Error correction in an inference task

Dian van Elst
3624145
04-07-2013

7,5 ECTS
Begeleider Eindhoven: R.H. Cuijpers
Begeleider Utrecht: J.M. Broersen

1 Introduction

Many applications of robots need a certain level of navigation and interaction in or with a world. If, for example, a robot wants to help a human with accomplishing a certain task, it must be able to interact with the human and navigate to the same spot as the human. If the location of the human is known, navigating towards him or her is not such a big deal. But what if the human is walking to a certain goal and the robot wants to end up at the same spot? The robot can follow the human, although it is not necessary to take the same route. When mirroring the human's route the robot is perceived as less intelligent and likable compared to navigating to the correct spot himself, as found by Bruna (2011). Therefore, it is preferred to infer the human's goal and set it as the robot's goal.

As Bekkering, Wohlschläger & Gattis (2000) stated, the goal of a goal-based action is often more important than the route or way used to achieve the goal. Imitating behavior, here reaching the same location as the human, involves the decomposing of motor patterns into their constituent components, which are later reconstructed into an action pattern - the goal of the human's route. The observation of the human's movements is decomposed and translated to a certain goal or action. The goal, the end location of the robot, is more important than the way he got there.

The human mirror system plays an important role in observing the actions of others. Mirror neurons, originally discovered in the premoter cortex of monkeys, discharge both when performing a certain action and when observing others perform that same action. These neurons have also been discovered in humans (Cattaneo & Rizzolatti, 2009). Bekkering et al. (2009) recap how we integrate observed actions with existing actions in our own personal motor repertoire, using the mirror system. In this way we can understand the actions of others and we are able to detect errors in ongoing behavior.

The detection of errors enables us to engage in corrective and adaptive behaviors. Our action system is flexible enough to not only imitate actions, but also to perform goal-directed behavioral adjustments depending on the context of the interaction. Our ability to detect and correct errors may enable us to interact efficiently in joint actions.

Moreover, Van Schie, Mars, Coles & Bekkering (2004) found that the activation of the mirror neurons reflects what the observer would have done and not the representation of the action from the observing perspective. Iacoboni et al. (2005) found that the activity of mirror neurons is altered by the intention of the observed actions. Therefore the intentions are important during action observations. Inferring the goal of an action has the advantage of allowing the observing party to achieve the same goal, without using the same means. (Erlhagen et al., 2005).

This information is used to infer a goal from action observations. Often a Bayesian network is used for such purposes. Bayesian statistics offer a very easy and flexible framework for modeling decisions with uncertainty. Essential is Bayes' rule, which turns likelihoods into posterior probability distributions. This is modulated by the prior probability distribution (prior beliefs). The main assumption fis that the human brain rather makes decisions and inferences based on posterior probability distributions than based on likelihoods. It can then make statistically optimal decisions. Therefore, the brain must have knowledge about likelihoods and prior beliefs. Often prior probability distributions are assumed to be Gaussian.

Another more philosophical view, is Levi's (Nordmann, 1988) view of human decision making. He says that the available body of knowledge determines which events or hypotheses are serious options and which we will rule out immediately. We treat our knowledge as infallible, until the contrary is proven. If new observations introduce any inconsistencies in our body of knowledge, we must reconsider what we considered infallibly true. After weighing the new information against the risk of error, we may modify or reaffirm our old belief.

The above is consistent with Bekkering et al. (2009), for they both consider us capable of noticing errors and revising our decisions if necessary. This results in better and more efficient interaction.

1.1 Relevance for AI

Within the field of Artificial Intelligence decision making is a relevant field and finds applications in different domains. For example Artificial Neural Networks, Intelligent Agents and probabilistic reasoning models are used to assist decision making (Philips-Wren & Jain, 2006). Robots are often looked at as Intelligent Agents, for they observe the world through sensors and act autonomously. The following model will use Bayes, which is a probabilistic reasoning model and generally can be extended to a neural network, or a neural field as Cuijpers & Erlhagen (2008) showed.

As mentioned before, robots are perceived as more intelligent when they infer a goal instead of purely mirroring movements. Error detection is an important human feature for efficient interaction in joint actions. Therefore it is expected that robots will perform more efficient in joint action if they are able to detect and correct their errors. They might also be perceived as more intelligent, which is the essential goal of artificial intelligence.

1.2 Problem description

Suppose there is a situation in which a robot wants to interact with a human, who is moving around. It wants to end up at the same location as the human. He wants to infer the human's goal, based on observations. For this purpose he will use Bayes and the likelihoods of the possible locations. After inferring the correct goal he will find his route towards the location. to do so he must avoid obstacles and this can be accomplished in many different ways, which will not be discussed in detail.

As mentioned before the robot will, as humans do, consider his choice infallibly true, until evidence comes along which proves otherwise. Steinhauser & Yeung (2010) found that the detection of errors, through continuously observing the world, is crucial for optimal performance in humans. Besides detecting errors, compensating errors is vital for effective goal-based behavior. They considered error detection to be a decision process where observations in the world were evaluated against an internal criterion to decide whether an error had occurred or not. They described 2 fundamental stages of error detection:

- Accumulating observations (input)
- Reaching a decision (output)

They found different processes in the human brain responsible for these different stages. Therefore, it is likely that such a process must be implemented in the robot. Besides the decision process, the human mirror system plays an important role. The robot will reason about the human as if it is reasoning about himself, mirroring the human's actions. The mirroring of the actions will not be exact, but the robot will use different means to accomplish the same goals.

The aim of the model described in this thesis is to improve the robot's behavior in the situation described before, by implementing the possibility to detect error and correct his errors. Questionable is when the robot has made a wrong choice and when an error should be corrected.

First a description of the model will be given as well as some implementation details. The hypotheses formulated in the model description will be tested and the results will be described. Finally, the results will be discussed and a conclusion will be formulated.

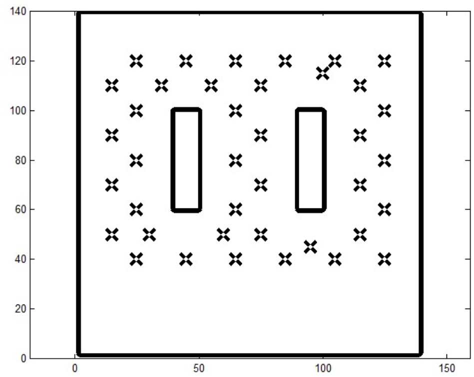

Figure 1: The world with targets and obstacles

2 Model

The described problem is illustrated by means of an inference task, where the location of another robot has to be deduced by a robot. The model is largely based on the model described in Cuijpers et al. (2006), with the addition of navigation and error detection. The model is simplified concerning the actions. The actions are degraded to simple locations (targets). No extra information about the targets is known, therefore no intentions can be assigned to a target. This simplifies the original model to a certain extent.

The original model described a cooperation task between two agents. They had to perform a construction task. A model had to be built from a collection of elementary components, like nuts, bolts, slats, etc. The agents knew what to do with the components and how to do that. They had to infer the goal of the other agent and perform some complementary action to help the building agent.

In our model knowledge is needed about the world and possible targets to infer anything about the other robot. The world consists of obstacles and two robots, as seen in Figure 1. One robot (referred to as Robot) navigates to a certain goal A, while avoiding obstacles and other robots. The other robot (referred to as Actor) has the task of inferring Robot's goal and navigating towards it.

Robot and Actor are in essence identical, which implies that they can reason about any action as if they have performed it themselves. This corresponds to the human mirror system as mentioned in the introduction. The robots have two sensors and are able to detect obstacles and each other. The robot's sensors both have a view direction of 30° in front of them, which makes a total view angle of 60°, with a small blind spot where the sensor fields meet. The range of the sensors is 20 on a world of 140 x 140 - about 14% of the world. This information is used to avoid collisions.

Those collisions are avoided by calculating a repelling value if an obstacle of any type is detected. The heading of the robots is defined by the repelling values of obstacles and robots and the attracting value of the target. The robot will always want to move towards the target, but the route is highly dependent on the obstacles.

Additionally, Actor has knowledge about the possible targets in the world and is aware of Robot's location at any time.

2.1 Likelihoods

As in Cuijpers et al. (2006) the likelihood of each target is calculated using observation o_t, which considers the locations of Actor ($\vec{x_e}$) and the Robot ($\vec{x_c}$) at time t. The observation is the distance between a possible target and Actor, defined by:

$$o_t = |\vec{x_c}(t) - \vec{x_e}(t)| \tag{1}$$

The likelihood that target t_n is the goal of Robot is defined as:

$$p(o_t|t_n) = \frac{1}{\sqrt{2\pi\sigma^2}}\text{exp} - \frac{o_t^2}{2\sigma^2} \tag{2}$$

where σ denotes the uncertainty. Equation (1) and (2) are defined as in Cuijpers et al. (2006).

Then Bayes' rule is applied. This is the first stage of the decision process, according to Steinhauser & Yeung (2010). Observations are accumulated into evidence.

2.2 Making a decision

After calculating the likelihoods for each target, the second stage of the decision process is reached: a decision must be made at some point. Just like humans, not all targets are considered. The nonsensical targets are disregarded. Relevant targets are the targets with the highest and the second highest probability.

The best target t can be defined as:

$$t_{best} = \arg\max_t p(o_t|t_n) \tag{3}$$

A decision must not be made too soon and not too late, to maximize the reward of reaching the correct target. The decision can be made if the likelihood of $p(o_t|t_{best})$ exceeds some predetermined level. Also, the likelihood ratio $t_{best}/t_{secondbest}$ must exceed the predetermined value α (Cuijpers et al., 2006) . A decision can be reached if:

$$p(o_t|t_{best}) > \alpha p(o_t|t_{secondbest}) \tag{4}$$

2.3 Detection of errors

However, knowledge and choices are not infallible and knowledge can change. The observations of the world change constantly and therewith the likelihoods. As described in the introduction, error detection is a process like decision making where observations are constantly evaluated to check whether they exceed a certain internal criterion (Steinhauser & Yeung, 2010).

What if Actor picked the wrong target? He can change his choice, but at a certain cost. If his pick is changed too soon, it might be a false negative. The target could have been the goal, but the likelihood might have dropped under the influence of one or more obstacles. We must prevent the changing of goal, when no good reason is present. Awareness of the influences of obstacles is important. Taking that into account, it is easy to see that Actor shouldn't switch rapidly but consider his options thoroughly. This must all be expressed in our definition of the internal criterion, a value like α before.

The essence of making a new decision is the same as for the first decision. However, the value of α will probably need to be another value, or even be defined differently. One possibility is to make α dependent on the cost of switching target:

$$\alpha = switchCost * threshold \tag{5}$$

Four different hypothesis were created concerning the definition of α, based on the general idea of Eq. (5).

	switchCost	**threshold**	α
static5	1	5	5
static1	$\frac{1}{threshold}$	5	1
dynamic5	$\frac{d_t}{d_{max}}$	5	$\frac{5d_t}{d_{max}}$
dynamic10	$\frac{d_t}{d_{max}}$	10	$\frac{10d_t}{d_{max}}$

The first hypothesis makes use of the same value for the threshold as α during the process of making the first decision. This is the baseline or normal hypothesis, as used for any decision in this model. Whereas hypothesis static1 reduces Eq. (4) to a mere comparison, thus Actor may change his pick at any time, given a better likelihood. This hypothesis is a minimum value and is the lowest possible threshold. Both can be called static values for they do not depend on anything in the world or anything that changes over time.

However, the latter two are dynamic in the sense that they depend on distances in the world. d_t describes the distance between the two targets under consideration and d_{max} is the maximum possible distance. Hypothesis dynamic5 uses the same threshold as the first decision process, whereas hypothesis dynamic10 uses the value 10 as a threshold to preserve the characteristics of the switchCost function. The α value should never drop under 1, otherwise Actor will only choose if the secondbest target is better than the best target, which is not possible by definition of the best and secondbest targets. Therefore, if the α value drops under 1, it is rounded to 1.

The internal criterion against which the observations are tested must be based on something. Just a threshold, as static5, is probably not sufficient for good performance. If the threshold is too low, the robot will choose too often and his path will not be the best or shortest path towards the goal. So we need a threshold that is not too low and neither too high.

The assumption is that it is easier to choose between two targets if they lie further apart. The likelihoods will differ more and the threshold can be higher to prevent error. If the targets are closer it is harder to make a choice, so the threshold must be lower to make a choice. For close targets the threshold must be low, but if the threshold is always low a lot of unnecessary errors will be detected. So a dynamically changing threshold should prevent unnecessary errors and ensure that the threshold is low enough if the targets are close.

The dynamic10 is fully based on the distances in the world and is not influenced by the threshold value. The dynamic5 uses the same threshold as during the first decision. Dynamic5 has a low adaptive value, while dynamic10 will produce higher values. Questionable is what the difference will be. Expected is that dynamic10 will perform better, because it doesn't change the observation or evidence.

2.4 Implementation details

The model is implemented using 38 locations, which are used as random instantiations of the robots' and targets' locations. First, as Bekkering, Wohlschläger & Gattis (2000) explained, the motor patterns are decomposed into their constituent components. Which is modeled as the observation of the location of Robot. Then, the components must be reconstructed to an action. Therefore, the likelihood is calculated and Bayes' rule is applied. When a decision is made, the chosen target is the new action or goal. The heading and velocity of Actor is updated and new observations are made. The decision is re-evaluated with each step and if an error is detected the target will be changed.

3 Results

3.1 First result set

The described implementation was run 1000 times for each proposed hypothesis with 2 random targets. Below we will compare the success of the hypotheses, the alpha values are described and the time they spent on reaching the goal. Besides that, the time they need to decide to revise their decision is shown, the number of switches and the deviation from the straight path are described. To analyze the results we will use Anova measures.

3.1.1 Success

An important measure is the success of the hypotheses. It is a global rate of the performance of Actor. This is calculated as the total number of runs that were successful, divided by the total number of runs. Actor is successful if he reaches the same goal as Robot. Figure 2 shows the percentage of successful runs. As expected, the static5 hypotheses performed worse than all the others. The other three hypotheses performed equally good.

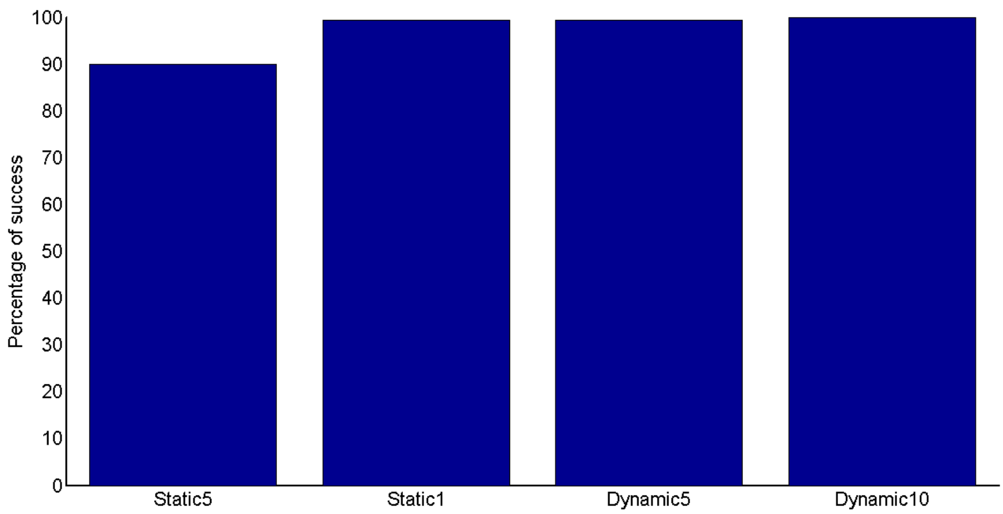

Figure 2: Measure of successfulness

3.1.2 Time

The total time needed to reach a goal can also be seen as a rate of performance. The quicker the robot is the better. Making and correcting errors will probably cost time and time is valuable. One unit of time is one check of the likelihoods for errors and an adjustment of the heading and velocity.

Figure 3 shows the time spent by all hypotheses. There a no significant differences: $F(3, 3996) = 0.82, p = 0.48$.

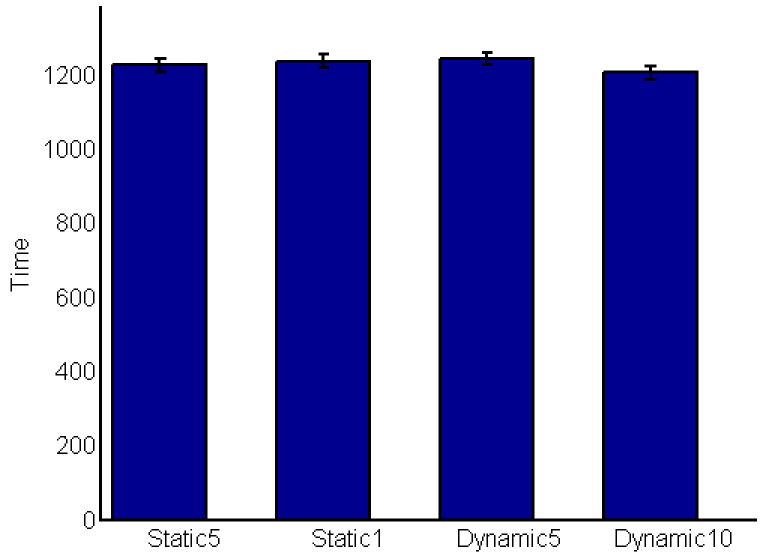

Figure 3: Average time spent to reach goal

3.1.3 Alpha

The α values for both static hypotheses are clear and constant. However, the dynamic hypotheses output different values for each run. Figure 4 shows the mean and standard deviation for all α values. The mean of the α values for the Dynamic5 hypothesis is 1.62, with a standard deviation of 0.58. Whereas the mean of the values for dynamic10 lies at 3.16, with a standard deviation of 1.38. The switchCost values for dynamic5 lie closer around 1.5. The α values of dynamic5 are closer to static1 and dynamic10 is closer to static5.

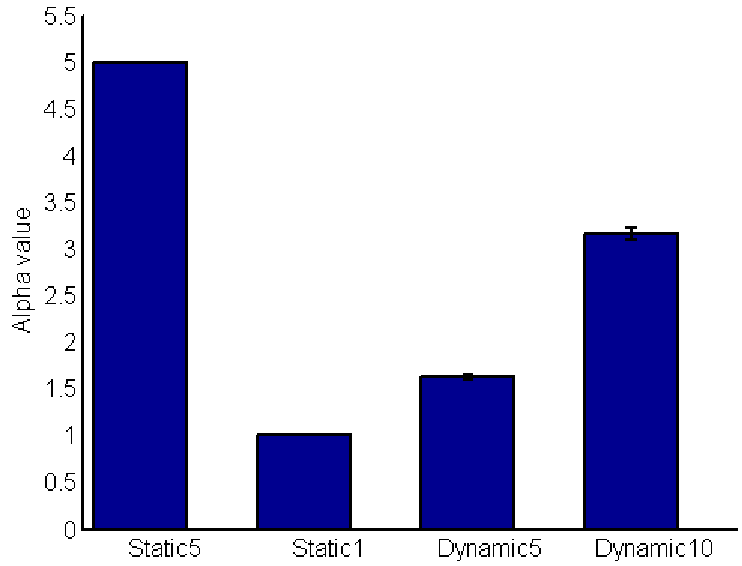

Figure 4: The average alpha values

3.1.4 Switches

In Figure 5 the number of switches made during a run can be seen. The values are expressed in a percentage of the total number of runs. If zero switches were made, the first decision was correct. About fifty to sixty percent of the decisions are made correctly the first time. The static5 hypotheses (Mean = 1.40, Standard deviation = 0.54) changes the decision less frequently than the other hypotheses. Significant results are found: $F(3, 3996) = 12.25, p < 0.0001$. The differences between the static5 hypotheses and the static1 (M = 1.54, SD = 0.61), dynamic5 (M = 1.52, SD = 0.57) and dynamic10 (M = 1.48, SD = 0.58) are significant, with respectively $p < 0.0001, p < 0.0001$ and $p = 0.0011$. Between the dynamic5 and dynamic10 hypotheses no significant difference is found. Neither is there a significant difference between the number of switches made by static1 and dynamic5.

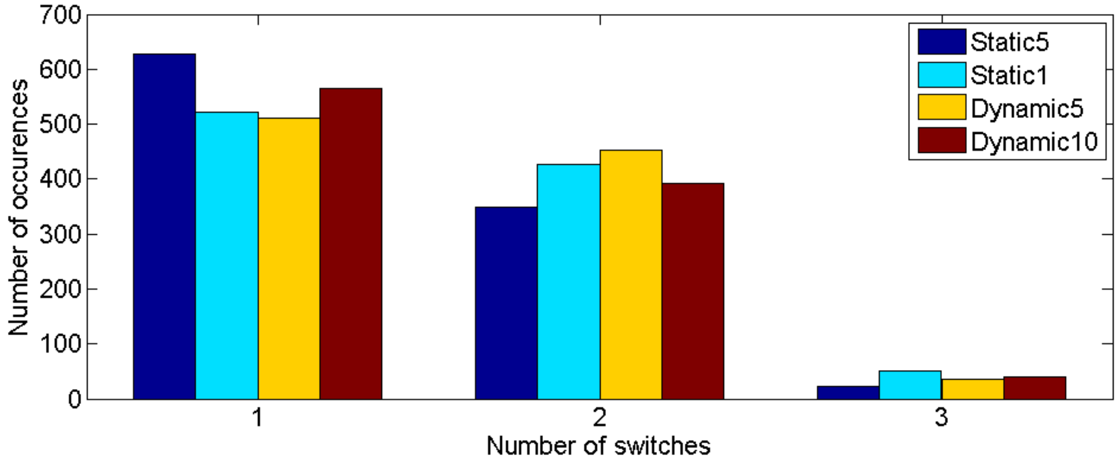

Figure 5: Number of switches made. Zero means no switches were made

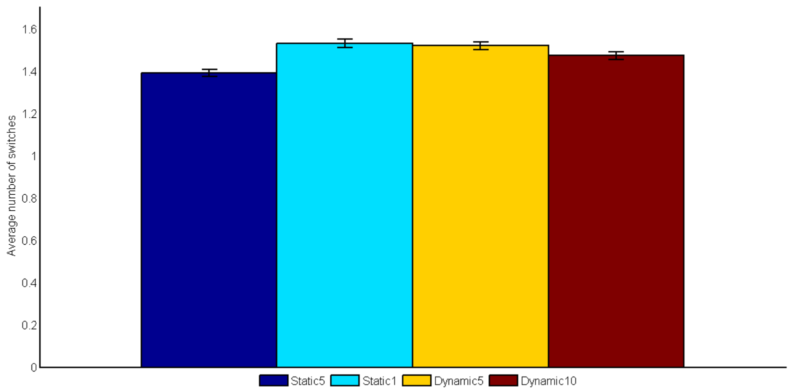

Figure 6: Average number of switches made

3.1.5 Decision time

Figure 7 shows the time needed to make a decision, or to correct an error. The time for the first decision is not significantly different for any hypothesis ($F(3, 3996) = 0.34, p = 0.80$), since they all use the same mechanism to make the first decision. For the second decision there is a significicant difference: $F(3, 1768) = 9.14, p < 0.001$. The static1 hypothesis (M = 36.71, SD = 19.31)significantly makes the quickest new decision (second decision), as expected from its α value, compared to static5 (M = 42.49, SD = 15.45), dynamic5 (M = 40.79, SD = 18.81) and dynamic10 (M = 41.91, SD = 18.85), with respectively $p < 0.0001, p < 0.0001$ and $p < 0.0001$. The other hypotheses have no significant differences. For the third decision, no significant differences can be observed: $F(3, 146) = 0.46, p = 0.52$. Noticeable is the difference in the time needed for spotting the first error and the second error. Spotting the second error is done much quicker.

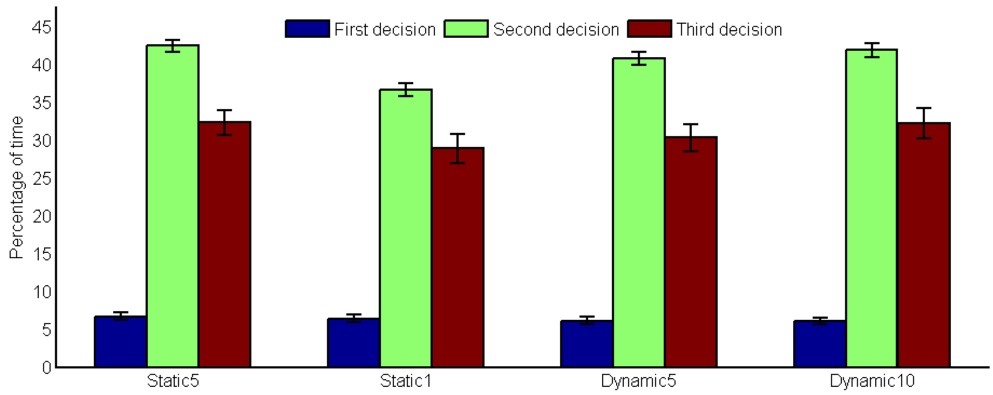

Figure 7: Time it took to make a decision

3.1.6 Path

Another measure is the maximum deviation from the best possible path. The best possible path starts at the initial location of Actor and goes in a straight line to Robot's goal. Obstacles are not taken into account. The maximum deviation is the distance to the point furthest away from the path.

Figure 8 describes the maximum deviation in general (all) and divided per number of switches. If no switches were made, the deviation is about equal for all hypotheses. No significant difference can be found: $F(3, 2182) = 023, p = 0.88$. In general (all) there is no significant difference between dynamic5 (M = 27.02, SD = 19.59),dynamic10 (M = 25.26, SD = 18.85), static1(M = 25.69, SD = 0.60) and static5(M = 25.43, SD = 0.62) is found, $F(3, 3900) = 1.67, p = 0.17$.

If one or more errors are detected, the deviation from the path is significantly different: $F(3, 1693) = 2.65, p < 0.0472$. For the static5 hypothesis (M = 38.38, SD = 18.47) the deviation is significantly larger than dynamic10 (M = 34.86, SD = 19.10) and static1 (M = 35.10, SD = 19.07), with both $p = 0.01$. However, the difference with dynamic5 is not significant. The static5 hypothesis probably takes a slightly longer route when switching.

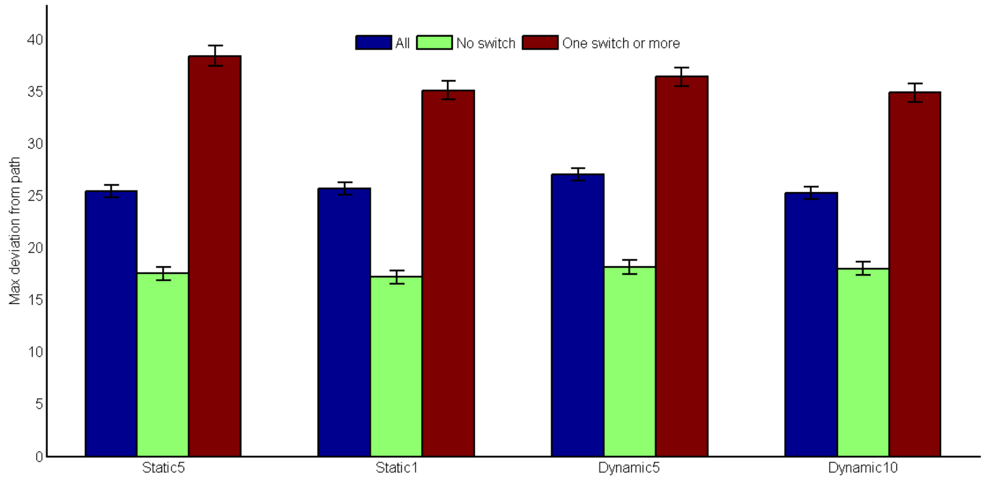

Figure 8: Deviation from path, per number of switches

3.2 Second run

Another run of 1000 was performed, with 5 random targets for each run. More possible targets will probably result in more errors detected and more errors made by the robot. The more choices, the harder it gets to decide. With these results we can show that the results are not dependent on the number of targets and that the hypotheses do not work better or worse depending on the targets.

3.2.1 Success

The success rate of static5 drops dramatically to 70%, whereas the other stay around 99 %. With more options the static5 hypothesis is outperformed by all other hypotheses.

3.2.2 Time

Figure 9 shows the time spent by the different hypotheses. In time there was a significant difference is noticeable: $F(3, 3996) = 3.23, p = 0.0214$. Dynamic10 (M = 1398.0 SD = 581.6) takes more time than dynamic5 (M = 1342.7, SD = 567.65), static1(M=1325.3, SD = 555.4) and static5(M = 1336.4, SD = 575.4).

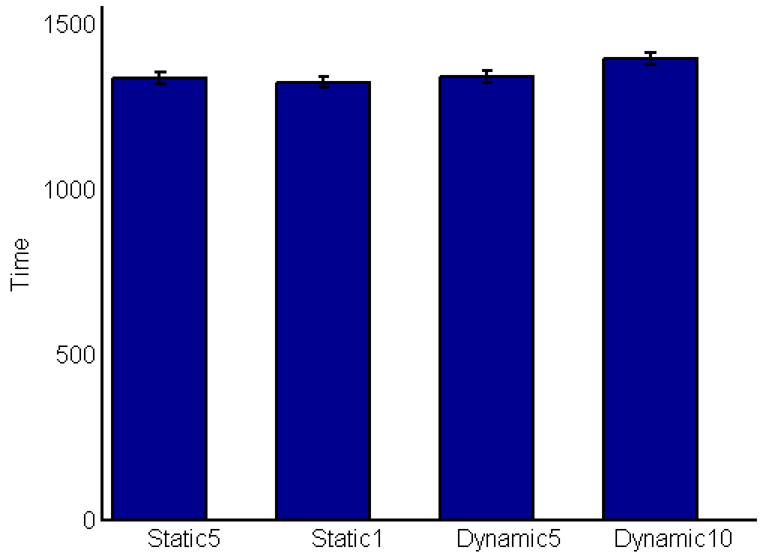

Figure 9: Average time spent to reach goal

3.2.3 Alpha

The α values stay the same, as shown in Figure 10. For dynamic10 the mean is 3.16 and the standard deviation is 0.07 and for dynamic5 M = 1.63, SD = 0.03.

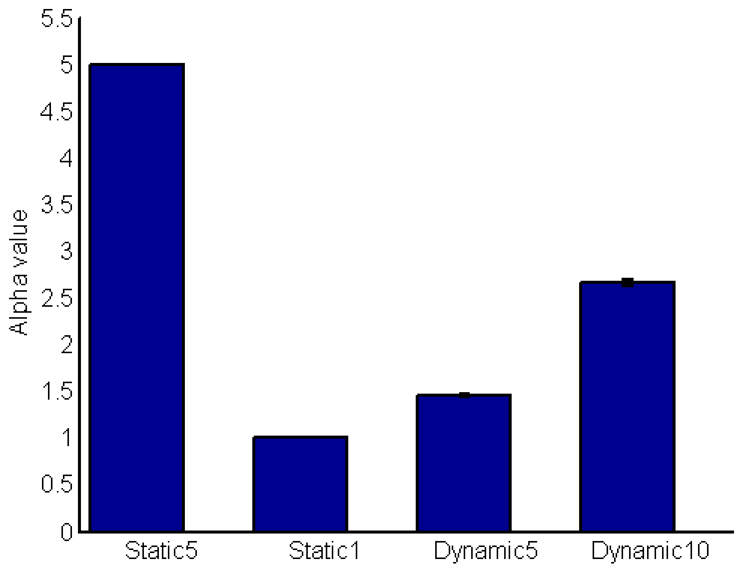

Figure 10: The α values

3.2.4 Switches

For the switches in Figure 11 and 12 a significant difference is apparent: $F(3, 3996) = 59.52, p < 0.0001$. The static5 hypothesis (M = 1.76, SD = 0.78) makes significantly less switches than static1(M =2.25, SD = 1.07), dynamic5 (M = 2.24, SD = 1.09) and dynamic10 (M = 2.27, SD = 1.07), with $p < 0.0001$. Just like the other result set.

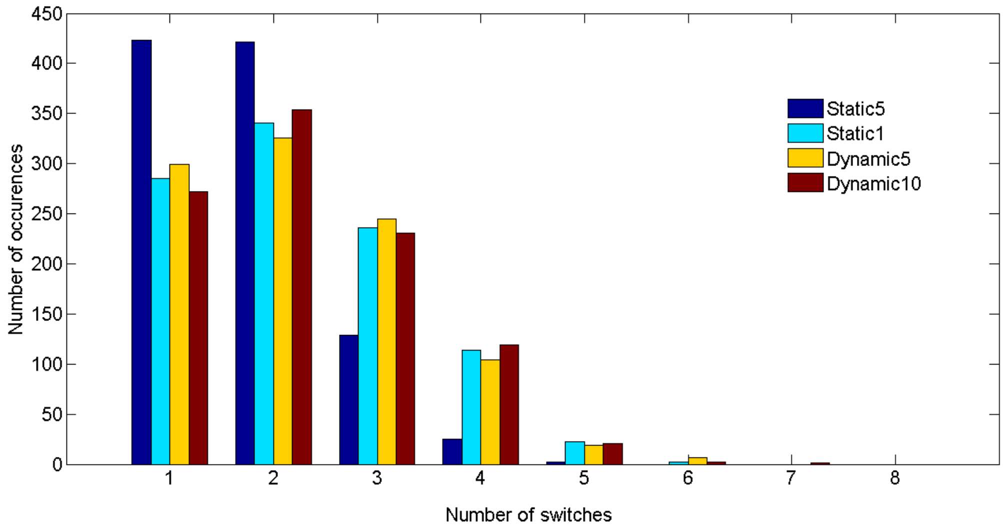

Figure 11: Number of switches made. Zero means no switches were made

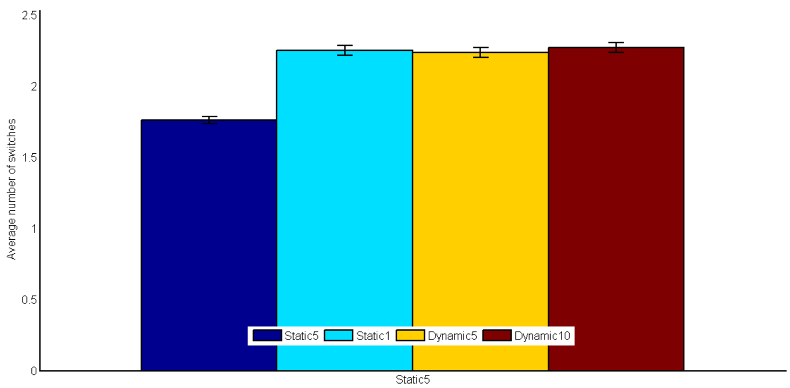

Figure 12: Average number of switches made

3.2.5 Decision time

Figure 13 shows the decision times. There is no significant difference for the first decisions made: $F(3, 3996) = 2.14, p = 0.09$. For the second decision a significant difference is noticeable, $F(3, 2717) = 34.26, p < 0.0001$. The dynamic10 (M = 31.1, SD = 17.45) hypothesis makes quicker decisions than static5 (M = 36.94, SD = 15.37), but is slower than dynamic5 (M = 27.43, SD = 16.01) and static1 (M = 28.05, SD = 16.84). For the third decision there are significant differences: $F(3, 1275) = 4.11, p = 0.006$. The static5 (M = 26.90. SD = 9.92) takes longer than static1 (M = 23.27, SD = 13.46), dynamic5 (M = 24.45, SD = 12.79) and dynamic10 (M = 24.73, SD = 12.11).

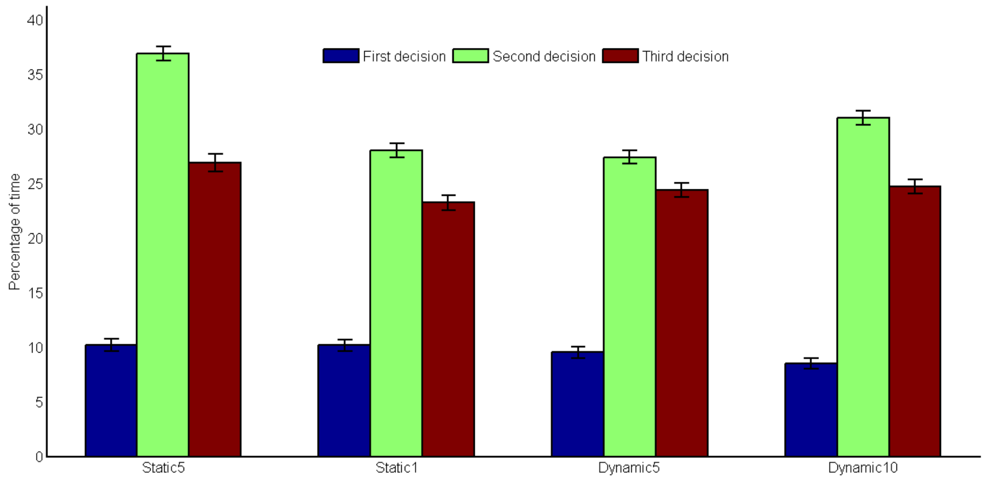

Figure 13: Time it took to make a decision

3.2.6 Path

In general the deviation from the path has significant differences: $F(3, 3900) = 2.62, p = 0.049$. The deviation from the path of static1 (M=30, SD = 19.40) is smaller than static5(M = 32.02, SD = 20.84) and dynamic10 (M = 32.45, SD = 19.85). However, there is no significant difference between the static1 and dynamic5 hypotheses. If no errors are detected there are no differences:$F(3, 1244) = 0.08, p = 0.97$. If at least one error is apparent no differences are significant: $F(3, 2500) = 1.02, p = 0.38$.

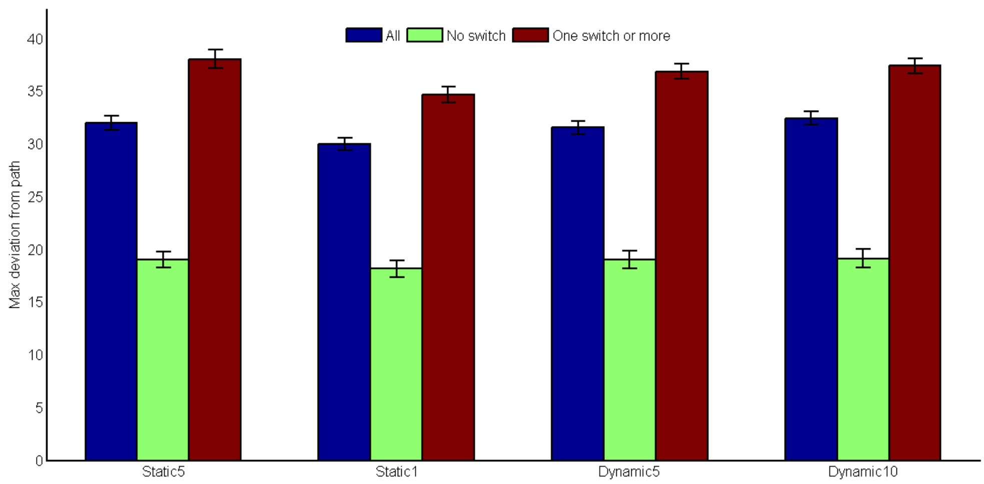

Figure 14: Deviation from path, per number of switches

4 Discussion

It was expected that static5 would perform the worst and that static1 would make the most switches. For dynamic5 and dynamic10 it was expected that they would perform better than both static hypotheses.

4.1 Static5

From the rate of success, which is about 90%, the conclusion can be drawn that static5 is not such a good hypothesis for the α value. The static5 hypothesis has the biggest deviation from the optimal path to the goal. This can be explained by the fact that about 10% of the runs do not reach the foreseen goal. This enlarges the deviation, because the path the robot chose has another goal. The number of switches made by static5 is the lowest, which could explain the inferior overall performance. If the robot is too ignorant to its errors, they will not be corrected and he will fail.

We can conclude that the α value of static5 is too high. The robot switches too little and he ends up in the wrong spot too often. Therefore, we will not consider static5 in the rest of the discussion.

4.2 Dynamic10

The overall performance of the dynamic10 hypothesis is good, if we take the success rate of almost 100% into account. however, not much differences with dynamic5 and static1 can be determined. The number of switches that dynamic10 makes is smaller than the static1 hypothesis. The success rate for static1 and dynamic10 are equal. One could say that dynamic10 performs slightly better than static1. However, there is no difference noticeable between dynamic10 and dynamic5. The maximum deviation from the path is smaller for dynamic10 than for dynamic5. Dynamic10 is on one result better than static1 and on one better than dynamic5. However, these are no convincing results. For the second result set the difference is that dynamic10 takes more time to reach the goal. He also takes more time to reach its decisions. This makes dynamic10 worse than dynamic5 and static1.

We can conclude that the α values of dynamic10 are successful on few targets, but performs worse when more targets are added.

4.3 Dynamic5

Dynamic5 is as good in its overall performance as dynamic10. However, on the deviation from the path this hypothesis performs worse. For all other results no significant differences could be found. For the second result set no other results are found.

We can conclude that dynamic5 is neither worse nor better than static1.

4.4 Static1

Static1 has the same success rates as both dynamic hypotheses. The hypothesis makes its decisions quicker than the others, but does not perform any worse. Although the others consider their decision more thoroughly, it makes no difference. Apparently is the time need to make a decision not of big influence on the total time needed to reach the goal. It was expected that static1 would make more switches, but no significant difference is apparent. For the second result set the static1 has a smaller deviation from the path than the other hypotheses. This is a surprising result because static1 does not make less switches. The deviation is apparently not fully dependent on the number of switches or the time at which the decision is made. On what is the deviation then also dependent? This question remains unanswered.

4.5 General observations

The α values fluctuate more for the dynamic10 hypothesis. This is a better reflection of the distance of the world, because multiplying the switchCost by 10 does not change the character of the values. However, the results show that it makes no difference compared with the other hypotheses.

The time a robot needs to detect the first error takes more time than the second error (third decision made). We can only speculate about the reasons for that. Most probably the first error spotted was not an error, because an obstacle was in the way. Robot was on his way to another target, because it had to avoid an obstacle. Actor thinks the he guessed wrong, but then Robot changes his route again because he has passed the obstacle. Actor changes his choice again. Since Robot is now closer to his goal, the likelihood will be higher for a target and drop for the other target. This might explain why Actor detects his error more quickly.

From the number of switches made and the corresponding success rate we can infer that the number of switches must not be too low, but we do not want them to be too high either. A threshold of 5 is too high, for the success rate drops terribly. An α value of 1 is acceptable, as are the dynamic α values. However, we would prefer to see a lower number of switches. With our hypothesized values we do not achieve that goal.

Why are our α values not good enough? A difference between the static1 and dynamic hypotheses was expected. Yet, no such conclusions can be drawn from the given data.

We discussed two stages of decision making:

- Evidence collection
- Reaching a decision

In both stages improvements might have to be made to improve the error detection.

We used the internal criterion found in humans by Steinhauser & Yeung (2010) as a basis for our assumptions on which we based the hypotheses to reach a decision. Questionable is whether this human process they found can be easily converted to a computational use. They found proof for their hypothesis in the human brain, but there might be other processes involved in creating the internal criterion. It might be complex and defined by other processes.

Our α values might have been dependent on the wrong world parameters. Maybe more or other knowledge is needed. Such as the angle of the robot with the target, or his heading. Only the locations of the target may not provide enough exact information, which results into the robot being indecisive and leads to a lot of false errors. The robot acts the same for a threshold of 1 and all dynamic thresholds.

The thresholds may lie too close to 1 to make a real difference. This could explain the dynamic5 behavior, for the average α is 1.6, this is close to 1. However, dynamic10 has an average α value of 3.16, so this behavior cannot easily be explained. When avoiding an obstacle the likelihood ratio will exceed α eventually, even with a higher threshold.

Furthermore, the robot does not take into account any information about the obstacles in the world. Indirectly it knows that obstacles are present, for it avoids them, but it does not take into account any detected obstacle information. This would be a quite complex process, but it has valuable information. If Actor knows the location of the obstacles it can predict the Robot's path better. It might detect that Robot is taking a detour, because of the obstacle. It can then decide to stick to it's old goal prediction and look into the likelihoods again when Robot is out of the obstacle's influence zone.

We only consider the two best targets, like Levi (1988) assumed we never consider all possible options. The problem most possibly lies not within this assumption made. Considering a third option will not influence the choice, for that likelihood is lower and it will not be ruled out by the threshold.

Levi also spoke about our body of knowledge. One option is adding more information, as mentioned above. Another option is to change the representation of the knowledge. Can targets be represented as simple locations or must other information be added to improve error detection? Cuijpers et al. (2006) added action information to the possibilities and Iacoboni et a. (2005) say that intentions are important in action observations. This information changes the stage of evidence collection and that indirectly changes the process of reaching a decision.

To notice errors it is necessary to check whether observations introduce inconsistencies in the body of knowledge. This seems to work, although too much errors are detected. As Levi said, the information must be weighed against the risk of error just like the internal criterion defined by Steinhauser & Yeung. This risk of error can probably defined like the internal criterion, which we tried to formulate. If we adjust these values, make them dependent on more or other information, better results might appear.

5 Conclusion

First, we can conclude that the inference does work. Three hypotheses have a 99% success rate, which is considered good enough. The work by Cuijpers et al. (2006) is not only applicable to construction tasks, but to navigation tasks as well. The computational interpretation of the human mirror neuron system seems to be correct.

The detection of errors is eventually successful, but is rather inefficient. The robot changes his mind often and seems to be quite indecisive. It comes down to the robot switching his choice if a target is better. Depending on the hypothesis the robot takes more or less time to decide. As humans we would consider that non-intelligent behavior.

We can conclude that small differences are noticeable between the hypotheses, but dynamic5 and static1 have no real differences. For now, the threshold values around 1 and 1,5 perform best. The assumptions made in the model are partly correct and others might be questionable. The error detection mechanism is working, but not as good as needed to make efficient interaction possible in joint actions. If the values of the internal criterion are improved, as discussed in the discussion, the model might work better.

For the field of Artificial Intelligence this means that improving the error detection mechanisms, will improve the perceived intelligence of a robot. Robots are able of detecting error, as shown, but rather inefficient. Some work still has to be done.

For future research the information used by the robot to infer his likelihoods could be improved. Only he location of the robot is scarce and may be too little information to base any decision on. Converting the targets to actions with intentions might also help the decision making, this also ads information. Another option is to base the internal criterion formulated as α values on more information and give a better distribution of the threshold values.

References

Bruna, M.T. (2011). The benefits of using high-level goal information for robot navigation, *Master thesis.*

Bekkering, H., Wohlschläger, A., & Gattis, M (2000). Imitation of gestures in children is goal-directed. *Quarterly Journal of Experimental Psychology, 53A*, 53-164.

Bekkering, H., De Bruijn, E. R. A., Cuijpers, R. H., Newman-Norlund, R., Van Schie, H T., Meulenbroek, R. (2009). Joint action: neurocognitive mechanisms supporting human interaction. *Topics in Cognitive Science,*1(2),340-352.

Cattaneo, L., Rizzolatti, G. (2009). The mirror neuron system. *Archives of Neurology,* 557-560.

Cuijpers, R. H. & Erlhagen, W. (2008), Implementing Bayes' rule with neural fields. *Artificial Neural Networks - ICANN 2008*, 228-237.

Cuijpers, R. H., Schie, H. T. van, van Koppen, M., Erlhagen, W. & Bekkering, H. (2006). Goals and means in action observation: A computational approach. *Neural Networks, 19*, 311-322.

Erlhagen, W., Mukovskiy, A., Bicho, E., Panin, G., Kiss, C. & Knoll, A. Van Schie, H., & Bekkering, H. (2005). Action understanding and imitation learning in a robot-human task. *Artificial Neural Networks: Biological Inspirations – ICANN 2005*, 261-268

Iacoboni, M., Moinar-Szakacs, I., Gallese, V., Buccino, G., Mazziotta, J. C. & Rizzolatti, G. (2005). Grasping the intentions of others with ones own mirror neuron system. *PLoS Biology, 3*(3), e79.

Nordmann, A (1988). Decisions and revisions: philosophical essays on knowledge and value by Isaac Levi; Hard choices: decision making under unresolved conflict by Isaac Levi. *Isis* Vol. 79, No.2, 310-312.

Philips-Wren, G., Jain, L. (2006). Artificial Intelligence for Decision Making. *Knowledge-based intelligent information and engineering systems*, 531-536.

Van Schie, H. T., Mars, R. B., Coles, M. G. H. & Bekkering, H. (2004). Modulation of activity in the medial frontal and motor cortices during error observation. *Nature Neuroscience, 7*(5), 549-554.

Steinhauser, M. & Yeung, N. (2010). Decision processes in human performance monitoring. *Journal of Neuroscience,*30, 15643-15653.