



High-fidelity performance metrics for generative models in PyTorch

Anton Obukhov <anton.obukhov@vision.ee.ethz.ch>

Computer Vision Lab, ETH Zurich

How to quantify the diversity of a set of images?

Generative models are trained to produce samples of high variety. **Inception Score (IS)** is large for diverse sample sets, where diversity is measured as statistics of the Inception model activations of all images.

$$\text{IS} \left(\begin{matrix} \text{cat} & \text{cat} \\ \text{cat} & \text{cat} \end{matrix} \right) > \text{IS} \left(\begin{matrix} \text{cat} & \text{cat} \\ \text{cat} & \text{cat} \end{matrix} \right), \quad \text{FID} \left(\begin{matrix} \text{cat} & \text{cat} & \text{cat} & \text{cat} \\ \text{cat} & \text{cat} & \text{cat} & \text{cat} \end{matrix} \right) < \text{FID} \left(\begin{matrix} \text{cat} & \text{cat} & \text{cat} & \text{cat} \\ \text{cat} & \text{cat} & \text{cat} & \text{cat} \end{matrix} \right)$$

Fréchet Inception Distance (FID) measures the degree of similarity between two datasets (e.g., real and generated) and assigns lower values to a pair of datasets with higher diversity from similar distributions.

Kernel Inception Distance (KID) is similar to FID but allows for easier comparison across a wider range of evaluation protocols.

Why another implementation?

Reference implementations in TF rely on framework-specific functions, take tens of minutes even on small datasets, pull in heavy dependencies.

Reimplementations in PyTorch exist but do not follow all canonical computations closely, resulting in large discrepancies.

As a result, researchers and adopters wishing to compare against state-of-the-art generative modeling are required to stick to reference implementations and cope with inconveniences.

torch-fidelity provides fast, precise, and accessible metrics in PyTorch.

```
> pip install torch-fidelity
```

Features of torch-fidelity

- Epsilon-exact precision as compared to reference implementations;
- Provides API and a standalone console app for usage in scripts;
- Caching layer avoids recomputations across runs and metrics;
- Suitable for frequent evaluations in the training loop;
- Extensibility beyond image modality;
- Interchangeable feature extractor;
- Supports IS, FID, and KID;
- Few external dependencies;
- Weekly unit testing;
- Apache 2 license.

torch-fidelity
is looking for
cloud CI plans
for open source
with GPU

References



<https://github.com/toshas/torch-fidelity>



<https://twitter.com/AntonObukhov1>



<https://obukhov.ai>



ETH zürich

