



Giga-voxel computational morphogenesis for structural design

Aage, Niels; Andreassen, Erik; Lazarov, Boyan Stefanov; Sigmund, Ole

Published in:
Nature

Link to article, DOI:
[10.1038/nature23911](https://doi.org/10.1038/nature23911)

Publication date:
2017

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Aage, N., Andreassen, E., Lazarov, B. S., & Sigmund, O. (2017). Giga-voxel computational morphogenesis for structural design. *Nature*, 550(7674), 84-86. <https://doi.org/10.1038/nature23911>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Giga-voxel computational morphogenesis for structural design

Niels Aage^{1,2}, Erik Andreassen¹, Boyan S. Lazarov¹ & Ole Sigmund¹

¹Department of Mechanical Engineering, Technical University of Denmark, Nils Koppels Allé, Building 404, 2800 Kongens Lyngby, Denmark.

²Centre for Acoustic-Mechanical Micro Systems, Technical University of Denmark, 2800 Kongens Lyngby, Denmark.

Intelligent distribution of material is used to maximize the performance and minimize the cost of industrial products ranging from hearing aids to automobiles and aeroplanes. Historically, human intuition and insight have driven the evolution of mechanical design, recently assisted by computer-aided design approaches. The computer-aided approach known as topology optimization enables unrestricted design freedom and shows great promise with regards to weight savings, but its applicability has so far been limited to the design of single components or simple structures. That is, optimizing full assemblies, or large structures, require orders of magnitude higher resolution than permitted by current state-of-the-art methods^{1,2}. Here we report a computational morphogenesis tool, implemented on a supercomputer, that produces designs with giga-voxel resolution—more than two orders of magnitude higher than previously reported. Such resolution provides insights into the optimal distribution of material within a structure that were hitherto unachievable owing to the challenges of upscaling existing modelling and optimization frameworks. As an illustrative example, we apply the tool to the design of the internal structure of a full-scale aeroplane wing. The optimized full wing has unprecedented structural detail at length scales ranging from tens of meters to millimetres and, intriguingly, shows remarkable similarity to naturally occurring bone structures in, for example, bird beaks. We estimate that our optimized design corresponds to a reduction in mass compared to currently used aeroplane wing designs of 2–5 per cent, which translates into a reduction in fuel consumption of about 40–200 tonnes per year per aeroplane. Our morphogenesis process is generally applicable, not only to mechanical design, but also to flow systems³, antennas⁴, nano-optics⁵ and micro-systems^{6,7}.

Computational morphogenesis is used in engineering to determine the best possible shapes and material distributions for prescribed structural objectives. A very common goal is to minimize structural weight, subject to constraints on deflections and mechanical stresses. Similar goals are drivers of animal and plant evolution. For plants, the optimal ratio of loading capacity to weight ensures efficient use of limited resources; for animals, it minimizes energy consumption for walking

and running. Whereas efficient structures in nature generally result from slow genetic evolution, in engineering fast solutions that also take manufacturing limitations into account are necessary. Topology optimization^{1,2} works by redistributing material within a predetermined design domain, not unlike natural bone-growth processes. Here, the design domain is defined as the geometric volume in which material can be distributed, such as the space above the building lot for a high-rise building, or the internal part of a car, less the space reserved for passengers, the engine and so on. Through an iterative, deterministic process that involves modelling of the physical system, gradient computations and mathematical programming-based design updates, structural material is gradually redistributed towards the optimal design (Extended Data Fig. 1).

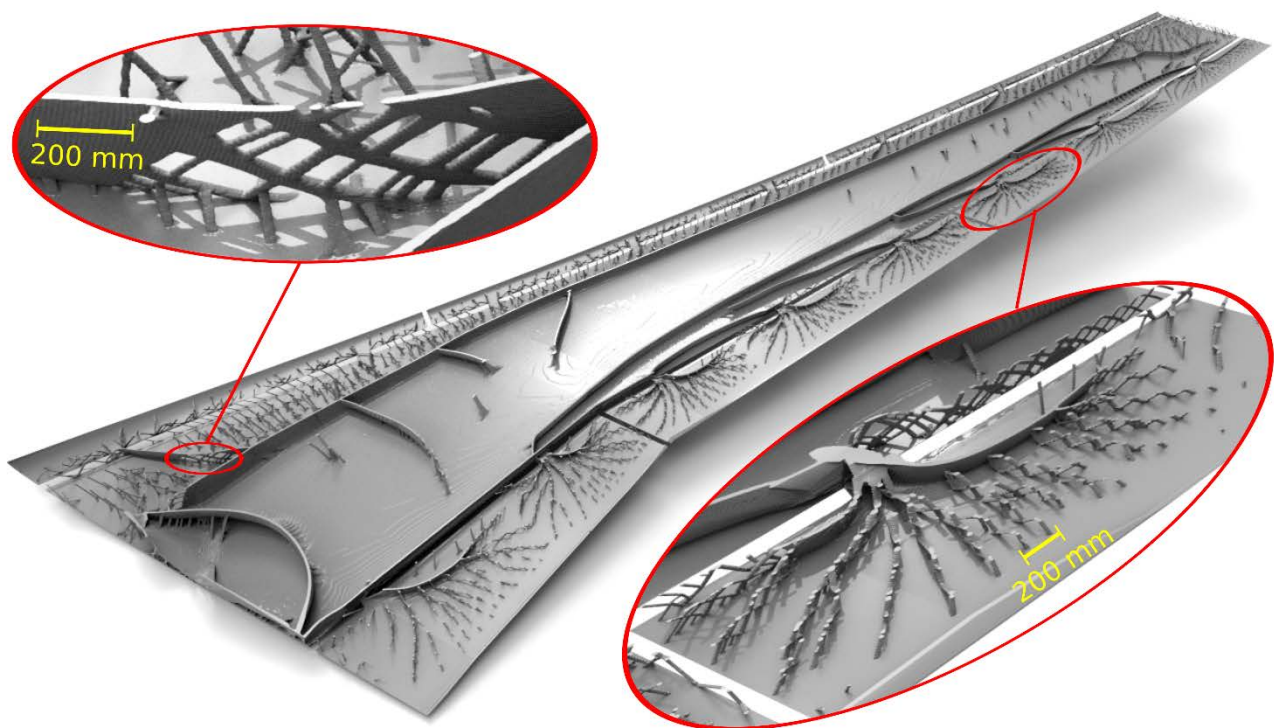


Figure 1 | Optimized wing structure. The result of the giga-voxel morphogenesis process applied to full-scale aircraft wing design is shown after 400 steps of the morphogenesis procedure. Fixed upper layers have been removed to reveal the internal details. Apart from the thin, fixed outer layer, no a priori assumptions were made about the internal geometry of the wing structure. Hence, all of the intricate details, such as curved spars, truss and wall structures, that can be observed have appeared spontaneously as a result of the optimization process.

Traditional stiffness-based topology optimization of components is well developed and used routinely in all major mechanical engineering industries, including the automotive⁸ and aerospace⁹ industries. In the automotive industry, this optimization typically results in weight savings of 20%–

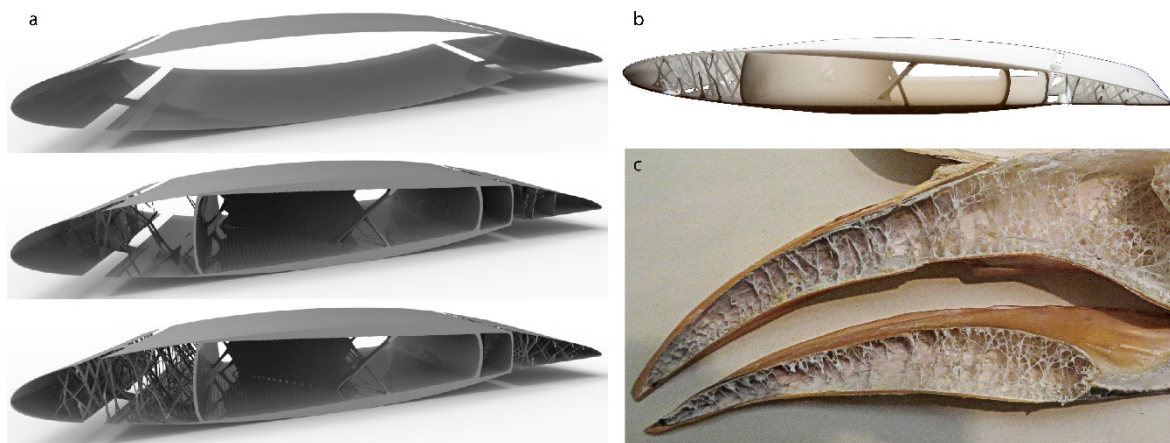


Figure 2 | Morphogenesis evolution, 3D printing and comparison to the hornbill bird beak. **a**, Steps of the morphogenesis process for the full-wing structure, demonstrated for a slice around the third connector from the wing root (see Extended Data Fig. 2 for precise location). The top figure corresponds to the initial design, in which only the outer aerofoil layer is given. The middle figure shows a snapshot during the morphogenesis process. The bottom figure shows the final design. **b**, A piece of the wing printed using additive manufacturing reveals additional details of and insight into the intricate but very regular structure of the optimized wing. The dimensions of the wing section are 300 mm × 100 mm × 45 mm. **c**, Hornbill bird beak with internal structure similar to that of the optimized aircraft wing. Picture taken by Jessica Miller-Camp at the Natural History Museum (London, UK) exhibit.

40% on structural parts and reduced development times compared to designs obtained by intuition or trial and error by skilled engineers, and eliminates the need to build prototypes. Despite its successes, the applicability of the topology optimization approach is still limited to the design of components or smaller structures. The main obstacle in applying it to large structures is resolution. Just as the visual quality of a television screen is limited by the number of pixels, the geometric quality and capability of the topology optimization process is limited by the number of voxels (three-dimensional equivalent of pixels) that are used to parameterize the design problem. Current state-of-the-art implementations limit voxel numbers to a few million and hence make their application to the detailed design of high-rise buildings, oil rigs or large aeroplanes impossible. This provokes the question of whether giga-voxel-resolution computational morphogenesis would lead to radical design changes. If so, we expect these new designs to result in substantial weight savings and reduced environmental impact.

We investigate this question using high-performance computing and demonstrate the potential of giga-voxel resolution computational morphogenesis on a full-scale wing-design

problem for an aircraft similar to a Boeing 777 (Fig. 1; see also Extended Data Figs 2 and 3, and Supplementary Figs 1 and 2 for high-resolution images). The extension to being able to solve problems with more than one billion voxels requires several non-trivial challenges to be overcome, including the efficient solution of huge linear equation systems, tailoring of optimization algorithms, and data transfer and visualization (see Methods). Obtaining the optimized designs presented here required access to massive computational resources with run times of 1–5 days on 8,000 CPUs.

The case study that we present considers the design of the load-carrying internal structure for a full aircraft wing based on the publicly available NASA Common Research Model^{10,11}. This wing model is similar to the wing of a Boeing 777 and has a half wing span (length of a single wing) of 27 m and an inner and outer chord (the ‘widths’ of the wing at the root and the tip) of 12 m and 1.5 m, respectively. The fixed-shape wing is discretized into 1.1 billion voxels (corresponding to a mesh of $1,216 \times 3,456 \times 256$ elements) with a maximum element size of 0.8 cm at the root, which is reasonable for a wing structure of this size¹². From the NASA model we included two aerodynamic load scenarios, corresponding to angles of attack of 0° and 4° at Mach 0.85 and an altitude of 10,700 m. The outer layer of voxels is fixed to be solid, representing the outer layer of the aerofoil; material can then be freely distributed in the internal part of the wing with no further restrictions (Fig. 2a, top). The optimized design considering a simplified full-wing structure including flaps is shown in Fig. 1, with the top half of the outer layer removed to reveal the optimized interior structures (the full wing is shown in Extended Data Fig. 2). Remarkably, all of the internal structural details appeared spontaneously as a result of the morphogenetic process. A first and striking observation is the extremely intricate details and features that come out of the giga-voxel-resolution design studies revealing structural details covering multiple length scales. That is, Fig. 1 shows the presence of structural members covering the entire wingspan, i.e. measuring tens of meters, as well as truss-like structures at the leading and trailing edges with dimensions in the mm range (highlighted in the zoom inserts of Fig. 1). In addition, expected structural features do not occur; the structure does not have the standard straight-wing box feature, made up of two spars, straight ribs and reinforced bottom and top outer layers. Instead, on closer inspection and study of the cross-section (Fig. 2a), the spars of the wing box are revealed to be curved to optimize torsional stiffness (the optimal cross-section of a pure torsion rod is circular). Furthermore, the irregular wing box is reinforced by curved and diagonally placed ribs, which are not seen in conventional aircraft wing designs. A conservative estimate of the weight savings that

are achievable compared to currently used designs is 2%–5%, corresponding to 200–500 kg per wing. Conversion of this weight reduction to fuel savings indicates that a potential saving of 40–200 tonnes of fuel per year is possible (see Methods for the analysis behind these predictions). The unconventional structural details provide new insight and will help aircraft designers to develop lighter and stronger wings in the future.

The outcome of the optimization process depends strongly on the choice of loading cases. To illustrate this aspect we investigate the influence of varying loading scenarios (Extended Data Fig. 3). Optimizing for single aerodynamic loads results in open-walled wing boxes (Extended Data Fig. 3a, b) and hence designs that have very low torsional stiffness. On the other hand, optimizing for two (or more) load situations results in closed and torsionally stiff wing boxes (Extended Data Fig. 3c, d). In practice, engineers will first optimize for what are expected to be the most critical load cases. Subsequently, the resulting wing structure will be checked for all design load cases and, if it turns out that some load cases that were not considered originally are critical for the new design, they will be included in a new optimization run.

When comparing the results presented here with the state-of-the-art in engineering design, we note that currently available methods and software have limited design resolutions of approximately five million voxels and hence a minimum feature size of 5 cm for the wing example. This feature size is too large to capture typical wing features and hence design studies based on coarse discretizations produce blurred and unclear designs with no spar- or rib-type formations. This issue has been discussed in previous studies of full-scale wing-structure optimization^{9,13,14,15}. In ref. 15, the lack of appearance of classical features such as ribs and spars is attributed to inappropriate physics, too few load cases and incorrect boundary conditions, or incomplete formulation of the optimization problem. In contrast, on the basis of our high-resolution study, we suggest that the real problem stems from a lack of resolution. With a giga-voxel framework, we can solve engineering design challenges that hitherto were considered insurmountable.

The optimized wing designs from our giga-voxel-resolution case study are subject to limitations and manufacturing challenges. In the analyses we included only linear elasticity; we leave the inclusion of anisotropic and nonlinear materials, dynamic effects, buckling, aero-elasticity, and so on, for future studies. Also, the fabrication of the complex design is currently infeasible. However, we imagine that large-scale additive manufacturing (3D printing) will be a possibility in the future, and we illustrate the potential use of additive manufacturing by printing a

small section of the wing design (Fig. 2b). Despite these limitations, we expect that our designs will be of interest to structural engineers, who can use our case study for inspiration for future, simplified designs based on the intriguing curved and diagonal ribs seen in the optimized designs.

As a last remark on our case study, it is striking that the optimized designs resemble natural bone structures, such as the hornbill bird beak structure shown in Fig. 2c. This is not a coincidence. Wolff's law states that animal bone adapts to applied loads¹⁶, and computational morphogenesis works in a similar way. Both structures have closed outer shells (for aerodynamic or food containment purposes) and need to be stiff against longitudinal bending and applied surface pressure, thus resulting in similarly looking internal reinforcement structures. In other words, the outcome of the computational morphogenesis approach can be compared to those of evolutionary design processes in nature, with one key difference—it is obtained in a small fraction of the time, owing to the efficient, deterministic optimization approach.

The methodology that we have presented is directly applicable to similar morphogenesis problems in other engineering disciplines, as well as in architecture and industrial design. Structures relevant to this approach are ships, offshore structures, wind turbine blades and towers, power masts, high-rise buildings, bridges and other complex built structures. Giga-voxel-resolution capability is expected to revolutionize design outcomes and yield new insights into optimal design. The computing time and resources reported here provide a severe bottleneck and will, for the time being, limit the use to early design stages in research-intensive (large) companies and academia. However, procedures for reducing the number of degrees of freedom by model reduction^{17,18}, reducing the number of iterations, and increasing the granularity by deeper hierarchy¹⁹, are being developed and are expected to result in substantial reductions in computing requirements in the near future. In addition, improvements are needed in peripheral computational tools such as meshing algorithms and visualization, as discussed in Methods.

Received 4 January; accepted 4 August 2017.

Supplementary Information is available in the online version of the paper.

Acknowledgements This work was funded by the Villum Foundation through the NextTop project and a PRACE (Partnership for Advanced Computing in Europe) grant TopWING giving access to the Curie supercomputer (GENCI@CEA, France). Access to, and efficient support from, the technical staff at Curie is highly appreciated. We also acknowledge access to and support from the Visualization Cluster at Copenhagen University through T. Haugbølle and Å. Norlund, and discussions with A. Horsewell and J. J. Thomsen from the Technical University of Denmark.

Author Contributions N.A. contributed to the original idea, method development, implementation, supercomputing, visualization, renderings and manuscript preparation. E.A. contributed to the original idea, method development, implementation, visualization and manuscript preparation. B.S.L. contributed to mesh mapping and manuscript editing. O.S. contributed to the original idea, method development, analytical studies and manuscript preparation.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations. Correspondence and requests for materials should be addressed to N.A. (naage@mek.dtu.dk).

METHODS

General

Among the many topology optimization methodologies available, we work with the so-called density approach. We first describe the basic methodology and, second, the details of extending this approach to the giga-voxel-resolution regime.

Topology optimization

The structural design problem considered here, that is, stiffness maximization, is solved on the basis of an assumption of static, linear elastic behaviour. Hence, the governing physics corresponds to the linear Navier–Lamé equations with appropriate boundary conditions. The partial differential equation is solved using the finite element method, using eight-node hexahedral elements, which results in a linear system of equations:

$$K(\mathbf{r})\mathbf{u} = \mathbf{F}$$

Here $K(\mathbf{r})$ is the stiffness matrix, with properties that depend on the spatial coordinate \mathbf{r} , \mathbf{u} the displacement vector and \mathbf{F} the load vector. The stiffness dependence on the spatial coordinate is introduced to facilitate the structural optimization procedure presented next. By smoothly interpolating between a reference solid (here aluminium with Young's modulus $E_{\text{solid}} = 70$ GPa) and a very weak material ($E_{\text{void}} = 10^{-9}E_{\text{solid}}$) using the 'solid isotropic material with penalization' (SIMP)^{1,2} model, we are able to use gradient-based optimization methods to solve the stiffness maximization problem. In the following, we assign a single density variable to each element in the mesh and state the SIMP interpolation scheme as

$$E(\rho) = E_{\text{void}} + \rho^p (E_{\text{solid}} - E_{\text{void}})$$

Here $\rho \in [0, 1]$ is a continuous, relative density and $p \geq 1$ is a penalization parameter that enforces almost discrete final designs with ρ being either 0 or 1 through penalization of the stiffness for intermediate densities. The final ingredient, before stating the mathematical optimization problem, is to introduce compliance (C) as the objective function for structural design. Compliance is the work done by the external forces, $C = \mathbf{F}^T \mathbf{u}$, and is inversely proportional to structural stiffness. Hence, to maximize structural stiffness we pose the following minimization problem:

$$\min_{\mathbf{x} \in \square^n} \phi = \sum_i^L \mathbf{F}_i^T \mathbf{u}_i \quad (1)$$

subject to:

$$\begin{aligned} K(\mathbf{x})\mathbf{u}_i &= \mathbf{F}_i, \quad i = 1 \dots L \\ \frac{V(\mathbf{x})}{V^*} - 1 &\leq 0 \\ 0 \leq x_j &\leq 1, \quad j = 1 \dots n \end{aligned} \quad (2)$$

Here \mathbf{x} is a vector of mathematical design variables with length corresponding to the number of design elements and ϕ is a sum of multiple compliances arising from L loading scenarios. The three constraints are as follows: the first ensures mechanical equilibrium; the second poses a limit on the amount of available material, with $V(\mathbf{x})$ being the volume of the current structure and V^* the total amount of available material; and the third constitutes box constraints on the design variables. Owing to inherent issues with numerical stability and mesh dependence², we use a filtering technique to relate the mathematical design variables \mathbf{x} to the physical density ρ (refs 20, 21). The filter operation is a convolution-type operator, which can be stated for the discretized model as

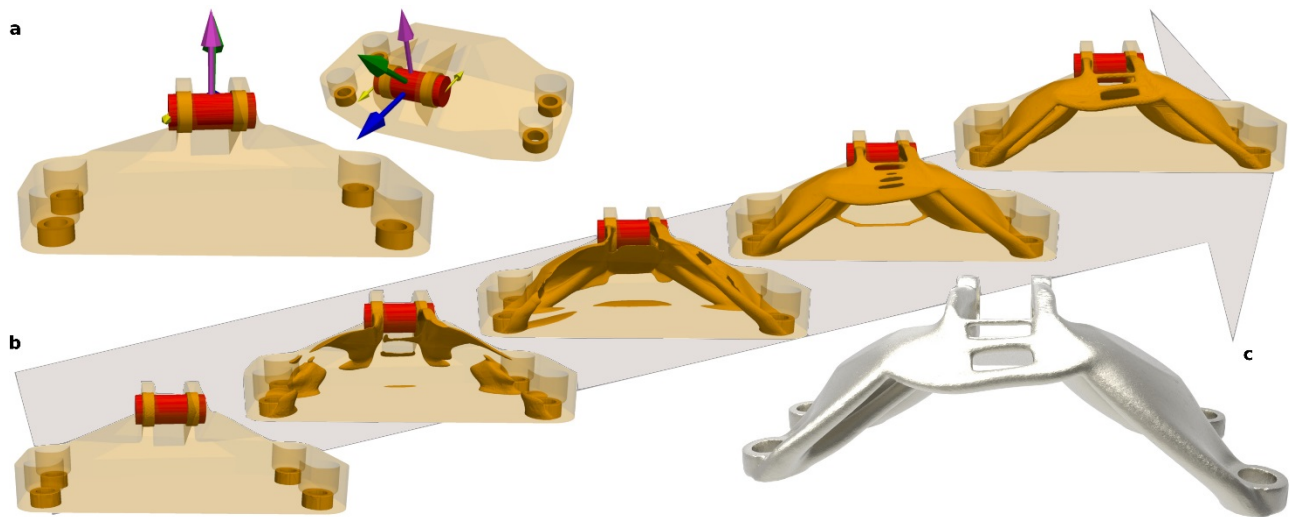
$$\rho_e = \frac{\sum_{i \in N_e} w(r_i) v_i x_i}{\sum_{i \in N_e} w(r_i) v_i}$$

Here $N_e = \{i \mid r_i - r_e \leq R\}$ is a neighbourhood set, with R being the filter radius, r_i the spatial (centre) coordinate of element i , r_e the spatial (centre) coordinate of element e , $w(r_i) = \max\{0, R - (r_i - r_e)\}$ the weighting function and v_i the volume of element i . To obtain the highest level of structural detail, we use a filter radius corresponding to a single voxel, i.e. a neighbourhood of 27 voxels. The sensitivities of the objective function are determined efficiently without additional analysis using the adjoint method, and the filtering operation is included by the chain rule²². The optimization problem is solved using a nested approach and an optimality criteria algorithm^{2,23} (simple MATLAB code available at <http://www.topopt.dtu.dk>). Owing to the massive increase in

computational power over the past decades, the design problem in equations (1)–(2) can now be solved interactively even on hand-held devices such as smartphones or tablets for ‘small’ values of n (up to several thousand).

Example problem

To demonstrate the iterative procedure of the topology optimization method, we present the solution to a simple part-design problem (Extended Data Fig. 1). We consider the design of a jet-engine bracket used in the 2013 GrabCAD challenge. Extended Data Fig. 1a shows the design domain together with supports and a number of applied mechanical loading cases. Extended Data Fig. 1b shows snapshots of the material distribution during the 200-step morphogenesis process, starting from a blank design.



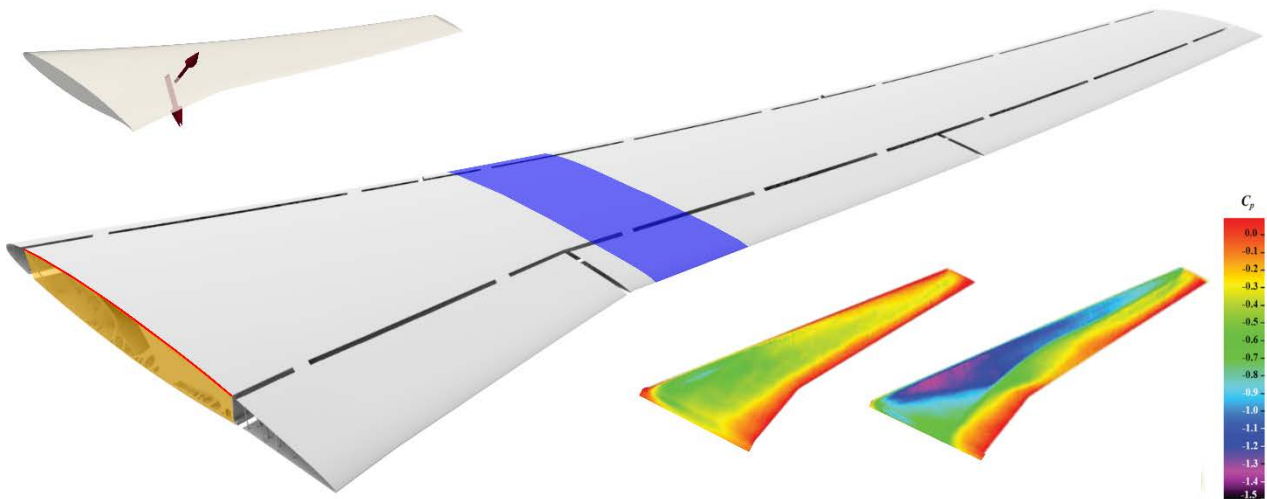
Extended Data Figure 1 | Schematic of the morphogenesis process. The figure illustrates the design of a jet engine bracket by morphogenesis. **a**, The design domain and the four load scenarios (coloured arrows). **b**, Snapshots from the design evolution. **c**, Rendering of the final titanium bracket design.

An animation of the iteration process is provided in Supplementary Video 1. The optimized design, obtained by topology optimization, saves an additional 10% in weight compared to the winner of the GrabCAD challenge, which had 638 participants world-wide. Interestingly, even though we optimized for only compliance, the resulting design easily satisfies the strength requirement set in the competition. This is partly due to the filtering process, which not only enforces a minimum length scale, but also avoids sharp corners and thereby stress concentrations. The example demonstrates that compliance-based topology optimization generally, but not always,

results in structures with favourable strength response. Remarkably, but unsurprisingly, the lack of initial geometric constraints and bindings to computer-aided design parameterizations result in the morphogenesis process producing rather organic looking structures that provide optimal force paths from applied forces to the supports, with features reminiscent of human and animal bone structure.

Meshing for giga-voxel resolution

Generating a body-fitted mesh with more than one billion finite elements is not a trivial task for standard mesh tools; hence, we use alternative procedures. In this work we have used a structured, logical mesh, which is then mapped, slice-by-slice, to the geometry of the NASA Common Research Model. Although a shell-element discretization of the outer shell would increase accuracy of the model, on the basis of experience and tests of various element types, we do not expect that a mixed solid/shell formulation will change the resulting morphologies in any substantial way.



Extended Data Figure 2 | Wing design model problem. Rendering of the full wing model is shown. All load cases use symmetry conditions at the root (orange) with only the top-most layer of elements being fully fixed (red). The bottom right insets show the pressure coefficients C_p of the two aerodynamic load cases (adapted from figure 10 in ref. 11) and the top left inset indicates the position and direction of the simulated engine loads for the third load case (purple arrows). The section highlighted in blue marks the cut-out region used in Fig. 2a, b.

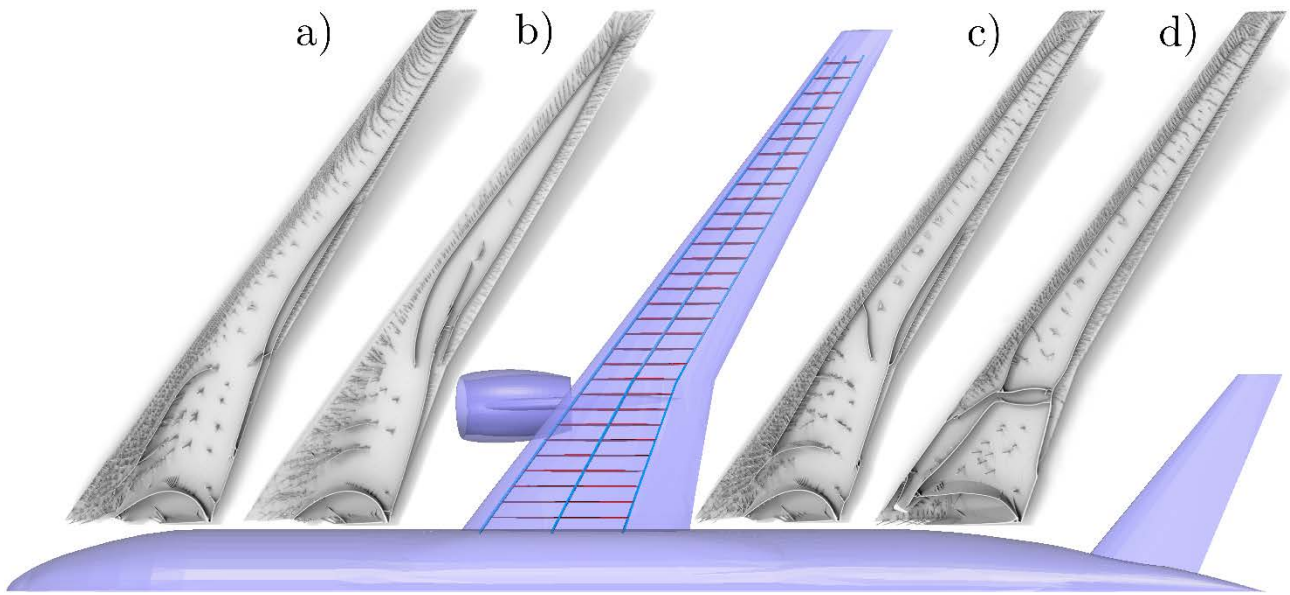
Ensuring structural details

Owing to the non-convexity of the stiffness-penalized optimization problem in equations (1)–(2), any gradient-based solution method is likely to end at a local minimum. To ensure that the designs produced are of high quality, that is, strong local minimum, we use a continuation strategy for the penalization parameter in the SIMP interpolation. This is motivated by two factors. First, equations (1)–(2) with $p = 1$ is a convex problem. However, without stiffness penalization it will not result in

pure solid-void designs, also called a 0–1 solutions, and so intermediate densities will dominate the optimized design. Second, it is of key interest that this study produces optimized designs with a high level of structural detail. That is, a limitation on structural details will inadvertently result in a reduced design space and thus lead to sub-optimal structures. Starting the optimization procedure with, for example, $p = 3$ will quickly force the design into a 0–1 structure, but with very little and crude detail. Both of these issues are alleviated by a continuation strategy in which the penalization parameter is slowly raised in steps of 0.25 from 1 to 3 [AUTHOR: The final value of the penalization parameter is 3], distributed over a total of 400 design cycles.

Giga-voxel-scale solution procedure

By far the most costly and challenging part of the topology optimization process is the solution of the huge linear systems that originate from the finely discretized finite-element problem. For a one-billion-element discretization, the linear systems have more than three billion unknowns. Despite

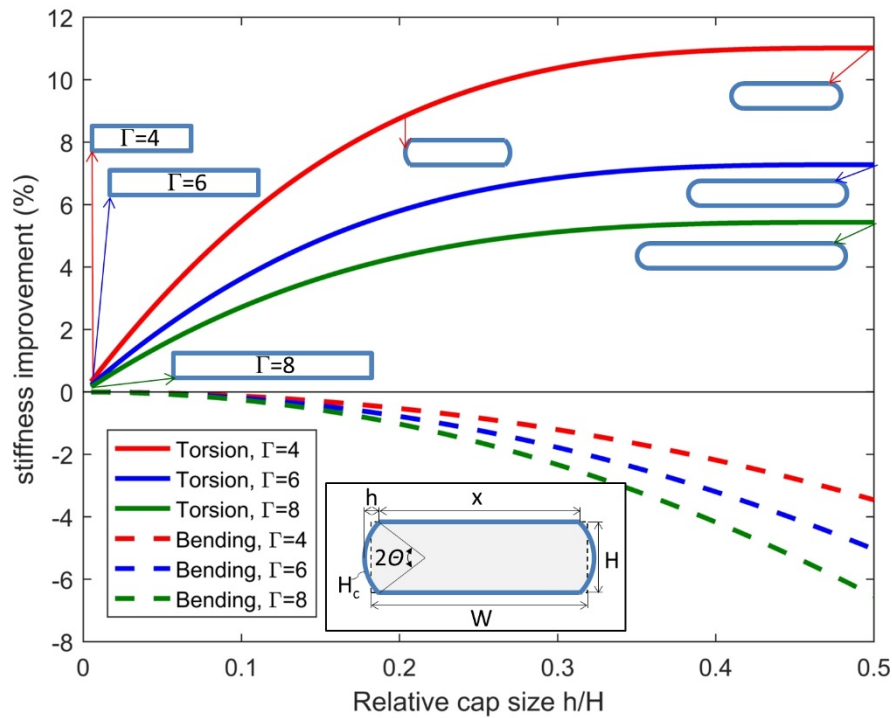


Extended Data Figure 3 | Influence of multiple load cases. a–d, Top view of the NASA Common Research Model (blue), along with the optimized wing-box structures for varying loading cases (grey). The NASA Common Research Model shows a standard wing-box design with spars (dark blue) and ribs (red). The optimized designs are for a single load at 0° incidence (**a**), a single load at 4° incidence (**b**), two loads at 0° and 4° incidence (**c**), and three loads, including engine load and two aerodynamic loads at 0° and 4° incidence (**d**). The load cases are shown in Extended Data Fig. 2.

sparsity, they cannot be solved using direct solvers owing to storage and memory requirements, and so require iterative approaches that make use of efficient preconditioners. Construction of a powerful and robust linear solver and preconditioner is hard for several reasons. First, the equation

system is ill-conditioned owing to high stiffness contrast. Second, depending on the design problem, high aspect ratios of the design domain and ill-formed elements may challenge convergence further. Therefore, solution strategies that work well for coarser problems are inadequate when the problem is scaled up²⁴. The solution methodology proposed next is based on the large-scale topology optimization framework presented in ref. 25.

We applied Krylov-subspace-equivalent methods as the overall tool for both the solver and preconditioning steps. These are chosen because of their explicit dependence on (sparse) matrix–vector products and, hence, very efficient parallel computer implementation. The key to a robust and efficient linear solver lies in the construction of a good preconditioner—that is, one that is fast to compute and to apply. Doing so can require long development times, which are often not reported²⁶. The present methodology is built as an add-on to the numerical library PETSc (ref. 27). The result is a very robust multi-level preconditioner strategy that uses Galerkin-projected multigrid with successive over-relaxation preconditioned Chebyshev smoothers down to a coarse level of

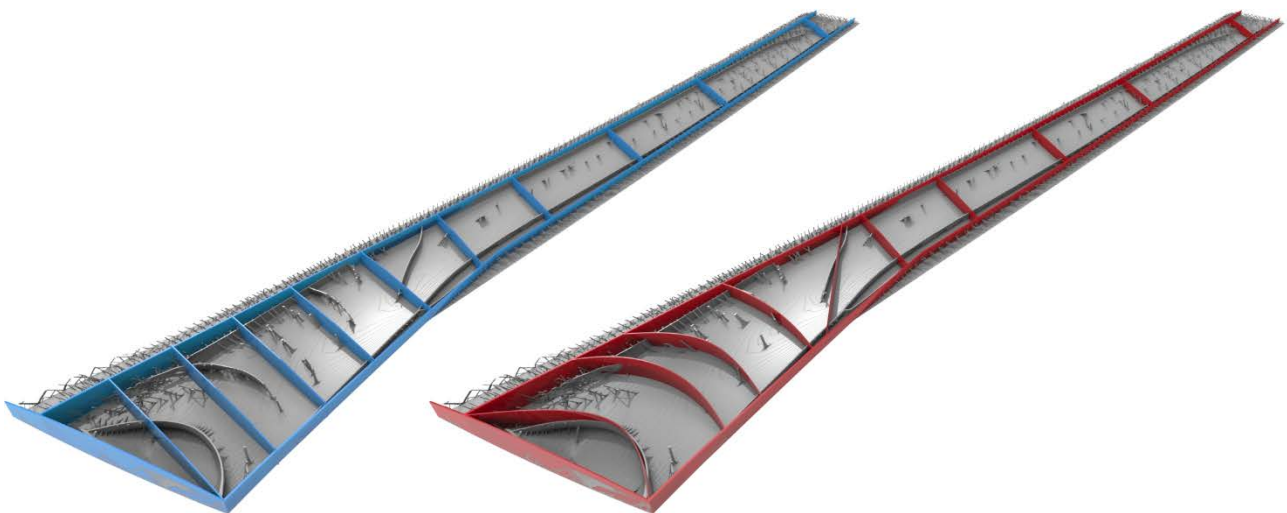


Extended Data Figure 4 | Curved spars. Stiffness improvement (per cent) is shown as function of the dimensions of the curved spar ($h/H = 0$ indicates a straight spar and $h/H = 1/2$ indicates a half-circular spar; some approximate shapes are sketched, indicated by arrows). Red, blue and green lines correspond to torsion boxes with aspect ratios of 4, 6 and 8, respectively. Solid lines indicate increased torsion stiffness and dashed lines decreased bending stiffness. The insert shows the parameterization of the cross-section of the wing box.

approximately six million degrees of freedom using four geometric multigrid levels. The resulting coarse-grid problem is then solved using an algebraic-multigrid, preconditioned conjugate-gradient method. As the outer solver we use the flexible generalized minimum-residual approach²⁸ owing to the slightly varying preconditioner that arises from the smoothers. The ultra-large-scale problems considered require massive computer storage capability, and visualization is also a challenge when the designer wishes to inspect the results of the morphogenesis process.

Computations

The giga-voxel-scale optimization problem, and the solution strategy that we have presented, requires a massive amount of computational resources. All results presented here were generated on the Curie Supercomputer (GENCI@CEA) in France. The optimal number of processors capable of solving the problem was found by numerical experiments and memory requirements to be around 8,000.



Extended Data Figure 5 | Non-traditional ribs. Simplified wing-box models were used to estimate the stiffness gain from non-traditional wing rib geometries. Both images show the optimized wing box for two aerodynamic load cases (compare with Extended Data Fig. 3c) in grey. The optimized design is overlaid in blue (left), showing a wing box with conventional straight ribs, or red (right), showing a wing box with unconventional ribs. Thicknesses of ribs are tailored to ensure equal mass.

Estimate of potential weight savings

The wing study presented is qualitative. To determine an exact figure of weight savings would require more detailed studies and tens of thousands of engineering hours for one specific aircraft. However, some simpler, partly analytical studies can be performed to obtain an estimate of potential weight savings. The topology optimization process resulted in two major geometrical findings: first,

wing-box spars are rounded instead of straight; and second, the few wing ribs that appear have odd geometries, represented by curved and diagonally placed ribs.

First, we perform an analytical study to estimate the stiffness improvement (or weight saving) that is achievable by rounded instead of straight wing spars. Here the wing box is assumed to be straight, with constant wall thickness and rectangular cross-section with aspect ratio $\Gamma = L/H$ (for definitions of parameters see inset in Extended Data Fig. 4). The torsional stiffness of the wing box is proportional to the ratio of the squared cross-sectional area and its perimeter (this is why a circular cross-section represents the optimal torsion rod). The torsional stiffness can therefore be improved by substituting the straight spars with circular ones. Assuming constant perimeter, the relative torsional stiffness improvement for a cross-section with circular end caps of relative width h/H can be found analytically, resulting in a simply derivable, but complicated expression. The resulting stiffness improvement for three different wing-box aspect ratios is shown in Extended Data Fig. 4. The improvement is largest for smaller aspect ratios (for which the torsional stiffness difference between a circular and a square cross-section is the largest). For all aspect ratios, cross-sections with half-circular ($h/H = 1/2$) end caps are the stiffest. The improved torsional stiffness comes at the cost of decreased bending stiffness. The curved shear webs (spars) are less efficient in carrying shear loads. This time, again assuming constant material volume, a circular end cap results in a longer and thinner shear web, resulting in a substantial decrease in stiffness. The analytical study for bending assumes a thin-walled prismatic wing box with ratio of height to span length of 20. The optimal spar curvature for the full-wing structure thus represents a compromise between increase in torsional stiffness and decrease in bending resistance. Estimating cap sizes from the wing design studies to be $h/H = 0.12$ and the wing-box aspect ratio to be around $\Gamma = 6$ (see cross-sectional plot in Fig. 2a (bottom) and Fig. 2b), the expected torsional stiffness improvement or similar weight saving is 4% (see Extended Data Fig. 4). At the same time, the loss in bending

Load scenario:	Conventional design [J]	Optimized Design [J]	Improvement
Load case 1:	511,479	506,756	0.92%
Load case 2:	15,112,593	14,987,090	0.83%

Extended Data Table 1 | Performance of non-traditional and traditional rib structures

Table showing the compliances for two aerodynamic load cases for the post-processed designs presented in Extended Data Fig. 5. The improvements obtained by the curved ribs are stated per cent.

stiffness is only 0.5%. Depending on the actual distribution between bending and torsional forces, the increase in stiffness or corresponding weight saving may be 1%–4%. Considering that a full

aircraft wing can weigh 10 tonnes (ref. 29) and that 100–200 kg of fuel may be saved per kilogram of saved weight per year for a commercial aeroplane³⁰, this weight saving is important—even a small percentage can have a huge influence.

Second, we perform a numerical study on a crudely simplified wing-box model to estimate the stiffness improvements achieved by curved and diagonal ribs placements compared to conventional straight ribs. As a basis for the study, we use the optimized wing box shown in Extended Data Fig. 3c as inspiration for an engineering interpretation. The result of this post-processing step is shown in Extended Data Fig. 5, in which the two reinforced wing-box models are fitted to have the same mass. The study is performed on a mesh of 141 million elements (430 million degrees of freedom), which is fully capable of capturing the details of the post-processed wing-box designs, and the numerical experiment is conducted using 640 cores on a small cluster. To quantify the performance improvement of the optimized rib/spar design compared to the conventional rib/spar design, we subject both models to the two aerodynamical load scenarios as described and depicted in Extended Data Fig. 2, that is, the same two load cases used in the initial topology optimization. The numerical results are collected in Extended Data Table 1. From the results in Extended Data Table 1, we observe that even a crude interpretation of the optimized wing-box design leads to a performance improvement of 0.8%–0.9% for the two load scenarios.

Finally, we summarize the possible weight savings identified through the study of two of the main features observed in the optimized wing structure. The weight of both wings is about 20 tonnes, and our work indicates that a total weight reduction of 2%–5% is possible using both the curved spars and the new rib layout. With regard to weight reduction, this corresponds to a total saving of 400–1,000 kg, corresponding to an estimated reduction in fuel consumption of 40–200 tonnes per year^{29,30}.

Data and code availability

The basis code used for this work is publically available at https://github.com/topopt/TopOpt_in_PETSc (ref. 25) and the results, that is, the optimized designs in the form of STL files, are available from the corresponding author on reasonable request.

References

1. Bendsøe, M. P. & Kikuchi, N. Generating optimal topologies in structural design using a homogenization method. *Comput. Methods Appl. Mech. Eng.* 71, 197–224 (1988).

2. Bendsøe, M. P. & Sigmund, O. *Topology Optimization — Theory, Methods and Applications* (Springer, 2004).
3. Alexandersen, J., Sigmund, O. & Aage, N. Large scale three-dimensional topology optimisation of heat sinks cooled by natural convection. *Int. J. Heat Mass Transfer* 100, 876–891 (2016).
4. Nomura, T., Sato, K., Taguchi, K., Kashiwa, T. & Nishiwaki, S. Structural topology optimization for the design of broadband dielectric resonator antennas using the finite difference time domain technique. *Int. J. Numer. Methods Eng.* 71, 1261–1296 (2007).
5. Jensen, J. S. & Sigmund, O. Topology optimization for nano-photonics. *Laser Photonics Rev.* 5, 308–321 (2011).
6. Sardan, O. et al. Rapid prototyping of nanotube-based devices using topology-optimized microgrippers. *Nanotechnology* 19, 495503 (2008).
7. Dühning, M., Mortensen, N. A. & Sigmund, O. Plasmonic versus dielectric enhancement in thin-film solar cells. *Appl. Phys. Lett.* 100, 211914 (2012).
8. Cavazzuti, M. et al. High performance automotive chassis design: a topology optimization based approach. *Struct. Multidiscipl. Optim.* 44, 45–56 (2011).
9. Zhu, J., Zhang, W. & Xia, L. Topology optimization in aircraft and aerospace structures design. *Arch. Comput. Methods Eng.* 23, 1–28 (2015).
10. Vassberg, J. C., DeHaan, M. A., Rivers, S. M. & Wahls, R. A. Development of a common research model for applied CFD validation studies. In *26th AIAA Applied Aerodynamics Conference* 6919 (AIAA, 2008).
11. Bell, J. H. Pressure-sensitive paint measurements on the NASA Common Research Model in the NASA 11-ft transonic wind tunnel. In *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition* 1128 (AIAA, 2011).
12. Kennedy, G. & Martins, J. A comparison of metallic and composite aircraft wings using aerostructural design optimization. In *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference* 5475 (AIAA, 2012).

13. Rao, J., Kiran, S., Kamesh, J., Padmanabhan, M. & Chandra, S. Topology optimization of aircraft wing. *J. Aerosp. Sci. Technol.* 61, 402–414 (2009).
14. Chin, T. & Kennedy, G. Large-scale compliance-minimization and buckling topology optimization of the undeformed Common Research Model wing. In 57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference 0939 (AIAA, 2016).
15. Stanford, B. & Dunning, P. Optimal topology of aircraft rib and spar structures under aeroelastic loads. *J. Aircr.* 52, 1298–1311 (2015).
16. Wolff, J. *The Law of Bone Remodeling* (Springer, 1986) [transl.]
17. Hesthaven, J., Rozza, G. & Stamm, B. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations* (Springer, 2015).
18. Groen, J. P. & Sigmund, O. Homogenization-based topology optimization for high-resolution manufacturable microstructures. *Int. J. Num. Methods Eng.* <https://doi.org/10.1002/nme.5575> (2017).
19. Alexandersen, J. & Lazarov, B. Topology optimisation of manufacturable microstructural details without length scale separation using a spectral coarse basis preconditioner. *Comput. Methods Appl. Mech. Eng.* 290, 156–182 (2015).
20. Bourdin, B. Filters in topology optimization. *Int. J. Numer. Methods Eng.* 50, 2143–2158 (2001).
21. Bruns, T. E. & Tortorelli, D. A. Topology optimization of non-linear elastic structures and compliant mechanisms. *Comput. Methods Appl. Mech. Eng.* 190, 3443–3459 (2001).
22. Tortorelli, D. A. & Michaleris, P. Design sensitivity analysis: overview and review. *Inverse Probl. Eng.* 1, 71–105 (1994).
23. Sigmund, O. A 99 line topology optimization code written in Matlab. *Struct. Multidiscipl. Optim.* 21, 120–127 (2001).
24. Evgrafov, A., Rupp, C. J., Maute, K. & Dunn, M. L. Large-scale parallel topology optimization using a dual-primal substructuring solver. *Struct. Multidiscipl. Optim.* 36, 329–345 (2008).

25. Aage, N., Andreassen, E. & Lazarov, B. S. Topology optimization using PETSc: an easy-to-use, fully parallel, open source topology optimization framework. *Struct. Multidiscipl. Optim.* 51, 565–572 (2015).
26. Amir, O., Aage, N. & Lazarov, B. S. On multigrid-CG for efficient topology optimization. *Struct. Multidiscipl. Optim.* 49, 815–829 (2014).
27. Balay, S. et al. PETSc Users Manual. Technical Report No. ANL-95/11, Revision 3.6, (Argonne National Laboratory, 2015).
28. Saad, Y. *Iterative Methods for Sparse Linear Systems* (SIAM, 2003).
29. Kenway, G. K. W., Kennedy, G. J. & Martins, J. R. R. A. Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations. *AIAA J.* 52, 935–951 (2014).
30. Helms, H. & Lambrecht, U. The potential contribution of light-weighting to reduce transport energy consumption. *Int. J. Life Cycle Assess.* 12, 58–64 (2007).