



Audio/Speech Quality Evaluation Toolkit

While audio researchers mainly rely on human subjective evaluations to understand the performance of their systems, many objective evaluations exist [1]. However, there exists no single easy-to-use Python codebase with these objective metrics (note: there are some codebases that cover fundamental metrics, such as torchmetrics, but we want a package that contains more advanced metrics such as those mentioned in the reference). These metrics are either dispersed across unmaintained repositories, only exist in Matlab or C, or are proprietary. We want to change that!

The goal of this project is to create a codebase containing many, if not all, objective audio evaluation metrics written in Python, and have it follow the KISS principle: Easy to setup, easy to use. To this end, the ideal candidate will be eager to understand existing codebases, implement metrics from paper descriptions, and understand their fundamental underpinnings.

Requirements: Strong programming skills. Ideally with a background in signal processing.

Weekly meetings will be scheduled to address questions, discuss progress, and brainstorm future ideas.

Contact

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

References

- [1] M. Torcoli, T. Kastner, and J. Herre, “Objective measures of perceptual audio quality reviewed: An evaluation of their application domain dependence,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1530–1541, 2021.