MSc thesis in Geomatics

Impact of Building Model Voxel Resolution in Wind Simulations using RANS

Bart Manden

2025

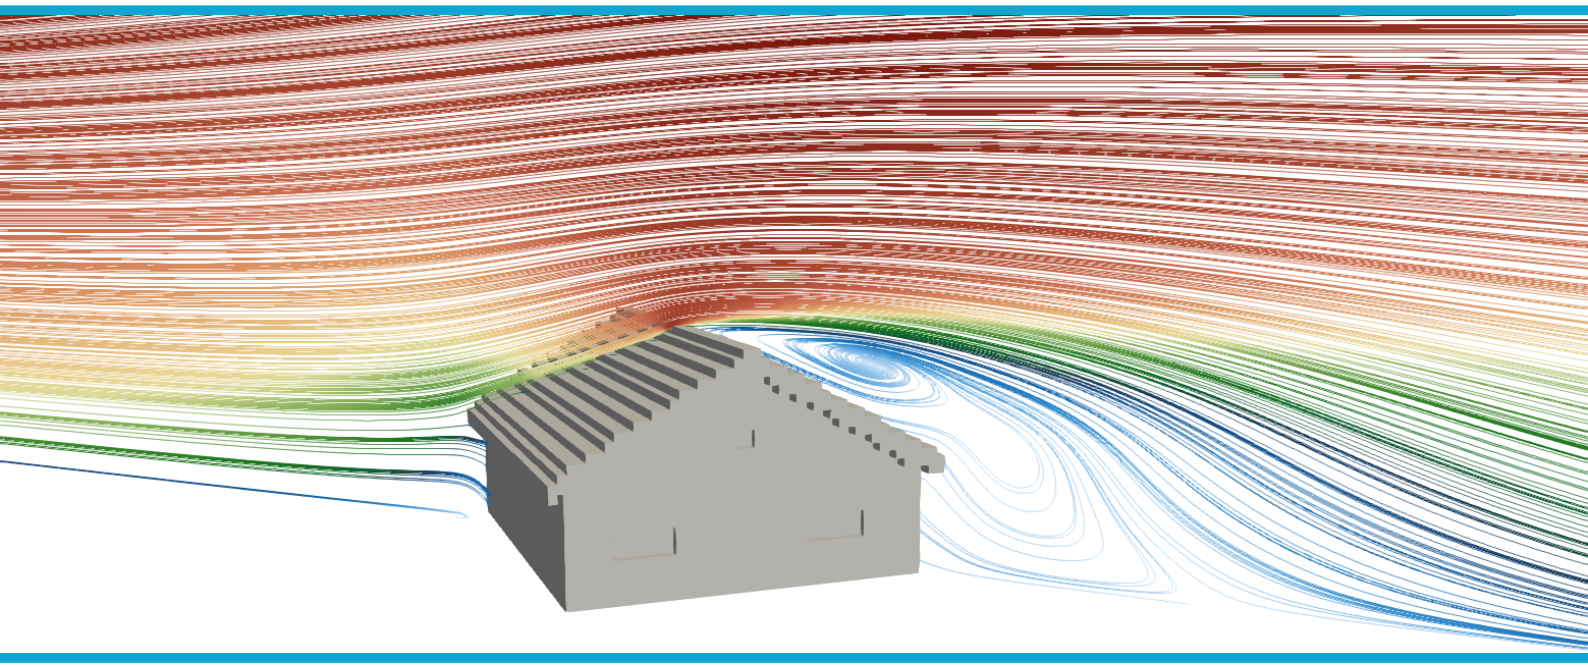

TU Delft

MSc thesis in Geomatics

Impact of Building Model Voxel Resolution in Wind Simulations using RANS

Bart Manden

July 2025

A thesis submitted to the Delft University of Technology in partial fulfillment of the requirements for the degree of Master of Science in Geomatics

The work in this thesis was carried out in the:

3D geoinformation group
Delft University of Technology

Supervisors:	Clara García-Sánchez
	Akshay Patil
	Jasper van der Vaart
Co-reader:	Nadine Hobeika

Abstract

Computational Fluid Dynamics (CFD) is widely used to analyse wind flow around buildings; however, creating detailed input geometries and corresponding meshes can be a time-consuming process. This thesis investigates voxelization as a means to simplify building models for their use in CFD and analyses the impact of voxel resolution on simulation accuracy.

Three building geometries with varying roof shapes and footprints were converted from detailed continuous models into voxel models with increasingly finer voxel resolutions. The voxelized models were compared to a non-voxelized LoD 3.2 model to assess accuracy under four key wind directions (90°, 45°, 22.5°, and 0°).

The CFD simulations were performed using OpenFOAM's RANS solver with a k–ϵ turbulence model. Due to its higher computational efficiency compared to other turbulence-resolving frameworks, the RANS approach enabled a large number of simulations while maintaining sufficient accuracy for urban CFD applications. A grid-independence test was conducted using the Grid Convergence Index (GCI) method for one model. The resulting grid-independent mesh was then scaled for the other models, ensuring that all simulations remained grid-independent.

The results show that coarse voxel resolutions (1 m and 0.5 m) significantly increase the size of the building geometry and leads to large velocity differences compared to the non-voxelized model. Sloped roofs were most affected by voxelization, as these models showed greater velocity differences than those with rounded roofs.

Wind direction also plays a significant role in voxelization accuracy. While the 90°, 22.5°, and 0° wind directions showed similar results across voxel resolutions, the 45° direction produced notable velocity differences. An exception was observed for the model with a rounded roof, which showed more consistent results across all wind directions.

Overall, the velocity difference between non-voxelized and voxelized models decreases as voxel size decreases. However, below a voxel size of 0.1 m, the reduction in velocity difference stagnates, indicating that smaller voxel sizes offer limited additional benefit to CFD accuracy.

Acknowledgements

I would like to thank my supervisors, Clara García Sánchez, Akshay Patil, and Jasper van der Vaart, for their guidance and support throughout this thesis. I am especially grateful for the timely help and information you provided whenever I encountered difficulties. I also wish to thank Nadine Hobeika for co-reading this thesis and providing valuable feedback.

I acknowledge the use of computational resources provided by the DelftBlue supercomputer through the Delft High Performance Computing Centre [Delft High Performance Computing Centre , DHPC]. I also acknowledge the use of AI-based tools for assistance with spelling, grammar, and for drafting code used in the plots; all conclusions and interpretations are solely my own.

Contents

List of Figures

List of Tables

Acronyms

1. Introduction

This chapter begins by outlining the motivation for the study, highlighting the broader context and research gap. It then presents the research questions that guide the work, and concludes with a concise overview of the overall thesis structure.

1.1. Motivation

As a large fraction of the world's population currently lives in cities, and with more people continuing to migrate towards urban areas [Mitkov et al., 2024], cities become increasingly vulnerable due to their high population density. Furthermore, cities have a significant impact on local meteorology; for example, by slowing down breezes through increased surface roughness, or by enhancing sea breezes due to thermal effects [Masson et al., 2020]. Numerical modelling approaches such as computational fluid dynamics (CFD) can play an important role in analysing urban meteorology and in developing adaptation strategies for urban climates [Masson et al., 2020]. However, applying CFD to entire cities introduces major data and computational challenges.

To create detailed CFD simulations for complete cities, a substantial amount of input data is required. The preprocessing of this data into a suitable format for use in CFD simulations demands considerable time and effort [Pađen et al., 2024]. As explained by van der Vaart et al. [2024a], errors regularly occur in city data that would render the model unusable for CFD applications. A few examples of such errors are: missing or intersecting faces, non-watertight solids, and non-planar surfaces. Even when CFD compatible models are available, more preprocessing is needed before a CFD simulation can be started. A mesh must then be created around the model for CFD use; this step often requires the most human resources. During the meshing, difficulties can arise due to small details in the model that is to be meshed [Kortelainen, 2009]. Consequently, many studies have explored geometry abstraction methods to reduce this time intensive step.

This bottleneck mainly applies to body-fitted meshes, where the mesh sits on the surfaces from the model. Although this approach results in accurate meshes, the process is slow and often requires adaptations to the geometry of the model [Jindal et al., 2007]. A different approach to the body-fitted mesh is the immersed boundary method. In contrast to the body-fitted mesh, the immersed boundary method does not require a model inside the mesh but instead uses body forces in the cells to represent the model [Vanky et al., 2024]. This allows the complete mesh to be a structured grid, which is easily generated. The geometry representation of the model however is altered to follow the shape of the grid, effectively creating a voxel based abstracted geometry from the model. A draw back of using this method compared to the body-fitted method are high local velocity differences which reach up to 50% of the reference velocity [Vanky et al., 2024].

Another approach to reduce the difficulties of small details in body-fitted meshes is to abstract the geometry before creating the CFD mesh. One such method is to use the LoD framework, which represents the same building geometry at increasing levels of detail [van der Vaart et al., 2024b; Heazel, 2021]. Common LoD's that can be found in 3D city models are LoD 1 (Extruded footprint to the building's height) and LoD 2 (Extruded footprint but with roof shape) [Sanchez et al., 2021]. As LoD 2 removes roof details such as overhangs, it reduces some of the complex geometries that can create difficulties for body-fitted meshes. The building geometry, however, still leads to partially cut cells within the mesh, increasing meshing complexity and preprocessing time. Using LoD 1 would resolve this as the shape of the geometry would always be prismatic. Research by Sanchez et al. [2021] and Patil and García-Sánchez [2025] shows that CFD simulation results can vary noticeably when using LoD 1 versus LoD 2. While Patil and García-Sánchez [2025] reported an average directional wind velocity difference of less than 20%, while Sanchez et al. [2021] found local velocity differences of up to 50%, with LoD 2 producing more realistic outcomes. While the LoD framework provides standardized abstraction levels, few studies explore simplifications beyond this concept.

Current strategies for reducing the meshing times in CFD studies are either using simplified buildings such as LoD 1 or the immersed boundary method which both have high local differences compared to more detailed building models such as LoD 2. A geometric abstraction method that retains more building details than LoD 1 or 2 while fitting efficiently to a mesh, as in immersed boundary methods, is currently missing. This motivates exploring new abstraction strategies that can reduce preprocessing time without compromising flow accuracy.

Such an abstraction strategy could be using voxelized models in a body-fitted mesh. With voxelization, 3D geometries are represented in a uniform grid using cubic cells called voxels. Depending on voxel size, geometric details of the original building model can be retained or discarded. As the voxel grid can align directly with a structured mesh, voxelized models can potentially eliminate much of the manual meshing process. This provides the opportunity to create models that retain geometric details comparable to LoD 2 or LoD 3 while also fitting to the mesh similar to the immersed boundary method. This makes voxelization a promising direction to bridge the gap between city data models and CFD simulations.

The absence of research in the CFD field concerning building abstractions using voxelization creates a promising opportunity for further investigation. This thesis aims to address that gap, by exploring the effects of voxel resolution on the CFD results for wind velocity. By comparing wind velocity differences between LoD 3 models and their voxelized versions, this thesis investigates to what extent voxelized models can reproduce LoD 3-based CFD results.

1.2. Research questions

To maximize the use of the automatic generation of voxelized models from Industy Foundation Classes (IFC) models in CFD simulations, it is essential that the outcome of the simulations deviates as little as possible from simulations obtained from higher-detailed models (see figure 1.1). The goal of this thesis is to analyse the behaviour of flows around voxelized models using CFD, and to compare these to flows around higher-detail, non-voxelized models. To achieve this, IFC models will be voxelized multiple times using different voxel

resolutions. From the IFC model, a LoD 3.2 model will also be created to serve as the comparative high-detail model. To quantify the impact of voxelization, the following research question will be addressed:

- *How does the voxel resolution within building voxelized models influence the accuracy and computational efficiency of wind flow simulations using Computational Fluid Dynamics (CFD)?*

To be able to answer this research question, it must first be defined which aspects are evaluated to quantify the influence of voxelization. For this thesis, the aspects will be limited to flow velocity and structure, as these are the main aspects of urban CFD research as explained in Section 1.1. To quantify how the simulations are impacted, two main aspects will be researched: the impact of geometric shapes and whether certain regions (such as wakes or corners) are impacted. This limited scope gives rise to four secondary research questions derived from the overarching main research question:

- Are some geometric shapes more sensitive to voxelization than others?
- To what extent does voxelization impact key CFD output parameters, such as velocity and flow structure?
- How sensitive are CFD outcomes to changes in voxel resolution across different regions of a building model?
- What is the minimum voxel resolution required to achieve an optimal balance between accuracy and computational efficiency?

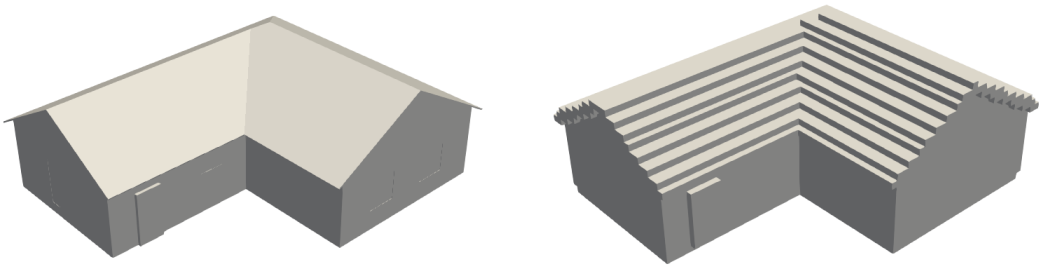

Figure 1.1.: Comparison between a higher detail model (LoD 3.2) left, and a voxelized model right.

1.3. Thesis structure

The thesis is divided into four chapters, which cover the following:

Chapter 2 presents related work. It is divided into two main sections: the first section discusses different urban geometry abstraction methods, while the second section discusses the CFD pipeline, including solver types, mesh creation, and validation.

Chapter 3 covers the methodology. It discusses the model voxelization and the domain setup for all models. This is followed by a description of the mesh validation procedure to be used, as well as the setup of the CFD simulations. Lastly, it explains how the data from the simulations will be analysed.

Chapter 4 presents the results and their analysis. It consists of four sections: the first section reports the results of the mesh validation, the second section presents the simulation results for each building, the third section discusses computational performance, and the final section summarises the main findings from the CFD simulations.

Chapter 5 provides the conclusion, in which an attempt will be made to answer the research question. It will also include recommendations for future work.

2. Related Work

This chapter reviews scientific work relevant to the thesis, focusing on two main areas: the geometric abstraction of buildings and the application of CFD software.

2.1. Geometry Abstraction

The first area, the abstraction of building geometry, encompasses techniques that reduce the complexity of urban models while preserving their essential characteristics. In this section, two key approaches are addressed: LoD, which represents the depth of detail in buildings, and voxelization, which converts continuous geometry into discrete grid-based representations.

2.1.1. Level of Detail in 3D Models

The Level of Detail (LoD) is a central concept in 3D city modelling, enabling the differentiation of multi-scale representations of 3D building models [Biljecki et al., 2016]. The LoD of a building model is influenced by both the data acquisition method (e.g., aerial photogrammetry or terrestrial laser scanning) and the modelling approach (e.g., manual versus automated processes) [van der Vaart et al., 2024b].

CityGML 3.0 defines four standard LoDs, ranging from LoD0 (footprint or roofprint) to LoD3 (3D mesh). However, Biljecki et al. [2016] proposed a more nuanced framework, introducing 16 LoDs by subdividing four of the original CityGML levels. This refined hierarchy provides a more continuous scale of geometric complexity and supports more precise use cases. An overview of these 16 LoDs is shown in Figure 2.1.

The different LoDs can be applied in various contexts depending on the requirements. For example, 3D city models used for CFD simulations are usually in LoD1 or LoD2 [Sanchez et al., 2021], while LoD3 models, with more details such as windows and doors, are used for estimating heat losses [Geiger et al., 2018].

Simplified 3D city models often lack critical geometric attributes, such as volume and surface area, which is lost during the generations of the models. This happens because every building model is somewhat simplified when compared to the real-world counterpart, and the amount of simplification is often unknown [van der Vaart et al., 2024a]. Since these simplified buildings (e.g., LoD1) do not preserve enough detail to recover the real-world structure, any derived analysis may deviate significantly from reality. This highlights a limitation of highly simplified models when used as data sources in analytics such as CFD applications, where geometric accuracy can strongly influence simulation results.

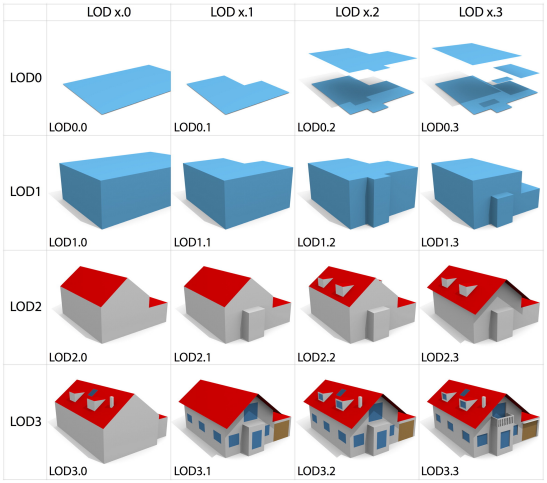

Figure 2.1.: Example of the 16 LoDs for a residential building [Biljecki et al., 2016]

2.1.2. Voxelization of 3D Models

An alternative method for abstracting 3D building models is voxelization, where continuous 3D models are converted into a structured grid of uniform cubic volumes (voxels) [Ridzuan et al., 2022]. Unlike low level LoD generalizations, where simplification often obscures the original building geometry beyond recovery, voxelization retains a controllable level of geometric accuracy depending on the chosen voxel size [Ridzuan et al., 2022].

Voxelization has been widely used in various domains such as medical imaging and computer gaming. In fields more closely related to this thesis —namely, 3D building modelling and geosciences— voxelization offers a way to represent urban environments in a discrete, regular format [Ridzuan et al., 2022]. One example is presented by Hinks et al. [2013], who propose an automated method for converting urban point cloud data into voxelized 3D urban geometries. Although this method has limitations, such as handling low-density point clouds or occlusions [Hinks et al., 2013], it shows potential for transforming city-scale point cloud data into usable urban models that could subsequently be applied in the CFD field.

2.2. Geometry abstraction for CFD

The trade-off between geometric abstraction and CFD accuracy is an important topic in CFD modelling. While geometry simplification can significantly reduce preprocessing times, it also alters the flow characteristics such as flow separation and turbulence generation. Therefore, understanding how different abstraction methods —such as LoD-based simplifications or voxel representations— affect CFD outcomes is essential for selecting appropriate modelling strategies in city-scale CFD studies.

To create a CFD simulation, a mesh in or around the desired geometry needs to be created. During the meshing the continuos domain is divided into discrete cells. Two main methods for the meshing are the most used in CFD, which are so-called body-fitted meshes and Cartesian meshes [Yu et al., 2018]. With body-fitted meshes, the cells of the mesh are reshaped

to match the shape of the used geometry. For Cartesian meshes, particularly the immersed boundary method, the mesh is unaffected by the geometry and instead simulates the solid walls of the geometry by including force calculations near the geometry [Li et al., 2024]. A hybrid approach, known as the cut cell method, also exists. In this method, a grid similar to that used in the immersed boundary approach is employed, but instead of using only complete cells from the Cartesian grid, cells that intersect the geometry are cut to conform to the geometry [Karalit, 2023]. Geometry abstraction therefore is particularly useful for body-fitted meshes as the complexity of the model affects the complexity of the mesh.

The previously discussed LoD method is already being used in CFD. With Sanchez et al. [2021] using LoD 1.3 and LoD 2.2 (see figure 2.1) for a comparative study on the accuracy of said LoD's. While the LoD 1.3 would fit nicely to a Cartesian grid as sloped roofs are simplified to flat surfaces, the study showed significant differences in velocity in the CFD simulations when compared to the higher detailed LoD 2.2 model. A similar study performed by Patil and García-Sánchez [2025] compared LoD 1.2 to LoD 2.2, where a lower average directional velocity difference was found (under 20%) between the two LoD's. Using the higher LoD's in turn often lead to higher meshing times as the meshes become increasingly complex to generate.

A different approach to model abstraction for CFD is to simplify the geometry to conform to a Cartesian grid, which is effectively voxelization of the geometry. Villi and Carli [2014] simplified the human body using this approach for the use in CFD for indoor flow. Three simplifications were used ranging from a rectangular box to distinct features including legs, arms and the head. The results from the CFD simulations were compared to experimental data, showing that the voxelized models over predicted the wind flow up to 55% of the inlet velocity.

2.3. Computational Fluid Dynamics (CFD)

2.3.1. Governing Equations

CFD simulations for an urban environment solve the fundamental conservation equations: mass (continuity) and momentum (Navier–Stokes), often under the incompressible assumption for low-speed atmospheric flows [Toja-Silva et al., 2018]. For turbulent flows, which are typical for urban regions, additional terms emerge from Reynolds averaging (for RANS) or filtering (for LES), requiring closure via turbulence models [Toja-Silva et al., 2018].

2.3.2. Turbulence Modelling

There are a number of approaches to deal with turbulence predictions in urban CFD:

- **RANS (Reynolds-Averaged Navier–Stokes):** Modelling all the effects from turbulence, it is the most common approach in urban wind flow studies due to its relatively low computational cost. Models such as standard $k-\epsilon$, realizable $k-\epsilon$, and $k-\omega$ SST are frequently used [Toja-Silva et al., 2018; Xiong et al., 2022a].

- **LES (Large Eddy Simulation):** Since it resolves larger scale turbulences, it captures more detail in unsteady and transient flow features, especially in wakes and separated flow regions. But LES is more computationally expensive and often requires finer meshes, high temporal resolution, and careful inflow turbulence specification [Toja-Silva et al., 2018].
- **Hybrid / Blending methods:** Sometimes used to combine the advantages of RANS near walls with LES in separated flow zones (e.g. Detached-Eddy Simulation, or RANS-LES blending) in literature such as on tall buildings [Wijesooriya et al., 2023].

Comparative studies show that for many urban wind flow applications, RANS with a well-chosen turbulence model gives reasonable results for mean velocity and pressure fields, though not always for turbulence fluctuations or scalar transport (e.g. pollutant dispersion) [Xiong et al., 2022b; Wang and McNamara, 2006].

2.3.3. Grid Convergence Index

Mesh validation is essential to ensure that CFD results are independent of the numerical resolution. A standard approach is to perform a mesh independence or grid convergence study, where key solution variables (e.g., mean velocity, pressure coefficients) are evaluated on successively refined meshes.

A widely recommended procedure is the grid convergence index (GCI), introduced by Roache (1994), which provides a quantitative estimate of the discretisation uncertainty based on Richardson extrapolation [Roache, 1994, 1997]. The GCI expresses the relative error between solutions on different grids and helps assess whether the mesh is sufficiently fine for the quantities of interest.

Studies in the building aerodynamics literature have applied GCI or similar refinement strategies to verify that predictions such as pedestrian wind speed or façade pressure are not biased by grid resolution [Xiong et al., 2022b; Toja-Silva et al., 2018]. Proper use of the GCI requires at least three systematically refined meshes and calculation of the observed order of accuracy. When the GCI falls below an acceptable threshold (often a few percent), the solution can be considered grid-independent.

3. Methodology

The aim of this chapter is to analyse the methodology used in this thesis to address the research questions as stated in Section 1.2. The methodology is divided into four separate sections: The first section introduces the models that were used in the CFD simulations, and how these models were voxelized. The second section introduces the domain that was used for the CFD simulations, where minimum size and refinement regions are discussed. The third section discusses the minimum mesh refinement required for the domain, where the aim was to create a mesh independent setup. The fourth and final section explains the setup that was used for all the simulations that were performed to address the research questions from Section 1.2.

3.1. Model voxelization

In this thesis three building models were used – each with different geometric characteristics – to assess the effects of voxelization on CFD simulation results. Figure 3.1 shows the IFC models of the three buildings. All three models had either a sloped or rounded roof. These models were chosen because voxelization most strongly affects non-rectangular geometries. Buildings with flat roofs undergo minimal geometric change when voxelized and therefore provide limited insight.

(a) A-40 building (b) FZK-Haus building (c) Institute building

Figure 3.1.: IFC models of the three buildings used during the research.

The first building was the A-40 building, which is a one story building with an asymmetric topology. This building was chosen specifically due to its asymmetry, especially the asymmetry of the sloped roof. This way the influence – if any – of asymmetry could be quantified when comparing it to one of the symmetric models.

The second building was the FZK-Haus building [Institute for Automation and Applied Informatics (IAI) and Karlsruhe Institute of Technology (KIT), a], which is again a one story building. This building is very similar to the A-40 building, with the key difference that the FZK-Haus building is symmetric.

The final building that was used for the CFD simulations was the Institute building [Institute for Automation and Applied Informatics (IAI) and Karlsruhe Institute of Technology (KIT),

b], a three story symmetric building. This building was significantly larger than the other two buildings, which was part of the reason that this building was chosen. The main reason that this building was used however, was because of its rounded roof. It was interesting to see if there were significant differences between the outcomes of the sloped roof buildings and the rounded roof building.

The voxelized models used in this thesis were generated using the IFCEnvelopeExtractor [Vaart, 2025] software. This software was originally designed to produce LoD building representations from IFC data rather than pure voxel grids. The software follows a hybrid approach that combines geometry filtering, simplification, and voxel-based abstraction. In the first step, the program selected the relevant IFC entity types (e.g., walls, roofs, doors, etc.) to include in the model. Here doors and windows were simplified to reduce geometric complexity while preserving their overall shape. Subsequently, a voxelization step was performed in which building components were discretized according to a user-defined voxel size. Depending on the chosen level of detail, this voxelization could range from a full volumetric representation (LoD 3) to a column-based voxelization where vertical elements were extruded through the building height (LoD 1/2).

After this process, the software constructed a signed volume field by converting each building element into a triangulated mesh and testing for intersections with the voxel grid. Because of this intersection-based approach, voxel boundaries could slightly overestimate object volumes—typically by one voxel—reflecting the transition from linear to volumetric precision. Finally, to identify enclosed building volumes, the software applied a region-growing algorithm that merged neighboring voxels belonging to the same object. This step could be computationally expensive, especially for complex or finely detailed geometries.

Six voxel sizes were chosen for the abstraction of the IFC models, which were: 1-meter, 0.5-meter, 0.25-meter, 0.1-meter, 0.05-meter and 0.025-meter. Next to the created voxel models, for each building a LoD 3.2 (see Figure 2.1) model was generated. This LoD 3.2 model was used as the reference 'truth' model to which the voxelized models were compared. This LoD3.2 model was still an abstraction of the actual building as designed in an IFC model. However, the LoD 3.2 – as it falls within the highest LoD family – contains a lot of detail with the LoD3.2 only leaving out features that are smaller than 1 meter [Biljecki et al., 2016]. This way smaller voxel sizes like the 0.025-meter and 0.01-meter voxels could capture smaller details like overhangs and windows which would not be present when choosing an LoD 2 model. In Figure 3.2 the results of the voxelization using different sized voxels are shown, as well as the LoD 3.2 truth model.

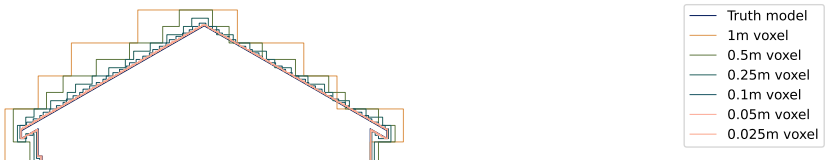

Figure 3.2.: Effect of voxelization using differently sized voxels, on the sloped roof of the FZK-Haus building.

This process was automated by the IFCEnvelopeExtractor [Vaart, 2025] software, where the only required inputs were an IFC model and the desired voxel size or LoD. It should be noted however that for the A-40 building some manual preprocessing was required. This

was due to the IFC model, in which the walls and the roof were not connected leading to a non-watertight geometry.

3.2. CFD Domain design

The next step in the pipeline is the setup of the CFD domain. To do this, the guidelines as described by Blocken [2015] have been used, to ensure a valid domain size was used. Blocken [2015] describes two factors that influence the size of the domain:

1. Distance between the building model and the domain boundaries
2. Blockage ratio

Both of these factors specify how large the domain should be compared to the building, to limit artificial acceleration due to boundary edges being to close to the building. This artificial acceleration occurs when one of the domain boundaries that is not a physical boundary – like the top and side boundaries – is so close to the building that the flow is strongly contracted by these boundaries [Blocken, 2015]. For the first factor – the distance between the building and the boundaries – the minimum distance is based on the maximum height of the building. Figure 3.3 visualizes these distances, where the inlet, top and side boundaries of the domain should be at least 5 times the maximum height from the building away from the building. For the outlet, a mimimum distance of 15 times the height of the building is specified, to allow the wake to fully develop before exiting the domain [Blocken, 2015].

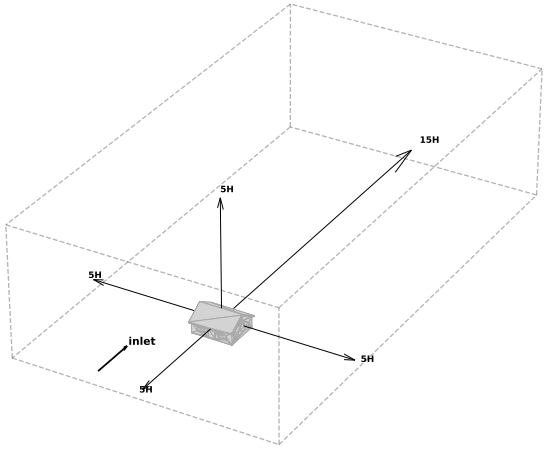

Figure 3.3.: Minimum domain distance away from the model, as specified by Blocken [2015]

As the width of a building is not taken into account with the previous method, it could happen that a very wide building is still close to the boundary while it complies with the 5 $H_{\max}$ rule. To prevent this, the second factor – the blockage ratio – uses the area of the building and compares that to the area of the inlet. Two different blockage ratio's have been defined by Blocken [2015], a less strict ratio where the blockage should be less than 3%:

$$\text{Blockage ratio} = \frac{A_{\text{building}}}{A_{\text{domain}}} < 3\% \tag{3.1}$$

And a more strict blockage ratio where the length and height of the building is compared to the length and height of the domain. Here the blockage ratio is defined as:

$$\begin{aligned} \mathrm{BR}_L &= \frac{L_{\text{building}}}{L_{\text{domain}}} < 17\% \\ \mathrm{BR}_H &= \frac{H_{\text{building}}}{H_{\text{domain}}} < 17\% \end{aligned} \tag{3.2}$$

When complying with the first rule, where the height of the domain is 6 times the max height of the building, automatically the second rule of this stricter blockage ratio is also complied with. And when complying with both of these stricter blockage ratios, the first blockage ratio is complied with as: $BR_H * BR_L = BR_A$ and $0.17 * 0.17 = 0.0289 \approx 2.9\% < 3\%$.

Combining this into one rule set that must be followed to comply with the domain size as described by Blocken [2015], leads to the following:

1. The inlet and top domain boundary should be minimum $5H_{\text{max}}$ away from the building
2. The outlet should be a minimum of $15H_{\text{max}}$ away from the building
3. For the lateral boundaries $\frac{L_{\text{building}}}{L_{\text{domain}}} < 17\%$

These setup applies to a rectangular domain where only one wind direction is specified. As in this thesis four wind directions have been used (0°, 22.5°, 45° and 90°). The described domain would have to be made four times for each wind direction and seven times for each voxelized model, resulting in 28 domains that would have to be made. To decrease the amount of domains that would have to be made, a circular domain can be constructed for each voxelized model reducing the total amount of domains to 21.

As the three previous defined rules still apply to this circular domain, either the second rule or the third rule defines the radius of this circular domain. As none of the buildings that have been used for this Thesis are exceptionally wide, the second rule is defining for the radius of the domain. Where the radius of the domains is $0.5L + 15H_{\text{max}}$ where L is either the width or length of the building, depending on which one is larger. Because of this, the inlet is also $15H_{\text{max}}$ away from the building. This means that the domain is over defined, but as it reduces the total amount of domains that have to be made it is acceptable. An example of the resulting domain can be seen in Figure 3.4 where the domain of the FZK-Haus building is shown.

3.3. Mesh validation

The created domain, as explained in the previous section, needs to be meshed in the next step. The meshing allows to model the airflow around the model, by solving the equations in each volumetric cell. To create the mesh the fully parallel, split hex mesh generator *snappyHexMesh* from OpenFOAM is used [OpenFOAM, 2019]

The meshing involves two main steps: First the determination of the level of meshing needed around the surfaces i.e.how fine does the mesh get when going closer to the surface. Secondly, the created mesh needs to be validated to quantify the level of uncertainty. This is

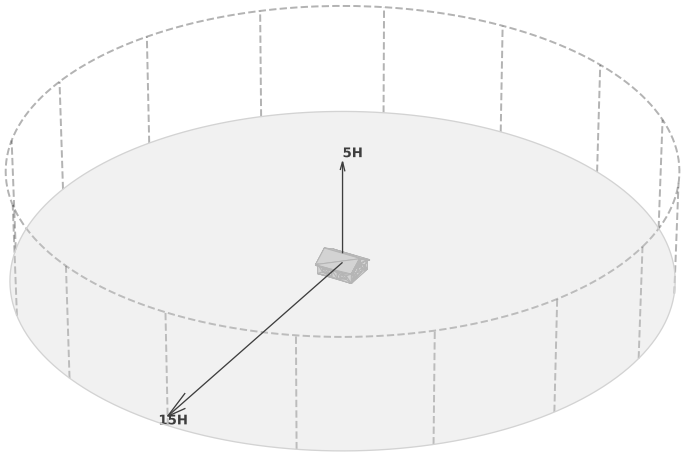

Figure 3.4.: Example of the circular domain from the FZK-Haus, as has been used during this Thesis

done to make sure that the simulations are grid independent i.e.the outcome of the simulations is not influenced by the level of grid refinement [Roache, 1994].

The first step – How fine does the mesh need to be at near-surface locations – is defined in the guidelines from Blocken [2015]. Firstly for building models, at least 10 cells per cube root of the building volume should be used. Additionally to the minimum of 10 cells, the area of interest should be located in the 3rd or more cell. This all is done so that the characteristics of the wind flow on roofs and walls is captured fully Blocken [2015].

The second step would be to evaluate the performance from the created mesh by doing a grid independency analysis. For a grid independence study, a widely accepted and recommended method is to use Richardson extrapolation for discretization error estimation [Celik et al., 2008]. A method that utilizes Richardson extrapolation is that of Roache [1994], where a GCI is introduced to provide a way of uniform reporting on grid convergence studies. As this method is, according to Celik et al. [2008], the recommended way of doing a grid convergence study due to its evaluation over a large number of cases. This method has been used in this Thesis to perform the grid independency test.

The following method for calculating the GCI is defined by Celik et al. [2008]. The GCI method uses (a minium of) three separate meshes from the same case, where the only difference between the meshes is its level of mesh refinement. All three meshes should be systematically refined where the refinement between two subsequent meshes is greater than 1.3, calculated as: $r = h_{\text{coarse}}/h_{\text{fine}}$. Here h is the representative cell size of the mesh, which in the case of this Thesis can be calculated from the 3 dimensional grid as followed:

$$h = \left[\frac{1}{N}\sum_{i=1}^{N}(\Delta V_i)\right]^{1/3} \tag{3.3}$$

With ΔV_i being the total volume of the domain, and N being the total number of cells for the domain.

As mentioned by Celik et al. [2008], the previously explained refinement should be done equally throughout the domain regardless of the mesh structure from the domain. This means that –as explained previously – the highly refined areas near the surfaces will be refined as much (minimum 1.3 times) as the coarser areas further away from the surfaces. This refinement of the areas near the surfaces can increase the amount total amount of cells rapidly, making it key to not start with a to fine mesh to prevent extreme computation costs.

Continuing with the method from Celik et al. [2008], the apparent order of convergence p can be calculated using the the refinement ratios of the three meshes. Using $r_{21} = h_2/h_1$ and $r_{32} = h_3/h_2$ with h_1, h_2 and h_3 being the fine, medium and coarse meshes respectively. Then p can be calculated using the iterative equation 3.4

$$\begin{aligned} p &= \frac{1}{ln(r_{21})} \left| ln \left| \epsilon_{32}/\epsilon_{21} \right| + q(p) \right| \\ q(p) &= ln(\frac{r_{21}^p - s}{r_{32}^p - s}) \\ s &= 1 * sgn(\epsilon_{32}/\epsilon_{21}) \end{aligned} \tag{3.4}$$

With $\epsilon_{32} = \phi_3 - \phi_2$ and $\epsilon_{21} = \phi_2 - \phi_1$, ϕ_k is the solution on a sampled location in the domain for the k^{th} mesh.

As the last step from Celik et al. [2008], the final error estimates can be calculated using the following equations three equations:

1.
$$e_a^{21} = \left| \frac{\phi_1 - \phi_2}{\phi_1} \right| \tag{3.5}$$

2.
$$e_{ext}^{21} = \left| \frac{\phi_{ext}^{21} - \phi_1}{\phi_{ext}^{21}} \right| \tag{3.6}$$

with:

$$\phi_{ext}^{21} = (r_{21}^p \phi_1 - \phi_2)/(r_{21}^p - 1) \tag{3.7}$$

3.
$$GCI_{fine} = \frac{1.25 e_a^{21}}{r_{21}^p - 1} \tag{3.8}$$

The same equations can then be used to calculate the GCI for the medium resolution grid. Using ϕ_{ext}^{32} instead of ϕ_{ext}^{21} which can be calculated in the same way using ϕ_2 and ϕ_3 instead of ϕ_1 and ϕ_2.

The resulting error estimation can be used to determine the level of grid refinement to be used. As the GCI method is mesh specific, it should be calculated for each domain separately. In the case of this Thesis that would result in 63 separate meshes that would need to be

created. All of these meshes would then have to be simulated and sampled to be able to calculate the error estimates. As this would take a lot of time and computational power, a different approach was used for this Thesis. The GCI method was used on the domain of one model and the resulting mesh would then be scaled to the domain of the other models. When scaling, it is ensured that the cell sizes for each part of the domain are of the same size across different domains. This way, a lot of time is saved while still maintaining the grid independence across all the models.

While this does ensure comparable cell sizes, local wind flow could still be impacted by the mesh when more complex building geometries are used. As this thesis uses relatively simple building geometries with comparable complexity across the models, this was considered acceptable.

3.4. Simulations

To run the simulations that have been mentioned before, the open-source set of C^{++} libraries, OpenFOAM [OpenFOAM, 2019], has been used. The solver that was used to run the CFD simulations is the Reynolds-averaged Navier–Stokes equations (RANS) solver simpleFoam, which uses the following governing equations:

$$\begin{aligned} \frac{\partial U_j}{\partial x_j} &= 0, \\ U_j \frac{\partial U_i}{\partial x_j} &= -\frac{1}{\rho}\frac{\partial p}{\partial x_i} + \frac{\partial}{\partial x_j}\left[(\nu + \nu_t)\left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i}\right)\right] \end{aligned} \tag{3.9}$$

Here, ρ is the constant fluid density, ν the molecular kinematic viscosity, U_i the Reynolds-averaged velocity component, and p the mean static pressure. The spatial coordinates are denoted by x_j, with $j = 1, 2, 3$ corresponding to the x, y, and z directions, respectively, and $\frac{\partial}{\partial x_j}$ denotes the partial derivative with respect to x_j. The first equation represents the incompressible continuity condition, ensuring a divergence-free mean velocity field, while the second equation describes the momentum balance. The first term on the left-hand side of the second equation represents the convective transport of momentum, the first term on the right-hand side corresponds to the pressure-gradient force, and the final term accounts for viscous and turbulent diffusion, where the turbulent viscosity ν_t is related to the mean rate of strain, $\left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i}\right)$.

The RANS solver is a reliable solver to choose when performing CFD in an urban environment, when the main objective of the study is to access pedestrian wind comfort [Janssen et al., 2013]. Although pedestrian wind comfort is not necessarily the objective of the research in this thesis, it is one of the fields where voxelization of the geometries can well be used. The main reason for choosing the RANS solver over other approaches like the Large Eddy Simulation (LES) however, is that the RANS approaches are faster [Janssen et al., 2013]. This is especially useful for cases where a large number of simulations has to be done, like in this thesis. A drawback of using RANS instead of LES is that LES methods prove to be more accurate than RANS methods, especially in the case of areas where unsteady features of the airflow field are more prominent [Janssen et al., 2013].

RANS approaches have multiple turbulence models that can be used, a few examples are: the $k-\epsilon$ model, the $k-\epsilon$ RNG model, the $k-\omega$ model and non-linear models [Setaih et al., 2014]. From these models the $k-\epsilon$ turbulence model is used the most for engineering calculations [Yang and Shih, 1993]. This turbulence model was therefore also used for the simulations in this thesis. The equations for k and ϵ are:

$$\begin{aligned} U_j \frac{\partial k}{\partial x_j} &= \frac{\partial}{\partial x_j}\left[\left(\nu + \frac{\nu_t}{\sigma_k}\right)\frac{\partial k}{\partial x_j}\right] + P_k - \varepsilon, \\ U_j \frac{\partial \varepsilon}{\partial x_j} &= \frac{\partial}{\partial x_j}\left[\left(\nu + \frac{\nu_t}{\sigma_\varepsilon}\right)\frac{\partial \varepsilon}{\partial x_j}\right] + C_{1\varepsilon}\frac{\varepsilon}{k}P_k - C_{2\varepsilon}\frac{\varepsilon^2}{k} \end{aligned} \tag{3.10}$$

with

$$\nu_t = C_\mu \frac{k^2}{\varepsilon}, \quad P_k = \nu_t \left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i}\right)\frac{\partial U_i}{\partial x_j} \tag{3.11}$$

In the k–ε model, k represents the turbulent kinetic energy and ε is its rate of dissipation. The turbulent viscosity, ν_t, is related to k and ε, while P_k represents the production of turbulent kinetic energy due to the mean velocity gradients. The constants used in these equations are listed in Table 3.1.

Symbol	Description	Value
C_μ	Empirical constant for turbulent viscosity	0.09
$C_{1\varepsilon}$	Production coefficient in ε-equation	1.44
$C_{2\varepsilon}$	Dissipation coefficient in ε-equation	1.92
σ_k	Turbulent Prandtl number for k	1.0
σ_ε	Turbulent Prandtl number for ε	1.3

Table 3.1.: Model constants for the k–ε turbulence model.

Table 3.2 summarizes the domain and boundary-layer setup used in the simulations. U_{ref}, Z_{ref}, Z_0, κ, and E were specified to define the inlet and wall boundary conditions. At solid boundaries, near-wall turbulence quantities were modelled using standard k-ε wall functions.

Symbol	Description	Value
U_{ref}	Reference velocity	5 m/s
Z_{ref}	Reference height	10 m
Z_0	Surface roughness length	0.03 m
κ	von Kármán constant	0.41
E	Empirical model constant	9.8

Table 3.2.: Wall function parameters

For this thesis, for each building model 28 simulations were made which results in a total of 84 simulations. To run multiple simulations at the same time, the DelftBlue supercomputer was used [Delft High Performance Computing Centre , DHPC]. The total CPU hours that have been utilized can bee seen in Table 3.3

Type	CPU Hours
Computing nodes	25,916
Visual nodes	316
Total	26,232

Table 3.3.: Used CPU hours on DelftBlue per node type

3.5. Data analysis

To be able to answer the research question as stated in section 1.2, the simulations will need to be processed further as the raw data does little to provide insight. For this, two main methods have been used: The first is using a point cloud to sample the data in a structured way, the second method involves constructing stream graphs to analyse the wake structures.

For the data sampling, a rectangular box has been created that is three times the height of the model in the x and y directions and extends to 1.5 time the height of the model in the z direction. This box is then sampled with a spacing of 0.2 meters for the x and y directions, to capture enough detail while not creating excessive computation times.

This point cloud has been used to create horizontal planes that show velocity magnitude differences between the voxel models and the truth model for each building. As these planes have been used to analyse the overall similarity between the voxel models and the truth models, a higher resolution point cloud provides little extra insight as detailed analysis can better be done trough other methods like stream graphs and velocity profiles. the z-spacing of the sampling points is larger with a spacing of 0.5 meters to further reduce the computation times.

To examine wake structures across models, two-dimensional streamlines were generated in ParaView [Ahrens et al., 2005] to visualize recirculation zones and flow separation. To be able to distinguish the different components of the wake structure, a 2D plane has been used that is oriented along the wind direction. This way it is easier to analyse the wake, as a 3D stream graph becomes highly cluttered.

4. Results and Analysis

This chapter presents the results obtained from the CFD simulations and evaluates the differences between the voxelized and real case for the three buildings. The objective is to assess the quality of the simulation outcomes, confirm mesh independence through a GCI analysis, and interpret simulation outcomes. This chapter provides a detailed presentation of the results obtained from the CFD simulations. The presented results aim to support the findings discussed later in section 4.4.

4.1. Grid Convergence Index (GCI)

To assess the mesh independence of the CFD simulations, as explained in section 3.3, the GCI has been computed based on the results obtained from three constantly refined grids. The GCI was computed for one representative configuration, in this case the FZK-Haus (see Figure 3.1). This test case was selected based on its typical geometric characteristics. The results from these calculations will be presented in this section and the grid independence will be evaluated.

For this analysis, key field variables—velocity, pressure, turbulent kinetic energy, and epsilon—were sampled at 200 randomly distributed points throughout the domain. For each point, the GCI was computed using values from the three grid levels, and the median GCI was taken to summarize convergence behavior for each variable. This approach reduces the influence of local anomalies or outliers in the field data.

Table 4.1.: Metrics of the three grids used for the GCI calculations, the mesh volume remains nearly constant ensuring mainly spatial refinement.

Grid Level	Total Cells (x10^6)	Mesh Volume (x$10^6 m^3$)	Av. Cell Length Δh_r (m)
Coarse	6.136	1.638	0.64
Medium	17.148	1.642	0.46
Fine	44.444	1.644	0.33

The three sets of grids used had a total of 6, 20 and 40 million cells for coarse, medium and fine meshes respectively. The refinement ratio between the grids is approximately $r \approx 1.4$, with average cell lengths of 0.64 m, 0.46 m, and 0.33 m for the coarse, medium, and fine meshes, respectively. Table 4.1 summarizes these mesh characteristics used in the GCI analysis. The cell length values were derived using the volumetric definition:

$$h = \left[\frac{1}{N} \sum_{i=1}^{N} (\Delta V_i) \right]^{1/3} \tag{4.1}$$

The calculated order of accuracy ranges from 2.80 to 6.63 with an average across all variables of approximately $p \approx 3.87$, consistent with expectations for a second order discretization scheme. As shown in Table 4.2, GCI values remained below 3% for most variables, indicating satisfactory grid convergence. Exceptions were observed for vertical velocity and pressure, which exhibited higher GCI values.

Table 4.2.: Grid Convergence Index (GCI) results for the medium grid refinement.

Metric	ε	k	p	U_x	U_y	U_z
Apparent order p	2.80	2.87	3.62	6.63	3.01	4.63
Relative Approx. Error (%)	3.26	0.24	8.13	12.21	0.44	15.51
Extrapolated Relative Error (%)	1.87	0.15	4.18	1.67	0.21	5.81
GCI (%)	2.29	0.19	5.12	2.08	0.27	7.34

Given the consistent domain setup across simulations with different buildings, the mesh resolution for these cases have been derived from the test case. This was done by scaling the mesh according to the new building height—as this defines the domain size—ensuring that the mesh cells differ as little as possible from the original cell size (see Figure 4.1 for a visual comparison between scaled grids). This approach was chosen as performing a full GCI analysis for every scenario would become computationally exhaustive, and the grid topology remains consistent.

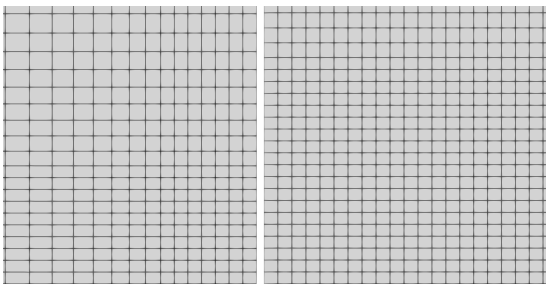

Figure 4.1.: Comparison between two scaled grids: the original FZK-Haus (left) and the taller Institute (right). The mesh resolution is adjusted based on building height while maintaining similar cell sizes. The domain is scaled when moving away from the building. As a result, refinement begins at different absolute heights—visible here as larger cells appearing earlier in the FZK-Haus mesh (left).

4.2. Flow field results

This section presents a detailed analysis of the CFD simulation results for the three building models. The analysis begins with an evaluation of the geometric differences between voxelized models. Subsequently, regional velocity variations are examined for each voxel resolution and wind direction. Finally, representative wind flow patterns are compared across multiple wind directions for selected resolutions.

4.2.1. Geometric and velocity differences

A-40

The first building that will be analysed is the A-40 building, which is a simple one story building with an asymmetric layout. A 3D representation of the building can be seen in figure 4.2a

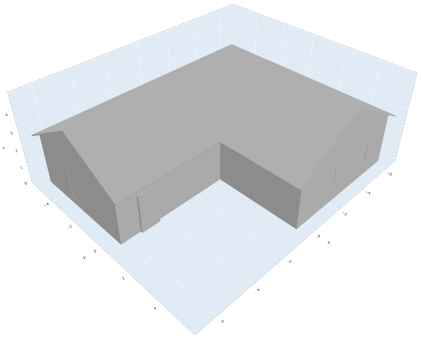

(a) 3D view of the A-40 building

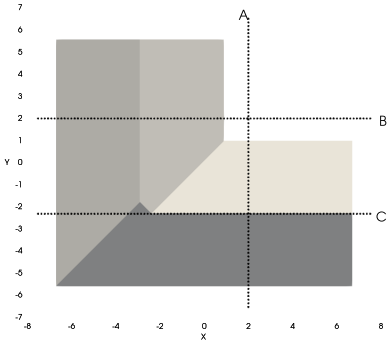

(b) Overview of slice locations

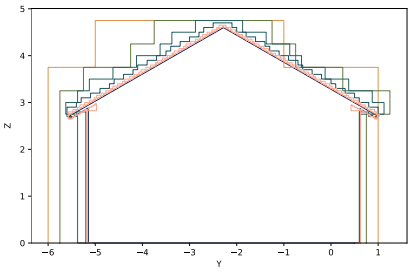

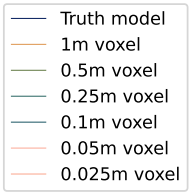

(c) Slice A

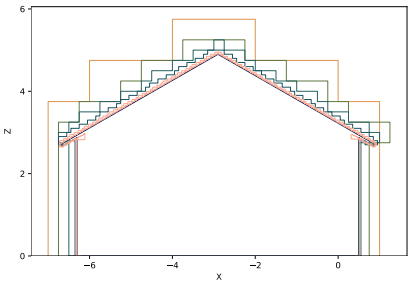

(d) Slice B

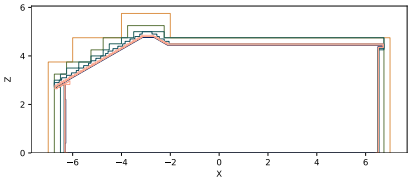

(e) Slice C

Figure 4.2.: Geometric difference between voxel resolutions, visualized in 2D slices.

As mentioned in chapter 3, the buildings were voxelized using six different levels of resolution. Specifically, voxel sizes of 1-meter, 0.5-meter, 0.25-meter, 0.1-meter, 0.05-meter, and 0.025-meter were used to represent the building geometry in varying degrees of detail. The

effect of voxel size on the building's shape is illustrated in Figure 4.2, where 2D slices of the models are compared to highlight the geometric differences introduced by each resolution.

As can be seen in Figure 4.2, the 1-meter voxel results in a significant increase in the building's dimensions—both horizontally and vertically—compared to the non-voxelized truth model. This discrepancy is also apparent for the 0.5-meter, 0.25-meter, and 0.1-meter voxel resolutions. When comparing the 0.25-meter voxel with the 0.1-meter voxel, it can be observed that although the 0.25-meter voxel introduces larger steps in the building's geometry, both models reach approximately the same overall height and have a similar final building width.

At even higher voxel resolutions, namely 0.05 meters and 0.025 meters, the voxelized representations closely follow the geometry of the truth model. These higher-resolution models also converge to nearly the same final height and width as the original, non-voxelized geometry.

Another notable observation is that each voxelized model tends to be one voxel length taller and one voxel length wider on at least one side of the building. The exceptions to this are the models voxelized with the 1-meter and 0.5-meter resolutions. However, even these coarser voxel models deviate from the truth model by approximately one voxel length in certain locations (see the roof in Figure 4.2d).

These geometric deviations are not only visual but also influence the simulated flow behaviour. To assess this impact, the CFD results are compared across voxel resolutions. Some of the simulation results, as described in Chapter 3, are shown in Figure 4.3. In this figure, the results of the simulations for the three coarsest meshes are compared to those of the truth model. A square region around the models, with a size equal to three times the height of the truth model, is sampled with a spacing of 0.2 meters between sample points. The resulting data point clouds for the 1-meter, 0.5-meter, and 0.25-meter voxel models are then compared to the data point cloud of the truth model. The velocity error—normalized by the inlet wind speed of 5 m/s—is plotted for wind directions of 0°, 22.5° 45°, and 90° relative to north. The data is taken at a height of 3.2 meters, which corresponds to the point where the roof begins. This height was chosen as the roof is the part of the building that is most altered and subsequently will create the biggest velocity differences.

As shown in Figure 4.3, the largest differences occur for the 45° wind direction. Interestingly, in this case the differences in the wake appears to increase as the voxel size decreases. While the 1-meter voxel model exhibits negative differences in the wake, both the 0.5-meter and 0.25-meter voxel models show positive differences. This is in sharp contrast to the other wind directions, where the differences decreases with decreasing voxel size. The increasing differences for the smaller voxel sizes in the 45° wind direction can be explained by the difference in roof height between the voxelized models. As seen most clearly in Figure 4.2d, the height of the left-hand roof in the 1-meter voxel model is significantly greater than in the 0.5-meter or 0.25-meter voxel models, with all voxelized roofs being higher than the truth model. The increased height of the voxelized models compared to the truth model ranges from approximately 0.85 m for the 1-meter voxel model to about 0.35 m for both the 0.5-meter and 0.25-meter voxel models at the highest part of the building (Figure 4.2b)

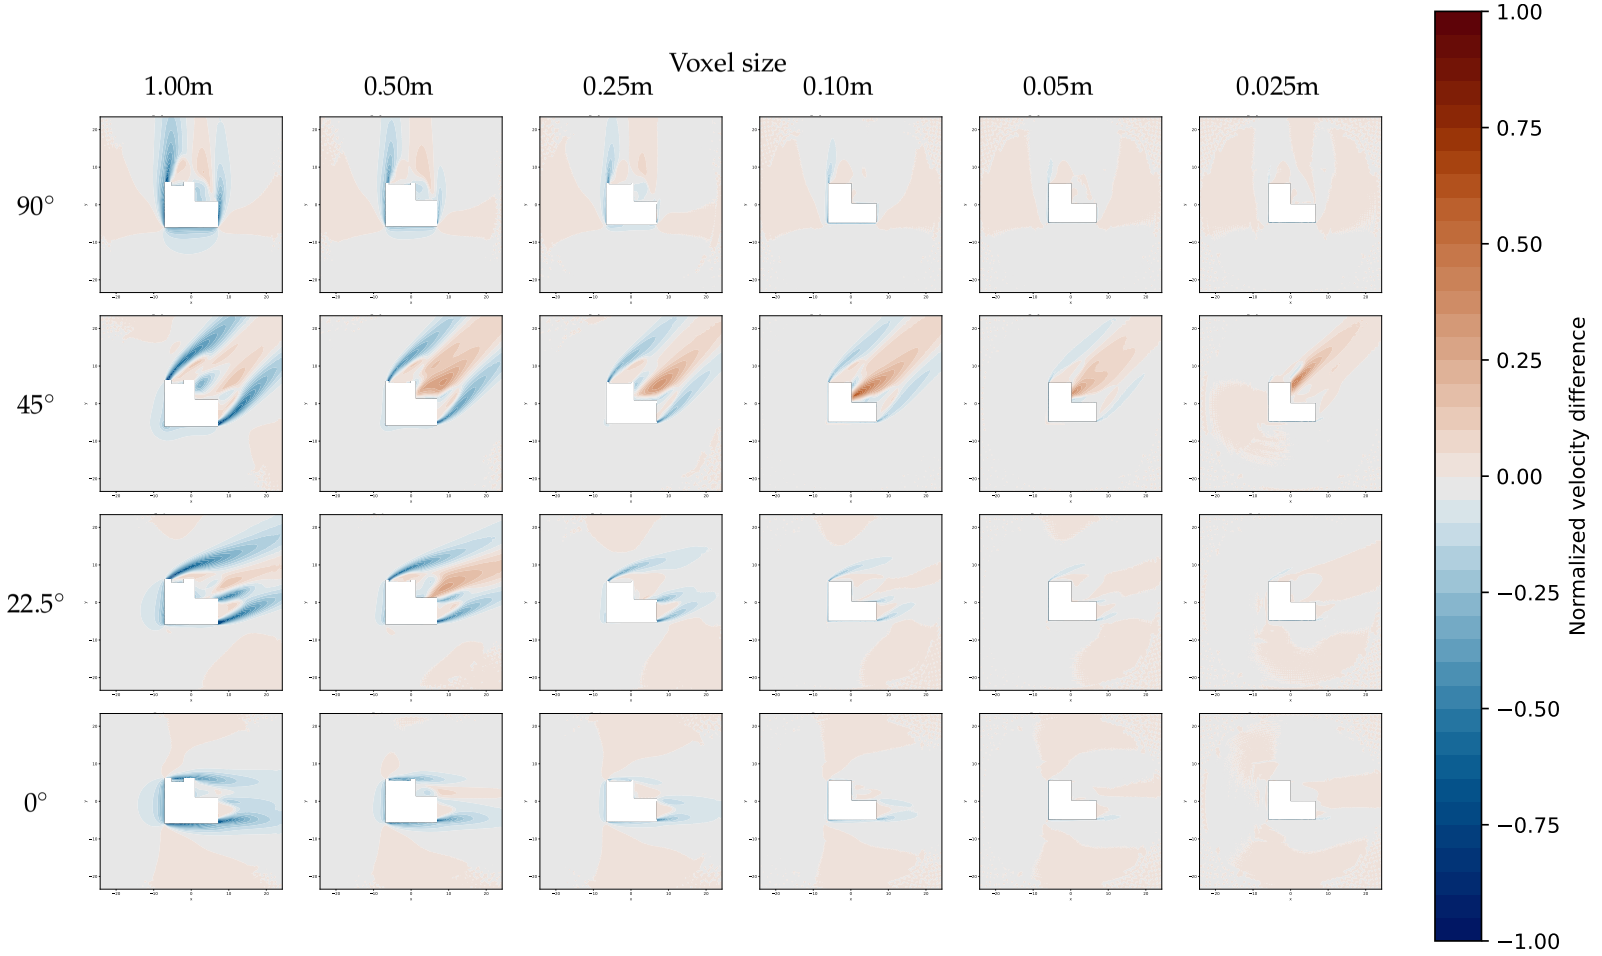

Figure 4.3.: Normalized velocity magnitude difference ($\frac{\|U_{\text{voxel}}\|-\|U_{\text{truth}}\|}{U_{\text{ref}}}$), relative to 5 m/s free stream velocity, at 3.2 m height for Building A-40. Columns show voxel sizes, rows show wind directions. Colorbar shown at right.

FZK-Haus

Building on the observations from the A-40 case, the next model — the FZK-Haus — offers a useful contrast to the A-40 model due to its geometric symmetry. A 3D representation can be seen in Figure 4.4; as the building is perfectly symmetrical, only one side is shown. Due to the symmetric nature of this building—which makes it less geometrically complex—it is expected to produce fewer velocity differences in the wake, particularly for the 45° diagonal wind direction.

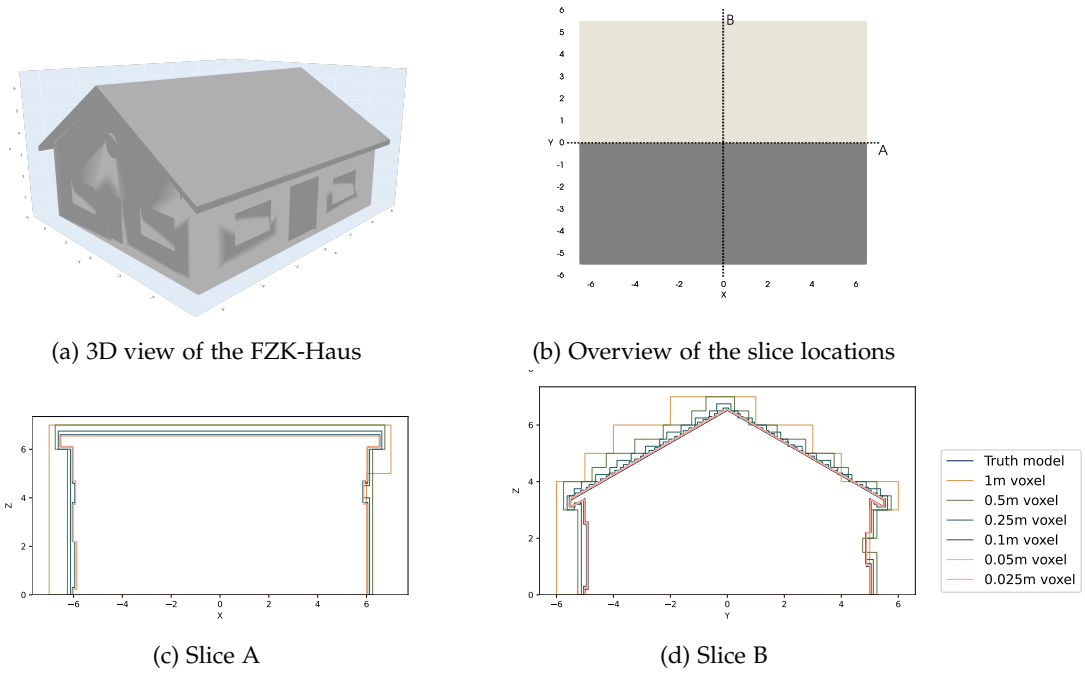

(a) 3D view of the FZK-Haus

(b) Overview of the slice locations

(c) Slice A

(d) Slice B

Figure 4.4.: Geometric difference between resolutions, visualized in 2D slices.

In the same figure, the 2D slices are also shown. As with the A-40 building, an increase in height and overall building size can be observed for the coarser voxel sizes, i.e., 1-meter, 0.5-meter, and 0.25-meter. In contrast to the A-40, however, the difference in height between these voxel sizes and the truth model is not as large. The greatest difference in height is 0.16 m for both the 1-meter and 0.5-meter voxel models. The other voxel models are all slightly lower than the truth model, but all fall within 0.1 m of its height. Because the differences between the voxel models and the truth model are smaller compared to those observed for the A-40, the resulting velocity differences may also be smaller. However, as the increase in building width is similar to that of the A-40, the velocity differences occurring at the building corners are still expected to be comparable to those observed for the A-40.

Interestingly, Figures 4.4c and 4.4d show that, in the coarser voxel models (0.25 m–1 m), not all windows visible in the truth model are represented. Although this effect is less significant than the increase in building width, it could still influence flow behavior for larger, window-dense buildings.

The results of the wind simulations are shown in Figure 4.5, using the same setup as for the A-40 building. For the FZK-Haus, the results follow the expected trend of decreasing differences with decreasing voxel size. In contrast to the A-40 building, no anomalies in the wake velocity differences are observed for any of the wind directions when the voxel size is reduced.

Both the length of the wake differences region and the magnitude of the corner-effect differences decrease substantially as the voxel size decreases. This reduction is more pronounced than for the A-40 building, particularly for the 45° wind direction, where the A-40 building retained larger difference regions. The largest differences for the FZK-Haus occur at the leading corners, while wake differences are negligible even for the 0.5-meter voxel size when the wind is parallel or perpendicular to the building faces. The corner differences are likely caused by the overestimation of the building footprint during voxelization.

Because the voxelized models for the FZK-Haus more closely match the geometry of the truth model compared to those of the A-40 building, the overall differences are smaller. Again the 45° wind direction results in the highest velocity differences between the voxel models and the truth model. Although the area where the differences occur is similar to that observed in the A-40 buildings simulations (Figure 4.3), the intensity of the velocity difference is lower for the FZK-Haus from the 0.5-meter voxel and smaller.

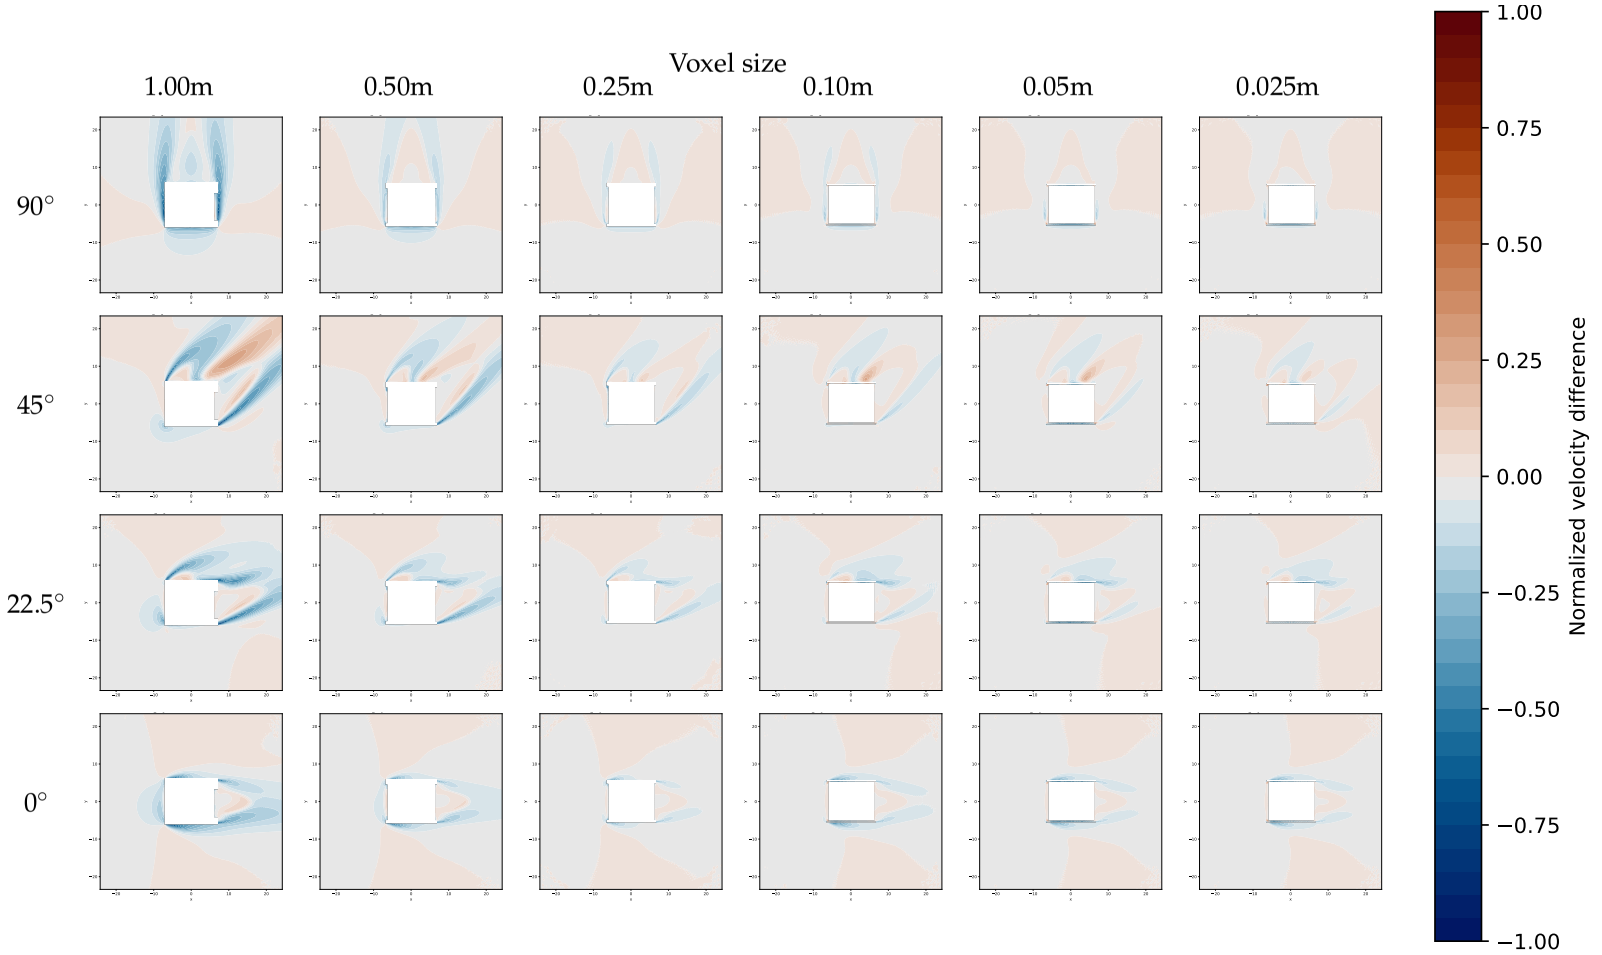

Figure 4.5.: Normalized velocity magnitude difference ($\frac{\|U_{\text{voxel}}\|-\|U_{\text{truth}}\|}{U_{\text{ref}}}$), relative to 5 m/s free stream velocity, at 3.2 m height for Building FZK-Haus. Columns show voxel sizes, rows show wind directions. Colorbar shown at right.

Institute

Finally, to examine how voxelization affects a larger and more complex geometry, the Institute building is analysed. Its most obvious difference compared to the other two buildings is its size: the Institute is more than twice as tall and has a footprint nearly six times larger. Another notable difference is the roof shape: the Institute has a rounded roof, unlike the more traditional sloped roofs of the other two buildings. The overall shape of the building is, like the FZK-Haus, nearly symmetric, as shown in Figure 4.6.

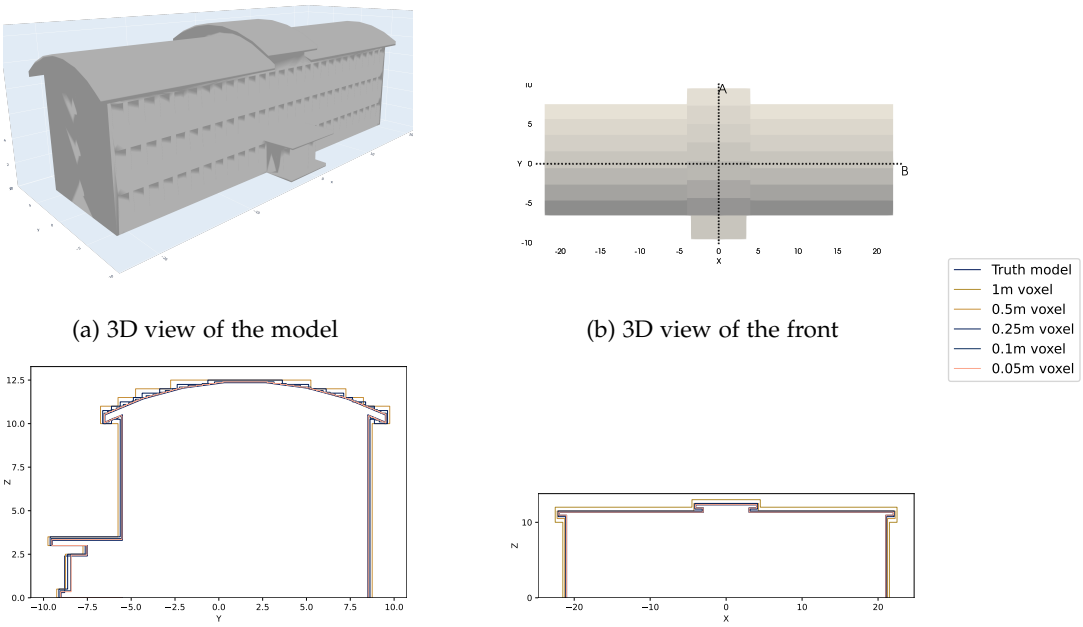

(a) 3D view of the model

(b) 3D view of the front

(c) Slice A

(d) Slice B

Figure 4.6.: Geometric differences between resolutions, visualised in 2D slices.

Analysing the slices shown in Figures 4.6c and 4.6d, the relative difference between the truth model and the voxelized models is much smaller than for the other two buildings. The largest difference in building height occurs between the 1-meter voxel model and the truth model, amounting to 0.65 meters, while the differences for the 0.5-meter and 0.25-meter voxel models are only 0.15 meters. Although this deviation is greater than that observed for the FZK-Haus (a maximum of 0.16 meters), it remains less than for the A-40 building (a maximum of 0.85 meters). When these values are normalised relative to the building's height, the Institute shows the smallest increase: 5%, compared to 10% for the A-40 and 13% for the FZK-Haus.

While its relative geometric differences are smaller, the Institute's rounded roof introduces a new consideration: flow separation may be more strongly affected by voxelization. Because the roof curvature has a large radius, voxelization -—particularly at coarser resolutions—- introduces larger steps in the surface profile. These abrupt steps may influence the roof's flow separation points, with separation expected to occur at the end of the roof for the smooth (truth) model and earlier, still above the roof, for the voxelized representations.

The results of the simulations are shown in Figure 4.7. The setup is the same as for the previous two buildings, ensuring direct comparability. One key difference, however, is that here the slices are taken at approximately 9 meters instead of 3 meters, since the roof begins at approximately this height.

The simulations reveal that the largest differences occur for the 90° wind direction, unlike the other two buildings where the maximum differences appeared at 45°. For the Institute, the 90° direction corresponds to the broad side of the building, producing a much larger wake difference. Whereas the wake differences for the FZK-Haus and A-40 buildings were contained within three times the building height, this is not the case for the Institute. This can be explained by the width–height ratios: the FZK-Haus and A-40 buildings had widths approximately two and 2.5 times their heights respectively, whereas the Institute's width is 3.5 times its height. As a result, the wake differences extend further downstream. This is also evident for the 0° wind direction, where the wake difference is minimal –—almost non-existent—- since this side of the building is only 1.5 times its height.

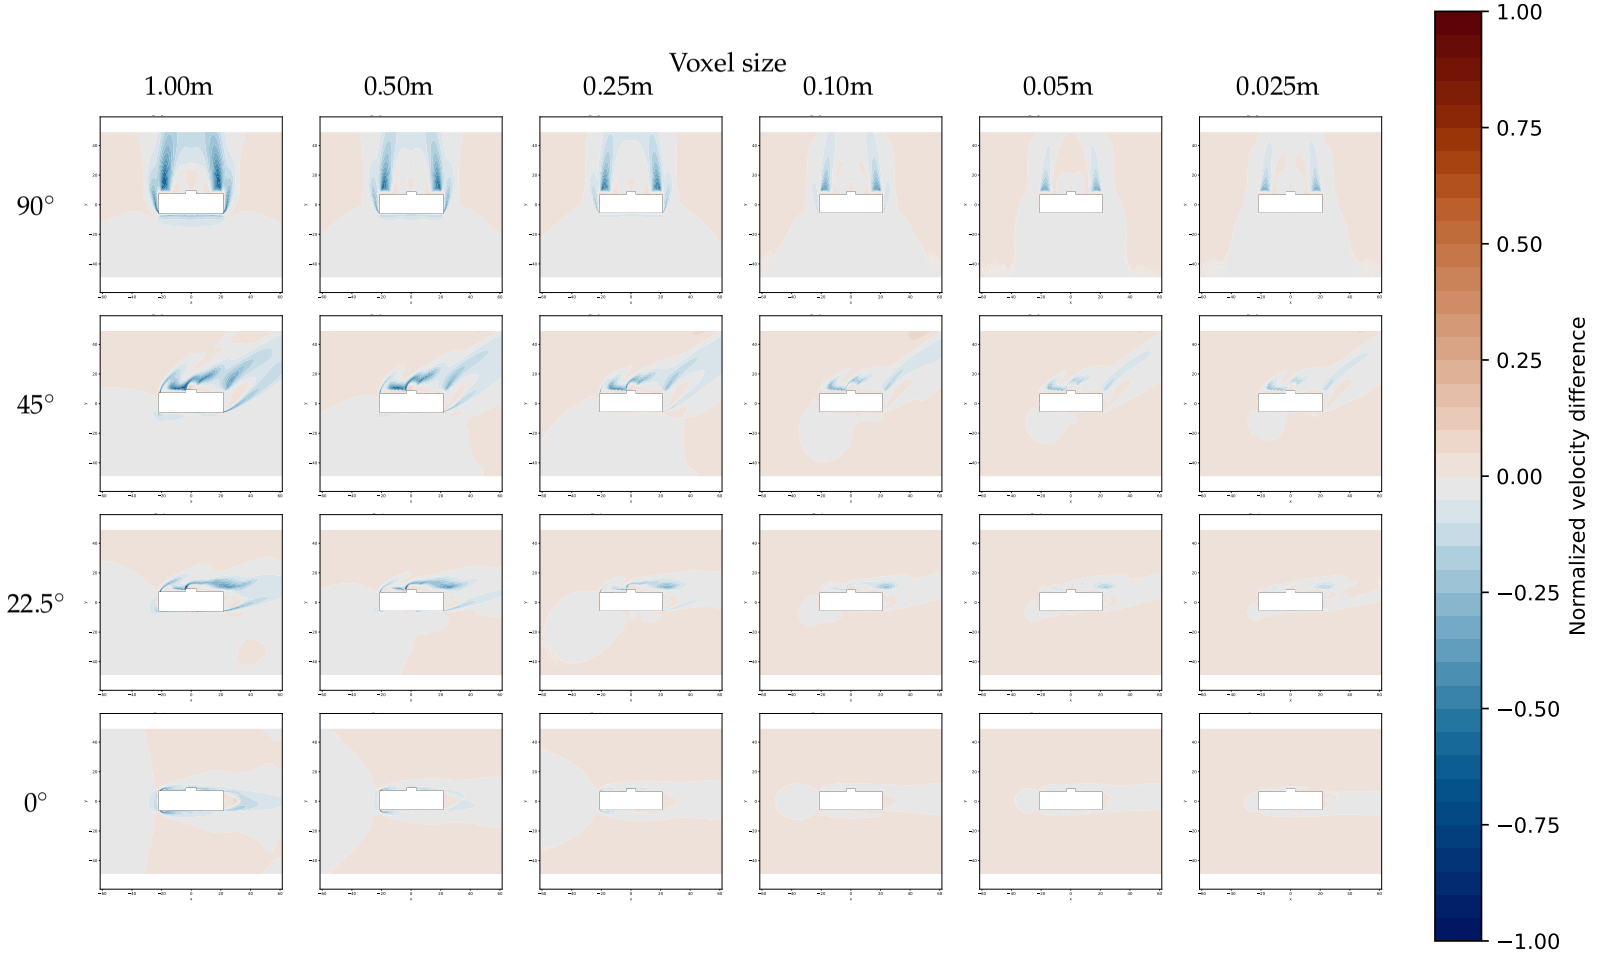

Figure 4.7.: Normalized velocity magnitude difference ($\frac{\|U_{\text{voxel}}\|-\|U_{\text{truth}}\|}{U_{\text{ref}}}$), relative to 5 m/s free stream velocity, at 8.7 m height for Building Institute. Columns show voxel sizes, rows show wind directions. Colorbar shown at right.

4.2.2. Wind flow structure

A-40

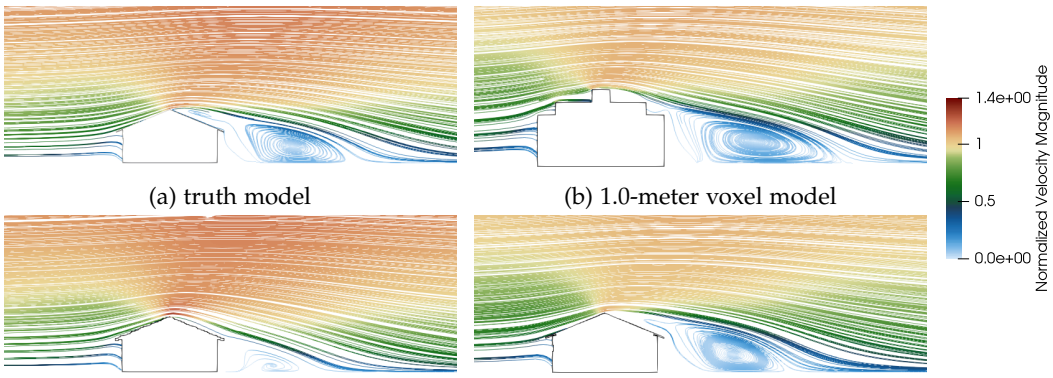

(a) truth model (b) 1.0-meter voxel model

(c) 0.1-meter voxel model (d) 0.025-meter voxel model

Figure 4.8.: A stream tracer along the 45° wind direction, trough the center of the A-40 building. Velocity is normalized by the inlet velocity (5 m/s).

Figure 4.8 shows a stream tracer comparison of the A-40 building for the 45° wind direction. The four panels show the truth model, 1.0-meter voxel model, 0.1-meter voxel model, and 0.025-meter voxel model. Comparing the angled wind flow over the two roofs, the high positive differences observed in Figure 4.3 can be explained. Due to the increased roof angle caused by voxelization of the sloped roof, the airflow experiences a steeper downward deflection after passing over the roof crest in the 0.1-meter voxel model. This steeper downward deflection leads to the differences observed in Figure 4.3, as the higher wind speeds are deflected closer to the ground in the voxel model than in the truth model, resulting in larger deviations at the sampling height.

Figure 4.9 presents a deeper analysis for the A-40 building, comparing the truth model and the 0.1-meter voxel model across three wind directions: 45°, 22.5°, and 0°. For both the 0° and 22.5° wind directions, the truth and voxel models produce similar wake structures with only slight variations. In contrast, at 45°, a clear deviation appears, with a substantially different wake structure for the voxelized model.

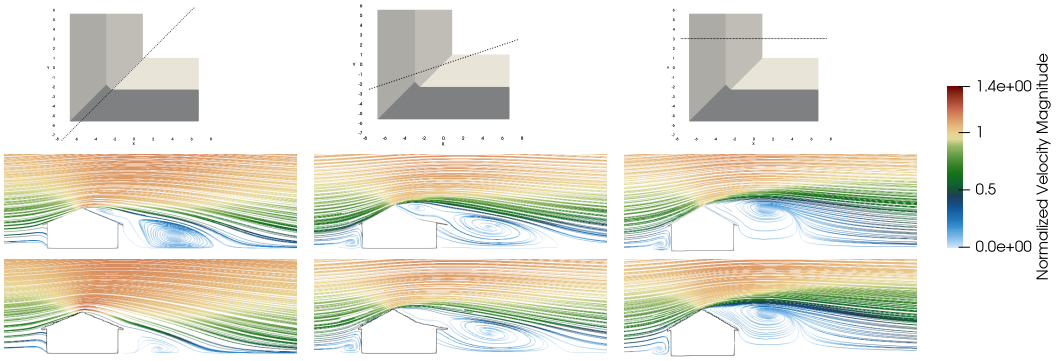

(a) 45° wind direction (b) 22.5° wind direction (c) 0° wind direction

Figure 4.9.: A comparison between the 45°, 22.5° and 0° wind directions for the A-40 building. The top row shows the truth model and the bottom row shows the 0.1-meter voxel model. Velocity is normalized by the inlet velocity (5 m/s)

The deviation at 45° is linked to the geometry of the ridge line. Although the 0.1-meter voxel model is identical for all wind directions, the ridge is intersected at different angles, changing its effective geometry. Specifically, at 45°, the ridge becomes wider by a factor of $\sqrt{2}$ compared to the 0° case since the voxel elements are cut diagonally. In addition, the roof slope appears lower for the 45° direction than for 0°, giving the flow a gentler rise toward the ridge and producing a wider horizontal ridge section. This allows the flow to remain attached longer and shifts the separation point downstream in the voxel model compared to the truth model.

For 0° and 22.5° wind directions, the ridge is narrower and the roof slope steeper. In these cases, the flow separates directly at the ridge for the voxelized and truth models, leading to similar wake structures and consequently much smaller velocity differences than in the 45° case.

FZK-Haus

Figure 4.10 shows a slice through the center of the FZK-Haus along the 45° wind direction, visualizing the wind flow through stream tracers for the truth model and the 1.0-meter, 0.1-meter, and 0.025-meter voxel models. The figure highlights several effects of voxelization. For the truth model, the flow separates cleanly at the top of the roof ridge, forming a single coherent recirculation zone. In the voxelized models, however, the separation point shifts to the sides of the highest voxels, and the shape of the recirculation zone is strongly dependent on voxel size. In the 1.0-meter voxel model, the stair-stepped geometry of the roof alters the main recirculation zone, producing multiple smaller eddies rather than one large coherent eddy. This difference in recirculation zones is voxel-size specific and disappears for the finer 0.025-meter model, which maintains a single recirculation zone similar to the truth model.

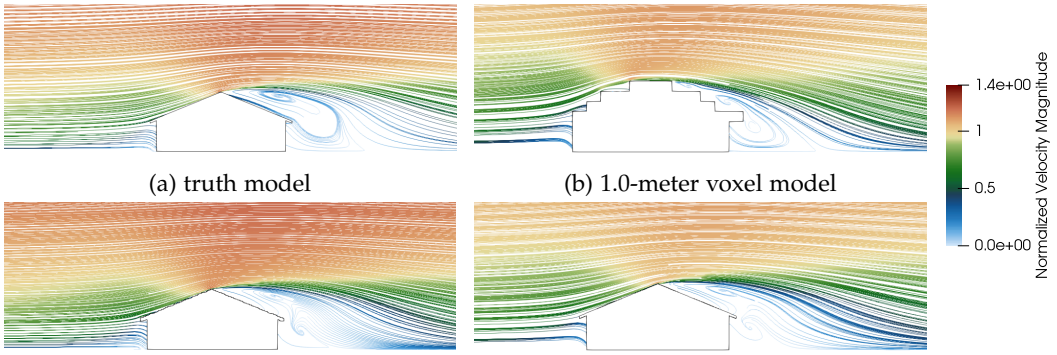

(a) truth model (b) 1.0-meter voxel model

(c) 0.1-meter voxel model (d) 0.025-meter voxel model

Figure 4.10.: A stream tracer along the 45° wind direction, trough the center of the FZK-Haus building. Velocity is normalized by the inlet velocity (5 m/s).

For the 0.1-meter voxel model, the change in recirculation occurs due to a different factor. Although the ridge of the voxel model is flattened, it spans only a single 0.1-meter voxel. This shifts the separation point only slightly downstream, to the trailing edge of the voxel (about 0.05 m). This minor shift is not the main cause of the altered wake structure. Both the truth model and 0.1-meter voxel model have wakes of similar overall size, but the main eddy in the truth model is located above the roof, while in the voxelized model the eddy is displaced closer to the ground behind the building. Analysing the recirculation zone of the 0.1-meter voxel model shows that the stepped roof geometry interferes with the flow recirculation. Whereas the smooth roof of the truth model supports a recirculation zone above the roof, the stepped roof of the voxel model interrupts this recirculation and redirects the flow parallel to the roof. In the stream tracers, this can be seen as lines that stop at the roof surface; since this is a 2D slice at the end of a 3D line, the flow at this point either goes into or out of the plane, parallel to the roof. This indicates that the stepped geometry does not just shift the separation location but fundamentally alters the three-dimensional recirculation patterns over the building.

To investigate the effect of wind direction in more detail, Figure 4.11 compares the flow for the truth model and the 0.1-meter voxel model at three wind directions: 45°, 22.5°, and 90°. This comparison allows a closer look at why the 45° case exhibits larger deviations compared to the other directions, even for the finer voxel models.

The larger velocity difference for the 45° direction is comparable to that of the A-40 building, although the magnitude of the velocity differences is lower than for the A-40 building. The cause of the velocity differences for the FZK-Haus differ from that of the A-40 building. For the FZK-Haus, the wake discrepancies between the truth model and the voxel model arise mainly from the interruption of the recirculating flow by the voxelized roof. Figure 4.11 illustrates why this disruption occurs for the 45° direction, but not for the 22.5° or 90° cases.

At 22.5°, the wind passes over only part of the roof's height (starting and ending halfway up the slope), while spanning more than the full roof width. This creates a very gentle effective slope, meaning the ridge is much less pronounced compared to the 45° and 90° directions. As a result, the flow does not strongly separate at the ridge but instead remains attached, fully separating at the roof's trailing edge. Since both the truth model and voxel model have nearly identical wall geometries (the voxelized version being only slightly wider; see Figure 4.4c), the resulting recirculation zones are also nearly identical.

For the 90° wind direction, the recirculation zone also develops above the roof, but it is not significantly affected by voxelization. A possible explanation is that the roof slope at 90° is steeper than at 45°. Combined with the fact that the effective voxel width is smaller at 90° (due to the way the voxels are intersected), the flow in the recirculation zone above the roof does not become horizontal enough to interact with the vertical faces of the voxel steps. Consequently, the recirculation remains largely uninterrupted by voxelization effects.

Therefore, voxelization only disrupts the wake for the 45° wind direction, while for both 22.5° and 90° the flow separates and reattaches in a manner very similar to the truth model.

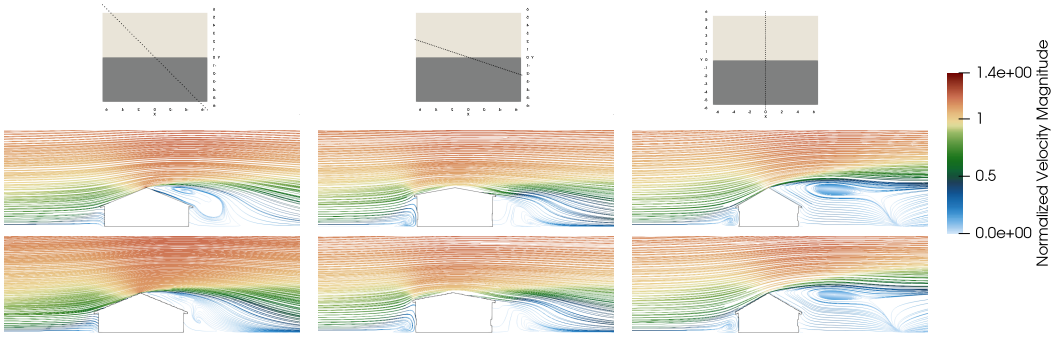

(a) 45° wind direction (b) 22.5° wind direction (c) 90° wind direction

Figure 4.11.: A comparison between the 45°, 22.5° and 90° wind directions for the FZK-Haus building. The top row shows the truth model and the bottom row shows the 0.1-meter voxel model. Velocity is normalized by in the inlet velocity (5 m/s).

Institute

The biggest difference between the building geometry of the Institute building and the FZK-Haus and A-40 buildings is the roof shape. Thus, looking at the velocity differences for the Institute building, it seems that the voxelization of a rounded roof does not impact the outcome of simulations as much as for a sloped roof.

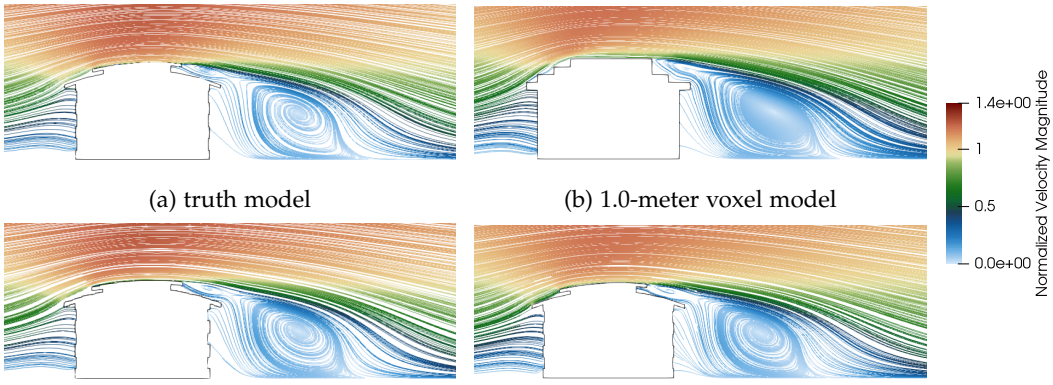

(a) truth model (b) 1.0-meter voxel model

(c) 0.1-meter voxel model (d) 0.025-meter voxel model

Figure 4.12.: A stream tracer along the 45° wind direction, trough the center of the Institute building. Velocity is normalized by the inlet velocity.

This can be well seen when comparing Figures 4.10 and 4.12 with each other. Where the flow separation point shifted significantly for the 1-meter voxel model of the FZK-Haus building, this is not the case with the 1-meter voxel mdoel of the Institute building. As this slice is taken trough the center of the building, the slice includes the heightened middle portion of the building. Due to this heightened portion the flow separation point is where the roof steps down from the higher part of the roof to the lower part of the roof. Altough the roof of the 1-meter voxel model is not curved like the roof of the truth model is, the length of the upper parts of the roofs are roughly the same. Therefore the flow separation point – which is after the upper part of the roof – is roughly the same for both models as well. There are however subtle differences between the two wake areas: where the truth model has a small eddy forming on the lower part of the roof, this doesn't happen for the 1-meter voxel model where another step in roof prevents this. And because the overall building height for the 1-meter voxel is heigher than that of the truth model (Figure 4.6), the height of the recirculation zone is higher for the 1-meter voxel model than that of the truth model.

As the Institute building has the layered roof only in the middle part of the buidling, Figure 4.13 shows the flow development over the front part of the roof that is continuously curved. On this part of the building, it can be seen that the coarser voxel models – 1-meter and 0.5-meter – have a different flow separation point from the truth model and the fine 0.1-meter voxel model. Both the coarser voxel models have the flow separation at the trailing edge of the upper voxels, instead of at the trailing edge of the entire roof. As the flow separates at an earlier stage – and thus at a higher point – the resulting eddy is larger for the 1-meter and 0.5-meter voxel models than it is for the truth model. The 0.1-meter voxel model follows the pattern of the truth model more, with a separation point at the edge of the roof and a similar sized eddy.

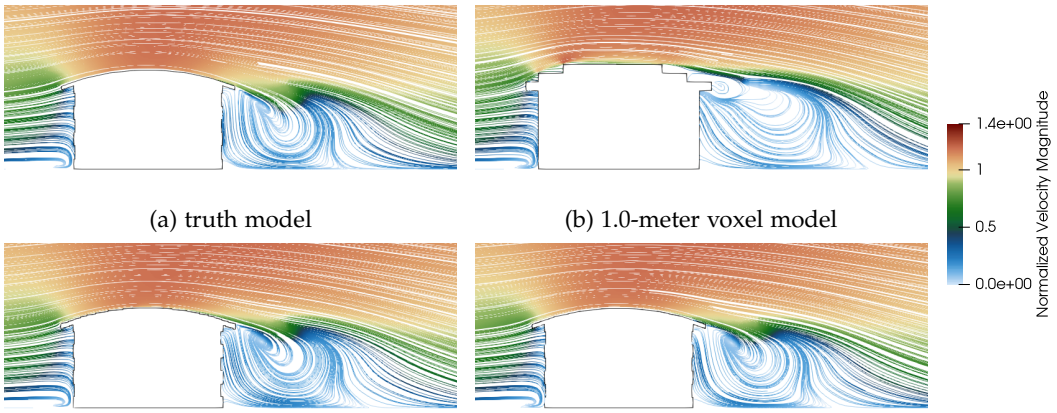

(a) truth model

(b) 1.0-meter voxel model

(c) 0.1-meter voxel model

(d) 0.025-meter voxel model

Figure 4.13.: A stream tracer along the 45° wind direction, trough the front part of the Institute building. Velocity is normalized by the inlet velocity (5 m/s)

A comparison between the 45°, 22.5°, and 90° wind directions is shown in Figure 4.14. For all three wind directions, the 0.1-meter voxel model produces a wake structure that is similar to that of the truth model. The largest difference in wake structure occurs for the 90° wind direction, where the voxel model exhibits a slightly larger wake.

This observation corresponds with the previously discussed velocity differences (Figure 4.7). The increased wake area for the 90° wind direction appears to be caused by an earlier flow separation in the voxel model. As the curvature of the roof increases toward the roof edges, the flow in the truth model separates from the roof before reaching the edge. In the voxel model, however, the discretization introduces smaller horizontal between the roof segments near to the edge. These smaller horizontal spaces between steps lead to an earlier separation of the flow compared to the truth model.

For the 45° wind direction, this earlier separation is not observed. The same geometric effect that influenced the previously analysed sloped-roof models also applies here: at 45°, the effective roof slope becomes gentler due to the diagonal intersection of the voxels. In contrast to the sloped-roof buildings, however, this gentler slope allows the flow to follow the curvature of the Institute's rounded roof until it reaches the trailing edge for both the truth and voxel models. As a result, the wake structures are nearly identical, and the corresponding velocity differences (Figure 4.7) remain small for the 45° wind direction.

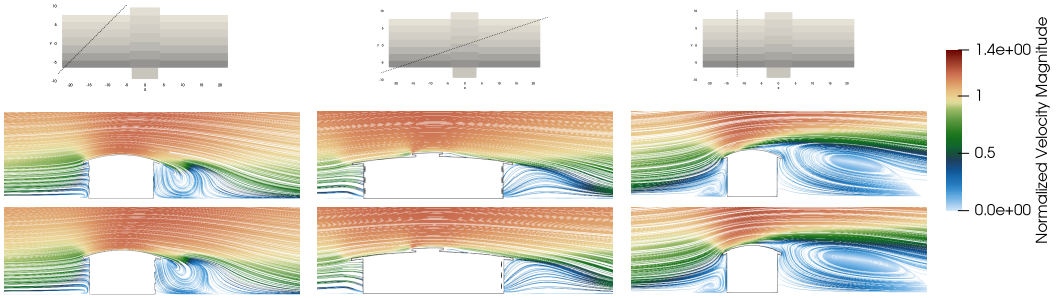

(a) 45° wind direction (b) 22.5° wind direction (c) 90° wind direction

Figure 4.14.: A comparison between the 45°, 22.5° and 90° wind directions for the Institute building. The top row shows the truth model and the bottom row shows the 0.1-meter voxel model. Velocity is normalized by the inlet velocity (5 m/s)

4.3. Computational performance

In this section, the computational performance of the different voxel models is compared to that of the truth model. As explained in Chapter 3, the simulation meshes are based on the truth model mesh. This means that simulation times for the voxelized models are not directly representative, since the meshes for coarser voxelizations are still over-refined and therefore result in longer than necessary computation times.

To quantify computational performance in this thesis, the focus is on the runtime of the voxelization program itself. Figure 4.15 shows the voxelization times for six voxel sizes for each building, along with the truth model times. A sharp increase in computation time is evident for the smallest voxel sizes (0.05 m and 0.025 m), with a smaller but still noticeable increase at 0.1 m. By contrast, the three coarsest voxel sizes (1.0 m, 0.5 m, and 0.25 m) all have similarly low runtimes. Because the truth models are also generated by the same program, their computation times are indicated by the horizontal transparent lines in the background. The truth models were generated with a voxel size of 0.3 m, chosen as a balance between preserving fine details (such as windows) and keeping runtimes low.

The comparison shows that for the A-40 and FZK-Haus buildings, the voxelization time of the 0.1 m model lies just within the range of the truth model computation time. For the larger Institute building, however, the 0.1 m voxelization requires nearly four times as long as the truth model. More generally, voxelization times rise steeply as voxel size decreases below 0.25 m, eventually surpassing the truth model runtimes for all three buildings. For the smaller A-40 and FZK-Haus, this crossover occurs at or below 0.1 m, while for the Institute building the truth model is already faster than the voxelization at 0.25 m.

Based on computation time alone, the 0.25 m voxel size appears to provide the best trade-off. The increase in time compared to the 0.5 m voxel is small, while the 0.25 m voxel model remains significantly faster to generate than the LOD3.2 truth model.

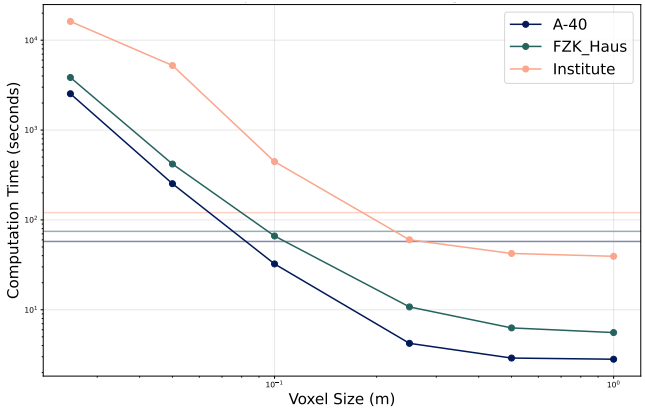

Figure 4.15.: Computation times of the three buildings per voxel size. The computation time of the truth model is showed as the horizontal semi-transparent line.

4.4. Key findings

When comparing the results for the three buildings, some clear patterns emerge. Velocity differences for the models with sloped roofs—as shown in Figures 4.3 and 4.5—are significantly larger than for the Institute model with a rounded roof. These differences are especially pronounced for the 45° wind direction. For the other wind directions, the velocity differences between the voxel models and the truth models become very small when the voxel sizes are reduced, whereas the 45° wind direction remains more sensitive, particularly for the A-40 and FZK-Haus models.

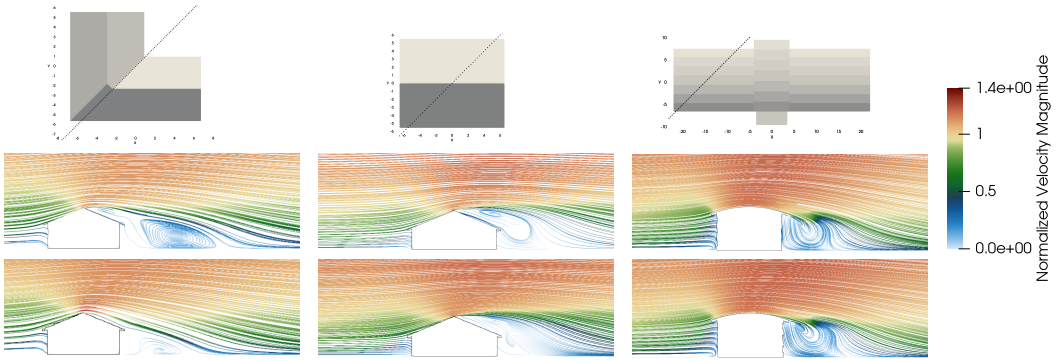

(a) A-40 (b) FZK-Haus (c) Institute

Figure 4.16.: A comparison of the wake development trough stream tracers using the 45° wind direction. The top row is the truth model and the bottom row is the 0.1-meter voxel model. The velocity if normalized by the inlet velocity (5 m/s).

The first finding — where the sloped roof models seem to be impacted more strongly by voxelization than the model with the rounded roof — can best be seen through the stream tracer plots. Figure 4.16 shows a comparison between the stream tracers of the truth model

and the 0.1-meter voxel model for all three buildings. The figure shows a clear change in wake structure for the sloped roof models, while the wakes in the rounded roof model are nearly identical. This highlights the influence of roof geometry on how sensitive the wake flow is to voxelization.

For the A-40 building, changes to the ridge geometry caused by voxelization result in noticeable shifts in flow separation and wake regions. In contrast, the FZK-Haus exhibits a different mechanism: the stepped roof geometry interrupts recirculation zones above the roof, leading to wake structures that diverge from the truth model under specific wind directions.

The influence of voxelization on wake behavior is also linked to how the voxel geometry presents itself at different wind directions. At 45°, the roof voxels create relatively large horizontal sections, which make the stepped roof more pronounced and strongly alter separation behavior compared to the truth model. At 22.5°, the voxels become wider in the horizontal direction, effectively flattening the overall slope of the roof, which appears to reduce the impact of voxelization as the flow remains more attached. At 0° and 90°, the horizontal and vertical extents of the voxels are equal, so the horizontal terraces are small and do not form large enough platforms for recirculating flow to enter, resulting in flow patterns more comparable to the truth model.

Finally, the time required to generate voxel models increases sharply as voxel size decreases, with the smallest voxels (0.05 m and 0.025 m) requiring the longest runtimes. For the smaller A-40 and FZK-Haus buildings, the 0.1 m voxel models are generated in a time comparable to the truth model, whereas for the larger Institute building the 0.1 m voxelization takes nearly four times as long. This demonstrates that computational cost is sensitive both to voxel size and building size, and highlights the practical trade-off between model detail and runtime. However, this behaviour is software-specific, and using a different algorithm capable of parallel computation could reduce these computation times even for the smallest voxel sizes tested in this thesis.

5. Conclusion

In this thesis three building models, each with different geometric properties, were voxelized and subsequently used in CFD simulations. This was done to address the main research question:

- *How does the voxel resolution within urban voxelized models influence the accuracy and computational efficiency of wind flow simulations using Computational Fluid Dynamics (CFD)?*

To answer this overarching question, four sub-questions were formulated and addressed.

The first sub-question asked:

- To what extent does voxelization impact key CFD output parameters, such as velocity and flow structure?

The analysis in Chapter 4 shows a clear trend: as voxel size decreases, velocity differences between the LoD model and voxel models generally decrease as well. At voxel sizes of 0.1 m or smaller, velocity differences are usually minor. With the biggest mean velocity difference being 0.32 m/s for the FZK-Haus with a 0.1-meter voxel and the smallest mean velocity difference being 0.082 m/s for the A-40 model with a 0.025-meter voxel. An important exception occurs for the 45° wind direction for the buildings with sloped roofs (A-40 and FZK-Haus), where significant differences persist even at relatively finer voxel resolutions. Here the biggest mean velocity difference is 0.5 m/s for the A-40 building and the smallest mean velocity difference is 0.23 m/s.

The higher velocity difference for the 45° wind direction is linked to how the voxels are intersected at this wind direction. Because of the 45° orientation, the horizontal extent of each voxel step increases, allowing the flow to interact more strongly with the stepped geometry. Within the recirculation zone, this can obstruct the returning flow, leading to a noticeably altered wake structure. Additionally, the enlarged horizontal sections near the ridge modify the ridge geometry itself, allowing portions of the flow to remain more horizontal and reducing the amount of separation compared to the truth model.

This links directly to the second sub-question:

- Are some geometric shapes more sensitive to voxelization than others?

The results demonstrate that geometry strongly influences sensitivity to voxelization. The rounded roof of the Institute model is more robust: wake structures remain relatively consistent across voxel sizes and closely match the truth model (see Figure 4.12). In contrast, the sloped-roof buildings (A-40 and FZK-Haus) show much stronger sensitivity. For these, voxelization leads to delayed flow separation or interruptions of the recirculation zone above the roof. These effects are particularly pronounced for coarse voxel sizes but diminish with finer voxel resolutions (see Figures 4.8 and 4.10).

These geometric effects also depend on wind direction. As discussed previously, the 45° wind direction produces effective roof and voxel shapes that enhance the impact of voxelization on wake formation. At 22.5°, however, the effective slope of the slanted roofs becomes shallow enough to delay flow separation, thereby reducing the voxelization influence on wake structures.

This is also part of the third sub-question:

- How sensitive are CFD outcomes to changes in voxel resolution across different regions of an urban model?

Two regions were identified as most sensitive to voxel resolution. The first is the building footprint. With coarse voxels (1 m or 0.5 m), the footprint is artificially widened, shifting flow reattachment zones and producing negative velocity differences at building corners. This effect diminishes as voxel resolution increases and the footprint aligns more closely with the truth model. The second region is the roof geometry. Because none of the roofs in this study are flat, they are consistently misrepresented at coarse voxel sizes, leading to the largest velocity differences in the wake. At fine resolutions, footprint effects largely disappear, but roof-related wake differences remain the dominant source of velocity differences (see Figure 4.3).

The final sub-question was:

- What is the minimum voxel resolution required to achieve an optimal balance between accuracy and computational efficiency?

Computation times reveal a steep increase as voxel size decreases below 0.25 m (see Figure 4.15). For all three buildings, the 0.25 m voxel models are faster to generate than the LoD3.2 truth models, while the 0.1 m voxel models already exceed the truth model's computation time for the Institute. In terms of accuracy, velocity differences reduce consistently with finer voxel sizes but show diminishing returns below 0.1 m. Taken together, this suggests that the practical balance point lies between 0.1 m and 0.25 m, though this depends on building size and wind direction.

Returning to the main research question, the findings show that voxel resolution has a clear influence on both CFD accuracy and computational efficiency. Coarse voxel models can be generated quickly, but at the cost of large deviations in flow structure and velocity compared to the truth models. Conversely, fine voxel models approach the accuracy of the LoD3.2 models but with computation times that may exceed them, particularly for larger buildings.

This makes the effectiveness of voxelization highly case-dependent. For smaller regions of interest—containing only a limited number of buildings—the relatively small time savings in preprocessing may not justify the loss in accuracy. For larger cases, such as city-scale simulations, however, the cumulative time savings per building can become substantial, making voxelization an attractive approach despite moderate accuracy losses.

5.1. Future Work

While this study provides insights into the influence of voxel resolution, several aspects for future research remain. Six of such aspect have been listed below.

First, all simulations in this thesis were conducted with a RANS turbulence model. Investigating voxelization effects under Large Eddy Simulation (LES) would provide more detailed information on the influence of voxel resolution on turbulent structures, especially in wake and recirculation regions.

Second, the models in this study were placed on an unobstructed plane. In realistic urban environments, obstacles such as trees, vegetation, or surrounding buildings may interact with voxelized geometries in ways not captured here. Exploring such scenarios would be valuable for assessing the robustness of voxelization in applied contexts.

Third, only four wind directions were tested, with results showing that 45° is particularly problematic for sloped roofs. Extending the analysis to finer directional increments (e.g., every 10°) could reveal additional orientations where voxelization produces unfavourable deviations.

Fourth, this study assessed computational efficiency using meshes derived from the LoD3.2 models, which were then applied to the voxel models. It is likely, however, that voxel models could be meshed more coarsely while still satisfying grid convergence requirements, leading to additional efficiency gains not quantified in this thesis.

Fifth, in this thesis all buildings were aligned with the Cartesian axes of the voxel grid. In reality, many buildings have non-rectangular footprints or orientations that are not grid-aligned. Such cases would produce jagged or staircase-like wall representations after voxelization. Analysing the effects of voxelization on non-rectangular or rotated buildings would therefore provide crucial insights into how alignment and footprint shape influence voxelization accuracy.

Finally, the study focused on buildings with relatively simple roof geometries. More complex shapes — such as pyramid roofs or highly detailed structures — may respond differently to voxelization and could help refine the understanding of which architectural forms are most sensitive.

A. Reproducibility self-assessment

A.1. Marks for each of the criteria

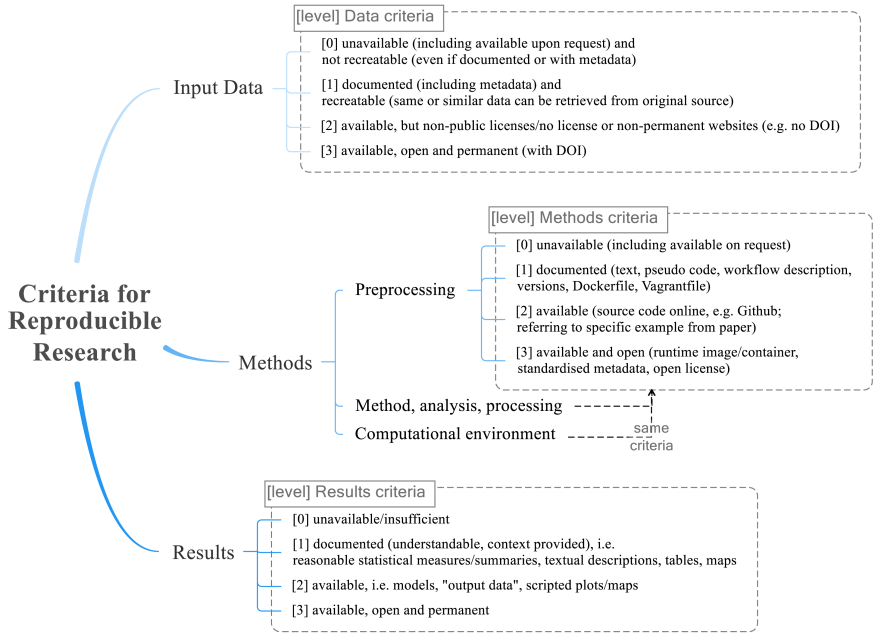

Figure A.1.: Reproducibility criteria to be assessed.

Criteria	Grade	Explenation
Input data	0-3	The FZK-Haus and Institute buildings are open source and have unrestricted use. The A-40 building however is not an open-source model
Preprocessing	3	The preprocessing uses only open source software available on Github with a GPL-3.0 license.
Methods	3	The method uses only open source software (OpenFOAM and ParaView) and the setup can be recreated using the methodology from this thesis.
Computational environment	3	Altough in this thesis DelftBlue is used which requires acces, the used software can be downloaded on any device (altough it can take longer depending on the hardware)
Results	2-3	The code that has been used to create the data in the results is available on github, making it possible to recreate the shown results.

The GitHub can be found here.

Bibliography

Ahrens, J., Geveci, B., and Law, C. (2005). Paraview: An end-user tool for large data visualization. In *Visualization Handbook*. Elsevier.

Biljecki, F., Ledoux, H., and Stoter, J. (2016). An improved lod specification for 3d building models. *Computers, Environment and Urban Systems*, 59:25–37.

Blocken, B. (2015). Computational fluid dynamics for urban physics: Importance, scales, possibilities, limitations and ten tips and tricks towards accurate and reliable simulations. *Building and Environment*, 91:219–245. Paper about CFD in general, the one used in the windmoddeling course.

Celik, I., Ghia, U., Roache, P. J., and Frietas, C. J. (2008). Procedure for estimation and reporting of uncertainty due to discretization in cfd applications. *Journal of Fluids Engineering*, 130(7):078001.

Delft High Performance Computing Centre (DHPC) (2024). DelftBlue Supercomputer (Phase 2). https://www.tudelft.nl/dhpc/ark:/44463/DelftBluePhase2.

Geiger, A., Benner, J., Häfele, K.-H., and Hagenmeyer, V. (2018). Thermal energy simulation of buildings based on the citygml energy application domain extension. In *BauSIM2018–7. Deutsch-Österreichische IBPSA-Konferenz: Tagungsband. Hrsg.: P. Von Both*, pages 295–302.

Heazel, C. (2021). Ogc city geography markup language (citygml) 3.0 conceptual model users guide. OGC User Guide 20-066, Open Geospatial Consortium. Copyright © 2021 Open Geospatial Consortium. To obtain additional rights of use, visit http://www.opengeospatial.org/legal/.

Hinks, T., Carr, H., Truong-Hong, L., and Laefer, D. F. (2013). Point cloud data conversion into solid models via point-based voxelization. *Journal of Surveying Engineering*, 139(2):72–83.

Institute for Automation and Applied Informatics (IAI) and Karlsruhe Institute of Technology (KIT). Fzk haus. Industry Foundation Classes 4.

Institute for Automation and Applied Informatics (IAI) and Karlsruhe Institute of Technology (KIT). Office building. Industry Foundation Classes 4.

Janssen, W., Blocken, B., and van Hooff, T. (2013). Pedestrian wind comfort around buildings: Comparison of wind comfort criteria based on whole-flow field data for a complex case study. *Building and Environment*, 59:547–562.

Jindal, S., Khalighi, B., Johnson, J. P., Chen, K.-H., and Iaccarino, G. (2007). The immersed boundary cfd approach for complex aerodynamics flow predictions. *SAE Transactions*, 116:50–60.

Karalit (2023). The immersed boundary method versus the cut-cell method. Technical report, Karalit. Accessed: 2025-10-22.

Kortelainen, J. (2009). *Meshing Tools for Open Source CFD: A Practical Point of View*. Number VTT-R-02440-09 in VTT Research Report. VTT Technical Research Centre of Finland, Finland.

Li, X., Yang, M., Bi, L., Xu, R., Luo, C., Yuan, S., Yuan, X., and Tang, Z. (2024). An efficient cartesian mesh generation strategy for complex geometries. *Computer Methods in Applied Mechanics and Engineering*, 418:116564.

Masson, V., Heldens, W., Bocher, E., Bonhomme, M., Bucher, B., Burmeister, C., de Munck, C., Esch, T., Hidalgo, J., Kanani-Sühring, F., Kwok, Y., Lemonsu, A., Lévy, J., Maronga, B., Pavlik, D., Petit, G., See, L., Schoetter, R., Tornay, N., Votsis, A., and Zeidler, J. (2020). City-descriptive input data for urban climate models: Model requirements, data sources and challenges. *Urban Climate*.

Mitkov, R., Pantusheva, M., Petrova-Antonova, D., Naserentin, V., and Logg, A. (2024). The role of computational fluid dynamics within city digital twins: Opportunities and challenges. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*.

OpenFOAM (2019). OpenFOAM: The open source cfd toolbox. https://www.openfoam.org. Version 7.

Patil, A. and García-Sánchez, C. (2025). Quantifying the impact of urban geometric detail for urban air mobility risk forecasting. *Sustainable Cities and Society*, 132:106750.

Pađen, I., Peters, R., García-Sánchez, C., and Ledoux, H. (2024). Automatic high-detailed building reconstruction workflow for urban microscale simulations. *Building and Environment*, 265:111978.

Ridzuan, N., Ujang, U., Azri, S., Yusoff, I. M., and Choon, T. L. (2022). Voxelization techniques: Data segmentation and data modelling for 3d building models. In *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume XLVIII-4/W3-2022, pages —, Castelo Branco, Portugal. The 7th International Conference on Smart City Applications, Copernicus Publications.

Roache, P. J. (1994). Perspective: A method for uniform reporting of grid refinement studies.

Roache, P. J. (1997). Quantification of uncertainty in computational fluid dynamics. *Annual Review of Fluid Mechanics*, 29:123–160.

Sanchez, C. G., Vitalis, S., Paden, I., and Stoter, J. (2021). The impact of level of detail in 3d city models for cfd-based wind flow simulations. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences - ISPRS Archives*, pages 67–72. Short introduction to the effects of different LoDs on the outcome of a CFD. Only for LoD 1.3 and 2.2.

Setaih, K., Hamza, N., Mohammed, M. A., Dudek, S., and Townshend, T. (2014). Cfd modeling as a tool for assessing outdoor thermal comfort conditions in urban settings in hot arid climates. *Journal of Information Technology in Construction (ITcon)*, 19:248–269.

Toja-Silva, F., Kono, T., Peralta, C., López-García, , and Chen, J. (2018). A review of computational fluid dynamics (cfd) simulations of the wind flow around buildings for urban wind energy exploitation. *Journal of Wind Engineering and Industrial Aerodynamics*, 180:66–87.

Vaart, J. (2025). Ifcenvelopeextractor. https://github.com/jaspervdv/IFC_BuildingEnvExtractor. Version 0.2.5, GitHub repository.

van der Vaart, J., Stoter, J., Agugiaro, G., Arroyo Ohori, K., Hakim, A., and El Yamani, S. (2024a). Enriching lower lod 3d city models with semantic data computed by the voxelisation of bim sources. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, X-4/W5-2024:297–308.

van der Vaart, J., Stoter, J., Diakité, A., Biljecki, F., Ohori, K. A., and Hakim, A. (2024b). Assessment of the lod specification for the integration of bim-derived building models in 3d city models. In Kolbe, T. H., Donaubauer, A., and Beil, C., editors, *Recent Advances in 3D Geoinformation Science*, pages 171–191, Cham. Springer Nature Switzerland.

Vanky, P., Mark, A., Hunger, F., Villamor Saucedo, G., Haeger-Eugensson, M., Bennetsen, J. C., Tarraso, J., Adelfio, M., Sasic Kalagasidis, A., and Sardina, G. (2024). Evaluation of an immersed boundary numerical framework to address the wind field in complex urban topographies. *Building and Environment*, 266:112036.

Villi, G. and Carli, M. D. (2014). Detailing the effects of geometry approximation and grid simplification on the capability of a cfd model to address the benchmark test case for flow around a computer simulated person. *Building Simulation*, 7(1):35–55.

Wang, X. and McNamara, K. F. (2006). Evaluation of cfd simulation using rans turbulence models for building effects on pollutant dispersion. *Environmental Fluid Mechanics*, 6(3):181–202.

Wijesooriya, K., Mohotti, D., Lee, C.-K., and Mendis, P. (2023). A technical review of computational fluid dynamics (cfd) applications on wind design of tall buildings and structures: Past, present and future. *Journal of Building Engineering*, 74:106828.

Xiong, M., Chen, B., Zhang, H., and Qian, Y. (2022a). Study on accuracy of cfd simulations of wind environment around high-rise buildings: A comparative study of k-epsilon turbulence models based on polyhedral meshes and wind tunnel experiments. *Applied Sciences*, 12(14):7105.

Xiong, M., Chen, B., Zhang, H., and Qian, Y. (2022b). Study on accuracy of cfd simulations of wind environment around high-rise buildings: A comparative study of k-epsilon turbulence models based on polyhedral meshes and wind tunnel experiments. *Applied Sciences*, 12(14):7105.

Yang, Z. and Shih, T. (1993). New time scale based k-epsilon model for near-wall turbulence. *AIAA Journal*, 31:1191–1198.

Yu, F., Zhang, H., Li, Y., Xiao, J., Class, A., and Jordan, T. (2018). Voxelization-based high-efficiency mesh generation method for parallel cfd code gasflow-mpi. *Annals of Nuclear Energy*, 117:277–289.

Colophon

This document was typeset using LaTeX, using the KOMA-Script class `scrbook`. The main font is Palatino.