Bring your Python Machine Learning projects to production scale with Weaviate

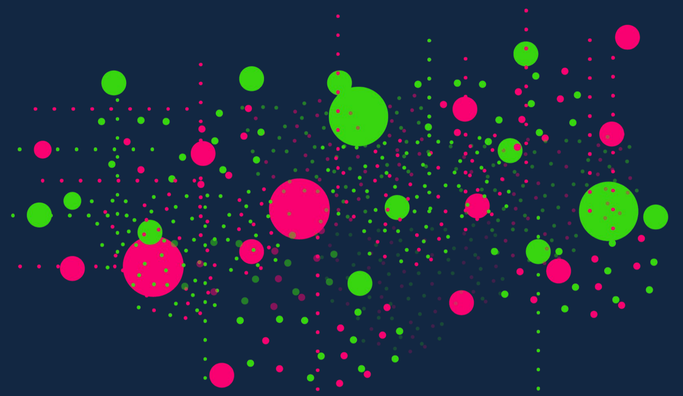

Weaviate is an open-source vector database enabling similarity search

Store and search through billions of data objects stored as vectors, using your own NLP/ML model or any out-of-the-box Huggingface🤗 transformer model.

Introduction video & demo

Weaviate
Open Source

- GitHub
- Documentation
- Python Client

Do you want to scale your NLP project?

Machine Learning projects in Python are often hard to scale. If you want to apply your own machine learning model in production and with large datasets at scale, then Weaviate is a great solution. **Weaviate supports custom ML and NLP models** (for example a model using PyTorch, TensorFlow or Keras).

Python Client

There is a full **Python client** available for Weaviate, which enables to manage your ML data, with access to all **RESTful** and **GraphQL API** endpoints.

Data storage and retrieval with Weaviate

With Weaviate, you can scale your Machine Learning projects to production from your Python scripts. You can stay focused on your Machine Learning project, while Weaviate takes care of efficient (vector) data storage and retrieval!

Get started

- Use out-of-the-box ML models
- Use transformer NLP & image models
- Use Question Answering models
- Use your own **custom ML model**

tds — Get started with the Python client" on Towards Data Science

Get started in Google Colab

Do you want to store and search through your (ML) data?

Weaviate stores data as **vectors** computed by Machine Learning. You can use **your own ML model** or one of the **out-of-the-box models** for

- Data storage in a high-dimensional vector space
- Text and image vectorization
- Data similarity search based on natural language
- Question Answering
- Zero-shot, *k*nn and other types of classification

Weaviate:

- builds upon the **ANN** (approximate nearest neighbor) algorithm **HSNW** to store and retrieve vectors effectively (within milliseconds).
- is a next-generation database and search engine, because it **combines similarity (vector) search with traditional (scalar) search**. This makes Weaviate unique compared to similar solutions.
- supports various structured and **unstructured data** types: from numbers and text to images and more to come.
- Has a Python client that is easy to integrate and easy to use

More information

For questions or more information:
- Visit our website semi.technology
- Join the community on Slack
- Send me an email: laura@semi.technology

SeMI Technologies