Abstract

Artificial Intelligence models are proving extremely useful and are rapidly spreading across the industrial landscape; however, their robustness remains an open challenge that requires further investigation. The use of Machine Learning models—mostly data-driven—poses issues in high-risk domains where model reliability is essential, where understanding the reasoning behind a model's decision is crucial, or in situations where data is scarce or noisy, making purely data-driven approaches unreliable. These concerns have given rise to various research directions, all aiming to improve the trustworthiness and reliability of Machine Learning models. One of the most promising approaches focuses on enhancing models' robustness by incorporating domain knowledge of various kinds into predominantly data-driven Machine Learning models. These techniques are referred as *Knowledge-Informed* Machine Learning.

This Ph.D. project has been dedicated to advancing research in this area by organizing existing techniques and pushing the boundaries of the field through new applications, with the ultimate goal of accelerating technology transfer to industry. This final thesis presents the key outcomes of this work. Firstly, we report a comprehensive study and classification of existing Knowledge-Informed Machine Learning techniques. To date, in industrial contexts, these techniques have been applied almost exclusively to predictive modeling. For the first time, a formal analysis of these techniques is provided, comparing them with traditional Machine Learning approaches and offering an extensive overview of those already adopted in industry. Having laid these foundations, the thesis moves on to its core focus: the development of new practical applications and the expansion of research into emerging directions. Two representative case studies are presented. The first is a practical implementation of a Knowledge-Informed Predictive Model, specifically a power flow analysis using physics to inform a Machine Learning model. In this case study, existing iterative models

for solving the underlying physical equations were highly inefficient, making them unsuitable for real-time control applications. To address this, a surrogate model was developed by integrating multiple knowledge-informing techniques, enabling accurate predictions while significantly outperforming traditional models in terms of both efficiency and reliability. The second application explores a Knowledge-Informed Generative Model, a class of models still in its early stages but already showing potential for industrial use. In this case, the generative modeling framework, specifically diffusion models, was employed to tackle an inverse design problem applied to mechanical metamaterials. A completely novel technique was introduced, which we demonstrated to be highly effective in scenarios where strict adherence to design constraints is crucial. Both applications demonstrate the advantages of using Knowledge-Informed Machine Learning to enhance model robustness and validate the research carried out throughout this doctoral project. This manuscript concludes by outlining the ongoing developments and what we believe to be the future prospects of this line of research.