Future Generation Computer Systems 174 (2026) 107935

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Towards a widespread usage of computational fluid dynamics simulations for automated virtual nasal surgery planning

Mario Rüttgers [a], Moritz Waldmann [b], Fabian Hübenthal [b], Klaus Vogt [c], Makoto Tsubokura [d,e], Sangseung Lee [a], Andreas Lintermann [f,*]

[a] Data-Driven Fluid Engineering (DDFE) Laboratory, Inha University, Incheon, South Korea
[b] Institute of Aerodynamics and Chair of Fluid Mechanics (AIA), RWTH Aachen University, Germany
[c] Faculty of Medicine, Center of Experimental Surgery, University of Latvia, 1586 Riga, Latvia
[d] Department of Computational Science, Graduate School of System Informatics, Kobe University, Japan
[e] Complex Phenomena Unified Simulation Research Team, RIKEN Center for Computational Science, Kobe, Japan
[f] Jülich Supercomputing Centre (JSC), Forschungszentrum Jülich GmbH, Germany

ARTICLE INFO

Keywords:
Septoplasty
Turbinectomy
Non-linear regression

ABSTRACT

Efficient computational approaches are crucial for advancing computational fluid dynamics (CFD)-based automated planning in nasal surgeries, such as septoplasties and turbinectomies. This study introduces a hybrid lattice-Boltzmann and level-set method to address the trade-off between computational cost and automation. By interpolating geometry variations in discrete steps between pre-surgical and target states, the approach achieves computational efficiency with only 21 surface variations per intervention. Previous methods rely on more costly coupling strategies, such as reinforcement learning or thermal modeling, which may still be appropriate for complex planning scenarios involving multiple intervention sites or thermal flow analysis. In contrast, the presented method reduces complexity while retaining key predictive capabilities, making it particularly suitable for widespread, time-sensitive clinical use focused on a single surgical intervention. Fluid mechanical metrics, including pressure loss and volume flow rate balance, are evaluated alongside tissue removal volume to recommend optimized surgical plans. Case studies on three patients demonstrate tissue savings of 12–25% without compromising key flow parameters. Additionally, a non-linear regression model trained on as few as 11 CFD simulations predicts pressure loss and flow rates with errors below 4%, and reduces computational costs by 50%. The proposed framework represents a significant step toward making CFD-based virtual nasal surgery planning more accessible and practical.

1. Introduction

Computational fluid dynamics (CFD) simulations are experiencing an increasing popularity for evaluating surgeries of respiratory diseases. In [1], it was stated that neither rhinometry nor computed tomography (CT) can adequately quantify nasal airflow pattern changes following surgery. The study demonstrated the feasibility of assessing changes in nasal airflow dynamics following partial middle turbinate resection using CFD techniques. Flow simulations were conducted for quasi-steady laminar nasal airflow with the Fluent©[1] flow solver and meshes with 1.7–1.9 million cells. The 3D models of the pre- and post-surgical nasal cavities were reconstructed from CT data. It was found that the partial middle turbinate resection results in a shift of regional airflow towards the area of the removal with a resultant decreased airflow velocity, decreased wall-shear stress, and increased local air pressure. Similarly, in [2], the fluid mechanical effects of miniscrew-assisted rapid maxillary expansion (MARME) treatment on respiratory flow and the breathing capability were analyzed numerically. In addition to the fluid mechanical quantities mentioned in the previous study, the heating capability was investigated with a thermal lattice-Boltzmann (LB) method. The results confirmed that the respiratory resistance and the average wall-shear stress decrease after the treatment, whereas the heating capability deteriorates. In [3], the steady inspiratory airflow of eight patients with chronic nasal obstruction was studied pre- and post-operatively by analyzing the heat transfer from the mucous membrane. Therefore, CFD calculations based on patient-specific cone beam CT images were performed. The patients had enlarged inferior turbinates and radiofrequency thermal ablation treatment was applied to the anterior parts of the inferior turbinates on both sides. The OpenFOAM [4] flow

* Corresponding author.
E-mail address: a.lintermann@fz-juelich.de (A. Lintermann).
[1] https://www.ansys.com/products/fluids/ansys-fluent.

https://doi.org/10.1016/j.future.2025.107935
Received 7 February 2025; Received in revised form 29 April 2025; Accepted 25 May 2025
Available online 11 June 2025

solver was used for CFD calculations with a laminar and incompressible flow assumption and meshes containing between 4–6 million cells. It was found that the heat transfer from the anterior parts of the inferior turbinates, where surgical interventions were performed, decreased significantly. In [5,6], the physical quantities used to evaluate surgeries were extended by relative humidity, an important indicator for the air conditioning ability of the nose [7]. Both studies reported streams of much colder air in the nasopharynx following turbinate reduction, which matched with each patient's perception.

Apart from estimating an already conducted surgery's outcome, recently, CFD simulations have also been employed for virtually planning nasal surgeries. In [8], virtual surgeries and post-operative results were compared to the pre-operative breathing condition for two patients with nasal obstructions requiring septoplasty. By adding air voxels or performing millimetric translations in the segmentation files of CT data, new shapes were created aiming to simulate a realistic surgical procedure. No statistical difference between the results of the virtual surgery and the post-operative flow regarding heat flux, wall-shear stress, total pressure, and temperature were observed. In contrast, pre-operative measurements were significantly different for the heat flux, total pressure, and the temperature, but not for the wall-shear stress. In [9], CFD and 3D virtual surgery techniques were used to enhance the precision of nasal surgery and to optimize patients' future surgery outcomes. Laminar CFD simulations were carried out with OpenFOAM using meshes that consist of around 9 million tetrahedral cells. The virtual surgeries were conducted using an integrated module within Flowgy©,[2] a software enabling 3D virtual surgery to be performed directly on the patient's 3D model. It was found that both integration of CFD and 3D virtual surgery techniques in otolaryngology can substantially reduce variability in surgical planning and decision-making, ultimately leading to improved patient outcomes. In [10], the implications of congenital nasal pyriform aperture stenosis (CNPAS) on neonatal nasal airflow were investigated with CFD simulations. The virtual surgery was conducted by manual modifications of the segmentation file, resulting in a geometry that mirrors a post-operatively corrected patient. In the CNPAS model, airflow dynamics underwent discernible alterations, with a marked pressure drop around the nasal valve and diminished velocities.

The aforementioned studies about virtual surgery planning have in common that only single virtually planned 3D models based on manual modifications of the segmentation files generated from the pre-surgical CT images were investigated. In contrast, the work in [11] proposed a method that goes beyond the analysis of single modifications. It considers multiple geometry variations that are between the pre-surgical 3D model and a surgeon's initial idea for a desired surgical outcome. Two surgical interventions were analyzed, i.e., the straightening of a deviated septum near the nostrils and the removal of a bony spur further inside the nose. The pre-surgical 3D model was generated with the help of the machine learning (ML)-based pipeline described in [12], where convolutional neural networks (CNNs) segment the upper airways and prepare in- and outflow regions for boundary condition prescription. To obtain the geometry that represents the surgeon's initial idea, the pre-surgical segmentation file was modified by the surgeon using the open-source software 3DSlicer [13] and then by-passed into the pipeline in [12]. The geometry variations were computed with a level-set (LS) approach that interpolates between the pre-surgical state and the state based on the surgeon's initial idea. The approach allowed to monitor the simulation results while the pre-surgical state was iteratively adjusted towards the surgeon's initial idea. It was reported that in the course of the virtual surgery the nasal resistance was reduced by 25.3%, whereas the heating capability was retained despite the geometry modification.

[2] https://www.flowgy.com.

The method in [11] has two drawbacks: (i) It is expensive, since the LS field needs to be computed for each time step. In the current study, it is observed that computing a time step of the coupled LB-LS approach can take a factor of 2.5 longer than using the LB method alone. (ii) If a user prefers to monitor changes at more than one location, they can only be monitored sequentially. However, there might be a combination of changes at both locations that yields a better surgical outcome than the outcome of the states that are monitored when changing the surfaces sequentially, one after another.

In [14], both challenges were addressed by coupling the LB method to a reinforcement learning (RL) algorithm. First, the RL agent can change the LS field with discrete actions in a pre-defined time step interval, e.g., every 25,000 time steps. Second, if two separate surgery locations are investigated, the RL agent learns to find the optimal combination of changes at the two locations based on feedback in terms of the time-averaged pressure loss and temperature increase between the inlets (nostrils) and outlet (pharynx) after each modification. The method was demonstrated for two patients, the first one suffering from a deviated septum and a bony spur, and the second one from enlarged turbinates. The simulation domain of the first patient was resolved by about $110 \cdot 10^6$ cells, and the domain of the second patient by about $220 \cdot 10^6$ cells, using mesh resolutions of $\Delta x = 0.1\,\mathrm{mm}$ to accurately resolve narrow channels and thin boundary layers [15,16]. For the first time, large-scale CFD simulations with grids on the order of hundreds of million cells were coupled to an RL algorithm. For equal weights on the pressure loss and temperature increase, the algorithm proposed a slightly weaker correction of the deviated septum, compared to the surgeon's plan. It also suggested to keep the bony spur in case of the first patient. In case of the second patient, it recommended to nearly completely remove the inferior turbinate and moderately reduce the middle turbinate. However, training only in a single RL environment entails the risk of ending up in local optima, and the number of possible surface variations and the large grids make this automated approach computationally expensive.

To overcome these limitations, the method described in [14] was improved in [17] by training the RL agents in multiple environments in parallel and replacing some of the expensive CFD computations by predictions of a Gaussian process regression (GPR) model. The former allows the agents to share their experiences in pre-defined intervals and jointly search for the global optimum. For the latter, the automated virtual surgery procedure is subdivided into two stages. In the first stage, the GPR model does not interact and feedback for the RL algorithm is solely provided by the CFD solver [14]. When a pre-defined number of CFD computations have been computed, they are used as training data for the GPR model in the second stage. That is, whenever the RL agent explores a new geometry variation, the GPR model is trained with the so far computed CFD data. With the help of the uncertainty quantification of the GPR model it is decided whether to use the GPR prediction or a new CFD computation to determine the pressure loss and temperature increase. It was found that employing a parallel RL algorithm improves the reliability of the surgery planning tool in finding the global optimum. However, parallel training also led to a larger number of geometry variations that need to be computed by the CFD solver. This overhead was compensated by replacing some of the computations with the GPR algorithm, i.e., around 6% of the computations could be saved without significantly degrading the predictions' accuracy. Nevertheless, the approach is still expensive and can be taken into consideration for specific patients that need to undergo complex surgical interventions, i.e., at two or more separate locations inside of the nasal cavity. There is still a need for a cheaper automated approach that is similarly accurate and can be employed to a wide set of applications, e.g., for receiving first estimates of a surgical plan for a single surgical intervention location.

The current study presents such a method. That is, fluid mechanical properties are computed based on converged simulations in discrete interpolation steps only for single interventions. If an intervention

involves geometry changes in several locations, they are carried out simultaneously. The low number of surface variations makes the coupling to the parallel RL algorithm obsolete, which is why it is excluded here. Furthermore, instead of expensive thermal computations that can increase computational costs by a factor of 2, the investigation focuses on the pressure loss between the inflow (nostrils) and outflow (pharynx) regions, and balanced volume flow rates through the left and right nasal passages.

It should be noted that thermal effects such as volume expansion due to density changes or viscosity changes are not negligible. Their physiological relevance has been addressed in detail in [2,11,14,17]. The corresponding tools remain available and can be employed when a deeper analysis is required, e.g., including thermal flow and optimization for two separate surgery locations. In contrast, the presented method aims to reduce complexity while retaining key predictive capabilities, making it particularly suitable for a widespread, time-sensitive clinical use that focuses on single surgical interventions. Ultimately, the choice of whether or not to include temperature should depend on the specific needs and constraints of the patient or physician.

It is further investigated, how replacing some of the CFD simulations by predictions from non-linear regression models can save computational resources and, hence, speed up surgery planning. This investigation includes the analysis of interpolations between the pre-surgical geometry and the geometry based on the surgeon's initial plan, generated by the regression model.

An important factor that has not yet been sufficiently studied is the marginal benefit of a planned surgical intervention. Whereas the previously mentioned studies focus on finding geometry variations that guarantee optimal fluid mechanical parameters, the volume of the removed tissue has not yet been taken into consideration. The proposed method considers the marginal benefit of, e.g., further decreasing pressure loss, by setting it into relation to the additional tissue that needs to be removed. For example, it might be the case that the benefit from an acceptable pressure loss to the global optimum is small, such that an additional removal of tissue is not justifiable. A surgeon's responsibility is to remove as less tissue as possible, since her/his actions are mostly irreversible. This is analyzed for three patients in this work.

The manuscript is structured as follows. The medical data of the patients are detailed in Section 2, and the computational methods are presented in Section 3. Section 4 shows the results of the automated virtual surgery planning, Section 5 provides a summary and discussion, and Section 6 concludes the study.

2. Medical data

Anonymized CT data of three patients are used. The patients gave informed consent for inclusion of the data in the current study. The first patient suffers from a deviated septum and a bony spur (Case A), as shown in Fig. 1, which has been investigated in [11,14,17]. The bounds of a surgeon's initial idea of the desired post-operative result are illustrated by grey 3D regions inside of the transparent-blue representations of the nasal cavity. In the close-up views, representative cross-sections of the nasal cavity are juxtaposed to their corresponding CT images. Apart from the boundary of the pre-surgical state, which is visualized by the black contour, Fig. 1 also shows the boundary of the geometry based on a surgeon's initial idea of the post-operative result, highlighted by the red contour. This idea is an extension of the right nasal passage alongside a reduction of the left nasal passage in the vicinity of the nostrils, and an extension of the left nasal passage at the middle turbinate downstream. The CT data of the first patient are composed of 119 axial slices with 512×512 pixels each. The pixel spacing is 0.5 mm, and the space between the axial slices is 0.7 mm. The patient has been treated at Klinika Headline, Riga, Latvia, and the CT data have been obtained by a PHILLIPS Brilliance 16 CT scanner at Diamed ARS MEDICAL, Riga, Latvia.

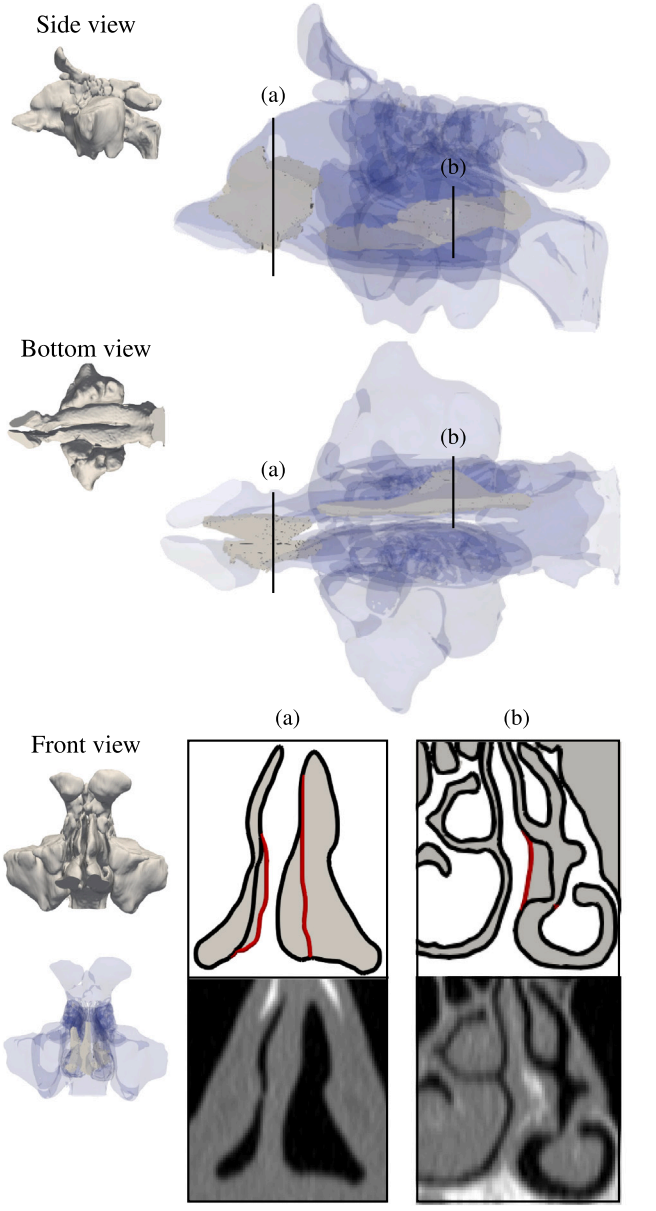

Fig. 1. The nasal cavity of the first patient, suffering from a deviated septum and a bony spur (Case A), which was investigated in [11,14,17]. The bounds of a surgeon's initial idea of the desired post-operative result are illustrated by the grey 3D regions inside of the transparent-blue representations of the nasal cavity. The close-ups (a) and (b) on the right side correspond to the lines depicted in the 3D images above. They illustrate boundaries at representative cross-sectional areas, which are juxtaposed to the corresponding pre-surgical CT images below. The black and red contours represent the boundaries of the pre-surgical structure and the geometry based on a surgeon's initial idea of the desired post-operative result.

The second patient has been diagnosed enlarged inferior and middle turbinates (Case B), illustrated in Fig. 2, which has been analyzed in [14,17]. The bounds of a total turbinectomy, a complete removal of the turbinates, are illustrated by the grey 3D regions inside of the transparent-blue representations of the nasal cavity. A representative cross-section of the nasal cavity is juxtaposed to its corresponding CT image. The boundary of the pre-surgical state and the total turbinectomy are highlighted by the black and red contours. The CT recordings of the second patient have 103 axial slices, again with 512×512 pixels each. The pixel spacing is 0.326 mm, and the space between the axial slices is 1.0 mm. The patient has been treated at HNOmedic, Bornheim,

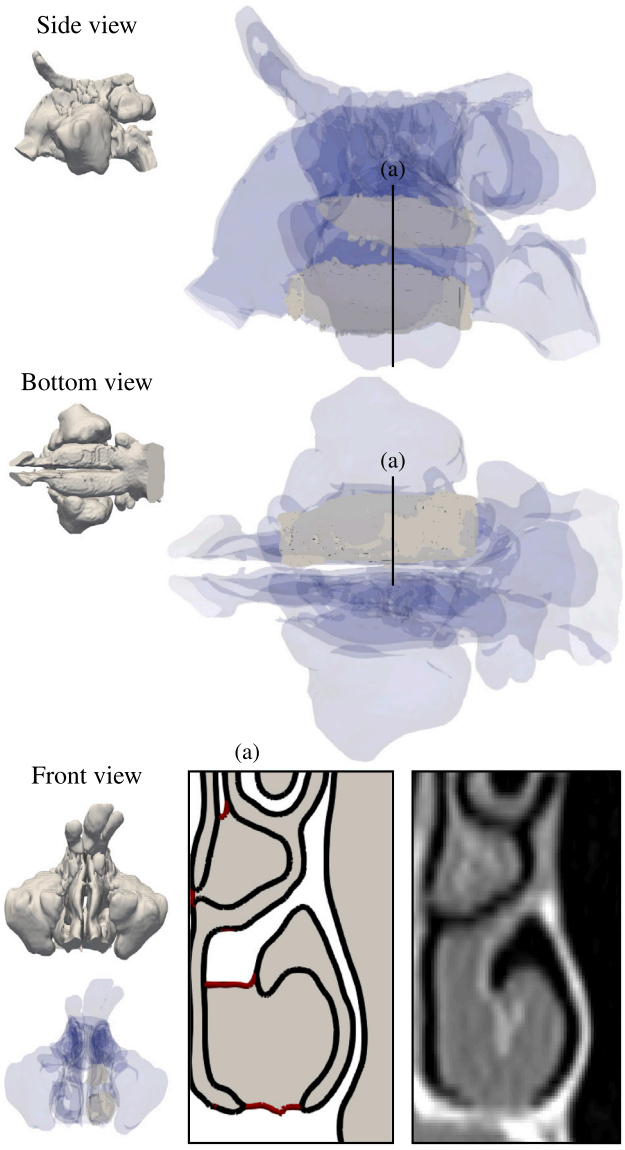

Fig. 2. The nasal cavity of the second patient, suffering from enlarged inferior and middle turbinates (Case B), which was investigated in [14,17]. The bounds of a total turbinectomy are illustrated by the grey 3D regions inside of the transparent-blue representations of the nasal cavity. The close-up on the left side in (a) illustrates boundaries at a representative cross-sectional area, which is juxtaposed to the corresponding pre-surgical CT image on the right. The black and red contours represent the pre-surgical structure and the geometry based on the total turbinectomy. The location of the cross-section is displayed above in the 3D images.

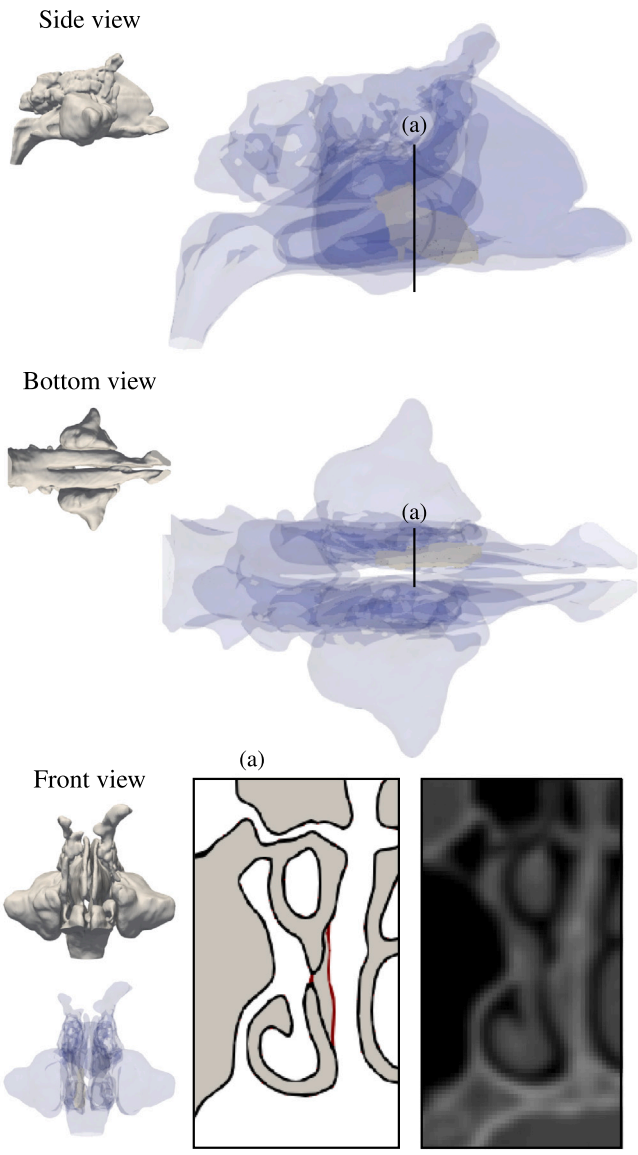

Fig. 3. The nasal cavity of the third patient, suffering from a deviated septum that blocks the right nasal passage (Case C). The bounds of a surgeon's initial idea of the desired post-operative result are illustrated by the grey 3D regions inside of the transparent-blue representations of the nasal cavity. The close-up on the left side of (a) illustrates boundaries at a representative cross-sectional area, which is juxtaposed to the corresponding pre-surgical CT image on the right. The black and red contours represent the boundary of the pre-surgical structure and the geometry based on a surgeon's initial idea of the desired post-operative result. The location of the cross-section is displayed above in the 3D images.

Germany, and the CT data have been obtained by a SIEMENS Somatom Emotion 16 CT scanner at Radiomedicum, Frankfurt, Germany.

The third patient suffers from a deviated septum (Case C), as shown in Fig. 3. Similar to the other patients, the bound of a surgeon's initial idea of the desired post-operative result are illustrated by the grey 3D regions inside of the transparent-blue representations of the nasal cavity, a representative cross-section of the nasal cavity is juxtaposed to its corresponding CT image, and the boundary of the pre-surgical state and the surgeon's initial idea of the post-operative result are stressed by the black and red contours. The CT recordings of the second patient have 146 axial slices, again with 512×512 pixels each. The pixel spacing is 0.465 mm, and the space between the axial slices is 0.7 mm. The patient has been treated at Uniklinikum Aachen, Germany, and the CT data have been obtained by a SIEMENS Somatom Definition AS CT scanner.

For all patients, the 3D model of the pre-surgical upper airway is generated based on the digital imaging and communications in medicine (DICOM) files of the CT data with the following steps:

1. Pre-filtering of the CT images to better identify the air–tissue interface
2. Segmentation of the upper airway by a 2D CNN
3. Segmentation of the nostrils by a 3D CNN
4. Generating a 3D nasal cavity model from the segmentation
5. Identifying inflow areas and their normal vectors by a 2D CNN
6. Identifying the outflow area and its normal vector with the centerline at the pharynx
7. Cutting the 3D model at the inflow and outflow regions and close the holes with flat surfaces

8. Distributing the 3D model's triangles to inflow, wall, and outflow regions

Pre-filtering is done with a convolutional filter for edge detection [18] and a gradient anisotropic diffusion filter [19] to reduce image noise without removing edges. The CNNs for steps 2 and 3 have been trained in a supervised manner to segment all voxels that represent the human airway. Two steps are required, since the most challenging segmentation task is separating ambient air from the airway at the nostrils. In step 4, a triangulated 3D model is generated with a marching cubes algorithm [20], which is smoothed by a windowed sync filter [21]. In steps 5–7, flat surfaces of the inflow and outflow regions are generated. This is necessary to provide clear extrapolation directions at the inlets and outlet for the boundary conditions described in Section 3.2. The final step 8 is required to automatically define different boundary segments and start the simulation. The complete pipeline and more details are described in [12].

For the first and the third patient, surgery planning of a septoplasty is investigated, and for the second patient, the surgical potential of a turbinectomy is analyzed. More details about the medical data, types of surgeries, and motivation from a medical background are given in [11,14,17].

3. Computational methods

The numerical methods used to simulate the respiratory flows are described in Section 3.1, the flow and boundary conditions for the CFD simulations are presented in Section 3.2, and the LS approach for obtaining surface variations is explained in Section 3.3. This is followed by providing details on the non-linear regression model in Section 3.4.

3.1. Numerical flow simulations

Unstructured, hierarchical, uniformly-refined Cartesian meshes are created using the massively parallel mesh generator from multiphysics Aerodynamisches Institut Aachen (m-AIA) [22]. The mesh generation relies on an octree structure formed through iterative subdivision of an initial cube that encloses the region of interest (ROI) [23], i.e., in this case the nasal cavity. Starting with the initial cube, successive refinements divide each cube into eight sub-cubes over a specified number of refinement levels. The octree structure emerges from the parent–child relationships between cubes and sub-cubes. Cells outside the ROI are removed to optimize the mesh.

Below a specific refinement level, the cells are ordered along a Z-curve [24]. Above this level, the domain is partitioned using a Hilbert space-filling curve decomposition method [25]. The Hilbert curve preserves spatial locality, meaning that adjacent cells in the computational domain are likely assigned to the same process due to the curve's self-similar structure. As a result, neighboring subdomains tend to reside on the same compute node or on adjacent nodes, enhancing process locality within and across nodes. Since the distribution is balanced from the beginning and stays fixed, node imbalance does not appear. In any case, when adaptive mesh refinement is required, node imbalance is treated by an efficient load-balancing strategy presented in [26]. The final mesh is efficiently stored in Network Common Data Form (NetCDF) format using parallel I/O routines [27].

CFD simulations are conducted with the LB method of the m-AIA solver framework. It is well known that LB methods feature some major advantages over finite-element or finite-volume methods when highly intricate geometries are considered [28]. Since the main idea of this study is to demonstrate a tool that can be employed efficiently on a widespread basis, the functionalities that have been developed in [11,14,17] need to be reduced to save computational efforts. For this purpose, the computation of the temperature distribution is neglected and the simulation only considers the velocity and pressure fields, which is sufficient to estimate the condition of a patient and potential surgery outcomes.

Table 1
Geometry and simulation parameters.

Case	Re_P	d_P [m]	A_P [m^2]	A_i [m^2]	A_r [m^2]	No. of cells
A	960	$11.45 \cdot 10^{-3}$	$20.13 \cdot 10^{-5}$	$13.56 \cdot 10^{-5}$	$15.10 \cdot 10^{-5}$	$110 \cdot 10^6$
B	766	$16.19 \cdot 10^{-3}$	$32.44 \cdot 10^{-5}$	$6.04 \cdot 10^{-5}$	$7.13 \cdot 10^{-5}$	$220 \cdot 10^6$
C	1390	$10.54 \cdot 10^{-3}$	$11.63 \cdot 10^{-5}$	$7.06 \cdot 10^{-5}$	$5.08 \cdot 10^{-5}$	$185 \cdot 10^6$

The LB module solves the discretized form of the Boltzmann equation with the Bhatnagar–Gross–Krook (BGK) approximation of the right-hand side collision process [29], i.e.,

$$f_i(\boldsymbol{x} + \boldsymbol{\xi_i}\delta t, t + \delta t) - f_i(\boldsymbol{x}, t) = -\omega(f_i(\boldsymbol{x}, t) - f_i^{eq}(\boldsymbol{x}, t)) \tag{1}$$

is solved for the particle probability distribution functions (PPDFs) f_i at neighboring fluid cells at locations $\boldsymbol{x} + \boldsymbol{\xi_i}\delta t$. They are functions of the location vector $\boldsymbol{x} = (x_1, x_2, x_3)^T$, the discrete molecular velocity vector $\boldsymbol{\xi_i} = (\xi_1, \xi_2, \xi_3)^T$, and the time and time increment t and δt. The collision frequency is expressed by ω. The advantage over the methods in [11,14,17], where a second set of PPDFs is used to model the temperature in form of a passive scalar transport equation, is that the exclusion of the temperature distribution requires to only solve for one set of PPDFs, i.e., reducing the computational effort.

The discretization is based on the D3Q27 model [30], with $i \in \{1, \dots, Q\}$ directions in 3D for f_i and $Q = 27$. The D3Q27 model is chosen because it conserves rotational invariance in laminar flow and produces much smaller errors in transitional or turbulent regimes compared to other three-dimensional lattice discretization models such as the D3Q19 model [31,32]. The discrete Boltzmann–Maxwellian distribution function is defined as

$$f_i^{eq} = wc_i\rho\left(1 + \frac{\boldsymbol{\xi_i}\cdot\boldsymbol{u}}{c_s^2} + \frac{1}{2}\left(\frac{\boldsymbol{\xi_i}\cdot\boldsymbol{u}}{c_s^2}\right)^2 - \frac{\boldsymbol{u}\cdot\boldsymbol{u}}{2c_s^2}\right), \tag{2}$$

with the isothermal speed of sound $c_s = 1/\sqrt{3}$, density ρ, fluid velocity vector $\boldsymbol{u} = (u, v, w)^T$ with the velocity components u, v, and w in the x_1-, x_2- and x_3-directions, and weight coefficients wc_i [30]. The macroscopic variables ρ and $\boldsymbol{u}$ can be computed by

$$\rho = \sum_{i=1}^{Q} f_i, \tag{3}$$

$$\rho\boldsymbol{u} = \sum_{i=1}^{Q} \boldsymbol{\xi_i} \cdot f_i. \tag{4}$$

The static pressure p_{stat} is obtained from the density by $p_{stat} = c_s^2\rho$.

3.2. Flow and boundary conditions

For the simulations, the following boundary conditions are employed. At the inlets, i.e., at the nostrils of the nasal cavity, the equation of Saint-Venant and Wantzel is applied [15] to compute ρ, and the velocity is extrapolated from the inner cells. At the pharynx outlet, the pressure is iteratively adapted to fit a prescribed volume flux of $\dot{V}_P = 250\,\mathrm{mL/s}$ [15], yielding steady inflow conditions at inspiration. For the first patient, this corresponds to a Reynolds number of $Re_P = (\dot{V}_P \cdot d_P)/(A_P \cdot \nu) = 960$, calculated from the hydraulic diameter of the pharynx $d_P = 11.45 \cdot 10^{-3}$ m, the pharyngeal cross-sectional area $A_P = 20.13 \cdot 10^{-5}$ m^2, and the kinematic viscosity of air $\nu = 1.63 \cdot 10^5$ m^2/s. For the second patient it is $Re_P = 766$ with $d_P = 16.19 \cdot 10^{-3}$ m and $A_P = 32.44 \cdot 10^{-5}$ m^2, and for the third patient it is $Re_P = 1390$ with $d_P = 10.54 \cdot 10^{-3}$ m and $A_P = 11.63 \cdot 10^{-5}$ m^2. The resulting simulation parameters are summarized in Table 1. The velocity is extrapolated from the inner cells. An interpolated bounce-back scheme is used to satisfy the no-slip condition at the inner walls of the nasal cavity [33].

The simulation domains of the three patients are resolved by about $110 \cdot 10^6$, $220 \cdot 10^6$, and $185 \cdot 10^6$ cells, using mesh resolutions of $\Delta x = 0.1$ mm to accurately resolve narrow channels and thin boundary layers. The decision is based on the mesh refinement study presented in [15], which was conducted with the D3Q27 model of m-AIA's LB method. Simulation results for nasal cavity flows with a cell size of $\Delta x = 0.093569$ mm and $92.6 \cdot 10^6$ cells were compared to the results with a cell size of $\Delta x = 0.046784$ mm and $724 \cdot 10^6$ cells. The same volume flux at the pharynx of $\dot{V}_P = 250\,\mathrm{mL/s}$ as in the current study was prescribed. The various resolutions resulted in only a small difference of the total pressure loss of 2.1‰ and 0.8‰ for the left and right nasal cavity. For a more detailed comparison of velocity profiles and the power spectral density for the different resolutions, the reader is referred to [15].

3.3. Level-set approach for varying the surface

Automated surgery planning is realized with a coupled LB-LS approach. In the LS method, the geometry is represented by the signed distance function φ, the LS field. That is, for each cell of the computational mesh, the minimum distance to the geometry surface is calculated. If the cell center is located outside the geometry, the value of the corresponding cell is multiplied by -1, resulting in $\varphi < 0$. Hence, the surface of the geometry is represented by $\varphi = 0$. To perform a virtual surgery, three LS fields are stored. The first LS field φ_1 is calculated using the linear interpolation [11]

$$\varphi_1 = (1 - \alpha)\varphi_2 + \alpha\varphi_3, \tag{5}$$

with the second and the third LS fields φ_2 and φ_3 containing information on the pre-surgical shape and the shape generated based on the surgeon's initial plan. The factor $\alpha \in [0, 1]$ defines the interpolation between the pre-surgical state ($\alpha = 0$) and the state based on the surgeon's plan ($\alpha = 1$).

The LS field φ_1 is changed in discrete steps of $|\Delta\alpha| = 0.05$ to propose geometry variations that account for the precision of standard surgical tools [14]. Unlike the investigations in [14,17], where for each patient combinations of two surgical interventions were analyzed, the current study analyzes one intervention per patient. This yields a reduction from 441 potential combinations of surface variations that have to be explored in the optimization procedure to only 21 geometry variations per patient. Thus, computational costs for recommending a surgical intervention are drastically reduced, which paves the way for a widespread usage of the surgery planning tool.

To avoid surgically impossible combinations of α, the surgeon must guarantee that the modifications of the pre-surgical segmentation in 3D Slicer (see Section 2) yield a surgically possible 3D model, and that all states between the pre-surgical state and the state based on the surgeon's modification are surgically feasible. After φ_1 is calculated, it is transferred to the CFD solver, where the boundary cells, i.e., cells containing the isosurface at $\varphi_1 = 0$, are determined. Examples of the LS fields φ_1 at $\alpha = 0.6$, φ_2, φ_3, and the corresponding contours for $\varphi_1 = 0$ (green), $\varphi_2 = 0$ (black), and $\varphi_3 = 0$ (red) are shown for the third patient in Fig. 4. To allow for an accurate computation of the nasal cavity's surface, the grid of the maximum bound of the virtual surgery (light grey) is thickened by an outer layer (dark grey). Inside the dotted rectangle, the LS fields are computed for every cell. Outside the rectangle, the LS fields are only needed to define the nasal cavity's surface. They are hence only computed along the boundary. This reduces the computational effort at the first time step when the LS fields are computed for the first time. During the virtual surgery, only the LS field in the dotted rectangle is updated.

In contrast to [11], where α, and hence φ_1, is modified every time step, here α and φ_1 are only updated every $N = 25{,}000$ time steps (δt) to guarantee that time-averaging is performed on a converged, yet still unsteady flow field. This saves computing time compared to the original method in [11]. Details about the choice for $N = 25{,}000$ time steps and the numerical threshold for convergence are provided in Section 3.4.

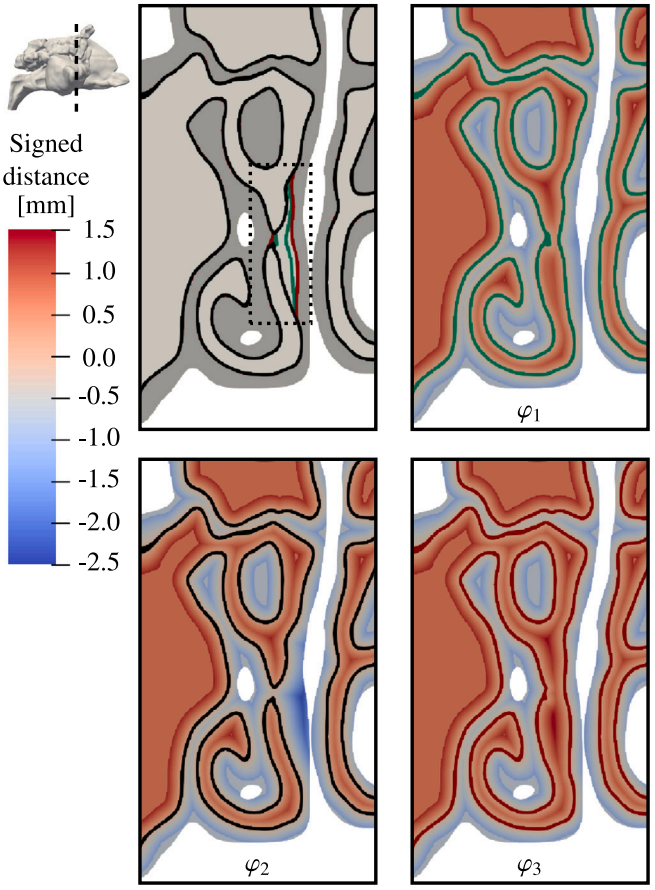

Fig. 4. LS fields φ_1 at $\alpha = 0.6$, φ_2, and φ_3 and the corresponding contours for $\varphi_1 = 0$ (green), $\varphi_2 = 0$ (black), and $\varphi_3 = 0$ (red) of the third patient for the same cross-sectional area that is shown in Fig. 3. The light grey area stands for the maximum bound of the virtual surgery, and the dark grey area shows the thickened mesh.

3.4. Non-linear regression with polynomial features

Non-linear regression models are essential for capturing complex relationships between independent variables and a dependent variable that cannot be adequately described by a simple linear function. The independent variable is α, and two dependent variables are predicted. The first one is the normalized pressure loss Δp_{norm} between the inlets (nostrils) and the outlet (pharynx)

$$\Delta p_{norm} = \frac{\Delta p(\alpha) - \Delta p(\alpha = 1)}{\Delta p(\alpha = 0) - \Delta p(\alpha = 1)}. \tag{6}$$

The pressure loss Δp is defined as

$$\Delta p = p_{out} - \frac{1}{2}(p_{in,l} + p_{in,r}), \tag{7}$$

with the total pressure at the outlet p_{out}, and at the two inlets $p_{in,l}$ and $p_{in,r}$ being defined as follows:

$$p_{out} = \frac{1}{N_a} \sum_{i=N-N_a}^{N} \left(\frac{1}{H_{out}} \sum_{j=0}^{H_{out}} p^i_{tot,j} \right), \tag{8}$$

$$p_{in,l} = \frac{1}{N_a} \sum_{i=N-N_a}^{N} \left(\frac{1}{H_{in,l}} \sum_{j=0}^{H_{in,l}} p^i_{tot,j} \right), \tag{9}$$

$$p_{in,r} = \frac{1}{N_a} \sum_{i=N-N_a}^{N} \left(\frac{1}{H_{in,r}} \sum_{j=0}^{H_{in,r}} p^i_{tot,j} \right). \tag{10}$$

Here, $H_{in,l}$, $H_{in,r}$, and H_{out} are the number of boundary cells at the inflow and outflow regions. A feedback loop takes $N = 25{,}000$ time

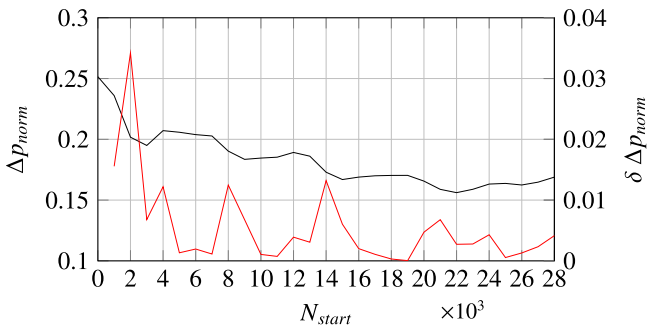

Fig. 5. Δp_{norm} for the geometry change from $\alpha = 0.4$ to $\alpha = 0.35$ of Case B with different start time steps $N_{start} \in \{0, 1{,}000, 2{,}000, \ldots, 28{,}000\}$ of the averaging period.

steps and the feedback is averaged over the last $N_a = 10{,}000$ time steps. Note that p_{tot} is the total pressure, expressed as the sum of the static pressure p_{stat} and the dynamic pressure $p_{dyn} = \rho u_{mag}^2/2$, where u_{mag} stands for the velocity magnitude.

The rational behind choosing $N = 25{,}000$ and $N_a = 10{,}000$ is explained with the help of Fig. 5. The figure shows Δp_{norm} for the geometry change from $\alpha = 0.4$ to $\alpha = 0.35$ of Case B with different start time steps $N_{start} \in \{0, 1{,}000, 2{,}000, \ldots, 28{,}000\}$ of the averaging period. This geometry change has been selected as it features the largest change in active/inactive cells among all changes for the three patients. The quantity Δp_{norm} is chosen as criterion as it is known to adapt slower to changes in the flow configuration than, e.g., the velocity field. The velocity is determined by local momentum exchange, while density, being a conserved quantity, requires time to equilibrate through slower modes [34]. This difference arises from the LB method's separation of scales, where momentum adjusts rapidly but mass conservation leads to slower global density adaptation. The numerical threshold for convergence is defined with the help of $\delta\Delta p_{norm}$, which is the absolute change of Δp_{norm} for two consecutive averaging periods with a gap of $\Delta N_{start} = 1{,}000$. For $0 < N_{start} < 15{,}000$, the flow field adapts to the change from $\alpha = 0.4$ to $\alpha = 0.35$ and $\delta\Delta p_{norm}$ fluctuates. After that, Δp_{norm} starts to converge. For the simulations of the current study, the convergence threshold is defined as $\delta\Delta p_{norm} < 0.01$ for 10 consecutive averaging periods. This is achieved at $N_{start} = 15{,}000$ for the first time.

The second dependent variable is the normalized volume flow rate $\dot{V}_{norm} = \dot{V}_l/\dot{V}_r$, where $\dot{V}_l$ and $\dot{V}_r$ represent the volume flow rates through the left and right nostrils from a patient's view. They are defined by

$$\dot{V}_l = \frac{A_l}{N_a} \sum_{i=N-N_a}^{N} \left(\frac{1}{H_{in,l}} \sum_{j=0}^{H_{in,l}} (u_j^i \cdot n_{l,x_1} + v_j^i \cdot n_{l,x_2} + w_j^i \cdot n_{l,x_3}) \right), \tag{11}$$

$$\dot{V}_r = \frac{A_r}{N_a} \sum_{i=N-N_a}^{N} \left(\frac{1}{H_{in,r}} \sum_{j=0}^{H_{in,r}} (u_j^i \cdot n_{r,x_1} + v_j^i \cdot n_{r,x_2} + w_j^i \cdot n_{r,x_3}) \right). \tag{12}$$

with $(n_{l,x_1}, n_{l,x_2}, n_{l,x_3})^T$ and $(n_{r,x_1}, n_{r,x_2}, n_{r,x_3})^T$ representing the normal vectors, and A_l and A_r the areas of the left and right inflow regions (see Table 1).

Among non-linear regression methods, polynomial regression offers a straightforward yet powerful approach by leveraging polynomial transformations of the input data. It involves modeling the relationship between Δp_{norm} or $\dot{V}_{norm}$, here denoted as the function ϕ, and α as an n-degree polynomial

$$\phi \in \{\Delta p_{norm}, \dot{V}_{norm}\} = \beta_0 + \beta_1\alpha + \beta_2\alpha^2 + \cdots + \beta_n\alpha^\zeta + \epsilon, \tag{13}$$

where β_i are the coefficients to be estimated, ϵ represents the error term, and ζ is the degree of the polynomial that determines the complexity of the model, allowing it to capture various levels of non-linearity. The method enables the modeling of non-linear relationships by transforming the input features into polynomial terms, such as $\alpha, \alpha^2, \alpha^3, \ldots, \alpha^\zeta$, and interaction terms for multivariate data. The regression task is realized with the non-linear regression tool with polynomial features of the Scikit-learn Python library [35]. First, polynomial terms are generated with the *polynomial features transformer*. After feature transformation, a linear regression model is applied to the expanded feature space. Although the model remains linear in the transformed space, it effectively represents a non-linear relationship in the original input space.

Table 2
Training data for combinations CB 1, CB 2, and CB 3.

Combination	$\phi(\alpha = \ldots)$
CB 1	0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0
CB 2	0.0, 0.15, 0.3, 0.45, 0.6, 0.75, 0.9, 1.0
CB 3	0.0, 0.2, 0.4, 0.6, 0.8, 1.0

Training data is used to fit the regression model and test data to evaluate the model's performance. The training data are computed with the LB method. In principle, cross-validation should be performed to assess the risk of overfitting and to better separate variance from bias in such a small dataset. However, for this study, classical cross-validation techniques involving randomized or shuffled splits cannot be applied. The 21 available samples, defined by pairs of $\phi \in \{\Delta p_{norm}, \dot{V}_{norm}\}$ and interpolation parameter α, follow a physically continuous sequence of deformations between the pre-surgical shape and the surgeon's initial plan. Random reordering or omission of samples would break this continuity and lead to large, abrupt geometric changes between training points. Such jumps can cause numerical instabilities in the CFD solver, as newly activated cells (due to the LS change) are initialized with zero velocities. If too many cells are changed at a time, these gradients may lead to divergence in the simulations.

Including more than 21 simulations would not make sense, since surface variations based on a smaller step size, or even a continuous action space, could not be realized by a surgeon due to the lacking precision of standard surgical tools. Three different combinations (CB 1, CB 2, and CB 3) of training and test data are investigated. Table 2 shows the training data for each combination.

4. Results

In Section 4.1, surgery planning solely based on CFD simulations is demonstrated for the three patients described in Section 2. In Section 4.2, the potential of replacing some of the CFD computations with predictions from the regression model is analyzed. The simulations were run on the central processing unit (CPU) and graphics processing unit (GPU) partitions of the Jülich Research on Exascale Cluster Architectures (JURECA-DC), Forschungszentrum Jülich [36,37]. Each node of the GPU partition is equipped with four NVIDIA A100 GPUs and two AMD EPYC 7742 CPUs with 64 cores clocked at 2.25 GHz and 512 GB DDR4 memory. Each node of the CPU partition contains two AMD EPYC 7742 processors without acceleration.

The flexibility of m-AIA allows to run the LB method on either CPUs or GPUs. To showcase this flexibility, Cases A and B were run on 16 CPU nodes, i.e., on 2,048 cores, and Case C on 8 GPU nodes, i.e., on 32 GPUs. For Cases A and B, the solver was compiled with the GCC/13.3.0 compiler and OpenMPI/5.0.5 module of the HPC system. For Case C, the solver was compiled with the NVHPC/ 23.7-CUDA-12 compiler and OpenMPI/4.1.5. Since the regression model only consumes a negligible fraction of the total compute, it can run on the same partition as the CFD simulation.

4.1. Surgery planning based on CFD simulations

The results presented in this section are obtained solely by CFD computations. That is, a simulation first runs with $\alpha = 1.0$ for 150,000 time steps δt until a converged flow field is reached, followed by

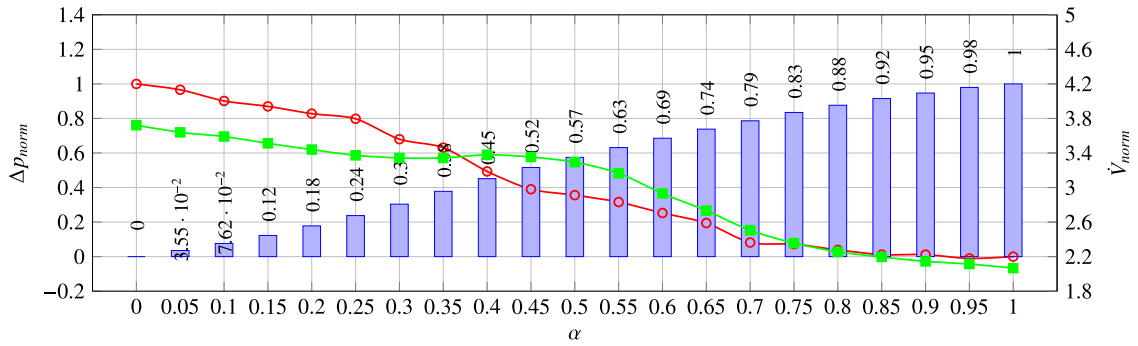

Fig. 6. Δp_{norm} (red), $\dot{V}_{norm}$ (green), and Vol_{norm} (violet) in the course of the virtual surgery for Case A.

reductions of α by $\Delta\alpha = -0.05$ each $N = 25{,}000$ time steps. Choosing $\alpha = 1.0$ as the starting point and $\Delta\alpha = -0.05$ as the decrement is important, since geometry variations with $\Delta\alpha < 0$ converge faster than those with $\Delta\alpha > 0$. This comes from the fact that with $\Delta\alpha > 0$ usually more previously inactive grid cells become active than the other way round. A newly activated cell is initialized with $\rho = 1$ and $\boldsymbol{u} = (0, 0, 0)^T$, which creates large gradients to those cells that were already active before. If the number of activated cells is large, this can lead to numerical instabilities. In contrast, if previously active cells become inactive, the new near-wall flow field needs to be adjusted only slightly, because ρ and $\boldsymbol{u}$ of the new boundary cells were computed under near-wall conditions before the geometry change.

For Case A, Fig. 6 shows Δp_{norm} (red), which is assigned to the left vertical axis, $\dot{V}_{norm}$ (green) following the right vertical axis, and the normalized volume Vol_{norm} (violet bars), which is removed in the course of the virtual surgery for $\alpha \in \{0.0, 0.05, \ldots, 1.0\}$.

The normalized removed volume Vol_{norm} is defined as

$$Vol_{norm} = \frac{Vol(\alpha) - Vol(\alpha = 0)}{Vol(\alpha = 1) - Vol(\alpha = 0)}, \tag{14}$$

where Vol stands for the complete volume of the nasal cavity. Although Δp_{norm} is continuously reduced from $\alpha = 0.0$ to $\alpha = 1.0$, there is only a marginal gain between $\Delta p_{norm}(\alpha = 0.8) = 0.012$ and $\Delta p_{norm}(\alpha = 1.0) = 0.0$. The pre-surgical geometry is characterized with $\dot{V}_{norm}(\alpha = 0.0) = 3.722$ by a large deviation between $\dot{V}_l$ and $\dot{V}_r$. The quantity $\dot{V}_{norm}$ reaches a plateau between $\alpha = 0.25$ and $\alpha = 0.45$, before it continues to decrease steadily. Again, there is only a marginal change between $\dot{V}_{norm}(\alpha = 0.8) = 2.255$ and $\dot{V}_{norm}(\alpha = 1.0) = 2.067$. Based on these results, the surgeon can conclude that $\alpha = 0.8$ is already satisfying, considering that between $Vol_{norm}(\alpha = 0.8)$ and $Vol_{norm}(\alpha = 1.0)$ 12.0% of the tissue volume that was planned to be removed can be saved. In an irreversible intervention like nasal surgeries it lies in the responsibility of the surgeon to remove only as much tissue as needed.

The difference between the pre-surgical structure (black), the geometry based on a surgeon's initial idea of the desired post-operative result (red), and the geometry with $\alpha = 0.8$ is illustrated for the cross-sectional areas I–V in Fig. 7. Two cross-sections (I, II) are located in the anterior part of the nasal cavity, and three cross-sections (III–V) near the bone spur. The difference between the boundaries of $\alpha = 0.8$ and $\alpha = 1.0$ is particularly evident near the bone spur. Instead of the relatively flat surface at $\alpha = 1.0$, a curvy residue remains at $\alpha = 0.8$, which, according to the results in Fig. 6, does not seem to significantly increase the pressure loss.

The differences in the pressure distribution between $\alpha = 0.0$ and $\alpha = 1.0$ have been extensively studied and visualized in [14,17]. However, the velocity fields did only play a minor role in the previous studies, which is why an example of the normalized velocity magnitude u_{mag}/u_{in} for the cross-sectional area II of Case A is provided in Fig. 8. The inflow velocity u_{in} is defined as

$$u_{in} = \frac{1}{2}\left(\frac{\dot{V}_l}{A_l} + \frac{\dot{V}_r}{A_r}\right). \tag{15}$$

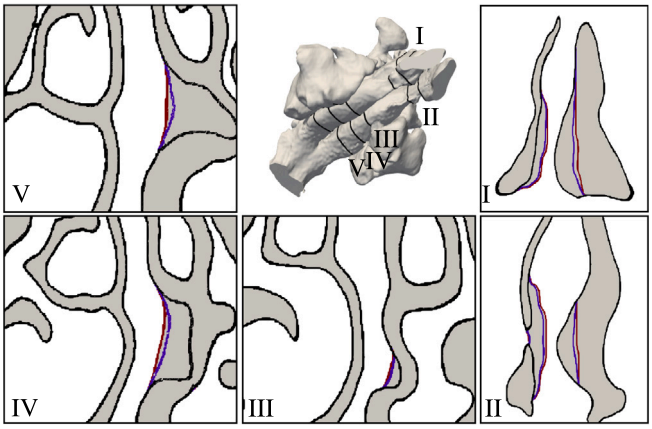

Fig. 7. Cross-sectional areas I–V of Case A that show the contours of the pre-surgical structure (black), the geometry based on a surgeon's initial idea of the desired post-operative result (red), and the geometry with $\alpha = 0.8$ (blue).

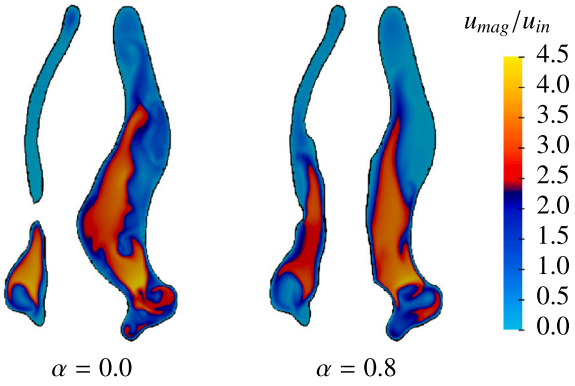

Fig. 8. u_{mag}/u_{in} at the cross-sectional area II of Case A for the pre-surgical structure (left) and the geometry with $\alpha = 0.8$ (right).

The left part of the figure shows how the incoming air is hindered by the narrowed right passage of the pre-surgical nasal cavity to spread evenly. Instead, it is trapped in the lower part of the right nasal passage. In contrast, the left nasal passage for $\alpha = 0.0$ is widened and more extensively supplied with air. The right part of Fig. 8 reveals a more balanced flow through the left and right nasal passages. However, there is still a gap quantified by a factor of $\dot{V}_{norm}(\alpha = 0.8) = 2.255$.

Overall, the spatial differences between the geometries defined by $\alpha = 0.8$ and $\alpha = 1.0$ are not large, e.g., the difference between the blue and red contours in Fig. 7 are sometimes hardly visible. This shows that the surgeon's initial idea goes into the right direction, and it underlines that the proposed method does not only help to

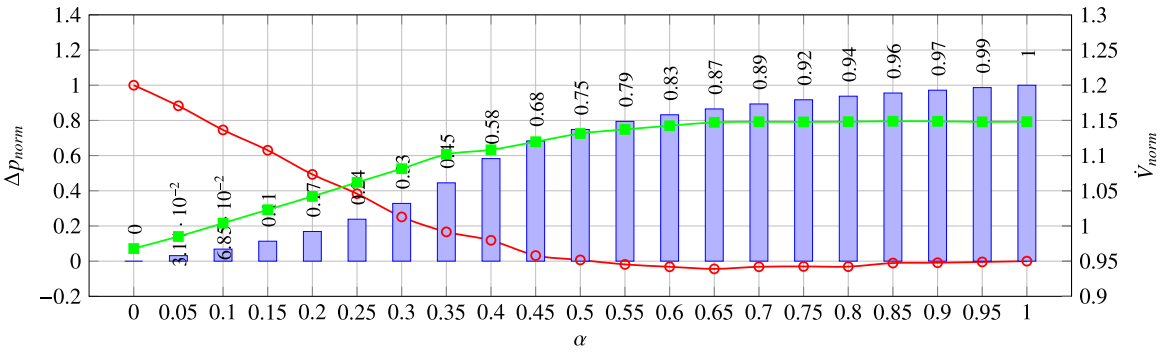

Fig. 9. Δp_{norm} (red), $\dot{V}_{norm}$ (green), and Vol_{norm} (violet) in the course of the virtual surgery for Case B.

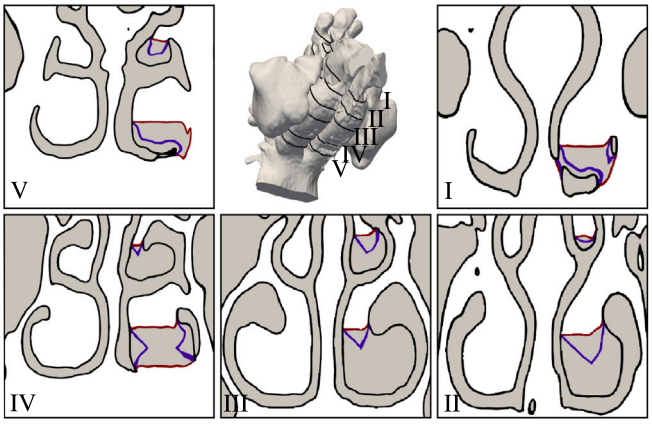

Fig. 10. Cross-sectional areas I–V of Case B that show the contours of the pre-surgical structure (black), the geometry based on the maximum possible intervention (total turbinectomy) (red), and the geometry with $\alpha = 0.5$ (blue).

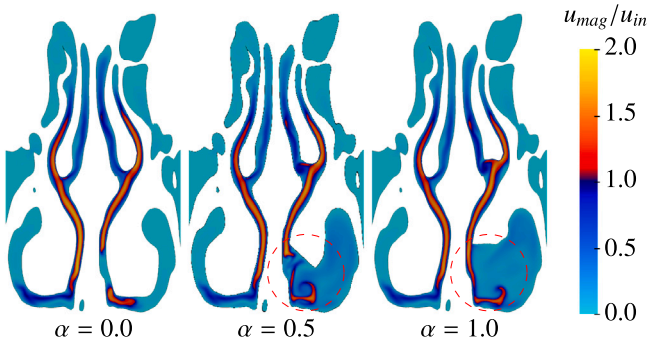

Fig. 11. u_{mag}/u_{in} at the cross-sectional area II of Case B for the pre-surgical structure (left), the geometry with $\alpha = 0.5$ (center), and the geometry based on the maximum possible intervention (right).

suggest surgical interventions, but also to provide valuable feedback on a surgeon's initial thoughts. This helps to further develop a surgeon's skills and complement her/his experience.

Fig. 9 shows Δp_{norm} (red), $\dot{V}_{norm}$ (green), and Vol_{norm} (violet bars) for $\alpha \in \{0.0, 0.05, \ldots, 1.0\}$ of Case B. Enlarging the turbinates clearly leads to a reduced pressure loss. The loss Δp_{norm} reaches its global minimum at $\alpha = 0.65$, with a value of $\Delta p_{norm}(\alpha = 0.65) = -0.044$, indicating a lower normalized pressure drop compared to the reference configuration $\alpha = 1.0$, where $\Delta p_{norm} = 0.0$. However, a pressure loss nearly identical to that at $\alpha = 1.0$ is already achieved at $\alpha = 0.5$, with $\Delta p_{norm}(\alpha = 0.5) = -0.019$. If a surgeon decides $\Delta p_{norm}(\alpha = 0.5) = -0.019$ to be sufficiently low, 25% less tissue in comparison to a total turbinectomy (which is represented by the surface variations with $\alpha = 1.0$), would be removed.

The volume flow rate through the right nasal passage is slightly increased compared to the one through the left nasal passage in the pre-surgical state. The widening of the left nasal passage increases $\dot{V}_l$, which leads to an increase in the volume flow rate ratio from $\dot{V}_{norm}(\alpha = 0.0) = 0.968$ to $\dot{V}_{norm}(\alpha = 1.0) = 1.148$. In contrast to Vol_{norm} for Case A in Fig. 6, where a nearly steady increase between $Vol_{norm}(\alpha = 0.0)$ and $Vol_{norm}(\alpha = 1.0)$ is observed, the course of Vol_{norm} for Case B in Fig. 9 shows a turning point at $\alpha = 0.4$. The strongest changes of Vol_{norm} are observed between $\alpha = 0.2$ and $\alpha = 0.6$.

The difference between the pre-surgical structure (black), the geometry based on the maximum possible intervention (total turbinectomy - red), and the geometry with $\alpha = 0.5$ is illustrated for the cross-sections I–V in Fig. 10. The results reveal that with $\alpha = 0.5$ much less tissue is removed compared to $\alpha = 1.0$ in both the middle and inferior turbinates, i.e., the blue contours are far away from a total turbinectomy represented by the red contours.

Similar to Case A, the differences in the pressure distribution has been extensively visualized in [14,17]. Therefore, Fig. 11 focuses on the velocity field and its effect on $\dot{V}_{norm}$. The left part of the figure shows that in the pre-surgical geometry, incoming airflow passes through the portion of the nasal passage surrounded by the middle turbinate but does not reach the area near the inferior turbinate. Especially in the left passage, where the turbinates are diagnosed to be enlarged, incoming air is largely prevented from passing through in the vicinity of the inferior turbinate. In contrast, the center part of Fig. 11 shows that in the course of removing tissue from the inferior turbinate and widening the left nasal passage, incoming air is distributed more evenly to the left nasal passage. However, the largest velocity magnitudes are concentrated in the small region highlighted by the dashed red circle. Here, a jet forms as a result of the inflow through the narrowed nasal valve. The jet persists under a further widening of the left nasal passage, i.e., shifting from $\alpha = 0.5$ to $\alpha = 1.0$. This is illustrated in the right part of Fig. 11. Interestingly, the velocity field away from this jet does not undergo any decisive change. This explains why $\dot{V}_{norm}$ experiences no major change between $\alpha = 0.5$ and $\alpha = 1.0$.

In contrast to Case A, where the surgeon's initial plan was aligned with the numerically optimized outcome, the results for Case B indicate that a total turbinectomy should be avoided. Depending on clinical priorities, such as minimizing pressure loss while maintaining balanced left and right volume flow rates, an intervention with $\alpha < 0.5$ (for example, $\alpha = 0.25$) may represent a suitable compromise.

Fig. 12 shows Δp_{norm} (red), $\dot{V}_{norm}$ (green), and Vol_{norm} (violet bars) for $\alpha \in \{0.0, 0.05, \ldots, 1.0\}$ of Case C. Straightening the deviated septum clearly leads to a reduced pressure loss. In contrast to the steady decrease of Δp_{norm} in Case A (see Fig. 6), the overall decrease of Δp_{norm} in Case C is accompanied by local fluctuations. As in Case A, there is only a marginal gain in reducing the pressure loss between $\Delta p_{norm}(\alpha = 0.8) = 0.040$ and $\Delta p_{norm}(\alpha = 1.0) = 0.0$. However, 20% of tissue removal can be saved when choosing $\alpha = 0.8$ over $\alpha = 1.0$. A major difference to the two other cases is that the virtually planned septoplasty in Case C has only a minor influence on $\dot{V}_{norm}$, i.e., $\dot{V}_{norm}(\alpha = 0.0) = 0.734$

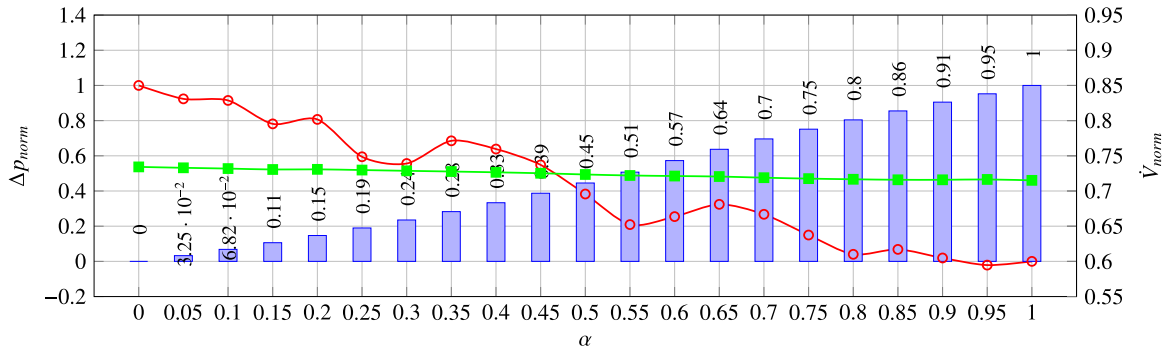

Fig. 12. Δp_{norm} (red), $\dot{V}_{norm}$ (green), and Vol_{norm} (violet) in the course of the virtual surgery for Case C.

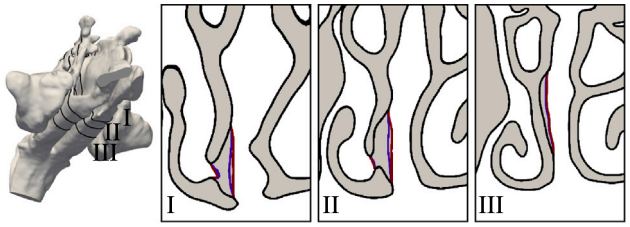

Fig. 13. Cross-sectional areas I–III of Case C that show the contours of the pre-surgical structure (black), the geometry based on a surgeon's initial idea of the desired post-operative result (red), and the geometry with $\alpha = 0.8$ (blue).

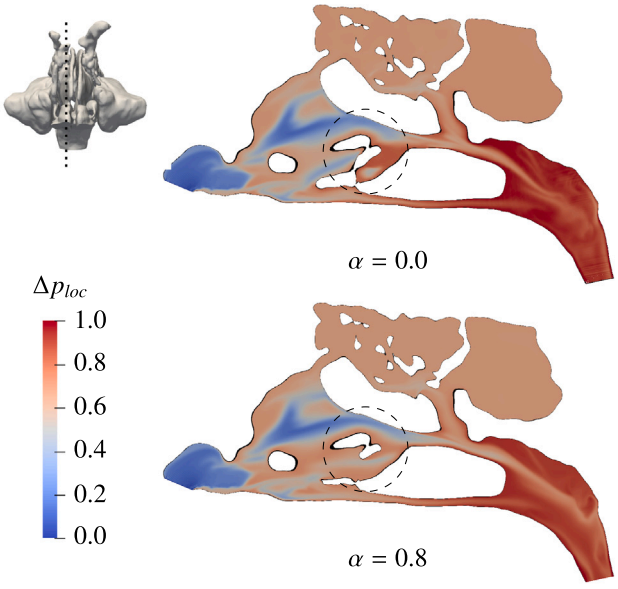

Fig. 14. $\Delta p_{norm,I}$ of Case C at the cross-sectional area highlighted by the dotted line for the pre-surgical structure (top), and the geometry with $\alpha = 0.8$ (bottom).

and $\dot{V}_{norm}(\alpha = 1.0) = 0.715$ deviate by only 2.6%. Similar to Case A, a nearly steady increase between $Vol_{norm}(\alpha = 0.0)$ and $Vol_{norm}(\alpha = 1.0)$ is observed.

The difference between the pre-surgical structure (black), the geometry based on a surgeon's initial idea of the desired post-operative result (red), and the geometry with $\alpha = 0.8$ is illustrated for the cross-sectional areas I–III in Fig. 13. The main purpose of the procedure, i.e., widening the blocked right nasal passage, can clearly be seen.

The surgery planning outcomes show no major influence on $\dot{V}_{norm}$. Therefore, instead of visually analyzing velocity fields (as done in Cases A and B), the pressure distribution is analyzed in Fig. 14 in terms of the local pressure loss normalized by $\Delta p(\alpha = 0)$ and $\Delta p(\alpha = 0.8)$ respectively (Δp_{loc}).

The upper part of the figure indicates a large increase in Δp_{loc}, which is caused by the interrupted nasal passage. This is highlighted by the dashed black circle. In the lower part of Fig. 14, it is shown how straightening the septum and thereby reconnecting the nasal passage leads to a reduced pressure loss. This effect is not only visible in the vicinity of the dashed black circle, but also in the downstream part of the nasal cavity up to the outflow region at the pharynx.

Overall, similar to Case A, the spatial differences between the geometries defined by $\alpha = 0.8$ and $\alpha = 1.0$ are not large. This again demonstrates the potential of the tool to also provide positive feedback and confirm a surgeon's initially planned procedures with valuable physics-based assessments.

4.2. Surgery planning based on CFD simulations and a non-linear regression model

The results presented in this section are partly based on computations by the LB solver and partly on predictions by a non-linear regression model. The regression model predicts the dependent variables $\phi \in \{\Delta p_{norm}, \dot{V}_{norm}\}$ as a function of α. Every prediction saves the computation of 25,000 time steps. Note that Vol_{norm} does not need to be predicted, since it can be easily calculated within one time step without time-averaging. The purpose of this section is to generalize regression-based predictions, i.e., to determine the parameter ζ in a way that reliable predictions for all patients are achieved. The graphs of Δp_{norm} and $\dot{V}_{norm}$ in Section 4.1 are not linear or parabolic. Therefore, polynomials with $\zeta < 3$ are skipped.

For the Cases A–C, regression-based predictions are evaluated with the mean absolute percentage error (MAPE)

$$e_{MAPE} = \frac{1}{N_{test}} \sum_{i=0}^{N_{test}} \left(\frac{|\phi_{sim}(i) - \phi_{test}(i)|}{|\phi_{max} - \phi_{min}|} \right), \tag{16}$$

and the maximum absolute percentage error (MaxAPE)

$$e_{MaxAPE} = Maximum \left(\frac{|\phi_{sim}(i) - \phi_{test}(i)|}{|\phi_{max} - \phi_{min}|} \right). \tag{17}$$

The subscripts *pred* and *sim* mark regression-based predictions and ground truth from CFD simulations, the subscripts *min* and *max* represent the minimum and maximum values of the corresponding data set, and N_{test} is the size of test data.

Fig. 15 shows e_{MAPE} and e_{MaxAPE} of regression-based predictions with different degrees of the polynomial (ζ) and training-test data combinations CB 1 (red), CB 2 (green), and CB 3 (blue) for Case A. The predictions of Δp_{norm} in Fig. 15(a) reveal that all predictions have an error of $e_{MAPE} < 5\%$. Averaging e_{MAPE} for all degrees of polynomials yields $\bar{e}_{MAPE}(CB1) = 0.0246$, $\bar{e}_{MAPE}(CB2) = 0.0285$, and $\bar{e}_{MAPE}(CB3) = 0.0304$. This shows the prediction accuracy generally deteriorates when the number of training data decreases.

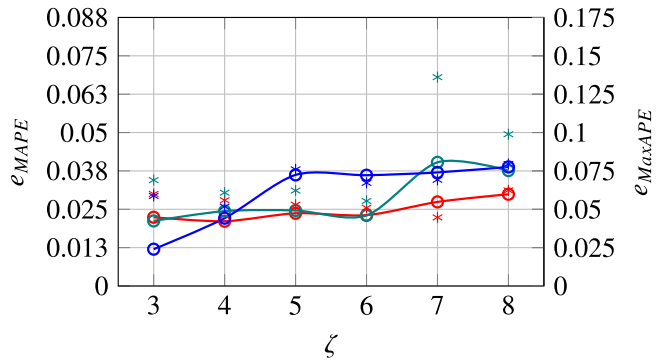

(a) Predictions of Δp_{norm}.

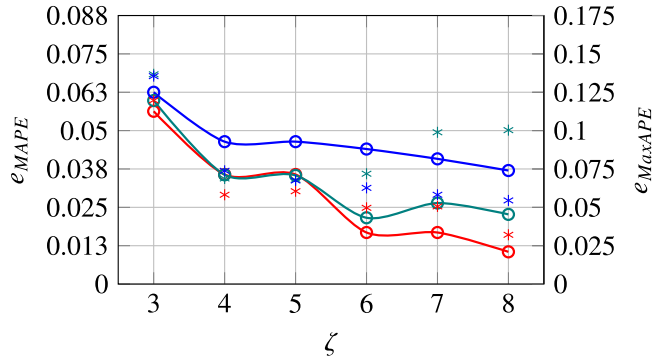

(b) Predictions of $\dot{V}_{norm}$.

Fig. 15. e_{MAPE} (solid line, assigned to the left vertical axis) and e_{MaxAPE} (*, assigned to the right vertical axis) of regression-based predictions with different degrees of the polynomial (ζ) and training-test data combinations CB 1 (red), CB 2 (green), and CB 3 (blue) for Case A.

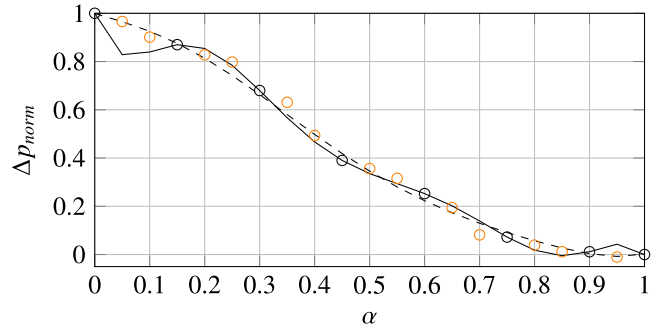

Fig. 16. Regression-based predictions for Case A of Δp_{norm} with $\zeta = 6$ (dashed line) and $\zeta = 7$ (solid line) compared to the corresponding training data (black circles) and test data (orange circles) for CB 2.

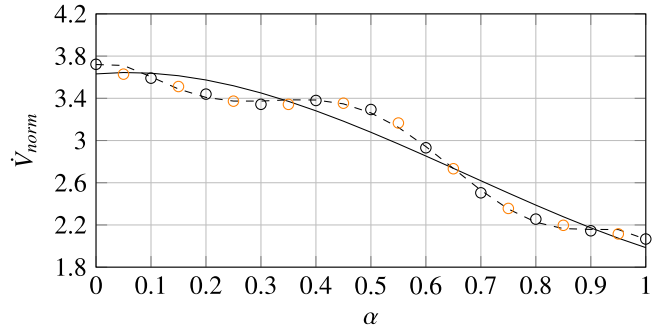

Fig. 17. Regression-based predictions for Case A of $\dot{V}_{norm}$ with $\zeta = 6$ (dashed line) and $\zeta = 3$ (solid line) compared to the corresponding training data (black circles) and test data (orange circles) for CB 1.

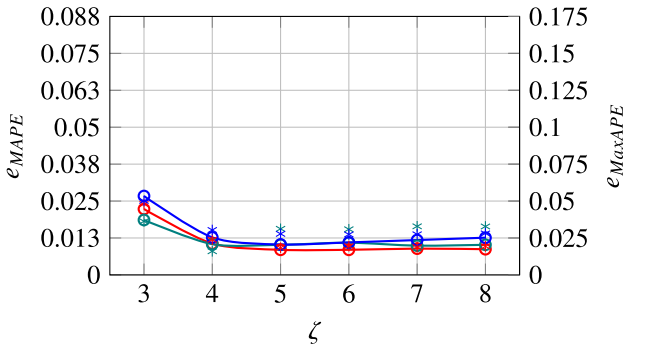

(a) Predictions of Δp_{norm}.

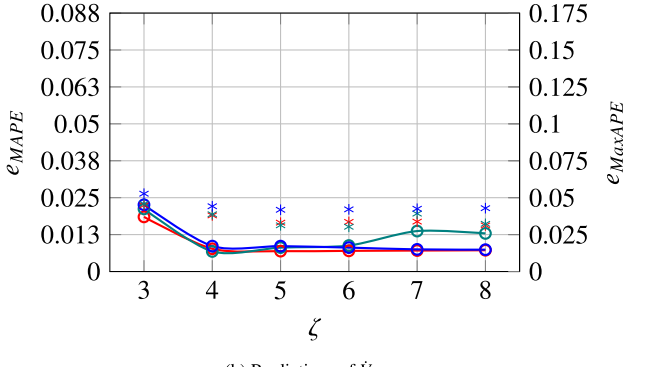

(b) Predictions of $\dot{V}_{norm}$.

Fig. 18. e_{MAPE} (solid line, assigned to the left vertical axis) and e_{MaxAPE} (*, assigned to the right vertical axis) of regression-based predictions with different degrees of the polynomial (ζ) and training-test data combinations CB 1 (red), CB 2 (green), and CB 3 (blue) for Case B.

Looking at the individual training and test data combinations, the prediction accuracy tends to worsen with increasing ζ. For example, CB 2 shows with $e_{MaxAPE}(\zeta = 7) = 13.61\%$ and $e_{MaxAPE}(\zeta = 8) = 9.89\%$ large errors at high ζ values, which is caused by the well-known overfitting phenomenon. This is demonstrated in Fig. 16 for the predictions of Δp_{norm} with $\zeta = 6$ (dashed line) and $\zeta = 7$ (solid line). The regression model of this example is trained with the training-test data split CB 2, and the corresponding training data (black circles) and test data (orange circles) are illustrated in the figure. Although both graphs intersect with all training points, only the graph with $\zeta = 6$ manages to accurately predict the test data. The graph with $\zeta = 7$ misses some of the test data, e.g., at $\alpha = 0.05$, $\alpha = 0.1$, or $\alpha = 0.95$.

The predictions of $\dot{V}_{norm}$ in Fig. 15(b) confirm the overall trend for $\bar{e}_{MAPE}$ that has been reported earlier in Fig. 15(a), i.e., $\bar{e}_{MAPE}(CB1) = 0.0286 < \bar{e}_{MAPE}(CB2) = 0.0336 < \bar{e}_{MAPE}(CB3) = 0.0462$. The figure also indicates relatively large errors for $\zeta = 3$, even for CB 1, which has the largest number of training data. The reasons for this are explained with the help of Fig. 17, which contains predictions of $\dot{V}_{norm}$ with $\zeta = 3$ (solid line) and $\zeta = 6$ (dashed line). The regression model of this example is trained with the training-test data split CB 1, and the corresponding training data (black circles) and test data (orange circles) are illustrated in the figure. In contrast to the overfitted regression model trained with $\zeta = 7$ shown in Fig. 16, here $\zeta = 3$ is not complex enough to reproduce $\dot{V}_{norm}$. Thus, looking only at Case A, $3 < \zeta < 7$ yields the best results.

Fig. 18 shows e_{MAPE} and e_{MaxAPE} of regression-based predictions as functions of ζ and CB 1 (red), CB 2 (green), and CB 3 (blue) for Case B. Both types of errors are for the predictions of Δp_{norm} (Fig. 18(a)) and $\dot{V}_{norm}$ (Fig. 18(b)) below 2.5%. This shows that the distribution of the pressure and volume flow rates (see Fig. 9) are in general easier to predict compared to those of Case A (see Fig. 6). However, slightly increased errors are observed for $\zeta = 3$, $\zeta = 7$, and $\zeta = 8$, which again gives the impression that $3 < \zeta < 7$ yields the best results.

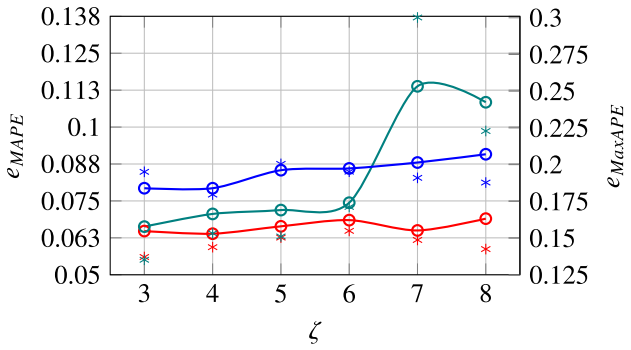

(a) Predictions of Δp_{norm}.

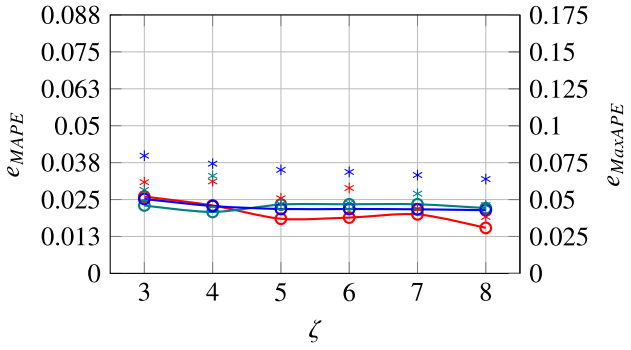

(b) Predictions of $\dot{V}_{norm}$.

Fig. 19. e_{MAPE} (solid line, assigned to the left vertical axis) and e_{MaxAPE} (*, assigned to the right vertical axis) of regression-based predictions with different degrees of the polynomial (ζ) and training-test data combinations CB 1 (red), CB 2 (green), and CB 3 (blue) for Case C.

Fig. 19 shows e_{MAPE} and e_{MaxAPE} of regression-based predictions with different degrees of the polynomial (ζ) and training-test data combinations CB 1 (red), CB 2 (green), and CB 3 (blue) for Case C. Whereas the regression model seems to have no problems in predicting $\dot{V}_{norm}$ (Fig. 19(b)), Fig. 19(a) reveals that the predictions of Δp_{norm} are the most challenging ones. This comes from the largely fluctuating pressure loss visible in Fig. 9. Whereas training with CB 1 yields a still acceptable $\bar{e}_{MAPE}(CB1) = 0.0663$, training with CB 2 ($\bar{e}_{MAPE}(CB2) = 0.0842$) and CB 3 ($\bar{e}_{MAPE}(CB3) = 0.0848$) is not preferable for replacing some of the CFD computations. Therefore, if non-linear regression is employed to save computational resources, it is recommended to train the regression model with CB 1 and replace 50% of the CFD computations.

To assess the physical relevance of the regression errors, it is analyzed whether the reported prediction inaccuracies could influence the surgical decision-making process. Since the proposed approach aims to support virtual surgery planning, it is critical to understand if deviations between predicted and simulated outcomes may lead to different clinical interpretations. The prediction errors for $\dot{V}_{norm}$ are generally lower and follow the trends observed in the CFD simulations shown in Figs. 6, 9, and 12, making them less likely to affect the overall interpretation. In contrast, when predicting Δp_{norm}, large errors are observed for Case C with CB 2 and $\zeta = 7$ or $\zeta = 8$. Fig. 20 helps to show how those may lead to contradictory interpretations of the virtual surgery outcome. With $\zeta = 7$ (dashed line) and $\zeta = 8$ (solid line), Δp_{norm} at $\alpha = 0.05$ is larger than based on the simulation at the pre-surgical state ($\alpha = 0$). This might lead to a wrong conclusion that widening the nasal airways increases the pressure loss. Furthermore, at $\alpha = 0.95$, the regression-based pressure loss is smaller than Δp_{norm} based on the simulation. If a surgeon comes to the conclusion that $\alpha = 0.95$ is the desired surface variation, much more tissue would be removed than suggested in Section 4.1, cf. Fig. 12, where $\alpha = 0.8$ was found to be the maximum meaningful surface variation. While these examples

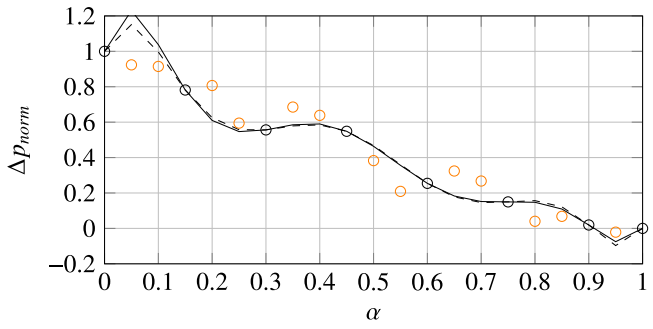

Fig. 20. Regression-based predictions for Case C of Δp_{norm} with $\zeta = 7$ (dashed line) and $\zeta = 8$ (solid line) compared to the corresponding training data (black circles) and test data (orange circles) for CB 2.

Table 3
$\hat{e}_{MAPE}(\zeta \in \{4,5,6\})$ for predicting Δp_{norm} and $\dot{V}_{norm}$ with CB 1.

ϕ	ζ	e_{MAPE} for Cases A, B, C	$\hat{e}_{MAPE}$
Δp_{norm}	4	0.0211, 0.0106, 0.0639	0.0317
Δp_{norm}	5	0.0237, 0.0085, 0.0664	0.0329
Δp_{norm}	6	0.0231, 0.0085, 0.0685	0.0334
$\dot{V}_{norm}$	4	0.0357, 0.0077, 0.0230	0.0221
$\dot{V}_{norm}$	5	0.0357, 0.0069, 0.0185	0.0204
$\dot{V}_{norm}$	6	0.0168, 0.0070, 0.0189	0.0142

highlight the importance of understanding the impact of errors, a more comprehensive clinical interpretation of these trends would require further validation, which is, however, beyond the scope of this study.

The question remains which of $\zeta = 4$, $\zeta = 5$, or $\zeta = 6$ delivers the best results. To give a first answer to this question, $\hat{e}_{MAPE}(\zeta \in \{4,5,6\})$ is defined, which is the averaged error of the training runs with CB 1 for Cases A, B, and C shown in Table 3. According to the table, $\zeta = 4$ is recommended for predicting Δp_{norm} and $\zeta = 6$ for predicting $\dot{V}_{norm}$. However, since this is based on the mean of just three samples these results should be interpreted with caution. An extended study with more patients and methods like bootstrapping is required to provide a strong recommendation.

5. Summary and discussion

The complexities involved in planning nasal surgeries, such as septoplasties or turbinectomies, highlight the need for innovative and efficient approaches to enable the widespread adoption of CFD-based automated surgery planning. Methods of previous works were either automated but computationally expensive, or cheaper methods that were not automated and solely analyzed single solutions based on manual geometry modifications. The current study bridges the gap between these two drawbacks. Three patients are investigated, the first patient suffering from a deviated septum accompanied by a bony spur, the second patient from enlarged turbinates in the left nasal passage, and the third patient from a deviated septum.

Two types of investigations were carried out. The first investigation presented a combined LB-LS method that could be used to efficiently compute fluid mechanical parameters of geometry variations that lie between the pre-surgical geometry and a surgeon's initial plan for the desired surgical outcome. In contrast to previous works, where the geometry was varied and the LS field needed to be recomputed each time step, in the proposed method the geometry is varied in discrete interpolation steps between the two reference states. After each step, the CFD computations were run for 25,000 time steps until a converged flow field was reached. Furthermore, while previous works considered multiple surgical intervention locations, resulting in up to 441 potential geometry variations requiring flow field computations, this study focuses on single interventions, reducing the number of

surface variations to 21. The fluid mechanical parameters included the pressure loss between the inflow (nostrils) and outflow (pharynx) regions, and the balance between the volume flow rates through the left and right nasal passages. An additional criterion that can be taken into account by a surgeon is the volume of the removed tissue for each geometry variation.

For the first patient, a geometry variation was recommended that saves 12% of the volume initially planned to be removed by the surgeon. In an irreversible intervention like nasal surgery, it is the surgeon's responsibility to remove only as much tissue as needed. At the same time, the proposed geometry variation causes no meaningful deterioration in terms of pressure loss or balanced volume flow rates compared to the surgeon's initial plan. For the second patient, it was found that a total turbinectomy should be avoided. Considering only the pressure loss, it can be maintained with a geometry variation that saves 25% of the volume compared to completely removing the turbinates. These savings can be even greater if the surgeon places more emphasis on balanced volume flows. This is because a jet develops after the incoming air streams through the narrowed nasal valve, and its intensification does not change significantly when the nasal passage is widened beyond a certain size. For the third patient, a geometry variation was recommended that saves 20% of the volume initially planned to be removed by the surgeon, without degrading pressure loss or balanced volume flow rates.

The second investigation shows how some of the fluid mechanical parameters that are normally extracted from CFD simulations can be replaced by predictions of a non-linear regression model. Scenarios were analyzed where the regression model was trained with data from 11 (CB 1), 8 (CB 2), and 6 (CB 3) CFD simulations. While the predictions of the regression model trained with CB 2 or CB 3 sometimes caused large inaccuracies, the predictions with CB 1 yielded averaged errors below 4% for predictions of the pressure loss, and averaged errors below 3% for the volume flow rates.

Employing the regression model with CB 1 saves 50% of the CFD computations and hence also 50% of the computational resources. This reduced the run time for Case A using 16 CPU nodes, i.e., 2,048 cores, from around 18 to 9h, and for Case B from around 24 to 12h. To show the flexibility in terms of employing the proposed tool on CPUs or GPUs, Case C was also run on GPUs. Here, surgery planning took around 20h on 8 GPU nodes, i.e., using 32 GPUs, without and 10h with the regression model.

Acquiring both pre-surgical CT data and expert-modified post-surgical models that reflect a surgeon's initial plan is difficult. Hence, the number of patients in this study is limited to three. Importantly, this work does not aim to perform a clinical evaluation or assess long-term surgical outcomes. Instead, the focus is on demonstrating a numerical framework and monitoring techniques for virtual nasal surgery planning. Within this scope, the selected cases are sufficient to illustrate the feasibility and potential of the proposed algorithms. The regression results should, however, be interpreted with care, as generalization to broader patient populations requires a validation on larger datasets. Future work will address this by expanding the patient cohort and evaluating model robustness across varying anatomical and clinical conditions.

As this study uses a standard regression approach to predict pressure loss and volume flow rate from geometric variation, regularization techniques such as ridge regression or Bayesian regression could provide advantages in controlling model complexity and mitigating overfitting, especially in low-data regimes. However, given the structured and deterministic way in which training samples are generated in this study (via CFD simulations at controlled deformation levels) overfitting was primarily addressed through careful design of the training/test split and performance evaluation on unseen intermediate geometries. The inclusion of regularization methods may be beneficial in future studies with larger patient datasets or more complex geometric parametrizations.

The most important criteria for a widespread usage of an automated virtual nasal surgery planning tool are accuracy, affordability, and practicability. The results of the current investigation have already shown a solid foundation for all three criteria. However, further investigations can still be made on the way towards a widespread usage. A validation of the LB solver of the m-AIA framework for simulating respiratory flows to examine the accuracy has so far only been done with flow at expiration through a simplified lung model in [38]. That is, stereo-particle-image velocimetry measurements were performed in multiple parallel measurement planes for a Reynolds number of $Re = 700$ based on the hydraulic diameter of the trachea to determine the highly three-dimensional flow in a realistic transparent silicone lung model, and a comparison with the LB method in terms of the absolute velocity and the out-of-plane velocity component show a good agreement. However, the inflow region of the lung model is located at the pharynx and flow through the nasal cavity is not considered. For an increased credibility, further validations are needed that include flow through the nasal cavity and, e.g., inspiration or unsteady inflow conditions of a breathing cycle.

Automated nasal surgery planning could be made more affordable by entirely replacing CFD simulations with ML methods. This could be attempted by trying to predicting the 3D flow and pressure field in the nasal cavity based on geometric information as input to the ML model, or by directly predicting target quantities such as the pressure loss or the ratio of left and right volume flow rates. However, such attempts are challenging because they require a large dataset of nasal cavities and the corresponding ground truth generated with CFD simulations [39].

The tool's practicability could be enhanced by a more user-friendly graphical user interface (GUI) for creating the 3D model of the surgeon's initial surgery idea. Currently, the user loads a binary segmentation file of the nasal cavity into the software 3DSlicer [13]. While the segmentation file is generated automatically with the pipeline described in [12], manual interaction is necessary for creating the binary segmentation file of the surgeon's initial plan. That is, the user adds or removes voxels with paint or erase functionalities of the segment editor, either slice by slice, or with a spherical tool in a 3D viewer. An improvement that should be considered for future improvements is replacing this tedious process by combining non-uniform rational B-splines (NURBS) [40] for geometric flexibility and isogeometric analysis (IGA) [41,42] for simulation-readiness. Doing so it would be possible to create a user-friendly system where a nasal cavity 3D model can be modified interactively with a mouse click-and-drag mechanism.

Overall, the presented method offers a promising approach for integrating CFD, RL, and regression techniques in virtual nasal cavity surgery planning. Since the study does not include validation against real surgical outcomes, the results should currently be considered as a technical feasibility study. The application of such computational tools in clinical decision processes will require further investigation through clinical collaborations, larger datasets, and regulatory approval processes.

6. Conclusion

This study addressed the challenges of automated virtual nasal surgery planning by bridging the gaps in previous methods, providing a solution that balances accuracy, affordability, and practicability. The combined LB-LS method introduced here enables efficient computation of fluid mechanical parameters for geometry variations, reducing the computational cost while maintaining accuracy. Recommendations for all three patients demonstrate the method's ability to save tissue volume while preserving critical flow characteristics, such as pressure loss and balanced volume flows. Additionally, incorporating a non-linear regression model further reduces computational resources by replacing certain CFD simulations, achieving significant cost savings without compromising accuracy.

While the results lay a strong foundation for the widespread adoption of CFD-based automated surgery planning, further validation is needed, particularly including flow through the nasal cavity. Future efforts could explore replacing CFD simulations entirely with ML models, though this will require a large dataset of nasal geometries and simulations. Lastly, enhancing the tool's usability through a more intuitive graphical user interface, leveraging NURBS and IGA, would streamline the creation and modification of 3D nasal models, enabling surgeons to interactively plan surgeries with greater efficiency.

CRediT authorship contribution statement

Mario Rüttgers: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Moritz Waldmann:** Writing – review & editing, Validation, Software, Methodology, Conceptualization. **Fabian Hübenthal:** Writing – review & editing, Validation, Software, Methodology, Conceptualization. **Klaus Vogt:** Writing – review & editing. **Makoto Tsubokura:** Writing – review & editing. **Sangseung Lee:** Writing – review & editing. **Andreas Lintermann:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The research leading to these results has been funded by the German Research Foundation within the Walter Benjamin fellowship RU 2771/1-1 and by the HANAMI project within the European Union Horizon Europe Programme - Grant Agreement Number 101136269 under the call HORIZON-EUROHPC-JU-2022-INCO-04. The authors gratefully acknowledge the computing time granted by the JARA Vergabegremium and provided on the JARA Partition part of the supercomputer JURECA [36,37] at *Forschungszentrum Jülich*.

Data availability

Data will be made available on request.

References

[1] K. Zhao, P. Malhotra, D. Rosen, P. Dalton, E.A. Pribitkin, Computational fluid dynamics as surgical planning tool: a pilot study on middle turbinate resection, Anat. Rec. 11 (297) (2014) 2187–2195.
[2] M. Waldmann, A. Lintermann, Y.-J. Choi, W. Schröder, Analysis of the effects of MARME treatment on respiratory flow using the Lattice–Boltzmann method, in: New Results in Numerical and Experimental Fluid Mechanics XII, in: Notes on Numerical Fluid Mechanics and Multidisciplinary Design, vol. 142, Springer, Berlin, Heidelberg, 2020, pp. 853–863.
[3] J. Ormiskangas, O. Valtonen, T. Harju, M. Rautiainen, I. Kivekäs, Computational fluid dynamics assessed changes of nasal airflow after inferior turbinate surgery, Respir. Physiol. Neurobiol. 302 (2022) 103917.
[4] H.G. Weller, G. Tabor, H. Jasak, C. Fureby, A tensorial approach to computational continuum mechanics using object orientated techniques, Comput. Phys. 12 (1998) 620–631.
[5] J. Siu, K. Inthavong, J. Dong, Y. Shang, R.G. Douglas, Nasal air conditioning following total inferior turbinectomy compared to inferior turbinoplasty – A computational fluid dynamics study, Clin. Biomech. 81 (2021) 105237.
[6] S. Ito, M. Rüttgers, M. Waldmann, A. Lintermann, Wet-Surface Modeling in Lattice-Boltzmann Simulations for Evaluating Surgery Impacts on the Humidity Transfer in Nasal Flows, in: Proceedings of the 35th Parallel CFD International Conference 2024, in: Schriften des Forschungszentrums Jülich IAS Series, 69, Forschungszentrum Jülich GmbH Zentralbibliothek, Verlag, 2025, pp. 158–162.
[7] T. Keck, R. Leiacker, A. Heinrich, S. Kühnemann, G. Rettinger, Humidity and temperature profile in the nasal cavity, Rhinology 38 (4) (2000) 167–171.
[8] T. Radulesco, L. Meister, G. Bouchet, A. Varoquaux, J. Giordano, P. Dessi, P. Perrier, J. Michel, Computational fluid dynamics and septal deviations: virtual surgery can predict postsurgery results. A preliminary study including 2 patients, Clin. Otolaryngol. 45 (2019).
[9] M.A. Burgos, L. Rosique, F. Piqueras, C. García-Navalón, M.A. Sevilla-García, D. Hellín, F. Esteban, Reducing variability in nasal surgery outcomes through computational fluid dynamics and advanced 3D virtual surgery techniques, Heliyon 10 (5) (2024) e26855.
[10] E. Moreddu, L. Meister, M. Médale, R. Nicollas, Computational fluid dynamics (CFD), virtual rhinomanometry, and virtual surgery for neonatal congenital nasal pyriform aperture stenosis, Int. J. Pediatr. Otorhinolaryngol. 182 (2024) 112025.
[11] M. Waldmann, M. Rüttgers, A. Lintermann, W. Schröder, Virtual surgeries of nasal cavities using a coupled Lattice-Boltzmann–level-set approach, J. Eng. Sci. Med. Diagn. Ther. 5 (3) (2022).
[12] M. Rüttgers, M. Waldmann, W. Schröder, A. Lintermann, A machine-learning-based method for automatizing Lattice–Boltzmann simulations of respiratory flows, Appl. Intell. 52 (2022) 9080–9100.
[13] R. Kikinis, S.D. Pieper, K.G. Vosburgh, 3D slicer: A platform for subject-specific image analysis, visualization, and clinical support, in: F.A. Jolesz (Ed.), Intraoperative Imaging and Image-Guided Therapy, Springer New York, New York, NY, 2014, pp. 277–289.
[14] M. Rüttgers, M. Waldmann, K. Vogt, J. Ilgner, W. Schröder, A. Lintermann, Automated surgery planning for an obstructed nose by combining computational fluid dynamics with reinforcement learning, Comput. Biol. Med. 173 (2024) 108383.
[15] A. Lintermann, M. Meinke, W. Schröder, Fluid mechanics based classification of the respiratory efficiency of several nasal cavities, Comput. Biol. Med. 43 (11) (2013) 1833–1852.
[16] A. Lintermann, G. Eitel-Amor, M. Meinke, W. Schröder, Lattice–Boltzmann solutions with local grid refinement for nasal cavity flows, in: New Results in Numerical and Experimental Fluid Mechanics VIII, in: Notes on Numerical Fluid Mechanics and Multidisciplinary Design, vol. 121, Springer, Berlin, Heidelberg, 2013, pp. 583–590.
[17] M. Rüttgers, F. Hübenthal, M. Tsubokura, A. Lintermann, Parallel reinforcement learning and Gaussian process regression for improved physics-based nasal surgery planning, in: R. Wyrzykowski, J. Dongarra, E. Deelman, K. Karczewski (Eds.), Parallel Processing and Applied Mathematics, Springer Nature Switzerland, Cham, 2025, pp. 79–96.
[18] G. Eitel, R.K. Freitas, A. Lintermann, M. Meinke, W. Schröder, Numerical simulation of nasal cavity flow based on a Lattice–Boltzmann method, in: Andreas Dillmann, Gerd Heller, Michael Klaas, Hans-Peter Kreplin, Wolfgang Nitsche, Wolfgang Schröder (Eds.), New Results in Numerical and Experimental Fluid Mechanics VII, in: Notes on Numerical Fluid Mechanics and Multidisciplinary Design, vol. 112, Springer, Berlin, Heidelberg, 2010, pp. 513–520.
[19] P. Perona, J. Malik, Scale-space and edge detection using anisotropic diffusion, IEEE Trans. Pattern Anal. Mach. Intell. 12 (7) (1990) 629–639.
[20] W. Lorensen, H. Cline, Marching cubes: A high resolution 3D surface construction algorithm, ACM SIGGRAPH Comput. Graph. 21 (1987) 163.
[21] G. Taubin, T. Zhang, G. Golub, Optimal surface smoothing as filter design, 1996.
[22] A. Lintermann, S. Schlimpert, J.H. Grimmen, C. Günther, M. Meinke, W. Schröder, Massively parallel grid generation on HPC systems, Comput. Methods Appl. Mech. Engrg. 277 (2014) 131–153.
[23] D. Hartmann, M. Meinke, W. Schröder, An adaptive multilevel multigrid formulation for Cartesian hierarchical grid methods, Comput. & Fluids 37 (9) (2008) 1103–1125.
[24] G.M. Morton, A Computer Oriented Geodetic Data Base and a New Technique in File Sequencing, International Business Machines Company, 1966.
[25] H. Sagan, Hilbert's space-filling curve, in: Space-Filling Curves, Springer New York, New York, NY, 1994, pp. 9–30.
[26] A. Niemöller, M. Schlottke-Lakemper, M. Meinke, W. Schröder, Dynamic load balancing for direct-coupled multiphysics simulations, Comput. & Fluids 199 (2020) 104437.
[27] J. Li, W.-K. Liao, A. Choudhary, R. Ross, R. Thakur, W. Gropp, R. Latham, A. Siegel, B. Gallagher, M. Zingale, Parallel netCDF: A high-performance scientific I/O interface, in: SC '03: Proceedings of the 2003 ACM/IEEE Conference on Supercomputing, 2003, 39–39.
[28] S. Geller, M. Krafczyk, J. Tölke, S. Turek, J. Hron, Benchmark computations based on Lattice–Boltzmann, finite element and finite volume methods for laminar flows, Comput. & Fluids 35 (8) (2006) 888–897, Proceedings of the First International Conference for Mesoscopic Methods in Engineering and Science.
[29] X. He, L.-S. Luo, Theory of the lattice Boltzmann method: From the Boltzmann equation to the lattice Boltzmann equation, Phys. Rev. E 56 (6) (1997) 6811–6817.
[30] Y.H. Qian, D. D'Humières, P. Lallemand, Lattice BGK models for Navier–Stokes equation, Europhys. Lett. 17 (6) (1992) 479–484.
[31] Y. Kuwata, K. Suga, Anomaly of the lattice Boltzmann methods in three-dimensional cylindrical flows, J. Comput. Phys. 280 (2015) 563–569.
[32] A.T. White, C.K. Chong, Rotational invariance in the three-dimensional lattice Boltzmann method is dependent on the choice of lattice, J. Comput. Phys. 230 (16) (2011) 6367–6378.

[33] M. Bouzidi, M. Firdaouss, P. Lallemand, Momentum transfer of a Boltzmann–Lattice fluid with boundaries, Phys. Fluids 13 (11) (2001) 3452–3459.
[34] T. Krüger, H. Kusumaatmaja, A. Kuzmin, O. Shardt, G. Silva, E.M. Viggen, The Lattice Boltzmann method - principles and practice, 2016.
[35] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, J. Mach. Learn. Res. 12 (2011) 2825–2830.
[36] P. Thörnig, JURECA: Data Centric and Booster Modules implementing the Modular Supercomputing Architecture at Jülich Supercomputing Centre, J. Large-Scale Res. Facil. 7 (A182) (2021).
[37] D. Krause, P. Thörnig, JURECA: Modular supercomputer at Jülich Supercomputing Centre, J. Large-Scale Res. Facil. 4 (2018) A132.
[38] F. Schröder, A. Lintermann, M. Klaas, W. Schröder, Experimental and numerical investigation of the three-dimensional flow at expiration in the upper human airways, Int. J. Fluid Eng. (2014).
[39] H. Calmet, J. Calafell, R. Sarma, M. Rüttgers, A. Lintermann, G. Houzeaux, Creating a Virtual Population of the Human Nasal Cavity for Velocity-Based Predictions of Respiratory Flow Features Using Graph Convolutional Neural Networks, in: Proceedings of the 35th Parallel CFD International Conference 2024, in: Schriften des Forschungszentrums Jülich IAS Series, 69, Forschungszentrum Jülich GmbH Zentralbibliothek, Verlag, 2025, pp. 80–83.
[40] L. Piegl, W. Tiller, The NURBS Book, Springer, Berlin, Heidelberg, 1997.
[41] T.J.R. Hughes, J.A. Cottrell, Y. Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement, Comput. Methods Appl. Mech. Engrg. 194 (39–41) (2005) 4135–4195.
[42] J.A. Cottrell, T.J.R. Hughes, Y. Bazilevs, Isogeometric Analysis: Toward Integration of CAD and FEA, John Wiley & Sons, Chichester, UK, 2009.