

Structural bioinformatics

BioStructures.jl: read, write and manipulate macromolecular structures in Julia

Joe G. Greener^{1,*}, Joel Selvaraj² and Ben J. Ward³

¹Department of Computer Science, University College London, London WC1E 6BT, UK, ²School of Computer Science and Engineering, Vellore Institute of Technology, Vellore 632014, Tamil Nadu, India and ³The Earlham Institute, Norwich Research Park, Norwich NR4 7UZ, UK

*To whom correspondence should be addressed.

Associate Editor: Arne Elofsson

Received on March 5, 2020; revised on April 20, 2020; editorial decision on May 6, 2020; accepted on May 7, 2020

Abstract

Summary: Robust, flexible and fast software to read, write and manipulate macromolecular structures is a prerequisite for productively doing structural bioinformatics. We present BioStructures.jl, the first dedicated package in the Julia programming language for dealing with macromolecular structures and the Protein Data Bank. BioStructures.jl builds on the lessons learned with similar packages to provide a large feature set, a flexible object representation and high performance.

Availability and implementation: BioStructures.jl is freely available under the MIT license. Source code and documentation are available at <https://github.com/BioJulia/BioStructures.jl>. BioStructures.jl is compatible with Julia versions 0.6 and later and is system-independent.

Contact: j.greener@ucl.ac.uk

1 Introduction

Open source software packages to parse files from the Protein Data Bank (PDB) (Berman *et al.*, 2000) and manipulate macromolecular structures exist in many languages (Grant *et al.*, 2006; Goto *et al.*, 2010; Hamelryck and Manderick, 2003; Lafita *et al.*, 2019; Loriot *et al.*, 2010; Stajich *et al.*, 2002). Such packages must strike a balance between a powerful and useful representation of molecules, fast performance, easy integration with other tools and tolerance to the ambiguities in PDB data.

Julia is a high-performance, dynamically-typed, open source programming language (Bezanson *et al.*, 2017). Since its first release in 2012 it has grown rapidly in popularity, particularly in the scientific computing community, with version 1.0 being released in 2018. To date it has over 13 million downloads and over 3000 packages registered for community use. In particular, the ability to write performant code in a high-level language means that Julia can solve the ‘two-language problem’ of having to prototype code in one language and then write a performant version in another language. BioStructures.jl is a Julia package to read, write and manipulate macromolecular structures. Whilst other Julia packages have provided functionality related to structural bioinformatics (Greener *et al.*, 2017; Zea *et al.*, 2017), BioStructures.jl is the first dedicated package and contains all the main features that structural bioinformaticians need to be productive in Julia. It is designed to be used for standard structural analysis tasks, interactive data analysis and to act as a platform on which others can build to create more specific tools. BioStructures.jl is part of BioJulia, an organization that provides bioinformatics infrastructure for the Julia language.

2 Features

BioStructures.jl has the following features:

- Read in PDB, mmCIF and MMTF (Bradley *et al.*, 2017) files into a hierarchical representation of structure. The parsers have been tested on the whole PDB, only throwing errors on a small number of known ambiguous cases.
- Write out PDB, mmCIF and MMTF files. The ability to read and write freely between these file formats is not available in many similar packages.
- Read mmCIF and MMTF files into a dictionary, e.g. allowing access to header information. MMTF files are decoded with the related package MMTF.jl.
- Iterate over structures at various levels, e.g. iterate over atoms in a residue or residues in a chain.
- Select various structural elements using pre-defined or custom selectors, e.g. collect all C β atoms (Ca in the case of glycine) from standard residues.
- Retrieve amino acid sequences and integrate with the broader BioJulia ecosystem, e.g. allowing fast sequence alignments.
- Spatial calculations including distances, bond angles, dihedral angles, contact maps and distance maps. Contact and distance maps can be plotted.




Fig. 1. Performance of structural bioinformatics tasks in 15 packages (Bakan *et al.*, 2011; Gowers *et al.*, 2016; Grant *et al.*, 2006; Goto *et al.*, 2010; Hamelryck and Manderick, 2003; Hirsh *et al.*, 2015; Ireland and Martin, 2020; Kunzmann and Hamacher, 2018; Lafita *et al.*, 2019; Loriot *et al.*, 2010; Stajich *et al.*, 2002; Zea *et al.*, 2017) covering seven programming languages. Comparison should be treated with caution since each package does something slightly different and may use a different object representation or do less error checking, e.g. MIToS does not read files into a hierarchical representation of structure. The tasks are reading a small (1CRN) and a large (1HTQ) PDB entry (Gajda, 2013) in the PDB, mmCIF and MMTF formats; counting the number of alanine residues in adenylate kinase (1AKE); calculating the distance between residues 50 and 60 of chain A in adenylate kinase; and calculating the Ramachandran ϕ/ψ angles in adenylate kinase. In each case, the mean time of the fastest implementation for each software that makes use of the provided API is given. Tasks are not implemented in packages where there is no obvious API for implementation. Times for Julia packages are measured after JIT compilation. Packages are ordered by increasing time to parse PDB 1CRN, with BioStructures first. See <https://github.com/jgreener64/pdb-benchmarks> for more details. The version of the benchmarks presented here is archived at Zenodo with doi:10.5281/zenodo.3753016

- Superimposition of structures and calculation of the RMSD.
- Download files and data from the RCSB PDB including functions to maintain a local copy of the PDB.
- Interoperability with the broader Julia ecosystem, e.g. exporting to a data frame or creating a graph of contacting residues.
- Visualization of molecular structures in a pop-up window or Jupyter notebook using the related package Bio3DView.jl (<https://github.com/jgreener64/Bio3DView.jl>), which is a wrapper around 3Dmol.js (Rego and Koes, 2015).
- Easy installation with Julia's package manager.
- Comprehensive test suite, continuous integration build testing and a benchmark suite to test for performance regressions.
- Thorough online documentation and in-code docstrings.
- Fully open source with a permissive MIT license.
- Faster than similar packages at most tasks. Our benchmarks, summarized in Figure 1 and described further at <https://github.com/jgreener64/pdb-benchmarks>, indicate that the package has competitive or superior performance to 14 other commonly used packages from both interpreted and compiled languages. For example, parsing the small PDB entry 1CRN takes 0.76 ms/1.9 ms/1.1 ms in the PDB/mmCIF/MMTF formats after just-in-time (JIT) compilation on a standard desktop computer. It does this whilst using a hierarchical structure representation, allowing variation between models in a structure, and accounting for alternative locations at the atom and residue levels (see below). These features take time to execute but increase the utility and flexibility of the package.

3 Design considerations

BioStructures.jl is heavily influenced by the Bio.PDB module of Biopython (Hamelryck and Manderick, 2003), the design of which has proved effective. The structure object has a hierarchical type system of the form ProteinStructure—Model—Chain—AbstractResidue—AbstractAtom. Atoms with alternative locations are stored in a DisorderedAtom container and residues with alternative locations (i.e. point mutations with different residue names) are stored in a DisorderedResidue container. Function calls fall back to the default atom or residue, so alternative locations can be ignored if the user is not interested in them, but building alternative locations into the type system allows correct representation of many more aspects of the PDB. Whilst BioStructures.jl retains the flexibility of Bio.PDB, its implementation in Julia allows it to have superior speed.

Acknowledgement

The authors would like to thank the BioJulia, Julia and Biopython communities for discussion, assistance and support.

Financial Support: none declared.

Conflict of Interest: none declared.

References

- Bakan, A. *et al.* (2011) ProDy: protein dynamics inferred from theory and experiments. *Bioinformatics*, 27, 1575–1577.
- Berman, H.M. *et al.* (2000) The Protein Data Bank. *Nucleic Acids Res.*, 28, 235–242.
- Bezanson, J. *et al.* (2017) Julia: a fresh approach to numerical computing. *SIAM Rev.*, 59, 65–98.
- Bradley, A.R. *et al.* (2017) MMTF—an efficient file format for the transmission, visualization, and analysis of macromolecular structures. *PLoS Comput. Biol.*, 13, e1005575.
- Gajda, M.J. (2013) hPDB—Haskell library for processing atomic biomolecular structures in Protein Data Bank format. *BMC Res. Notes*, 6, 483, <https://bmresnotes.biomedcentral.com/articles/10.1186/1756-0500-6-483>.
- Goto, N. *et al.* (2010) BioRuby: bioinformatics software for the Ruby Programming Language. *Bioinformatics*, 26, 2617–2619.
- Gowers, R. *et al.* (2016) MDAnalysis: a Python Package for the rapid analysis of molecular dynamics simulations. In: Benthall, S. and Rostrup, S. (eds), *Proceedings of the 15th Python in Science Conference*. pp. 102–109, Austin, TX.
- Grant, B.J. *et al.* (2006) Bio3d: an R package for the comparative analysis of protein structures. *Bioinformatics*, 22, 2695–2696.
- Greener, J.G. *et al.* (2017) Predicting protein dynamics and allostery using multi-protein atomic distance constraints. *Structure*, 25, 546–558.
- Hamelryck, T. and Manderick, B. (2003) PDB file Parser and structure class implemented in Python. *Bioinformatics*, 19, 2308–2310.
- Hirsh, L. *et al.* (2015) The Victor C Library for protein representation and advanced manipulation. *Bioinformatics*, 31, 1138–1140.
- Ireland, S.M. and Martin, A.C.R. (2020) Atomium—a Python Structure Parser. *Bioinformatics*, 36, 2750–2754.
- Kunzmann, P. and Hamacher, K. (2018) Biotite: a unifying open source computational biology framework in Python. *BMC Bioinformatics*, 19, 346.
- Lafita, A. *et al.* (2019) BioJava 5: a community driven open-source bioinformatics library. *PLoS Comput. Biol.*, 15, e1006791.
- Loriot, S. *et al.* (2010) ESBTL: efficient PDB Parser and data structure for the structural and geometric analysis of biological macromolecules. *Bioinformatics*, 26, 1127–1128.
- Rego, N. and Koes, D. (2015) 3Dmol.js: molecular visualization with WebGL. *Bioinformatics*, 31, 1322–1324.
- Stajich, J.E. *et al.* (2002) The Bioperl Toolkit: Perl modules for the life sciences. *Genome Res.*, 12, 1611–1618.
- Zea, D.J. *et al.* (2017) MIToS.jl: mutual information tools for protein sequence analysis in the Julia language. *Bioinformatics*, 33, 564–565.