AFRL-AFOSR-VA-TR-2022-0386

Supermaneuverable Autonomous Pursuit - Peregrine Falcon Versus Pigeon Inspired UAVs

Scott Delp
LELAND STANFORD JUNIOR UNIVERSITY
450 SERRA MALL
STANFORD, CA,
US

07/14/2022
Final Technical Report

Air Force Research Laboratory
Air Force Office of Scientific Research
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
20220714	Final	START DATE 20180915	END DATE 20210914

4. TITLE AND SUBTITLE
Supermaneuverable Autonomous Pursuit - Peregrine Falcon Versus Pigeon Inspired UAVs

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
	FA9550-18-1-0525	61102F
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Scott Delp

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
LELAND STANFORD JUNIOR UNIVERSITY
450 SERRA MALL
STANFORD, CA
US

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
Air Force Office of Scientific Research 875 N. Randolph St. Room 3112 Arlington, VA 22203	AFRL/AFOSR RTB2	AFRL-AFOSR-VA-TR-2022-0386

12. DISTRIBUTION/AVAILABILITY STATEMENT
A Distribution Unlimited: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
A peregrine falcon pursuing a pigeon in flight provides compelling inspiration for unmanned aerial vehicle autonomy, flight control, grasping, perching, and rudderless morphing flight system design, because of the multifunctional capabilities of the falcon. The Highly
Maneuverable Autonomous UAV DESI project involved a close collaboration among two Stanford university labs (Lentink and Cutkosky) and their industry partner (Skydio), to develop aerial vehicles capable of rudderless morphing flight in turbulence, perching on
complex surfaces such as tree branches and catching airborne targets such as small UAVs in flight. To translate falcon capabilities to the UAV realm, we focused our bioinspired robotic multifunctional structures and autonomous pursuit research on four key challenges during two years of collaborative development. First, we developed the Stereotyped Nature inspired Aerial Grasper, SNAG, a robotic leg and end effector mounted to a quadrotor that can perch on a range of dirt, moss and lichen covered branches and can catch objects dynamically. Next, we developed the first rudderless morphing flyer, PigeonBot 2, a biohybrid rudderless morphing flyer with underactuated soft feathered wings and tail. It not only uncovered the previously unknown principles of morphing rudderless flight, it also demonstrated them in a high turbulence in a wind tunnel and during extensive autonomous flight testing outdoors. Then, we focused on achieving in-flight dynamic grasping with a quadcopter that can grasp target airborne drones in flight and recover from the collision that triggers the grasp. Finally,
we demonstrate an aerial grasping mechanism integrated onto the autonomous Skydio X2 platform. We then used the Skydio vision-based autonomy engine to test grasping targets in flight. This work inspired the development of a trajectory generator that combines differential
flatness with trajectory optimization to satisfy state constraints more directly and quickly. With these four major collaborative accomplishments we have demonstrated how the basic outcomes of this DESI project enable the developing of more versatile and robust aerial
vehicles generally.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	UU	39

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
BYUNG LEE	426-8483

Super-Maneuverable Autonomous Pursuit:

Peregrine Falcon vs Pigeon Inspired UAVs

B.-L. ("Les") Lee ***PM***

Jean-Luc Cambier ***Co-PM***

Frederick Leve ***Co-PM***

David Lentink ***PI*** Stanford University

Mark Cutkosky ***Co-PI*** Stanford University

Hayk Martiros ***Industry Partner*** Skydio

Project Outcome Overview

The project involved a close collaboration among two university labs at Stanford (Lentink and Cutkosky) and Industry (Skydio) to develop agile rudderless morphing fliers, demonstrate perching on complex surfaces such as tree branches and catch mobile airborne targets (small drones) in flight. The team's four main accomplishments & four flying demonstrators are:

1. Stereotyped Nature Inspired Aerial Grasper (SNAG)
Bird inspired stereotyped perching on complex surfaces including real tree branches—covered by dirt, lichen and moss—and dynamically grasping complex objects.
Science Robotics cover paper

2. Rudderless Morphing Flyer (PigeonBot 2)
First 'rudderless morphing flyer' with underactuated morphing wings and tail consisting of hierarchically soft surfaces and demonstrated robust flight control in high turbulence.
Invited by Science Robotics (in prep.).

3. Aerial Drone Grasper
Capable of grasping small airborne drones in flight and flight control recovery after impact harnessing a lightweight gripper that uses collision energy to securely grasp drones.
RA-L and IROS paper

4. Integrated Autonomous Platform
Collaborative platform integrating the autonomous Skydio X2 with the Stanford gripper to leverage vision-based autonomy for aerial grasping of small airborne drones.
Prototype demonstration

Abstract

A peregrine falcon pursuing a pigeon in flight provides compelling inspiration for unmanned aerial vehicle autonomy, flight control, grasping, perching, and rudderless morphing flight system design, because of the multifunctional capabilities of the falcon. The Highly Maneuverable Autonomous UAV DESI project involved a close collaboration among two Stanford university labs (Lentink and Cutkosky) and their industry partner (Skydio), to develop aerial vehicles capable of rudderless morphing flight in turbulence, perching on complex surfaces such as tree branches and catching airborne targets such as small UAVs in flight. To translate falcon capabilities to the UAV realm, we focused our bioinspired robotic multifunctional structures and autonomous pursuit research on four key challenges during two years of collaborative development. First, we developed the Stereotyped Nature inspired Aerial Grasper, SNAG, a robotic leg and end effector mounted to a quadrotor that can perch on a range of dirt, moss and lichen covered branches and can catch objects dynamically. Next, we developed the first rudderless morphing flyer, PigeonBot 2, a biohybrid rudderless morphing flyer with underactuated soft feathered wings and tail. It not only uncovered the previously unknown principles of morphing rudderless flight, it also demonstrated them in a high turbulence in a wind tunnel and during extensive autonomous flight testing outdoors. Then, we focused on achieving in-flight dynamic grasping with a quadcopter that can grasp target airborne drones in flight and recover from the collision that triggers the grasp. Finally, we demonstrate an aerial grasping mechanism integrated onto the autonomous Skydio X2 platform. We then used the Skydio vision-based autonomy engine to test grasping targets in flight. This work inspired the development of a trajectory generator that combines differential flatness with trajectory optimization to satisfy state constraints more directly and quickly. With these four major collaborative accomplishments we have demonstrated how the basic outcomes of this DESI project enable the developing of more versatile and robust aerial vehicles generally.

I Aerial Grasping and Perching:

SNAG (Stereotyped Nature-inspired Aerial Gripper)

Leveraging recent findings on how birds take off, land, and grasp, we developed a biomimetic robot that can dynamically perch on complex arboreal surfaces, grasp irregular objects, and release them. To determine the range of design, kinematic, and control parameters that are sufficient for perching success, we launched the robot at tree branches. The results corroborate a mathematical model, which shows that larger animals and robots must accommodate disproportionately larger angular momenta to achieve similar performance. Closed-loop balance control after contact serves an important role in maximizing the range of parameters sufficient for perching; similar concepts apply to grasping in mid-air. These findings advance our understanding of the avian perching apparatus and highlight new design concepts that enable robots to perch on natural surfaces for longitudinal environmental monitoring.

Stereotyped behaviors and mechanisms inspired by birds

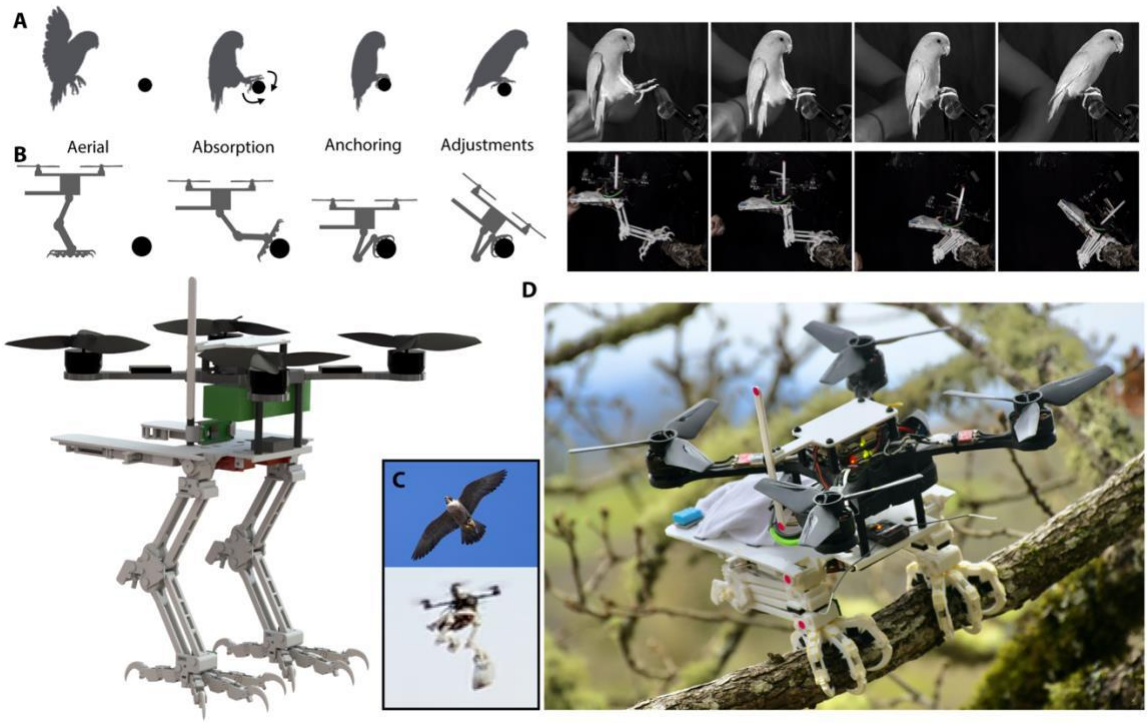

Figure G1 SNAG is a bird-inspired robotic leg and end effector, which enables aerial robots to take off and land on complex surfaces as well catch objects in the air.

Leveraging our recent discoveries explaining how birds dynamically grasp a wide range of complex surfaces to perch (1), we developed the stereotyped nature-inspired aerial grasper (SNAG) for aerial robots (Fig. G1B). Like birds, SNAG harnesses stereotyped, passive, and active control behaviors across a diverse set of perches to land. SNAG integrates bird-inspired mechanisms and dimensions into its two legs that work together to grasp the perch when landing. They absorb impact energy to reconcile the momentum difference between the end effectors and the surface. During landing, SNAG also balances to stabilize itself and can safely release from the surface to take off. Experimental testing and modeling enabled us to explore the robot's perching sufficiency region: the multidimensional space of variables, including hardware design, kinematic, behavior, and perch characteristics, in which the robot succeeds at perching. Using the robot, we illustrate how the notable diversity in modern bird foot morphology is associated with modest variations in perching performance differences across natural surfaces. Because the perching process during landing and takeoff is analogous to that of catching and releasing objects in flight, we also explored and demonstrated SNAG's ability to catch objects (Fig. G1C). Last, we show how SNAG can land on and take off from tree branches in a forest, highlighting its use as a low-cost sensor for studying natural ecosystems (Fig. G1D).

Mechanism Design

SNAG's capabilities are primarily enabled by the design of its leg mechanism. SNAG comprises a bird-inspired bipedal foot and leg system (Fig. G2). It is mounted on a quadrotor aerial platform to control its flight. The grasping mechanism structure consists of three- dimensional (3D) printed segments. Jointed subcomponents are primarily printed in place to facilitate fast iteration, assembly, and repair. SNAG can be mounted on a variety of aerial platforms. Similar to the legs and feet of birds, the grasping mechanism handles surface variability upon contact for perching robustly, enabling the aerial approach phase to be stereotyped for all surfaces during landing. Similar to some birds of prey, SNAG can also dynamically catch objects with the same legs and feet used for perching. Movie 1 (link) demonstrates SNAG's core functionality and many of its key biomimetic design features. The first design iterations relied on more traditional engineering solutions, which underperformed. Only after 20 design iterations that increasingly mimicked bird legs and feet did we reach a design that could reliably land like birds do.

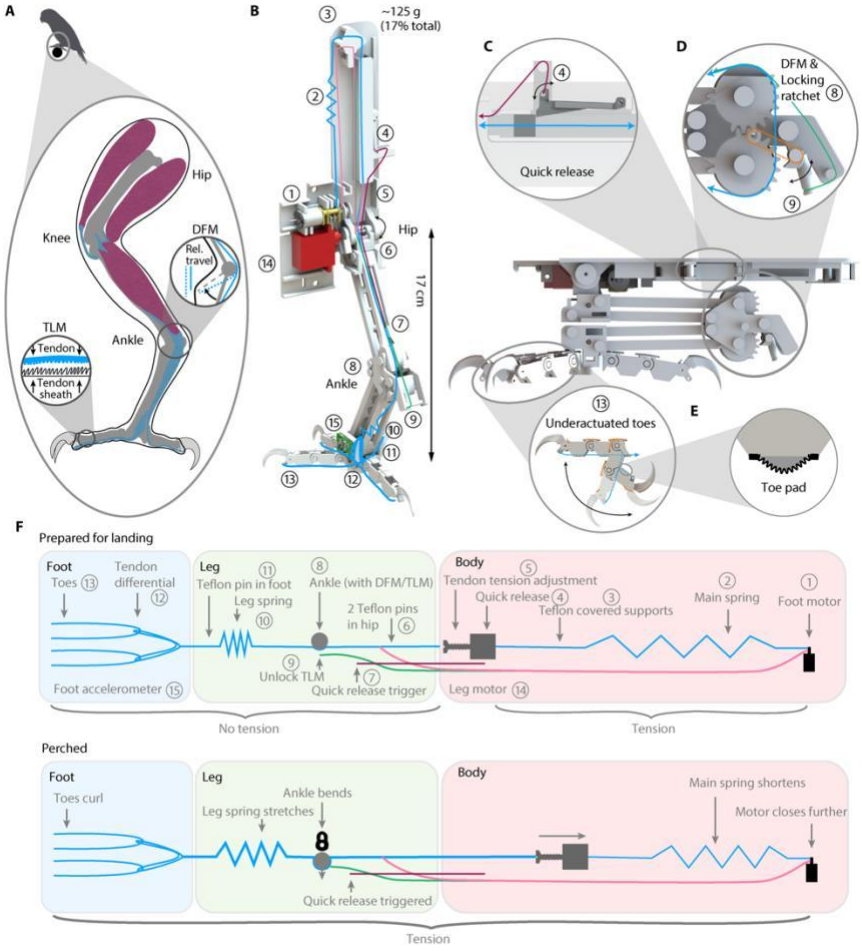

Figure G2 Technical overview of the mechanical design of the key mechanisms of the SNAG gripper.

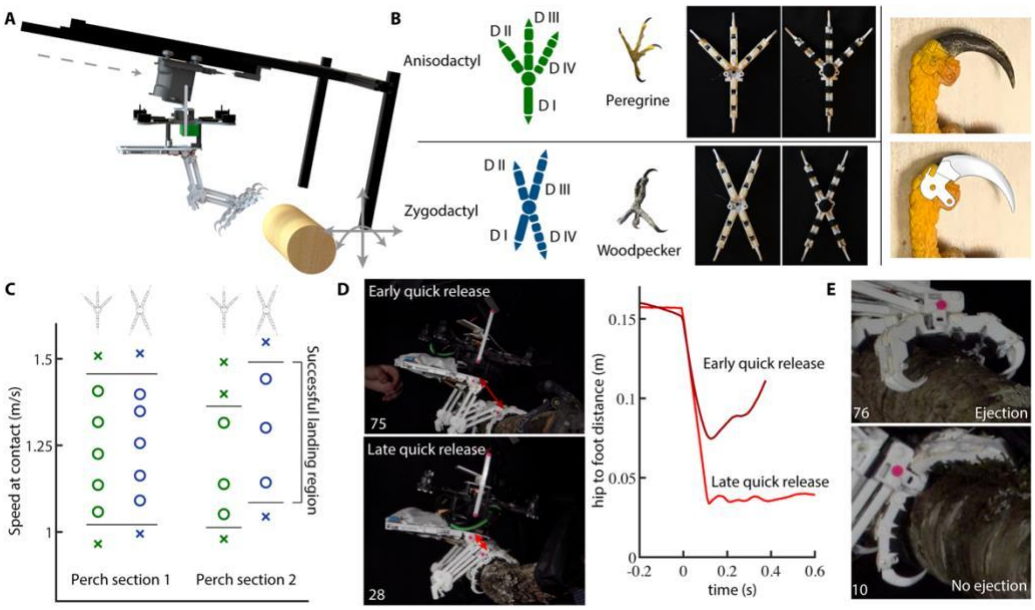

Figure G3 Experimental investigation of perching sufficiency: The effect of toe arrangement, quick-release trigger timing, and foot ejection

Results

To study how different mechanism designs and impact parameters affect perching, we constructed a setup that could produce consistent collisions between the robot and a perch. The robot was mounted to a slide on a rail that mechanically constrained the robot to have only one degree of freedom until the slide mechanically triggered the robot to be released a few millimeters away from contact with the perch (Fig. G3A). The robot was propelled by an elastic band stretched to a recorded distance to precisely control the impact speed. The rail could be rotated to test different velocity directions and translated to test different impact angles.

Focusing on the peregrine falcon anisodactyl toe arrangement and closed-loop balance control for the rest of the experiments, we study the effects of leg orientation and foot placement on the perching sufficiency region. At first foot-surface contact, there are three primary parameters that influence perching performance: the leg angle (leg), the impact angle (impact), and foot misalignment (foot) (Fig. G3A). If the foot misalignment is too high, then the toes will curl too far before contacting the perch, resulting in a failed landing (Fig. G3E). Higher leg and impact angles will result in more angular momentum over the center of the perch. If the angular momentum is too high, then the robot fails to land successfully and violates the upper angular momentum constraint in our model. Both the magnitude and the direction of the robot's velocity

on contact shape the perching sufficiency region (Fig. G6B). If the robot approaches the perch at shallow velocity angles, then the robot violates the angular momentum constraints of the sufficiency region model and fails to land experimentally when the impact speed is too high or too low. Too low a speed may also cause the legs to not collapse fully, which increases the pitch-back moment and can result in failure. To account for this effect, our model places a lower bound on the acceptable linear momentum. On the other hand, if the robot drops vertically onto the perch, then gravity is sufficient to collapse the leg fully even at an initial velocity of 0 m/s. In this case, the robot gains speed as the leg collapses. If the robot lands too hard on the perch, however, then components can break. Therefore, our model also incorporates an upper bound on the linear momentum, which we did not probe experimentally. In addition to kinematic, hardware, and behavior tradeoffs, the properties of the perch itself can change the size of the perching sufficiency region (Fig. G4C). We tested three diameters of an oak tree: 38 mm (1.5 inches), where the feet can wrap around most of the perch; 64 mm (2.5 inches), where the feet wrap about halfway around the perch; and 165 mm (6.5 inches), where the feet wrap less than a quarter of the way around the perch. Parrotlets can accommodate perches of similar proportions relative to their feet (1). SNAG can land on each diameter. On small-diameter perches, SNAG's feet can fully wrap the surface, causing more tendon length to be drawn from the feet. Consequently, leg stiffness and grip force are reduced. On the other hand, larger diameter perches, which tend to have larger asperities, prevent the foot from closing. This amplifies the leg stiffness and the grip force. We also tested a medium diameter alder tree (Alnus) with relatively smooth bark and some moss on top. Smoother bark makes the claws more likely to eject the foot, but we found that SNAG successfully lands on this slippery perch. Further, although many tree branches are nearly horizontal in nature, there are also many that are angled in a vertical plane. Our experiments demonstrate how the independent passive energy absorption

from each leg enables SNAG to accommodate these angular variations (Fig. G4D). However, a small amount of foot misalignment can cause the robot to fall.

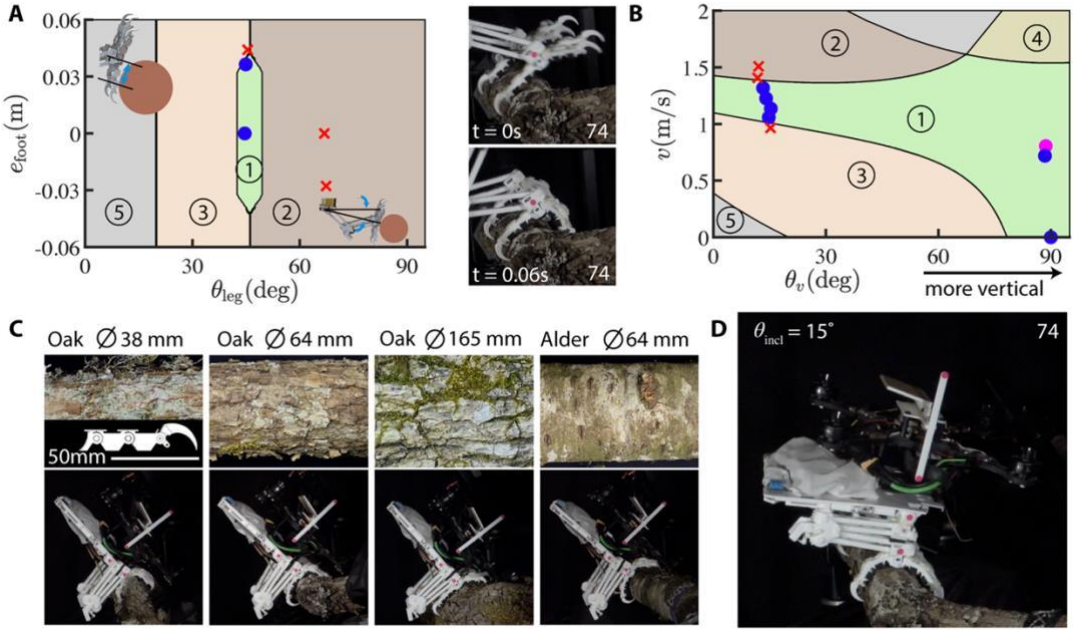

Figure G4 Experimental investigation of perching sufficiency: The effect of leg orientation, foot placement, impact velocity, and surface conditions

Last, SNAG is able to catch a wide variety of objects harnessing the same hardware as for perching (Fig. G5). Catching is, in many ways, analogous to perching. The main differences are the velocity of the grasper and the object, the textures of the object, and the forces imparted on the grasper. Our laboratory experiments show that SNAG can catch and release objects of similar size and weight to the prey of peregrine falcons (60) using peregrine falcon–inspired feet (Fig. G5). We also found that SNAG can catch objects during outdoor flight (Fig. G1D).

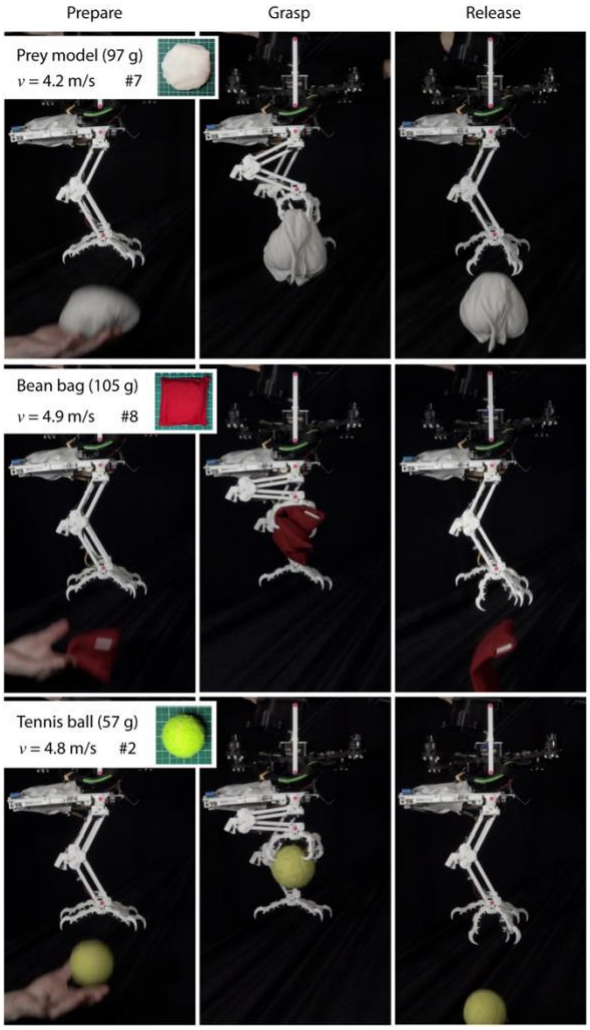

Figure G5 Experimental investigation of SNAG's catching ability

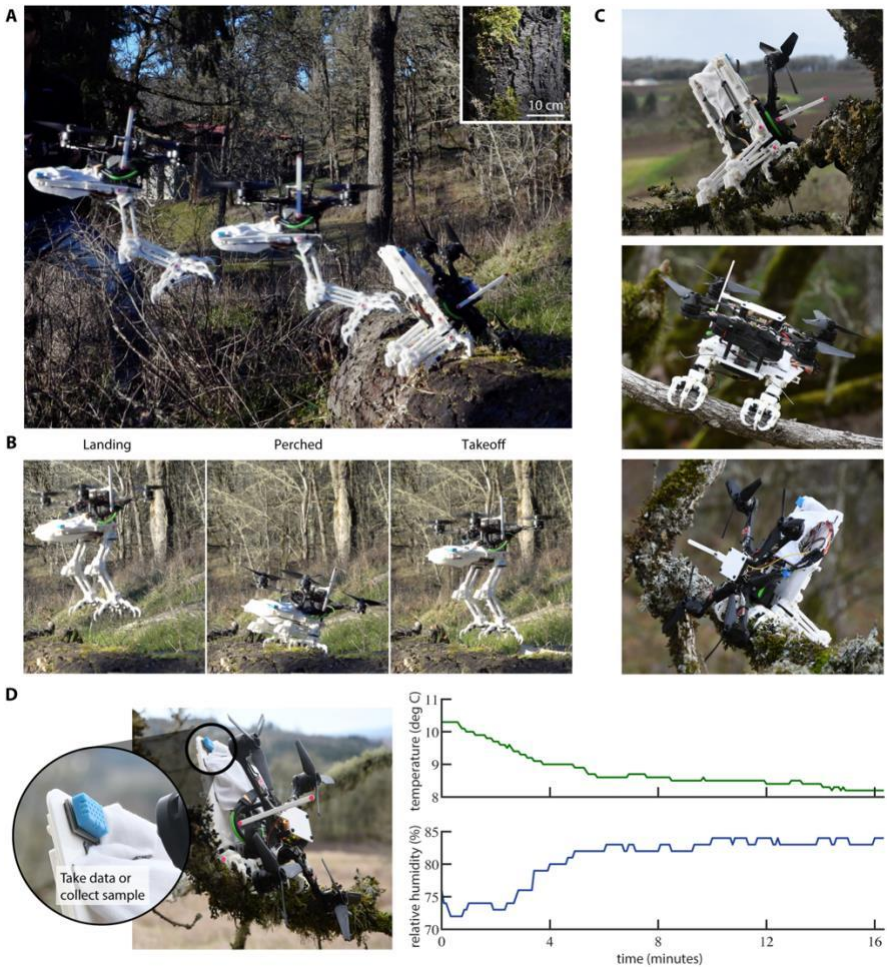

Figure G6 SNAG enables environmental monitoring in forests

Beyond The Lab

To demonstrate how SNAG can be used to enable quadcopters to monitor the natural environment at low energetic cost, we tested its perching performance in a forest (Fig. G6, A to C). For these tests, the robot was manually flown by the pilot with a remote controller (RC). We dynamically perched SNAG on the surface of a Douglas fir (Pseudotsuga menziesii) during a bird-inspired near-horizontal landing (Fig. G6A). Despite high foot error causing the front toes to curl prematurely, the robot perches successfully, demonstrating its robust bird-like performance. Tests with a vertical approach show that it can take off from the ground, land on a tree surface, rest perched on the surface, and take off to return (Fig. G6B). While perched, SNAG can characterize the microclimate by recording temperature and humidity for environmental research (Fig. G6, A to D).

In Flight Aerial Grasping Quadrotor and Mechanism

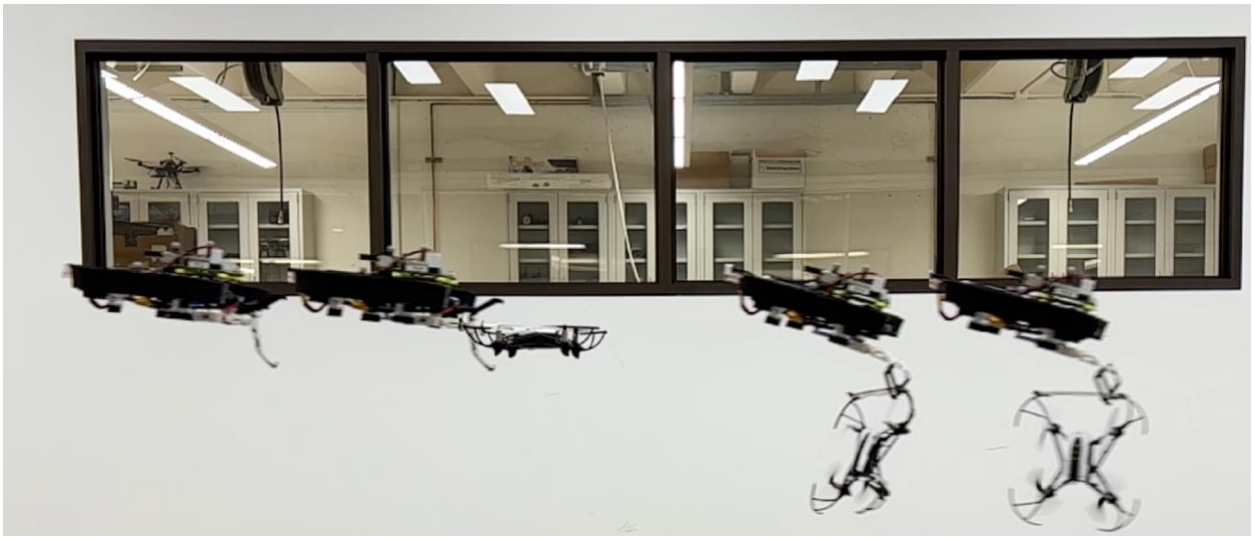

Figure A1 Capture sequence of the aerial grasping robot.

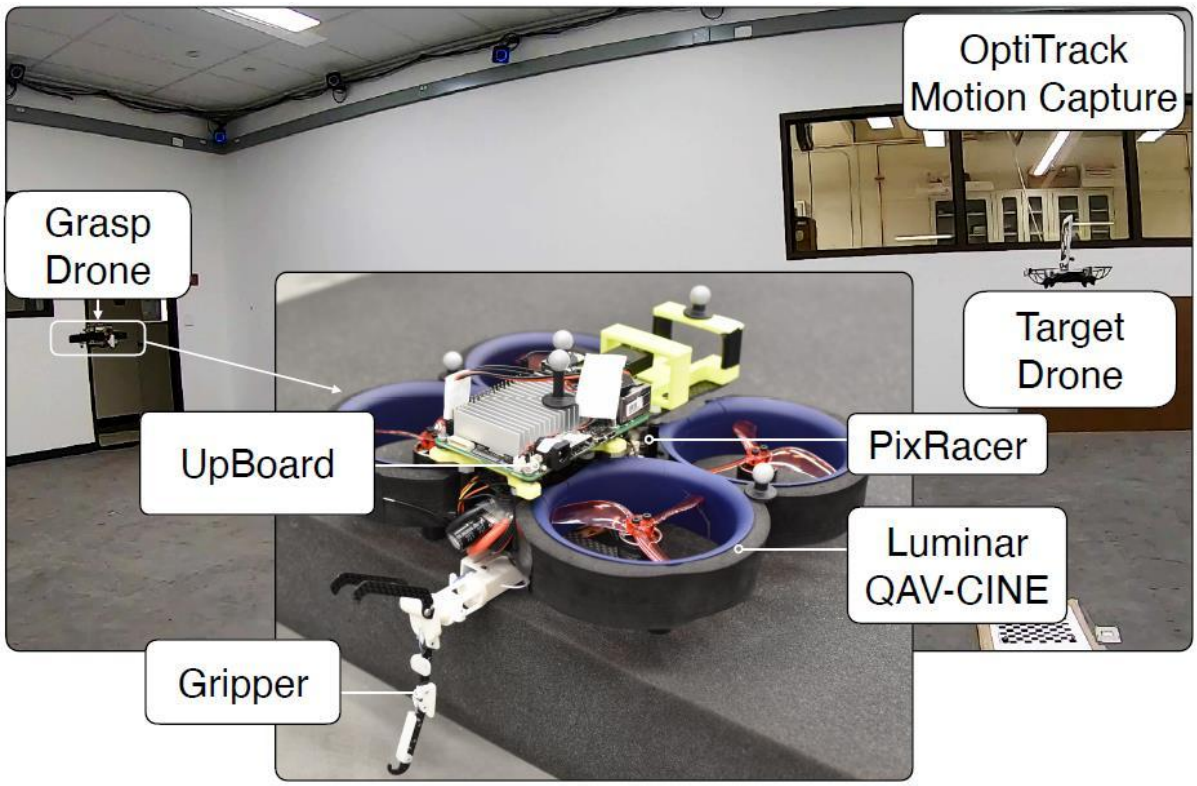

Figure A2 The grasping quad platform and flight test arena.

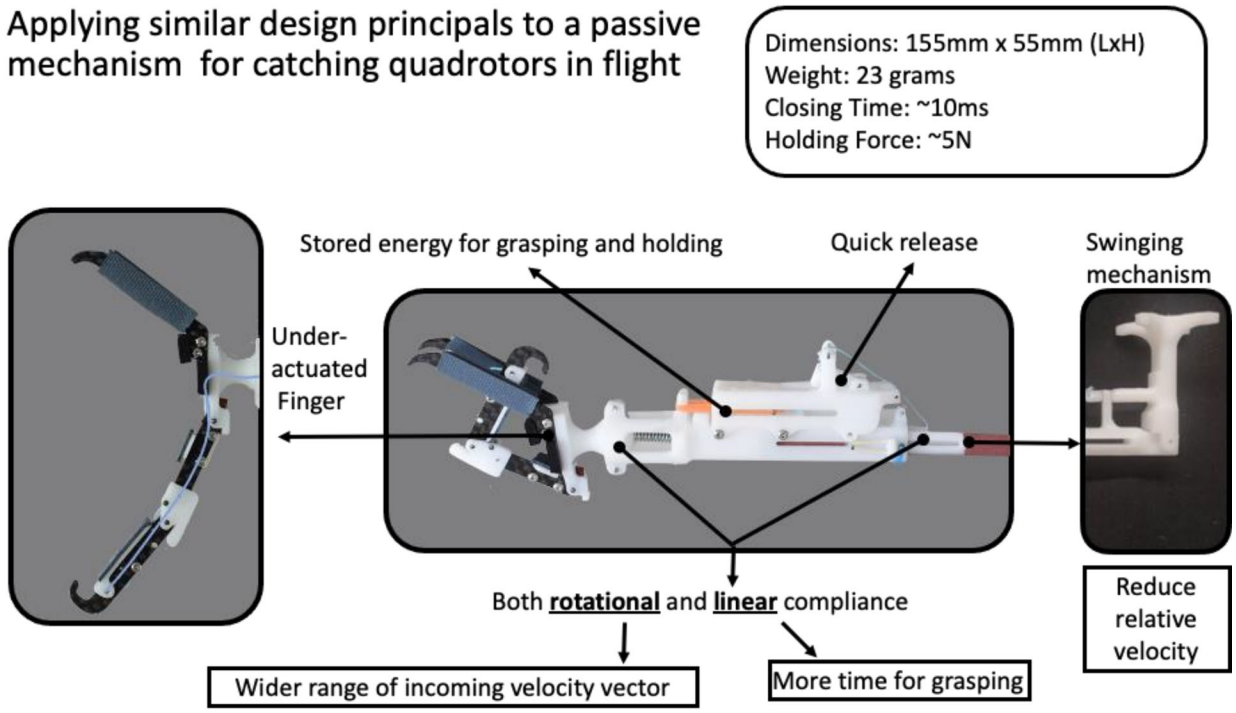

Figure A3 Mechanism overview of the gripper on the aerial grasping robot.

Based on some of the key principles learned from SNAG, we developed a grasping gripper to be used with a quadrotor to capture a flying target (Figs. A1, A2, A3). The gripper has the following features to accomplish this task: (1) The gripper is lightweight (~25 grams), which is important as the payload of a quadcopter is limited. (2) The gripper is completely passive, thus has a very fast closing time (~10ms) upon impact with the target. This enables a large range of relative velocity between the target drone and the main drone, up to 3 m/s. (3) The gripper has two configurations, rigid and swinging after capture, mimicking the natural motion of a falcon after catching a prey to minimize the disturbance to system stability. Testing this gripper on a force-torque testing platform (ATI), along with theoretical modelling, yields a sufficiency region, similar in concept to the SNAG. However, most drone flight controllers operate in velocity control mode rather than force/torque control, so experiments on a custom drone launching platform were conducted to translate this sufficiency region from force space to velocity space. This enables exploration of the drone linear and angular velocities required at the point of impact, including upper and lower bounds in all three directions, to be specified for the drone trajectory controller.

This, in combination with vision-based target tracking detailed in the latter section, form the basis for successful aerial capture of flying targets with quadrotors.

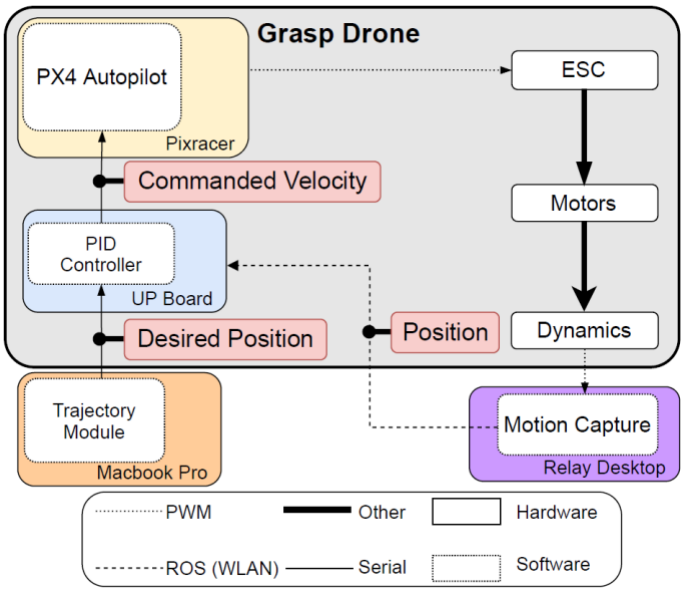

Figure A4 High level control system design for the aerial grasping robot.

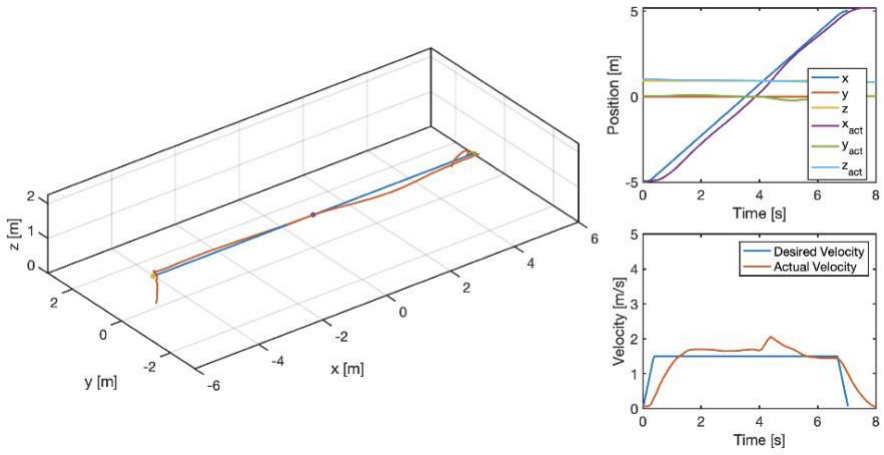

Figure A5 Example flight test scenario. A combined position and velocity signal is sent to the PID controller.

II Autonomous and Rudderless Aerial Vehicle with Morphing Wings and Tails:

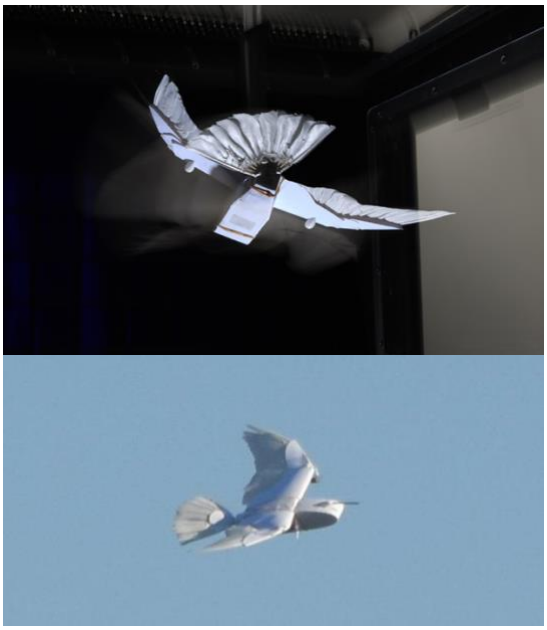

Birds such as falcons and pigeons can achieve unparalleled agility during flight without conventional airplane control surfaces and vertical tails. To study this, we applied recent findings on how birds underactuate wing feather motion to a new biomimetic flying robot modeled after the common pigeon (*Columba livia*). The new robot adds a morphing rudderless tail, in total underactuating 52 feathers via 8 active degrees of freedom. Through a variable-turbulence wind tunnel experiment, we found that morphing-based roll control is most robust when wing- and tail-morphing degrees of freedom are synergized in concert. In addition, we found gain scheduling each control input for the gross morphing configuration important for stability without a vertical tail. We successfully applied our wind tunnel-tuned feedback and feedforward control scheme to free flight, demonstrating how our learnings in the lab extend to real-world flows. This work expands our understanding of highly dynamic bird flight control, and how to extend such control to small rudderless UAVs.

A New Autonomous Rudderless Morphing Flyer Platform

To study how avian tails contribute to feather-based morphing flight control, we overhauled our previous PigeonBot design to include a morphing tail with four actively controlled degrees of freedom. While incorporating the morphing tail, we also made significant fuselage design changes to reduce the tail moment arm to match a real bird as much as possible. In PigeonBot 1, we used a conventional horizontal and vertical tail with an unrealistically large moment arm to the center of gravity to guarantee flight stability. However, we aim to use PigeonBot 2 for bird flight stability and control studies, which requires moving the tail forward. Fig. PB1 depicts the new vehicle planform overlaid on the planform of a stock dove (*Columba oenas*) from (6), which is in the same family, Columbidae, as the pigeon PigeonBot 2 is designed to mimic.

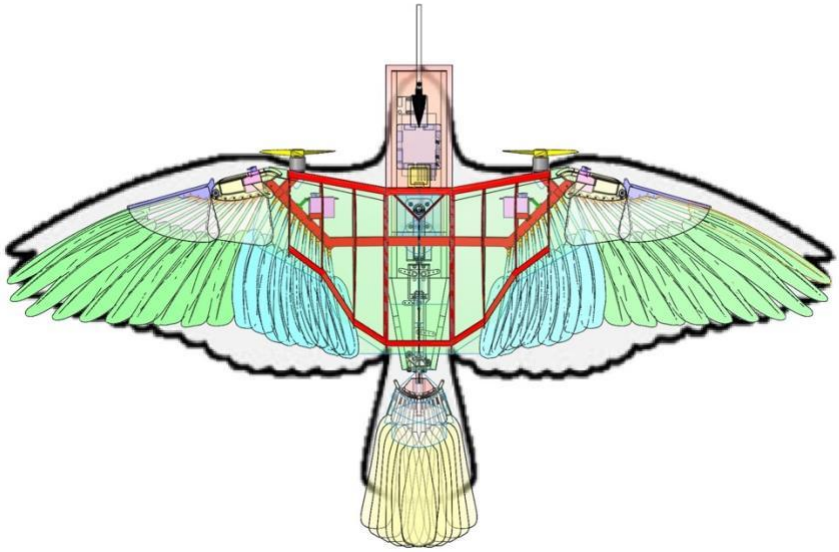

Figure PB1: PigeonBot 2 planform versus stock dove. PigeonBot 2's entire vehicle planform closely matches that of a stock dove (*Columba oenas*) from (6), which is in the same family, Columbidae, as the pigeon PigeonBot 2 is designed to mimic.

The final fabricated PigeonBot 2 consists of 52 real pigeon feathers: 40 wing feathers and 12 tail feathers (Fig. PB2). The wing has the same four actively controlled degrees of freedom as in PigeonBot 1, while the tail has the same four actively controlled degrees of freedom as in the morphing tail robot. Together, this robot represents the most comprehensive mechanical bird model to date and is additionally designed to be capable of untethered flight with a flying mass of 300 g, below the 316-358 g weight of real pigeons (7,8) and a center of gravity at the same location of Columbidae (9).

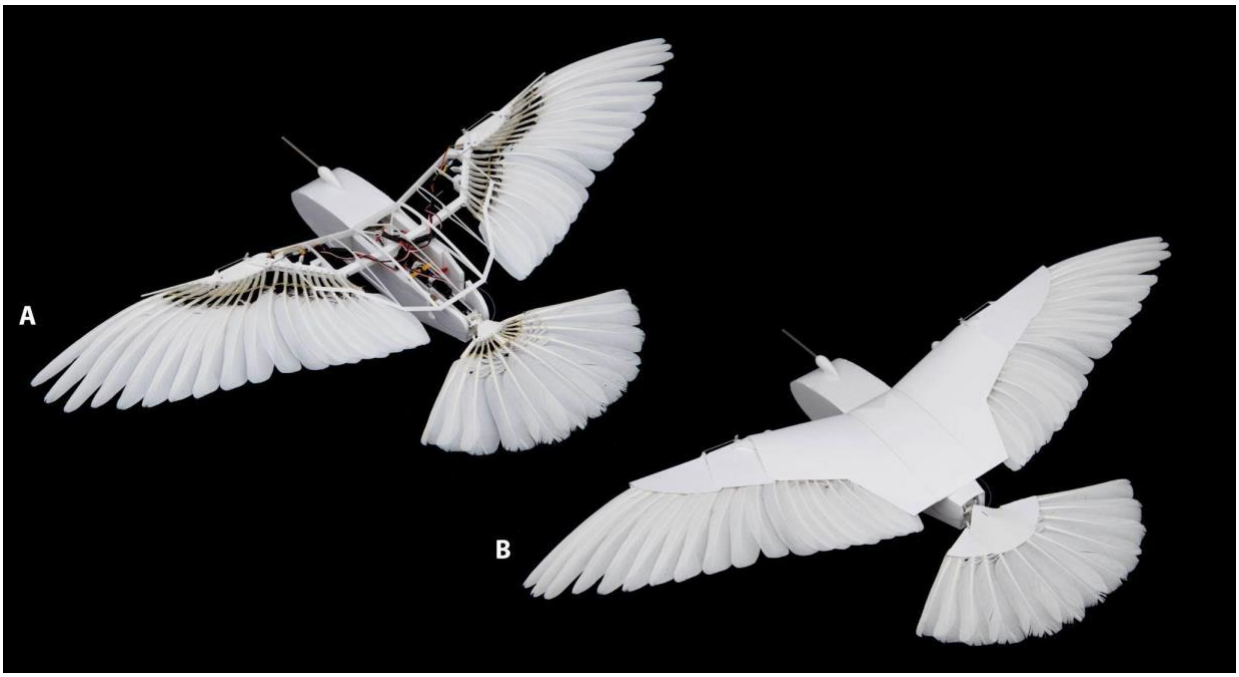

Figure PB2: Rudderless Morphing Flyer, PigeonBot 2, assembled, shown without rachis coverings (A) and with rachis coverings in flight configuration (B). In total, PigeonBot 2 has 52 feathers, actuated by eight actively controlled degrees of freedom, and weighs 300 g.

Rudderless Morphing Tail Controller

Following hardware prototyping, we designed and tuned the onboard controller to achieve stable flight in the wind tunnel. In order to have the hardware capability to control nine servos independently, we implemented a microcontroller (Teensy 4.0) as an 'output mixer' (Fig. PB3). In this hardware control scheme, we still used the PixRacer for sensor fusion and high level pitch/roll/yaw control, but we then send these high level rotational controls to the output mixer to be mapped to the appropriate servomotor. This configuration also enables easier implementation of more complex feed-forward terms to compensate for known dynamics such as adverse yaw when using asymmetric morphing for roll control (3). In this example, we can feed-forward lateral tilt to generate a yawing moment to counter the adverse yaw without relying on feedback control.

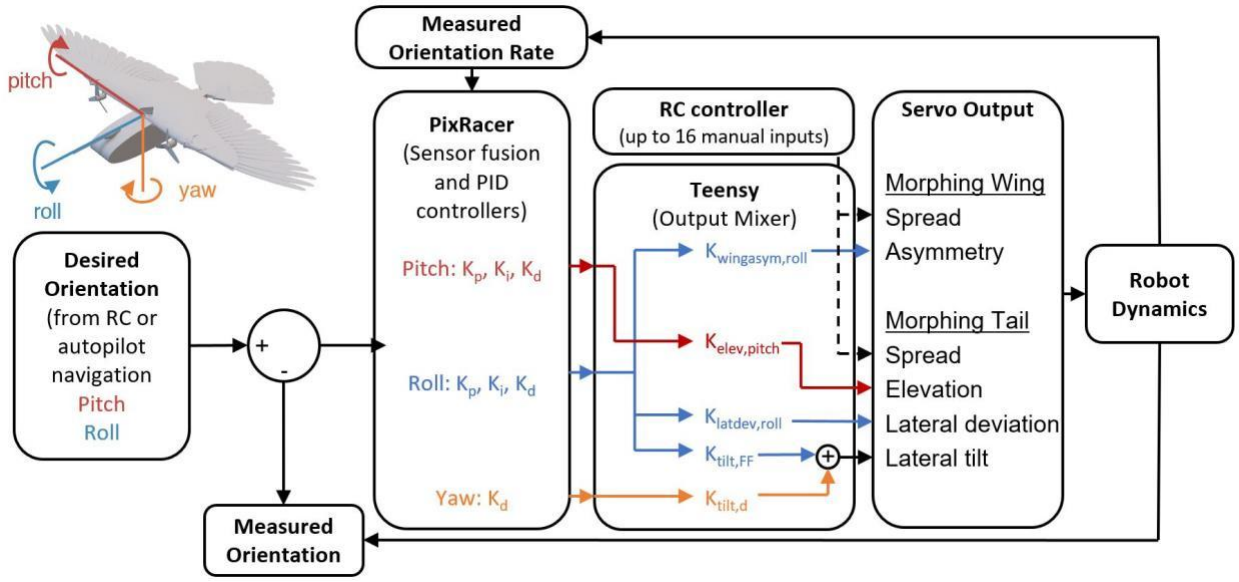

Figure PB3: PigeonBot 2 control block diagram. To control PigeonBot 2's eight active degrees of freedom, we simplify its 4 wing degrees of freedom into spread and asymmetry commands. We then apply our flight test results from PigeonBot 1 to implement feed-forward control terms that will result in favorable flight handling. One example of this is implemented with $K_{tilt,FF}$, which automatically mixes in lateral tail tilt during wing asymmetry to compensate for adverse yaw.

The resulting implemented control system performs similarly to birds when perturbed in the hand, in which a synergy of wing and tail movements actuate to compensate for roll, pitch, and yaw deflections (10,11). This simple demo shows the scientific potential for such a platform, in which one could program our robot to exhibit the same wing and tail morphing responses with appropriate gain and phase as measured in real birds, and evaluate if such a control system alone is sufficient for stabilizing bird flight.

Control and stability testing in the wind tunnel and outdoors

To test the performance of the morphing-based control system, we evaluated the rotational dynamics of PigeonBot 2 in a variable turbulence wind tunnel mounted on a ball-joint located at the robot center-of-gravity. We found an inherent Dutch roll instability mode without a vertical table that was able to be suppressed with a yaw damper mapped to lateral tilt and wing asymmetry mapped to a roll feedback controller. To evaluate robustness, we generated variable turbulence

intensities and evaluated PigeonBot 2's ability to maintain stability in turbulence (Fig. PB4). Using our morphing controller, PigeonBot was able to remain stable up to turbulence intensities of 20%, despite being perturbed to roll and yaw angles of ~30°. While this result is not sufficient for flying in the worst atmospheric turbulence intensities in cities that can be up to 40\% (12), it is a promising step toward developing small flying vehicles that can approach the stability of real birds.

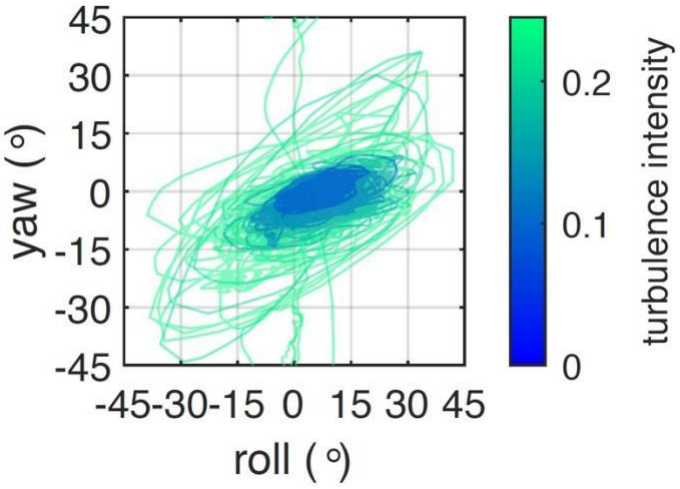

Figure PB4: Active morphing control enables robust stability in turbulence. PigeonBot 2 is capable of stabilizing in turbulence intensity levels up to 20%, encountering aerodynamic perturbations that destabilize PigeonBot 2 up to ~30° in roll and yaw.

We also evaluate PigeonBot 2's response to desired sinusoidal trajectories (Fig. PB5). We generated a sequence of roll inputs with gradually increasing frequency at an amplitude of ~25°. PigeonBot 2 is capable of tracking these large roll angles up to a frequency of 1 Hz, producing roll rates up to 160°. This roll rate is on par with fighter aircraft such as the F-18 that can achieve ~120° and the F-16 that can achieve ~240°. PigeonBot's roll tracking performance is minimally affected by gross wing span, as tests with the wing and tail in an intermediate pose versus at fully spread poses yielded similar bandwidth results. PigeonBot 2 uses only ~20% of maximum wing asymmetry, or ~20° of wrist angle asymmetry during these tests, which suggests that the falloff in performance is caused by insufficient speed of the hobby-grade servomotors used. As actuators continue to gain performance for the same weight, vehicles that use similar control strategies will be able to achieve even more aggressive maneuverability while remaining stable in extreme turbulence.

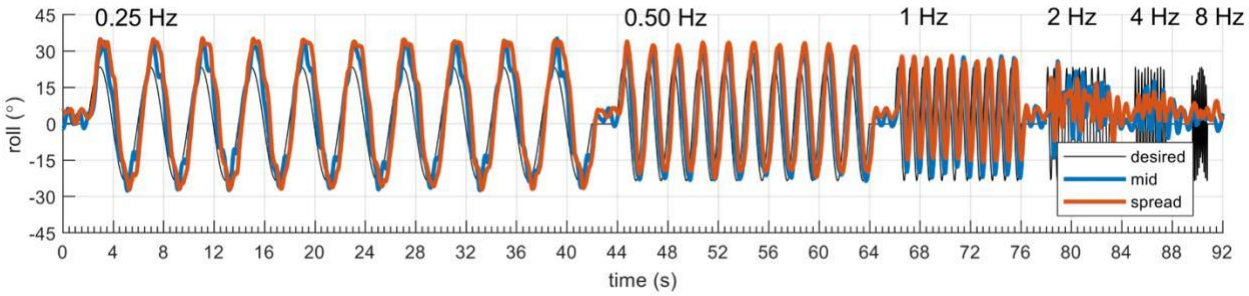

Figure PB5: Morphing-based roll control up to 1 Hz. By utilizing many morphing degrees of freedom, PigeonBot 2 can track desired large roll angle trajectories up to a bandwidth of 1 Hz with its wing and tail in an intermediately spread planform, and at fully spread planforms.

We then implemented a gain scheduling scheme that interpolated the best control gains for each discrete morphing configuration tested in the wind tunnel across the continuous morphing configuration space on PigeonBot 2, and deployed the robot in untethered free flight in outdoor flow (flight video). We programmed PigeonBot 2 to autonomously fly circular "loiter" flight patterns while grossly morphing the wing and tail into different configurations, and allowing the control system to automatically apply the appropriate controls to maintain the desired flight trajectory. While a more detailed flight study analysis is underway, we were pleased to find that our learnings in the wind tunnel translated to real world flow, and PigeonBot 2 was able to achieve fully controlled and stable rudderless flight using only feather morphing for control.

III Fast Trajectory Planning for Aggressive Maneuvers and Simulation

Existing methods for quadrotor trajectory planning can be divided into two classes: (i) differential flatness based methods, and (ii) trajectory optimization methods. Differential flatness methods are computationally fast, but they can only indirectly satisfy state and input constraints. Trajectory optimization methods can directly handle state and input constraints, but are too slow to solve online in using Model Predictive Control (MPC) methods. We explored a hybrid approach that uses differential flatness as a warm-start to an Augmented Lagrangian iLQR trajectory optimization algorithm. We introduced several modifications to the iLQR method to yield real-time re-solve performance. Our method satisfies state and input constraints, gives a closed-loop feedback policy rather than an open-loop plan, and is fast enough to re-solve online in an MPC loop with onboard compute resources. Our method can also solve minimum time trajectory problems.

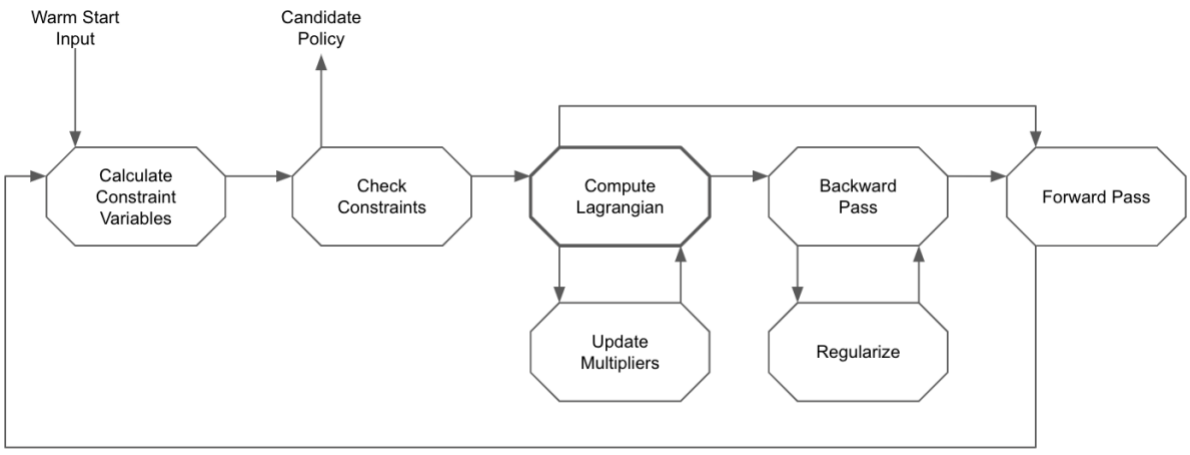

Figure T1 Flowchart of the trajectory solver

The approach is illustrated in the diagram above:

1) An initial warm-start trajectory is generated using differential flatness-based methods. These can have unsatisfied state and input constraints.
2) The constraints are computed along the entire trajectory
3) The constraints are checked along the entire trajectory. If satisfied, we output a trajectory (using what was computed in the backward pass).
4) We compute the Lagrangian. This encodes both the objective and constraints of the trajectory through weighted terms.

5) A trajectory with feedback is generated using the Riccati equation in the backward pass.
6) The trajectory is then rolled forward in time using the forward pass.

The constraints are then re-computed and re-checked in an iterative process until they are satisfied. Note the importance of the multiplier update step, which is used to increase the weights of the constraints until convergence, and the regularization step, to ensure updates on the non-linear dynamics and constraints remain reasonable.

Results

An example solution where a drone is tasked to fly through a small gate (state constraints) while not saturating motor commands (input constraints) can be found in figure T2.

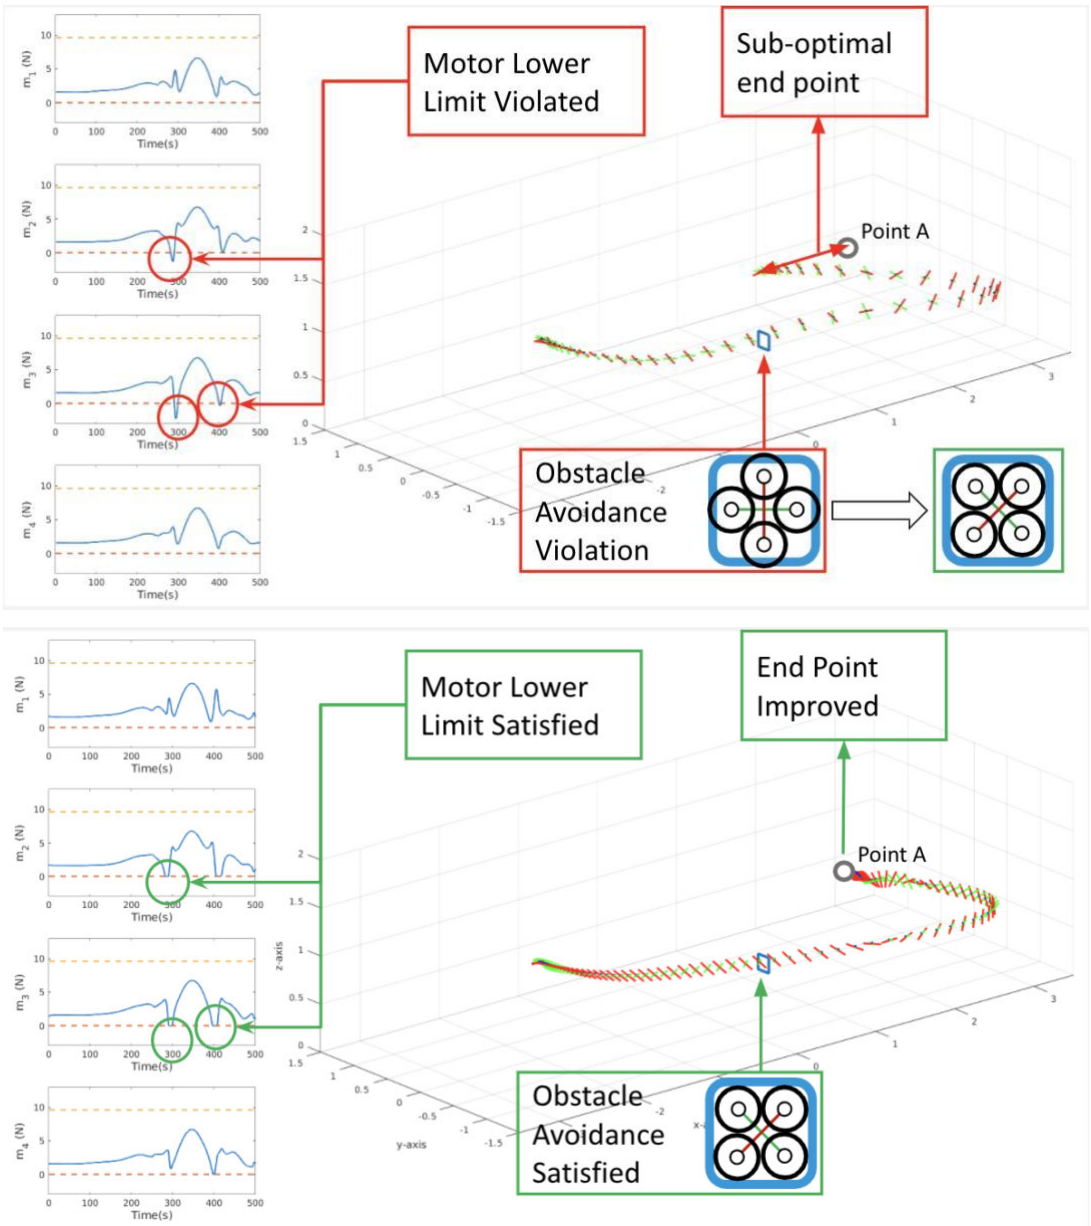

Figure T2 Comparison of the performance of the solver.

A comparison of the solve times (MATLAB) with differential flatness (DF) and AL-iLQR can be found below:

Method	Solver Outcome	Simulation Outcome	Solve Time
Pure Differential Flatness	solved	crash (gate)	0.114s
Pure AL-iLQR	solved	success	14.427s
DF + AL-iLQR	solved	success	0.643s

Similar results were observed in two other trajectories:

Horizontal Gate				
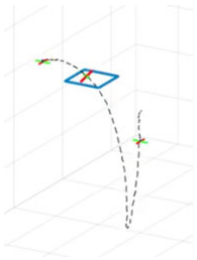	**Method**	**Solver Outcome**	**Solver Outcome**	**Solve Time (MATLAB)**
	Pure Differential Flatness	solved	crash (motor lim.)	0.115s
	Pure AL-iLQR	failed	n/a	n/a
	DF + AL-iLQR	solved	success	6.506s

Vertical Slit				
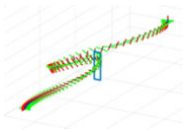	**Method**	**Solver Outcome**	**Simulation Outcome**	**Solve Time (MATLAB)**
	Pure Differential Flatness	solved	crash (gate)	0.121s
	Pure AL-iLQR	failed	n/a	n/a
	DF + AL-iLQR	solved	success	6.315s

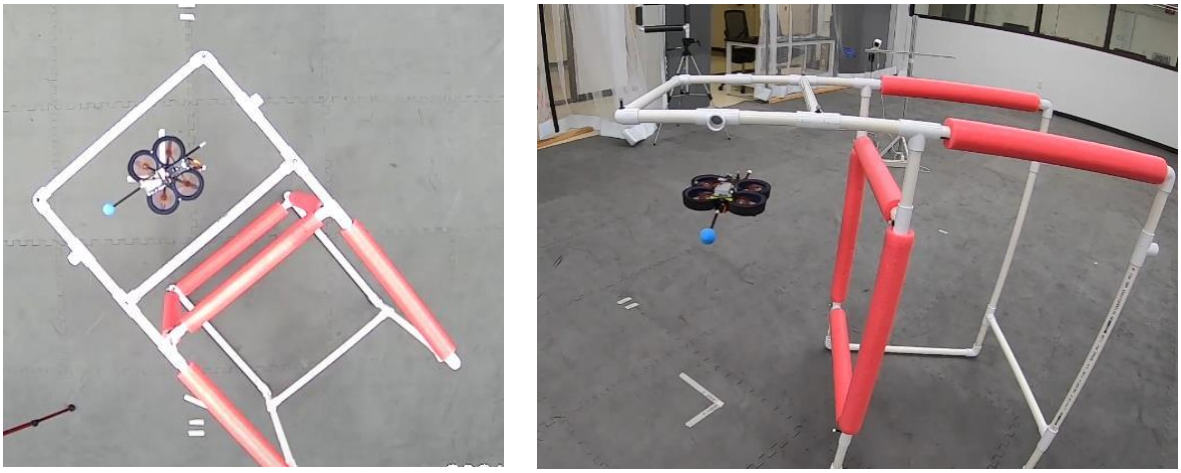

Figure T3 Flight testing of the solver flying through obstacle constraints.

Experimental verification was performed on simple gate pass-throughs like the one illustrated above where the drone must fly through a horizontal gate in a specific orientation (or risk collision). It was found that while the experiments mirror simulation results for low-speed trajectories, it does not carry over to high-speed ones. The cause of this is that the linearization points along the trajectory can be highly sensitive to deviations; successful tracking using the feedback policy only works within certain tubes along the trajectory. High-speed trajectories is one situation where the volume of these tubes is severely decreased. Likewise, flips and spins were found to have similar problems. Current work is now focused on trying to determine the exact regions of these tubes along the trajectories.

Morphing Wing Robot Simulation

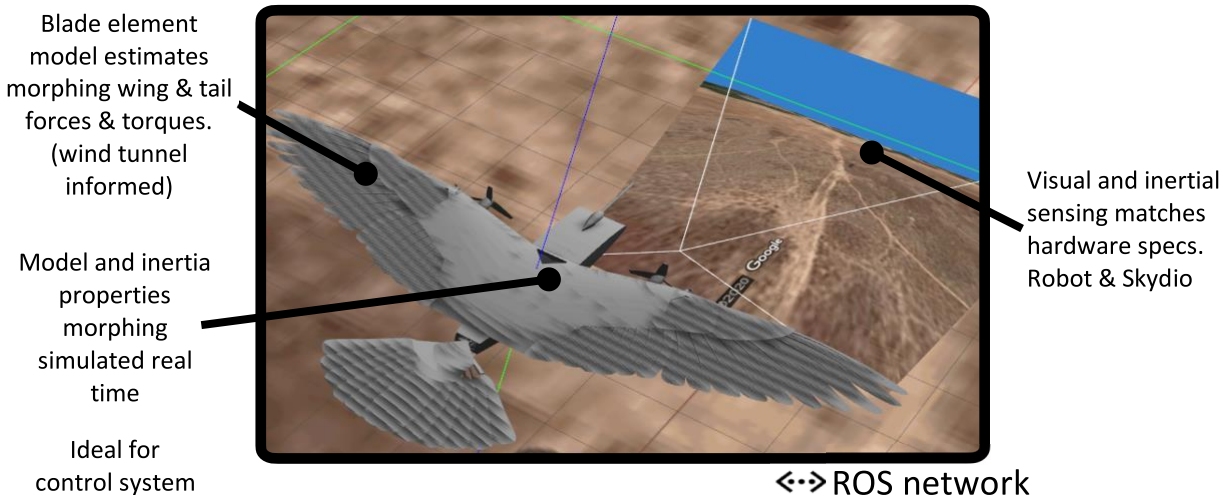

Figure C1 Visual depiction of the simulation infrastructure and environment used for control development.

To enable control development during the project, we developed an end-to-end simulation of the morphing wing robots presented earlier. We estimate the lift and drag forces of the aerodynamic surfaces in MATLAB using a blade element model. The total force acting on the robot at each time step is estimated, sent to Gazebo (Simulation environment for robotics), and then the simulation is iterated forward. This environment provided the ability to test different hardware and software configurations before testing on the real hardware platforms.

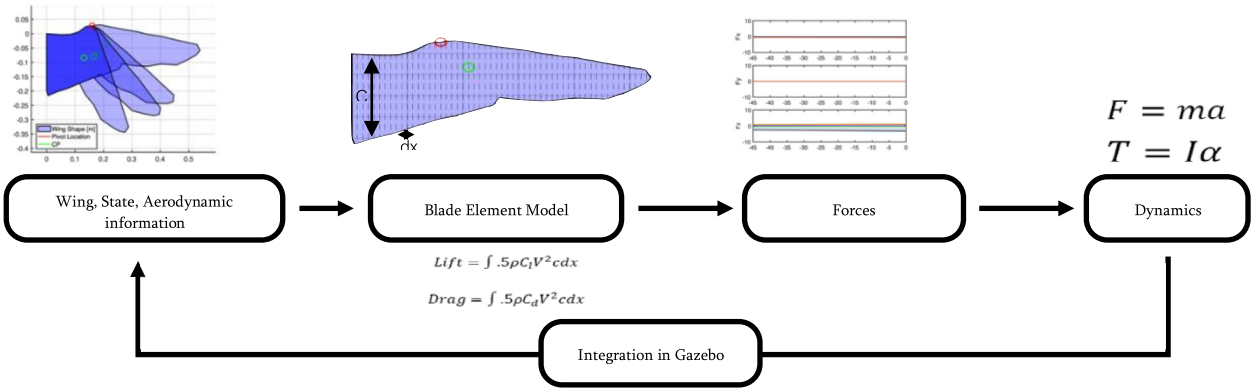

Figure C2 Simulation loop used to compute the aerodynamic forces acting on the morphing wing robot.

IV Vision based localization, Tracking and Control for UAVs:

Figure S1 The integrated Skydio X2 vehicle with the Stanford gripper and other necessary modifications.7

This program marked close collaboration between the Stanford team and Skydio's autonomy and hardware teams. Notably, we successfully integrated the aerial grasping mechanism onto the Skydio X2 vehicle, leveraged Skydio's Autonomy Engine for the grasping task, and incorporated two control and planning concepts for aerial grasping. We are integrating our platforms and control with Skydio's Autonomy Engine to create a pipeline to visually detect, track, and filter the positions of target drones. We are also incorporating Skydio's motion planning system, which synthesizes the desired trajectory with obstacle avoidance. This approach supports autonomous missions by avoiding non-target obstacles. Data collected from the heavily instrumented Skydio X2 platform provides information useful for tracking the dynamics of an aerial grasping action with abrupt dynamic forces and changes in inertia. We can estimate the collision forces to assess quickly whether a grasping attempt is successful. The superior Skydio vision capabilities, even with respect to most off-board motion capture systems, (e.g. 1 gimbaled subject camera + 6 navigation cameras) enable consistent detection and subject tracking as compared to a monocular or stereo camera configuration. Ongoing work includes improving the robustness of visual detections, dynamic pursuit, and use of more generalized visual detection schemes.

Integrating the Stanford Mechanism onto the Skydio Platform

Scaling the gripper designed by the Stanford team to the Skydio platform required solving various engineering problems. These included designing a larger mechanism, adding propeller guards, determining the new mass properties, and adjusting the vision masks. Following these adjustments, we successfully flight tested the X2 with a target captured in its gripper. This marked the first step in developing the Skydio X2 into an aerial grasping robot platform.

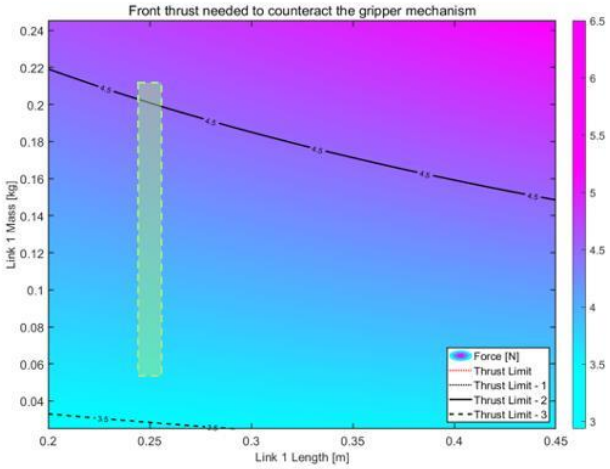

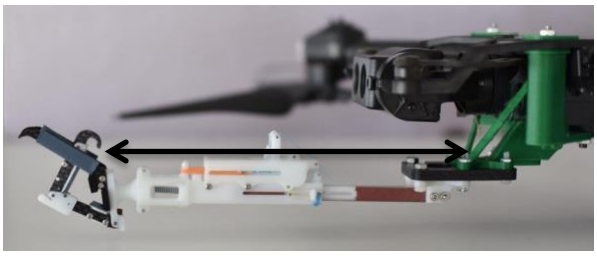

Figure S2 An example of the calculations to determine the gripper dimensions while staying within the capabilities of the mechanism.

Figure S3 The Skydio X2 flying with a captured target drone. This demonstrates the capabilities of the X2 to adapt to the aerial grasping scenario.

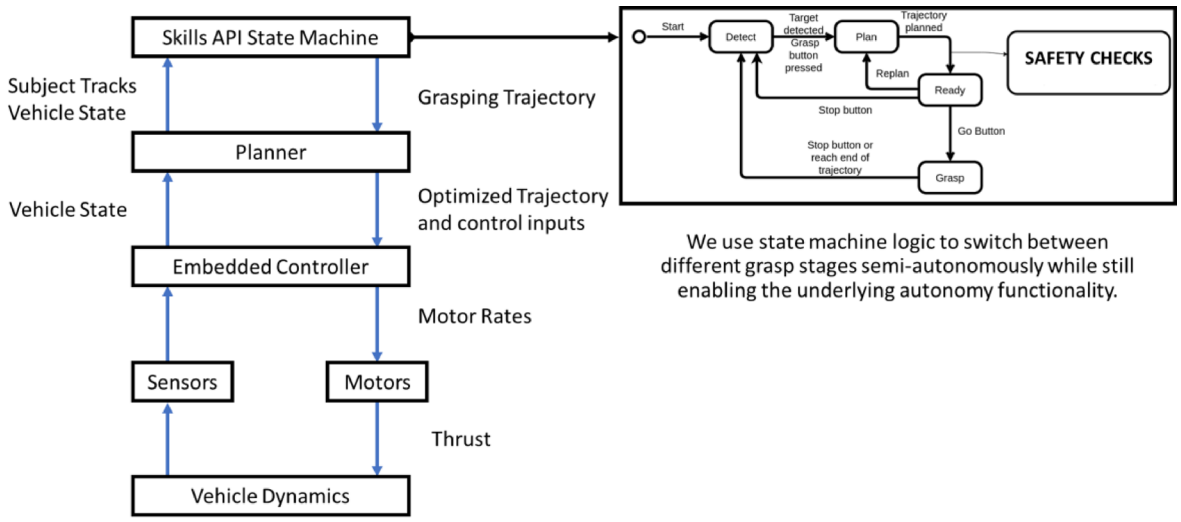

Figure S4 The high-level architecture of the Skydio motion planning system. The grasping skill sits on top of the planner stack.

Detection, Interface, and Control Design

The grasping task, "Skill", is built on top of Skydio's existing infrastructure (figure S4). This enables quicker development and iteration, without the need for extensive additional development.

For visual detection and tracking of the target drone, we use Skydio's subject tracker. This system visually detects, filters, and tracks the position of the target drone. Then, Skydio's Skills API enables us to use the subject tracker to plan and fly a trajectory. This pipeline enables other benefits, such as existing within the Skydio mobile app and being able to use its functionality. This includes augmented reality markers to show the user where the robot estimates the target is.

The grasping interface is shown in figure S5. Notably, the interface shows the target location, planned trajectory, distance to target, and gives an indication to the user what the robot is doing.

Figure S5 Example screenshot of the grasping skill in the Skydio app.

Offset Point Tracking

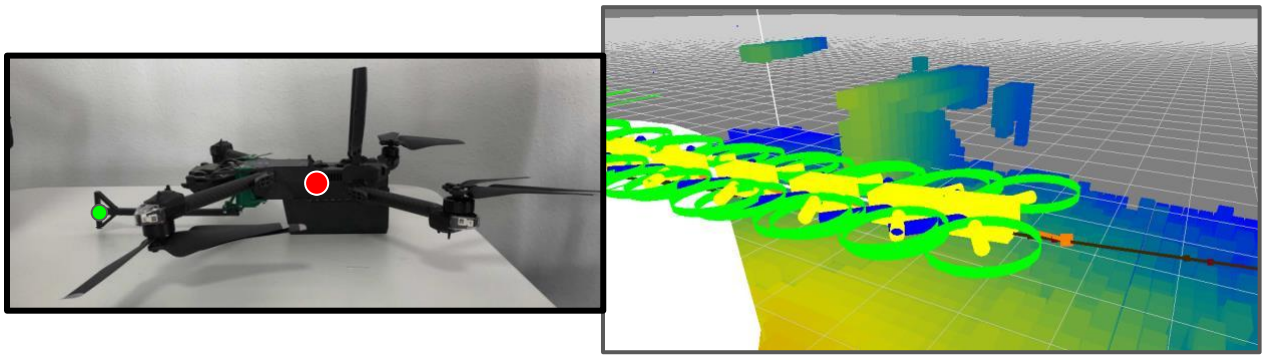

Figure S6 To enable aerial grasping, we needed to modify the motion planning system to track an offset body point (green) rather than the vehicle center (red).

Critical to the grasping effort is the ability to fly the opening of the gripper to the desired position, rather than controlling the vehicle center. To accomplish this, we redefine the trajectory tracking objective in the Skydio autonomy engine's planner. We define a body offset point, and measure the error in position, velocity, and acceleration to that point. Planning the grasping problem like this allows us to incorporate all other features of the planner, such as obstacle avoidance. We flight tested this planner change and found it enabled us to hit the target consistently with good trajectory tracking performance.

Figure S6 demonstrates the planner following an offset point. The orange depicts the opening of the gripper. Rather than the center of the yellow region following the brown trajectory, the orange point is, and therefore, the gripper is following the desired trajectory.

Trajectory Design for Aerial Grasping

Flying a trajectory along the offset body point enables significant progress towards the aerial grasping goal. The next key concept for aerial grasping is being thoughtful about the velocity profile of the trajectory. To increase the likeliness of the grasper being at the desired position, we design a constant velocity profile with a buffer region about the target. This attempts to keep the pitch as constant as possible, minimizing the error at the controlled collision point. This approach is not unlike Falcons, which follow smooth velocity trajectories in their swooping maneuvers.

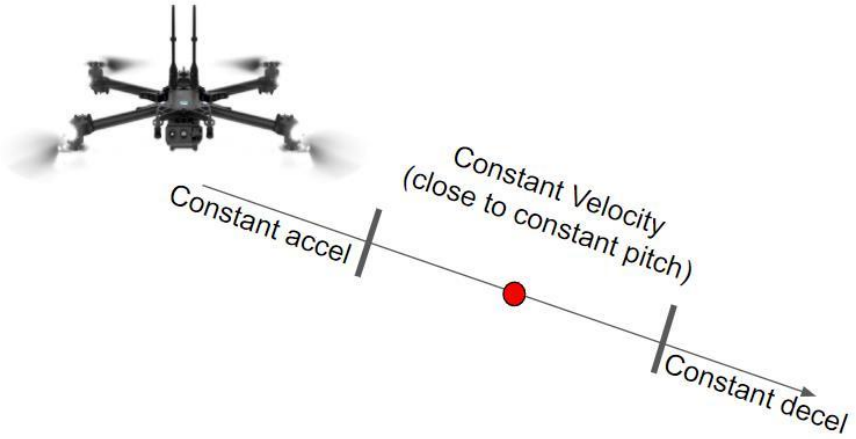

Figure S7 Mockup of the velocity trajectory the pursuit drone follows. A critical takeaway for the control of quick aerial grasping robots is to have a region of constant velocity surrounding the target.

Flight Test

Integrating autonomy, design, and control enabled significant flight testing of the integrated Stanford-Skydio Aerial Grasping Robot. We can consistently detect the target, plan a trajectory, hit the target, and fly with the added mass. Not only does this demonstrate the feasibility of the grasping problem using the Skydio platform, but also the capability of the Skydio vehicle.

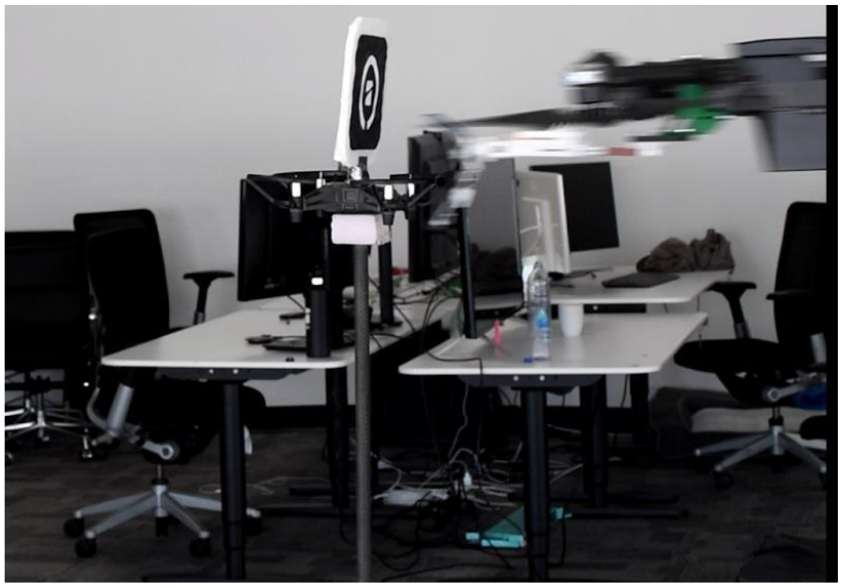

Figure S8 The X2 with gripper hitting the target.

PigeonBot Tracking

To detect and track the PigeonBot target drone, we expand Skydio's system for dynamic person and car following. The system is a combination of GPU image processing, deep learning, graph matching, and deep learning. It is tightly coupled with onboard state estimation and depth perception for 3D tracking. Finally, the overall architecture adapted to be target-agnostic, with the ability to train on additional objects an important open area for Skydio.

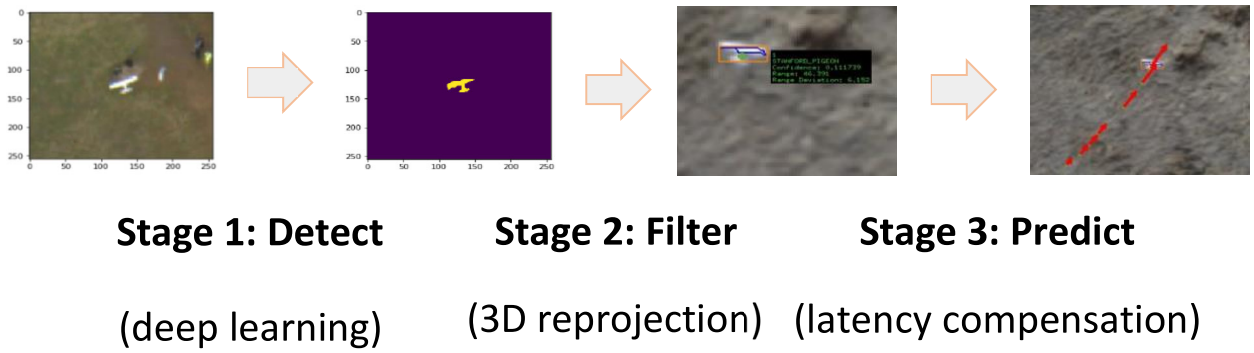

Figure S9 Deep learning tracking pipeline developed to detect and track the PigeonBot Target drone.

Full System Integration

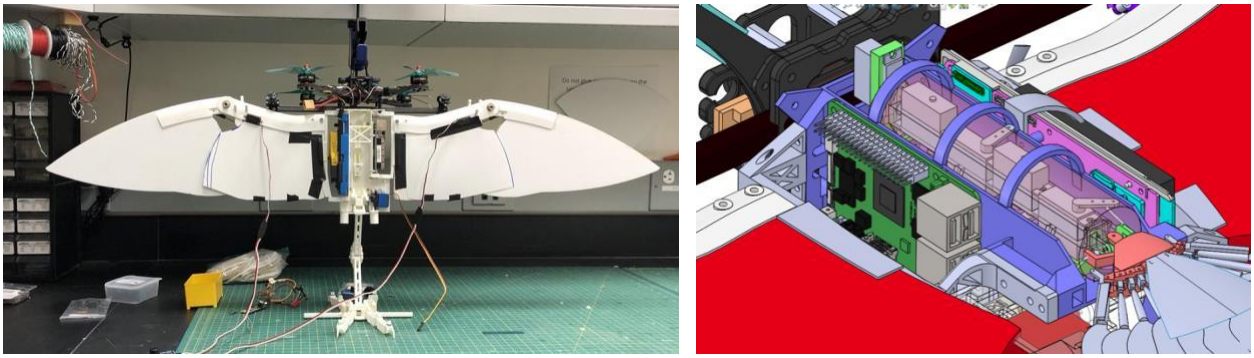

Figure S10 A prototype FalconBot drone with Skydio electronics incorporated into its design.

A key accomplishment of the project is the full system integration of the Skydio electronic hardware into a Stanford designed robot. This was done by using a development version of the Skydio 2 Main Logic Board (MLB) and two cameras on a stiff carbon beam. In doing so, we demonstrated the feasibility of expanding the Skydio platform beyond traditional quadrotor platforms.

Hardware

Skydio spent 2018-2020 designing, manufacturing, and optimizing a cutting edge flight computer for intelligent drones, which we refer to as the main logic board, or MLB. It is designed to support running multiple cutting-edge robotics algorithms in parallel, including deep learning and computer vision programs. The MLB is built around a Nvidia Tegra TX2 compute platform, as well as a number of auxiliary processors capable of performing real-time control tasks, and contains integrated support for various sensors, including inertial measurement units, a GPS transceiver, a barometer, and high-bandwidth links for up to seven high resolution color cameras. This flight computer is used in the Skydio 2, 2+, and X2.

For the DESI falcon vehicle, we designed a custom set of adapters to enable the MLB to operate in the configuration needed for the unique geometry of the bioinspired falcon form factor. This includes a power adapter board to integrate the MLB with the battery system of the falcon vehicle, as well as custom adapter boards to extend the length of the wires connecting the cameras to the MLB.

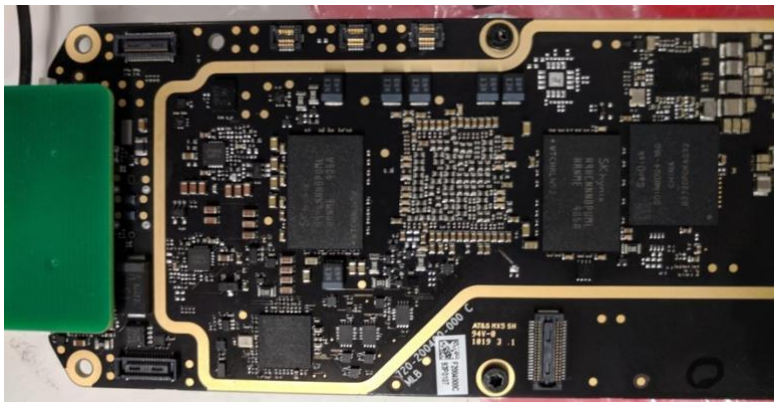

Figure S11 The Skydio MLB development kit

We installed this customized MLB in a test configuration onto a quadcopter platform that simulates the eventual intended sensor configuration of the falcon vehicle, with two parallel fisheye cameras pointing 45° down. After making the necessary adjustments to the software running onboard the MLB, we were able to successfully capture and store synchronized data from the two cameras, the inertial measurement unit, and the onboard barometer.

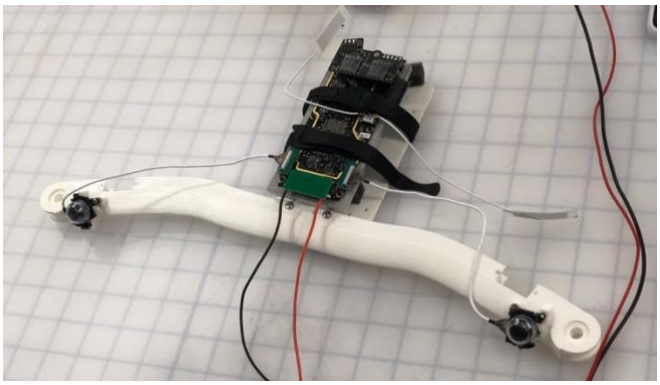

Figure S12 The development mule used to test the feasibility and rigidity of mounting the Skydio cameras on a fixed wing platform.

Calibration

For the advanced computer vision algorithms running on the MLB to reliably function, they need highly accurate calibration information for the cameras and other sensors mounted to the vehicle. Skydio has developed a custom system to generate this information and used this calibration system to analyze the quadcopter test platform described above. The resulting calibration information was used to run the vision algorithms described below.

Additionally, we manually created several necessary components of the vehicle calibration information for the test quadcopter platform that cannot be automatically inferred during the calibration process.

Localization

Skydio has developed an integrated localization system that uses measurements from all available sensors to precisely estimate the position, orientation, and velocity of the vehicle during the flight. During the last year, this system has been updated to work with data coming from rolling-shutter fisheye cameras like the cameras mounted on the falcon robot. Additionally, we have made significant improvements to support flight using GPS in environments where visual tracking is impossible, particularly over large bodies of water or at very high altitudes. We anticipate these improvements being critical for the falcon robot, given the large range of altitudes and environments that natural birds can operate in.

Using the data captured by the test platform pictured above, we have successfully adapted the Skydio localization system to operate successfully with the novel physical configuration of the falcon robot. The video below shows this system successfully tracking visual features to localize itself.

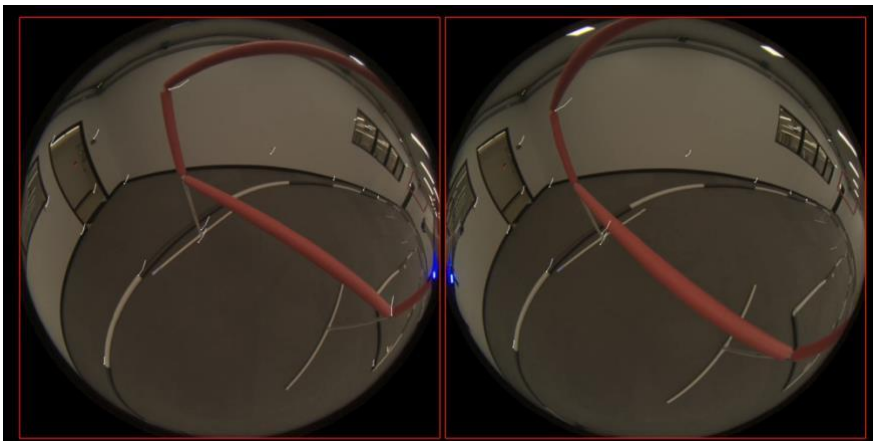

Figure S13 Fisheye camera frame from the wing mounted cameras

Mapping

Skydio has invested significant research and engineering resources towards vision-based estimation of stereo depth in unstructured environments. The general approach has been a combination of robust engineering principles and novel deep neural-network learning techniques. Skydio has, in the last few years, surpassed the state-of-the-art in real-time stereo depth estimation, enabling revolutionary obstacle avoidance and intelligent path planning behaviors.

This mapping system has been adapted to function reliably with data from rolling-shutter cameras, correctly accounting for both translational and rotational movement of the camera while the frame is being captured.

Again, we adapted this system to function on the falcon robot quadcopter test platform, and the video below shows the high-quality depth information obtained from a real flight. This demonstrates both the functioning of the stereo vision system on this novel vehicle geometry, and demonstrates the tolerance of the system to the high-frequency vibrations induced by the propellers during flight.

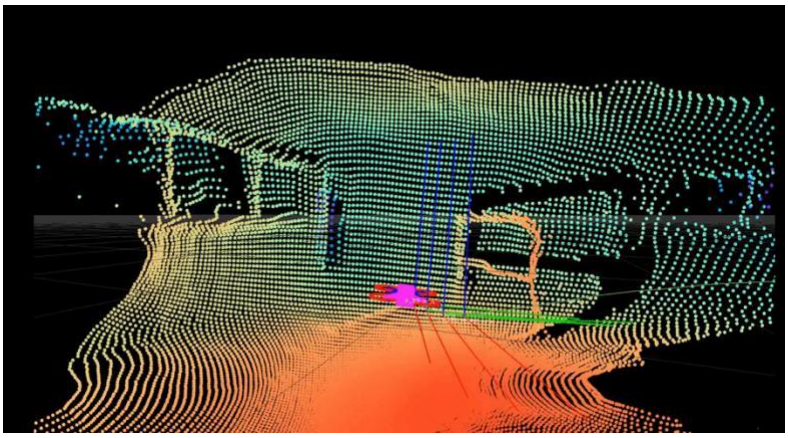

Figure S14 The obstacle avoidance system's view of the flight-testing environment.

Motion Planning

During the past years, Skydio has rebuilt our motion planning system from scratch to more reliably generate high-quality planned trajectories that stay within the physical capabilities of the vehicle. Additionally, these changes have made the system significantly more generalizable to non-quadcopter flying vehicles, like the falcon robot.

Cited References

[1] W.R.T. Roderick, M.R. Cutkosky and D. Lentink, "Bird-inspired dynamic grasping and perching in arboreal environments," *Science Robotics*. 6, no. 61 (2021).

[2] Roderick, William RT, Diana D. Chin, Mark R. Cutkosky, and David Lentink. "Birds land reliably on complex surfaces by adapting their foot-surface interactions upon contact." *Elife* 8 (2019): e46415.

[3] Chang, Eric, Laura Y. Matloff, Amanda K. Stowers, and David Lentink. "Soft biohybrid morphing wings with feathers underactuated by wrist and finger motion." *Science Robotics* 5, no. 38 (2020).

[4] Matloff, Laura Y., Eric Chang, Teresa J. Feo, Lindsie Jeffries, Amanda K. Stowers, Cole Thomson, and David Lentink. "How flight feathers stick together to form a continuous morphing wing." *Science* 367, no. 6475 (2020): 293-297.

[5] Howell, Taylor A., Brian E. Jackson, and Zachary Manchester. "Altro: A fast solver for constrained trajectory optimization." In 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 7674-7679. IEEE, 2019.

[6] A. L. R. Thomas, G. K. Taylor, Animal Flight Dynamics I. Stability in Gliding Flight. *Journal of Theoretical Biology*. **212**, 399–424 (2001).

[7] R. L. Nudds, G. K. Taylor, A. L. R. Thomas, Tuning of Strouhal number for high propulsive efficiency accurately predicts how wingbeat frequency and stroke amplitude relate and scale with size and flight speed in birds. *Proceedings of the Royal Society of London. Series B: Biological Sciences*. **271**, 2071–2076 (2004).

[8] B. W. Tobalske, K. P. Dial, Flight kinematics of black-billed magpies and pigeons over a wide range of speeds. *The Journal of Experimental Biology*. **199**, 263–280 (1996).

[9] G. K. Taylor, A. L. R. Thomas, Animal Flight Dynamics II. Longitudinal Stability in Flapping Flight. *Journal of Theoretical Biology*. **214**, 351–370 (2002).

[10] K. L. McArthur, J. D. Dickman, State-dependent sensorimotor processing: gaze and posture stability during simulated flight in birds. *Journal of Neurophysiology*. **105**, 1689–1700 (2011).

[11] J. D. Delius, F. W. Vollrath, Rotation compensating reflexes independent of the labyrinth: Neurosensory correlates in pigeons. *J. Comp. Physiol.* **83**, 123–134 (1973).

[12] M. Roth, Review of atmospheric turbulence over cities. *Quarterly Journal of the Royal Meteorological Society*. **126**, 941–990 (2000).

Published DESI Project Research

DESI PhD Theses and students graduated: 2

1. Chang E. Avian-inspired Rudderless Flight with Morphing Wings and Tails (Doctoral dissertation, Stanford University).
2. Roderick WR. Takeoff and Landing in Birds as a Source of Inspiration for Designing Aerial Robots (Doctoral dissertation, Stanford University).

Publications: 2

3. Roderick, W. R.T., M. R. Cutkosky, and D. Lentink. 2021. "Bird-Inspired Dynamic Grasping and Perching in Arboreal Environments." ***Science Robotics*** 6 (61). Cover Article https://doi.org/10.1126/SCIROBOTICS.ABJ7562.
   [*This article received world-wide media attention and was featured in a Science AAAS video feature. The Science Robotics cover paper media attention is listed here: https://robotics.altmetric.com/details/117933589 The paper was categorized as "Among the highest-scoring outputs from this source (#18 of 458)"*].
4. Martiros, Hayk, Aaron Miller, Nathan Bucki, Bradley Solliday, Ryan Kennedy, Jack Zhu, Teo Tomic, et al. 2022. "SymForce: Symbolic Computation and Code Generation for Robotics." In ***Proceedings of Robotics: Science and Systems*.** https://github.com/symforce-org/symforce.

Publication at International Symposia: 1

5. Chang, Eric, and David Lentink. "Biohybrid morphing tail aerial robot." *9th International Symposium on Adaptive Motion of Animals and Machines (AMAM 2019)*. 2019.

Publication in Revision: 1

6. Chen, A. G., K. Hoffmann, D. Lentink., M. R. Cutkosky. 2022. "Aerial Grasping and the Velocity Sufficiency Region" ***RA-L and IROS***

Publication Submitted: 1

7. Hoffmann, K., Chen, A.G., M. R. Cutkosky, D. Lentink. 2022. "Bird-inspired robotics principles as a framework for new smart and robotic aerospace material development"

Publication in Preparation: 1

8. Chang, E., and Lentink, D. Rudderless morphing flight (invited by Science Robotics).