

Overhead the Albatross Hangs Motionless Upon the Air with Dynamic Soaring

Antonia Bronars
bronars@mit.edu

Rebecca Jiang
rhjiang@mit.edu

Siddharth Nayak
sidnayak@mit.edu

Abstract—Albatrosses are capable of travelling thousands of kilometres daily with very little fuel supplies. They utilise a flight strategy called dynamic soaring which helps them to extract propulsive energy from the horizontal wind shear layers formed a few metres above the ocean surface. This allows them to fly for hours or even days without flapping their wings. We investigate this flight behaviour from a trajectory optimization point of view and try to emulate these zero-cost trajectories numerically. We also analyse the variation in these trajectories caused by changing the wind parameters. Finally we use Linear-Quadratic Regulator (LQR) control for tracking these trajectories, to account for errors and noise in the dynamics. The code for reproducing the experiments in this paper can be found [here](#).

I. INTRODUCTION

Dynamic soaring is a flight maneuver utilised by birds such as the wandering albatross (*Diomedea exulans*) to gain energy by repeatedly crossing the boundaries between wind layers of different velocities. Such wind layers of different velocities are generally found close to obstacles and surfaces. Dynamic soaring is widely used in remote controlled gliders and UAVs to travel long distances with limited energy consumption and to achieve high speeds. Therefore, dynamic soaring has the potential to improve the energy efficiency of robotic gliders. An albatross-like glider which utilises the wind stream gradients could be used to survey oceans with a virtually infinite range.

The motion of an albatross over long distance is known to be periodic, and it was our goal through this project to find and track a periodic trajectory that obeys the strongly nonlinear and underactuated dynamics of the albatross flight. Trajectory optimization provides a suitable framework for this problem. One can use trajectory optimization methods to find trajectories which have zero energy cost and satisfy simplified model dynamics of the bird. In this work, we use direct collocation optimization methods to find these trajectories under different wind conditions and time periods of the trajectory cycles.

Optimal trajectories for the full albatross dynamics can be computed offline, which is important given the highly nonlinear and underactuated nature of the bird's dynamics. Also offline, this nominal trajectory can be linearized and the gains of a corresponding finite horizon Linear-Quadratic Regulator (LQR) controller can be computed. Wrapping an LQR controller around the nominal trajectory allows for

disturbance rejection in the presence of noise. At runtime, it is only necessary for the algorithm to interpolate between the pre-computed gains of the finite horizon LQR controller in order to drive the albatross's state to the nominal trajectory. This framework provides a fast, tractable means of controlling the nonlinear, underactuated dynamics of albatross flight with zero energy cost.

We discuss about previous works related to dynamic soaring in Section II. We define the dynamics model for the glider and the wind profile in Section III. Then we formulate the trajectory optimization problem as a non-linear program in Section IV. Following this in Section VI, we discuss about the usage of LQR Control for stabilising the trajectories about the nominal paths obtained via trajectory optimization in the presence of noise and errors within the dynamics of the model. Section V contains the results for all of our experiments.

II. RELATED WORK

The study of dynamic soaring was first carried out by Leonardo da Vinci [1] around 400 years before the invention of the first aircraft by the Wright brothers. But he did not model dynamic soaring formally with mathematical equations. Later, Lord Rayleigh [2] modelled the different parts of the albatross flight pattern and proposed the requirement of a non-uniform wind by describing a two-layer system with the two layers moving at different speeds.

Many of the works related to analysing albatross flight patterns have been focused on the energy management [3] and the metabolic costs [4] of the albatrosses. These works have influenced the design of long-endurance fixed-wing UAVs. Richardson [5] utilises dynamic soaring for marginally increasing the speed of the glider in each cycle of the periodic trajectory and making the glider reach up to a record speed of 780km/h. They also suggest that due to the compressibility of air at higher speeds, the highest achievable speed of the gliders will level out at 980km/h.

On the other hand, Bousquet et al. [6] analyse the flight patterns for different shear thickness for the wind layers. They model their optimization to find the minimum wind strength required to sustain the zero-energy cost trajectories. They evaluate their model for different shear layer thickness and show that the trajectories vary a lot under these different

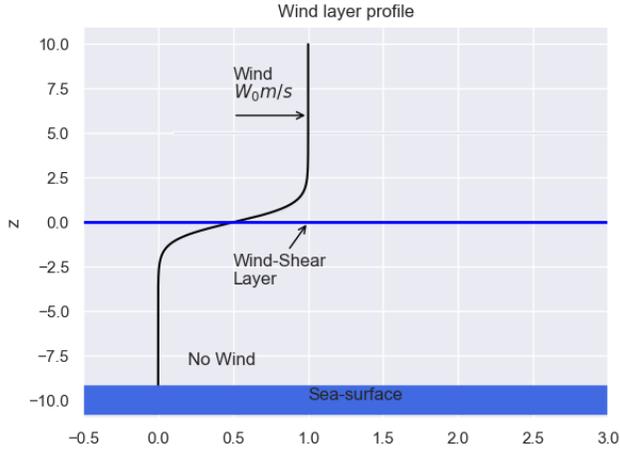


Fig. 1. The sigmoid wind profile and the wind shear layer. The sea-level is at $z = -10$, the wind shear boundary at $z = 0$

conditions specifically when the shear layers are quite thin. These trajectories are shown to be consistent with GPS recordings of albatrosses. We base our trajectory optimization formulation and simplified dynamic model of the albatross heavily on this work.

III. DYNAMICS

We model a 3-dimensional version of the glider as a point mass (as commonly used in the dynamic soaring literature). Specifically we model the dynamics of the glider as defined in [6]. The state of the glider is denoted by $\mathbf{X} = (x, y, z, V, \psi, \gamma)$ where x, y, z are the spatial Cartesian coordinates, V is the albatross airspeed and ψ and γ are the relative air-heading angle and air-relative flight path angle respectively. The control inputs are the lift coefficient and roll angle $\mathbf{u} = [c_L, \phi]$. Additionally, we use a sigmoid wind-profile as $W(z) = \frac{W_0}{1 + \exp(-z/\delta)}$ flowing in the y direction, where z is the altitude. Here W_0 is the free stream wind speed at a reference altitude ($z = 10$) and δ is the shear layer thickness. As shown in Fig 1, note that $z = 0$ is not the sea-level and rather is called the “wind-shear boundary layer.” The velocity of the wind should be zero just above the sea-level surface and thus we use $z = -10$ as the sea-level surface. As noted in [6], the sigmoid wind profile is a more robust way to approximate a wide class of wind-fields. For example, the sigmoid wind model is equivalent to Rayleigh’s wind model [2] (step function) in the limit $\delta \rightarrow 0$

Parameter	Value
g	9.8 kg/m^2
m	9.5 kg
$c_{D,0}$	0.01
S	0.65 m^2
ρ	1.2 kg/m^3
f_{max}	40

TABLE I

PARAMETER VALUES USED IN THIS WORK

The equations of motion for the glider are:

$$\dot{x} = V \cos \gamma \cos \psi \quad (1)$$

$$\dot{y} = V \cos \gamma \sin \psi - W \quad (2)$$

$$\dot{z} = V \sin \gamma \quad (3)$$

$$m\dot{V} = -D - mg \sin \gamma + m\dot{W} \cos \gamma \sin \psi \quad (4)$$

$$mV\dot{\gamma} = L \cos \phi - mg \cos \gamma - m\dot{W} \sin \gamma \sin \psi \quad (5)$$

$$mV\dot{\psi} \cos \gamma = L \sin \phi + m\dot{W} \cos \psi \quad (6)$$

$$c_D = c_{D,0} + kc_L^2 \quad (7)$$

$$L = \frac{1}{2}c_L\rho SV^2 \quad (8)$$

$$D = \frac{1}{2}c_D\rho SV^2 \quad (9)$$

The coefficient $c_{D,0}$ represents the system’s drag when no lift is generated and the parameter k is the coefficient of additional drag generated due to the lift. Denoting $f \equiv c_L/c_D$, the glider’s lift-to-drag ratio (or *finesse*), we can relate k as $k^{-1} = 4f_{max}^2 c_{D,0}$. The parameters used for our experiments are in Table I. The above equations are a simplified point-mass model of the dynamics of the albatross flights and do not take into account the aerodynamic slip. In order to increase the fidelity of the study to the true albatross dynamics, one could take this slip into account, as it can have significant effects in the high-speed regime. To do so, one can consider the flat-plate model equations as defined by Cory and Tedrake [7]. We do not consider this model in our study, as a first order approximation of the albatross dynamics is sufficient for our study of nonlinear trajectory optimization, trajectory linearization, and trajectory stabilization. We leave the inclusion of more complex dynamic models to future work.

IV. TRAJECTORY OPTIMIZATION

We formulate the trajectory optimization problem in two different forms:

- Travelling Trajectories: these trajectories are S-shaped with alternating turn directions.
- Non-travelling (Loitering) Trajectories: these trajectories are O-shaped with a constant turn direction.

Both of the above formulations are non-convex and energetically equivalent, only differing slightly in the constraints used to optimize them. We want the motion to be periodic as the albatrosses repeat the same pieces of trajectories while flying long distances. So we just try to find the trajectory for one single time period T . We divide the trajectory into N collocation points. We define $dt = T/(N - 1)$. Here we

denote the dynamics defined in Section III in a compact form as $\dot{\mathbf{X}} = g(\mathbf{X}, \mathbf{u})$. Since we are not minimising any cost, and are instead solving a constraint satisfaction problem, we use a dummy cost for the objective in the formulation. We constrain the model to follow the equations of motion.

A. Travelling Trajectories

Since travelling trajectories are S-shaped, the periodicity constraints are enforced by making z, V, ψ, γ to be the same at the beginning and the end of the time period.

$$\begin{aligned} & \min_{\mathbf{X}[0], \dots, \mathbf{X}[N], \mathbf{u}[0], \dots, \mathbf{u}[N]} \quad \mathbf{0}^T \mathbf{X} \\ & \text{s.t.} \quad z[0] = 0 \\ & \quad V[n], c_L[n] > 0 \\ & \quad z[n] > \text{sea-level} \\ & \quad -\pi < \psi[n] < \pi \\ & \quad -\pi/2 < \gamma[n] < \pi/2 \\ & (z[N], V[N], \psi[N], \gamma[N]) = (z[0], V[0], \psi[0], \gamma[0]) \\ & \mathbf{X}[n+1] = \mathbf{X}[n] + g(\mathbf{X}[n], \mathbf{u}[n]) \cdot dt \quad \forall n \in \{0, \dots, N-1\} \end{aligned}$$

B. Non-travelling Trajectories

The non-travelling, circular trajectories differ from the travelling ones in three constraints:

- The initial and the final air-heading angle ψ now differ by 2π instead of 0
- The air-heading angle ψ is now restricted to be within $\pm 3\pi$ instead of $\pm\pi$
- The initial and the final x coordinates are restricted to be equal. Note that this constraint is not strictly required. This constraint just helps to maintain realistic values for c_L while maintaining the O-shaped feature of the trajectory.

$$\begin{aligned} & \min_{\mathbf{X}[0], \dots, \mathbf{X}[N], \mathbf{u}[0], \dots, \mathbf{u}[N]} \quad \mathbf{0}^T \mathbf{X} \\ & \text{s.t.} \quad z[0] = 0 \\ & \quad V[n], c_L[n] > 0 \\ & \quad z[n] > \text{sea-level} \\ & \quad -3\pi < \psi[n] < 3\pi \\ & \quad -\pi/2 < \gamma[n] < \pi/2 \\ & (x[N], z[N], V[N], \gamma[N]) = (x[0], z[0], V[0], \gamma[0]) \\ & \quad \psi[N] = \psi[0] + 2\pi \\ & \mathbf{X}[n+1] = \mathbf{X}[n] + f(\mathbf{X}[n], \mathbf{u}[n]) \cdot dt \quad \forall n \in \{0, \dots, N-1\} \end{aligned}$$

Note that since we constrain $z[0] = z[N]$ and $V[0] = V[N]$, the total energy (kinetic + potential) change after a time-period is zero. We use the “trust-constr” optimizer [8] from the python `scipy` optimization package [9] to optimize the above formulations. We typically use $N = 50$ time steps, leading to $\mathcal{O}(500)$ variables and constraints. Our Python implementation converged in $\mathcal{O}(1-5)$ minutes on a MacBook Air 2015.

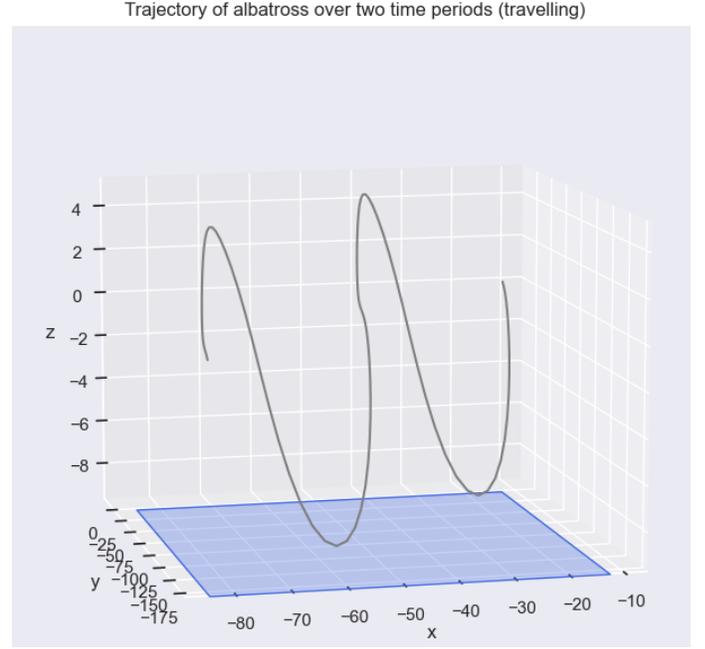


Fig. 2. The trajectory over two time periods for the “travelling” formulation. The blue layer represents the sea-level.

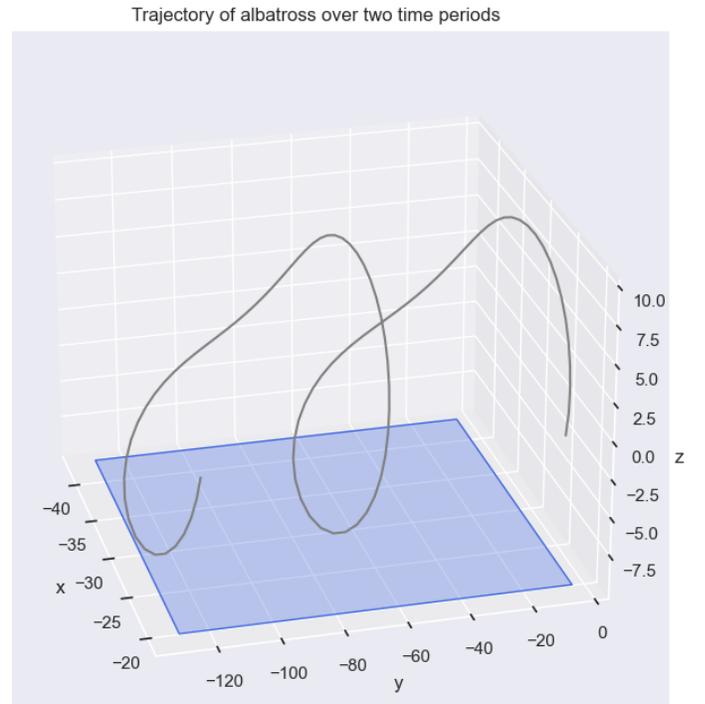


Fig. 3. The trajectory over two time periods for the “non-travelling” formulation.

C. Initializing the Optimizer

It is necessary to specify an initial guess to the optimizer. Part of the spirit of this project is to attempt to recover the intuition and natural phenomena of albatross flight without artificially inserting it, and thus we do not attempt to model actual albatross behavior in order to form the initial guess. Instead, however, we form an initial guess that is dynamically consistent, while violating most other constraints. To do so, we simulate a trajectory forward in time from a (fairly arbitrarily selected) initial condition of $\mathbf{X}(0) = [0, 0, 0, 10, \pi/8, 0]$ which a (similarly arbitrary) constant control input of $\mathbf{u} = [1.5, \pi/8]$. We use these control values and the resulting trajectory to form the initial guess.

V. OPTIMIZATION EXPERIMENTS

We perform experiments with the trajectory optimization by varying the time-period (T), wind-strength (W_0) and shear layer thickness (δ). Our default values for all experiments are: $T = 7s, W_0 = 7.8m/s, \delta = 12m$.

Figures 2 and 3 show the trajectories (over two time periods) obtained with the default parameters for the travelling and the non-travelling cases. In both cases, the shear layer thickness is sufficiently high for the trajectories to be somewhat three dimensional (characteristic of cases with sufficiently thick shear layers, where the albatross has enough wind to move without much turning). Note that for both trajectories, the shear layer exists at z height zero, and penetrations through the shear layer at $z = 0$ are the means by which the albatross gains lift and dynamic soaring occurs.

Figure 4 shows the variation of the kinetic energy, potential energy and the total mechanical energy over a duration of one time period of a travelling trajectory. The change in the total mechanical energy over the time-period is zero, as constrained by the optimizer. This ensures that the periodic motion of the albatross over multiple periods will have constant energy. The energy of the albatross dips during the period as a result of losses due to drag, and it re-gains energy by penetrating the shear layer at $z = 0$. Instead of constraining $V[0] = V[n]$, one can constrain $V[n]$ to be slightly higher than $V[0]$ if we want the glider to gain speed over time.

A. Varying the time-period

We vary the time-period of the trajectories as $T = \{5s, 7s, 12s\}$ while keeping the other parameters at the default value for both the travelling and the non-travelling case. For the travelling trajectories, in Figure 5, for $T = 12s$, the trajectory almost completes two cycles of the upward-downward motion in the allotted time, suggesting that $T = 12s$ is quite large for the characteristic albatross flight path. The highest altitude reached is much higher for $T = 7s$ than $T = 5s$. For non-travelling trajectories in Figure 6 the path obtained for $T = 5s$ is almost circular, in the sense that it does not have to turn more than once. Whereas for $T = 7s$ and $T = 12s$ the glider has to turn more than once to achieve the

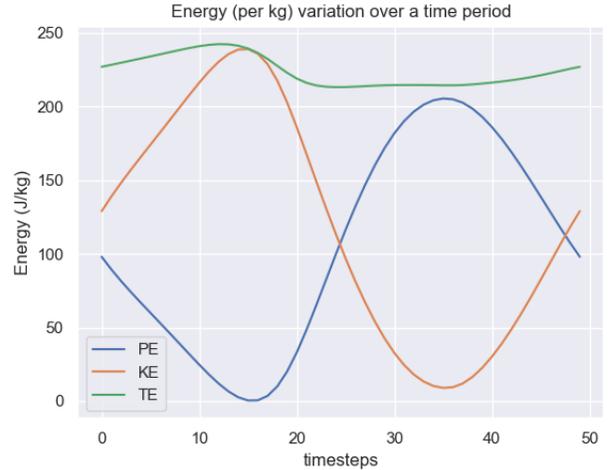


Fig. 4. Variation of the kinetic energy, potential energy and the total mechanical energy over a travelling trajectory

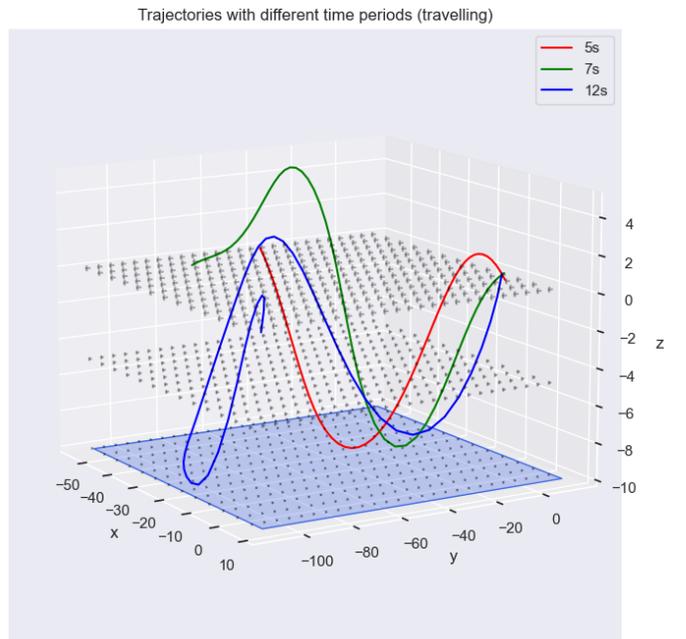


Fig. 5. Travelling trajectories with different time periods $T = \{5s, 7s, 12s\}$. The arrows represent the wind flow.

zero cost trajectory. Moreover, for $T = 12s$, the path is similar to the travelling case with its upward-downward motion.

B. Varying the wind strength

We vary the wind strength as $W_0 = \{5m/s, 7.8m/s, 12m/s\}$. Interestingly, for both the travelling and non-travelling cases, shown in Figures 7 and 8 respectively, there is a strictly negative correlation between W_0 and maximum altitude achieved by the albatross. The trajectories with higher wind speeds also exhibit more curvature in both the travelling and non-travelling cases, traversing larger distances in the x -direction, which may be

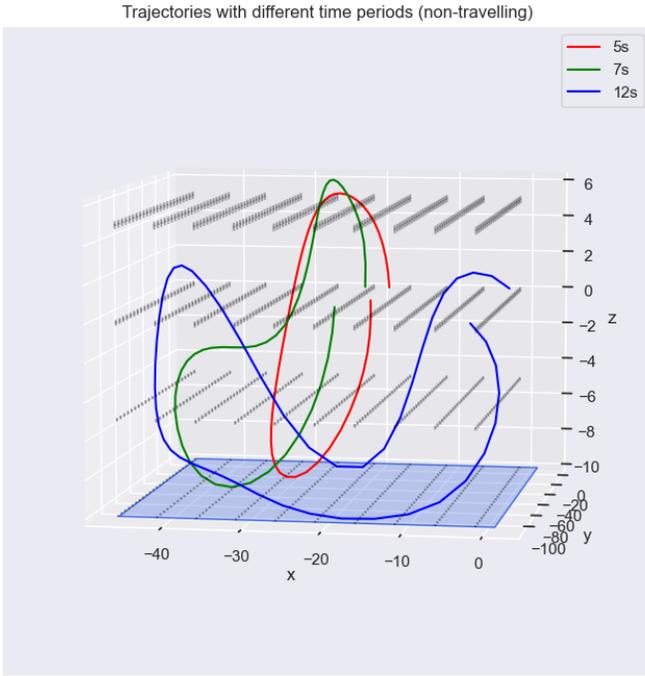


Fig. 6. Non-travelling trajectories with different time periods $T = \{5s, 7s, 12s\}$

explained by the fact that the higher wind speeds reduce the need to travel parallel to the wind, in the y -direction, as there is sufficient energy to expend more travel distance orthogonal to this direction.

C. Varying the shear layer thickness

We vary the shear layer thickness as $\delta = \{1m, 3m, 7m, 12m\}$ while keeping the other parameters at the default value for both the travelling and the non-travelling case. The thickness of the shear layer impacts the amount of lift the albatross has access to during its flight. For the travelling trajectories, the albatross flight is composed of more, smaller angle arcs as the thickness of the shear layer decreases. This is in line with what is seen in the flight of real albatross birds when flying in thin shear layers, according to [6]. For the non-travelling (or loitering) flight, the trajectories tend to decrease in height as the shear layer decreases in thickness. The smaller angle arc pattern is not seen for the non-travelling trajectories because of the way the dynamics of the flight type is constrained (in non-travelling flight, alternating turn directions is not permitted). Therefore, the primary trend that can be seen is a smaller range of motion as a result of access to less energy from the shear layer.

VI. TRACKING THE OPTIMIZED TRAJECTORY

The direct collocation approach results in approximate satisfaction of the system dynamics, due to the collocation method as well as the tolerance in the constraints during optimization. Additionally, when noise is added to perturb the system dynamics, open-loop execution of the optimized

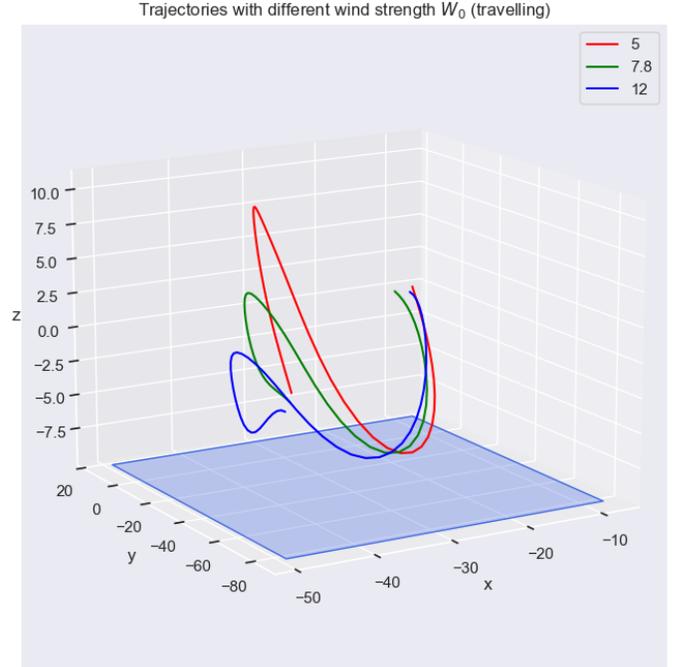


Fig. 7. Travelling trajectories with different wind-strengths $W_0 = \{5m/s, 7.8m/s, 12m/s\}$

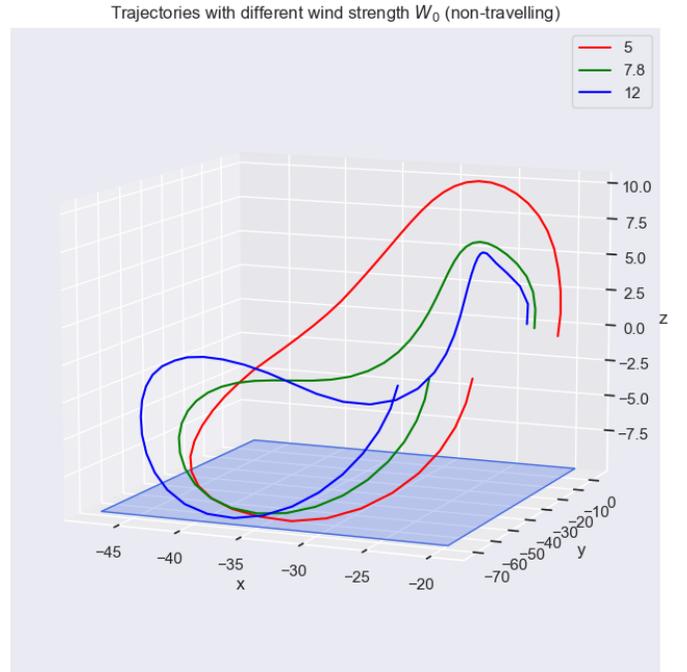


Fig. 8. Travelling trajectories with different wind-strengths $W_0 = \{5m/s, 7.8m/s, 12m/s\}$

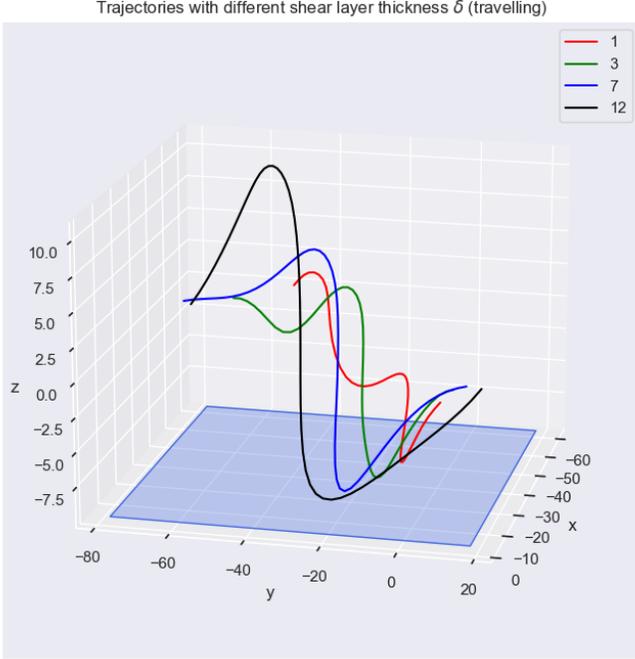


Fig. 9. Travelling trajectories with different wind shear layer thickness $\delta = \{1m, 3m, 7m, 12m\}$

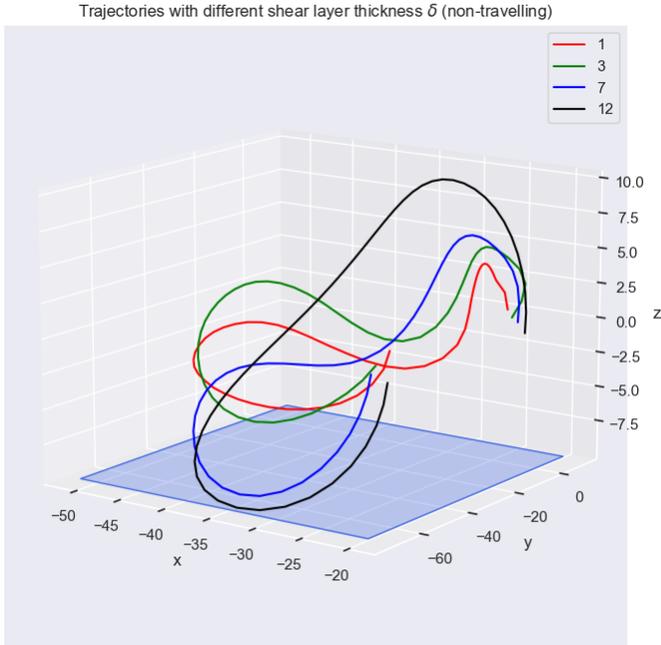


Fig. 10. Travelling trajectories with different wind shear layer thickness $\delta = \{1m, 3m, 7m, 12m\}$

control profile can result in poor tracking of the optimized trajectory. We model additive zero-mean Gaussian noise, in numerical experiments, with standard deviations: $[\sigma_{\dot{x}}, \dots, \sigma_{\dot{y}}] = [0.1 \text{ m/s}, 0.1 \text{ m/s}, 0.1 \text{ m/s}, 1 \text{ m/s}^2, \pi/20 \text{ rad/s}, \pi/20 \text{ rad/s}]$. We then use a time-varying finite horizon LQR scheme to control the system dynamics about the optimized trajectory in the presence of noise.

A. LQR formulation

We formulate a finite-horizon linear quadratic regulator to control the error coordinates, $\bar{\mathbf{X}}(t) = \mathbf{X}(t) - \mathbf{X}_0(t)$, where $\mathbf{X}_0(t)$ is the nominal trajectory, in this case the optimized trajectory. In order to use LQR, we use a time-varying linearization of the dynamics, $\dot{\bar{\mathbf{X}}}(t) = A(\mathbf{X}(t))\bar{\mathbf{X}}(t) + B(\mathbf{X}(t))\bar{\mathbf{u}}(t)$. We assume that the nominal trajectory satisfies these dynamics, $\dot{\mathbf{X}}_0(t) = A(\mathbf{X}_0(t))\mathbf{X}_0(t) + B(\mathbf{X}_0(t))\mathbf{u}_0(t)$, and that $\mathbf{X}(t) \approx \mathbf{X}_0(t)$ and thus $A(\mathbf{X}(t)) \approx A(\mathbf{X}_0(t))$, $B(\mathbf{X}(t)) \approx B(\mathbf{X}_0(t))$, and thus we can use the same linearization in the error coordinates, evaluated at the nominal trajectory:

$$\dot{\bar{\mathbf{X}}}(t) = \dot{\bar{\mathbf{X}}}(t) - \dot{\mathbf{X}}_0(t) = A(t)\bar{\mathbf{X}}(t) + B(t)\bar{\mathbf{u}}(t) \quad (10)$$

where we have written $A(\mathbf{X}_0(t)), B(\mathbf{X}_0(t))$ as $A(t), B(t)$ because the nominal profile $\mathbf{X}_0(t)$ is known. $\bar{\mathbf{u}}(t)$ is the deviation from the nominal control profile at time t , an error coordinate.

The $A(t)$ and $B(t)$ matrices are the Jacobians of g with respect to \mathbf{X} and \mathbf{u} , evaluated at \mathbf{X}_0 and \mathbf{u}_0 , where g , defined in Section III, is the dynamics transition function $\dot{\mathbf{X}} = g(\mathbf{X}, \mathbf{u})$.

In the error-coordinates, we can define a finite-horizon LQR to minimize the cost in Equation 11.

$$J = \bar{\mathbf{X}}^T(T)Q_f\bar{\mathbf{X}}(T) + \int_{t=0}^{t=T} \bar{\mathbf{X}}^T(t)Q\bar{\mathbf{X}}(t) + \bar{\mathbf{u}}^T(t)R\bar{\mathbf{u}}(t) \quad (11)$$

with Q_f, Q , and R being cost weight matrices, design choices, selected in numerical experiments as I_6, I_6 , and I_2 respectively.

This results in a linear feedback law, $\bar{\mathbf{u}}^*(t) = -R^{-1}B(t)^T S(t)\bar{\mathbf{X}}(t)$, where $S(t)$ satisfies the continuous-time differential Riccati equation,

$$-\dot{S}(t) = S(t)A(t) + A^T(t)S(t) - S(t)B(t)R^{-1}B(t)^T S(t) + Q \quad (12)$$

Equation 12 is solved backward in time from a final condition, $S(T) = Q_f$ [10]. The actual control inputs can be found simply by taking $\mathbf{u}^*(t) = \bar{\mathbf{u}}^*(t) + \mathbf{u}_0(t)$.

B. Results

Figure 11 shows the performance of the LQR compared with the nominal trajectory as well as open-loop execution of the nominal control inputs. Clearly, the tracking of the nominal trajectory is not perfect, but is certainly better than the open-loop trajectory.

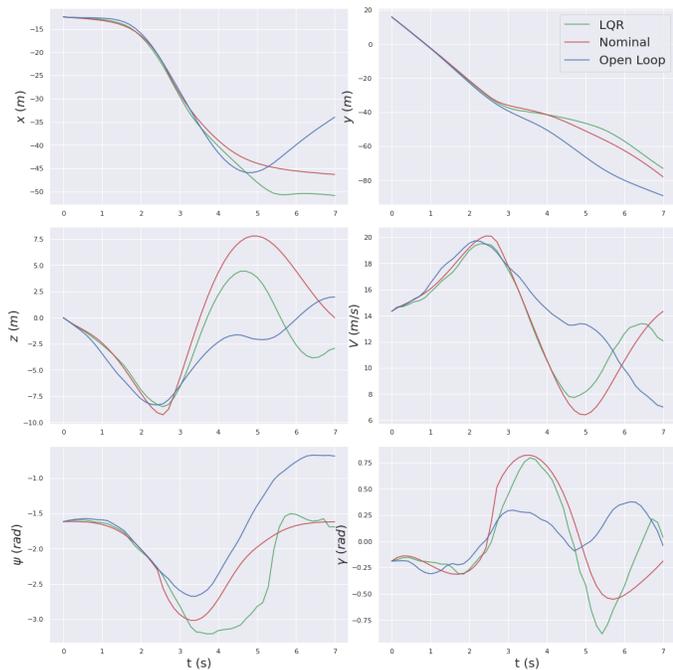


Fig. 11. Comparisons of tracking an optimized trajectory with noisy dynamics. LQR: LQR tracking of the nominal trajectory, with noisy dynamics. Nominal: Optimized trajectory with approximate dynamics (no noise). Open Loop: Open-loop simulation of the nominal (optimized) control inputs with noisy dynamics.

In some state components, particularly ψ , the LQR trajectory is more jerky than the open-loop trajectory. An intuitive explanation for this is that the LQR controller is applying control actions to correct for noisy disturbances, and noisy disturbances are jerky. The open-loop trajectory, on the other hand, while also subject to noisy disturbances, is not additionally subject to jerky control actions.

VII. CONCLUSION

We have largely managed to replicate the key observations from Bousquet et al. [6]. The observations throughout Section V are consistent with physics and intuition, as well as albatross behavior observed in nature. An interesting extension of this work, from the nature-observation perspective, would be to study how albatrosses handle noise, such as unexpected wind gusts, and whether this behavior shares any characteristics with the LQR control implemented.

The LQR results demonstrate the ability to overcome inexact modeling of the dynamics in the trajectory optimization, as well as noise injection in the system. It would be interesting to continue this study to investigate how different collocation techniques in the optimization routine influence the performance of the open-loop control trajectory. Perhaps a higher-order collocation technique such as a higher-order Runge Kutta method, or a spline method, would yield more accurate dynamics.

Additionally, while we worked within an optimization framework, we did not take advantage of the freedom to have a cost function. As noted in Section V, many of the trajectories found contained excessive motion, or several sequences of up-and-down motion within one period. If we had penalized some measure of movement, such as arc length travelled, or the time-period, perhaps we would see less of this behavior. It would be interesting to study which cost functions result in trajectories that are most similar to actual albatross trajectories.

A final direction of future work we would be interested in perusing is using Lyapunov analysis to quantify the stability of the trajectory. It could be informative to compute regions of convergence surrounding the trajectory to gain an understanding of both the magnitude of disturbance in state that the trajectory is stable to, and regions of the trajectories that are especially vulnerable to disturbances. It would be interesting to compare the results of a Lyapunov analysis to data collected on the actual trajectories of albatross flight in varying wind conditions.

ACKNOWLEDGMENT

We would like to thank the 6.832-Underactuated Robotics Spring 2021 course staff for helpful discussions and their comments on our project scope and progress.

REFERENCES

- [1] Philip Richardson. Leonardo da vinci's discovery of the dynamic soaring by birds in wind shear. *Notes and Records: the Royal Society journal of the history of science*, 73:20180024, 10 2018.
- [2] Rayleigh. The soaring of birds. *Nature*, 27(701):534–535, Apr 1883.
- [3] John Arnould, D.R. Briggs, P.A. Prince, and A.G. Wood. The foraging behaviour and energetics of wandering albatrosses brooding chicks. *Antarctic Science*, 8:229 – 236, 09 1996.
- [4] G. Sachs, J. Traugott, A. Nesterova, and F. Bonadonna. Experimental verification of dynamic soaring in albatrosses. *Journal of Experimental Biology*, 216:4222 – 4232, 2013.
- [5] Philip Richardson. High-speed dynamic soaring. *R/C Soaring Digest*, 29:36–49, 04 2012.
- [6] Gabriel Bousquet, Michael Triantafyllou, and Jean-Jacques Slotine. Optimal dynamic soaring consists of successive shallow arcs. *Journal of The Royal Society Interface*, 14:20170496, October 2017.
- [7] Rick Cory and Russ Tedrake. Experiments in fixed-wing uav perching. 08 2008.
- [8] Andrew R. Conn, Nicholas I. M. Gould, and Philippe L. Toint. *Trust Region Methods*. Society for Industrial and Applied Mathematics, 2000.
- [9] Eric Jones, Travis Oliphant, and Pearu Peterson. SciPy: open source scientific tools for Python, 2001.
- [10] Russ Tedrake. *Underactuated Robotics: Algorithms for Walking, Running, Swimming, Flying, and Manipulation (Course Notes for MIT 6.832)*. Downloaded on 05/12/2021 from <http://underactuated.mit.edu/>.