Data-based Optimization Assignment 1

Tobias Wejbora - 1785486

February 2026

Question 1:

The plot of the steady-state performance mapping $Q_J(x)$ over the domain $0 < x < 240\,\text{rad}$ can be seen below:

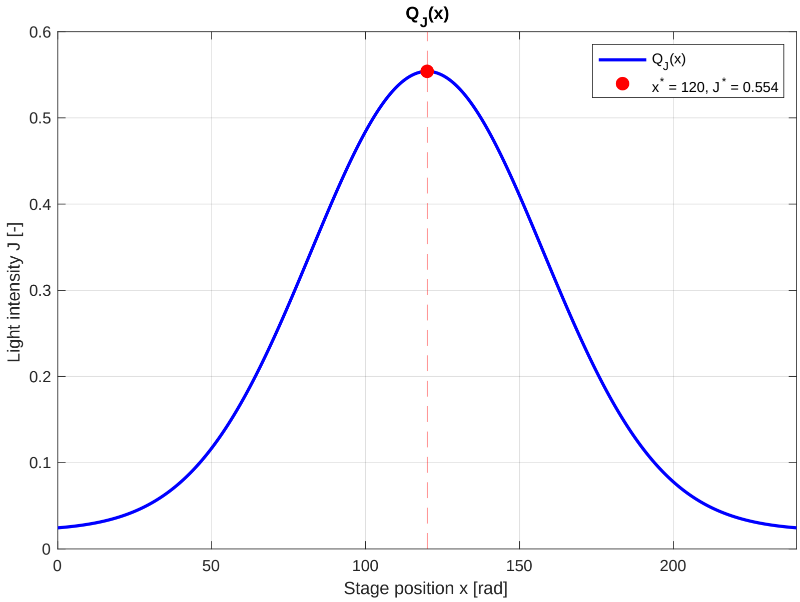

Figure 1: Performance mapping $Q_J(x)$ with global maximum at $x^* = 120\,\text{rad}$ ($J^* = 0.554$).

The assumption for global optimum convergence is that the $Q_J(x)$ must be unimodal(one local maximum) over the operating domain, making it also the global maximum. Otherwise the ESC would follow the gradient and most likely get trapped at some local optimum, depending on the initial condition, if $Q_J(x)$ was multimodal. The gradient $Q'_J(x)$ shows us that there is one x position where the gradient is zero, being at $x = 120\,\text{rad}$. Thus, Q_J possesses a single extremum in the domain

$$Q'_J(x) = -2 \cdot 3.5 \times 10^{-4} \cdot 0.533\,(x - 120)\,\exp\Big(-3.5 \times 10^{-4}(x - 120)^2\Big) \tag{1}$$

The gradient can also be plotted.

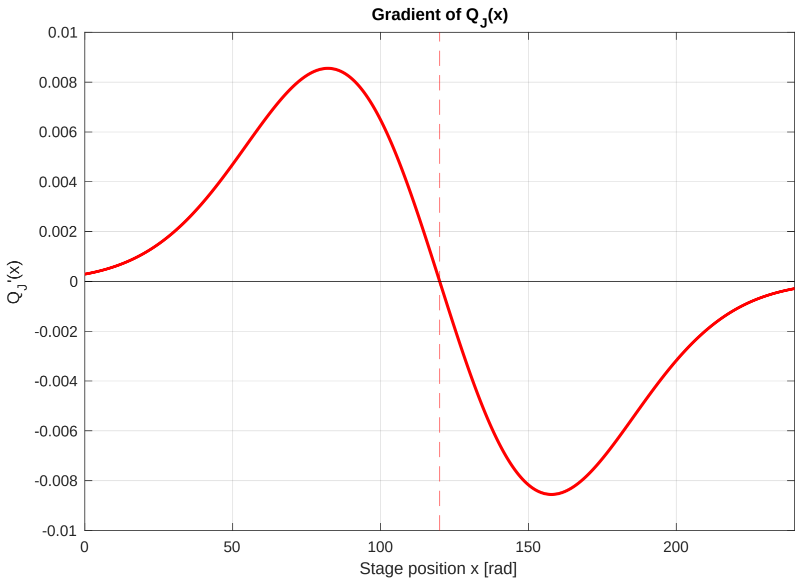

Figure 2: Gradient $Q'_J(x)$ with a single zero crossing at x^*.

Since $Q_J(x)$ is a Gaussian function, the double derivative at $Q''_J(120) < 0$ meaning it is strictly concave near the peak and as $|x - 120|$ increases, $Q_J(x)$ will strictly decrease. There is only one critical point in the entire domain. This means that Q_J is unimodal and the assumption is satisfied. The ESC will converge to the global optimum $x^* = 120\,\text{rad}$ no matter the initial condition x_0.

Question 2:

The ESC feedback loop was implemented in the Simulink model as seen in fig. 3. First a high pass filter of the form $H_{\text{HP}}(s) = \frac{s}{s+\omega_{\text{HP}}}$, then the demodulation signal being a cosine wave with amplitude $\frac{2}{a}$ and frequency ω, multiplied with the high-pass filtered intensity signal. Next the low-pass $H_{\text{LP}}(s) = \frac{\omega_{\text{LP}}}{s+\omega_{\text{LP}}}$, then an integrator with gain (k/s) with initial condition $x_0 = 10\,\text{rad}$. Then the dither signal being a cosine wave with amplitude a and frequency $\omega = 2\pi \cdot 1\,\text{Hz}$ added to $\hat{x}$.

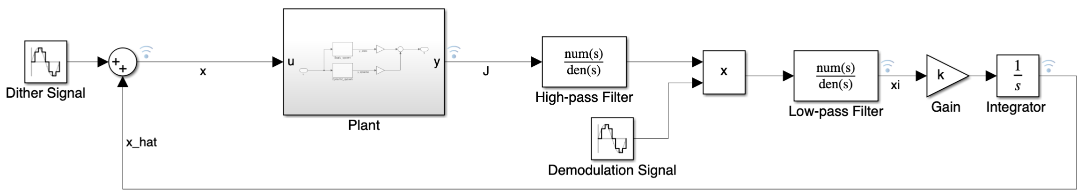

Figure 3: Implemented ESC loop in Simulink

Firstly, the filter cutoff frequency ω_f was tuned. The filters must pass the dither-frequency content while rejecting everything else, so $\omega_f \ll \omega$. The first value tried was $\omega_f = 0.5\,\text{rad/s}$, which can be tuned by looking at the gradient estimate ξ. If it was noisy at steady state, ω_f is too high, and if there is slow convergence even with large enough k then ω_f is too low. After testing, ω_f was set to $1.95\,\text{rad/s}$. Next the dither amplitude was tuned, starting at $a = 2\,\text{rad}$ going up to $a = 20\,\text{rad}$. While tuning, its important to look at how the steady-state is reacting and seeing if the oscillation around the optimum were too large since it has to have $0.1\,\text{rad}$ accuracy. Lower values of a were ideal and reducing a to $4\,\text{rad}$ made the oscillation smaller while ξ was still large enough for the ramp up of the objective function so the ESC can make it to the peak. Lastly, the integrator gain k was tuned. It was best to start on the low side at $k = 500$ and to increase it gradually. At low values ($k = 500$–2000) the ESC converged but it just took a long time. As k increased, the ESC converged faster but this also introduced overshoot before settling back to the optimum. At the end, $k = 5000$ was a good balance as it converged in about 17.3 s and the steady-state error was below 0.1 rad. Figure 4 shows the ESC convergence with the optimal parameters.

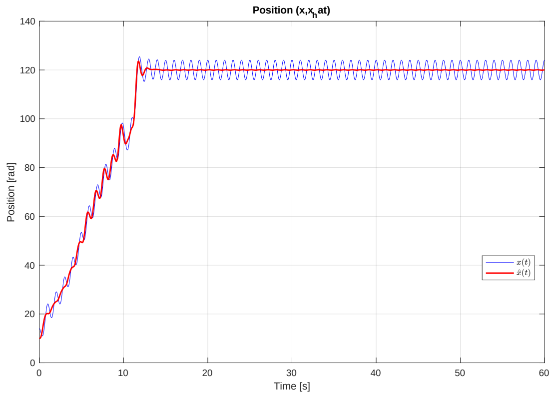

(a) Stage position $x(t)$ and $\hat{x}(t)$.

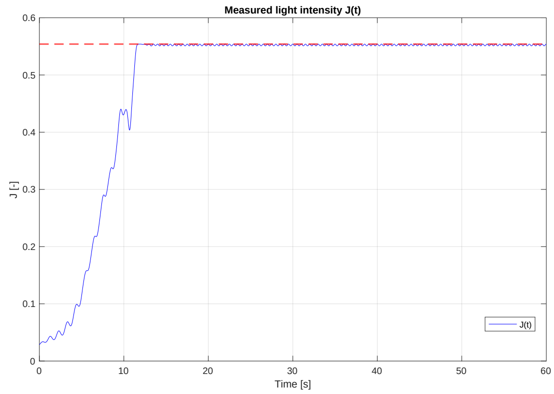

(b) Measured intensity $J(t)$.

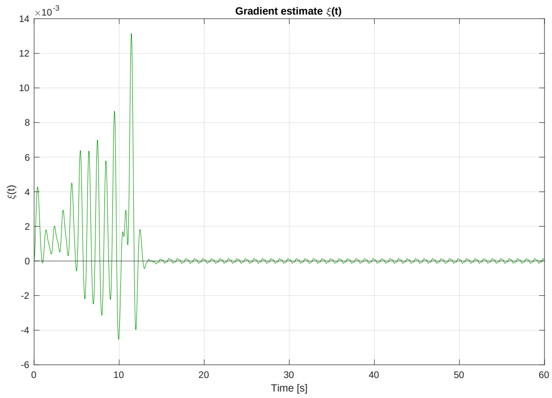

(c) Gradient estimate $\xi(t)$.

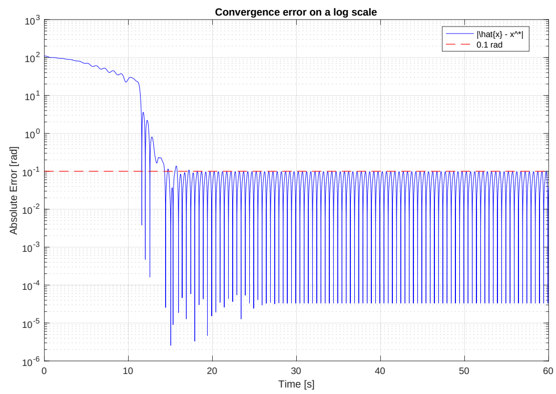

(d) Convergence error on a log scale.

Figure 4: ESC convergence with the optimal parameters.

So to conclude, the optimal parameters were $k = 5000$, $a = 4$ rad, $\omega_f = 1.95$ rad/s, achieving convergence in approximately 17.3 s with a steady-state error below 0.1 rad.

Question 3:

In the Simulink file a 1-D lookup table block was used to run the measured objective function, as shown in Figure 5.

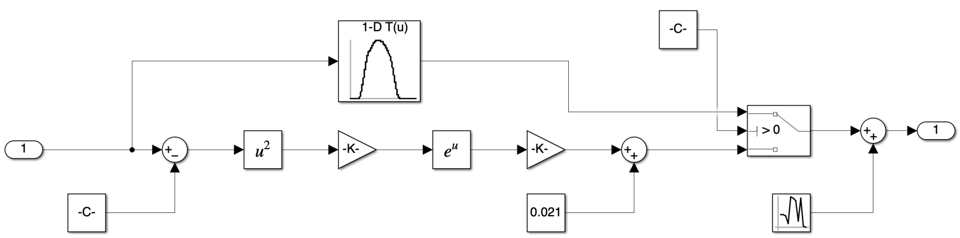

Figure 5: Implemented lookup table block scheme in static system.

In Figure 6(6a), the experimental measurement done on the physical setup is plotted with the analytical $Q_J(x)$. The measured data does for the most part have the Gaussian shape that the analytical function does, however, the experimental data seems to have an offset to the left, meaning the peak position is also more to the left, at around 98 rad. Also, the slope of the experimental $Q_J(x)$ is larger meaning that more light intensity is seen earlier on. A lookup table was made from the experimental data so that the discrete encoder mapping is transformed into a continuous objective function. Through this, values can be interpolated by the simulink block. First, the measured intensity values are binned into position intervals of $\Delta x = 0.5$ rad and then the average within each bin is found. Figure 6(6b) compares the ESC convergence using the analytical Q_J and the measured Q_J, both using the Q2 tuning parameters. The analytical function obviously converges to $x^* = 120$ rad as expected. However, the measured objective function does not converge. The ESC initially moves toward the measured peak near $x \approx 98$ rad, but at $x \approx 67$ rad it diverges and goes to a negative position around -60 rad. The main issue seems to be with the high filter cutoff frequency $\omega_f = 1.95$ rad/s. This is pretty fast and here the gradient estimate ξ reacts quickly to the steep slope of the measured objective function, but the filters also have phase lag. By the time the filters realize that $\hat{x}$ is approaching the peak, it's already too late and it starts reversing. The gradient turns negative and brings $\hat{x}$ negative where the lookup table has no gradient information to save it. Also the small dither amplitude $a = 4$ rad makes it more sensitive to sharp local slopes in the objective function.

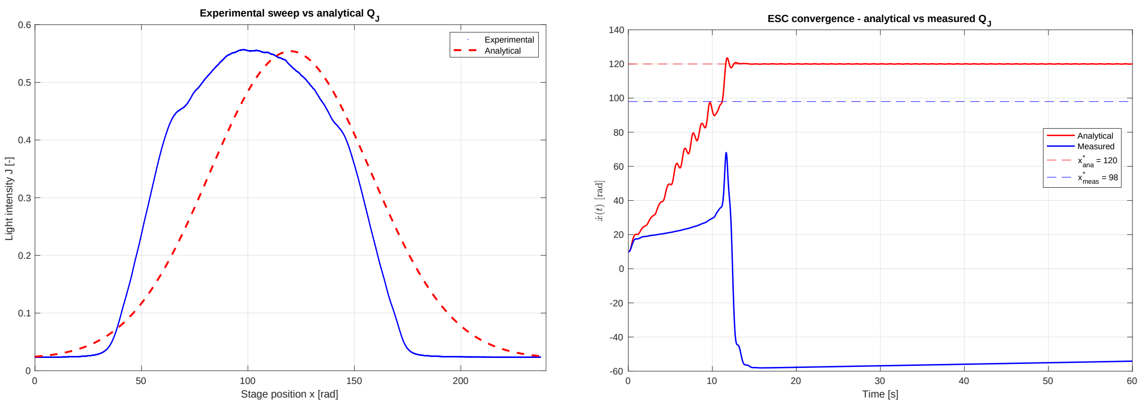

(a) Experimental sweep vs analytical $Q_J(x)$.

(b) ESC convergence: analytical vs measured Q_J.

Figure 6: Q3: Measured vs analytical objective function and ESC convergence comparison.

Question 4:

Firstly, for this question a baseline was set by first running the parameters from Q2 parameters ($k = 5000$, $a = 4$, $\omega_f = 1.95$) with disturbances enabled. This immediately showed that the steady-state error was above 0.1 rad. Looking at the gradient estimate, the measurement noise seems to gets amplified and looking at the high-pass filter and it being multiplied with the dither signal, it most likely picks up noise components at or near the dither frequency. Starting from the Q2 values, I first tried simply reducing k to lower the noise amplification. This helped the steady-state error but made convergence very slow. I then increased the dither amplitude from $a = 4$ to $a = 6$, thinking that a larger dither signal would make gradient be more robust against the noise. This did work so with $a = 6$, $k = 5000$ could be kept and the requirement of 0.1 rad was still met. After trying different values of ω_f, $\omega_f = 1.95$ still gave the best convergence speed which is the same value from question 2. Thus the final parameters are

$$k = 5000, \quad a = 6\,\text{rad}, \quad \omega_f = 1.95\,\text{rad/s} \tag{2}$$

These values got a convergence time of 17.3 s and a steady-state error of 0.086 rad. The timescale separation condition is also met since $\omega_f = 1.95 \ll \omega = 2\pi \approx 6.28$ rad/s. Tuning for the measured Q_J was harder. When trying the analytical parameters first it make the system unstable. This can be caused by the different shape of the measured Q_J since it has a steeper ascent and decent. Reducing ω_f was key to stabilising the system since a lower filter cut-off helps smooth out the gradient spikes before they reach the integrator. After trying different combinations:

$$k = 5250, \quad a = 7\,\text{rad}, \quad \omega_f = 1.01\,\text{rad/s} \tag{3}$$

This got a a convergence time of 15.9 s and a steady-state error of 0.054 rad, converging exactly to the peak intensity. The timescale separation is also met($\omega_f = 1.01 \ll \omega \approx 6.28\,\text{rad/s}$). It was also noticed that both k and a could be slightly increased compared to the analytical case ($k = 5250$ vs 5000, $a = 7$ vs 6). The larger dither helps averages the gradient which smooths the steep accent and thus giving the ESC a better gradient estimate. Table 1 compares the the analytical vs the measured objective function. Interestingly, the measured case converges slightly faster even though it has a more difficult objective function shape. This could be mainly because the distance from the initial condition to the optimum is shorter: $x_0 = 10 \to x^* = 98$. However, both cases still meet the 0.1 rad steady-state requirement. There was a clear tuning trade-off, being convergence speed

Table 1: Q4: Comparison of ESC parameters and performance with disturbances.

	Analytical Q_J	Measured Q_J (LUT)
k	5000	5250
a [rad]	6	7
ω_f [rad/s]	1.95	1.01
x^* [rad]	120	98.0
Convergence time [s]	17.3	15.9
Steady-state error [rad]	0.086	0.054
J at steady state	0.554	0.557

and steady-state accuracy. There were 3 main things that could be changed in the tuning process and they each affected the system in different ways.

- The filter cut-offs have an effect since increasing ω_f makes the ESC respond faster but lets more noise through to the integrator. So more steady-state oscillations. For example, when tuning the measured objective function, increasing ω_f from about 0.5 to 1.0 rad/s almost halved the convergence time, but steady-state error doubled.
- Also the integrator gain changed things since a higher k speeds up convergence but amplifies the signals reaching the integrator, including noise. Through trial and error going over $k \approx 7000$ in the analytical objective function caused divergence because of the noise.
- Lastly, the dither amplitude: increasing a gives a better signal to noise ratio clearly seen in the gradient estimate giving faster convergence. But, it also increases the oscillation of $x(t)$ around $\hat{x}(t)$ at steady state caused by the dither.

Question 5:

To get the ESC to work, the dynamic system G has to act like a static map at the frequencies we are operating at. So G is just the static performance map Q_J stacked on top of the closed-loop tracking system $T(s)$(reference r to the actual stage position x). We need $|T(0)| = 1$ so that the stage actually ends up at steady state ($x = r$). Also the closed-loop bandwidth (ω_{BW}) needs to be way faster than our dither frequency ω. If $\omega_{\text{BW}} \gg \omega$, the dither signal goes right through the system without losing amplitude ($|T(j\omega)| \approx 1$) which is what we want. When those points are met,$x(t) \approx r(t)$ at the frequencies we care about. Then the ESC algorithm just sees $J \approx Q_J(r)$, which is exactly the static map assumption we need for the gradient estimation to work. This can be verified with matlab, using the `dcgain(T)` command, the result is exactly 1.0000 (because there is an integrator in the plant $P(s)$).Also using the `bandwidth(T)` command, the system gave 67.3 rad/s or 10.7 Hz. At the 1 Hz dither frequency ($\omega \approx 6.28\,\text{rad/s}$), the `bode(T, w)` command was used to get the exact values, which gave a magnitude of -0.66 dB and a phase of $-9.3°$. Since 1 Hz is well within the 10.7 Hz bandwidth, we can say that the system basically behaves statically

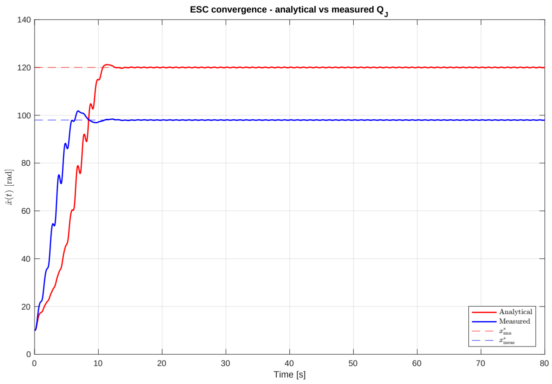

(a) $\hat{x}(t)$: analytical vs measured Q_J.

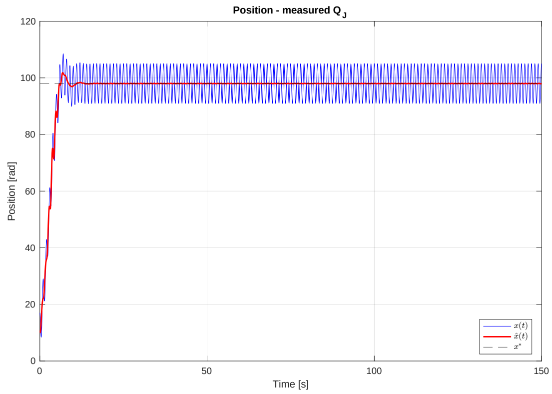

(b) Stage position $x(t)$ and $\hat{x}(t)$ (measured Q_J).

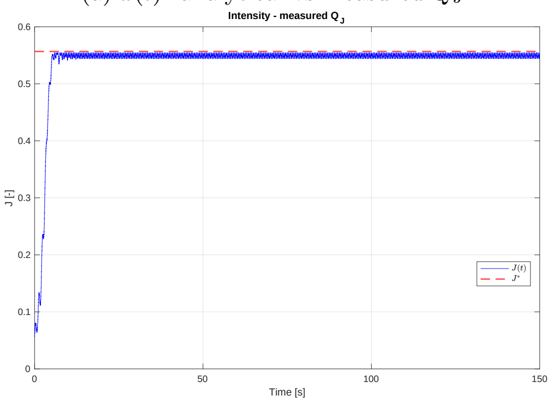

(c) Noisy intensity $J(t)$ (measured Q_J).

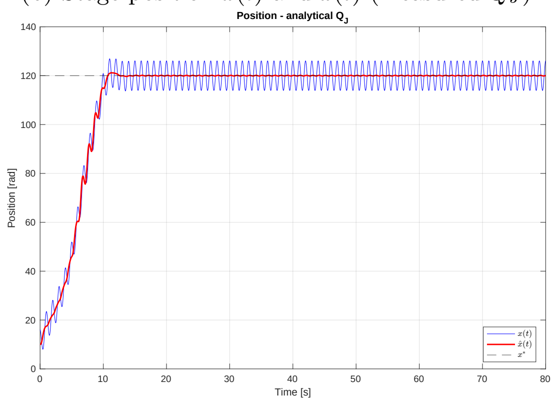

(d) Stage position $x(t)$ and $\hat{x}(t)$ (analytical Q_J).

Figure 7: Q4: ESC convergence comparison between analytical and measured objective function.

for the ESC. There was a note in the assignment regarding the discrete time implementation. Here an assumption was made that the frequencies we care about are so small (like our 1 Hz dither) compared to the sample rate of 4 kHz limit, the digital steps basically look continuous anyway.

Question 6:

Figure 8 shows the Bode plot of $T(s)$. Looking at below the bandwidth frequency, the magnitude is approximately flat at 0 dB and the phase is basically 0° .There is also a resonance peak at around 10^4 rad/s, which is the mechanical dynamics of the plant. Based on the bode plot, I chose to stick with the dither frequency of $f = 1\,\text{Hz}$ ($\omega \approx 6.28\,\text{rad/s}$) from question 5. There are multiple reasons why this dither frequency is fine. Firstly, at 1 Hz, we have $|T(j\omega)| = -0.66\,\text{dB} \approx 0\,\text{dB}$. The dither passes through the closed loop with not much attenuation ($|T| = 0.93$, so only a 7% loss). Also the phase lag is $\angle T(j\omega) = -9.3^\circ \approx 0^\circ$. Which is quite small so the demodulation signal in phase with the plant. Lastly, $\omega \ll \omega_{\text{BW}}$, so the bandwidth is about 10.7 times larger than the dither. This creates good margin to assume the system acts like a static map. We can choose a higher dither frequency (for example, $f = 5\,\text{Hz}$, where $|T| = -2.1\,\text{dB}$ and $\angle T = -19.4^\circ$) and this dither frequency would technically increase the ESC convergence speed, but they would bring more attenuation and phase lag. This would thus risk messing up the gradient estimate.

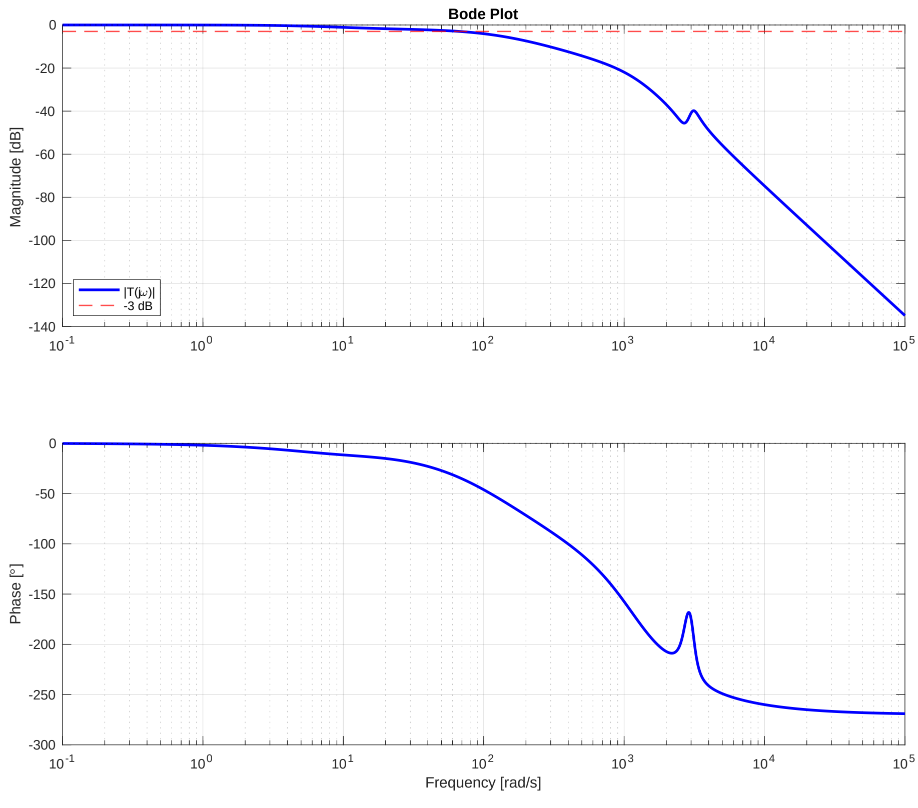

Figure 8: Bode plot of the closed-loop transfer function $T(s) = x(s)/r(s)$.

Question 7:

Firstly in this question `disturbances = true` and `dynamics = true`. The dither frequency is used from Q6 ($\omega = 2\pi f_{\text{Q6}}$). The dither frequency from Q6 will be used since it makes sure that $|T(j\omega)| \approx 1$ and $\angle T(j\omega) \approx 0^\circ$ and we will then have a good tracking at the dither frequency for the static-map assumption. The values were re-calibrated using the same approach as in Q2. The optimal parameters found are:

$$k = 3000, \quad a = 7.5\,\text{rad}, \quad \omega_f = 1.8\,\text{rad/s}, \quad \omega = 2\pi(1) \approx 6.3\,\text{rad/s} \tag{4}$$

These parameters make the ESC converge in approximately 20.8 seconds. When the dynamic plant is enabled, the closed-loop transfer function $T(s)$ introduces new effects compared to the static case. These are the following:

- Even though we picked a dither frequency in the flat part of the bode plot, the physical system still delays the signal a tiny bit. This phase lag makes the demodulation signal delayed with the actual light intensity changes. So since they aren't perfectly lined up, our gradient estimate ends up being weaker.
- If the dither frequency is too close to the bandwidth frequency then the physical stage just can't keep up and the dither modulation becomes less strong and more pointless. Therefore, the stage isn't moving as far as the reference told it to go and the light intensity doesn't change as much. This makes the signal compress compared to the noise, which ruins the gradient estimate.
- The plant dynamics can also be prone to excitation by for example if large steps are taken in $\hat{x}$, which could cause some vibration or resonance in $x(t)$ that takes time to decay back.
- Another important thing to mention is that the filter speeds have to be much slower than our dither, which also has to be slower than the bandwidth of the system. We discussed how the system adds delays meaning that we have less margin of error to work with. Therefore we might have to use slower filtering to keep the system from going unstable, which makes the tracking slower.

Figure 9 shows the optimum ESC variables with the dynamic plant.

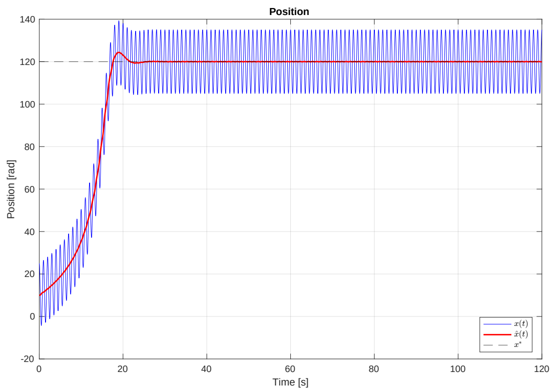

(a) Stage position $x(t)$ and $\hat{x}(t)$.

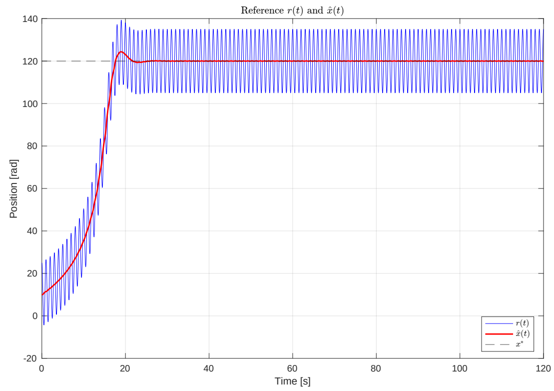

(b) Reference $r(t)$ and $\hat{x}(t)$.

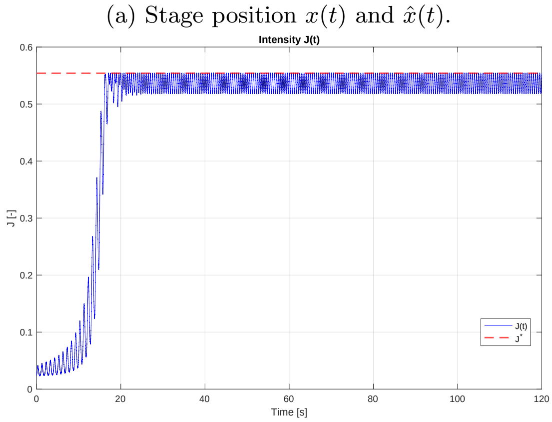

(c) Intensity $J(t)$.

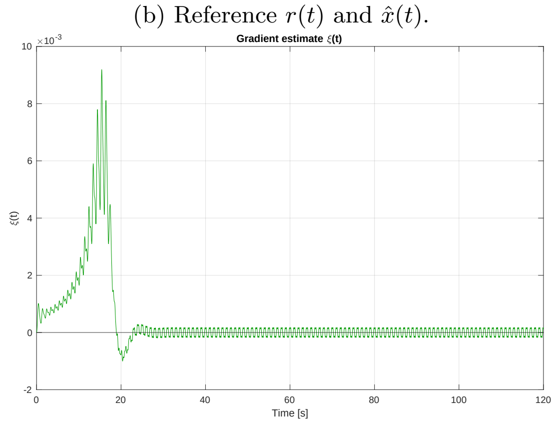

(d) Gradient estimate $\xi(t)$.

Figure 9: ESC convergence with plant dynamics and noise.