




MACHINE LEARNING

NOVETTA NLP



A high-level framework for testing, running, and deploying state-of-the-art natural language processing models.

NovettaNLP provides access to cutting-edge natural language processing (NLP) capabilities through a standardized API, empowering users to perform named entity recognition, text classification, and question answering.

By standardizing the input/output data and function calls, regardless of the model used in the backend, NovettaNLP simplifies implementation of NLP algorithms. New, high-performing models are integrated seamlessly into NovettaNLP, with all functionality made available through an easy-to-use, production-ready python package.

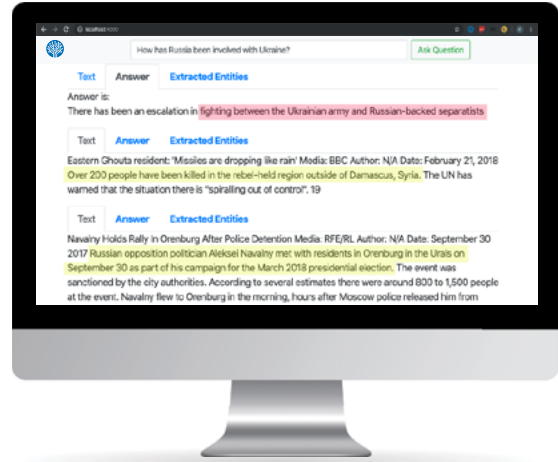
KEY FEATURES

- Pre-trained models for NLP capabilities such as named entity recognition, text classification, and question answering with top-rated accuracy
- Latest pre-trained multilingual language models
- Advanced training and optimization framework for NLP tasks
- Unified API for simple use and integration of complex NLP models
- Incorporation of CI/CD DevSecOps (e.g., Kubernetes, Jenkins)

BENEFITS

- Find persons of interest in large volumes of text
- Evaluate trends in social media
- Retrieve semantically-relevant text through dynamic queries
- Employ high-level modules for efficient state-of-the-art NLP experimentation and R&D
- Lower the barriers to entry in conducting and applying NLP technology

Empowering users to employ state-of-the-art NLP algorithms.



HOW IT WORKS

NovettaNLP is built on top of two open source libraries, Zalando Research's Flair and Hugging Face Transformers, which provide access to pre-trained language models and language-based task models. NovettaNLP provides a clear and easily accessible interface that remains stable over time, even if the underlying libraries change. As these libraries update, the inputs and outputs to NovettaNLP remain consistent - all changes to the underlying code will remain hidden from the user.

NovettaNLP also provides a REST service that stands up endpoints for inference on pre-trained and custom models on a backend API. This allows for simple inference calls and UI implementations after installing the package and building the Docker app. This can be run on CPUs, though NVIDIA GPUs significantly speed up training and inference.

STRUCTURE

