Politecnico di Torino

ScuDo
Scuola di Dottorato ~ Doctoral School
WHAT YOU ARE, TAKES YOU FAR

Doctoral Dissertation
Doctoral Program in Computer and Control Engineering (38^{th}cycle)

Conversational AI and Summarization of Digital Sources

By

Giuseppe Gallipoli

******

Supervisor(s):
Prof. Luca Cagliero, Supervisor
Prof. Paolo Garza, Co-Supervisor
Prof. Elena Baralis, Co-Supervisor

Doctoral Examination Committee:
Prof. Luigi Di Caro, Referee, Università degli Studi di Torino, Italy
Prof. Serge Sharoff, Referee, University of Leeds, United Kingdom
Prof. Silvia Chiusano, Politecnico di Torino, Italy

Politecnico di Torino
2026

Abstract

Conversational AI has rapidly evolved into a central paradigm for enabling natural and effective human–computer interaction. Despite substantial progress, the development of reliable conversational systems remains challenging. The emergence of Large Language Models (LLMs) has significantly advanced the field; however, new challenges related to computational cost, controllability, and adaptation have arisen. This thesis investigates key research problems in Conversational AI, with a primary focus on text summarization. Its main goal is to leverage domain and task characteristics to design tailored approaches that both advance the state of the art and identify the limitations of general-purpose systems. It explores four main research directions: conversation summarization, machine unlearning for text summarization, text style transfer, and the application of LLMs to complex understanding tasks, specifically figurative language understanding and math word problem solving. Across these domains, the thesis introduces novel pipelines, architectural designs, and training objectives. In addition, it proposes new problem formulations and extends existing settings to more realistic scenarios.

In conversation summarization, this work introduces an entailment-enriched representation and a claim-specific pre-summarization strategy to capture cross-document interactions and reduce redundancy. Extensive evaluations show that combining argument representation with redundancy-aware preprocessing leads to consistent improvements across datasets and models, including LLMs.

In machine unlearning, this thesis explores methods to selectively remove undesired information from summarization models without requiring full retraining. After formalizing the novel task of machine unlearning for summarization, it evaluates in-context unlearning as a cost-effective solution and introduces token-level unlearning, which enables fine-grained control over retained and forgotten content. Additionally, a new dataset and a context-aware unlearning framework leveraging spatio-temporal

metadata are proposed. The results demonstrate that these approaches yield effective and scalable solutions for controllable summarization.

In unsupervised text style transfer, the thesis proposes a CycleGAN architecture that enforces cycle consistency at the sequence level to improve content preservation and incorporates classifier-guided generation to strengthen style control. Building on this framework, this work extends style transfer to mixed-style and multi-style settings. Experimental results show that the proposed methods outperform both traditional approaches and modern LLMs, particularly in content preservation.

Finally, the thesis investigates the application of LLMs to two complex understanding tasks. In figurative language understanding, it evaluates multiple adaptation strategies for textual entailment and explanation generation, showing that lightweight fine-tuning enables smaller models to outperform larger proprietary systems. In math word problem solving, it proposes a lightweight fine-tuning framework that augments cross-entropy loss with task-specific alternatives to align the training objective with the structural requirements of mathematical reasoning tasks.

Overall, our contributions highlight the importance of integrating domain knowledge and task characteristics into model design and training. At the same time, this work underscores the need for improved efficiency, controllability, and interpretability, particularly in real-world applications. In conclusion, while LLMs have transformed the landscape of Conversational AI, they do not constitute a universal solution. Their effective deployment depends on careful adaptation to domain- and task-specific challenges and constraints. This thesis contributes to this direction by providing both theoretical and practical insights, paving the way for more robust, scalable, and human-aligned conversational systems.