




Bring Your Own Idea: Music Dataset

Generative AI is the newest sensation in a list of interesting applications of machine learning. Interestingly, however, mainly the textual domain with LLMs and the visual domain with Stable Diffusion et al. have had a major impact. Why does the audio domain not have the same groundbreaking results? As it turns out, one significant hurdle for machine learning for music is the lack of large datasets. To this end we have created the largest publicly available music dataset to date: <https://huggingface.co/datasets/DISCOX/DISCO-10M>

While the dataset is currently the most-downloaded music dataset on Hugging Face, we ourselves have yet to utilize it. This is where you come in: Do you have an interesting application or idea that can take advantage of a large music dataset?

We are looking for a student that has an interesting idea what could be done with this dataset, for example questions that could be analyzed and answered, specific new insights that could be gathered by analyzing song embeddings, or even more out-of-the-box ideas. If you want to work with this data and have a cool idea, let's talk!



Requirements: Ability to work independently and determined to obtain results, creative thinking, knowledge with Machine Learning and Python.

We will have weekly meetings to address questions, discuss progress and think about future ideas.

Contact

- Luca Lanzendörfer : lanzendoerfer@ethz.ch, ETZ G93
- Florian Grötschla : fgroetschla@ethz.ch, ETZ G63