

Training: Python Academy
Fast NumPy Processing with Cython

TRAINING GOALS:

The course targets medium level to experienced Python programmers who want to break through the limits of Python performance. A basic understanding of the C language is helpful but not required. Basic understanding of Cython and NumPy as provided in the courses Fast Code with the Cython Compiler and Numerical Calculations with NumPy is necessary.

NumPy and SciPy come with a broad set of high-level functionality that allows to express complex computational algorithms concisely and efficiently. However, in many cases, sequential operations on NumPy arrays introduce a considerable overhead. This can happen when arrays are unnecessarily being copied during an operation that does not work in-place, but also due to lacking CPU cache locality when large arrays are being traversed multiple times in a row. In both cases, Cython can provide a substantial speed-up by expressing algorithms more efficiently.

The main features that make Cython so attractive for NumPy users are its ability to access and process the arrays directly at the C level, and the native support for parallel loops based on the OpenMP compiler infrastructure. To work efficiently with arrays and other memory buffers, Cython has native syntax support for the Python buffer protocol, which allows C extensions (like NumPy or image processing libraries) to grant foreign code direct access to their internal data buffers.

CONSPECT:

- Use of Python's buffer interface from Cython code.
 - Directly accessing data buffers of other Python extensions.
 - Retrieving meta data about the buffer layout.
 - Setting up efficient memory views on external buffers.
- Implementing fast Cython loops over NumPy arrays.
 - Looping over NumPy exported buffers.
 - Implementing a simple image processing algorithm
 - Using "fused types" (simple templating) to implement an algorithm once and run it efficiently on different C data types.
- Use of parallel loops to make use of multiple processing cores.
 - Building modules with OpenMP.
 - Processing data in parallel.
 - Speeding up an existing loop using OpenMP threads.

REQUIREMENTS:

- Knowledge of Python language.
- Basic understanding of NumPy and Cython language.

Difficulty level



CERTIFICATE:

The participants will obtain certificates signed by Python Academy.

TRAINER:

Authorized Python Academy Trainer.