



Multi-Objective Evolutionary Optimization of Wind Turbine Airfoils Incorporating Leading-Edge Roughness Insensitivity

Ryan Cameron¹ and Matthew Lackner¹

¹University of Massachusetts Amherst, Amherst, MA

Correspondence: Ryan Cameron (rtcameron@umass.edu)

Abstract. Wind turbine airfoil design has historically targeted three objectives: high lift coefficient (C_L), high lift-to-drag ratio (C_L/C_D), and insensitivity to leading-edge roughness (LER). The airfoils developed in the 1980s and 1990s for these objectives remain in widespread use today, yet the trade-offs among these competing goals have never been systematically mapped using modern global optimization methods. This paper develops a multi-objective evolutionary strategy (ES) to compute Pareto-optimal airfoil sets that reveal these trade-offs explicitly. The ES is initialized from a symmetric NACA airfoil rather than an existing wind turbine design, uses a Chebyshev-based CST parameterization, and is rigorously tuned via parameter studies. At the outer bounds of the Pareto front, the optimized airfoils improve upon the DU 93-210 reference by 87% in C_L , 26% in C_L/C_D , and reduce $\Delta\bar{C}_{L,LER}$ to near zero. These gains are driven largely by increased camber; when a practical camber constraint of 5% is applied, improvements of 33% in C_L and 16% in C_L/C_D are still realized with near-complete LER insensitivity. A key finding is that, counter to conventional design wisdom, aft-loading correlates with reduced LER insensitivity for the optimized airfoils: roughness forces boundary-layer transition near the leading edge, producing a thicker turbulent boundary layer that separates earlier in the aft recovery region, negating the expected benefit of aft-loaded lift distributions. The framework, including the tuned ES and optimized airfoil database, is made available for public use.

1 Introduction

Research on airfoils for wind turbine (WT) applications began in the 1980s with the aim of maximizing power production for horizontal-axis wind turbines by optimizing local two-dimensional aerodynamic forces. Researchers quickly converged on three requirements: high lift-to-drag ratio, C_L/C_D , sufficiently high lift coefficient, C_L , and insensitivity to leading-edge roughness (LER) (Timmer and Bak, 2023). The first two objectives follow directly from the physics: sectional torque and bending moment scale with $C_L c$ and C_L/C_D , where c is the chord length, and in blade element momentum theory increasing both metrics improves power capture efficiency. The need to consider LER arises from the operating environment. Wind turbines operate within the atmospheric boundary layer for 15 to 25 years, and over time the blades accumulate dirt, debris, and insects, and experience surface erosion. These effects degrade the leading-edge geometry and reduce aerodynamic performance (Ehrmann et al., 2013; Ramsay et al., 1995). The airfoils developed during this period, notably the DU, FFA, NREL S, and NACA 6-digit families, remain in widespread use in modern blade designs, despite significant changes in turbine scale, operating environment, and design requirements.



Despite the clear importance of these three objectives, the trade-offs among them have never been systematically mapped using modern global optimization methods. Early airfoil design for wind turbines relied on inverse methods, which begin with target aerodynamic properties such as lift slope and pressure coefficient (C_p) distributions at a specified angle of attack, α , and invert a two-dimensional aerodynamic solver such as XFOIL (Drela, 1989) to obtain a shape that achieves those targets. This approach is limited by the fidelity of the aerodynamic model and, critically, produces a single shape for a given set of targets rather than revealing trade-offs across competing objectives. The airfoils developed in this era also established a widely adopted design heuristic: aft-loaded airfoils are more insensitive to LER because lift generation is shifted away from the roughness-affected leading edge (Timmer and Bak, 2023). Whether this principle holds across the full range of feasible airfoil geometries has not been tested.

The objective of this research is to develop and evaluate a multi-objective global optimization framework that quantifies trade-offs among C_L/C_D , C_L , and LER insensitivity, subject to structural and manufacturing constraints. We compute Pareto fronts and analyze how airfoil geometry varies as these competing objectives are traded against one another. In addition to establishing a design tool, this analysis reveals that the relationship between aft-loading and LER insensitivity reverses at high camber levels, contradicting the conventional design heuristic.

We use an evolutionary strategy (ES) (Back, 1996) because it accommodates heterogeneous objectives and constraints and performs well in nonconvex, irregular design spaces. Early ES-based airfoil studies targeted specific aerodynamic characteristics, for example a desired C_p distribution, followed by an inverse process to find the shape (De Falco et al., 1996, 1999). More recently, ES variants have been applied to aerodynamic objectives with structural constraints to improve specific sectional airfoils (Hansen, 2018). A notable difference in our work is the initialization strategy: the search begins from a symmetric NACA airfoil at the required design thickness, rather than from an existing WT airfoil. This choice avoids biasing the optimizer toward a known local optimum and allows the ES to explore a broader region of the design space.

The remainder of the paper is organized as follows. Section 2 outlines the CST parameterization and details the ES implementation, objective and constraint formulations, hyperparameter tuning, and convergence behavior. Section 3 presents results for a baseline aerodynamic case and a three-objective case that adds LER insensitivity, including off-design sensitivity checks.

The optimized airfoil database and the tuned ES implementation are made available for public use.

2 Methods

2.1 Class Shape Transformation Background

Airfoil optimization requires a compact, well-conditioned set of design variables. Using the raw Cartesian coordinates of the surface as the design variables can introduce 200 to 400 parameters, which is computationally expensive and difficult to manage. A reduced-dimensional parameterization that captures the dominant shape features is therefore necessary.

Bezier curves and cubic B-splines are traditional parameterizations. They are robust but not tailored to airfoils, so most combinations of control points produce nonphysical shapes. For optimizers that rely on randomized search, this inefficiency can degrade performance and cause divergence. The Class Shape Transformation (CST) addresses this issue by imposing an



airfoil class function on any candidate solution (Kulfan and Bussioletti, 2012). The original CST used a Bernstein polynomial
 60 basis; later work found that orthogonal bases fit airfoil shapes more effectively (Doronina et al., 2025). In this study we use
 Chebyshev polynomials as the basis functions. With the class function layered on the basis, most candidate solutions map to
 physically plausible airfoil geometries.

Recently, there have been concerns that the CST is overly restrictive and limits the range of feasible perturbations to an airfoil
 shape. Data-driven parameterizations built from curated airfoil databases offer one way to regularize the design space while
 65 capturing perturbations across scales (Doronina et al., 2025). These approaches can be effective, especially for mutation-based
 searches. However, the simplicity and structure of the CST formulation make it well suited for enforcing design constraints.
 We return to this point in Section 2.3.2.

2.2 Class Shape Transformation Formulation

Using the CST reduces the dimensionality of the optimization from m discrete surface points that define an airfoil to n weights
 70 that define the parameterization. The CST expresses the normalized airfoil ordinate as a product of a class function, $C(x)$, and
 a shape function, $S(x)$:

$$\psi(x) = C(x) \cdot S(x) \tag{1}$$

$$\psi(x) = \frac{y}{c} \tag{2}$$

The class function in Equation 3 represents a basic airfoil shape. In this work we set $N_1 = 0.5$ and $N_2 = 1$. This basic shape
 75 is then perturbed by a weighted sum of basis functions. Note that all x -coordinates are normalized to the chord length, c , so
 that $x \in [0, 1]$.

The class and shape functions and the Chebyshev polynomials are (5):

$$C(x) = x^{N_1}(1-x)^{N_2} \tag{3}$$

$$80 \quad S(x) = \sum_{i=0}^n w_i T_i(x) \tag{4}$$

$$T_i = \begin{cases} T_0(x) & = 1 \\ T_1(x) & = x \\ T_n(x) & = 2xT_{n-1}(x) - T_{n-2}(x) \end{cases} \tag{5}$$

Expressing the fit as a system of linear equations simplifies the numerical implementation. Because the airfoil requires
 separate parameterizations for the upper and lower surfaces, the weights, \bar{w} , are arranged in an $n \times 2$ matrix, with the first



85 column for the upper surface and the second for the lower surface. The polynomial formulation can be written as a matrix
 product:

$$\bar{\psi} = \bar{C}\bar{T} \cdot \bar{w} \quad (6)$$

If the $m \times n$ matrix $\bar{C}\bar{T}$ is precomputed using a consistent vector of x-coordinates, x^* , then $\bar{\psi}$ depends only on \bar{w} . The
 resulting overdetermined system is solved using linear least squares, or any standard method for approximating overdetermined
 90 systems.

2.3 Evolutionary Strategy

Evolutionary strategies aim to mimic Darwin's theory of evolution by using selection, recombination, and mutation to search
 complex design spaces. Our implementation follows a (N_p, N) strategy that runs for $N_G = 250$ generations, starting from an
 initial population I_0 . At each generation we evaluate fitness, select N_p parents using non-dominated sorting (Deb et al., 2002),
 95 apply crossover and mutation to produce N offspring, evaluate the new population, and repeat until a specific stopping criterion
 is met.

The implemented ES is described by Algorithm 1 and Equations 7 - 12. Here, s is the selection operator, which uses the
 non-dominated sorting algorithm to choose N_p parents. The genetic operators ω_{Θ_1} and ω_{Θ_2} are the crossover and mutation
 functions, respectively.

$$100 \quad \Phi = f(\bar{w}) : \bar{w} \rightarrow \mathbb{R}^k \quad (7)$$

$$s : I^N \rightarrow I^{N_p} \quad (8)$$

$$\omega_{\Theta_i} = \begin{cases} \omega_{\Theta_1} : I^{N_p} \rightarrow I^N \\ \omega_{\Theta_2} : I^N \rightarrow I^N \end{cases} \quad (9)$$

$$\Psi = I^N \rightarrow I^N \quad (10)$$

$$\Psi(P(t)) = \omega_{\Theta_2}(\omega_{\Theta_1}(s(P(t)), \sigma_m, \lambda_m)) \quad \forall 0 < t < 250 \quad (11)$$

$$105 \quad P(t+1) = \Psi(P(t)) \quad (12)$$

The population, $I_0 = P(0)$ initializes the optimizer, so it needs to be carefully chosen in accordance with the goals of the
 process. Here, the ES is seeded with a symmetric NACA airfoil of a specified design thickness, τ_{design} . We select a general
 symmetric airfoil in order to assess what geometric characteristics differentiate a WT airfoil from traditional designs, and to
 ensure that we do not initialize the problem near a local optimum. Each member of the initial population, $\bar{a}_{i,0}$, is generated by
 110 uniformly randomizing the weights of the seed airfoil within $\pm\sigma_m \cdot \bar{w}$. More discussion on the hyperparameter σ_m is in Section
 2.3.3.



Algorithm 1 Basic structure and process of our ES

```

t ← 0
I0N ← [w̄*, ..., w̄*]
P0 = ωΘ2(I0N, σm, λm = 1.0) {λm = 1.0 signifies every individual is mutated}
Evaluate Φ(P0)
while t ≤ NG do
    P(t + 1) = Ψ(t)
    Evaluate Φ(P(t + 1))
    Update Pareto front
end while
    
```

Variable	Value
Re	1.5×10^6
N_{crit}	9 (clean), 1 (rough)
max iterations	1000
α_{design}	7

Table 1. XFOIL parameters

The fitness function in the ES maps the CST weight matrix to k objectives, $\Phi : \bar{w} \rightarrow \mathbb{R}^k$. The objectives included in Φ are defined in Section 2.3.1. For this work, $\Phi = f(\bar{w})$ is evaluated with XFOIL under the conditions listed in Table 1. We use a viscous boundary-layer model with 160 surface panels distributed by half-cosine spacing. If a single XFOIL call exceeds a 20 second timeout, Φ returns zero for all objectives, which discourages numerically unstable shapes. We use $Re = 1.5 \times 10^6$ because XFOIL's viscous-inviscid coupling becomes unreliable at higher Reynolds numbers. Extending these results to $Re = 15 - 20 \times 10^6$ representative of modern offshore turbines requires higher-fidelity solvers and is the subject of ongoing work. Since XFOIL is also unreliable in stall and post-stall regimes, the design angle of attack, $\alpha_{design} = 7^\circ$ is used to ensure that the majority of airfoils created will be evaluated within the pre-stall region. The sensitivity of the ES to this angle is assessed in Section 3.5.

Parent selection is based on non-dominated sorting (Deb et al., 2002). If the current non-dominated set contains more than N_p members, we randomly sample N_p parents for the next generation. Offspring are created by a weighted average of two parents, with weights drawn from a uniform distribution and re-sampled for each new individual. After crossover, each offspring undergoes mutation with probability λ_m . When selected, mutation perturbs the CST weights using the same uniform scheme and scale σ_m as in the initialization step. The process repeats until the stopping criterion of N_G generations is met. This procedure yields a Pareto-optimal set of airfoils for the specified objectives and constraints.



2.3.1 Design Objectives

We adopt a multi-objective formulation and compute the Pareto front rather than collapsing objectives into a single weighted score. This approach reveals how changes in airfoil geometry map to trade-offs in performance. We evaluate two sets of objectives:

- Baseline aerodynamic case: maximize C_L and maximize $\frac{C_L}{C_D}$.
- LER-insensitive case: retain the two aerodynamic objectives and add an insensitivity objective that minimizes the relative lift drop between clean and rough surfaces, shown in Equation 13.

$$\Delta \bar{C}_{L,LER} = \frac{C_{L,clean} - C_{L,rough}}{C_{L,clean}} \quad (13)$$

Minimizing $\Delta \bar{C}_{L,LER}$ promotes shapes that maintain lift under leading-edge degradation, while the aerodynamic objectives drive efficiency and lift at the design condition. The resulting Pareto fronts allow us to quantify trade-offs and examine the associated geometric trends. The rough condition is modeled in XFOIL by reducing N_{crit} from 9 to 1 (Hansen, 2018), which describes the turbulence present in the flow field using an e^n model. Reducing this value near zero promotes higher levels of turbulence in the freestream flow, which forces the boundary layer to transition from laminar to turbulent further towards the leading edge of the airfoil. Physically, reducing N_{crit} raises the freestream turbulence level rather than modeling surface roughness geometry directly; however, the dominant aerodynamic effect of leading-edge roughness is premature boundary-layer transition, which the N_{crit} reduction captures (Ehrmann et al., 2013). Hansen (2018) used the same approach for airfoil optimization and found results consistent with higher-fidelity roughness models.

2.3.2 Design Constraints

We enforce two constraints that reflect structural and manufacturing considerations: a global maximum design thickness and minimum thickness requirements at both 70 percent and 80 percent chord.

The global maximum design thickness is enforced directly in the CST framework. The upper and lower surfaces are defined as $\bar{\psi}_{top}$ and $\bar{\psi}_{bot}$. The thickness profile, $\bar{\tau}$, is then defined in Equation 17 as the difference between the two surfaces, which can be written using the same CST matrix and a thickness weight vector \bar{l} .

$$\bar{\tau} = \bar{\psi}_{top} - \bar{\psi}_{bot} \quad (14)$$

$$= \bar{C}\bar{T} \cdot \bar{w}_{top} - \bar{C}\bar{T} \cdot \bar{w}_{bot} \quad (15)$$

$$= \bar{C}\bar{T} \cdot (\bar{w}_{top} - \bar{w}_{bot}) \quad (16)$$

$$= \bar{C}\bar{T} \cdot \bar{l} \quad (17)$$



The design thickness is enforced by simply finding the current maximum value of $\bar{\tau}$, and then calculating a scaling parameter ϵ that multiplies the entire profile to meet that design thickness. Because the system is linear, ϵ can be directly applied to the CST weights, \bar{w} as well.

The second constraint requires minimum thickness at selected chordwise positions to preserve structural capacity and manufacturability. Specifically, we enforce $\bar{\tau}|_{[0.7,0.8]} \geq [0.4, 0.2]\tau_{design}$, where τ_{design} is the specified design thickness. These are inequality constraints, so we implement them numerically within the ES. After mutation, each individual is screened for feasibility. Feasible individuals proceed unchanged. If an individual violates the constraint, it is discarded and a new mutant is generated to replace it. This step is repeated until the entire population satisfies the constraints.

Often, the maximum camber of the airfoil is constrained for manufacturing and structural purposes as well (Karcher et al., 2025), but we have decided not to enforce this in the optimization process in order to observe the aerodynamic design tradeoffs with as few constraints as possible. In general, increased camber raises both $C_{L,max}$ and drag, but C_L/C_D at the design angle still improves, so we expect the ES to favor high-camber shapes. Rather than constraining camber during optimization, we allow it to vary freely and filter the results in post-processing (Section 3.4).

2.3.3 Hyperparameters

The basic structure of the ES is given in Algorithm 1, but performance and accuracy are determined by a set of tunable hyperparameters. For this study we identified four key parameters (Table 2). Each was assigned five levels, and the resulting design space was sampled using an Orthogonal Fractional Factorial (OFF) (Mills et al., 2015) methodology. To keep computation tractable, we filtered the space to 70 total combinations for testing.

We assessed sensitivity by examining variation in mean fitness for each parameter, following the procedure in (Mills et al., 2015). The analysis uses absolute effects, E_i , to compare parameters and to indicate where fine tuning is warranted. The optimal value for each parameter is then selected based on the highest average fitness.

We define the set of test cases S_{ij} in Equation 18 for parameter x_i and level l_j . For each case k in that set, y_{ijk} is the average fitness across all generations.

$$S_{ij} = \{x \in S |_{x_i=l_j}\} \quad (18)$$

From those values, a single average fitness for parameter x_i at level l_j is found, \bar{y}_{ij} in Equation 19.

$$\bar{y}_{ij} = \frac{\sum_{k=1}^N y_{ijk}}{N} \quad (19)$$

The absolute effect of parameter x_i is the range of \bar{y}_{ij} across levels, E_i , as defined in Equation 20.

$$E_i = \max\{\bar{y}_{ij}\} - \min\{\bar{y}_{ij}\} \quad (20)$$



Name	Description	Levels	Optimal
N	Population size	[20,40,60,80,100]	80
N_p	Number of parents	[0.05,0.07,0.1,0.15,0.2]* N	8
λ_m	Mutation rate	[20,22.5,25,27.5,30]%	27.5%
σ_m	Mutation variability	[20,25,27.5,30,32.5,35]%	25%

Table 2. Hyperparameters and the most fine grained levels they were tested at for tuning the algorithm. λ_m and σ_m were refined from a more general set of levels.

Parameters with large E_i are candidates for refined searches with adjusted level ranges. After any refinement, the analysis is repeated and optimal levels are chosen as those with the highest \bar{y}_{ij} .

185 Table 2 reports the final hyperparameters along with the initial tested levels. For each parameter we computed E_i for the two strictly aerodynamic objectives, C_L and $\frac{C_L}{C_D}$; the effects appear in Figure A1. In the initial study, the mutation scale σ_m was by far the most sensitive factor for both objectives, followed by population size. We also found the optimal mutation probability λ_m at the boundary of its tested range, which motivated a second, refined search. In the refined study we expanded the range for λ_m and tightened the range for σ_m . The refined results showed that λ_m was most influential for the max C_L objective, while σ_m dominated for $\frac{C_L}{C_D}$.

190 Although the initial study suggested the same optimal levels across objectives, the refined analysis yielded two different optima for σ_m : 32.5 percent for max C_L and 25.0 percent for max $\frac{C_L}{C_D}$. Because σ_m had greater influence on $\frac{C_L}{C_D}$, we selected $\sigma_m = 25$ percent as the final setting.

2.3.4 Convergence

The final hyperparameter is the stopping condition, ι . Although stopping can be defined by a monitored statistic, we use a fixed 195 limit and set $N_G = 250$. Figure 1 shows the evolution of the upper quantile of fitness across generations for each objective, indicating that 250 generations are sufficient for convergence. C_L and $\frac{C_L}{C_D}$ for the DU 93-210 airfoil (a 21% thick WT airfoil) are plotted as horizontal dashed lines for comparison.

For hyperparameter tuning we held the random seed fixed to enable fair comparisons of absolute effects. For convergence assessment, we varied the seed across $N_T = 500$ independent realizations, all initialized from the same base airfoil.

200 Because each realization yields a different Pareto front, we summarize performance with scalar metrics and with statistics on the fronts themselves. First, we extract the endpoints of each run's Pareto front: maximum C_L and maximum C_L/C_D . We bin these across realizations, and plot their distributions in Figure 2. The narrow spreads indicate that the outer envelope of the Pareto set is well bounded.

205 Finally, we estimate how many independent ES runs are needed to approach the global optimum for the defined objectives and constraints. We form a global front by taking the union of all points from the 500 realizations and compute the Pareto front of that union. For subsets of size n drawn at random from the 500 runs, we compute a semi-global Pareto front using the union




Figure 1. Upper quantile fitness across the entire population for each generation. Evaluated for C_L and $\frac{C_L}{C_D}$ objectives.




Figure 2. Outer envelope distributions across all independent realizations of the optimization algorithm.

of all points in the subset. A Euclidean distance from the origin is calculated for each member of the semi-global front, and we quantify the optimality of the subset through the root mean square (RMS) of the distances. To compare across objectives, each coordinate is normalized by the global maxima so that the reference front lies in $[0, 1]^2$. For each subset size, we repeat the random draw 1000 times to obtain a mean RMS and its standard error. Figure 3 shows the RMS distance versus the number of realizations combined, which provides practical guidance on how many ES runs are needed to achieve a desired approximation accuracy to the global Pareto front. The data are also given in Table 3, showing that after 10 iterations of the ES, minimal gains are made compared to the extra compute time.

3 Results

215 This section reports results from the ES optimization for two cases: a baseline with purely aerodynamic objectives (C_L and C_L/C_D), and a WT case that adds leading-edge roughness insensitivity as a third objective. These results are further narrowed



Iterations	Mean [%]	Err [%]
1	93.060	0.296
10	98.982	0.005
50	99.193	0.004
100	99.301	0.001
500	100.000	0.000

Table 3. Standard error for the number of independent runs of the ES used for a single problem.




Figure 3. A semi-global Pareto front is calculated using number of iterations of the ES on the x-axis, and compared to the global Pareto front (500 iterations).

down by setting a maximum camber filter in order analyze camber's effects on performance. The section concludes with an off-design check at $\alpha_{design} = 6^\circ$ to verify that the optimized airfoils are not sensitive to the specific design angle of attack.

Because the ES explores a broad design space initialized from a symmetric airfoil, many individuals achieve high fitness
 220 in one objective but are impractical as wind turbine sections. The following filters remove such individuals: a minimum lift requirement ensures adequate torque contribution at the design angle of attack, the lift-to-drag bounds exclude both low-efficiency shapes and numerically suspect outliers, and the upper limit on $\Delta\bar{C}_{L,LER}$ discards airfoils whose performance collapses entirely under leading-edge roughness.

- $C_L(\bar{a}_i, \alpha_{design}) \geq 0.6$
- 225 - $60 \leq \frac{C_L}{C_D}(\bar{a}_i, \alpha_{design}) \leq 200$
- $\Delta\bar{C}_{L,LER}(\bar{a}_i, \alpha_{design}) \leq 1$

Each population is an aggregation of $N = 10$ independent realizations of the ES, producing $N_{\bar{a}} = 201,600$ unique individuals across all 250 generations. The filters above are applied to every individual regardless of the generation in which it was produced.




Figure 4. Shapes of the reference WT airfoils.

Name	C_L	C_L/C_D	$\Delta \bar{C}_{L,LER}$
DU 93-210	1.2740	137.5070	13.17%
NREL S809	0.9000	70.6542	3.34 %
NACA 63(3)-421	1.1760	132.3137	12.56%
FFA-W3-211	1.1773	126.9479	6.13 %

Table 4. ES fitness evaluations of the reference WT airfoils.

230 3.1 Reference WT Airfoils

Our results are compared with existing WT airfoils of the same 21% design thickness: the DU 93-210 (Timmer and van Rooij, 2003), NREL S809 (Tangler and Somers, 1995), FFA W3-211 (Björck, 1990), and NACA 63(3)-421 (Abbott and Von Doenhoff, 2012). Except for the NACA 6-digit airfoil, each was designed specifically for WT applications with LER insensitivity as a design consideration. Table 4 lists the performance of each airfoil in the objective space used by the ES, and Figure 4 shows their shapes.

The data here show the main tradeoff in WT airfoil design: the DU and NACA six digit airfoils both prioritize traditional aerodynamic performance at the expense of LER insensitivity, while the NREL and FFA airfoils have greater insensitivity to LER but worse baseline aerodynamic performance. Shapes for these airfoils are shown in Figure 4.

The data illustrate the central tradeoff in WT airfoil design: the DU and NACA 6-digit airfoils prioritize aerodynamic performance at the expense of LER sensitivity, while the NREL and FFA airfoils achieve better roughness insensitivity at the cost of lower baseline lift and efficiency.

3.2 Aerodynamic Optimization (Two Objectives)

Each ES run differs only through the random seed. Figure 5 shows results for the baseline aerodynamic case, where the objectives are to maximize C_L and C_L/C_D . This case establishes a reference against which we compare optimizations that also include LER insensitivity.

The baseline optimization shows that for pure performance metrics, there is substantial improvement over both the initial airfoil and a comparable, dedicated WT airfoil like the DU 93-210. Of course, the DU airfoil was not designed just for purely aerodynamic conditions, as it also must limit LER effects. Comparing with the DU airfoil shows that the majority of improvement in performance is in C_L , which increases by over 88% increase in some cases. For the same value of C_L/C_D as the DU airfoil, there is an increase of 80% in design lift from 1.27 to 2.29.

The baseline Pareto front shows substantial improvement over both the NACA 0021 initialization and the DU 93-210, a dedicated WT airfoil whose design also accounts for LER insensitivity. The largest gains are in lift: for the same C_L/C_D as




Figure 5. Solutions for a two objective aerodynamic optimization, maximizing C_L and C_L/C_D . The right panel shows lift polars for each shape that is displayed in top left.

the DU airfoil, design lift increases by 80%, from 1.27 to 2.29. At the extreme of the front, C_L improves by over 88% relative to the DU airfoil (Table 5).

255 Figure 5 also highlights the extreme camber of the Pareto-optimal airfoils. Because the fitness evaluation operates at a single design angle of attack, the ES increases lift most effectively by increasing camber. A consequence is a sharp drop-off in lift during stall and post-stall conditions, visible in the polars compared to the DU airfoil (though XFOIL's reliability in post-stall is limited (Van Rooij, 1996)). For non-pitch-regulated turbines this is a practical concern, as turbulence, structural deflections, and other perturbations can push the airfoil in and out of the stall region.

260 It should be noted that forces are not calculated in XFOIL over the full range of α during the fitness evaluation to minimize computation time. One of the properties of the ES is quick convergence on a general optimal region, and then very slow convergence to the final Pareto front. Due to this slow final convergence, it is desirable to keep fitness evaluation computation time to a minimum.

We also compute the aerodynamic center (AC), defined in Equation 21, where $C_{m,LE}$ is the moment coefficient about the leading edge. The AC indicates the chordwise location at which the pitching moment is zero, and is a measure of where the lift is distributed along the chord. The farther aft the AC, the more lift is generated in the aft portion of the airfoil.




Figure 6. The Pareto front solution is sampled from max C_L to min C_L , and the chosen airfoils are evenly spaced along the front.

$$x_{AC} = -\frac{C_{mLE}}{C_L} \quad (21)$$

Figure 6 shows shapes and C_p distributions for airfoils sampled at equal intervals along the front, from the minimum- C_L (maximum- C_L/C_D) end to the maximum- C_L (minimum- C_L/C_D) end. Apart from the lowest-lift airfoil, every solution
 270 drives the trailing-edge thickness to the structural minimum, confirming the expected tradeoff between structural margin and aerodynamic performance. The progression of x_{AC} along the front shows no correlation between aft-loading and position on the Pareto front, which is expected for an optimization that does not account for LER effects. All optimized airfoils have x_{AC} values farther aft than the DU airfoil ($x_{AC} = 0.2$).

3.3 LER Insensitivity Optimization (Three Objectives)

275 The next case adds LER insensitivity as a third objective by minimizing $\Delta\bar{C}_{L,LER}$ (Equation 13). Because visualizing a three-dimensional Pareto front directly is difficult, we isolate unique two objective tradeoffs in post-processing: for each pair of objectives, we project the full 3D front onto the corresponding plane and recompute the Pareto front using all individuals. Figure 7 shows the resulting projections, with the full 3D front in orange. Figure 8 shows sampled shapes and lift polars from each isolated front.

280 The C_L vs. $\Delta\bar{C}_{L,LER}$ front lies close to the two objective front in the same plane, indicating that no airfoils are uniquely optimized for C_L/C_D alone, and instead airfoils with poor lift are also poor in efficiency. In the left panel of Figure 7, the




Figure 7. Pareto fronts for each unique two objective combination. Full three objective Pareto front is still shown in orange.




Figure 8. Solution shapes and their corresponding lift polars are shown for each unique two objective Pareto front: top) C_L vs C_L/C_D , middle) C_L vs $\Delta\bar{C}_{L,LER}$, bottom) C_L/C_D vs $\Delta\bar{C}_{L,LER}$.



front peaks near $C_L \approx 1.5$; below that value, the optimizer cannot reduce drag enough to reach the front, so all orange points below that threshold are driven entirely by LER performance. This tight grouping suggests a floor on achievable drag for these shapes, whether set by the fluid physics or by XFOIL's solver limitations.

285 Figure 7 also compares against traditional WT airfoils of similar thickness: the DU 93-210, NACA 63-421, FFA-W3-211, and NREL S809. The optimized airfoils achieve substantially higher C_L and C_L/C_D at comparable $\Delta\bar{C}_{L,LER}$. Compared to the Pareto front population as a whole, however, the traditional airfoils exhibit high LER insensitivity, reinforcing it as a priority in traditional design approaches.

To examine how geometric and aerodynamic properties vary, we sample airfoils at evenly spaced positions along each front. 290 Each objective is normalized to $[0, 1]$ and then the polar angle of each member's objective vector from the origin is measured. Percentages in the legends of Figures 6, 8, and 9 refer to quantiles of this angle.

Although the optimizer evaluates each airfoil at a single α , the full pre-stall and stall behavior is important to assess. The top and middle row of Figure 8 show lift polars on the C_L vs. C_L/C_D and C_L vs. $\Delta\bar{C}_{L,LER}$ fronts where the airfoils are highly cambered. Gaps near stall indicate unconverged XFOIL cases, and the minimal post-stall decline in lift, while aerodynamically 295 attractive, is likely a solver artifact rather than physical behavior. Capturing fully separated flows accurately would require higher-fidelity CFD tools.

In contrast, the bottom row of Figure 8 shows airfoils along the C_L/C_D vs. $\Delta\bar{C}_{L,LER}$ front, where camber is less extreme and the lift polars are more conventional. Even though C_L is not explicitly optimized on this front, the three highest-lift airfoils still exceed the DU airfoil in design lift. At the maximum C_L/C_D , the top two airfoils are nearly identical in pre-stall behavior, 300 differing only in stall angle. Examining their shapes, the 75% airfoil has a lower maximum height, thinner trailing section, and more rounded leading edge. Such small geometric differences producing a measurable shift in stall angle demonstrates that the optimizer resolves sufficient detail in the design space to isolate key performance sensitivities.

Figure 9 shows the aerodynamic center position and pressure distributions for the sampled airfoils along each isolated Pareto front, evaluated under both clean and rough conditions. Traditionally, WT airfoils were designed to be aft-loaded on the premise 305 that shifting lift generation away from the roughness-affected leading edge would preserve performance under LER (Timmer and van Rooij, 2003). The optimized airfoils show the opposite trend: as C_L (and $\Delta\bar{C}_{L,LER}$) increases along the front, the aerodynamic center moves aft while LER performance deteriorates.

The mechanism is visible in the C_p distributions. In the clean case, natural transition occurs well downstream of the leading edge, so the boundary layer experiences a shorter turbulent development length and reaches the aft adverse pressure gradient 310 with smaller momentum and displacement thickness. Under roughness, transition is forced near the leading edge, and the turbulent boundary layer grows over nearly the entire chord. By the time it reaches the aft recovery region, the boundary layer is thicker (larger momentum and displacement thickness) with reduced margin to withstand the strong adverse pressure gradient imposed by the high-camber geometry. The result is notably earlier separation: the C_p distribution flattens in the aft region, eliminating the lift contribution that aft-loading was intended to provide. Aft-loading therefore amplifies, rather than 315 mitigates, the penalty of LER on these high-camber shapes.




Figure 9. C_p distributions and aerodynamic center of pressure for three objective optimization. The legends in C_p plots are ordered by increasing: left) C_L , middle) C_L , and right) $\frac{C_L}{C_D}$. Also corresponds to decreasing: left) $\frac{C_L}{C_D}$, and increasing middle) $\Delta \bar{C}_{L,LER}$ and right) $\Delta \bar{C}_{L,LER}$.




Figure 10. Boundary layer thickness, Θ , and local displacement thickness, δ^* along the chord for the maximum lift and minimum roughness insensitivity airfoils.

Figure 10 provides direct evidence of this mechanism, comparing the boundary-layer momentum thickness, Θ , and displacement thickness, δ^* , over the upper surface for two airfoils: the maximum- C_L shape and the minimum- $\Delta\bar{C}_{L,LER}$ shape. On the high-lift, high-camber airfoil (blue lines), the rough boundary layer (dashed line) starts thicker and grows more rapidly than in the clean condition (solid line), as seen by the larger Θ . Separation is identified by the rapid increase in δ^* and the coincident discontinuity in Θ , occurring at approximately 85% and 40% of the chord for the clean and rough cases, respectively. The large difference in separation location confirms that roughness causes separated flow over nearly half the chord on this airfoil. On the roughness-optimized airfoil, however, the clean and rough boundary layers grow at similar rates along the full chord, with no abrupt separation. This contrasting behavior illustrates why low-camber shapes are insensitive to roughness: the milder adverse pressure gradient does not amplify the difference between the clean and rough boundary layers.

An apparent exception reinforces this interpretation. On the C_L vs. C_L/C_D isolated front (Figure 8), the minimum- C_L airfoil has the lowest camber and a correspondingly flat C_p distribution along the entire chord (Figure 9). Although roughness still causes earlier separation on this airfoil, the pressure difference between the attached and separated regions is small, so the lift lost to separation is also small. This pushes x_{AC} farther aft than the next airfoil up the front, despite better LER performance, confirming that separation extent, not aerodynamic center position alone, governs roughness sensitivity.

It is important to reconcile these results with the established design principle that aft-loading improves LER insensitivity (e.g. Timmer and van Rooij (2003)). Traditional WT airfoils such as the DU and FFA families have maximum camber ratios




Figure 11. Pareto front combinations after results are filtered for maximum camber $\leq 5\%$ of the chord.

below 3%, where the adverse pressure gradient in the aft region is mild enough that even a roughness-thickened boundary layer remains attached over most of the chord. In this low-camber regime, aft-loading does shift lift generation away from the leading edge, and the conventional wisdom holds. The optimized airfoils in this study occupy a higher-camber regime (5–18%) where the adverse pressure gradient is too severe for the boundary layer to remain attached. The benefit of aft-loading is therefore contingent on the boundary layer remaining attached through the aft region, a condition that breaks down as camber increases. Aft-loading and LER insensitivity are thus not universally correlated; the sign and magnitude of the relationship depends on the camber level, and by extension on the Reynolds number and roughness severity that determine how far downstream the boundary layer can remain attached.

3.4 Role of Camber in Aft-Loading and Roughness Sensitivity

Wind turbine blades may be constrained in the amount of camber that is feasible due to manufacturing and structural considerations. To assess performance within practical limits, the Pareto front from Figure 7 is filtered for airfoils with a maximum camber ratio of 5% (for comparison, the DU 93-210 has a maximum camber of 2.8%). The filtered solutions are shown in Figures 11 and 12.

Figure 11 shows that the camber filter substantially reduces the maximum achievable lift, from $C_L = 2.38$ to 1.70, which still represents a 33% improvement over the DU 93-210. On the C_L vs. C_L/C_D front, only two non-dominated solutions remain, indicating that the eliminated airfoils are predominantly high-lift, high-camber shapes. The $\Delta\bar{C}_{L,LER}$ vs. C_L front confirms this, as the filtered airfoils are now much closer in roughness performance to the traditional WT airfoils. The $\Delta\bar{C}_{L,LER}$ vs. C_L/C_D front (right panel) is largely unchanged, because constraining camber primarily eliminated the airfoils that perform poorly in LER sensitivity.

To compare directly with an existing design, we select optimized airfoils whose $\Delta\bar{C}_{L,LER}$ falls within 5% of the FFA-W3-211, which occupies a middle ground between LER insensitivity and lift. Figure 13 shows the matched airfoils, ordered by camber quantile.




Figure 12. Solution shapes and their corresponding lift polars after results are filtered for maximum camber $\leq 5\%$: top) C_L vs C_L/C_D , middle) C_L vs $\Delta\bar{C}_{L,LER}$, bottom) C_L/C_D vs $\Delta\bar{C}_{L,LER}$.




Figure 13. Shapes and lift polars of the LER matched airfoils from the full Pareto solution.



As expected, the highest-camber airfoil (12%) achieves the best lift but might face manufacturing and structural limitations. More broadly, this result indicates that high camber does not necessarily correlate with worse LER performance. Camber drops sharply to 5% at the upper quartile and decreases steadily to 4.1% at the lowest. These lower-camber airfoils compensate with thinner trailing edges and greater upper-surface curvature, both of which boost C_L and C_L/C_D . Taken together, among airfoils with equivalent LER insensitivity, higher camber generally yields better aerodynamic performance but at the cost of structural and manufacturing feasibility. The ordering is not perfectly monotonic across all five matched airfoils, so camber alone does not fully determine the tradeoff.

Section 3.3 showed that the high camber airfoils favored by the optimizer contradict the established design principle that aft-loading improves LER performance. To test this directly, we filter the original three-objective population for successively lower maximum camber values, from 10% down to 2%, recompute the Pareto front at each level, and isolate the two-objective fronts as before. Figure 14 shows the C_p distributions for the 5%-filtered set and the trend in x_{AC} vs. $\Delta\bar{C}_{L,LER}$ as the maximum camber is reduced. Trend lines are computed by first-order linear regression.

As the maximum allowed camber is reduced toward values typical of existing WT airfoils, the negative correlation between aft-loading and LER performance weakens but does not reverse (Figure 14). Even at a 2% camber limit, the trend line slope is essentially zero, meaning that no camber threshold is found at which aft-loading begins to improve LER insensitivity. This persistence is partly an artifact of the roughness model: at $N_{crit} = 1$, boundary-layer transition is forced very close to the leading edge on every airfoil regardless of camber, so even low-camber shapes experience a longer turbulent run than they would under clean conditions. At higher N_{crit} values representative of less severe roughness, transition on low-camber airfoils would occur farther downstream and the aft-loading benefit may be recovered. The C_p distributions in Figure 14 support this interpretation: as the camber filter is tightened, the difference between clean and rough C_p in the aft region diminishes and separation moves farther downstream, reducing but not eliminating the LER penalty associated with aft-loading.

3.5 Off-Design Scenario

To verify that the optimized airfoils are not overly sensitive to the chosen design angle of attack, we repeat the three-objective optimization at $\alpha_{design} = 6^\circ$ and evaluate the resulting Pareto front at $\alpha = 7^\circ$ for direct comparison with the baseline case.

Figure 15 overlays the off-design Pareto front on the baseline for all three isolated projections. The two fronts are closely aligned in both extent and density. As expected, the off-design airfoils do not exceed the baseline front, but they remain close enough to confirm that the Pareto-optimal set is not overly sensitive to the chosen design angle of attack.

4 Conclusions and Future Work

This paper developed a multi-objective evolutionary optimization framework for wind turbine airfoil design that systematically maps the trade-offs among lift, lift-to-drag ratio, and leading-edge roughness insensitivity. The framework integrates a Chebyshev-based CST parameterization with a rigorously tuned ES and produces Pareto-optimal airfoil sets rather than single-point designs. Notably, the ES is initialized from a symmetric NACA 0021 airfoil rather than an existing WT design, which




Figure 14. Top plot shows the AC position vs LER performance, for different maximum camber limits. Middle and bottom rows show difference in C_p distributions at the design angle of attack, for clean and rough calculations. C_p distributions are taken from a filtered dataset for max camber of 5%.




Figure 15. Pareto front for an off design optimization ($\alpha_{design} = 6^\circ$) overlaid on top of the optimal points in Figure 7, all points evaluated at $\alpha = 7^\circ$.

Case		Optimal Objective	C_L	C_L/C_D	$\Delta\bar{C}_{L,LER}$	Max camber	TE Thickness
NACA 0021		-	0.6558	72.1973	4.27%	0	0.0914
	DU	-	1.2740	137.5070	13.17%	0.0283	0.0583
2 Objective	value	C_L	2.4050	68.7942	-	0.1713	0.0234
		C_L/C_D	1.8359	176.2444	-	0.0844	0.0315
	% change	C_L	88.77%	-49.97%	-	505.30%	-59.86%
		C_L/C_D	44.10%	28.17%	-	198.23%	-45.97%
3 Objective	value	C_L	2.3807	61.4422	40.56%	0.1832	0.0376
		C_L/C_D	1.7575	172.6112	25.81%	0.0761	0.0492
		$\Delta\bar{C}_{L,LER}$	0.9674	60.1214	0.01%	0.0357	0.0553
	% change	C_L	86.87%	-55.32%	207.97%	547.35%	-35.51%
		C_L/C_D	37.95%	25.53%	95.98%	168.90%	-15.61%
		$\Delta\bar{C}_{L,LER}$	-24.07%	-56.28%	-99.92%	26.15%	-5.15%
3 Objective, camber $\leq 5\%$	value	C_L	1.6976	147.3361	8.47%	0.0468	0.0308
		C_L/C_D	1.5996	159.3386	11.43%	0.0498	0.0362
		$\Delta\bar{C}_{L,LER}$	0.9674	60.1214	0.01%	0.0358	0.0553
	% change	C_L	33.25%	7.15%	-35.72%	65.30%	-47.22%
		C_L/C_D	25.55%	15.88%	-13.22%	76.02%	-37.95%
		$\Delta\bar{C}_{L,LER}$	-24.07%	-56.28%	-99.92%	26.48%	-5.10%

Table 5. Summary of results. "Optimal Objective" refers to the camber and thickness at the most optimal point of the listed objective. Trailing edge (TE) thickness is taken at $x = 0.8$, where the thickness constraint of a minimum of 10% of the design thickness is enforced.

allows the optimizer to explore the design space without bias. The ES highlights that while certain existing designs lie close



to one slice of the Pareto front, expanding into a multi-objective space breaks out of local maxima/minima that traditional WT designs occupy.

At the outer bounds of the three-objective Pareto front (Table 5), the optimized airfoils improve upon the DU 93-210 by 87%
390 in C_L , 26% in C_L/C_D , and reduce $\Delta\bar{C}_{L,LER}$ to near zero. These airfoils involve high camber (up to 18% of chord) that may be impractical for manufacturing and structural reasons. When a 5% maximum camber constraint is applied in post-processing, the filtered front still yields improvements of 33% in C_L , 16% in C_L/C_D , and near-complete LER insensitivity, demonstrating that substantial gains are achievable within practical geometric limits. The framework provides a designer with a continuum of non-dominated airfoils, enabling explicit trade-offs between aerodynamic performance, roughness insensitivity, and structural
395 feasibility.

The most significant finding is that aft-loading, traditionally associated with improved LER insensitivity, correlates with *reduced* LER performance across the optimized airfoil set. The mechanism stems from roughness forcing transition near the leading edge, which causes the turbulent boundary layer to grow over nearly the entire chord. By the time it reaches the aft recovery region, it has accumulated sufficient momentum deficit that it cannot negotiate the strong adverse pressure gradient
400 imposed by high-camber geometries. The resulting early separation eliminates the lift contribution that aft-loading was intended to preserve. Traditional WT airfoils such as the DU 93-210 and FFA-W3-211 operate at sufficiently low camber that the adverse pressure gradient remains mild and the boundary layer stays attached under roughness, so the conventional wisdom holds in that regime. The transition between these two regimes represents a design boundary that, to our knowledge, has not been previously identified. Furthermore, because the roughness model used here ($N_{crit} = 1$) represents a severe condition, the camber threshold
405 at which the aft-loading benefit recovers likely depends on roughness severity and Reynolds number.

Several limitations motivate future work. XFOIL is reliable only in the pre-stall regime, so a full analysis of the Pareto-optimal airfoils with higher-fidelity CFD is needed to verify drag predictions, evaluate post-stall behavior, and further investigate the physical mechanisms underlying the interaction between aft-loading and LER insensitivity. The Reynolds number in this study ($Re = 1.5 \times 10^6$) is at the lower end of the range for modern turbines; extending to $Re = 15\text{--}20 \times 10^6$, representative
410 of outboard stations on large offshore rotors, will require higher-fidelity solvers. Because the transition behavior that drives this interaction is Reynolds-number dependent, the critical camber threshold may shift at higher Re . Similarly, evaluating the Pareto-optimal airfoils at intermediate N_{crit} values would clarify how roughness severity affects the relationship and identify the conditions under which the conventional design principle is recovered. Finally, the ES exhibits slower convergence near the global Pareto front, and hybrid strategies that switch to gradient-based refinement in the late stages of optimization may
415 improve efficiency.



Appendix A: Supplemental Figures




Figure A1. Absolute effects for a) the first pass parameters tuning, and b) the refined parameter tuning carried out only on parameters λ_m and σ_m




Figure A2. Shapes and polars for off design optimization at $\alpha = 6^\circ$

Author contributions. Ryan Cameron contributed to the conceptualization, developed the methodology, wrote the necessary software, performed the formal analysis, and wrote the original draft. Matthew Lackner contributed to the conceptualization, procured funding, edited and reviewed the draft, and supervised the research.

420 *Competing interests.* The authors declare that they have no conflict of interest.



Disclaimer. This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Acknowledgements. This program is supported: by the Massachusetts Clean Energy Center (MassCEC) Workforce program through funding that increases access to clean energy careers, education, and training, and with funding from the Massachusetts Department of Energy Resources (DOER). This material is based in part by work initially supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Wind Energy Technologies Office (WETO) Award Number DE-EE0011269, and continuing support from the Massachusetts Clean Energy Center and the Maryland Energy Administration. The views expressed herein do not necessarily represent the views of the U.S. Department of Energy, the United States Government, the Massachusetts Clean Energy Center or the Maryland Energy Administration.



435 References

- Abbott, I. H. and Von Doenhoff, A. E.: Theory of wing sections: including a summary of airfoil data, Courier Corporation, ISBN 978-0-486-13499-4, 2012.
- Back, T.: Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms, Oxford University Press, ISBN 978-0-19-535670-0, 1996.
- 440 Björck, A.: Coordinates and Calculations for the FFA-W1-xxx, FFA-W2-xxx and FFA-W3-xxx Series of Airfoils for Horizontal Axis Wind Turbines, Aeronautical Research Institute of Sweden, 1990.
- De Falco, I., Del Balio, R., Della Cioppa, A., and Tarantino, E.: Breeder genetic algorithms for airfoil design optimisation, in: Proceedings of IEEE International Conference on Evolutionary Computation, pp. 71–75, <https://doi.org/10.1109/ICEC.1996.542336>, 1996.
- De Falco, I., Della Cioppa, A., Iazzetta, A., and Tarantino, E.: A new mutation operator for evolutionary airfoil design, *Soft Computing*, 3,
445 44–51, <https://doi.org/10.1007/s005000050090>, 1999.
- Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE Transactions on Evolutionary Computation*, 6, 182–197, <https://doi.org/10.1109/4235.996017>, 2002.
- Doronina, O. A., Lee, B., Grey, Z. J., and Glaws, A.: Aerodynamic Sensitivities over Separable Shape Tensors, *AIAA Journal*, 63, 2707–2720, <https://doi.org/10.2514/1.J064749>, 2025.
- 450 Drela, M.: XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils, *Low Reynolds Number Aerodynamics*, 54, 1–12, https://doi.org/10.1007/978-3-642-84010-4_1, 1989.
- Ehrmann, R. S., White, E. B., Maniaci, D. C., Chow, R., Langel, C. M., and Van Dam, C. P.: Realistic Leading-Edge Roughness Effects on Airfoil Performance, in: 31st AIAA Applied Aerodynamics Conference, American Institute of Aeronautics and Astronautics, San Diego, CA, <https://doi.org/10.2514/6.2013-2800>, 2013.
- 455 Hansen, T. H.: Airfoil optimization for wind turbine application, *Wind Energy*, 21, 502–514, <https://doi.org/10.1002/we.2174>, 2018.
- Karcher, C. J., Maniaci, D. C., Kelley, C., Hsieh, A., deVelder, N., and Gupta, A.: Design of a Preliminary Family of Airfoils for High Reynolds Number Wind Turbine Applications, American Institute of Aeronautics and Astronautics, ISBN 978-1-62410-723-8, <https://doi.org/10.2514/6.2025-0840>, 2025.
- Kulfan, B. and Bussolletti, J.: "Fundamental" Parameteric Geometry Representations for Aircraft Component Shapes, in: 11th
460 AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, American Institute of Aeronautics and Astronautics, <https://doi.org/10.2514/6.2006-6948>, [_eprint: https://arc.aiaa.org/doi/pdf/10.2514/6.2006-6948](https://arc.aiaa.org/doi/pdf/10.2514/6.2006-6948), 2012.
- Mills, K. L., Filliben, J. J., and Haines, A. L.: Determining Relative Importance and Effective Settings for Genetic Algorithm Control Parameters, *Evolutionary Computation*, 23, 309–342, https://doi.org/10.1162/EVCO_a_00137, 2015.
- Ramsay, R. F., Hoffman, M. J., and Gregorek, G. M.: Effects of grit roughness and pitch oscillations on the S809 airfoil, Tech. Rep. NREL/TP-
465 442-7817, National Renewable Energy Lab. (NREL), Golden, CO (United States), <https://doi.org/10.2172/205563>, 1995.
- Tangler, J. L. and Somers, D. M.: NREL airfoil families for HAWTs, Tech. Rep. NREL/TP-442-7109, 10106095, <https://doi.org/10.2172/10106095>, 1995.
- Timmer, W. and Bak, C.: Aerodynamic characteristics of wind turbine blade airfoils, in: *Advances in Wind Turbine Blade Design and Materials*, pp. 129–167, Elsevier, <https://doi.org/10.1016/B978-0-08-103007-3.00011-2>, 2023.
- 470 Timmer, W. A. and van Rooij, R. P. J. O. M.: Summary of the Delft University Wind Turbine Dedicated Airfoils, *Journal of Solar Energy Engineering*, 125, 488–496, <https://doi.org/10.1115/1.1626129>, 2003.

<https://doi.org/10.5194/wes-2026-72>
Preprint. Discussion started: 28 April 2026
© Author(s) 2026. CC BY 4.0 License.



Van Rooij, R.: Modification of the boundary layer calculation in RFOIL for improved airfoil stall prediction, 1996.