Proceedings of the 4th International Conference on Applied Physics and Mathematical Modeling
DOI: 10.54254/2753-8818/2026.37878

Aerodynamic Optimisation in Car Exterior Design: Improving Energy Efficiency Through Formula 1 Design Principles

Yuanchen Zhao

Wuxi Dipont School of Arts and Science, Wuxi, China
003503@nkcswx.cn

Abstract. Exterior shape is one of the few levers a designer can pull to cut energy use without adding a single kilogram to the vehicle. Body contours govern how air attaches to the surface, where the boundary layer lets go, and how much static pressure is recovered before the flow leaves the tail. Those three behaviours set the drag area (CdA), and drag area sets power demand once speed climbs. This paper does not argue for transplanting racing hardware onto road cars. It uses Formula 1 aerodynamics, in particular the downforce-drag trade-off and adaptive devices such as the Drag Reduction System (DRS), as a conceptual template for condition-dependent flow control. Comparing passive shaping against active devices shows that the payoff is strongly speed-dependent: gains are modest in city traffic and substantial above roughly 100 km/h. Active hardware carries its own bill, though. Actuators add mass, draw current, complicate assembly and durability, and eat into cabin and luggage space. Morphing geometries and non-mechanical flow control look promising for the next generation of electric vehicles, but only where the fluidic gain outweighs the mass, energy and computational penalty they bring with them.

Keywords: Aerodynamics, Energy efficiency, Drag area, Active flow control, Vehicle geometry

1. Introduction

Transport is under pressure to decarbonise, and the pressure is not abstract. Internal combustion engine (ICE) vehicles still emit at the tailpipe in everyday use, while electric vehicles (EVs) run into range anxiety and the stubborn ceiling of battery gravimetric energy density. At highway speed, the dominant cost is not the drivetrain but the air: aerodynamic drag consumes between 50% and 70% of tractive energy. That figure explains a familiar complaint, namely that an EV which comfortably covers its rated range in the city can lose a fifth of it on a long motorway run. Adding cells or re-engineering the powertrain works, but both are expensive in the currencies that matter: curb weight, unit cost, and structural complexity. Refining the outer surface is the one intervention that buys range without adding mass.

Electrification has also rewritten trade-offs inherited from the combustion era. An ICE car needs a large frontal grille because the radiator must reject a great deal of waste heat, and pushing air

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2753-8818/2026.37878

through the engine bay costs momentum. An electric powertrain is thermally less demanding overall, yet far more demanding locally, since high-voltage battery modules need tight temperature control during fast charging and sustained high-speed cruising. The result is not less cooling but differently distributed cooling, which is why low-drag shaping on an EV is a coupled fluid-thermal problem rather than an exercise in smoothing panels. Get the coupling wrong, and the savings evaporate on a hot day, or after twenty minutes at 130 km/h.

The discipline itself has moved through three broad phases. Early attempts borrowed directly from aeronautics: the Rumpler Tropfenwagen and the Chrysler Airflow wrapped teardrop profiles around passenger compartments, cutting drag but compromising packaging and crosswind stability, which is part of why neither sold well. Wind tunnels then made the field empirical, delivering measured drag coefficients (Cd) and wake surveys, though at high operating cost and usually at reduced scale. Computational Fluid Dynamics (CFD) changed the resolution of the question again, letting designers read local pressure gradients across the whole surface. With that resolution came a shift in ambition: away from a single optimised static shape and toward flow managed in real time, an idea motorsport had exploited for years.

This paper asks how far exterior optimisation actually improves energy efficiency, and under what conditions. The argument runs along one physical chain: geometry governs boundary layer attachment, separation, and rear pressure recovery. These govern drag area (CdA). And CdA governs power demand. Formula 1 enters as a conceptual framework, through the downforce-drag trade-off and active mechanisms such as DRS, rather than as a parts catalogue for road vehicles. Section 2 sets out the fluid dynamic fundamentals, including CdA, wake behaviour, and thermal management. Section 3 reviews passive and active techniques. Section 4 turns to production vehicles and the compromises they reveal. Section 5 looks ahead to morphing surfaces and predictive control, and Section 6 concludes.

2. Fundamentals of vehicle aerodynamics

Energy savings cannot be read off the drag coefficient alone. The force a car must overcome depends on drag area, CdA, which multiplies geometric efficiency by frontal reference area [1]. A slippery shape on a tall, wide body may still perform poorly. Shrinking frontal area therefore counts as much as refining contours. Nor is a low Cd sufficient on its own. Lift coefficient (CL), crosswind sensitivity, and axle load distribution must be managed together if the car is to stay stable and brake predictably at speed [2].

Pressure over the body produces vertical force, which decides lift and pitch behaviour. Rear axle lift is the dangerous case: as the rear unloads, tyre grip falls, and the car becomes nervous in fast corners and under emergency braking. Balancing front-to-rear lift distribution is thus as important as trimming drag. Much of that balance is decided underneath, where proximity to the road alters static pressure beneath the floor (the familiar ground effect). Controlled pressure recovery along the floorpan is what prevents lift from changing abruptly when the car pitches or rolls.

Forward geometry sets up everything downstream. Hood curvature, windshield rake, and A-pillar radii determine where the flow attaches and how cleanly it travels over the front of the car [3]. A high-pressure stagnation region forms at the bumper, and abrupt changes of section behind it can separate the flow early and generate pressure drag [4]. Gentle transitions keep the boundary layer attached and preserve total pressure [5].

The rear does most of the damage. Moving aft, the boundary layer works against an adverse pressure gradient, and a roofline that falls away too steeply separates the flow before the tailgate can recover any pressure [6]. What remains is a low-pressure wake and a front-to-rear pressure

DOI: 10.54254/2753-8818/2026.37878

imbalance that typically forms the single largest component of drag [7]. A long, gradual taper fixes this in theory. In practice, rear headroom, luggage volume, and overall length place hard limits on how far the roofline can be drawn out [8].

Cooling openings set thermal need directly against drag. Air admitted through the front fascia is routed to heat exchangers and battery packs, and forcing it through constricted ducts destroys momentum, the penalty called cooling drag [9]. Active grille shutters and shaped ducting keep that penalty proportional to actual thermal demand.

3. Aerodynamic optimisation techniques

Reducing drag in practice means combining patient surface work with hardware that can change state [10]. The passive half of the job is to keep the boundary layer attached for as long as the body allows: flush door handles, sealed wheel rims, full underbody panels, and a rear that tapers as far as packaging permits [8]. None of these features is dramatic on its own. Collectively, they suppress local separation, cut parasitic drag, and weaken vortex shedding. Replacing mirrors with camera pods illustrates how narrow the margins can be. The aerodynamic case is straightforward, since the pods shrink frontal reference area and clean up the turbulent field around the A-pillars, but the engineering case is not. Approval varies by market, displays add latency, the hardware costs more, and the processing draws current continuously [10].

Wheels deserve separate attention, because they are the most turbulent region on an otherwise well-behaved car. A rotating wheel throws out strong outwash vortices and sits in its own low-pressure pocket. On a typical passenger car, the wheels and wheel wells account for as much as a quarter of total drag. Enclosing the rims, adding deflectors ahead of the front tyres, and flattening the floorpan all attack the same shear layers. Combined with grille shutters that close when cooling demand drops, these measures amount to an aerodynamic package that adapts to speed and thermal load rather than one tuned for a single operating point.

Appendages are widely misunderstood. A rear spoiler is rarely fitted to smooth the flow or to cut drag in isolation. Its job is to fix the separation line, reduce rear axle lift, and keep the car stable at highway speed, and it often costs a little drag to do so. Lift management, in other words, takes priority over drag in the region where the two conflict. The diffuser works by a different mechanism than its popular description suggests. It is not a channel that speeds air up and out, but an expanding duct in which fast-moving underbody air decelerates gradually and recovers static pressure before meeting the wake. Whether that recovery is worth having depends on expansion angle, ride height, and how the underbody flow interacts with what is coming off the roof.

Speed decides how much any of this is worth. Drag force grows with the square of velocity, $F_d = \frac{1}{2}\rho C_d A v^2$, while the power required to overcome it grows with the cube, $P_d = \frac{1}{2}\rho C_d A v^3$. The consequence is easy to state and easy to forget. Above roughly 100 km/h, where air resistance already accounts for 50% to 70% of tractive energy, a 10% reduction in C_dA is worth about 5% to 7% of highway range on an electric vehicle [8, 10]. Repeat the same modification in urban traffic and the benefit largely disappears, because rolling resistance and the energy spent accelerating mass now dominate the budget. Any honest assessment of an aerodynamic device therefore has to state the drive cycle on which it was evaluated, since the same component can look excellent on a motorway loop and worthless in a stop-start cycle. Formula 1 is instructive here rather than directly applicable. DRS opens a slot in the rear wing on designated straights, trading downforce for speed only where cornering grip is not needed. The underlying idea, that the optimal configuration changes with operating condition, transfers to road cars even though the hardware does not.

DOI: 10.54254/2753-8818/2026.37878

4. Case studies and practical applications

Production cars show what integrated aerodynamic work is actually worth. The Mercedes-Benz EQE (Cd = 0.22) and the Hyundai Ioniq 6 (Cd = 0.21) reach those numbers not through a single feature but through the accumulation of small ones: smooth front fascias, flush handles, covered underbodies, and rear sections tapered to delay separation. Concepts show the ceiling. The Mercedes-Benz VISION EQXX reaches Cd = 0.17 with an active rear diffuser and a narrowed rear track, a configuration that would be hard to justify in a car that also has to carry a family and its luggage. The translation into energy is direct enough to serve as a rule of thumb: at a steady 120 km/h, each 0.01 removed from Cd is worth roughly 2.5% to 3.0% of range, or about 0.4 to 0.6 kWh per 100 km (Howell & Le Good, 2008). Active grille shutters contribute separately, closing the cooling path when thermal demand is low and cutting cooling drag by up to 8% without letting powertrain temperatures climb. Real-world numbers do drift from certified ones, however. Published Cd figures come from a tunnel with steady, aligned flow, whereas a car on a motorway meets yaw angles from crosswinds and from the wakes of surrounding traffic, and drag at yaw can exceed the zero-yaw value appreciably. A shape optimised only at zero yaw can therefore look better on paper than it behaves on the road.

Borrowing more aggressive ideas from motorsport is where the compromises begin. Deployable rear spoilers, dynamic front splitters, and motorised underbody flaps do exactly what they promise, switching the car between a low-drag cruising configuration and a high-downforce stable one. The trouble is everything attached to them. Actuators, position sensors, wiring, and the structure needed to mount them all add mass, and mass is charged at a different rate depending on where the car is driven in an urban cycle dominated by rolling resistance and repeated acceleration; hardware that only earns its keep above 100 km/h can plausibly consume more energy than it saves.

Durability is the less discussed risk. Shutters, spoilers, and movable diffusers rely on electric actuators, position feedback, and seals that must survive thermal cycling, road salt, grit, and ice. The failure mode that matters is not a device stuck closed but one stuck open: drag rises, range falls, and the driver may have no indication beyond a worse trip average. Warranty exposure across a ten-year service life is a genuine constraint on how much active hardware a manufacturer is willing to fit, and it explains why many production systems are deliberately simpler than what is technically possible.

Surface design also carries acoustic consequences that have become more visible with electrification. Wind noise around mirrors, A-pillars, and door seals used to sit beneath engine noise. Remove the engine, and it does not. Reducing high-frequency turbulence in those regions now contributes directly to perceived cabin quality, a criterion that sometimes aligns with drag reduction and sometimes competes with it.

Packaging is the constraint that never goes away. An extended, aggressive rear taper keeps flow attached longer, and it also removes rear headroom and luggage volume (Barnard, 2009). Long rear overhangs worsen departure angles and make tight parking harder, which customers notice daily and which drag reduction does not compensate for. Sealed underbodies bring a thermal version of the same problem, trapping heat around the battery pack and brakes under sustained load and then requiring internal ducting that gives part of the aerodynamic gain straight back. Design teams end up arbitrating between efficiency, interior space, crash structure, cost, and everyday usability, and the aerodynamicist rarely wins every argument.

DOI: 10.54254/2753-8818/2026.37878

5. Future trends and research directions

Exterior design is moving away from fixed panels and toward surfaces that change shape, with flow that is actively managed rather than merely accommodated. Three developments drive this: smart materials, embedded sensing, and control algorithms cheap enough to run onboard.

Morphing skins are the most visible of the three. Shape-memory alloys (SMA), electro-active polymers (EAP), and flexible elastomeric composites allow a surface to deform continuously, without the panel gaps and heavy hinges that rigid flaps require, and panel gaps are themselves a source of drag, which is part of why mechanical solutions give back some of what they gain. A rear section could extend at motorway speed into something close to an optimal boat-tail, then retract in town so the car remains a normal length in a parking space. Durability and actuation energy are the open questions here, not the aerodynamic principle.

Sensing and control are advancing just as quickly. Conventional CFD is far too slow to sit inside a control loop and therefore lives in the design office. Physics-informed neural networks (PINNs) and deep-learning-accelerated solvers change that arithmetic, and combined with pressure taps distributed across the body, they make it possible to predict separation as it develops rather than after the fact. The actuator can then be non-mechanical: pulsed synthetic micro-jets or dielectric barrier discharge (DBD) plasma devices re-energise the boundary layer with no moving parts and no hinge mass. Whether their power draw is repaid by the drag they remove is precisely the question such systems must answer before they reach production.

Prediction can extend beyond the car itself. With vehicle-to-everything (V2X) communication and cloud-side forecasting, a control unit can read navigation data, wind forecasts, and road gradient, then set body panels or ride height before the gust or the climb arrives rather than reacting once it has. Freight benefits most obviously. Automated platooning depends on coordinated flow management to hold short following distances safely, and the trailing vehicles then run in a wake that has already been partly organised for them.

New architectures loosen constraints that shaped cars for a century. An autonomous electric vehicle need not preserve a driver's sightlines or a conventional windshield angle, which opens the door to teardrop cabins and enclosed wheel arches that would be unacceptable in a driven car. Heavy commercial EVs can combine platooning with adaptive trailer fairings to reduce wake interference between closely spaced trucks. In both cases, the gain comes from removing a requirement rather than from adding a device, historically the most reliable way to make a vehicle slippery. Test procedures will have to catch up as well. A vehicle whose shape changes with speed, temperature, and traffic cannot be characterised by a single Cd measured in a tunnel, and certification cycles that assume fixed geometry will misstate what such systems actually deliver.

6. Conclusion

This paper set out to trace how exterior aerodynamic work translates into energy efficiency, and the chain proved consistent: geometry determines boundary layer attachment and pressure recovery, those determine CdA, and CdA determines the power a vehicle must produce to hold speed. Reading passive shaping, appendages, and motorsport-derived active concepts against that chain. supports a straightforward conclusion. Aerodynamic refinement remains one of the few ways to cut consumption and extend EV range without adding mass, cost, or a larger battery.

Moving ideas from Formula 1 into production is a different exercise from admiring them. Racing demonstrates the value of condition-dependent flow control convincingly, but a road car has to

DOI: 10.54254/2753-8818/2026.37878

weigh that value against curb weight, tooling and unit cost, and hardware that must keep working for a decade in weather no racing car ever sees.

The limits are worth stating plainly. Low-drag geometry competes with interior volume, overall length, and the styling preferences customers actually buy on. Active devices such as deployable spoilers, splitters, and motorised shutters add actuators, complexity, and mass. In low-speed use, where rolling resistance and inertia dominate, that mass and the current the systems draw can cancel the savings they earn on the motorway. The case for active aerodynamics is therefore conditional rather than general, and it depends on how a given vehicle is actually driven.

Future work should concentrate where the penalty currently sits: lightweight morphing materials, reliable real-time pressure sensing, and wake control that a production control unit can run. A body that adapts its shape to the flow without carrying heavy hardware to do it would resolve most of the trade-offs described here, and would let designers pursue high-speed efficiency and low-speed usability together instead of choosing between them.

References

[1] Cembalo, A., Coirault, P., Borée, J., Dumand, C., & Mercère, G. (2024). Active control of road vehicle's drag for varying upstream flow conditions using a recursive subspace-based predictive control methodology. *Control Engineering Practice, 152*, 106071. https://doi.org/10.1016/j.conengprac.2024.106071

[2] Ravelli, U., & Ravelli, S. (2026). An overview of drag reduction methods in road cars. *Fluid Dynamics & Materials Processing, 22*(2), 1. https://doi.org/10.32604/fdmp.2026.075225

[3] Abood, M. S., & Hussain, I. Y. (2024). Ground effect of an inverted double element wing diffuser on a sedan car. *Heliyon, 10*(8), e29435. https://doi.org/10.1016/j.heliyon.2024.e29435

[4] Rijns, S., Teschner, T.-R., Blackburn, K., & Brighton, J. (2024). Performance analyses of active aerodynamic load balancing designs on high-performance vehicles in cornering conditions. *Physics of Fluids, 36*(8), 085199. https://doi.org/10.1063/5.0222780

[5] Seo, J. (2025). Component-level energy consumption and range analysis of battery electric vehicles under urban and highway driving conditions. *Applied Thermal Engineering, 274*, 126797. https://doi.org/10.1016/j.applthermaleng.2025.126797

[6] Keirsbulck, L., Cadot, O., Lippert, M., Boussemart, D., Basley, J., Delprat, S., & Paganelli, S. (2024). Underbody flow control for base drag reduction of a real car model. *Journal of Wind Engineering and Industrial Aerodynamics, 252*, 105822. https://doi.org/10.1016/j.jweia.2024.105822

[7] Piechna, J. R., Kurec, K., Broniszewski, J., Remer, M., Piechna, A., Kamieniecki, K., & Bibik, P. (2022). Influence of the car movable aerodynamic elements on fast road car cornering. *Energies, 15*(3), 689. https://doi.org/10.3390/en15030689

[8] Bao, D., Borée, J., Sicot, C., & Roebroeck, C. (2023). Salient features of wheel-vehicle aerodynamic interactions: Consequences for drag. *Journal of Wind Engineering and Industrial Aerodynamics, 236*, 105366. https://doi.org/10.1016/j.jweia.2023.105366

[9] Qin, P., Ricci, A., & Blocken, B. (2024). CFD simulation of aerodynamic forces on the DrivAer car model: Impact of computational parameters. *Journal of Wind Engineering and Industrial Aerodynamics, 248*, 105711. https://doi.org/10.1016/j.jweia.2024.105711

[10] Ma, X., Li, J., Zhao, J., & Chen, J. (2024). Aerodynamic characteristics of the race car in pitch and roll attitude. *International Journal of Numerical Methods for Heat & Fluid Flow, 35*(1), 330–357. https://doi.org/10.1108/HFF-05-2024-0375