Small Wind-Turbine Blade: Design Project Final Report

Members: Gregory Sharma, Jacob Gonzalez, Michael Alexiadis, Ata Zavaro

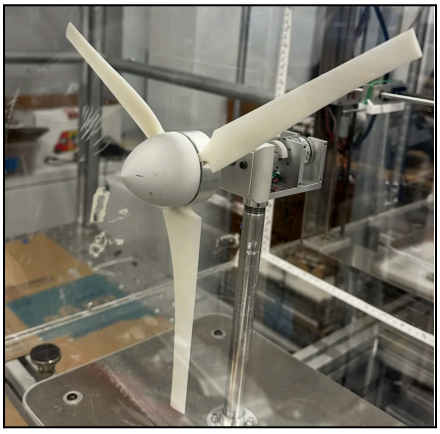

Executive Summary

The project designed and optimized a small wind turbine blade for low-wind conditions, then compared modeled predictions with experimental performance to evaluate design assumptions. To break this project into something more approachable, certain assumptions were specified. These assumptions simplify the aerodynamic analysis allowing efficient evaluation of blade geometries using steady, 2D flow, negligible friction and 3D effects, and theoretical Cl/Cd data at Re ≈ 50,000. While some real-world effects are ignored due to their assumed insignificance in the overall design, this approach gives our team a starting point.

With this, the objective is to maximize performance by achieving a target operating point of 1739 RPM, 3.31 N·cm of torque, and 6.02 W of power. Blade parameters, including pitch, twist, and chord distribution, are optimized using differential evolution, with low–Reynolds-number airfoils such as the NACA 4412 selected for their superior aerodynamic efficiency.

Design Process

On a long-term project with multiple parts, requirements, and engineers, it was critical to establish a workflow that distributed work efficiently and respected the required sequence of

design steps. To visualize this process, Figure 1 illustrates the parallel workflow used in this project, where the numerical optimizer (code) and the physical blade design (CAD) were developed concurrently and merged once results were available.

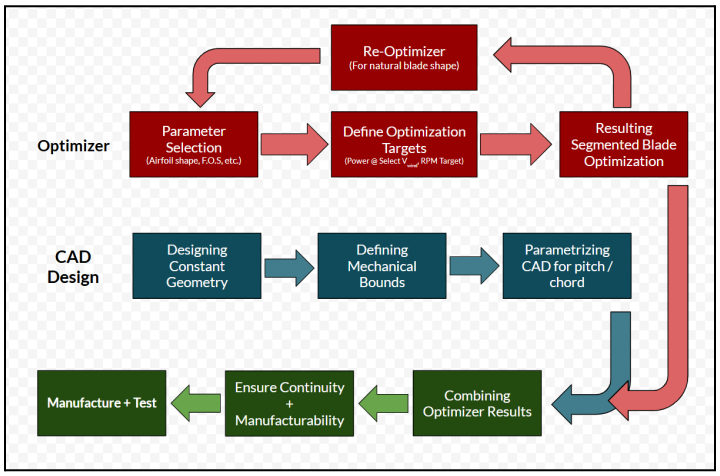

Figure 1 (Above)

The blade geometry was parameterized using pitch, chord, airfoil shape, and span, variables that correspond directly to how the blade is modeled in CAD software like Fusion. The optimizer outputs these parameters, which are then translated into a physical blade by the CAD. To handle the continuous variation of these properties along the blade span, the blade was discretized into 10 spanwise segments, each defined by its chord, twist, and airfoil placement. This segmentation mirrors the CAD modeling approach, where the blade is built by controlling the size, position, and orientation of an airfoil at discrete locations. Interpolation between segments is a CAD decision that would be discussed later in this section.

Overall, this setup creates a direct link between the optimizer and CAD: the optimizer determines the best aerodynamic parameters at discrete points, and the CAD converts these into a physical model that can be manufactured. The NACA 4412 was selected for superior lift-to-drag performance at low Reynolds numbers, outperforming alternatives like S1223 and S822. By fixing the airfoil and using segmented interpolation, the design workflow allows iterative optimization while ensuring the outputs are compatible with CAD modeling and realistic manufacturing constraints.

With the workflow broken down as mentioned above, the explained rationale behind specific design choices will follow the same format.

Optimizer:

The design process was centered on developing a numerical opimization framework to maximize power output within the project's constraints. We used a Python-based optimizer using differential evolution to iteratively refine blade geometry and operating parameters.

Engineering Model and Formulas

The model integrates aerodynamic performance calculation using a simplified blade element theory (BET) approach with structural analysis based on beam bending theory. The blade is discretized into 9 segments along 10 radial stations from the hub (r = 0.0254 m) to the tip (r = 0.1778 m, fixed blade length of 0.1524 m). Aerodynamic forces are computed per segment, summed to obtain total torque and power, and used to evaluate bending stresses. (Appendix A for in-depth equations)

Optimization Algorithm: Differential Evolution

We employed differential evolution (DE), a population-based global optimization algorithm suitable for non-linear, multi-variable problems like blade design. DE evolves a group of candidate solutions over generations to minimize the objective function, which focuses on maximizing power output while respecting constraints. The optimized variables are 12 in total: 10 pitch angles (θ_0 to θ_9, one at each radial station from hub to tip), the target RPM (RPM_target), and the root chord length (c_root).

How DE Works:

1. **Initialization**: Create a population of 144 random vectors of parameters generated within bounds like pitch (0°-45°), RPM (100-2000), and root chord (0.0254-0.0508 m)
2. **Mutation**: For each candidate, form a mutant by adjusting the best current solution with a scaled difference from two random others (scaling factor 0-2, adaptive in SciPy)
3. **Crossover**: Blend the mutant and original candidate component-by-component, randomly selecting from each based on a probability (0-1)
4. **Selection**: Compare the trial (from crossover) to the original using the objective function; replace if better, ensuring "survival of the fittest".
5. **Iteration**: Repeat steps 2-4 for up to 200 generations or until convergence (e.g., low population variance), gradually refining solutions within bounds.

We used SciPy's `differential_evolution` function, employing deferred updating for batched processing, a seed value of 42 to ensure reproducibility, and the `polish=True` option for a final local refinement step. The objective function assesses candidate designs by approximating the

expected power output through discrete sampling over wind speeds. In particular, it is formulated as

$$f(\theta_0, \theta_1, \ldots, \theta_9, \mathrm{RPM}_{\mathrm{target}}, c_{\mathrm{root}}) = -\sum_i P(U_i; \theta_0, \ldots, \theta_9, \mathrm{RPM}_{\mathrm{target}}, c_{\mathrm{root}}) \cdot f_{\mathrm{Weibull}}(U_i) \cdot \Delta U + \mathrm{penalties}$$

where $P(U_i)$ denotes the power at wind speed U_i, and ΔUrepresents the interval for numerical integration across sampled wind speeds from 1 to 20 m/s. This expression does not admit a closed-form polynomial representation; instead, it embodies a procedural computational framework that incorporates established aerodynamic models, such as lift and drag forces derived from NACA 4412 polar data augmented with Viterna-Corrigan extrapolation. By minimizing the negated weighted sum, the formulation effectively maximizes the average power output.

To optimize efficiently, only the root chord is a variable; the full chord distribution follows a quadratic taper:

$$c(r) = c_{\mathrm{tip}} + (c_{\mathrm{root}} - c_{\mathrm{tip}}) \left(\frac{R - r}{R - r_{\mathrm{hub}}} \right)^2$$

Where $c_{tip} = 0.5 \times c_{root}$, R is the maximum radius (hub radius + blade length), and r_{hub} is the hub radius.

Post-optimization, the torque brake is set to the computed Q at the optimal RPM and mean wind speed

Constraints and Penalties

Design respects geometric, structural, and operational limits; constraint violations are penalized to guide DE toward feasible solutions

Violations were penalized with a large value (1e6) added to the objective if stress or torque limits were exceeded at any windspeed, or the peak power RPM was outside the safe range. This soft-constraint approach via penalties guided the differential evolution algorithm toward feasible regions of the design space.

Post-Optimization Enhanced Analysis

Following optimization, the raw pitch angles often exhibited non-physical spikes (e.g., jumping from 5° to 10° to 2° to 7° across radial stations), which are impractical for

manufacturing. Post-optimization pitch angles were smoothed using a quadratic fit to ensure manufacturable, monotonic distributions

Results

The optimizer converged after approximately 1 hour yielding an expected power of 6.02 W under the Weibull distribution, under an optimal RPM of 1739 at mean wind speed (4.59 m/s), maximum stress was 0.944 MPa (<36.67 MPa limit), and torque per blade 0.011 N·m.

Physical Design:

Project constraints established clear bounds for the final blade. In CAD, a defining volume represented these limits, while parametric variables, the pitch and chord at each segment, were added to the Fusion library. This allowed rapid iteration once optimization results were available and enabled quick updates to the physical model as needed.

The hub connection was carefully shaped with lofts and fillets to match simulated bending stresses and maintain the factor of safety, avoiding stress concentrations. By blending the cross section to create a simple, manufacturable connection, the design preserved the performance predicted by the model while remaining practical to produce.

The largest difference between the CAD model and the simulation results comes from modifications to the airfoil. With the attachment piece added, even at a 2 in chord, the NACA 4412 is much thinner than the attachment. To address this, the trailing edge was truncated beyond the region contributing most lift, as shown in the XFoil graph (Figure 3), and each cross-section was scaled to match the chord, increasing blade thickness. This preserves the front half of the airfoil, which generates the majority of lift, while strengthening the connection to the attachment and reducing bending stress, maintaining both structural integrity and aerodynamic performance. To create a trailing edge, the truncated portion was rounded off for consistency.

The combination of all these design choices and iterations can all be seen in the results shown in Figure 4.

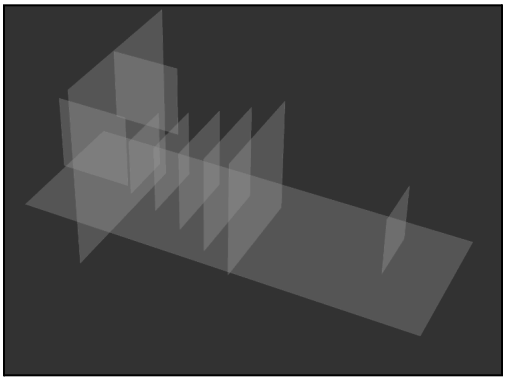

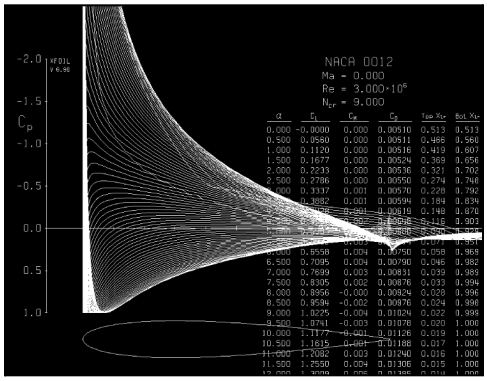

Figure 2 (Top Left) Figure 3 (Top Right):
A sample airfoil in XFoil showing the distribution of the Coefficient of pressure along the airfoil length. Note it spikes at the leading edge where the pressure difference and force component is the largest).

Figure 4 (Above)

Assessment of Design's Performance:

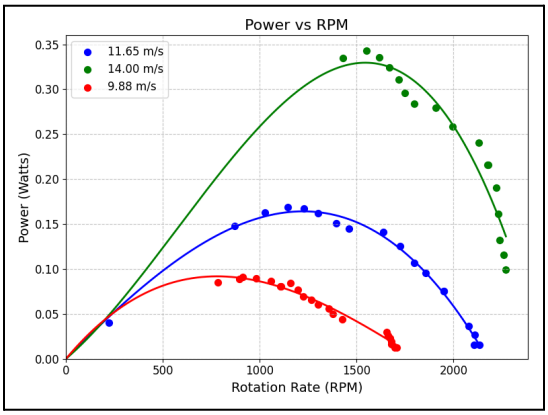

Power Curves for 3 Recorded Windspeeds - Figure 5

To better understand our blade's overall performance under intended operating conditions, we tested it at varying wind speeds, capturing data representative of different power outputs. We started at 9.8 m/s, the minimum speed at which the blades would spin (not at the ideal 4.59 m/s), then increased the fan frequency to two higher discrete speeds. At each speed, we collected data across a range of torque levels and RPMs until stall, yielding power curves. We compared the results to expected outcomes from our objectives and optimizer's design. The figures above illustrate the trends as wind speeds and RPMs increased.

Ultimately, while our testing yielded valuable data, it revealed that the blade design underperformed at lower than expected wind speeds, including our ideal of 4.59 m/s, where it failed to spin. Performance improved at higher speeds, which misaligns with our design objectives.

Results and Reflections

Windspeed(m/s)	Max Power (W)	RPM	Torque (N-cm)
9.8	0.091	914	0.095
11.55	0.169	1145	0.141
14.0	0.343	1551	0.212

Key Power Characteristic Results - Figure 6

At the design windspeed of 4.59 m/s, the blades did not spin. Testing began at 9.8 m/s, yielding a max power of 0.09 W, far below the expected 6 W. We repeated the same process for two distinct windspeeds, 11.55 m/s and 14 m/s. This yielded peak power at higher RPMs as seen above, giving us a maximum power of 0.343W, which was 6% of our expected max power. However, we did experience this peak power at 1551 RPM, which was within our expected and required max RPM range under 2000 RPM and close to 1739 RPM. Despite not fully meeting our objectives, we were still able to proceed with collecting enough data to characterize the design.

Our lack of performance at the lower end of wind speeds initially came to us as a mystery, given that the design was optimized and predicted to perform better at lower wind speeds. This was until we realized a major flaw in our design, that being the pitch of our blades being in the wrong direction, restricting flow and requiring a forceful accumulation of pressure on the wrong side for it to spin at all. This was unfortunately an error caused by miscommunication and a lack of understanding in how our blade should’ve come together, but we moved forward with collecting data with what we had. A note of improvement here would be to design the blades for the appropriate conditions and orientation.

As expected, our design did not fail under under use and survived the testing conditions we operated under. This can be attributed to the design being led by a large factor of safety of 58, which was much higher than the required 1.5. Our blades were designed to withstand high stress at the most fragile areas.

Optimizer Improvements for Future Prototypes:

Improvements can address the current differential evolution's probabilistic nature, hyperparameter sensitivity, error amplification, and computational cost through algorithmic refinements, scaling, and enhanced model fidelity.

Algorithmic Refinements

- **Increase Population Size and Iteration Limits**: Increase the population from 12 to 60–120 (5–10 per design variable) and raise iterations from 200 to ~1000 to enhance search space exploration and avoid local maxima.
- **Hyperparameter Optimization**:Tune mutation factor (F), crossover rate (CR), and penalty coefficients to reduce run-to-run variability, acknowledging DE's inherently probabilistic behavior.
- **Adopt Enhanced DE Variants**: Implement advanced DE algorithms, such as those incorporating multiple mutation strategies to enhance global search capabilities
- **Refine Penalty Handling in Weibull Integration**: Adjust Weibull penalties to avoid over-penalizing designs at extreme, unlikely wind speeds

Model and Data Improvements

- **Incorporate Experimental Data**: Use experimental data for Cl, Cd, and hub friction to refine inputs and improve model accuracy.
- **Account for Advanced Physics**: Include 3D aerodynamics and frictional losses to improve model realism and prevent unphysical optimizations.

Conclusion

This project underscores that effective engineering design depends not only on sound theory and optimization but also on clear team communication and coordination. While our aerodynamic model and differential evolution optimizer yielded a theoretically efficient and structurally robust blade, experimental results revealed how communication breakdowns and critical changes led to suboptimal performance. From a group dynamics standpoint, the process enabled us to divide and conquer tasks, producing a refined, high-quality model and blade. However, we learned the vital need for cross-checking to ensure smooth transitions across divided work, as evidenced by the pitch orientation error.

Although the blade fell short of its intended aerodynamic goals, it proved structurally sound, validating our conservative design approach, and positioned us for success by enabling clear identification of error sources and multiple improvement points. Sound optimization techniques and rigorous experimental procedures were key in pinpointing these areas precisely. Ultimately, the project reinforced that successful outcomes arise from aligning technical rigor with strong team processes, emphasizing the need to continually verify optimization results against physical intuition, CAD interpretation, and assembly realities to bridge analytical success and experimental performance.

APPENDIX:

Optimizer Assumptions:

- Aerodynamics: 2D airfoil data only (no 3D effects like tip losses, radial flow, or spanwise variation)
- Steady-state uniform inflow (no turbulence, gusts, or yaw)
- Inviscid flow except for viscous drag captured in Cd
- Reynolds number fixed at Re = 50,000 (low-speed approximation)
- No induction factors (simplified BET, not full BEM).
- Structural: Euler-Bernoulli beam theory for bending only (no shear, torsion, buckling, or fatigue);
- Isotropic material with constant properties; stress limited to 36.67 MPa (55 MPa ultimate strength / 1.5 safety factor).
- No mechanical losses (e.g., bearing friction, generator efficiency assumed 100%)

- Wind speeds follow idealized Weibull distribution

These assumptions prioritize computational efficiency for optimization while capturing essential physics, though they may overestimate performance compared to real-world experiments.

Optimizer Equations:
Aerodynamic Calculations:

For each segment at mid-radius r_mid, the local relative velocity W is:

$$W = \sqrt{U^2 + (\omega r_{\text{mid}})^2}$$

where U is the freestream wind speed, ω is the angular velocity (ω = RPM × 2π / 60).

The flow angle β and angle of attack α are:

$$\beta = \tan^{-1}\left(\frac{\omega r_{\text{mid}}}{U}\right), \quad \phi = 90^\circ - \beta, \quad \alpha = \phi - \theta_{\text{mid}}$$

where θ_mid is the local pitch angle.

Lift and drag coefficients (Cl, Cd) are interpolated from airfoil polar data (NACA 4412 at Re = 50,000). For angles outside the measured range (-9.25° to 14°), we applied Viterna-Corrigan extrapolation for post-stall (high α, mimicking flat-plate behavior with Cl → 0 at 90° and Cd increasing) and simple flat-plate theory for low α (Cl = sin(2α), Cd increasing with |α|).

Normal and tangential forces per segment are:

$$dF_n = \frac{1}{2}\rho W^2 c_{\text{mid}} dr \, (Cl \sin\beta + Cd \sin\phi)$$

$$dF_t = \frac{1}{2}\rho W^2 c_{\text{mid}} dr \, (Cl \cos\beta - Cd \cos\phi)$$

where ρ = 1.225 kg/m³ (air density), c_mid is the local chord, and dr is the segment length.

Total torque Q (per blade) is computed by summing the contributions from each blade segment:

$$Q = B \sum r_{\text{mid}} dF_t$$

with B = 1 for per-blade analysis (scaled to 3 blades in post-processing). Power is then:

$$P = Q\omega$$

During optimization, angular velocity ω (derived from RPM) is treated as an input variable, assuming a torque brake maintains constant rotational speed. For a given candidate RPM (thus ω), torque Q is calculated based on the aerodynamic forces at that fixed ω across various wind speeds U. This yields P(U) at the target ω. To evaluate operational feasibility, we also sweep a range of RPM values (100 to 2000) at the mean wind speed to compute Q and P vs. RPM, identifying the RPM where P peaks.

Expected power is integrated over the Weibull wind distribution (shape k = 5, scale c = 5 m/s):

$$P_{\text{expected}} = \int_1^{20} P(U) f(U) dU \approx \sum P(U_i) f(U_i) \Delta U$$

where f(U) = (k/c) (U/c)^{k-1} exp(-(U/c)^k), discretized over 20 wind speeds.

Structural Calculations:

Bending moment at station k is:

$$M(k) = \sum_{\text{seg}=k}^{N-1} dF_n(\text{seg})(r_{\text{mid}}(\text{seg}) - r_k)$$

The moment of inertia accounts for pitch rotation:

$$I_\theta = I_y^{\text{nd}} \cos^2\theta + I_z^{\text{nd}} \sin^2\theta$$

where I_y^nd and I_z^nd are non-dimensional second moments for NACA 4412 (computed numerically via trapezoidal integration over the airfoil profile). Scaled I = I_θ × chord^4.

Maximum stress is:

$$\sigma = \frac{M \cdot c}{I}, \quad c = c_{\max}^{\mathrm{nd}} \times \mathrm{chord}$$

with c_{max}^{nd} the non-dimensional extreme fiber distance.

Chord distribution follows a quadratic taper:

$c(r) = c_{tip} + (c_{root} - c_{tip})[(R - r)/(R - r_{hub})]^2$, with $c_{tip} = 0.5c_{root}$