


RESEARCH ARTICLE

# Comparing flow-based and anatomy-based features in the data-driven study of nasal pathologies

Andrea Schillaci<sup>1</sup>, Kazuto Hasegawa<sup>1,2</sup>, Carlotta Pipolo<sup>3</sup>, Giacomo Boracchi<sup>4</sup> and Maurizio Quadrio<sup>1,\*</sup> 

<sup>1</sup>Department of Aerospace Science and Technologies, Politecnico di Milano, 20156 Milano, Italy

<sup>2</sup>Department of Mechanical Engineering, Keio University, 3-14-1 Hiyoshi, Kohoku-ku, Yokohama, Japan

<sup>3</sup>Otolaryngology Unit, Asst Santi Paolo e Carlo, Department of Health Sciences, University of Milan, 20142 Milano, Italy

<sup>4</sup>DEIB, Politecnico di Milano, 20133 Milano, Italy

\*Corresponding author. E-mail: [maurizio.quadrio@polimi.it](mailto:maurizio.quadrio@polimi.it)

**Received:** 29 August 2023; **Revised:** 6 February 2024; **Accepted:** 5 March 2024

**Keywords:** Nasal cavities; Computational fluid dynamics; Dimensionality reduction; Functional maps

## Abstract

In several problems involving fluid flows, computational fluid dynamics (CFD) provides detailed quantitative information and allows the designer to successfully optimize the system by minimizing a cost function. Sometimes, however, one cannot improve the system with CFD alone, because a suitable cost function is not readily available; one notable example is diagnosis in medicine. The application considered here belongs to the field of rhinology; a correct air flow is key for the functioning of the human nose, yet the notion of a functionally normal nose is not available and a cost function cannot be written. An alternative and attractive pathway to diagnosis and surgery planning is offered by data-driven methods. In this work, we consider the machine learning study of nasal impairment caused by anatomic malformations, with the aim of understanding whether fluid dynamic features, available after a CFD analysis, are more effective than purely geometric features at the training of a neural network for regression. Our experiments are carried out on an extremely simplified anatomic model and a correspondingly simple CFD approach; nevertheless, they show that flow-based features perform better than geometry-based ones and allow the training of a neural network with fewer inputs, a crucial advantage in fields like medicine.

## Impact Statement

Machine learning (ML) algorithms and computational fluid dynamics (CFD) techniques are often discussed together in the recent scientific literature in fluid mechanics. However, ML is always used as a tool to perform better/cheaper/simpler/faster CFD. In this work, we explore the potential of the inverse approach, in which standard CFD provides useful information to an ML model. In an idealized three-dimensional problem, flow features restricted to the boundary of the computational domain and derived from CFD are shown to be more informative than the geometry of the boundary itself, leading to a better ML classifier, which can be trained with fewer labelled data.

The application described herein is of the medical type and concerns rhinology, where large amounts of accurately labelled data are not readily available. Since the larger information content of flow-based features derives from the nonlinear relationship between geometry and the corresponding flow field, the present result is relevant to other flow problems addressed with CFD where the lack of a clearly defined cost function suggests a data-driven approach.

## 1. Introduction

With the continuous development of computing hardware and software, computational fluid dynamics (CFD) is becoming increasingly useful in several applications, from industry to health. CFD, ranging from the cheaper and lower-fidelity flow models like the Reynold-averaged Navier–Stokes (RANS) equations to the opposite extreme of the highly accurate direct numerical simulation, is a useful tool to improve the design of industrial systems, by e.g. increasing the aerodynamic efficiency of an airplane, reducing the aerodynamic drag of a vehicle or enhancing the mixing in a fluidic system.

Sometimes, however, the CFD solution, albeit informative, does not explicitly provide an immediate means to improve the system. This is often the case in the medical field. The specific example considered in this work concerns the air flow in the human nose. The nasal cavities are placed between the external ambient and the lungs, and play the role of a connecting duct; in addition, they serve a number of other functions, which include smell, filtering and humidifying the incoming air, and heating/cooling it to the correct temperature. Most of these functions are directly driven by the anatomical shape of the nasal cavities. In fact, nowadays, ear, nose and throat (ENT) surgeons routinely take their surgical decisions mostly based on the analysis of computed tomography (CT) scans, which provide a detailed view of the anatomy of the nasal cavities. In principle, the nose flow can be well described by CFD, which is indeed increasingly useful to support ENT doctors in their diagnosis (Moreddu *et al.* 2019; Tjahjono *et al.* 2023) and to improve our understanding of the complex physics of the nose flow (Calmet *et al.* 2019; Farnoud *et al.* 2020). Yet, the basic questions routinely asked by ENT doctors (e.g. whether to perform a surgery on a given patient and where) cannot be straightforwardly answered by CFD alone. Designing a surgery can be considered to be akin to a shape optimization problem. Unfortunately, the mathematical and numerical tools available for shape optimization cannot be deployed to solve the nose problem, because the goal is not self-evident, owing to the lack of a functionally normal reference nose. A huge anatomical variability among healthy anatomies is present, which makes the discrimination between healthy and pathological cases far from obvious.

Several studies have attempted to develop a robust workflow for a CFD analysis of the nose flow (Quadrio *et al.* 2014; Tretiakow *et al.* 2020) and to understand what is a healthy airflow (Zhao & Jiang 2014; Borojeni *et al.* 2020) via multi-patient analyses. However, as seen from a clinical perspective, the present state of affairs remains unsatisfactory. There is evidence that the rate of failure of certain surgical corrections is extremely high: for the correction of septal deviations, for example, more than 50 % of the patients report poor postoperative satisfaction ratings (Rhee *et al.* 2003; Sundh & Sunnergren 2015; Tsang *et al.* 2018). Although there is general agreement (Inthavong *et al.* 2019) that CFD has a significant potential for improved surgery planning, this potential remains, as yet, largely untapped.

An interesting approach to diagnose pathologies and suggest surgeries relies on the use of machine learning (ML) techniques. In medicine, ML is mainly considered as an image-processing tool. The central question that we are going to address in this paper is whether CFD-computed flow information can be useful in diagnosis and surgery planning, and whether it can be more effective than the geometrical information embedded in the CT scan. The possibility that CFD-based features outperform geometrical features is rooted in the nature of the highly nonlinear Navier–Stokes equations, which link anatomy to the ensuing CFD solution.

The ML in fluid mechanics has recently seen a huge activity and recorded significant progresses (see e.g. the review by Vinuesa & Brunton 2022); however, very little information is available in the literature regarding the combined use of ML and CFD in rhinology, if exception is made for our own preliminary study (Schillaci *et al.* 2021). The use of artificial intelligence and machine learning techniques in rhinology has been limited so far to classification of images derived from CT scans (Crowson *et al.* 2020). In closing their paper, Lin, Hsieh & Hsieh (2020) simply mention that putting together CFD and ML would be an interesting future avenue for research. A very recent study by Jin *et al.* (2023) uses CFD and ML approaches, but only one at a time, and there is no attempt to combine them in any way.

The main goal of the present work is thus to answer the question whether flow features derived from CFD of the nasal airflow can be useful for an ML-based analysis. A positive answer would carry general

interest in all those situations where ML is used in the context of flow systems to replace optimization, since a cost function is not readily available.

To answer the question above, we consider synthetic (but realistic) physiologic and impaired nasal anatomies, created with CAD, and simple CFD solutions of the flow within them (computed by solving the Reynolds-averaged Navier–Stokes equations and a standard turbulence model). An inference model made by a standard neural network is trained to understand whether each synthetic nasal anatomy is affected by a pathology and to predict its severity. Two alternative approaches are employed. One resembles the approach currently followed by ENT doctors for their diagnosis and is solely based upon geometric/anatomic information. The other, instead, relies on flow features extracted from CFD.

By comparing the performance of the two types of features, we intend to understand whether CFD carries potential advantages over anatomy alone, like e.g. increased accuracy or the need for less observations, which will need to be evaluated together with the additional computational cost and complexity. In particular, the ability to train the ML model with information derived from fewer patients would be extremely important, owing to the difficulty to obtain highly informative and accurately labelled training data in health-related applications.

The structure of the work is as follows. After this introduction, § 2 describes the generation of a suitable set of geometries, the set-up of CFD simulations, the extraction of information from the CFD solution and the neural network used for regression. Results presented in § 3 are used to critically discuss the comparison between geometrical and CFD features. Lastly, concluding remarks are put forward in § 4.

## 2. Methods

This section illustrates the entire workflow and describes in § 2.1 how the parametric CAD geometries of the noses are created, in § 2.2 how the CFD simulations are set up, in § 2.3 how different anatomies and flow solutions are compared, and in § 2.4 how a neural network is designed and trained to predict nasal pathologies.

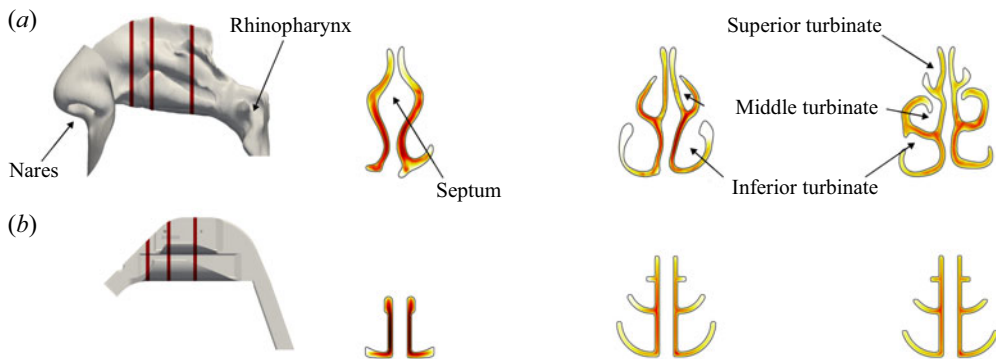
A CAD-based, synthetic reference nose model is built first. The model presents the fundamental advantage of being parametric; changing the numerical values of a small set of geometric parameters allows us to introduce anatomical variability, related to both physiological inter-subject differences and pathological conditions. Using this parametric nose model, 200 distinct anatomic shapes are generated. For each shape, a point-to-point mapping is computed between each nose and the reference nose model.

A CFD simulation is then carried out for each geometry. Thanks to the previously computed mapping, both flow and geometrical features of each geometry can be mapped back to the reference one, so that any (flow or geometrical) feature of any model can be observed on the reference one. At this point, a neural network is trained to perform a regression on the geometrical parameters of each nose and to predict the value of the pathological parameters.

### 2.1. The anatomies

All the anatomies considered in the present work originate from a reference CAD-based simplified model of the human nasal cavities, already introduced by Schillaci *et al.* (2021), which lends itself to a simple geometrical parametrization. The use of simplified model geometries for the study of the flow in the human nasal cavities is not new: for example, Liu *et al.* (2009) employed a model obtained by averaging together the CT scans of 30 patients. Our CAD-based approach relates more closely to that by Naftali *et al.* (1998), who used a CAD nose-like model to reproduce the essential features of the nasal cavities. However, to the best of our knowledge, we are the first to build a fully parametrized model that is required for the generation of a large and controlled dataset.

The reference CAD nose model is meant to represent a physiologic anatomy and mimics the major anatomical structures of a real nose. A qualitative comparison between the nasal cavities of a real



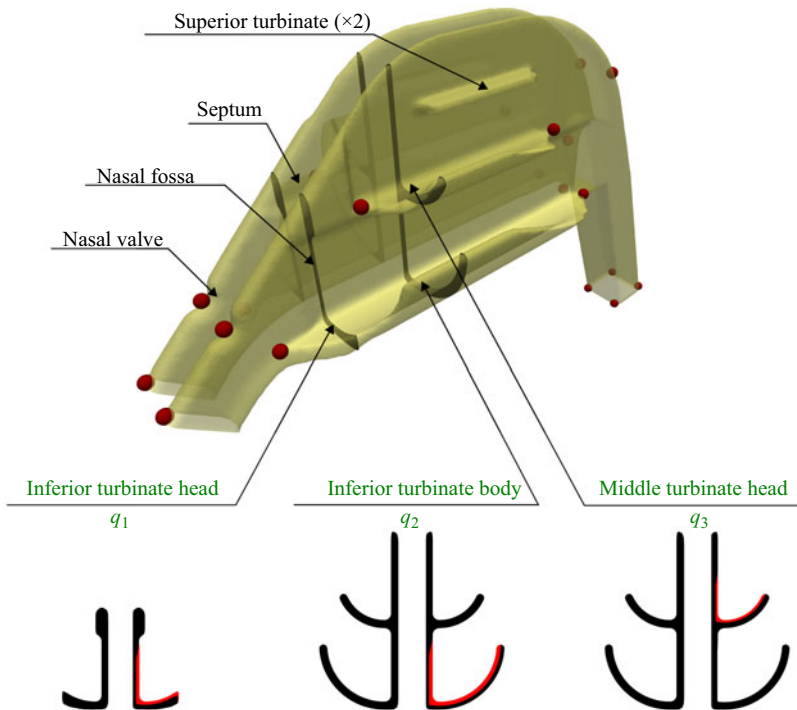
**Figure 1.** Comparison between (a) a real nasal anatomy and (b) the simplified CAD model used in the present work. The three-dimensional view on the left marks the three coronal sections plotted on the right. The cross-sections with the colourmap represent the CFD solution in terms of the magnitude of the velocity vector. The sections in panel (a) are from a time-averaged LES solution (not reported in this paper) and those in panel (b) are from a RANS solution. Although the CAD model is highly simplified, the major flow features in the model are in line with those for the real anatomy.

patient and our simplified baseline model can be seen in Figures 1a and 1b, respectively. The CAD model, developed in collaboration with a group of ENT surgeons, is designed to reconcile the opposite requirements of simplicity and clinical significance. Its planar or constant-curvature surfaces are indeed highly idealized, but the model replicates in a quantitatively accurate way the crucial anatomical features, such as the dimensions of the septum between the two fossae, the hook-like structure of the inferior and middle turbinates, and the thickness of the passageway in the most critical areas of the nasal fossae. This is an essential prerequisite to provide deformations with clinical significance. The comparison (shown in Figure 1) of the flow field computed with a high-fidelity approach on a patient-specific anatomy confirms the suitability of the present model.

The nasal geometry begins anteriorly with the nares, and ends posteriorly with the rhinopharynx and the hypopharynx. The nasal septum, a thin structure lying approximately on the median plane, separates the nasal cavities in two halves, the left and right nasal fossae. Each fossa has a cross-sectional shape that changes significantly along the nasal vestibules, developing as a narrow channel of convoluted shape, medially bounded by the septum; the particular conformation of the fossae is considered to improve humidification and heat exchange. Three long and curled bony structures, the (inferior, middle and superior) turbinates, define the cross-sectional shape of the fossae. The turbinates unfold roughly parallel to the flow and are attached to the lateral walls of the nose. The inferior turbinate is the largest, running almost the entire way from the vestibule to the rhinopharynx.

The parametric CAD model contains eight numerical parameters, whose range of variation is meant to account at the same time for the physiological and pathological variabilities of real anatomies. Three parameters,  $q_1$ ,  $q_2$  and  $q_3$ , correspond to three clinically sensitive regions of the nasal cavities, and describe the intensity of selected pathologies, related to an hypertrophy of the turbinate; the remaining five parameters are clinically harmless and represent the physiological variability among healthy anatomies. The parameters describe anatomical changes or their arbitrary combinations; their numerical values are expressed in millimetres and are changed in steps of 0.05 mm. In the healthy reference anatomy, all the parameters are set to zero. Figure 2 shows where in the model the parameters act to modify the reference nose.

Ninety-nine extra healthy anatomies are created from the reference one by varying the values of the five healthy parameters. They alter (see Figure 2) the vertical and longitudinal position of the superior turbinate, the axial position of the nasal valve, the thickness of the septum in the area of the nasal valve and the thickness of the septum in correspondence to the head of the inferior turbinate. Varying



**Figure 2.** Nose model and position of its parametric modifications. Top: three-dimensional view, with black labels indicating physiological variations and green labels indicating pathologies. The red points are landmarks used for functional mapping (see text). Two cross-sections (corresponding to the two most anterior ones plotted in Figure 1) are plotted on the bottom, and highlight in red the geometry changes induced by the pathological parameters  $q_1$ ,  $q_2$  and  $q_3$ .

the parameters produces localized and relatively small changes in the geometry; the strongest change leads to a 14% reduction in cross-sectional area in the most affected section, while the smallest change decreases the area by 0.7% only.

One hundred pathological anatomies are created by varying the values of the remaining three pathology-related parameters. They are designed to mimic in a quantitatively reliable way a common condition called turbinate hypertrophy, a swelling of the turbinates which produces a constriction of the meati, up to a point where the airway may, in the most severe cases, become completely obstructed. The obvious consequences are reduced nasal patency and difficulty to breathe. Such hypertrophies have a number of causes, from allergic rhinitis to inflammation of the sinuses, and affect one or more turbinates in either of the nasal fossae. The reference CAD shape is modified by altering the values of the three pathological parameters; their numerical values are set under tight supervision of ENT doctors and remain within clinically meaningful values to ensure that deformations are realistic, notwithstanding the idealized model. Parameter  $q_1$  varies between 0 and 0.7 mm, and mimics an hypertrophy of the head (anterior portion) of the inferior turbinate; parameter  $q_2$  varies between 0 and 0.7 mm, and mimics an hypertrophy of the body (intermediate portion) of the inferior turbinate; parameter  $q_3$  varies between 0 and 0.55 mm, and mimics an hypertrophy of the head of the middle turbinate. All the pathologies are applied to the right fossa. The geometrical changes induced by non-zero pathological parameters are visualized in the bottom part of figure 2.

An important point to stress is that the parametrization of the geometry provides an unambiguous label for each case. The label consists in numerical values of the three pathological parameters: in other words, for each case, it is precisely known which pathology is at play and how much is its severity. This

characteristic, that will be essential when training the inference model, is impossible to obtain when working with real anatomies.

## 2.2. Simulations

A CFD solution is computed for each of the 200 distinct anatomies. In view of the simplified geometries employed here, a basic RANS flow model is employed; the finite-volume package OpenFOAM (Weller *et al.* 1998), with its SIMPLE solver and a first-order discretization are used to arrive quickly at a converged solution. In doing so, we follow standard modelling and discretization choices, which are summarized below.

The computational grid is generated with the utilities available in OpenFOAM. An uniform background mesh with cubic cells of side length 1 mm is created first; the mesh is then refined further and adapted to the surface. The final mesh is rather coarse, in line with the RANS approach, and consists of approximately 1.1 millions cells. Values in the nasal flow literature range between 0.1 up to 44 millions cells (Inthavong *et al.* 2018), with the finest meshes being generally used in highly resolved LES studies (Calmet *et al.* 2016; Covello *et al.* 2018); meanwhile, typical, modern RANS simulations range between 1 and 4 millions elements (Liu *et al.* 2007; Wen *et al.* 2008). Although its significance in the present problem is limited, we mention that the average value of the wall-normal distance of the first cell measured in local viscous units is 0.6, with minimum and maximum at 0.004 and 1.9, respectively. No wall treatment is used, neither in terms of boundary conditions (i.e. no wall functions) nor damping functions. The mesh adopts no cell layers near the wall and relies on the decrease of cell size induced by the snappyhexmesh tool in presence of irregular boundaries to provide near-wall refinement.

The dataset is built for a steady inspiration, which is considered as the most clinically representative breathing condition. The inspiration is driven by a pressure difference of  $20 \text{ m}^2 \text{ s}^{-2}$  between the inlet and the outlet sections, which for the reference geometry corresponds to a flow rate of approximately  $178 \text{ ml s}^{-1}$ . This value corresponds to an inspiration at rest or at mild physical activity (Wang, Lee & Gordon 2012). As in the vast majority of CFD studies in this field (Radulesco *et al.* 2019), the nasal walls are considered as rigid, thus neglecting the compliancy of the tissues (which is very small at this low breathing rate) and the modifications of the erectile mucosal tissue during the nasal cycle, known to cyclically alter the shape of the passageways over a time scale of a few hours (Patel *et al.* 2015). The velocity has zero gradient at the inlet and outlet sections. The no-slip boundary condition for the velocity components is applied at the walls, whereas for pressure, a zero-gradient condition is enforced. The turbulence model of choice is  $k - \omega - SST$  (Menter, Kuntz & Langtry 2003), which is commonly employed in such simulations (Li *et al.* 2017). The model solves two additional differential equations, one for the turbulent kinetic energy  $k$  and one for the turbulent frequency  $\omega$ . At the inlet,  $k$  is set by assuming just 1% turbulent intensity, zero gradient is used at the outlet and  $k = 0$  is set at the walls. For the turbulence frequency  $\omega$ , at the inlet,  $\omega = 1 \text{ s}^{-1}$  is prescribed, at the outlet, a null gradient is imposed and at the wall, the value is set as by Menter (1994).

The outcome of a typical simulation for the healthy anatomy is compared in Figure 1 with the temporally averaged solution obtained on a real anatomy with an higher-fidelity (and significantly more expensive) CFD method, namely large eddy simulation (LES), where a WALE turbulence model is used (Ducros *et al.* 1999), for the same flow conditions. The LES mesh is of approximately 12.8 millions cells. Although the comparison clearly has to be intended in a qualitative sense only, it is seen that most of the flow passes through the same regions, in particular, in the meatus between the septum and the inferior and middle turbinate. This confirms the suitability of the simplified model for the purpose of the present work.

We stress once again that, in the present study, seeking the highest fidelity in the numerical solution is not our primary concern. Hence, the solution method (the RANS equations), the numerical schemes (first order) and the quality of the mesh (fairly coarse) are all standard, and meant to generate quickly and cheaply a dataset of reasonable size. Since each case carries approximately 100 MB of data, the full dataset has a total size of 20 GB. To put these numbers into perspective, the Google Open Images



Dataset V6 dataset (Kuznetsova *et al.* 2020) consists of approximately  $9 \times 10^6$  images which equates to approximately 561 GB of data (including labels). In other words, our dataset is only one order of magnitude smaller, but consists of four orders of magnitude less observations, which emphasizes the high dimensionality of a typical CFD dataset.

### 2.3. Functional maps

A correspondence needs to be determined between the reference nose and each of the remaining models in the dataset. This correspondence is computed with a tool derived from computational geometry and called a functional map (FM) (Ovsjanikov *et al.* 2012); functional mapping is briefly introduced below and then specialized to the implementation employed here (Melzi *et al.* 2019).

Functional mapping provides an efficient method to estimate the correspondence between two shapes, as well as the correspondence of functions represented on them. It is a relatively new tool that has been introduced for solving shape classification problems (Magnet *et al.* 2023). Rather than directly estimating point-to-point correspondence between shapes, a FM registers functional spaces defined over the two shapes. The multi-scale basis for the function space on each shape is given by the finite (truncated) set of eigenfunctions  $\Phi_j$ ,  $j = 1 \dots N$  of its Laplace–Beltrami operator. Once the basis is known, any function  $f$  defined on the shape is approximated by the following linear combination of eigenfunctions:

$$f = \sum_{j=1}^N \gamma_j \Phi_j. \quad (2.1)$$

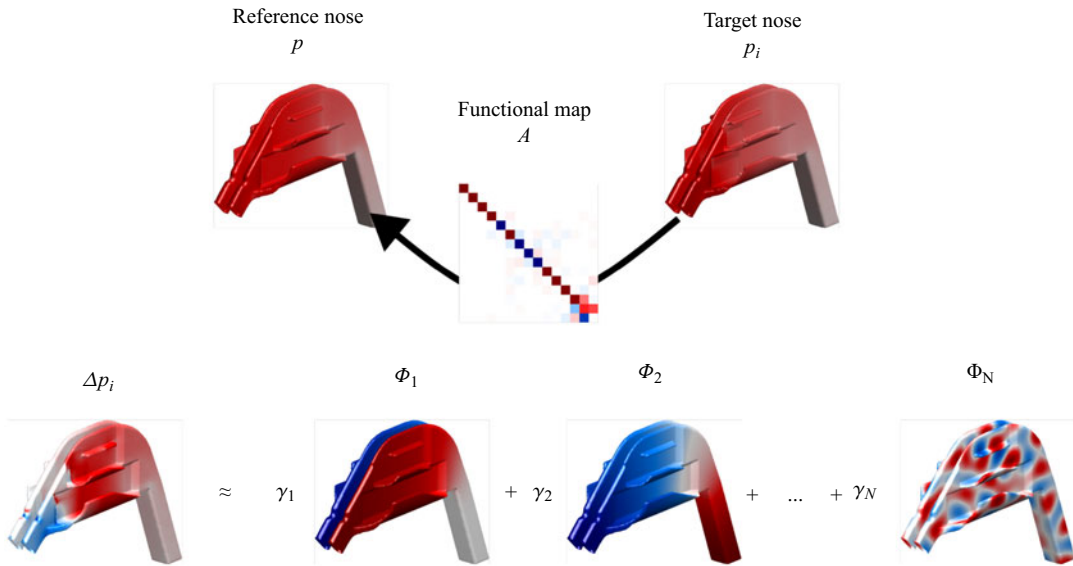
In general, the functional mapping between two shapes is described by a matrix  $A$ , whose elements describe how each eigenfunction on one shape is expressed as a linear combination of the eigenfunctions on the other shape. We refer the interested readers to the original paper by Ovsjanikov *et al.* (2012) or to the recent contribution by Magnet & Ovsjanikov (2023) for detailed information on functional mapping.

Given two shapes, the functional map between them is computed by solving a least squares minimization problem. To retrieve more precise maps, in our case, twenty landmarks are identified for each of the shapes and used as descriptors. Landmarks, shown in figure 2 for the reference nose, are points selected on the geometry because of their anatomical significance; they are often used in ENT practice (see e.g. Denour *et al.* 2020) to help deal with different anatomies in the context of CT analysis or registration. As for the basis, the eigenfunctions of the Laplace–Beltrami operator used here possess the convenient characteristic of bringing out the dominating ‘frequencies’ of the shape; therefore, they naturally provide a multi-scale representation of the geometry. Note that each nasal geometry has its own basis, but the more the two geometries are similar, the more their bases are similar. Hence, the matrix  $A$  becomes more diagonally dominant when two geometries are alike.

The specific FM implementation used here is called zoom-out and has been introduced by Melzi *et al.* (2019). Instead of computing the map for the full set of basis functions, with zoom-out, one computes first a smaller map and a smaller matrix that involves fewer basis elements (say, 10); the map is then iteratively extended, by adding rows and columns to  $A$ , up to the desired size. We have determined that, with the shapes of interest, truncating the modal expansion to  $N = 150$  eigenfunctions provides satisfactory results. Since the mapping is always between the reference shape and every other shape, 199 maps in total are computed.

Computing the correspondence between a generic shape and the reference one using FM involves the following steps:

1. compute (once) the truncated Laplace–Beltrami base on the reference shape;
2. position (once) the 20 landmarks on the reference shape;
3. compute the truncated Laplace–Beltrami base on the generic shape;
4. position the 20 landmarks on the generic shape;



**Figure 3.** Difference of a wall-based flow quantity (pressure  $p$  in this figure) between the reference and the generic  $i$ th anatomies. Pressure  $p_i$  on the boundary of the  $i$ th nose is mapped to the baseline nose as  $\hat{p}_i$ . The difference  $\Delta p_i = p - \hat{p}_i$  is expressed as a linear combination of the eigenfunctions of the reference nose with coefficients  $\gamma_j, j = 1 \dots N$ .

5. compute the functional map matrix  $A$  by solving a least-squares optimization; and
6. convert the matrix  $A$  into a point-to-point map.

Once the matrix  $A$  is available for the generic  $i$ th nose, the workflow, graphically sketched in figure 3, starts from the corresponding CFD solution, from which relevant flow quantities (for example, the pressure field  $p_i$  or the skin-friction field  $\tau_i$ ) are computed at the wall. Via FM, the wall field, say  $p_i$ , is transported back to the reference nose to yield the field  $\hat{p}_i$ ; the difference field  $\Delta p_i = p - \hat{p}_i$  can be expanded by using the Laplace–Beltrami basis  $\Phi$  of the reference anatomy and the corresponding coefficients  $\gamma_j$ , which will be used later in the regression. Note that eigenfunctions of the reference nose only are involved in the latter expansion; moreover, only wall-based quantities are mapped, thus eliminating any issue regarding the continuity equation. Vector fields are transformed component-wise.

### 2.4. The classifier

Given the dataset with  $\ell = 200$  observations, a vector of input features (be it geometrical or derived from the flow solution) must be associated to a target value which describes the pathology through the numerical values of each of the pathology parameters  $q$ . Carrying out the proper association is the task of a regressor, usually implemented as some kind of neural network (NN). The raw data from CFD for each observation are available in each cell of the discretized domain. Since the cardinality of raw data is much larger than  $\ell$ , a process of dimensionality reduction of the input, called feature extraction, is necessary to balance the number of observations with the number of inputs. The outcome of the feature extraction process depends on the specific experiment; hence, feature extraction will be described later in § 3, where the various experiments are discussed.

We train two different NN models, depending on the input data: a multi-layer perceptron (MLP) (Goodfellow, Bengio & Courville 2016), when the input is a feature vector, and a convolutional neural network (CNN) (Lecun *et al.* 1998), when the input is the matrix  $A$ , which we treat as spatial data. CNN is chosen because of its ability to capture patterns in maps; it is widely used in the field of image



recognition and it has also been applied to fluid dynamics in recent years (see e.g. Fukami, Fukagata & Taira 2020), due to its capability to deal with spatially coherent information.

In general, designing the architecture of an NN involves several choices, e.g. deciding on the number of hidden layers and nodes, the activation function, and the loss function. Our MLP is a regression network; it has an input layer, whose number of nodes is equal to the length of the feature vector, three hidden layers (with 30, 20 and 10 nodes each) and an output layer with one node only. The activation function is the hyperbolic tangent for all the nodes, except for the output nodes, which has a linear activation function. Since the goal is to predict the numerical values of the parameters which quantify the severity of the pathology, we adopt the mean square error as the loss function. Lastly, the optimization algorithm, which updates weights and biases of the NN, is the classic Levenberg–Marquardt one (Lera & Pinzolas 2002). Reduction in the number of inputs is obtained by extracting the 20 most informative features via the Lasso method (Tibshirani 1996). Note that this implies a reordering of the features, so that the retained ones do not coincide with the lowest frequencies.

Our CNN has a rather standard architecture. The functional map, i.e. matrix  $A$ , is passed into a convolutional block, which consists of a  $3 \times 3$  convolution layer, then a dropout layer randomly deactivates 20% of the weights in its layer. Afterwards, the data are normalized by a batch normalization layer and then fed into a hyperbolic tangent layer; finally, data are reduced in size through a maxpooling layer with a  $2 \times 2$  filter. The dropout and batch normalization layers are applied to avoid overfitting. After the convolution block, data are flattened into a vector and input into a fully-connected (FC) block, consisting of 20 perceptrons, dropout layers, batch normalization and hyperbolic tangent layers. The size of the FC block is halved until a single output node remains. The weights of the CNN are optimized with the widely used Adam method (Kingma & Ba 2017), owing to its efficiency and stability.

For both NN architectures, a reliable assessment of the error over the entire dataset is obtained with the  $k$ -fold cross-validation method (James *et al.* 2021). The dataset is partitioned over  $k = 5$  folds: each has 140 cases used for training, 20 for validation and 40 for testing. The 40 simulations used for testing do not overlap over the five folds, so that the performance is assessed over the whole dataset, albeit by training five different NN (with the same architecture). To avoid the potential bias of considering just a specific run, this operation is run 100 times per feature and per pathology, and eventually the average absolute error is computed.

### 3. Experiments


Results of the regression experiments are now presented in comparative form for geometrical and flow features. The goal of the experiments consists in retrieving the numerical value of the three pathological parameters.

In consideration of the relatively small size of the database, instead of training a single NN to predict the three parameters in a single attempt, we opt for training one NN for each of the three parameters.

#### 3.1. Geometrical features

Multiple options exist to select geometrical features. In this work, we consider two features, identified with G1 and G2. The geometry-based feature G1 is simply the displacement between each point on a certain nose model and the corresponding point on the reference shape. Feature G2, instead, is the matrix  $A$  obtained via FM when registering each nose with the reference one. Owing to the different nature of G1 and G2, a different NN architecture is used, namely an MLP for G1 and a CNN for G2.

Feature G1 (shown in Figure 4) is simply the distance between points on the reference nose and the corresponding ones on the modified shape. Since the modified shape is only varied through local changes, and thus does not need registration, unchanged points lead to zero displacement. The  $i$ th nose is mapped on the reference nose, so that the pointwise distance is a scalar field defined on the surface of the reference shape. This field is then decomposed into the Laplace–Beltrami basis  $\Phi_j$  of the reference



**Figure 4.** Pointwise distance (feature G1) for (a) a healthy nose and (b) a pathological nose. The numerical values of the three pathological parameters are reported above each panel. Panel (b) shows large values of G1 which highlight a non-pathological parameter (namely the position of the superior turbinate) and non-zero but smaller values in the area interested by the pathological parameter  $q_2$  (which mimics hypertrophy of the body of the inferior turbinate). The colourmap units are in millimetres.

nose, arranged in a rectangular matrix, and the vector of Laplace–Beltrami coefficients  $\gamma_j$ . The ensuing overdetermined linear system is solved by using the Lasso method with a penalization constant. A careful choice of the constant leads to a solution with 20 coefficients only. Note, once again, that the 20 extracted coefficients do not necessarily correspond to the 20 lower-order eigenfunctions.

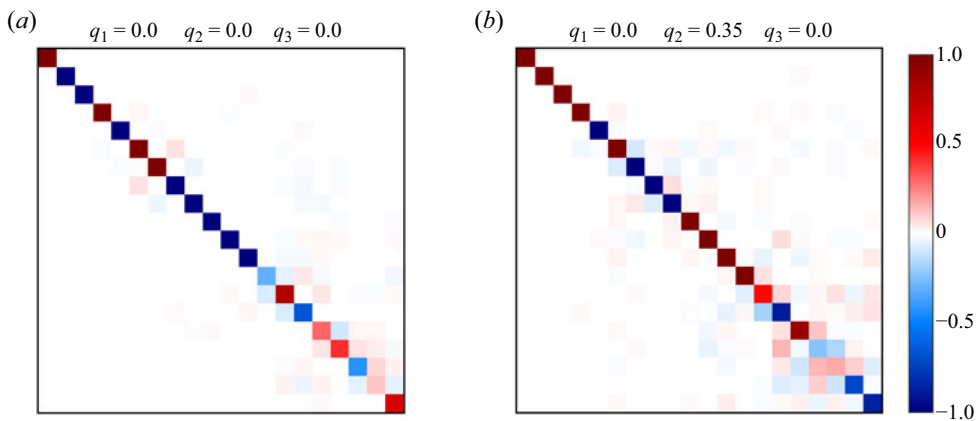
Feature G1 is effective in spotting differences introduced by changes in the parameters. Figure 4 indeed shows that G1 peaks exactly where the geometrical variations have been introduced. A quantity like G1 is expected to be subject to a small amount of noise, since each nose undergoes its own meshing process. Even though a certain portion the surface is unaltered, its points would modify their position slightly when a different mesh is computed. Nevertheless, Figure 4 shows that the noise remains more than acceptable, so that the deterministic geometrical changes are clearly highlighted; in this case, where  $q_2$  is non-zero, G1 peaks at the superior turbinate (as a consequence of the non-zero value of a non-pathological parameter) and is also large in correspondence of the inferior meatus, defined by the inferior turbinate, which is directly affected by  $q_2$ .

Feature G2, instead, involves the functional map between the two shapes, expressed via its matrix  $A$ . Owing to the usual need to avoid overfitting, the CNN is not given the full matrix  $A$ , but only a square sub-matrix  $A_1$  of size 20. Hence, we are only comparing the first 20 eigenfunctions between the two anatomies. We stress again that  $A$ , since it is based on geometry only, does not contain information concerning the flow field and is computed without the need of a CFD solution.

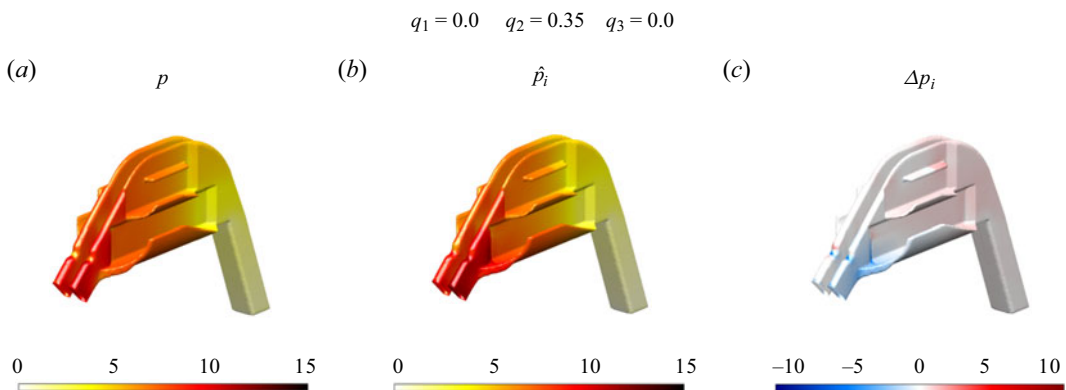
Since the reference shape is healthy, the matrix  $A$  appears to have a more diagonal structure when the actual shape corresponds to an healthy nose, compared with pathological cases. This can be confirmed by looking at Figure 5, which portrays in graphical form the structure of the sub-matrix  $A_1$ : each square represents an element, whose value is encoded in its colour after normalization to unit maximum. A diagonal matrix implies that each eigenfunction in one shape is fully described by the corresponding eigenfunction in the other shape. When off-diagonal elements are non-zero, one eigenfunction on one shape becomes a linear combination of several eigenfunctions on the other.

### 3.2. Flow features

Flow features are restricted at the wall. This choice is supported by clinical considerations; the feeling of discomfort is conveyed by nerve terminations residing in the mucosal lining (Sozansky & Houser



**Figure 5.** Functional map (feature G2), represented by the sub-matrix  $A_1$ , for (a) an healthy nose and (b) a pathological nose; each element is colour-coded after normalization at unitary maximum. The numerical values of the three pathological parameters are reported above each panel. Healthy anatomies tend to produce more diagonal maps, whereas pathologies alter the diagonal structure of the matrix.




**Figure 6.** Wall pressure  $p$  (feature F1): (a)  $p$  on the reference anatomy; (b)  $\hat{p}_i$  from a pathological anatomy (hypertrophy of the inferior turbinate,  $q_2 = 0.35$  mm) mapped back on the baseline; (c)  $\Delta p_i$ . Colourmap units are in pascals.


2014). The most straightforward wall-based features one can resort to in an incompressible flow are wall pressure  $p$  and the magnitude  $\tau$  of the wall shear stress (Bewley & Protas 2004). As before, the difference of either quantity between the reference case and each mapped-back case is computed and expanded in the reference basis: input to the NN are the first 20 coefficients of the expansion. In sharp contrast to geometry-based features that do not need CFD, these features are based on the CFD solution; however, they do not retain direct information about the generic shape. Only the eigenfunctions of the Laplace–Beltrami basis on the reference shape are used in the procedure.

Figure 6 illustrates the first flow-based feature F1, given by wall-pressure; it shows the wall-pressure field  $p$  for the baseline nose in panel a, the transformed pressure field  $\hat{p}_i$  for the  $i$ th nose, affected by a severe hypertrophy of the head of the inferior turbinate in panel b and the corresponding difference  $\Delta p_i$  represented on the baseline nose in panel c. Although the two pressure fields  $p$  and  $\hat{p}_i$  appear very similar, the difference field shows values of approximately 5 Pa, which is a clinically significant value whose order of magnitude captures the variance in nasal resistance for patients with hypertrophy.

Analogously, figure 7 describes the second flow-based feature F2 for the magnitude of the wall shear stress  $\tau$ .



**Figure 7.** Wall shear stress  $\tau$  (feature F2): (a)  $\tau$  on the reference anatomy; (b)  $\hat{\tau}_i$  from a pathological anatomy (hypertrophy of the inferior turbinate,  $q_2 = 0.35$  mm) mapped back on the baseline; (c)  $\Delta\tau_i$ . Colourmap units are in pascals.



**Figure 8.** Coefficients of the first 20 Laplace–Beltrami eigenfunctions of the expansion for the wall friction, for (a) the 100 healthy noses and (b) the 100 pathological ones. Nose ID is on the horizontal axis and mode number on the vertical axis (starting from the top).

It is instructive to observe how wall-based flow quantities are reconstructed using the Laplace–Beltrami basis. As an example, for the wall shear stress feature F2, **Figure 8** compares the 100 healthy noses in panel a with the 100 pathological ones in panel b: for each nose, the coefficients of the first 20 eigenfunctions of the Laplace–Beltrami expansion of the baseline anatomy are plotted, with colour indicating the magnitude of the coefficient. The informative nature of the flow-based coefficients is apparent. It can be noticed that the pathological cases involve contributions from a higher number of modes when compared with the healthy cases. Furthermore, certain eigenfunctions (e.g. modes 2, 3, 7, 9) are minimally involved in the description of the flow solution for healthy noses, but become important to reconstruct the fields pertaining to pathological cases.

### 3.3. Performance and discussion

Results of the regression experiments are presented in **Table 1**, in terms of the average error over the whole dataset. Distinct experiments are carried out for each pathology. The first, striking and most

**Table 1.** Average value (in millimetres) of the test error computed over 100 runs for each of the three pathological parameters  $q_1$ ,  $q_2$  and  $q_3$ . The test error of one run contains the average results of the five NN derived from  $k$ -fold cross-validation with  $k = 5$ . Flow-based features show a consistently lower error than geometry-based features.

Feature	NN type	$q_1$	$q_2$	$q_3$
G1: distance	MLP	0.156	0.116	0.094
G2: matrix $A_1$	CNN	0.148	0.117	0.099
F1: pressure	MLP	0.032	0.023	0.062
F2: wall shear stress	MLP	0.019	0.019	0.041

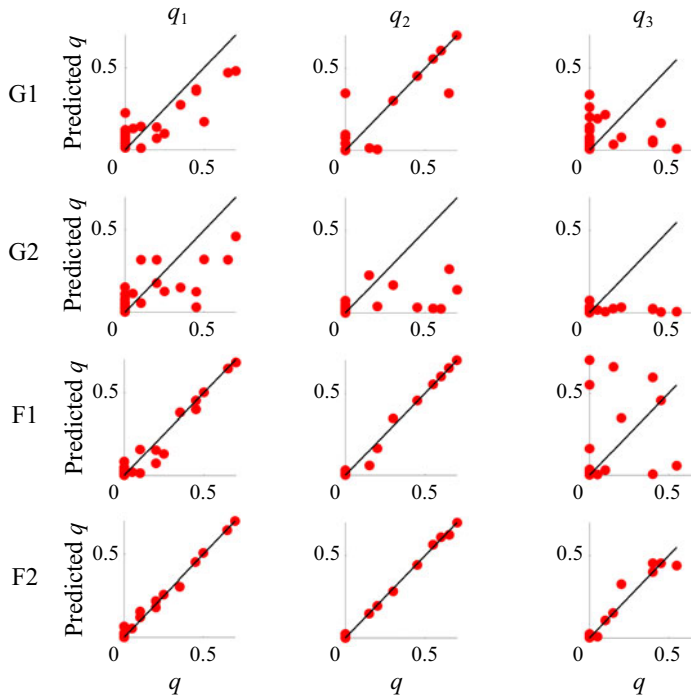
important observation is that geometry-based features produce an error which is approximately one order of magnitude larger than that obtained with flow-based features. Within flow-based features, the wall shear stress achieves a better performance over pressure in a consistent way over the set of pathologies. The pathology parameter  $q_3$  (hypertrophy of the middle turbinate) turns out to be the hardest to predict via flow features (although its error when geometric features are used is the lowest among the three pathologies): this is reasonable, because an hypertrophy of the middle turbinate affects the anatomy in a region that is not crucial in terms of distribution of the flow (see Figure 1). However, even for such an unfavourable situation, flow features do perform significantly better than the geometric ones.

The reason for the superiority of flow-based features can be traced back to the pathology being a geometrical modification which is typically quite localized in space, at least in our simplified model. As such, pathologies are small-scale features that tend to be visible only as higher-order modes of the Laplace–Beltrami basis. However, higher-order modes might be *per se* quite noisy and dependent on the geometrical discretization. This difficulty could be further emphasized by the pathologies considered here, which aim at being clinically faithful and, as a consequence, sometimes involve geometric modifications which are small in absolute terms, often a fraction of a millimetre. If pathologies are examined in terms of the corresponding flow field, instead, these small changes are passed through the ‘filter’ of the Navier–Stokes equations: a small and localized geometrical change in a sensitive position leads to a larger, more easily identified modification of the flow field. This, in turn, tends to appear within lower-order modes of the Laplace–Beltrami basis, and thus becomes easier to capture.

Figure 9 provides a closer look at the results by focusing on a single experiment, where one NN is tested on 1/5 of the dataset (40 noses). The specific test set and NN are chosen randomly, and are thus representative of the whole set. The figure shows the correlation between the predicted (vertical axis) and true (horizontal axis) labels for each feature and pathology parameter, and emphasizes how the geometric features G1 and G2 do not perform particularly well; it seems that the limited anatomical variability considered in the present work is already enough to throw off the model. It should be noted, though, that the model prediction is not meaningless; non-zero predictions are often associated with non-zero ground truth, in particular, for the parameter  $q_2$  (hypertrophy of the body of the inferior turbinate), which involves a larger surface. Both G1 and G2 are unable to predict the values of hypertrophy on the middle turbinate. Flow features F1 and F2, instead, demonstrate good regression capabilities, especially when used to predict pathologies of the inferior turbinate. For the middle turbinate, the prediction accuracy decreases. Between the two flow features, wall shear stress has a small, but significant edge over pressure.

#### 4. Conclusions and outlook

This work has introduced and discussed a novel interaction between CFD and ML. The key conclusion is that CFD-computed information may harbour extremely informative features, and may thus be useful to ML in the execution of classification and regression tasks.



**Figure 9.** Performance of the various features (rows) in one experiment for predicting the three pathological parameters (columns). Ground truth on the horizontal axis and predicted value on the vertical axis.

The problem of interest is the flow in the human nose and the classification of anatomic pathologies, in view of clinical decisions concerning functional surgery of the human upper airways. In this context, the strategic objective is learning to automatically discriminate physiological inter-subject anatomic variations from variations related to a pathological condition. Two major difficulties in this endeavour consist in the large anatomical variability which exists across healthy noses, which makes it challenging to discriminate between physiologic and impaired conditions judging from anatomy alone, and in the cost of obtaining the large number of annotated observations that is typically required to train ML algorithms. The latter issue, in particular, renders the standard approach of building a deep neural network (which, in principle, could learn directly from anatomies) highly impractical.

We have shown that an alternate solution strategy is possible: when using features extracted from the flow field computed with CFD, the training of a neural network becomes substantially easier in comparison to equivalent networks that rely on geometry-based features. The nonlinearity of the Navier–Stokes equations, together with the convective nature of the flow, is such that extracting significant information from the flow field is more effective than looking at the small anatomical changes that are behind that information. While a very large number of annotated CT scans (hence, a large amount of purely geometrical information) could, in principle, lead to a successful ML classification procedure, relying on the computed flow field is an interesting and effective alternative whenever, as in the medical field, annotated CT scans are difficult to obtain and thus necessarily available in limited quantity. Although the present model only considers localized geometrical deformations, it is found that CFD-based features outperform both small-scale geometrical features like G1 and the large-scale ones conveyed by G2.

The present work and the ensuing conclusions are obviously limited by the extreme simplification of the anatomical model and by the corresponding low-fidelity CFD approach, based on RANS simulations only. It is important to keep in mind that this work does not aim to introduce a clinically usable tool; in a realistic setting, the parametrization of the whole geometry would be impossible and alternative ways



for representing pathologies would be needed. Moreover, the obtained results remain somewhat specific, as the list of features considered here is finite, and the link between pathologies and flow-based features has been discussed only in terms of the reduced representation obtained via functional mapping.

However, thanks to the careful design of the reference nose and of its pathologies, which are both representative of clinically significant conditions, we are confident that the main conclusions are robust and will continue to apply even when the underlying anatomies become more realistic or, eventually, are derived from CT scans. The results of this work are motivating our ongoing research efforts for the classification of nasal pathologies, where instead of a nose model, real patient-specific anatomies are considered. Additional difficulties are encountered, like e.g. the need to avoid a full parametrization of the anatomy, but the present results motivate our attempt to design of a procedure based on CFD-computed features. At the same time, these conclusions may be of general interest and pave the way to using fluid mechanical features as input to improved ML methods.

**Supplementary material.** Data sets generated during the current study are available from the corresponding author on reasonable request.

**Acknowledgement.** M.Q. acknowledges support from the PRIN 2022 project, cod.2022BYA5AF CUP D53D2300343006. Computing time was provided by the CINECA Italian Supercomputing Center through the IS CRA-B projects ONOSE-AN and ONOSE-AI.

**Declaration of interests.** The authors declare no conflict of interest.

## References

- BEWLEY, T. & PROTAS, B. 2004 Skin friction and pressure: the ‘footprints’ of turbulence. *Physica D* **196**, 28–44.
- BOROJENI, A.A.T., GARCIA, G.J.M., MOGHADDAM, M.G., FRANK-ITO, D.O., KIMBELL, J.S., LAUD, P.W., KOENIG, L.J. & RHEE, J.S. 2020 Normative ranges of nasal airflow variables in healthy adults. *Intl J. CARS* **15** (1), 87–98.
- CALMET, H., GAMBARUTO, A.M., BATES, A.J., VÁZQUEZ, M., HOUZEAUX, G. & DOORLY, D.J. 2016 Large-scale CFD simulations of the transitional and turbulent regime for the large human airways during rapid inhalation. *Comput. Biol. Med.* **69**, 166–180.
- CALMET, H., INTHAVONG, K., EGUZKITZA, B., LEHMKUHL, O., HOUZEAUX, G. & VÁZQUEZ, M. 2019 Nasal sprayed particle deposition in a human nasal cavity under different inhalation conditions. *PLoS ONE* **14** (9), e0221330.
- COVELLO, V., PIPOLO, C., SAIBENE, A., FELISATI, G. & QUADRIO, M. 2018 Numerical simulation of thermal water delivery in the human nasal cavity. *Comput. Biol. Med.* **100**, 62–73.
- CROWSON, M.G., RANISAU, J., ESKANDER, A., BABIER, A., XU, B., KAHMKE, R.R., CHEN, J.M. & CHAN, T.C.Y. 2020 A contemporary review of machine learning in otolaryngology–head and neck surgery. *Laryngoscope* **130** (1), 45–51.
- DENOUR, E., ROUSSEL, L.O., WOO, A.S., BOYAJIAN, M. & CROZIER, J. 2020 Quantification of nasal septal deviation with computed tomography data. *J. Craniofac. Surg.* **31** (6), 1659–1663.
- DUCROS, F., FERRAND, V., NICLOUD, F., WEBER, C., DARRACQ, D., GACHERIEU, C. & POINOT, T. 1999 Large-eddy simulation of the shock/turbulence interaction. *J. Comput. Phys.* **152** (2), 517–549.
- FARNOUD, A., TOFIGHIAN, H., BAUMANN, I., GARCIA, G.J.M., SCHMID, O., GUTHEIL, E. & RASHIDI, M.M. 2020 Large eddy simulations of airflow and particle deposition in pulsating bi-directional nasal drug delivery. *Phys. Fluids* **32** (10), 101905.
- FUKAMI, K., FUKAGATA, K. & TAIRA, K. 2020 Assessment of supervised machine learning methods for fluid flows. *Theor. Comput. Fluid Dyn.* **34** (4), 497–519.
- GOODFELLOW, I., BENGIO, Y. & COURVILLE, A. 2016 *Deep Learning*. MIT Press.
- INTHAVONG, K., CHETTY, A., SHANG, Y. & TU, J. 2018 Examining mesh independence for flow dynamics in the human nasal cavity. *Comput. Biol. Med.* **102**, 40–50.
- INTHAVONG, K., DAS, P., SINGH, N. & SZNITMAN, J. 2019 In silico approaches to respiratory nasal flows: a review. *J. Biomech.* **97**, 109434.
- JAMES, G., WITTEN, D., HASTIE, T. & TIBSHIRANI, R. 2021 *An Introduction to Statistical Learning: With Applications in R*. Springer.
- JIN, X., LU, Y., REN, X., GUO, S., JIN, D., LIU, B., BAI, X. & LIU, J. 2023 Exploring the influence of nasal vestibule structure on nasal obstruction using CFD and machine learning method. *Med. Engng Phys.* **117**, 103988.
- KINGMA, D.P. & BA, J. 2017 Adam: a method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
- KUZNETSOVA, A., ROM, H., ALLDRIN, N., UJLINGS, J., KRASIN, I., KAMALI, S., POPOV, S., MALLOCI, M., KOLESNIKOV, A., DUERIG, T. & FERRARI, V. 2020 The open images dataset V4. *Intl J. Comput. Vis.* **128** (7), 1956–1981.
- LECUN, Y., BOTTOU, L., BENGIO, Y. & HAFNER, P. 1998 Gradient-based learning applied to document recognition. *Proc. IEEE* **86** (11), 2278–2324.
- LERA, G. & PINZOLAS, M. 2002 Neighborhood based Levenberg–Marquardt algorithm for neural network training. *IEEE Trans. Neural Netw.* **13** (5), 1200–1203.

- LI, C., JIANG, J., DONG, H. & ZHAO, K. 2017 Computational modeling and validation of human nasal airflow under various breathing conditions. *J. Biomech.* **64**, 59–68.
- LIN, H.-F., HSIEH, Y.-C. & HSIEH, Y.-L. 2020 Factors affecting location of nasal airway obstruction. In *2020 IEEE Eurasia Conference on IOT, Communication and Engineering (ECICE)*, pp. 21–24. IEEE.
- LIU, Y., JOHNSON, M.R., MATIDA, E.A., KHERANI, S. & MARSAN, J. 2009 Creation of a standardized geometry of the human nasal cavity. *J. Appl. Physiol.* **106** (3), 784–795.
- LIU, Y., MATIDA, E.A., GU, J. & JOHNSON, M.R. 2007 Numerical simulation of aerosol deposition in a 3-D human nasal cavity using RANS, RANS/EIM, and LES. *J. Aerosol Sci.* **38** (7), 683–700.
- MAGNET, R., BLOCH, K., TAVERNE, M., MELZI, S., GEOFFROY, M., KHONSARI, R.H. & OVSIJANIKOV, M. 2023 Assessing craniofacial growth and form without landmarks: a new automatic approach based on spectral methods. *J. Morphol.* **284** (8), e21609.
- MAGNET, R. & OVSIJANIKOV, M. 2023 Scalable and efficient functional map computations on dense meshes. *Comput. Graph. Forum* **42** (2), 89–101.
- MELZI, S., REN, J., RODOLÀ, E., SHARMA, A., WONKA, P. & OVSIJANIKOV, M. 2019 ZoomOut: spectral upsampling for efficient shape correspondence. *ACM Trans. Graph.* **38** (6), 155.
- MENTER, F.R. 1994 Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA J.* **32** (8), 1598–1605.
- MENTER, F.R., KUNTZ, M. & LANGTRY, R. 2003 Ten years of industrial experience with the SST turbulence model. *Turbul. Heat Mass Transfer* **4**, 8.
- MOREDDU, E., MEISTER, L., PHILIP-ALLIEZ, C., TRIGLIA, J.M., MEDALE, M. & NICOLLAS, R. 2019 Computational fluid dynamics in the assessment of nasal obstruction in children. *Eur. Ann. Otorhinolary.* **136**(2), 87–92.
- NAFTALI, S., SCHROTER, R.C., SHINER, R.J. & ELAD, D. 1998 Transport phenomena in the human nasal cavity: a computational model. *Ann. Biomed. Engng* **26** (5), 831–839.
- OVSIJANIKOV, M., BEN-CHEN, M., SOLOMON, J., BUTSCHER, A. & GUIBAS, L. 2012 Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* **31** (4), 30.
- PATEL, R.G., GARCIA, G.J.M., FRANK-ITO, D.O., KIMBELL, J.S. & RHEE, J.S. 2015 Simulating the nasal cycle with computational fluid dynamics. *Otolaryngol. Head Neck Surg.* **152** (2), 353–360.
- QUADRIO, M., PIPOLO, C., CORTI, S., LENZI, R., MESSINA, F., PESCI, C. & FELISATI, G. 2014 Review of computational fluid dynamics in the assessment of nasal air flow and analysis of its limitations. *Eur. Arch. Otorhinolary.* **271** (9), 2349–2354.
- RADULESCO, T., MEISTER, L., BOUCHET, G., GIORDANO, J., DESSI, P., PERRIER, P. & MICHEL, J. 2019 Functional relevance of computational fluid dynamics in the field of nasal obstruction: a literature review. *Clin. Otolaryngol.* **44** (5), 801–809.
- RHEE, J.S., BOOK, D.T., BURZYNSKI, M. & SMITH, T.L. 2003 Quality of life assessment in nasal airway obstruction. *Laryngoscope* **113** (7), 1118–1122.
- SCHILLACI, A., QUADRIO, M., PIPOLO, C., RESTELLI, M. & BORACCHI, G. 2021 Inferring functional properties from fluid dynamics features. In *2020 25th International Conference on Pattern Recognition (ICPR), Milan, Italy*, pp. 4091–4098. IEEE.
- SOZANSKY, J. & HOUSER, S.M. 2014 The physiological mechanism for sensing nasal airflow: a literature review. *Intl Forum. Allergy Rhinol.* **4** (10), 834–838.
- SUNDH, C. & SUNNERGREN, O. 2015 Long-term symptom relief after septoplasty. *Eur. Arch. Otorhinolaryngol.* **272** (10), 2871–2875.
- TIBSHIRANI, R. 1996 Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. B (Methodol.)* **58** (1), 267–288.
- TIJAHONO, R., SALATI, H., INTHAVONG, K. & SINGH, N. 2023 Correlation of nasal mucosal temperature and nasal patency—a computational fluid dynamics study. *Laryngoscope* **133** (6), 1328–1335.
- TRETIKOW, D., TESCH, K., MEYER-SZARY, J., MARKIET, K. & SKOREK, A. 2020 Three-dimensional modeling and automatic analysis of the human nasal cavity and paranasal sinuses using the computational fluid dynamics method. *Eur. Arch. Otorhinolaryngol.* **278**, 1443–1453.
- TSANG, C.L.N., NGUYEN, T., SIVESIND, T. & CERVIN, A. 2018 Long-term patient-related outcome measures of septoplasty: a systematic review. *Eur. Arch. Otorhinolaryngol.* **275** (5), 1039–1048.
- VINUESA, R. & BRUNTON, S.L. 2022 Enhancing computational fluid dynamics with machine learning. *Nat. Comput. Sci.* **2** (6), 358–366.
- WANG, D.Y., LEE, H.P. & GORDON, R. 2012 Impacts of fluid dynamics simulation in study of nasal airflow physiology and pathophysiology in realistic human three-dimensional nose models. *Clin. Exp. Otorhinolaryngol.* **5** (4), 181–187.
- WELLER, H.G., TABOR, G., JASAK, H. & FUREBY, C. 1998 A tensorial approach to computational continuum mechanics using object-oriented techniques. *Comput. Phys.* **12** (6), 620–631.
- WEN, J., INTHAVONG, K., TU, J. & WANG, S. 2008 Numerical simulations for detailed airflow dynamics in a human nasal cavity. *Respir. Physiol. Neurobiol.* **161** (2), 125–135.
- ZHAO, K. & JIANG, J. 2014 What is normal nasal airflow? A computational study of 22 healthy adults. *Intl Forum. Allergy Rhinol.* **4** (6), 435–446.