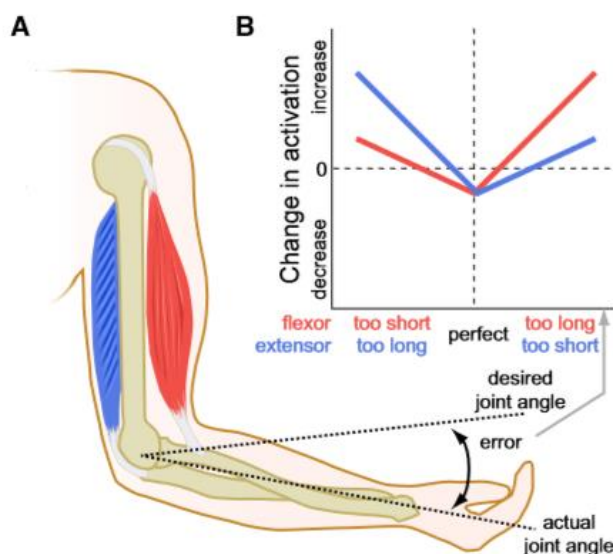


Figure 2 Graphical overview of the V-shape learning scheme. Adapted from Burdet et al. (2013). **A**, The Biceps and Triceps are used in this example to explain the V-shaped learning algorithm. There is a desired joint angle and an actual joint angle. Depending on the difference in muscle length between desired and actual position, the muscles will be more active in the next trial.

Franklin et al. (2008) show that by combining these three principles the CNS learns stable, accurate and efficient movements. They also state that the model is general, given that it works in both stable and unstable environments. However, they only tested on unstable environments, where the instability comes from external forces. Instability can also arise from internal forces such as, task geometry and signal-dependent noise (Selen et al. 2009). We will focus on this part to see if the

model is also equipped to work with this kind of instability. In order to do so, we first need to understand how and where this instability is coming from. Therefore, we need to look at the physiological and geometrical properties of the muscles.

Muscles

When you look at the muscles in the arm it is noticeable that there are a lot of different muscles, as is shown in **figure 3**. For modeling purposes these muscles are often lumped together into six functionally different muscles, of which four mono-articular and two bi-articular muscles. **Table 1** shows an overview of the six muscles used, with some characteristics of these muscles.

Name of muscles	Mono- or bi-articular	joint	Flexor or extensor
Pectoralis Major	Mono articular	Shoulder	Flexor
Posterior Deltoid	Mono articular	Shoulder	Extensor
Brachioradialis	Mono articular	Elbow	Flexor
Triceps Lateralis	Mono articular	Elbow	Extensor
Biceps Branchii	Bi articular	Shoulder and elbow	Flexor
Triceps Longus	Bi articular	Shoulder and elbow	Extensor

Table 1 Overview of the six muscles with name, mono- or biarticular muscles, which joint it works with and if it is a flexor or extensor

These muscles are all located in the upper arm. Mono-articular muscles are muscles that work on one joint (in this scenario this would either be the shoulder or the elbow). Bi-articular muscles work on two joints (in this scenario both the shoulder and the elbow). When increasing muscle activation, more force will be generated as well as an increase in stiffness. By co-contracting opposing muscles the net force will remain the same, however the stiffness in the arm will increase. For instance, by co-contracting the flexor and extensor of the elbow the stiffness will increase, which will make it harder to displace the hand in certain directions.




Figure 3 Graphical overview from the muscles in the arm. Adapted from Burdet et al. (2013). In this experiment the simulated muscles are encircled. **A**, anterior view of the muscles in the right arm. **B**, posterior view of the right arm

The simulation will use a planar 2-joint 6-muscle arm model as shown in **figure 4**. This figure shows a graphical overview of the two joints (shoulder and elbow) and the six arm muscles that play a role in the experiment.




Figure 4 Graphical representation of the planar two-joint six-muscle model, adapted from Franklin et al. (2008). In this figure you can see the two joints and the six muscles that were recorded and modulated in both the experiment by Franklin et al. (2008) as well as the experiment by Selen et al. (2009)

In the experiment by Franklin et al. (2008) the subjects had to make a forward reaching movement in a force field. This force field was either stable or unstable. The learning model proposed by Franklin et al. (2008) was able to replicate the behavior in both fields. To see whether the model is general, we implemented a task in which the subjects had to perform a different movement. This task is described in the paper by Selen et al. (2009).

Task used in Selen et al. (2009)

In the task described by Selen et al. (2009), the subject works with a robotic manipulandum. This manipulandum is programmed to simulate four different objects in different trials, all with different curvatures. The task of the subject is to generate forces onto such object, i.e. they have to push against it. The subjects move a simulated ball (by moving the vBOT, Howard et al., 2009) towards the given object and exerts a force onto the object. This force is depicted as a red arrow. The bigger the force, the longer the arrow becomes. The goal for the subject is to create enough force so that the tip of the arrow is located inside a target area, represented by a small square, for a specified amount of time (1.5-2.5 seconds). The more curved the objects are, the less stable the interaction is. **Figure 5** shows an overview of the experimental setup. Selen et al. (2009) show that the neuromuscular system controls the endpoint stiffness of the arm by trading off stability and noise. The higher the muscle activity, the more noise will arise. This is also referred to as signal-dependent noise (Hamilton et al., 2004).

For their experiment, they used four different objects. The first object is a cup. In this scenario the subject only had to focus on exerting enough force to get the tip of the arrow into the target area. The reason for this is that the hand cannot slip of the object.

The other three objects are all circles with different curvatures. The biggest circle (yellow object in figure 4) has a curvature of 0.67cm^{-1} , which is equal to a radius of 1.5cm. The next circle (orange object in figure 4) has a curvature of 2.00cm^{-1} , which is equal to a radius of 0.5cm. The last circle (red object in figure 4) has a curvature of 4.00cm^{-1} , which is equal to a radius of 0.25cm.




Figure 5 Graphical overview of the task. Adapted from Selen et al. (2009). The subjects hand was represented by the blue ball. By moving the V-bot the subject was able to move the ball as well. By pressing the ball against the simulated object, a red arrow would emerge from the object. By exerting more force onto the object the arrow became longer. The goal for the subject was to exert enough force to place the tip of the arrow inside the target area, (represented by the small square). There were four possible objects that could be simulated, seen on the right. In this example the second object from the top was used (the orange circle with a curvature of 2.00cm^{-1}).

Figure 6 shows the results from the experiment done in the paper by Selen et al. (2009). Here they show the stiffness ellipses, a measure of the resistance of the hand to be moved in the different directions, of the given objects for every subject in both experimental sessions. There are a few things to notice from these results. First, when the curvature of the object increases the size of the stiffness ellipse also increases. Second, with increasing curvature the direction of the stiffness ellipse turns towards the direction of the instability (in this case that would be the vertical axis). The key objective of the current study is to see whether we can replicate these modulations of the stiffness ellipses as a function of object curvature using the V-shaped learning algorithm.




Figure 6 Stiffness ellipses from the experiment by Selen et al. (2009). The ellipses are color coded. In the experiment by Selen et al. they also researched whether the order in which the objects were presented had an impact on what the endpoint stiffness of the subjects were, therefore they did two sessions in which one the order was randomized.

Simulation

As described above, the task by Selen et al. (2009) needed to be implemented in the model by Franklin et al. (2008). As a starting point we used the original Matlab code for the V-Shaped learning algorithm, provided by Franklin et al. (2008).

The first step in the implementation process was to try and freeze the hand on a certain position. In the original simulation the hand made a forward reaching movement, whereas in the task by Selen et al. (2009) the hand needed to generate a force onto an object that was located to the left of the hand. After this was accomplished we started working on simulating objects. The first object that was simulated was a wall. This object was used as a stable condition to help us acquire a better understanding of the model.

We implemented the force production task by defining positions. We decided to place the hand 1cm inside of the object (to the left of the wall). The object would function as a spring. The further the subject goes into the object the greater the force. The starting point was also the goal position. Being inside the object meant that there would be a force field that would exert force on the hand to push it out of the object. The further the hand was inside the object the higher the forces were to try and push the hand out. At the starting position (1cm inside the object) the amount of force was 10 Newton. This was the maximum amount of force that would be exerted onto the hand during the trials. The direction of the force was orthogonal to the edge of the object. Given that we were dealing with the wall the direction would always be parallel to the horizontal axis and orthogonal to the wall.

After this was accomplished, the simulated wall was replaced by a simulated circle. The idea behind the simulation is the same. The goal of the simulation is to remain at the start position. This required an activation of the muscles because an external force is being exerted onto the hand, which is pushing the hand out of the object. The amount of force is still the same as with the wall situation, 10N when the hand is 1cm inside and the closer the hand gets to the wall the lower the amount of

force becomes until the edge of the circle is reached. There the force is 0N. Depending on the position of the hand, the force is being exerted in different directions. Given that simulation is working with circular objects, the external force will always be orthogonal to the edge of the circle (see figure 7).

To make sure that the results from the simulation were comparable to the results discovered in the paper by Selen et al. (2008) we tried to keep the parameters as close as possible to the parameters used in the paper. This meant that the position of the hand was located directly in front of the body midline, at a distance of 35 cm. We assumed that an average shoulder width is around 40cm. **Figure 7** shows a graphical setup of how the simulation will work.

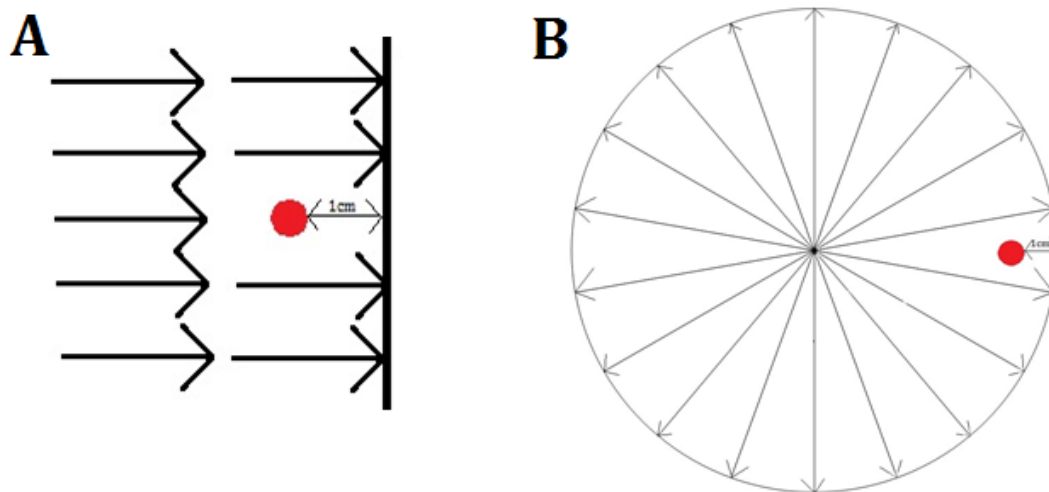


Figure 7 Graphical overview of simulation setup. A, the hand is placed 1 cm inside of the wall. At this point, there is a force of 10N being exerted onto the hand. The direction of the force will always be orthogonal to the wall. This force gradually becomes less the closer the hand gets to the wall. B, the hand is also placed 1cm inside of the object. The object is circular and therefore the direction of the force will differ depending on the position of the hand.

In pilot simulations, it became clear that for the smallest objects that were used in Selen et al. (2009) our simulation was not capable of stabilizing these environments. Therefore, we chose to increase the size of these objects, and increased the original radii with a factor 10. We used the objects with the following radii, 15cm, 5cm and 2.5cm.




Figure 8 Experimental setup, positions of the hand similar to the experiment described in the paper by Selen et al. (2009) and used in the simulation we programmed.

The hand gets pushed out of the object. This results in an error. During the current and following trials the simulation will increase force to the left and stiffness in the arm. When the hand gets to the desired position the stiffness will decrease, but the net force will remain. Each object will be simulated with 200 trials. Each trial exists out of 650 time samples.

Results

Hand paths

Figure 9 shows results from the simulations for the different object curvatures. Each panel shows the hand positions of the first 10 and the last 10 trials. The first 10 trials are of importance, because the simulation has not learned anything yet. The hand paths are displayed with a red color. The last 10 trials should be optimized after the learning period. These are indicated by the blue lines. The top row is the overview of the entire object with the hand paths, and the bottom row is a close up on the hand paths itself. In the close ups of the hand paths is visible that the red paths (before learning) are mostly to the right of the blue paths (after learning), which mean that after learning the simulation more able get closer to the desired position.

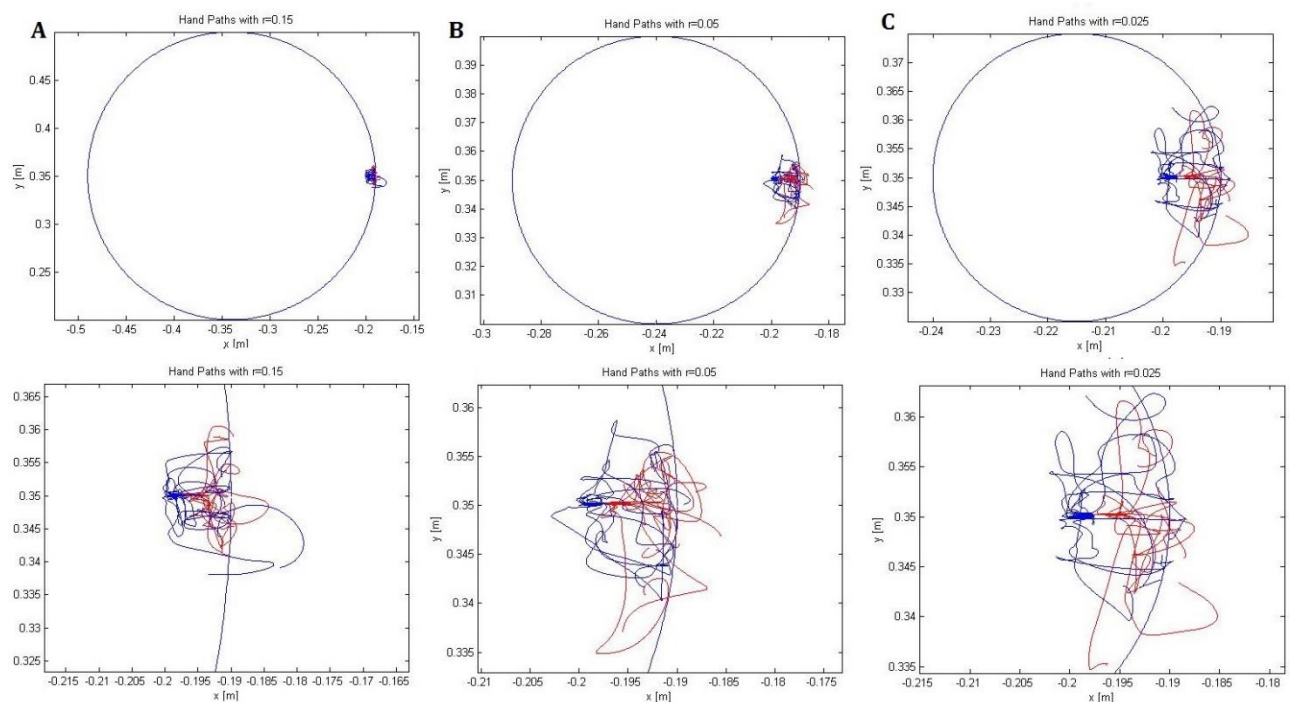


Figure 9 Plots of the hand paths that the simulation makes. The red lines are the first ten trials and the blue lines are the last ten trials. **A**, these are the hand paths from the biggest circle, namely with a radius of 15cm. In the plot is visible that the hand paths are close and confined to close space. **B**, these are the hand paths for the 5 cm radius circle. **C**, these are the hand paths for the smallest circle with a radius of 2.5cm

Stiffness Ellipses

We also looked at the stiffness ellipses. In the paper by Selen et al. (2009) they discovered that when the curvature of an object increases (and therefore gets more instable), the size of the stiffness ellipses increases. They also found that the direction of the stiffness ellipse turns towards the direction of the instability. This size change and rotation is something we looked for as well.

Figure 10a presents the stiffness ellipses for the three different objects used. For the object with the smallest curvature (green, and thus the biggest radius), the stiffness ellipse is the smallest and the direction of the ellipse is horizontal. This is because the radius is so big that it is not hard to remain on the surface. Therefore, the stiffness of the arm is mainly focused on pushing against the object to go to the goal position. The second ellipse (yellow) is for radius of 5cm. Here you see that the ellipse is slightly bigger, which means that the arm increased stiffness to perform the task. Another thing to

notice is that the direction of the ellipse changed. The object has become smaller (from 15 to 5cm) and therefore it gets harder to not slip of the object. To compensate from the decrease in surface the arm needs to stiffen its muscles more in the direction of the instability, which is the vertical axis. The last ellipse is for the smallest radius (red). This ellipse is a lot bigger than the other two, which means that the stiffness in the arm is also higher than the stiffness of the other two. Also here you can see that the direction of the ellipse is turned more toward the vertical axis, and therefore the instability. **Figure 10b** shows the object used with the radii and curvatures attached to each object. They are color coded to match the ellipses. **Figure 10c** shows a figure in which the size of each stiffness ellipse is plotted. It is clear that there is an increase in size when the curvature increases (and thus gets more unstable). **Figure 10d** shows a figure in which the orientation of the ellipses are plotted. The total rotation of the ellipses is around 45 degrees, which is almost double the rotation found by Selen et al. (2009) (maximum of 23 degrees).

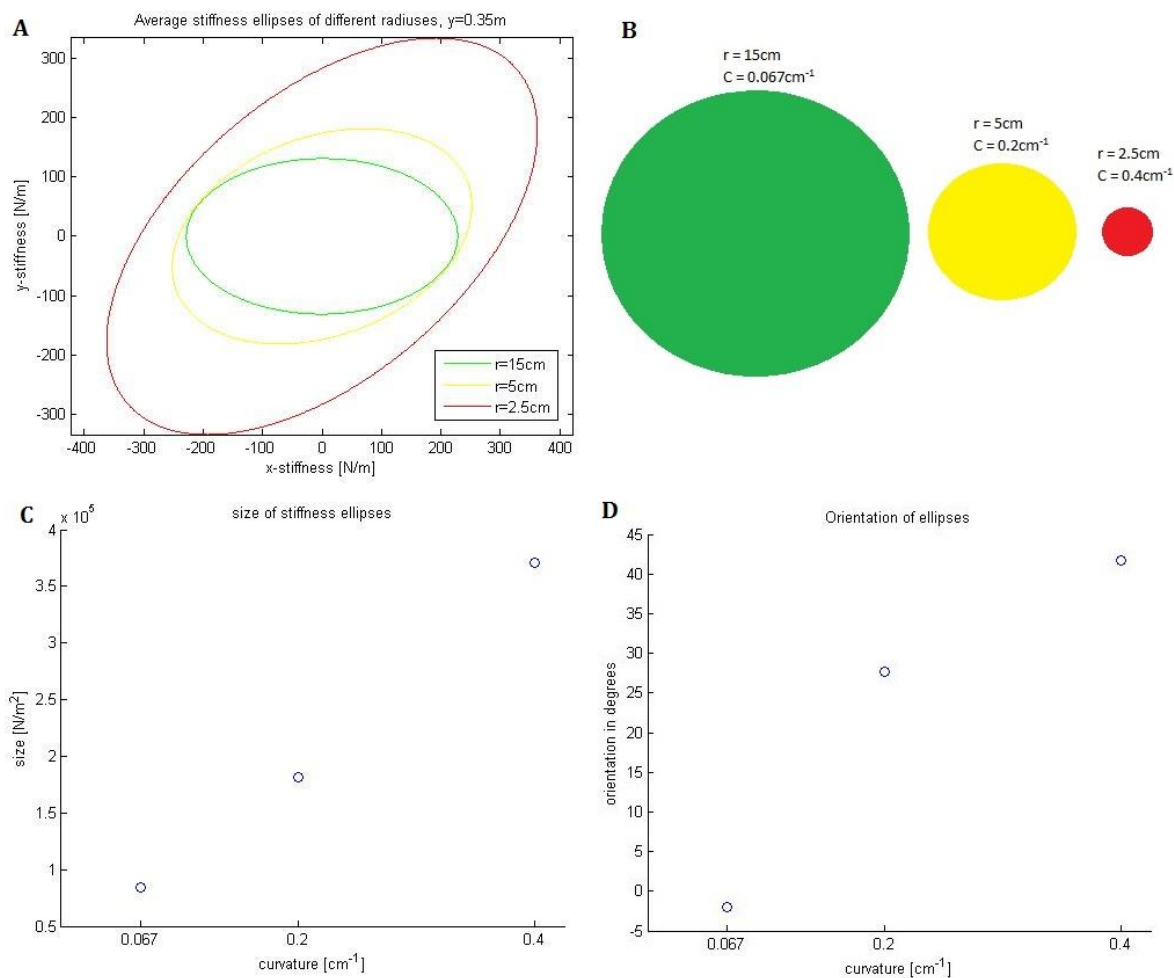


Figure 10A, Plot of average stiffness ellipses for the three different curvatures. *10B*, the different objects with radius and curvature of that object. *10C*, plot of the increasing size of the stiffness ellipse when the curvature increases (and thus gets more unstable). *10D*, plot that shows the orientation of each stiffness ellipse. Here it is visible that the ellipse turn counterclockwise, with a total rotation of around 45 degrees.

Something to notice is that the rotation is counterclockwise, whereas the rotation of the ellipses in Selen et al. (2009) paper was clockwise.

To investigate this further, we looked at the positioning of the hand and see if this has any impact on the stiffness ellipses and the direction of the rotation.

The position of the hand in the main study was at -20cm in the horizontal axis and 35cm in the

vertical axis. To investigate we differed the position of the hand with 5cm below 35 and 5cm above 35. Thus, one hand position of (-20, 30) and one hand position of (-20,40).

Figure 11 shows the stiffness ellipses for two different hand positions. **Figure 11a** is a plot in which the hand was placed 30cm from the body. This is five cm closer to the body in comparison to the original placement of the hand. In comparison with the stiffness ellipses of the hand position (-20,35), these ellipses are smaller and longer. Looking at the difference in rotation there is a small difference but the rotation is still counterclockwise. **Figure 11b** shows a plot in which the hand was placed 40cm from the body. This is five cm further from the body in comparison to the original placement of the hand. These ellipses are a lot wider and in fact turns clockwise.

A reason for this could be that the task is performed too far away from the body. It could be possible that because of the distance the entire arm has to stiffen up to perform the task successfully. Tasks are performed with more ease and less stress on the body when the task is performed more close to the body (Occupational Health Department, 2007).

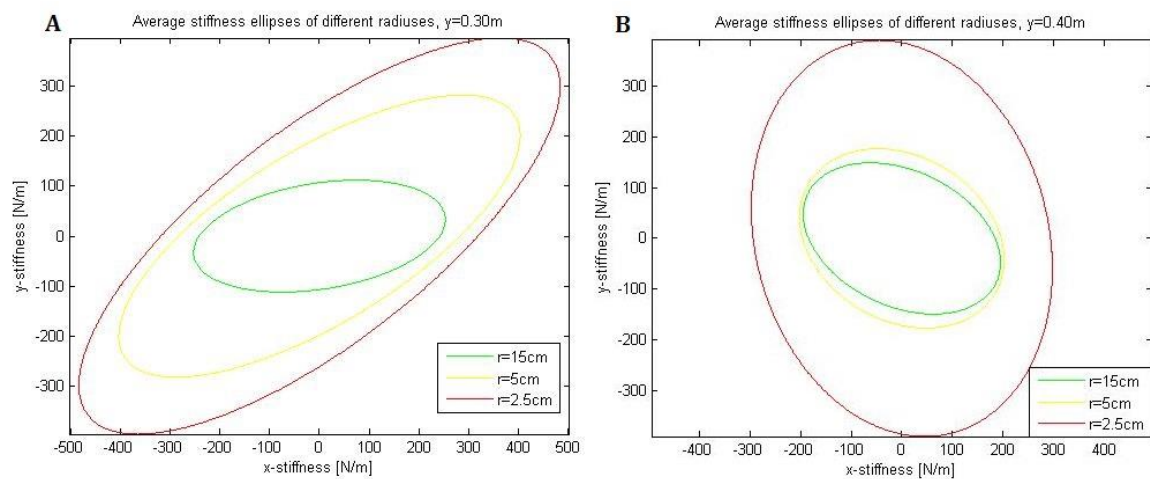


Figure 11A, Plot of average stiffness ellipses for the three different curvatures. The hand is placed at a distance of 30cm from the body, instead of the standard used which was 35cm **B**, Plot of average stiffness ellipses for the three different curvatures. The hand is placed at a distance of 40cm from the body, instead of the standard used which was 35cm.

Muscle activation

The stiffness ellipse is ultimately a reflection of the activation of the six muscles in our model. Therefore, we also looked at the individual muscle activations of each of the six muscles. As mentioned earlier, we looked at the shoulder flexor and extensor, the biarticular flexor and extensor, and the elbow flexor and extensor.

Figure 12 shows the simulated muscle activities of these six different muscles for the three different objects.

When looking at the figure there are some noticeable things to see. First, when we look at the shoulder muscles, you see that the muscle activity stays around the same height. That makes sense, because the shoulder flexor is the muscle that ensures the 10 Newton force that is required to get to the target place. Second, when we look at the biarticular muscles we see an increase when the curvature increases. Both the flexor and the extensor increase, which results in increase in stiffness to compensate for the instable environment. And lastly, when looking at the elbow flexor and extensor, there is an increase in both muscles when the curvatures increases. This also makes sense,

because when the curvature increases the environment becomes more unstable. The instability is noticeable in the vertical axis, because that is where the hand could potentially slip off. To minimize this potential movement the muscle activity of the elbow muscles have to increase, and therefore, reducing the possibility of a change of hand position in the vertical direction.

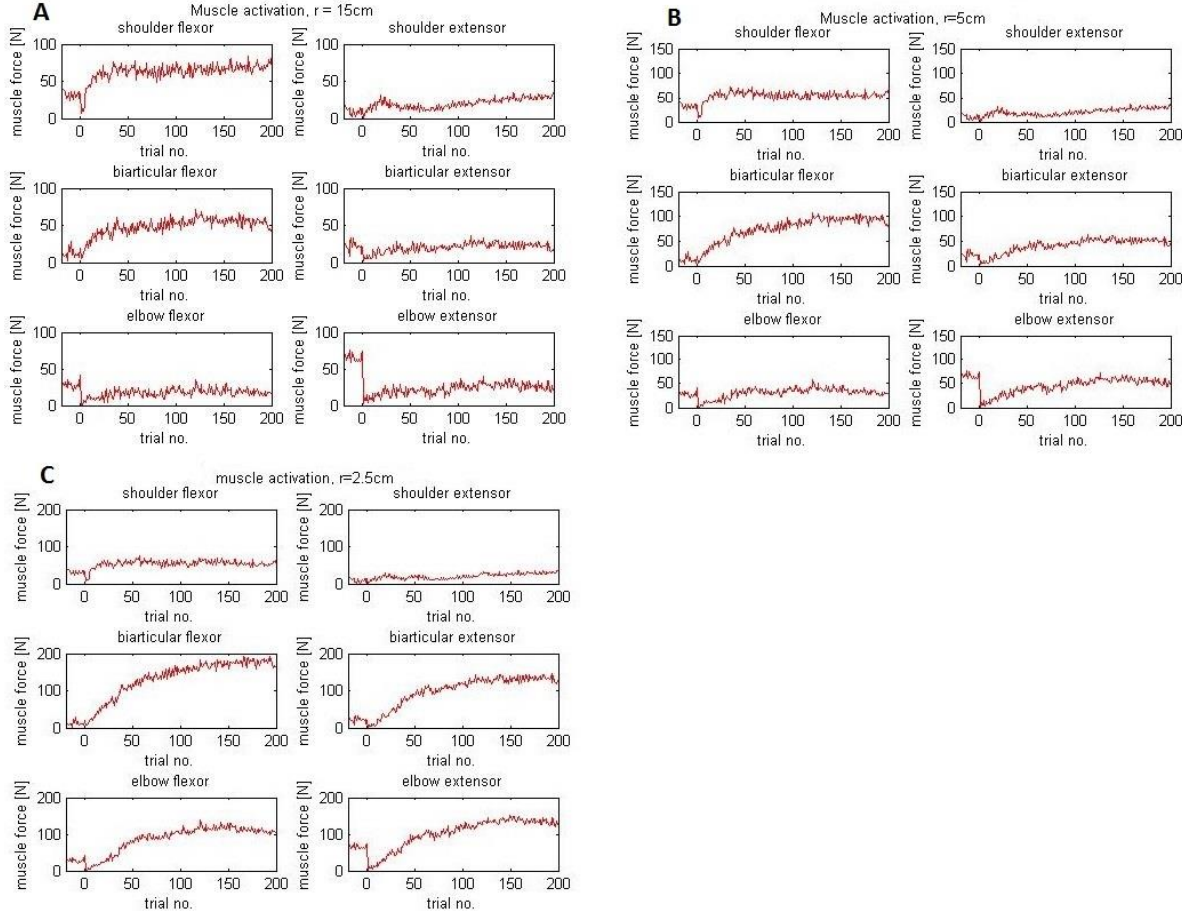


Figure 12 overview of muscle activity. **A**, muscle activity of each of the six muscles when pushing against circle with 15cm radius. **B**, muscle activity of each of the six muscles when pushing against circle with 5cm radius. **C**, muscle activity of each of the six muscles when pushing against circle with 2.5cm radius

Discussion

In this study, we investigated the computational model of motor learning that was proposed by Franklin et al. (2008). They proposed that this model is a general model. We tested to see if this model was able to predict stability control on curved objects.

The simulations showed that when the curvatures increases, the stiffness ellipses rotates towards the direction of the instability. This matches with the results from Selen et al. (2009). When you look more closely at these results there is in fact a difference. In the results by Selen et al. (2009), the rotation was clockwise, whereas in our results the rotation was counterclockwise.

As was shown earlier in **figure 11** it became clear that the model is very sensitive to the position of the hand. Given that we kept the positioning of the hand as close as possible to the positioning used by Selen et al. (2009) in their experiment, there may be a different reason for the mirrored rotation of the direction of the stiffness ellipses.

A reason could be the way we implemented the experiment. Initially it was a force driven task. The goal for the subjects was to push with a certain amount of force such that the tip of the arrow would be located inside a target area. Whereas with our experiment the task became a position driven task. In the simulation the hand was placed inside the object and the goal was to stay in that position. The difference between these two tasks is that in the position driven task we are dealing with a tolerance zone. As shown in **figure 13** as long as the hand is somewhere in the tolerance zone the simulation thinks it is alright. The width is also represented in the target area. As long as the tip of the arrow is somewhere in the target area it will suffice. When we look at height in **figure 13A** there are no extra restraints on where in the circle the hand needs to be. Thus, it might as well be in the very top of the tolerance circle. If we look at **figure 13B**, the height matters a lot. If the subjects differ just a bit from the middle line it will have great consequences on where the tip of the arrow might end. It is therefore extra important to stabilize the arm as much as possible to remain precisely at the right position to the circle.

To truly know if and how this effects the results it needs to be investigated by completely rewriting the code by Franklin et al. (2008) to work with forces instead of locations and see if this changes the results from the simulations.

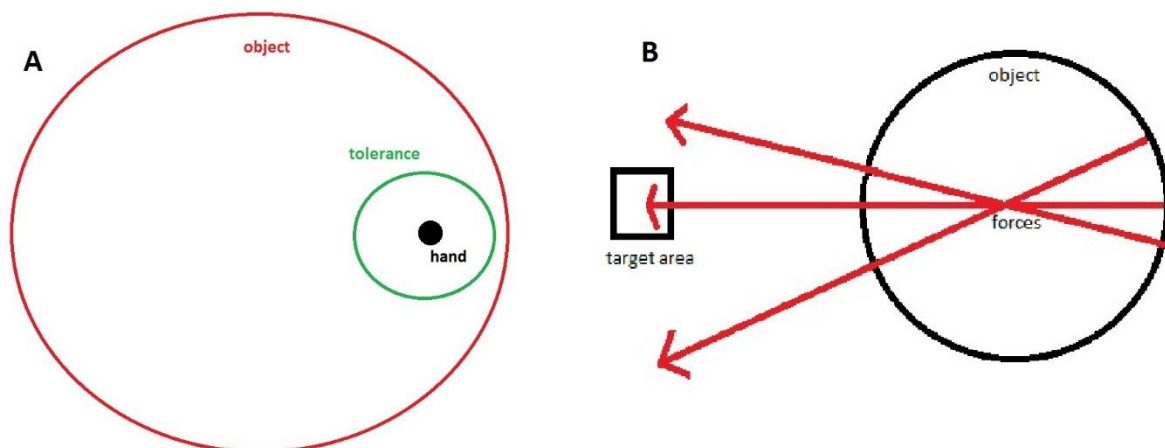



Figure 13A Overview of the position driven task that was done in the simulation. The red circle represents the object, the green circle represents the tolerance area where the hand is allowed to be in and the black dot is the position of hand itself. 13B, overview of the force driven task described by Selen et al. (2009). Here the black circle is the object, the black square is the target area where the tip of the arrow needs to be, and the red arrows that represent the force from different places around the circle.

Another reason could be that we used different radii than in the original experiment. This was the case, because of the way we implemented the task. We implemented the task in a way such that the hand is placed 1cm inside of the object and forces are gradually pushing the hand out. The problem is that when the radii get smaller than 1 cm the hand would either be placed in the left half of the object or even outside of the object. This happened with the two of the object radii used by Selen et al. (2009), namely the 0.5cm and the 0.25cm radius. **Figure 14** shows the results from these runs. 14A is still relatively speaking possible, but when looking at 14B and 14C it is clear that it is impossible to work with these parameters.



*Figure 14 Overview of what happens to the hand positions in the small radii (that were used by Selen et al. (2009)). **A**, biggest object with a radius of 1.5cm. **B**, object with a radius of 0.5cm. Here is visible that the hand would be placed to the left of the object. **C**, object with a radius of .25cm. Here the starting position is not even in the neighborhood of the object.*

As an alternative, we tried to place the hand 1mm inside of the circle (instead of 1cm). The problem here is that the gradual decrease of the 10N outwards was not gradual enough. Therefore, the hand was charged outwards with such an amount of force that the simulation was not able to recover. Thus, we chose to use bigger objects. By multiplying the original radii with a factor 10, the ratio between the three objects should remain the same. In retrospect, what we might have done is making the desired position time dependent so that every trial starts with 0N external force and the desired position just moves slowly towards the inside of the object.

And finally, another possibility is that the model deems all the muscles equal. There is one general V-Shape that applies to all of the muscles. The same goes for the coefficient of variation. However, this is not the case in the real world. Research has shown that every muscle has its own coefficient of variation (Hamilton et al., 2004). The results from this research are shown in **figure 15**. Here is visible that each muscle has their own slope. The smaller the muscle the higher the coefficient of variation is. To see whether this has any significant effect, both the coefficient of variation and the V-shape need to be determined for all of the six muscles and implemented in the model.

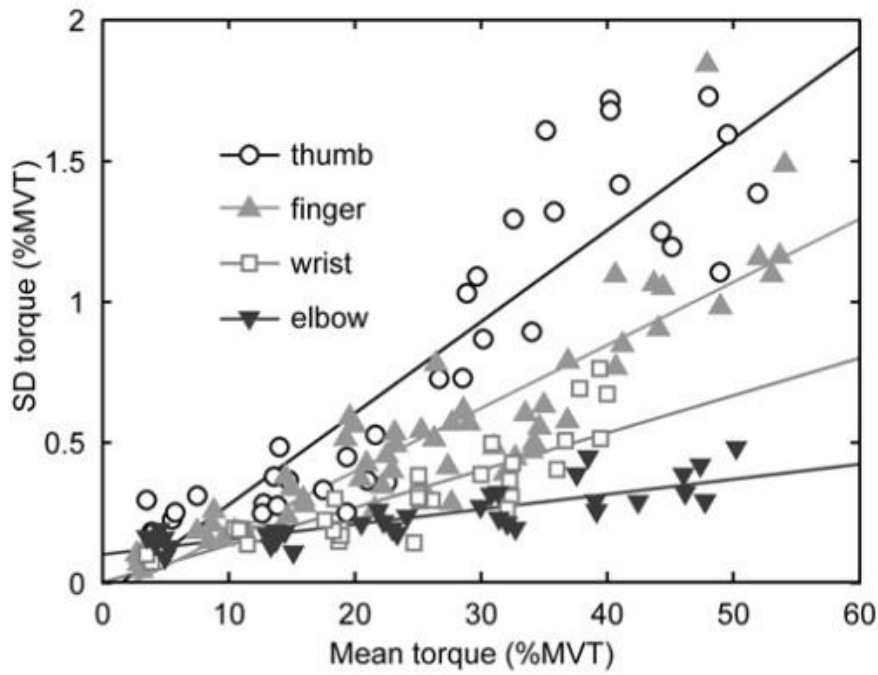


Figure 15 Plot that shows that every muscle has their own slope of how fast the muscle noise increases when muscle activity increases. Adapted from Hamilton et al. (2004)

In conclusion, we have shown that the learning model proposed by Franklin et al. (2008) generalizes to other tasks. Further research is needed to see how the details of the task can be implemented, but these initial results are promising

References

- Burdet, E., Osu, R., Franklin, D. W., Milner, T. E., & Kawato, M. (2001). The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature*, 414(6862), 446–449. <https://doi.org/10.1038/35106566>
- Burdet, E., Franklin, D.W., & Milner, T.E. (2013). *Human Robotics: Neuromechanics and motor control*. Neuroscience/Robotics. MIT Press, 2013.
- Emken, J. L., Benitez, R., Sideris, A., Bobrow, J. E., & Reinkensmeyer, D. J. (2007). Motor Adaptation as a Greedy Optimization of Error and Effort. *Journal of Neurophysiology*, 97(6), 3997–4006. <https://doi.org/10.1152/jn.01095.2006>
- Franklin, D. W., Burdet, E., Tee, K. P., Osu, R., Chew, C. M., Milner, T. E., & Kawato, M. (2008). CNS learns stable, accurate, and efficient movements using a simple algorithm. *The journal of neuroscience*, 28(44), 11165-11173.
- Hamilton A.F., Jones, K.E., & Wolpert, D.M. (2004). The scaling of motor noise with muscle strength and motor unit number in humans. *Exp Brain Res*, 157:417–430
- Hogan, N. (1984). Adaptive control of mechanical impedance by coactivation of antagonist muscles *IEEE Transactions on Automatic Control* 29(8):681, 0018-9286
- Hogan, N., (1984). Impedance Control: An Approach to Manipulation, American Control Conference, San Diego, CA, USA, 1984, pp. 304-313.
- Howard, I.S., Ingram, J.N., & Wolpert D.M. (2009). A modular planar robotic manipulandum with end-point torque control. *J Neurosci Methods*, 181:199– 211.
- Ikeura, R., & Inooka, H. (1995). Variable impedance control of a robot for cooperation with a human. *Proceedings - IEEE International Conference on Robotics and Automation*, 3, 3097–3102. <https://doi.org/10.1109/ROBOT.1995.525725>
- Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical neural-network model for control and learning of voluntary movement. *Biological Cybernetics*, 57(3), 169–185. <https://doi.org/10.1007/BF00364149>
- Occupational Health Department. *Moving and Handling Techniques*. Occupational Health Department, Imperial College, page 13, 2007
- Selen, L. P., Franklin, D. W., & Wolpert, D. M. (2009). Impedance control reduces instability that arises from motor noise. *The Journal of Neuroscience*, 29(40), 12606-12616.