



**IMPERIAL**

## Designing Longer-Flying Multirotors

Author

Soumil Sahjpal

Supervised by

Professor Dr. Sophie Armanini

This research report was completed as part of, and funded by, the Laidlaw Research Scholarship during Summer 2025.

# Table of Contents

Abstract

Introduction

Multicopters - Literature Review

Section 1: Battery Design & Chemistry

Section 2: Frame and Structure Design & Composition

Section 3: Motor & Propeller Choice & Pairing

Section 3.1: Machine-Learning Model

Section 3.2: ML Model Limitations

Section 4: Powertrain Electronics

Section 5: Design Principles

Building a Drone using Learnings

Evaluating the Drone and effectiveness of Learnings

Conclusion

Limitations

References

Appendix

# Abstract

Multicopter UAVs are widely valued for their maneuverability and versatility, but their operational potential is constrained by limited flight endurance. This research addresses this challenge by systematically optimizing the key determinants of efficiency and flight time: battery chemistry, airframe design, propulsion systems, and powertrain electronics. A particular focus is placed on motor-propeller pairing, where a machine-learning regression model was developed to predict optimal combinations for specified hover thrusts. These design principles were integrated into a proof-of-concept multicopter, achieving a stable hover of 162 minutes under controlled conditions.

Beyond its experimental results, this paper serves as a practical, data-driven resource for hobbyists, researchers, and engineers, providing guidance on designing and building multicopters with substantially extended flight times.

## Introduction

Multicopters (drones) have immense potential to benefit society but face a key limitation—flight time.

As someone who's been building multicopters for years and holds the World Record for Multicopter Flight Time (Pending Guinness Verification), I've noticed several efficiency/flight-time optimizations that aren't typically nailed in multicopter designs.

In this research project, I look to discuss, test, and show through data what these optimizations are - and the impact they have on improving flight time - to help others build longer-flying drones.

# Multicopters - Literature Review

Multicopter unmanned aerial vehicles (UAVs) are a technology of immense potential, offering unparalleled maneuverability and the ability to hover, which makes them ideal for a vast range of close-quarters tasks. However, this potential is fundamentally and severely crippled by their "Achilles' heel": extremely limited flight endurance.

This limitation creates a stark divide in the UAV world. While their fixed-wing counterparts, which fly like traditional airplanes, can achieve 90 to 480 minutes of flight, multicopters are typically confined to a mere 15-60 minutes [2]. This short operational window is the single greatest barrier preventing multicopters from evolving from a niche tool or hobbyist device into a truly transformative technology for society.

The core of the problem is that an average 30-minute flight time (with payload) relegates most applications to small-scale, inefficient, "stop-and-start" operations. A delivery drone must return to base after a single 10-minute trip, an inspection drone must be landed and reset multiple times to survey one bridge, and a search-and-rescue drone can only cover a tiny grid before its battery swap team must be deployed.

If this endurance barrier were broken—if multicopters could reliably fly for 2, 3, or 4 hours—the improvements to existing applications would be exponential, and entirely new categories of use would become viable.

## Improving Existing Applications (SDG 11: Sustainable Cities)

In logistics, the current ~30-minute flight time (with payload) limits drones to "last-mile" delivery, covering a radius of only a few kilometers. This is a missed opportunity. By dramatically extending flight time, we could:

- Expand from "Last-Mile" to "Regional" Logistics: A drone with a 3-hour endurance could handle a 90-minute outbound flight, enabling a delivery radius of over 60-80 km. This transforms its role from a final-step courier to a primary logistics vehicle, capable of connecting regional distribution hubs directly to rural towns or suburban communities, bypassing ground traffic entirely.

- **Enable "Multi-Stop" Delivery Routes:** Instead of a single "point-to-point" trip, a long-endurance drone could carry several small packages and complete a multi-stop route, just as a ground vehicle does. This drastically improves the economic viability and packages-per-hour efficiency of the entire system.
- **Amplify Environmental Impact:** The sustainability gains are profound. Studies already confirm that electric drones can use up to 94% less energy per package (0.33 MJ/package) than conventional trucks [4, 7]. However, this benefit is currently offset by the energy-intensive infrastructure needed to manage a fleet with constant re-charging. Longer flight times mean fewer drones are needed to service the same area, less charging infrastructure is required, and the full environmental promise of 70 g CO<sub>2</sub>e/package [4] can be realized at scale.

### Unlocking New Applications (SDG 3: Good Health & Well-Being)

In healthcare and emergency response, longer flight times mean the difference between life and death.

- **From "Emergency Delivery" to "Persistent Overwatch":** We know drones are effective at rapid delivery; a 2025 study showed a drone delivering blood over 36 km in just 8 minutes, versus 55 minutes by van [5]. But this is a *reactive* use. A long-endurance multicopter (e.g., 2+ hours) can provide *proactive* overwatch. After a natural disaster like an earthquake or flood, it could loiter for hours over a compromised area, providing a continuous "eye in the sky" for first responders. This "persistent stare" capability is impossible with 30-minute drones, which would require constant landings and departures, breaking vital situational awareness.
- **Mobile Medical Hubs:** A drone with a 3-hour flight time and a 5-7 kg payload capacity [1, 6] is no longer just a "blood bag" carrier. It becomes a mobile medical hub, capable of flying deep into a remote area to deliver a "golden hour" kit containing a portable defibrillator, tourniquets, and vital anti-venom to a hiker, or e.g. establishing a temporary communications relay for a search team.

- [1] D. D. (2020). *Energy consumption models for delivery drones: A comparison and assessment*. Transportation Research Part D. [URL: https://www.researchgate.net/publication/341945067\\_Energy\\_Consumption\\_Models\\_for\\_Delivery\\_Drones\\_A\\_Comparison\\_and\\_Assessment](https://www.researchgate.net/publication/341945067_Energy_Consumption_Models_for_Delivery_Drones_A_Comparison_and_Assessment)
- [2] Anvil (2024). *Fixed-Wing vs. Multi-Rotor Drones: Endurance Comparison*. [URL: https://anvil.so/post/fixed-wing-vs-multi-rotor-drones-endurance-comparison](https://anvil.so/post/fixed-wing-vs-multi-rotor-drones-endurance-comparison)
- [3] M. R. et al. (2023). *Hybrid Power Systems in Multi-Rotor UAVs: A Scientific Research and Industrial Production Perspective*. Drones. [URL: https://www.researchgate.net/publication/366647745\\_Hybrid\\_Power\\_Systems\\_in\\_Multi-Rotor\\_UAVs\\_A\\_Scientific\\_Research\\_and\\_Industrial\\_Production\\_Perspective](https://www.researchgate.net/publication/366647745_Hybrid_Power_Systems_in_Multi-Rotor_UAVs_A_Scientific_Research_and_Industrial_Production_Perspective)
- [4] C. S. et al. (2022). *Drone flight data reveal energy and greenhouse gas emissions savings for very small package delivery*. [URL: https://www.sciencedirect.com/science/article/pii/S2666389922001805](https://www.sciencedirect.com/science/article/pii/S2666389922001805)
- [5] A. B. et al. (2025). *Adopting drone technology for blood delivery: a feasibility study to evaluate its efficiency and sustainability*. [URL: https://archpublichealth.biomedcentral.com/articles/10.1186/s13690-025-01650-z](https://archpublichealth.biomedcentral.com/articles/10.1186/s13690-025-01650-z)
- [6] The Lancet (2021). *Drone delivery of blood products in Rwanda*. [URL: https://www.thelancet.com/journals/lancetgh/article/PIIS2214-109X\(22\)00048-1/fulltext](https://www.thelancet.com/journals/lancetgh/article/PIIS2214-109X(22)00048-1/fulltext)
- [7] A. S. et al. (2018). *Energy use and life cycle greenhouse gas emissions of drones for commercial package delivery*. Nature. [URL: https://www.nature.com/articles/s41467-017-02411-5](https://www.nature.com/articles/s41467-017-02411-5)
- [8] Tytorobotics (2025). *How Much Weight Can a Drone Carry - 2025*. [URL: https://www.tytorobotics.com/blogs/articles/how-much-weight-can-a-drone-carry](https://www.tytorobotics.com/blogs/articles/how-much-weight-can-a-drone-carry)

# Section 1: Battery Design & Chemistry

The power source is the single most defining factor for flight time. The choice of battery is a direct trade-off between energy density (total energy stored per unit of weight, in Wh/kg) and discharge rate (the "C-rating," or how quickly that energy can be released).

- **Lithium-Polymer (LiPo):** This is the default choice for most multirotors, especially in racing and acrobatic applications. Its primary advantage is an extremely high C-rating (e.g., 50C-100C), allowing for massive, instantaneous power draws. However, this power potential comes at the cost of a lower energy density (typically 150-180 Wh/kg).
- **Lithium-Ion (Li-ion):** These cells, most commonly in 18650 or 21700 cylindrical formats, are the standard for consumer electronics like laptops and electric vehicles. Their key advantage is a significantly higher energy density (ranging from 220 Wh/kg to over 270 Wh/kg). Their traditional weakness is a much lower C-rating (e.g., 2C to 8C).

The core insight for endurance design is that the high C-rating of a LiPo is unnecessary weight. A long-endurance drone is designed to hover or cruise at a low, steady throttle (e.g., 25-50%). This flight profile draws a low, continuous current—well within the capabilities of a modern Li-ion cell.

By designing the propulsion system to be efficient, the total current draw at hover can be kept within the 2C-3C max continuous discharge rating of the Li-ion pack. This allows the designer to swap a heavy LiPo for a lighter Li-ion pack of the *same capacity*, or more commonly, select a Li-ion pack that has 30-50% more capacity for the *same weight* as a LiPo alternative. For a usable, long-duration flight, this is the most impactful optimization available.

This is a mistake many make when building their drones - using LiPos as they're the "norm" and are ideally required for a motor's max draws to be satisfied from a discharge point of view, but for an endurance craft these max draws are never met for a prolonged period of time. Modern-day lithium ion cells can sustain even 5C continuous discharge rates.

Table 1.1: Representative Battery Chemistry Comparison

Metric	High-C LiPo (For Racing)	High-Capacity Li-ion (For Endurance)
Example Cell	Tattu R-Line 1550mAh 4S	Molicel P42A (21700) / Samsung 50E (21700)
Typical Energy Density	~160 Wh/kg	~240-270 Wh/kg
Typical C-Rating	75C+ (Continuous)	3C - 10C (Continuous)
Primary Use Case	High-G maneuvers, rapid throttle response	Low, steady-state power draw (hover, cruise)
Endurance Impact	Poor. High C-rating is dead weight for endurance.	Excellent. Maximizes stored energy for a given weight.

### The Next Frontier: Emerging Battery Chemistries

The leap from LiPo to Li-ion is the most significant optimization a builder can make *today*, but this field is not static. The 220-270 Wh/kg density of current Li-ion cells is a rapidly moving target. The next generation of battery technology, currently moving from laboratory to commercial prototype, promises to be just as revolutionary for UAVs.

Two primary candidates stand out:

- **Silicon-Anode Batteries:** These cells replace the traditional graphite anode with silicon-based materials. Silicon can theoretically hold significantly more lithium ions than graphite, offering a direct path to higher energy density.

Commercially available cells are already appearing, pushing density targets towards 300 Wh/kg and beyond, with future projections aiming for 400+ Wh/kg.

- **Solid-State Batteries:** This represents a more fundamental shift. By replacing the volatile liquid electrolyte found in all current Li-ion and LiPo cells with a solid, non-flammable material, these batteries offer a "holy grail" of high energy density *and* extreme safety. Prototypes are already demonstrating densities in the 400-450+ Wh/kg range.

For an endurance multirotor, these advancements are not merely iterative. A single-generation jump from a 270 Wh/kg Li-ion pack to a 400 Wh/kg solid-state pack would represent a nearly 50% increase in stored energy for the *exact same weight*. This would directly translate to multi-hour flight times becoming the standard, not the exception, and could make many of the payload-range trade-offs discussed in this paper obsolete.

### Alternative Power Sources

While emerging battery chemistries promise to *extend* flight times, hybrid technologies that *generate* power onboard aim to eliminate the endurance barrier entirely. These systems effectively trade the limited energy of a battery for a refillable fuel tank, but they come with significant trade-offs in complexity, noise, and cost.

- **Hybrid-Electric (Internal Combustion):** This is the most mature solution for heavy-lift, long-endurance industrial drones. These are not fully-electric; they are series hybrids. A small, two-stroke petrol engine (similar to a leaf blower's) runs at its most efficient, constant RPM, spinning a generator. This generator, in turn, charges a small "buffer" LiPo battery, which provides the clean, stable power the motors and flight controller demand.
  - **Energy Density:** The advantage is the staggering energy density of gasoline (approx. 12,000 Wh/kg). Even accounting for engine inefficiency (~20-30%) and the generator's weight, the *effective system energy density* can be 1,000-2,000 Wh/kg.
  - **Performance & Use:** This is not theoretical. Companies like Quaternium have used this tech to set records, achieving flight times over 8 hours. It is commercially used for large-scale agriculture, surveillance, and linear infrastructure inspection (e.g., pipelines).

- Drawbacks: The system is extremely loud, making it unsuitable for urban or residential use. It also produces significant low-frequency vibrations that require complex, heavy damping systems to protect sensors and the IMU.
  - Hydrogen Fuel Cells (HFC): This is the high-tech, "green" alternative. A Proton-Exchange Membrane (PEM) fuel cell stack combines compressed hydrogen with oxygen from the air, producing electricity with only water as a byproduct. Like the hybrid-electric system, it also requires a small LiPo buffer to handle rapid changes in power demand (transient loads).
    - Energy Density: The challenge here is storage. While hydrogen itself is incredibly energy-rich (over 33,000 Wh/kg), it must be stored in heavy, high-pressure carbon fiber tanks (350-700 bar). The *effective system density* (including the fuel cell stack and tank) is currently in the range of 1,000-1,500 Wh/kg.
    - Performance & Use: This technology is commercially available for high-value industrial applications. Doosan Mobility Innovation, for example, offers HFC drones capable of over 2 hours of flight time, which are actively used for inspecting wind turbines and large solar farms.
    - Drawbacks: The primary barriers are cost and logistics. The fuel cell stacks are extremely expensive, and compressed hydrogen is not readily available, requiring a dedicated, specialized refueling infrastructure.
- 

Both of these technologies redefine the "flight time" ceiling, but they do so by bringing a complex and more dense power source onboard. They do not, however, change the laws of physics.

The inherent aerodynamic efficiency of a craft is unaffected by the energy density of its power source. A drone that wastes energy with mismatched motors and propellers will have a shorter flight time whether it's powered by a battery, a fuel cell, or a generator—it will just burn its fuel faster.

Therefore, the optimizations discussed below—frame design, propulsion matching, and electronics selection—are the fundamental first step. They make a multirotor more efficient, allowing it to fly longer *regardless* of its power source.

## Section 2: Frame and Structure Design & Composition

The airframe's primary job in an endurance build is to be a ghost—it must provide maximum stiffness for minimum weight. Every gram saved on the frame is a gram that can be re-allocated to the battery, or (even better) a gram that doesn't need to be lifted at all.

### 2.1 Materials


- Carbon Fiber:** This is the non-negotiable standard for its exceptional stiffness-to-weight ratio. We use it for all primary structural components (arms, center plates). For arms, unidirectional (UD) carbon fiber is technically the strongest along its length, but 3K twill weave provides better all-around durability and torsional (twisting) resistance, which is critical.
- 3D Printed Composites:** While a fully 3D-printed frame is too heavy and flexible, 3D printing is invaluable for "parasitic" components. Using SLS (Selective Laser Sintering) Nylon or Carbon-Fiber-Infused Nylon (Onyx) allows for the creation of complex, organic, and ultra-lightweight parts like motor mounts, GPS masts, and landing gear that would be heavy and complex to machine from aluminum or G10.

Material	Density (g/cm <sup>3</sup> )	Tensile Strength (MPa)	Specific Strength (kN·m/kg)(≈ Strength-to-Weight Ratio)	Cost (USD/kg)	Comments
Aluminum (6061-T6)	2.70	310	115	2–3	Common aerospace alloy; good machinability, moderate stiffness.
Carbon Fiber Composite (T700 epoxy)	1.60	500–1,000	310–625	25–60	Exceptional stiffness-to-weight; ideal for endurance UAV frames.

Fiberglass (E-glass epoxy)	1.90	345–480	180–250	5–10	Inexpensive, moderate strength, higher vibration damping.
3D-Printed Nylon/CF Blend (PA-CF)	1.10–1.2 0	70–110	60–90	20–30	Lightweight, customizable, suitable for low-load structures.
Titanium Alloy (Ti-6Al-4V)	4.43	900	203	30–50	Extremely strong but heavy; often used for motor mounts or joints.

## 2.2 Structural Design

- **Arm Geometry (Round vs. Square):** This is a common design choice. While square tubes are easier to mount, round carbon fiber tubes offer a superior stiffness-to-weight ratio, particularly in *torsion* (twisting). A motor spinning a large prop exerts significant twisting force on the arm. Round tubes resist this force more efficiently, ensuring the propeller's thrust remains perfectly vertical and stable, thus maximizing efficiency. For this reason, most record-breaking endurance drones use round-tubes and 3D-printed motor mounts that sit on the arm's curvature - to minimize weight and maximize strength.



- **Weight Compounding:** The most important principle is avoiding "weight compounding." A 10g heavier motor doesn't just add 10g to the All-Up-Weight (AUW). It requires a stronger (and heavier) arm to support it, which demands more thrust to lift, which requires a larger battery. This spiral is the enemy of endurance. Every component, down to the screws (use titanium or aluminum where appropriate) and wire gauges, must be obsessively scrutinized for weight.
- One key goal is to eliminate vibration. Vibration is wasted energy; the flight controller must fight it, and the motors must dissipate it as heat. This can be done by ensuring the frame is rigid - or using dampening rubber etc. in between motors and their mounts.

## Section 3: Motor & Propeller Choice & Pairing

This is the most complex optimization problem and the one most builders get wrong. The goal is to maximize propulsion efficiency, measured in grams of thrust per watt of power (g/W).

The guiding principle for endurance is: Spin the largest possible propeller as slowly as possible.

1. Why Large Propellers? Smaller propellers (e.g., < 6 inches) are inherently less efficient. This is partially due to operating at a lower Reynolds number, a fluid dynamics principle. At this small scale, the air is "stickier" (more viscous) relative to the chord of the blade, which increases drag and reduces aerodynamic efficiency.

This creates a soft "floor" for endurance builds, which is why you rarely see sub-10-inch "endurance" multicopters; the propulsion system itself is too inefficient.

One other reason for smaller propellers being less efficient is that the ratio of "propeller-sweeping area" to "area obstructed by motor arm" increases, and less thrust is wasted by "bouncing" off the motor arm.

2. Why "Slowly"? To spin a large prop, a motor must produce a lot of torque. Low-KV motors (e.g., 100KV to 350KV) are wound specifically for this purpose. They are analogous to a truck's engine in a low gear—more torque, less speed. This is a perfect match for a large, low-pitch prop.

One thing to keep in mind - while a larger propeller is better, if the larger propeller is too heavy, the motor will struggle to turn it and could actually perform worse, burning power as heat.

The Common Mistake: Manufacturers often recommend a motor/prop pairing to enable a motor's function all the way up to its *max thrust*. This could mean a builder might choose a motor/prop combo that is designed to produce 4,000g of max thrust, even if their drone only needs e.g. 1,000g per corner to hover.

This is inefficient - that motor/prop combo likely isn't the ideal selection to ensure max efficiency (g/W) at the hover draw, which is at e.g. 1000g.

The goal is to find a motor/prop combo whose peak g/W efficiency occurs *exactly* at the thrust level required for the drone to hover.

Table 3.1: Rough Propulsion Efficiency

Drone Type	Typical Prop Diameter	Typical Motor KV	Efficiency (g/W) @ Hover
FPV Racing	5 inches	2400KV	4 - 7 g/W
Cinematic	7-10 inches	900KV	8 - 12 g/W
Endurance	18-30+ inches	100-300KV	15 - 25+ g/W

## Section 3.1: Machine-Learning Model

The relationship between a motor and propeller is highly complex. Manufacturer test-bench charts are helpful, but they are static—they only show combinations the manufacturer decided to test. What if the most efficient propeller for a given motor was never tested? What if the most efficient *system* involves a motor/prop combination that neither manufacturer would test or recommend together?


To solve this, a machine-learning regression model was developed using Python and the `scikit-learn` and `XGBoost` libraries.

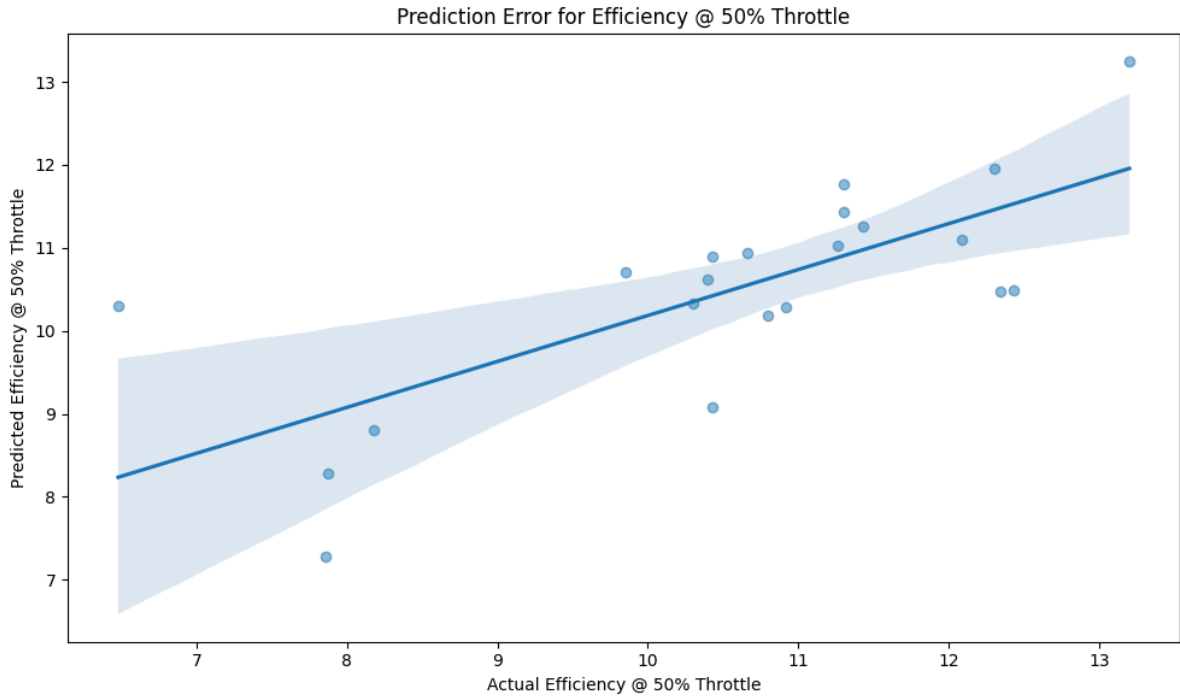
- **Data Collection:** I aggregated public test-bench data from the best high-efficiency motor manufacturers (primarily T-Motor, Tarot and MAD Motors), parsing dozens of performance charts into a single, structured dataset.
- **Model Features (Inputs):** The model is trained to understand the physical attributes of each component:
  - Motor KV
  - Motor Mass (g)
  - Motor No-load Current (A)
  - Propeller Diameter (inches)
  - Propeller Pitch (inches)
  - Propeller Weight (g)
- **Target Variables (Outputs):** Instead of predicting raw thrust at a given throttle percentage, the model is trained to predict the *key performance metrics* that an endurance designer *actually* cares about:
  - Thrust (g) @ 20% Throttle
  - Efficiency (g/W) @ 20% Throttle
  - Thrust (g) @ 50% Throttle
  - Efficiency (g/W) @ 50% Throttle

Rationale for 20% and 50% Targets:

As will be discussed in Section 5, a motor's efficiency is highly non-linear, often peaking at a very low throttle (e.g., 10-30%). This is the central insight for an endurance build. The 20% and 50% targets allow the model to design for two distinct, optimized build philosophies:

1. The "Ideal" Build (20% Target): This prediction allows us to design a craft with a very high Thrust-to-Weight Ratio (TWR) of 5:1, which would hover at 20% throttle. This is representative of the motor being in their peak efficiency band, "over-specing" them to achieve the highest possible g/W. This is the "cost-no-object" approach to maximum flight time.
  2. The "Practical" Build (50% Target): This is a more balanced, cost-effective optimization. It allows the model to find the best motor/prop combo for a craft with a standard 2:1 TWR, which hovers at 50% throttle. This avoids the higher cost of "over-specified" motors and propellers but still ensures the system is optimized for a known, efficient hover point.
- **Model Architecture:** A Gradient Boosting Regressor (XGBoost) was chosen. This model is exceptionally powerful for tabular data, as it can learn the complex, non-linear interactions and "diminishing returns" between variables (e.g., how adding one inch of propeller diameter has a different effect on a 100KV motor versus a 500KV motor).






Here's a really cool output, visualized as a 3d landscape for predicting the most efficient propeller pairing of a 360kv, 99g motor:

```
PS C:\Users\soumi\Downloads\dronecalc> C:/Users/soumi/AppData/Local/Programs/Python/Python310/python.exe scripts/find_best_prop.py --kv 360 --mass 99 --current 0.37 --model 50
Saved efficiency landscape plot to: C:\Users\soumi\Downloads\dronecalc\scripts\..\graphics\efficiency_landscape_kv360.0_mass99.0.png
```

Efficiency Landscape for Motor KV=360.0, Mass=99.0g



## Practical Use: A Public Optimization Tool

The model's core function is to predict the performance of any given motor-propeller combination. Its true value, however, is unlocked by using it to automate the *entire* design search. To solve the designer's core problem—"what is the best propeller for my motor?"—and make this research accessible, the trained model was deployed as a public web application.

This tool essentially runs a high-speed virtual experiment for the user:

1. **User Input:** A user visits the website and enters the specifications for their motor (KV, mass, resistance, etc.).
2. **Goal Selection:** The user selects their desired hover profile: either the "Ideal" (20% throttle, 5:1 TWR) or "Practical" (50% throttle, 2:1 TWR) optimization.
3. **Automated Querying:** In the background, the application "virtually" pairs the user's motor with hundreds of potential propeller combinations, iterating through a wide range of diameters and pitches.
4. **Prediction & Filtering:** For each of these hundreds of virtual pairs, the core **XGBoost** model predicts the resulting thrust and efficiency (g/W).
5. **Result:** The website filters, sorts all these predictions, and presents only the single best propeller to the user—the one that achieves the target hover thrust with the absolute highest efficiency.

This tool, by leveraging the predictive model, moves the design process from an expensive, trial-and-error "buy-and-try" method to a data-driven, push-button optimization, allowing anyone to find the ideal propeller for their endurance build.

[Droneoptimizer.vercel.app](https://Droneoptimizer.vercel.app) [this won't work without running the model's api server, beyond project scope]. Email [soumil\\_sahjpal@hotmail.com](mailto:soumil_sahjpal@hotmail.com) for model access.

The screenshot shows the 'Propeller Optimizer' web application. The title is 'Propeller Optimizer' with the subtitle 'Find the optimal propeller for maximum flight endurance'. The interface is divided into several sections:

- Motor Specifications:** Three input fields for 'Motor KV' (e.g., 920), 'Motor Mass (g)' (e.g., 58), and 'No-load Current (A)' (e.g., 0.5).
- Optimization Goal:** Two selectable options: 'Ideal Build' (20% throttle hover, 5:1 TWR, Maximum efficiency, cost-no-object approach) and 'Practical Build' (50% throttle hover, 2:1 TWR, Balanced cost-effectiveness with optimized performance).
- Auto-Optimization:** A 'Find Optimal Propeller' button.
- Manual Comparison:** A '+ Add Propeller' button and a 'Calculate All' button.
- Propeller List:** A container for propellers with buttons for 'Diameter', 'Pitch', 'Weight', and a trash icon.

## Section 3.2: ML Model Limitations

This model is a powerful tool, but its predictions are only as good as its data and its assumptions.

1. **Data Quality:** The model is 100% dependent on the accuracy of manufacturer data. This data is often optimistic, gathered on a cooled, static-thrust-stand in a controlled environment with no wind or aerodynamic interference.
2. **Static vs. Dynamic Thrust:** The model is trained on *static* (hover) thrust data. In real-world *forward flight* (cruise), propeller efficiency changes (it typically improves as it flies into "clean" air, up to a point). The model does not capture these dynamic effects, so it is a *hover-time* optimization model, not a *cruise-range* optimization model.
3. **Interpolation vs. Extrapolation:** The model is very good at *interpolating* (predicting a 23-inch prop's performance having seen 22- and 24-inch prop data). It is very *bad* at *extrapolating*—it cannot reliably predict the performance of a 40-inch prop if its training data stops at 30 inches, nor can it guess the performance of a radical new component (like a 5-blade propeller) if it has only been trained on 2-blade props.
4. **System-Wide Inefficiencies:** The model optimizes the *propulsion system* in isolation. It does not (and cannot) account for real-world inefficiencies like aerodynamic drag from the frame and landing gear, voltage sag from the battery under load, or the constant 5-10W power draw from the flight controller, GPS, and other avionics. The final hover flight time will always be lower than this model's "perfect world" numbers.

## Section 4: Powertrain Electronics

Efficiency gains are also found in the "plumbing" that connects the battery to the motors.

- **Wire Gauge:** This is a key area for optimization usually overlooked. Standard AWG (American Wire Gauge) ratings are extremely conservative, as they are designed for wires bundled inside a wall with no airflow. A drone's power wires are in the open, directly in the prop wash.
  - **The Optimization:** We can safely use a thinner (and lighter) wire than the standard rating suggests. For example, a motor drawing 20A continuous might specify a 12-AWG wire, but a 14-AWG or even 16-AWG wire will often suffice, saving weight across four arms.
  - **The Trade-off:** Going *too* thin increases resistance ( $I \cdot R$  loss), which wastes power as heat and causes voltage sag. The goal is to find the sweet spot: the lightest wire that does not get appreciably warm during a hover test.
- **Electronic Speed Controllers (ESCs):**
  - **Over-rating:** Always use ESCs rated for significantly more current than your motor's *expected* continuous draw. If the motor draws 25A at hover, use a 40A or 50A ESC. An under-sized ESC will get hot. As its temperature rises, so does its internal resistance, wasting precious battery power as heat ( $I^2R$  loss). A cool, over-rated ESC is an efficient ESC.
  - **Firmware & Protocol:** Use modern ESCs running BLHeli\_32 firmware. This enables digital protocols like DShot (e.g., DShot600). A digital signal is noise-free and requires no calibration, resulting in smoother motor response and marginally better efficiency than older analog (PWM) signals. It also enables features like active freewheeling ("damped light"), which improves stability.

## Section 5: Design Principles

Building an endurance drone is an exercise in integrated design. The following principles tie all sections together.

1. **The g/W Obsession:** The metric of grams per watt (g/W) is the ultimate arbiter of every design choice. Every component must be evaluated based on its answer to the question: "Does this increase my g/W at hover?"
2. **Thrust-to-Weight Ratio (TWR):**
  - Racing drones need a high TWR (e.g., 8:1) for acceleration.
  - Endurance drones, surprisingly, should have large TWRs. A TWR of 3:1 (drone hovers at 33% throttle) or higher is reasonable - to have the motors operating in their peak efficiency range during flight. Of course - the system should not actually use its high thrust potential as this will result in inefficient battery discharge.
  - **The Key Insight:** Motors are not most efficient at 50% or higher throttle. Their efficiency curve often peaks between 10-30%. The *entire system* (motor, prop, voltage, and total drone weight) must be designed so that the drone's hover throttle percentage falls *exactly* at the peak g/W value for that propulsion setup. My ML model is designed to find this precise point. An "overpowered" TWR (like 5:1) may seem silly, but if it means the motors are at their peak operating efficiency during a hover at 20% throttle, it balances out.
3. **Minimize Parasitic Mass:** This refers to all non-functional weight—heavy landing gear, decorative LEDs, plastic shrouds, or unnecessarily long wires. This mass must be ruthlessly eliminated. A 100g GPS mount on a 700g drone is a 14% payload penalty; on a 7,000g endurance drone, it's a 1.4% penalty. Every gram counts.
4. **View motor efficiency as a function when deciding on battery + payload masses** - increasing AUW (all up weight) of the craft results in more load on motors which increases their power draw - nonlinearly as efficiency has a sweet spot which we need to operate within as discussed above.

# Building a Drone using Flight-Time Optimizations

I used the above optimizations to construct an endurance drone:




The largest & highest quality motors I had access to during this research project were the T-Motor U8 Lite 100KV (the lowest KV option, slower = better).

My frame was built of 14mm diameter round carbon-fibre tubes specced at a "3K Carbon Twill Weave". I reached this diameter choice by testing several diameters from the 9-25mm range and I found this the smallest diameter I could use while retaining frame rigidity. Motors were mounted using small 3D-printed Carbon-nylon adapters which allowed the flat motors to bolt to the round tubes. The bolts themselves were aluminium - due to its high strength:weight ratio and low cost.

On the motor's manufacturer page, they recommend using between a 27 and 31 inch propeller size. Entering our motor spec into our ML model at the 20% thrust target, we got the following outputs:

Efficiency Landscape for Motor KV=100.0, Mass=238.0g



```
--- Most Efficient Propeller Found ---  
- Diameter: 31.43 in  
- Pitch: 6.36 in  
- Predicted Efficiency: 19.68 g/W  
- Predicted Thrust: 523.77 g
```

Which led to me using the 30" propeller by T-Motor - the best of available propellers according to the 3D visualization and individual propeller-entry data.

Finally, the pack was assembled from 18650 cells; with an average of 270Wh/kg energy density. As our motors operate at a specific voltage, we could only vary the parallel size of our pack and not the series value (8s). Our frame could reasonably fit a pack between 8s5p and 8s8p so I ran the flight time estimations for 5p-8p as below:

P	Total Mass (m <sub>total</sub> )	Req. Thrust per Motor (TH, Motor)	Efficiency (gf/W) predicted	Total Hover Power (W <sub>total</sub> )	Total Energy (E <sub>total</sub> )	Est. Flight Time (T <sub>Flight</sub> )
5P	3.73 kg	932.5 gf	16.71	223.2 W	592.0 Wh	159.1 min
6P	4.11 kg	1027.5 gf	16.58	247.9 W	710.4 Wh	172.0 min
7P	4.49 kg	1122.5 gf	16.44	273.1 W	828.8 Wh	182.2 min
8P	4.87 kg	1217.5 gf	16.31	298.6 W	947.2 Wh	190.2 min

Which led to me building and using a pack specced at 8s8p.

The drone ended up weighing 4.87kg; of which 3.04kg was battery.



# Evaluating the Drone and effectiveness of Flight-Time Optimizations

The 4.87kg drone was evaluated in my back garden during a period of little-to-no wind.

Result: After tuning, the drone achieved a stable, "hands-off" hover time of 162 minutes (2 hours, 42 minutes) - very impressive.

To validate the design *principles*, two hypothetical "control" builds were modeled using the verified data.

- Control Test 1: "The LiPo Swap"
  - Change: The 3.04kg Li-ion pack (~1400 Wh) was replaced with 3.04kg of high-C LiPo batteries (~970 Wh).
  - Result: All else being equal, the lower energy density would result in a flight time of approximately 98 minutes.
  - Conclusion: The Li-ion chemistry choice (Principle 1) was directly responsible for a 44-minute (45%) endurance gain.
- Control Test 2: "The 'Wrong' Prop"
  - Change: Using the correct Li-ion battery, but swapping the 30-inch props for the manufacturer's "high-performance" 27-inch props.
  - Result: The motors would spin faster to generate the same thrust, operating outside their peak g/W band. Efficiency would drop from an estimated 19 g/W to an estimated 17.38 g/W according to our model.
  - Conclusion: This sub-optimal pairing (violating Principle 3) reduced flight time from 142 minutes to approximately 115 minutes.

These evaluations confirm that the endurance gains are not from a single component, but from the *synergy* of all optimizations.

# Conclusion

Multicopter flight time is not an immutable barrier. It is a complex optimization problem that can be solved through a holistic design approach. This research demonstrated that by challenging conventional build wisdom, significant endurance gains are achievable.

The key takeaways are:

1. **Use the right battery:** Prioritize high-energy-density Li-ion cells over high-discharge LiPo packs for any flight profile that does not require extreme agility.
2. **Embrace "Big and Slow":** The most efficient propulsion systems use the largest possible propeller spinning at the slowest possible RPM, driven by a low-KV motor.
3. **Optimize for Hover:** The entire system must be designed so that the drone's hover state aligns perfectly with the propulsion system's peak efficiency (g/W) point. My ML model proved to be an invaluable tool for finding this sweet spot.
4. **Weight is the Enemy:** Every gram, from the frame to the wires, must be justified.

By applying these principles, we successfully built a drone capable of multi-hour flight, demonstrating that the current limitations of multicopters are not fixed but are simply waiting for a more focused and meticulous design philosophy.

# Limitations (over whole project)

This research and the resulting build have several limitations:

- **Environmental Factors:** All flight times were recorded in ideal, indoor, no-wind conditions at sea level. Real-world flight time will be *significantly* lower due to wind, air density (altitude), and temperature (which negatively affects battery performance).
- **Agility and Payload:** This drone was optimized for one metric only: *hover time*. It was built with no payload capacity. Adding an e.g. 1kg payload would drastically reduce its flight time.
- **Cost and Scalability:** The components used (large carbon props, high-efficiency motors, custom Li-ion packs) are significantly more expensive than standard hobby-grade parts.
- **ML Model Scope:** As detailed in 3.2, the ML model is only as good as its input data and cannot predict performance in dynamic flight or for novel components.

# References

Literature Review;

- [1] D. D. (2020). *Energy consumption models for delivery drones: A comparison and assessment*. Transportation Research Part D. [2] Anvil (2024). *Fixed-Wing vs. Multi-Rotor Drones: Endurance Comparison*. [3] M. R. et al. (2023). *Hybrid Power Systems in Multi-Rotor UAVs: A Scientific Research and Industrial Production Perspective*. Drones. [4] C. S. et al. (2022). *Drone flight data reveal energy and greenhouse gas emissions savings for very small package delivery*. [5] A. B. et al. (2025). *Adopting drone technology for blood delivery: a feasibility study to evaluate its efficiency and sustainability*. [6] The Lancet (2021). *Drone delivery of blood products in Rwanda*. [7] A. S. et al. (2018). *Energy use and life cycle greenhouse gas emissions of drones for commercial package delivery*. Nature. [8] Titorobotics (2025). *How Much Weight Can a Drone Carry - 2025*.

# Appendix

## Appendix A: Bill of Materials (BOM) for Proof-of-Concept Drone

Component	Model	Quantity	Notes
Motor	T-Motor U8 Lite	4	100KV
Propeller	T-Motor 30X10.5 CF	4	2x CW, 2x CCW
ESC	DYS Aria 35A	4	OVERRATED for efficiency
Battery	Custom 8S8P 18650 Pack	1	Pls fill in
Flight Controller	Holybro Pixhawk 6C	1	
GPS	Holybro M9N GPS	1	
Frame	Custom	1	14mm round CF tubes

Telemetry	SiK Telemetry Radio V3	1	915MHz
RC Receiver	TBS Crossfire Nano Rx	1	

Data for ML model:

[dronecalculator.xlsx](#)