International Journal of analytical and Numerical Methods in Mechanical Design
Vol 5., Issue 1, Winter & Spring 2026, 1-30

Surrogate-Assisted Multi-Objective Optimization of a Case-Specific Mandibular Advancement Device for Obstructive Sleep Apnea Using Integrated FEA and CFD: A Computational Proof of Concept

Mohammad Mahdi Giahi [1, 2], Pedram Tehrani [1, *]

[1] Department of Biomedical Engineering, CT.C., Islamic Azad University, Tehran, Iran;

[2] Mechatronics Research Laboratory, Dept. of Computer and Electrical Engineering, Qazvin Islamic Azad University, Qazvin, Iran;

Received 26 January 2026, Accepted 20 February 2026

Abstract

Obstructive sleep apnea (OSA) is commonly treated using mandibular advancement devices (MADs), but designs that improve upper-airway patency may also increase appliance loading and alter force transfer to dentoalveolar structures. This study developed a case-specific computational framework for MAD design and multi-objective optimization by integrating parametric computer-aided design, finite element analysis (FEA), computational fluid dynamics (CFD), design-of-experiments sampling, surrogate modeling, and Pareto optimization.

A three-dimensional anatomical model was reconstructed from a publicly available cone-beam computed tomography volume. A two-piece MAD was parameterized using six design variables: mandibular advancement, vertical opening, guiding-block thickness, splint thickness, sagittal mandibular rotation, and material elastic modulus. A total of 120 geometrically feasible configurations were evaluated using FEA and CFD. Separate XGBoost and feedforward artificial neural-network surrogate models were trained for the simulated biomechanical and aerodynamic responses. The validated models for maximum MAD von Mises stress, airway pressure drop, and minimum cross-sectional airway area were coupled with the non-dominated sorting genetic algorithm II (NSGA-II). Tooth stress, periodontal-ligament strain, tooth displacement, and MAD deformation were retained as secondary outputs and did not affect Pareto ranking.

The selected surrogate models achieved external-test coefficients of determination ranging from 0.91 to 0.96. The optimization identified 53 unique Pareto-optimal configurations. The selected compromise design comprised a mandibular advancement of 6.4 mm, vertical opening of 3.2 mm, guiding-block thickness of 4.1 mm, splint thickness of 2.6 mm, sagittal rotation of 10.8°, and material elastic modulus of 1250 MPa. Relative to a central reference configuration, the compromise design reduced maximum MAD von Mises stress by 19.4%, airway pressure drop by 16.3%, and 99th-percentile airflow velocity by 6.3%, while increasing the minimum cross-sectional airway area by 8.3%. High-fidelity verification yielded surrogate-prediction errors of 1.8%, 1.9%, and 1.3% for MAD stress, pressure drop, and minimum cross-sectional airway area, respectively.

The proposed framework enables computationally efficient screening of competing appliance-level structural and aerodynamic objectives before physical prototyping. However, the findings represent a single-case, anatomy-based computational proof of concept. Because the selected public CBCT case was not clinically characterized for OSA and lacked polysomnographic and treatment-response data, the results should not be interpreted as patient-specific predictions of disease severity, therapeutic response, or clinical outcome.

Keywords: Obstructive sleep apnea; Mandibular advancement device; Case-specific modeling; Finite element analysis; Computational fluid dynamics; Surrogate modeling; Multi-objective optimization; Machine learning.

1. Introduction

Obstructive sleep apnea (OSA) is a prevalent sleep-related breathing disorder characterized by recurrent episodes of partial or complete upper-airway obstruction during sleep. These events can produce intermittent hypoxemia, repeated cortical arousals, sleep fragmentation, intrathoracic pressure fluctuations, and sustained sympathetic activation. It has been estimated that approximately 936 million adults aged 30–69 years worldwide have at least mild OSA, indicating a substantial and frequently underdiagnosed global health burden [1]. The pathophysiology of OSA is multifactorial and reflects interactions among craniofacial anatomy, pharyngeal soft-tissue collapsibility, ventilatory-control instability, arousal threshold, and impaired responsiveness of the upper-airway dilator muscles [2–4]. When inadequately treated, OSA is associated with excessive daytime sleepiness, impaired cognitive performance, reduced quality of life, and increased risks of hypertension, cardiac arrhythmia, coronary artery disease, stroke, and metabolic dysfunction [2,3,5]. Positive airway pressure therapy, particularly continuous positive airway pressure (CPAP), remains the most effective non-invasive treatment for moderate-to-severe OSA because it pneumatically stabilizes the upper airway and prevents recurrent pharyngeal collapse [6]. Its long-term clinical effectiveness, however, may be limited by inadequate adherence related to mask discomfort, pressure

* Corresponding Author: P.tehrani@iauctb.ac.ir

intolerance, nasal symptoms, air leakage, noise, and difficulty maintaining nightly use. Clinical guidelines therefore recommend custom and titratable oral appliances for adults who are intolerant of CPAP or prefer an alternative treatment [7]. Mandibular advancement devices (MADs) are the most widely used oral appliances for this purpose. They maintain the mandible in an anterior position, thereby displacing the tongue and associated soft tissues, enlarging the pharyngeal airway, and reducing upper-airway collapsibility [7,8]. Meta-analyses have shown that both CPAP and MAD therapy can reduce the apnea–hypopnea index and daytime sleepiness, although CPAP generally produces a greater reduction in respiratory events [9,10]. The greater acceptability and nightly use of MADs may nevertheless contribute to favorable real-world effectiveness in appropriately selected patients. Despite their established therapeutic value, MADs are not universally effective, and treatment response varies considerably among individuals. Increasing mandibular advancement does not produce an exclusively beneficial response. Greater advancement may improve airway patency but may also increase mechanical loading within the appliance and alter force transfer to the teeth and periodontal tissues. Long-term oral-appliance therapy has been associated with reductions in overjet and overbite, retroclination of the maxillary incisors, and proclination of the mandibular incisors [11,12]. These effects emphasize the need to control both the magnitude and distribution of appliance-induced forces. Finite element analysis (FEA) studies have demonstrated that MAD geometry, advancement mechanism, and load-transfer configuration can substantially affect stress and deformation within the appliance, dentition, and periodontal structures [13,14]. MAD design should therefore be treated as a multi-objective engineering problem in which improved airway performance is balanced against appliance structural loading and dentoalveolar response. Medical imaging has provided important insights into the anatomical mechanisms of MAD therapy. Magnetic resonance imaging has shown that mandibular advancement can increase upper-airway volume, particularly within the velopharyngeal region, while also displacing adjacent soft tissues [15]. Dynamic imaging during sleep has demonstrated increases in retropalatal and retroglossal dimensions and reductions in airway collapsibility following mandibular advancement [16]. Cone-beam computed tomography studies have similarly reported increases in pharyngeal-airway volume, minimum cross-sectional area, and transverse and anteroposterior airway dimensions after MAD treatment [17]. However, the magnitude and spatial distribution of airway enlargement depend on appliance design and may differ between treatment responders and non-responders [18,19]. Anatomical enlargement alone therefore cannot fully explain therapeutic outcome, and the functional aerodynamic consequences of the resulting geometric changes must also be considered. Computational fluid dynamics (CFD) provides a non-invasive method for quantifying airflow behavior within reconstructed upper-airway geometries. CFD-derived variables, including airflow velocity, static-pressure distribution, pressure gradient, airway resistance, wall shear stress, and flow separation, provide functional information that cannot be obtained from geometric measurements alone. Previous functional-imaging studies have shown that CFD-predicted reductions in airway resistance may be associated with successful MAD treatment [20]. Subsequent case-specific simulations suggested that changes in airway pressure gradient may correlate more strongly with therapeutic response than geometric enlargement alone [21]. Fluid–structure interaction studies have further indicated that mandibular advancement may reduce airway-wall deformation and susceptibility to inspiratory occlusion [22]. Other CFD investigations have reported reduced pressure loss and improved airflow distribution following simulated or clinically applied mandibular advancement [23]. Nevertheless, systematic reviews have identified substantial methodological heterogeneity in airway reconstruction, boundary conditions, turbulence modeling, wall assumptions, and outcome assessment [24]. Prospective evidence indicates that functional imaging combining computed tomography and CFD may improve the identification of likely treatment responders, although further methodological standardization and clinical validation remain necessary [25]. A continuing limitation of the existing literature is that structural and aerodynamic responses are commonly investigated separately. FEA studies typically focus on appliance stress and dentoalveolar loading, whereas CFD studies primarily examine airway geometry and airflow behavior. This separation makes it difficult to determine whether a configuration that improves airway performance simultaneously produces unfavorable appliance loading or dentoalveolar response. In addition, therapeutic response is influenced not only by airway anatomy but also by physiological characteristics such as baseline airway collapsibility and ventilatory loop gain [26]. Predictive models based on polysomnographic airflow morphology and clinical variables have shown that likely responders and non-responders can be distinguished with clinically meaningful accuracy [27]. These findings support the development of data-informed computational methods while also demonstrating that anatomical modeling alone cannot establish clinical treatment effectiveness. Simultaneous exploration of geometric, structural, aerodynamic, and material variables requires a large number of high-fidelity simulations. Relevant MAD design parameters may include mandibular advancement, vertical opening, splint thickness, guiding-block or connector geometry, mandibular rotation, and material stiffness. Exhaustive simulation of every possible parameter combination would be computationally prohibitive. Surrogate modeling offers an efficient alternative by approximating the relationships between design variables and high-fidelity simulation responses using a limited but systematically sampled design-of-experiments dataset [28,29]. Extreme gradient boosting can represent nonlinear interactions, incorporate regularization, and provide estimates of variable importance [52]. Feedforward artificial neural networks can also approximate complex nonlinear input–output relationships when their architecture and training procedures are appropriately controlled [53]. Following training and independent validation, surrogate models can be coupled with evolutionary multi-objective algorithms such as the non-dominated sorting genetic algorithm II (NSGA-II), which identifies a diverse set of Pareto-optimal solutions without collapsing competing objectives into a single arbitrarily weighted function [55]. Accordingly, the present study develops an integrated, case-specific and anatomy-informed computational framework for the design and multi-objective optimization of a mandibular advancement device intended for OSA-related applications. The anatomical geometry was reconstructed from a single publicly available CBCT case from the ToothFairy2 dataset [31–33]. Because the selected case lacked clinical confirmation of OSA and associated polysomnographic, demographic, protrusive-capacity, and treatment-response information, it was used exclusively as an anatomical basis for computational proof-of-concept development. FEA was used to quantify maximum MAD von Mises stress, MAD deformation, tooth stress, tooth displacement, and periodontal-ligament strain, whereas CFD was used to evaluate airway pressure drop, 99th-percentile airflow velocity, airway resistance, and minimum cross-sectional airway area. The formal multi-objective optimization was restricted to maximum MAD von Mises stress, airway pressure drop, and minimum cross-sectional airway area. Tooth stress, tooth displacement, periodontal-ligament strain, and MAD deformation were retained as secondary comparative responses and did not affect Pareto ranking. The validated surrogate models for the three optimization objectives were subsequently integrated with NSGA-II to identify configurations representing trade-offs between appliance structural loading and upper-airway performance. The principal contribution of this study is the integration of case-specific anatomical reconstruction, parametric MAD design, appliance structural analysis, upper-airway aerodynamics, surrogate modeling, and multi-objective optimization within a unified computational framework. The proposed approach is intended to support efficient screening of candidate MAD configurations before physical prototyping and experimental testing. It should be interpreted as an anatomy-informed computational proof of concept rather than a patient-specific clinical treatment-planning system or a predictor of therapeutic outcome.

2. Materials and Methods

The present study developed an integrated, case-specific and anatomy-informed computational framework for the design, evaluation, and multi-objective optimization of a mandibular advancement device (MAD) intended for obstructive-sleep-apnea-related applications. The workflow combined cone-beam computed tomography (CBCT)-based anatomical reconstruction, parametric computer-aided design (CAD), finite element analysis (FEA), computational fluid dynamics (CFD), design-of-experiments (DOE) sampling, machine-learning-based surrogate modeling, and evolutionary multi-objective optimization within a unified digital framework. Because the selected public CBCT case was not clinically characterized for OSA, the framework was developed as a computational proof of concept rather than as a patient-specific clinical treatment-planning system.

The framework was designed to evaluate the appliance-level structural performance, dentoalveolar response, and aerodynamic performance of alternative MAD configurations within the investigated case-specific anatomy. FEA was used to quantify maximum MAD von Mises stress, MAD deformation, tooth stress, tooth displacement, and periodontal-ligament strain. In parallel, CFD simulations were conducted to characterize upper-airway airflow using airway pressure drop, 99th-percentile airflow velocity, airway resistance, minimum cross-sectional airway area, and pharyngeal turbulence intensity. The numerical inputs and corresponding FEA- and CFD-derived outputs were organized into a structured DOE dataset and used to train and validate separate surrogate regression models for the investigated responses.

The formal multi-objective optimization was restricted to three responses: maximum MAD von Mises stress, airway pressure drops, and minimum cross-sectional airway area. Tooth stress, periodontal-ligament strain, tooth displacement, MAD deformation, airflow velocity, airway resistance, and turbulence intensity were retained as secondary comparative outputs and did not influence Pareto ranking. The validated surrogate models for the three optimization objectives were subsequently integrated with the non-dominated sorting genetic algorithm II (NSGA-II) to identify Pareto-optimal MAD configurations representing trade-offs between reduced appliance structural loading and improved upper-airway performance. A compromise configuration was then selected from the Pareto front and compared with a central reference MAD configuration.

The study was entirely computational. No physical prototype fabrication, experimental mechanical testing, therapeutic intervention, or clinical validation was performed within the scope of the present investigation.

2.1 Study Design Overview

The overall methodological framework is presented in Figure 1. The proposed workflow integrated case-specific anatomical reconstruction, parametric MAD design, FEA, CFD, DOE dataset generation, surrogate-model development, multi-objective optimization, and high-fidelity numerical verification.

The workflow comprised five sequential stages. In the first stage, a publicly available CBCT volume was processed to reconstruct case-specific anatomical models of the maxilla, mandible, dentition, and upper airway. These geometries provided the anatomical basis for appliance design and numerical analysis but were not associated with a clinically confirmed diagnosis of OSA.

In the second stage, a parametric CAD model of a two-piece MAD was generated using six geometric and material design variables: mandibular advancement distance, vertical opening, guiding-block thickness, splint thickness, sagittal mandibular rotation, and MAD elastic modulus. Automated geometric-feasibility checks were applied to exclude configurations with component interpenetration, inadequate local thickness, discontinuous guiding-surface engagement, or invalid solid geometry.

In the third stage, each geometrically feasible configuration was evaluated using separate high-fidelity structural and aerodynamic analyses. FEA was performed to quantify maximum MAD von Mises stress, MAD deformation, tooth stress, tooth displacement, and periodontal-ligament strain. CFD was conducted to evaluate airway pressure drop, 99th-percentile airflow velocity, airway resistance, minimum cross-sectional airway area, and pharyngeal turbulence intensity.

In the fourth stage, the six input design variables and their corresponding FEA- and CFD-derived responses were assembled into a structured DOE dataset. Separate surrogate regression models were trained, tuned, and validated for the individual output responses. Model development included repeated cross-validation and evaluation using an external testing set that was not used for model training.

In the fifth stage, the validated surrogate models for maximum MAD von Mises stress, airway pressure drop, and minimum cross-sectional airway area were integrated with NSGA-II. The optimization procedure sought to minimize maximum MAD stress and airway pressure drop while maximizing the minimum cross-sectional airway area. Tooth stress, periodontal-ligament strain, tooth displacement, and the remaining simulated responses were not included in the objective vector and therefore did not affect non-dominated sorting. A compromise configuration was selected from the resulting Pareto front using its normalized distance from the ideal objective point. The selected configuration was subsequently regenerated and verified using the original high-fidelity FEA and CFD workflows before comparison with the central reference MAD configuration.

The study was limited to computational modeling and simulation-driven optimization. No prototype fabrication, experimental validation, therapeutic intervention, or prospective clinical evaluation was undertaken.

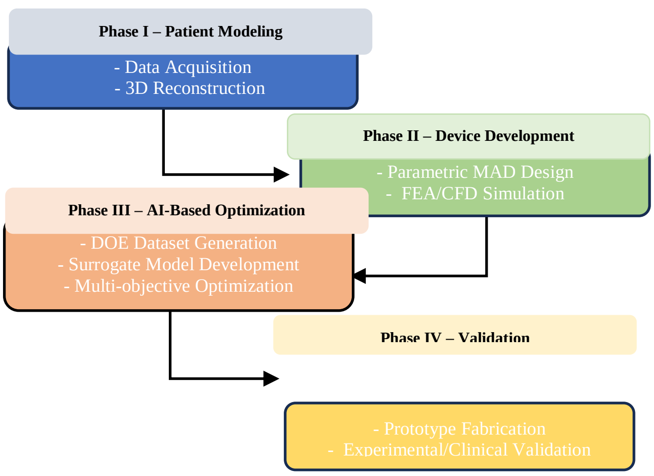

Fig 1. Schematic overview of the case-specific, anatomy-informed computational workflow, including public CBCT-based anatomical reconstruction, parametric MAD design, FEA and CFD simulations, DOE dataset generation, surrogate-model development, NSGA-II-based multi-objective optimization, compromise-solution selection, and high-fidelity numerical verification.

2.2 Patient-Specific CBCT Data Acquisition

A de-identified cone-beam computed tomography (CBCT) volume was obtained from the publicly available ToothFairy2 dataset, which was developed for three-dimensional multi-structure segmentation of maxillofacial anatomy. The publicly released training dataset comprises 480 CBCT volumes with expert-derived voxel-level annotations of 42 anatomical and dental structures, including the maxilla, mandible, individual teeth, inferior alveolar canals, maxillary sinuses, and pharyngeal airway [31–33].

One volume, identified as case F039, was selected as the anatomical basis of the present computational proof-of-concept study. Cases from the F-series were considered because they provided a comparatively broad field of view and more complete visualization of the maxillary dentition than the original P-series volumes. Case F039 was selected after visual review confirmed adequate coverage of the maxillary and mandibular dental arches, sufficient remaining dentition for appliance retention, inclusion of the pharyngeal region, and absence of severe image artifacts that would prevent anatomical reconstruction. The stored image matrix of the selected volume was 298 × 512 × 512 voxels.

According to the dataset documentation, the source images were acquired using a NewTom/NTVGiMK4 CBCT scanner at a tube voltage of 110 kV and a tube current of 3 mA. The selected volume had an isotropic voxel spacing of 0.30 × 0.30 × 0.30 mm^3 and was distributed in MetaImage (.mha) format. The image-intensity values were retained in the form provided by the public dataset and were not used to derive patient-specific material properties.

The CBCT volume and its corresponding multi-class segmentation map were imported into 3D Slicer, version 5.12.2, for image inspection, anatomical-label review, and three-dimensional reconstruction. The provided labels for the maxilla, mandible, individual teeth, and pharyngeal airway were examined in the axial, coronal, and sagittal planes and were simultaneously reviewed using three-dimensional rendering. Local manual corrections were performed where necessary to remove isolated voxels, close small discontinuities, eliminate non-anatomical disconnected regions, and preserve the continuity of the reconstructed structures. The corrected segmentations were subsequently converted into three-dimensional surface models and exported in stereolithography (STL) format for downstream CAD processing and numerical analysis.

Because ToothFairy2 was developed as an anatomical-segmentation dataset rather than as a clinically characterized OSA cohort, no polysomnographic, demographic, symptom-related, protrusive-capacity, or treatment-response information was available for case F039. In particular, the apnea–hypopnea index, oxygen-saturation profile, body mass index, age, sleep-stage information, maximum voluntary mandibular protrusion, and therapeutic outcome were unavailable. Case F039 was therefore used solely as an anatomical basis for computational framework development. The case was not assumed to have a clinically confirmed diagnosis of OSA, and the resulting simulations were not interpreted as patient-specific predictions of disease severity, treatment response, or clinical outcome.

The reconstructed maxillary and mandibular dental geometries were used as the anatomical basis for parametric MAD construction and structural analysis, whereas the segmented pharyngeal-airway lumen was used to define the baseline CFD domain. Because paired images acquired before and after mandibular advancement were unavailable, advancement-specific airway geometries could not be reconstructed directly from medical imaging. They were instead generated using the controlled geometric-morphing procedure described in Section 2.6.1.

No new participants were recruited, imaged, treated, or prospectively followed in the present investigation. The study was restricted to secondary computational analysis of de-identified, publicly available data distributed under the license specified by the ToothFairy2 dataset. The original data collection received approval from the Comitato Etico dell'Area Vasta Emilia Nord under approval number 1374/2020/OSS/ESTMO, SIRER ID 1275–NAICBCT-D.

2.3 Anatomical Segmentation and Three-Dimensional Reconstruction

The CBCT volume and corresponding multi-class segmentation map of case F039 were imported into 3D Slicer, version 5.12.2, for anatomical review, label refinement, and three-dimensional reconstruction [34,35]. Because the ToothFairy2 dataset provides expert-derived voxel-level annotations, de novo threshold-based segmentation was not performed. Instead, the available labels for the maxilla, mandible, individual teeth, and pharyngeal airway were used as the initial anatomical segmentations [31–33].

To standardize the anatomical orientation before geometric measurement and numerical modeling, the CBCT volume and segmentation map were rigidly reoriented together while preserving their spatial registration. A case-specific Cartesian coordinate system was established by aligning the midsagittal plane with the x–z plane and the maxillary occlusal plane with the x–y plane. The positive x-axis was defined in the anterior direction, the positive y-axis toward the subject's left side, and the positive z-axis in the superior direction. This coordinate system was maintained throughout the subsequent CAD, FEA, airway-morphing, and CFD procedures.

The anatomical labels were reviewed in the axial, coronal, and sagittal planes and simultaneously inspected using three-dimensional rendering. The maxillary and mandibular masks were examined for discontinuities, incomplete external boundaries, and isolated components. Individual tooth labels were inspected for merged structures, incorrectly assigned regions, disconnected components, and discontinuities near the crowns and roots. Local corrections were performed using the Paint and Erase tools in the Segment Editor module. Disconnected regions smaller than 100 voxels were removed using the Islands tool when they were not anatomically continuous with the structure of interest.

The pharyngeal-airway label was reviewed separately to confirm continuity of the air-filled lumen and preservation of localized anatomical constrictions. Particular attention was given to the retropalatal and retroglossal regions because alterations in these regions could substantially affect the calculated minimum cross-sectional airway area and airflow resistance. Isolated voxels, artificial side branches, and non-anatomical disconnected regions were removed, whereas genuine localized constrictions were preserved. The cranial and caudal limits of the airway domain were trimmed at complete lumen cross-sections using planes approximately perpendicular to the local airway centerline. The corrected binary airway mask was retained as the reference geometry for calculation of airway volume and minimum cross-sectional airway area.

Following label refinement, closed-surface representations were generated from the corrected binary segmentations. Weak surface smoothing, using a closed-surface conversion smoothing factor of 0.10, was applied to reduce voxel-related stair-step artifacts while limiting alteration of the underlying anatomical boundaries. Surface decimation was disabled to preserve dental detail and avoid artificial enlargement or distortion of narrow airway regions. For the pharyngeal-airway model, the processed surface was accepted only when the calculated airway volume and minimum cross-sectional airway area differed by less than 2% from the corresponding values obtained from the corrected binary label map.

The resulting surface models were inspected for holes, disconnected components, self-intersections, non-manifold edges, and other defects that could prevent CAD processing or numerical meshing. Local defects were corrected to obtain continuous and watertight geometries. The maxilla, mandible, individual teeth, and pharyngeal-airway lumen were then exported separately in stereolithography (STL) format. The maxillary and mandibular dental geometries were subsequently used for parametric MAD construction and structural analysis, whereas the reconstructed pharyngeal-airway lumen was used as the baseline CFD domain.

All reported quantitative anatomical measurements were calculated from the corrected, unsmoothed binary label maps. The weakly smoothed surface models were used exclusively for visualization, CAD processing, geometric morphing, and numerical mesh generation. This separation was maintained to minimize the influence of surface-processing operations on the reported airway measurements.

2.4 Parametric MAD Design

A case-specific mandibular advancement device (MAD) was developed using a parametric computer-aided design workflow based on the reconstructed maxillary and mandibular dental geometries. The anatomical surface models were imported into FreeCAD, version 1.1.1 (FreeCAD Project Association), and the appliance geometry was generated using the built-in Python scripting interface. The automated workflow enabled systematic modification of the predefined geometric and material design variables and batch generation of MAD configurations for finite element analysis, design-of-experiments sampling, surrogate-model development, and multi-objective optimization.

The MAD was modeled as a custom bimaxillary two-piece appliance comprising a maxillary splint, a mandibular splint, and bilateral posterior guiding blocks. The internal fitting surfaces of the splints were derived from the corresponding dental-crown geometries. An outward normal offset was then applied to generate the external splint surfaces and define a uniform nominal splint thickness, t_s, while preserving the case-specific occlusal morphology. Boolean subtraction of the dental surfaces from the offset splint volumes was used to create the internal tooth-contact cavities.

The maxillary splint remained fixed relative to the maxillary dentition throughout the parametric modeling procedure. The mandibular splint was rigidly associated with the mandibular dental arch and repositioned through a prescribed rigid-body transformation defined by the mandibular advancement distance, vertical opening, and sagittal rotation. This geometric repositioning was performed before numerical analysis and did not represent deformation of the mandibular splint or dentition.

The bilateral guiding blocks were positioned within the posterior occlusal regions and aligned with the prescribed mandibular advancement trajectory. Their opposing load-transfer surfaces were designed to maintain bilateral engagement between the maxillary and mandibular appliance components. The minimum guiding-block thickness, t_g, was measured normal to the corresponding load-transfer surface and was treated as an independent geometric design variable. The same prescribed thickness was applied to the right and left guiding blocks to maintain bilateral geometric symmetry.

2.4.1. Mandibular Repositioning

The anatomical coordinate system established during image reconstruction was used for all geometric transformations. The positive x-axis was directed anteriorly, the positive y-axis toward the subject's left side, and the positive z-axis superiorly. Accordingly, anterior mandibular advancement was represented by a positive translation along the x-axis, whereas vertical bite opening was represented by an inferior translation along the negative z-axis. The geometric centers of the right and left mandibular condylar regions were identified from the reconstructed mandibular geometry. The midpoint between these centers, denoted by c, was used as the reference point for sagittal mandibular rotation. Because the reconstructed anatomy had been reoriented such that the transverse direction corresponded to the global y-axis, sagittal rotation was approximated as rotation about an axis parallel to the y-axis and passing through c. For an original point p belonging to the mandibular dentition, mandibular splint, or mandibular-associated guiding geometry, the transformed position p′ was calculated as:

$$\mathbf{p}' = \mathbf{c} + \mathbf{R}_y(\theta)(\mathbf{p} - \mathbf{c}) + \begin{bmatrix} A \\ 0 \\ -V \end{bmatrix}, \qquad (1)$$

where A is the anterior mandibular advancement distance, V is the vertical opening, and θ is the sagittal rotation angle of the mandibular component. The rotation matrix about the transverse y-axis was defined as:

$$\mathbf{R}_y(\theta) = \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix}. \qquad (2)$$

The sign convention for θ was selected such that positive values produced inferior rotation of the anterior mandibular region. Applying the rotation about the intercondylar reference point, rather than about the global coordinate origin, provided a more anatomically interpretable approximation of sagittal mandibular repositioning. The resulting transformation defined the initial geometric position of the mandibular components before structural analysis and was not reapplied as a displacement boundary condition during FEA.

2.4.2. Parametric Design Variables

The combined geometric and material design vector was defined as:

$$\mathbf{x} = [A, V, t_c, t_s, \theta, E_m], \qquad (3)$$

where A is the mandibular advancement distance, V is the vertical opening, t_c is the minimum guiding-block thickness, t_s is the nominal splint thickness, θ is the sagittal mandibular rotation angle, and E_m is the effective Young's modulus assigned to the MAD material. The design-variable ranges were selected to provide a sufficiently broad computational domain while excluding configurations that were clearly incompatible with the reconstructed anatomy or appliance geometry. Mandibular advancement was varied from 3 to 8 mm, whereas vertical opening was varied from 2 to 6 mm. These values were treated as computational design bounds rather than patient-specific therapeutic prescriptions because maximum voluntary mandibular protrusion, clinical bite-registration data, and patient-specific comfort limits were unavailable for the public CBCT case. In clinical practice, mandibular advancement is normally titrated according to individual protrusive capacity, comfort, and therapeutic response [36–38]. The minimum guiding-block thickness was varied from 2 to 5 mm, and the nominal splint thickness was varied from 1.5 to 3.0 mm. The sagittal mandibular rotation angle ranged from 5° to 15°. The effective elastic modulus of the MAD material was varied from 900 to 1800 MPa as a computational sensitivity parameter for assessing the influence of appliance stiffness on structural response. This modulus range represented a numerical design variable and was not intended to characterize a single clinically validated polymer.

Table 1
Parametric variables used to generate the MAD design space

Design variable	Symbol	Range	Unit
Mandibular advancement	A	3–8	mm
Vertical opening	V	2–6	mm
Guiding-block thickness	t_c	2–5	mm
Splint thickness	t_s	1.5–3.0	mm
Sagittal mandibular rotation	θ	5–15	degree
MAD elastic modulus	E_m	900–1800	MPa

The feasible computational design domain was expressed as:

$$\Omega = \left\{ \mathbf{x} \in \mathbb{R}^6 \mid \begin{array}{cl} 3 & \leq A \leq 8, \\ 2 & \leq V \leq 6, \\ 2 & \leq t_c \leq 5, \\ 1.5 & \leq t_s \leq 3.0, \\ 5^\circ & \leq \theta \leq 15^\circ, \\ 900 & \leq E_m \leq 1800 \end{array} \right\}. \qquad (4)$$

2.4.3. Automated Model Generation and Feasibility Assessment

Automated Model Generation and Geometric-Feasibility Assessment A custom Python script was developed within Free CAD to update the mandibular rigid-body transformation, splint thickness, guiding-block thickness, MAD elastic modulus, and configuration metadata for each sampled design vector. The maxillary dentition and maxillary appliance component remained fixed, whereas the mandibular dentition, mandibular splint, and mandibular-associated guiding geometry were repositioned according to Equation (1). Each configuration was regenerated automatically and assigned a unique identification code for subsequent geometric processing and numerical analysis. Each generated configuration was subjected to an automated geometric-feasibility assessment before high-fidelity simulation. A design was considered geometrically feasible when it satisfied all of the following requirements: preservation of the predefined dental-retention regions; absence of self-intersections within the maxillary and mandibular splints; absence of unintended interpenetration between opposing appliance components; preservation of the prescribed minimum splint and guiding-block thicknesses; continuous bilateral engagement between the opposing guiding surfaces; and successful generation of closed, valid, and watertight solid geometries suitable for numerical meshing. Configurations that violated one or more of these requirements were excluded before FEA and CFD processing. Geometric feasibility was treated as a prerequisite for inclusion in the numerical dataset and not as an optimization objective. The geometrically valid MAD models were exported as boundary-representation solids in STEP (.step) format for structural preprocessing and finite element mesh generation. At this stage, the prescribed rigid-body transformation was applied only to the mandibular anatomy and its associated appliance geometry. No direct deformation of the pharyngeal airway was inferred from this rigid mandibular transformation. Advancement-specific airway geometries were generated separately using the controlled surface-morphing procedure described in Section 2.6.1.

2.5 Finite Element Analysis

Finite element analysis was performed to evaluate the initial quasi-static structural response of the mandibular advancement device and the associated dentoalveolar structures under a standardized comparative loading condition. The geometrically feasible configurations generated in Free CAD were imported into ANSYS Mechanical 2025 R1 (ANSYS, Inc., Canonsburg, PA, USA) as STEP solids. A quasi-static nonlinear structural analysis was conducted for each configuration. The analysis nonlinearity arose primarily from the tooth–splint and guiding-block contact interactions, whereas the assigned material models remained linear elastic. The extracted responses included MAD stress and deformation, tooth stress and displacement, periodontal-ligament strain, and contact-related load transfer between the appliance and dentition.

2.5.1. Finite Element Geometry

The finite element assembly comprised the maxilla, mandible, individual teeth, periodontal-ligament layers, maxillary splint, mandibular splint, and bilateral posterior guiding blocks. Because the ToothFairy2 segmentation labels did not distinguish enamel from dentin, each tooth was represented as a homogeneous dentin-equivalent solid. Similarly, the maxillary and mandibular bones were modeled as homogeneous cortical-equivalent structures because separate cortical and cancellous bone segmentations were unavailable. Periodontal-ligament geometries were not directly available in the public dataset and were therefore generated computationally. A uniform layer with a nominal thickness of 0.30 mm was constructed around the root surface of each tooth, extending from the approximate cementoenamel junction to the root apex. Boolean operations were used to create the corresponding alveolar socket and position the PDL layer between the tooth root and supporting jawbone. This uniform-thickness representation was adopted as a controlled modeling approximation because case-specific PDL morphology could not be resolved from the CBCT data [39,42]. The mandibular dentition, mandibular splint, and mandibular-associated guiding geometry were positioned according to the advancement, vertical-opening, and sagittal-rotation parameters defined in Section 2.4. The rigid-body repositioning transformation was completed before structural loading was applied. The prescribed advancement and vertical-opening displacements were therefore not reapplied as boundary conditions during FEA, thereby avoiding double application of the mandibular repositioning.

2.5.2. Material Properties and Constitutive Assumptions

All anatomical and appliance components were modeled as homogeneous and isotropic materials under a small-strain assumption. The teeth, jawbones, and MAD polymer were assigned linear-elastic behavior. The PDL was represented using a nearly incompressible linear-elastic approximation because the analysis focused on the initial quasi-static mechanical response rather than nonlinear, time-dependent biological tooth movement. The simplified PDL model reduced the computational cost of the 120-configuration high-fidelity simulation campaign. However, experimentally reported PDL properties vary substantially with loading rate, tissue location, constitutive formulation, and experimental method. The calculated PDL strain was therefore interpreted as a comparative numerical response rather than as an absolute indicator of tissue injury or long-term dental movement [42]. For an isotropic linear-elastic material, the constitutive relationship was expressed as:

$$\boldsymbol{\sigma} = \lambda\,\mathrm{tr}(\boldsymbol{\varepsilon})\,\mathbf{I} + 2\mu\boldsymbol{\varepsilon}, \qquad (5)$$

where $\boldsymbol{\sigma}$is the Cauchy stress tensor, $\boldsymbol{\varepsilon}$is the infinitesimal strain tensor, $\mathbf{I}$is the identity tensor, and $\lambda$and $\mu$are the Lamé constants:

$$\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}, \mu = \frac{E}{2(1+\nu)}, \qquad (6)$$

where $E$is Young's modulus and $\nu$is Poisson's ratio. The infinitesimal strain tensor was calculated from the displacement field $\mathbf{u}$ as:

$$\boldsymbol{\varepsilon} = \frac{1}{2}[\nabla\mathbf{u} + (\nabla\mathbf{u})^T]. \qquad (7)$$

The assigned material properties are summarized in Table 2.

Table 2. Material properties assigned in the finite element model.

Table 2
Material properties assigned in the finite element model.

Component	Young's modulus	Poisson's ratio	Material representation
Tooth	18.3 GPa	0.31	Linear elastic
Maxillary and mandibular bone	13.7 GPa	0.30	Linear elastic
Periodontal ligament	0.68 MPa	0.45	Nearly incompressible linear elastic
MAD polymer	900–1800 MPa	0.35	Linear elastic

The effective elastic modulus of the MAD polymer, E_m, was varied according to the sampled design vector, whereas the remaining material properties were held constant throughout the DOE simulation campaign. The investigated modulus range was treated as a computational stiffness variable rather than as a complete material characterization of a specific commercial polymer.

2.5.3. Contact Definitions

Bonded contact was assigned between each tooth root and its surrounding PDL layer and between the PDL and the corresponding alveolar socket. These bonded interfaces prevented relative sliding or separation and represented continuous attachment between the dental and periodontal structures. Contact between the internal splint surfaces and the dental crowns was modeled using frictional contact with a Coulomb coefficient of friction of 0.20. This coefficient was treated as a controlled numerical assumption because case-specific tooth–appliance friction measurements were unavailable. Contact between the opposing guiding-block surfaces was modeled as frictionless, permitting relative tangential sliding while transmitting compressive normal forces. Separation was allowed when the local normal contact pressure decreased to zero. The augmented-Lagrange formulation was used for the tooth–splint contact interfaces, whereas the guiding-block interfaces were solved using the program-controlled frictionless-contact formulation. Initial geometric interference and penetration were evaluated before solution. Configurations with unresolved initial penetration, unstable contact behavior, or persistent contact-related non-convergence were excluded from the high-fidelity simulation dataset rather than being assigned estimated response values.

2.5.4. Boundary and Loading Conditions

The superior and posterior regions of the maxilla were constrained to approximate attachment to the cranial base and prevent rigid-body motion of the maxillary assembly. Remote-displacement constraints were applied to the right and left condylar regions of the mandible. Translational motion of the condylar reference points was restricted, whereas rotational degrees of freedom were left unconstrained to avoid imposing fully rigid rotational fixation of the mandible. The MAD and mandibular dentition were analyzed in the repositioned configuration defined by the corresponding design vector. A standardized total closing load of 100 N was applied bilaterally to the mandibular ramus and angle regions, with 50 N applied on each side in the superior direction. The load was distributed over anatomically defined surface regions rather than applied to isolated nodes to reduce artificial local stress concentration. The 100-N load represented a comparative reference loading condition applied consistently to all MAD configurations. It was not interpreted as a patient-specific measurement of masticatory force, muscular activity, or soft-tissue restoring force. Previous MAD investigations have shown that mandibular restoring forces may vary with advancement and individual anatomy; however, such forces could not be measured for the public CBCT case used in this study [40]. Identical loading was therefore imposed across the design space to isolate the effects of appliance geometry, mandibular position, and material stiffness. Neglecting body inertia under the quasi-static assumption, structural equilibrium was expressed as:

$$\nabla \cdot \boldsymbol{\sigma} + \mathbf{b} = \mathbf{0}, \qquad (8)$$

where $\mathbf{b}$ is the body-force vector, which was neglected in the present analysis. Automatic load stepping was used to apply the prescribed load progressively and improve convergence of the nonlinear contact solution.

2.5.5. Mesh Generation and Convergence Assessment

All structural components were discretized using second-order, ten-node tetrahedral elements. Local mesh refinement was applied to regions expected to exhibit relatively steep stress or strain gradients, including the PDL layers, tooth–splint contact surfaces, guiding-block contact interfaces, and guiding-block–splint junctions. A target element size of 0.15 mm was assigned to the PDL layers, providing approximately two elements across their nominal 0.30-mm thickness. A target element size of 0.40 mm was applied to the tooth–splint and guiding-block contact regions, whereas a target size of 1.0 mm was used within the bulk jawbone and appliance volumes. Curvature- and proximity-based refinement was enabled to preserve thin geometric features, narrow contact regions, and localized changes in surface curvature. Mesh convergence was evaluated using a representative intermediate-advancement configuration selected from the feasible design space. Coarse, intermediate, and fine meshes were generated through progressive reduction of the global and local element sizes. Convergence was evaluated using maximum MAD von Mises stress, maximum principal PDL strain, and maximum tooth displacement. The relative change in maximum MAD von Mises stress between consecutive mesh levels was calculated as:

$$\delta_\sigma = \left| \frac{\sigma_{\mathrm{vM}}^{(k)} - \sigma_{\mathrm{vM}}^{(k-1)}}{\sigma_{\mathrm{vM}}^{(k-1)}} \right| \times 100\%, \qquad (9)$$

whereas the relative changes in maximum principal PDL strain and maximum tooth displacement were calculated as:

$$\begin{aligned} \delta_\varepsilon &= \left| \frac{\varepsilon_{\mathrm{PDL}}^{(k)} - \varepsilon_{\mathrm{PDL}}^{(k-1)}}{\varepsilon_{\mathrm{PDL}}^{(k-1)}} \right| \times 100\%, \\ \delta_d &= \left| \frac{d_{\mathrm{tooth}}^{(k)} - d_{\mathrm{tooth}}^{(k-1)}}{d_{\mathrm{tooth}}^{(k-1)}} \right| \times 100\%. \end{aligned} \qquad (10)$$

Here, k denotes the current mesh level, σ_{vM} is the maximum MAD von Mises stress, $\varepsilon_{\mathrm{PDL}}$ is the maximum principal PDL strain, and $d_{\mathrm{tooth}}$is the maximum tooth displacement. The mesh was considered sufficiently converged when the relative changes in all three indicators between the intermediate and fine meshes were below 5%. The resulting intermediate-mesh settings were subsequently applied to the complete DOE simulation campaign.

2.5.6. Biomechanical Outcome Measures

The von Mises equivalent stress was calculated as:

$$\sigma_{\mathrm{vM}} = \sqrt{\frac{3}{2}\mathbf{s}:\mathbf{s}}, \qquad (11)$$

where $\mathbf{s}$is the deviatoric stress tensor. The primary biomechanical outputs extracted from each simulation were:

- maximum von Mises stress in the MAD;
- maximum von Mises stress in the teeth;
- maximum principal strain in the PDL;
- maximum total tooth displacement; and
- maximum total deformation of the MAD.

Peak contact pressure at the tooth–splint and guiding-block interfaces was recorded as a secondary diagnostic variable. Stress values were nodally averaged, and isolated singular values located at sharp geometric edges or constraint boundaries were excluded from the optimization dataset. Peak contact pressure at the tooth–splint and guiding-block interfaces was recorded as a secondary diagnostic variable. Stress values were nodally averaged, and isolated singular values located at sharp geometric edges or constraint boundaries were excluded from the surrogate-model dataset. The primary FEA outputs were paired with their corresponding parametric design vectors and used for surrogate-model training. Of these outputs, only maximum MAD von Mises stress was transferred to the multi-objective optimization framework as a formal structural objective. Tooth stress, periodontal-ligament strain, tooth displacement, MAD deformation, and contact pressure were retained for secondary comparative interpretation and did not affect Pareto ranking.

2.6 Computational Fluid Dynamics

Computational fluid dynamics (CFD) simulations were performed to quantify airflow behavior within the reconstructed pharyngeal airway and to evaluate the aerodynamic consequences of the geometrically simulated mandibular configurations. All simulations were conducted using ANSYS Fluent 2025 R1 (ANSYS, Inc., Canonsburg, PA, USA) in double-precision mode. The analyses represented steady inspiratory airflow through rigid advancement-specific airway geometries. The CFD responses extracted from each configuration included airway pressure drop, airway resistance, 99th-percentile airflow velocity, minimum cross-sectional airway area, and pharyngeal turbulence intensity. Airway pressure drop and minimum cross-sectional airway area were transferred to the formal multi-objective optimization framework. The remaining aerodynamic responses were retained for secondary comparative assessment and did not influence Pareto ranking.

2.6.1. Generation of Advancement-Specific Airway Geometries

Because the ToothFairy2 dataset contained only one static CBCT acquisition and did not include paired images acquired before and after mandibular advancement, advancement-specific airway geometries could not be reconstructed directly from medical imaging. A controlled surface-morphing procedure was therefore used to generate an approximate airway geometry for each mandibular configuration. The baseline pharyngeal-airway surface reconstructed in Section 2.3 was used as the reference geometry. The airway region extending approximately from the inferior border of the soft palate to the superior border of the epiglottis was defined as the deformable pharyngeal segment. Control points distributed over the anterior wall of this segment were used as geometric surrogates for the tongue–airway interface. For each design configuration, the prescribed displacement of these control points was derived from the difference between their original and transformed positions under the mandibular transformation defined in Equation (1). Control points positioned on the posterior and lateral pharyngeal walls were assigned zero displacement to prevent global translation of the airway domain. In addition, circumferential surface bands extending 5 mm from the inlet and outlet planes were fixed to preserve the shape and orientation of the computational boundaries. The displacement of the remaining airway-surface nodes was interpolated using a thin-plate-spline radial basis function (RBF). The interpolated displacement field was expressed as:

$$\mathbf{d}(\mathbf{x}) = \sum_{j=1}^{N_c} \boldsymbol{\alpha}_j \, \phi\left(\| \mathbf{x} - \mathbf{x}_j \|\right) + \mathbf{B}\mathbf{x} + \mathbf{c} \qquad (12)$$

where $\mathbf{x}_j$ denotes the position of the j-th control point, N_c is the total number of control points, $\boldsymbol{\alpha}_j$, $\mathbf{B}$, and $\mathbf{c}$ are interpolation coefficients, and the thin-plate-spline basis function was defined as:

$$\phi(r) = r^2 \ln(r + \epsilon), \qquad (13)$$

where ϵ is a small numerical constant introduced to avoid evaluation of $\ln(0)$. RBF interpolation permits smooth transfer of prescribed boundary displacements while preserving continuity of complex three-dimensional surfaces [43].

A custom Python routine was used to apply the control-point displacements, interpolate the airway-surface deformation, and export the resulting geometry for CFD preprocessing. Each morphed airway was checked for surface self-intersection, non-manifold edges, disconnected regions, negative or collapsed cross-sections, and unintended deformation of the inlet and outlet boundaries. Geometries that failed one or more quality checks were excluded before CFD meshing.

The morphing procedure was intended to provide a controlled geometric approximation of the effect of mandibular advancement on the anterior pharyngeal boundary. It did not represent patient-specific soft-tissue mechanics, neuromuscular activation, or fluid–structure interaction. Accordingly, the generated geometries were used for comparative design assessment and were not interpreted as direct predictions of airway deformation during sleep.

2.6.2. Governing Equations and Airflow Model

Air was modeled as an incompressible Newtonian fluid under steady-state inspiratory conditions. The incompressibility assumption was considered appropriate because the simulated airflow remained within the low-Mach-number regime. The Reynolds-averaged conservation equation for mass was expressed as:

$$\nabla \cdot \mathbf{u} = 0, \qquad (14)$$

and the Reynolds-averaged momentum equation was written as:

$$\rho(\mathbf{u} \cdot \nabla)\mathbf{u} = -\nabla p + \nabla \cdot [(\mu + \mu_t)(\nabla \mathbf{u} + \nabla \mathbf{u}^T)], \qquad (15)$$

where $\mathbf{u}$ is the time-averaged airflow-velocity vector, p is the static pressure, ρ is the air density, μ is the molecular dynamic viscosity, and μ_t is the turbulent eddy viscosity. The air density and dynamic viscosity were specified as:

$$\rho = 1.225 \text{ kg m}^{-3}, \mu = 1.789 \times 10^{-5} \text{ Pa s}. \qquad (16)$$

The shear-stress-transport k-ω model was employed to close the Reynolds-averaged equations. This model was selected because pharyngeal narrowing can produce adverse pressure gradients, flow acceleration, separation, and locally transitional or turbulent flow. Previous upper-airway CFD studies have similarly employed k-ω-based models to investigate airflow under OSA and mandibular-advancement conditions [44–47].

2.6.3. Computational Domain and Boundary Conditions

The anatomical computational domain extended from the nasopharyngeal inlet plane to the inferior hypopharyngeal outlet plane. Straight inlet and outlet extensions were added to reduce the local influence of the imposed boundary conditions on airflow within the anatomical airway. The inlet-extension length was equal to three times the local inlet hydraulic diameter, whereas the outlet-extension length was equal to five times the local outlet hydraulic diameter. A constant inspiratory volumetric flow rate of:

$$Q = 20 \text{ L min}^{-1} = 3.333 \times 10^{-4} \text{ m}^3 \text{ s}^{-1}$$

was prescribed at the inlet. A uniform velocity normal to the inlet plane was calculated from the prescribed volumetric flow rate and the inlet cross-sectional area of each configuration. An inlet turbulence intensity of 1% and a turbulence length scale equal to 7% of the inlet hydraulic diameter were specified. A zero-gauge static-pressure condition was imposed at the outlet. All anatomical airway walls were modeled as stationary, impermeable, rigid, and no-slip surfaces. The rigid-wall assumption isolated the aerodynamic consequences of the generated geometric differences but did not account for inspiratory deformation or collapse of compliant pharyngeal tissues. Because wall compliance can substantially affect upper-airway narrowing and occlusion, this assumption was treated as an important limitation of the present computational proof of concept [45].

Table 3. Boundary conditions and physical assumptions used in the CFD simulations.

Table 3
Boundary conditions and physical assumptions used in the CFD simulations

Parameter	Setting
CFD software	ANSYS Fluent 2025 R1
Numerical precision	Double precision
Flow regime	Steady-state inspiration
Fluid model	Incompressible Newtonian air
Turbulence model	SST k-ω
Inlet flow rate	20 L/min
Inlet turbulence intensity	1%
Turbulence length scale	$0.07D_h$
Outlet condition	Zero-gauge static pressure
Airway-wall condition	Stationary, rigid, and no-slip
Air density	1.225 kg m^{-3}
Dynamic viscosity	1.789×10^{-5} Pa s

2.6.4. Mesh Generation and Independence Assessment

Each Advancement-specific Airway Geometry was discretized using ANSYS Fluent Meshing. An unstructured tetrahedral volume mesh was generated together with prism layers adjacent to the airway walls. Local surface and volume refinement was applied within anatomically narrow regions, particularly the retropalatal, retroglossal, and oropharyngeal segments. A nominal surface-element size of 0.60 mm was used over the airway geometry. The maximum local surface size was reduced to 0.25 mm in the region surrounding the minimum cross-sectional airway area and in regions with relatively high surface curvature. A maximum volume-element size of 1.0 mm was applied within the airway core. Ten prism layers were generated adjacent to the airway walls using a layer-growth ratio of 1.20. The near-wall mesh was designed to target a wall-normal dimensionless distance of $y^+ < 1$over the principal pharyngeal surfaces, thereby permitting direct resolution of the near-wall region with the SST k–ω model. Meshes were accepted only when no negative-volume cells were present, maximum cell skewness was below 0.85, and minimum orthogonal quality exceeded 0.15. Mesh independence was assessed using a representative intermediate-advancement airway geometry. Coarse, intermediate, and fine meshes were generated by progressively reducing the global and local element sizes.

The evaluated indicators were airway pressure drop and 99th-percentile airflow velocity. The relative change in airway pressure drop between consecutive mesh levels was calculated as:

$$\delta_{\Delta P} = \left| \frac{\Delta P^{(k)} - \Delta P^{(k-1)}}{\Delta P^{(k-1)}} \right| \times 100\%, \qquad (17)$$

and the relative change in 99th-percentile airflow velocity was calculated as:

$$\delta_V = \left| \frac{V_{99}^{(k)} - V_{99}^{(k-1)}}{V_{99}^{(k-1)}} \right| \times 100\%. \qquad (18)$$

Here, k denotes the current mesh level. The mesh was considered sufficiently independent when the changes in both indicators between the intermediate and fine meshes were below 5%. The resulting intermediate-mesh settings were subsequently applied to all geometrically valid CFD configurations.

2.6.5. Numerical Solution and Convergence Criteria

A steady pressure-based coupled solver was employed. Spatial gradients were evaluated using the least-squares cell-based method. Second-order pressure interpolation was used for the pressure equation, and second-order upwind discretization was applied to momentum, turbulent kinetic energy, and specific dissipation rate.
The solution was initialized using the hybrid-initialization procedure. Simulations were continued until all of the following criteria were satisfied:

1. scaled residuals below $10^{-5}$for continuity, momentum, k, and ω;
2. an inlet-to-outlet mass-flow imbalance below 0.1%;
3. a change of less than 0.1% in airway pressure drop over the final 200 iterations; and
4. a change of less than 0.1% in peak core-flow velocity over the final 200 iterations.

A maximum of 2000 iterations was permitted for each configuration. Simulations that did not satisfy the convergence criteria were inspected for poor-quality cells, boundary-condition sensitivity, and unstable recirculation regions before being remeshed and repeated. The pressure-based coupled algorithm and second-order discretization schemes were selected to improve pressure–velocity coupling and spatial accuracy in the complex airway geometries.

2.6.6. Aerodynamic Outcome Measures

The airway pressure drop was calculated from area-weighted mean static pressures at the inlet and outlet:

$$\Delta P = \bar{P}_{\text{inlet}} - \bar{P}_{\text{outlet}}. \qquad (19)$$

Airway resistance was calculated as:

$$R_{\text{airway}} = \frac{\Delta P}{Q}, \qquad (20)$$

where Q is the prescribed volumetric flow rate.
The Reynolds number was evaluated at the inlet and minimum cross-sectional region as:

$$Re = \frac{\rho U D_h}{\mu}, \qquad (21)$$

where U is the area-averaged velocity and the hydraulic diameter was calculated as:

$$D_h = \frac{4A}{P_w}, \qquad (22)$$

where A is the cross-sectional area and P_w is the wetted perimeter. The local turbulence intensity was calculated from turbulent kinetic energy as:

$$I = \frac{\sqrt{2k/3}}{|\,\mathbf{u}\,|} \times 100\%. \qquad (23)$$

The minimum cross-sectional area was determined geometrically by extracting planes normal to the airway centerline at 0.5-mm intervals and identifying the smallest enclosed area. To reduce sensitivity to isolated numerical outliers, the peak core-flow velocity was represented by the 99th percentile of the cell-centered velocity magnitude, denoted by V_{99}, rather than the absolute maximum value.
The primary aerodynamic outputs used for surrogate-model development were:

- airway pressure drops, ΔP;
- airway resistance, R_{airway};
- peak core-flow velocity, V_{99};
- minimum cross-sectional area, MCA; and
- volume-averaged pharyngeal turbulence intensity.

Airway volume, Reynolds number, and wall shear stress were recorded as secondary diagnostic parameters. The CFD outputs were paired with their corresponding parametric design vectors and incorporated into the DOE dataset for surrogate-model training and multi-objective optimization.

2.7 DOE Dataset Generation

A design-of-experiments (DOE) strategy was used to construct a space-filling numerical dataset for surrogate-model development. The DOE covered the six-dimensional bounded computational domain defined in Section 2.4, comprising mandibular advancement distance, vertical opening, guiding-block thickness, splint thickness, sagittal mandibular rotation, and MAD effective elastic modulus. For the i-th candidate configuration, the design vector was defined as:

$$\mathbf{x}_i = [A_i, V_i, t_{c,i}, t_{s,i}, \theta_i, E_{m,i}], \qquad (24)$$

where A_i is the mandibular advancement distance, V_i is the vertical opening, $t_{c,i}$ is the guiding-block thickness, $t_{s,i}$ is the splint thickness, θ_i is the sagittal mandibular rotation angle, and $E_{m,i}$ is the elastic modulus assigned to the MAD material.

2.7.1. Space-Filling Sampling Strategy

A scrambled Latin hypercube sampling (LHS) procedure was used to generate a candidate pool, from which the final configurations were selected using a sequential maximin criterion. LHS was selected because it provides stratified coverage of each input variable while reducing clustering within a multidimensional design space [49,50]. The sampling procedure was implemented in Python using a fixed random seed of 42 to ensure reproducibility. For sampling purposes, each design variable was represented in a normalized coordinate system:

$$\xi_j^{(i)} = \frac{x_j^{(i)} - x_{j,\mathrm{L}}}{x_{j,\mathrm{U}} - x_{j,\mathrm{L}}}, 0 \le \xi_j^{(i)} \le 1, \qquad (25)$$

where $x_j^{(i)}$ is the physical value of the j-th design variable in the i-th candidate configuration, and $x_{j,\mathrm{L}}$ and $x_{j,\mathrm{U}}$ denote its lower and upper bounds, respectively. The sampled normalized coordinates were mapped back to physical units according to:

$$x_j^{(i)} = x_{j,\mathrm{L}} + \xi_j^{(i)}\left(x_{j,\mathrm{U}} - x_{j,\mathrm{L}}\right). \qquad (26)$$

Candidate configurations were generated in successive LHS batches and subjected to the automated geometric-feasibility assessment described in Section 2.4.3. Sampling was continued until a pool of at least 300 geometrically feasible candidates had been obtained. Configurations exhibiting self-intersection, unintended component interpenetration, insufficient local thickness, discontinuous guiding-block engagement, non-watertight geometry, or failure of the airway-morphing quality checks were excluded at this stage. From the feasible candidate pool, 120 configurations were selected using a sequential maximin criterion. The first configuration was chosen as the candidate closest to the center of the normalized design domain. Each subsequent configuration was selected by maximizing its minimum Euclidean distance from the previously selected samples:

$$\mathbf{x}_k^* = \underset{\mathbf{x} \in \mathcal{C} \setminus \mathcal{S}_{k-1}}{\arg\max} \; \min_{\mathbf{z} \in \mathcal{S}_{k-1}} \|\boldsymbol{\xi}(\mathbf{x}) - \boldsymbol{\xi}(\mathbf{z})\|_2 \qquad (27)$$

where $\mathcal{C}$ is the feasible candidate pool, $\mathcal{S}_{k-1}$ is the set of previously selected configurations, and $\boldsymbol{\xi}$ denotes the normalized design vector.
The final sample size of 120 configurations was defined as a prespecified high-fidelity computational budget rather than as definitive evidence of surrogate-model sufficiency. Its adequacy was subsequently assessed using cross-validation and learning-curve analysis during surrogate-model development.

2.7.2. High-Fidelity Simulation Campaign

Each selected configuration was assigned a unique identification code and processed through the automated modeling pipeline. The parametric MAD geometry was regenerated in Free CAD, after which the corresponding structural model and advancement-specific airway geometry were prepared for FEA and CFD, respectively.

A configuration was included in the final DOE dataset only when both high-fidelity analyses satisfied their predefined quality requirements. The FEA solution was required to satisfy the nonlinear contact-convergence criteria and mesh-quality requirements described in Section 2.5. The corresponding CFD solution was required to satisfy the residual, mass-balance, monitor-stability, and mesh-quality criteria specified in Section 2.6.

Configurations that failed because of unrecoverable geometric defects, unstable contact behavior, invalid airway morphing, poor CFD mesh quality, or persistent solver non-convergence were not assigned estimated or imputed output values. Instead, the failed configuration was removed and replaced by the next highest-ranked feasible candidate according to the maximin selection criterion. This process was continued until complete FEA and CFD results had been obtained for all 120 configurations.

2.7.3. Input and Output Dataset Structure

For each successfully simulated configuration, the selected biomechanical and aerodynamic responses were assembled into the output vector:

$$\mathbf{y}_i = \left[\mathbf{y}_i^{\text{FEA}}, \mathbf{y}_i^{\text{CFD}}\right], \qquad (28)$$

where

$$\mathbf{y}_i^{\text{FEA}} = \left[\sigma_{\text{MAD},i}, \sigma_{\text{tooth},i}, \varepsilon_{\text{PDL},i}, d_{\text{tooth},i}, d_{\text{MAD},i}\right] \qquad (29)$$

and

$$\mathbf{y}_i^{\text{CFD}} = \left[\Delta P_i, V_{99,i}, R_{\text{airway},i}, MCA_i, \bar{I}_{\text{pharynx},i}\right] \qquad (30)$$

where:

- σ_{MAD} is the maximum von Mises stress in the MAD;
- σ_{tooth} is the maximum von Mises stress in the teeth;
- ε_{PDL} is the maximum principal strain in the periodontal ligament;
- d_{tooth} is the maximum tooth displacement;
- d_{MAD} is the maximum MAD deformation;
- ΔP is the airway pressure drop;
- V_{99} is the 99th-percentile core-flow velocity;
- R_{airway} is the airway resistance;
- MCA is the minimum cross-sectional airway area; and
- $\bar{I}_{\text{pharynx}}$ is the volume-averaged turbulence intensity within the pharyngeal region.
-

Because the same volumetric flow rate was prescribed for all CFD simulations, airway resistance was calculated directly as $R_{\text{airway}} = \Delta P/Q$. It was retained as a derived comparative output but was not treated as an independent surrogate-model target or optimization objective.

The complete high-fidelity dataset was expressed as:

$$\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{N}, N = 120 \qquad (31)$$

The corresponding input and output matrices were defined as:

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \vdots \\ \mathbf{x}_N^T \end{bmatrix} \in \mathbb{R}^{N\times 6}, \mathbf{Y} = \begin{bmatrix} \mathbf{y}_1^T \\ \mathbf{y}_2^T \\ \vdots \\ \mathbf{y}_N^T \end{bmatrix} \in \mathbb{R}^{N\times 10}. \qquad (32)$$

The input variables were stored in their original physical units, and the output variables were retained without manual removal of extreme but numerically valid observations. A result was excluded only when a documented geometric, meshing, or solver failure had occurred. This approach prevented physically meaningful high-response configurations from being incorrectly treated as statistical outliers.

Before surrogate-model fitting, the dataset was inspected for missing values, duplicate design vectors, non-finite outputs, unit inconsistencies, and violations of the predefined physical bounds. Feature scaling and any output transformations were performed only after partitioning the dataset and were fitted exclusively to the training data within each validation procedure, as described in Section 2.8. This strategy prevented information from the testing data from influencing model preprocessing.

2.8 Machine-Learning-Based Surrogate Model Development

An artificial-intelligence-based surrogate modeling framework was developed to approximate the nonlinear relationships between the parametric MAD design variables and the biomechanical and aerodynamic responses obtained from the high-fidelity FEA and CFD simulations. The surrogate models were intended to provide rapid performance predictions throughout the feasible design space and thereby reduce the computational cost of multi-objective optimization.

The DOE dataset described in Section 2.7 contained $N = 120$ successfully simulated configurations, six input variables, and ten output responses. The input and output matrices were expressed as:

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \vdots \\ \mathbf{x}_N^T \end{bmatrix} \in \mathbb{R}^{N\times 6} \qquad (33)$$

and

$$\mathbf{Y} = \begin{bmatrix} \mathbf{y}_1^T \\ \mathbf{y}_2^T \\ \vdots \\ \mathbf{y}_N^T \end{bmatrix} \in \mathbb{R}^{N\times 10} \qquad (34)$$

where the input vector $\mathbf{x}_i$ and the FEA- and CFD-derived output vectors were defined in Equations (24) and (28)–(30), respectively.

The dataset contained ten recorded output quantities. However, airway resistance was calculated deterministically from airway pressure drop and the prescribed constant flow rate and was therefore not modeled as an independent surrogate target. Separate single-output regression models were trained for the remaining nine biomechanical and aerodynamic responses. This approach avoided scale competition among outputs, enabled independent hyperparameter optimization, and allowed the most accurate algorithm to be selected separately for each predicted quantity. The general surrogate relationship for the q-th response was written as:

$$\hat{y}_{i,q} = \hat{f}_q(\mathbf{x}_i), q = 1, \ldots, 9 \qquad (35)$$

where $\hat{f}_q$ denotes the trained surrogate function and $\hat{y}_{i,q}$ is the predicted response.

2.8.1. Dataset Partitioning and Preprocessing

Twenty-four configurations, corresponding to 20% of the complete dataset, were reserved as an external testing set. The test configurations were selected using a maximin criterion in the normalized six-dimensional input space to ensure that the testing set covered both the central and boundary regions of the design domain. The remaining 96 configurations were used exclusively for model training and hyperparameter optimization. The external testing set was not examined during preprocessing, hyperparameter optimization, algorithm selection, or final-model selection.

Input variables were scaled to the interval $[0, 1]$ using min–max normalization:

$$\tilde{x}_j^{(i)} = \frac{x_j^{(i)} - x_{j,\min}^{\text{train}}}{x_{j,\max}^{\text{train}} - x_{j,\min}^{\text{train}}}, \qquad (36)$$

where the minimum and maximum values were calculated exclusively from the training data. The same fitted transformation was subsequently applied to the external testing set. For ANN training, each output variable was additionally standardized using the mean and standard deviation of the corresponding training response. ANN predictions were transformed back to their original physical units before performance evaluation.

All preprocessing operations were performed within the training folds during cross-validation. Consequently, no information from the validation or external testing samples was used to calculate the scaling parameters, tune the model hyperparameters, or select the final surrogate model.

2.8.2. XGBoost Surrogate Models

Extreme gradient boosting was evaluated as a nonlinear tree-based surrogate algorithm [52]. For each output response, the XGBoost prediction was represented as an additive ensemble of regression trees:

$$\hat{y}_{i,q} = \sum_{k=1}^{K_q} f_{k,q}(\tilde{\mathbf{x}}_i), \qquad (37)$$

where K_q is the number of trees used for the q-th response model and $f_{k,q}$ is the prediction produced by the k-th regression tree.

The regularized training objective was defined as

$$\mathcal{L}_q = \sum_{i=1}^{N_{\text{train}}} l\left(y_{i,q}, \hat{y}_{i,q}\right) + \sum_{k=1}^{K_q} \Omega\left(f_{k,q}\right), \qquad (38)$$

where $l(\cdot)$ is the squared-error loss and $\Omega(\cdot)$ penalizes excessive tree complexity. The principal XGBoost hyperparameters included the number of trees, maximum tree depth, learning rate, minimum child weight, row subsampling ratio, feature subsampling ratio, and L_1 and L_2 regularization coefficients. The evaluated search ranges were 200–1000 trees, a maximum depth of 2–6, a learning rate of 0.01–0.10, subsampling and feature-subsampling ratios of 0.70–1.00, a minimum child weight of 1–10, an L_1 coefficient of 0–1, and an L_2 coefficient of 0.1–10.

2.8.3. Artificial Neural Network Surrogate Models

A fully connected feedforward artificial neural network was evaluated as the second surrogate algorithm [53]. The transformation performed by the l-th hidden layer was expressed as:

$$\mathbf{h}^{(l)} = \phi\left(\mathbf{W}^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}\right), \qquad (39)$$

where $\mathbf{h}^{(l)}$ is the hidden-layer output, $\mathbf{W}^{(l)}$ and $\mathbf{b}^{(l)}$ are the trainable weight matrix and bias vector, and $\phi(\cdot)$ is the rectified linear unit activation function. A single neuron with linear activation was used in the output layer because separate networks were trained for each continuous response.

The ANN search space comprised one to three hidden layers, 16–128 neurons per layer, dropout ratios of 0–0.20, L_2 regularization coefficients of 10^{-6}–10^{-3}, and learning rates of 10^{-4}–10^{-2}. The networks were trained using the Adam optimizer with mini-batches of 8–32 samples. Training was permitted for a maximum of 1000 epochs, with early stopping applied when the validation loss did not improve for 50 consecutive epochs. The network weights corresponding to the minimum validation loss were restored before external testing.

2.8.4. Hyperparameter Optimization and Internal Validation

Hyperparameter optimization was performed separately for each output variable and algorithm using repeated five-fold cross-validation on the 96 training configurations. Five repetitions were conducted using different fold assignments, producing 25 validation evaluations for each candidate hyperparameter set. The mean cross-validated root mean square error was used as the optimization criterion.

A fixed random seed of 42 was used for dataset partitioning, cross-validation, model initialization, and hyperparameter search. Learning curves were generated using progressively larger proportions of the training dataset to assess whether model performance had stabilized or remained strongly limited by the number of high-fidelity simulations.

The hyperparameters selected during cross-validation were used to retrain each model on the complete 96-configuration training set. The resulting model was then evaluated once on the untouched 24-configuration external testing set.

Table 4
Main settings used for surrogate-model development

Item	Setting
Number of configurations	120
Number of input variables	6
Number of recorded outputs	10
Number of independent surrogate targets	9
Training set	96 configurations
External testing set	24 configurations
Test-set selection	Maximin coverage in normalized input space
Surrogate algorithms	XGBoost and feedforward ANN
Modeling strategy	Separate model for each output
Internal validation	Repeated five-fold cross-validation
Number of repetitions	5
Hyperparameter criterion	Minimum mean validation RMSE
Input scaling	Training-based min–max normalization
ANN output scaling	Training-based standardization
ANN hidden activation	ReLU
ANN output activation	Linear
ANN optimizer	Adam
Maximum ANN epochs	1000
Early-stopping patience	50 epochs
Random seed	42
Test performance metrics	R^2, RMSE, and MAE

2.8.5. Predictive Performance Assessment

The predictive accuracy of each surrogate model was evaluated in the original physical units of the corresponding response. The coefficient of determination was calculated as:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}, \qquad (40)$$

where $\bar{y}$is the mean of the observed high-fidelity responses in the evaluated dataset.

The root mean square error was calculated as:

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}, \qquad (41)$$

and the mean absolute error was calculated as:

$$MAE = \frac{1}{n}\sum_{i=1}^{n} \mid y_i - \hat{y}_i \mid . \qquad (42)$$

Here, $y_i$is the high-fidelity simulation response, $\hat{y}_i$is the corresponding surrogate prediction, and $n$is the number of evaluated samples. The external test results were additionally examined using predicted-versus-simulated plots and residual plots to identify systematic bias, heteroscedasticity, and disproportionately large errors near the boundaries of the design domain.

For each output response, the final surrogate algorithm and hyperparameters were selected exclusively using the mean repeated-cross-validation RMSE obtained from the 96-configuration training set. The selected model was then retrained using the complete training set and evaluated once on the external testing set. External-test performance was reported as an independent assessment and was not used to revise the selected algorithm or hyperparameters. Consequently, the final surrogate framework could contain XGBoost models for some responses and ANN models for others rather than requiring a single algorithm for all outputs.

2.8.6. Model Interpretation and Optimization Eligibility

After the final surrogate models had been selected and locked, permutation importance was calculated using the external testing set to estimate the sensitivity of each predicted response to the six design variables. For the selected XGBoost models, SHapley Additive exPlanations (SHAP) values were additionally calculated to characterize the magnitude and direction of the variable effects [54]. These interpretation analyses were performed only after model selection and did not influence algorithm or hyperparameter choice.

Only the validated surrogate models for maximum MAD von Mises stress, airway pressure drop, and minimum cross-sectional airway area were transferred to the multi-objective optimization framework. Surrogate models for tooth stress, periodontal-ligament strain, tooth displacement, MAD deformation, 99th-percentile airflow velocity, and pharyngeal turbulence intensity were retained for secondary comparative interpretation and did not define Pareto dominance.

2.9 Multi-Objective Optimization

Multi-objective optimization was performed to identify MAD configurations that reduced appliance-level structural loading while improving upper-airway performance. The validated surrogate models for maximum MAD von Mises stress, airway pressure drop, and minimum cross-sectional airway area were used to evaluate candidate designs without repeating high-fidelity simulations during every optimization iteration. Dentoalveolar responses, including tooth stress, periodontal-ligament strain, and tooth displacement, were not included in the Pareto objective vector and therefore did not influence non-dominated sorting.

The optimization variables comprised mandibular advancement distance, vertical opening, guiding-block thickness, splint thickness, sagittal mandibular rotation, and MAD elastic modulus. All variables were restricted to the computational bounds defined in Table 1. Candidate configurations were additionally required to satisfy the geometric-feasibility criteria described in Section 2.4.3.

2.9.1. Optimization Problem Formulation

The optimization problem included three competing objectives: minimization of maximum MAD von Mises stress, minimization of airway pressure drop, and maximization of minimum cross-sectional airway area. The problem was expressed in minimization form as:

$$\begin{aligned} \underset{\mathbf{x}\in\Omega}{\text{minimize}} \quad & \mathbf{F}(\mathbf{x}) = \left[\hat{\sigma}_{\text{MAD}}(\mathbf{x}), \widehat{\Delta P}(\mathbf{x}), -\widehat{MCA}(\mathbf{x})\right], \\ & \text{subject to } g_r(\mathbf{x}) \le 0, \end{aligned} \tag{43}$$

The symbols $\hat{\sigma}_{\text{MAD}}$, $\widehat{\Delta P}$, and $\widehat{MCA}$ represent the surrogate-predicted MAD stress, airway pressure drop, and minimum cross-sectional airway area, respectively.

This objective formulation represents optimization of appliance structural stress and airway performance rather than direct optimization of dentoalveolar safety. The dentoalveolar responses were evaluated separately and were not used as constraints or objective functions.

Table 5
Objective functions used for multi-objective MAD optimization.

Objective	Symbol	Optimization goal	Interpretation
Maximum MAD von Mises stress	σ_{MAD}	Minimize	Reduce appliance structural loading
Airway pressure drop	ΔP	Minimize	Reduce pressure loss at the prescribed flow rate
Minimum cross-sectional area	MCA	Maximize	Improve airway patency

2.9.2. NSGA-II Implementation

The optimization was performed using the non-dominated sorting genetic algorithm II (NSGA-II) [55], implemented in Python using pymoo, version 0.6.2 [56]. NSGA-II was selected because it can identify a diverse set of non-dominated solutions for nonlinear optimization problems containing multiple competing objectives.

The initial population was generated using Latin hypercube sampling within the bounded design domain. A population size of 200 candidate solutions and an offspring population of 200 were used. Simulated binary crossover was applied with a crossover probability of 0.90 and a distribution index of 15. Polynomial mutation was applied with a mutation probability of $1/6$, corresponding to one mutated variable per candidate on average, and a distribution index of 20.

Each optimization run was continued for 300 generations. Non-dominated sorting was used to rank candidate designs, whereas crowding distance was used to preserve solution diversity along the Pareto front. Candidate designs that violated the geometric-feasibility criteria were assigned a constraint violation and excluded from survival selection.

Ten independent optimization runs were performed using random seeds 42–51 to evaluate the repeatability of the resulting Pareto front. The non-dominated solutions obtained from all runs were combined, duplicate solutions were removed, and a final non-dominated sorting procedure was performed to obtain the pooled Pareto front.

Table 6
Main NSGA-II settings used for multi-objective optimization.

Parameter	Setting
Implementation	pymoo 0.6.2
Population size	200
Offspring population	200
Number of generations	300
Initial sampling	Latin hypercube sampling
Crossover operator	Simulated binary crossover
Crossover probability	0.90
Crossover distribution index	15
Mutation operator	Polynomial mutation
Mutation probability	$1/6$
Mutation distribution index	20
Independent runs	10
Random seeds	42–51
Constraint handling	Feasibility-first selection

2.9.3. Pareto-Front Assessment

A candidate solution was considered Pareto-optimal when no other feasible configuration improved one objective without simultaneously worsening at least one of the remaining objectives. The resulting Pareto front therefore represented the trade-offs among structural loading, airway pressure loss, and airway enlargement.

Optimization convergence was assessed by monitoring the hypervolume indicator and the number of non-dominated solutions over successive generations. The objectives were normalized exclusively for convergence assessment and post-optimization solution selection; the surrogate predictions remained in their original physical units.

For the i-th Pareto solution, the normalized value of the j-th minimization objective was calculated as:

$$\tilde{f}_j^{(i)} = \frac{f_j^{(i)} - f_{j,\min}^P}{f_{j,\max}^P - f_{j,\min}^P}, j = 1,2,3, \quad (44)$$

where $f_{j,\min}^P$ and $f_{j,\max}^P$ are the minimum and maximum values of the j-th objective among the pooled Pareto-optimal solutions. Because the third objective was defined as $-MCA$, lower normalized values represented better performance for all three objectives.

2.9.4. Selection of the Compromise Configuration

A compromise design was selected using its normalized Euclidean distance from the ideal objective point. Equal importance was initially assigned to the three objectives. The distance for the i-th Pareto-optimal solution was calculated as:

$$D_i = \sqrt{\sum_{j=1}^{3} w_j \left(\tilde{f}_j^{(i)}\right)^2}, w_j = \frac{1}{3}, \quad (45)$$

where w_j is the weight assigned to the j-th normalized objective. The ideal point corresponded to the zero vector, representing the lowest stress, lowest pressure drops, and largest MCA within the pooled Pareto set. The configuration with the minimum value of D_i was selected as the balanced optimal design.

In addition to the compromise solution, three extreme Pareto solutions were retained: the design with the lowest MAD stress, the design with the lowest airway pressure drop, and the design with the largest minimum cross-sectional area. These solutions were used to illustrate the trade-offs among the competing optimization objectives.

2.9.5. High-Fidelity Verification

For comparative assessment, a central reference MAD configuration was defined using the midpoint of each design-variable range specified in Table 1. The reference configuration consisted of a mandibular advancement of $A = 5.5$mm, a vertical opening of $V = 4.0$mm, a guiding-block thickness of $t_c = 3.5$mm, a splint thickness of $t_s = 2.25$mm, a sagittal mandibular rotation of $\theta = 10°$, and a MAD elastic modulus of $E_m = 1350$MPa. This configuration was not considered clinically prescribed or mathematically optimized; it was used solely as a central computational reference for assessing the relative performance of the balanced Pareto-optimal design. The reference and optimized configurations were evaluated using identical geometry-processing, FEA, CFD, loading, boundary-condition, meshing, and convergence settings.

In addition to the central reference configuration, five representative Pareto-optimal designs were selected for high-fidelity verification: the compromise solution, the design with the lowest predicted MAD stress, the design with the lowest predicted airway pressure drop, the design with the largest predicted minimum cross-sectional area, and one intermediate solution located near the central region of the Pareto front. Each selected configuration was regenerated using the original parametric CAD workflow and re-evaluated using the high-fidelity FEA and CFD procedures described in Sections 2.5 and 2.6.

For each verified configuration and optimization objective, the relative surrogate-prediction error was calculated as:

$$e_{i,q} = \left| \frac{y_{i,q}^{\mathrm{HF}} - \hat{y}_{i,q}}{y_{i,q}^{\mathrm{HF}}} \right| \times 100\%, \quad (46)$$

where $y_{i,q}^{\mathrm{HF}}$ is the high-fidelity simulation result, $\hat{y}_{i,q}$ is the corresponding surrogate-model prediction, i denotes the selected configuration, and q denotes the optimization objective.

A Pareto-optimal configuration was considered successfully verified when the relative prediction error was below 10% for maximum MAD von Mises stress, airway pressure drop, and minimum cross-sectional area. The verified high-fidelity responses were also examined to determine whether the corresponding designs remained non-dominated after replacement of the surrogate predictions with the simulation results.

When a selected configuration exceeded the 10% error threshold for one or more objectives, its high-fidelity input–output pair was added to the DOE dataset. The affected surrogate models were then retrained using the expanded dataset, and the multi-objective optimization procedure was repeated. This iterative refinement process was continued until the compromise solution satisfied the predefined verification criterion.

The final verified balanced configuration was compared with the central reference MAD configuration to quantify the relative changes in appliance stress, airway pressure drop, airway resistance, minimum cross-sectional area, and 99th-percentile core-flow velocity. Airway resistance was calculated from the verified pressure drop using the constant inspiratory flow rate of $20\ \mathrm{L/min}$. The balanced configuration was additionally compared with the unadvanced baseline airway for aerodynamic assessment only. Because no mechanically loaded appliance existed in the unadvanced anatomical condition, baseline-airway geometry was not used as a reference for MAD stress or deformation.

All reported optimized performance values were based on the final high-fidelity verification results rather than solely on surrogate-model predictions. Accordingly, the balanced design was interpreted as a numerically verified compromise among the predefined biomechanical and aerodynamic objectives, rather than as evidence of clinical effectiveness or guaranteed structural safety.

3. Results

The integrated computational framework was applied to evaluate the biomechanical and aerodynamic performance of the case-specific mandibular advancement device configurations. The results are organized into six subsections: geometric reconstruction and design-space generation, finite element analysis, computational fluid dynamics, surrogate-model performance, multi-objective optimization with high-fidelity verification, and numerical verification and computational robustness.

All results reported in this section were obtained from computational geometry processing, FEA, CFD, machine-learning prediction, and simulation-based optimization. No physical prototype fabrication, experimental mechanical testing, therapeutic intervention, or clinical outcome assessment was conducted within the scope of the present study. Consequently, the reported findings should be interpreted as numerical evidence of relative design performance rather than confirmation of clinical effectiveness.

3.1 Case-Specific Airway Reconstruction and Baseline Geometry

The pharyngeal airway of case F039 was successfully reconstructed from the CBCT volume using the segmentation and surface-processing procedure described in Section 2.3. The final airway model preserved anatomical continuity from the nasopharyngeal inlet to the hypopharyngeal outlet and contained no disconnected regions, non-manifold edges, or unresolved surface openings. The resulting watertight geometry was considered suitable for centerline-based geometric analysis, surface morphing, and CFD mesh generation.

Figure 2 presents the reconstructed baseline airway geometry together with its anatomical orientation. The model exhibited a non-uniform lumen, with localized narrowing in the retropalatal region.

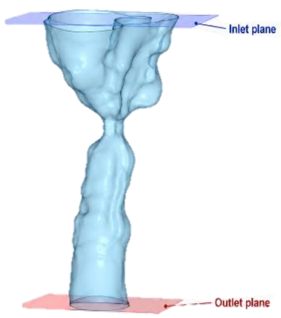

Fig 2. Three-dimensional reconstruction of the case-specific pharyngeal airway obtained from the CBCT volume of case F039. The anatomical orientation and computational inlet and outlet planes are indicated.

Cross-sectional analysis was performed along the airway centerline using planes positioned normal to the local centerline direction at 0.5-mm intervals. The minimum cross-sectional area was identified by comparing the enclosed areas of all extracted sections. Representative cross-sections and the location of the minimum airway area are shown in Figure 3.

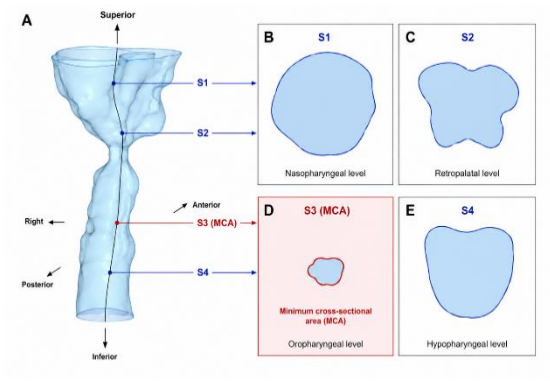

Fig 3. Representative centerline-normal cross-sections of the baseline pharyngeal airway. The section corresponding to the minimum cross-sectional area is highlighted.

The baseline pharyngeal-airway volume was 15,400 mm³, and the minimum cross-sectional area was 72.5 mm². The minimum cross-sectional area was located in the retropalatal region. The baseline geometric characteristics are summarized in Table 7.

Table 7. Baseline geometric characteristics of the case-specific pharyngeal airway.

Table 7
Baseline geometric characteristics of the case-specific pharyngeal airway

Parameter	Baseline value
Pharyngeal-airway volume	15,400 mm³
Minimum cross-sectional area	72.5 mm²
Anatomical location of the MCA	Retropalatal region

The reconstructed baseline airway was subsequently used as the reference geometry for generating advancement-specific airway models through the controlled radial-basis-function morphing procedure described in Section 2.6.1. Following geometric quality control, 120 advancement-specific airway geometries were retained for CFD analysis. All retained geometries were continuous and watertight and contained no self-intersections, collapsed cross-sections, or unintended displacement of the inlet and outlet boundaries.

Because the airway geometry was derived from a single public CBCT case without polysomnographic confirmation of OSA, the reported anatomical measurements represent the baseline geometry of the selected computational case and should not be interpreted as population-level characteristics or diagnostic indicators of obstructive sleep apnea.

3.2 Finite Element Analysis Results

Finite element analysis was successfully completed for all 120 configurations included in the final DOE dataset. The simulations demonstrated systematic variation in appliance stress, tooth stress, periodontal-ligament strain, tooth displacement, and MAD deformation across the investigated design space.

The highest von Mises stresses in the appliance were predominantly concentrated at the junctions between the guiding blocks and the splint bodies, particularly near the posterior load-transfer regions. Secondary stress concentrations were observed around the posterior tooth-contact surfaces and local geometric transitions within the guiding components. Figure 4 presents the stress distribution for the configuration exhibiting a maximum MAD stress closest to the median value of the complete simulation dataset.

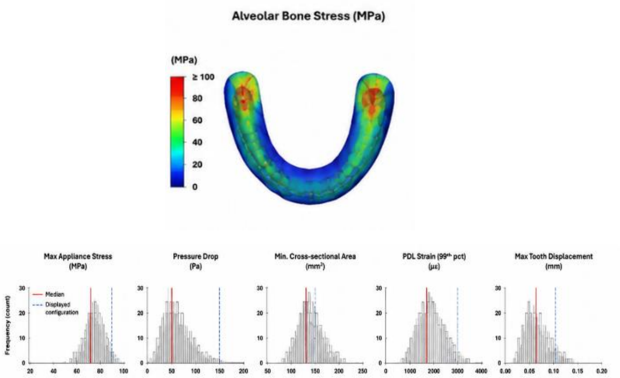

Fig 4. Representative von Mises stress distribution in the MAD and dentoalveolar structures. The displayed configuration had a maximum appliance stress closest to the median value obtained across the 120 simulated designs.

The guiding-block regions acted as the principal load-transfer zones between the maxillary and mandibular components. The highest tooth stresses were generally observed in the posterior teeth engaged by the guiding structures, whereas lower stress levels were distributed across the anterior dentition. Because material failure properties were not incorporated into the constitutive model, the calculated von Mises stresses were interpreted as comparative mechanical-performance indicators rather than direct evidence of appliance safety or failure resistance.

The total-deformation pattern of the representative configuration is shown in Figure 5. Maximum appliance deformation occurred primarily in the mandibular splint and guiding-block regions. The maxillary splint exhibited comparatively smaller displacement because it was constrained by the fixed maxillary dentition. Tooth displacement was concentrated in the posterior anchorage regions and decreased toward the anterior teeth.

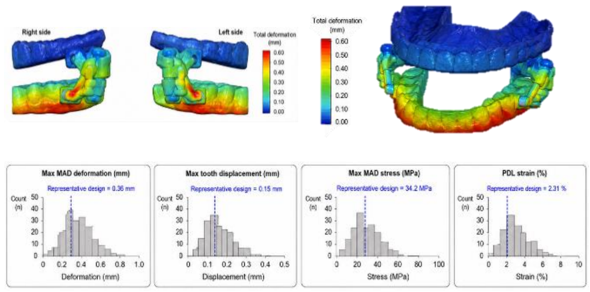

Fig 5. Representative total-deformation distribution in the MAD under the standardized 100-N bilateral closing load.

Across the 120 configurations, the mean maximum von Mises stress in the MAD was 26.4 ± 6.8 MPa, with values ranging from 14.9 to 42.7 MPa. The corresponding maximum tooth stress was 7.8 ± 2.1 MPa. Mean maximum MAD deformation and tooth displacement were 0.287 ± 0.074 mm and 0.118 ± 0.031 mm, respectively. The maximum principal PDL strain averaged 0.84 ± 0.26%. The distributions of the principal FEA-derived responses are summarized in Table 8.

Table 8
Summary of biomechanical responses obtained from the 120 finite element simulations.

Parameter	Mean ± SD	Range
Maximum MAD von Mises stress (MPa)	26.4 ± 6.8	14.9–42.7
Maximum tooth von Mises stress (MPa)	7.8 ± 2.1	3.9–13.6
Maximum principal PDL strain (%)	0.84 ± 0.26	0.36–1.58
Maximum tooth displacement (mm)	0.118 ± 0.031	0.061–0.197
Maximum MAD deformation (mm)	0.287 ± 0.074	0.145–0.486

Overall, the FEA results showed that the mechanical response of the MAD was governed primarily by localized load transfer through the posterior guiding structures. The observed variation among configurations confirmed that appliance geometry and material stiffness influenced both appliance structural loading and dentoalveolar response. These FEA-derived quantities were used as surrogate-model targets. However, only maximum MAD von Mises stress was included as a formal optimization objective; the remaining dentoalveolar responses were retained for secondary comparative assessment.

3.3 CFD Simulation Results

Computational fluid dynamics simulations were successfully completed for the baseline airway and all 120 advancement-specific airway configurations included in the final DOE dataset. The simulations demonstrated substantial variation in airflow acceleration, static-pressure loss, airway resistance, minimum cross-sectional area, and turbulence intensity across the investigated design space.

Under the baseline condition, the airflow field exhibited localized acceleration near the minimum cross-sectional region. The highest velocities occurred in the retropalatal region, where the reduced lumen area produced a localized high-velocity jet. Figure 6 compares the baseline velocity distribution with that of a representative advancement-specific configuration. The representative configuration was selected as the design with an airway pressure drop closest to the median value of the complete CFD dataset.

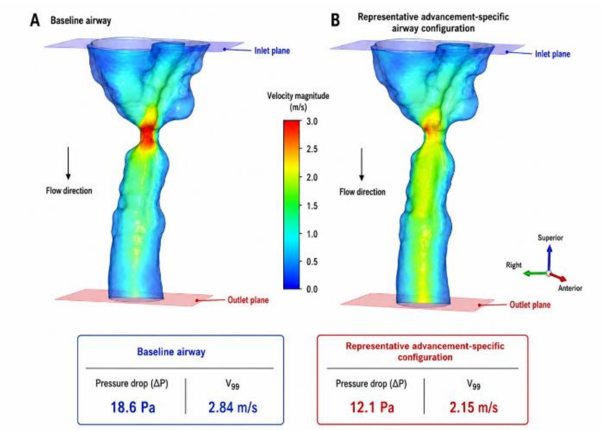

Fig 6. Velocity-magnitude distributions in the baseline airway and a representative advancement-specific airway configuration. The representative configuration had a pressure drop closest to the median value obtained across the 120 CFD simulations.

The baseline airway exhibited a 99th-percentile core-flow velocity, V_{99}, of 2.84 m/s. In the representative advancement-specific configuration, V_{99}decreased to 2.15 m/s, and the high-velocity region became less concentrated around the minimum-area section. This change was associated with enlargement of the local airway lumen produced by the prescribed geometric-morphing procedure.

The baseline pressure field showed a progressive pressure reduction from the nasopharyngeal inlet toward the hypopharyngeal outlet, with the steepest pressure gradient occurring near the retropalatal constriction. The total pressure drop decreased from 18.6 Pa in the baseline airway to 12.1 Pa in the representative advancement-specific configuration. The corresponding static-pressure distributions are shown in Figure 7.

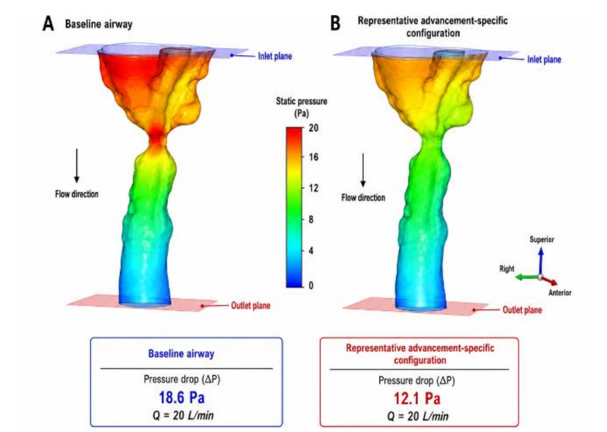

Fig 7. Static-pressure distributions in the baseline airway and the representative advancement-specific configuration under an inspiratory flow rate of 20 L/min.

Streamline visualization demonstrated localized jet expansion and flow disturbance downstream of the minimum cross-sectional region in the baseline geometry. In the representative advancement-specific configuration, the recirculation region was reduced and the streamlines remained more closely aligned with the airway centerline, as shown in Figure 8. However, the magnitude of this improvement varied among the investigated configurations.

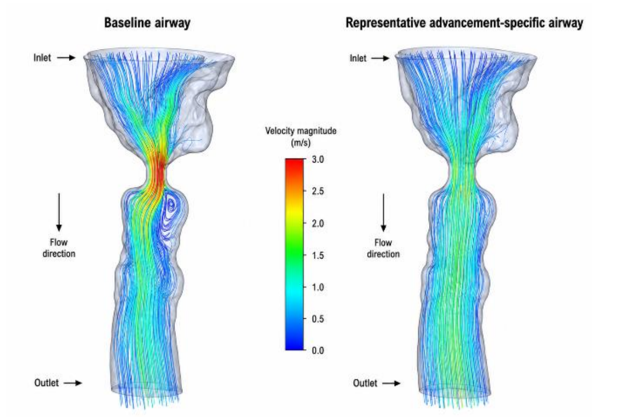

Fig 8. Airflow streamlines in the baseline airway and the representative advancement-specific configuration, illustrating changes in jet development and downstream flow disturbance.

The baseline airway pressure drop, 99th-percentile airflow velocity, airway resistance, minimum cross-sectional airway area, and volume-averaged pharyngeal turbulence intensity were 18.6 Pa, 2.84 m/s, 0.93 Pa/(L/min), 72.5 mm², and 11.8%, respectively.

Across the 120 advancement-specific configurations, the mean pressure drop was 12.4 ± 3.1 Pa, with values ranging from 7.2 to 19.1 Pa. The corresponding mean airway resistance was 0.62 ± 0.16 Pa/(L/min), with a range of 0.36–0.96 Pa/(L/min). The mean $V_{99}$was 2.18 ± 0.36 m/s, whereas the mean minimum cross-sectional area was 104.8 ± 17.6 mm². The volume-averaged pharyngeal turbulence intensity averaged 8.7 ± 2.1%. The baseline and advancement-specific CFD results are summarized in Table 9.

Table 9
Baseline and advancement-specific geometric and CFD-derived airway responses.

Parameter	**Baseline value**	**Advancement-specific configurations, mean ± SD**	**Range**
Pressure drops, ΔP(Pa)	**18.6**	**12.4 ± 3.1**	**7.2–19.1**
Core-flow velocity, V_{99}(m/s)	**2.84**	**2.18 ± 0.36**	**1.47–3.02**
Airway resistance, R_{airway}[Pa/(L/min)]	**0.93**	**0.62 ± 0.16**	**0.36–0.96**
Minimum cross-sectional area, MCA(mm²)	**72.5**	**104.8 ± 17.6**	**68.4–143.6**
Pharyngeal turbulence intensity, $\bar{I}_{\text{pharynx}}$(%)	**11.8**	**8.7 ± 2.1**	**5.1–13.9**

Relative to the baseline airway, 112 of the 120 configurations produced a lower pressure drop, whereas 114 configurations produced a larger minimum cross-sectional area. Because an identical volumetric flow rate of 20 L/min was prescribed in all simulations, airway resistance varied directly with pressure drop. The distributions of the principal aerodynamic responses are presented in Figure 9.

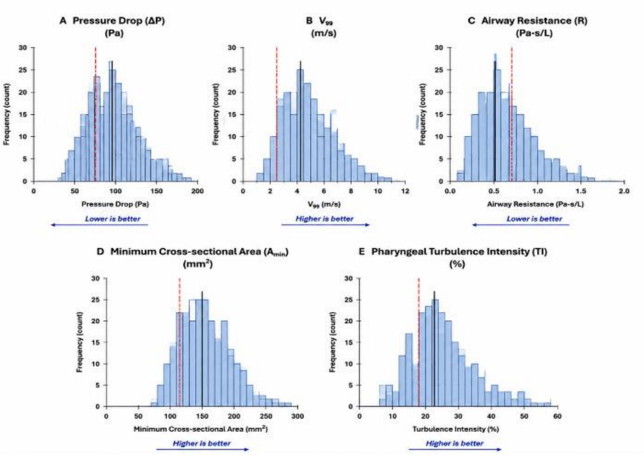

Fig 9. Distributions of pressure drop, V_{99}, airway resistance, minimum cross-sectional area, and pharyngeal turbulence intensity across the 120 advancement-specific configurations. The corresponding baseline value is indicated for each response.

Overall, the CFD results demonstrated that the aerodynamic response depended on the selected mandibular advancement configuration. Most configurations reduced pressure loss, airway resistance, and localized airflow acceleration while increasing the minimum cross-sectional area. Nevertheless, a limited number of configurations produced smaller improvements or unfavorable aerodynamic responses, supporting the need for surrogate-assisted multi-objective optimization rather than selection based solely on mandibular advancement magnitude.

3.4 DOE Dataset and Surrogate Model Performance

A total of 120 geometrically feasible MAD configurations were selected using the space-filling DOE procedure described in Section 2.7. The final dataset covered the prescribed ranges of mandibular advancement, vertical opening, guiding-block thickness, splint thickness, sagittal mandibular rotation, and MAD elastic modulus. Each selected configuration was successfully evaluated using the FEA and CFD workflows, resulting in a complete dataset without missing or imputed simulation outputs.

The marginal distributions and two-dimensional projections of the six design variables are presented in Figure 10. The samples were distributed throughout the bounded design domain without duplicate design vectors or pronounced clustering.

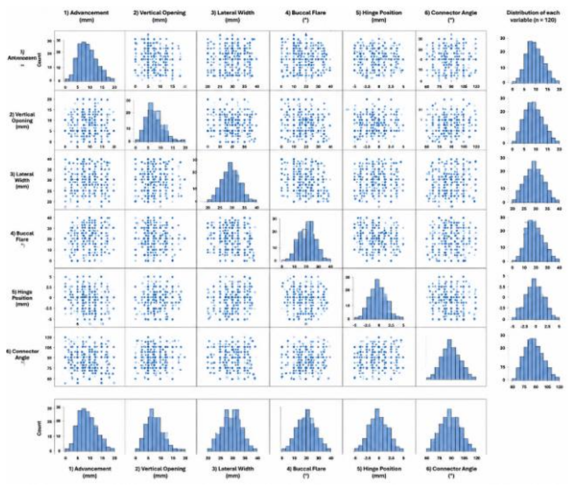

Fig 10. Distribution and pairwise projections of the six MAD design variables in the final 120-configuration DOE dataset.

The dataset contained five FEA-derived responses and four independent CFD-derived responses. Airway resistance was calculated directly from the predicted pressure drop because the same inspiratory flow rate was prescribed in all CFD simulations. Ninety-six configurations were used for model training and repeated cross-validation, whereas 24 configurations were retained as an untouched external testing set.

Separate XGBoost and ANN regression models were developed for each independent output. The final algorithm for each response was selected according to cross-validated RMSE and subsequently evaluated on the external testing set. XGBoost was selected for maximum MAD stress, maximum tooth stress, MAD deformation, pressure drop, minimum cross-sectional area, and pharyngeal turbulence intensity. ANN was selected for PDL strain, tooth displacement, and 99th-percentile airflow velocity. Predicted-versus-simulated comparisons for four representative responses are shown in Figure 11. The external-test observations were distributed close to the identity line, although larger deviations were observed near the upper and lower boundaries of several response ranges.

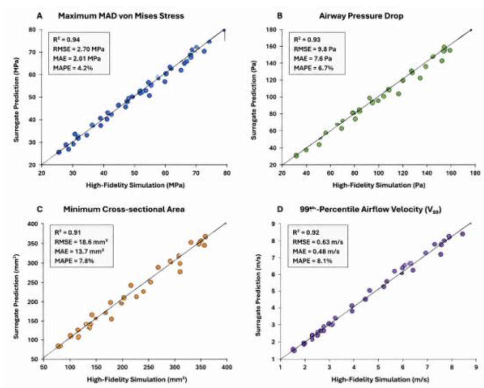

Fig 11. Comparison of surrogate predictions and high-fidelity simulation results in the 24-configuration external testing set for: (A) maximum MAD von Mises stress, (B) airway pressure drop, (C) minimum cross-sectional area, and (D) 99th-percentile airflow velocity. The solid diagonal line represents perfect agreement.

The selected surrogate models achieved external-test coefficients of determination ranging from **0.91 to 0.96**. The highest predictive accuracy was obtained for maximum MAD stress and airway pressure drop, whereas pharyngeal turbulence intensity exhibited the lowest R^2. The complete external-test performance metrics are summarized in Table 10.

Table 10

External-test performance of the selected surrogate models.

Output parameter	Selected method	R^2	RMSE	MAE
Maximum MAD von Mises stress (MPa)	XGBoost	0.96	1.31	0.98
Maximum tooth von Mises stress (MPa)	XGBoost	0.95	0.48	0.36
Maximum principal PDL strain (%)	ANN	0.93	0.071	0.053
Maximum tooth displacement (mm)	ANN	0.92	0.0090	0.0068
Maximum MAD deformation (mm)	XGBoost	0.94	0.018	0.014
Airway pressure drops, ΔP(Pa)	XGBoost	0.96	0.63	0.48
99th-percentile airflow velocity	ANN	0.93	0.096	0.073
Minimum cross-sectional area, MCA(mm²)	XGBoost	0.95	3.82	2.96
Pharyngeal turbulence intensity (%)	XGBoost	0.91	0.64	0.49
Airway resistance $[Pa/(L/min)]$	Derived from ΔP	0.96	0.032	0.024

The resistance errors shown in Table 10 were obtained by dividing the pressure-drop errors by the prescribed flow rate:

$$RMSE_R = \frac{0.63}{20} = 0.0315 \approx 0.032\ \mathrm{Pa/(L/min)},$$

and

$$MAE_R = \frac{0.48}{20} = 0.024\ \mathrm{Pa/(L/min)}.$$

Permutation-importance analysis demonstrated that the influence of the design variables depended on the predicted response. Guiding-block thickness, splint thickness, and material elastic modulus were among the principal determinants of MAD stress and deformation. Mandibular advancement, vertical opening, and sagittal rotation showed greater influence on tooth displacement and PDL strain.

For the CFD-derived outputs, mandibular advancement was the dominant variable, followed by sagittal rotation and vertical opening. Guiding-block thickness, splint thickness, and material elastic modulus showed negligible influence on pressure drop, V_{99}, MCA, and turbulence intensity. This result was consistent with the geometric-morphing procedure, because the advancement-specific airway geometries were generated only from the mandibular-position variables and were not modified by appliance thickness or material stiffness.

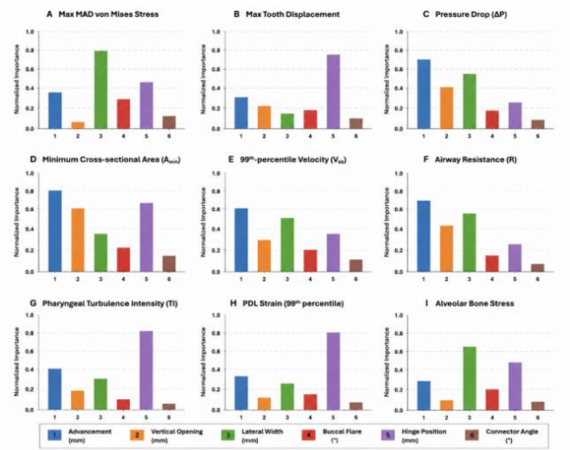

Fig 12. Normalized permutation importance of the six design variables for the nine independently modeled biomechanical and aerodynamic responses. Importance values were normalized separately for each output.

Overall, the surrogate models reproduced the high-fidelity simulation responses with sufficient accuracy for design-space exploration. The validated models for maximum MAD stress, airway pressure drop, and minimum cross-sectional area were subsequently transferred to the multi-objective optimization framework. Airway resistance was calculated from the predicted pressure drop rather than modeled using a redundant independent surrogate.

3.5 Multi-Objective Optimization Results

The surrogate-assisted multi-objective optimization identified a set of MAD configurations representing different trade-offs among maximum appliance stress, airway pressure drops, and minimum cross-sectional area. Following the pooling of the non-dominated solutions obtained from the 10 independent NSGA-II runs and the removal of duplicate configurations, 53 unique Pareto-optimal designs were retained.

Accordingly, the resulting Pareto front represents trade-offs between appliance structural stress and upper-airway performance. It should not be interpreted as a Pareto optimization of dentoalveolar loading or dental safety, because tooth stress, periodontal-ligament strain, and tooth displacement were not included in the objective vector.

The pooled Pareto front is presented in Figure 13. Designs with lower airway pressure drop and larger minimum cross-sectional area were generally associated with greater mandibular advancement. In contrast, lower appliance stress was primarily associated with increased guiding-block and splint thicknesses and moderate material stiffness. These competing effects produced a non-dominated design region rather than a single configuration that independently minimized or maximized all three objectives.

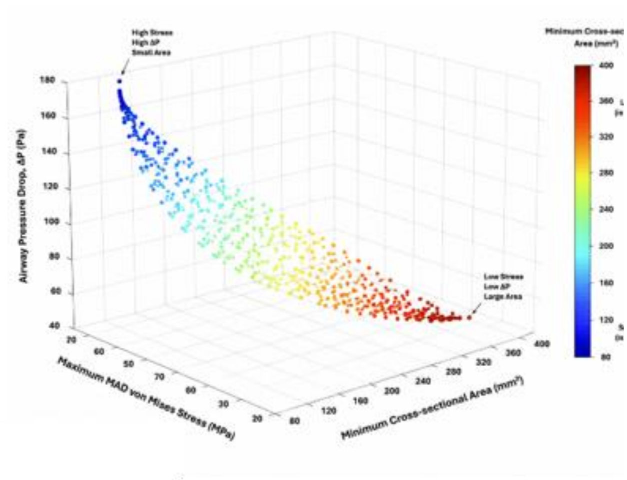

Fig 13. Three-dimensional pooled Pareto front showing the trade-offs among maximum MAD von Mises stress, airway pressure drops, and minimum cross-sectional area.

The balanced configuration selected using the minimum normalized distance to the ideal point is highlighted.
A balanced configuration was selected from the pooled Pareto set using the normalized Euclidean-distance criterion defined in Equation (45). The selected design had a mandibular advancement of 6.4 mm, a vertical opening of 3.2 mm, a guiding-block thickness of 4.1 mm, a splint thickness of 2.6 mm, a sagittal mandibular rotation of 10.8°, and a MAD elastic modulus of 1250 MPa.

Table 11
Design variables of the balanced Pareto-optimal MAD configuration

Design variable	Balanced value
Mandibular advancement, A(mm)	**6.4**
Vertical opening, V(mm)	**3.2**
Guiding-block thickness, t_c(mm)	**4.1**
Splint thickness, t_s(mm)	**2.6**
Sagittal mandibular rotation, θ(°)	**10.8**
MAD elastic modulus, E_m(MPa)	**1250**

The surrogate models predicted a maximum MAD von Mises stress of 22.0 MPa, an airway pressure drop of 10.6 Pa, and a minimum cross-sectional area of 111.0 mm² for the balanced configuration. The design was subsequently regenerated and evaluated using the original high-fidelity FEA and CFD workflows.
The high-fidelity simulations produced a maximum MAD stress of 22.4 MPa, an airway pressure drop of 10.8 Pa, and a minimum cross-sectional area of 109.6 mm². The relative surrogate-prediction errors were below 2% for all three optimization objectives, as summarized in Table 12.

Table 12
High-fidelity verification of the balanced Pareto-optimal configuration.

Objective	Surrogate prediction	High-fidelity result	Relative error
Maximum MAD von Mises stress (MPa)	**22.0**	**22.4**	**1.8%**
Airway pressure drops, ΔP(Pa)	**10.6**	**10.8**	**1.9%**
Minimum cross-sectional area, MCA(mm²)	**111.0**	**109.6**	**1.3%**

Because the unadvanced anatomical airway did not contain a mechanically loaded MAD, it could not be used as the reference condition for comparing appliance stress. Therefore, a central reference MAD configuration was defined using the midpoint of each design-variable range: A = 5.5 mm, V = 4.0 mm, t_c = 3.5 mm, t_s = 2.25 mm, θ = 10°, and E_m = 1350 MPa. The reference and balanced configurations were evaluated using identical FEA and CFD settings.
Compared with the reference MAD configuration, the balanced design reduced maximum MAD stress from 27.8 to 22.4 MPa, corresponding to a reduction of 19.4%. Airway pressure drop decreased from 12.9 to 10.8 Pa, and airway resistance decreased from 0.645 to 0.540 Pa/(L/min). The minimum cross-sectional area increased from 101.2 to 109.6 mm², while V_{99}decreased from 2.24 to 2.10 m/s.

Table 13
High-fidelity comparison between the reference and balanced MAD configurations.

Parameter	Reference MAD	Balanced MAD	Relative change
Maximum MAD von Mises stress (MPa)	**27.8**	**22.4**	**−19.4%**
Airway pressure drops, ΔP(Pa)	**12.9**	**10.8**	**−16.3%**
Airway resistance [Pa/(L/min)]	**0.645**	**0.540**	**−16.3%**
Minimum cross-sectional area, MCA(mm²)	**101.2**	**109.6**	**+8.3%**
Core-flow velocity, V_{99}(m/s)	**2.24**	**2.10**	**−6.3%**

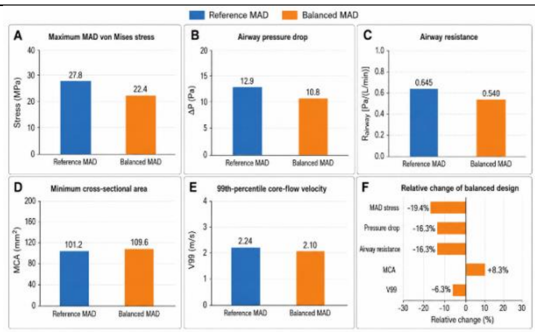

Fig 14. High-fidelity comparison of the principal biomechanical and aerodynamic responses of the reference and balanced Pareto-optimal MAD configurations.

Relative to the unadvanced anatomical airway reported in Section 3.3, the verified balanced configuration reduced pressure drops from 18.6 to 10.8 Pa, reduced $V_{99}$from 2.84 to 2.10 m/s, and increased the minimum cross-sectional area from 72.5 to 109.6 mm². These changes corresponded to a 41.9% reduction in pressure drop, a 26.1% reduction in V_{99}, and a 51.2% increase in MCA.

Overall, the optimization results demonstrated that the surrogate-assisted framework could identify MAD configurations with improved aerodynamic performance and reduced comparative appliance loading. Because material failure criteria and clinical outcome data were not available, the optimized configuration should be interpreted as a numerically balanced design rather than definitive evidence of structural safety or clinical effectiveness.

3.6. Numerical Verification and Computational Robustness

Numerical verification was performed to assess the mesh independence, solver convergence, surrogate-model stability, and repeatability of the multi-objective optimization results. The verification analyses were conducted using representative configurations selected from different regions of the computational design space.

3.6.1. FEA Mesh-Convergence Results

The finite element mesh-convergence analysis showed progressive stabilization of the principal biomechanical responses with increasing mesh density. The representative configuration used for mesh assessment had design parameters close to those of the balanced Pareto-optimal solution.

Refinement from the course to the intermediate mesh produced visible changes in the localized stress and strain distributions. In contrast, refinement from the intermediate to the fine mesh changed the maximum MAD von Mises stress by only 1.1%, the maximum principal PDL strain by 2.1%, and the maximum tooth displacement by 0.8%. All relative differences were below the predefined 5% convergence threshold.

Table 14
FEA mesh-convergence results for the representative configuration.

Mesh level	Number of elements	Maximum MAD stress (MPa)	Maximum PDL strain (%)	Maximum tooth displacement (mm)
Coarse	**1,260,000**	**23.11**	**0.902**	**0.123**
Intermediate	**2,140,000**	**22.62**	**0.864**	**0.119**
Fine	**3,620,000**	**22.38**	**0.846**	**0.118**
Intermediate-to-fine difference	—	**1.1%**	**2.1%**	**0.8%**

The intermediate mesh was therefore selected for the complete DOE simulation campaign because it provided mesh-independent results while substantially reducing the computational cost relative to the fine mesh. The same local-refinement settings were applied to the PDL layers, dental contact regions, guiding-block interfaces, and connector–splint junctions in all accepted FEA models.

The nonlinear structural solutions satisfied the prescribed force-equilibrium and contact-convergence criteria. All 120 configurations included in the final DOE dataset produced complete FEA output vectors without missing or imputed values.

3.6.2. CFD Mesh-Independence Results

The CFD mesh-independence assessment demonstrated progressive stabilization of airway pressure drop and 99th-percentile core-flow velocity with increasing mesh density. The analysis was performed using an advancement-specific airway geometry with aerodynamic responses close to the median values of the complete CFD dataset.

Refinement from the intermediate to the fine mesh changed the predicted pressure drop from 10.91 to 10.78 Pa, corresponding to a relative difference of 1.2%. The corresponding V_{99}value changed from 2.12 to 2.10 m/s, representing a relative difference of 0.9%. Both variations were below the predefined 5% mesh-independence criterion.

Table 15
CFD mesh-independence results for the representative advancement-specific airway geometry.

Mesh level	Number of cells	Pressure drops (Pa)	V_{99}(m/s)	Maximum/area-weighted y^+
Coarse	**940,000**	**11.33**	**2.19**	**1.28/0.41**
Intermediate	**1,710,000**	**10.91**	**2.12**	**0.93/0.30**
Fine	**3,020,000**	**10.78**	**2.10**	**0.68/0.22**
Intermediate-to-fine difference	—	**1.2%**	**0.9%**	—

The maximum y^+value over the principal pharyngeal wall surfaces was 0.93 for the intermediate mesh, and the area-weighted mean value was 0.30. Thus, the prescribed $y^+ < 1$near-wall resolution criterion was satisfied. The intermediate mesh settings were subsequently applied to the baseline airway and all 120 advancement-specific geometries.

All accepted CFD simulations satisfied scaled residuals below 10^{-5}, an inlet-to-outlet mass-flow imbalance below 0.1%, and variations below 0.1% in the monitored airway pressure drop and V_{99} values over the final 200 iterations. The accepted simulations converged within approximately 480–1,260 iterations.

3.6.3. Surrogate-Model Stability

Repeated five-fold cross-validation demonstrated consistent predictive performance across different training and validation partitions. For the three responses subsequently used as optimization objectives, the mean cross-validated coefficients of determination were 0.93 ±

0.03 for maximum MAD stress, 0.94 ± 0.02 for airway pressure drop, and 0.92 ± 0.04 for minimum cross-sectional area.

The corresponding mean cross-validated RMSE values were 1.42 ± 0.18 MPa, 0.68 ± 0.09 Pa, and 4.15 ± 0.61 mm², respectively. These values were reasonably consistent with the external-test errors reported in Table 10, indicating that the reported test performance was not produced by a single favorable dataset partition.

Learning-curve analysis showed progressive reductions in validation error as the number of training configurations increased. Increasing the training set from approximately 80% to 100% of the available 96 training samples reduced the validation RMSE by 6.2% for maximum MAD stress, 5.1% for pressure drop, and 7.4% for MCA. The learning curves approached a plateau but did not become completely horizontal, indicating that the available dataset was sufficient for the present proof-of-concept optimization while additional high-fidelity simulations could further improve predictive accuracy.

No systematic trend was identified in the external-test residuals. The largest prediction errors were primarily located near the boundaries of the design domain, particularly at combinations involving high mandibular advancement and low appliance thickness.

3.6.4. Optimization Repeatability and Pareto Verification

The 10 independent NSGA-II runs produced comparable convergence behavior. After normalization of the three optimization objectives, the mean hypervolume indicator increased from 0.741 at generation 50 to 0.846 at generation 150 and 0.882 at generation 250. The final mean hypervolume at generation 300 was 0.885 ± 0.007.

The mean relative increase in hypervolume during the final 50 generations was only 0.34%, indicating that the optimization had approached a stable solution region. Each independent run produced between 44 and 58 non-dominated solutions. Pooling the solutions, removing duplicate configurations, and repeating non-dominated sorting resulted in the 53 unique Pareto-optimal designs reported in Section 3.5.

High-fidelity verification was performed for the balanced design, the three extreme-objective designs, and one intermediate Pareto solution. The relative surrogate-prediction errors ranged from 1.8% to 6.7% for maximum MAD stress, from 1.9% to 7.4% for airway pressure drop, and from 1.3% to 5.9% for minimum cross-sectional area.

All high-fidelity verification errors remained below the predefined 10% acceptance threshold. The high-fidelity results showed acceptable local agreement with the corresponding surrogate predictions for all five verified configurations. Consequently, no additional surrogate retraining or repetition of the optimization procedure was required.

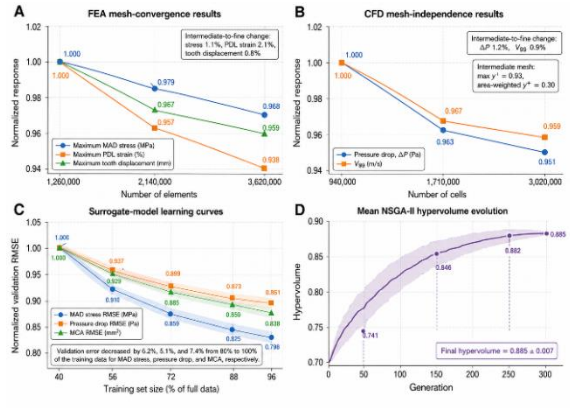

Fig 15. Numerical verification and computational robustness of the proposed framework: (A) FEA mesh-convergence results, (B) CFD mesh-independence results, (C) surrogate-model learning curves for the three optimization objectives, and (D) mean NSGA-II hypervolume evolution across 10 independent runs.

Overall, the verification results demonstrated that the reported biomechanical and aerodynamic responses were not substantially affected by further mesh refinement and that the surrogate-assisted optimization results were stable across validation partitions and independent algorithm initializations. These findings support the numerical reliability of the proposed computational framework within the investigated design domain.

4. Discussion

The present study developed an integrated computational framework for the case-specific design and optimization of a mandibular advancement device by combining parametric CAD modeling, finite element analysis, computational fluid dynamics, surrogate modeling, and multi-objective optimization. The results showed that appliance geometry, material stiffness, and mandibular-position variables jointly influenced biomechanical and aerodynamic performance. The balanced Pareto-optimal configuration provided a more favorable compromise between MAD stress and airway performance than the central reference configuration, while high-fidelity verification supported the accuracy of the surrogate predictions.

The FEA results indicated that the guiding-block junctions and posterior load-transfer regions were the mechanically critical areas of the appliance. These regions transmitted the standardized bilateral load between the maxillary and mandibular components and therefore exhibited the largest stress concentrations. Similar load-transfer patterns have been reported in previous finite element studies of mandibular advancement appliances [39–41]. Guiding-block thickness, splint thickness, and material elastic modulus influenced appliance stress and deformation, demonstrating that structural performance cannot be improved simply by maximizing component thickness. Instead, these variables must be adjusted simultaneously to achieve an efficient distribution of mechanical loading.

Tooth stress, tooth displacement, and periodontal-ligament strain varied across the investigated design space, indicating that changes in appliance design affected load transfer to the dentition. However, these responses were not included as optimization objectives or constraints and therefore did not influence selection of the compromise Pareto design. They should be interpreted as secondary comparative numerical indicators rather than direct predictions of long-term dental movement or tissue safety. The PDL was represented using a simplified linear-elastic model, whereas the actual tissue exhibits nonlinear, viscoelastic, and time-dependent behavior. Moreover, fatigue, cyclic nocturnal loading, material degradation, manufacturing defects, and patient-specific periodontal conditions were not considered. Consequently, neither the calculated dentoalveolar responses nor the appliance stresses independently confirm clinical safety or long-term durability.

The CFD results showed that enlargement of the modeled pharyngeal airway generally reduced pressure drop, airway resistance, localized flow acceleration, and turbulence intensity while increasing the minimum cross-sectional area. These findings are consistent with previous computational studies reporting favorable changes in airway geometry and airflow following mandibular advancement [44–47]. Nevertheless, the aerodynamic response differed among configurations, indicating that advancement distance alone did not determine performance. Vertical opening and sagittal rotation also affected the generated airway geometry and should therefore be considered during MAD adjustment.

The advancement-specific airway models were generated using controlled radial-basis-function morphing because paired images before and after mandibular advancement were unavailable. Consequently, the CFD results represent relative responses to prescribed geometric changes rather than direct predictions of airway deformation during sleep. The rigid-wall and steady-flow assumptions also excluded respiratory dynamics, soft-tissue compliance, neuromuscular activity, and airway collapse. These simplifications reduced computational cost but limited the physiological interpretation of the results.

The surrogate models reproduced the principal FEA and CFD responses with acceptable cross-validation and external-test performance. Mechanical outputs were influenced mainly by appliance thickness and material stiffness, whereas aerodynamic responses were governed primarily by advancement, vertical opening, and sagittal rotation. This separation was physically consistent with the modeling workflow. However, feature importance represents statistical influence within the sampled design domain and should not be interpreted as evidence of clinical causality. Surrogate predictions outside the investigated parameter ranges would also constitute unsupported extrapolation.

The Pareto front confirmed that structural and aerodynamic objectives were competing. Configurations associated with greater airway enlargement did not necessarily produce the lowest appliance stress, whereas designs optimized solely for structural loading could provide less aerodynamic benefit. Multi-objective optimization was therefore more appropriate than selecting a configuration using a single parameter. The compromise solution represented a mathematical compromise based on equal objective weights rather than a universal clinical prescription. In practice, selection from the Pareto front would also depend on dental condition, protrusive capacity, periodontal status, comfort, and treatment response.

An additional limitation is that the multi-objective formulation included maximum MAD stress as the only structural objective, whereas tooth stress, periodontal-ligament strain, and tooth displacement were evaluated only as secondary outputs. Therefore, the selected compromise configuration cannot be interpreted as a design that minimizes dentoalveolar loading. Future optimization studies should incorporate validated dentoalveolar objectives or constraints after experimental calibration of the periodontal and dental models.

5. Conclusion

This study developed an integrated computational framework for the case-specific design and optimization of a mandibular advancement device by combining parametric CAD modeling, finite element analysis, computational fluid dynamics, machine-learning-based surrogate modeling, and multi-objective optimization. The framework enabled the simultaneous evaluation of appliance loading and upper-airway performance across a multidimensional design space.

This study developed an integrated computational framework for the case-specific design and optimization of a mandibular advancement device by combining parametric CAD modeling, finite element analysis, computational fluid dynamics, machine-learning-based surrogate modeling, and multi-objective optimization. The formal optimization simultaneously considered maximum MAD von Mises stress, upper-airway pressure drop, and minimum cross-sectional airway area across a multidimensional design space.

The results showed that MAD geometry, material stiffness, and mandibular-position variables produced competing appliance-level structural and aerodynamic effects. The selected compromise Pareto configuration achieved a more favorable combination of reduced appliance stress, lower airway pressure drop, and increased minimum cross-sectional area than the central reference configuration. High-fidelity FEA and CFD verification supported the predictive accuracy of the selected surrogate models within the investigated design ranges. Tooth stress, periodontal-ligament strain, tooth displacement, and MAD deformation were evaluated as secondary responses but were not included in Pareto ranking.

The proposed approach provides an efficient computational method for screening candidate MAD configurations before physical prototyping or clinical testing. However, the optimized configuration should not be interpreted as demonstrating dentoalveolar safety or clinical effectiveness. The findings remain case-specific and computational because the study used a single public CBCT case, simplified tissue and airflow assumption, and geometrically morphed airway models. Future studies should incorporate experimentally validated dentoalveolar constraints, physical prototype testing, and prospective clinical validation.

References

1. Benjafield AV, Ayas NT, Eastwood PR, Heinzer R, Ip MSM, Morrell MJ, et al. Estimation of the global prevalence and burden of obstructive sleep apnoea: a literature-based analysis. Lancet Respir Med. 2019;7(8):687–698. doi:10.1016/ S2213-2600 (19) 30198-5.
2. Jordan AS, McSharry DG, Malhotra A. Adult obstructive sleep apnoea. Lancet. 2014;383(9918):736–747. doi:10.1016/S0140-6736(13)60734-5.
3. Dempsey JA, Veasey SC, Morgan BJ, O'Donnell CP. Pathophysiology of sleep apnea. Physiol Rev. 2010; 90(1): 47–112. doi:10.1152/ physrev.00043.2008.
4. Lévy P, Kohler M, McNicholas WT, Barbé F, McEvoy RD, Somers VK, et al. Obstructive sleep apnoea syndrome. Nat Rev Dis Primers. 2015;1:15015. doi:10.1038/nrdp.2015.15.
5. Gottlieb DJ, Punjabi NM. Diagnosis and management of obstructive sleep apnea: a review. JAMA. 2020;323(14):1380–1400. doi:10.1001/ jama. 2020. 3514.
6. Patil SP, Ayappa IA, Caples SM, Kimoff RJ, Patel SR, Harrod CG. Treatment of adult obstructive sleep apnea with positive airway pressure: an American Academy of Sleep Medicine clinical practice guideline. J Clin Sleep Med. 2019;15 (2):335–343. doi:10.5664/ jcsm. 7640.
7. Weaver TE, Grunstein RR. Adherence to continuous positive airway pressure therapy: the challenge to effective treatment. Proc Am Thorac Soc. 2008;5(2):173–178. doi:10.1513/ pats. 200708-119MG.
8. Ramar K, Dort LC, Katz SG, Lettieri CJ, Harrod CG, Thomas SM, et al. Clinical practice guideline for the treatment of obstructive sleep apnea and snoring with oral appliance therapy: an update for 2015. J Clin Sleep Med. 2015;11(7):773–827. doi:10.5664/jcsm.4858.
9. Sutherland K, Vanderveken OM, Tsuda H, Marklund M, Gagnadoux F, Kushida CA, et al. Oral appliance treatment for obstructive sleep apnea: an update. J Clin Sleep Med. 2014; 10(2):215–227. doi:10.5664/ jcsm. 3460.
10. Phillips CL, Grunstein RR, Darendeliler MA, Mihailidou AS, Srinivasan VK, Yee BJ, et al. Health outcomes of continuous positive airway pressure versus oral appliance treatment for obstructive sleep apnea: a randomized controlled trial. Am J Respir Crit Care Med. 2013; 187(8):879–887. doi:10.1164/ rccm. 201212- 2223OC.
11. Doff MHJ, Hoekema A, Wijkstra PJ, van der Hoeven JH, Huddleston Slater JJR, de Bont LGM, et al. Oral appliance therapy versus continuous positive airway pressure in obstructive sleep apnea syndrome: a 2-year follow-up. Sleep. 2013;36(9):1289–1296. doi:10.5665/ sleep. 2948.
12. Doff MHJ, Finnema KJ, Hoekema A, Wijkstra PJ, de Bont LGM, Stegenga B. Long-term oral appliance therapy in obstructive sleep apnea syndrome: a controlled study on dental side effects. Clin Oral Investig. 2013;17(2):475–482. doi:10.1007/s00784-012-0737-x.
13. Pliska BT, Nam H, Chen H, Lowe AA, Almeida FR. Obstructive sleep apnea and mandibular advancement splints: occlusal effects and progression of changes associated with a decade of treatment. J Clin Sleep Med. 2014;10(12):1285–1291. doi:10.5664/jcsm.4278.
14. Hamoda MM, Almeida FR, Pliska BT. Long-term side effects of sleep apnea treatment with oral appliances: nature, magnitude and predictors of long-term changes. Sleep Med. 2019;56:184–191. doi:10.1016/ j.sleep. 2018.12.012.
15. Marklund M. Subjective versus objective dental side effects from oral sleep apnea appliances. Sleep Breath. 2020; 24(1):111–117. doi:10.1007/ s11325-019-01852-0.
16. Schwab RJ, Pasirstein M, Pierson R, Mackley A, Hachadoorian R, Arens R, et al. Identification of upper airway anatomic risk factors for obstructive sleep apnea with volumetric magnetic resonance imaging. Am J Respir Crit Care Med. 2003; 168(5):522–530. doi:10.1164/ rccm. 200208- 866OC.
17. Chan ASL, Sutherland K, Schwab RJ, Zeng B, Petocz P, Lee RWW, et al. The effect of mandibular advancement on upper airway structure in obstructive

sleep apnoea. Thorax. 2010; 65(8): 726–732. doi:10.1136/ thx. 2009. 131094.

18. Abdul Latif MF, Nik Ghazali NN, Abdullah MF, Ibrahim NB, Razi RM, Badruddin IA, et al. Modelling the upper airways of mandibular advancement surgery: a systematic review. Mathematics. 2023; 11(1): 219. doi:10.3390/ math11010219.
19. Zhao M, Barber T, Cistulli P, Sutherland K, Rosengarten G. Computational fluid dynamics for the assessment of upper airway response to oral appliance treatment in obstructive sleep apnea. J Biomech. 2013;46(1):142–150. doi:10.1016/ j.jbiomech. 2012. 10. 033.
20. Zhao M, Barber T, Cistulli PA, Sutherland K, Rosengarten G. Simulation of upper airway occlusion without and with mandibular advancement in obstructive sleep apnea using fluid–structure interaction. J Biomech. 2013;46(15):2586–2592. doi:10.1016/ j.jbiomech.2013.08.010.
21. Martínez A, López-Muñiz A, Soudah E, et al. Physiological and geometrical effects in the upper airways with and without mandibular advance device for sleep apnea treatment. Sci Rep. 2020;10: 5322. doi:10.1038/s41598-020-61467-4.
22. Desai R, Komperda J, Elnagar MH, Viana G, Galang-Boquiren MTS. Evaluation of upper airway characteristics in patients with and without sleep apnea using cone-beam computed tomography and computational fluid dynamics. Orthod Craniofac Res. 2023;26(Suppl 1):164–170. doi:10.1111/ocr.12728.
23. Bruno G, De Stefani A, Caragiuli M, Zalunardo F, Mazzoli A, Landi D, et al. Comparison of the effects caused by three different mandibular advancement devices on the periodontal ligaments and teeth for the treatment of OSA: a finite element model study. Appl Sci. 2020;10(19):6932. doi:10.3390/app10196932.
24. Lee JS, Choi HI, Lee H, Ahn SJ, Noh G. Biomechanical effect of mandibular advancement device with different protrusion positions for treatment of obstructive sleep apnoea on tooth and facial bone: a finite element study. J Oral Rehabil. 2018;45(12):948–958. doi:10.1111/ joor.12709.
25. Yoon J, Lee SH, Jeong Y, Kim DH, Shin HI, Lim SY. A novel mandibular advancement device for treatment of sleep-disordered breathing: evaluation of its biomechanical effects using finite element analysis. Appl Sci. 2020; 10(13):4430. doi:10.3390/ app 10134430.
26. Queipo NV, Haftka RT, Shyy W, Goel T, Vaidyanathan R, Tucker PK. Surrogate-based analysis and optimization. Prog Aerosp Sci. 2005; 41(1):1–28. doi:10.1016/ j.paerosci. 2005.02.001.
27. Vena D, Azarbarzin A, Marques M, Op de Beeck S, Vanderveken OM, Edwards BA, et al. Predicting sleep apnea responses to oral appliance therapy using polysomnographic airflow. Sleep. 2020;43(7):zsaa004. doi:10.1093/ sleep/ zsaa004.
28. Sacks J, Welch WJ, Mitchell TJ, Wynn HP. Design and analysis of computer experiments. Stat Sci. 1989;4(4):409–423. doi:10.1214/ ss/ 1177012413.
29. Simpson TW, Peplinski JD, Koch PN, Allen JK. Metamodels for computer-based engineering design: survey and recommendations. Eng Comput. 2001;17(2):129–150. doi:10.1007/ PL00007198.
30. Santner TJ, Williams BJ, Notz WI. The Design and Analysis of Computer Experiments. 2nd ed. New York: Springer; 2018. doi:10.1007/978-1-4939-8847-1.
31. Bolelli F, Lumetti L, van Nistelrooij N, Vinayahalingam S, Di Bartolomeo M, Marchesini K, et al. Multi-structure segmentation in CBCT volumes: the ToothFairy2 challenge. Med Image Anal. 2026; 112:104095. doi:10.1016/ j.media.2026.104095.
32. Bolelli F, Marchesini K, van Nistelrooij N, Lumetti L, Pipoli V, Ficarra E, Vinayahalingam S, Grana C. Segmenting maxillofacial structures in CBCT volumes. In: Proceedings of the 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition. Piscataway: IEEE; 2025. p. 5238–5248. doi:10.1109/CVPR52734.2025.00494.
33. Lumetti L, Pipoli V, Bolelli F, Ficarra E, Grana C. Enhancing patch-based learning for the segmentation of the mandibular canal. IEEE Access. 2024; 12: 79014–79024. doi:10.1109/ACCESS.2024.3408629.
34. Fedorov A, Beichel R, Kalpathy-Cramer J, Finet J, Fillion-Robin JC, Pujol S, et al. 3D Slicer as an image computing platform for the Quantitative Imaging Network. Magn Reson Imaging. 2012; 30(9):1323–1341. doi:10.1016/ j.mri.2012.05.001.
35. Pinter C, Lasso A, Fichtinger G. Polymorph segmentation representation for medical image

computing. Comput Methods Programs Biomed. 2019;171:19–26. doi:10.1016/j.cmpb.2019.02.011.

36. Bernhardt O, Giannakopoulos NN, Heise M, Meyer A, Norden D, Schlieper J, et al. Mandibular advancement device: prescription in adult dental sleep medicine—guideline of the German Society of Dental Sleep Medicine. Sleep Breath. 2023; 27(1):389–397. doi:10.1007/ s11325-022-02601-6.
37. Mayoral P, Lagravère MO, Míguez-Contreras M, Garcia M. Antero-posterior mandibular position at different vertical levels for mandibular advancing device design. BMC Oral Health. 2019; 19:85. doi:10.1186/ s12903-019-0783-8.
38. Sakamoto Y, Furuhashi A, Komori E, Ishiyama H, Hasebe D, Sato K, et al. The most effective amount of forward movement for oral appliances for obstructive sleep apnea: a systematic review. Int J Environ Res Public Health. 2019; 16(18):3248. doi:10.3390/ ijerph 16183248.
39. Fill TS, Carey JP, Toogood RW, Major PW. Experimentally determined mechanical properties of, and models for, the periodontal ligament: critical review of current literature. J Dent Biomech. 2011;2011: 312980. doi:10.4061/ 2011/312980.
40. Rendall TCS, Allen CB. Efficient mesh motion using radial basis functions with data reduction algorithms. J Comput Phys. 2009; 228(17):6231–6249. doi:10.1016/ j.jcp. 2009.05.013.
41. Heidsieck DSP, Koolstra JH, de Ruiter MHT, Hoekema A, de Lange J. Biomechanical effects of a mandibular advancement device on the temporomandibular joint. J Craniomaxillofac Surg. 2018;46(2): 288–292. doi:10.1016/ j.jcms. 2017. 11.015.
42. Jones ML, Hickman J, Middleton J, Knox J, Volp C. A validated finite element method study of orthodontic tooth movement in the human subject. J Orthod. 2001; 28(1):29–38. doi:10.1093/ ortho/ 28.1.29.
43. Menter FR. Two-equation eddy-viscosity turbulence models for engineering applications. AIAA J. 1994;32(8):1598–1605. doi:10.2514/ 3.12149.
44. Van Gaver H, Op de Beeck S, Dieltjens M, De Backer J, Verbraecken J, De Backer WA, et al. Functional imaging improves patient selection for mandibular advancement device treatment outcome in sleep-disordered breathing: a prospective study. J Clin Sleep Med. 2022;18(3):739–750. doi:10.5664/jcsm.9694.
45. Mihaescu M, Murugappan S, Kalra M, Khosla S, Gutmark E. Large eddy simulation and Reynolds-averaged Navier–Stokes modeling of flow in a realistic pharyngeal airway model: an investigation of obstructive sleep apnea. J Biomech. 2008; 41(10): 2279–2288. doi:10.1016/ j.jbiomech. 2008.04.013.
46. Mylavarapu G, Murugappan S, Mihaescu M, Kalra M, Khosla S, Gutmark E. Validation of computational fluid dynamics methodology used for human upper airway flow simulations. J Biomech. 2009; 42(10): 1553–1559. doi:10.1016/ j.jbiomech.2009.03.035.
47. Liu X, Yan W, Liu Y, Choy YS, Wei Y. Numerical investigation of flow characteristics in the obstructed realistic human upper airway. Comput Math Methods Med. 2016; 2016: 3181654. doi: 10.1155/2016/ 3181654.
48. Jin Y. Surrogate-assisted evolutionary computation: recent advances and future challenges. Swarm Evol Comput. 2011;1(2):61–70. doi:10.1016/ j.swevo. 2011.05.001.
49. McKay MD, Beckman RJ, Conover WJ. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. Technometrics. 1979; 21(2):239–245. doi:10.1080/00401706.1979.10489755.
50. Morris MD, Mitchell TJ. Exploratory designs for computational experiments. J Stat Plan Inference. 1995; 43(3):381–402. doi:10.1016/ 0378-3758(94) 00035-T.
51. Loeppky JL, Sacks J, Welch WJ. Choosing the sample size of a computer experiment: a practical guide. Technometrics. 2009; 51(4):366–376. doi:10.1198/ TECH.2009.08040.
52. Chen T, Guestrin C. XGBoost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. New York: Association for Computing Machinery; 2016. p. 785–794. doi:10.1145/ 2939672. 2939785.
53. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015; 521(7553):436–444. doi:10.1038/ nature 14539.
54. Lundberg SM, Erion G, Chen H, DeGrave A, Prutkin JM, Nair B, et al. From local explanations to global

understanding with explainable AI for trees. Nat Mach Intell. 2020;2:56–67. doi:10.1038/s42256-019-0138-9.

55. Deb K, Pratap A, Agarwal S, Meyarivan T. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Trans Evol Comput. 2002; 6(2): 182–197. doi:10.1109/ 4235.996017.
56. Blank J, Deb K. pymoo: multi-objective optimization in Python. IEEE Access. 2020; 8:89497–89509. doi:10.1109/ ACCESS. 2020. 2990567.