Python meets C/C++ Exercises

Before you start:
- create a suitable directory for this exercise
- download the zipped material from `http://www.physik.uzh.ch/~python/python/lecture_c++/`

Exercise 1: Root Finder

Use the Newton method ($x_n = x_{n-1} - f(x_{n-1})/f'(x_{n-1})$ with $f : \mathbb{R} \to \mathbb{R}$) to find a root (or more generic for any value $a \in \mathbb{V}_f$ in the value domain of f a value of x such that $f(x) = a$).

Design the algorithm first in pure Python and then in cython. Compare the performance.

Tip: You can start with a simple version and then improve the functionality, *e.g.* defining the function or the precision of the derivative calculation at runtime.

Exercise 2: Differential Equation Solver

Implement the functionality of a differential equation solver taking as arguments the starting value of the independent, $t = t_0$, and the dependent variable, $x = x_0$, as well as the actual equation as a function $x' = f(x, t)$.
Create first the pure Python and then the cython implementation. You can take your favourite solving algorithm (Euler ($x(t+\Delta t) = x(t) + x'(t)\Delta t$), Runge-Kutta, . . . ; `en.wikipedia.org/wiki/Numerical_methods_for_ordinary_differential_equations`).

Tip: Create it first for a single equation and afterwards expand it to an array of coupled equations also allowing for higher-orders. The output of the solver could be an array of $x(t)$ at the different time steps or just the value of x at the final time.
You can use as a test bed for the single equation $dT/dt + C(T - T_0) = 0$ corresponding to the temperature change of an object in an environment with ambient temperature T_0. For a set of functions you can use the differential equation of a damped oscillator $d^2x/dt^2 + \beta dx/dt + \omega^2 x = 0$ (*i.e.* with the second trivial equation $x' = dx/dt$).

Exercise 3: Wrap a Class in cython

The folder `wrapping` contains a C++ class `Vector` and its derived class `UnitVector`. Have a look at the corresponding files. Some of the features are commented out at the moment.
Write a wrapper in cython for the `Vector` class (*cf.* lecture). Compile the code to a shared object and test it in Python.

Include now also the `coordinates` function using `std::vector`.

Include the addition operator.

Also try to include all the other operators.

Finally, write also a wrapper for the `UnitVector` class.

Write some test functions in Python and measure the CPU time. If you are still not bored about wrapping C++ classes, write the same test functions in C++ and compare the used CPU time.

Exercise 4: Wrap a Class with SWIG

The folder `wrapping` contains a C++ class `Vector` and its derived class `UnitVector`. Have a look at the corresponding files. Some of the features are commented out at the moment.
Write a SWIG interface file for the `Vector` class (*cf.* lecture). Compile a shared object and test it in Python.

Include sequentially the excluded functions and operators, modify the interface file and recompile the shared object.

Write some test functions in Python to measure the CPU time, compare it to – if done – the cython wrapping and the same test functions in pure C++.

Exercise 5: Exception Handling in cython

This exercise has the goal to play a bit with the exception behaviour in cython.

You can find in the folder `exception` for this part the file `except1.pyx` a function throwing an error if its input is zero. Otherwise it returns the input value. Compile this cython code running the corresponding setup script `setup_except1.py` as we have learnt in the lecture. (Be aware that you might need to change the shared object name due to additional extensions!) Afterwards you can use the `runExcept1.py` script to run the function in an except clause with `$ python runExcept1.py <input number>`. What happens if the input is 0, 1 or 2? And why? Add the `except *` statement in the definition of `pythonError` in `except1.pyx`. Recompile the file and run it again. What happens now?
Now change it to `except 1` and recompile it. What is now the outcome and why?

Now we want also study the handling of C++ exceptions. You can find a bunch of functions raising C++ exceptions in `except.h`. The file `except2.pyx` wraps them. So compile this cython code with the `setup_except2.py` file and import the created module in a Python session. Try out what happens if you call the different wrapper functions. How can you change this annoying behaviour (seen in the lecture)? How can you force the wrapper functions to throw a specific Python error? Try it out!

In the last part of this exercise we want to discuss a specific difference between Python and C. You can find in the corresponding exercise folder the file `div.c`. Compile this program with `$ g++ div.c -o div`. This program takes two numbers as arguments and calculates the ratio

(*e.g.* `$ ./div 3 4`). What happens if you make a division by zero? Now start Python and perform a division by zero. What happens?

Create a function in cython that calculates the ratio of two numbers. What is the output in case of a division by zero here?

Cython has a way to force C behaviour on this issue. The `cython` module has a decorator `@cython.cdivision(bool)`, which can be called before the function. Try it out!

Exercise 6: Playing with STL in cython

Write an algorithm that finds all anagrams (*i.e.* words with the same set of letters) based on a Python dictionary (Python) or a `std::map` in cython. Use the `/usr/share/dict/ words` file as input. What is the set of letter allowing to build to largest number of words? (If you struggle with the algorithm, have a look at `STL/anagram.py`!)

Compare the Python and the cython implementation in terms of spent CPU time. Why does the observed result occur? How can you improve the speed of the cython version?

Exercise 7: NumPy *vs.* cython

One topic which was not discussed in the lecture, is the use of arrays in cython. You can find in the back-up slides some information about this aspect. Typically, this is treated with typed memoryviews allowing an efficient memory access. The declaration of arrays is done as, *e.g.* for a two-dimensional double array `double[:,::1]`. The `:1` part indicates the fastest changing index, *i.e.* in `double[:,::1]` corresponds to a C-contiguous and `double[::1,:]` to a F-contiguous array. `double[:,:]` is compatible with C- and F-contiguous arrays.

Write a function that returns the element-wise squared version of the input array. Compare it to the NumPy and the Python calculation in terms of CPU time.

Write a function that takes two two-dimensional arrays and calculates the matrix product. Again compare it to the NumPy matrix calculation in terms of speed.

If you still are interested in arrays in cython, play around with the different alignments and test the impact on the speed of your code.

Exercise 8: Fortran

For those of you who still do it the Old School way: The folder `fortran` contains some files including a README that can be used as a starting point for interfacing Python code with Fortran77 and Fortran90 code. Have fun and play around!

Exercise 9: Auto-Regressive Model

We want to simulate an auto-regressive time series X_t of the form

$$X_t = c + \alpha X_{t-1} + \varepsilon_t$$

where ε_t is a uniform random number between -1 and 1, $X_0 = 0$, c is a constant and $|\alpha| < 1$.

Create an algorithm returning the n-th value of the series. Use pure python as well as cython. How large is the gain in CPU time?

Tip: You can import into cython C's standard library random number generator for positive integers and the maximum value wiht `from libc.stdlib cimport rand, RAND_MAX`.