Zgłoszenie tematu pracy dyplomowej **:: STUDIA II STOPNIA ::**

rok akademicki 2025/26

Promotor:	**dr hab. Jozef Kapusta, prof. UKEN**
Temat pracy magisterskiej (j. polski, j.angielski):	Retrieval-Augmented Generation for the University Environment *Generowanie wspomagane wyszukiwaniem wiedzy w środowisku uniwersyteckim*
Zakres pracy i oczekiwane rezultaty praktyczne:	Retrieval-Augmented Generation (RAG) is an artificial intelligence technique that combines a large language model (LLM) with an external knowledge base in order to improve the accuracy and relevance of responses. Instead of relying solely on the LLM's training data, RAG retrieves relevant information from a given data source—such as internal documents or a database—and uses it to augment the original query, helping to generate more up-to-date, factual, and context-specific answers. The aim of the thesis is to design, implement, and evaluate a Retrieval-Augmented Generation (RAG) system for a university environment that can answer questions based on guidelines, study regulations, web content, and internal documents, while providing references to the sources used. The work will cover the full pipeline, from document preprocessing and embedding creation to retrieval and answer generation. It will also assess the potential improvement over traditional full-text search using selected evaluation metrics. The expected outcome is a validated prototype accompanied by evidence-based recommendations for safe and effective deployment within the university's information ecosystem.
Aspekt naukowy, problemowy, innowacyjny pracy:	implementation methods natural language, implementation machine learning methods.
*Oprogramowanie, język programowania, środowisko systemowe:	
*Środowisko uruchomieniowe	Local machine with Python and required libraries, optionally Google Colab for GPU usage
Dodatkowe wymagania i uwagi:	English language
*Literatura:	• Bird, S., Klein E., and Loper, E. (2009). Natural Language Processing with Python - Analyzing Text with the Natural Language Toolkit. O'Reilly Media. • Bengfort, B., Ojeda, T., Bilbro, R. (2018). Applied Text Analysis with Python: Enabling Language - Aware Data Products with Machine Learning, O'Reilly Media, 332 p.

	• Lewis, P., Yih, W. T., Pihur, V., & Schaer, S. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS 2020). • Gao, Y., Ma, X., Lin, J., Zhang, X., Yang, Y., Su, Y., Dong, K., Zhu, K., Wen, W., Li, J., Ding, Z., Ma, B., & Lin, H. (2024). Retrieval-Augmented Generation for Large Language Models: A Survey. ACM Computing Surveys, 56(4), 1–45. (Preprint available on arXiv: https://arxiv.org/abs/2312.10997) • Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)

*pola opcjonalne