Autore

IMBRIANI, MANUELA

Titolo

Hybrid Methods for Medical Image Analysis

Relatori

Gina Belmonte

Riassunto

Medical image analysis stands at a critical juncture where pure deep learning approaches, despite achieving impressive accuracy metrics, face fundamental barriers to clinical adoption due to their lack of interpretability and inability to provide spatial reasoning guarantees. This thesis investigates hybrid neuro-symbolic approaches that combine the pattern recognition capabilities of neural networks with the formal spatial reasoning of symbolic methods to address both performance requirements and clinical acceptability constraints in medical imaging applications.

The work is structured around two complementary investigations that together establish both the necessity and viability of hybrid integration.

First, to understand the fundamental limitations that motivate hybrid integration, we conduct systematic benchmark experiments on controlled synthetic spatial reasoning tasks with mathematically precise ground truth generated through VoxLogicA specifications. This evaluation framework isolates specific aspects of spatial understanding—metric computation, topological connectivity, and configuration generalization—through three experiments that reveal fundamental architectural limitations. The first experiment demonstrates that neural networks fail dramatically at pure topological reasoning (maze navigation), with performance degrading as resolution increases. The second experiment shows that hybrid metric-topological tasks can be learned successfully when geometric constraints guide connectivity operations. The third experiment establishes that position-invariant learning is achievable but at the cost of increased variance, indicating incomplete generalization. Synthesis across experiments establishes that networks excel at metric spatial operations but struggle with pure topological reasoning. Hybrid tasks combining metric and topological components can be learned when geometric guidance constrains the search space. Position-invariant learning is possible through training diversity but with reduced consistency across spatial configurations.

The complete experimental framework, including dataset generation pipelines, training protocols, and evaluation tools, is released as open-source software to promote reproducibility and enable community-driven research into hybrid approaches for medical image analysis.

Second, we begin exploring a methodological framework for clinical deployment through a hybrid pipeline for ASPECTS scoring, a critical tool for acute stroke assessment requiring precise identification of early ischemic changes across standardized brain regions. This pipeline integrates nnU-Net, a state-of-the-art self-configuring neural network, with VoxLogicA spatial model checking, which provides formal verification capabilities through declarative logical specifications. The pipeline design illustrates the potential of combining learned pattern recognition with explicit logical specifications of anatomical definitions and spatial relationships to provide both improved performance and essential interpretability for time-critical clinical decisions. As an initial application of the hybrid approach, the architectural framework provides a concrete example of how symbolic spatial reasoning can refine and verify neural network outputs through specifications that encode clinical domain knowledge in human-understandable form.

These findings establish clear design guidelines for medical image analysis systems. Pure neural approaches suffice when tasks involve primarily local pattern recognition. Hybrid neuro-symbolic integration becomes essential when topological correctness must be guaranteed, interpretability is required for clinical accountability, or tasks involve complex spatial relationships that exceed the inherent capabilities of convolutional architectures. The convergent evidence validates that effective medical image analysis benefits from combining the complementary strengths of both paradigms: neural networks contribute learned feature extraction and pattern recognition, while symbolic methods provide formal spatial reasoning guarantees and complete interpretability through declarative specifications.