Understanding and Improving Fully Autonomous Aerial Systems

Paper No: 426

Abstract

Fully autonomous aerial systems (FAAS) fly complex missions guided wholly by software. They are cheaper and safer than unmanned aerial systems piloted by humans; however, compute hardware and software settings affect where FAAS fly and how quickly they complete missions. Poor choices slow down missions and use hardware inefficiently, increasing costs. This paper presents a model-driven approach to manage FAAS missions. We fly real FAAS missions, profile compute and aircraft resource usage and model expected demands. However, profiling numerous hardware and software settings also inflates costs. We propose autonomy cubes: data structures that characterize whole flight areas for FAAS missions. With autonomy cubes, we can profile resource usage across settings without repeatedly flying missions. We combine expected resource demands to model missions completed, i.e., mission throughput. We validated our model by creating FAAS, measuring mission throughput across many system settings and comparing to our model predictions. Our FAAS benchmarks execute realistic missions: autonomous photography, search and rescue, and agricultural scouting. And they use well-known software, e.g., A* Search and DJI flight control. Our model predicted throughput with 4% error across mission, software and hardware settings. Competing approaches yielded 10–24% error. We used our model to study (1) GPU acceleration, scale up, and scale out, (2) onboard, edge and cloud computing, (3) energy and monetary budgets, and (4) software driven GPU management. We found that model-driven management can boost mission throughput by 10X and reduce costs by 87%.

1 Introduction

Unmanned aerial systems (UAS) hover, fly to waypoints and perform defined actions, e.g., landing and takeoff. In addition to rotors and motors, these aircraft carry computer systems, cameras, batteries, etc. They can access high, vast or unsafe places and capture detailed images and sensor readings.

Photographers, farmers and first responders pilot UAS via remote control or smart phone [2, 9, 18, 23, 30]. These end users decide where the UAS flies, when it senses data and when missions are complete. Complex missions link sensed data and completion, i.e., the mission is complete if the data satisfies user goals. Complex missions require dynamic flight paths, making them hard to pilot. Piloting mistakes can have severe consequences. For example, flying UAS in restricted areas risks human lives. Other common mistakes, e.g., flying to unneeded waypoints, degrade mission throughput (i.e., the number of missions completed). Aerial systems that require less human piloting are needed [23, 29]. There is growing support for software development kits that control aircraft. Da-Jiang Innovations (DJI) aircraft support software control from iOS, Android and Linux devices [27]. Pixhawk and Aerostack also provide platforms for software control [24–26].

Fully autonomous aerial systems (FAAS) fly dynamic missions defined by software. End users do not pilot FAAS. Instead, they provide goals, constraints and software that execute missions. FAAS take off, sense data at runtime and compute flight paths likely to achieve goals. FAAS software can prohibit risky piloting practices, e.g., flying in restricted zones, making FAAS an innately safer option. FAAS compute hardware, aircraft and software determine mission throughput— not human piloting. Good settings can be replicated if, for example, a FAAS achieves high mission throughput.

Recent unmanned aircraft carry sophisticated processors onboard. For example, the DJI Mavic carries the Myriad 2, a system on chip that includes streaming vector engine processors, hardware accelerators, multiple RISC cores and 2 MB on-chip memory. However, FAAS will increase computational demands significantly [2, 21, 29]. If aircraft surrender battery capacity to onboard processors and carry heavier payloads for data storage, flight times will suffer. Instead of using onboard resources, FAAS workloads may run on land using networked messages to control the aircraft; an architectural approach called edge computing. Another choice connects aircraft to fast cloud data centers. Latency, processing capacity and cost differ among these choices. End users must

understand how these factors affect mission throughput.

Autonomous systems combine many independent software components. Many components support settings that trade compute demand for energy savings. It is hard to predict the effects of these settings on mission throughput because they affect where FAAS fly, how many compute resources they require, and the effects are mission specific. For example, lightweight image classifiers can lower compute demand and save energy, but FAAS also fly to more waypoints which can negate savings. End users could test each setting and measure mission throughput directly. However, long complex missions and many software settings make exhaustive testing impractical.

This paper presents a modeling approach that predicts mission throughput. Our approach profiles energy demands for aircraft and compute. We model missions as a sequence of waypoints. The waypoint that exhausts battery capacity defines mission throughput. It is hard to profile energy demands across system settings that affect autonomous decisions; we call these autonomy settings. Autonomy settings change flight paths, affecting which data is sensed during mission execution and ultimately energy demands. We propose *autonomy cubes,* data structures that can supply sensed data for any flight path—inspired by data cubes [12].

To validate the model, we created FAAS that perform the following complex missions: (1) autonomously capture high quality photographs of human faces, (2) search and map defined areas for first responders, and (3) scout crop fields for representative samples. Our FAAS use path finding and AI approaches found in prior research [21, 25]. Each FAAS supports a wide range of software and hardware settings. Toggling these settings can increase waypoints per mission by 4X and compute demands by 35X.

We collected 122 autonomy cubes, flying in 56 locations and capturing 20970 data readings. We also flew 145 actual FAAS missions and measured ground-truth mission throughput. FAAS used DJI Spark and Mavic Pro aircraft. We compared our modeling approach to Aerostack [25] and Autoware [15, 21]. These approaches use reference traces from prior missions to model energy usage. Our model predicted throughput with 4% error. When trace and mission settings differed on multiple dimensions, Aerostack and Autoware yielded error up to 5X and 10X larger than our approach.

Consider end users that buy aircraft, software and hardware separately and then combine them to make a FAAS. After purchasing aircraft and software, these end users would like to purchase compute hardware that will provide high throughput. Under a cost budget, these end users may have to maximize throughput per dollar. End users can use our models to answer these questions. We used our modeling approach to compare onboard, edge and cloud architectures. Onboard compute and storage reduced flight time by 50%, significantly degrading mission throughput. Edge computing eventually provided best mission throughput and throughput per dollar. However,
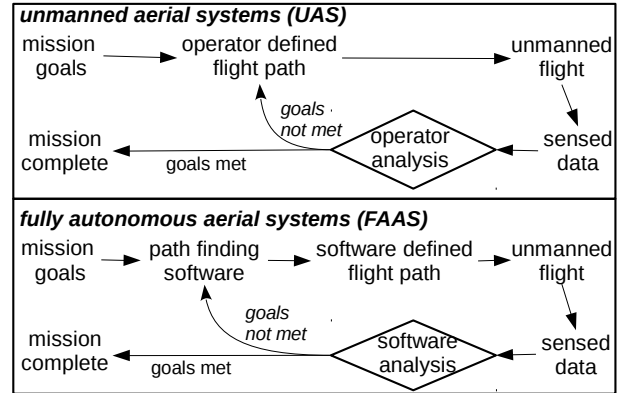

Fig. 1: With FAAS, humans set mission goals but do not pilot.

we observed that larger aircraft could boost onboard architectures. We also used our modeling approach to compare scale out, scale up and GPU approaches to meet compute demands. We found that the best approach depended on (1) autonomy settings and (2) energy capacity. GPU improve throughput, but deplete batteries quickly. Scaling up cores on chip provided a reliable approach to increase throughput.

Some end users will control aircraft, processor and software [23]. These end users can adjust all facets of FAAS to find high throughput settings. We studied software and hardware co-design. We set up an adaptive policy that toggles between deep, compute intensive AI models and less precise but energy efficient AI models. An edge system running CPU and GPU can power off the GPU for less precise models. We compared approaches that toggle GPU states during missions and between missions. Toggling GPU states GPU during a mission provided higher mission throughput.

Our contributions are as follows:

- Our modeling approach driven by autonomy cubes precisely predicts mission throughput. We show that autonomy settings have large effects on FAAS flight path and energy usage.

- We built three FAAS (autonomous photography, search and rescue, and agricultural scouting). We measure their mission throughput directly and validate our modeling approach.

- Model-driven analysis yields insights that boost throughput (up to 10X) and lower costs (87%).

The remainder of this paper is as follows. Section 2 provides an overview of FAAS. Section 3 presents an energy model driven by autonomy cubes. Section 4 describes the implementation of FAAS. Section 5 validate our model and studies FAAS workloads. Section 6 uses our modeling approach to guide FAAS system management. Section 7 discusses related work. Section 8 presents conclusions.

2 Fully Autonomous Aerial Systems

Autonomous systems perform complex tasks in vaguely defined areas without receiving commands from humans. By this definition, UAS are *not* autonomous. Humans decide (1) where to fly, (2) when to sense data and (3) when a mission is complete. Figure 1 depicts UAS workflow. Humans set high-level mission goals, e.g., take a great photo of a human target. Then, they pilot the aircraft to waypoints by (1) using remote control devices, (2) making gestures or (3) providing a list of GPS coordinates. At each waypoint, the UAS senses data from its surroundings, e.g. detailed images or GPS data. After studying data, humans decide if their goals are met. If not, they choose new waypoints and repeat.

Figure 1 also depicts workflow for FAAS. Humans set goals, but all remaining work is done by software. The system is capable of completing multiple missions without a human issuing commands. To remove humans from the loop, software must decide when a mission is complete, meaning both human end users and software can understand mission goals. There is a semantic gap between true goals and goals that can be expressed in software. Today, autonomous systems require end users to translate their vague, high-level goals into mathematical equations (called utility functions). If the mission is not complete, FAAS software must also choose the next waypoint.

An Example FAAS Mission: Crop fields are vast. Scouting reports can miss subtle problems, e.g., over crowding or crop disease. Human piloted UAS have transformed scouting. Companies, like Fly The Farm [20], fly over fields and capture detailed images. These images inform farmers, guiding the application of fertilizer and pesticides. However, UAS pilots charge $1–$5 per acre. One scouting report can cost nearly 3% of net profits for corn fields [4, 11, 30]. By eschewing human pilots, FAAS can lower costs.

Figure 2 depicts flight area and two agricultural scouting missions. Flight area comprises waypoints where the aircraft can fly. Each cell in Figure 2 represents a waypoint (here, a GPS location). Following the blue line, our FAAS flight controller directs the aircraft to a waypoint (A1) and captures a detailed image. FAAS software analyzes the image, counting corn crops. If the image contains enough crops, the mission is complete. (Note, the image may be fed into subsequent analysis, such as yield modeling [18].) If not, FAAS path finding software chooses a new waypoint (B1). This process repeats and the mission completes at waypoint B2.

FAAS visit only some waypoints on each mission. As shown in Figure 2 path finding software (A* search versus nearest neighbors) changes where FAAS fly and how quickly missions complete. The settings represented by the red line choose a longer flight path.
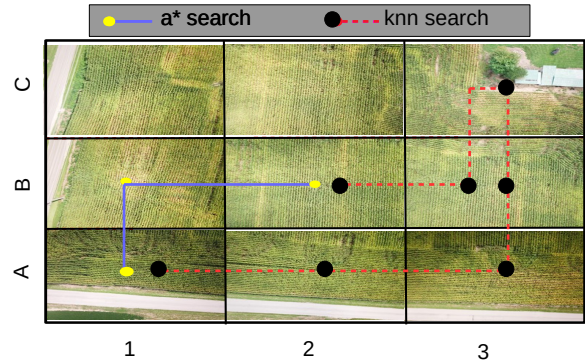

Fig. 2: Two agricultural scouting missions. Each image represents an allowed waypoint where the aircraft could have flown (i.e., flight area). Lines represent actual waypoints visited. Both missions begin at A1.

3 Performance Modeling

Figure 3 outlines our approach to model mission throughput. Model inputs relate to autonomy (goals & workload), compute architecture and aircraft. Our approach has four stages, shown as boxes in Figure 3. First, we collect all data that could be sensed during a mission, i.e., an *autonomy cube*. Autonomy cubes precisely model flight paths. Given a flight path, the next stages profile compute and aircraft workloads. Finally, bottleneck analysis predicts whether aircraft or compute exhausts batteries first.

3.1 Autonomy profiling

FAAS are hard to model, because their flight path depends on *which* data is sensed at runtime. Two FAAS flying only a few feet apart can differ on the utility of their sensed data. Their flight paths could diverge, affecting energy usage per mission and ultimately mission throughput. We mitigate variation caused by spatial displacement by modeling expected mission throughput averaged over many runs. Path finding, AI models, utility functions and other autonomy settings have systemic and non-linear effects on flight path.

Competing Approaches: Before detailing our approach, it helps to explain how recent work models autonomous systems [15,21,26]. Autoware [15,21] uses a long, representative trace from an self driving car. This trace suffices for research because safety concerns constrain driving maneuvers and execution environments. Aerostack [26] creates multiple traces where users change the environment between traces. This approach captures a wide range of maneuvers, but, if autonomy settings change, flight paths will diverge from previous traces. Further, even if some tested settings behave like the new settings, irrelevant tested settings dilute their impact.

We use *autonomy cubes*, a data structure that captures all images that can be sensed during a mission within user defined constraints. Autonomy cubes represent a principled ideal for

Fig. 3: Modeling mission throughput for multiple architectures, aircraft and autonomy settings.

benchmarking; they can be used to compute flight paths across any FAAS autonomy setting. By depicting captured images at each allowed GPS location, Figure 2 presents a simple 2-D autonomy cube.

Defining autonomy cubes: Shown in Equation 1, A waypoint x is a multi-dimensional point. Dimensions can abstract (1) GPS or grid positions (e.g., Figure 2), (2) aircraft poses (e.g., gimbal positions), and localized data (e.g., altimeter and compass readings). FAAS fly in a discrete and finite space, where $\forall i : 0 < d_i < D_i$.

$$x = (d_0, ..., d_K) \tag{1}$$

$$\mathbf{fa_i} \in \mathbf{FA} : \{x_m\} \rightarrow \{x_n\} \tag{2}$$

$$\mathbf{FS} : \{x, fa_i, y\} \rightarrow \{1|0\} \ \ \text{if} \ \ fa_i(x) = y \tag{3}$$

$$\mathbf{ug} : \{x, fa, y\} \rightarrow r \in \mathbb{R} : r = \frac{u(ai(x))}{u(ai(y))} \circ FS \tag{4}$$

As shown in Equation 2, we abstract flight actions *FA* as a set of functions. Precisely, a flight action applies a set force. The force moves a hovering aircraft along six degrees of freedom. Actions are calibrated offline. Each action fa_i moves aircraft from one expected waypoint to another. Due to spatial displacement, actual and expected positions may differ slightly. For example, wind can apply unexpected force, moving the aircraft away from its expected position. Note, waypoints reachable by any combination of actions define flight area.

A single step along a flight path has a starting point, action and ending point. The flight step function *FS* (Equation 3) indicates valid steps where the action leads to the ending point. At each waypoint, FAAS senses its surroundings, transforms sensed data using AI models (*ai*) and computes utility ($0 < u(x) <= 1$) of its current state. Each valid flight step, i.e., $FS(...) \rightarrow 1$, has utility gain. Referring to Figure 2, the first flight step for the blue line is: $\{[A, 1], \text{FlyNorth}, [B, 1]\}$. The AI models for this mission include A* (i.e., ai_{A*}).

We now define autonomy cube, $\forall x \in \mathbf{x} \ ac = \cup (x, ai(x))$, where $\mathbf{x}$ represents the set of all reachable waypoints. Equation 4 shows that *ac* allows FAAS to compute utility gain for any valid flight step. As shown in Equation 5, a flight path is

a sequence of *N* valid flight steps. Autonomous systems aim to maximize total utility gain **TG**, i.e., the product of utility gains acquired at each step in the path.

$$\text{fp} = \{x_n, fa_{f(n)}, y_n\}_N$$
$$\mathbf{TG} : \{fp\} \rightarrow r \in \mathbb{R} : r = \prod_n ug(x_n, fa_{f(n)}, y_n) \tag{5}$$

Using autonomy cubes: Autonomy cubes can be used to simulate a FAAS mission. For this paper, we constrain flight areas to an n-dimensional mesh of waypoints, e.g., a building or crop field. Each dimension corresponds to a FAAS flight action. Each action has an inverse action that is expected to return the aircraft to its original position. Supported actions include x,y, z translation and pitch, yaw, roll and gimbal pitch.
Capturing autonomy cubes: Quadcopters support 6 controllable degrees of freedom, meaning they can use pitch, yaw and roll to fly in any direction along an x,y,z coordinate system [22]. Unlike cars and fixed wing planes, the coordinate system can be explored in any direction relatively quickly, requiring at most rotation and thrust. However, they only move forward in time. We exploit quadcopter maneuverability to capture sensed data before it changes, i.e., we transform time into discrete blocks based on how frequently the utility of sensed data changes.

Equation 6 defines scene persistence *P* as the minimum discrete time slots *t* such that a hovering aircraft percieves qualitatively similiar utility. As shown in Equation 7, to capture an autonomy cube, the aircraft flies to each waypoint in **x** before *P* seconds have elapsed. We use the slowest flight action to model shifting between waypoints.

$$P = \text{Latency}(fa_{hover}(x)) \times t :$$
$$\arg \min t: ug(x_0, \circ_t fa_{hover}) < K_2 \tag{6}$$

$$\frac{\text{Max}_i(\text{Latency}(fa_i(x))) * |\mathbf{x}|}{s} < P \tag{7}$$

Examining Equation 7, we observe four techniques to scale our approach.

1. Increase scene persistence: Scenes can be tweaked manually so that key features change less frequently. For example, we have tested our FAAS benchmarking suite using mannequins in place of fidgety humans, and farm land.

4

2. Shrink flight area: We can shrink the total area where aircraft can fly or allow fewer flight actions. Of course, fewer flight actions degrades total utility gain.

3. Speedup flight actions: We could also reduce flight time going between waypoints. Scheduling autonomy cube flight paths by prioritizing waypoints with the shortest flight time relative to the current waypoint would minimize delay per flight step.

4. Use swarms to partition cubes: Finally, multiple quadcopters can be deployed at once, allowing each to capture a fraction ($\frac{1}{S}$) of the flight area. We have used swarms to capture autonomy cubes used with our FAAS. However, for workloads that require tight maneuvering (autonomous photography), partitioning presents several research challenges. First, partitioning to minimize expected delay per action is challenging. Seconding, partitioning should consider the effects of battery capacity. Partitioning on search and rescue (partitioning by rooms) and agricultural sampling (by field region) are much more feasible. Finally, aircraft flying in the same region may interfere with each other.

We used swarms comprising 2 & 3 aircraft to partition flight area along the vertical axis (y-axis) for our autonomous photography benchmark. As expected, we were able to cover up 3X more flight area in the best case. However, we also observed anomalies unique to aerial systems. Placing aircraft immediately under each other (i.e., partitioning y while strictly keeping x & z the same) affected wind patterns, creating suctions. The aircraft crashed into each other. Partitions that worked well slightly offset the aircraft in the x & z dimensions.

3.2 Aircraft profiling

In the second stage of our modeling approach, we profile latency and energy functions for flight actions on an input aircraft. In Equation 8, the *LatencyA* function estimates latency using the average delay of flight actions executed at multiple, sampled points. Conceptually, we do the same for energy and hover power. This profiling is done offline.

$$\text{LatencyA} = \{fa_i(x)\} \rightarrow r \in \mathbb{R} :$$
$$r = \frac{\sum_{n'} \overline{\text{LatencyA}(fa_i((x_{n'}))}}{N'}, N' << |\mathbf{x}| \tag{8}$$

3.3 Compute profiling

Compute latency and energy vary depending on the content of data sensed at a waypoint. Unlike UAV actions, which have a tight latency distribution, compute latency has more variance. Scheduling fluctuations, unpredictable threading overhead and model timing, and network interference all cause compute timings to vary. Compute latency was profiled online, requiring a varied set of execution environments and

conditions. experimental latency numbers for all individual components of our benchmark were compiled into normal distributions, truncated to the third standard deviation. Our model uses these distributions to predict compute latency for offline missions.

$$\text{LatencyC} = \{ai(x)\} \rightarrow (\mu, \sigma) \in \mathbb{R} :$$
$$\sigma = \frac{\sum_{n'} \overline{\text{Latency}(ai((x_{n'}))}}{N'}, N' << |\mathbf{x}| \tag{9}$$

3.4 Throughput modeling

Recall, FAAS compute their flight path fp_i at runtime. As shown in Equation 10, each flight step $fs_{[i,n]} \in fp_i$ is informed by data observed during execution.

$$\widehat{\mathbf{ug}_{TD}} : \{x_n, fa_{f(n)}, y_n\}_N \rightarrow \{x_{n+1}, fa_{f(n+1)}, y_{n+1}\} \tag{10}$$

Specifically, FAAS compute expected utility gain $\widehat{\mathbf{ug}}$, using past and training data to infer the effects of flight actions. Flight path fp_i is the result of iterative invocations of expected utility gain given an autonomous cube, i.e., $fp_i = \{\widehat{\mathbf{ug}} \circ_N ac\}$. To model throughput, we assume we access to $\widehat{\mathbf{ug}}$ and ac.

We model energy used by the aircraft with two terms. First for every step along a flight path, we sum energy used for the corresponding action. Second, multiply latency for compute by power used when hovering. The inverse applies to compute. When both compute and aircraft have distinct energy sources with known storage capacities (C_{air} and C_{cmp}), mission throughput is computed by looking at the number of missions completed before exhausting one energy sources, see Equations 11.

$$E_{air}(N) : \sum_n (\text{EnA}(fa_n) + \text{LatencyC}(ai(x_n)) \times \text{PwrA}(fa_{hover}))$$
$$E_{cmp}(N) : \sum_n (\text{EnC}(ai(x_n)) + \text{LatencyA}(fa_n) \times \text{PwrC}(idle))$$
$$\text{tput} = \min N : C_{air} - E_{air}(N) = 0 \text{ or } C_{cmp} - E_{cmp}(N) = 0 \tag{11}$$

4 Implementing FAAS

Table 1 decomposes FAAS and presents a layered, systems view of their components. This section presents each layer and compares system settings for 3 FAAS.

FAAS missions: We implemented (1) autonomous photography, (2) search and rescue, and (3) agricultural scouting. For autonomous photography, the FAAS positions itself and takes high-quality portraits of human faces. It autonomously explores its flight area. This workload was inspired by computational photography and SkyDio [1, 2].

FAAS mission	autonomous photography	search and rescue	agricultural scouting
utility functions and flight constraints	1) prioritize high utility (util) 2) prefer short missions (tput) 3) good mix		1) high utility req. 2) max waypoints (10, 20 or 30)
flight area	2x2x3x3 hyper cube near subject	15 hyper cubes in a target area	n x m grid over crop field
flight actions	translate x, y and z, gimbal pitch		translate x and y
path finding algorithm	1) choose from K-nearest neighbors 2) use A* search to avoid local optima 3) use energy aware A* (EA*) which prioritizes low power movements		
AI classifier accuracy & complexity	1) Integer precision (int) fast but less accurate, 2) floating point precision (fp), 3) deep most accurate but require GPU (all)		
execution context	edge systems, cloud, or onboard		
arch. support	scale up, scale out, gpu acceleration		

Table 1: Layered implementation and system settings.

The search and rescue FAAS extends autonomous photography. It searches multiple areas for humans. During emergencies or disasters, it could help first responders discover victims. The FAAS navigates the aircraft between areas, e.g., rooms in a building, and also explores areas thoroughly.

As discussed, agricultural scouting is commercially viable today. This FAAS takes aerial images of a crop field. For this work, we had access to a corn field, so our missions produce detailed images of corn and planting rows. Scouting as a workload kernel underlies aerial surveillance and military target detection.

Flight area and flight actions: Autonomous photography covers a 2x2x3x3 hyper cube. The aircraft can translate X, Y and Z axes and rotate the camera. We have collected 110 autonomy cubes for this benchmark. Search and rescue explores 15 2x2x3x3 hyper cubes and supports the same actions.

Agricultural scouting covers a 75-acre crop field. The flight area is a 55x43 grid. The aircraft can translate X and Y dimensions only. In total, we have collected 122 autonomy cubes (20970 images) across (1) all 3 FAAS, (2) diverse settings: outdoors, indoors, raining and windy, and (3) with multiple targets: humans and corn. Capturing a cube took roughly 11 minutes for autonomous photography cubes and 4 hours for agricultural scouting.

Utility functions and constraints: Photography and crop analysis use different utility functions. A good portrait contains a centered, bright and crisp face [5]. We created a utility function using the following features: face detection, face location in the image, image brightness and size of the facial bounding box. A good picture of a crop field avoids blur, contrasts crops and soil and does not include extraneous objects. Our utility function here considers glare, image brightness and corn crops counted. For all utility functions, each feature is weighted and the whole function is normalized.

End users set thresholds. When utility exceeds the threshold, the FAAS mission is complete. Our photography FAAS support 3 thresholds. A high threshold encourages the FAAS to explore its flight area. As a result, missions are longer. A low threshold encourages the FAAS to land quickly. We label this setting as high throughput. Finally, the default settings aims for a medium threshold that provides good mix.

Flight area bounds FAAS flight path spatially. Flight actions that cause the aircraft to leave that area are not executed. If the aircraft battery falls below 10% of its capacity, our FAAS lands immediately. The Max Waypoints setting bounds flight path temporally. After exceeding this threshold, the mission completes.

Path finding: By default, 4000 training data entries are used to decide where to fly. The K-Nearest Neighbors (KNN) algorithm finds 9 entries with utility features nearest to the sensed image. By default, we implement greedy path finding. The expected utility gain is the mean gain observed by nearest neighbors grouped by flight action. This approach takes the flight action with the largest expected utility gain.

A* Search improves greedy KNN with a linear heuristic to model the whole flight area, choosing a flight action along the best expected path. Energy-aware A* Search weights flight actions according to aircraft profiles. It produces flight paths that prefer low energy actions. A* Search and its energy-aware variant are well studied and have been used in recent research [6, 15, 21].

AI models: Each FAAS characterizes sensed data into a vector with up to 64 dimensions. Each dimension represents the output of an AI classifier. We distinguish classifiers by compute demand and support any subset of these groups. Integer models include OpenCV local binary pattern, cascade models using only integer data types, and RGB image classifiers. These models are lightweight, fast and imprecise. Floating point models include DLIB histogram of gradients. The are more precise than integer models but also slower to compute. Deep models include DLIB's convolutional neural network (CNN) for face recognition and our custom CNN for crop recognition. We execute deep models only when GPUs are available (i.e., not on CPUs).

Execution context and architectural support: Our suite is composed of micro services. Each micro service provides basic functions, e.g., issuing aircraft commands, data sensing, data storage, running sensed data through an AI model, querying path finding algorithms etc. Micro services exchange messages using Californium UDP CoAP clients and servers [19]. Our platform supports scale up, scale out and GPU-based compute architectures. It also supports edge and cloud computing resources.

arch	2c	2x2c	4c	4c	4c	4c	2c	2c	2c	2c	2c
gpu	no	no	no	yes	no	yes	no	no	no	no	no
AI models	fp	fp	fp	all	fp	all	int	int	fp	int	int
path find	knn	knn	knn	knn	a*	a*	ea*	knn	knn	a*	a*
waypoints	10	10	10	10	10	10	6	15	10	15	6
goal	mix	mix	mix	mix	tput	tput	tput	util	util	util	util

Fig. 4: Our model predicts mission throughput precisely. Baseline setting is highlighted

5 Workload and model evaluation

Our FAAS provide ground truth. We can measure mission throughput directly with real aircraft, goals, software settings and compute hardware. This section first compares our model predictions to observed throughput. Then, we compare competing modeling approaches. Finally, we isolate compute and aircraft profiles, and characterize these workloads.

5.1 Model validation

We flew each benchmark under system settings in Table 1. Our FAAS uses the DJI Android SDK to control the aircraft via WiFi connected laptop (edge device). Our platform can also run software components across multiple devices or on the cloud. Edge devices run Ubuntu Linux 18.04.

Each test started with fully charged aircraft and edge batteries. We then flew missions until one of the batteries fell below the safe landing threshold. Observed mission throughput is the number of missions completed. We repeated each test 6 times and report mean throughput.

Unless noted otherwise, we used the DJI Spark aircraft [8]. Its body is roughly 6 square inches. It weighs 300 grams. We observed that it can hover for 11–13 minutes without recharging its 16 Wh lithium ion battery. Also unless otherwise noted, we use edge architecture setup, because it is easier to change architecture settings. We tested edge devices with the following compute architectures.

- **2c:** HP G6 laptop; 2-core i5 7200u processor; 3.1 GHz; 3 MB cache; 4 GB DDR4 RAM; 500 GB hard drive.

- **2x2c:** 2 HP G6 laptops using 1 Gbps ethernet router. One laptop runs flight control, pulls images from the aircraft and computes integer AI models. The other laptop runs path finding algorithms and floating point models.

- **4c:** 4-core i5 7300u processor; 3.5 Ghz Ghz; 3 MB cache; 4 GB DDR4 RAM; 500 GB hard drive.

- **4ci7:** 4-core i7 7500u processor; 3.5 Ghz; 4 MB cache; 8 GB RAM; 256 GB SSD.


Fig. 5: Depicting Autoware versus our approach.

- **4c+gpu:** 4ci7 connects to an NVIDIA 1080 Ti.

- **2c+gpu:** GPU is connected to 2c.

Prediction accuracy: Recall, our modeling approach predicts expected mission throughput, i.e., an average over many missions. For autonomous photography and search and rescue, our approach uses autonomy cubes to produce 50 mission flight paths for each autonomy setting. Note, compute hardware settings do not affect flight paths. Autonomy settings include AI models, path finding and utility functions. For each flight path, Section 3 describes the workflow to predict mission throughput. Agricultural scouting covers a larger area. We have fewer cubes. Here, we generate 6 flight paths for each autonomy setting. Autonomy cubes were implemented by a micro service that returns an image from a cube waypoint in place of the aircraft camera. FAAS software interacts with the micro service as it would with the aircraft.

Figure 4 compares predicted and observed mission throughput for autonomous photography. We shorten the names of mission goal parameters to mix, util and tput for space. We also shorten integer and floating point settings for AI models to int and fp. Our tests cover every autonomy setting supported. Mean absolute percent error (i.e., $\frac{|\text{pred}-\text{obs}|}{\text{obs}}$) was 4%. We found that GPU, goals and path finding settings affected throughput by up to 1.8X, 1.75X and 1.71X in isolation. Combined, settings had complex effects. For example, adding a GPU sped up throughput by 1.15X under 4c, KNN and util. However, under 4c, A* and tput speed up was 1.23X–a 7% improvement. Util and KNN missions spent more time hovering. Energy used hovering lessened whole system speedup gained by adding a GPU.

Competing modeling approaches: We also studied modeling approaches inspired by recent research. In *Autoware* [15, 16, 21], researchers use ROSBAG recordings from a real, long-running self-driving car. We mimicked this approach by collecting long traces over multiple missions. For autonomous photography and search and rescue, we combined 100 mission flight paths. Scouting used 12 missions.

Autoware does not consider autonomy settings. As such, this approach does not model flight path well. Figure 5 depicts the problem using examples from Figure 2. Autoware profiles compute workloads on new hardware. However, Autoware can not acquire data outside of the trace. If autonomy settings change where FAAS would fly, Autoware doesn't have access to the sensed data and profiles using available data. Figure 5 highlights the problem: A* missions complete faster than
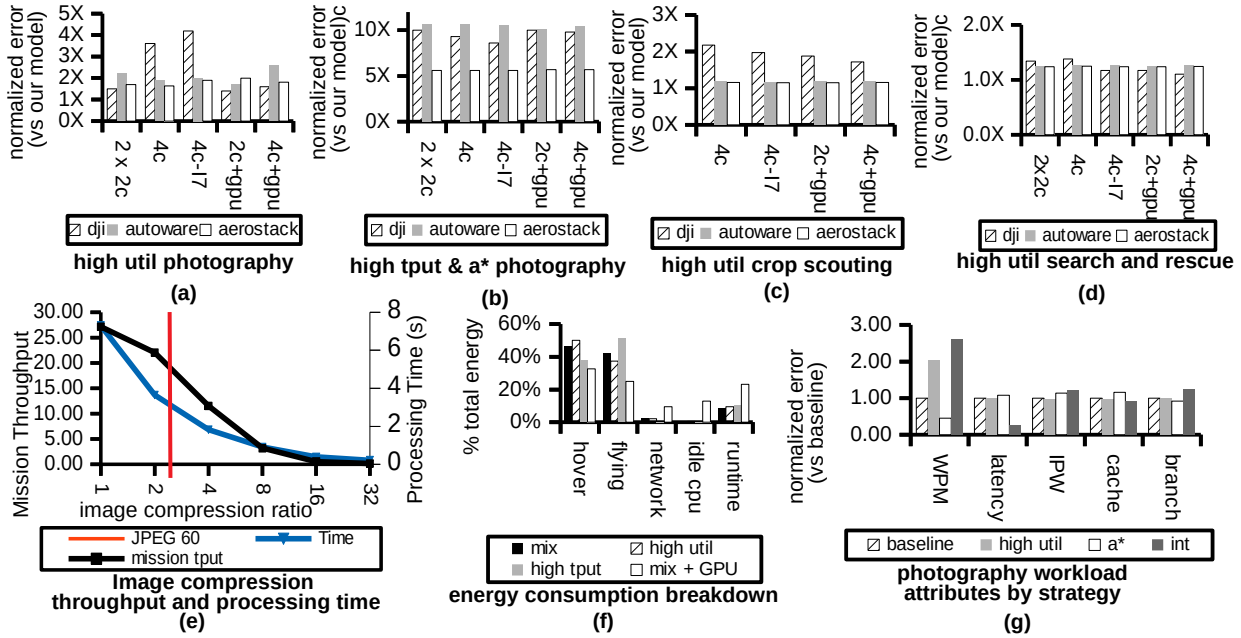
Fig. 6: (a – d) Comparing trace and cube driven modeling approaches. (e) poor image quality degrades mission. (f) Autonomy settings shift compute and aircraft energy demands. (g) Autonomy settings affect common architecture counters.

KNN and Mix missions. As a result, Autoware over estimates the total compute workload.

Aerostack flies autonomous aircraft in a wide range of settings by manually inserting obstacles [26]. This approach improves Autoware's methodology, because traces include data from multiple settings. We mimicked this approach by creating 3 long running traces for each setting.

Figures 6 (a – d) compare our approach, Autoware and Aerostack. We also compare a simple modeling approach driven by data collected from DJI and Intel (*DJI*). This approach ignores autonomy and uses flight time and aggregate cycles per second to model throughput as $\left(\text{Speedup} * \frac{\text{MaxWaypoints}*\text{WaypointTime}}{\text{FlightTime}}\right)$. Autoware and Aerostack traces used missions conducted under baseline setting. Aerostack traces toggled waypoints (15) and A* for multiple traces.

Across all workloads, settings, and architectures, competing approaches incease relative error from 1.2x–10x. Workloads with high flight overhead and lower detail sensed data experienced less error than low overhead workloads. In autonomous photography, where subtle differences in pathing can cause massive differences in sensed data, sees between 1.7X to over 10X error when using other approaches.

In Figures 6 (a,c & d), we used a setting close to the reference trace: we changed mix to util. In these graphs, Autoware and Aerostack avoid inflating error by 2X. Given our model predicts thoughput with 4%, these results are not too bad. However, Figures 6(b) makes 2 major changes: we changed mix to tput and knn to A*. As shown earlier, these settings affect throughput greatly. DJI inflates relative error 10X, Autoware by up to 20X and Aerostack by 5%. These results suggest that benchmarking must account for flight path— and, more broadly, software settings related to autonomy.

Changing aircraft: To assure the validity of our modeling approach, we created and validated models for the DJI Mavic Pro as well as the Spark. The DJI Mavic Pro is a 734g personal UAV, roughly 12 inches in length. It has a 43 Wh lithium ion battery and a maximum hover time of 23-25 minutes. Mavic, with more powerful motors and processors, requires more energy to run than Spark. Across the flight components we modeled, Mavic consumes 45-55% more energy than Spark. Validation through 5 fully autonomous missions provided an average error of 3% for our Mavic model.

Image Quality: Figure 6(e) describes the effects of image quality on throughput. Recent UAS work suggests using high compression ratios [3, 10] (such as JPEG60) or low resolution images to speed up detection. As shown, processing times decrease with compression ratio. However, image quality degrade object detection. As a result, aircraft explore more waypoints, possibly without producing valuable outcomes.

Figure 6(e) shows the decrease in throughput as image quality degrades using DLIB's facial recognition CNN. At the default quality of the DJI Spark camera (12 megapixels), our FAAS can complete 27 missions per charge. At lower image quality (3 megapixels [3, 10]), mission throughput has degraded 62%. This result shows that end-to-end metrics are critical in autonomous systems— results driven by processing time alone can miss whole system impacts.

5.2 Workload study

Figure 6(f) reports the impact of aircraft hover, flying, networked data transfer, idle compute and runtime software on total system energy. The aircraft accounts for 58-90%. The use of GPU increases the impact of compute by 4.6X. Table 6(g) delves into the architectural metrics affecting compute latency on facial recognition workloads. These data were collected on the 4c hardware using the Linux Perf tool. We observe that autonomy settings affect waypoints per mission (WPM). Integer models are too imprecise, causing the FAAS to visit many waypoints. However, integer models execute efficiently on general purpose processors, reducing the frequency of cache and branch misses by 25%. This setting provided the lowest latency, speeding up runtime by 4X.

Under A* search, the runtime executes more instructions per waypoint (IPW) before encountering cache misses than baseline setting. However, despite the lower cache miss rate, it also incurs more branch misses and executes more IPW (i.e., instructions spent computing utility gain for a sequence of actions). The net result is a 29% slowdown.

6 System Management

Our model can help FAAS end users: (1) manage compute hardware, (2) assess tradeoffs between tightly and loosely coupled aircraft, software and hardware, and (3) adapt hardware and software at runtime.

6.1 Managing compute resources

Our modeling approach uses autonomy settings to construct realistic flight paths. Flight paths and autonomy cubes yield representative compute workloads. These workloads can be tested without flying the aircraft. Consider an end user that owns a commodity aircraft. This end user may ask, *which hardware resources will provide good throughput?* Reusing flight paths across competing hardware solves this problem.

When upgrading compute resources, there are 3 options. With *scale out*, compute resources are replicated and the workload is balanced across them. Upgrading from our **2c** to **2x2c** setups reflects scale out. *Scale up* replaces resources with faster or more energy efficient resources, e.g., **2c i5** to **4c i7**. Finally, workload targeted *accelerators* can augment existing resources, e.g., **2c+gpu**.

Figure 7(a) plots speedup achieved by scaling out, scaling up, and adding GPU using the autonomous photography FAAS. Speedup is $\frac{\text{tput}_{\text{new}}}{\text{tput}_{\text{old}}}$. For this plot, the denominator is from a **2c** processor running on a device that has 2 Wh battery. Under 2 Wh battery, only scale up provides speedup greater than the increase in system cost. If the upgrade includes a 20 Wh battery, scale out and scale up are worthy investments. GPU speedup does not match its 9X cost increase. However, a GPU provides greatest increase in throughput.

6.2 Comparing onboard, edge and cloud

DJI software development kits support edge architecture where tablets run AI software and control aircraft remotely [7]. For developers, these devices offer one hop, low latency access to the aircraft and powerful compute. Further, developers can procure resources as needed.

Processors located onboard could provide lower latency, but there is a downside: onboard devices take energy from the aircraft, decreasing flight time. Note, flight time decreases for two reasons. First, and most directly, processors use energy for vision processing, path finding, etc. Second, more subtely, their weight increases thrust needed to take off, hover and fly. Small aircraft simply can not move enough air to carry an Nvidia 1080 Ti (1041 g). Even larger unmanned aerial vehicles would notice decreased in flight time.

The cloud is also an option. Elastic cloud services could dynamically provision resources, allowing end users to lease hardware on demand and avoid over provisioning. The downside is that slow network latency reduces responsiveness.

We extended our aircraft profiles to model flight time given added payload. The relationship between flight time and payload weight depends on nominal thrust and aircraft weight [22]. Specifically, we modeled flight time lost to carry Intel i5 CPU, DDR4 RAM and SSD. We compared three aircraft: (1) Spark, a 300g UAV that can carry 500g; (2) Mavic, a 734g UAV that can carry 1300g; And (3) Matrice 100, a 2400g enterprise UAV that carries 3600g. For Spark, onboard CPU and RAM would degrade flight time by 20%. The full compute system would degrade flight time by 50%. For Mavic and Matrice, the full compute system onboard would degrade flight time by 10%.

We updated our aircraft profiles to get onboard throughput. We increased energy needs for each flight action in proportion to flight time degradation caused by onboard payload. Then at each waypoint, we subtracted compute energy from aircraft capacity. For cloud throughput, we deployed **2x2c** set up using an AWS micro instance as the second processor. This led to a 12% throughput degradation due to moving images between the edge and cloud.

Figure 7(b) explores the relationship between throughput per dollar and aircraft lifetime (measured in missions). This figure uses the Spark aircraft. Throughput per dollar of the cloud system remains static. We used AWS on-demand micro instances for pricing. Onboard and edge systems have overhead cost that cloud systems do not, but minimal maintenance costs, meaning they experience gains in total throughput per dollar as the system is used.

The crossover point is where onboard and edge systems become more cost effective compared to cloud systems. Using our 4c configuration with a high throughput autonomy setting and a DJI Spark, an edge system would become cost effective after only 10 missions. Moving the system onboard takes 5X as long to cross over.
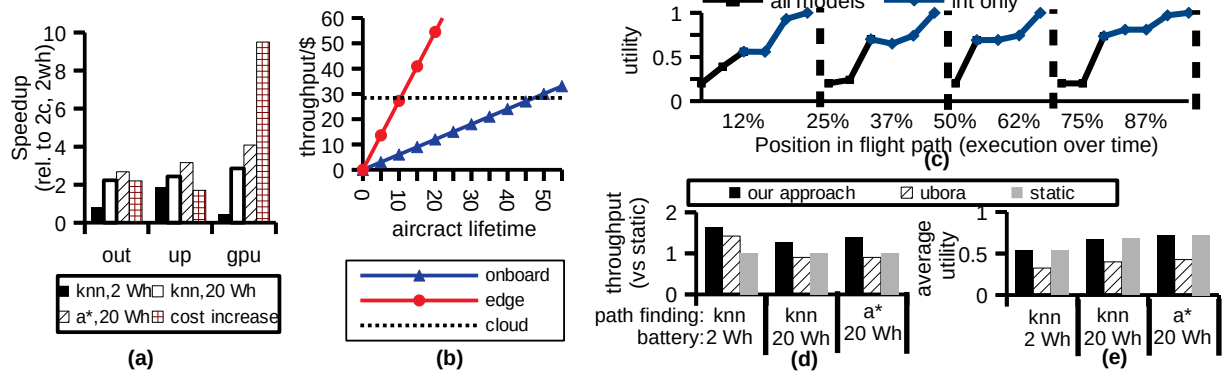
Fig. 7: Model-driven management of autonomous photography FAAS: (a) speedup and cost by architecture optimization, (b) efficiency of onboard, edge, and cloud systems, (c) utility of images captured across FAAS missions with adaptive model switching, (d-e) a comparison of throughput and utility across adaptive switching approaches.

6.3   Adaptive hardware-workload co-design

End users may have many options as to which AI models they choose to deploy on their FAAS. Our benchmarks can switch between multiple models that vary in (1) recognition accuracy and (2) latency. Highly accurate models are needed to detect distant or dark objects. Less complex models suffice for clear, crisp images. However, highly accurate DNN with DLIB (deep models) require a costly, power hungry GPU. We also use the OpenCV LBP cascade classifier (int models) which, when run on a 2 core laptop, has lower latency than DNN, but also lower accuracy. Deep models can find small, unclear faces in large, noisy images, but as images become clearer, it's performance converges with that of int models.

As the performance of Deep and Int models converge, it is prudent to turn off the GPU and use only the faster Int models. This approach conserves edge battery and increases throughput by decreasing feature extraction latency.

Figure 7(c) depicts an experimentally obtained example mission sequence where the GPU is duty cycled. We set a utility threshold of 0.5, turning off the GPU and using Int models only after a 0.5 utility image was found. **All** signifies waypoints where deep models were computed whereas **Int** signifies waypoints where only Int models were computed. For waypoints occurring after the duty cycle threshold, Int models and deep models performed similarly, finding images at comparable utility and choosing the same paths.

Figure 7(d-e) explore the differences in throughput and utility of duty cycling GPU using 3 different policies:

- **Our Approach:** Assigns a user defined threshold for duty cycling. Once one image in a flight path exceeds that threshold using the DNN model, the GPU is turned off and the LBP model is used for facial recognition.

- **Ubora:** Mimicks adaptive quality management in recent research [17]. Each mission is treated as a query. GPU and Int models are toggled once at the start of each mission.

Average utility taken over the flight path is compared to a duty cycle threshold. If average utility exceeds the threshold, GPU is turned off until average utility falls below the threshold.

- **Static:** Uses deep models for all feature extraction, with no GPU duty cycling.

Figure 7(d) shows a 1.3X gain from using our duty cycling approach as compared to the static approach, and a 1.4X gain as compared to Ubora when using either A* configuration. The A* configurations both have large enough edge batteries such that they are bottlenecked by the UAV battery, so gains or losses in throughput are entirely dependent on execution time savings during feature extraction, which are realized by the GPU configuration. The Ubora approach sees a decrease in throughput as compared to both others. Using a cumulative utility threshold allows for the Ubora approach to miss local utility spikes in a high variance workload like UAV data collection. In our test configurations, Ubora duty cycled the GPU either too early or too late. Duty cycling too late (after integer models and deep models converge) causes Ubora to function like our approach, but with more GPU usage. Duty cycling too early potentially switches to integer models before accuracy converges, taking more waypoints on average to meet utility goal. This affect can be seen in Figure 7(e), where Ubora sees considerably lower average utility than both our approach and the static approach. The failings of the Ubora approach in this context, contrasted with the success of our simpler approach, demonstrate that while duty cycling models and hardware in FAAS workloads can be advantageous, one must carefully choose their duty cycling approach.

Average utility across autonomy setting is also important. As our architecture, models, and path-finding algorithms improve, so does average utility. Our KNN configuration sees 0.83X lower image utility as compared to A* using the same models. When transitioning to a deep model on the GPU

Fig. 8: Efficiency of onboard, edge, and cloud systems

configuration, we see a 1.06X improvement in average utility which can be attributed to the higher accuracy of the deep model. Using our low battery configuration (where edge battery is a throughput bottleneck), we see that our approach makes a 1.65X improvement over a non-duty cycling approach, and a 1.15X improvement over Ubora.

6.4 Speedup for autonomous photography

Our model-driven approaches found highest throughput using 4-core Intel I7 with GPU while using low utility threshold, adaptive duty cycling and A* search. This setting completed 34 missions using Spark without recharging. Compared to 2-core Intel I5 using greedy KNN to find high utility images, the 4-core setting sped up compute latency by 15X. Looking deeper, the following changes were significant:

- **Autonomy settings:** A* search and lower threshold reduced waypoints per mission, providing 4.1X speedup.
- **Using a GPU:** Up to 2.25X over other approaches.
- **Software driven power management:** 1.3X throughput increase over static GPU usage.

Combined, the best settings yield a 10.2X increase in mission throughput compared to the 2-core setting mentioned above.

6.5 End-to-end savings for crop scouting

Our scouting FAAS covers roughly 1-acre per mission and completes 15 missions per hour. 1 FAAS would require 10 hours to scout a 150-acre field. However, multiple FAAS can work in parallel to scout the whole field faster. Given a deadline, we can estimate total hardware and software cost for all parallel FAAS. In contrast, UAS require human piloting. Based on first hand experience, we assume human pilots can execute 11 lawnmower missions for $20 per hour [20]. We also model cost for UAS equipment such as batteries, compute resources, and aircraft.

Figure 8 shows the cost for parallel UAS and baseline FAAS to achieve x mission throughput per hour. Baseline FAAS (2c, autoware, KNN) outperforms human piloted UAS. additional equipment and labor costs inflate UAS costs. Our model-driven approaches improve mission throughput significantly, gaining 6X and 4.2X on UAS and baseline. Adaptive GPU power cycling provides further improvements. Model-driven, adaptive FAAS reduce costs by 87% compared to

human-piloted UAS.

7 Related Work

FAAS choose their flight path at runtime similar to self-driving cars. Workload settings that affect their flight path can change energy usage and throughput significantly. We quantified this and proposed autonomy cubes to capture representative traces when settings change. Autoware is a project designed to make autonomous driving more open [15]. Autoware presents open algorithms, libraries, and consumer hardware components for autonomous driving, many of which are applicable to FAAS. It's motion planning design, as referenced in section 5, was improved upon in this paper. Lin et. al extend Autoware to study accelerators [21]. Object detection and tracking for self-driving cars can be sped up 169X using consumer grade hardware, but compute speedup can reduce driving range. This result parallels our observations with mission throughput. Aerostack [26] presents an open source, component based software architecture for aerial robotics, emphasizing full autonomy. Aerostack's design influenced the design and implementation of our own FAAS software.

Other recent studies explore accerlation and edge devices. Sirius [14], an intelligent personal assistant examined the affects of FPGAs, CPUs, GPUs, and coprocessors on a series of personal assistant benchmarks. In-situ AI [28] is an autonomous computing framework for IoT and studied GPUs, FPGAs, and analytical models to speed up model training. DjiNN [13] is a DNN service framework for data centers, exploring various acceleration options.

8 Conclusion

Unmanned aircraft are changing industries from agriculture to surveillance and photography. Fully autonomous aerial systems are piloted by software—eschewing costly and mistake prone human piloting. Software and hardware settings affect where these systems fly and when missions complete. Recent benchmarking papers use few settings, e.g., from prior traces, but extrapolate throughput broadly. This paper presents a modeling approach that can model flight paths across autonomy settings. Autonomy cubes provide sensed data for any reachable waypoint, enabling our approach. We have collected autonomy cubes for real FAAS executing diverse missions across a wide range of settings. Our model predicts FAAS throughput within 4%. We used our model to evaluate system mangement problems and uncovered insights that can improve throughput 10X and reduce costs 87%.

References

[1] E. Ackerman. Skydio announces sdk to make world's cleverest drone even cleverer. https://spectrum.ieee.org/.

[2] Brendan Barry, Cormac Brick, Fergal Connor, David Donohoe, David Moloney, Richard Richmond, Martin O'Riordan, and Vasile Toma. Always-on vision processing unit for mobile applications. *IEEE Micro*, 35(2), 2015.

[3] Behzad Boroujerdian, Hasan Genc, Srivatsan Krishnan, Wenzhi Cui, Aleksandra Faust, and Vijay Reddi. Mavbench: Micro aerial vehicle benchmarking. In *MICRO*, 2018.

[4] T. Brechman. 3 examples showing why crop scouting pays, even in an off year. www.indianaprairiefarmer.com, 2016.

[5] CNN. 7 tips for taking better selfies. https://www.cnn.com/2013/12/11/tech/mobile/selfie-photo-tips/, 2013.

[6] Carmelo Di Franco and Giorgio C Buttazzo. Energy-aware coverage path planning of uavs. In *ICARSC*, pages 111–117, 2015.

[7] DJI. Prerequisites-dji mobile sdk documentation. https://developer.dji.com/, 2018.

[8] DJI. Spark specs. https://www.dji.com/spark/info, 2018.

[9] Matthias Faessler, Flavio Fontana, Christian Forster, Elias Mueggler, Matia Pizzoli, and Davide Scaramuzza. Autonomous, vision-based flight and live dense 3d mapping with a quadrotor micro aerial vehicle. *Journal of Field Robotics*, 33(4), 2016.

[10] H. Genc, Y. Zu, T. Chin, M. Halpern, and V. J. Reddi. Flying iot: Toward low-power vision in the sky. *IEEE Micro*, 37(6):40–51, November 2017.

[11] D. Gomez-Candon, A. De Castro, and F. Lopez-Grandos. Assessing the accuracy of mosaics from unmanned aerial vehicle (uav) imagery for precision agriculture purposes in wheat. In *Remote Sensing*, 2014.

[12] Jim Gray, Surajit Chaudhuri, Adam Bosworth, Andrew Layman, Don Reichart, Murali Venkatrao, Frank Pellow, and Hamid Pirahesh. Data cube: A relational aggregation operator generalizing group-by, cross-tab, and sub-totals. *Data mining and knowledge discovery*, 1997.

[13] J. Hauswald, Y. Kang, M. A. Laurenzano, Q. Chen, C. Li, T. Mudge, R. G. Dreslinski, J. Mars, and L. Tang. Djinn and tonic: Dnn as a service and its implications for future warehouse scale computers. In *2015 ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA)*, June 2015.

[14] Johann Hauswald, Michael A Laurenzano, Yunqi Zhang, Cheng Li, Austin Rovinski, Arjun Khurana, Ronald G Dreslinski, Trevor Mudge, Vinicius Petrucci, Lingjia Tang, et al. Sirius: An open end-to-end voice and vision personal assistant and its implications for future warehouse scale computers. In *ASPLOS*, 2015.

[15] Shinpei Kato, Eijiro Takeuchi, Yoshio Ishiguro, Yoshiki Ninomiya, Kazuya Takeda, and Tsuyoshi Hamada. An open approach to autonomous vehicles. *IEEE Micro*, 35(6), 2015.

[16] Shinpei Kato, Shota Tokunaga, Yuya Maruyama, Seiya Maeda, Manato Hirabayashi, Yuki Kitsukawa, Abraham Monrroy, Tomohito Ando, Yusuke Fujii, and Takuya Azumi. Autoware on board: enabling autonomous vehicles with embedded systems. In *International Conference on Cyber-Physical Systems*, 2018.

[17] Jaimie Kelley, Christopher Stewart, Nathaniel Morris, Devesh Tiwari, Yuxiong He, and Sameh Elnikety. Obtaining and managing answer quality for online data-intensive services. In *ACM Transactions on Modeling and Performance Evaluation of Computing Systems*, 2017.

[18] Sami Khanal, John Fulton, Nathan Douridas, Andrew Klopfenstein, and Scott Shearer. Integrating aerial images for in-season nitrogen management in a corn field. *computers and electronics in agriculture*, 148, 2018.

[19] Matthias Kovatsch, Martin Lanter, and Zach Shelby. Californium: Scalable cloud services for the internet of things with coap. In *International Conference on the Internet of Things*, 2014.

[20] Meg Kummerow. Fly the farm. http://www.flythefarm.com.au/, 2018.

[21] Shih-Chieh Lin, Yunqi Zhang, Chang-Hong Hsu, Matt Skach, Md E Haque, Lingjia Tang, and Jason Mars. The architectural implications of autonomous driving: Constraints and acceleration. In *ASPLOS*, 2018.

[22] Robert Mahony, Vijay Kumar, and Peter Corke. Multirotor aerial vehicles. *IEEE Robotics and Automation magazine*, 20(32), 2012.

[23] Lockheed Martin. The future of autonomy isn't human-less. it's human more. https://www.lockheedmartin.com/en-us/capabilities/autonomous-unmanned-systems.html, 2018.

[24] Lorenz Meier, Petri Tanskanen, Lionel Heng, Gim Hee Lee, Friedrich Fraundorfer, and Marc Pollefeys. Pixhawk: A micro aerial vehicle design for autonomous flight using onboard computer vision. *Autonomous Robots*, 33(1-2), 2012.

[25] Jose Luis Sanchez-Lopez, Ramón A Suárez Fernández, Hriday Bavle, Carlos Sampedro, Martin Molina, Jesus Pestana, and Pascual Campoy. Aerostack: An architecture and open-source software framework for aerial robotics. In *International Conference on Unmanned Aircraft Systems*, 2016.

[26] Jose Luis Sanchez-Lopez, Martin Molina, Hriday Bavle, Carlos Sampedro, Ramón A Suárez Fernández, and Pascual Campoy. A multi-layered component-based approach for the development of aerial robotic systems: the aerostack framework. *Journal of Intelligent & Robotic Systems*, 88, 2017.

[27] Popular Science. Dji wants you to develop software for their drones. https://www.popsci.com/, 2015.

[28] M. Song, K. Zhong, J. Zhang, Y. Hu, D. Liu, W. Zhang, J. Wang, and T. Li. In-situ ai: Towards autonomous and incremental deep learning for iot systems. In *High Performance Computer Architecture*, 2018.

[29] Kimon P Valavanis and George J Vachtsevanos. Future of unmanned aviation. In *Handbook of unmanned aerial vehicles*. Springer, 2015.

[30] Deepak Vasisht, Zerina Kapetanovic, Jongho Won, Ranveer Chandra, Anish Kapoor, Sudipta Sinha, Madhusudhan Sudarshan, and Sean Strätman. Farmbeats: An iot platform for data-driven agriculture. In *NSDI*, 2017.