

Developer's Guide: Setting up & Running Fairsense-AI on Local

Fairsense-AI is built keeping in mind the sustainability practices for efficient LLM inference locally with the support for both CPU & GPU based systems.

Use of Local LLMs

Running large language models (LLMs) locally provides significant advantages, whether for research, experimentation, or building advanced applications. However, the process of configuring the necessary environment and setting up LLMs on local systems is often complex and resource intensive.

Why Use Local LLMs?

- Privacy: Keep your data secure by running models locally.
- Speed: No cloud dependencies mean faster responses.
- Customization: Fine-tune models for your specific needs.
- Cost Efficiency: Save on cloud service fees.
- Offline Access: Use models without an internet connection.
- Hardware Optimization: Efficiently use your local system resources.

Steps to Set up Development Environment

1. Prerequisites

Before starting, ensure you have the following:

- A computer with disk space and RAM for model files.
- Python 3.9 or later installed.

2. Setting of Inference For CPU

For CPU systems, we make use of [Ollama](#), a tool wrapped around [llama.cpp](#), which simplifies using LLMs locally with an easy-to-use interface.

A. Download Ollama

Ollama works on macOS, Windows, and Linux. Choose one of these options to download it:

- Visit the Ollama GitHub page for installation instructions:

<https://github.com/ollama>

- Go to Ollama's official website and download the installer for macOS or Windows:

<https://ollama.com>

B. Install Ollama

Use Python's package manager, pip, to install Ollama. Open a terminal (or command prompt) and run:

```
pip install ollama
```


A. Install Torch with Cuda Support

```
pip install torch torchvision torchaudio --index-url  
https://download.pytorch.org/whl/cu117
```

2. Installation of the requirements

All the necessary python libraries can be installed by running the following command

```
pip install fair-sense-ai
```

Fairsense-AI makes use of **local inference** with LLMs, tailored for bias analysis. It efficiently processes text/image, analyzes text/image provided by the user for bias and offers recommendations on mitigating bias for more fair and inclusive content.

Local Inference in Fairsense-AI

The tool identifies the system specification and defines the runtime based on the system. The script provides a complete platform for analyzing biases in textual and image data, generating AI governance insights, and visualizing AI safety risks. It uses advanced AI models for analysis and Gradio for an interactive user interface.

The **FairsenseGPURuntime** is designed to enable GPU-based execution of text generation tasks using a Hugging Face language model. It leverages the computational power of GPUs for faster and more efficient inference.

Running the file in GPU system

Perform the following steps to run the fairsense interface

1. Navigate to the fairsenseai folder
2. Set the PYTHONPATH using the following command after navigating to the folder
`export PYTHONPATH=.`

```
(venv) mchettiar@gpu053:~/documents$ cd fairsenseai/  
(venv) mchettiar@gpu053:~/documents/fairsenseai$ export PYTHONPATH=.  
(venv) mchettiar@gpu053:~/documents/fairsenseai$
```


3. Run the fairsenseai/app.py file

```
python fairsenseai/app.py
```

```
(venv) mchettiar@gpu053:~/documents/fairsenseai$ cd .  
(venv) mchettiar@gpu053:~/documents$ python fairsenseai/app.py  
Starting FairsenseRuntime with file system access.  
Loading models...  
Device set to use cuda:0  
Models loaded successfully.  
The new embeddings will be initialized from a multivariate normal distribution that has old embeddings' mean and covariance. As described in this article: https://nlp.stanford.edu/~johnhew/vocab-expansion.html. To disable this, use `mean_resizing=False`  
* Running on local URL: http://127.0.0.1:7860  
* Running on public URL: https://1f1fd83a95df578bf3.gradio.live  
  
This share link expires in 72 hours. For free permanent hosting and GPU upgrades, run `gradio deploy` from the terminal in the working directory to deploy to Hugging Face Spaces (https://huggingface.co/spaces)
```

It can be observed that the models are loaded into the GPU with the use of cuda.

Executing the file launches gradio interface for Fairsense-AI which can be accessed using the public URL provided



The **FairsenseCPURuntime** is designed for executing text generation tasks on a CPU. It interfaces with Ollama to generate responses efficiently without requiring a GPU.

Running the file in CPU system

Perform the following steps to run the fairsense interface

1. Navigate to the fairsenseai folder
2. Set the PYTHONPATH using the following command after navigating to the folder
`export PYTHONPATH=.`

```
(test_env) mukundchettiar@VI-C02G9CFEML7H fair_sense_ai % cd fairsenseai
(test_env) mukundchettiar@VI-C02G9CFEML7H fairsenseai % export PYTHONPATH=.
(test_env) mukundchettiar@VI-C02G9CFEML7H fairsenseai %
```


3. Return back to the project directory and run the fairsenseai/app.py file
`python fairsenseai/app.py`

```
(test_env) mukundchettiar@VI-C02G9CFEML7H fairsenseai % cd ..
(test_env) mukundchettiar@VI-C02G9CFEML7H fair_sense_ai % python fairsenseai/app.py
Starting FairsenseRuntime with file system access.
Loading models...
Device set to use cpu
Models loaded successfully.
* Running on local URL: http://127.0.0.1:7860
* Running on public URL: https://761648ebd260d810c4.gradio.live

This share link expires in 72 hours. For free permanent hosting and GPU upgrades, run `gradio deploy` from the terminal in the working directory to
deploy to Hugging Face Spaces (https://huggingface.co/spaces)
```

It can be observed that the models are loaded into the CPU now.

Executing the file launches gradio interface for Fairsense-AI which can be accessed using the public URL provided



Summary

Fairsense-AI makes use of local LLMs with the ability to support both CPU and GPU devices and aids in efficient analysis of bias for text and images. It processes the data (text/image), uses local LLM to analyze for bias and highlights any bias present it through optimized inference for local runtime. Prepared By

- **Name:** Shaina Raza, PhD shaina.raza@vectorinstitute.ai
 Marcelo Lotif marcelo.lotif@vectorinstitute.ai
 Tahniat Khan tahniat.khan@vectorinstitute.ai
 Mukund Sayeeganesh Chettiar mukund.chettiar@vectorinstitute.ai
 Matin Yousefabadi matin.yousefabadi@vectorinstitute.ai
- **Affiliation:** Vector Institute for Artificial Intelligence

This tutorial is presented for practical use based on existing methods and open resources (llama.cpp, ollama) for efficient evaluation and optimization of LLMs. It incorporates

strategies for reducing carbon emissions and computational costs while highlighting use cases for secure local inference. Special thanks to the llama.cpp, ollama contributors whose work inspired this adaptation.