Package: hollr (via r-universe)

November 15, 2024

Type Package

Title Unified Framework for Chat Completion and Text Annotation with Local and OpenAI Language Models

Version 1.0.0

Maintainer Jason Timm <JaTimm@salud.unm.edu>

Description Enables chat completion and text annotation with local and OpenAI <https://openai.com/> language models, supporting batch processing, multiple annotators, and consistent output formats.

License MIT + file LICENSE

Encoding UTF-8

LazyData false

Depends R (>= 3.5)

Imports data.table, httr, jsonlite, pbapply, reticulate

RoxygenNote 7.3.1

URL https://github.com/jaytimm/hollr, https://jaytimm.github.io/hollr/

BugReports https://github.com/jaytimm/hollr/issues

Config/pak/sysreqs libpng-dev libssl-dev python3

Repository https://jaytimm.r-universe.dev

RemoteUrl https://github.com/jaytimm/hollr

RemoteRef HEAD

RemoteSha d10060bc169ebba9aabb80551214ab61314f1379

Contents

hollr *LLM Completions*

Description

This function generates text using either the OpenAI API or a local model.

Usage

```
hollr(
  id,
  user_message = "",
  annotators = 1,
  model,
  temperature = 1,
  top_p = 1,
  max_tokens = NULL,
  max_length = 1024,
  max_new_tokens = NULL,
  system_message = "",
  force_json = TRUE,
  flatten_json = TRUE,
  max_attempts = 10,
  openai_api_key = Sys.getenv("OPENAI_API_KEY"),
  openai_organization = NULL,
  cores = 1,
  batch_size = 1,
  extract_json = TRUE
)
```

Arguments

`id` A unique identifier for the request.

`user_message` The message provided by the user.

`annotators` The number of annotators (default is 1).

`model` The name of the model to use.

`temperature` The temperature for the model's output (default is 1).

`top_p` The top-p sampling value (default is 1).

`max_tokens` The maximum number of tokens to generate (default is NULL).

`max_length` The maximum length of the input prompt (default is 1024, for local models only).

`max_new_tokens` The maximum number of new tokens to generate (default is NULL, for local models only).

`system_message` The message provided by the system (default is ”).

`force_json` A logical indicating whether the output should be JSON (default is TRUE).

`flatten_json` A logical indicating whether the output should be converted to DF (default is TRUE).

`max_attempts` The maximum number of attempts to make for generating valid output (default is 10).

`openai_api_key` The API key for the OpenAI API (default is retrieved from environment variables).

`openai_organization`
The organization ID for the OpenAI API (default is NULL).

`cores` The number of cores to use for parallel processing (default is 1).

`batch_size` The number of batch_size to process (default is 1, only for local models).

`extract_json` A logical indicating whether to extract and clean JSON strings from the response (default is TRUE).

Value

A data.table containing the generated text and metadata.

Examples

```
## Not run:
hollr(id = "example_id",
user_message = "What is the capital of France?",
model = "gpt-3.5-turbo",
openai_api_key = "your_api_key")

## End(Not run)
```

`hollr_local_batches` *Batch Processing for Local LLM*

Description

This function generates text in batches using a local model.

Usage

```
hollr_local_batches(
  id,
  user_message = "",
  annotators = 1,
  model,
  temperature = 1,
  top_p = 1,
  max_new_tokens = 100,
  max_length = NULL,
  system_message = "",
  batch_size = 10,
  extract_json = TRUE
)
```

Arguments

`id` A unique identifier for the request.

`user_message` The message provided by the user.

`annotators` The number of annotators (default is 1).

`model` The name of the model to use.

`temperature` The temperature for the model's output (default is 1).

`top_p` The top-p sampling value (default is 1).

`max_new_tokens` The maximum number of new tokens to generate (default is 100).

`max_length` The maximum length of the input prompt (default is NULL).

`system_message` The message provided by the system (default is ”).

`batch_size` The number of messages to process in each batch (default is 10).

`extract_json` A logical indicating whether to extract and clean JSON strings from the response (default is TRUE).

Value

A list containing the generated responses for each batch.

Examples

```
## Not run:
hollr_local_batches(id = "example_id",
user_message = "Hello, how are you?",
model = "local_model")

## End(Not run)
```

Index