

Biopython Tutorial and Cookbook

Jeff Chang, Brad Chapman, Iddo Friedberg, Thomas Hamelryck, Michiel de Hoon, Peter Cock

Last Update—15 June 2008

Contents

1	Introduction	5
1.1	What is Biopython?	5
1.1.1	What can I find in the Biopython package	5
1.2	Installing Biopython	6
1.3	FAQ	6
2	Quick Start – What can you do with Biopython?	8
2.1	General overview of what Biopython provides	8
2.2	Working with sequences	8
2.3	A usage example	9
2.4	Parsing sequence file formats	10
2.4.1	Simple FASTA parsing example	10
2.4.2	Simple GenBank parsing example	11
2.4.3	I love parsing – please don't stop talking about it!	11
2.5	Connecting with biological databases	11
2.6	What to do next	12
3	Sequence objects	13
3.1	Sequences and Alphabets	13
3.2	Sequences act like strings	14
3.3	Slicing a sequence	14
3.4	Turning Seq objects into strings	15
3.5	Nucleotide sequences and (reverse) complements	15
3.6	Concatenating or adding sequences	16
3.7	MutableSeq objects	16
3.8	Transcribing and Translation	17
3.9	Working with directly strings	18
4	Sequence Input/Output	19
4.1	Parsing or Reading Sequences	19
4.1.1	Reading Sequence Files	19
4.1.2	Iterating over the records in a sequence file	20
4.1.3	Getting a list of the records in a sequence file	21
4.1.4	Extracting data	21
4.2	Parsing sequences from the net	23
4.2.1	Parsing GenBank records from the net	23
4.2.2	Parsing SwissProt sequences from the net	24
4.3	Sequence files as Dictionaries	25
4.3.1	Specifying the dictionary keys	25
4.3.2	Indexing a dictionary using the SEGUID checksum	26

4.4	Writing Sequence Files	27
4.4.1	Converting between sequence file formats	28
4.4.2	Converting a file of sequences to their reverse complements	28
5	Sequence Alignment Input/Output	31
5.1	Parsing or Reading Sequence Alignments	31
5.1.1	Single Alignments	31
5.1.2	Multiple Alignments	34
5.1.3	Ambiguous Alignments	35
5.2	Writing Alignments	37
5.2.1	Converting between sequence alignment file formats	38
6	BLAST	42
6.1	Running BLAST locally	42
6.2	Running BLAST over the Internet	43
6.3	Saving BLAST output	44
6.4	Parsing BLAST output	45
6.5	The BLAST record class	47
6.6	Deprecated BLAST parsers	49
6.6.1	Parsing plain-text BLAST output	49
6.6.2	Parsing a file full of BLAST runs	51
6.6.3	Finding a bad record somewhere in a huge file	51
6.7	Dealing with PSIBlast	53
7	Accessing NCBI's Entrez databases	54
7.1	Entrez Guidelines	54
7.2	EInfo: Obtaining information about the Entrez databases	55
7.3	ESearch: Searching the Entrez databases	57
7.4	EPost	57
7.5	ESummary: Retrieving summaries from primary IDs	57
7.6	EFetch: Downloading full records from Entrez	58
7.7	ELink	59
7.8	EGQuery: Obtaining counts for search terms	60
7.9	ESpell: Obtaining spelling suggestions	60
7.10	Examples	60
7.10.1	Searching and downloading Entrez Nucleotide records	60
7.10.2	Finding the lineage of an organism	62
7.10.3	Using the history and WebEnv	62
8	Swiss-Prot, Prosite, Prodoc, and ExPASy	64
8.1	Bio.SwissProt: Parsing Swiss-Prot files	64
8.1.1	Parsing Swiss-Prot records	64
8.1.2	Parsing the Swiss-Prot keyword and category list	66
8.2	Bio.Prosite: Parsing Prosite records	66
8.3	Bio.Prosite.Prodoc: Parsing Prodoc records	68
8.4	Bio.ExPASy: Accessing the ExPASy server	68
8.4.1	Retrieving a Swiss-Prot record	68
8.4.2	Searching Swiss-Prot	69
8.4.3	Retrieving Prosite and Prodoc records	69

9 Cookbook – Cool things to do with it	71
9.1 PubMed	71
9.1.1 Sending a query to PubMed	71
9.1.2 Retrieving a PubMed record	71
9.2 GenBank	72
9.2.1 Retrieving GenBank entries from NCBI	73
9.2.2 Parsing GenBank records	74
9.2.3 Iterating over GenBank records	74
9.3 Dealing with alignments	75
9.3.1 Clustalw	75
9.3.2 Calculating summary information	76
9.3.3 Calculating a quick consensus sequence	77
9.3.4 Position Specific Score Matrices	77
9.3.5 Information Content	79
9.3.6 Translating between Alignment formats	80
9.4 Substitution Matrices	80
9.4.1 Using common substitution matrices	80
9.4.2 Creating your own substitution matrix from an alignment	80
9.5 BioSQL – storing sequences in a relational database	82
9.6 BioCorba	82
9.7 Going 3D: The PDB module	82
9.7.1 Structure representation	82
9.7.2 Disorder	87
9.7.3 Hetero residues	88
9.7.4 Some random usage examples	88
9.7.5 Common problems in PDB files	89
9.7.6 Other features	91
9.8 Bio.PopGen: Population genetics	91
9.8.1 GenePop	91
9.8.2 Coalescent simulation	93
9.8.3 Other applications	96
9.8.4 Future Developments	99
9.9 InterPro	99
10 Advanced	100
10.1 The SeqRecord and SeqFeature classes	100
10.1.1 Sequence ids and Descriptions – dealing with SeqRecords	100
10.1.2 Features and Annotations – SeqFeatures	101
10.2 Regression Testing Framework	104
10.2.1 Writing a Regression Test	104
10.3 Parser Design	105
10.4 Substitution Matrices	105
10.4.1 SubsMat	105
10.4.2 FreqTable	108
11 Where to go from here – contributing to Biopython	110
11.1 Maintaining a distribution for a platform	110
11.2 Bug Reports + Feature Requests	111
11.3 Contributing Code	111

12 Appendix: Useful stuff about Python	112
12.1 What the heck is a handle?	112
12.1.1 Creating a handle from a string	112

Chapter 1

Introduction

1.1 What is Biopython?

The Biopython Project is an international association of developers of freely available Python (<http://www.python.org>) tools for computational molecular biology. The web site <http://www.biopython.org> provides an online resource for modules, scripts, and web links for developers of Python-based software for life science research.

Basically, we just like to program in python and want to make it as easy as possible to use python for bioinformatics by creating high-quality, reusable modules and scripts.

1.1.1 What can I find in the Biopython package

The main Biopython releases have lots of functionality, including:

- The ability to parse bioinformatics files into python utilizable data structures, including support for the following formats:
 - Blast output – both from standalone and WWW Blast
 - Clustalw
 - FASTA
 - GenBank
 - PubMed and Medline
 - Expasy files, like Enzyme, Prodoc and Prosite
 - SCOP, including ‘dom’ and ‘lin’ files
 - UniGene
 - SwissProt
- Files in the supported formats can be iterated over record by record or indexed and accessed via a Dictionary interface.
- Code to deal with popular on-line bioinformatics destinations such as:
 - NCBI – Blast, Entrez and PubMed services
 - Expasy – Prodoc and Prosite entries
- Interfaces to common bioinformatics programs such as:

- Standalone Blast from NCBI
 - Clustalw alignment program.
- A standard sequence class that deals with sequences, ids on sequences, and sequence features.
 - Tools for performing common operations on sequences, such as translation, transcription and weight calculations.
 - Code to perform classification of data using k Nearest Neighbors, Naive Bayes or Support Vector Machines.
 - Code for dealing with alignments, including a standard way to create and deal with substitution matrices.
 - Code making it easy to split up parallelizable tasks into separate processes.
 - GUI-based programs to do basic sequence manipulations, translations, BLASTing, etc.
 - Extensive documentation and help with using the modules, including this file, on-line wiki documentation, the web site, and the mailing list.
 - Integration with other languages, including the Bioperl and Biojava projects, using the BioCorba interface standard (available with the biopython-corba module).

We hope this gives you plenty of reasons to download and start using Biopython!

1.2 Installing Biopython

All of the installation information for Biopython was separated from this document to make it easier to keep updated. The instructions cover installation of python, Biopython dependencies and Biopython itself. It is available in pdf (<http://biopython.org/DIST/docs/install/Installation.pdf>) and html formats (<http://biopython.org/DIST/docs/install/Installation.html>).

1.3 FAQ

1. *Why doesn't Bio.SeqIO work? It imports fine but there is no parse function etc.*
You need Biopython 1.43 or later. Older versions did contain some related code under the Bio.SeqIO name which has since been deprecated - and this is why the import “works”.
2. *Why doesn't Bio.SeqIO.read() work? The module imports fine but there is no read function!*
You need Biopython 1.45 or later.
3. *Why doesn't Bio.Blast work with the latest plain text NCBI blast output?*
The NCBI keep tweaking the plain text output from the BLAST tools, and keeping our parser up to date is an ongoing struggle. We recommend you use the XML output instead, which is designed to be read by a computer program.
4. *Why isn't Bio.AlignIO present? The module import fails!*
You need Biopython 1.46 or later.
5. *Why doesn't Bio.Entrez.read() work? The module imports fine but there is no read function!*
You need Biopython 1.46 or later.

6. *I looked in a directory for code, but I couldn't seem to find the code that does something. Where's it hidden?*

One thing to know is that we put code in `__init__.py` files. If you are not used to looking for code in this file this can be confusing. The reason we do this is to make the imports easier for users. For instance, instead of having to do a “repetitive” import like `from Bio.GenBank import GenBank`, you can just import like `from Bio import GenBank`.

Chapter 2

Quick Start – What can you do with Biopython?

This section is designed to get you started quickly with Biopython, and to give a general overview of what is available and how to use it. All of the examples in this section assume that you have some general working knowledge of python, and that you have successfully installed Biopython on your system. If you think you need to brush up on your python, the main python web site provides quite a bit of free documentation to get started with (<http://www.python.org/doc/>).

Since much biological work on the computer involves connecting with databases on the internet, some of the examples will also require a working internet connection in order to run.

Now that that is all out of the way, let's get into what we can do with Biopython.

2.1 General overview of what Biopython provides

As mentioned in the introduction, Biopython is a set of libraries to provide the ability to deal with "things" of interest to biologists working on the computer. In general this means that you will need to have at least some programming experience (in python, of course!) or at least an interest in learning to program. Biopython's job is to make your job easier as a programmer by supplying reusable libraries so that you can focus on answering your specific question of interest, instead of focusing on the internals of parsing a particular file format (of course, if you want to help by writing a parser that doesn't exist and contributing it to Biopython, please go ahead!). So Biopython's job is to make you happy!

One thing to note about Biopython is that it often provides multiple ways of "doing the same thing." To me, this can be frustrating since I often want to just know the one right way to do something. However, this is also a real benefit because it gives you lots of flexibility and control over the libraries. The tutorial helps to show you the common or easy ways to do things so that you can just make things work. To learn more about the alternative possibilities, look into the Cookbook section (which tells you some cool tricks and tips) and the Advanced section (which provides you with as much detail as you'd ever want to know!).

2.2 Working with sequences

Disputedly (of course!), the central object in bioinformatics is the sequence. Thus, we'll start with a quick introduction to the Biopython mechanisms for dealing with sequences, the `Seq` object, which we'll discuss in more detail in Chapter 3.

Most of the time when we think about sequences we have in my mind a string of letters like 'AGTACACTGGT'. You can create such `Seq` object with this sequence as follows - the ">>>" represents the python prompt followed by what you would type in:

```

>>> from Bio.Seq import Seq
>>> my_seq = Seq("AGTACTGGT")
>>> my_seq.alphabet
Alphabet()
>>> print my_seq.tostring()
AGTACTGGT

```

What we have here is a sequence object with a *generic* alphabet - reflecting the fact we have *not* specified if this is a DNA or protein sequence (okay, a protein with a lot of Alanines, Glycines, Cysteines and Threonines!). We'll talk more about alphabets in Chapter 3.

In addition to having an alphabet, the `Seq` object differs from the python string in the methods it supports. You can't do this with a plain string:

```

>>> my_seq
Seq('AGTACTGGT', Alphabet())
>>> my_seq.complement()
Seq('TCATGTGACCA', Alphabet())
>>> my_seq.reverse_complement()
Seq('ACCAGTGTACT', Alphabet())

```

The next most important class is the `SeqRecord` or Sequence Record. This holds a sequence (as a `Seq` object) with additional annotation including an identifier, name and description. The `Bio.SeqIO` module for reading and writing sequence file formats works with `SeqRecord` objects, which will be introduced below and cover in more detail by Chapter 4.

This covers the basic features and uses of the Biopython sequence class. Now that you've got some idea of what it is like to interact with the Biopython libraries, it's time to delve into the fun, fun world of dealing with biological file formats!

2.3 A usage example

Before we jump right into parsers and everything else to do with Biopython, let's set up an example to motivate everything we do and make life more interesting. After all, if there wasn't any biology in this tutorial, why would you want you read it?

Since I love plants, I think we're just going to have to have a plant based example (sorry to all the fans of other organisms out there!). Having just completed a recent trip to our local greenhouse, we've suddenly developed an incredible obsession with Lady Slipper Orchids (if you wonder why, have a look at some [Lady Slipper Orchids photos on Flickr](#), or try a [Google Image Search](#)).

Of course, orchids are not only beautiful to look at, they are also extremely interesting for people studying evolution and systematics. So let's suppose we're thinking about writing a funding proposal to do a molecular study of Lady Slipper evolution, and would like to see what kind of research has already been done and how we can add to that.

After a little bit of reading up we discover that the Lady Slipper Orchids are in the Orchidaceae family and the Cypripedioideae sub-family and are made up of 5 genera: *Cypripedium*, *Paphiopedilum*, *Phragmipedium*, *Selenipedium* and *Mexipedium*.

That gives us enough to get started delving for more information. So, let's look at how the Biopython tools can help us. We'll start with sequence parsing in Section 2.4, but the orchids will be back later on as well - for example we'll extra data from Swiss-Prot from certain orchid proteins in Chapter 8, search PubMed for papers about orchids in Section 9.1, extract sequence data from GenBank in Section 9.2.1, and work with ClustalW multiple sequence alignments of orchid proteins in Section 9.3.1.

2.4 Parsing sequence file formats

A large part of much bioinformatics work involves dealing with the many types of file formats designed to hold biological data. These files are loaded with interesting biological data, and a special challenge is parsing these files into a format so that you can manipulate them with some kind of programming language. However the task of parsing these files can be frustrated by the fact that the formats can change quite regularly, and that formats may contain small subtleties which can break even the most well designed parsers.

We are now going to briefly introduce the `Bio.SeqIO` module – you can find out more in Chapter 4. We'll start with an online search for our friends, the lady slipper orchids. To keep this introduction simple, we're just using the NCBI website by hand. Let's just take a look through the nucleotide databases at NCBI, using an Entrez online search (<http://www.ncbi.nlm.nih.gov:80/entrez/query.fcgi?db=Nucleotide>) for everything mentioning the text `Cypripedioideae` (this is the subfamily of lady slipper orchids).

When this tutorial was originally written, this search gave us only 94 hits, which we saved as a FASTA formatted text file ([ls_orchid.fasta](#); also available online [here](#)) and as a GenBank formatted text file ([ls_orchid.gbk](#); also available online [here](#)).

If you run the search today, you'll get hundreds of results! When following the tutorial, if you want to see the same list of genes, just download the two files above or copy them from `docs/examples/` in the Biopython source code. In Section 2.5 we will look at how to do a search like this from within python.

2.4.1 Simple FASTA parsing example

If you open the lady slipper orchids FASTA file in your favourite text editor, you'll see that the file starts like this:

```
>gi|2765658|emb|Z78533.1|CIZ78533 C.irapeanum 5.8S rRNA gene and ITS1 and ITS2 DNA
CGTAACAAGGTTTCCGTAGGTGAACCTGCGGAAGGATCATTGATGAGACCGTGAATAAACGATCGAGTG
AATCCGGAGGACCGGTGTACTCAGCTCACCGGGGGCATTGCTCCCGTGGTGACCCTGATTTGTTGTTGGG
...
```

It contains 94 records, each has a line starting with ">" (greater-than symbol) followed by the sequence on one or more lines. Now try this in python:

```
from Bio import SeqIO
handle = open("ls_orchid.fasta")
for seq_record in SeqIO.parse(handle, "fasta") :
    print seq_record.id
    print repr(seq_record.seq)
    print len(seq_record)
handle.close()
```

You should get something like this on your screen:

```
gi|2765658|emb|Z78533.1|CIZ78533
Seq('CGTAACAAGGTTTCCGTAGGTGAACCTGCGGAAGGATCATTGATGAGACCGTGG...CGC', SingleLetterAlphabet())
740
...
gi|2765664|emb|Z78439.1|PBZ78439
Seq('CATTGTTGAGATCACATAATAATTGATCGAGTTAATCTGGAGGATCTGTTTACT...GCC', SingleLetterAlphabet())
592
```

Notice that the FASTA format does not specify the alphabet, so `Bio.SeqIO` has defaulted to the rather generic `SingleLetterAlphabet()` rather than something DNA specific.

2.4.2 Simple GenBank parsing example

Now let's load the GenBank file instead - notice that the code to do this is almost identical to the snippet used above for a FASTA file - the only difference is we changed the filename and the format string:

```
from Bio import SeqIO
handle = open("ls_orchid.gbk")
for seq_record in SeqIO.parse(handle, "genbank") :
    print seq_record.id
    print repr(seq_record.seq)
    print len(seq_record)
handle.close()
```

This should give:

```
Z78533.1
Seq('CGTAACAAGGTTTCCGTAGGTGAACCTGCGGAAGGATCATTGATGAGACCGTGG...CGC', IUPACAmbiguousDNA())
740
...
Z78439.1
Seq('CATTGTTGAGATCACATAATAATTGATCGAGTTAATCTGGAGGATCTGTTTACT...GCC', IUPACAmbiguousDNA())
592
```

This time `Bio.SeqIO` has been able to choose a sensible alphabet, IUPAC Ambiguous DNA. You'll also notice that a shorter string has been used as the `seq_record.id` in this case.

2.4.3 I love parsing – please don't stop talking about it!

Biopython has a lot of parsers, and each has its own little special niches based on the sequence format it is parsing and all of that. While the most popular file formats have parsers integrated into `Bio.SeqIO`, for some of the rarer and unloved file formats there is either no parser at all, or an old parser which has not been linked in yet.

Chapter 4 covers `Bio.SeqIO` in more detail. Please also check the wiki page (<http://biopython.org/wiki/SeqIO>) for the latest information, or ask on the mailing list. The wiki page should includes an up to date list of supported file types, and more examples including writing sequences to a file, and converting between file formats.

If you are interested in sequence alignments, then look at the `Bio.AlignIO` module introduced in Chapter 5. Again, there is a wiki page for this too (<http://biopython.org/wiki/AlignIO>).

The next place to look for information about specific parsers and how to do cool things with them is in the Cookbook, Section 9 of this Tutorial. If you don't find the information you are looking for, please consider helping out your poor overworked documentors and submitting a cookbook entry about it! (once you figure out how to do it, that is!)

2.5 Connecting with biological databases

One of the very common things that you need to do in bioinformatics is extract information from biological databases. It can be quite tedious to access these databases manually, especially if you have a lot of repetitive work to do. Biopython attempts to save you time and energy by making some on-line databases available from python scripts. Currently, Biopython has code to extract information from the following databases:

- [ExpASY](#) – See Chapter 8 in the Cookbook for more information.
- [Entrez from NCBI](#) – See Section 7.10.

- [PubMed from NCBI](#) – See Section [9.1](#) in the Cookbook for example code detailing how to use this.
- [SCOP](#)

The code in these modules basically makes it easy to write python code that interact with the CGI scripts on these pages, so that you can get results in an easy to deal with format. In some cases, the results can be tightly integrated with the Biopython parsers to make it even easier to extract information.

2.6 What to do next

Now that you've made it this far, you hopefully have a good understanding of the basics of Biopython and are ready to start using it for doing useful work. The best thing to do now is to start snooping around in the source code and looking at the automatically generated documentation.

Once you get a picture of what you want to do, and what libraries in Biopython will do it, you should take a peak at the Cookbook, which may have example code to do something similar to what you want to do.

If you know what you want to do, but can't figure out how to do it, please feel free to post questions to the main biopython list (biopython@biopython.org). This will not only help us answer your question, it will also allow us to improve the documentation so it can help the next person do what you want to do.

Enjoy the code!

Chapter 3

Sequence objects

Biological sequences are arguably the central object in Bioinformatics, and in this chapter we'll introduce the Biopython mechanism for dealing with sequences, the `Seq` object. In Chapter 4 on Sequence Input/Output (and Section 10.1), we'll see that the `Seq` object is also used in the `SeqRecord` object, which combines the sequence information with any annotation.

Sequences are essentially strings of letters like `AGTACACTGGT`, which seems very natural since this is the most common way that sequences are seen in biological file formats.

There are two important differences between the `Seq` object and standard python strings. First of all the `Seq` object has a slightly different set of methods to a plain python string (for example, a `reverse_complement()` method used for nucleotide sequences). Secondly, the `Seq` object has an important attribute, `alphabet`, which is an object describing what the individual characters making up the sequence string “mean”, and how they should be interpreted. For example, is `AGTACACTGGT` a DNA sequence, or just a protein sequence that happens to be rich in Alanines, Glycines, Cysteines and Threonines?

3.1 Sequences and Alphabets

The alphabet object is perhaps the important thing that makes the `Seq` object more than just a string. The currently available alphabets for Biopython are defined in the `Bio.Alphabet` module. We'll use the IUPAC alphabets (<http://www.chem.qmw.ac.uk/iupac/>) here to deal with some of our favorite objects: DNA, RNA and Proteins.

`Bio.Alphabet.IUPAC` provides basic definitions for proteins, DNA and RNA, but additionally provides the ability to extend and customize the basic definitions. For instance, for proteins, there is a basic `IUPACProtein` class, but there is an additional `ExtendedIUPACProtein` class providing for the additional elements “Asx” (asparagine or aspartic acid), “Sec” (selenocysteine), and “Glx” (glutamine or glutamic acid). For DNA you've got choices of `IUPACUnambiguousDNA`, which provides for just the basic letters, `IUPACAmbiguousDNA` (which provides for ambiguity letters for every possible situation) and `ExtendedIUPACDNA`, which allows letters for modified bases. Similarly, RNA can be represented by `IUPACAmbiguousRNA` or `IUPACUnambiguousRNA`.

The advantages of having an alphabet class are two fold. First, this gives an idea of the type of information the `Seq` object contains. Secondly, this provides a means of constraining the information, as a means of type checking.

Now that we know what we are dealing with, let's look at how to utilize this class to do interesting work. You can create an ambiguous sequence with the default generic alphabet like this:

```
>>> from Bio.Seq import Seq
>>> my_seq = Seq("AGTACACTGGT")
>>> my_seq
```

```
Seq('AGTACTGGT', Alphabet())
>>> my_seq.alphabet
Alphabet()
```

However, where possible you should specify the alphabet explicitly when creating your sequence objects - in this case an unambiguous DNA alphabet object:

```
>>> from Bio.Seq import Seq
>>> from Bio.Alphabet import IUPAC
>>> my_seq = Seq('AGTACTGGT', IUPAC.unambiguous_dna)
>>> my_seq
Seq('AGTACTGGT', IUPACUnambiguousDNA())
>>> my_seq.alphabet
IUPACUnambiguousDNA()
```

3.2 Sequences act like strings

In many ways, we can deal with Seq objects as if they were normal python strings, for example getting the length, or iterating over the elements:

```
from Bio.Seq import Seq
from Bio.Alphabet import IUPAC
my_seq = Seq('GATCGATGGCCTATATAGGATCGAAAATCGC', IUPAC.unambiguous_dna)
for index, letter in enumerate(my_seq) :
    print index, letter
print len(letter)
```

You can access elements of the sequence in the same way as for strings (but remember, python counts from zero!):

```
>>> print my_seq[0] #first element
>>> print my_seq[2] #third element
>>> print my_seq[-1] #last element
```

The Seq object has a `.count()` method, just like a string:

```
>>> len(my_seq)
32
>>> my_seq.count("G")
10
>>> float(my_seq.count("G") + my_seq.count("C")) / len(my_seq)
0.46875
```

While you could use the above snippet of code to calculate a GC%, note that Biopython does have some GC functions already built in, see the `Bio.SeqUtils` module.

3.3 Slicing a sequence

A more complicated example, let's get a slice of the sequence:

```
>>> from Bio.Seq import Seq
>>> from Bio.Alphabet import IUPAC
>>> my_seq = Seq('GATCGATGGCCTATATAGGATCGAAAATCGC', IUPAC.unambiguous_dna)
>>> my_seq[4:12]
Seq('GATGGGCC', IUPACUnambiguousDNA())
```

Two things are interesting to note. First, this follows the normal conventions for python strings. So the first element of the sequence is 0 (which is normal for computer science, but not so normal for biology). When you do a slice the first item is included (i. e. 4 in this case) and the last is excluded (12 in this case), which is the way things work in python, but of course not necessarily the way everyone in the world would expect. The main goal is to stay consistent with what python does.

The second thing to notice is that the slice is performed on the sequence data string, but the new object produced is another `Seq` object which retains the alphabet information from the original `Seq` object.

Also like a python string, you can do slices with a start, stop and *stride* (the step size, which defaults to one). For example, we can get the first, second and third codon positions of this DNA sequence:

```
>>> my_seq[0:3]
Seq('GCTGTAGTAAG', IUPACUnambiguousDNA())
>>> my_seq[1:3]
Seq('AGGCATGCATC', IUPACUnambiguousDNA())
>>> my_seq[2:3]
Seq('TAGCTAAGAC', IUPACUnambiguousDNA())
```

Another stride trick you might have seen with a python string is the use of a -1 stride to reverse the string. You can do this with a `Seq` object too:

```
>>> my_seq[::-1]
Seq('CGCTAAAAGCTAGGATATATCCGGTAGCTAG', IUPACUnambiguousDNA())
```

3.4 Turning Seq objects into strings

If you are really do just need a plain string, for example to print out, write to a file, or insert into a database, then this is very easy to get:

```
>>> my_seq.tostring()
'GATCGATGGGCCTATATAGGATCGAAAATCGC'
```

3.5 Nucleotide sequences and (reverse) complements

For nucleotide sequences, you can easily obtain the complement or reverse complement of a `Seq` object:

```
>>> my_seq
Seq('GATCGATGGGCCTATATAGGATCGAAAATCGC', IUPACUnambiguousDNA())
>>> my_seq.complement()
Seq('CTAGCTACCCGGATATATCCTAGCTTTTAGCG', IUPACUnambiguousDNA())
>>> my_seq.reverse_complement()
Seq('GCGATTTTCGATCCTATATAGGCCCATCGATC', IUPACUnambiguousDNA())
```

In all of these operations, the alphabet property is maintained. This is very useful in case you accidentally end up trying to do something weird like take the (reverse)complement of a protein sequence:

```
>>> protein_seq = Seq("EVRNAK", IUPAC.protein)
>>> dna_seq = Seq("ACGT", IUPAC.unambiguous_dna)
>>> protein_seq.complement()
Traceback (most recent call last):
  File "<stdin>", line 1, in ?
  File "/usr/local/lib/python2.4/site-packages/Bio/Seq.py", line 108, in complement
    raise ValueError, "Proteins do not have complements!"
ValueError: Proteins do not have complements!
```


3.6 Concatenating or adding sequences

Naturally, you can in principle add any two Seq objects together - just like you can with python strings to concatenate them. However, you can't add sequences with incompatible alphabets, such as a protein sequence and a DNA sequence:

```
>>> protein_seq + dna_seq
Traceback (most recent call last):
  File "<stdin>", line 1, in ?
  File "/usr/local/lib/python2.4/site-packages/Bio/Seq.py", line 42, in __add__
    raise TypeError, ("incompatable alphabets", str(self.alphabet),
TypeError: ('incompatable alphabets', 'IUPACProtein()', 'IUPACUnambiguousDNA()')
```

If you *really* wanted to do this, you'd have to first give both sequences generic alphabets:

```
>>> from Bio.Alphabet import generic_alphabet
>>> protein_seq.alphabet = generic_alphabet
>>> dna_seq.alphabet = generic_alphabet
>>> protein_seq + dna_seq
Seq('EVRNAKACGT', Alphabet())
```

Here is an example of adding a generic nucleotide sequence to an unambiguous IUPAC DNA sequence, resulting in an ambiguous nucleotide sequence:

```
>>> from Bio.Seq import Seq
>>> from Bio.Alphabet import generic_nucleotide
>>> from Bio.Alphabet import IUPAC
>>> nuc_seq = Seq('GATCGATGC', generic_nucleotide)
>>> dna_seq = Seq('ACGT', IUPAC.unambiguous_dna)
>>> nuc_seq
Seq('GATCGATGC', NucleotideAlphabet())
>>> dna_seq
Seq('ACGT', IUPACUnambiguousDNA())
>>> nuc_seq + dna_seq
Seq('GATCGATGCACGT', NucleotideAlphabet())
```

3.7 MutableSeq objects

Just like the normal python string, the Seq object is "read only", or in python terminology, not mutable. Apart from the wanting the Seq object to act like a string, this is also a useful default since in many biological applications you want to ensure you are not changing your sequence data:

```
>>> my_seq[5] = "G"
Traceback (most recent call last):
  File "<stdin>", line 1, in ?
AttributeError: 'Seq' instance has no attribute '__setitem__'
```

However, you can convert it into a mutable sequence (a MutableSeq object) and do pretty much anything you want with it:

```
>>> mutable_seq = my_seq.tomutable()
>>> print mutable_seq
MutableSeq('GATCGATGGGCCTATATAGGATCGAAAATCGC', IUPACUnambiguousDNA())
```

```

>>> mutable_seq[5] = "T"
>>> print mutable_seq
MutableSeq('GATCGTTGGGCCTATATAGGATCGAAAATCGC', IUPACUnambiguousDNA())
>>> mutable_seq.remove("T")
>>> print mutable_seq
MutableSeq('GACGTTGGGCCTATATAGGATCGAAAATCGC', IUPACUnambiguousDNA())
>>> mutable_seq.reverse()
>>> print mutable_seq
MutableSeq('CGCTAAAAGCTAGGATATATCCGGGTTGCAG', IUPACUnambiguousDNA())

```

3.8 Transcribing and Translation

Now that the nature of the sequence object makes some sense, the next thing to look at is what kind of things we can do with a sequence. The Bio directory contains two useful modules to transcribe and translate a sequence object. These tools work based on the alphabet of the sequence.

For instance, let's suppose we want to transcribe a DNA sequence:

```

>>> from Bio.Seq import Seq
>>> from Bio.Alphabet import IUPAC
>>> my_seq = Seq("GATCGATGGGCCTATATAGGATCGAAAATCGC", IUPAC.unambiguous_dna)

```

This contains an unambiguous alphabet, so to transcribe we would do the following:

```

>>> from Bio import Transcribe
>>> transcriber = Transcribe.unambiguous_transcriber
>>> my_rna_seq = transcriber.transcribe(my_seq)
>>> print my_rna_seq
Seq('GAUCGAUGGGCCUAUAUAGGAUCGAAAUCGC', IUPACUnambiguousRNA())

```

The alphabet of the new RNA Seq object is created for free, so again, dealing with a Seq object is no more difficult than dealing with a simple string.

You can also reverse transcribe RNA sequences:

```

>>> transcriber.back_transcribe(my_rna_seq)
Seq('GATCGATGGGCCTATATAGGATCGAAAATCGC', IUPACUnambiguousDNA())

```

To translate our DNA object we have quite a few choices. First, we can use any number of translation tables depending on what we know about our DNA sequence. The translation tables available in biopython were taken from information at <ftp://ftp.ncbi.nlm.nih.gov/entrez/misc/data/gc.prt>. So, you have tons of choices to pick from. For this, let's just focus on two choices: the Standard translation table, and the Translation table for Vertebrate Mitochondrial DNA. These tables are labeled with id numbers 1 and 2, respectively. Now that we know what tables we are looking to get, we're all set to perform a basic translation. First, we need to get our translators that use these tables. Since we are still dealing with our unambiguous DNA object, we want to fetch translators that take this into account:

```

>>> from Bio import Translate
>>> standard_translator = Translate.unambiguous_dna_by_id[1]
>>> mito_translator = Translate.unambiguous_dna_by_id[2]

```

Once we've got the proper translators, it's time to go ahead and translate a sequence:

```

>>> my_seq = Seq("GCCATTGTAATGGGCCGCTGAAAGGGTGCCCGA", IUPAC.unambiguous_dna)
>>> standard_translator.translate(my_seq)
Seq('AIVMGR*KGAR', IUPACProtein())
>>> mito_translator.translate(my_seq)
Seq('AIVMGRWKGAR', IUPACProtein())

```

Notice that the default translation will just go ahead and proceed blindly through a stop codon. If you are aware that you are translating some kind of open reading frame and want to just see everything up until the stop codon, this can be easily done with the `translate_to_stop` function:

```

>>> standard_translator.translate_to_stop(my_seq)
Seq('AIVMGR', IUPACProtein())

```

Similar to the transcriber, it is also possible to reverse translate a protein into a DNA sequence:

```

>>> my_protein = Seq("AVMGRWKGGRAAG", IUPAC.protein)
>>> standard_translator.back_translate(my_protein)
Seq('GCTGTTATGGGTCGTTGGAAGGGTGGTCGTGCTGCTGGT', IUPACUnambiguousDNA())

```

3.9 Working with directly strings

To close this chapter, for those you who *really* don't want to use the sequence objects, there are a few module level functions in `Bio.Seq` which will accept plain python strings (or `Seq` objects or `MutableSeq` objects):

```

>>> from Bio.Seq import reverse_complement, transcribe, back_transcribe, translate
>>> my_string = "GCTGTTATGGGTCGTTGGAAGGGTGGTCGTGCTGCTGGTTAG"
>>> reverse_complement(my_string)
'CTAACCAGCAGCACGACCACCCTTCCAACGACCCATAACAGC'
>>> transcribe(my_string)
'GCUGUUAUGGGUCGUUGGAAGGGUGGUCGUGCUGCUGGUUAG'
>>> back_transcribe(my_string)
'GCTGTTATGGGTCGTTGGAAGGGTGGTCGTGCTGCTGGTTAG'
>>> translate(my_string)
'AVMGRWKGGRAAG*'

```

You are however, encouraged to work with the `Seq` object by default.

Chapter 4

Sequence Input/Output

In this chapter we'll discuss in more detail the `Bio.SeqIO` module, which was briefly introduced in Chapter 2. This is a relatively new interface, added in Biopython 1.43, which aims to provide a simple interface for working with assorted sequence file formats in a uniform way.

The “catch” is that you have to work with `SeqRecord` objects - which contain a `Seq` object (as described in Chapter 3) plus annotation like an identifier and description. We'll introduce the basics of `SeqRecord` object in this chapter, but see Section 10.1 for more details.

4.1 Parsing or Reading Sequences

The workhorse function `Bio.SeqIO.parse()` is used to read in sequence data as `SeqRecord` objects. This function expects two arguments:

1. The first argument is a *handle* to read the data from. A handle is typically a file opened for reading, but could be the output from a command line program, or data downloaded from the internet (see Section 4.2). See Section 12.1 for more about handles.
2. The second argument is a lower case string specifying sequence format – we don't try and guess the file format for you! See <http://biopython.org/wiki/SeqIO> for a full listing of supported formats.

This returns an *iterator* which gives `SeqRecord` objects. Iterators are typically used in a for loop.

Sometimes you'll find yourself dealing with files which contain only a single record. For this situation Biopython 1.45 introduced the function `Bio.SeqIO.read()`. Again, this takes a handle and format as arguments. Provided there is one and only one record, this is returned as a `SeqRecord` object.

4.1.1 Reading Sequence Files

In general `Bio.SeqIO.parse()` is used to read in sequence files as `SeqRecord` objects, and is typically used with a for loop like this:

```
from Bio import SeqIO
handle = open("ls_orchid.fasta")
for seq_record in SeqIO.parse(handle, "fasta") :
    print seq_record.id
    print repr(seq_record.seq)
    print len(seq_record.seq)
handle.close()
```

The above example is repeated from the introduction in Section 2.4, and will load the orchid DNA sequences in the FASTA format file `ls_orchid.fasta`. If instead you wanted to load a GenBank format file like `ls_orchid.gb` then all you need to do is change the filename and the format string:

```
from Bio import SeqIO
handle = open("ls_orchid.gb")
for seq_record in SeqIO.parse(handle, "genbank") :
    print seq_record.id
    print seq_record.seq
    print len(seq_record.seq)
handle.close()
```

Similarly, if you wanted to read in a file in another file format, then assuming `Bio.SeqIO.parse()` supports it you would just need to change the format string as appropriate, for example “swiss” for SwissProt files or “embl” for EMBL text files. There is a full listing on the wiki page (<http://biopython.org/wiki/SeqIO>).

4.1.2 Iterating over the records in a sequence file

In the above examples, we have usually used a for loop to iterate over all the records one by one. You can use the for loop with all sorts of Python objects (including lists, tuples and strings) which support the iteration interface.

The object returned by `Bio.SeqIO` is actually an iterator which returns `SeqRecord` objects. You get to see each record in turn, but once and only once. The plus point is that an iterator can save you memory when dealing with large files.

Instead of using a for loop, can also use the `.next()` method of an iterator to step through the entries, like this:

```
from Bio import SeqIO
handle = open("ls_orchid.fasta")
record_iterator = SeqIO.parse(handle, "fasta")

first_record = record_iterator.next()
print first_record.id
print first_record.description

second_record = record_iterator.next()
print second_record.id
print second_record.description

handle.close()
```

Note that if you try and use `.next()` and there are no more results, you’ll either get back the special Python object `None` or a `StopIteration` exception.

One special case to consider is when your sequence files have multiple records, but you only want the first one. In this situation the following code is very concise:

```
from Bio import SeqIO
first_record = SeqIO.parse(open("ls_orchid.gb"), "genbank").next()
```

A word of warning here – using the `.next()` method like this will silently ignore any additional records in the file. If your files have *one and only one* record, like some of the online examples later in this chapter, or a GenBank file for a single chromosome, then use the new `Bio.SeqIO.read()` function instead. This will check there are no extra unexpected records present.

4.1.3 Getting a list of the records in a sequence file

In the previous section we talked about the fact that `Bio.SeqIO.parse()` gives you a `SeqRecord` iterator, and that you get the records one by one. Very often you need to be able to access the records in any order. The Python *list* data type is perfect for this, and we can turn the record iterator into a list of `SeqRecord` objects using the built-in Python function `list()` like so:

```
from Bio import SeqIO
handle = open("ls_orchid.gbk")
records = list(SeqIO.parse(handle, "genbank"))
handle.close()

print "Found %i records" % len(records)

print "The last record"
last_record = records[-1] #using Python's list tricks
print last_record.id
print repr(last_record.seq)
print len(last_record.seq)

print "The first record"
first_record = records[0] #remember, Python counts from zero
print first_record.id
print repr(first_record.seq)
print len(first_record.seq)
```

Giving:

```
Found 94 records
The last record
Z78439.1
Seq('CATTGTTGAGATCACATAATAATTGATCGAGTTAATCTGGAGGATCTGTTTACT...GCC', IUPACAmbiguousDNA())
592
The first record
Z78533.1
Seq('CGTAACAAGGTTTCCGTAGGTGAACCTGCGGAAGGATCATTGATGAGACCGTGG...CGC', IUPACAmbiguousDNA())
740
```

You can of course still use a for loop with a list of `SeqRecord` objects. Using a list is much more flexible than an iterator (for example, you can determine the number of records from the length of the list), but does need more memory because it will hold all the records in memory at once.

4.1.4 Extracting data

Suppose you wanted to extract a list of the species from the [ls_orchid.gbk](#) GenBank file. Let's have a close look at the first record in the file and see where the species gets stored:

```
from Bio import SeqIO
record_iterator = SeqIO.parse(open("ls_orchid.gbk"), "genbank")
first_record = record_iterator.next()
print first_record
```

That should give something like this:

```

ID: Z78533.1
Name: Z78533
Description: C.irapeanum 5.8S rRNA gene and ITS1 and ITS2 DNA.
/source=Cypripedium irapeanum
/taxonomy=['Eukaryota', 'Viridiplantae', 'Streptophyta', ..., 'Cypripedium']
/keywords=['5.8S ribosomal RNA', '5.8S rRNA gene', 'internal transcribed spacer', 'ITS1', 'ITS2']
/references=[...]
/accessions=['Z78533']
/data_file_division=PLN
/date=30-NOV-2006
/organism=Cypripedium irapeanum
/gi=2765658
Seq('CGTAACAAGGTTTCCGTAGGTGAACCTGCGGAAGGATCATTGATGAGACCGTGG...CGC', IUPACAmbiguousDNA())

```

The information we want, *Cypripedium irapeanum*, is held in the annotations dictionary under 'source' and 'organism', which we can access like this:

```
print first_record.annotations["source"]
```

or:

```
print first_record.annotations["organism"]
```

In general, 'organism' is used for the scientific name (in latin, e.g. *Arabidopsis thaliana*), while 'source' will often be the common name (e.g. thale cress). In this example, as is often the case, the two fields are identical.

Now let's go through all the records, building up a list of the species each orchid sequence is from:

```

from Bio import SeqIO
handle = open("ls_orchid.gb")
all_species = []
for seq_record in SeqIO.parse(handle, "genbank") :
    all_species.append(seq_record.annotations["organism"])
handle.close()
print all_species

```

Another way of writing this code is to use a list comprehension (introduced in Python 2.0) like this:

```

from Bio import SeqIO
all_species = [seq_record.annotations["organism"] for seq_record in \
                SeqIO.parse(open("ls_orchid.gb"), "genbank")]
print all_species

```

In either case, the result is:

```
['Cypripedium irapeanum', 'Cypripedium californicum', ..., 'Paphiopedilum barbatum']
```

Great. That was pretty easy because GenBank files are annotated in a standardised way. Now, let's suppose you wanted to extract a list of the species from your FASTA file, rather than the GenBank file. The bad news is you will have to write some code to extract the data you want from the record's description line - if the information is in the file in the first place!

For this example, notice that if you break up the description line at the spaces, then the species is there as field number one (field zero is the record identifier). That means we can do this:

```

from Bio import SeqIO
handle = open("ls_orchid.fasta")
all_species = []
for seq_record in SeqIO.parse(handle, "fasta") :
    all_species.append(seq_record.description.split()[1])
handle.close()
print all_species

```

This gives:

```
['C.irapeanum', 'C.californicum', 'C.fasciculatum', 'C.margaritaceum', ..., 'P.barbatum']
```

The concise alternative using list comprehensions (Python 2.0 or later) would be:

```

from Bio import SeqIO
all_species == [seq_record.description.split()[1] for seq_record in \
    SeqIO.parse(open("ls_orchid.fasta"), "fasta")]
print all_species

```

In general, extracting information from the FASTA description line is not very nice. If you can get your sequences in a well annotated file format like GenBank or EMBL, then this sort of annotation information is much easier to deal with.

4.2 Parsing sequences from the net

In the previous section, we looked at parsing sequence data from a file handle. We hinted that handles are not always from files, and in this section we'll use handles to internet connections to download sequences.

Note that just because you *can* download sequence data and parse it into a `SeqRecord` object in one go doesn't mean this is always a good idea. In general, you should probably download sequences *once* and save them to a file for reuse.

4.2.1 Parsing GenBank records from the net

Section 7.6 talks about the Entrez EFetch interface in more detail, but for now let's just connect to the NCBI and get a few orchid proteins from GenBank using their GI numbers.

First of all, let's fetch just one record. Remember, when you expect the handle to contain one and only one record, use the `Bio.SeqIO.read()` function:

```

from Bio import Entrez
from Bio import SeqIO
handle = Entrez.efetch(db="protein", rettype="genbank", id="6273291")
seq_record = SeqIO.read(handle, "genbank")
handle.close()
print "%s with %i features" % (seq_record.id, len(seq_record.features))

```

Expected output:

```
gi|6273291|gb|AF191665.1|AF191665 with 3 features
```

The NCBI will also let you ask for the file in other formats, for example as a FASTA file. If you don't care about the annotations and features in a GenBank file, this is a better choice to download because it's a smaller file:


```

from Bio import Entrez
from Bio import SeqIO
handle = Entrez.efetch(db="protein", rettype="fasta", id="6273291")
seq_record = SeqIO.read(handle, "fasta")
handle.close()
print "%s with %i features" % (seq_record.id, len(seq_record.features))

```

Expected output:

```
gi|6273291|gb|AF191665.1|AF191665 with 0 features
```

Now let's fetch several records. This time the handle contains multiple records, so we must use the `Bio.SeqIO.parse()` function:

```

from Bio import Entrez
from Bio import SeqIO
handle = Entrez.efetch(db="protein", rettype="genbank", \
                      id="6273291,6273290,6273289")
for seq_record in SeqIO.parse(handle, "genbank") :
    print seq_record.id, seq_record.description[:50] + "..."
    print "Sequence length %i," % len(seq_record),
    print "%i features," % len(seq_record.features),
    print "from: %s" % seq_record.annotations['source']
handle.close()

```

That should give the following output:

```

AF191665.1 Opuntia marenae rpl16 gene; chloroplast gene for c...
Sequence length 902, 3 features, from: chloroplast Opuntia marenae
AF191664.1 Opuntia clavata rpl16 gene; chloroplast gene for c...
Sequence length 899, 3 features, from: chloroplast Grusonia clavata
AF191663.1 Opuntia bradtiana rpl16 gene; chloroplast gene for...
Sequence length 899, 3 features, from: chloroplast Opuntia bradtiana

```

See Chapter 7 for more about the `Bio.Entrez` module, and make sure to read about the NCBI guidelines for using Entrez (Section 7.1).

4.2.2 Parsing SwissProt sequences from the net

Now let's use a handle to download a SwissProt file from ExPASy, something covered in more depth in Chapter 8. As mentioned above, the `Bio.SeqIO.read()` function is included in Biopython 1.45 or later.

```

from Bio import ExPASy
from Bio import SeqIO
handle = ExPASy.get_sprot_raw("023729")
seq_record = SeqIO.read(handle, "swiss")
handle.close()
print seq_record.id
print seq_record.name
print seq_record.description
print repr(seq_record.seq)
print "Length %i" % len(seq_record)
print seq_record.annotations['keywords']

```

Assuming your network connection is OK, you should get back:

```
023729
CHS3_BROFI
Chalcone synthase 3 (EC 2.3.1.74) (Naringenin-chalcone synthase 3).
Seq('MAPAMEEIRQAQRAEGPAAVLAIGTSTPPNALYQADYPDYYFRITKSEHLTELK...GAE', ProteinAlphabet())
Length 394
['Acyltransferase', 'Flavonoid biosynthesis', 'Transferase']
```

4.3 Sequence files as Dictionaries

The next thing that we'll do with our ubiquitous orchid files is to show how to index them and access them like a database using the Python *dictionary* datatype (like a hash in Perl). This is very useful for large files where you only need to access certain elements of the file, and makes for a nice quick 'n dirty database.

You can use the function `SeqIO.to_dict()` to make a `SeqRecord` dictionary (in memory). By default this will use each record's identifier (i.e. the `.id` attribute) as the key. Let's try this using our GenBank file:

```
from Bio import SeqIO
handle = open("ls_orchid.gb")
orchid_dict = SeqIO.to_dict(SeqIO.parse(handle, "genbank"))
handle.close()
```

Since this variable `orchid_dict` is an ordinary Python dictionary, we can look at all of the keys we have available:

```
>>> print orchid_dict.keys()
['Z78484.1', 'Z78464.1', 'Z78455.1', 'Z78442.1', 'Z78532.1', 'Z78453.1', ..., 'Z78471.1']
```

We can access a single `SeqRecord` object via the keys and manipulate the object as normal:

```
>>> seq_record = orchid_dict["Z78475.1"]
>>> print seq_record.description
P.supardii 5.8S rRNA gene and ITS1 and ITS2 DNA
>>> print repr(seq_record.seq)
Seq('CGTAACAAGGTTTCCGTAGGTGAACCTGCGGAAGGATCATTGTTGAGATCACAT...GGT', IUPACAmbiguousDNA())
```

So, it is very easy to create an in memory "database" of our GenBank records. Next we'll try this for the FASTA file instead.

4.3.1 Specifying the dictionary keys

Using the same code as above, but for the FASTA file instead:

```
from Bio import SeqIO
handle = open("ls_orchid.fasta")
orchid_dict = SeqIO.to_dict(SeqIO.parse(handle, "fasta"))
handle.close()
print orchid_dict.keys()
```

This time the keys are:

```
['gi|2765596|emb|Z78471.1|PDZ78471', 'gi|2765646|emb|Z78521.1|CCZ78521', ...
..., 'gi|2765613|emb|Z78488.1|PTZ78488', 'gi|2765583|emb|Z78458.1|PHZ78458']
```

You should recognise these strings from when we parsed the FASTA file earlier in Section 2.4.1. Suppose you would rather have something else as the keys - like the accession numbers. This brings us nicely to `SeqIO.to_dict()`'s optional argument `key_function`, which lets you define what to use as the dictionary key for your records.

First you must write your own function to return the key you want (as a string) when given a `SeqRecord` object. In general, the details of function will depend on the sort of input records you are dealing with. But for our orchids, we can just split up the record's identifier using the "pipe" character (the vertical line) and return the fourth entry (field three):

```
def get_accession(record) :
    """Given a SeqRecord, return the accession number as a string

    e.g. "gi|2765613|emb|Z78488.1|PTZ78488" -> "Z78488.1"
    """
    parts = record.id.split("|")
    assert len(parts) == 5 and parts[0] == "gi" and parts[2] == "emb"
    return parts[3]
```

Then we can give this function to the `SeqIO.to_dict()` function to use in building the dictionary:

```
from Bio import SeqIO
handle = open("ls_orchid.fasta")
orchid_dict = SeqIO.to_dict(SeqIO.parse(handle, "fasta"), key_function=get_accession)
handle.close()
print orchid_dict.keys()
```

Finally, as desired, the new dictionary keys:

```
>>> print orchid_dict.keys()
['Z78484.1', 'Z78464.1', 'Z78455.1', 'Z78442.1', 'Z78532.1', 'Z78453.1', ..., 'Z78471.1']
```

Not too complicated, I hope!

4.3.2 Indexing a dictionary using the SEGUID checksum

To give another example of working with dictionaries of `SeqRecord` objects, we'll use the SEGUID checksum function (added in Biopython 1.44). This is a relatively recent checksum, and collisions should be very rare (i.e. two different sequences with the same checksum), an improvement on the CRC64 checksum.

Once again, working with the orchids GenBank file:

```
from Bio import SeqIO
from Bio.SeqUtils.CheckSum import seguid
for record in SeqIO.parse(open("ls_orchid.gb"), "genbank") :
    print record.id, seguid(record.seq)
```

This should give:

```
Z78533.1 JUEoWn6DPhgZ9nAyowsgtoD9TTo
Z78532.1 MN/s0q9zDoCVEEc+k/IFwCNF2pY
...
Z78439.1 H+JfaShya/4yyAj7IbMqgNkxdxQ
```

Now, recall the `Bio.SeqIO.to_dict()` function's `key_function` argument expects a function which turns a `SeqRecord` into a string. We can't use the `seguid()` function directly because it expects to be given a `Seq` object (or a string). However, we can use python's `lambda` feature to create a "one off" function to give to `Bio.SeqIO.to_dict()` instead:

```

from Bio import SeqIO
from Bio.SeqUtils.CheckSum import seguid
seguid_dict = SeqIO.to_dict(SeqIO.parse(open("ls_orchid.gbk"), "genbank"),
                             lambda rec : seguid(rec.seq))
record = seguid_dict["MN/s0q9zDoCVEEc+k/IFwCNF2pY"]
print record.id
print record.description

```

That should have retrieved the record Z78532.1, the second entry in the file.

4.4 Writing Sequence Files

We've talked about using `Bio.SeqIO.parse()` for sequence input (reading files), and now we'll look at `Bio.SeqIO.write()` which is for sequence output (writing files). This is a function taking three arguments: some `SeqRecord` objects, a handle to write to, and a sequence format.

Here is an example, where we start by creating a few `SeqRecord` objects the hard way (by hand, rather than by loading them from a file):

```

from Bio.Seq import Seq
from Bio.SeqRecord import SeqRecord
from Bio.Alphabet import generic_protein

rec1 = SeqRecord(Seq("MMYQQGCFAGGTVLRRLAKDLAENNRGARVLVVCSEITAVTFRGPPSETHLDSMVGQALFGD" \
                    + "GAGAVIVGSDPDLSEVERPLYELVWTGATLLPDSEGAIDGHLREVGLTFHLLKDVPLGLISK" \
                    + "NIEKSLKEAFTPLGISDWNSTFWIAHPGGPAILDQVEAKLGLKKEEMRATREVLSEYGNM" \
                    + "SSAC", generic_protein),
                 id="gi|14150838|gb|AAK54648.1|AF376133_1",
                 description="chalcone synthase [Cucumis sativus]")

rec2 = SeqRecord(Seq("YPDYYFRITNREHKAELKEKFQRMCDKSMIKKRYMYLTEEILKENPSMCEYMAPSLDARQ" \
                    + "DMVVVEIPKLGKEAAVKAIKEWQQ", generic_protein),
                 id="gi|13919613|gb|AAK33142.1|",
                 description="chalcone synthase [Fragaria vesca subsp. bracteata]")

rec3 = SeqRecord(Seq("MVTVEEFRAQCAEGPATVMAIGTATPSNCVDQSTYPDYYFRITNSEHKVELKEKFKRMC" \
                    + "EKSMIKKRYMHLTEEILKENPNICAYMAPSLDARQDIVVVEVPKLGKEAAQKAIKEWQQP" \
                    + "KSKITHLVFCTTSGVDMPGCDYQLTKLLGLRPSVKRFMMYQQGCFAGGTVLRMAKDLAEN" \
                    + "NKGARVLVVCSEITAVTFRGPNTHLDSL VGQALFGDGAAAVIIGSDPIPEVERPLFELV" \
                    + "SAAQTLTPDSEGAIDGHLREVGLTFHLLKDVPLGLISKNIEKSLVEAFQPLGISDWNLSLFW" \
                    + "IAHPGGPAILDQVELKLGKQEKLKATR KVLVLSNYGNMSSACVLFILDEM RKASAKEGLGT" \
                    + "TGEGLWGVLFVFGPGLTVETVVLHSVAT", generic_protein),
                 id="gi|13925890|gb|AAK49457.1|",
                 description="chalcone synthase [Nicotiana tabacum]")

my_records = [rec1, rec2, rec3]

```

Now we have a list of `SeqRecord` objects, we'll write them to a FASTA format file:

```

from Bio import SeqIO
handle = open("my_example.faa", "w")
SeqIO.write(my_records, handle, "fasta")
handle.close()

```

And if you open this file in your favourite text editor it should look like this:

```
>gi|14150838|gb|AAK54648.1|AF376133_1 chalcone synthase [Cucumis sativus]
MMYQQGCFAGGTVLRRLAKDLAENNRGARVLVVCSEITAVTFRGPSETHLDSMVGQALFGD
GAGAVIVGSDPDLSEVERPLYELVWTGATLLPDSEGAIDGHLREVGLTFHLLKDVPLISK
NIEKSLKEAFTPLGISDWNSTFWIAHPGGPAILDQVEAKLGLKEEKMRATREVLSEYGNM
SSAC
>gi|13919613|gb|AAK33142.1| chalcone synthase [Fragaria vesca subsp. bracteata]
YPDYYFRITNREHKAELKEKFRMCDKSMIKKRYMYLTEEILKENPSMCEYMAPSLDARQ
DMVVVEIPKLGKEAAVKAKEWGQ
>gi|13925890|gb|AAK49457.1| chalcone synthase [Nicotiana tabacum]
MVTVEEFRAQAEGPATVMAIGTATPSNCVDQSTYYPDYFRITNSEHKVELKEKFKRMC
EKSMIKKRYMHLTEEILKENPNICAYMAPSLDARQDIVVVEVPKLGKEAAQKAIKEWGQP
KSKITHLVFCTTSGVDMPGCDYQLTKLLGLRPSVKRFMMYQQGCFAGGTVLRMAKDLAEN
NKGARVLVVCSEITAVTFRGPNTHLDSLQALFGDGAAAVIIGSDPIPEVERPLFELV
SAAQTLTPDSEGAIDGHLREVGLTFHLLKDVPLISKNIEKSLVEAFQPLGISDWNLSLFW
IAHPGGPAILDQVELKGLKQEKLKATRKVLSNYGNMSSACVLFILDEMRSASAKEGLGT
TGEGLWGVLFVGFPGPLTVETVVLHSVAT
```

4.4.1 Converting between sequence file formats

In previous example we used a list of `SeqRecord` objects as input to the `Bio.SeqIO.write()` function, but it will also accept a `SeqRecord` iterator like we get from `Bio.SeqIO.parse()` – this lets us do file conversion very succinctly. For this example we'll read in the GenBank format file [ls_orchid.gb](#) and write it out in FASTA format:

```
from Bio import SeqIO
in_handle = open("ls_orchid.gb", "r")
out_handle = open("my_example.fasta", "w")
SeqIO.write(SeqIO.parse(in_handle, "genbank"), out_handle, "fasta")
in_handle.close()
out_handle.close()
```

You *can* in fact do this in one line, by being lazy about closing the file handles. This is arguably bad style, but it is very concise:

```
from Bio import SeqIO
SeqIO.write(SeqIO.parse(open("ls_orchid.gb"), "genbank"), open("my_example.faa", "w"), "fasta")
```

4.4.2 Converting a file of sequences to their reverse complements

Suppose you had a file of nucleotide sequences, and you wanted to turn it into a file containing their reverse complement sequences. This time a little bit of work is required to transform the `SeqRecords` we get from our input file into something suitable for saving to our output file.

To start with, we'll use `Bio.SeqIO.parse()` to load some nucleotide sequences from a file, then print out their reverse complements using the `Seq` object's built in `.reverse_complement()` method (see Section 3.5):

```
from Bio import SeqIO
in_handle = open("ls_orchid.gb")
for record in SeqIO.parse(in_handle, "genbank") :
    print record.id
    print record.seq.reverse_complement().tostring()
in_handle.close()
```

Now, if we want to save these reverse complements to a file, we'll need to make `SeqRecord` objects. For this I think its most elegant to write our own function, where we can decide how to name our new records:

```
from Bio.SeqRecord import SeqRecord

def make_rc_record(record) :
    """Returns a new SeqRecord with the reverse complement sequence"""
    rc_rec = SeqRecord(seq = record.seq.reverse_complement(), \
                       id = "rc_" + record.id, \
                       name = "rc_" + record.name, \
                       description = "reverse complement")

    return rc_rec
```

We can then use this to turn the input records into reverse complement records ready for output. If you don't mind about having all the records in memory at once, then the python `map()` function is a very intuitive way of doing this:

```
from Bio import SeqIO

in_handle = open("ls_orchid.fasta", "r")
records = map(make_rc_record, SeqIO.parse(in_handle, "fasta"))
in_handle.close()

out_handle = open("rev_comp.fasta", "w")
SeqIO.write(records, out_handle, "fasta")
out_handle.close()
```

This is an excellent place to demonstrate the power of list comprehensions (added to Python 2.0) which in their simplest are a long-winded equivalent to using `map()`, like this:

```
records = [make_rc_record(rec) for rec in SeqIO.parse(in_handle, "fasta")]
```

Now list comprehensions have a nice trick up their sleeves, you can add a conditional statement:

```
records = [make_rc_record(rec) for rec in SeqIO.parse(in_handle, "fasta") if len(rec)<700]
```

That would create an in memory list of reverse complement records where the sequence length was under 700 base pairs. However, if you are using Python 2.4 or later, we can do exactly the same with a generator expression - but with the advantage that this does not create a list of all the records in memory at once:

```
records = (make_rc_record(rec) for rec in SeqIO.parse(in_handle, "fasta") if len(rec)<700)
```

If you like compact code, and don't mind being lax about closing file handles, we can reduce this to one long line:

```
from Bio import SeqIO
SeqIO.write((make_rc_record(rec) for rec in \
             SeqIO.parse(open("ls_orchid.fasta", "r"), "fasta") if len(rec) < 700), \
            open("rev_comp.fasta", "w"), "fasta")
```

Personally, I think the above snippet of code is a little too compact, and I find the following much easier to read:

```
from Bio import SeqIO
records = (make_rc_record(rec) for rec in \
           SeqIO.parse(open("ls_orchid.fasta", "r"), "fasta") \
           if len(rec) < 700)
SeqIO.write(records, open("rev_comp.fasta", "w"), "fasta")
```

or, for Python 2.3 or older,

```
from Bio import SeqIO
records = [make_rc_record(rec) for rec in \
          SeqIO.parse(open("ls_orchid.fasta", "r"), "fasta") \
          if len(rec) < 700]
SeqIO.write(records, open("rev_comp.fasta", "w"), "fasta")
```

Chapter 5

Sequence Alignment Input/Output

In this chapter we'll discuss the `Bio.AlignIO` module, which is very similar to the `Bio.SeqIO` module from the previous chapter, but deals with `Alignment` objects rather than `SeqRecord` objects. This is new interface in Biopython 1.46, which aims to provide a simple interface for working with assorted sequence alignment file formats in a uniform way.

Note that both `Bio.SeqIO` and `Bio.AlignIO` can read and write sequence alignment files. The appropriate choice will depend largely on what you want to do with the data.

5.1 Parsing or Reading Sequence Alignments

We have two functions for reading in sequence alignments, `Bio.AlignIO.read()` and `Bio.AlignIO.parse()` which following the convention introduced in `Bio.SeqIO` are for files containing one or multiple alignments respectively.

Using `Bio.AlignIO.parse()` will return an *iterator* which gives `Alignment` objects. Iterators are typically used in a for loop. Examples of situations where you will have multiple different alignments include resampled alignments from the PHYLIP tool `seqboot`, or multiple pairwise alignments from the EMBOSS tools `water` or `needle`, or Bill Pearson's FASTA tools.

However, in many situations you will be dealing with files which contain only a single alignment. In this case, you should use the `Bio.AlignIO.read()` function which return a single `Alignment` object.

Both functions expect two mandatory arguments:

1. The first argument is a *handle* to read the data from, typically an open file (see Section 12.1).
2. The second argument is a lower case string specifying sequence coformat. As in `Bio.SeqIO` we don't try and guess the file format for you! See <http://biopython.org/wiki/AlignIO> for a full listing of supported formats.

There is also an optional `seq_count` argument which is discussed in Section 5.1.3 below for dealing with ambiguous file formats which may contain more than one alignment.

5.1.1 Single Alignments

As an example, consider the following annotation rich protein alignment in the PFAM or Stockholm file format:

```
# STOCKHOLM 1.0
#=GS COATB_BPIKE/30-81 AC P03620.1
#=GS COATB_BPIKE/30-81 DR PDB; 1if1 ; 1-52;
```



```

#=GS Q9TOQ8_BPIKE/1-52 AC Q9TOQ8.1
#=GS COATB_BPI22/32-83 AC P15416.1
#=GS COATB_BPM13/24-72 AC P69541.1
#=GS COATB_BPM13/24-72 DR PDB; 2cpb ; 1-49;
#=GS COATB_BPM13/24-72 DR PDB; 2cps ; 1-49;
#=GS COATB_BPZJ2/1-49 AC P03618.1
#=GS Q9TOQ9_BPF1/1-49 AC Q9TOQ9.1
#=GS Q9TOQ9_BPF1/1-49 DR PDB; 1nh4 A; 1-49;
#=GS COATB_BPIF1/22-73 AC P03619.2
#=GS COATB_BPIF1/22-73 DR PDB; 1ifk ; 1-50;
COATB_BPIKE/30-81 AEPNAATNYATEAMDSLKTQAIDLISQTWPVVTTVVVAGLVIRLFKFFSSKA
#=GR COATB_BPIKE/30-81 SS -HHHHHHHHHHHHHHHH--HHHHHHHH--HHHHHHHHHHHHHHHHHHHHHHHHHHHH---
Q9TOQ8_BPIKE/1-52 AEPNAATNYATEAMDSLKTQAIDLISQTWPVVTTVVVAGLVIKLFFKFFVSR
COATB_BPI22/32-83 DGTSTATSYATEAMNSLKTQATDLIDQTWPVVTSVAVAGLAIRLFKFFSSKA
COATB_BPM13/24-72 AEGDDP...AKAAFNSLQASATEYIGYAWAMVVVIVGATIGIKLFFKFFSKA
#=GR COATB_BPM13/24-72 SS ---S-T...CHCHHHHCCCTCCCTTCHHHHHHHHHHHHHHHHHHHHHHHCTT--
COATB_BPZJ2/1-49 AEGDDP...AKAAFDSLQASATEYIGYAWAMVVVIVGATIGIKLFFKFFASKA
Q9TOQ9_BPF1/1-49 AEGDDP...AKAAFDSLQASATEYIGYAWAMVVVIVGATIGIKLFFKFFSKA
#=GR Q9TOQ9_BPF1/1-49 SS -----...-HHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHHH--
COATB_BPIF1/22-73 FAADDATSQAKAAFDSLTAQATEMSGYAWALVVLVVGATVGIKLFKFFVSR
#=GR COATB_BPIF1/22-73 SS XX-HHHH--HHHHHH--HHHHHH--HHHHHHHHHHHHHHHHHHHHHHHHHHHH---
#=GC SS_cons XHHHHHHHHHHHHHHHHCHHHHHHHHHCHHHHHHHHHHHHHHHHHHHHHHHHHHHHC--
#=GC seq_cons AEssss...AptAhDSLpspAT-hIu.sWshVsslVsAsluIKLFFKFFsSKA
//

```

This is the seed alignment for the Phage_Coat_Gp8 (PF05371) PFAM entry, downloaded as a compressed archive from <http://pfam.sanger.ac.uk/family/alignment/download/gzipped?acc=PF05371&alnType=seed>. We can load this file as follows (assuming it has been saved to disk as "PF05371_seed.sth" in the current working directory):

```

from Bio import AlignIO
alignment = AlignIO.read(open("PF05371_seed.sth"), "stockholm")
print alignment

```

This code will print out a summary of the alignment:

```

SingleLetterAlphabet() alignment with 7 rows and 52 columns
AEPNAATNYATEAMDSLKTQAIDLISQTWPVVTTVVVAGLVIRL...SKA COATB_BPIKE/30-81
AEPNAATNYATEAMDSLKTQAIDLISQTWPVVTTVVVAGLVIKL...SRA Q9TOQ8_BPIKE/1-52
DGTSTATSYATEAMNSLKTQATDLIDQTWPVVTSVAVAGLAIRL...SKA COATB_BPI22/32-83
AEGDDP---AKAAFNSLQASATEYIGYAWAMVVVIVGATIGIKL...SKA COATB_BPM13/24-72
AEGDDP---AKAAFDSLQASATEYIGYAWAMVVVIVGATIGIKL...SKA COATB_BPZJ2/1-49
AEGDDP---AKAAFDSLQASATEYIGYAWAMVVVIVGATIGIKL...SKA Q9TOQ9_BPF1/1-49
FAADDATSQAKAAFDSLTAQATEMSGYAWALVVLVVGATVGIKL...SRA COATB_BPIF1/22-73

```

You'll notice in the above output the sequences have been truncated. We could instead write our own code to format this as we please by iterating over the rows as SeqRecord objects:

```

from Bio import AlignIO
alignment = AlignIO.read(open("PF05371_seed.sth"), "stockholm")
print "Alignment length %i" % alignment.get_alignment_length()
for record in alignment :
    print "%s - %s" % (record.seq, record.id)

```

This will give the following output:

```
Alignment length 52
AEPNAATNYATEAMDSLKTQAIDLISQTWPVVTTVVVAGLVIRLFKKFSSKA - COATB_BPIKE/30-81
AEPNAATNYATEAMDSLKTQAIDLISQTWPVVTTVVVAGLVIKLFKKFVSRA - Q9TOQ8_BPIKE/1-52
DGTSTATSYATEAMNSLKTQATDLIDQTWPVVTSVAVAGLAIRLFKKFSSKA - COATB_BPI22/32-83
AEGDDP---AKAAFNSLQASATEYIGYAWAMVVVIVGATIGIKLFKKFTSKA - COATB_BPM13/24-72
AEGDDP---AKAAFDSLQASATEYIGYAWAMVVVIVGATIGIKLFKKFASKA - COATB_BPZJ2/1-49
AEGDDP---AKAAFDSLQASATEYIGYAWAMVVVIVGATIGIKLFKKFTSKA - Q9TOQ9_BPF1/1-49
FAADDATSQAKAAFDSLTAQATEMSGYAWALVVLVVGATVGIKLFKKFVSRA - COATB_BPIF1/22-73
```

Did you notice in the raw file above that several of the sequences include database cross-references to the PDB and the associated known secondary structure? Try this:

```
for record in alignment :
    if record.dbxrefs :
        print record.id, record.dbxrefs
```

giving:

```
COATB_BPIKE/30-81 ['PDB; 1if1 ; 1-52;']
COATB_BPM13/24-72 ['PDB; 2cpb ; 1-49;', 'PDB; 2cps ; 1-49;']
Q9TOQ9_BPF1/1-49 ['PDB; 1nh4 A; 1-49;']
COATB_BPIF1/22-73 ['PDB; 1ifk ; 1-50;']
```

To have a look at all the sequence annotation, try this:

```
for record in alignment :
    print record
```

Sanger provide a nice web interface at <http://pfam.sanger.ac.uk/family?acc=PF05371> which will actually let you download this alignment in several other formats. This is what the file looks like in the FASTA file format:

```
>COATB_BPIKE/30-81
AEPNAATNYATEAMDSLKTQAIDLISQTWPVVTTVVVAGLVIRLFKKFSSKA
>Q9TOQ8_BPIKE/1-52
AEPNAATNYATEAMDSLKTQAIDLISQTWPVVTTVVVAGLVIKLFKKFVSRA
>COATB_BPI22/32-83
DGTSTATSYATEAMNSLKTQATDLIDQTWPVVTSVAVAGLAIRLFKKFSSKA
>COATB_BPM13/24-72
AEGDDP...AKAAFNSLQASATEYIGYAWAMVVVIVGATIGIKLFKKFTSKA
>COATB_BPZJ2/1-49
AEGDDP...AKAAFDSLQASATEYIGYAWAMVVVIVGATIGIKLFKKFASKA
>Q9TOQ9_BPF1/1-49
AEGDDP...AKAAFDSLQASATEYIGYAWAMVVVIVGATIGIKLFKKFTSKA
>COATB_BPIF1/22-73
FAADDATSQAKAAFDSLTAQATEMSGYAWALVVLVVGATVGIKLFKKFVSRA
```

Assuming you download and save this as file "PF05371_seed.faa" then you can load it with almost exactly the same code:

```
from Bio import AlignIO
alignment = AlignIO.read(open("PF05371_seed.faa"), "fasta")
print alignment
```

All that has changed in this code is the filename and the format string. You'll get the same output as before, the sequences and record identifiers are the same. However, as you should expect, if you check each `SeqRecord` there is no annotation nor database cross-references because these are not included in the FASTA file format.

Note that rather than using the Sanger website, you could have used `Bio.AlignIO` to convert the original Stockholm format file into a FASTA file yourself (see below).

With any supported file format, you can load an alignment in exactly the same way just by changing the format string. For example, use “phylip” for PHYLIP files, “nexus” for NEXUS files or “emboss” for the alignments output by the EMBOSS tools. There is a full listing on the wiki page (<http://biopython.org/wiki/AlignIO>).

5.1.2 Multiple Alignments

The previous section focused on reading files containing a single alignment. In general however, files can contain more than one alignment, and to read these files we must use the `Bio.AlignIO.parse()` function.

Suppose you have a small alignment in PHYLIP format:

```
      5      6
Alpha   AACAAC
Beta    AACCCC
Gamma   ACCAAC
Delta   CCACCA
Epsilon CCAAAC
```

If you wanted to bootstrap a phylogenetic tree using the PHYLIP tools, one of the steps would be to create a set of many resampled alignments using the tool `bootseq`. This would give output something like this, which has been abbreviated for conciseness:

```
      5      6
Alpha   AAACCA
Beta    AAACCC
Gamma   ACCCCA
Delta   CCCAAC
Epsilon CCCAAA

      5      6
Alpha   AAACAA
Beta    AAACCC
Gamma   ACCCAA
Delta   CCCACC
Epsilon CCCAAA

      5      6
Alpha   AAAAAC
Beta    AAACCC
Gamma   AACAAC
Delta   CCCCCA
Epsilon CCCAAC
...

      5      6
Alpha   AAAACC
Beta    ACCCCC
Gamma   AAAACC
Delta   CCCCAA
Epsilon CAAACC
```

If you wanted to read this in using `Bio.AlignIO` you could use:

```
from Bio import AlignIO
alignments = AlignIO.parse(open("resampled.phy"), "phylip")
for alignment in alignments :
    print alignment
    print
```

This would give the following output, again abbreviated for display:

```
SingleLetterAlphabet() alignment with 5 rows and 6 columns
AAACCA Alpha
AAACCC Beta
ACCCCA Gamma
CCCAAC Delta
CCCAAA Epsilon
```

```
SingleLetterAlphabet() alignment with 5 rows and 6 columns
AAACAA Alpha
AAACCC Beta
ACCCAA Gamma
CCCACC Delta
CCCAAA Epsilon
```

```
SingleLetterAlphabet() alignment with 5 rows and 6 columns
AAAAAC Alpha
AAACCC Beta
AACCAAC Gamma
CCCCCA Delta
CCCAAC Epsilon
```

...

```
SingleLetterAlphabet() alignment with 5 rows and 6 columns
AAAACC Alpha
ACCCCC Beta
AAAACC Gamma
CCCCAA Delta
CAAACC Epsilon
```

As with the function `Bio.SeqIO.parse()`, using `Bio.AlignIO.parse()` returns an iterator. If you want to keep all the alignments in memory at once, which will allow you to access them in any order, then turn the iterator into a list:

```
from Bio import AlignIO
alignments = list(AlignIO.parse(open("resampled.phy"), "phylip"))
last_align = alignments[-1]
first_align = alignments[0]
```

5.1.3 Ambiguous Alignments

Many alignment file formats can explicitly store more than one alignment, and the division between each alignment is clear. However, when a general sequence file format has been used there is no such block

structure. The most common such situation is when alignments have been saved in the FASTA file format. For example consider the following:

```
>Alpha
ACTACGACTAGCTCAG--G
>Beta
ACTACCGCTAGCTCAGAAG
>Gamma
ACTACGGCTAGCACAGAAG
>Alpha
ACTACGACTAGCTCAGG--
>Beta
ACTACCGCTAGCTCAGAAG
>Gamma
ACTACGGCTAGCACAGAAG
```

This could be a single alignment containing six sequences (with repeated identifiers). Or, judging from the identifiers, this is probably two different alignments each with three sequences, which happen to all have the same length.

What about this next example?

```
>Alpha
ACTACGACTAGCTCAG--G
>Beta
ACTACCGCTAGCTCAGAAG
>Alpha
ACTACGACTAGCTCAGG--
>Gamma
ACTACGGCTAGCACAGAAG
>Alpha
ACTACGACTAGCTCAGG--
>Delta
ACTACGGCTAGCACAGAAG
```

Again, this could be a single alignment with six sequences. However this time based on the identifiers we might guess this is three pairwise alignments which by chance have all got the same lengths.

This final example is similar:

```
>Alpha
ACTACGACTAGCTCAG--G
>XXX
ACTACCGCTAGCTCAGAAG
>Alpha
ACTACGACTAGCTCAGG
>YYY
ACTACGGCAAGCACAGG
>Alpha
--ACTACGAC--TAGCTCAGG
>ZZZ
GGACTACGACAATAGCTCAGG
```

In this third example, because of the differing lengths, this cannot be treated as a single alignment containing all six records. However, it could be three pairwise alignments.

Clearly trying to store more than one alignment in a FASTA file is not ideal. However, if you are forced to deal with these as input files `Bio.AlignIO` can cope with the most common situation where all the alignments have the same number of records. One example of this is a collection of pairwise alignments, which can be produced by the EMBOSS tools `needle` and `water` – although in this situation, `Bio.AlignIO` should be able to understand their native output using “emboss” as the format string.

To interpret these FASTA examples as several separate alignments, we can use `Bio.AlignIO.parse()` with the optional `seq_count` argument which specifies how many sequences are expected in each alignment (in these examples, 3, 2 and 2 respectively). For example, using the third example as the input data:

```
for alignment in AlignIO.parse(handle, "fasta", seq_count=2) :
    print "Alignment length %i" % alignment.get_alignment_length()
    for record in alignment :
        print "%s - %s" % (record.seq, record.id)
    print
```

giving:

```
Alignment length 19
ACTACGACTAGCTCAG--G - Alpha
ACTACCGCTAGCTCAGAAG - XXX
```

```
Alignment length 17
ACTACGACTAGCTCAGG - Alpha
ACTACGGCAAGCACAGG - YYY
```

```
Alignment length 21
--ACTACGAC--TAGCTCAGG - Alpha
GGACTACGACAATAGCTCAGG - ZZZ
```

Using `Bio.AlignIO.read()` or `Bio.AlignIO.parse()` without the `seq_count` argument would give a single alignment containing all six records for the first two examples. For the third example, an exception would be raised because the lengths differ preventing them being turned into a single alignment.

If the file format itself has a block structure allowing `Bio.AlignIO` to determine the number of sequences in each alignment directly, then the `seq_count` argument is not needed. If it is supplied, and doesn't agree with the file contents, an error is raised.

Note that this optional `seq_count` argument assumes each alignment in the file has the same number of sequences. Hypothetically you may come across stranger situations, for example a FASTA file containing several alignments each with a different number of sequences – although I would love to hear of a real world example of this. Assuming you cannot get the data in a nicer file format, there is no straight forward way to deal with this using `Bio.AlignIO`. In this case, you could consider reading in the sequences themselves using `Bio.SeqIO` and batching them together to create the alignments as appropriate.

5.2 Writing Alignments

We've talked about using `Bio.AlignIO.read()` and `Bio.AlignIO.parse()` for alignment input (reading files), and now we'll look at `Bio.AlignIO.write()` which is for alignment output (writing files). This is a function taking three arguments: some `Alignment` objects, a handle to write to, and a sequence format.

Here is an example, where we start by creating a few `Alignment` objects the hard way (by hand, rather than by loading them from a file):

```
from Bio.Align.Generic import Alignment
from Bio.Alphabet import IUPAC, Gapped
```

```

alphabet = Gapped(IUPAC.unambiguous_dna)

align1 = Alignment(alphabet)
align1.add_sequence("Alpha", "ACTGCTAGCTAG")
align1.add_sequence("Beta", "ACT-CTAGCTAG")
align1.add_sequence("Gamma", "ACTGCTAGDTAG")

align2 = Alignment(alphabet)
align2.add_sequence("Delta", "GTCAGC-AG")
align2.add_sequence("Epsilon", "GACAGCTAG")
align2.add_sequence("Zeta", "GTCAGCTAG")

align3 = Alignment(alphabet)
align3.add_sequence("Eta", "ACTAGTACAGCTG")
align3.add_sequence("Theta", "ACTAGTACAGCT-")
align3.add_sequence("Iota", "-CTACTACAGGTG")

my_alignments = [align1, align2, align3]

```

Now we have a list of `Alignment` objects, we'll write them to a PHYLIP format file:

```

from Bio import AlignIO
handle = open("my_example.phy", "w")
SeqIO.write(my_alignments, handle, "phylip")
handle.close()

```

And if you open this file in your favourite text editor it should look like this:

```

3 12
Alpha      ACTGCTAGCT AG
Beta       ACT-CTAGCT AG
Gamma      ACTGCTAGDT AG
3 9
Delta      GTCAGC-AG
Epsilon    GACAGCTAG
Zeta       GTCAGCTAG
3 13
Eta        ACTAGTACAG CTG
Theta      ACTAGTACAG CT-
Iota       -CTACTACAG GTG

```

Its more common to want to load an existing alignment, and save that, perhaps after some simple manipulation like removing certain rows or columns.

5.2.1 Converting between sequence alignment file formats

Converting between sequence alignment file formats with `Bio.AlignIO` works in the same way as converting between sequence file formats with `Bio.SeqIO` – we load generally the alignment(s) using `Bio.AlignIO.parse()` and then save them using the `Bio.AlignIO.write()`.

For this example, we'll load the PFAM/Stockholm format file used earlier and save it as a Clustal W format file:

```

from Bio import AlignIO
alignments = AlignIO.parse(open("PF05371_seed.sth"), "stockholm")
handle = open("PF05371_seed.aln", "w")
AlignIO.write(alignments, handle, "clustal")
handle.close()

```

The `Bio.AlignIO.write()` function expects to be given multiple alignment objects. In the example above we gave it the alignment iterator returned by `Bio.AlignIO.parse()`.

In this case, we know there is only one alignment in the file so we could instead have used `Bio.AlignIO.read()` but notice we have to pass this alignment to `Bio.AlignIO.write()` as a single element list:

```

from Bio import AlignIO
alignment = AlignIO.read(open("PF05371_seed.sth"), "stockholm")
handle = open("PF05371_seed.aln", "w")
AlignIO.write([alignment], handle, "clustal")
handle.close()

```

Either way, you should end up with the same new Clustal W format file “PF05371_seed.aln” with the following content:

CLUSTAL X (1.81) multiple sequence alignment

```

COATB_BPIKE/30-81      AEPNAATNYATEAMDSLKTQAIDLISQTPVVTTVVVAGLVIRLFKFFSS
Q9TOQ8_BPIKE/1-52     AEPNAATNYATEAMDSLKTQAIDLISQTPVVTTVVVAGLVIKLFKFFVS
COATB_BPI22/32-83    DGTSTATSYATEAMNSLKTQATDLIDQTPVVTSVAVAGLAIRLFKFFSS
COATB_BPM13/24-72    AEGDDP---AKAAFNSLQASATEYIGYAWAMVVVIVGATIGIKLFKFFTS
COATB_BPZJ2/1-49     AEGDDP---AKAAFDSLQASATEYIGYAWAMVVVIVGATIGIKLFKFFAS
Q9TOQ9_BPF1/1-49     AEGDDP---AKAAFDSLQASATEYIGYAWAMVVVIVGATIGIKLFKFFTS
COATB_BPIF1/22-73    FAADDATSQAKAAFDSLTAQATEMSGYAWALVVLVVGATVGIKLFKFFVS

```

```

COATB_BPIKE/30-81      KA
Q9TOQ8_BPIKE/1-52     RA
COATB_BPI22/32-83    KA
COATB_BPM13/24-72    KA
COATB_BPZJ2/1-49     KA
Q9TOQ9_BPF1/1-49     KA
COATB_BPIF1/22-73    RA

```

Alternatively, you could make a PHYLIP format file which we'll name “PF05371_seed.phy”:

```

from Bio import AlignIO
alignment = AlignIO.read(open("PF05371_seed.sth"), "stockholm")
handle = open("PF05371_seed.phy", "w")
AlignIO.write([alignment], handle, "phylip")
handle.close()

```

This time the output looks like this:

```

7 52
COATB_BPIK AEPNAATNYA TEAMDSLKTQ AIDLISQTPV VTTVVVAGL VIRLFKFFSS
Q9TOQ8_BPI AEPNAATNYA TEAMDSLKTQ AIDLISQTPV VTTVVVAGL VIKLFKFFVS
COATB_BPI2 DGTSTATSYA TEAMNSLKTQ ATDLIDQTPV VTSVAVAGL AIRLFKFFSS

```



```

COATB_BPM1 AEGDDP---A KAAFNSLQAS ATEYIGYAWA MVVVIVGATI GIKLFFKKFSTS
COATB_BPZJ AEGDDP---A KAAFDSLQAS ATEYIGYAWA MVVVIVGATI GIKLFFKKFAS
Q9TOQ9_BPF AEGDDP---A KAAFDSLQAS ATEYIGYAWA MVVVIVGATI GIKLFFKKFSTS
COATB_BPIF FAADDATSQA KAAFDSLTAQ ATEMSGYAWA LVVLVVGATV GIKLFFKKFVS

```

```

KA
RA
KA
KA
KA
KA
RA

```

One of the big handicaps of the PHYLIP alignment file format is that the sequence identifiers are strictly truncated at ten characters. In this example, as you can see the resulting names are still unique - but they are not very readable. In this particular case, there is no clear way to compress the identifiers, but for the sake of argument you may want to assign your own names or numbering system. This following bit of code manipulates the record identifiers before saving the output:

```

from Bio import AlignIO
alignment = AlignIO.read(open("PF05371_seed.sth"), "stockholm")
name_mapping = {}
for i, record in enumerate(alignment) :
    name_mapping[i] = record.id
    record.id = "seq%i" % i
print name_mapping

handle = open("PF05371_seed.phy","w")
AlignIO.write([alignment], handle, "phylip")
handle.close()

```

This code using a python dictionary to record a simple mapping from the new sequence system to the original identifier:

```
{0: 'COATB_BPIKE/30-81', 1: 'Q9TOQ8_BPIKE/1-52', 2: 'COATB_BPI22/32-83', ...}
```

Here is the new PHYLIP format output:

```

7 52
seq0      AEPNAATNYA TEAMDSLKTQ AIDLISQTP VTTTIVVAGL VIRLFFKKFSS
seq1      AEPNAATNYA TEAMDSLKTQ AIDLISQTP VTTTIVVAGL VIKLFFKKFVS
seq2      DGTSTATSYA TEAMNSLKTQ ATDLIDQTP VVTSVAVAGL AIRLFFKKFSS
seq3      AEGDDP---A KAAFNSLQAS ATEYIGYAWA MVVVIVGATI GIKLFFKKFSTS
seq4      AEGDDP---A KAAFDSLQAS ATEYIGYAWA MVVVIVGATI GIKLFFKKFAS
seq5      AEGDDP---A KAAFDSLQAS ATEYIGYAWA MVVVIVGATI GIKLFFKKFSTS
seq6      FAADDATSQA KAAFDSLTAQ ATEMSGYAWA LVVLVVGATV GIKLFFKKFVS

```

```

KA
RA
KA
KA
KA
KA
RA

```

In general, because of the identifier limitation, working with PHYLIP file formats shouldn't be your first choice. Using the PFAM/Stockholm format on the other hand allows you to record a lot of additional annotation too.

Chapter 6

BLAST

Hey, everybody loves BLAST right? I mean, geez, how can get it get any easier to do comparisons between one of your sequences and every other sequence in the known world? But, of course, this section isn't about how cool BLAST is, since we already know that. It is about the problem with BLAST – it can be really difficult to deal with the volume of data generated by large runs, and to automate BLAST runs in general.

Fortunately, the Biopython folks know this only too well, so they've developed lots of tools for dealing with BLAST and making things much easier. This section details how to use these tools and do useful things with them.

Dealing with BLAST can be split up into two steps, both of which can be done from within Biopython. Firstly, running BLAST for your query sequence(s), and getting some output. Secondly, parsing the BLAST output in python for further analysis. We'll start by talking about running the BLAST command line tools locally, and then discuss running BLAST via the web.

6.1 Running BLAST locally

Running BLAST locally (as opposed to over the internet, see Section 6.2) has two advantages:

- Local BLAST may be faster than BLAST over the internet;
- Local BLAST allows you to make your own database to search for sequences against.

Dealing with proprietary or unpublished sequence data can be another reason to run BLAST locally. You may not be allowed to redistribute the sequences, so submitting them to the NCBI as a BLAST query would not be an option.

Biopython provides lots of nice code to enable you to call local BLAST executables from your scripts, and have full access to the many command line options that these executables provide. You can obtain local BLAST precompiled for a number of platforms at <ftp://ftp.ncbi.nlm.nih.gov/blast/executables/>, or can compile it yourself in the NCBI toolbox (<ftp://ftp.ncbi.nlm.nih.gov/toolbox/>).

The code for dealing with local BLAST is found in `Bio.Blast.NCBISTandalone`, specifically in the functions `blastall`, `blastpgp` and `rpsblast`, which correspond with the BLAST executables that their names imply.

Let's use these functions to run a `blastall` against a local database and return the results. First, we want to set up the paths to everything that we'll need to do the BLAST. What we need to know is the path to the database (which should have been prepared using `formatdb`, see <ftp://ftp.ncbi.nlm.nih.gov/blast/documents/formatdb.html>) to search against, the path to the file we want to search, and the path to the `blastall` executable.

On Linux or Mac OS X your paths might look like this:

```
>>> my_blast_db = "/home/mdehoon/Data/Genomes/Databases/bsubtilis"
# I used formatdb to create a BLAST database named bsubtilis
# (for Bacillus subtilis) consisting of the following three files:
# /home/mdehoon/Data/Genomes/Databases/bsubtilis.nhr
# /home/mdehoon/Data/Genomes/Databases/bsubtilis.nin
# /home/mdehoon/Data/Genomes/Databases/bsubtilis.nsq
```

```
>>> my_blast_file = "m_cold.fasta"
# A FASTA file with the sequence I want to BLAST
```

```
>>> my_blast_exe = "/usr/local/blast/bin/blastall"
# The name of my BLAST executable
```

while on Windows you might have something like this:

```
>>> my_blast_db = r"C:\Blast\Data\bsubtilis"
# Assuming you used formatdb to create a BLAST database named bsubtilis
# (for Bacillus subtilis) consisting of the following three files:
# C:\Blast\Data\bsubtilis\bsubtilis.nhr
# C:\Blast\Data\bsubtilis\bsubtilis.nin
# C:\Blast\Data\bsubtilis\bsubtilis.nsq
>>> my_blast_file = "m_cold.fasta"
>>> my_blast_exe = r"C:\Blast\bin\blastall.exe"
```

The FASTA file used in this example is available [here](#) as well as [online](#).

Now that we've got that all set, we are ready to run the BLAST and collect the results. We can do this with two lines:

```
>>> from Bio.Blast import NCBIStandalone
>>> result_handle, error_handle = NCBIStandalone.blastall(my_blast_exe, "blastn",
                                                         my_blast_db, my_blast_file)
```

Note that the Biopython interfaces to local blast programs returns two values. The first is a handle to the blast output, which is ready to either be saved or passed to a parser. The second is the possible error output generated by the blast command. See Section 12.1 for more about handles.

The error info can be hard to deal with, because if you try to do a `error_handle.read()` and there was no error info returned, then the `read()` call will block and not return, locking your script. In my opinion, the best way to deal with the error is only to print it out if you are not getting `result_handle` results to be parsed, but otherwise to leave it alone.

This command will generate BLAST output in XML format, as that is the format expected by the XML parser, described in Section 6.4. For plain text output, use the `align_view='0'` keyword. To parse text output instead of XML output, see the Section 6.6 below. However, parsing text output is not recommended, as the BLAST plain text output changes frequently, breaking our parsers.

If you are interested in saving your results to a file before parsing them, see Section 6.3. To find out how to parse the BLAST results, go to Section 6.4

6.2 Running BLAST over the Internet

The first step in automating BLASTing is to make everything accessible from Python scripts. So, Biopython contains code that allows you to run the WWW version of BLAST (<http://www.ncbi.nlm.nih.gov/BLAST/>) directly from your Python scripts. This is very nice, especially since otherwise BLAST can be a real pain to deal with from scripts, especially with the whole BLAST queue thing and the separate results page.

The code to deal with the WWW version of BLAST is found in the `Bio.Blast.NCBIWWW` module, and the `qblast` function. Let's say we want to BLAST info we have in a FASTA formatted file against the database. First, we need to get the info in the FASTA file.

The easiest way to do this is to use the `Bio.SeqIO` module (see Chapter 4). In this example we'll use the `Bio.SeqIO.read` function to turn a FASTA file containing a single entry into a `SeqRecord` object:

```
>>> from Bio import SeqIO
>>> record = SeqIO.read(open("m_cold.fasta"), format="fasta")
```

Now we can take the sequence as a plain string from the `SeqRecord` and run BLAST on it. The code to do the simplest possible BLAST (a simple `blastn` of the FASTA file against all of the non-redundant databases) is:

```
>>> from Bio.Blast import NCBIWWW
>>> result_handle = NCBIWWW.qblast("blastn", "nr", record.seq.tostring())
```

The first three arguments to our `NCBIWWW.qblast` function are non-optional:

- The first argument is the blast program to use for the search, as a lower case string. The options and descriptions of the programs are available at http://www.ncbi.nlm.nih.gov/BLAST/blast_program.html. Currently `qblast` only works with `blastn`, `blastp`, `blastx`, `tblast` and `tblastx`.
- The second argument specifies the databases to search against. Again, the options for this are available on the NCBI web pages at http://www.ncbi.nlm.nih.gov/BLAST/blast_databases.html.
- The third argument is a string containing your query sequence. This can either be the sequence itself (as above), the sequence in fasta format, or an identifier like a GI number.

The `qblast` function also take a number of other option arguments which are basically analogous to the different parameters you can set on the BLAST web page. We'll just highlight a few of them here:

- The `qblast` function can return the BLAST results in various formats, which you can choose with the optional `format_type` keyword: "HTML", "Text", "ASN.1", or "XML". The default is "XML", as that is the format expected by the parser, described in section 6.4 below.
- The argument `expect` sets the expectation or e-value threshold.

For more about the optional BLAST arguments, we refer you to the NCBI's own documentation, or that built into Biopython:

```
>>> from Bio.Blast import NCBIWWW
>>> help(NCBIWWW.qblast)
```

After you have set the search options, you are all ready to BLAST. Biopython takes care of worrying about when the results are available, and will pause until it can get the results and return them.

6.3 Saving BLAST output

Before parsing the results, it is often useful to save them into a file so that you can use them later without having to go back and re-blast everything. I find this especially useful when debugging my code that extracts info from the BLAST files, but it could also be useful just for making backups of things you've done.

If you don't want to save the BLAST output, you can skip to section 6.4. If you do, read on.

We need to be a bit careful since we can use `result_handle.read()` to read the BLAST output only once – calling `result_handle.read()` again returns an empty string. First, we use `read()` and store all of the information from the handle into a string:

```
>>> blast_results = result_handle.read()
```

Next, we save this string in a file:

```
>>> save_file = open("my_blast.xml", "w")
>>> save_file.write(blast_results)
>>> save_file.close()
```

After doing this, the results are in the file `my_blast.xml` and the variable `blast_results` contains the BLAST results in a string form. However, the `parse` function of the BLAST parser (described in 6.4) takes a file-handle-like object, not a plain string. To get a handle, there are two things you can do:

- Use the Python standard library module `cStringIO`. The following code will turn the plain string into a handle, which we can feed directly into the BLAST parser:

```
>>> import cStringIO
>>> result_handle = cStringIO.StringIO(blast_results)
```

- Open the saved file for reading. Duh.

```
>>> result_handle = open("my_blast.xml")
```

Now that we've got the BLAST results, we are ready to do something with them, so this leads us right into the parsing section.

6.4 Parsing BLAST output

As mentioned above, BLAST can generate output in various formats, such as XML, HTML, and plain text. Originally, Biopython had a parser for BLAST plain text and HTML output, as these were the only output formats supported by BLAST. Unfortunately, the BLAST output in these formats kept changing, each time breaking the Biopython parsers. As keeping up with changes in BLAST became a hopeless endeavor, especially with users running different BLAST versions, we now recommend to parse the output in XML format, which can be generated by recent versions of BLAST. Not only is the XML output more stable than the plain text and HTML output, it is also much easier to parse automatically, making Biopython a whole lot more stable.

Though deprecated, the parsers for BLAST output in plain text or HTML output are still available in Biopython (see Section 6.6). Use them at your own risk: they may or may not work, depending on which BLAST version you're using.

You can get BLAST output in XML format in various ways. For the parser, it doesn't matter how the output was generated, as long as it is in the XML format.

- You can use Biopython to run BLAST locally, as described in section 6.1.
- You can use Biopython to run BLAST over the internet, as described in section 6.2.
- You can do the BLAST search yourself on the NCBI site through your web browser, and then save the results. You need to choose XML as the format in which to receive the results, and save the final BLAST page you get (you know, the one with all of the interesting results!) to a file.
- You can also run BLAST locally without using Biopython, and save the output in a file. Again, you need to choose XML as the format in which to receive the results.

The important point is that you do not have to use Biopython scripts to fetch the data in order to be able to parse it.

Doing things in one of these ways, you then need to get a handle to the results. In Python, a handle is just a nice general way of describing input to any info source so that the info can be retrieved using `read()` and `readline()` functions. This is the type of input the BLAST parser (and most other Biopython parsers) take.

If you followed the code above for interacting with BLAST through a script, then you already have `result_handle`, the handle to the BLAST results. For example:

```
>>> from Bio import SeqIO
>>> record = SeqIO.read(open("m_cold.fasta"), format="fasta")
>>> from Bio.Blast import NCBIWWW
>>> result_handle = NCBIWWW.qblast("blastn", "nr", record.seq.tostring())
```

If instead you ran BLAST some other way, and have the BLAST output (in XML format) in the file `my_blast.xml`, all you need to do is to open the file for reading:

```
>>> result_handle = open("my_blast.xml")
```

Now that we've got a handle, we are ready to parse the output. The code to parse it is really quite small:

```
>>> from Bio.Blast import NCBIXML
>>> blast_records = NCBIXML.parse(result_handle)
```

To understand what `NCBIXML.parse` returns, there are two things that you need to keep in mind:

- The BLAST output may contain the output of more than one BLAST search. This will for example be the case if you ran BLAST locally on a Fasta file containing more than one sequence. For each sequence, the BLAST parser will return one BLAST record.
- The BLAST output may therefore be huge.

To be able to handle these situations, `NCBIXML.parse` returns an iterator (just like `Bio.SeqIO.parse`). In plain English, an iterator allows you to step through the BLAST output, retrieving BLAST records one by one for each BLAST search:

```
>>> blast_record = blast_records.next()
# ... do something with blast_record
>>> blast_record = blast_records.next()
# ... do something with blast_record
>>> blast_record = blast_records.next()
# ... do something with blast_record
>>> blast_record = blast_records.next()
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
StopIteration
# No further records
```

Or, you can use a for-loop:

```
>>> for blast_record in blast_records:
...     # Do something with blast_record
```

Note though that you can step through the BLAST records only once. Usually, from each BLAST record you would save the information that you are interested in. If you want to save all returned BLAST records, you can convert the iterator into a list:




Figure 6.1: Class diagram for the Blast Record class representing all of the info in a BLAST report

The PSIBlast record object is similar, but has support for the rounds that are used in the iteration steps of PSIBlast. The class diagram for PSIBlast is shown in Figure 6.2.

6.6 Deprecated BLAST parsers

Older versions of Biopython had parsers for BLAST output in plain text or HTML format. Over the years, we discovered that it is very hard to maintain these parsers in working order. Basically, any small change to the BLAST output in newly released BLAST versions tends to cause the plain text and HTML parsers to break. We therefore recommend parsing BLAST output in XML format, as described in section 6.4. Whereas the plain text and HTML parsers are still available in Biopython, use them at your own risk. They may or may not work, depending on which BLAST versions you're using.

6.6.1 Parsing plain-text BLAST output

The plain text BLAST parser is located in `Bio.Blast.NCBISTandalone`.

As with the XML parser, we need to have a handle object that we can pass to the parser. The handle must implement the `readline()` method and do this properly. The common ways to get such a handle are to either use the provided `blastall` or `blastpgp` functions to run the local blast, or to run a local blast via the command line, and then do something like the following:

```
>>> result_handle = open("my_file_of_blast_output.txt")
```

Well, now that we've got a handle (which we'll call `result_handle`), we are ready to parse it. This can be done with the following code:

```
>>> from Bio.Blast import NCBISTandalone
>>> blast_parser = NCBISTandalone.BlastParser()
>>> blast_record = blast_parser.parse(result_handle)
```

This will parse the BLAST report into a Blast Record class (either a Blast or a PSIBlast record, depending on what you are parsing) so that you can extract the information from it. In our case, let's just use print out a quick summary of all of the alignments greater than some threshold value.

```
>>> E_VALUE_THRESH = 0.04
>>> for alignment in b_record.alignments:
...     for hsp in alignment.hsps:
...         if hsp.expect < E_VALUE_THRESH:
...             print '****Alignment****'
...             print 'sequence:', alignment.title
...             print 'length:', alignment.length
...             print 'e value:', hsp.expect
...             print hsp.query[0:75] + '...'
...             print hsp.match[0:75] + '...'
...             print hsp.sbjct[0:75] + '...'
```

If you also read the section 6.4 on parsing BLAST XML output, you'll notice that the above code is identical to what is found in that section. Once you parse something into a record class you can deal with it independent of the format of the original BLAST info you were parsing. Pretty snazzy!

Sure, parsing one record is great, but I've got a BLAST file with tons of records – how can I parse them all? Well, fear not, the answer lies in the very next section.




Figure 6.2: Class diagram for the PSIBlast Record class.

6.6.2 Parsing a file full of BLAST runs

Of course, local blast is cool because you can run a whole bunch of sequences against a database and get back a nice report on all of it. So, Biopython definitely has facilities to make it easy to parse humongous files without memory problems.

We can do this using the blast iterator. To set up an iterator, we first set up a parser, to parse our blast reports in Blast Record objects:

```
>>> from Bio.Blast import NCBIStandalone
>>> blast_parser = NCBIStandalone.BlastParser()
```

Then we will assume we have a handle to a bunch of blast records, which we'll call `result_handle`. Getting a handle is described in full detail above in the blast parsing sections.

Now that we've got a parser and a handle, we are ready to set up the iterator with the following command:

```
>>> blast_iterator = NCBIStandalone.Iterator(blast_handle, blast_parser)
```

The second option, the parser, is optional. If we don't supply a parser, then the iterator will just return the raw BLAST reports one at a time.

Now that we've got an iterator, we start retrieving blast records (generated by our parser) using `next()`:

```
>>> blast_record = blast_iterator.next()
```

Each call to `next` will return a new record that we can deal with. Now we can iterate through this records and generate our old favorite, a nice little blast report:

```
>>> for b_record in b_iterator :
...     E_VALUE_THRESH = 0.04
...     for alignment in b_record.alignments:
...         for hsp in alignment.hsps:
...             if hsp.expect < E_VALUE_THRESH:
...                 print '****Alignment****'
...                 print 'sequence:', alignment.title
...                 print 'length:', alignment.length
...                 print 'e value:', hsp.expect
...                 if len(hsp.query) > 75:
...                     dots = '...'
...                 else:
...                     dots = ''
...                 print hsp.query[0:75] + dots
...                 print hsp.match[0:75] + dots
...                 print hsp.sbjct[0:75] + dots
```

The iterator allows you to deal with huge blast records without any memory problems, since things are read in one at a time. I have parsed tremendously huge files without any problems using this.

6.6.3 Finding a bad record somewhere in a huge file

One really ugly problem that happens to me is that I'll be parsing a huge blast file for a while, and the parser will bomb out with a `ValueError`. This is a serious problem, since you can't tell if the `ValueError` is due to a parser problem, or a problem with the BLAST. To make it even worse, you have no idea where the parse failed, so you can't just ignore the error, since this could be ignoring an important data point.

We used to have to make a little script to get around this problem, but the `Bio.Blast` module now includes a `BlastErrorParser` which really helps make this easier. The `BlastErrorParser` works very

similar to the regular `BlastParser`, but it adds an extra layer of work by catching `ValueErrors` that are generated by the parser, and attempting to diagnose the errors.

Let's take a look at using this parser – first we define the file we are going to parse and the file to write the problem reports to:

```
>>> import os
>>> blast_file = os.path.join(os.getcwd(), "blast_out", "big_blast.out")
>>> error_file = os.path.join(os.getcwd(), "blast_out", "big_blast.problems")
```

Now we want to get a `BlastErrorParser`:

```
>>> from Bio.Blast import NCBIStandalone
>>> error_handle = open(error_file, "w")
>>> blast_error_parser = NCBIStandalone.BlastErrorParser(error_handle)
```

Notice that the parser take an optional argument of a handle. If a handle is passed, then the parser will write any blast records which generate a `ValueError` to this handle. Otherwise, these records will not be recorded.

Now we can use the `BlastErrorParser` just like a regular blast parser. Specifically, we might want to make an iterator that goes through our blast records one at a time and parses them with the error parser:

```
>>> result_handle = open(blast_file)
>>> iterator = NCBIStandalone.Iterator(result_handle, blast_error_parser)
```

We can read these records one a time, but now we can catch and deal with errors that are due to problems with Blast (and not with the parser itself):

```
>>> try:
...     next_record = iterator.next()
... except NCBIStandalone.LowQualityBlastError, info:
...     print "LowQualityBlastError detected in id %s" % info[1]
```

The `.next()` method is normally called indirectly via a `for`-loop. Right now the `BlastErrorParser` can generate the following errors:

- **ValueError** – This is the same error generated by the regular `BlastParser`, and is due to the parser not being able to parse a specific file. This is normally either due to a bug in the parser, or some kind of discrepancy between the version of BLAST you are using and the versions the parser is able to handle.
- **LowQualityBlastError** – When BLASTing a sequence that is of really bad quality (for example, a short sequence that is basically a stretch of one nucleotide), it seems that Blast ends up masking out the entire sequence and ending up with nothing to parse. In this case it will produce a truncated report that causes the parser to generate a `ValueError`. `LowQualityBlastError` is reported in these cases. This error returns an `info` item with the following information:
 - `item[0]` – The error message
 - `item[1]` – The id of the input record that caused the error. This is really useful if you want to record all of the records that are causing problems.

As mentioned, with each error generated, the `BlastErrorParser` will write the offending record to the specified `error_handle`. You can then go ahead and look and these and deal with them as you see fit. Either you will be able to debug the parser with a single blast report, or will find out problems in your blast runs. Either way, it will definitely be a useful experience!

Hopefully the `BlastErrorParser` will make it much easier to debug and deal with large Blast files.

6.7 Dealing with PSIBlast

We should write some stuff to make it easier to deal directly with PSIBlast from scripts (i. e. output the align file in the proper format from an alignment). I need to look at PSIBlast more and come up with some good ways of going this...

Chapter 7

Accessing NCBI's Entrez databases

Entrez (<http://www.ncbi.nlm.nih.gov/Entrez>) is a data retrieval system that provides users access to NCBI's databases such as PubMed, GenBank, GEO, and many others. You can access Entrez from a web browser to manually enter queries, or you can use Biopython's `Bio.Entrez` module for programmatic access to Entrez. The latter allows you for example to search PubMed or download GenBank records from within a Python script.

The `Bio.Entrez` module makes use of the Entrez Programming Utilities, consisting of eight tools that are described in detail on NCBI's page at <http://www.ncbi.nlm.nih.gov/entrez/utils/>. Each of these tools corresponds to one Python function in the `Bio.Entrez` module, as described in the sections below. This module makes sure that the correct URL is used for the queries, and that not more than one request is made every three seconds, as required by NCBI.

The output returned by the Entrez Programming Utilities is typically in XML format. To parse such output, you have several options:

1. Use `Bio.Entrez`'s parser to parse the XML output into a Python object;
2. Use the DOM (Document Object Model) parser in Python's standard library;
3. Use the SAX (Simple API for XML) parser in Python's standard library;
4. Read the XML output as raw text, and parse it by string searching and manipulation.

For the DOM and SAX parsers, see the Python documentation. The parser in `Bio.Entrez` is discussed below.

For sequence databases, the Entrez Programming Utilities can also generate output in other formats (such as the Fasta and GenBank file format). This can then be parsed into a `SeqRecord` using `Bio.SeqIO` (see Chapter 4, and the example below).

7.1 Entrez Guidelines

Before using Biopython to access the NCBI's online resources (via `Bio.Entrez` or some of the other modules), please read the [NCBI's Entrez User Requirements](#). If the NCBI finds you are abusing their systems, they can and will ban your access!

To paraphrase:

- For any series of more than 100 requests, do this at weekends or outside USA peak times. This is up to you to obey.
- Use the <http://eutils.ncbi.nlm.nih.gov> address, not the standard NCBI Web address. Biopython uses this web address.

- Make no more than one request every 3 seconds. This is automatically enforced by Biopython.
- Use the optional email parameter so the NCBI can contact you if there is a problem. In the examples below, we have used `email="A.N.Other@example.com"` for illustrative purposes. The `example.com` address is a reserved domain name specifically for documentation (RFC 2606). Please DO NOT use a random email – it's better not to give an email at all.
- If you are using Biopython within some larger software suite, use the `tool` parameter to specify this. The `tool` parameter will default to Biopython.

For large queries, the NCBI also recommend using their session history feature (the WebEnv session cookie string). This is only slightly more complicated.

In conclusion, be sensible with your usage levels. If you plan to download lots of data, consider other options. For example, if you want easy access to all the human genes, consider fetching each chromosome by FTP as a GenBank file, and importing these into your own BioSQL database (see Section 9.5).

7.2 EInfo: Obtaining information about the Entrez databases

EInfo provides field index term counts, last update, and available links for each of NCBI's databases. In addition, you can use EInfo to obtain a list of all database names accessible through the Entrez utilities:

```
>>> from Bio import Entrez
>>> handle = Entrez.einfo(email="A.N.Other@example.com")
>>> result = handle.read()
```

The variable `result` now contains a list of databases in XML format:

```
>>> print result
<?xml version="1.0"?>
<!DOCTYPE eInfoResult PUBLIC "-//NLM//DTD eInfoResult, 11 May 2002//EN"
"http://www.ncbi.nlm.nih.gov/entrez/query/DTD/eInfo_020511.dtd">
<eInfoResult>
<DbList>
  <DbName>pubmed</DbName>
  <DbName>protein</DbName>
  <DbName>nucleotide</DbName>
  <DbName>nuccore</DbName>
  <DbName>nucgss</DbName>
  <DbName>nucest</DbName>
  <DbName>structure</DbName>
  <DbName>genome</DbName>
  <DbName>books</DbName>
  <DbName>cancerchromosomes</DbName>
  <DbName>cdd</DbName>
  <DbName>gap</DbName>
  <DbName>domains</DbName>
  <DbName>gene</DbName>
  <DbName>genomeprj</DbName>
  <DbName>gensat</DbName>
  <DbName>geo</DbName>
  <DbName>gds</DbName>
  <DbName>homologene</DbName>
  <DbName>journals</DbName>
```



```

    <DbName>mesh</DbName>
    <DbName>ncbisearch</DbName>
    <DbName>nlmcatalog</DbName>
    <DbName>omia</DbName>
    <DbName>omim</DbName>
    <DbName>pmc</DbName>
    <DbName>popset</DbName>
    <DbName>probe</DbName>
    <DbName>proteinclusters</DbName>
    <DbName>pcassay</DbName>
    <DbName>pccompound</DbName>
    <DbName>pcsubstance</DbName>
    <DbName>snp</DbName>
    <DbName>taxonomy</DbName>
    <DbName>toolkit</DbName>
    <DbName>unigene</DbName>
    <DbName>unists</DbName>
</DbList>
</eInfoResult>

```

Since this is a fairly simple XML file, we could extract the information it contains simply by string searching. Using Bio.Entrez's parser instead, we can directly parse this XML file into a Python object:

```

>>> from Bio import Entrez
>>> handle = Entrez.einfo(email="A.N.Other@example.com")
>>> record = Entrez.read(handle)

```

Now `record` is a dictionary with exactly one key:

```

>>> record.keys()
[u'DbList']

```

The values stored in this key is the list of database names shown in the XML above:

```

>>> record["DbList"]
['pubmed', 'protein', 'nucleotide', 'nuccore', 'nucgss', 'nucest',
 'structure', 'genome', 'books', 'cancerchromosomes', 'cdd', 'gap',
 'domains', 'gene', 'genomeprj', 'gensat', 'geo', 'gds', 'homologene',
 'journals', 'mesh', 'ncbisearch', 'nlmcatalog', 'omia', 'omim', 'pmc',
 'popset', 'probe', 'proteinclusters', 'pcassay', 'pccompound',
 'pcsubstance', 'snp', 'taxonomy', 'toolkit', 'unigene', 'unists']

```

For each of these databases, we can use EInfo again to obtain more information:

```

>>> handle = Entrez.einfo(db="pubmed", email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
>>> record["DbInfo"]["Description"]
'PubMed bibliographic record'
>>> record["DbInfo"]["Count"]
'17989604'
>>> record["DbInfo"]["LastUpdate"]
'2008/05/24 06:45'

```

Try `record["DbInfo"].keys()` for other information stored in this record.

7.3 ESearch: Searching the Entrez databases

To search any of these databases, we use `Bio.Entrez.esearch()`. For example, let's search in PubMed for publications related to Biopython:

```
>>> from Bio import Entrez
>>> handle = Entrez.esearch(db="pubmed", term="biopython", email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
>>> record["IdList"]
['16403221', '16377612', '14871861', '14630660', '12230038']
```

In this output, you see five PubMed IDs (16403221, 16377612, 14871861, 14630660, 12230038), which can be retrieved by `EFetch` (see section 7.6).

You can also use `ESearch` to search GenBank. Here we'll do a quick search for the *rpl16* gene in *Opuntia*:

```
>>> handle = Entrez.esearch(db="nucleotide", term="Opuntia and rpl16",
                             email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
>>> record["Count"]
'9'
>>> record["IdList"]
['57240072', '57240071', '6273287', '6273291', '6273290',
 '6273289', '6273286', '6273285', '6273284']
```

Each of the IDs (57240072, 57240071, 6273287...) is a GenBank identifier. See section 7.6 for information on how to actually download these GenBank records.

As a final example, let's get a list of computational journal titles:

```
>>> handle = Entrez.esearch(db="journals", term="computational",
                             email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
>>> record["Count"]
'16'
>>> record["IdList"]
['30367', '33843', '33823', '32989', '33190', '33009', '31986',
 '34502', '8799', '22857', '32675', '20258', '33859', '32534',
 '32357', '32249']
```

Again, we could use `EFetch` to obtain more information for each of these journal IDs.

`ESearch` has many useful options — see the [ESearch help page](#) for more information.

7.4 EPost

`EPost` posts a list of UIs for use in subsequent search strategies; see the [EPost help page](#) for more information. It is available from Biopython through `Bio.Entrez.epost()`.

7.5 ESummary: Retrieving summaries from primary IDs

`ESummary` retrieves document summaries from a list of primary IDs (see the [ESummary help page](#) for more information). In Biopython, `ESummary` is available as `Bio.Entrez.esummary()`. Using the search result above, we can for example find out more about the journal with ID 30367:

```

>>> from Bio import Entrez
>>> handle = Entrez.esummary(db="journals", id="30367", email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
>>> record[0]["Id"]
'30367'
>>> record[0]["Title"]
'Computational biology and chemistry'
>>> record[0]["Publisher"]
'Pergamon,'

```

7.6 EFetch: Downloading full records from Entrez

EFetch is what you use when you want to retrieve a full record from Entrez. For the *Opuntia* example above, we can download GenBank record 57240072 using `Bio.Entrez.efetch`:

```

>>> handle = Entrez.efetch(db="nucleotide", id="57240072", rettype="genbank",
                           email="A.N.Other@example.com")
>>> print handle.read()
LOCUS      AY851612                892 bp    DNA        linear    PLN 10-APR-2007
DEFINITION Opuntia subulata rpl16 gene, intron; chloroplast.
ACCESSION  AY851612
VERSION    AY851612.1  GI:57240072
KEYWORDS   .
SOURCE     chloroplast Austrocylindropuntia subulata
  ORGANISM Austrocylindropuntia subulata
            Eukaryota; Viridiplantae; Streptophyta; Embryophyta; Tracheophyta;
            Spermatophyta; Magnoliophyta; eudicotyledons; core eudicotyledons;
            Caryophyllales; Cactaceae; Opuntioideae; Austrocylindropuntia.
REFERENCE  1 (bases 1 to 892)
  AUTHORS  Butterworth,C.A. and Wallace,R.S.
  TITLE    Molecular Phylogenetics of the Leafy Cactus Genus Pereskia
            (Cactaceae)
  JOURNAL  Syst. Bot. 30 (4), 800-808 (2005)
REFERENCE  2 (bases 1 to 892)
  AUTHORS  Butterworth,C.A. and Wallace,R.S.
  TITLE    Direct Submission
  JOURNAL  Submitted (10-DEC-2004) Desert Botanical Garden, 1201 North Galvin
            Parkway, Phoenix, AZ 85008, USA
FEATURES   Location/Qualifiers
   source   1..892
            /organism="Austrocylindropuntia subulata"
            /organelle="plastid:chloroplast"
            /mol_type="genomic DNA"
            /db_xref="taxon:106982"
   gene     <1..>892
            /gene="rpl16"
   intron   <1..>892
            /gene="rpl16"
ORIGIN
1 cattaaagaa gggggatgcg gataaatgga aaggcgaag aaagaaaaa atgaatctaa
61 atgatatacg attccactat gtaaggtctt tgaatcatat cataaaagac aatgtaataa

```

```

121 agcatgaata cagattcaca cataattatc tgatatgaat ctattcatag aaaaaagaaa
181 aaagtaagag cctccggcca ataaagacta agagggttgg ctcaagaaca aagttcatta
241 agagctccat tgtagaattc agacctaatc attaatacaag aagcgaatggg aacgatgtaa
301 tccatgaata cagaagattc aattgaaaaa gatcctaatag atcattggga aggatggcgg
361 aacgaaccag agaccaattc atctattctg aaaagtgata aactaatcct ataaaactaa
421 aatagatatt gaaagagtaa atattcgccc gcgaaaattc cttttttatt aaattgctca
481 tattttattt tagcaatgca atctaataaa atatatctat acaaaaaaat atagacaaac
541 tatatatata taatatattt caaatctcct tatataccca aatataaaaa tatctaataa
601 attagatgaa tatcaaagaa tctattgatt tagtgtatta ttaaagtat atcttaattc
661 aatattatta ttctattcat ttttattcat tttcaaattt ataatatatt aatctatata
721 ttaatttata attctattct aattcgaatt caatttttaa atattcatat tcaattaaaa
781 ttgaaatttt ttcattcgcg aggagccgga tgagaagaaa ctctcatgtc cggttctgta
841 gtagagatgg aattaagaaa aaaccatcaa ctataacccc aagagaacca ga

```

```
//
```

The argument `rettype="genbank"` lets us download this record in the GenBank format. Alternatively, you could for example use `rettype="fasta"` to get the Fasta-format; see the [EFetch Help page](#) for other options. The available formats depend on which database you are downloading from.

If you fetch the record in one of the formats accepted by `Bio.SeqIO` (see Chapter 4), you can directly parse it into a `SeqRecord`:

```

>>> from Bio import Entrez, SeqIO
>>> handle = Entrez.efetch(db="nucleotide", id="57240072",rettype="genbank",
                           email="A.N.Other@example.com")
>>> record = SeqIO.read(handle, "genbank")
>>> print record
ID: AY851612.1
Name: AY851612
Description: Opuntia subulata rpl16 gene, intron; chloroplast.
/sequence_version=1
/source=chloroplast Austrocyllindropuntia subulata
....

```

By default you get the output in XML format, which you can parse using the `Bio.Entrez.read()` function:

```

>>> from Bio import Entrez
>>> handle = Entrez.efetch(db="nucleotide", id="57240072",
                           email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
>>> record[0]["GBSeq_definition"]
'Opuntia subulata rpl16 gene, intron; chloroplast'
>>> record[0]["GBSeq_source"]
'chloroplast Austrocyllindropuntia subulata'
....

```

7.7 ELink

For help on ELink, see the [ELink help page](#). ELink is available from Biopython through `Bio.Entrez.elink()`.

7.8 EGQuery: Obtaining counts for search terms

EGQuery provides counts for a search term in each of the Entrez databases. This is particularly useful to find out how many items a search will return before actually performing the search with ESearch (see the example in 7.10.1 below).

In this example, we use `Bio.Entrez.egquery()` to obtain the counts for “Biopython”:

```
>>> handle = Entrez.egquery(term="biopython",
                             email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
>>> record["eGQueryResult"][0]["DbName"]
'pubmed'
>>> record["eGQueryResult"][0]["Count"]
'5'
```

See the [EGQuery help page](#) for more information.

7.9 ESpell: Obtaining spelling suggestions

ESpell retrieves spelling suggestions. In this example, we use `Bio.Entrez.espell()` to obtain the correct spelling of Biopython:

```
>>> from Bio import Entrez
>>> handle = Entrez.espell(term="biopythooon", email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
>>> record["Query"]
'biopythooon'
>>> record["CorrectedQuery"]
'biopython'
```

See the [ESpell help page](#) for more information.

7.10 Examples

7.10.1 Searching and downloading Entrez Nucleotide records

Here we’ll show a simple example of performing a remote Entrez query. In section 2.3 of the parsing examples, we talked about using NCBI’s Entrez website to search the NCBI nucleotide databases for info on *Cypripedioideae*, our friends the lady slipper orchids. Now, we’ll look at how to automate that process using a Python script. In this example, we’ll just show how to connect, get the results, and parse them, with the Entrez module doing all of the work.

First, we use EGQuery to find out the number of results we will get before actually downloading them:

```
>>> from Bio import Entrez
>>> handle = Entrez.egquery(term='Cypripedioideae', email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
>>> for row in record['eGQueryResult']:
...     if row['DbName']=='nuccore':
...         print row['Count']
814
```

So, we expect to find 814 Entrez Nucleotide records. If you find some ridiculously high number of hits, you may want to reconsider if you really want to download all of them, which is our next step:

```
>>> from Bio import Entrez
>>> handle = Entrez.esearch(db='nucleotide', term='Cypripedioideae', retmax=814,
                           email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
```

Here, `record` is a Python dictionary containing the search results and some auxiliary information. Just for information, let's look at what is stored in this dictionary:

```
>>> print record.keys()
[u'Count', u'RetMax', u'IdList', u'TranslationSet', u'RetStart', u'QueryTranslation']
```

First, let's check how many results were found:

```
>>> print record['Count']
'814'
```

which is the number we expected. The 814 results are stored in `record['IdList']`:

```
>>> print len(record['IdList'])
814
```

Let's look at the first five results:

```
>>> print record['IdList'][:5]
['187237168', '187372713', '187372690', '187372688', '187372686']
```

We can download these records using `efetch`. While you could download these records one by one, to reduce the load on NCBI's servers, it is better to fetch a bunch of records at the same time, shown below. However, in this situation you should ideally be using the history feature described later in Section [7.10.3](#).

```
>>> idlist = ",".join(record['IdList'][:5])
>>> print idlist
187237168,187372713,187372690,187372688,187372686
>>> handle = Entrez.efetch(db='nucleotide', id=idlist, retmode='xml',
                           email="A.N.Other@example.com")
>>> records = Entrez.read(handle)
>>> print len(records)
5
```

Each of these records corresponds to one GenBank record.

```
>>> print records[0].keys()
[u'GBSeq_moltype', u'GBSeq_source', u'GBSeq_sequence',
 u'GBSeq_primary-accession', u'GBSeq_definition', u'GBSeq_accession-version',
 u'GBSeq_topology', u'GBSeq_length', u'GBSeq_feature-table',
 u'GBSeq_create-date', u'GBSeq_other-seqids', u'GBSeq_division',
 u'GBSeq_taxonomy', u'GBSeq_references', u'GBSeq_update-date',
 u'GBSeq_organism', u'GBSeq_locus', u'GBSeq_strandedness']
```

```
>>> print records[0]['GBSeq_primary-accession']
DQ110336
```

```
>>> print records[0]['GBSeq_other-seqids']
['gb|DQ110336.1|', 'gi|187237168']
```

```
>>> print records[0]['GBSeq_definition']
```

```
Cypripedium calceolus voucher Davis 03-03 A maturase (matR) gene, partial cds;
mitochondrial
```

```
>>> print records[0]['GBSeq_organism']
Cypripedium calceolus
```

You could use this to quickly set up searches – but for heavy usage, see Section [7.10.3](#).

7.10.2 Finding the lineage of an organism

Staying with the same organism, let's now find its lineage. First, we search the Taxonomy database for Cypripedioideae. We find exactly one accession number:

```
>>> handle = Entrez.esearch(db="Taxonomy", term="Cypripedioideae",
                           email="A.N.Other@example.com")
>>> record = Entrez.read(handle)
>>> record["IdList"]
['158330']
>>> record["IdList"][0]
'158330'
```

Now, we use `efetch` to download this entry in the Taxonomy database and to parse it:

```
>>> handle = Entrez.efetch(db="Taxonomy", id="158330", retmode='xml')
>>> records = Entrez.read(handle)
```

Again, this record stores lots of information:

```
>>> records[0].keys()
[u'Lineage', u'Division', u'ParentTaxId', u'PubDate', u'LineageEx',
 u'CreateDate', u'TaxId', u'Rank', u'GeneticCode', u'ScientificName',
 u'MitoGeneticCode', u'UpdateDate']
```

We can get the lineage directly from this record:

```
>>> records[0]['Lineage']
'cellular organisms; Eukaryota; Viridiplantae; Streptophyta; Streptophytina;
 Embryophyta; Tracheophyta; Euphyllophyta; Spermatophyta; Magnoliophyta;
 Liliopsida; Asparagales; Orchidaceae'
```

7.10.3 Using the history and WebEnv

Often you will want to make a series of linked queries. Most typically, running a search, perhaps refining the search, and then retrieving detailed search results. You *can* do this by making a series of separate calls to Entrez. However, the NCBI prefer you to take advantage of their history support.

For example, suppose we want to search and download all Orchid `rpl16` nucleotide sequences, and store them in a FASTA file. We could naively combine the example code for `Bio.Entrez.esearch()` (Section [7.3](#)) to get a list of GI numbers, and then repeatedly call `Bio.Entrez.efetch()` (Section [7.6](#)) to download them all. You could reduce the number of queries by asking for the records in batches (see Section [7.10.1](#)). That would probably be better, but is still not what the NCBI encourage.

The approved approach is to run the search with the history feature. Then, we can fetch the results by reference to the search results - which the NCBI can anticipate and cache.

```

from Bio import Entrez
search_handle = Entrez.esearch(db="nucleotide",term="Opuntia and rpl16",
                              usehistory="y", email="history.user@example.com")
search_results = Entrez.read(search_handle)
search_handle.close()

gi_list = search_results["IdList"]
count = int(search_results["Count"])
assert count == len(gi_list)

session_cookie = search_results["WebEnv"]
query_key = search_results["QueryKey"]

```

In addition to the GI numbers of the sequences found in the search, because we have asked to use the history feature the XML search results also include `WebEnv` and `QueryKey` values which are used to refer to these search results. Having stored these values in variables `session_cookie` and `query_key` we can use them as parameters to `Bio.Entrez.efetch()` instead of giving the GI numbers as identifiers.

While for small searches you might be OK downloading everything at once, its better download in batches. You use the `retstart` and `retmax` parameters to specify which range of search results you want returned (starting entry using zero-based counting, and maximum number of results to return). For example,

```

batch_size = 3
out_handle = open("orchid_rpl16.fasta", "w")
for start in range(0,count,batch_size) :
    end = min(count, start+batch_size)
    print "Going to download record %i to %i" % (start+1, end)
    fetch_handle = Entrez.efetch(db="nucleotide", rettype="fasta",
                                retstart=start, retmax=batch_size,
                                webenv=session_cookie, query_key=query_key,
                                email="history.user@example.com")

    data = fetch_handle.read()
    fetch_handle.close()
    out_handle.write(data)
out_handle.close()

```

And finally, don't forget to include your *own* email address in the Entrez calls.

Chapter 8

Swiss-Prot, Prosite, Prodoc, and ExPASy

8.1 Bio.SwissProt: Parsing Swiss-Prot files

8.1.1 Parsing Swiss-Prot records

Swiss-Prot (<http://www.expasy.org/sprot>) is a hand-curated database of protein sequences. In Section 4.2.2, we described how to extract the sequence of a Swiss-Prot record as a `SeqRecord` object. Alternatively, you can store the Swiss-Prot record in a `Bio.SwissProt.SProt.Record` object, which in fact stores the complete information contained in the Swiss-Prot record. In this Section, we describe how to extract `Bio.SwissProt.SProt.Record` objects from a Swiss-Prot file.

To parse a Swiss-Prot record, we first get a handle to a Swiss-Prot record. There are several ways to do so, depending on where and how the Swiss-Prot record is stored:

- Open a Swiss-Prot file locally:

```
>>> handle = open("myswissprotfile.dat")
```
- Open a gzipped Swiss-Prot file:

```
>>> import gzip
>>> handle = gzip.open("myswissprotfile.dat.gz")
```
- Open a Swiss-Prot file over the internet:

```
>>> import urllib
>>> handle = urllib.urlopen("http://www.somelocation.org/data/someswissprotfile.dat")
```
- Open a Swiss-Prot file over the internet from the ExPASy database (see section 8.4.1):

```
>>> from Bio import ExPASy
>>> handle = ExPASy.get_sprot_raw(myaccessionnumber)
```

The key point is that for the parser, it doesn't matter how the handle was created, as long as it points to data in the Swiss-Prot format.

We can use `Bio.SeqIO` as described in Section 4.2.2 to get file format agnostic `SeqRecord` objects. Alternatively, we can get `Bio.SwissProt.SProt.Record` objects which are a much closer match to the underlying file format, using following code.

To read one Swiss-Prot record from the handle, we use the function `read()`:

```
>>> from Bio import SwissProt
>>> record = SwissProt.read(handle)
```

This function should be used if the handle points to exactly one Swiss-Prot record. It raises a `ValueError` if no Swiss-Prot record was found, and also if more than one record was found.

We can now print out some information about this record:

```
>>> print record.description
CHALCONE SYNTHASE 3 (EC 2.3.1.74) (NARINGENIN-CHALCONE SYNTHASE 3).
>>> for ref in record.references:
...     print "authors:", ref.authors
...     print "title:", ref.title
...
authors: Liew C.F., Lim S.H., Loh C.S., Goh C.J.;
title: "Molecular cloning and sequence analysis of chalcone synthase cDNAs of
Bromheadia finlaysoniana.";
>>> print record.organism_classification
['Eukaryota', 'Viridiplantae', 'Embryophyta', 'Tracheophyta', 'Spermatophyta',
'Magnoliophyta', 'Liliopsida', 'Asparagales', 'Orchidaceae', 'Bromheadia']
```

To parse a file that contains more than one Swiss-Prot record, we use the `parse` function instead. This function allows us to iterate over the records in the file. For example, let's parse the full Swiss-Prot database and collect all the descriptions. The full Swiss-Prot database, downloaded from ExPASy on 4 December 2007, contains 290484 Swiss-Prot records in a single gzipped-file `uniprot_sprot.dat.gz`.

```
>>> import gzip
>>> input = gzip.open("uniprot_sprot.dat.gz")
>>> from Bio import SwissProt
>>> records = SwissProt.parse(input)
>>> descriptions = []
>>> for record in records:
...     description = record.description
...     descriptions.append(description)
...
>>> len(descriptions)
290484
>>> descriptions[:3]
['104 kDa microneme/rhoptry antigen precursor (p104).',
'104 kDa microneme/rhoptry antigen precursor (p104).',
'Protein 108 precursor.']
```

It is equally easy to extract any kind of information you'd like from Swiss-Prot records. To see the members of a Swiss-Prot record, use

```
>>> dir(record)
['_doc__', '__init__', '__module__', 'accessions', 'annotation_update',
'comments', 'created', 'cross_references', 'data_class', 'description',
'entry_name', 'features', 'gene_name', 'host_organism', 'keywords',
'molecule_type', 'organelle', 'organism', 'organism_classification',
'references', 'seqinfo', 'sequence', 'sequence_length',
'sequence_update', 'taxonomy_id']
```

8.1.2 Parsing the Swiss-Prot keyword and category list

Swiss-Prot also distributes a file `keywlist.txt`, which lists the keywords and categories used in Swiss-Prot. The file contains entries in the following form:

```
ID  2Fe-2S.
AC  KW-0001
DE  Protein which contains at least one 2Fe-2S iron-sulfur cluster: 2 iron
DE  atoms complexed to 2 inorganic sulfides and 4 sulfur atoms of
DE  cysteines from the protein.
SY  Fe2S2; [2Fe-2S] cluster; [Fe2S2] cluster; Fe2/S2 (inorganic) cluster;
SY  Di-mu-sulfido-diiron; 2 iron, 2 sulfur cluster binding.
GO  GO:0051537; 2 iron, 2 sulfur cluster binding
HI  Ligand: Iron; Iron-sulfur; 2Fe-2S.
HI  Ligand: Metal-binding; 2Fe-2S.
CA  Ligand.
//
ID  3D-structure.
AC  KW-0002
DE  Protein, or part of a protein, whose three-dimensional structure has
DE  been resolved experimentally (for example by X-ray crystallography or
DE  NMR spectroscopy) and whose coordinates are available in the PDB
DE  database. Can also be used for theoretical models.
HI  Technical term: 3D-structure.
CA  Technical term.
//
ID  3Fe-4S.
...
```

The entries in this file can be parsed by the `parse` function in the `Bio.SwissProt.KeyWList` module. Each entry is then stored as a `Bio.SwissProt.KeyWList.Record`, which is a Python dictionary.

```
>>> from Bio.SwissProt import KeyWList
>>> handle = open("keywlist.txt")
>>> records = KeyWList.parse(handle)
>>> for record in records:
...     print record['ID']
...     print record['DE']
```

This prints

```
2Fe-2S.
Protein which contains at least one 2Fe-2S iron-sulfur cluster: 2 iron atoms
complexed to 2 inorganic sulfides and 4 sulfur atoms of cysteines from the
protein.
...
```

8.2 Bio.ProSITE: Parsing ProSITE records

ProSITE is a database containing protein domains, protein families, functional sites, as well as the patterns and profiles to recognize them. ProSITE was developed in parallel with Swiss-Prot. In Biopython, a ProSITE record is represented by the `Bio.ProSITE.Record` class, whose members correspond to the different fields in a ProSITE record.

In general, a Prosite file can contain more than one Prosite records. For example, the full set of Prosite records, which can be downloaded as a single file (`prosite.dat`) from ExPASy, contains 2073 records in (version 20.24 released on 4 December 2007). To parse such a file, we again make use of an iterator:

```
>>> from Bio import Prosite
>>> handle = open("myprositefile.dat")
>>> records = Prosite.parse(handle)
```

We can now take the records one at a time and print out some information. For example, using the file containing the complete Prosite database, we'd find

```
>>> from Bio import Prosite
>>> handle = open("prosite.dat")
>>> records = Prosite.parse(handle)
>>> record = records.next()
>>> record.accession
'PS00001'
>>> record.name
'ASN_GLYCOSYLATION'
>>> record.pdoc
'PDOC00001'
>>> record = records.next()
>>> record.accession
'PS00004'
>>> record.name
'CAMP_PHOSPHO_SITE'
>>> record.pdoc
'PDOC00004'
>>> record = records.next()
>>> record.accession
'PS00005'
>>> record.name
'PKC_PHOSPHO_SITE'
>>> record.pdoc
'PDOC00005'
```

and so on. If you're interested in how many Prosite records there are, you could use

```
>>> from Bio import Prosite
>>> handle = open("prosite.dat")
>>> records = Prosite.parse(handle)
>>> n = 0
>>> for record in records: n+=1
...
>>> print n
2073
```

To read exactly one Prosite from the handle, you can use the `read` function:

```
>>> from Bio import Prosite
>>> handle = open("mysingleprositerecord.dat")
>>> record = Prosite.read(handle)
```

This function raises a `ValueError` if no Prosite record is found, and also if more than one Prosite record is found.

8.3 Bio.Prosite.Prodoc: Parsing Prodoc records

In the Prosite example above, the `record.pdoc` accession numbers 'PDOC00001', 'PDOC00004', 'PDOC00005' and so on refer to Prodoc records, which contain the Prosite Documentation. The Prodoc records are available from ExPASy as individual files, and as one file (`prosite.doc`) containing all Prodoc records.

We use the parser in `Bio.Prosite.Prodoc` to parse Prodoc records. For example, to create a list of all Prodoc accession numbers, you can use

```
>>> from Bio.Prosite import Prodoc
>>> handle = open("prosite.doc")
>>> records = Prodoc.parse(handle)
>>> accessions = [record.accession for record in records]
```

Again a `read()` function is provided to read exactly one Prodoc record from the handle.

8.4 Bio.ExPASy: Accessing the ExPASy server

Swiss-Prot, Prosite, and Prodoc records can be downloaded from the ExPASy web server at <http://www.expasy.org>. Six kinds of queries are available from ExPASy:

`get_prodoc_entry` To download a Prodoc record in HTML format

`get_prosite_entry` To download a Prosite record in HTML format

`get_prosite_raw` To download a Prosite or Prodoc record in raw format

`get_sprot_raw` To download a Swiss-Prot record in raw format

`sprot_search_ful` To search for a Swiss-Prot record

`sprot_search_de` To search for a Swiss-Prot record

To access this web server from a Python script, we use the `Bio.ExPASy` module.

8.4.1 Retrieving a Swiss-Prot record

Let's say we are looking at chalcone synthases for Orchids (see section 2.3 for some justification for looking for interesting things about orchids). Chalcone synthase is involved in flavanoid biosynthesis in plants, and flavanoids make lots of cool things like pigment colors and UV protectants.

If you do a search on Swiss-Prot, you can find three orchid proteins for Chalcone Synthase, id numbers O23729, O23730, O23731. Now, let's write a script which grabs these, and parses out some interesting information.

First, we grab the records, using the `get_sprot_raw()` function of `Bio.ExPASy`. This function is very nice since you can feed it an id and get back a handle to a raw text record (no html to mess with!). We can use `Bio.SwissProt.read` to pull out the Swiss-Prot record, or `Bio.SeqIO.read` to get a `SeqRecord`. The following code accomplishes what I just wrote:

```
>>> from Bio import ExPASy
>>> from Bio import SwissProt

>>> accessions = ["O23729", "O23730", "O23731"]
>>> records = []

>>> for accession in accessions:
```

```

...     handle = ExpASy.get_sprot_raw(accession)
...     record = SwissProt.read(handle)
...     records.append(record)

```

If the accession number you provided to `ExpASy.get_sprot_raw` does not exist, then `SwissProt.read(handle)` will raise a `ValueError`. You can catch `ValueException` exceptions to detect invalid accession numbers:

```

>>> for accession in accessions:
...     handle = ExpASy.get_sprot_raw(accession)
...     try:
...         record = SwissProt.read(handle)
...     except ValueError:
...         print "WARNING: Accession %s not found" % accession
...     records.append(record)

```

8.4.2 Searching Swiss-Prot

Now, you may remark that I knew the records' accession numbers beforehand. Indeed, `get_sprot_raw()` needs either the entry name or an accession number. When you don't have them handy, you can use one of the `sprot_search_de()` or `sprot_search_ful()` functions.

`sprot_search_de()` searches in the ID, DE, GN, OS and OG lines; `sprot_search_ful()` searches in (nearly) all the fields. They are detailed on <http://www.expasy.org/cgi-bin/sprot-search-de> and <http://www.expasy.org/cgi-bin/sprot-search-ful> respectively. Note that they don't search in TrEMBL by default (argument `trembl`). Note also that they return html pages; however, accession numbers are quite easily extractable:

```

>>> from Bio import ExpASy
>>> import re

>>> handle = ExpASy.sprot_search_de("Orchid Chalcone Synthase")
>>> # or:
>>> # handle = ExpASy.sprot_search_ful("Orchid and {Chalcone Synthase}")
>>> html_results = handle.read()
>>> if "Number of sequences found" in html_results:
...     ids = re.findall(r'HREF="/uniprot/(\w+)"', html_results)
... else:
...     ids = re.findall(r'href="/cgi-bin/niceprot\.pl\?(\w+)"', html_results)

```

8.4.3 Retrieving Prosite and Prodoc records

Prosite and Prodoc records can be retrieved either in HTML format, or in raw format. To parse Prosite and Prodoc records with Biopython, you should retrieve the records in raw format. For other purposes, however, you may be interested in these records in HTML format.

To retrieve a Prosite or Prodoc record in raw format, use `get_prosite_raw()`. Although this function has `prosite` in the name, it can be used for Prodoc records as well. For example, to download a Prosite record and print it out in raw text format, use

```

>>> from Bio import ExpASy
>>> handle = ExpASy.get_prosite_raw('PS00001')
>>> text = handle.read()
>>> print text

```

To retrieve a Prosite record and parse it into a `Bio.Prosite.Record` object, use

```
>>> from Bio import ExpPASy
>>> from Bio import Prosite
>>> handle = ExpPASy.get_prosite_raw('PS00001')
>>> record = Prosite.read(handle)
```

Finally, to retrieve a Prodoc record and parse it into a `Bio.Prosite.Prodoc.Record` object, use

```
>>> from Bio import ExpPASy
>>> from Bio.Prosite import Prodoc
>>> handle = ExpPASy.get_prosite_raw('PDOC00001')
>>> record = Prodoc.read(handle)
```

For non-existing accession numbers, `ExpPASy.get_prosite_raw` returns a handle to an empty string. When faced with an empty string, `Prosite.read` and `Prodoc.read` will raise a `ValueError`. You can catch these exceptions to detect invalid accession numbers.

The functions `get_prosite_entry()` and `get_prodoc_entry()` are used to download Prosite and Prodoc records in HTML format. To create a web page showing one Prosite record, you can use

```
>>> from Bio import ExpPASy
>>> handle = ExpPASy.get_prosite_entry('PS00001')
>>> html = handle.read()
>>> output = open("myprositerecord.html", "w")
>>> output.write(html)
>>> output.close()
```

and similarly for a Prodoc record:

```
>>> from Bio import ExpPASy
>>> handle = ExpPASy.get_prodoc_entry('PDOC00001')
>>> html = handle.read()
>>> output = open("myprodocrecord.html", "w")
>>> output.write(html)
>>> output.close()
```

For these functions, an invalid accession number returns an error message in HTML format.

Chapter 9

Cookbook – Cool things to do with it

9.1 PubMed

The `Bio.PubMed` module uses `Bio.Entrez` internally to access the NCBI. *Please* remember to read and observe the NCBI's guidelines on using their Entrez online facilities responsibly. If you are found to be abusing their servers, they can and will block your access! See Section 7.1.

9.1.1 Sending a query to PubMed

If you are in the Medical field or interested in human issues (and many times even if you are not!), PubMed (<http://www.ncbi.nlm.nih.gov/PubMed/>) is an excellent source of all kinds of goodies. So like other things, we'd like to be able to grab information from it and use it in python scripts.

Querying PubMed using Biopython is extremely painless. To get all of the article ids for articles having to do with orchids (see section 2.3 for our motivation), we only need the following three lines of code:

```
from Bio import PubMed

search_term = 'orchid'
orchid_ids = PubMed.search_for(search_term)
```

This returns a python list containing all of the orchid ids

```
['11070358', '11064040', '11028023', '10947239', '10938351', '10936520',
'10905611', '10899814', '10856762', '10854740', '10758893', '10716342',
...]
```

With this list of ids we are ready to start retrieving the records, so follow on ahead to the next section.

9.1.2 Retrieving a PubMed record

The previous section described how to get a bunch of article ids. Now that we've got them, we obviously want to get the corresponding Medline records and extract the information from them.

The interface for retrieving records from PubMed should be very intuitive to python programmers – it models a python dictionary. To set up this interface, we need to set up a parser that will parse the results that we retrieve. The following lines of code get everything set up:

```
from Bio import PubMed
from Bio import Medline
```



```
rec_parser = Medline.RecordParser()
medline_dict = PubMed.Dictionary(parser = rec_parser)
```

What we've done is create a dictionary like object `medline_dict`. To get an article we access it like `medline_dict[id_to_get]`. What this does is connect with PubMed, get the article you ask for, parse it into a record object, and return it. Very cool!

Now let's look at how to use this nice dictionary to print out some information about some ids. We just need to loop through our ids (`orchid_ids` from the previous section) and print out the information we are interested in:

```
for oid in orchid_ids[0:5]:
    cur_record = medline_dict[oid]
    print 'title:', cur_record.title.rstrip()
    print 'authors:', cur_record.authors
    print 'source:', cur_record.source.strip()
    print
```

The output for this looks like:

```
title: Sex pheromone mimicry in the early spider orchid (ophrys sphegodes):
patterns of hydrocarbons as the key mechanism for pollination by sexual
deception [In Process Citation]
authors: ['Schiestl FP', 'Ayasse M', 'Paulus HF', 'Lofstedt C', 'Hansson BS',
'Ibarra F', 'Francke W']
source: J Comp Physiol [A] 2000 Jun;186(6):567-74
```

Especially interesting to note is the list of authors, which is returned as a standard python list. This makes it easy to manipulate and search using standard python tools. For instance, we could loop through a whole bunch of entries searching for a particular author with code like the following:

```
search_author = 'Waits T'

for our_id in our_id_list:
    cur_record = medline_dict[our_id]

    if search_author in cur_record.authors:
        print "Author %s found: %s" % (search_author,
                                       cur_record.source.strip())
```

The PubMed and Medline interfaces are very mature and nice to work with – hopefully this section gave you an idea of the power of the interfaces and how they can be used.

9.2 GenBank

The GenBank record format is a very popular method of holding information about sequences, sequence features, and other associated sequence information. The format is a good way to get information from the NCBI databases at <http://www.ncbi.nlm.nih.gov/>.

The `Bio.GenBank` module uses the NCBI's Entrez interface to run searches and fetch data. *Please* remember to read and observe the NCBI's guidelines on using their Entrez online facilities responsibly. If you are found to be abusing their servers, they can and will block your access! See Section 7.1.

9.2.1 Retrieving GenBank entries from NCBI

One very nice feature of the GenBank libraries is the ability to automate retrieval of entries from GenBank. This is very convenient for creating scripts that automate a lot of your daily work. In this example we'll show how to query the NCBI databases, and to retrieve the records from the query - something touched on in Section 4.2.1.

First, we want to make a query and find out the ids of the records to retrieve. Here we'll do a quick search for our favorite organism, *Opuntia*. We can do quick search and get back the GIs (GenBank identifiers) for all of the corresponding records:

```
from Bio import GenBank
```

```
gi_list = GenBank.search_for("Opuntia AND rp116")
```

`gi_list` will be a list of all of the GenBank identifiers that match our query:

```
["6273291", "6273290", "6273289", "6273287", "6273286", "6273285", "6273284"]
```

Now that we've got the GIs, we can use these to access the NCBI database through a dictionary interface. For instance, to retrieve the information for the first GI, we'll first have to create a dictionary that accesses NCBI:

```
ncbi_dict = GenBank.NCBIDictionary("nucleotide", "genbank")
```

Now that we've got this, we do the retrieval:

```
gb_record = ncbi_dict[gi_list[0]]
```

In this case, `gb_record` will be GenBank formatted record:

```
LOCUS      AF191665      902 bp    DNA          PLN          07-NOV-1999
DEFINITION Opuntia marenae rp116 gene; chloroplast gene for chloroplast
            product, partial intron sequence.
ACCESSION  AF191665
VERSION    AF191665.1  GI:6273291
...
```

In this case, we are just getting the raw records. We can also pass these records directly into a parser and return the parsed record. For instance, if we wanted to get back `SeqRecord` objects with the GenBank file parsed into `SeqFeature` objects we would need to create the dictionary with the `GenBank FeatureParser`:

```
record_parser = GenBank.FeatureParser()
ncbi_dict = GenBank.NCBIDictionary("nucleotide", "genbank",
                                   parser = record_parser)
```

Now retrieving a record will give you a `SeqRecord` object instead of the raw record:

```
>>> gb_seqrecord = ncbi_dict[gi_list[0]]
>>> print gb_seqrecord
<Bio.SeqRecord.SeqRecord instance at 0x102f9404>
```

Using these automated query retrieval functionality is a big plus over doing things by hand. Although the module should obey the NCBI's three-second rule, the NCBI have other recommendations like avoiding peak hours. See Section 7.1.

For more information of formats you can parse GenBank records into, please see section 9.2.2.

9.2.2 Parsing GenBank records

While GenBank files are nice and have lots of information, at the same time you probably only want to extract a small amount of that information at a time. The key to doing this is parsing out the information. Biopython provides GenBank parsers which help you accomplish this task. Right now the GenBank module provides the following parsers:

1. RecordParser – This parses the raw record into a GenBank specific Record object. This object models the information in a raw record very closely, so this is good to use if you are just interested in GenBank records themselves.
2. FeatureParser – This parses the raw record in a SeqRecord object with all of the feature table information represented in SeqFeatures (see section 10.1 for more info on these objects). This is best to use if you are interested in getting things in a more standard format. If you use Bio.SeqIO (Chapter 4) to read a GenBank file, it will call this FeatureParser for you.

Depending on the type of GenBank files you are interested in, they will either contain a single record, or multiple records. Each record will start with a LOCUS line, various other header lines, a list of features, and finally the sequence data, ending with a // line.

Dealing with a GenBank file containing a single record is very easy. For example, let's use a small bacterial genome, *Nanoarchaeum equitans* *Kin4-M* (RefSeq NC_005213, GenBank AE017199) which can be downloaded from the NCBI here (only 1.15 MB):

```
from Bio import GenBank
feature_parser = GenBank.FeatureParser()
gb_record = feature_parser.parse(open("AE017199.gb"))
# now do something with the record
print "Name %s, %i features" % (gb_record.name, len(gb_record.features))
print repr(gb_record.seq)
```

Or, using Bio.SeqIO instead (see Chapter 4):

```
from Bio import SeqIO
gb_record = SeqIO.read(open("AE017199.gb"), "genbank")
print "Name %s, %i features" % (gb_record.name, len(gb_record.features))
print repr(gb_record.seq)
```

Either should give the following output:

```
Name AE017199, 1107 features
Seq('TCTCGCAGAGTTCTTTTTGTATTAACAAACCCAAAACCCATAGAATTTAATGA...TTA', IUPACAmbiguousDNA())
```

9.2.3 Iterating over GenBank records

For multi-record GenBank files, the most common usage will be creating an iterator, and parsing through the file record by record. Doing this is very similar to how things are done in other formats, as the following code demonstrates, using an example file [cor6.6.gb](#) which is included in the BioPython source code under the Tests/GenBank/ directory:

```
from Bio import GenBank
feature_parser = GenBank.FeatureParser()
gb_iterator = GenBank.Iterator(open("cor6.6.gb"), feature_parser)
for cur_record in gb_iterator :
    print "Name %s, %i features" % (cur_record.name, len(cur_record.features))
    print repr(cur_record.seq)
```

Or, using `Bio.SeqIO` instead (see Chapter 4):

```
from Bio import SeqIO
for cur_record in SeqIO.parse(open("cor6_6.gb"), "genbank") :
    print "Name %s, %i features" % (cur_record.name, len(cur_record.features))
    print repr(cur_record.seq)
```

This just iterates over a GenBank file, parsing it into `SeqRecord` and `SeqFeature` objects, and prints out the `Seq` objects representing the sequences in the record.

As with other formats, you have lots of tools for dealing with GenBank records. This should make it possible to do whatever you need to with GenBank.

9.3 Dealing with alignments

It is often very useful to be able to align particular sequences. I do this quite often to get a quick and dirty idea of relationships between sequences. Consequently, it is very nice to be able to quickly write up a python script that does an alignment and gives you back objects that are easy to work with. The alignment related code in Biopython is meant to allow python-level access to alignment programs so that you can run alignments quickly from within scripts.

9.3.1 Clustalw

Clustalx (<http://www-igbmc.u-strasbg.fr/BioInfo/ClustalX/Top.html>) is a very nice program for doing multiple alignments. Biopython offers access to alignments in clustal format (these normally have a `*.aln` extension) that are produced by Clustalx. It also offers access to clustalw, which is the command line version of clustalx.

We'll need some sequences to align, such as `opuntia.fasta` (also available online [here](#)) which is a small FASTA file containing seven orchid gene DNA sequences, which you can also find from `Doc/examples/` in the Biopython source distribution.

The first step in interacting with clustalw is to set up a command line you want to pass to the program. Clustalw has a ton of command line options, and if you set a lot of parameters, you can end up typing in a huge ol' command line quite a bit. This command line class models the command line by making all of the options be attributes of the class that can be set. A few convenience functions also exist to set certain parameters, so that some error checking on the parameters can be done.

To create a command line object to do a clustalw multiple alignment we do the following:

```
import os
from Bio.Clustalw import MultipleAlignCL

cline = MultipleAlignCL(os.path.join(os.getcwd(), "opuntia.fasta"))
cline.set_output("test.aln")
```

First we import the `MultipleAlignCL` object, which models running a multiple alignment from clustalw. We then initialize the command line, with a single argument of the fasta file that we are going to be using for the alignment. The initialization function also takes an optional second argument which specifies the location of the `clustalw` executable. By default, the commandline will just be invoked with `'clustalw,'` assuming that you've got it somewhere on your `PATH`.

The second argument sets the output to go to the file `test.aln`. The `MultipleAlignCL` object also has numerous other parameters to specify things like output format, gap costs, etc.

We can look at the command line we have generated by invoking the `__str__` member attribute of the `MultipleAlignCL` class. This is done by calling `str(cline)` or simply by printing out the command line with `print cline`. In this case, doing this would give the following output:

```
clustalw ./opuntia.fasta -OUTFILE=test.aln
```

Now that we've set up a simple command line, we now want to run the commandline and collect the results so we can deal with them. This can be done using the `do_alignment` function of `Clustalw` as follows:

```
from Bio import Clustalw

alignment = Clustalw.do_alignment(cline)
```

What happens when you run this if that Biopython executes your command line and runs `clustalw` with the given parameters. It then grabs the output, and if it is in a format that Biopython can parse (currently only `clustal` format), then it will parse the results and return them as an alignment object of the appropriate type. So in this case since we are getting results in the default `clustal` format, the returned `alignment` object will be a `ClustalAlignment` type.

Once we've got this alignment, we can do some interesting things with it such as get `seq_record` objects for all of the sequences involved in the alignment:

```
all_records = alignment.get_all_seqs()

print "description:", all_records[0].description
print "sequence:", all_records[0].seq
```

This prints out the description and sequence object for the first sequence in the alignment:

```
description: gi|6273285|gb|AF191659.1|AF191
sequence: Seq('TATACATTAAAGAAGGGGGATGCGGATAAATGGAAAGGCGAAAGAAAGAAAAAATGAAT
...', IUPACAmbiguousDNA())
```

You can also calculate the maximum length of the alignment with:

```
length = alignment.get_alignment_length()
```

Finally, to write out the alignment object in the original format, we just need to access the `__str__` function. So doing a `print alignment` gives:

```
CLUSTAL X (1.81) multiple sequence alignment

gi|6273285|gb|AF191659.1|AF191      TATACATTAAAGAAGGGGGATGCGGATAAATGGAAAGGCGAAAGAAAGAA
gi|6273284|gb|AF191658.1|AF191      TATACATTAAAGAAGGGGGATGCGGATAAATGGAAAGGCGAAAGAAAGAA
...
```

This makes it easy to write your alignment back into a file with all of the original info intact.

If you want to do more interesting things with an alignment, the best thing to do is to pass the alignment to an alignment information generating object, such as the `SummaryInfo` object, described in section [9.3.2](#).

9.3.2 Calculating summary information

Once you have an alignment, you are very likely going to want to find out information about it. Instead of trying to have all of the functions that can generate information about an alignment in the alignment object itself, we've tried to separate out the functionality into separate classes, which act on the alignment.

Getting ready to calculate summary information about an object is quick to do. Let's say we've got an alignment object called `alignment`. All we need to do to get an object that will calculate summary information is:

```
from Bio.Align import AlignInfo
summary_align = AlignInfo.SummaryInfo(alignment)
```

The `summary_align` object is very useful, and will do the following neat things for you:

1. Calculate a quick consensus sequence – see section [9.3.3](#)
2. Get a position specific score matrix for the alignment – see section [9.3.4](#)
3. Calculate the information content for the alignment – see section [9.3.5](#)
4. Generate information on substitutions in the alignment – section [9.4](#) details using this to generate a substitution matrix.

9.3.3 Calculating a quick consensus sequence

The `SummaryInfo` object, described in section [9.3.2](#), provides functionality to calculate a quick consensus of an alignment. Assuming we've got a `SummaryInfo` object called `summary_align` we can calculate a consensus by doing:

```
consensus = summary_align.dumb_consensus()
```

As the name suggests, this is a really simple consensus calculator, and will just add up all of the residues at each point in the consensus, and if the most common value is higher than some threshold value (the default is `.3`) will add the common residue to the consensus. If it doesn't reach the threshold, it adds an ambiguity character to the consensus. The returned consensus object is `Seq` object whose alphabet is inferred from the alphabets of the sequences making up the consensus. So doing a `print consensus` would give:

```
consensus Seq('TATACATNAAAGNAGGGGATGCGGATAAATGGAAAGGCGAAAGAAAGAAAAAATGAAT
...', IUPACAmbiguousDNA())
```

You can adjust how `dumb_consensus` works by passing optional parameters:

the threshold This is the threshold specifying how common a particular residue has to be at a position before it is added. The default is `.7`.

the ambiguous character This is the ambiguity character to use. The default is `'N'`.

the consensus alphabet This is the alphabet to use for the consensus sequence. If an alphabet is not specified then we will try to guess the alphabet based on the alphabets of the sequences in the alignment.

9.3.4 Position Specific Score Matrices

Position specific score matrices (PSSMs) summarize the alignment information in a different way than a consensus, and may be useful for different tasks. Basically, a PSSM is a count matrix. For each column in the alignment, the number of each alphabet letters is counted and totaled. The totals are displayed relative to some representative sequence along the left axis. This sequence may be the consensus sequence, but can also be any sequence in the alignment. For instance for the alignment,

```
GTATC
AT--C
CTGTC
```

the PSSM is:

```

      G A T C
G 1 1 0 1
T 0 0 3 0
A 1 1 0 0
T 0 0 2 0
C 0 0 0 3

```

Let's assume we've got an alignment object called `c_align`. To get a PSSM with the consensus sequence along the side we first get a summary object and calculate the consensus sequence:

```

summary_align = AlignInfo.SummaryInfo(c_align)
consensus = summary_align.dumb_consensus()

```

Now, we want to make the PSSM, but ignore any N ambiguity residues when calculating this:

```

my_pssm = summary_align.pos_specific_score_matrix(consensus,
                                                chars_to_ignore = ['N'])

```

Two notes should be made about this:

1. To maintain strictness with the alphabets, you can only include characters along the top of the PSSM that are in the alphabet of the alignment object. Gaps are not included along the top axis of the PSSM.
2. The sequence passed to be displayed along the left side of the axis does not need to be the consensus. For instance, if you wanted to display the second sequence in the alignment along this axis, you would need to do:

```

second_seq = alignment.get_seq_by_num(1)
my_pssm = summary_align.pos_specific_score_matrix(second_seq
                                                chars_to_ignore = ['N'])

```

The command above returns a PSSM object. To print out the PSSM as we showed above, we simply need to do a `print my_pssm`, which gives:

```

      A   C   G   T
T  0.0  0.0  0.0  7.0
A  7.0  0.0  0.0  0.0
T  0.0  0.0  0.0  7.0
A  7.0  0.0  0.0  0.0
C  0.0  7.0  0.0  0.0
A  7.0  0.0  0.0  0.0
T  0.0  0.0  0.0  7.0
T  1.0  0.0  0.0  6.0
...

```

You can access any element of the PSSM by subscripting like `your_pssm[sequence_number][residue_count_name]`. For instance, to get the counts for the 'A' residue in the second element of the above PSSM you would do:

```

>>> print my_pssm[1]["A"]
7.0

```

The structure of the PSSM class hopefully makes it easy both to access elements and to pretty print the matrix.

9.3.5 Information Content

A potentially useful measure of evolutionary conservation is the information content of a sequence.

A useful introduction to information theory targetted towards molecular biologists can be found at <http://www.lecb.ncifcrf.gov/~toms/paper/primer/>. For our purposes, we will be looking at the information content of a consensus sequence, or a portion of a consensus sequence. We calculate information content at a particular column in a multiple sequence alignment using the following formula:

$$IC_j = \sum_{i=1}^{N_a} P_{ij} \log \left(\frac{P_{ij}}{Q_i} \right)$$

where:

- IC_j – The information content for the j -th column in an alignment.
- N_a – The number of letters in the alphabet.
- P_{ij} – The frequency of a particular letter i in the j -th column (i. e. if G occurred 3 out of 6 times in an alignment column, this would be 0.5)
- Q_i – The expected frequency of a letter i . This is an optional argument, usage of which is left at the user's discretion. By default, it is automatically assigned to $0.05 = 1/20$ for a protein alphabet, and $0.25 = 1/4$ for a nucleic acid alphabet. This is for getting the information content without any assumption of prior distributions. When assuming priors, or when using a non-standard alphabet, you should supply the values for Q_i .

Well, now that we have an idea what information content is being calculated in Biopython, let's look at how to get it for a particular region of the alignment.

First, we need to use our alignment to get a alignment summary object, which we'll assume is called `summary_align` (see section 9.3.2) for instructions on how to get this. Once we've got this object, calculating the information content for a region is as easy as:

```
info_content = summary_align.information_content(5, 30,
                                                chars_to_ignore = ['N'])
```

Wow, that was much easier than the formula above made it look! The variable `info_content` now contains a float value specifying the information content over the specified region (from 5 to 30 of the alignment). We specifically ignore the ambiguity residue 'N' when calculating the information content, since this value is not included in our alphabet (so we shouldn't be interested in looking at it!).

As mentioned above, we can also calculate relative information content by supplying the expected frequencies:

```
expect_freq = {
    'A' : .3,
    'G' : .2,
    'T' : .3,
    'C' : .2}
```

The expected should not be passed as a raw dictionary, but instead by passed as a `SubsMat.FreqTable` object (see section 10.4.2 for more information about `FreqTables`). The `FreqTable` object provides a standard for associating the dictionary with an `Alphabet`, similar to how the Biopython `Seq` class works.

To create a `FreqTable` object, from the frequency dictionary you just need to do:


```

from Bio.Alphabet import IUPAC
from Bio.SubsMat import FreqTable

e_freq_table = FreqTable.FreqTable(expect_freq, FreqTable.FREQ,
                                   IUPAC.unambiguous_dna)

```

Now that we've got that, calculating the relative information content for our region of the alignment is as simple as:

```

info_content = summary_align.information_content(5, 30,
                                                e_freq_table = e_freq_table,
                                                chars_to_ignore = ['N'])

```

Now, `info_content` will contain the relative information content over the region in relation to the expected frequencies.

The value return is calculated using base 2 as the logarithm base in the formula above. You can modify this by passing the parameter `log_base` as the base you want:

```

info_content = summary_align.information_content(5, 30, log_base = 10
                                                chars_to_ignore = ['N'])

```

Well, now you are ready to calculate information content. If you want to try applying this to some real life problems, it would probably be best to dig into the literature on information content to get an idea of how it is used. Hopefully your digging won't reveal any mistakes made in coding this function!

9.3.6 Translating between Alignment formats

One thing that you always end up having to do is convert between different formats. You can do this using the `Bio.AlignIO` module, see Section 5.2.1.

9.4 Substitution Matrices

Substitution matrices are an extremely important part of everyday bioinformatics work. They provide the scoring terms for classifying how likely two different residues are to substitute for each other. This is essential in doing sequence comparisons. The book "Biological Sequence Analysis" by Durbin et al. provides a really nice introduction to Substitution Matrices and their uses. Some famous substitution matrices are the PAM and BLOSUM series of matrices.

Biopython provides a ton of common substitution matrices, and also provides functionality for creating your own substitution matrices.

9.4.1 Using common substitution matrices

9.4.2 Creating your own substitution matrix from an alignment

A very cool thing that you can do easily with the substitution matrix classes is to create your own substitution matrix from an alignment. In practice, this is normally done with protein alignments. In this example, we'll first get a biopython alignment object and then get a summary object to calculate info about the alignment. The file containing [protein.aln](#) (also available online [here](#)) contains the Clustalw alignment output.

```

from Bio import Clustalw
from Bio.Alphabet import IUPAC
from Bio.Align import AlignInfo

```

```
# get an alignment object from a Clustalw alignment output
c_align = Clustalw.parse_file("protein.aln", IUPAC.protein)
summary_align = AlignInfo.SummaryInfo(c_align)
```

Sections 9.3.1 and 9.3.2 contain more information on doing this.

Now that we've got our `summary_align` object, we want to use it to find out the number of times different residues substitute for each other. To make the example more readable, we'll focus on only amino acids with polar charged side chains. Luckily, this can be done easily when generating a replacement dictionary, by passing in all of the characters that should be ignored. Thus we'll create a dictionary of replacements for only charged polar amino acids using:

```
replace_info = summary_align.replacement_dictionary(["G", "A", "V", "L", "I",
                                                    "M", "P", "F", "W", "S",
                                                    "T", "N", "Q", "Y", "C"])
```

This information about amino acid replacements is represented as a python dictionary which will look something like:

```
{('R', 'R'): 2079.0, ('R', 'H'): 17.0, ('R', 'K'): 103.0, ('R', 'E'): 2.0,
 ('R', 'D'): 2.0, ('H', 'R'): 0, ('D', 'H'): 15.0, ('K', 'K'): 3218.0,
 ('K', 'H'): 24.0, ('H', 'K'): 8.0, ('E', 'H'): 15.0, ('H', 'H'): 1235.0,
 ('H', 'E'): 18.0, ('H', 'D'): 0, ('K', 'D'): 0, ('K', 'E'): 9.0,
 ('D', 'R'): 48.0, ('E', 'R'): 2.0, ('D', 'K'): 1.0, ('E', 'K'): 45.0,
 ('K', 'R'): 130.0, ('E', 'D'): 241.0, ('E', 'E'): 3305.0,
 ('D', 'E'): 270.0, ('D', 'D'): 2360.0}
```

This information gives us our accepted number of replacements, or how often we expect different things to substitute for each other. It turns out, amazingly enough, that this is all of the information we need to go ahead and create a substitution matrix. First, we use the replacement dictionary information to create an Accepted Replacement Matrix (ARM):

```
from Bio import SubsMat
my_arm = SubsMat.SeqMat(replace_info)
```

With this accepted replacement matrix, we can go right ahead and create our log odds matrix (i. e. a standard type Substitution Matrix):

```
my_lom = SubsMat.make_log_odds_matrix(my_arm)
```

The log odds matrix you create is customizable with the following optional arguments:

- `exp_freq_table` – You can pass a table of expected frequencies for each alphabet. If supplied, this will be used instead of the passed accepted replacement matrix when calculate expected replacements.
- `logbase` - The base of the logarithm taken to create the log odd matrix. Defaults to base 10.
- `factor` - The factor to multiply each matrix entry by. This defaults to 10, which normally makes the matrix numbers easy to work with.
- `round_digit` - The digit to round to in the matrix. This defaults to 0 (i. e. no digits).

Once you've got your log odds matrix, you can display it prettily using the function `print_mat`. Doing this on our created matrix gives:

```
>>> my_lom.print_mat()
D   6
E  -5   5
H -15 -13  10
K -31 -15 -13   6
R -13 -25 -14  -7   7
   D   E   H   K   R
```

Very nice. Now we've got our very own substitution matrix to play with!

9.5 BioSQL – storing sequences in a relational database

BioSQL is a joint effort between the **OBF** projects (BioPerl, BioJava etc) to support a shared database schema for storing sequence data. In theory, you could load a GenBank file into the database with BioPerl, then using Biopython extract this from the database as a record object with features - and get more or less the same thing as if you had loaded the GenBank file directly as a SeqRecord using `Bio.SeqIO` (Chapter 4).

Biopython's BioSQL module is currently documented at <http://biopython.org/wiki/BioSQL> which is part of our wiki pages.

9.6 BioCorba

Biorcorba does some cool stuff with CORBA. Basically, it allows you to easily interact with code written in other languages, including Perl and Java. This is all done through an interface which is very similar to the standard biopython interface. Much work has been done to make it easy to use knowing only very little about CORBA. You should check out the biorcorba specific documentation, which describes in detail how to use it.

9.7 Going 3D: The PDB module

Biopython also allows you to explore the extensive realm of macromolecular structure. Biopython comes with a PDBParser class that produces a Structure object. The Structure object can be used to access the atomic data in the file in a convenient manner.

9.7.1 Structure representation

A macromolecular structure is represented using a structure, model chain, residue, atom (or SMCRA) hierarchy. Fig. 9.7.1 shows a UML class diagram of the SMCRA data structure. Such a data structure is not necessarily best suited for the representation of the macromolecular content of a structure, but it is absolutely necessary for a good interpretation of the data present in a file that describes the structure (typically a PDB or MMCIF file). If this hierarchy cannot represent the contents of a structure file, it is fairly certain that the file contains an error or at least does not describe the structure unambiguously. If a SMCRA data structure cannot be generated, there is reason to suspect a problem. Parsing a PDB file can thus be used to detect likely problems. We will give several examples of this in section 9.7.5.1.

Structure, Model, Chain and Residue are all subclasses of the Entity base class. The Atom class only (partly) implements the Entity interface (because an Atom does not have children).

For each Entity subclass, you can extract a child by using a unique id for that child as a key (e.g. you can extract an Atom object from a Residue object by using an atom name string as a key, you can extract a Chain object from a Model object by using its chain identifier as a key).




Figure 9.1: UML diagram of the SMCRA data structure used to represent a macromolecular structure.

Disordered atoms and residues are represented by `DisorderedAtom` and `DisorderedResidue` classes, which are both subclasses of the `DisorderedEntityWrapper` base class. They hide the complexity associated with disorder and behave exactly as `Atom` and `Residue` objects.

In general, a child Entity object (i.e. `Atom`, `Residue`, `Chain`, `Model`) can be extracted from its parent (i.e. `Residue`, `Chain`, `Model`, `Structure`, respectively) by using an id as a key.

```
child_entity=parent_entity[child_id]
```

You can also get a list of all child Entities of a parent Entity object. Note that this list is sorted in a specific way (e.g. according to chain identifier for `Chain` objects in a `Model` object).

```
child_list=parent_entity.get_list()
```

You can also get the parent from a child.

```
parent_entity=child_entity.get_parent()
```

At all levels of the SMCRA hierarchy, you can also extract a *full id*. The full id is a tuple containing all id's starting from the top object (`Structure`) down to the current object. A full id for a `Residue` object e.g. is something like:

```
full_id=residue.get_full_id()
```

```
print full_id
```

```
("1abc", 0, "A", ("", 10, "A"))
```

This corresponds to:

- The `Structure` with id "1abc"
- The `Model` with id 0
- The `Chain` with id "A"
- The `Residue` with id (" ", 10, "A").

The `Residue` id indicates that the residue is not a hetero-residue (nor a water) because it has a blank hetero field, that its sequence identifier is 10 and that its insertion code is "A".

Some other useful methods:

```
# get the entity's id
```

```
entity.get_id()
```

```
# check if there is a child with a given id
```

```
entity.has_id(entity_id)
```

```
# get number of children
```

```
nr_children=len(entity)
```

It is possible to delete, rename, add, etc. child entities from a parent entity, but this does not include any sanity checks (e.g. it is possible to add two residues with the same id to one chain). This really should be done via a nice Decorator class that includes integrity checking, but you can take a look at the code (`Entity.py`) if you want to use the raw interface.

9.7.1.1 Structure

The Structure object is at the top of the hierarchy. Its id is a user given string. The Structure contains a number of Model children. Most crystal structures (but not all) contain a single model, while NMR structures typically consist of several models. Disorder in crystal structures of large parts of molecules can also result in several models.

9.7.1.1.1 Constructing a Structure object A Structure object is produced by a PDBParser object:

```
from Bio.PDB.PDBParser import PDBParser

p=PDBParser(PERMISSIVE=1)

structure_id="1fat"

filename="pdb1fat.ent"

s=p.get_structure(structure_id, filename)
```

The PERMISSIVE flag indicates that a number of common problems (see 9.7.5.1) associated with PDB files will be ignored (but note that some atoms and/or residues will be missing). If the flag is not present a PDBConstructionException will be generated during the parse operation.

9.7.1.1.2 Header and trailer You can extract the header and trailer (simple lists of strings) of the PDB file from the PDBParser object with the `get_header` and `get_trailer` methods.

9.7.1.2 Model

The id of the Model object is an integer, which is derived from the position of the model in the parsed file (they are automatically numbered starting from 0). The Model object stores a list of Chain children.

9.7.1.2.1 Example Get the first model from a Structure object.

```
first_model=structure[0]
```

9.7.1.3 Chain

The id of a Chain object is derived from the chain identifier in the structure file, and can be any string. Each Chain in a Model object has a unique id. The Chain object stores a list of Residue children.

9.7.1.3.1 Example Get the Chain object with identifier “A” from a Model object.

```
chain_A=model["A"]
```

9.7.1.4 Residue

Unsurprisingly, a Residue object stores a set of Atom children. In addition, it also contains a string that specifies the residue name (e.g. “ASN”) and the segment identifier of the residue (well known to X-PLOR users, but not used in the construction of the SMCRA data structure).

The id of a Residue object is composed of three parts: the hetero field (hetfield), the sequence identifier (resseq) and the insertion code (icode).

The hetero field is a string : it is “W” for waters, “H.” followed by the residue name (e.g. “H_FUC”) for other hetero residues and blank for standard amino and nucleic acids. This scheme is adopted for reasons described in section 9.7.3.1.

The second field in the Residue id is the sequence identifier, an integer describing the position of the residue in the chain.

The third field is a string, consisting of the insertion code. The insertion code is sometimes used to preserve a certain desirable residue numbering scheme. A Ser 80 insertion mutant (inserted e.g. between a Thr 80 and an Asn 81 residue) could e.g. have sequence identifiers and insertion codes as followed: Thr 80 A, Ser 80 B, Asn 81. In this way the residue numbering scheme stays in tune with that of the wild type structure.

Let's give some examples. Asn 10 with a blank insertion code would have residue id ('' ', 10, '' '). Water 10 would have residue id (''W'', 10, '' '). A glucose molecule (a hetero residue with residue name GLC) with sequence identifier 10 would have residue id (''H_GLC'', 10, '' '). In this way, the three residues (with the same insertion code and sequence identifier) can be part of the same chain because their residue id's are distinct.

In most cases, the hetflag and insertion code fields will be blank, e.g. ('' ', 10, '' '). In these cases, the sequence identifier can be used as a shortcut for the full id:

```
# use full id

res10=chain(["", 10, ""])

# use shortcut

res10=chain[10]
```

Each Residue object in a Chain object should have a unique id. However, disordered residues are dealt with in a special way, as described in section [9.7.2.3.2](#).

A Residue object has a number of additional methods:

```
r.get_resname() # return residue name, e.g. "ASN"
r.get_segid() # return the SEGID, e.g. "CHN1"
```

9.7.1.5 Atom

The Atom object stores the data associated with an atom, and has no children. The id of an atom is its atom name (e.g. "OG" for the side chain oxygen of a Ser residue). An Atom id needs to be unique in a Residue. Again, an exception is made for disordered atoms, as described in section [9.7.2.2](#).

In a PDB file, an atom name consists of 4 chars, typically with leading and trailing spaces. Often these spaces can be removed for ease of use (e.g. an amino acid C α atom is labeled ".CA." in a PDB file, where the dots represent spaces). To generate an atom name (and thus an atom id) the spaces are removed, unless this would result in a name collision in a Residue (i.e. two Atom objects with the same atom name and id). In the latter case, the atom name including spaces is tried. This situation can e.g. happen when one residue contains atoms with names ".CA." and "CA.", although this is not very likely.

The atomic data stored includes the atom name, the atomic coordinates (including standard deviation if present), the B factor (including anisotropic B factors and standard deviation if present), the altloc specifier and the full atom name including spaces. Less used items like the atom element number or the atomic charge sometimes specified in a PDB file are not stored.

An Atom object has the following additional methods:

```
a.get_name() # atom name (spaces stripped, e.g. "CA")
a.get_id() # id (equals atom name)
a.get_coord() # atomic coordinates
a.get_bfactor() # B factor
a.get_occupancy() # occupancy
a.get_altloc() # alternative location specifie
```

```

a.get_sigatm()      # std. dev. of atomic parameters
a.get_siguij()     # std. dev. of anisotropic B factor
a.get_anisou()     # anisotropic B factor
a.get_fullname()   # atom name (with spaces, e.g. ".CA.")

```

To represent the atom coordinates, siguij, anisotropic B factor and sigatm Numpy arrays are used.

9.7.2 Disorder

9.7.2.1 General approach

Disorder should be dealt with from two points of view: the atom and the residue points of view. In general, we have tried to encapsulate all the complexity that arises from disorder. If you just want to loop over all C α atoms, you do not care that some residues have a disordered side chain. On the other hand it should also be possible to represent disorder completely in the data structure. Therefore, disordered atoms or residues are stored in special objects that behave as if there is no disorder. This is done by only representing a subset of the disordered atoms or residues. Which subset is picked (e.g. which of the two disordered OG side chain atom positions of a Ser residue is used) can be specified by the user.

9.7.2.2 Disordered atoms

Disordered atoms are represented by ordinary Atom objects, but all Atom objects that represent the same physical atom are stored in a DisorderedAtom object. Each Atom object in a DisorderedAtom object can be uniquely indexed using its altloc specifier. The DisorderedAtom object forwards all uncaught method calls to the selected Atom object, by default the one that represents the atom with with the highest occupancy. The user can of course change the selected Atom object, making use of its altloc specifier. In this way atom disorder is represented correctly without much additional complexity. In other words, if you are not interested in atom disorder, you will not be bothered by it.

Each disordered atom has a characteristic altloc identifier. You can specify that a DisorderedAtom object should behave like the Atom object associated with a specific altloc identifier:

```

atom.disordered\_select("A") # select altloc A atom

print atom.get_altloc()
"A"

atom.disordered\_select("B")      # select altloc B atom
print atom.get_altloc()
"B"

```

9.7.2.3 Disordered residues

9.7.2.3.1 Common case The most common case is a residue that contains one or more disordered atoms. This is evidently solved by using DisorderedAtom objects to represent the disordered atoms, and storing the DisorderedAtom object in a Residue object just like ordinary Atom objects. The DisorderedAtom will behave exactly like an ordinary atom (in fact the atom with the highest occupancy) by forwarding all uncaught method calls to one of the Atom objects (the selected Atom object) it contains.

9.7.2.3.2 Point mutations A special case arises when disorder is due to a point mutation, i.e. when two or more point mutants of a polypeptide are present in the crystal. An example of this can be found in PDB structure 1EN2.

Since these residues belong to a different residue type (e.g. let's say Ser 60 and Cys 60) they should not be stored in a single Residue object as in the common case. In this case, each residue is represented by one Residue object, and both Residue objects are stored in a DisorderedResidue object.

The DisorderedResidue object forwards all uncaught methods to the selected Residue object (by default the last Residue object added), and thus behaves like an ordinary residue. Each Residue object in a DisorderedResidue object can be uniquely identified by its residue name. In the above example, residue Ser 60 would have id "SER" in the DisorderedResidue object, while residue Cys 60 would have id "CYS". The user can select the active Residue object in a DisorderedResidue object via this id.

9.7.3 Hetero residues

9.7.3.1 Associated problems

A common problem with hetero residues is that several hetero and non-hetero residues present in the same chain share the same sequence identifier (and insertion code). Therefore, to generate a unique id for each hetero residue, waters and other hetero residues are treated in a different way.

Remember that Residue objects have the tuple (hetfield, resseq, icode) as id. The hetfield is blank (" ") for amino and nucleic acids, and a string for waters and other hetero residues. The content of the hetfield is explained below.

9.7.3.2 Water residues

The hetfield string of a water residue consists of the letter "W". So a typical residue id for a water is ("W", 1, " ").

9.7.3.3 Other hetero residues

The hetfield string for other hetero residues starts with "H_" followed by the residue name. A glucose molecule e.g. with residue name "GLC" would have hetfield "H_GLC". Its residue id could e.g. be ("H_GLC", 1, " ").

9.7.4 Some random usage examples

Parse a PDB file, and extract some Model, Chain, Residue and Atom objects.

```
from PDBParser import PDBParser

parser=PDBParser()

structure=parser.get_structure("test", "1fat.pdb")
model=structure[0]
chain=model["A"]
residue=chain[1]
atom=residue["CA"]
```

Extract a hetero residue from a chain (e.g. a glucose (GLC) moiety with resseq 10).

```
residue_id=("H_GLC", 10, " ")
residue=chain[residue_id]
```

Print all hetero residues in chain.

```

for residue in chain.get_list():
    residue_id=residue.get_id()
    hetfield=residue_id[0]
    if hetfield[0]=="H":
    print residue_id

```

Print out the coordinates of all CA atoms in a structure with B factor greater than 50.

```

for model in structure.get_list():
    for chain in model.get_list():
        for residue in chain.get_list():
            if residue.has_id("CA"):
                ca=residue["CA"]
                if ca.get_bfactor(>50.0):
                    print ca.get_coord()

```

Print out all the residues that contain disordered atoms.

```

for model in structure.get_list()
for chain in model.get_list():
    for residue in chain.get_list():
        if residue.is_disordered():
            resseq=residue.get_id()[1]
            resname=residue.get_resname()
            model_id=model.get_id()
            chain_id=chain.get_id()
            print model_id, chain_id, resname, resseq

```

Loop over all disordered atoms, and select all atoms with altloc A (if present). This will make sure that the SMCRA data structure will behave as if only the atoms with altloc A are present.

```

for model in structure.get_list()
for chain in model.get_list():
    for residue in chain.get_list():
        if residue.is_disordered():
            for atom in residue.get_list():
                if atom.is_disordered():
                    if atom.disordered_has_id("A"):
                        atom.disordered_select("A")

```

Suppose that a chain has a point mutation at position 10, consisting of a Ser and a Cys residue. Make sure that residue 10 of this chain behaves as the Cys residue.

```

residue=chain[10]
residue.disordered_select("CYS")

```

9.7.5 Common problems in PDB files

9.7.5.1 Examples

The PDBParser/Structure class was tested on about 800 structures (each belonging to a unique SCOP superfamily). This takes about 20 minutes, or on average 1.5 seconds per structure. Parsing the structure of the large ribosomal subunit (1FKK), which contains about 64000 atoms, takes 10 seconds on a 1000 MHz PC.

Three exceptions were generated in cases where an unambiguous data structure could not be built. In all three cases, the likely cause is an error in the PDB file that should be corrected. Generating an exception in these cases is much better than running the chance of incorrectly describing the structure in a data structure.

9.7.5.1.1 Duplicate residues One structure contains two amino acid residues in one chain with the same sequence identifier (resseq 3) and icode. Upon inspection it was found that this chain contains the residues Thr A3, . . . , Gly A202, Leu A3, Glu A204. Clearly, Leu A3 should be Leu A203. A couple of similar situations exist for structure 1FFK (which e.g. contains Gly B64, Met B65, Glu B65, Thr B67, i.e. residue Glu B65 should be Glu B66).

9.7.5.1.2 Duplicate atoms Structure 1EJG contains a Ser/Pro point mutation in chain A at position 22. In turn, Ser 22 contains some disordered atoms. As expected, all atoms belonging to Ser 22 have a non-blank altloc specifier (B or C). All atoms of Pro 22 have altloc A, except the N atom which has a blank altloc. This generates an exception, because all atoms belonging to two residues at a point mutation should have non-blank altloc. It turns out that this atom is probably shared by Ser and Pro 22, as Ser 22 misses the N atom. Again, this points to a problem in the file: the N atom should be present in both the Ser and the Pro residue, in both cases associated with a suitable altloc identifier.

9.7.5.2 Automatic correction

Some errors are quite common and can be easily corrected without much risk of making a wrong interpretation. These cases are listed below.

9.7.5.2.1 A blank altloc for a disordered atom Normally each disordered atom should have a non-blanc altloc identifier. However, there are many structures that do not follow this convention, and have a blank and a non-blank identifier for two disordered positions of the same atom. This is automatically interpreted in the right way.

9.7.5.2.2 Broken chains Sometimes a structure contains a list of residues belonging to chain A, followed by residues belonging to chain B, and again followed by residues belonging to chain A, i.e. the chains are “broken”. This is correctly interpreted.

9.7.5.3 Fatal errors

Sometimes a PDB file cannot be unambiguously interpreted. Rather than guessing and risking a mistake, an exception is generated, and the user is expected to correct the PDB file. These cases are listed below.

9.7.5.3.1 Duplicate residues All residues in a chain should have a unique id. This id is generated based on:

- The sequence identifier (resseq).
- The insertion code (icode).
- The hetfield string (“W” for waters and “H_” followed by the residue name for other hetero residues)
- The residue names of the residues in the case of point mutations (to store the Residue objects in a DisorderedResidue object).

If this does not lead to a unique id something is quite likely wrong, and an exception is generated.

9.7.5.3.2 Duplicate atoms All atoms in a residue should have a unique id. This id is generated based on:

- The atom name (without spaces, or with spaces if a problem arises).
- The altloc specifier.

If this does not lead to a unique id something is quite likely wrong, and an exception is generated.

9.7.6 Other features

There are also some tools to analyze a crystal structure. Tools exist to superimpose two coordinate sets (SVDSuperimposer), to extract polypeptides from a structure (Polypeptide), to perform neighbor lookup (NeighborSearch) and to write out PDB files (PDBIO). The neighbor lookup is done using a KD tree module written in C++. It is very fast and also includes a fast method to find all point pairs within a certain distance of each other.

A Polypeptide object is simply a UserList of Residue objects. You can construct a list of Polypeptide objects from a Structure object as follows:

```
model_nr=1
polypeptide_list=build_peptides(structure, model_nr)

for polypeptide in polypeptide_list:
    print polypeptide
```

The Polypeptide objects are always created from a single Model (in this case model 1).

9.8 Bio.PopGen: Population genetics

Bio.PopGen is a new Biopython module supporting population genetics, available in Biopython 1.44 onwards.

The medium term objective for the module is to support widely used data formats, applications and databases. This module is currently under intense development and support for new features should appear at a rather fast pace. Unfortunately this might also entail some instability on the API, especially if you are using a CVS version. APIs that are made available on public versions should be much stabler.

9.8.1 GenePop

GenePop (<http://genepop.curtin.edu.au/>) is a popular population genetics software package supporting Hardy-Weinberg tests, linkage disequilibrium, population differentiation, basic statistics, F_{st} and migration estimates, among others. GenePop does not supply sequence based statistics as it doesn't handle sequence data. The GenePop file format is supported by a wide range of other population genetic software applications, thus making it a relevant format in the population genetics field.

Bio.PopGen provides a parser and generator of GenePop file format. Utilities to manipulate the content of a record are also provided. Here is an example on how to read a GenePop file (you can find example GenePop data files in the Test/PopGen directory of Biopython):

```
from Bio.PopGen import GenePop

handle = open("example.gen")
rec = GenePop.parse(handle)
handle.close()
```

This will read a file called example.gen and parse it. If you do print rec, the record will be output again, in GenePop format.

The most important information in rec will be the loci names and population information (but there is more – use help(GenePop.Record) to check the API documentation). Loci names can be found on rec.loci_list. Population information can be found on rec.populations. Populations is a list with one element per population. Each element is itself a list of individuals, each individual is a pair composed by individual name and a list of alleles (2 per marker), here is an example for rec.populations:

```
[
  [
```

```

    ('Ind1', [(1, 2), (3, 3), (200, 201)],
    ('Ind2', [(2, None), (3, 3), (None, None)],
],
[
    ('Other1', [(1, 1), (4, 3), (200, 200)],
]
]

```

So we have two populations, the first with two individuals, the second with only one. The first individual of the first population is called Ind1, allelic information for each of the 3 loci follows. Please note that for any locus, information might be missing (see as an example, Ind2 above).

A few utility functions to manipulate GenePop records are made available, here is an example:

```

from Bio.PopGen import GenePop

#Imagine that you have loaded rec, as per the code snippet above...

rec.remove_population(pos)
#Removes a population from a record, pos is the population position in
# rec.populations, remember that it starts on position 0.
# rec is altered.

rec.remove_locus_by_position(pos)
#Removes a locus by its position, pos is the locus position in
# rec.loci_list, remember that it starts on position 0.
# rec is altered.

rec.remove_locus_by_name(name)
#Removes a locus by its name, name is the locus name as in
# rec.loci_list. If the name doesn't exist the function fails
# silently.
# rec is altered.

rec_loci = rec.split_in_loci()
#Splits a record in loci, that is, for each loci, it creates a new
# record, with a single loci and all populations.
# The result is returned in a dictionary, being each key the locus name.
# The value is the GenePop record.
# rec is not altered.

rec_pops = rec.split_in_pops(pop_names)
#Splits a record in populations, that is, for each population, it creates
# a new record, with a single population and all loci.
# The result is returned in a dictionary, being each key
# the population name. As population names are not available in GenePop,
# they are passed in array (pop_names).
# The value of each dictionary entry is the GenePop record.
# rec is not altered.

```

GenePop does not support population names, a limitation which can be cumbersome at times. Functionality to enable population names is currently being planned for Biopython. These extensions won't break compatibility in any way with the standard format. In the medium term, we would also like to support the GenePop web service.

9.8.2 Coalescent simulation

A coalescent simulation is a backward model of population genetics with relation to time. A simulation of ancestry is done until the Most Recent Common Ancestor (MRCA) is found. This ancestry relationship starting on the MRCA and ending on the current generation sample is sometimes called a genealogy. Simple cases assume a population of constant size in time, haploidy, no population structure, and simulate the alleles of a single locus under no selection pressure.

Coalescent theory is used in many fields like selection detection, estimation of demographic parameters of real populations or disease gene mapping.

The strategy followed in the Biopython implementation of the coalescent was not to create a new, built-in, simulator from scratch but to use an existing one, SIMCOAL2 (<http://cmpg.unibe.ch/software/simcoal2/>). SIMCOAL2 allows for, among others, population structure, multiple demographic events, simulation of multiple types of loci (SNPs, sequences, STRs/microsatellites and RFLPs) with recombination, diploidy multiple chromosomes or ascertainment bias. Notably SIMCOAL2 doesn't support any selection model. We recommend reading SIMCOAL2's documentation, available in the link above.

The input for SIMCOAL2 is a file specifying the desired demography and genome, the output is a set of files (typically around 1000) with the simulated genomes of a sample of individuals per subpopulation. This set of files can be used in many ways, like to compute confidence intervals where which certain statistics (e.g., F_{st} or Tajima D) are expected to lie. Real population genetics datasets statistics can then be compared to those confidence intervals.

Biopython coalescent code allows to create demographic scenarios and genomes and to run SIMCOAL2.

9.8.2.1 Creating scenarios

Creating a scenario involves both creating a demography and a chromosome structure. In many cases (e.g. when doing Approximate Bayesian Computations – ABC) it is important to test many parameter variations (e.g. vary the effective population size, N_e , between 10, 50, 500 and 1000 individuals). The code provided allows for the simulation of scenarios with different demographic parameters very easily.

Below we see how we can create scenarios and then how simulate them.

9.8.2.1.1 Demography A few predefined demographies are built-in, all have two shared parameters: sample size (called `sample_size` on the template, see below for its use) per deme and deme size, i.e. subpopulation size (`pop_size`). All demographies are available as templates where all parameters can be varied, each template has a system name. The predefined demographies/templates are:

Single population, constant size The standard parameters are enough to specify it. Template name: `simple`.

Single population, bottleneck As seen on figure 9.2. The parameters are current population size (`pop_size` on template `ne3` on figure), time of expansion, given as the generation in the past when it occurred (`expand_gen`), effective population size during bottleneck (`ne2`), time of contraction (`contract_gen`) and original size in the remote past (`ne3`). Template name: `bottle`.

Island model The typical island model. The total number of demes is specified by `total_demes` and the migration rate by `mig`. Template name `island`.

Stepping stone model - 1 dimension The stepping stone model in 1 dimension, extremes disconnected. The total number of demes is `total_demes`, migration rate is `mig`. Template name is `ssm_1d`.

Stepping stone model - 2 dimensions The stepping stone model in 2 dimensions, extremes disconnected. The parameters are `x` for the horizontal dimension and `y` for the vertical (being the total number of demes `x` times `y`), migration rate is `mig`. Template name is `ssm_2d`.

In our first example, we will generate a template for a single population, constant size model with a sample size of 30 and a deme size of 500. The code for this is:




Figure 9.2: A bottleneck

```
from Bio.PopGen.SimCoal.Template import generate_simcoal_from_template

generate_simcoal_from_template('simple',
    [(1, [('SNP', [24, 0.0005, 0.0])])],
    [('sample_size', [30]),
     ('pop_size', [100])])
```

Executing this code snippet will generate a file on the current directory called `simple_100.300.par` this file can be given as input to SIMCOAL2 to simulate the demography (below we will see how Biopython can take care of calling SIMCOAL2).

This code consists of a single function call, lets discuss it parameter by parameter.

The first parameter is the template id (from the list above). We are using the id 'simple' which is the template for a single population of constant size along time.

The second parameter is the chromosome structure. Please ignore it for now, it will be explained in the next section.

The third parameter is a list of all required parameters (recall that the simple model only needs `sample_size` and `pop_size`) and possible values (in this case each parameter only has a possible value).

Now, lets consider an example where we want to generate several island models, and we are interested in varying the number of demes: 10, 50 and 100 with a migration rate of 1%. Sample size and deme size will be the same as before. Here is the code:

```
from Bio.PopGen.SimCoal.Template import generate_simcoal_from_template

generate_simcoal_from_template('island',
    [(1, [('SNP', [24, 0.0005, 0.0])])],
    [('sample_size', [30]),
     ('pop_size', [100]),
     ('mig', [0.01]),
     ('total_demes', [10, 50, 100])])
```

In this case, 3 files will be generated: `island_100.0.01.100.30.par`, `island_10.0.01.100.30.par` and `island_50.0.01.100.30.par`. Notice the rule to make file names: template name, followed by parameter values in reverse order.

A few, arguably more esoteric template demographies exist (please check the `Bio/PopGen/SimCoal/data` directory on Biopython source tree). Furthermore it is possible for the user to create new templates. That functionality will be discussed in a future version of this document.

9.8.2.1.2 Chromosome structure We strongly recommend reading SIMCOAL2 documentation to understand the full potential available in modeling chromosome structures. In this subsection we only discuss

how to implement chromosome structures using the Biopython interface, not the underlying SIMCOAL2 capabilities.

We will start by implementing a single chromosome, with 24 SNPs with a recombination rate immediately on the right of each locus of 0.0005 and a minimum frequency of the minor allele of 0. This will be specified by the following list (to be passed as second parameter to the function `generate_simcoal_from_template`):

```
[(1, [('SNP', [24, 0.0005, 0.0]])]
```

This is actually the chromosome structure used in the above examples.

The chromosome structure is represented by a list of chromosomes, each chromosome (i.e., each element in the list) is composed by a tuple (a pair): the first element is the number of times the chromosome is to be repeated (as there might be interest in repeating the same chromosome many times). The second element is a list of the actual components of the chromosome. Each element is again a pair, the first member is the locus type and the second element the parameters for that locus type. Confused? Before showing more examples lets review the example above: We have a list with one element (thus one chromosome), the chromosome is a single instance (therefore not to be repeated), it is composed of 24 SNPs, with a recombination rate of 0.0005 between each consecutive SNP, the minimum frequency of the minor allele is 0.0 (i.e, it can be absent from a certain population).

Lets see a more complicated example:

```
[
  (5, [
    ('SNP', [24, 0.0005, 0.0])
  ]),
  (2, [
    ('DNA', [10, 0.0, 0.00005, 0.33]),
    ('RFLP', [1, 0.0, 0.0001]),
    ('MICROSAT', [1, 0.0, 0.001, 0.0, 0.0])
  ])
]
```

We start by having 5 chromosomes with the same structure as above (i.e., 24 SNPs). We then have 2 chromosomes which have a DNA sequence with 10 nucleotides, 0.0 recombination rate, 0.0005 mutation rate, and a transition rate of 0.33. Then we have an RFLP with 0.0 recombination rate to the next locus and a 0.0001 mutation rate. Finally we have a microsatellite (or STR), with 0.0 recombination rate to the next locus (note, that as this is a single microsatellite which has no loci following, this recombination rate here is irrelevant), with a mutation rate of 0.001, geometric parameter of 0.0 and a range constraint of 0.0 (for information about this parameters please consult the SIMCOAL2 documentation, you can use them to simulate various mutation models, including the typical – for microsatellites – stepwise mutation model among others).

9.8.2.2 Running SIMCOAL2

We now discuss how to run SIMCOAL2 from inside Biopython. It is required that the binary for SIMCOAL2 is called `simcoal2` (or `simcoal2.exe` on Windows based platforms), please note that the typical name when downloading the program is in the format `simcoal2_x.y`. As such renaming of the binary after download is needed.

It is possible to run SIMCOAL2 on files that were not generated using the method above (e.g., writing a parameter file by hand), but we will show an example by creating a model using the framework presented above.


```

from Bio.PopGen.SimCoal.Template import generate_simcoal_from_template
from Bio.PopGen.SimCoal.Controller import SimCoalController

generate_simcoal_from_template('simple',
    [
        (5, [
            ('SNP', [24, 0.0005, 0.0])
        ]),
        (2, [
            ('DNA', [10, 0.0, 0.00005, 0.33]),
            ('RFLP', [1, 0.0, 0.0001]),
            ('MICROSAT', [1, 0.0, 0.001, 0.0, 0.0])
        ])
    ],
    [('sample_size', [30]),
     ('pop_size', [100])])

ctrl = SimCoalController('.')
ctrl.run_simcoal('simple_100_30.par', 50)

```

The lines of interest are the last two (plus the new import). Firstly a controller for the application is created. The directory where the binary is located has to be specified.

The simulator is then run on the last line: we know, from the rules explained above, that the input file name is `simple_100_30.par` for the simulation parameter file created. We then specify that we want to run 50 independent simulations, by default Biopython requests a simulation of diploid data, but a third parameter can be added to simulate haploid data (adding as a parameter the string '0'). SIMCOAL2 will now run (please note that this can take quite a lot of time) and will create a directory with the simulation results. The results can now be analysed (typically studying the data with Arlequin3). In the future Biopython might support reading the Arlequin3 format and thus allowing for the analysis of SIMCOAL2 data inside Biopython.

9.8.3 Other applications

Here we discuss interfaces and utilities to deal with population genetics' applications which arguably have a smaller user base.

9.8.3.1 FDist: Detecting selection and molecular adaptation

FDist is a selection detection application suite based on computing (i.e. simulating) a “neutral” confidence interval based on F_{st} and heterozygosity. Markers (which can be SNPs, microsatellites, AFLPs among others) which lie outside the “neutral” interval are to be considered as possible candidates for being under selection.

FDist is mainly used when the number of markers is considered enough to estimate an average F_{st} , but not enough to either have outliers calculated from the dataset directly or, with even more markers for which the relative positions in the genome are known, to use approaches based on, e.g., Extended Haplotype Heterozygosity (EHH).

The typical usage pattern for FDist is as follows:

1. Import a dataset from an external format into FDist format.

2. Compute average F_{st} . This is done by datacal inside FDist.
3. Simulate “neutral” markers based on the average F_{st} and expected number of total populations. This is the core operation, done by fdist inside FDist.
4. Calculate the confidence interval, based on the desired confidence boundaries (typically 95% or 99%). This is done by cplot and is mainly used to plot the interval.
5. Assess each marker status against the simulation “neutral” confidence interval. Done by pv. This is used to detect the outlier status of each marker against the simulation.

We will now discuss each step with illustrating example code (for this example to work FDist binaries have to be on the executable PATH).

The FDist data format is application specific and is not used at all by other applications, as such you will probably have to convert your data for use with FDist. Biopython can help you do this. Here is an example converting from GenePop format to FDist format (along with imports that will be needed on examples further below):

```
from Bio.PopGen import GenePop
from Bio.PopGen import FDist
from Bio.PopGen.FDist import Controller
from Bio.PopGen.FDist.Utils import convert_genepop_to_fdist

gp_rec = GenePop.parse(open("example.gen"))
fd_rec = convert_genepop_to_fdist(gp_rec)
in_file = open("infile", "w")
in_file.write(str(fd_rec))
in_file.close()
```

In this code we simply parse a GenePop file and convert it to a FDist record.

Printing an FDist record will generate a string that can be directly saved to a file and supplied to FDist. FDist requires the input file to be called infile, therefore we save the record on a file with that name.

The most important fields on a FDist record are: num_pops, the number of populations; num_loci, the number of loci and loci_data with the marker data itself. Most probably the details of the record are of no interest to the user, as the record only purpose is to be passed to FDist.

The next step is to calculate the average F_{st} of the dataset (along with the sample size):

```
ctrl = Controller.FDistController()
fst, samp_size = ctrl.run_datacal()
```

On the first line we create an object to control the call of FDist suite, this object will be used further on in order to call other suite applications.

On the second line we call the datacal application which computes the average F_{st} and the sample size. It is worth noting that the F_{st} computed by datacal is a *variation* of Weir and Cockerham’s θ .

We can now call the main fdist application in order to simulate neutral markers.

```
sim_fst = ctrl.run_fdist(npops = 15, nsamples = fd_rec.num_pops, fst = fst,
    sample_size = samp_size, mut = 0, num_sims = 40000)
```

npops Number of populations existing in nature. This is really a “guestimate”. Has to be lower than 100.

nsamples Number of populations sampled, has to be lower than npops.

fst Average F_{st} .

sample_size Average number of individuals sampled on each population.

mut Mutation model: 0 - Infinite alleles; 1 - Stepwise mutations

num_sims Number of simulations to perform. Typically a number around 40000 will be OK, but if you get a confidence interval that looks sharp (this can be detected when plotting the confidence interval computed below) the value can be increased (a suggestion would be steps of 10000 simulations).

The confusion in wording between number of samples and sample size stems from the original application. A file named out.dat will be created with the simulated heterozygosities and F_{st} s, it will have as many lines as the number of simulations requested.

Note that fdist returns the average F_{st} that it was *capable* of simulating, for more details about this issue please read below the paragraph on approximating the desired average F_{st} .

The next (optional) step is to calculate the confidence interval:

```
cpl_interval = ctrl.run_cplot(ci=0.99)
```

You can only call cplot after having run fdist.

This will calculate the confidence intervals (99% in this case) for a previous fdist run. A list of quadruples is returned. The first element represents the heterozygosity, the second the lower bound of F_{st} confidence interval for that heterozygosity, the third the average and the fourth the upper bound. This can be used to trace the confidence interval contour. This list is also written to a file, out.cpl.

The main purpose of this step is return a set of points which can be easily used to plot a confidence interval. It can be skipped if the objective is only to assess the status of each marker against the simulation, which is the next step...

```
pv_data = ctrl.run_pv()
```

You can only call cplot after having run datacal and fdist.

This will use the simulated markers to assess the status of each individual real marker. A list, in the same order than the loci_list that is on the FDist record (which is in the same order that the GenePop record) is returned. Each element in the list is a quadruple, the fundamental member of each quadruple is the last element (regarding the other elements, please refer to the pv documentation – for the sake of simplicity we will not discuss them here) which returns the probability of the simulated F_{st} being lower than the marker F_{st} . Higher values would indicate a stronger candidate for positive selection, lower values a candidate for balancing selection, and intermediate values a possible neutral marker. What is “higher”, “lower” or “intermediate” is really a subjective issue, but taking a “confidence interval” approach and considering a 95% confidence interval, “higher” would be between 0.95 and 1.0, “lower” between 0.0 and 0.05 and “intermediate” between 0.05 and 0.95.

9.8.3.1.1 Approximating the desired average F_{st} Fdist tries to approximate the desired average F_{st} by doing a coalescent simulation using migration rates based on the formula

$$N_m = \frac{1 - F_{st}}{4F_{st}}$$

This formula assumes a few premises like an infinite number of populations.

In practice, when the number of populations is low, the mutation model is stepwise and the sample size increases, fdist will not be able to simulate an acceptable approximate average F_{st} .

To address that, a function is provided to iteratively approach the desired value by running several fdist in sequence. This approach is computationally more intensive than running a single fdist run, but yields good results. The following code runs fdist approximating the desired F_{st} :

```
sim_fst = ctrl.run_fdist_force_fst(npops = 15, nsamples = fd_rec.num_pops,
    fst = fst, sample_size = samp_size, mut = 0, num_sims = 40000,
    limit = 0.05)
```

The only new optional parameter, when comparing with `run_fdist`, is `limit` which is the desired maximum error. `run_fdist` can (and probably should) be safely replaced with `run_fdist_force_fst`.

9.8.3.1.2 Final notes The process to determine the average F_{st} can be more sophisticated than the one presented here. For more information we refer you to the FDist README file. Biopython's code can be used to implement more sophisticated approaches.

9.8.4 Future Developments

The most desired future developments would be the ones you add yourself ;)

That being said, already existing fully functional code is currently being incorporated in `Bio.PopGen`, that code covers the applications `FDist` and `SimCoal2`, the `HapMap` and `UCSC Table Browser` databases and some simple statistics like F_{st} , or allele counts.

9.9 InterPro

The `Bio.InterPro` module works with files from the InterPro database, which can be obtained from the InterPro project: <http://www.ebi.ac.uk/interpro/>.

The `Bio.InterPro.Record` contains all the information stored in an InterPro record. Its string representation also is a valid InterPro record, but it is NOT guaranteed to be equivalent to the record from which it was produced.

The `Bio.InterPro.Record` contains:

- Database
- Accession
- Name
- Dates
- Type
- Parent
- Process
- Function
- Component
- Signatures
- Abstract
- Examples
- References
- Database links

Chapter 10

Advanced

10.1 The SeqRecord and SeqFeature classes

You read all about the basic Biopython sequence class in Chapter 3, which described how to do many useful things with just the sequence. However, many times sequences have important additional properties associated with them – as you will have seen with the `SeqRecord` object in Chapter 4. This section described how Biopython handles these higher level descriptions of a sequence.

10.1.1 Sequence ids and Descriptions – dealing with SeqRecords

Immediately above the Sequence class is the Sequence Record class, defined in the `Bio.SeqRecord` module. This class allows higher level features such as ids and features to be associated with the sequence, and is used throughout the sequence input/output interface `Bio.SeqIO`, described in Chapter 4. The `SeqRecord` class itself is very simple, and offers the following information as attributes:

seq – The sequence itself – A `Seq` object

id – The primary id used to identify the sequence. In most cases this is something like an accession number.

name – A “common” name/id for the sequence. In some cases this will be the same as the accession number, but it could also be a clone name. I think of this as being analagous to the LOCUS id in a GenBank record.

description – A human readable description or expressive name for the sequence. This is similar to what follows the id information in a FASTA formatted entry.

annotations – A dictionary of additional information about the sequence. The keys are the name of the information, and the information is contained in the value. This allows the addition of more “unstructured” information to the sequence.

features – A list of `SeqFeature` objects with more structured information about the features on a sequence. The structure of sequence features is described below in Section 10.1.2.

Using a `SeqRecord` class is not very complicated, since all of the information is stored as attributes of the class. Initializing the class just involves passing a `Seq` object to the `SeqRecord`:

```
>>> from Bio.Seq import Seq
>>> simple_seq = Seq("GATC")
>>> from Bio.SeqRecord import SeqRecord
>>> simple_seq_r = SeqRecord(simple_seq)
```

Additionally, you can also pass the id, name and description to the initialization function, but if not they will be set as strings indicating they are unknown, and can be modified subsequently:

```
>>> simple_seq_r.id
'<unknown id>'
>>> simple_seq_r.id = 'AC12345'
>>> simple_seq_r.description = 'My little made up sequence I wish I could
write a paper about and submit to GenBank'
>>> print simple_seq_r.description
My little made up sequence I wish I could write a paper about and submit
to GenBank
>>> simple_seq_r.seq
Seq('GATC', Alphabet())
```

Adding annotations is almost as easy, and just involves dealing directly with the annotation dictionary:

```
>>> simple_seq_r.annotations['evidence'] = 'None. I just made it up.'
>>> print simple_seq_r.annotations
{'evidence': 'None. I just made it up.'}
```

That's just about all there is to it! Next, you may want to learn about `SeqFeatures`, which offer an additional structured way to represent information about a sequence.

10.1.2 Features and Annotations – `SeqFeatures`

Sequence features are an essential part of describing a sequence. Once you get beyond the sequence itself, you need some way to organize and easily get at the more “abstract” information that is known about the sequence. While it is probably impossible to develop a general sequence feature class that will cover everything, the Biopython `SeqFeature` class attempts to encapsulate as much of the information about the sequence as possible. The design is heavily based on the GenBank/EMBL feature tables, so if you understand how they look, you'll probably have an easier time grasping the structure of the Biopython classes.

10.1.2.1 `SeqFeatures` themselves

The first level of dealing with Sequence features is the `SeqFeature` class itself. This class has a number of attributes, so first we'll list them and their general features, and then work through an example to show how this applies to a real life example, a GenBank feature table. The attributes of a `SeqFeature` are:

location – The location of the `SeqFeature` on the sequence that you are dealing with. The locations end-points may be fuzzy – section [10.1.2.2](#) has a lot more description on how to deal with descriptions.

type – This is a textual description of the type of feature (for instance, this will be something like 'CDS' or 'gene').

ref – A reference to a different sequence. Some times features may be “on” a particular sequence, but may need to refer to a different sequence, and this provides the reference (normally an accession number). A good example of this is a genomic sequence that has most of a coding sequence, but one of the exons is on a different accession. In this case, the feature would need to refer to this different accession for this missing exon.

ref_db – This works along with **ref** to provide a cross sequence reference. If there is a reference, **ref_db** will be set as `None` if the reference is in the same database, and will be set to the name of the database otherwise.

strand – The strand on the sequence that the feature is located on. This may either be '1' for the top strand, '-1' for the bottom strand, or '0' for both strands (or if it doesn't matter). Keep in mind that this only really makes sense for double stranded DNA, and not for proteins or RNA.

qualifiers – This is a python dictionary of additional information about the feature. The key is some kind of terse one-word description of what the information contained in the value is about, and the value is the actual information. For example, a common key for a qualifier might be “evidence” and the value might be “computational (non-experimental).” This is just a way to let the person who is looking at the feature know that it has not been experimentally (i. e. in a wet lab) confirmed.

sub_features – A very important feature of a feature is that it can have additional **sub_features** underneath it. This allows nesting of features, and helps us to deal with things such as the GenBank/EMBL feature lines in a (we hope) intuitive way.

To show an example of SeqFeatures in action, let's take a look at the following feature from a GenBank feature table:

```
mRNA          complement(join(<49223..49300,49780..>50208))
              /gene="F28B23.12"
```

To look at the easiest attributes of the SeqFeature first, if you got a SeqFeature object for this it would have its **type** of 'mRNA', a **strand** of -1 (due to the 'complement'), and would have None for the **ref** and **ref_db** since there are no references to external databases. The **qualifiers** for this SeqFeature would be a python dictionary that looked like {'gene' : 'F28B23.12'}.

Now let's look at the more tricky part, how the 'join' in the location line is handled. First, the location for the top level SeqFeature (the one we are dealing with right now) is set as going from '<49223' to '>50208' (see section 10.1.2.2 for the nitty gritty on how fuzzy locations like this are handled). So the location of the top level object is the entire span of the feature. So, how do you get at the information in the 'join?' Well, that's where the **sub_features** go in.

The **sub_features** attribute will have a list with two SeqFeature objects in it, and these contain the information in the join. Let's look at **top_level_feature.sub_features[0]**; the first **sub_feature**). This object is a SeqFeature object with a **type** of 'mRNA_join,' a **strand** of -1 (inherited from the parent SeqFeature) and a location going from '<49223' to '49300'.

So, the **sub_features** allow you to get at the internal information if you want it (i. e. if you were trying to get only the exons out of a genomic sequence), or just to deal with the broad picture (i. e. you just want to know that the coding sequence for a gene lies in a region). Hopefully this structuring makes it easy and intuitive to get at the sometimes complex information that can be contained in a SeqFeature.

10.1.2.2 Locations

In the section on SeqFeatures above, we skipped over one of the more difficult parts of Features, dealing with the locations. The reason this can be difficult is because of fuzziness of the positions in locations. Before we get into all of this, let's just define the vocabulary we'll use to talk about this. Basically there are two terms we'll use:

position – This refers to a single position on a sequence, which may be fuzzy or not. For instance, 5, 20, <100 and 3~5 are all positions.

location – A location is two positions that defines a region of a sequence. For instance 5..20 (i. e. 5 to 20) is a location.

I just mention this because sometimes I get confused between the two.

The complication in dealing with locations comes in the positions themselves. In biology many times things aren't entirely certain (as much as us wet lab biologists try to make them certain!). For instance,

you might do a dinucleotide priming experiment and discover that the start of mRNA transcript starts at one of two sites. This is very useful information, but the complication comes in how to represent this as a position. To help us deal with this, we have the concept of fuzzy positions. Basically there are five types of fuzzy positions, so we have five classes to deal with them:

ExactPosition – As its name suggests, this class represents a position which is specified as exact along the sequence. This is represented as just a number, and you can get the position by looking at the `position` attribute of the object.

BeforePosition – This class represents a fuzzy position that occurs prior to some specified site. In GenBank/EMBL notation, this is represented as something like '<13', signifying that the real position is located somewhere less than 13. To get the specified upper boundary, look at the `position` attribute of the object.

AfterPosition – Contrary to **BeforePosition**, this class represents a position that occurs after some specified site. This is represented in GenBank as '>13', and like **BeforePosition**, you get the boundary number by looking at the `position` attribute of the object.

WithinPosition – This class models a position which occurs somewhere between two specified nucleotides. In GenBank/EMBL notation, this would be represented as '(1.5)', to represent that the position is somewhere within the range 1 to 5. To get the information in this class you have to look at two attributes. The `position` attribute specifies the lower boundary of the range we are looking at, so in our example case this would be one. The `extension` attribute specifies the range to the higher boundary, so in this case it would be 4. So `object.position` is the lower boundary and `object.position + object.extension` is the upper boundary.

BetweenPosition – This class deals with a position that occurs between two coordinates. For instance, you might have a protein binding site that occurs between two nucleotides on a sequence. This is represented as '2^3', which indicates that the real position happens between position 2 and 3. Getting this information from the object is very similar to **WithinPosition**, the `position` attribute specifies the lower boundary (2, in this case) and the `extension` indicates the range to the higher boundary (1 in this case).

Now that we've got all of the types of fuzzy positions we can have taken care of, we are ready to actually specify a location on a sequence. This is handled by the **FeatureLocation** class. An object of this type basically just holds the potentially fuzzy start and end positions of a feature. You can create a **FeatureLocation** object by creating the positions and passing them in:

```
>>> from Bio import SeqFeature
>>> start_pos = SeqFeature.AfterPosition(5)
>>> end_pos = SeqFeature.BetweenPosition(8, 1)
>>> my_location = SeqFeature.FeatureLocation(start_pos, end_pos)
```

If you print out a **FeatureLocation** object, you can get a nice representation of the information:

```
>>> print my_location
[>5:(8^9)]
```

We can access the fuzzy start and end positions using the start and end attributes of the location:

```
>>> my_location.start
<Bio.SeqFeature.AfterPosition instance at 0x101d7164>
>>> print my_location.start
>5
>>> print my_location.end
(8^9)
```


If you don't want to deal with fuzzy positions and just want numbers, you just need to ask for the `nofuzzy_start` and `nofuzzy_end` attributes of the location:

```
>>> my_location.nofuzzy_start
5
>>> my_location.nofuzzy_end
8
```

Notice that this just gives you back the position attributes of the fuzzy locations.

Similarly, to make it easy to create a position without worrying about fuzzy positions, you can just pass in numbers to the `FeaturePosition` constructors, and you'll get back out `ExactPosition` objects:

```
>>> exact_location = SeqFeature.FeatureLocation(5, 8)
>>> print exact_location
[5:8]
>>> exact_location.start
<Bio.SeqFeature.ExactPosition instance at 0x101dcab4>
```

That is all of the nitty gritty about dealing with fuzzy positions in Biopython. It has been designed so that dealing with fuzziness is not that much more complicated than dealing with exact positions, and hopefully you find that true!

10.1.2.3 References

Another common annotation related to a sequence is a reference to a journal or other published work dealing with the sequence. We have a fairly simple way of representing a Reference in Biopython – we have a `Bio.SeqFeature.Reference` class that stores the relevant information about a reference as attributes of an object.

The attributes include things that you would expect to see in a reference like `journal`, `title` and `authors`. Additionally, it also can hold the `medline_id` and `pubmed_id` and a `comment` about the reference. These are all accessed simply as attributes of the object.

A reference also has a `location` object so that it can specify a particular location on the sequence that the reference refers to. For instance, you might have a journal that is dealing with a particular gene located on a BAC, and want to specify that it only refers to this position exactly. The `location` is a potentially fuzzy location, as described in section [10.1.2.2](#).

That's all there is too it. References are meant to be easy to deal with, and hopefully general enough to cover lots of usage cases.

10.2 Regression Testing Framework

Biopython has a regression testing framework originally written by Andrew Dalke and ported to PyUnit by Brad Chapman which helps us make sure the code is as bug-free as possible before going out.

10.2.1 Writing a Regression Test

Every module that goes into Biopython should have a test (and should also have documentation!). Let's say you've written a new module called `Biospam` – here is what you should do to make a regression test:

1. Write a script called `test_Biospam.py`

- This script should live in the `Tests` directory
- The script should test all of the important functionality of the module (the more you test the better your test is, of course!).

- Try to avoid anything which might be platform specific, such as printing floating point numbers without using an explicit formatting string.
2. If the script requires files to do the testing, these should go in the directory Tests/Biospam.
 3. Write out the test output and verify the output to be correct. There are two ways to do this:
 - (a) The long way:
 - Run the script and write its output to a file. On UNIX machines, you would do something like: `python test_Biospam.py > test_Biospam` which would write the output to the file `test_Biospam`.
 - Manually look at the file `test_Biospam` to make sure the output is correct. When you are sure it is all right and there are no bugs, you need to quickly edit the `test_Biospam` file so that the first line is: `'test_Biospam'` (no quotes).
 - copy the `test_Biospam` file to the directory Tests/output
 - (b) The quick way:
 - Run `python run_tests.py -g test_Biospam.py`. The regression testing framework is nifty enough that it'll put the output in the right place in just the way it likes it.
 - Go to the output (which should be in Tests/output/test_Biospam) and double check the output to make sure it is all correct.
 4. Now change to the Tests directory and run the regression tests with `python run_tests.py`. This will run all of the tests, and you should see your test run (and pass!).
 5. That's it! Now you've got a nice test for your module ready to check into CVS. Congratulations!

10.3 Parser Design

Many of the older Biopython parsers were built around an event-oriented design that includes Scanner and Consumer objects.

Scanners take input from a data source and analyze it line by line, sending off an event whenever it recognizes some information in the data. For example, if the data includes information about an organism name, the scanner may generate an `organism_name` event whenever it encounters a line containing the name.

Consumers are objects that receive the events generated by Scanners. Following the previous example, the consumer receives the `organism_name` event, and the processes it in whatever manner necessary in the current application.

This is a very flexible framework, which is advantageous if you want to be able to parse a file format into more than one representation. For example, both the `Bio.GenBank` and `Bio.SwissProt` modules use this to read their file formats as both generic `SeqRecord` objects and file-format-specific record objects.

More recently, many of the parsers added for `Bio.SeqIO` and `Bio.AlignIO` take a much simpler approach, but only generate a single object representation (`SeqRecord` and `Alignment` objects respectively).

10.4 Substitution Matrices

10.4.1 SubsMat

This module provides a class and a few routines for generating substitution matrices, similar to BLOSUM or PAM matrices, but based on user-provided data.

Additionally, you may select a matrix from `MatrixInfo.py`, a collection of established substitution matrices.

```
class SeqMat(UserDict.UserDict)
```

1. Attributes

- (a) `self.data`: a dictionary in the form of $\{(i_1, j_1):n_1, (i_1, j_2):n_2, \dots, (i_k, j_k):n_k\}$ where i, j are alphabet letters, and n is a value.
- (b) `self.alphabet`: a class as defined in `Bio.Alphabet`
- (c) `self.ab_list`: a list of the alphabet's letters, sorted. Needed mainly for internal purposes
- (d) `self.sum_letters`: a dictionary. $\{i_1: s_1, i_2: s_2, \dots, i_n: s_n\}$ where:
 - i. i : an alphabet letter;
 - ii. s : sum of all values in a half-matrix for that letter;
 - iii. n : number of letters in alphabet.

2. Methods

- (a) `__init__(self, data=None, alphabet=None, mat_type=NOTYPE, mat_name='', build_later=0)`:
 - i. `data`: can be either a dictionary, or another `SeqMat` instance.
 - ii. `alphabet`: a `Bio.Alphabet` instance. If not provided, construct an alphabet from `data`.
 - iii. `mat_type`: type of matrix generated. One of the following:
 - NOTYPE** No type defined
 - ACCREP** Accepted Replacements Matrix
 - OBSFREQ** Observed Frequency Matrix
 - EXPFREQ** Expsected Frequency Matrix
 - SUBS** Substitution Matrix
 - LO** Log Odds Matrix`mat_type` is provided automatically by some of `SubsMat`'s functions.
 - iv. `mat_name`: matrix name, such as "BLOSUM62" or "PAM250"
 - v. `build_later`: default false. If true, user may supply only alphabet and empty dictionary, if intending to build the matrix later. this skips the sanity check of alphabet size vs. matrix size.
- (b) `entropy(self, obs_freq_mat)`
 - i. `obs_freq_mat`: an observed frequency matrix. Returns the matrix's entropy, based on the frequency in `obs_freq_mat`. The matrix instance should be `LO` or `SUBS`.
- (c) `letter_sum(self, letter)`

Returns the sum of all values in the matrix, for the provided `letter`
- (d) `all_letters_sum(self)`

Fills the dictionary attribute `self.sum_letters` with the sum of values for each letter in the matrix's alphabet.
- (e) `print_mat(self, f, format="%4d", bottomformat="%4s", alphabet=None)`

prints the matrix to file handle `f`. `format` is the format field for the matrix values; `bottomformat` is the format field for the bottom row, containing matrix letters. Example output for a 3-letter alphabet matrix:

```
A 23
B 12 34
C 7 22 27
  A  B  C
```

The `alphabet` optional argument is a string of all characters in the alphabet. If supplied, the order of letters along the axes is taken from the string, rather than by alphabetical order.

3. Usage

The following section is layed out in the order by which most people wish to generate a log-odds matrix. Of course, interim matrices can be generated and investigated. Most people just want a log-odds matrix, that's all.

(a) Generating an Accepted Replacement Matrix

Initially, you should generate an accepted replacement matrix (ARM) from your data. The values in ARM are the counted number of replacements according to your data. The data could be a set of pairs or multiple alignments. So for instance if Alanine was replaced by Cysteine 10 times, and Cysteine by Alanine 12 times, the corresponding ARM entries would be:

```
('A','C'): 10, ('C','A'): 12
```

as order doesn't matter, user can already provide only one entry:

```
('A','C'): 22
```

A `SeqMat` instance may be initialized with either a full (first method of counting: 10, 12) or half (the latter method, 22) matrices. A full protein alphabet matrix would be of the size $20 \times 20 = 400$. A half matrix of that alphabet would be $20 \times 20 / 2 + 20 / 2 = 210$. That is because same-letter entries don't change. (The matrix diagonal). Given an alphabet size of N :

- i. Full matrix size: $N \times N$
- ii. Half matrix size: $N(N+1)/2$

The `SeqMat` constructor automatically generates a half-matrix, if a full matrix is passed. If a half matrix is passed, letters in the key should be provided in alphabetical order: ('A','C') and not ('C','A').

At this point, if all you wish to do is generate a log-odds matrix, please go to the section titled Example of Use. The following text describes the nitty-gritty of internal functions, to be used by people who wish to investigate their nucleotide/amino-acid frequency data more thoroughly.

(b) Generating the observed frequency matrix (OFM)

Use:

```
OFM = SubsMat._build_obs_freq_mat(ARM)
```

The OFM is generated from the ARM, only instead of replacement counts, it contains replacement frequencies.

(c) Generating an expected frequency matrix (EFM)

Use:

```
EFM = SubsMat._build_exp_freq_mat(OFM,exp_freq_table)
```

- i. `exp_freq_table`: should be a `FreqTable` instance. See section [10.4.2](#) for detailed information on `FreqTable`. Briefly, the expected frequency table has the frequencies of appearance for each member of the alphabet. It is implemented as a dictionary with the alphabet letters as keys, and each letter's frequency as a value. Values sum to 1.

The expected frequency table can (and generally should) be generated from the observed frequency matrix. So in most cases you will generate `exp_freq_table` using:

```
>>> exp_freq_table = SubsMat._exp_freq_table_from_obs_freq(OFM)
>>> EFM = SubsMat._build_exp_freq_mat(OFM,exp_freq_table)
```

But you can supply your own `exp_freq_table`, if you wish

- (d) Generating a substitution frequency matrix (SFM)

Use:

```
SFM = SubsMat._build_subs_mat(OFM,EFM)
```

Accepts an OFM, EFM. Provides the division product of the corresponding values.

- (e) Generating a log-odds matrix (LOM)

Use:

```
LOM=SubsMat._build_log_odds_mat(SFM[,logbase=10,factor=10.0,round_digit=1])
```

- i. Accepts an SFM.
- ii. `logbase`: base of the logarithm used to generate the log-odds values.
- iii. `factor`: factor used to multiply the log-odds values. Each entry is generated by $\log(\text{LOM}[\text{key}]) \times \text{factor}$ And rounded to the `round_digit` place after the decimal point, if required.

4. Example of use

As most people would want to generate a log-odds matrix, with minimum hassle, SubsMat provides one function which does it all:

```
make_log_odds_matrix(acc_rep_mat,exp_freq_table=None,logbase=10,  
                    factor=10.0,round_digit=0):
```

- (a) `acc_rep_mat`: user provided accepted replacements matrix
- (b) `exp_freq_table`: expected frequencies table. Used if provided, if not, generated from the `acc_rep_mat`.
- (c) `logbase`: base of logarithm for the log-odds matrix. Default base 10.
- (d) `round_digit`: number after decimal digit to which result should be rounded. Default zero.

10.4.2 FreqTable

`FreqTable.FreqTable(UserDict.UserDict)`

1. Attributes:

- (a) `alphabet`: A `Bio.Alphabet` instance.
- (b) `data`: frequency dictionary
- (c) `count`: count dictionary (in case counts are provided).

2. Functions:

- (a) `read_count(f)`: read a count file from stream `f`. Then convert to frequencies
- (b) `read_freq(f)`: read a frequency data file from stream `f`. Of course, we then don't have the counts, but it is usually the letter frequencies which are interesting.

3. Example of use: The expected count of the residues in the database is sitting in a file, whitespace delimited, in the following format (example given for a 3-letter alphabet):

```
A 35  
B 65  
C 100
```

And will be read using the `FreqTable.read_count(file_handle)` function.

An equivalent frequency file:

```
A 0.175
B 0.325
C 0.5
```

Conversely, the residue frequencies or counts can be passed as a dictionary. Example of a count dictionary (3-letter alphabet):

```
{ 'A': 35, 'B': 65, 'C': 100 }
```

Which means that an expected data count would give a 0.5 frequency for 'C', a 0.325 probability of 'B' and a 0.175 probability of 'A' out of 200 total, sum of A, B and C)

A frequency dictionary for the same data would be:

```
{ 'A': 0.175, 'B': 0.325, 'C': 0.5 }
```

Summing up to 1.

When passing a dictionary as an argument, you should indicate whether it is a count or a frequency dictionary. Therefore the `FreqTable` class constructor requires two arguments: the dictionary itself, and `FreqTable.COUNT` or `FreqTable.FREQ` indicating counts or frequencies, respectively.

Read expected counts. `readCount` will already generate the frequencies Any one of the following may be done to generate the frequency table (ftab):

```
>>> from SubsMat import *
>>> ftab = FreqTable.FreqTable(my_frequency_dictionary, FreqTable.FREQ)
>>> ftab = FreqTable.FreqTable(my_count_dictionary, FreqTable.COUNT)
>>> ftab = FreqTable.read_count(open('myCountFile'))
>>> ftab = FreqTable.read_frequency(open('myFrequencyFile'))
```

Chapter 11

Where to go from here – contributing to Biopython

11.1 Maintaining a distribution for a platform

We try to release Biopython to make it as easy to install as possible for users. Thus, we try to provide the Biopython libraries in as many install formats as we can. Doing this from release to release can be a lot of work for developers, and sometimes requires them to maintain packages they are not all that familiar with. This section is meant to provide tips to encourage other people besides developers to maintain platform builds.

In general, this is fairly easy – all you would need to do is produce the system specific package whenever we make a release. You should then check the package (of course!) to make sure it installs everything properly. Then you just send it to one of the main developers, they stick the package on the web site and just like that, you’ve contributed to Biopython! Snazzy.

Below are some tips for certain platforms to maybe get people started with helping out:

RPMs – RPMs are pretty popular package systems on some platforms. There is lots of documentation on RPMs available at <http://www.rpm.org> to help you get started with them. To create an RPM for your platform is really easy. You just need to be able to build the package from source (having a C compiler that works is thus essential) – see the Biopython installation instructions for more info on this.

To make the RPM, you just need to do:

```
python setup.py bdist_rpm
```

This will create an RPM for your specific platform and a source RPM in the directory `dist`. This RPM should be good and ready to go, so this is all you need to do! Nice and easy.

Windows – Windows products typically have a nice graphical installer that installs all of the essential components in the right place. We can use Distutils to create a installer of this type fairly easily.

You must first make sure you have a C compiler on your Windows computer, and that you can compile and install things (see the Biopython installation instructions for info on how to do this).

Once you are setup with a C compiler, making the installer just requires doing:

```
python setup.py bdist_wininst
```

Now you’ve got a Windows installer. Congrats!

Macintosh – We would love to find someone who wants to maintain a Macintosh distribution, and make it available in a Macintosh friendly format like bin-hex. This would basically include finding a way to compile everything on the Mac, making sure all of the code written by us UNIX-based developers works well on the Mac, and providing any Mac-friendly hints for us.

Once you've got a package, please test it on your system to make sure it installs everything in a good way and seems to work properly. Once you feel good about it, send it off to one of the biopython developers (write to our main list serve at biopython@biopython.org if you're not sure who to send it to) and you've done it. Thanks!

11.2 Bug Reports + Feature Requests

Getting feedback on the Biopython modules is very important to us. Open-source projects like this benefit greatly from feedback, bug-reports (and patches!) from a wide variety of contributors.

The main forums for discussing feature requests and potential bugs are the biopython development lists:

- biopython@biopython.org – An unmoderated list for discussion of anything to do with biopython.
- biopython-dev@biopython.org – A more development oriented list that is mainly used by developers (but anyone is free to contribute!).

Additionally, if you think you've found a bug, you can submit it to our bug-tracking page at <http://bugzilla.open-bio.org/>. This way, it won't get buried in anyone's Inbox and forgotten about.

11.3 Contributing Code

There are no barriers to joining biopython code development other than an interest in creating biology-related code in python. The best place to express an interest is on the biopython mailing lists – just let us know you are interested in coding and what kind of stuff you want to work on. Normally, we try to have some discussion on modules before coding them, since that helps generate good ideas – then just feel free to jump right in and start coding!

The main biopython release tries to be fairly uniform and interworkable, to make it easier for users. You can read about some of (fairly informal) coding style guidelines we try to use in biopython in the contributing documentation at <http://biopython.org/wiki/Contributing>. We also try to add code to the distribution along with tests (see section 10.2 for more info on the regression testing framework) and documentation, so that everything can stay as workable and well documented as possible. This is, of course, the most ideal situation, under many situations you'll be able to find other people on the list who will be willing to help add documentation or more tests for your code once you make it available. So, to end this paragraph like the last, feel free to start working!

Additionally, if you have code that you don't think fits in the distribution, but that you want to make available, we maintain Script Central (<http://biopython.org/wiki/Scriptcentral>) which has pointers to freely available code in python for bioinformatics.

Hopefully this documentation has got you excited enough about biopython to try it out (and most importantly, contribute!). Thanks for reading all the way through!

Chapter 12

Appendix: Useful stuff about Python

If you haven't spent a lot of time programming in python, many questions and problems that come up in using Biopython are often related to python itself. This section tries to present some ideas and code that come up often (at least for us!) while using the Biopython libraries. If you have any suggestions for useful pointers that could go here, please contribute!

12.1 What the heck is a handle?

Handles are mentioned quite frequently throughout this documentation, and are also fairly confusing (at least to me!). Basically, you can think of a handle as being a “wrapper” around text information.

Handles provide (at least) two benefits over plain text information:

1. They provide a standard way to deal with information stored in different ways. The text information can be in a file, or in a string stored in memory, or the output from a command line program, or at some remote website, but the handle provides a common way of dealing with information in all of these formats.
2. They allow text information to be read incrementally, instead of all at once. This is really important when you are dealing with huge text files which would use up all of your memory if you had to load them all.

Handles can deal with text information that is being read (e. g. reading from a file) or written (e. g. writing information to a file). In the case of a “read” handle, commonly used functions are `read()`, which reads the entire text information from the handle, and `readline()`, which reads information one line at a time. For “write” handles, the function `write()` is regularly used.

The most common usage for handles is reading information from a file, which is done using the built-in python function `open`. Here, we open a handle to the file `m_cold.fasta` (also available online [here](#)):

```
>>> handle = open("m_cold.fasta", "r")
>>> handle.readline()
">gi|8332116|gb|BE037100.1|BE037100 MP14H09 MP Mesembryanthemum ... \n"
```

Handles are regularly used in Biopython for passing information to parsers.

12.1.1 Creating a handle from a string

One useful thing is to be able to turn information contained in a string into a handle. The following example shows how to do this using `cStringIO` from the Python standard library:

```
>>> my_info = 'A string\n with multiple lines.'
>>> print my_info
A string
  with multiple lines.
>>> import cStringIO
>>> my_info_handle = cStringIO.StringIO(my_info)
>>> first_line = my_info_handle.readline()
>>> print first_line
A string

>>> second_line = my_info_handle.readline()
>>> print second_line
  with multiple lines.
```