UCL | Université Catholique de Louvain

Degree in Biomedical Engineering
2018-2019

Bachelor Thesis

"Characterization of feedback control strategies for human motor control"

Andrés Sánchez Pérez

Tutor & cotutor

Prof. Frédéric Crevecoeur
Prof. Jesús Andrés Picó Marco

Louvain la Neuve, 2019

Acknowledgements

I would like to express my most sincere gratefulness to Prof. Frédéric Crevecoeur for his commitment, aid and support throughout the realization of this thesis, without whom it would not have been possible. I would also like to thank Prof. Jesús Andrés Picó Marco for his help in the searching of an appropriate project for me in the field of control theory and neuroscience.

Content

1. Introduction

1.1. Objective

In the wide field covered by control theory, human motor control is still an area with the potential to be deeply explored, regarding the little knowledge we have about it. The huge research in medical technology that occurred during the last decades, considering the development of devices in bio-robotics such as exoskeletons, robotic arms and many others, made it necessary to keep gaining more insight into the functioning of the human body – not just from the anatomical and physiological standpoint, but also from the motor control point of view present in neuroscience. Henceforth, the focus in this thesis lies between decision-making in humans and feedback control strategies in human motor control.

It has been proposed that when humans face multiple options, for example, there are multiple apples on a table and they want to reach for one, there is a competition in the brain between different control laws, and the movements generated before a final commitment reflect a motor 'averaging' where the averaged movements are transient ones until one of the strategies wins upon the other, once a final decision is made.

The alternative hypothesis considers that instead of averaging distinct motor plans or controllers, the brain picks an intermediate option and produces a controller suitable for this option. In this case, a time-varying intermediate trade-off goal is set, and the control pursues it as to eventually drift towards one of this two options.

Therefore, in this project the main objective is to investigate these two alternative hypotheses in order to explore and challenge them from a computational standpoint. It is expected that the results of this project allow us to gain more profound knowledge in this hot topic in neuroscience and in the design of proper control strategies for such medical devices.

1.2. Control framework

For the implementation and the testing of both strategies, the coding environment of MATLAB was used using the Linear Quadratic Gaussian (LQG) control framework in agreement with current models of reaching control in humans. For doing so, testing was carried out on a model of a human arm, by adjusting biomechanical parameters and exploring its features. To this end, simulations involving decision-making and mechanical perturbations were assessed.

For these simulations, without going too deep to the generalised mathematical model which was extracted from [1] and [2], the control system had the following form:

$$z_{k+1} = Az_k + Bu_k + \xi_k$$

Where the variable z is the state vector at every time-step k contains the displacement, the velocity and the force in the x and y axis, thus the model in this project considers a movement with 2 degrees of freedom in a 2D plane. Therefore, the control vector u controls the state variables from each axis separately.

Moreover, in order to set some constraints in the simulations, a cost function was used with the general form:

$$J(x,u) = \sum_{k=1}^{N} (z_k^T Q_k z_k + u_k^T R u_k)$$

Where R is a constant and Q is the matrix used in order the be able to penalise each of the state variables separately, depending on the objective of each simulation.

Finally, as the Gaussian term in LQG defines by its name, the term ξ_k follows a normal distribution and is the noise introduced into the system to include perturbations. Analogue to the original mathematical model in [2], a series of Kalman gains were used to correct the state estimation at each time-step.

For instance, to obtain the experimental results, the simulations were done by adjusting the weights of each of each state variables and the number of time-steps, as well as including target displacements and multiple targets in order to answer the main question of this study.

2. Methodology

2.1. Validation of the model: target jump

The first stage consisted on validating the model to prove its correct performance moving towards a desired fixed target and a moving target. For all these simulations, some parameters were previously fixed, such as the mass displaced with a value of 1 kg, in agreement with [4], and the muscle time constant with a value of 0.066 s. The reason for this last value is that the contraction and relaxation of the muscles has a value between 0.06 and 0.09 s respectively, as it was shown in physiological studies mentioned in [4]. Moreover, initially each simulation had a running time of 1 s and a time-step of 0.001 s, but it is later changed in one case to prove some results discussed later. Lastly, the number of simulations done in each case was fixed to 5, due to the good reliability of the results.

For instance, in the first scenario the aimed target was: $(x, y) = (0.5, 0.4)\ m$ with a desired final velocity of $(x', y') = (0, 0)\ m/s$ without any target jump. Since the focus relies on the trajectory followed, the weights used on the forces were always zero, on the positions were always non-zero, and for the velocities it was changed between zero and non-zero values from one case to another. For each case, the control gains and the control vectors for each axis were obtained, as well as the variation of trajectories and velocities with time (see Fig. 1 and Fig.2). Figure 1 shows the results obtained considering the weights for the velocities non-zero and the same to the weights in the position, while figure 2 shows the results considering a weight in velocity with zero value. In both cases, the target was reached with an error <1%, but when the weight in velocity was set to zero, the final velocity was not the one desired. Also, the control vector shows a correction in figure 1 towards the end, as to make the mass slow down and reach zero final velocity. Considering a single target, the model had an acceptable performance thus further situations could be studied.

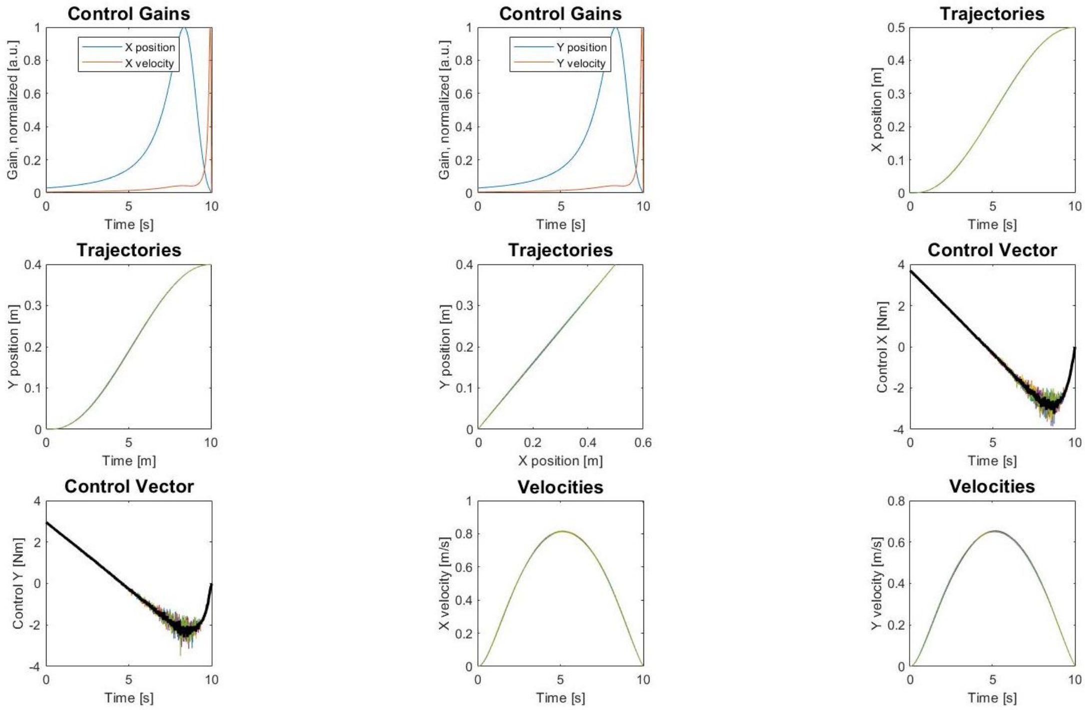

Fig. 1: *Performance with target at (0.5, 0.4) and non-zero weights in position and velocity:* (top left & top) Control gains of x and y positions and velocities. (top right & middle left) Variation of X and Y trajectories with time. (middle) Y position vs X position. (middle right & bottom left) Control vector of each axis. (bottom and bottom right) Variation of X and Y velocities with time.

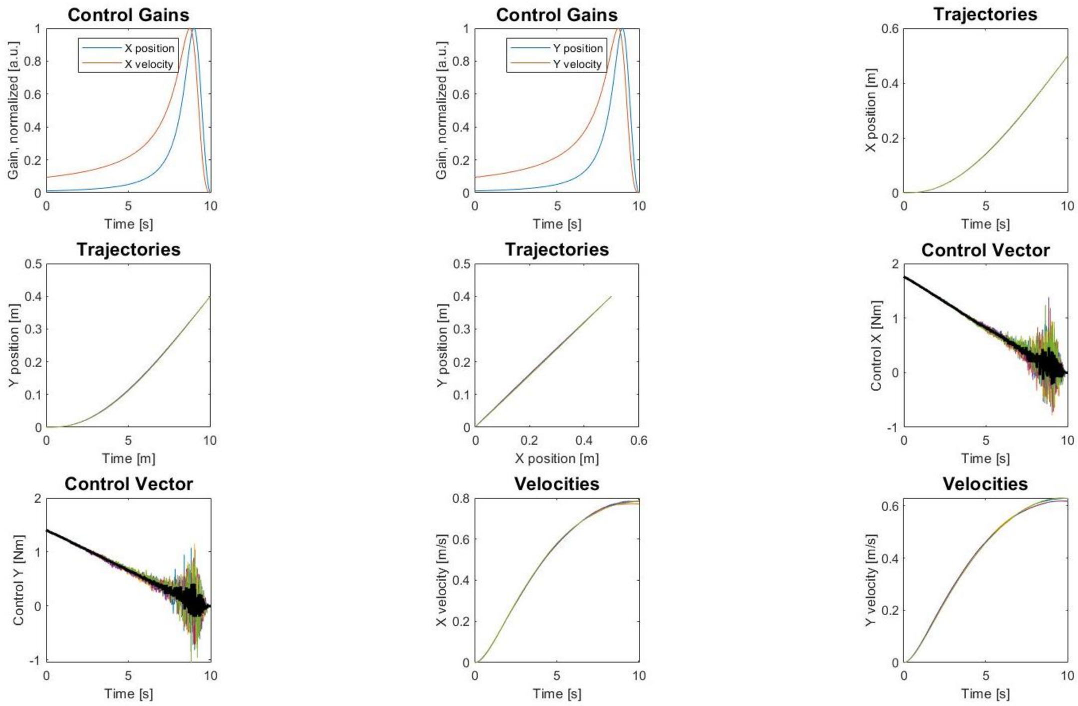

Fig. 2: *Performance with target at (0.5, 0.4). Non-zero weights in position and zero weights in velocity*

For the case of the moving target, its position is changed halfway during the simulation. With an analogue methodology and presentation of results as in the first case, figure 3 shows a target jump at 50% of the simulation from $(x, y) = (0.5, 0.4)\ m$ to $(x, y) = (-0.5, 0.4)\ m$ and non-zero weights in velocities. The desired final velocities were always kept to 0. The graphs show that in this case, the final position had an error <4% and a final velocity of 0 m/s.

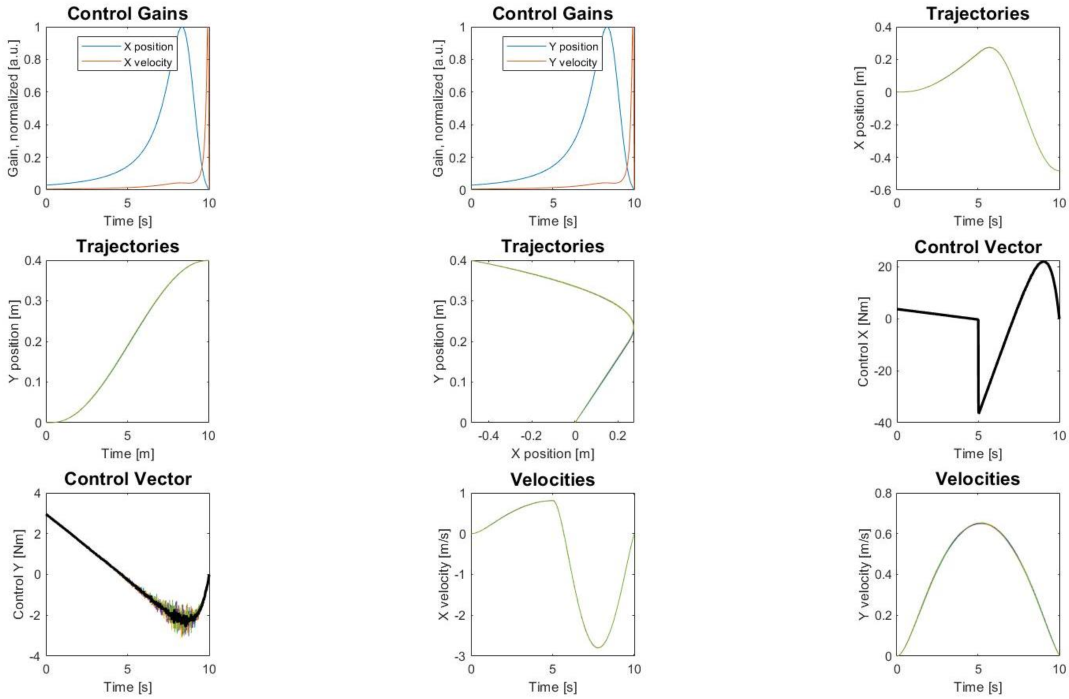

FIG. 3: *PERFORMANCE OF A TARGET JUMP AT 50% OF THE SIMULATION TIME, FROM DESIRED POSITION TARGET (0.5, 0.4) TO A NEW TARGET (-0.5, 0.4), WITH NON-ZERO WEIGHTS IN POSITION AND VELOCITY.*

On the other hand, figure 4 shows the same results as figure 3, except for the weight on velocity, which was zero. As expected, the final velocities were non-zero due to not penalising the velocity target, with the difference that the error in the final position was <1.5%. In the same way it happened with the fixed target, the control vector in the X axis, did a correction halfway through the simulation due to the target jump and another towards the end to slow down the mass and achieve a final velocity with zero value, for the case of figure 3. However, to compare the effect of the weights in relation to the moment there is a target jump (and consequently, a decision is taken), in two further experiments the weights in velocity are also changed from a non-zero to a zero value but considering a later target jump, at 75% of the simulation.

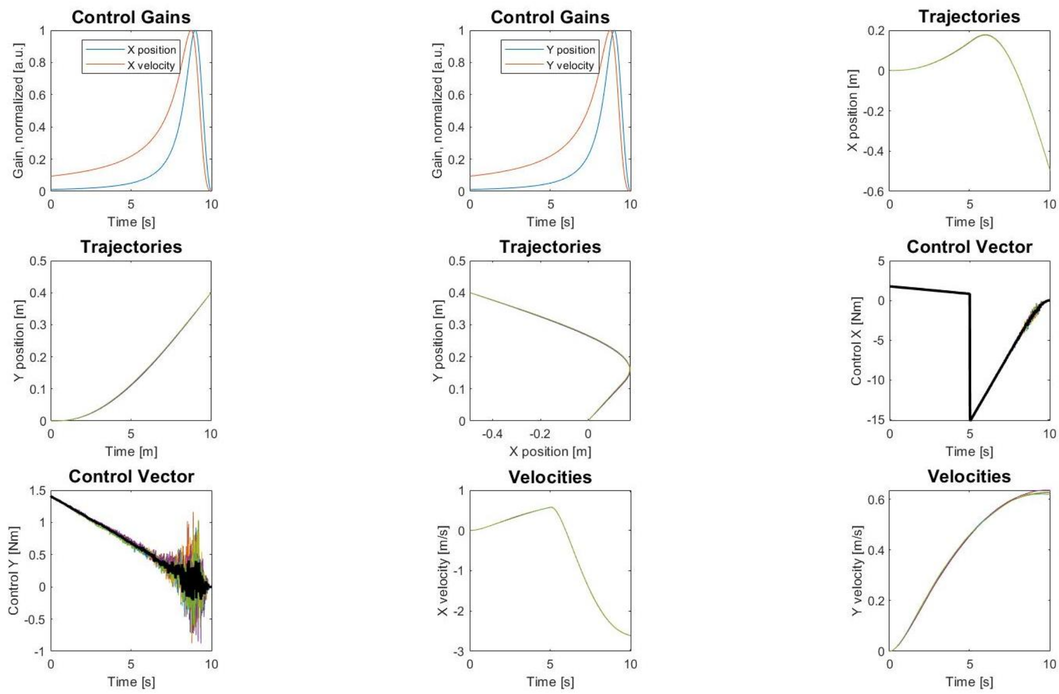

FIG. 4: *PERFORMANCE OF A TARGET JUMP AT 50% OF THE SIMULATION TIME, FROM DESIRED POSITION TARGET (0.5, 0.4) TO A NEW TARGET (-0.5, 0.4), WITH NON-ZERO WEIGHTS IN POSITION AND ZERO WEIGHTS IN VELOCITY.*

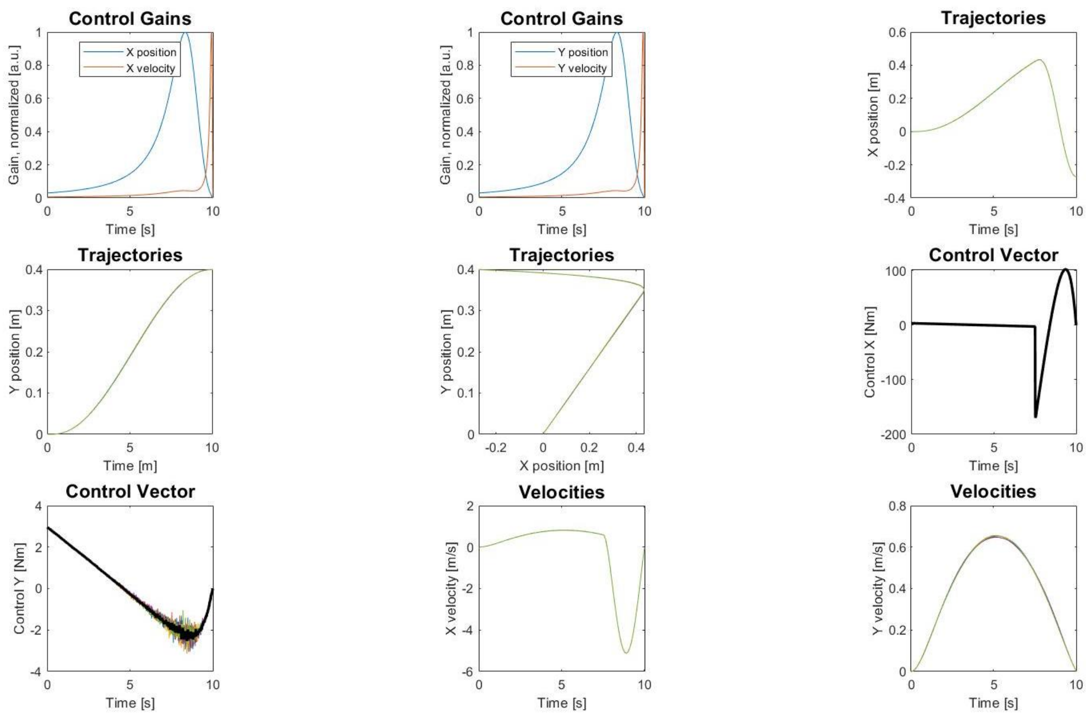

FIG. 5: *PERFORMANCE OF A TARGET JUMP AT 75% OF THE SIMULATION TIME, FROM DESIRED POSITION TARGET (0.5, 0.4) TO A NEW TARGET (-0.5, 0.4), WITH NON-ZERO WEIGHTS IN POSITION AND VELOCITY.*

Figure 5 shows the performance considering the weight in velocity being non-zero. Again, a correction towards the end by the control vector to achieve final velocity zero and a correction at 75% due to the target jump is done. Even though, the final velocity is the expected one, in this case the final X position was far from the desired one, with a huge error of <45%. Notice that the desired Y target was kept constant throughout all the experiments as a control group and used as a comparison to X, to avoid biased results due to a possible lack of robustness in the code. As a last comparison in target jumps, figure 6 shows the same jump at 75% of the simulation, but in this case, using a weight of zero in velocity. As it was expected, final velocity is non-zero, but interestingly, the final position had a much smaller error than in figure 5, with a value of only <10% compared to the previous <45%.

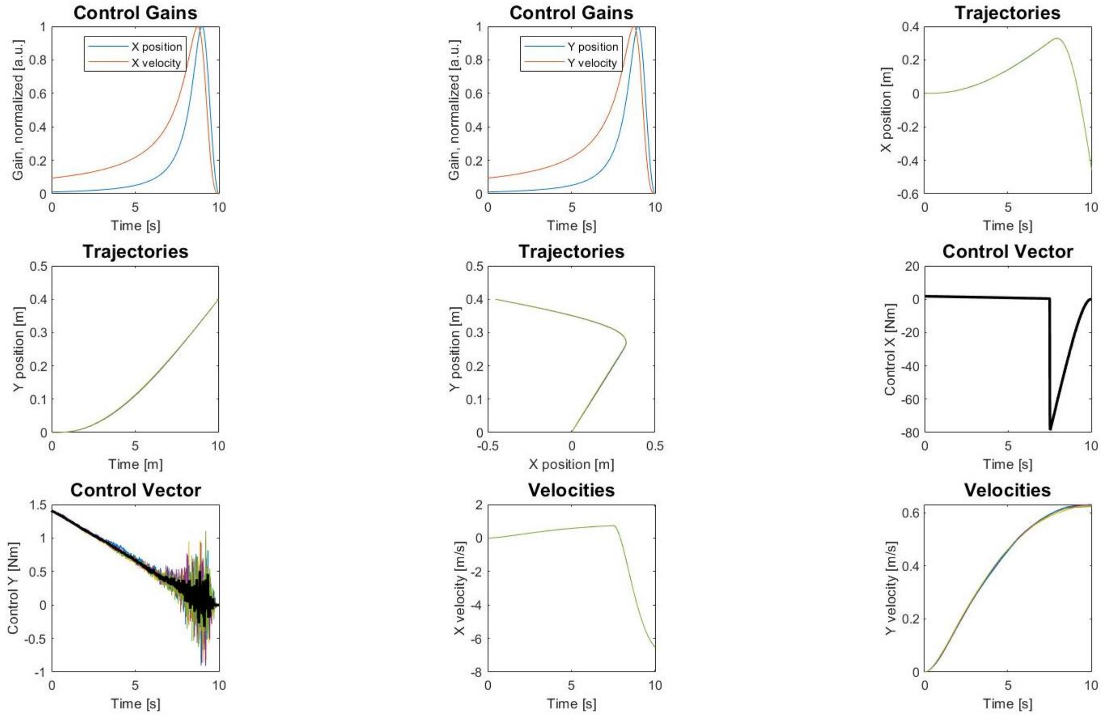

FIG. 6: *PERFORMANCE OF A TARGET JUMP AT 75% OF THE SIMULATION TIME, FROM DESIRED POSITION TARGET (0.5, 0.4) TO A NEW TARGET (-0.5, 0.4), WITH NON-ZERO WEIGHTS IN POSITION AND ZERO WEIGHTS IN VELOCITY.*

Once the first part of the study was done, considering there were only single targets and simulations were done regarding the time the decision was taken and the penalisation we apply to each state variable, further experiments including a second target were done, in order the address the main question of the study. However, questions regarding controller design and the effect of unexpected events, can already be assessed and are discussed later in this thesis.

2.2. Experiment: multiple targets

In this second stage of the project, to assess the functioning of the motor plan, a second target was included in such way, that initially the mass had two potential targets. To answer the main question of the thesis, the targets $(x, y) = (0.5, 0.4)\ m$ and $(x, y) = (-0.5, 0.4)\ m$ are simultaneously used. However, to assess the performance of the controller, during the simulation one of the targets is eliminated in order to leave only one potential target. In terms of the code, the second target is made equal to the first time at different times of the simulation. If there are 1000 time-steps, then the first target is eliminated at time-step 600, 700, 800 and 900 to observe the trajectory and correction in every case, as well as leaving both targets untouched until the end of the simulation in another case. Since the effect of the weight in velocity is already analysed in figures 3, 4, 5 and 6; figure 7 considers the weight in velocity equal to zero, as well and the desired final velocity $(x', y') = (0, 0)\ m/s$ like in the previous experiments.

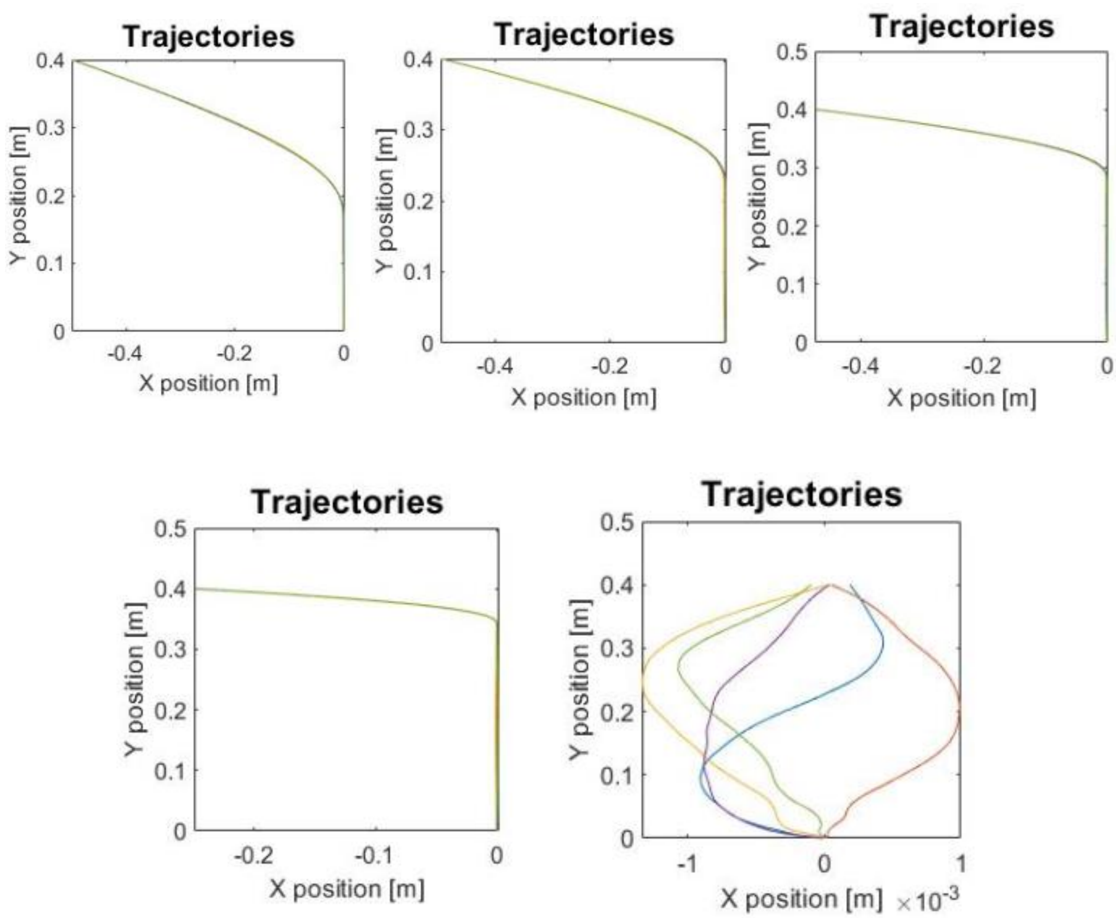

FIG. 7: *TRAJECTORY WITH TWO INITIAL POTENTIAL TARGETS (0.5, 0.4) AND (-0.5, 0.4) AND THE ELIMINATION OF THE FIRST TARGET AT DIFFERENT TIME-STEPS, LEAVING ONLY THE TARGET (-0.5, 0.4). EACH FIGURE SHOWS THE ELIMINATION OF THE TARGET AT TIME-STEPS 600, 700, 800, 900 AND WITHOUT CHANGE, FROM TOP-LEFT TO BOTTOM-RIGHT RESPECTIVELY. WEIGHTS IN VELOCITY ARE EQUAL TO ZERO.*

From the graphs in figure 7, it is observed that doing the elimination of the target in earlier time-steps, the final error in position is smaller. At time-step 600, the final relative error in position is <4% whilst at time-step 900, the final relative error is <54% which is quite big. At time-steps 700 and 800, the relative error remains small, <8% even though the change is done at 70% and 80% of the simulation. In the case of leaving the two targets untouched until the end, even though it looks like the trajectories are very random, it can be seen in the X axis that the maximum deviation is of the magnitude of 10^{-3}, which is almost negligible in comparison to the magnitude of the position of the targets. This means that leaving the targets untouched, the trajectory follows almost a straight line towards the intermediate target, which is $(x, y) = (0, 0.4)\ m$. The rest of the graphs also show this tendency towards the intermediate target before the elimination one of the targets, even though they correct the trajectory towards the only available target afterwards.

Even though the effect of the weight in velocity was already assessed in figures from 3 to 6, figure 8 simulates the same results as figure 7, but considering the weight imposed in velocity being non-zero and equal to the weight in position, in order to give more consistency to the results obtained previously during the single target jump experiment.

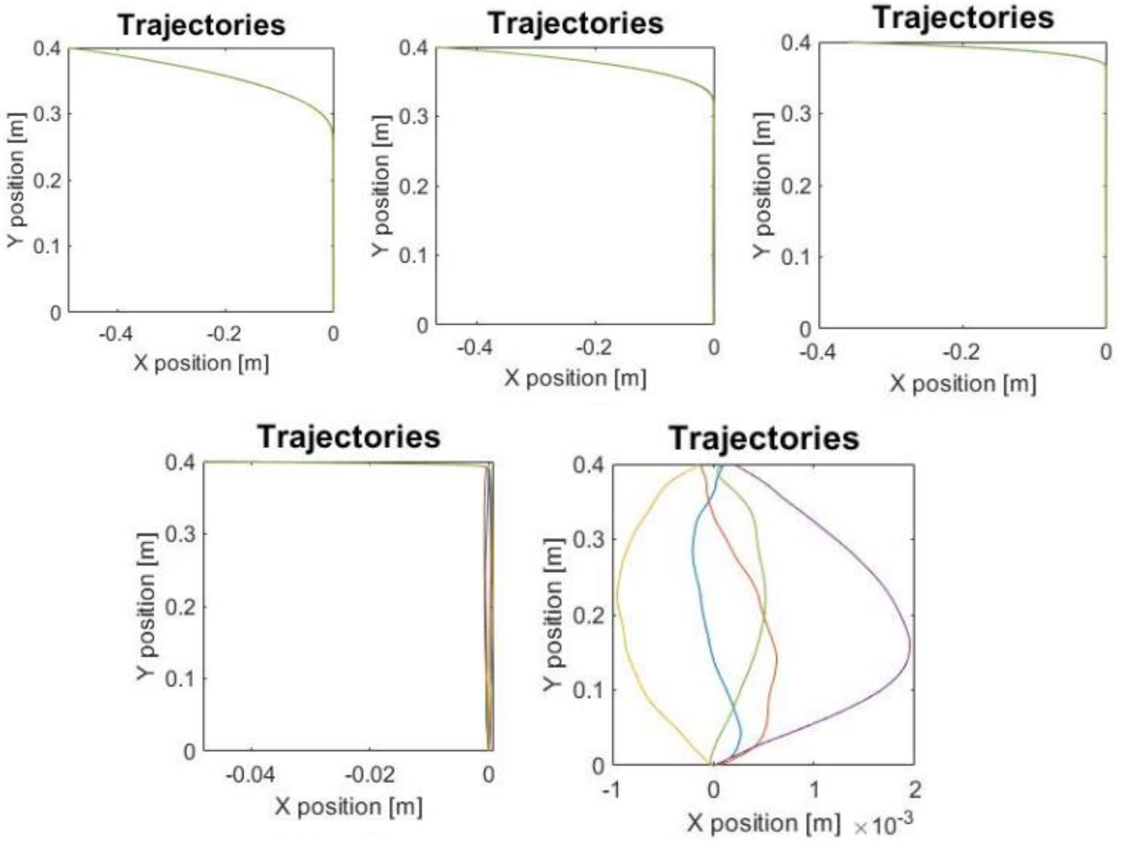

FIG. 8: *TRAJECTORY WITH TWO INITIAL POTENTIAL TARGETS (0.5, 0.4) AND (-0.5, 0.4) AND THE ELIMINATION OF THE FIRST TARGET AT DIFFERENT TIME-STEPS, LEAVING ONLY THE TARGET (-0.5, 0.4). EACH FIGURE SHOWS THE ELIMINATION OF THE TARGET AT TIME-STEPS 600, 700, 800, 900 AND WITHOUT CHANGE, FROM TOP-LEFT TO BOTTOM-RIGHT RESPECTIVELY. WEIGHTS IN VELOCITY ARE NON-ZERO AND EQUAL TO WEIGHT IN POSITION.*

As it is expected, results in figure 8 show less accuracy in the final position. Even though doing a change at time-step 600 cause a final error <4%, the error for time-steps 700, 800 and 900 are <8%, <28% and <90% respectively, which increase consistently in comparison to those obtained in figure 7. However, even though it's not shown in this last figure, final velocity obtained was the desired one, as opposed to the final velocity obtained in the results of figure 7.

Moreover, in this situation, a trajectory towards the intermediate point can be observed in each case, but with a slightly bigger error in accuracy as it can be seen in the last graph of figure 8. However, this error is almost negligible.

To be more critical with the results obtained, in order to answer the question about the two hypotheses, more experiments were done considering different targets as to analyse if the trajectory followed was consistent with the hypothesis of the intermediate motor plan or not. Therefore, different simulations were done changing these two potential targets each time and leaving them untouched until the end.

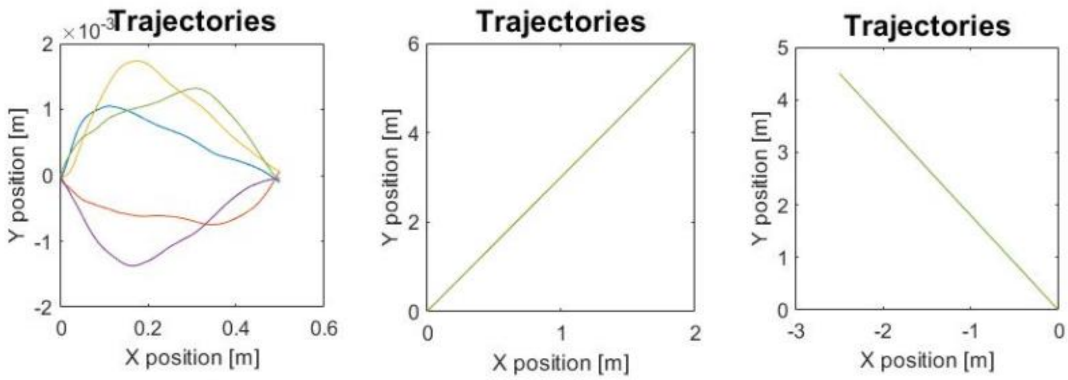

FIG. 9: *FROM LEFT TO RIGHT: THE FIRST GRAPH SHOWS THE TRAJECTORY HAVING THE POTENTIAL TARGETS (0.5, 0.4) AND (0.5, -0.4) SIMULTANEOUSLY UNTOUCHED UNTIL THE END OF THE SIMULATION. IN AN ANALOGUE WAY, THE SECOND GRAPH SHOWS THE TRAJECTORY WITH THE POTENTIAL TARGETS (6, 4) AND (-2, 8), AND THE THIRD GRAPH (-3, 4) AND (-2, 5). IN EACH CASE, THE WEIGHT IN VELOCITY IS EQUAL TO ZERO.*

Figure 9 shows the performance the controller regarding two different targets and leaving them untouched until the end of the simulation, to observe its performance. In the first case, the two potential targets were $(x, y) = (0.5, 0.4)\ m$ and $(x, y) = (0.5, -0.4)\ m$, and the position reached was exactly the intermediate point $(x, y) = (0.5, 0)\ m$ with an error <0.05%. In the second case, the two potential targets were $(x, y) = (6, 4)\ m$ and $(x, y) = (-2, 8)\ m$, and the position reached was again exactly the intermediate point $(x, y) = (-2, 6)\ m$. Finally, in the third case, the same result was obtained but this time considering the two potential targets $(x, y) = (-3, 4)\ m$ and $(x, y) = (-2, 5\ )\ m$. The position reached at the end of the simulation was $(x, y) = (-2.5, 4.5)\ m$ with an error <0.2%.

In each case, the trajectory followed a straight line towards the intermediate point of the two potential targets. It is important to consider too that the relative error between the expected final position and the actual one, throughout 5 simulations, was almost zero and therefore negligible.

With these experimental results being already enough to assess the main question of the study, one last experiment was done regarding the simulation time of all the experiments and the time of the target jump. The idea was to compare the performance of the controller considering two different running times: 1 second and 2 seconds of simulation.

During the experiments with multiple targets, the elimination of one potential at time-step 900 (out of 1000) was the one with the biggest error in comparison to the desired final position. Therefore, figure 10 compares the performance of the control eliminating the first potential target at time-step 900 out of 1000 time-steps, with the performance doing the elimination of the target at time step 1800 out 0f 2000 time-steps. In both cases, this elimination is done at 90% of the running time of the simulation, but in the second case, the running time and thus the number of time-steps is double to those in the first case.

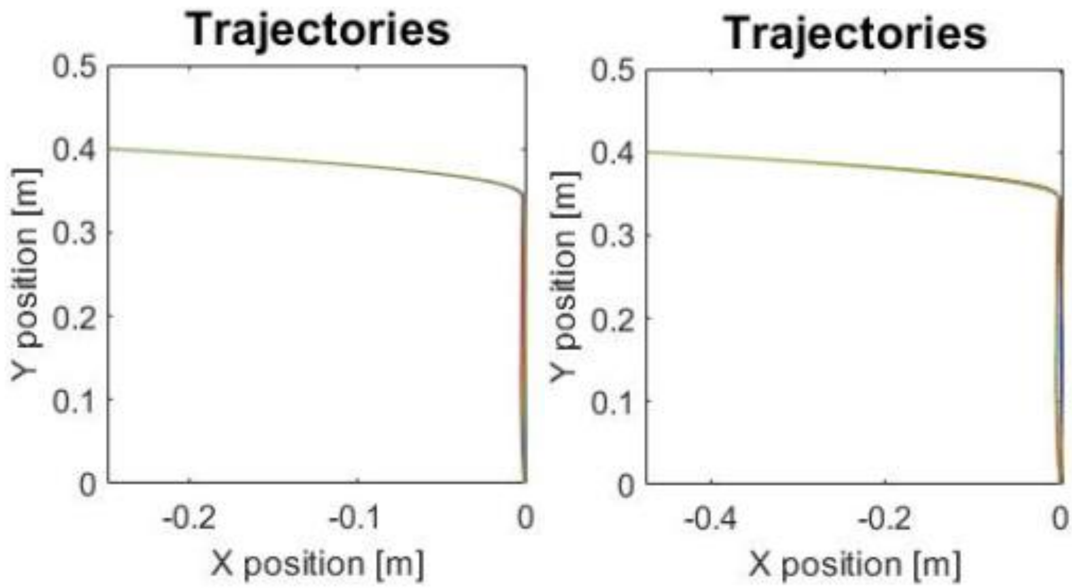

FIG. 10: *TRAJECTORY WITH TWO INITIAL POTENTIAL TARGETS (0.5, 0.4) AND (-0.5, 0.4) AND THE ELIMINATION OF THE FIRST TARGET LEAVING ONLY THE TARGET (-0.5, 0.4). THE FIRST GRAPH (LEFT) SHOWS THE TRAJECTORY FOLLOWED DOING THE ELIMINATION OF THE FIRST POTENTIAL TARGET AT TIME-STEP 900 OUT OF A TOTAL OF 1000. THE SECOND ONE (RIGHT) SHOWS THE TRAJECTORY WITH THE SAME ELIMINATION AT TIME-STEP 1800 OUT OF A TOTAL OF 2000. IN BOTH CASES, THE WEIGHT IN VELOCITY WAS KEPT TO ZERO.*

These results show that with a simulation time of 1 second, the relative error made in the final position reached was <54%, whilst for the simulation with a run-time of 2 seconds, the relative error made was <8%. The only differences between these two results are the run-time stablished and thus the total number of time-steps.

Regarding all these results obtained, they were later discussed considering other results obtained previously in other studies. The main points to discuss are the accuracy of the results and the "averaging" aspect, most related to the hypotheses discussed.

3. Discussion of results

3.1. Accuracy

Considering the results obtained previously, the first figures (Fig. 1 and 2) show a correct performance of the model in terms of accuracy. Firstly, the gains showed are not dependent on the targets set, due to the model being linear. However, they showed an expect increase towards the end of the simulation as the be able to correct the position and velocity when reaching the target.

Also, comparing both figures 1 and 2, the control vector showed a correction too towards the end when implementing a non-zero weight in velocity, seen in figure 1. This of course has the consequence that the final velocity was the desired one with a good accuracy. Since there was only a single target already stablished from the beginning of each simulation, the penalisation of velocity did not have a visible effect on the accuracy of position, as opposed to results observed later in other simulations.

Having all this in mind, it was proved that the model had a correct performance in terms of accuracy and was adequate to be used in further simulations with more targets.

Regarding the simulations in which a target jump was done, it could already be observed that the penalisation of velocity had a much bigger effect on the position accuracy. Figures 3 and 4 already show slightly this consequence, as the final error in position when the weight in velocity was set to zero, was slightly smaller. Figures 5 and 6, further show this consequence by doing the target jump later in the simulation. The target jump done at 75% of the simulation left less time for the controller to react and correct the trajectory, but in this situation, figure 6 showed a much smaller relative error, <10%, as opposed to the <45% relative error in position in figure 5. These results agree with the results obtained in the study done by Liu [4] where there was a trade-off between stability and accuracy. This means that the accuracy in final position wasn't just dependent on the time at which the target jump was done, but also the stability of the system. For instance, the experiment showed in figure 5 had a greater stability than for the experiment in figure 6, since the final velocity had to be controlled too.

Hence, when the weight in velocity was set to zero, the accuracy in position increased consistently with the trade-off of a lower stability in the system. If the constraint of reaching the target with a specific velocity was not considered, then the controller was able to reach the desired target with a lower relative error. However, it can be seen too that a late target jump lead to a bigger undershoot, comparing both figures 3 and 4, to figures 5 and 6. This is also consistent with the results obtained by Liu [4], where target jumps in later stages of the simulation lead to bigger undershoots.

Moreover, this concept of the time-undershoot being bigger un later target jumps was shown in the simulations regarding multiple targets. Figure 7 shows precisely two simultaneous potential targets and the elimination of one of them at different times of the simulation. For the target elimination at 60%, 70% and 80% of the simulation, the undershoot was relatively small in comparison to the target elimination at 90% of the simulation, from a relative error of 8% to <54% respectively. This is coherent with the results obtained also by Liu [4] where later target jumps resulted in a bigger undershoot. However, the differences in errors between Liu's results and the one obtained in this study, are all related to other factors, mainly the difference between the use of a computational model and the use of real human subjects. That is probably the reason why target jumps in very late stages, for example at 80% of the simulation, had a better position accuracy that target jumps at 50% with human subjects.

A simple explanation for this fast correction relies on the number of time steps used. It's evident that giving human subjects more time to react towards a target jump, would have an inherent increase in accuracy. For this computational model its was proven too. Figure 10 shows again this target elimination at 90% of the simulation time but having in mind a case where the simulation time was 1 second and another case where it was 2 seconds. In the second case, the accuracy of the trajectory was almost perfect even though this 'target jump' was done almost at the end of the simulation, and in both cases at 90% of the simulation This relies on the fact that there are double the number of time-steps and therefore the controller reacts twice as fast.

It's expected that a computational model and a human subject, both having to react towards the same target jump within the same time, would lead to differences in accuracy due to the faster reaction of the first one. This however means that a faster reaction has the consequence of a higher computational cost. Even though it's not part of this study, it is believed from this results that the correct parametrisation of this kind of models could resemble in a very similar way the performance and reaction time of a real human motor controller.

3.2. Averaging

Once the accuracy aspects were discussed, the study answers the main question proposed at the beginning concerning the averaging aspect of the results. Figures from 7 to 9 clearly show this performance.

Regarding figure 7 and 8, which have as a solely difference the penalisation in velocity, it is shown that there was always an average trajectory towards the intermediate point of both potential targets, before the correction due to the elimination of one of the targets. To prove that this average trajectory was indeed following the intermediate target as the unique potential target, they were both kept untouched until the end in one of the cases. The last graph in both figures 7 and 8 show that the trajectory went towards the intermediate the target in both cases. The fact that the relative error in both figures is almost negligible and very similar, shows that the controller considered that intermediate target as the only potential target and did not do and competition between two control plans considering the other two targets. If this was the case, figure 8 would have shown a much bigger relative error in position in comparison to figure 7. However, the same thing that happened in the first simulations with a single target, was observed in this case, where the weight in velocity didn't really have an effect in position accuracy as the target was set from the beginning. This clearly means that the controller had only one potential target from the beginning and not both simultaneously.

Figure 9 is simply a further experiment done to prove this point. Three cases were done with different multiple targets in each one, also having in mind different magnitudes in the target position. In every case, there was an average trajectory towards the intermediate target for which the relative error in position was <1% in every simulation done. These averaged trajectories agree with the results obtained by Chapman [5] were there was always an averaged, probabilistic motor plan as a function to the number and position of the potential targets.

4. Conclusions

To conclude, regarding the methodology and results obtained throughout this study, it would be sensible to say that the hypothesis describing a control plan towards an intermediate option, instead of an averaging of different motor plans, fits better human motor performance. Results clearly showed that the control strategy defined a controller towards a unique target even though there were simultaneous potential targets.

Moreover, even though these results were obtained from a computational standpoint which many times can be criticised for not being really representative of real human performance, they were coherent and agreed with results obtained in experimental trials done before in previous studies were real humans were the actual subject of the experiment.

As a result, this study wasn't just able to answer sensibly the main question proposed at the beginning but also opens new opportunities in computational modelling. This model was able to resemble in some way a real human-like performance but has also the potential to be further developed to be used in the development of medical devices and the area of bio-robotics. The fact that a human-like performance was achieved with such a simple model makes it possible to design and implement more complex control strategies, in robots for example, but achieving an even a better performance in comparison to a human being.

It has also to be mentioned that decision-making was simulated as well as the control strategy, which lies more in the area of neuroscience. Being able to simulate this kind of behaviour is one step forward in the understanding of human behaviour and thus its implementation in new medical devices.

It can be affirmed that this kind of studies are small steps forward in the area of neuroscience and human motor control but are necessary to follow the correct path. Computational models are nowadays gaining more leadership and are slowly becoming an essential tool in this field of study, and even though they have to be backed up with experimental models to demonstrate their reliability and validity, they have the potential to give us the ability to design very complex strategies in control theory and make us understand better human neurophysiologic behaviour.

5. References

[1] Crevecoeur F. (2019). Stochastic optimal control and Kalman filtering: Single-joint example. [PDF]

[2] Crevecoeur F. (2019). LQG control and applications to neural control of movement. [PDF]

[3] Crevecoeur F., Kurtzer I. (2018). Long-latency reflexes for inter-effector coordination reflect a continuous state feedback controller. *Journal of Neurophysiology* 120, 2466 – 2483.

[4] Liu D., Todorov E. (2007). Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *The Journal of Neuroscience* 27, 9354 – 9368.

[5] Craig S. Chapman, Jason P. Gallivan, Daniel K. Wood, Jennifer L. Milne, Jody C. Culham, Melvin A. Goodale. (2010). Reaching for the unknown: Multiple target encoding and real-time decision-making in a rapid reach task. *Cognition* 116, 168 – 176.