BSc Thesis Applied Mathematics

Detecting parameter change in dynamical systems using statistical methods

David de Haan

Supervisor: Prof. Dr. A.A. Stoorvogel

August 26, 2019

Department of Applied Mathematics
Faculty of Electrical Engineering,
Mathematics and Computer Science

UNIVERSITY OF TWENTE.

Preface

I want to thank Anton Stoorvogel for his guidance throughout the project, as well as his insightful view upon the subject of research. Furthermore, I want to thank my friends who helped me, and my family for their support throughout the making of this paper.

Detecting parameter change in dynamical systems using statistical methods

David de Haan*

August 26, 2019

Abstract

In this paper, different statistical methods for detecting parameter change in dynamical systems have been analysed. First, using the equations describing the system, different cases of errors have been identified. Using a t-test on the mean and a chi square test on the variance, we could detect these errors with a certain probability. The power on both of these tests is calculated to better these results. Furthermore, the time it takes for the system to reach a new steady state is analysed as well. Finally, the report gives an overview on the probabilities of detecting the right error in a dynamical system using statistical tests.

Keywords: fault detection, dynamical system, statistical methods, parameter change

1 Introduction

The world is getting more and more autonomous. Machines are being used more often and more efficient in factories, and vehicles are also getting more and more 'self-conscious'. Unfortunately, machines intend to break down after a certain period. Measurements by sensors may start to get unreliable and they should be replaced before they start to have an impact. An example of this can be seen when looking at the Boeing 737-MAX: during two of its flights over the past few months a sensor didn't work as intended, thus causing the plane to behave in the wrong way, which in these cases caused both planes to crash.

The detection and isolation of faults in dynamical systems by now is a popular subject of research. Many methods have already been developed for detecting these changes, but the demand for better performance is getting higher, such that systems are getting more reliable and more safe [3]. For this, two main factors on the detection of faults are optimised, namely the speed with which the error can be discovered (detection) and the determination of the location where the error takes place (isolation). This paper focuses on the latter.

Although different dynamical systems could vary a lot in how they behave, there still are systems that could often (partially) be described as a linear system [7]. This makes it easier to analyse for mistakes. For these systems, methods like the (adaptive) Kalman Filter have already been used and analysed multiple times [4, 6] to analyse and detect the mistake in a system, but there has not been much research done from a statistical point of view on the detection of faults in dynamical systems. Nevertheless, it seems interesting to determine an error with a certainty, such that it could be more reliable to detect what

*Email: d.dehaan@student.utwente.nl

sensor in a system failed.

The purpose of this paper is to analyse a linear dynamical system from a statistical point of view using statistical tests, as well as a combination of existing error detection methods in combination with statistical methods.

2 Constructing a first order motor

To start, we want to create equations describing a simple motor, which we can use to investigate its behaviour as several parameters change, in the following sections. We introduce a constant current (I) and a motor torque constant (K) as a tool of transforming this current into a force (F). In reality, this could be done by a gyrator. The resulting equation can be found in equation 1 [1].

$$F = KI \tag{1}$$

Using Newton's second law, we can then derive the following equation [1]:

$$F = J\frac{\mathrm{d}v}{\mathrm{d}t} + Dv \tag{2}$$

Where v denotes the (rotational) velocity, J the moment of inertia of the motor and D denotes the friction on the motor. When combining equation 1 and 2, this results in the first order differential that we will be using in the following sections, given in equation 3.

$$\frac{\mathrm{d}v}{\mathrm{d}t} = \frac{K}{J}I - \frac{D}{J}v \tag{3}$$

The chosen values for each of these constants can be found in appendix A. The system is monitored on it's velocity, i.e. $y = v$. To get an idea of how the system behaves, the machine is run for one second. This gives the image shown in figure 1.

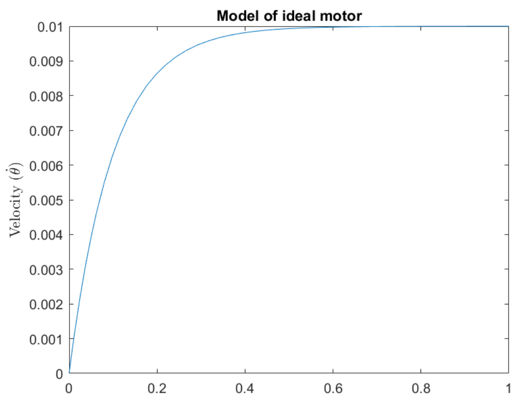

FIGURE 1: Continuous first order motor, reaching a steady state

As can be seen from figure 1, when under the influence of a constant current, the velocity slowly builds up until it reaches a certain (maximum) speed, which is a steady state. The system shown in figure 1 is a continuous system. This system should be discretised such that it is representable for our problem. To do this to our system, we assume that each step must be linear, thus creating the following equation:

$$\begin{aligned} a_k &= \frac{K}{J}I - \frac{D}{J}v_k \\ v_{k+1} &= v_k + a_k \cdot T \end{aligned} \tag{4}$$

Here, T represents the duration of one time step, which is a constant value. Note that there are better ways to discretise the system, but since the system is simple the equations are discretised like this to more easily elaborate on system behaviour. Now, this system is submitted to both a dynamical error ($V_k \sim N(0, P)$), as well as an observer based error ($W_k \sim N(0, Q)$), both of which are assumed to be additive white Gaussian noise. The variance of each of these errors are denoted by P and Q respectively. Both errors, V_k and W_k, are assumed to be independent (for each k). Adding these errors to equation 4, we get the following final equations that describe our system:

$$\begin{aligned} v_{k+1} &= v_k + \left(\frac{K}{J}I - \frac{D}{J}v_k\right)T + V_{k+1} \\ y_k &= v_k + W_k \end{aligned} \tag{5}$$

As can be seen from equation 5, the error V_k is added at the end of the equation, instead of being multiplied by the time constant. This is done so that V_k is not influenced by the size of the time constant. Finally, the system is again surveyed on its velocity. As can be seen in figure 2, the result shown in figure 1 can still be recognised.

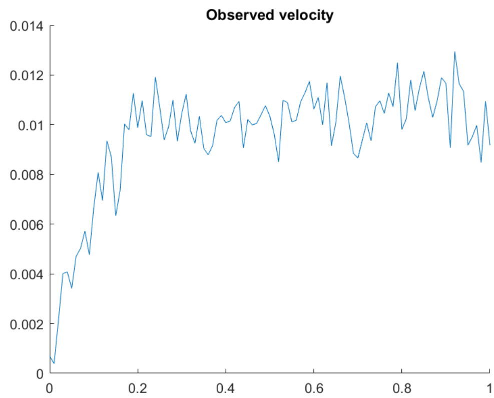

FIGURE 2: Discrete first order motor, subject to multiple errors

3 Introducing parameter change

Now that we have a simple first order motor, we analyse its behaviour as several components ‘break down’. First, we do this by simulating several parameter changes. Throughout this paper, we do this by changing the value of a parameter in our system after one third of our run, which in this case (see figure 3) corresponds to 0.33 seconds. In figure 3, the effects of several components failing have been shown.

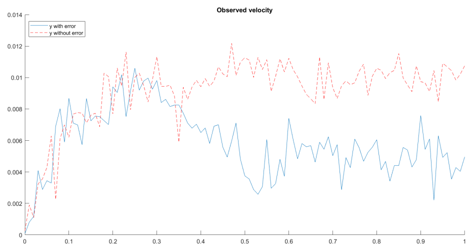

(A) Change in friction ($D_2 = 2D$)

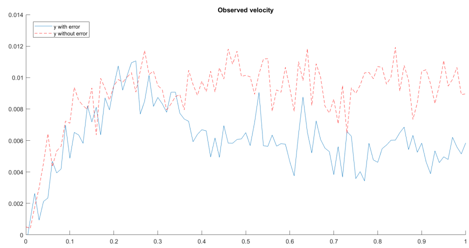

(B) Change in torque ($K_2 = K/2$)

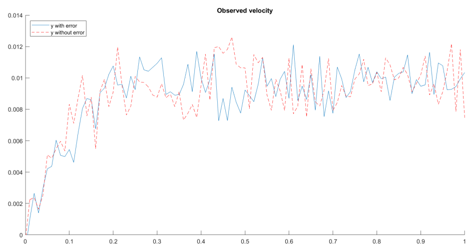

(C) Change in inertia ($J_2 = 2J$)

FIGURE 3: Effect of changes in several components, using 100 observations

When the results from figure 3 are compared to the observed velocity in figure 2, two characteristics stand out in particular. First, it can be seen that the faulty parameter of the moment of inertia of the system seems to have no influence on the mean of the steady state. This could be reasonable. Keeping equation 3 in mind, at the moment the system is in a steady state, a change in the moment of inertia has (almost) no influence on the change in velocity.

Another thing that stands out is that the adjusted value of both the friction as well as the torque shows a similar behaviour: both functions settle to a new steady state value: around half of the original value. However, it seems like the new steady state for the adjusted friction settles faster than the new steady state for the adjusted motor torque, which we will analyse in detail in section 4.2.

Now that we have an idea of what happens when changing certain parameters, we want to solidify these presumptions theoretically by analysing the equations describing the system. Since we notice a change in steady state, we analyse the value of y_k at a certain point, where the function has reached its steady state. We do this by looking at the expected value the function has taken. This can be found in equations 6 and 7.

$$
\begin{aligned}
E[y_k] &= E[v_k + W_k] = E[v_k] + E[W_k] = E[v_k] \\
&= E\left[v_{k-1} - \frac{D}{J}Tv_{k-1} + \frac{K}{J}IT + V_k\right] \\
&= E\left[\left(1 - \frac{D}{J}T\right)v_{k-1}\right] + E\left[\frac{K}{J}IT\right] + E\left[V_k\right]
\end{aligned}
\tag{6}
$$

For simplicity, we define $a = 1 - \frac{D}{J}T$ and $b = \frac{K}{J}IT$.

$$
\begin{aligned}
E[y_k] &= E\left[av_{k-1}\right] + E\left[b\right] + E\left[V_k\right] \\
&= aE\left[v_{k-1}\right] + b \\
&= aE\left[av_{k-2} + b + V_{k-1}\right] + b \\
&= a^2 E[v_{k-2}] + ab + b \\
&\vdots \\
&= a^k E[v_0] + a^{k-1}b + ... + ab + b \\
&= \left(a^{k-1} + ... + a + 1\right) b \\
&= b\sum_{i=0}^{k-1} a^i \\
&= b\left(\frac{1 - a^{k-1}}{1 - a}\right)
\end{aligned}
\tag{7}
$$

Note that this can only be done for $|a| < 1$. Since $a = 1 - \frac{D}{J}T$, we can choose our time step size T such that this is indeed the case. Furthermore, we notice that if our system has reached steady state, k has to be of a large enough value such that a^{k-1} goes to zero. Using this we see that our expected y_k is approximately equal to the following:

$$
E[y_k] = \frac{b}{1-a} = \frac{\frac{K}{J}IT}{1 - \left(1 - \frac{D}{J}T\right)} = \frac{KI}{D} \tag{8}
$$

Indeed, we see that our steady state is dependant on three factors: K, I and D. This is in line with the results found in figure 3, as we see that the steady state adjusts for both a change in motor torque K and friction D, but not for a change in inertia J. We also see that for a motor torque that gets twice as low, the steady state is the same as if the friction gets twice as high, which is also in line with the result found in equation 8.

Now, having found an expected value for y_k, we could oppose a statistical test which we could use to test our system. Seeing that we want to test on the mean of our system, we now now that this mean should be equal to $\frac{KI}{D}$ for non-adjusted parameters. Using the values from appendix A, we find that our mean μ equals 0.01, which is in line with our previously found results in figures 1 and 2.

To detect whether our system has indeed a mean of 0.01, we need a statistical test to find out if this is a valuable hypothesis for our system. The two most common known statistical tests on the mean are the z-test and the t-test[5]. Both of these tests assume that the sample data is normally distributed. Thus we need to check for that first. It is known that the system described has two errors, one dynamical as well as one observer based error. Since both errors are additive, white Gaussian noise, it follows that the noises are distributed via a normal distribution. Since the addition of two normally distributed random independent variables is still normally distributed, we can assume our system is normally distributed at steady state.

Now, we have two tests which are suitable: the z-test and the t-test. The major difference is that the z-test assumes variance is known and the sample size is large, while the t-test assumes variance is unknown and the sample size is not that large. Since the sample size can be chosen to be any value, we first mainly focus on whether we can find the variance of our system. Since v_k and W_k are independent of each other, we notice that equation 9 holds.

$$\begin{aligned}\mathrm{Var}(y_k) &= \mathrm{Var}(v_k + W_k)\\ &= \mathrm{Var}(v_k) + 2\mathrm{Cov}(v_k, W_k) + \mathrm{Var}(W_k)\\ &= E[v_k^2] - E[v_k]^2 + \mathrm{Var}(W_k)\end{aligned} \tag{9}$$

From equation 9, we see that the variance of our system is dependant on the squared expected value of v_k. To make this equation more easy, we shift our function, such that our steady state is equal to zero. This way, we not only get that $E[v_k]^2 = 0$, but our system could also be described by a more simple equation, which can be found in equation 10. In this equation we use that $a = 1 - \frac{D}{J}T$.

$$\begin{aligned}v_k &= av_{k-1} + V_k\\ &= a^2 v_{k-2} + aV_{k-1} + V_k\\ &\vdots\\ &= a^k v_0 + a^{k-1}V_1 + ... + aV_{k-1} + V_k\\ &= a^{k-1}V_1 + ... + aV_{k-1} + V_k\end{aligned} \tag{10}$$

Now, all that is left to do, is to square this equation and to find its expected value. This is given in equation 11.

$$\begin{aligned}E[v_k^2] &= E\left[\left(a^{k-1}V_1 + ... + aV_{k-1} + V_k\right)^2\right]\\ &= E\Big[a^{k-1}V_1\left(a^{k-1}V_1 + a^{k-2}V_2 + ... + V_k\right) + a^{k-2}V_2\left(a^{k-1}V_1 + ... + V_k\right) +\\ &\quad ... + aV_{k-1}\left(a^{k-1}V_1 + ... + V_k\right) + V_k\left(a^{k-1}V_1 + ... + V_k\right)\Big]\\ &= a^{k-1}E\left[V_1\sum_{i=0}^{k-1}a^iV_{k-i}\right] + a^{k-2}E\left[V_2\sum_{i=0}^{k-1}a^iV_{k-i}\right] + ... + E\left[V_k\sum_{i=0}^{k-1}a^iV_{k-i}\right]\end{aligned} \tag{11}$$

The expected values found in equation 11 seem very complicated. However, since the dynamical errors V_i are independent of each other, many terms can be removed, leaving only the few terms that can be found in equation 12. Here, the fact is used that the variance of each dynamical error V_i is the same and that for big values of k, which we have in steady state, a^{2k} gets very small.

$$
\begin{aligned}
E[v_k^2] &= (a^{k-1})^2 E\left[V_1^2\right] + (a^{k-2})^2 E\left[V_2^2\right] + ... + E\left[V_k^2\right] \\
&= a^{2k-2}\text{Var}(V_1) + a^{2k-4}\text{Var}(V_2) + ... + \text{Var}(V_k) \\
&= \text{Var}(V_k)\left(a^{2k-2} + a^{2k-4} + ... + a^2 + 1\right) \\
&= \text{Var}(V_k)\sum_{i=0}^{k-1}(a^2)^i = \text{Var}(V_k)\frac{1-a^{2(k-1)}}{1-a^2} \\
&\approx \text{Var}(V_k)\frac{1}{1-a^2} = \text{Var}(V_k)\frac{1}{1-\left(1-\frac{D}{J}T\right)^2}
\end{aligned}
\tag{12}
$$

From equation 12 we can indeed conclude that a variance can be found. However, it seems that our variance is dependent on some of our parameters as well. Both the friction D and the inertia J have an influence on it. From figure 3, it seems that these parameters do not have a big influence on the variance, although that could be explained by the fact that the variance of the observer based error W_k is ten times bigger than that of the dynamical error V_k, which makes a change in that harder to detect. Nevertheless, we can possibly use a change in variance to measure what parameter has changed.

One final thing that stands out from equation 12 is that the time step size T has an influence on the variance as well, which seems peculiar, as it seems illogical that the time step size should have any influence on the behaviour of our system. However, looking at equations 4 and 5, it is indeed logical to say that the time step size has a similar influence as the inertia on the behaviour of our system.

Coming back to a statistical test on the mean, we can say that we can calculate both a mean as well as a variance for our system. Therefore, we can still choose to use both a t-test as well as a z-test on the mean. Since we want our system to detect quickly if an error has occurred, the t-test is the preferred statistical test, as we want our sample size to be low. Using this we hopefully can detect a change in friction D and motor torque K.

Now that we know that our variance changes as well, we want to oppose a statistical test for that as well. Since our system has only one variance, there is only one preferred test, which is the chi-squared test on a single variance. Using this we hopefully can detect a change in friction D and inertia J.

Finally, using the equations we found for the mean and variance in this section, we can make an overview on how to detect what exact parameter has changed in our system, looking only at the mean and the variance in the steady state of our system.

Faulty parameter	Change in mean	Change in variance
Friction (D)	Yes	Yes
Motor torque (K)	Yes	No
Inertia (J)	No	Yes

Table 1: Influence of parameter change on the mean and the variance of the system

4 Testing on the mean

Now we want to use a t-test on the mean of our system. From section 3, we know that we want to test on the hypothesis H_0, whether our mean μ is equal to 0.01. For this test, we first need to find a T value for our test, where $T = \frac{\bar{Y}-\mu}{s/\sqrt{N}}$. In this equation, N denotes the amount of measurements used.

Since we know our system is normally distributed, we can find a distribution $\bar{Y}$ using measurements. First, we take a sample using a few values $y_1, y_2, ..., y_n$ and find a mean value $\bar{y}$ using our samples. The value of this can be calculated like shown in equation 13 [5].

$$\bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i \tag{13}$$

Using this, we can find a mean for our distribution Y. Furthermore, from equation 8, we found an expected value of y of the steady state for different values of D, K and J. Now, we try to find out whether our null hypothesis, $H_0 : \mu = 0.01$ should be accepted against $H_1 : \mu \neq 0.01$. Since it is unknown what the exact value of the variance σ^2 is, we could not simply see whether our measured values are normally distributed with certain mean μ and variance σ^2. However, we could still find a sample variance using equation 14.

$$s^2 = \frac{1}{n-1}\sum_{i=1}^{n} (y_i - \bar{y})^2 \tag{14}$$

We try to find out how trustworthy the tests are by first using it on a system that does not break down, compared to the steady state value μ of the normal continuous model. The system is being run for 5 seconds, the measurements used for the test are the ones at the end of the observations from the system, since the system has reached its steady state by then (see figure 4). This means the total amount of observations is twice as big as the ones used for the test. From each of these sample sets, a distribution Y_i is created. These distributions are used for the t-test. The results of the t-tests for data from 1000 runs are shown in table 2. Each run has a confidence level of 95%.

Total observations	# Samples	# H_0 rejected
500	25	174
100	10	82
50	5	88
25	5	74

TABLE 2: Amount of times H_0 was rejected, over 1000 runs

As can be seen, the amount of times H_0 gets rejected is higher than expected, as we would expect the test to be rejected 5% of the time. This can be explained by examining the variance of our system. Looking back at section 3, we discovered that the variance of our system should be equal to $Var(V_k)\frac{1}{1-\left(1-\frac{D}{J}T\right)^2}$ (see equation 12). In case of the first test in table 2, we would have a friction D of 0.1, an inertia J of 0.01 and a time step size

T of $\frac{5}{500} = 0.01$. That would mean the variance of v_k should approximately be equal to 5.26×10^{-7}. However, if we ignore the observer based error W_k for now, and calculate the variance of our system using the last 25 observations, we find that the variance is equal to 4.96×10^{-7}, which is much lower. This shows that the data from our system suffers from dependency: due to the fact that the variance in each state is dependent of the variance of each previous state, the total system variance is lower than it should be. The fact that for less total observations used in table 2 the number of times H_0 gets rejected is less supports this idea: the bigger the distance between each measuring point, the less its variance is influenced by its previous states. We could play around the dependency a little by sampling only a few observations out of the total amount of observations, thus making the dependency between two measurements less. We try this by sampling uniformly from the total amount of observations used, which is again the second half of all of the observations, see figure 4. Results of this test can be found in table 3.

(A) Sampling method used in table 2

(B) Sampling method used in table 3

FIGURE 4: Different sampling methods

Total observations	# Samples	$\#H_0$ rejected
500	50	170
500	40	84
500	10	48
500	5	52

TABLE 3: Amount of times H_0 was rejected, over 1000 runs

Indeed, it stands out that for a smaller amount of measurements used, the test rejects H_0 not as often and gives the expected results, which would mean that the dependency is little to negligible. For the following tests, we will use only 10 measurements, since it gives good results that are also quite trustworthy.

Another test is being conducted, this time by changing the steady state value of the continuous model to a wrong value (in this case 1.1 times too big), to get a grasp on how accurate the test is. This yields the following results:

Total observations	# Samples	$\#H_0$ rejected
500	10	446
100	10	634

TABLE 4: Amount of times H_0 was rejected, over 1000 runs

Indeed, the test gives expected results by rejecting H_0 more often, although the null hypothesis is still often accepted. We could make sure H_0 gets rejected more often by making our confidence level smaller, or by using more measurements. These options will be explored more in section 4.1. Now, we should be able to check if we can detect a fault by looking at different cases of errors in our system.

We start again at $H_0 : \mu = 0.01$ in case the system friction is twice as big as the friction of a normal working system. In this case, the steady state of the broken system is expected to be unequal to 0.01. We try again using a 95% confidence interval to see how often H_0 gets rejected.

Total observations	# Samples	$\#H_0$ rejected
500	10	1000
100	10	1000

Table 5: Amount of times H_0 was rejected for a disturbed friction, over 1000 runs

The results shown in table 5 are great results, as this means that H_0 would (almost) always get rejected if the new steady state reached by the system is twice as small by a friction that is twice as high.

Now, we test the same, but for an adjusted motor torque. As can be seen in table 6, we again get the desired results, as H_0 is always rejected in case the steady state value is twice as small as it would be normally, due to the motor torque being twice as small as normal.

Total observations	# Samples	$\#H_0$ rejected
500	10	1000
100	10	1000

Table 6: Amount of times H_0 was rejected, over 1000 runs

4.1 Power of the test

Now, we want our test to give the most trustworthy results as possible. This means that we want the test on $H_0 : \mu = 0.01$ to be accepted as often as possible when the system is behaving like intended, but also that the test is rejected as often as possible if the system is not behaving like intended. For the latter, we check the power of the test. The power of the test indicates the probability that H_0 gets rejected in the case the alternative hypothesis, H_1, is true. In this case, H_0 is assumed to be the hypothesis that the mean value of the system is equal to that of a system without failure (e.g. $\mu = 0.01$). H_1 is assumed to be the hypothesis that the steady state of the system is that of a continuous system with a friction that is twice as high or a motor torque that is twice as small (e.g. $\mu = 0.005$). In total, there are four decisions that can be made depending on the reality of two situations. These four cases can be found in figure 5. Using the power we try to find out how often a type II error takes place. A type II error is the error takes place if H_1 is true, but H_0 won't get rejected by the test. The probability of that happening corresponds to one minus the power (see figure 5).

To calculate the power of the test, we have to find the probability that H_0 is rejected in the case H_1 is true, i.e. $\mathrm{P}(H_0 \text{ is rejected}|H_1 \text{ is true})$. To calculate this chance, we first need to know what the boundaries are wherein H_0 would be accepted for a confidence level of 95%. We know that under H_0, we have a mean μ of 0.01. Furthermore, from 3, we can calculate the variance in this case, which is equal to 1.1×10^{-6} in the case $T = 0.1$. Furthermore, using our confidence level, as well as the degrees of freedom of our function, we can find a value t from our student test. In this case we use a value of $\alpha = 0.05$ and 10

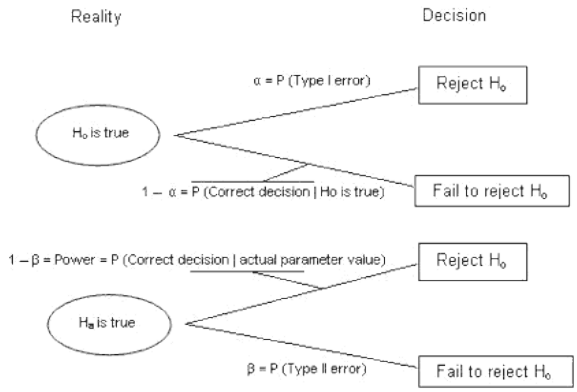

FIGURE 5: Possible errors in a statistical test [2]

measurements, giving 9 degrees of freedom.

Now, we can find boundaries for a 95% confidence interval of μ_0, using equation 15 [2]. In this equation, $\sigma_{\bar{x}}$ denotes the standard error [5], given by equation 16.

$$\begin{aligned} \text{Upper Bound} &= \mu_0 + t \cdot \sigma_{\bar{x}} \\ \text{Lower Bound} &= \mu_0 - t \cdot \sigma_{\bar{x}} \end{aligned} \tag{15}$$

$$\sigma_{\bar{x}} = \sqrt{\sigma^2/n} \tag{16}$$

Using the previously found numbers, we find that our region of acceptance of H_0 is $[0.0092, 0.0108]$. Thus, if a mean lies inside this interval, the test will accept the null hypothesis, while it would be rejected outside of this interval. Now, we need to find the probability that if we know the size of the error, how big the chance is that the measured mean falls inside that region. For that, we can simply look at the expected mean μ_1 and variance σ_1 under H_1, the alternative hypothesis, and measure the chance that a measurement in that normal distribution lies inside the region of acceptance of H_0. The equation describing this can be found in 17. In this equation, UB and LB describe the upper and lower bound of the region of acceptance, respectively.

$$\text{Power} = P\left(\frac{\text{LB} - \mu_1}{\sigma_1/\sqrt{n}}\right) + 1 - P\left(\frac{\text{UB} - \mu_1}{\sigma_1/\sqrt{n}}\right) \tag{17}$$

For example, if we take $H_1 : \mu = 0.02$, we get the result shown in table 7.

#Samples (n)	α	T	Power
10	0.05	0.01	1

TABLE 7: Power of the test with $\mu_1 = 0.02$

As expected, table 7 shows that for a large change in mean, the power is one, which means the chance of a type two error is equal to zero. Another example with $H_1 : \mu = 0.0108$ can be found in table 8.
From table 8 we again see an expected result: since we measure the test at $\mu_1 = 0.0108$, which is approximately equal to the value of the upper boundary of the region of acceptance under H_0, we see that about half of the tests are being accepted and half are being

#Samples (n)	α	T	Power
10	0.05	0.01	0.56

TABLE 8: Power of the test with $\mu_1 = 0.0108$

rejected, giving a power of approximately a half.

Now, we can test what can make our power even better for values in between the ones shown in tables 7 and 8. First, the amount of samples has been changed to different values, to see if that has any influence on the power of our test. The result can be found in figure 6.

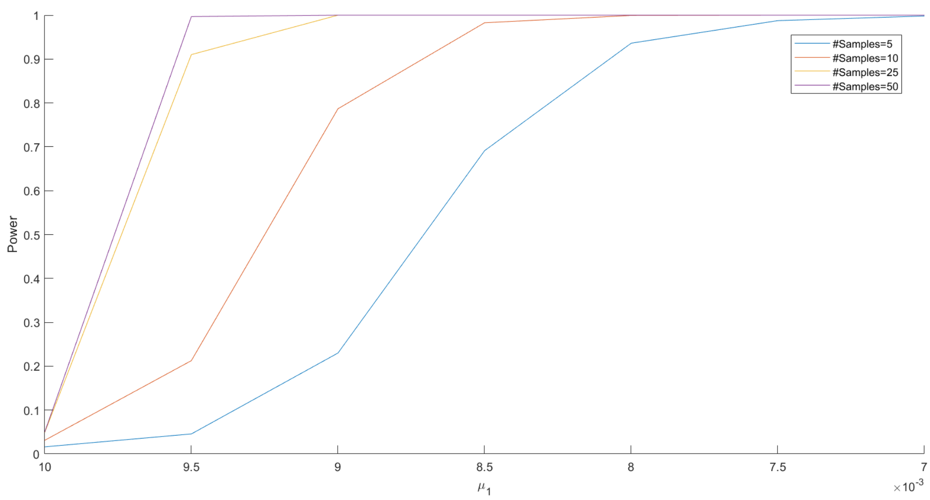

FIGURE 6: Power versus different numbers of samples, for $\alpha = 0.05$

Using more samples has a significant increase on the power of the system. For 50 samples, the power increases very rapidly, reaching a value of 1 for just a small change in mean. However, the calculation of this power can also be very time consuming, taking 10 times longer compared to using only 5 samples. Still, in case of a first order dynamical system, that does not take a lot of time. However, for a system that would be a lot more complicated, using a lot of samples is already getting quite time consuming. Especially since we need to simulate independency, it is useful to only get one sample per second, which would already mean that the system has to run for 50 seconds just so a more precise analysis can be given about the system.

Furthermore, for a large change in mean, the amount of sample does not yield a big difference in result: for example, 10 samples already have a power close to 1 for a μ smaller than 0.0085, which thus means that increasing the amount of samples would not better the results.

One other thing that is interesting to look at, is to see the influence on the power by changing the value of α to ensure a smaller region of acceptance, thus increasing the power as well. These results can be found in figure 7.

A bigger value of α does indeed have an impact on the power of the system, although it

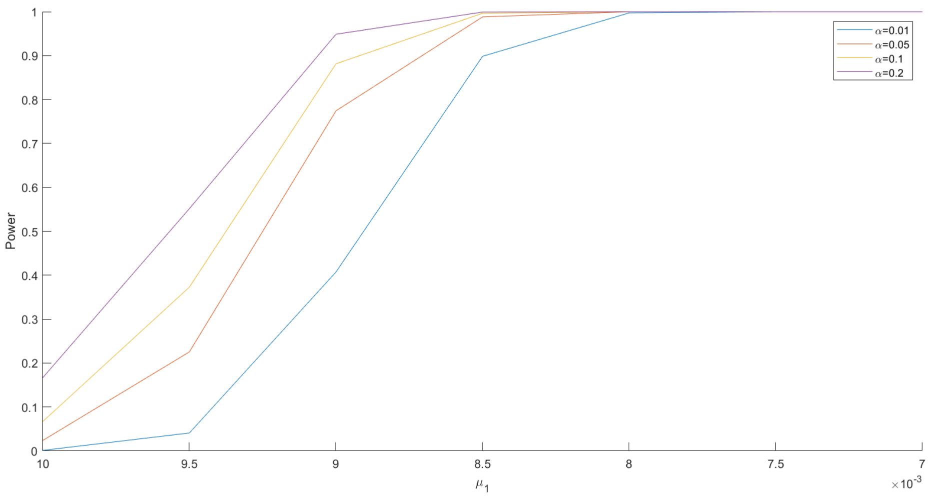

FIGURE 7: Power versus different values of α, for #samples = 25

seems to have less impact than increasing the number of samples. Other than that, the downside of increasing the value of α is that the chance of a type I error increases, which means that the test has a higher risk of rejecting the null hypothesis while it is actually true.

Choosing an optimal strategy based on the findings in figures 6 and 7 is heavily dependent on how big the error is. For a big error, the power of the test is often one, even for a small sample size. Next to that, still in case the error is big, the value for α could easily be decreased, such that the null hypothesis will less often get falsely rejected. For a smaller error however, it is probably more advisable to increase the sample size than to decrease the value of α, as increasing sample size has a bigger impact on the power.

An option for detecting an error in a dynamical system is to first test the system using a test with a low power and find out if there has been a big error, and increase the power if that does not yield any results, to detect smaller changes.

4.2 Analysis on breaking down speed

One other thing that we want to look at, is the speed with which the system adjusts to a new steady state. This may contain some clues as to whether the friction or the motor torque is failing. If we look at the formula of the continuous first order motor (equation 3), at steady state we have the following equation:

$$0 = \frac{K}{J}I - \frac{D}{J}v$$

Thus, at steady state we know that the two factors are in balance. From this we can see that if we were to make the friction twice as big, the speed will go down twice as fast as if we were to make the motor torque twice as small.

If we assume it is known when exactly the machine breaks, we can look at the speed with which it reaches a new steady state, and maybe we can conclude something out of that.

Thus we measure the distance between the last point where the system is still behaving normal and the first point after the system broke down. Again, we run the system over 5 seconds, using 100 measuring points. Over 1000 runs, this gives the results shown in table 9.

Faulty parameter	$\Delta y (\times 10^{-3})$
Friction (D)	-4.91
Motor torque (K)	-2.49

TABLE 9: Δy directly after a parameter changed in value

Indeed, the results match our expectations, the failing friction has a higher impact on the speed of the system than the motor torque does, close to twice as high.

Now, we could use this information to create a control system, where we could increase our current to something higher than $I = 0.1$, to detect at any given moment which of either the friction or the motor torque is failing in the same way we did for table 9. Note however, is that it is currently not possible with our system. As was stated in section 2, our system is described by an equation with a constant current. Thus, the system should first be changed to a second order system such that our current is variable and could thus be adjusted at any given moment. However, since we notice that adjusting our equation doesn't change the behaviour of our system [1], we let our system stay as a first order system, and say that we can just adjust our current. We do however want to raise attention that it is important to know that we would normally have to introduce more components to our system for it to be actually able to have a variable current (see [1]).

The plus side of being able to change our current is that we do not need the information of when exactly the system broke down. Thus, after a parameter has changed value and the system has settled to a new steady state, we increase our current to see how quickly it will change to another new steady state. The current is changed after two third of our simulation. So, again, we measure the distance between the first point where the system is acting according to a current $I = 0.1$ and the first point after it has been increased. Over 1000 runs for a system that runs 5 seconds using 100 measuring points, this gives the results shown in table 10.

Faulty parameter	New current (I)	$\Delta y (\times 10^{-3})$
Friction (D)	0.2	4.93
Motor torque (K)	0.2	2.51
Friction (D)	0.15	2.48
Motor torque (K)	0.15	1.25

TABLE 10: Δy directly after the current is increased

From table 10 we can conclude that it is possible to see what value changed looking at the difference between several observations. Nevertheless, we cannot use a statistical test over this, since we only use one measurement after changing the current. If we would want to use a statistical test, we could let our current constantly change, for example by making it a sine wave, and measure multiple times the difference between observations, and possibly using a test to determine if either the motor torque or friction changed.

5 Testing on the variance

Now that we know we can check for a change in mean, we want to do the same for the variance, as we want to use that to detect a change in inertia J, as well as detecting a change in friction D.

From section 3, we know that the variance of our system at a certain moment k in steady state is equal to the following:

$$Var(y_k) = \frac{1}{1 - \left(1 - \frac{D}{J}T\right)^2}\text{Var}(V_k) + \text{Var}(W_k) \tag{18}$$

From equation 18, we can see that both the friction as well as the inertia do have an influence on the variance of the system. Furthermore, what stands out, is that the size of each time step has an influence on the variance as well. Thus if we were to choose this bigger, we suspect that we could more clearly see a difference in variance. We test this first for an adjusted friction D versus a normal friction. In this case, we choose the adjusted friction to be twice as high (thus $D = 0.2$). We check the difference in variance by measuring the sample variances over 1000 runs with different values for the time step size. The total run time of the measurement is 5 seconds. The results are given in table 11.

Time step size (T)	#Samples	Adjusted Parameter	Variance (s^2)
0.05	10	None	$1.13 \cdot 10^{-6}$
0.05	10	Friction (D)	$1.10 \cdot 10^{-6}$
0.1	10	None	$1.09 \cdot 10^{-6}$
0.1	10	Friction (D)	$3.33 \cdot 10^{-5}$

Table 11: Measured sample variance for system with and without adjusted friction, measured over 1000 runs

From table 11, we can see that a change in parameter value does indeed change the variance of the system. It does seem like for a time step size of 0.1, the change in variance is more noticeable than it is for a time step size of 0.05. This makes sense looking at equation 18: all numbers are the same as if we were to calculate them theoretically. The most interesting case is the one where the time step size T is equal to 0.1 and the friction D is adjusted. Using equation 18, we would get that $\left(1 - \frac{D}{J}T\right) = -1$. Thus, we would have a denominator of zero, meaning our variance would become gigantic. Indeed, the measured variance is of order 10 times as high as we got from other measurements.

Another two interesting measurements are for $T = 0.05$ together with an adjusted friction, as well as $T = 0.1$ without one. Both of these have a variance of about 1.10×10^{-6}. If we were to use equation 18, we would also see that $\left(1 - \frac{D}{J}T\right) = 0$ for both of these cases. This means that the parameters are at such a state that they have no influence on the variance of the system, which would mean that the variance of the system would be exactly equal to the sum of the dynamical error and the observer based error. Furthermore, this means that if any parameter would be adjusted in this case, we would get a variance that is higher.

Therefore, we can say that the test we use on the sample variance of the system could be both a two-sided test, as well as a one-sided test in case we choose our values such that in a standard case, our variance is equal to the minimum of 1.1×10^{-6}. For this, we want to use $T = 0.1$ (see table 11). Thus, we get our null hypothesis to be $H_0 : \sigma^2 = 1.1 \times 10^{-6}$ against the hypothesis $H_1 : \sigma^2 > 1.1 \times 10^{-6}$. Like said in section 3, this is possible using a chi-square test on variance. We first test to see whether it works for our normal working system. We do this by running our system 1000 times, generating 100 measurements over 10 seconds (thus having $T = 0.1$), for normal values of friction. The samples used are again chosen uniformly from the second half of the observations. The results of these runs with a confidence level of 95% can be found in table 12.

Friction (D)	# Samples	# H_0 rejected
0.1	10	47

Table 12: Amount of times H_0 was rejected, over 1000 runs

Indeed, we notice that H_0 is rejected about 5% of the time, thus we can say with confidence that we have the right null hypothesis. Now we can test how often H_0 gets rejected for a friction that is higher. Using the same specifications as in the previous test, we get, with a confidence level of 95%, the results shown in table 13.

Friction	# Samples	# H_0 rejected
0.2	10	1000
0.19	10	247
0.18	10	121
0.15	10	75

Table 13: Amount of times H_0 was rejected, over 1000 runs

From table 13 we can see that indeed H_0 gets rejected always in the case the friction is twice as high. However, if the friction value is between 0.1 and 0.2, we get a lot of type two errors, seeing that H_0 gets rejected much less than it should be. Most likely, this is due to the fact that the change in variance is too little for our test to pick up. One thing that makes a big contribution to this, is the observer based error W_k. Since it is 10 times bigger compared to the dynamical error V_k, a small change in V_k does not have a big influence on the total variance of the system. Therefore, we try the test once more, but this time we subtract the observer based variance from our total variance. This means our null hypothesis changes to $H_0 : \sigma^2 = 10^{-7}$ against $H_1 : \sigma^2 > 10^{-7}$. The results are shown in table 14.

Friction	# Samples	Average variance	# H_0 rejected
0.19	10	6.6×10^{-7}	954
0.18	10	3.0×10^{-7}	748
0.15	10	1.4×10^{-7}	198
0.12	10	1.0×10^{-7}	70

Table 14: Amount of times H_0 was rejected, over 1000 runs

From table 14, we do notice an improvement in detecting a change in variance when there is no observer based error W_k. Nevertheless, for a small change in friction, H_0 still gets rejected many more times than it should be. Again, we could use trade offs like used in section 4.1 to hopefully get better results in this case.

5.1 Power of the test

Just like the t-test used on the mean of the system, a power can be calculated for the chi square test as well. To do so, we start again by calculating an interval for our region of acceptance for $H_0 : \sigma^2 = \sigma_0^2$. In this case, σ_0^2 denotes our theoretical variance. At first, we again look at $\sigma_0^2 = 1.1 \times 10^{-6}$, the theoretical variance in case are parameter values are normal and our time step size $T = 0.1$. We again decide to use a one-sided test, i.e. $H_1 : \sigma^2 > \sigma_0^2$. To find a region of acceptance, we use an inverse chi square test to find a critical value, with $\alpha = 0.05$ and 9 degrees of freedom. This gives a critical value x_{crit} of 16.92. Using this critical value, we can find a value for a boundary like shown in equation 19 [5].

$$\text{Upper Bound} = \frac{x_{crit}}{n-1}\sigma_0^2 \tag{19}$$

Using equation 19, we find that our upper bound lies at 2.07×10^{-6}. This explains why the results found in table 13 were bad. Since the upper bound of the variance test lies so high, the test still often decides to accept the null hypothesis. If we would not have an observer based error W_k, we would get an upper bound of 1.88×10^{-7}, which is already way better, but since a change in parameter has no big influence on the variance of the system, it would still often be rejected for small changes in variance, as can be seen from table 14 as well.

Now, to calculate the power of the chi square test, we define δ as the ratio between the two (theoretical) variances for both hypotheses, i.e., $\delta = \sigma_0^2/\sigma_1^2$. Using this ratio, we find that the power can be calculated by using the formula given in equation 20[8].

$$\text{Power} = 1 - \chi^2(x_{crit} \cdot \delta, n-1) \tag{20}$$

Using this, we can again look at the influence of changing the confidence interval or the amount of samples used for the test. We do this on a system where $\sigma_0^2 = 1.1 \times 10^{-6}$ Results of this can be found in figure 8.

As can be seen from figure 8, for a very low change in variance, the power is very low. Increasing the amount of samples does not help: only for a high change in variance it proves to be helpful. Increasing α to a very high number does help, but it also largely increases the chance of a type I error.

As said before, the main problem in measuring the variance is that the observer based error W_k is way higher than the dynamical error V_k. We can again subtract the observer based error from the measured variance as well as the theoretical variance, to get rid of that problem a bit. If we were to try to find the power for this system, we get the results shown in figure 9.

From figure 9, we can already see better results: the power reaches 1 pretty fast for 50 samples. The value of α in this case does not really make a big difference: as soon as the σ_1^2 is 3 times bigger than σ_0^2, the power is approximately one.

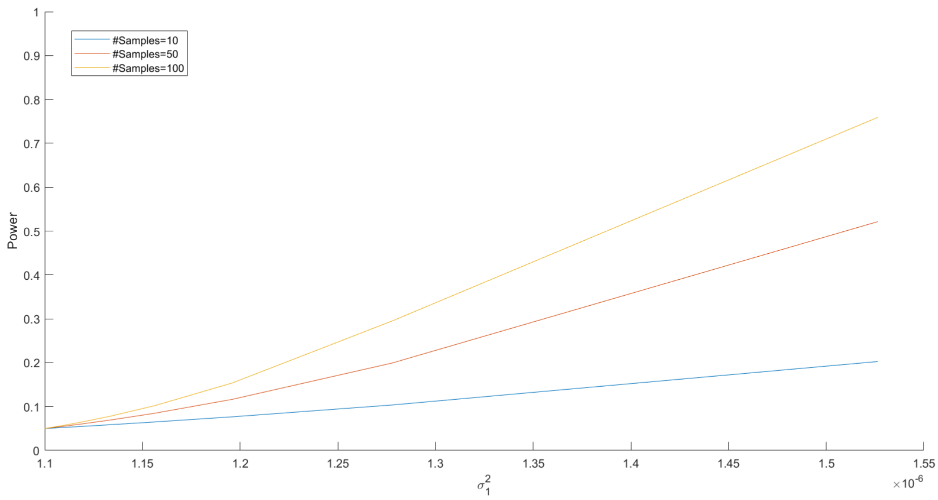

(A) Power versus different numbers of samples, for $\alpha = 0.05$

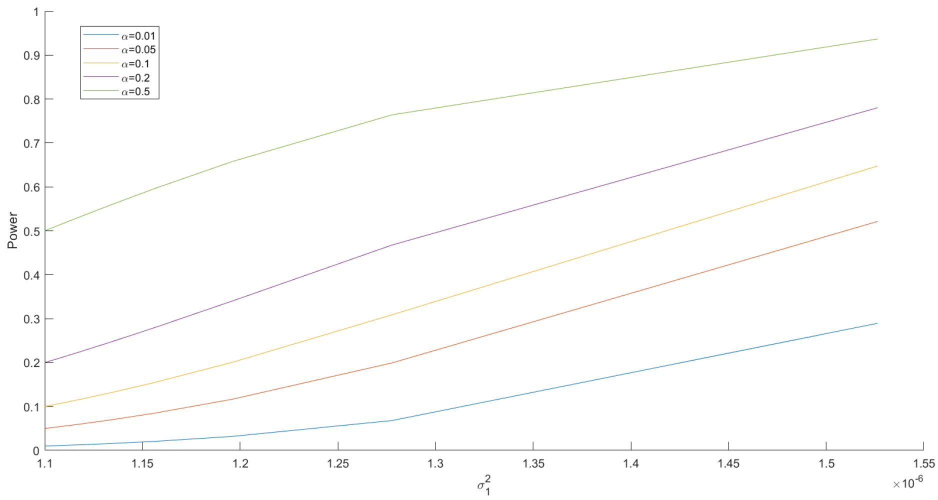

(B) Power versus different values of α, for #samples = 50

FIGURE 8: Power over different values of σ_1^2, with $\sigma_0^2 = 1.1 \times 10^{-6}$

6 Probability of detecting the right error

Now that we know how well the results of both the tests on the mean and variance of the system are, we can find a final overview of the chance of detecting an error in our system. We know, using the power of the system, what exactly those are, and how they could be improved. Therefore, we can give a total overview with probabilities on detecting an error. Like stated before, the probability of detecting an error does depend on the size of the error. Thus, we will analyse two situation: one where the error is relatively big and one where the error is relatively small. The values for the sample size and α are based on our previously found findings from figures 6, 7, 8 and 9. Furthermore, we will translate these errors in mean and variance into parameter values, using our findings in section 3, concluded in table 1.

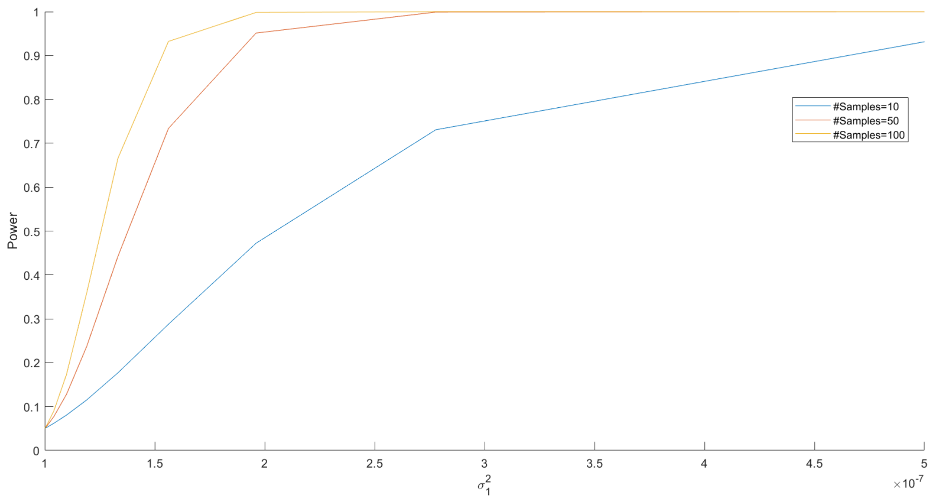

(A) Power versus different numbers of samples, for $\alpha = 0.05$

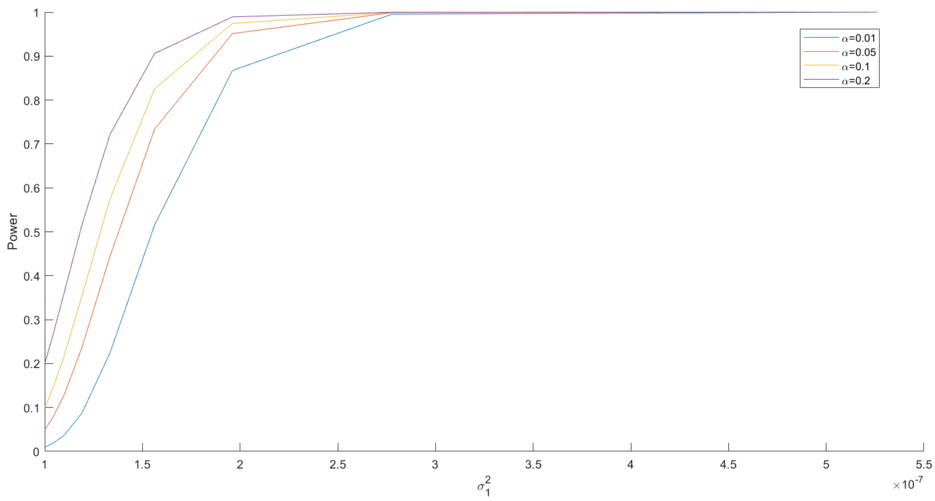

(B) Power versus different values of α, for #samples = 50

FIGURE 9: Power over different values of σ_1^2, with $\sigma_0^2 = 10^{-7}$

6.1 Big error

For a big change in system behaviour, our main concern is detecting an error in the variance, as it has been shown that detecting a big change in mean is very likely to go right. We assume that in this case our alternative mean μ_1 is twice as low as μ_0, while our alternative variance σ_1^2 is twice as high as σ_0^2. Furthermore, we use $T = 0.1$, giving again a variance σ_0^2 of 1.1×10^{-6} in case the parameter values are chosen such as in appendix A. As we have discovered that testing with variance is very difficult if the observer based variance W_k is not subtracted from the total variance, we again use that to change our test to only reflect the dynamical error, i.e., $\sigma_0^2 = 1 \times 10^{-7}$.

Next to that, we again have $\mu_0 = 0.01$. From section 4.1, we discovered that the power for this test is approximately one, even for a low value of α and a low amount of samples. Thus, we choose our α to be 0.01, making the probability of both a type I and type II error 0.01 as well. On the other hand, from our test on the variance and figure 9, we find that if we use 25 samples, together with a value of $\alpha = 0.12$, we get a power of 0.88. Thus, we get a probability of a type I and type II error to be 0.12. In total, we can say with the probabilities found in table 15 that the outcome of the tests is indeed the right one.

	No adjustments in μ $\mu = 0.01$	Adjusted μ $\mu \neq 0.01$
No adjustments in σ^2 $\sigma^2 = 1 \times 10^{-7}$	0.87	0.87
Adjusted σ^2 $\sigma^2 > 1 \times 10^{-7}$	0.87	0.87

TABLE 15: Probability of detecting the right error for big adjustments

From table 15, we can conclude that big adjustments in our system give trustworthy results. Translating this into parameter values, this would mean that the motor torque K would be twice as low than normal and the friction D would be twice as high, looking at the mean of the system. Looking at the variance, using equation 18 we see that that these measurements of variance are possible if our inertia J would be approximately 3.4 times as high, or our friction D would be 3.4 times as low. All of these adjustments are pretty large, but the measurements are quite certain than as well.

6.2 Small error

To see if our system is able to pick up small changes almost as reliably, we look at a certain case for that as well. In this case, under the alternative hypotheses, both the mean and variance are shifted only slightly of their original value. Here, we use that under the alternative hypothesis on the mean, $\mu = 0.009$ and that under the alternative hypothesis on the variance, $\sigma^2 = 1.3 \times 10^{-7}$ (once again subtracting the observer based variance W_k first). If we again want similar results to table 15, that is, having an equal probability on detecting each kind of mistake, we need to choose our α and sample size in such a way that the chance on getting a type I and type II error is equal for both tests.

For the t-test on the mean, we can see from figure 6 that for a sample size of again at least 25, we again have a very high power approximately equal to 1. Thus, we choose to lower the value of α as well, making it 0.02, making the chance of both a type I and type II error to be 0.02. However, for the variance, we can see from figure 9 that the power is not that high for a $sigma_1^2 = 1.3 \times 10^{-7}$. Since we use a sample size of 25, we discover that a value of $\alpha = 0.33$ gives an equal probability on a type I and type II error, which is then equal to 0.33 as well. This gives the final probabilities that can be found in table 16.

	No adjustments in μ $\mu = 0.01$	Adjusted μ $\mu \neq 0.01$
No adjustments in σ^2 $\sigma^2 = 1 \times 10^{-7}$	0.66	0.66
Adjusted σ^2 $\sigma^2 > 1 \times 10^{-7}$	0.66	0.66

TABLE 16: Probability of detecting the right error for small adjustments

From table 16, we can notice that detecting an error with small adjustments in mean and variance is still able to be picked up, although approximately 2 out of 3 times the tests will find the right error. In terms of parameter adjustments, a mean of 0.009 can be reached by decreasing the motor torque K by 0.001 or by increasing the friction D by the same value. In terms of variance, a variance of 1.3×10^{-7} still needs high adjustments to be reached: either the inertia J is doubled or the friction D is halved. This means that both of these values still need to change a lot to be able to detect them, with still a probability of only 67% to be certain.

7 Discussion

The results show that it is possible to detect a parameter change in a simple dynamical system using statistical methods. For a big change in parameter value this was easily measurable with a high certainty, but for a smaller change, especially a small change in variance, the test can quickly become less reliable. A possible idea to get better results is to amplify the variance in some way.

Furthermore, the power of the test was analysed and used to decrease the probability on getting a type I or type II error. We also touched upon the trade-off in changing the value of α, as well as the value of the sample size.

Analysis on the speed with which the system adjusts to a new steady state does help confirm what parameter value has been adjusted. No statistical analysis has been done on this, but that is definitely worth considering looking into in the future.

Finally it is worth saying that the statistical analysis on our system could be done on higher order systems as well, though it needs to be verified if this shows similar results.

8 Conclusion

In this paper we have discussed a number of methods to statistically check a change in parameter value in a simple linear dynamical system. Using basic tests on the mean and the variance of the system we could establish a difference between multiple cases of a broken system: one where the mean value changes, one where the variance changes and one where both change.

Using a two-sided t-test we could test on the mean of our system, results of which were pretty reliable. Using a chi square test we could test on the variance of our system. By utilising our time step size T, we could improve the reliability of these results, changing our two-sided test into a one-sided.

The power of the tests have been calculated and analysed. The power of the test is very high for a relatively small adjustment on the mean, which means it is easier to detect with certainty what error has occurred. However, on small adjustments in variance, it could still use improvement.

Finally, an analysis was done on the speed with which the system broke down to establish a difference between the fault a friction causes as oppose to the fault the motor torque causes. Here, we could see a clear difference in speed for the first time step after a parameter changed value. Using this information, we could build a control system to check at any given moment if either the motor torque or friction has changed value, by increasing the current in our system.

References

[1] http://ctms.engin.umich.edu/ctms/index.php?example=motorspeed&section=systemmodeling.

[2] https://apcentral.collegeboard.org/courses/ap-statistics/classroom-resources/power-in-tests-of-significance.

[3] E. Alcorta García and P.M. Frank. Deterministic nonlinear observer-based approaches to fault diagnosis: A survey. *Control Engineering Practice*, 5(5):663 – 670, 1997.

[4] Zoubin Ghahramani and Geoffrey E. Hinton. Parameter estimation for linear dynamical systems. 09 1996.

[5] D. Meijer. *Statistics, reader 626*, January 2016.

[6] E. A. Wan and R. Van Der Merwe. The unscented kalman filter for nonlinear estimation. In *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373)*, pages 153–158, Oct 2000.

[7] Alan S Willsky. A survey of design methods for failure detection in dynamic systems. *Automatica*, 12(6):601–611, 1976.

[8] Charles Zaiontz. Power of one sample variance testing.

Appendix

Appendix A: Parameter values

Parameter	Value
Motor torque constant (K)	0.01
Moment of inertia (J)	0.01
Friction (D)	0.1
Variance process noise (P)	10^{-7}
Variance observation noise (Q)	10^{-6}
Initial velocity (v)	0
Input current (I)	0.1