Summary

Quantum computing is often presented as a disruptive technology with the potential to accelerate problems that are hard for classical machines. Yet, in the current NISQ (Noisy Intermediate-Scale Quantum) era, the main challenge is no longer only the design of formally correct quantum algorithms, but their translation into practical workflows compatible with real hardware and meaningful application domains. This thesis addresses that challenge by investigating practical hybrid quantum-classical methodologies for neutral-atom quantum processors, with particular attention to analog computation, hardware-aware embedding, combinatorial optimization, and quantum machine learning.

A first contribution concerns the development of embedding strategies for neutral-atom platforms. Since, in analog processors, computation is tightly linked to geometry and distance-dependent interactions, mapping abstract problems onto the physical register is a central part of the algorithmic pipeline. This thesis studies methods to construct feasible and scalable embeddings for graph-based problems, showing how data representation and hardware constraints directly affect the quality and applicability of quantum routines.

Building on this foundation, the thesis explores hybrid approaches to quantum optimization. The main focus is on graph-based combinatorial problems, in particular vertex coloring and maximum-weighted independent set, which are relevant both from a theoretical viewpoint and for industrial applications such as Physical Cell Identity assignment and the design of space–terrestrial integrated networks. Several increasingly advanced hybrid frameworks are developed, combining classical optimization techniques with quantum subroutines on neutral-atom hardware. The results show that these approaches can improve solution quality while preserving resource efficiency, and underscore the importance of tight problem formulations, strong classical bounds, and hardware-aware branching or decomposition strategies.

The final major research direction concerns quantum machine learning. Two case studies are considered. The first investigates graph classification through quantum evolution kernels on neutral-atom processors, applied to protein classification. The results demonstrate that quantum dynamics can define informative similarity measures for graph-structured data, with hardware execution on Aquila's quantum processor achieving performance comparable to strong classical baselines. The second case study

extends beyond graph data and studies credit card default prediction through quantum reservoir computing. In this setting, the quantum system acts as a nonlinear feature generator within a hybrid learning pipeline. While hardware noise affects performance more significantly than in the kernel-based case, the results show that quantum-generated features remain useful and that classical post-processing can partially mitigate imperfections.

Overall, the thesis shows that the practical value of current quantum computing lies not in isolated quantum routines, but in carefully designed hybrid pipelines in which quantum and classical resources complement each other. The results highlight embedding as a fundamental problem, confirm the relevance of hardware-aware algorithm design, and show that neutral-atom platforms already support meaningful application-driven studies in both optimization and machine learning. These findings contribute to narrowing the gap between theoretical quantum computing and realistic industrial use cases, while also identifying the main limitations and opportunities for future research in the NISQ regime.

This Ph.D. thesis has been typeset by means of the TEX-system facilities. The typesetting engine was LuaLATEX. The document class was `toptesi`, by Claudio Beccari, with option `tipotesi=scudo`. This class is available in every up-to-date and complete TEX-system installation.