SOFTWARE METAPAPER

An Open Framework for the Reproducible Study of the Iterated Prisoner's Dilemma

Vincent Knight[1], Owen Campbell[2], Marc Harper[2], Karol M. Langner[3], James Campbell[1], Thomas Campbell[4], Alex Carney[1], Martin Chorley[1], Cameron Davidson-Pilon[2], Kristian Glass[2], Nikoleta Glynatsi[1], Tomáš Ehrlich[2], Martin Jones[2], Georgios Koutsovoulos[5], Holly Tibble[2], Jochen Muller[2], Geraint Palmer[1], Piotr Petunov[2], Paul Slavin[6], Timothy Standen[1], Luis Visintini[2] and Karl Molden[2]

[1] Development; Cardiff University, GB

[2] Development; Not affiliated

[3] Development; Google Inc., Mountain View, CA, US

[4] Development; St. Nicholas Catholic High School, Hartford, GB

[5] Development; The University of Edinburgh, GB

[6] Development; The University of Manchester, GB

Corresponding author: Vincent Knight (knightva@cf.ac.uk)

The Axelrod library is an open source Python package that allows for reproducible game theoretic research into the Iterated Prisoner's Dilemma. This area of research began in the 1980s but suffers from a lack of documentation and test code. The goal of the library is to provide such a resource, with facilities for the design of new strategies and interactions between them, as well as conducting tournaments and ecological simulations for populations of strategies.

With a growing collection of 139 strategies, the library is a also a platform for an original tournament that, in itself, is of interest to the game theoretic community.

This paper describes the Iterated Prisoner's Dilemma, the Axelrod library and its development, and insights gained from some novel research.

Keywords: Game Theory; Prisoners Dilemma; Python

(1) Overview

Introduction

Several Iterated Prisoner's Dilemma tournaments have generated much interest; Axelrod's original tournaments [2, 3], two 2004 anniversary tournaments [20], and the Stewart and Plotkin 2012 tournament [42], following the discovery of zero-determinant strategies. Subsequent research has spawned a number of papers (many of which are referenced throughout this paper), but rarely are the results reproducible. Amongst well-known tournaments, in only one case is the full original source code available (Axelrod's second tournament [3], in FORTRAN). In no cases is the available code well-documented, easily modifiable, or released with significant test suites.

To complicate matters further, a new strategy is often studied in isolation with opponents chosen by the creator of that strategy. Often such strategies are not sufficiently described to enable reliable recreation (in the absence of source code), with [40] being a notable counter-example.

In some cases, strategies are revised without updates to their names or published implementations [25, 26]. As such, the results cannot be reliably replicated and therefore have not met the basic scientific criterion of falsifiability.

This paper introduces a software package: the Axelrod-Python library. The Axelrod-Python project has the following stated goals:

- To enable the reproduction of Iterated Prisoner's Dilemma research as easily as possible
- To produce the de-facto tool for any future Iterated Prisoner's Dilemma research
- To provide as simple a means as possible for anyone to define and contribute new and original Iterated Prisoner's Dilemma strategies

The presented library is partly motivated by an ongoing discussion in the academic community about reproducible research [9, 16, 37, 38], and is:

- Open: all code is released under an MIT license;
- Reproducible and well-tested: at the time of writing there is an excellent level of integrated tests with 99.73% coverage (including property based tests: [28])
- Well-documented: all features of the library are documented for ease of use and modification
- Extensive: 139 strategies are included, with infinitely-many available in the case of parametrised strategies
- Extensible: easy to modify to include new strategies and to run new tournaments

Review of the literature

As stated in [6]: "*few works in social science have had the general impact of [Axelrod's study of the evolution of cooperation]*". In 1980, Axelrod wrote two papers: [2, 3] which describe a computer tournament that has been a major influence on subsequent game theoretic work [5, 6, 7, 8, 10, 11, 12, 13, 15, 18, 23, 24, 27, 32, 33, 34, 36, 41, 42]. As described in [6] this work has not only had impact in mathematics but has also led to insights in biology (for example in [41], a real tournament where Blue Jays are the participants is described) and in particular in the study of evolution.

The tournament is based on an iterated game (see [29] or similar for details) where two players repeatedly play the normal form game of (1) in full knowledge of each other's playing history to date. An excellent description of the *one shot* game is given in [13] which is paraphrased below:

Two players must choose between *Cooperate* (*C*) and *Defect* (*D*):

- If both choose *C*, they receive a payoff of *R* (**R**eward);
- If both choose *D*, they receive a payoff of *P* (**P**unishment);
- If one chooses *C* and the other *D*, the defector receives a payoff of *T* (**T**emptation) and the cooperator a payoff of *S* (**S**ucker).

and the following reward matrix results from the Cartesian product of two decision vectors $\langle C, D \rangle$,

$$\begin{pmatrix} R,R & S,T \\ T,S & P,P \end{pmatrix} \quad \text{such that } T > R > P > S \text{ and } 2R > T + S \quad (1)$$

The *game* of (1) is called the Prisoner's Dilemma. Specific numerical values of $(R, S, T, P) = (3, 0, 5, 1)$ are often used in the literature [2, 3], although any satisfying the conditions in 1 will yield similar results. Axelrod's tournaments (and further implementations of these) are sometimes referred to as Iterated Prisoner's Dilemma (IPD) tournaments. An incomplete representative overview of published tournaments is given in **Table 1**.

In [32] a description is given of how incomplete information can be used to enhance cooperation, in a similar approach to the proof of the Folk theorem for repeated games [29]. This aspect of incomplete information is also considered in [6, 24, 33] where "noisy" tournaments randomly flip the choice made by a given strategy. In [34], incomplete information is considered in the sense of a probabilistic termination of each round of the tournament.

As mentioned before, IPD tournaments have been studied in an evolutionary context: [12, 24, 36, 42] consider this in a traditional evolutionary game theory context.

These works investigate particular evolutionary contexts within which cooperation can evolve and persist. This can be in the context of direct interactions between strategies or population dynamics for populations of many players using a variety of strategies, which can lead to very different results. For example, in [24] a machine learning algorithm in a population context outperforms strategies described in [36] and [42] that are claimed to dominate any evolutionary opponent in head-to-head interactions.

Further to these evolutionary ideas, [8, 10] are examples of using machine learning techniques to evolve particular strategies. In [4], Axelrod describes how similar techniques are used to genetically evolve a high performing strategy from a given set of strategies. Note that in his original work, Axelrod only used a base strategy set of 12 strategies for this evolutionary study. This is noteworthy as the library now boasts over 139 strategies that are readily available for a similar analysis.

Implementation and architecture
Description of the Axelrod Python package

The library is written in Python (http://www.python.org/) which is a popular language in the academic community with libraries developed for a variety of uses including:

- Algorithmic Game Theory [30] (http://gambit.sourceforge.net/).
- Astrophysics [1] (http://www.astropy.org/);
- Data manipulation [31] (http://pandas.pydata.org/);
- Machine learning [35] (http://scikit-learn.org/);
- Mathematics [43] (http://www.sagemath.org/);
- Visualisation [17] (http://matplotlib.org/);

Furthermore, in [18] Python is described as an appropriate language for the reproduction of Iterated Prisoner's Dilemma tournaments due to its object oriented nature and readability.

The library itself is available at https://github.com/Axelrod-Python/Axelrod.

This is a hosted git repository. Git is a version control system which is one of the recommended aspects of reproducible research [9, 38].

As stated in the **Introduction**, one of the main goals of the library is to allow for the easy contribution of strategies. Doing this requires the writing of a simple Python class (which can inherit from other predefined classes). All components of the library are automatically tested using a combination of unit, property and integration tests. These tests are run as new features are added to the library to ensure compatibility (they are also run automatically using travis-ci.org). When submitting a strategy, a simple test is required which ensures the strategy

behaves as expected. Full contribution guidelines can be found in the documentation, which is also part of the library itself and is hosted using `readthedocs.org`. As an example, **Figures 1** and **2** show the source code for the Grudger strategy as well as its corresponding test.

You can see an overview of the structure of the source code in **Figure 3**. This shows the parallel collection of strategies and their tests. Furthermore the underlying engine for the library is a class for tournaments which lives in the `tournament.py` module. This class is responsible for coordinating the play of generated matches (from the `match.py` module). This generation of matches is the responsibility of a *match generator* class (in the `match_generator.py` module) which is designed in such a way as to be easily modifiable to create new types of tournaments. This is described further in a tutorial in the documentation which shows how to easily create a tournament where players only play each other with probability 0.5. This will be discussed further in the reuse section of this paper.

To date the library has had contributions from 26 contributors from a variety of backgrounds which are not solely academic. These contributions have been mostly in terms of strategies. One strategy is the creation of an undergraduate mathematics student with little prior knowledge of programming. Multiple other strategies were written by a 15 year old secondary school student. Both of these students are authors of this paper. As well as these strategy contributions, vital architectural improvements to the library itself have also been contributed.

(2) Availability
Operating system
The Axelrod library runs on all major operating systems: Linux, Mac OS X and Windows.

Year	Reference	Number of Strategies	Type	Source Code
1979	[2]	13	Standard	Not immediately available
1979	[3]	64	Standard	Available in FORTRAN
1991	[6]	13	Noisy	Not immediately available
2002	[41]	16	Wildlife	Not a computer based tournament
2005	[20]	223	Varied	Not available
2012	[42]	13	Standard	Not fully available

Table 1: An overview of a selection of published tournaments. Not all tournaments were 'standard' round robins; for more details see the indicated references.

```python
class Grudger(Player):
    """A player starts by cooperating however will defect if
        at any point the opponent has defected."""

    name = 'Grudger'
    classifier = {
        'memory_depth': float('inf'),  # Long memory
        'stochastic': False,
        'inspects_source': False,
        'manipulates_source': False,
        'manipulates_state': False
    }

    def strategy(self, opponent):
        """Begins by playing C, then plays D for the remaining
            rounds if the opponent ever plays D."""
        if opponent.defections:
            return D
        return C
```

Figure 1: Source code for the Grudger strategy.

```python
class TestGrudger(TestPlayer):

    name = "Grudger"
    player = axelrod.Grudger
    expected_classifier = {
        'memory_depth': float('inf'),  # Long memory
        'stochastic': False,
        'inspects_source': False,
        'manipulates_source': False,
        'manipulates_state': False
    }

    def test_initial_strategy(self):
        """
        Starts by cooperating
        """
        self.first_play_test(C)

    def test_strategy(self):
        """
        If opponent defects at any point then the player will defect forever
        """
        self.responses_test([C, D, D, D], [C, C, C, C], [C])
        self.responses_test([C, C, D, D, D], [C, D, C, C, C], [D])
```

Figure 2: Test code for the Grudger strategy.
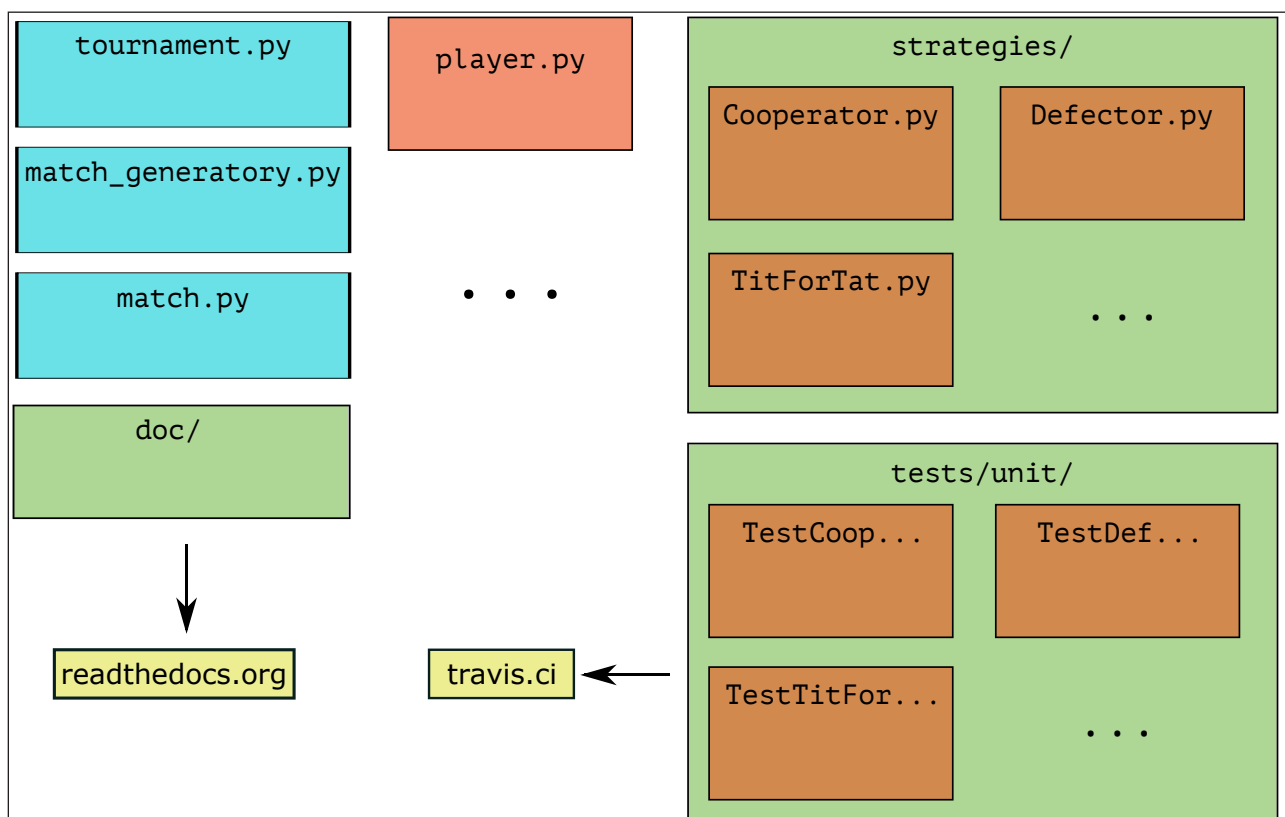
Figure 3: An overview of the source code.

Programming language
The library is continuously tested for compatibility with Python 2.7 and the two most recent python 3 releases.

Additional system requirements
There are no specific additional system requirements.

Support
Support is readily available in multiple forms:
- An online chat channel: `https://gitter.im/Axelrod-Python/Axelrod`.
- An email group: `https://groups.google.com/forum/#!topic/axelrod-python`.

Dependencies
The following Python libraries are required dependencies:

- Numpy 1.9.2
- Matplotlib 1.4.2 (only a requirement if graphical output is required)
- Tqdm 3.4.0
- Hypothesis 3.0 (only a requirement for development)

List of contributors
The names of all the contributors are not known: as these were mainly done through Github and some have not provided their name or responded to a request for further details. Here is an incomplete list:

- Owen Campbell
- Marc Harper
- Vincent Knight
- Karol M. Langner
- James Campbell
- Thomas Campbell
- Martin Jones
- Georgios Koutsovoulos
- Holly Tibble
- Jochen Müller
- Geraint Palmer
- Paul Slavin
- Alex Carney
- Martin Chorley
- Cameron Davidson-Pilon
- Kristian Glass
- Nikoleta Glynatsi
- Tomáš Ehrlich
- Timothy Standen
- Luis Visintini
- Karl Molden
- Jason Young
- Andy Boot
- Anna Barriscale

Software location
Archive
Name: Zenodo
Persistent identifier: 10.5281/zenodo.55509
Licence: MIT
Publisher: Vincent Knight
Version published: Axelrod: 1.2.0
Date published: 2016-06-13

Code repository
Name: Github
Identifier: `https://github.com/Axelrod-Python/Axelrod`
Licence: MIT
Date published: 2015-02-16

Reuse potential
The Axelrod library has been designed with sustainable software practices in mind. There is an extensive documentation suite: `axelrod.readthedocs.org/en/latest/`. Furthermore, there is a growing set of example Jupyter notebooks available here: `https://github.com/Axelrod-Python/Axelrod-notebooks`.

The availability of a large number of strategies makes this tool an excellent and obvious example of the benefits of open research which should positively impact the game theory community. This is evidently true already as the library has been used to study and create interesting and powerful new strategies.

Installation of the library is straightforward via standard python installation repositories (`https://pypi.python.org/pypi`). The package name is `axelrod` and can thus be installed by calling: `pip install axelrod` on all major operating systems (Windows, OS X and Linux).

Figure 4 shows a very simple example of using the library to create a basic tournament giving the graphical output shown in **Figure 5**.

New strategies, tournaments and implications
Due to the open nature of the library the number of strategies included has grown at a fast pace, as can be seen in **Figure 6**.

Nevertheless, due to previous research being done in an irreproducible manner with, for example, no source code and/or vaguely described strategies, not all previous tournaments can yet be reproduced. In fact, some of the early tournaments might be impossible to reproduce as the source code is apparently forever lost. This library aims to ensure reproducibility in the future.

One tournament that is possible to reproduce is that of [42]. The strategies used in that tournament are the following:

1. Cooperator
2. Defector
3. ZD-Extort-2
4. Joss: 0.9
5. Hard Tit For Tat
6. Hard Tit For 2 Tats
7. Tit For Tat
8. Grudger
9. Tit For 2 Tats
10. Win-Stay Lose-Shift
11. Random: 0.5
12. ZD-GTFT-2
13. GTFT: 0.33
14. Hard Prober
15. Prober
16. Prober 2
17. Prober 3
18. Calculator
19. Hard Go By Majority

This can be reproduced as shown in **Figure 8**, which gives the plot of **Figure 7**. Note that slight differences with the results of [42] are due to stochastic behaviour of some strategies.

In parallel to the Python library, a tournament is being kept up to date that pits all available strategies against each other. **Figure 9** shows the results from the full tournament which can also be seen (in full detail) here: `http://axelrod-tournament.readthedocs.org/`. Data sets are also available showing the plays of every match that takes place. Note that to recreate this

```
>>> import axelrod
>>> strategies = [s() for s in axelrod.demo_strategies]
>>> tournament = axelrod.Tournament(strategies)
>>> results = tournament.play()
>>> plot = axelrod.Plot(results)
>>> p = plot.boxplot()
>>> p.show()
```

Figure 4: A simple set of commands to create a demonstration tournament. The output is shown in Figure 5.
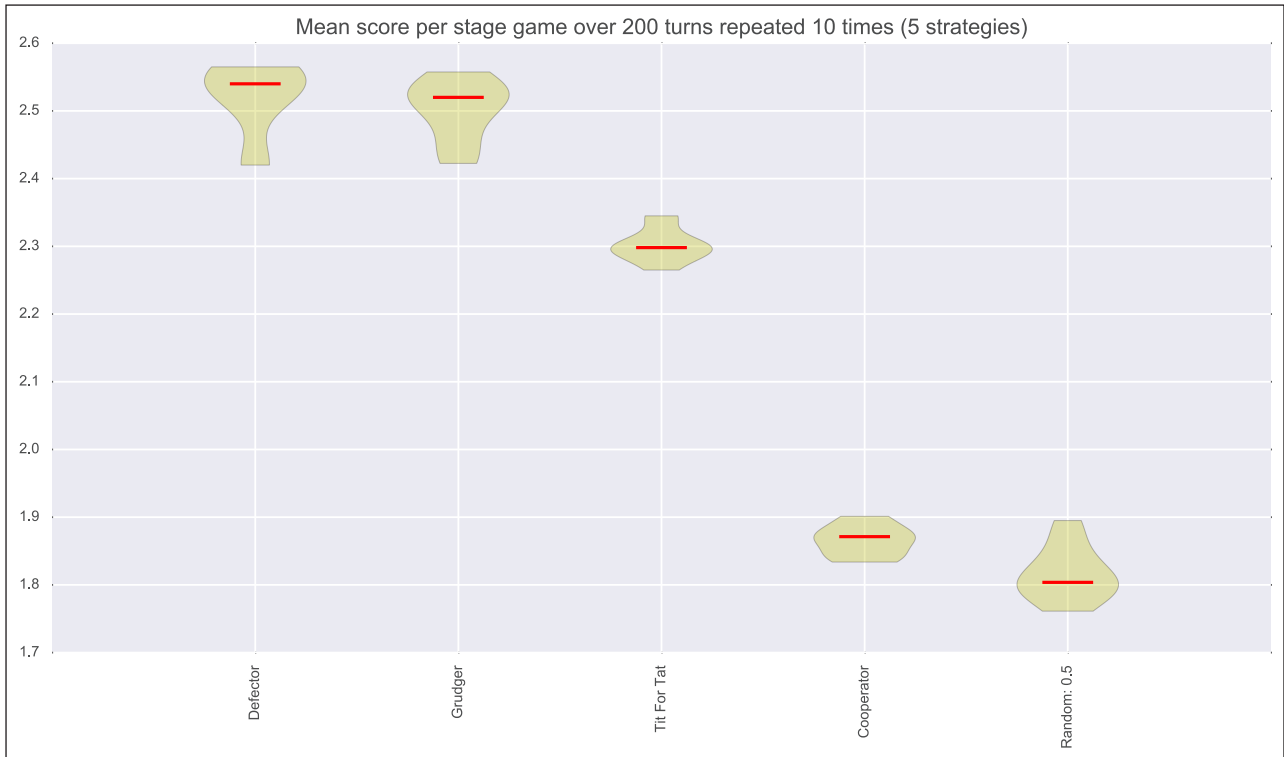
Figure 5: The results from a simple tournament.

tournament simply requires changing a single line of the code shown in **Figure 4**, changing:

```
>>> strategies = [s() for s in axelrod.
demo_strategies]}
```

to:

```
>>> strategies = [s() for s in axelrod.
ordinary_strategies]}.
```

The current winning strategy is new to the research literature: Looker Up. This is a strategy that maps a given set of states to actions. The state space is defined generically by m, n so as to map states to actions as shown in (2).

$$\underbrace{((C,D,D,D,C,D,D,C)}_{m \text{ first actions by opponent}}, \overbrace{((C,C),(C,C))}^{n \text{ last pairs of actions}}) \rightarrow D \qquad (2)$$

The example of (2) is an incomplete illustration of the mapping for $m = 8$, $n = 2$. Intuitively, this state space uses the initial plays of the opponent to gain some information about its intentions whilst still taking into account the recent play. The actual winning strategy is an instance of the framework for $m = n = 2$ for which a particle swarm algorithm has been used to train it. The second placed strategy was trained with an evolutionary algorithm [19]. In [21] experiments are described that evaluate how the second placed strategy behaves in environments other than those in which it was trained and it continues to perform strongly.

There are various other insights that have been gained from ongoing open research on the library, details can be found in [14]. These include:

· A closer look at zero determinant strategies, showing that extortionate strategies obtain a large number of wins: the number of times they outscore an opponent during a given match. *However* these do not perform particularly well from the overall tournament ranking point of view. This is relevant given the findings of [42] in which zero determinant strategies are shown to be able to perform better than any other strategy. This finding extends to noisy tournaments (which are also implemented in the library).
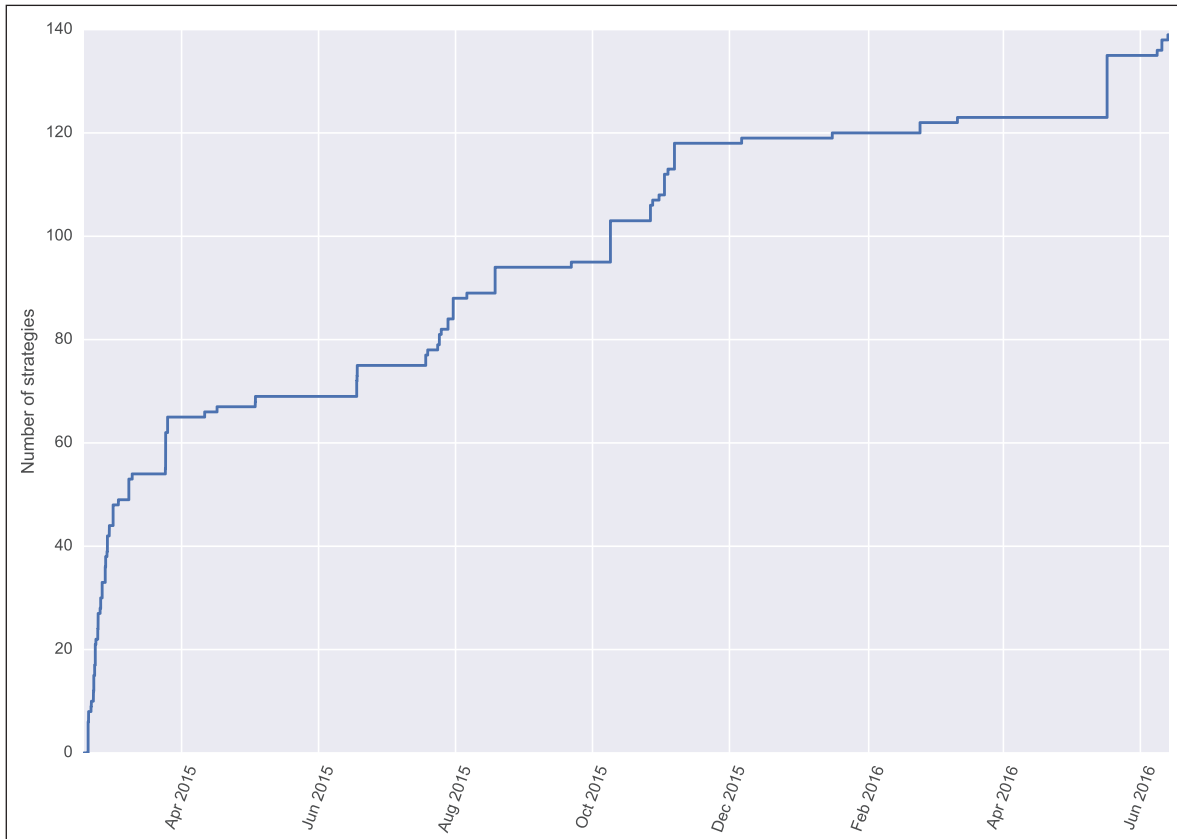
Figure 6: The number of strategies included in the library.
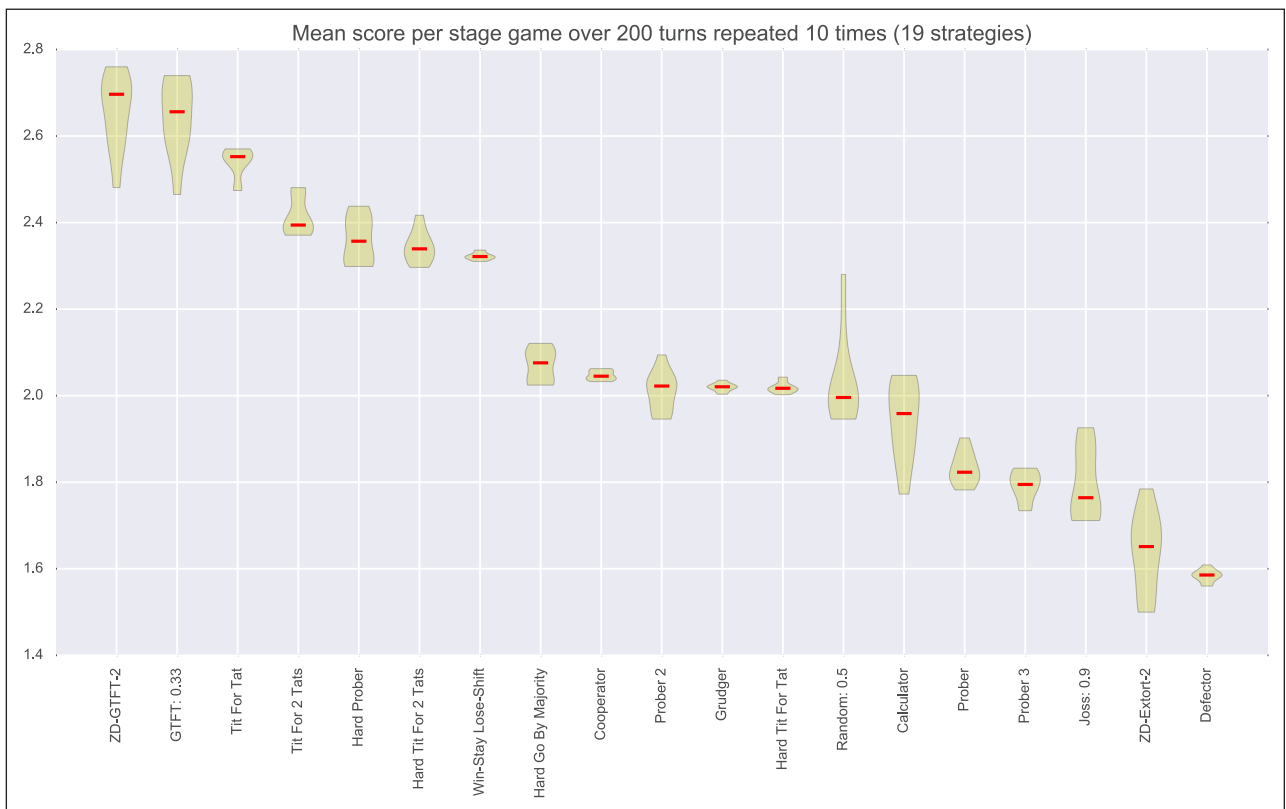
Figure 7: The results from [42], reproduced with the Axelrod library.

· This negative relationship between wins and performance does not generalise. There are some strategies that perform well, both in terms of matches won and overall performance: Back stabber, Double crosser, Looker Up, and Fool Me Once. These strategies continue to perform well in noisy tournaments, however

```
>>> import axelrod
>>> strategies = [axelrod.Cooperator(),
...                axelrod.Defector(),
...                axelrod.ZDExtort2(),
...                axelrod.Joss(),
...                axelrod.HardTitForTat(),
...                axelrod.HardTitFor2Tats(),
...                axelrod.TitForTat(),
...                axelrod.Grudger(),
...                axelrod.TitFor2Tats(),
...                axelrod.WinStayLoseShift(),
...                axelrod.Random(),
...                axelrod.ZDGTFT2(),
...                axelrod.GTFT(),
...                axelrod.HardProber(),
...                axelrod.Prober(),
...                axelrod.Prober2(),
...                axelrod.Prober3(),
...                axelrod.Calculator(),
...                axelrod.HardGoByMajority()]
>>> tournament = axelrod.Tournament(strategies)
>>> results = tournament.play()
>>> plot = axelrod.Plot(results)
>>> p = plot.boxplot()
>>> p.show()
```

Figure 8: Source code for reproducing the tournament of [42].

some of these have knowledge of the length of the game (Back stabber and Double crosser). This is not necessary to rank well in both wins and score as demonstrated by Looker Up and Fool Me Once.

· Strategies like Looker Up and Meta Hunter seem to be generally cooperative yet still exploit naive strategies. The Meta Hunter strategy is a particular type of Meta strategy which uses a variety of other strategy behaviours to choose a best action. These strategies perform very well in general and continue to do so in noisy tournaments.

Conclusion

This paper has presented a game theoretic software package that aims to address reproducibility of research into the Iterated Prisoner's Dilemma. The open nature of the development of the library has lead rapidly to the inclusion of many well known strategies, many novel strategies, and new and recapitulated insights.

The capabilities of the library mentioned above are not at all comprehensive, a list of the current abilities include:

· Noisy tournaments.
· Tournaments with probabilistic ending of interactions.
· Ecological analysis of tournaments.
· Moran processes.

· Morality metrics based on [39].
· Transformation of strategies (in effect giving an infinite number of strategies).
· Classification of strategies according to multiple dimensions.
· Gathering of full interaction history for all interactions.
· Parallelization of computations for tournaments with a high computational cost.

These capabilities are constantly being updated.

Acknowledgements

The authors would like to thank all contributors. Also, they thank Robert Axelrod himself for his well wishes with the library.

Competing Interests

The authors declare that they have no competing interests.

References

1. **Astropy Collaboration,** et al. "Astropy: A community Python package for astronomy". In: *Astronomy and Astrophysics* 558, A33 (Oct. 2013), A33. DOI: http://dx.doi.org/10.1051/0004-6361/201322068. arXiv: 1307.6212 [astro-ph.IM]
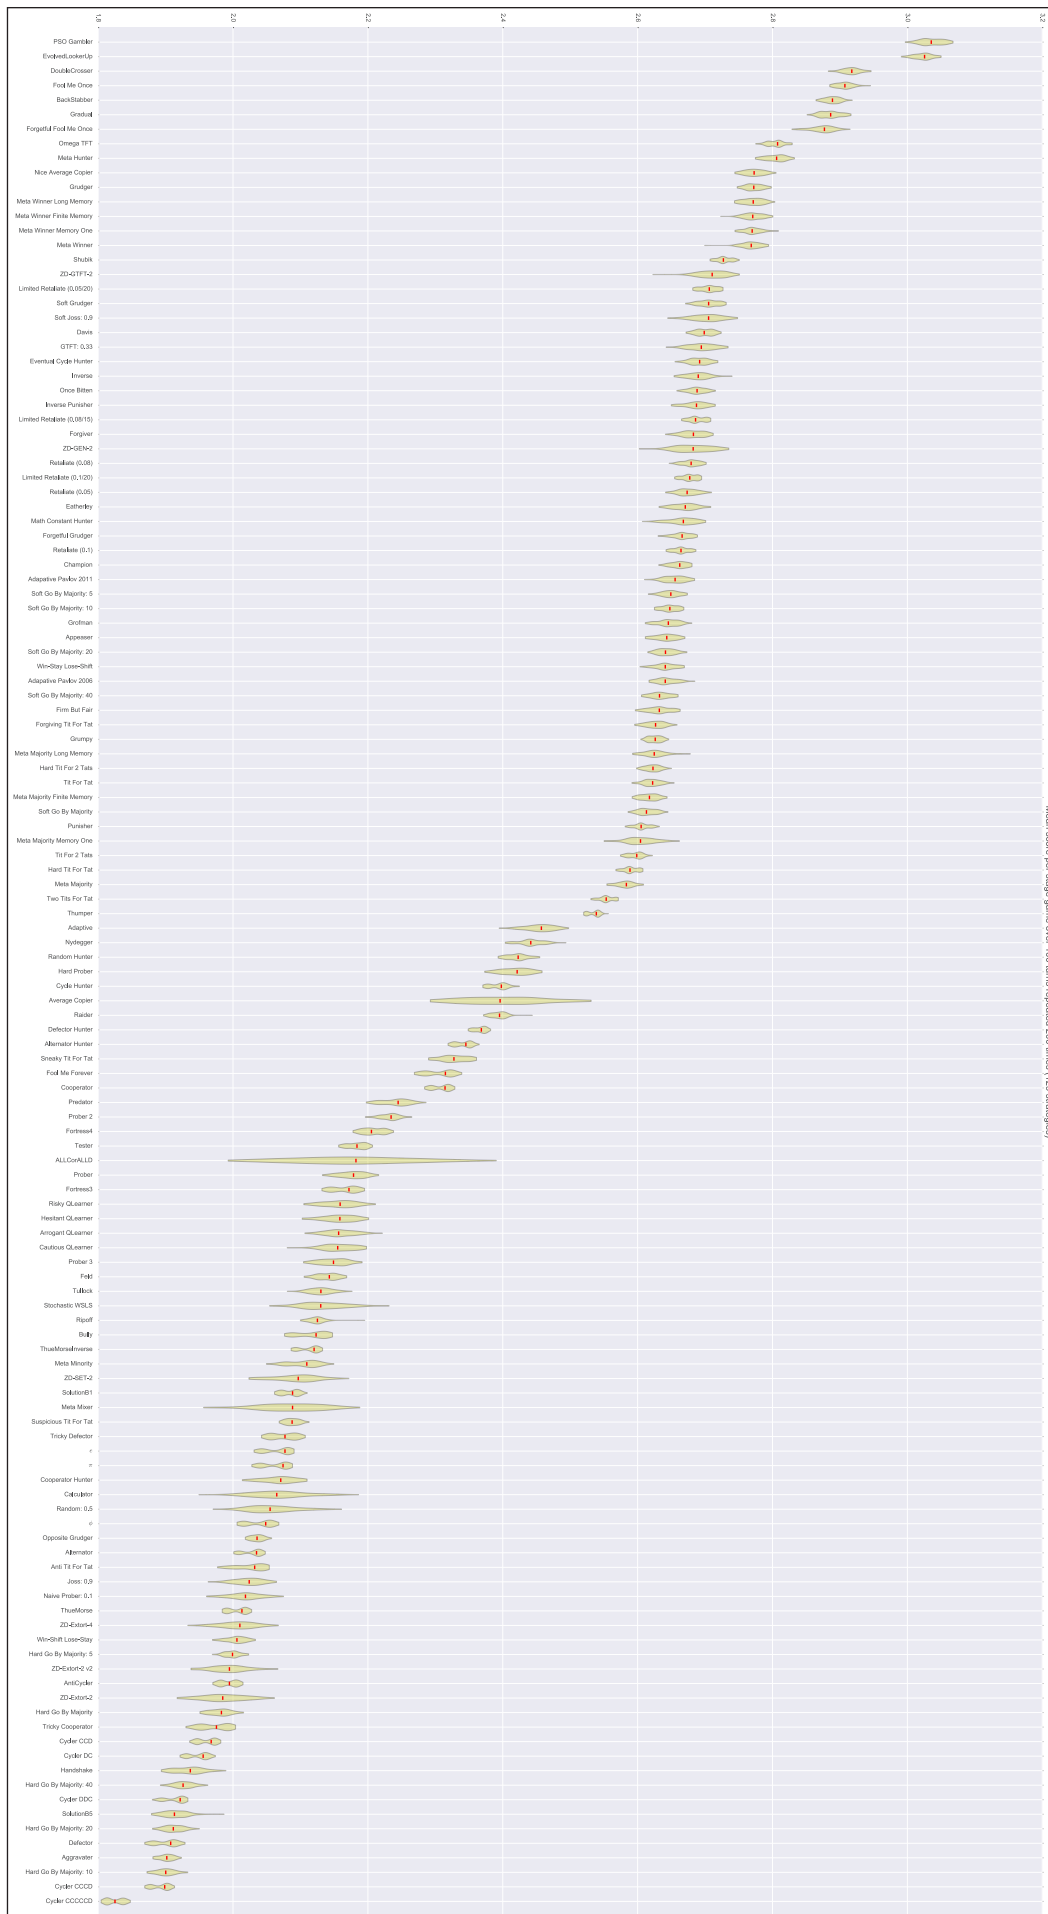
Figure 9: Results from the library tournament (2016-06-13).

2. **Axelrod, R** 1980 "Effective Choice in the Prisoner's Dilemma". In: *Journal of Conflict Resolution* 24.1, pp. 3–25. DOI: http://dx.doi.org/10.1177/002200278002400301

3. **Axelrod, R** 1980 "More Effective Choice in the Prisoner's Dilemma". In: *Journal of Conflict Resolution* 24.3, pp. 379–403. ISSN: 0022-0027. DOI: http://dx.doi.org/10.1177/002200278002400301

4. **Axelrod, R M** 2006 *The Evolution of Cooperation.* Basic books.

5. **Banks, J S** and **Sundaram, R K** 1990 "Repeated games, finite automata, and complexity". In: *Games and Economic Behavior* 2.2, pp. 97–117. ISSN: 08998256. DOI: http://dx.doi.org/10.1016/0899-8256(90)90024-O

6. **Bendor, J, Kramer, R M** and **Stout, S** 1991 "When in doubt . . .: Cooperation in a noisy prisoner's dilemma". In: *Journal of Conflict Resolution* 35.4, pp. 691–719. ISSN: 0022-0027. DOI: http://dx.doi.org/10.1177/0022002791035004007

7. **Boyd, R** and **Lorberbaum, J P** 1987 "No pure strategy is evolutionarily stable in the repeated Prisoner's Dilemma game". In: *Nature* 327, pp. 58–59. ISSN: 0028-0836. DOI: http://dx.doi.org/10.1038/327058a0

8. **Chellapilla, K** and **Fogel, D B** 1999 "Evolution, neural networks, games, and intelligence". In: *Proceedings of the IEEE* 87.9, pp. 1471–1496. ISSN: 00189219. DOI: http://dx.doi.org/10.1109/5.784222

9. **Crick, T,** et al 2014 "Share and Enjoy": Publishing Useful and Usable Scientific Models". arXiv: 1409.0367.

10. **David, B F** 1993 "Evolving Behaviors in the Iterated Prisoner's Dilemma". In: *Evol. Comput.* 1.1, pp. 77–97. ISSN: 1063-6560. DOI: http://dx.doi.org/10.1162/evco.1993.1.1.77

11. **Doebeli, M** and **Hauert, C** 2005 "Models of cooperation based on the Prisoner's Dilemma and the Snowdrift game". In: *Ecology Letters* 8.7, pp. 748–766. ISSN: 1461023X. DOI: http://dx.doi.org/10.1111/j.1461-0248.2005.00773.x

12. **Ellison, G** 1994 "Cooperation in the prisoner's dilemma with anonymous random matching". In: *Review of Economic Studies* 61.3, pp. 567–588. ISSN: 00346527. DOI: http://dx.doi.org/10.2307/2297904

13. **Gotts, N, Polhill, J** and **Law, A** 2003 "Agent-based simulation in the study of social dilemmas". In: *Artificial Intelligence Review* 19, pp. 3–92. ISSN: 0269-2821. DOI: http://dx.doi.org/10.1023/A:1022120928602

14. **Harper, M** 2015 *Marc Harper Codes.* http://marcharper.codes/2015-11-17/ipd2.html

15. **Hilbe, C, Nowak, M A** and **Traulsen, A** 2013 "Adaptive Dynamics of Extortion and Compliance". In: *PLoS ONE* 8.11, e77886. ISSN: 1932-6203. DOI: http://dx.doi.org/10.1371/journal.pone.0077886

16. **Hong, N P C,** et al 2015 "Top Tips to Make Your Research Irreproducible". pp. 5–6. arXiv: 1504.00062.

17. **Hunter, J D** 2007 "Matplotlib: A 2D graphics environment". In: *Computing In Science & Engineering* 9.3, pp. 90–95. DOI: http://dx.doi.org/10.1109/MCSE.2007.55

18. **Isaac, A** 2008 "Simulating Evolutionary Games: A Python-Based Introduction". In: *Journal of Artificial Societies and Social Simulation* 11.3, p. 8. ISSN: 14607425.

19. **Jones, M** 2015 *Evolving strategies for an Iterated Prisoner's Dilemma tournament.* http://mojones.net/evolving-strategies-for-an-iterated-prisoners-dilemma-tournament.html

20. **Kendall, G, Yao, X** and **Chong, S Y** 2007 *The iterated prisoners' dilemma: 20 years on.* World Scientific Publishing Co., Inc.

21. **Knight, V** 2015 http://vknight.org/unpeudemath/gametheory/2015/11/28/Experimenting-with-a-high-performing-evolved-strategy-in-other-environments.html.

22. **Koutsovoulos, G** 2016 *Optimising the LookerUp strategy for an Iterated Prisoner's Dilemma tournament.*

23. **Kraines, D** and **Kraines, V** 1989 "Pavlov and the prisoner's dilemma". In: *Theory and Decision* 26.1, pp. 47–79. ISSN: 00405833. DOI: http://dx.doi.org/10.1007/BF00134056

24. **Lee, C, Harper, M** and **Fryer, D** 2015 "The Art of War: Beyond Memory-one Strategies in Population Games". In: *Plos One* 10.3, e0120625. ISSN: 1932-6203. DOI: http://dx.doi.org/10.1371/journal.pone.0120625

25. **Li, J** 2007 "How to design a strategy to win an IPD tournament". In: *The iterated prisoners dilemma* 20, pp. 89–104. DOI: http://dx.doi.org/10.1142/9789812770684_0004

26. **Li, J, Hingston, P** and **Kendall, G** 2011 "Engineering design of strategies for winning iterated prisoner's dilemma competitions". In: *Computational Intelligence and AI in Games, IEEE Transactions on* 3.4, pp. 348–360. DOI: http://dx.doi.org/10.1109/tciaig.2011.2166268

27. **Lorberbaum, J P** 1994 "No strategy is evolutionarily stable in the repeated Prisoner's Dilemma game". In: *Journal of Theoretical Biology* 168.2, pp. 117–130. DOI: http://dx.doi.org/10.1006/jtbi.1994.1092

28. **MacIver, D R** 2016 *Hypothesis 3.0.3.* https://github.com/DRMacIver/hypothesis.

29. **Maschler, M, Solan, E** and **Zamir, S** 2013 *Game theory.* Cambridge University Press, p. 1003. ISBN: 9781107005488. DOI: http://dx.doi.org/10.1017/CBO9780511794216

30. **Mckelvey, R,** et al 2006 *Gambit: Software tools for game theory.* Tech. rep.

31. **McKinney, W** 2010 "Data Structures for Statistical Computing in Python". In: *Proceedings of the 9th Python in Science Conference*. Ed. by S. van der Walt and J. Millman. pp. 51–56.

32. **Milgrom, P, Roberts, J** and **Wilson, R** 1982 "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma". In: *Journal of Economic Theory* 252, pp. 245–252.

33. **Molander, P** 1985 "The optimal level of generosity in a selfish, uncertain environment". In: *The Journal of Conflict Resolution* 29.4, pp. 611–618. ISSN: 0022-0027. DOI: http://dx.doi.org/10.1177/0022002785029004004

34. **Murnighan, J K,** et al 1983 "Expecting Continued Play in Prisoner ' s Dilemma Games". 27.2, pp. 279–300.

35. **Pedregosa, F,** et al 2011 "Scikit-learn: Machine Learning in Python". In: *Journal of Machine Learning Research* 12, pp. 2825–2830.

36. **Press, W H** and **Dyson, F J** 2012 "Iterated Prisoner's Dilemma contains strategies that dominate any

evolutionary opponent". In: *Proceedings of the National Academy of Sciences* 109.26, pp. 10409–10413. ISSN: 0027-8424. DOI: http://dx.doi.org/10.1073/pnas.1206569109

37. **Prlić, A** and **Procter, J B** 2012 "Ten Simple Rules for the Open Development of Scientific Software". In: *PLoS Computational Biology* 8.12, e1002802. ISSN: 1553-7358. DOI: http://dx.doi.org/10.1371/journal.pcbi.1002802

38. **Sandve, G K,** et al 2013 "Ten Simple Rules for Reproducible Computational Research". In: *PLoS Computational Biology* 9.10, pp. 1–4. ISSN: 1553734X. DOI: http://dx.doi.org/10.1371/journal.pcbi.1003285

39. **Singer-Clark, T** 2014 "Morality Metrics On Iterated Prisoners Dilemma Players".

40. **Slany, W** and **Kienreich, W** 2007 "On some winning strategies for the iterated prisoners dilemma". In: *The iterated prisoners dilemma*, pp. 171–204.

41. **Stephens, D W, McLinn, C M** and **Stevens, J R** 2002 "Discounting and reciprocity in an Iterated Prisoner's Dilemma." In: *Science (New York, N.Y.)* 298.5601, pp. 2216–2218. ISSN: 00368075. DOI: http://dx.doi.org/10.1126/science.1078498

42. **Stewart, A J** and **Plotkin, J B** 2012 "Extortion and cooperation in the Prisoner's Dilemma". In: *Proceedings of the National Academy of Sciences* 109.26, pp. 10134–10135. ISSN: 0027-8424. DOI: http://dx.doi.org/10.1073/pnas.1208087109

43. **The Sage Developers.** *Sage Mathematics Software (Version 7.0).* http://www.sagemath.org. 7.0.