

CFD-ENABLED DESIGN OPTIMISATION OF INDUSTRIAL FLOWS: THEORY AND PRACTICE

Automated shape optimisation of a plane asymmetric diffuser using combined Computational Fluid Dynamic simulations and multi-objective Bayesian methodology

S.J. Daniels^a, A.A.M. Rahat^b, G. R. Tabor^a, J.E. Fieldsend^a, and R.M. Everson^a

^aCollege of Engineering, Mathematics and Physical Sciences, University of Exeter, UK;

^bSchool of Computing, Electronics and Mathematics, University of Plymouth, UK

ARTICLE HISTORY

Compiled November 7, 2019

ABSTRACT

An approach for shape optimisation of the flow through a diffuser is presented in this work. This multi-objective problem focuses on maximising the diffuser performance by simultaneously increasing the static pressure recovery across the geometry and the flow uniformity at the outflow. The hydrodynamic analysis of the geometry was conducted using the Computational Fluid Dynamics (CFD) software OpenFOAM, while a recently proposed multi-objective Bayesian approach was used for optimisation. The CFD and Bayesian methodology have been combined for fully automated operation using a Python-based framework. The proposed design parameterisation focuses on reshaping the diffuser in the expansion region. Catmull-Clark subdivision curves were employed to represent the shape of the diffuser wall; the influence of the number of control points (design points) for the curves on the optimum design was investigated. The optimal designs exhibit a reasonable performance improvement compared with the base design.

KEYWORDS

Buice Diffuser; Catmull-Clarke subdivision curves; flow uniformity index; pressure recovery; shape optimisation; multi-objective; Bayesian optimisation.

1. Introduction

The primary function of a diffuser device is to reduce the flow velocity, so as to increase the static pressure passing through the system. Thus, the physics of a diffuser is an integral part of many flow systems, such as to evenly distribute the air around a room in a Heating, Ventilation, and Air Conditioning (HVAC) application, or to reduce drag for the underside of a car. An improperly designed diffuser may lead to excessive consumption of pumping power, or may produce flow maldistribution downstream - a particularly undesirable feature in its application to the automotive industry. The design and performance of plane diffusers have been extensively investigated in the literature. Due to their simplicity, straight walls are commonly used for the diffuser's expanding section, and many studies have been carried out to highlight the parameters that affect the performance of this type of planar diffuser, e.g. length, area ratios,

and divergence angle. Such analysis has been conducted for both the laminar (e.g. Durst et al. (1974); Nabavi (2010); Suzuki et al. (2004); Tsui and Wang (1995)) and turbulent regimes (Kline et al. (1959); Waitman et al. (1961); Fox and Kline (1962); Reneau et al. (1967)). Indeed, turbulent flows through geometric expansions are of interest for numerous engineering applications, such as the design of turbomachines, combustion engines, heat-exchangers, vehicles, power plants and wind tunnels (e.g. Klein (1995); Göttlich (2011); Lan et al. (2009); Mehta and Bradshaw (1968)).

Designing the optimum shape of a diffuser has been a subject of investigation for many researchers over the last few decades. The majority of these contributions use Computational Fluid Dynamics (CFD) to assess the hydraulic performance of the design. Concurrently, a large number of these also include machine-learning or gradient-based techniques in the workflow, with the aim of efficiently identifying the optimum designs. Çlabuk and Modi (1992) applied adjoint optimisation (gradient-based method applied to CFD) to locate a local optimum shape for a two-dimensional diffuser in laminar flow. To identify other possible optimum geometries in the design space, modifications to the initial diffuser were made to steer the adjoint solution to a design with a greater efficiency. Svenningsen et al. (1996) also utilised the adjoint methodology in a three-dimensional diffuser, and quasi-analytical sensitivity analysis to explore the design space. They also applied this method to a two-dimensional diffuser, focusing on the contouring of the diverging wall. This resulted in a 5% improvement of efficiency, compared with that of the original straight-walled geometry.

The use of Genetic Algorithms (GAs) in combination with CFD has allowed a number of authors to explore the design of a diffuser more effectively. Ghosh et al. (2010) optimised a two-dimensional symmetric diffuser in a turbulent flow regime. Representation of the diverging walls of the diffuser was achieved using two methods: the first being from a series of control points of a non-uniform rational basic spline (NURBS), and second using fourth-order polynomials of which the coefficients were used as design variables. Mariotti et al. (2013, 2014, 2015) utilised GAs to optimise the diverging walls of a series of two-dimensional diffusers in laminar (Mariotti et al. (2013)) and turbulent (Mariotti et al. (2014)) flow regimes. Their work focused on the optimisation of cavities along the walls - represented by a series of control points for Bézier curves. It should be noted that despite the efficient representation, totalling 5 parameters, a substantial number of evaluations, between 150 and 175, were needed to obtain the 'optimum' design.

To complement the use of GAs to perform the optimisation, there has also been great interest in employing surrogate modelling to efficiently optimise the diffuser geometry. For example, Madsen et al. (2000) used a response surface technique to optimise the shape of a two-dimensional diffuser in the turbulent flow regime. The shape of the diverging wall was described using polynomials and B-splines. Later work, (Madsen and Langthjem (2001)) focused on optimising the divergence angle of the expanding region of the diffuser. This reduced the number of evaluations necessary to construct the surrogate model by incorporating a multi-fidelity approximation of the search space. More recently, Dehghani et al. (2015) optimised the diffuser expansion region, and subsequently focused their analysis into the flow separation characteristics along the optimum designs. They observed that the number of evaluations needed to optimise the geometry was approximately 25% of that of the equivalent study using GAs for the optimisation strategy.

The literature reviewed above has invariably focussed on single-objective optimisation of the diffuser. To the best of the authors' knowledge, a multi-objective optimisation of the diffuser design has not yet been reported in the literature. Moreover,

the frameworks used for automated optimisation in diffusers have typically employed commercial software to reconstruct the geometry and conduct the CFD analysis. The present work will outline an entirely open-source framework for the multi-objective optimisation of diffuser geometries.

1.1. *Paper outline*

The present work outlines a methodology for, and application of, a multi-objective optimisation of an asymmetric diffuser. To achieve this aim, the paper is structured as follows: §2 outlines the base geometry dimensions; this is followed by §3 which outlines the CFD methodology for assessing the flow through, and performance of, the diffuser design. The performance indicators for a diffuser are outlined in §4. The grid regeneration technique for each CFD evaluation in the optimisation run is outlined in §5; this proposed methodology is subsequently validated through a grid convergence study. §6 describes a recently-proposed multi-objective Bayesian methodology. This is closely followed by §7, which outlines the procedures (and parameterisation) of deforming the geometry for each evaluation.

In the second part of the work, a series of multi-objective shape optimisations are carried out in §8, varying the number of control points for the wall geometry. The computational details for this are outlined in §7.1. The profile of the diverging wall is optimised with the aim of maximising the pressure recovery and outflow uniformity. Finally, in §9, the observations, and suggestions for future work are summarised.

2. Base diffuser geometry

The present work focuses on the turbulent flow through an asymmetric two-dimensional diffuser, denoted in the literature as the ‘Obi’, ‘Buice’ or ‘Buice-Eaton’ Diffuser - named after the authors of a series of experimental investigations (Obi et al. (1993); Buice and Eaton (2000); Buice (1997)). From a CFD modelling perspective, the flow through this geometry is challenging - characterised by the boundary layer separation and reattachment along the diverging wall. As a result, this geometry has largely served as a benchmark test case for assessment of turbulence model performance in CFD. Several studies of this nature are documented in the literature, including the performance of Reynolds-Averaged Navier-Stokes (RANS) turbulence modelling (e.g. Kaltenbach et al. (1999); Wu et al. (2006); Apsley and Leschziner (2000); El-Behery and Hamed (2011); Iaccarino (2001); Herbst et al. (2007)) and Large-Eddy Simulations (Kaltenbach et al. (1999); Wu et al. (2006)). Thus, a well consolidated experimental and numerical database of benchmark data is available for this geometry.

The tested diffuser can be divided into three sections: an inflow channel having constant width, an asymmetric diverging channel, and an outflow channel (tailpipe) having again constant width (see Fig.1). These three parts are connected at sharp edges. The upstream channel is sufficiently long to obtain a fully developed turbulent channel flow at the inlet of the diverging section, as in El-Behery and Hamed (2011); Mariotti et al. (2014).




Figure 1.: Schematic of the diffuser geometry.

The diffuser is characterised by the following quantities: the inlet width, h (used here as reference length), the outlet width $k = 4.7h$, the length from the inlet to the beginning of the diverging part, $l_1 = 60h$, the length of the diffuser diverging part, $l_2 = 21h$, the length from the end of the diffuser diverging section to the outflow (tailpipe), $l_3 = 60h$, and the divergence angle, $\alpha = 10^\circ$. The diffuser area ratio, i.e. the ratio between the diffuser outlet and inlet cross-areas, is $AR = 4.7$. Moreover, the total diffuser length is $l = (l_1 + l_2 + l_3) = 141h$.

Note that in the proceeding figures, dimensionless coordinates are used, i.e. $X = x/h$ and $Y = y/h$ (capital letters are used for dimensionless parameters and lowercase letters for dimensional quantities).

3. Simulation set-up and numerical methodology

Assessment of the diffuser performance was conducted using the open-source CFD code OpenFOAM-4.x. The simulations are carried out with Reynolds number $Re = hu/\nu = 20,000$, where u is the streamwise-velocity component at the diffuser centreline at $X = -6.5$. This choice is due to the fact that at this cross-section a fully-developed turbulent flow has developed, as witnessed by a velocity profile characterised by a ratio between the centreline velocity and the bulk velocity $u_{cl}/u_b = 1.14$ (Dean (1978); Buice (1997)). The chosen Reynolds number of the inflow is consistent with experimental (Obi et al. (1993); Buice (1997); Buice and Eaton (2000)) and numerical (Kaltenbach et al. (1999); El-Behery and Hamed (2011); Mariotti et al. (2014)) cases from the literature.

Since the flow considered in this work has been observed to be steady (El-Behery and Hamed (2011); Mariotti et al. (2014)), steady-state simulations were carried out in the present work. The Finite Volume Method was employed to integrate the governing equations (Weller et al. (1998)). The second-order central difference scheme was used to discretise the diffusion terms, and the second-order upwind difference was adopted to approximate the convection term. For the pressure-velocity coupling, the SIMPLE algorithm (Patankar and Spalding (1972)) was used, with under-relaxation factors 0.7, 1, and 0.7 for the velocity, pressure, and turbulence quantities respectively. The generalised Geometric-Algebraic Multi-Grid solver was used for pressure, while the Gauss-Seidel linear solver was used for the remaining fields.

Regarding the boundary conditions, the characteristics of turbulence at the inflow were specified with turbulence intensity ($= u'_b/u_b$) equal to 1% (i.e. smooth flow), and turbulent length scale (\mathcal{L}/h) equal to 0.038. As mentioned above, the length of the inflow channel ensures that a fully developed turbulent flow was achieved by the diffuser reference position $X = -6.5$. A pressure boundary condition is imposed at the

outflow boundary ($X = 80$) and no-slip conditions were applied along the solid walls.

4. Diffuser performance indicators

The main purpose of the diffuser is to recover the kinetic energy through the inflow by increasing the pressure head. A typical measure of this conversion is given by the pressure recovery factor,

$$C_p = \frac{1}{\frac{1}{2}\rho u_{ref}^2} \left[\frac{1}{A_{out}} \int_{A_{out}} p_{out} dA_{out} - \frac{1}{A_{ref}} \int_{A_{ref}} p_{ref} dA_{ref} \right], \quad (1)$$

where p is the static pressure, and A denotes the cross-sectional area of the boundary. The subscript *out* denotes the outflow boundary section ($X = 81$). p_{ref} and u_{ref} are the static pressure and streamwise-velocity (bulk velocity, u_b) components at the diffuser reference (*ref*) section, taken at $X = -6.5$. Maximising C_p is conventionally chosen as the main objective (cost function) when optimising the diffuser geometry, and other turbomachinery apparatus (Daniels et al. (2018)). A simple model for diffuser pressure recovery can be achieved from inviscid, one-dimensional theory combined with empirical loss coefficients. The ideal (frictionless) performance of a diffuser follows from Bernoulli's theorem, which expresses the energy conservation of the flow along a streamline. From this, the resulting pressure recovery factor, $C_{p,ideal}$ is derived,

$$C_{p,ideal} = 1 - \left(\frac{1}{AR} \right)^2 = 1 - \left(\frac{h}{k} \right)^2. \quad (2)$$

For the diffuser considered in the present work (with Area Ratio $AR = 4.7$), $C_{p,ideal} = 0.9548$. A more suitable estimate for the pressure recovery can be attained by subtracting the frictional losses from this ideal performance. To this end, expressions for the frictional losses for the expansion (l_2 in Fig.1) and tailpipe (l_3 in Fig.1) sections are introduced:

$$C_{p,ideal-loss} = \underbrace{1 - \left(\frac{h}{k} \right)^2}_{\text{ideal recovery}} - \underbrace{K_{loss}}_{\text{expansion loss}} - \underbrace{\lambda \frac{l_3}{k} \left(\frac{h}{k} \right)^2}_{\text{tailpipe loss}}. \quad (3)$$

K_{loss} is the non-recoverable loss coefficient for a gradual expansion. This empirical coefficient can be obtained from various engineering databases or textbooks (e.g. Blevins (1992); ESDU:74015 (2007)), and is a function of the aspect ratio and length of the diverging section. Furthermore, according to Reid (1953), for the geometry considered, the divergence angle, α , is 5° for an equivalent symmetric diffuser - the corresponding value collocated in ESDU:74015 (2007) for K_{loss} is approximately 0.2. However, it should be noted that the asymmetric diffuser is generally considered to be less efficient than the equivalent symmetric design (Reid (1953)) and thus some degree of error is expected around the true value.

Finally, the third term in Eq.3 takes into account the static pressure loss along the straight tailpipe (l_3 in Fig.1). From the classical theory of pipe friction, the pressure recovery along this section may simply be described using the Darcy-Weisbach formula,

$$\Delta p = \lambda \frac{l_3}{k} \rho \frac{u_2^2}{2} = \lambda \frac{l_3}{k} \left(\frac{h}{k}\right)^2 \rho \frac{u_b^2}{2}, \quad (4)$$

where, u_2 represents the intake velocity of the tailpipe, and is derived from the continuity equation: $u_b = u_2(k/h)$. The dimensionless friction factor, λ , is found from the following approximation to the exact (but implicit) solution of Prandtl for a smooth pipe. The error generated from using this approximation is less than 1% for Reynolds numbers in the range below 10^8 (Madsen and Langthjem (2001); Eck (1988)),

$$\lambda = \frac{0.309}{\log(Re/7)^2}. \quad (5)$$

Overall, for the geometry considered, the resulting value for $C_{p,ideal-loss}$ is approximately 0.75. It can be seen from the above derivation that the main source of frictional losses occurs in the expansion region and is therefore the focus for optimisation. The low-fidelity model described above, Eq.3, does not take into account the recovery from flow redistribution across the expansion and tailpipe regions. As a consequence, for the excluded effects of the non-uniform velocity profiles at the outflow, the total kinetic energy in the low-fidelity model is larger than for a uniform outflow with corresponding mass-flux. Thus, CFD allows for a better estimation of the pressure recovery along the diffuser - able to simulate the flow separation and non-uniform distributions of velocity along the design.

Along with an efficient recovery of pressure, another desirable characteristic of a diffuser is to maintain the uniformity of the outflow. In the literature, this feature is usually assessed after an optimisation run (e.g. Dehghani et al. (2015)) rather than being used as a design objective. To measure this feature, the flow uniformity index, which is widely used in the automotive industry (e.g. Guhan et al. (2015)), was used to evaluate the distribution of velocity at the outflow, and is defined,

$$\gamma = 1 - \frac{1}{A_{out}} \int_{A_{out}} \frac{\phi}{2U_{ave}} dA_{out}, \quad (6)$$

where

$$\phi = ||U| - U_{ave}|, \quad U_{ave} = \frac{1}{A_{out}} \int_{A_{out}} |U| dA_{out}. \quad (7)$$

By this definition, the uniformity index is restricted to the range $\gamma \in [0, 1]$; in the present work, maximising γ is the second objective for optimisation. If $\gamma = 1$, the magnitudes and direction of local velocity are uniform across the outflow cross-section. For the present work, the uniformity index is expected to be high due to the extensive length of the tailpipe. Furthermore, as described above, increasing the uniformity of the outflow will aid the pressure recovery. As a result, the two performance objectives in the present work are expected to show a certain degree of correlation between them.

5. Automated geometry and grid regeneration

The geometry considered in this work was set up as a benchmark test case for the use of several institutions worldwide to study the accuracy of CFD. As a step to improve the insight of using CFD, and for an efficient optimisation procedure, the use of automated grid-regeneration techniques for capturing the key performance measures of the diffuser will be discussed in this section.

A fundamental requirement of an automated CFD grid adaptation procedure for shape optimisation is the ability to create a satisfactory mesh around a continuously changing domain. Two different methods can be used to perform such task: mesh deformation and mesh regeneration (Daniels et al. (2019)). The approach used in this work is to regenerate the mesh around the new design, rather than deforming the base mesh. Due to the large number of CFD simulations required during the optimisation run, the mesh resolution for each evaluation should satisfactorily obtain an accurate estimation of the cost functions within a reasonable time frame. For this purpose the automated meshing utility *cfMesh* was employed.

cfMesh-1.2.1 is an open-source library built on top of the OpenFOAM framework. To construct the CFD domain and grid, *cfMesh* requires a closed manifold-surface - typically a stereolithography file. A grid is constructed within the enclosed surface using a set of user-defined parameters required for the meshing algorithms; these also include parameters that define the boundary-layer mesh. The boundary layers are extruded from the chosen boundary faces of the manifold surface towards the interior. Moreover, *cfMesh* provides additional controls of boundary layer quality, intended for situations where a large number of layers is required, or when the thickness is required to vary smoothly; the majority of these parameters were kept as default. By experimentation, the most influential parameters needed for a mesh sensitivity test were reduced to a set of 3:

- *maxCellSize*: defines the maximum cell size generated in the grid;
- *localRefinement*: prescribes the surface cell size on a desired boundary (or patch) in the CFD domain;
- *maxFirstLayerThickness*: prescribes the first wall-normal cell height to a specified boundary.

For the present simulations, the ‘*maxFirstLayerThickness*’ parameter was fixed for the diffuser walls so the average distribution of $y_1^+ \leq 1$. Consequently, no wall models were used for turbulence. The $k-\omega$ Shear Stress Transport (SST) turbulence model was employed for all simulations in the present work, due to its robustness for predicting the flow separation along the Buice diffuser (El-Behery and Hamed (2011); Jiang et al. (2011); Mariotti et al. (2014)).

The chosen regions for refinement (or boundary-layer generation) were along the diffuser walls. A mesh sensitivity study was performed varying the global and local refinement parameters. For the sake of brevity, three mesh resolutions are presented here. The corresponding values are listed in Table 1. ‘Mesh A’ is the coarsest mesh, with a total number of cells equal to 165,007. ‘Mesh B’ has a smaller maximum cell-size than ‘Mesh A’, but with the same local refinement; the total number of cells was 651,494. Finally, ‘Mesh C’ has increased local refinement and maximum cell size to ‘Mesh B’; the total number of cells was approximately 4 million.

Mesh	Refinement		Cost functions	
	maxCellSize	localRefinement	C_p	γ
A	0.050	0.025	0.65292	0.94927
B	0.025	0.0250	0.70682	0.94396
C	0.010	0.0125	0.70682	0.94396

Table 1.: User-defined parameters varied during the mesh sensitivity study and resulting diffuser performance (C_p, γ).

Simulations using the numerical setup described in §3 were performed on the three meshes generated using cfMesh with input parameters outlined in Table 1. The convergence of the simulations was defined with the residuals for the flowfield (velocity, pressure, and turbulence quantities) and pressure recovery (C_p , Eq.1) falling below 10^{-7} . It can be seen in Table 1 that the cost functions show little sensitivity between ‘Mesh B’ and ‘Mesh C’. Furthermore, it can also be deduced from the convergence between ‘Mesh A’ and ‘Mesh B’ that the pressure recovery is more sensitive to the grid than the flow uniformity. Validation of the grid regeneration is performed by assessing the wall pressure coefficient along the diffuser, C_{pw} , which is defined,

$$C_{pw} = \frac{p_{wall} - p_{X=-1.69,wall}}{\frac{1}{2}\rho u_{X=-1.69}^2}, \quad (8)$$

where p_{wall} is the local static pressure on the wall. Fig.2 shows the profiles along the lower and upper walls. It can be seen that the pressure recovery for ‘Mesh A’ is consistently lower than the target data from the literature, while the results for ‘Mesh B’ and ‘Mesh C’ are nearly identical. It is also interesting to observe the development of pressure recovery along the diffuser, particularly when considering the geometrical changes considered later in this paper. To achieve this, a series of sample planes were placed along the diffuser. The performance quantities were attained from these planes using Eq.1, where *out* is synonymous with the position of the sample plane, and *ref* was sampled at $X = -6.5$. The resulting profiles can be seen in Fig.2, showing a similar trend to those for the wall pressure (Fig.2). Thus, it can be concluded that the predictions given by ‘Mesh B’ and ‘Mesh C’ satisfactorily agree with the reference data. Thus, for the sake of efficiency, the simulations of the optimisation run were performed using the grid parameters for ‘Mesh B’.

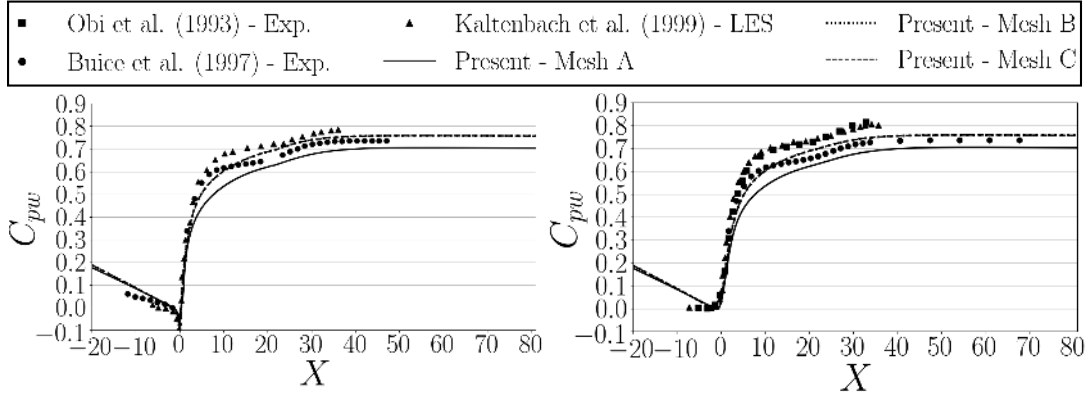


Figure 2.: Comparison between present predictions, LES results of Kaltenbach et al. (1999) and experimental results of Obi et al. (1993); Buice (1997) for wall pressure recovery, C_{pw} (Eq.8); left - diverging (bottom) wall; right - straight (top) wall.

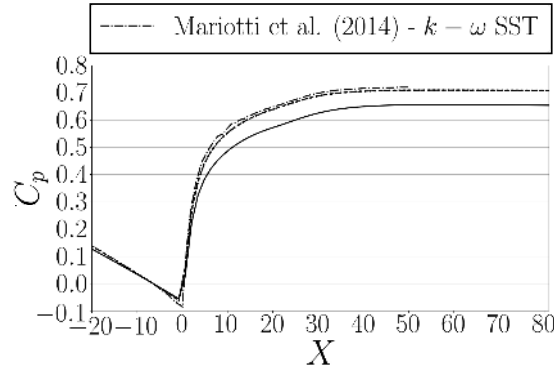


Figure 3.: Comparison between present predictions, and equivalent simulation results by Mariotti et al. (2014) for pressure recovery along the diffuser, C_p (Eq.1). The corresponding legend for the plots can be seen in Fig.3.

The positions of the points in the separated region where the streamwise velocity is equal to zero are compared in Fig.4 against data from the literature. It can be seen in this figure, that despite the grid resolution, the $k - \omega$ SST turbulence model has given a consistent prediction for the extent of the separated zone, which are in overall good agreement with the available experimental and numerical data.




Figure 4.: Position where streamwise velocity crosses zero compared with LES results of Kaltenbach et al. (1999) and experimental results of Obi et al. (1993); Buice (1997). The corresponding legend for the plots can be seen in Fig.3.

6. Multi-objective shape optimisation

In this paper, the primary goal is to simultaneously maximise the pressure recovery factor along the geometry and flow uniformity index at the outflow. Consider that a particular diffuser geometry (or shape) may be represented with an n -dimensional decision vector $\mathbf{d} \in \mathcal{D}$, where $\mathcal{D} \subseteq \mathbb{R}^n$ consists of all feasible shapes. Then the multi-objective optimisation problem may be expressed as¹:

$$\min_{\mathbf{d} \in \mathcal{D}} f_1(\mathbf{d}) = -C_p, \quad (9)$$

$$\min_{\mathbf{d} \in \mathcal{D}} f_2(\mathbf{d}) = -\gamma. \quad (10)$$

The objectives are potentially conflicting and to a certain degree anti-correlated (see §4): a shape that improves one objective results in a (near-)proportional decrease in performance for the other objective. As such, generally there is not a unique solution to this problem, but a range of shapes that trade-off between the objectives. The trade-off relationship is characterised by the notion of dominance (Coello Coello et al. 2007). A shape \mathbf{d} is said to dominate another shape \mathbf{d}' , denoted as $\mathbf{d} \prec \mathbf{d}'$, iff,

$$\begin{aligned} & f_1(\mathbf{d}) < f_1(\mathbf{d}') \text{ and } f_2(\mathbf{d}) \leq f_2(\mathbf{d}') \\ \text{or } & f_1(\mathbf{d}) \leq f_1(\mathbf{d}') \text{ and } f_2(\mathbf{d}) < f_2(\mathbf{d}'). \end{aligned} \quad (11)$$

The set of shapes that provide an optimal trade-off between the objectives is referred to as the Pareto set:

$$\mathcal{P} = \{\mathbf{d} \mid \mathbf{d}' \not\prec \mathbf{d} \ \forall \mathbf{d}', \mathbf{d}' \in \mathcal{D} \wedge \mathbf{d} \neq \mathbf{d}'\}. \quad (12)$$

The image of the Pareto set \mathcal{P} in the objective space is known as the Pareto front \mathcal{F} . It may not be possible to locate the exact Pareto set within a practical time limit, even if the objective functions were computationally cheap. Therefore, the overall goal of an effective optimisation approach is to generate a good approximation of the Pareto set denoted as $\mathcal{P}^* \subseteq \mathcal{D}$, and the associated estimated Pareto front is denoted as \mathcal{F}^* .

It should be noted that if two objectives are perfectly anti-correlated, then the knowledge about one objective is sufficient to compute the other. As a consequence, any alterations in the decision vector would result in proportionally equal changes in both objectives, and hence all solutions $\mathbf{d} \in \mathcal{D}$ will be mutually non-dominated and belong to the optimal trade-off front \mathcal{F} . This renders any multi-objective optimisation method inefficient and unnecessary. A better alternative is to determine a specific target for one objective (which sets the target for the other objective accordingly), and reduce the multi-objective problem into a single objective optimisation of achieving the desired target. Recall, as discussed in §4, the objectives considered here are not perfectly correlated due to complex interactions between flow considerations and unusual shapes, and thus there are some shapes that are dominated by others. Therefore, the problem here must still be treated as a multi-objective problem.

¹Maximising a function $f_i(\cdot)$ is equivalent to minimising $-f(\cdot)$

7. Spline representation for optimisation

The optimisation run in the present work alters the profile of the diverging (lower) wall. The methodology for changing the shape was chosen so as to be able to form a tapered shape, or form vortex cavities of varying width, depth, and position along the wall. For this purpose, Catmull-Clark subdivision curves (Stam (1998)) were utilised. As a curve may be completely defined with a finite set of control points, the diffuser design optimisation may then be considered as manipulating the positions of the control points to locate the optimal shape. Therefore, a vector of control point coordinates was used to represent the design space for m control points, where the i th control point $\mathbf{c}^i = [c_x^i, c_y^i]^\top \in \mathbb{R}^2$. The horizontal (x) and vertical (y) elements of the m control points are incorporated in a decision vector in the following manner: $d_{2(i-1)+4} = c_x^i$ and $d_{2(i-1)+5} = c_y^i$, where $i \in [1, m]$. Additionally, the optimisation is subject to all control points residing within the predefined space, and the resulting curve being a simple (non-intersecting) polygon.

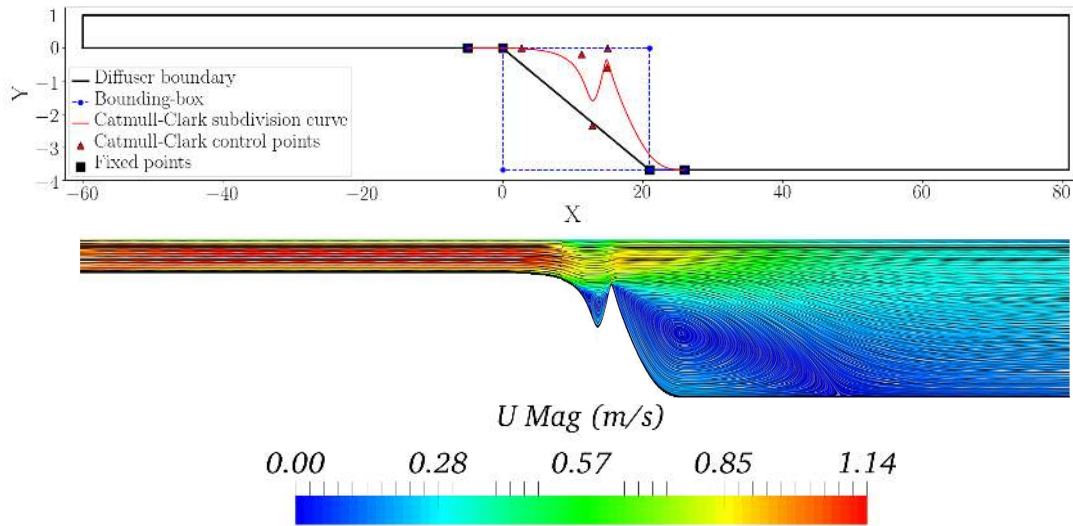


Figure 5.: Top: a schematic of the Catmull-Clark subdivision curve setup, and a randomly generated subdivision curve using five control points; bottom: velocity contour and streamlines of the resulting flowfield from the new (random) design.

To demonstrate the Catmull-Clark subdivision curve methodology, Fig.5 shows the use of five control points to generate the diverging wall of the diffuser. The control points (red triangles) are randomly generated within predefined bounding-boxes (circle markers, dotted lines). The bounding-box for the wall was defined to ensure that the lower wall does not fall below the outflow. At the same time, the box is constricted to be no higher than the inflow cross-section, but still allows for an extension of the backward-facing step. To ensure a smooth (unperturbed) transition between the inflow (l_1) and tailpipe (l_3) sections, fixed points are prescribed at either end of the subdivision curve (black squares). It should be noted that from a practical perspective, only a few iterations of subdivision usually results in a visually smooth curve. The iteration limit was set to 5 in the present work.

7.1. Computational details

The multi-objective Bayesian Optimisation methodology proposed in Rahat et al. (2017) was used to generate the results in this paper². The optimiser requires a few parameters to be set *a priori*. These settings are discussed below.

Firstly, it is essential to set the number of initial evaluations, so that a surrogate model may be constructed. Traditionally, this is set to $M = 11n - 1$, where n is the number of dimensions in the decision vector \mathbf{d} (Jones et al. (1998)). Here, $n = 2m$ and thus the number of initial samples was varied in accordance to the number of control points for the Catmull-Clarke subdivision curve (m). Usually, the boundary of the decision space is defined with a set of lower and upper limits corresponding to each dimension to enable Latin hypercube sampling and generate the initial set of decision vectors $D = \{\mathbf{d}^1, \dots, \mathbf{d}^M\}$. Here, additional constraints were imposed to ensure that no control points resided beyond the predefined bounding box as shown in Fig.5. Thus any Latin hypercube samples that violated these constraints was rejected. As a consequence, in one instance ($m = 4$), only 28 out of 87 initial samples were deemed feasible recasting $M = 28$. Note that the precise number of initial feasible solutions may vary as repeated Latin hypercube samples do not usually result in a unique set of shapes. These M feasible shapes were evaluated to acquire the function values and construct the training data set – which is used to train a Bayesian model relating decision vectors to the *hypervolume improvement* that measures the contribution of a solution in improving the estimation of the Pareto set. The other settings used in the MBO framework were: the maximum number of surrogate evaluations to maximise the expected improvement using Bi-POP-CMA-ES was set to $12000n$, the reference point $\mathbf{l} = [l_1, l_2]^T$ to compute the hypervolume was dynamically determined with $l_i = \max_{\mathbf{d} \in D} f_i(\mathbf{d}) + 0.1(\max_{\mathbf{d} \in D} f_i(\mathbf{d}) - \min_{\mathbf{d} \in D} f_i(\mathbf{d}))$, and the overall budget on the number of CFD simulations was limited to the number of initial samples plus $30n$. It should be noted that, while locating \mathbf{d}^{M+1} using Bi-POP-CMA-ES, each candidate solution is checked for feasibility – which is computationally cheap to determine. An infeasible solution \mathbf{d}' is deemed to have zero expected improvement (that is a utility function used to locate promising solutions in multi-objective Bayesian optimisation). This effectively discounts any infeasible solutions, and enables the evaluation of only feasible solutions using CFD.

On the CFD simulation side, a steady-state simulation was performed for each evaluation using the numerical setting outlined in §3. Each CFD simulation was assumed converged when all the residual tolerances had reduced to approximately 10^{-6} , which is more than sufficient for most engineering problems. The iterative convergence for C_p was also included in this criterion. All simulations were carried out on one node 16CPU 2x Intel Haswell E5-2640v3 2.6GHz cores. The wall-time for each evaluation was approximately 5 minutes.

8. Optimisation results and discussion

The details of the optimisation runs with increasing number of control points can be seen in Table 2. As described above, the number of evaluations increases proportionally with the number of control points. The number of initial samples reduces with the increasing number of control points as an increasing number of samples are rejected for violating the design constraints (i.e. outside the bounding box). It can be seen that the

²Python code for the MBO framework is available at: <http://bitbucket.org/arahat/gecco-2017>.

performance measures of the diffuser have significantly improved for each run, with a potential increase of $\sim 13\%$ for pressure recovery, and $\sim 2.5\%$ for the flow uniformity index. It is interesting to note that increasing the number of control points for the present spline representation has resulted in an increase of the maximum cost function values - although the improvement to diffuser performance beyond 3 control points is not significant. Thus, the designer has to consider the trade-off of computational cost with increasing number of control points to the incremental improvement of diffuser performance, as well as the manufacturability of the resulting design. This observation does not agree with those of Ghosh et al. (2010); Dehghani et al. (2015), who, having used Non-Uniform Rational Basis Splines (NURBS) as a representation for the lower wall, show that the performance of the diffuser reduces after a certain number of control points. From this it can be deduced that the application of Catmull-Clarke subdivision curves in the present work have allowed for a more comprehensive investigation of the design space.

No. control points (m)	Initial samples	Total evaluations	max C_p	max γ
Base design (0)	(-)	(-)	0.706819	0.943960
1	21	81	0.823481	0.956001
2	43	163	0.827739	0.958678
3	32	212	0.827912	0.960429
4	28	268	0.827988	0.967166
5	21	321	0.828988	0.967268
6	11	371	0.829632	0.968074

Table 2.: Details of the optimisation runs with increasing number of Catmull-Clarke subdivision curve control points, and the maximum values for the cost functions for each optimisation run.

The resulting objective space of the optimisation runs are shown in Figs.6-9. A few diffuser designs evaluated during the optimisation run are indicated in these diagrams for discussion. The data points in these diagrams are coloured according the evaluation number during the optimisation process (grey-scale). Regardless of the number of control points, these diagrams show a wide range of conventional and unusual designs for the diffuser. The base design is indicated by a green star. From initial impressions, when considering the first objective of the optimisation, $\min -C_p$, the designs that are unconventional (difficult to manufacture) are consistently located to the right-hand-side of the base geometry, i.e. a poorer pressure recovery. It can be seen that these designs largely possess sharp increases of curvature along the expansion section, resulting in aggressive changes in cross-sectional area along the geometry. The resulting regions of flow separation along the geometry have a detrimental impact on the pressure recovery and outflow uniformity. On the other hand, the more traditional diffuser designs, to the left-hand-side of the base geometry, are seen to be closer to (and on) the estimated Pareto front (\mathcal{F}^*) – indicated by the red squares.

As is the nature of multi-objective optimisation, the estimated optimal trade-off front leaves the final choice of optimal design to the engineer. However, it can be seen in Figs.6-9 that a small estimated Pareto set has been attained regardless of the number of control points. Based on the low-fidelity model outlined in §4, it was speculated that the performance indicators chosen for the present optimisation will show some correlation between them. Table 3 shows the correlation values between the

two objectives for each optimisation run. From this data, the two objectives clearly show an inherent relation, but arguably not to the extent of treating the problem as a single-objective problem.

No. control points (m)	Corr(C_p, γ)
1	0.880
2	0.786
3	0.803
4	0.797
5	0.755
6	0.871

Table 3.: Correlation coefficients between the diffuser performance measures with increasing number of control points.

In the present work, a few designs selected along \mathcal{F}^* are labelled ‘ $i-iv$ ’ (in Figs.6-9) for further examination. It can be seen that the optimal designs of the diverging wall show four distinct characteristics:

- (1) The designs consistently have a lower divergence angle than the base design;
- (2) a small recirculation region is formed either on the top or bottom wall in the tailpipe region;
- (3) no flow separation is formed along the expansion region (unlike in the base design);
- (4) all optimum designs show a slight expansion of the inflow pipe with an incremental increase of cross-sectional area into the divergence region.

Regarding the first point, it can be seen that the divergence angle for the optimum designs with different numbers of control points $m = 1$ to 4 are consistent with each other. The slight improvement of pressure recovery was achieved by increasing the number of control points, allowing for a smoother curvature at the tailpipe opening. Increasing the number of control points allows for the formation of vortex-cavities into the wall as seen for the optimal design with $m = 6$ control points (Fig.9). Finally, it is worth noting that the optimum geometries of this work are remarkably similar to the optimum shape attained from Mariotti et al. (2014) - who also showed that a smaller divergence angle with a small recirculation region in the wake increased the pressure recovery.

Regarding the second and third points, it has been reported many times in the literature that delaying or mitigating flow separation along the diffuser geometry increases the pressure recovery (e.g. Mariotti et al. (2013, 2014, 2015); Dehghani et al. (2015)). A noticeable trend in the present results is that the designs which perform poorly inherently have large recirculation regions, which contribute to the frictional losses in the expansion region. For example, the design ‘ vi ’ Fig.8 mimics a step-diffuser (or backward-facing step). According to ESDU:72011 (2007); ESDU:74015 (2007), the sudden expansion at the opening of the tailpipe increases K_{loss} (see §4) from 0.2 (base) to 0.45. The present results agree with this trend, as ‘ vi ’ has a worse performance than the base design. Conversely, (in relation to the final point), creating a small expansion at the inflow allows for the separation of the flow to be delayed (Dehghani et al. (2015)) or contribute to early separation of the flow - which has previously shown in an increase in the pressure recovery (Mariotti et al. (2013)). It should be noted

that the observations made of the flowfield in the present work (and indeed in similar publications) are limited to the fidelity of the CFD simulation used for each evaluation. The steady-state assumption, and choice of turbulence model ($k-\omega$ SST) in the present work, provides limited accuracy in cases where large flow separation regions are observed. Nevertheless, as described above, the diffuser designs which contain these flow features inherently perform poorly for both design objectives. As a result, the improved designs along the Pareto front contain flows with minimal flow separation and are thus satisfactorily captured through the present CFD methodology.

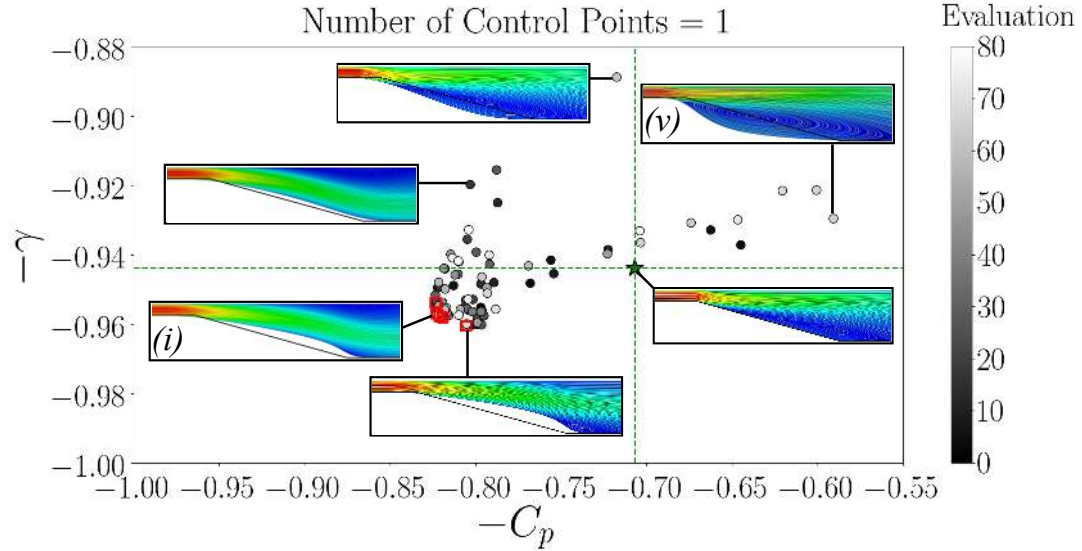


Figure 6.: The objective space of the two cost functions (Eqs.1 and 6), and the estimated Pareto front (\mathcal{F}^* , red squares) for 1 control point for the Catmull-Clarke subdivision curve.




Figure 7.: The objective space of the two cost functions (Eqs.1 and 6), and the estimated Pareto front (\mathcal{F}^* , red squares) for 3 control points for the Catmull-Clarke subdivision curve.

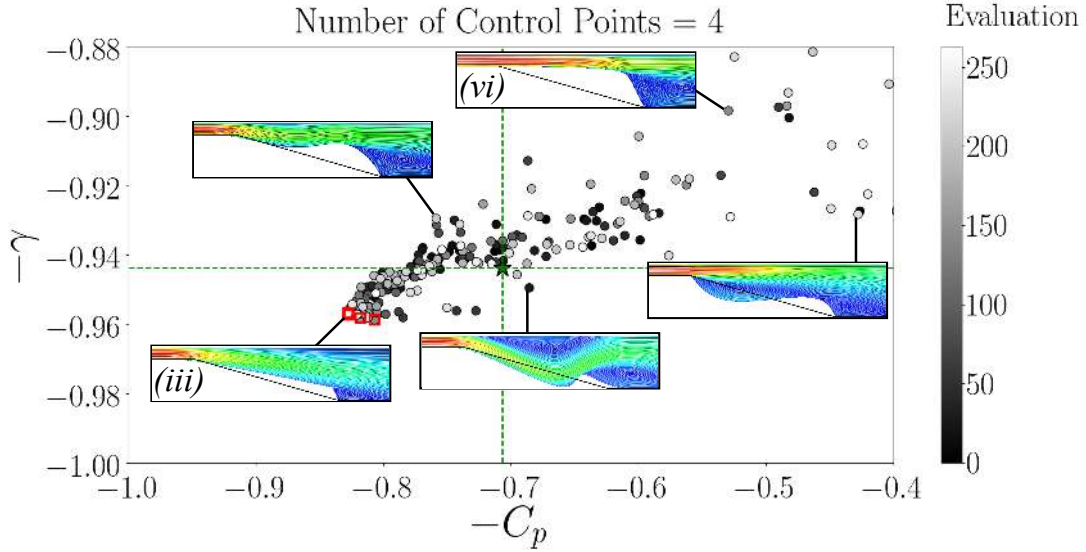


Figure 8.: The objective space of the two cost functions (Eqs.1 and 6), and the estimated Pareto front (\mathcal{F}^* , red squares) for 4 control points for the Catmull-Clarke subdivision curve.

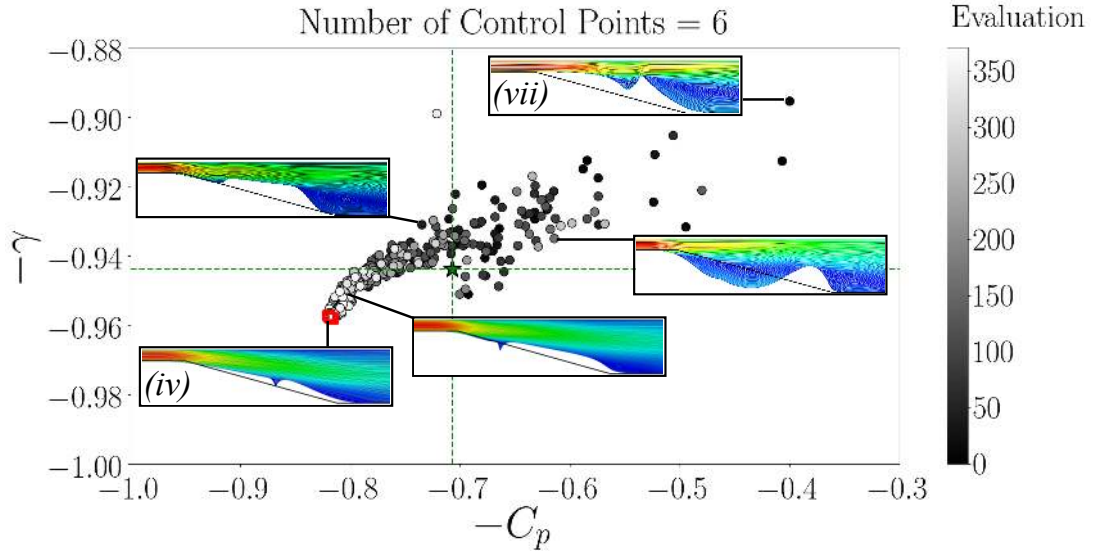


Figure 9.: The objective space of the two cost functions (Eqs.1 and 6), and the estimated Pareto front (\mathcal{F}^* , red squares) for 6 control points for the Catmull-Clarke subdivision curve.

The chosen spline representation in the present work allows for the consideration of efficient diffuser designs, such as the ‘bell-shaped’, ‘trumpet shaped’ and ‘inflected-wall shape’ diffusers (ESDU:76027 (2007)). Expansion of the diverging wall into a convex shape (bell) increases the longitudinal pressure gradient near the entry to the diffuser - where the flow is most able to endure a high pressure gradient without separation; the pressure gradient subsequently decreases to prevent or mitigate separation further downstream. As the bell-shaped diffuser increases the wall-angle at the entrance to the expansion region, it is important to avoid a sharp corner at this point, otherwise a large

separation region is formed and loss in pressure recovery results - as demonstrated in ‘*v*’ in Fig.6. A smaller wall-angle ($2 - 3^\circ$, ESDU:76027 (2007)) with a rounded entry corner avoids separation (Mariotti et al. (2013)); as described above, in the present work, all the identified optimum geometries have this design feature.

In contrast to the bell-shaped diffuser, the trumpet-shaped diffuser (concave) is designed to delay separation by reducing the pressure gradient early in the diffuser. However, the sudden increase in wall-angle further downstream (at the entrance of the tailpipe) may cause separation. The gradual increase in cross-sectional areas near the entry of the expansion region allows for the growth of the boundary layer (and reduction of effective area), which reduces the associated K_{loss} coefficient (ESDU:76027 (2007)). Most of the optimum designs identified in this work follow this design trend with a small recirculation region formed in the wake of the expansion region on the lower or upper wall (e.g. ‘*iii*’, Fig.8).

Finally, the inflected-wall design can be seen in ‘*i*’, Fig.6. As the flow develops along the length of the expanding section, it becomes increasingly less able to withstand the longitudinal pressure gradient (governed largely by the rate of change in cross-sectional area), resulting in an increased likelihood that the flow will separate. Curving the walls alters the pressure gradient characteristics and by extension the boundary-layer development. In accordance to ESDU:76027 (2007), the wall-angle connecting the expansion region to the tailpipe in ‘*i*’ does not exceed $17 - 18^\circ$.




Figure 10.: Diffuser designs found on the Pareto front with varying number of control points. Top: the design profile; bottom-left: the pressure recovery along the associated geometry; bottom-right: the velocity profiles at the outflow.

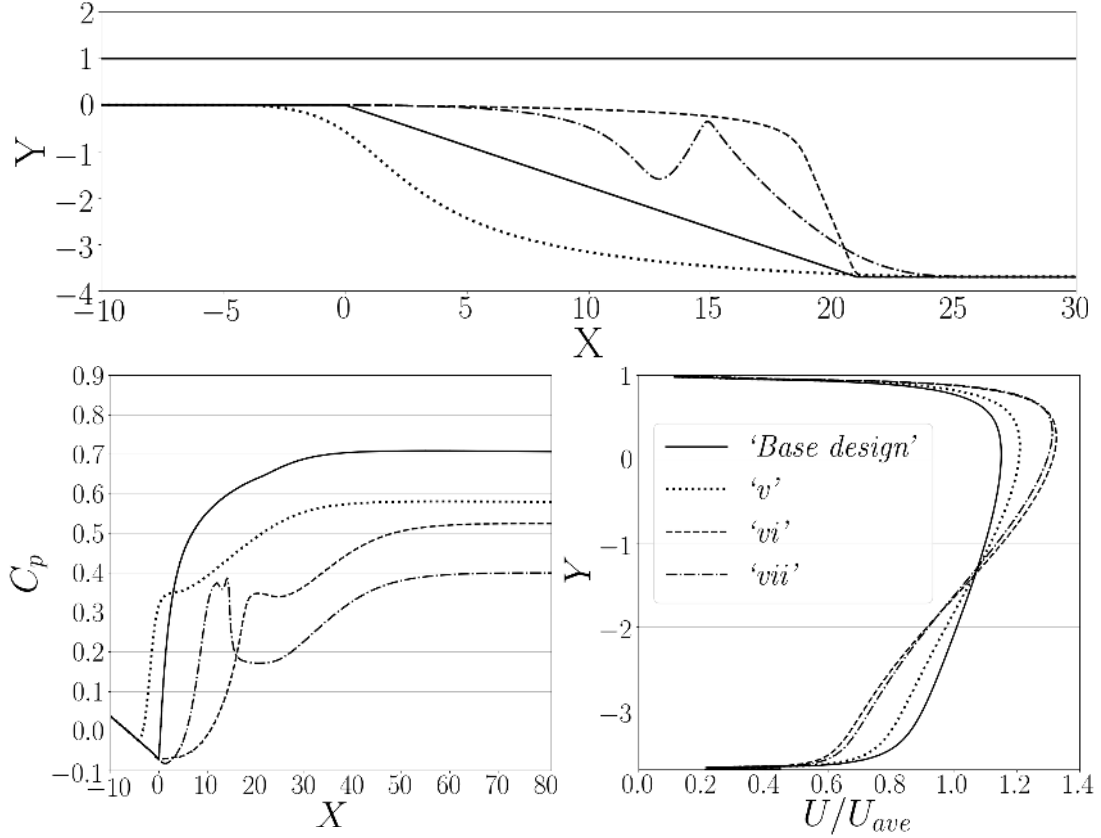


Figure 11.: Diffusers with a worse performance than the base diffuser with varying number of control points. Top: the design profile; bottom-left: the pressure recovery along the associated geometry; bottom-right: the velocity profiles at the outflow.

To demonstrate the optimum design characteristics for diffuser designs, Fig.10 shows the resulting performance indicators for the designs identified on the Pareto front. Regarding the pressure recovery, it can be seen that the expansion at the end of the inflow, and a smooth progression to the expansion section, has significantly improved the pressure recovery. The resulting ‘inflected’ and ‘trumpet’ shaped profiles of the diffuser have allowed a smooth expansion along the region. The shallow divergence angle (approximately $2-3^\circ$) of the wall has aided the pressure recovery (in accordance to Mariotti et al. (2013)). The resulting outflow velocity profiles are nearly symmetric around the centerline.

Finally, to demonstrate the characteristics of a poor diffuser design, Fig.11 shows the resulting performance indicators for the worse designs (compared to the base) identified in the objective spaces. It can be seen that in these designs, aggressive changes in cross-sectional areas significantly impairs the pressure recovery and flow uniformity.

9. Conclusions and future work

In the present work, the methodology and application of a multi-objective Bayesian optimisation to a two-dimensional asymmetrical ‘Buice’ diffuser has been performed. Reasonably high efficiency and performance has been achieved regardless of the num-

ber of design parameters, with a potential increase of $\sim 13\%$ for the pressure recovery, and $\sim 2.5\%$ for the flow uniformity at the outflow. These performance measures have shown to have some degree of correlation but nevertheless satisfy the case for a multi-objective run. The results of the optimisation run show consistently (regardless of design parameters) that the optimum asymmetrical diffuser design is that of a ‘trumpet-shape’ or ‘inflected-wall’.

The proposed fully automated methodology for optimisation based on the parameterisation using Catmull-Clarke subdivision curves, mesh regeneration using open-source software ‘*cfMesh*’, CFD simulations using OpenFOAM, and a recently proposed multi-objective Bayesian method (Rahat et al. (2017)) has been developed in a Python-based framework. The proposed methodology is also applicable for the optimisation of symmetrical diffuser designs.

Overall, the novel aspects of this work include:

- the application of a recently proposed multi-objective Bayesian optimisation method to a diffuser design;
- a proposed method for the automated reconstruction of the geometry and CFD grid for each evaluation;
- a review of the optimum diffuser designs to the current engineering codes and standards;
- the evaluation of the relation between the flow uniformity index and pressure recovery for diffusers.

9.1. *Future work*

This work naturally leads to the following topics of investigation involving optimisation of diffuser geometries:

- (1) To include the number of control points of the Catmull-Clarke subdivision curve as a design parameter;
- (2) to conduct a robustness study of the identified optimum designs over a wide range of inflow velocities;
- (3) to conduct a multi-fidelity study on the design using high-fidelity turbulence modelling techniques and transient simulations to capture unsteady flow phenomena in the designs;
- (4) to investigate the surrogate modelling of K_{loss} for unconventional wall designs (other than straight-wall);
- (5) to consider more constraints to the design, such as the uniformity of the inflow cross-section.

Acknowledgement(s)

The CFD simulations were performed on the ISCA HPC in the Advanced Computing Facility in the University of Exeter, UK.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant (reference number: EP/M017915/1) for the University of Exeter's College of Engineering, Mathematics, and Physical Sciences.

Nomenclature

A	Cross-sectional area (m^2)	λ	Frictional loss coefficient for straight pipe ($-$)
C_p	Pressure recovery factor ($-$)	k	tailpipe height (m)
C_{pw}	Wall pressure coefficient ($-$)	K_{loss}	Frictional loss coefficient for expanding sections ($-$)
p_w	Wall static pressure ($Kg\ m^{-1}s^{-2}$)	m	Number of control points ($-$)
\mathcal{F}^*	Estimated Pareto front	<i>out</i>	Outflow boundary
\mathcal{F}	Optimal Pareto front	<i>ref</i>	Reference position
\mathbf{l}	Reference vector for computing hypervolume	U	Velocity magnitude (ms^{-1})
\mathcal{P}^*	Estimated Pareto set	u_b	Bulk velocity (ms^{-1})
\mathcal{L}	Turbulent length scale (m)	u_c	Velocity along the inflow centerline (ms^{-1})
\mathcal{P}	Optimal Pareto set	U_{ave}	Average velocity magnitude (ms^{-1})
AR	Aspect Ratio ($-$)	X, Y	Nondimensional streamwise and vertical distances ($-$)
$C_{p,ideal}$	Ideal pressure recovery ($-$)	y_1^+	Nondimensional wall-normal distance ($-$)
h	Height of inflow section (m)		
l	Length of section (m)		
p	Static pressure ($Kg\ m^{-1}s^{-2}$)		
1, 2, 3	Index for diffuser sections (inflow, expansion, and tailpipe regions)		
γ	Flow uniformity index ($-$)		

References

- F. Durst, A. Melling, J. H. Whitelaw, Low reynolds number flow over a plane symmetric sudden expansion, *Journal of Fluid Mechanics* 64 (1974) 111–128.
- M. Nabavi, Three-dimensional asymmetric flow through a planar diffuser: Effects of divergence angle, reynolds number and aspect ratio, *International Communications in Heat and Mass Transfer* 37 (2010) 17–20.
- T. Suzuki, T. Colonius, S. Pirozzoli, Vortex shedding in a two-dimensional diffuser: theory and simulation of separation control by periodic mass injection, *Journal of Fluid Mechanics* 520 (2004) 187–213.
- Y. Tsui, C. Wang, Calculation of laminar separated flow in symmetric two-dimensional diffusers, *ASME. J. Fluids Eng.* 117 (1995) 612–616.
- S. Kline, D. Abbott, R. Fox, Optimum design of straight-walled diffusers, *ASME. J. Basic Eng.* 81 (1959) 321–331.
- B. Waitman, L. Reneau, S. Kline, Effects of inlet conditions on performance of two-dimensional subsonic diffusers, *ASME. J. Basic Eng.* 83 (1961) 349–360.
- R. Fox, S. Kline, Flow regimes in curved subsonic diffusers, *ASME. J. Basic Eng.* 84 (1962) 303–312.
- L. Reneau, J. Johnston, S. Kline, Performance and design of straight, two-dimensional diffusers, *ASME. J. Basic Eng.* 89 (1967) 141–150.
- A. Klein, Characteristics of combustor diffusers, *Progress in Aerospace Sciences* 31 (1995) 171–271.

- E. Göttlich, Research on the aerodynamics of intermediate turbine diffusers, *Progress in Aerospace Sciences* 47 (2011) 249–279.
- H. Lan, B. Armaly, J. Drallmeier, Turbulent forced convection in a plane asymmetric diffuser: Effect of diffuser angle, *ASME. J. Heat Transfer* 131 (2009) 071702–1– 071702–8.
- R. Mehta, P. Bradshaw, Design rules for small low speed wind tunnels, *The Aeronautical Journal* 83 (1968) 443–453.
- H. Çlabuk, V. Modi, Optimum plane diffusers in laminar flow, *Journal of Fluid Mechanics* 237 (1992) 373–393.
- K. Svenningsen, J. Madsen, N. Hassing, W. Päuker, Optimization of flow geometries applying quasi-analytical sensitivity analysis, *Appl. Math. Modelling* 20 (1996) 214–224.
- S. Ghosh, D. K. Pratihar, B. Maiti, P. K. Das, An evolutionary optimization of diffuser shapes based on cfd simulations, *International Journal For Numerical Methods In Fluids* 63 (2010) 1147–1166.
- A. Mariotti, A. Grozescu, G. Buresti, M. Salvetti, Separation control and efficiency improvement in a 2d diffuser by means of contoured cavities, *European Journal of Mechanics - B/Fluids* 41 (2013) 138–149.
- A. Mariotti, G. Buresti, M. Salvetti, Control of the turbulent flow in a plane diffuser through optimized contoured cavities, *European Journal of Mechanics - B/Fluids* 48 (2014) 254–265.
- A. Mariotti, G. Buresti, M. Salvetti, Use of multiple local recirculations to increase the efficiency in diffusers, *European Journal of Mechanics - B/Fluids* 50 (2015) 27–37.
- J. Madsen, W. Shyy, T. Haftka, Response surface techniques for diffuser shape optimization, *AIAA Journal* 38 (2000) 1512–1518.
- J. Madsen, M. Langthjem, Multifidelity response surface approximations for the optimum design of diffuser flows, *Optimization and Engineering* 2 (2001) 453–468.
- M. Dehghani, H. Ajam, S. Farahat, Automated diffuser shape optimization based on cfd simulations and surrogate modeling, *Journal of Applied Fluid Mechanics* 9 (2015) 2527–2535.
- S. Obi, K. Aoki, S. Masuda, Experimental and computational study of turbulent separating flow in a asymmetric plane diffuser, in: *Ninth Symposium on "Turbulent shear flows"* (1993), pp. 305.1 – 305.4.
- C. Buice, J. Eaton, Experimental investigation of flow through an asymmetric plane diffuser, *ASME. J. Fluids Eng.* 122 (2000) 433–435.
- C. Buice, Experimental investigation of flow through an asymmetric plane diffuser, Ph.D. thesis, Department of Mechanical Engineering, Stanford University, 1997.
- H. Kaltenbach, M. Fatica, R. Mittal, T. Lund, P. Moin, Study of flow in a planar asymmetric diffuser using large-eddy simulation, *J. of Fluid Mechanics* 390 (1999) 151–185.
- X. Wu, J. Schlüter, P. Moin, H. Pitsch, G. Iaccarino, F. Ham, Computational study on the internal layer in a diffuser, *J. of Fluid Mechanics* 550 (2006) 391–412.
- D. Apsley, M. Leschziner, Advanced turbulence modelling of separated flow in a diffuser, *Flow, Turbulence and Combustion* 63 (2000) 81–112.
- S. El-Behery, M. Hamed, A comparative study of turbulence models performance for separating flow in a planar asymmetric diffuser, *Computers Fluids* 44 (2011) 248–257.
- G. Iaccarino, Predictions of a turbulent separated flow using commercial cfd codes, *ASME. J. Fluids Eng.* 123 (2001) 819–828.
- A. H. Herbst, P. Schlatter, D. S. Henningson, Simulations of turbulent flow in a plane asymmetric diffuser, *Flow, Turbulence and Combustion* 79 (2007) 275–306.
- R. Dean, Reynolds number dependence of skin friction and other bulk flow variables in two-dimensional rectangular duct flow, *ASME. J. Fluids Eng.* 100 (1978) 215–223.
- H. G. Weller, G. Tabor, H. Jasak, C. Fureby, A tensorial approach to computational continuum mechanics using object-oriented techniques, *Computers in physics* 12 (1998) 620–631.
- S. Patankar, D. Spalding, A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows, *International Journal of Heat and Mass Transfer* 15 (1972) 1787–1806.
- S. Daniels, A. Rahat, R. Everson, G. Tabor, J. Fieldsend, A suite of computationally ex-

- pensive shape optimisation problems using computational fluid dynamics, in: International Conference on Parallel Problem Solving from Nature, 296–307.
- R. Blevins, Applied Fluid Dynamics Handbook, volume 2, Krieger Publishing Company, 1992.
- ESDU:74015, Performance in incompressible flow of plane-walled diffusers with single-plane expansion, Fluid Mechanics, Internal Flow, ESDU International plc, 2007.
- E. Reid, Performance Characteristics of Plane-wall Two-dimensional Diffusers, National Advisory Committee for Aeronautics, 1953.
- B. Eck, Technische Strömungslehre, 9th Edition, Springer Verlag: Berlin, 1988.
- C. O. A. Guhan, G. Arthanareeswaren, K. Varadarajan, Cfd study on pressure drop and uniformity index of three cylinder lcv exhaust system, Procedia Engineering 127 (2015) 1211–1218. International Conference on Computational Heat and Mass Transfer (CCHMT) - 2015.
- S. Daniels, A. Rahat, G. Tabor, J. Fieldsend, R. Everson, A Review of Shape Distortion Methods Available in the OpenFOAM Framework for Automated Design Optimisation, Springer, Cham, pp. 389–399.
- L. Jiang, G. Tabor, G. Gao, A new turbulence model for separated flows, International Journal of Computational Fluid Dynamics 25 (2011) 427–438.
- C. A. Coello Coello, G. B. Lamont, D. A. V. Veldhuizen, Evolutionary Algorithms for Solving Multi-Objective Problems, Springer, 2007.
- J. Stam, Exact evaluation of catmull-clark subdivision surfaces at arbitrary parameter values, in: Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '98, ACM, New York, NY, USA, 1998, pp. 395–404.
- A. A. M. Rahat, R. M. Everson, J. E. Fieldsend, Alternative infill strategies for expensive multi-objective optimisation, in: Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '17, ACM, New York, NY, USA, 2017, pp. 873–880.
- D. R. Jones, M. Schonlau, W. J. Welch, Efficient global optimization of expensive black-box functions, Journal of Global Optimization 13 (1998) 455–492.
- ESDU:72011, Flow through a sudden enlargement of area in a duct, Fluid Mechanics, Internal Flow, ESDU International plc, 2007.
- ESDU:76027, Introduction to design and performance data for diffusers, Fluid Mechanics, Internal Flow, ESDU International plc, 2007.