Integrating Task-Specific Models with Large Language Models for Cross-Lingual Transfer to Low-Resource Languages

Abstract

Cross-lingual transfer (CLT) seeks to enable NLP systems for low-resource languages by leveraging supervision from high-resource languages, offering a scalable path toward supporting the world's nearly 7,000 languages. This dissertation presents a series of hybrid frameworks that combine task-specific fine-tuned models with Large Language Models (LLMs), demonstrating that fine-tuning and in-context learning are complementary rather than competing paradigms for structured prediction.

The thesis advances the state of the art along four complementary directions: **(1)** novel fine-tuning architectures, **(2)** hybrid integration of supervised models with LLMs, **(3)** cross-lingual learning under noisy supervision, and **(4)** new multilingual benchmarks. For fine-tuning, we introduce **ZGUL**, a modular framework that ensembles source-language adapters using interpretable, typology-aware weights to achieve strong zero-resource transfer for sequence labeling. For hybrid prompting, we propose **SSP**, which converts noisy pseudo-labels from a fine-tuned model into high-quality in-context exemplars through optimization-based selection, enabling effective LLM prompting without target-language annotations. Under distant supervision, we develop **PARE**, a passage-level relation extraction model, and **HYDRE**, a hybrid inference-time framework that combines PARE's high-recall candidate generation with LLM reasoning via relation-aware exemplar retrieval, substantially improving monolingual relation extraction on English and cross-lingual relation extraction across three transfer settings. Finally, we introduce **INDICRE4**, a human-verified benchmark spanning four typologically diverse Indic languages.

The dissertation concludes by outlining how these hybrid, largely task-agnostic methods can be extended to open-ended cross-lingual generation and agentic, cost-aware multilingual systems that dynamically combine translation, task-specific models, and LLMs, moving toward scalable language technologies for underserved communities.