BIOINFORMATICS

Vol. 00 no. 00 2011
Pages 1–2

A Ruby API to query the Ensembl database for genomic features

Francesco Strozzi [1] and Jan Aerts [2*]

[1]Parco Tecnologico Padano, Via Einstein Loc. Cascina Codazza 26900 Lodi, Italy
[2]Faculty of Engineering - ESAT/SCD, Leuven University, Kasteelpark Arenberg 10 - bus 2446, 3001 Leuven, Belgium

Received on XXXXX; revised on XXXXX; accepted on XXXXX

Associate Editor: XXXXXXX

ABSTRACT

Summary The Ensembl database makes genomic features available via its Genome Browser. It is also possible to access the underlying data through a Perl API for advanced querying. We have developed a full-featured Ruby API to the Ensembl databases, providing the same functionality as the Perl interface. A single Ruby API is used to access different releases of the Ensembl databases and is also able to query multi-species databases.

Availability and Implementation Most functionality of the API is provided using the ActiveRecord pattern. The library depends on introspection to make it release-independent. The API is available through the Rubygem system and can be installed with the command `gem install ruby-ensembl-api`.

Contact jan.aerts@esat.kuleuven.be

1 INTRODUCTION

The Ensembl (Flicek *et al*., 2010) and UCSC (Fujita *et al*., 2010) genome browsers are the first point of call for a large community of genetics and genomics researchers. Both provide a graphical interface for browsing the genomes of a large number of species, displaying the location of genes, polymorphisms, repeats and regulatory regions. Each database can also be accessed directly via SQL and provides an interface for simple querying of the data: BioMart for Ensembl (Haider *et al*., 2009) and the Table Browser for UCSC. In addition, the Ensembl team provides a Perl API for advanced scripted access to the data (Flicek *et al*., 2010).

In recent years, the Python and Ruby scripting languages have gained significant ground in the bioinformatics community (see *e.g.* Goto *et al*., 2010, Cock *et al*., 2009, Aerts and Law, 2009), increasing the need for a programmable interface in these languages. In this paper, we describe a second API to the Ensembl database, focusing on the Ruby programming community.

2 IMPLEMENTATION

The data available in the Ensembl Genome Browser is stored in a set of MySQL relational databases and to a certain extent normalized. Every table covers one specific conceptual class of objects, such as ”genes” or ”transcripts”. In the Ruby API, the tables of the Core and Variation databases are linked to Ruby classes using the ActiveRecord pattern (http://en.wikipedia.org/wiki/Active_record_pattern); tuples within a table correspond to objects of that class.

As does the perl API, the Ruby library provides a `Slice` class which describes a region of the genome. Each class that defines a locus (*e.g.* gene, simple_feature, assembly_exception) includes the Sliceable mixin which provides additional methods such as returning the sequence of the object or its position. The Ruby and Perl Slice classes differ however significantly at the conceptual level. The slice for a feature (*e.g.* a gene) in the perl API is the complete seq_region this feature is located on, commonly the whole chromosome: *e.g.* ”chromosome 13” for the gene BRCA2. In contrast, the slice in the Ruby version is delimited by the boundaries of the feature: ”chromosome:GRCh37:13:32889611:32973347:1” for the same gene. As a result, methods such as `overlaps?`, `contains?` and `within?` are available for each object that implements Sliceable.

3 FEATURES

The user provides the species name in snake_case (*e.g.* ”homo_sapiens”) and an optional release number to connect to the Ensembl database. It is not necessary to make the distinction between Core or Variation; the code will internally open connections to either if necessary. In addition, the Ruby API is independent of the different Ensembl releases. Whereas the user needs the correct version of the Perl API to work with a given Ensembl database release, there is only a single Ruby interface that works for every release. Finally, the Ruby API is able to work with Ensembl Genomes databases, where multiple species are stored within the same database (*e.g.* bacterial, fungal and plant genomes; Kersey *et al*., 2010).

The Ruby Ensembl API provides - to our knowledge - the same functionality as the Perl API where concerning the Core and Variation databases. Class methods cover searching for records: every column in the table is available for querying by preceding it with `find_by_` (*e.g.* `my_gene = Gene.find_by_name("BRCA2")`). These class methods bypass the need for Adaptor objects as in the Perl API. Instance methods, on the other hand, work on records and for example provide access to specific data for a single column in a given row (*e.g.*

*to whom correspondence should be addressed

© Oxford University Press 2011.

```
require 'ensembl'
include Ensembl::Core
DBConnection.connect('homo_sapiens', 60)

my_gene = Gene.find_by_name('BRCA2')
puts my_gene.name + "\t" + my_gene.slice.to_s
my_variations = my_gene.slice.get_variation_features
puts my_variations.length.to_s + "\t" + my_variations[0].variation_name

my_locus = Slice.fetch_by_region('chromosome','13',32889000,33000000,1)
puts my_gene.slice.within?(my_locus)

slice = Slice.fetch_by_region('chromosome', '4', 10000, 100000)
slice.genes.each do |gene|
  puts gene.name
  gene.transcripts.each do |transcript|
    puts transcript.stable_id + "\t" + transcript.exons.length.to_s
  end
end
```

Figure 1: Example script using the Ruby Ensembl API

`my_gene.start`). As each method call returns an object, different method calls can be chained together. To obtain the start position of the first transcript for a gene `my_gene`, the user can invoke `my_gene.transcripts[0].start`. Functionality that is not automatically covered by the ActiveRecord pattern but that is present in the Perl API is also provided. This includes but is not limited to converting genomic positions between different coordinate systems (*e.g.* between chromosome, scaffold and contig) and SNP effect prediction. In addition, the user can ask the API what types of object are related to for example a gene with `Gene.reflect_on_all_associations(:has_many)`, thus providing additional real-time documentation for every class in the API.

The library provides two binaries. The `ensembl` command takes a species and release number as argument and drops the user in an interactive Ruby session for quick querying of the Ensembl database without the need to write one-off scripts. In addition, the `variation_effect_predictor` script takes a file containing known or novel SNPs/indels and annotates this list with a consequence type (*e.g.* essential_splice_site, stop_gained).

Future efforts will focus on extending the API to the Compara and FunctionalGenomics databases which provide data for multi-species comparisons and for functional as well as regulatory information.

3.1 Example

Figure 1 shows an example of using the Ruby Ensembl API. Lines 1 to 3 load the library and connect to the Core and Variation databases for human release 60. In lines 4 and 5, the BRCA2 gene is retrieved and gene name and location are printed. The `#slice` method creates a Slice object, which describes the locus itself and is serialized into a string. Lines 6 and 7 retrieve all variations within the gene and report on the total number of variations and the dbSNP accession of the first one. In lines 8 and 9, a locus is defined and the code checks if the BRCA2 gene is within that locus. Finally, in lines 10 to 16, a slice of the genome is selected directly. For every gene in this slice, the gene name in printed as well as the stable ID and number of exons for each transcript.

CONCLUSION

The application programming interface described here provides the functionality needed to query the Ensembl database using the Ruby programming language. This API significantly improves the way a user can interact with complex genomic data, providing powerful yet intuitive commands and methods. Although the user can write scripts to interact with the database, the readibility and terseness of Ruby code make interactive sessions (using the `ensembl` executable) in many cases a valid option as well. In combination with the popular Rails framework for building database-backed websites, the Ruby Ensembl API is also ideal for *e.g.* adding background information on candidate genes from the Ensembl database in applications geared at clinical geneticists.

From a library maintainer's perspective, the metaprogramming and introspection capabilities of the Ruby language and the ActiveRecord module allow for providing full functionality and easy maintenance with minimal effort. They enable the Ruby API to be very flexible and by definition virtually insensitive to adding or removing data columns in tables.

The API is available through the Rubygem system and can be installed with the command `gem install ruby-ensembl-api`. Source code is at http://github.com/jandot/ruby-ensembl-api. An extensive tutorial (with permission copied and modified from the Perl API) is also available at the GitHub website.

AUTHOR CONTRIBUTIONS

JA initiated the project and wrote the Core code and overall framework. FS created the Variation API and adapted the code for Ensembl Genomes multi-species databases. Both contributed to the manuscript.

ACKNOWLEDGMENTS

We thank the European Bioinformatics Institute for hosting JA under the Geek For A Week program, and specifically Glenn Proctor and Andreas Kahari. We thank Marc Hoeppner for useful discussions and Alejandro Sifrim for his contribution to the variation consequence calculation. Article processing charges are covered by SymBioSys II [grant number KUL PFV/10/016 SymBioSys].

REFERENCES

Aerts J and Law A (2009) An introduction to scripting in Ruby for biologists. *BMC Bioinf* **10**:221

Cock PJ, Antao T, Chang JT, Chapman BA, Cox CJ, Dalke A, Friedberg I, Hamelryck T, Kauff F, Wilczynski B and de Hoon MJ (2009) Biopython: freely available Python tools for computational molecular biology and bioinformatics. *Bioinformatics* **25**(11):1422-1423.

Flicek P, Amode MR, Barrell D, Beal K, Brent S, Chen Y, Clapham P, Coates G, Fairley S, Fitzgerald S *et al* (2010) Ensembl 2011. *Nucleic Acids Res*, Epub ahead of print

Fujita PA, Rhead B, Zweig AS, Hinrichs AS, Karolchik D, Cline MS, Goldman M, Barber GP, Clawson H, Coelho A *et al (2010) The UCSC Genome Browser database: update 2011. Nucleic Acids Res, Epub ahead of print.*

Goto N, Prins P, Nakao M, Bonnal R, Aerts J and Katayama T (2010) BioRuby: Bioinformatics software for the Ruby programming language. Bioinformatics **26***(20):2617-2619*

Haider S, Ballester B, Smedley D, Zhang J, Rice P and Kasprzyk A. BioMart Central Portal - unified access to biological data. Nucleic Acids Res **37***(2), W23-W27*

Kersey PJ, Lawson D, Birney E, Derwent PS, Haimel M, Herrero J, Keenan S, Kerhornou A, Koscielny G, Kahari A et al. (2010) Ensembl Genomes: extending Ensembl across the taxonomic space. Nucleic Acids Res **38***:D563-D569.*