

Bioinformatics M.Sc Module: BIOL60201

Project 1: A Python Proteome Analysis Pipeline

This project is designed to be completed by a group of 4 people, but 3 should still be able to manage it comfortably with some help. The project is divided into a series of Tasks, each of which will require some development of a simple Python program or will require some execution of the code supplied for groups smaller than 4. The Tasks will lead in to, or follow on from, another task, and hence it is essential that you think carefully about the formats your programs will be able to cope with, and that you draw up sensible specifications *before* you all start coding.

Overview

You will be given a genome sequence, from a bacterial organism, which has just been sequenced. An experimental group working on this organism wishes to carry out some proteomics experiments to identify proteins from the microbe. However, they need a set of proteins in a FASTA formatted database in order to search against with their mass spectra. They then want to know which is the best digestive enzyme to use, so that they get the best chance of identifying as many proteins as possible. In other words, they want to find the enzyme which produces the lowest number (on average) of peptides per protein, or indeed can unambiguously identify a protein. If a peptide mass is rare, or ideally totally unique (given the accuracy of the instrument), it will effectively be more diagnostic of the parent protein. However, their mass spectrometer can only measure peptides which have a mass (or strictly speaking, mass-to-charge ratio) between 1000 and 1500 Da.

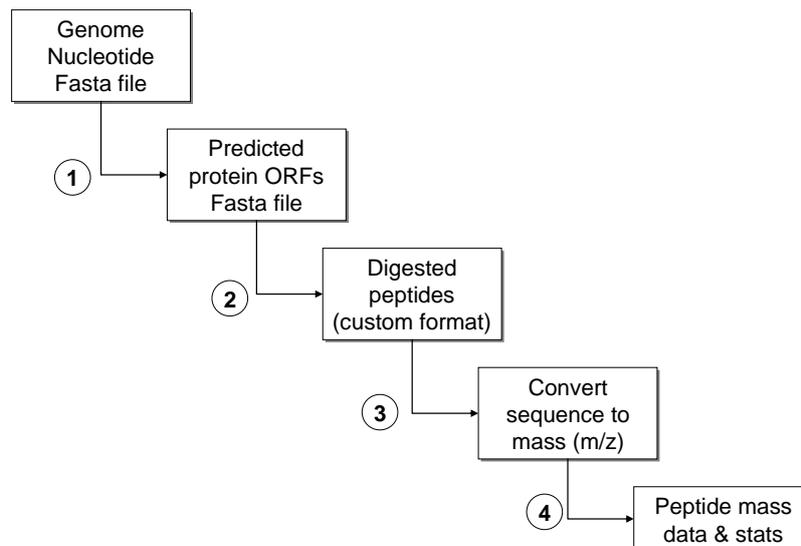


Figure 1. Flow chart of proteome project

The tasks are described in more detail below.

Task 1 – An ORF finder

You will need to write a Python program which can read in a Fasta format file containing nucleotide sequence (DNA sequence, containing ACGT) and find all the possible open reading frames (known as ORFs, beginning with ATG¹, and ending with a stop codon, TGA, TAA, TAG). You can also assume that

¹ Note: >90% of codons in bacterial genes start with ATG but around 10% don't. However, we'll assume they all do for the purposes of this exercise.

the longest ORF will be used, and not have to worry about any internal ones. You will have to worry about all 6 open reading frames though, and predict ORFs in all 6.

Your program should accept standard Fasta format as input, of the form:

```
>Genome sequence
AATTCAGTTACTTATCCCCTTATTGGCAGTATTAACGCATAGGACTGCG
ATAGACTTTACTAAGCTCAGTATCATTCCCTTATTGGCATAGTTACGCAT
CCACCAATAGACTTTACTATTCCCTTCATGGCATATTTACGCAGTAGGAC
ATAGACT
```

And should output protein ORFs with the following fasta format:

```
>orf_name frame length start
MAKSKDFPVKADFAAHVAIQSEFAHGV
>orf_name frame length start
MHILDECCA KSKDFPDFAAHVAYIQSEFAA
LLILPQWNSDAAHGV
```

Orf_name should be a unique name for the ORF, which might include the name of the genome/organism, frame in which the ORF is found, and some number (ie. if the organism is called *I.claudius* you could call your ORFs something CLAUD_F1_0017 (for the 17th ORF in frame 1 in *I.claudius*).

Output should be written to standard output.

Please note: like most genomes, it contains a few ambiguously called bases. This means that they are given the letter 'n' instead of A, C, G or T. Obviously, these can not be formally translated in an amino acid in most cases. I suggest you ignore all ORFs that contain a codon with a 'N' in it.

You should also think carefully about the minimum size of an ORF. I suggest you add an option which controls this, and I would use 50 amino acids as the minimum size, or perhaps even 100 amino acids (= 300 nucleotides).

Bonus tasks:

Error checking: spot improper formats, unknown amino acids, and correctly deal with "gaps" (there normally won't be any of course).

Add some command line options in, such as the ability to write to a specified output file instead of standard out with `-o filename`, or restriction to just one frame with `-f frame_number`.

Deal with more than one sequence, and more than one file supplied on the command line as arguments

Only taking longest ORF when the ones in different frames overlap – this is actually very tricky.

Task 2 – A protein digester

This python script should read in the output from task 1. ie. it should be able to cope with one or more protein sequences, and then digest them into smaller peptide sequences, reporting these back one

per line. The script should be able to cope with a user-defined choice of digesting enzyme, from the following set:

Trypsin: cuts at Lysine (Lys, K) or Arginine (Arg, R) unless the next amino acid is Proline (Pro, P).

Endoproteinase Lys-C: cuts at Lysine (Lys, K) unless the next amino acid is Proline (Pro, P).

Endoproteinase Arg-C: cuts at Arginine (Arg, R) unless the next amino acid is Proline (Pro, P).

V8 proteinase (Glu-C): cuts at Glutamic acid (Glu, E) unless the next amino acid is Proline (Pro, P)

The program should read in Fasta formatted protein sequence e.g.:

```
>protein1
ALTAMCMNVWEITYHKGSDVNRRAASFAQPPPQPPPPLLAIKPASDASD
```

And report the individual peptides after digestion (eg. For trypsin)

```
>protein1 peptide 1
ALTAMCMNVWEITYHK
>protein1 peptide 2
GSDVNR
>protein1 peptide 3
R
>protein1 peptide 4
ASFAQPPPQPPPPLLAIKPASDASD
```

Bonus tasks

Adapt your code so the program will also consider missed cleavages. These are instances when the enzyme misses out on a putative cutsite. You can allow up to 1 missed cleavage, for example, so that

```
AHLKNMILQEWMAALI
```

Would produce AHLKNMILQEWR and NMILQEWMAALI on top of the standard peptides AHLKNMILQEWR MAALI (ie. its inclusive, so that you calculate up to 1 missed cleavage). This should be supported by a command line option to allow the user to select the number of missed cleavages.

Add some error trapping so the program spots "unusual" (B,O,U,J,Z) or unknown amino acids (X)

Task 3 – A mass analyser

In this task you should read in a standard Fasta format file and convert the sequence information into masses, or strictly speaking, mass-to-charge values. You will find the masses of the standard amino acids to 4 decimal places here (http://www.matrixscience.com/help/aa_help.html). The program should cope with both monoisotopic and average masses, and should calculate the mass-to-charge value as if they were all ions of +1 charge (Hence, you need to add 1 to all values, for 1 proton). You also need to add the mass of water to the mass of the summed amino acid values, as these are for a single residue, so you need to add the masses of the terminating groups (e.g. H at the N-terminus and OH at the C-terminus). These values are: 18.0106 for monoisotopic masses, and 18.0153 for average masses. Your output should be something of the form.

Prot_name	peptide	mass-to-charge	z	p	sequence
ABC1_YEAST	1	1234.5678	1	0	ALTAMCMNVWEITYHK
ABC1_YEAST	2	894.3019	1	0	SSLPMAR
ABC1_YEAST	3	1031.6791	1	0	MNASFFAAQL

Where “peptide” is the number of the peptide in the protein, z is the charge of the ion (+1 by default), and p is the number of missed cleavages (you could calculate this yourself, or perhaps get Task 2 to pass this information on ?).

All of the options should be provided on the command line.

Bonus tasks

Also calculate an estimated average hydrophobicity for each peptide. Find a hydrophobicity scale on the internet (ie. 20 numbers, one for each amino acid) and use this to calculate a hydrophobicity $H\phi$ for each peptide (the average hydrophobicity for the peptide = sum of the individual values / number of amino acids). This term can be used to sort the peptides into bins in task 4.

Consider some standard post-translational modifications, such as blocking of reduced cysteines, oxidation of methionines and phosphorylation of Ser, Tyr and Thr.

The script should be able to report only N-terminal peptides (ie the first one in a protein) or just the C-terminal ones (the last one in a protein), selectable from the command line as an option.

Task 4 – the ion statistics calculator

This task should read in the output from task 3 and calculate some statistics and export the results, to the screen or a file. We recommend you allow the script to work in different ‘modes’, selected by the user on the command line. For example, in mode 1, as a minimum, it should be able to report the number of peptides in a given mass range i.e. calculate the average number of m/z values found within a m/z range. For another mode, mode 2 say, you might support a set of m/z bins specified by the user on the command line to calculate the number of peptides in a binned histogram. For the binned histogram data, the bin sizes should be user-definable and would allow you to model the mass accuracy of the instrument. In another mode, you might also adapt the script to consider sliding windows with a stepsize s and window size W, where (s < W). Finally, in another mode, you might support the consideration of how many actual proteins have been unambiguously identified (by virtue of containing at least one peptide whose m/z is unique in the entire protein given the mass accuracy supplied). The utility of this simple script will be self-evident when you read what you have to do in the next section. The whole team should be involved to help design the question this tool needs to be able to answer. So think carefully (as a team) as to what makes an enzyme more (or less) useful for proteomics? Although you may end up plotting the data yourself with other tools, or possibly using Python Notebooks, the script itself should just calculate the basic numbers and export CSV or tab-delimited text (not plots or pictures) which could be used by other tools.

Bonus tasks

Adapt the script so that it can consider hydrophobicity as calculated by task3. We suggest you might divide the peptides into sorted bins based on their hydrophobicity values (say 10 or 20 perhaps) and then consider the ‘uniqueness’ of the peptides in each bin, rather than in the entire proteome.

Adapt the script to cope with mass accuracy supplied as parts per million (ppm) instead of Da.

Adapt the script so that it is also to selectively count ions if they match a certain sequence pattern (ie. only ions which contain a Cys, or ions which start with a Met)

Overall aims

The overall aims of this project are to use your scripts to calculate the most effective identification system for a proteomics experiment under given (user-specified) conditions. A more effective system will be obtained if the proteomics researchers have a smaller number of peptides in the given mass (or more strictly, m/z) range, as any peptides they do find are more likely to be completely diagnostic. They might even be totally unambiguous – ie. there is only one protein which has a certain peptide with a certain m/z under certain conditions (Enzyme = Arg-c, missed cleavages = 1, N-terminal amino acid = Ser etc). You should conduct some focussed investigations with your programs, to answer the specific question:

Which is the best enzyme to use considering singly charged ions (+1), and m/z ratio between 1000 and 1500 Da, and up to 1 missed cleavage (ignore all post-translational modifications, additional sequence information etc.).

As well as more general investigations to see how the number of peptides per protein is affected by changing the full range of options.

Data: The genome sequence you need to work with will be uploaded on to the CoCalc environment, as well as the Blackboard site, in the file `genome.fasta`

Write-up: The write-up for this project is intended to be short, with two components – a group report on the pipeline as a whole (see below) and the individual programs alongside a man page description of what they are and how to use them. This should therefore be a single page of description for each script, which will be a Unix “Man page” style description of the operation of the script, its aims, options and results. This should be followed by the actual script itself (in a fixed width font please, which can be quite small if you like). The script should have sufficient comments to aid readability. Each person should prepare the “Man page” for their code, and the authors should be clearly marked.

You should also prepare (as a group) a short report (3 pages max) on which is the most effective enzyme, and why, as well as any general observations about what makes proteome protein identification easier (or harder) based on your observations.

Please submit all work via **Blackboard**. This should include all the individual reports (man page + script), as well as a single report per group. There should be two assignments on Blackboard to support this, one for each of you to submit individual work, and one to submit the group report.

Assessment

The marks for this work will be awarded as follows:

40 marks for each script awarded to the author – this **must be your own work**.

10 marks as a group mark for the report and general performance of the scripts.

Each person has a mark out of 50 by summing the individual and group marks

Deadline: 4:00 pm Thursday Dec 3rd

Project Submission Guideline and Requirements – please read carefully over the page

Project Submission Guideline and Requirements – please read carefully

Please read the following concerning your project submission. This describes the expectation on your code and how it will be assessed.

- 1) Submit the code and project report on Blackboard – I will make a submission area for individuals, and one for each group. Please submit the code/project in a text file or word doc so we can annotate them and return them to you for feedback.
- 2) Please also upload your code to CoCalc.com. You will each have been assigned a directory with your group name (i.e. Group 1 has a folder called “group01”), so you should be able to copy your code into this folder. It should be a standard alone piece of python code, please use the following naming convention: group1_task2.py. If you each have your final script in your own folder, when we collect it there will therefore be four pieces of code in the folder (plus man pages, and any other files). If you like, you can also include some of the actual output you have generated when running it. This might be useful for us when we are checking we have run it correctly.
- 3) Please upload a separate “man page” document to CoCalc in your relevant group folder – one for each task. This could be a plain text file or a word doc, but please use a fixed width font like Courier, so it looks like a real one.
- 4) We will evaluate the code by opening up a terminal window on CoCalc and running the code from the command line. We will expect it to be python3 compliant and run something like:

```
$ python3 group3_task2.py -m 1 -i input.fa -o results.out
```

or something very similar (this example is obviously for task2). There should be command line options for the user and we will also expect that there is some error reporting and/or usage if no options are supplied, or incorrect/incompatible parameters are given. We would not expect there to be a need for direct user interaction (ie. Responding to questions at the command line) when the scripts are run and all input/parameters are dealt with by the program from command line input. This would in principle allow the whole pipeline to be automated, run in the background, or scheduled as a job – a useful feature in bioinformatics.

- 5) Marks are awarded for the whole pipeline, based on functionality, ease of use, appropriate results, warnings/error reporting and indeed whether we can generate output from task 4 that can be plotted/analysed further.

For your information, this is the mark scheme we will use:

Individual scripts*Code features*

Appropriate use of basic constructs (strings, lists, dicts, loops, conditionals, etc)	10
Readability, comments, structure, modularity	10
Functionality, few or no errors, options, warnings, error trapping	15
<i>Man page</i> : correct use, description of function, options, example.....	5
Total individual mark	40

Group report marking scheme (marks possible)

Report (content, results, conclusions, presentation)	out of 5
Program suite functionality (does whole pipeline work? Options? etc).....	out of 5
Total group mark	out of 10
Overall mark	out of 50