

Main research topics @ KEG Natural Language Processing (NLP)


Knowledge Engineering Group

Computer Science, AC, UTC-N



rodica.potolea@cs.utcluj.ro
mihaela.dinsoreanu@cs.utcluj.ro
camelia.lemnaru@cs.utcluj.ro

raluca.portase@cs.utcluj.ro
richard.ardelean@cs.utcluj.ro
vlad.negru@cs.utcluj.ro
alex.lapusan@cs.utcluj.ro



Core research team & areas



Mihaela DINSOREANU



Rodica POTOLEA



Camelia LEMNARU


Natural Language Processing

Data Engineering

NeuroScience



Raluca PORTASE



Alex-Mihai LAPUSAN



Vlad NEGRU



Richard ARDELEAN



- Implementing specific **NLP pipelines** for solving various tasks (classification, sequence tagging, parsing, NLU tasks, etc)
- **Reasoning** with LLMs
- Characterizing and mitigating **biases** in LLMs
- Understanding and mitigating potential **attacks** on LLMs
- Transversal topics
 - *Multilinguality*
 - *Multitask learning*
 - *Few-shot* and *zero-shot* learning
 - *Imbalanced* data distributions (label, language)
 - *Resource-efficiency*





NLP: Understand text in multilingual, multitask contexts

- Real world challenges:
 - Lack of labeled data for target languages (or tasks)
 - Leverage unlabeled data (web scraped texts) in self-supervised solutions
 - Imbalance in the labeled data (language / task)
 - Cross lingual and domain transfer
 - Leverage interaction between tasks via joint learning
- Applications
 - Cooking recipes classification pipeline - industrial use
 - International research grant with AB Electrolux Italy
 - Bank conversations classification pipeline - PoC
 - POC grant: *VOITA - Sistem Informatic Integrat de Voice To Text Analytics*

Ideal



Real

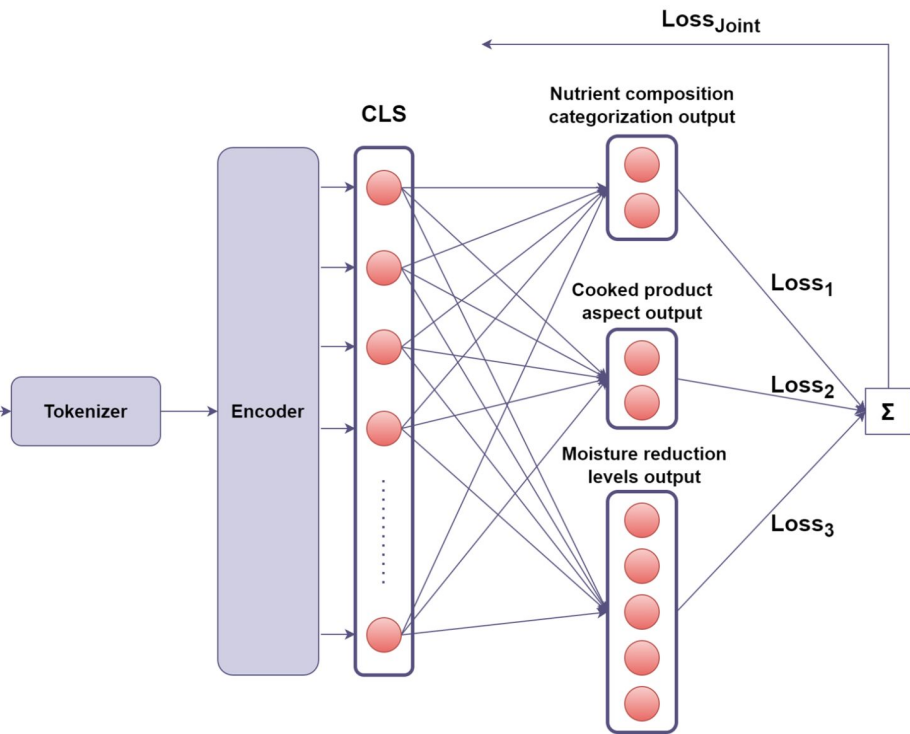




NLP: Understand text in multilingual, multitask contexts

- Recipe Classification Pipeline (AB Electrolux Italy)
 - Multilingual
 - Multitask
 - Adaptation to cooking domain

Chicken parmigiana. Preheat the oven to 200°C. Peel the potatoes and cut into slices. Boil the potatoes for 10 minutes in salted water and drain. In the meantime, toast the bread and crumble it over a deep plate. Grate the Parmesan cheese and add to the bread. In another deep plate, beat the egg with a fork. Place the flour in a third plate. Sprinkle the chicken with salt and pepper, then coat in the flour and egg, then cover completely with the bread mixture ... Bake for 15 minutes until golden brown on medium heat. Turn over.

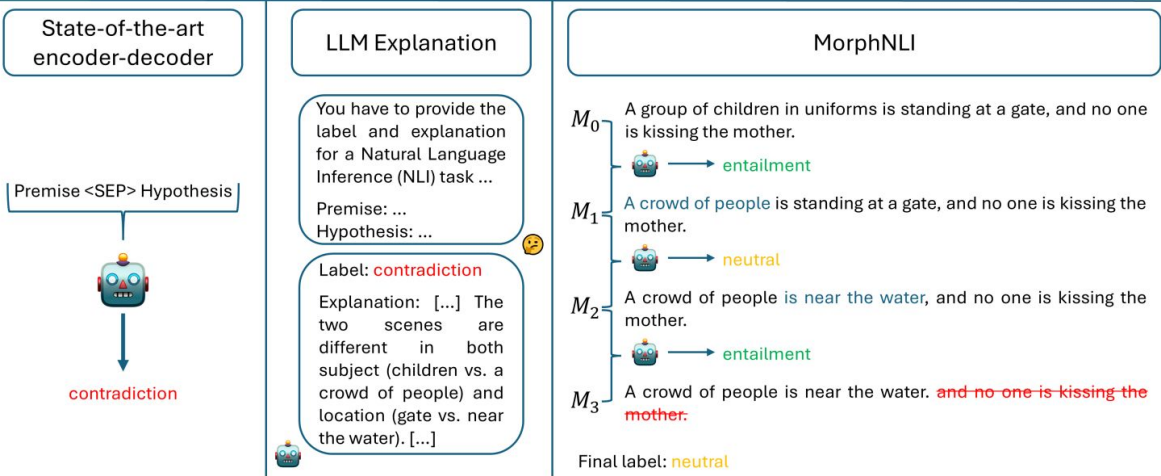


- Logic Reasoning or Pattern Matching?
 - Collaboration with Univ. of Arizona
- Models generally overfit on logic reasoning benchmarks
- *MorphNLI* - novel approach in breaking down logic reasoning problems

Premise: A group of children in uniforms is standing at a gate, and no one is kissing the mother.

Hypothesis: A crowd of people is near the water.

Gold label: **neutral**





Selected recent papers

- Vlad Andrei Negru, Robert Vacareanu, Camelia Lemnaru, Mihai Surdeanu, Rodica Potolea, [MorphNLI: A Stepwise Approach to Natural Language Inference Using Text Morphing](#), to be presented at NAACL 2025
- Negru, Vlad-Andrei; Suciu, Vasile; Lăpușan, Alex-Mihai; Lemnaru, Camelia; Dînșoreanu, Mihaela; Potolea, Rodica [Assessing language models' task and language transfer capabilities for sentiment analysis in dialog data](#), Computer Speech & Language, Volume 89, 2025, 101704, ISSN 0885-2308, 2025
- Toma, Sebastian-Antonio; Lemnaru, Camelia; Negru, Vlad Andrei; Potolea, Rodica, [Asking the Right Questions: Exploiting Hidden Interactions in a Generative Framework for Multilingual, Multitask Classification](#), ICNLS 2024, pp. 251-266, 2024
- Negru, Vlad-Andrei; Vacareanu, Robert; Lemnaru, Camelia; Surdeanu, Mihai; Potolea, Rodica, [From Text to Taste: Advancing Smart Appliances with Multilingual Recipe Interpretation](#), 2024 IEEE 40th International Conference on Data Engineering Workshops (ICDEW), pp 13-20, 2024
- Andrei-Victor Chisca, Andrei-Cristian Rad, Camelia Lemnaru, [Prompting fairness: Learning prompts for debiasing large language models](#), Proceedings of the Fourth Workshop on Language Technology for Equality, Diversity, Inclusion at EACL 2024, pp. 52-62, 2023
- Negru, Vlad-Andrei; Lemnaru, Camelia; Potolea, Rodica, [Multitask, Cross-Lingual Recipe Classification Using Joint Fine-Tuning Mechanisms](#), International Conference on Information Integration and Web Intelligence, pp. 547-552, Springer, 2023





Collaboration inquiry/offer

Open for **research collaborations** and **joint research proposals** on

Natural Language Processing:

- Applications of text classification, parsing, sequence tagging, etc.
- Multimodal, multi-task classification systems
- Integration of LLMs with Knowledge Graphs
- Reasoning with LLMs
- LLM bias and robustness

