

1. Results in KNIME

One simple example: translate a few words in French. Instead of translating, it seems to be *completing the sentence*.

Workflow response:

#	RowID	Task	Prompt	Response
1	Row0	Classification	Classify the following text into two appropriate categories: The recipe for the chocol...	Classify the following text into two appropriate categories: The recipe for the chocol...
2	Row1	Classification	Classify the following text into two appropriate categories: The stock market experi...	Classify the following text into two appropriate categories: The stock market experi...
3	Row2	Classification	Classify the following text into two appropriate categories: The latest smartphone m...	Classify the following text into two appropriate categories: The latest smartphone m...
4	Row3	Sentiment analysis	Give a sentiment analysis: I just got a promotion at work! I'm thrilled and can't wait t...	Give a sentiment analysis: I just got a promotion at work! I'm thrilled and can't wait t...
5	Row4	Sentiment analysis	Give a sentiment analysis: Unfortunately, my favorite team lost the championship ga...	Give a sentiment analysis: Unfortunately, my favorite team lost the championship ga...
6	Row5	Sentiment analysis	Give a sentiment analysis: The customer service at that restaurant was terrible. The ...	Give a sentiment analysis: The customer service at that restaurant was terrible. The ...
7	Row6	Text summarization	Summarize: The coral reefs of the Great Barrier Reef, located off the coast of Austral...	Summarize: The coral reefs of the Great Barrier Reef, located off the coast of Austra...
8	Row7	Text summarization	Summarize: In recent years, artificial intelligence has made significant advancement...	Summarize: In recent years, artificial intelligence has made significant advancement...
9	Row8	Text summarization	Summarize: The development of renewable energy sources, such as solar and wind ...	Summarize: The development of renewable energy sources, such as solar and wind ...
10	Row9	Text completion	The KNIME Analytics Platform	The KNIME Analytics Platform is an open source data analytics, reporting, and integ...
11	Row...	Text translation	Translate to french: I love you	Translate to french: I love you more than words can say

Standard configuration:

Prompt Hugging Face Hub LLMs

- This workflow shows how to connect to an LLM on Hugging Face Hub and send prompts to it.
- In order to run this workflow you need an access token for Hugging Face Hub. If you don't already have one, create an account on <https://huggingface.co/>, go to <https://huggingface.co/>
- To learn more about the workflow, click the left bar and check the description section.

Hugging Face Models

Note that the free Hugging Face hosted models may not be on par with the ones from OpenAI. You may be able to receive better results by altering the prompts given to the models.

Workflow Diagram:

```
graph LR; A[Credentials Configuration: Enter the Hugging Face Hub access token] --> B[HF Hub Authenticator: Authenticate with Hugging Face API Key]; B --> C[HF Hub LLM Connector: Get a text generating model from the hub];
```

Dialog - 3:90 - HF Hub LLM Connector (Get a text ge...)

Model Parameters

- Top k: 1
- Typical p: 0,95
- Repetition penalty: 1
- Max new tokens: 50
- Temperature: 0,2
- Top-p sampling: 0,15

Buttons: Cancel, Ok

2. Results in HuggingFace

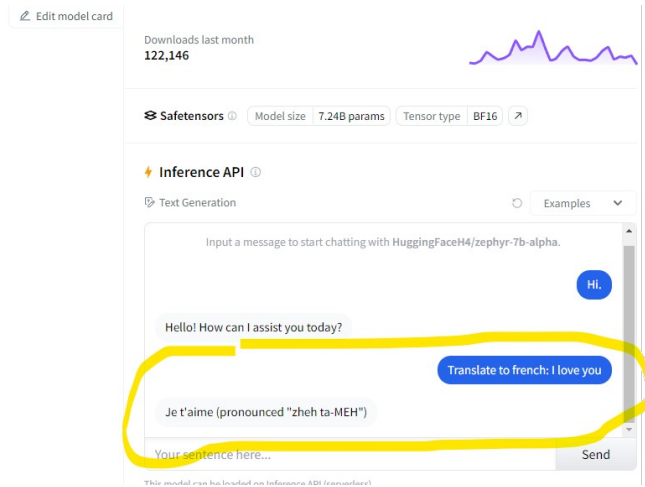
Huggingface response:



Model Card for Zephyr 7B Alpha

Zephyr is a series of language models that are trained to act as helpful assistants. Zephyr-7B- α is the first model in the series, and is a fine-tuned version of [mistralai/Mistral-7B-v0.1](#) that was trained on a mix of publicly available, synthetic datasets using [Direct Preference Optimization \(DPO\)](#). We found that removing the in-built alignment of these datasets boosted performance on [MT Bench](#) and made the model more helpful. However, this means that model is likely to generate problematic text when prompted to do so.

Model description

A screenshot of the HuggingFace model card for Zephyr-7B-alpha. The card shows the model's name, a download count of 122,146, and various technical specifications like model size (7.24B params) and tensor type (BF16). It also features an "Inference API" section with a chat interface. The chat interface shows a user input "Hello! How can I assist you today?" and a model response "Je t'aime (pronounced 'zheh ta-MEH')". A yellow circle highlights the model's response. There is also a button labeled "Translate to french: I love you".