The analysis of complex physical systems in industrial environments relies on high-fidelity numerical simulations, which are often computationally prohibitive, particularly in many-query scenarios such as design optimization, parametric studies, and uncertainty quantification. This limitation is further exacerbated by geometric variability, as each configuration typically requires the generation of a dedicated computational mesh. These challenges motivate the development of non-intrusive surrogate modeling strategies capable of leveraging available simulation data while remaining flexible with respect to both geometry representation and discretization.

This thesis investigates deep learning-based surrogate models for predicting physically relevant quantities in computational fluid dynamics (CFD) problems characterized by geometric variability. Classical approaches exhibit practical limitations, as they generally rely on a shared reference mesh or assume the availability of explicit geometric parametrizations, restricting their applicability in real-world industrial scenarios.

The objective of this work is to develop an end-to-end predictive framework capable of estimating fluid dynamic quantities of interest directly from geometrically varying configurations. By avoiding reliance on task-specific parametrizations, the proposed approach enables the integration of heterogeneous datasets and facilitates the reuse of previously generated designs, even when their underlying parametrization is unavailable or inconsistent.

In addition, the framework is designed to avoid dependencies on fixed geometric discretizations, such as volumetric or surface meshes. These representations often require costly and error-prone re-meshing or registration procedures to handle geometric variability across the dataset. Moreover, fixed discretizations introduce inherent resolution limitations, which may constrain the accuracy of the learned models. To address these issues, the proposed methodology adopts an implicit representation of the geometry based on signed distance functions (SDFs), allowing the domain to be modeled as a continuous function rather than a discrete set of values defined on mesh nodes or elements. This continuous formulation enables the model to capture geometric features with potentially arbitrary resolution and to extract information directly from the implicit description, bypassing the need for mesh alignment or interpolation steps.

The approach is implemented using Implicit Neural Representations (INRs), where geometry and solution fields are encoded through two separate models operating on point samples drawn from the underlying domain. A feed-forward neural network is subsequently trained to learn a mapping between the corresponding latent spaces, effectively linking the geometric representation to the associated physical solution. This formulation enables mesh-independent predictions and allows the evaluation of the predicted fields at arbitrary spatial locations. Moreover, the modularity of the framework allows to predict a wide range of outputs, including global integral quantities, such as lift and drag coefficients, as well as local high-fidelity scalar and vector fields, such as pressure and velocity distributions, defined over the entire spatial domain.

A key contribution of this work is the introduction of a regularization strategy for the INR-based encoders aimed at preserving the geometric structure within the latent space. This is achieved by enforcing consistency between distances in the geometry space and the corresponding distances in the latent representation, thereby promoting a coherent organization of the encoded samples.

The proposed framework is validated on three CFD test cases, demonstrating that the introduced regularization improves both predictive accuracy and generalization capabilities of the model, while maintaining robustness with respect to geometric variability.