3iA Côte d'Azur
Interdisciplinary Institute
for Artificial Intelligence

Ph.D. research topic

- Title of the proposed topic:
  "TableMind: Scalable Inference-Time Strategies for Structured Data Reasoning"
- Research axis of the 3IA: Fundamental of AI (axis 1)
- Supervisor:
  Paolo Papotti (EURECOM), papotti@eurecom.fr
- The laboratory and/or research group:
  Data Science Department, EURECOM

Apply by sending an email directly to the supervisor.
The application will include :

- Letter of recommendation of the supervisor indicated above
- Curriculum vitæ.
- Motivation Letter.
- Academic transcripts of a master's degree(s) or equivalent.
- At least, one letter of recommendation.
- Internship report, if possible.

- Description of the topic:

Recent advances in inference-time compute allocation for large language models (LLMs) demonstrate that strategic use of test-time resources can yield performance improvements exceeding traditional baselines - even enabling smaller models to surpass larger counterparts in FLOPs-matched evaluations. This paradigm shift, exemplified by models like OpenAI's GPT-o3 and DeepSeek-R1, underscores the growing importance of specialization in LLMs, particularly for complex reasoning tasks requiring multi-step problem-solving. Building on these insights, our proposal integrates recent inference-time strategies with tailored training methodologies to advance LLM capabilities in *structured data tasks* such as semantic parsing, tabular question answering (QA), tabular fact checking, and data analysis.

Our proposed framework leverages a hybrid approach, combining inference-time compute optimization with reinforcement learning (RL) and supervised fine-tuning (FT) pipelines inspired by breakthroughs in reasoning models. This aligns with our goal of creating models that adapt to structured data challenges while maintaining computational efficiency.

Specifically, our approach will incorporate:

1. Semantic Parsing
   a. Drawing from recent success stories, we will employ RL with dual rewards: (1) an accuracy reward (e.g., query execution correctness for parsing logic) and (2) a format reward enforced via LLM judges to ensure output structure (e.g., valid SQL syntax).

    b. We will generate synthetic datasets, e.g., using intermediate model checkpoints, to iteratively refine parsing accuracy, mimicking the emergent reasoning observed in RL-trained models.
2. Tabular Question Answering/Fact-checking
    a. Inspired by inference-time scaling techniques, such as those in GPT-o3, we will implement search-based methods (e.g., beam search, majority voting) guided by process-based verifiers. This enables robust exploration of tabular data relationships while filtering low-confidence responses.
    b. We will dynamically allocate compute resources based on query complexity—applying lightweight verification for simple facts (e.g., "maximum value in Column X") and intensive multi-step reasoning for complex queries (e.g., causal inference across tables).
3. Data Analysis
    a. We will use FT to instill domain-specific knowledge (e.g., statistical methods) and RL to refine analysis coherence. During inference, iterative refinement cycles (similar to current models' intermediate steps) generate insights while controlling costs via early termination for low-uncertainty tasks.

Our proposal will rigorously evaluate trade-offs between pretraining scale and inference compute, informed by findings that specialized models achieve parity with larger generalist models with training-inference balance. We extend this analysis to structured data tasks, testing hypotheses such as:

- Can RL-driven semantic parsers match the accuracy of 10× larger pretrained models with optimized inference loops?
- Does evolutionary search in tabular QA reduce hallucination rates compared to standard few-shot prompting and CoT?

By benchmarking against FLOPs-matched baselines, we quantify the value of inference-time compute allocation in data-centric settings.
Our research aims to bridge the gap between reasoning model advancements and structured data applications. By integrating RL-driven reasoning emergence, inference-time scaling, and synthetic data generation techniques, we aim to create self-improving agents capable of:

- Allocating resources dynamically based on problem difficulty, akin to human "thinking time."
- Iterative Refinement: Leveraging process-based verifiers to revise outputs, as in recent reward-based solutions.
- Efficient Specialization: Balancing FT for domain expertise and RL for structural integrity, reducing reliance on brute-force pretraining.

This approach not only advances LLM capabilities in data tasks but also informs the broader debate on compute-optimal scaling, offering a blueprint for high-performance, cost-effective AI systems in resource-constrained environments.

References:

1. [2408.03314] Scaling LLM Test-Time Compute Optimally can be More Effective than Scaling Model Parameters
2. [2501.09891] Evolving Deeper LLM Thinking
3. Mutual Reasoning Makes Smaller LLMs Stronger Problem-solvers
4. DeepSeek-V3 Technical Report
5. [2501.19393] s1: Simple test-time scaling