

Discussion of "50 Years of Classification and Regression Trees"

Rusch, Thomas; Zeileis, Achim

Published in:
International Statistical Review

DOI:
[10.1111/insr.12062](https://doi.org/10.1111/insr.12062)

Published: 01/01/2014

Document Version
Early version, also known as preprint

[Link to publication](#)

Citation for published version (APA):

Rusch, T., & Zeileis, A. (2014). Discussion of "50 Years of Classification and Regression Trees". *International Statistical Review*, 82, 361 - 367. <https://doi.org/10.1111/insr.12062>

To see the wood for the trees: Discussion of “50 Years of Classification and Regression Trees”

Thomas Rusch
WU Vienna

Achim Zeileis
Universität Innsbruck

Abstract

In this discussion paper, we argue that the literature on tree algorithms is very fragmented. We identify possible causes and discuss good and bad sides of this situation. Among the latter is the lack of free open-source implementations for many algorithms. We argue that if the community adopts a standard of creating and sharing free open-source implementations for their developed algorithms and creates easy access to these programs the bad sides of the fragmentation will be actively combated and will benefit the whole scientific community.

Keywords: classification and regression trees, open-source software, discussion

1 Introduction

We thank Wei-Yin Loh for this review paper [Loh, 2014]. He provides a much-needed guide to tree methods currently available as well as the main ideas behind them, indicative of his experience with and knowledge of this topic. His contribution proves to be very valuable in bringing structure into the vast interdisciplinary field of tree algorithms: We found 83 different tree induction algorithms for different response types listed in his paper and, along the lines of Loh’s disclaimer, this is not even an exhaustive list.

The availability of so many different algorithms for fitting tree-structured models directly relates to the main point of our discussion: The tree literature is highly fragmented. Loh hints at that issue already on the first page and we gladly take it up for discussion: There are so many recursive partitioning algorithms in the literature that it is nowadays very hard to see the wood for the trees.

In the remainder of our discussion paper we identify causes for and consequences of this fragmentation, discuss what we perceive to be advantages and disadvantages of the current state of the tree algorithm literature and offer suggestions that might improve the situation in the years ahead by retaining advantages and overcoming disadvantages.

2 The Fragmentation of Tree Algorithms

Currently, there is an abundance of different tree algorithms coming from different communities including statistics, machine learning and other fields. We believe that this fragmentation emerged from various causes and has a number of implications for the development and application of tree models. Some of them are in our opinion good, some are not so good and some are rather unfortunate.

2.1 The Good

The area of tree algorithms is a popular and fruitful field of research in statistics, computer science, and beyond. This leads to many people with different backgrounds contributing to the application and development of tree algorithms for various tasks. May this be to derive a set of if-then rules to make decisions, to analyze a large number of data relatively fast, to segment data, to detect or select important variables and interactions or to simply have an interpretable, visualizable, data-driven prediction machine with good performance, that is flexible and can be adapted easily to the problem at hand. An important contributing factor to their popularity is that recursive partitioning algorithms are easily adapted to different situations, as their core principles are easy to understand and intuitive. Most tree algorithms comprise a couple of similar steps, the difference between them entering at some point in the induction stage, where during development a concrete choice must be made – usually related to the loss function or measure of node impurity, split variable selection, split point selection, or pruning. Thus by, for example, changing the loss function used to measure node impurity, a new tailored algorithm can be easily invented for a given problem. Owing to this we now have tree algorithms for many types of problems and variables we might encounter, say, for online or dynamic data, longitudinal data, big sample sizes, substantive data models, various error structures and so forth. Looking at it this way, the fragmentation reflects in part the diversity of the problems encountered by the community of scholars working in this field, as well as the many ideas they have and different applications they face. Given that there is no free lunch [Wolpert, 1996], a rich diversity in algorithmic solutions is to be welcomed as no single solution will always lead to the best results. Thus, tree models were and are an active field of research and hopefully will remain so in the future.

2.2 The Bad

This abundance of tree algorithms also has its dark side. For one, not only is it hard to keep up to date with various developments but it is even harder to *choose* the “right” algorithm for a given problem. Take the case of regression trees for explaining and predicting a metric outcome and assume there is additive Gaussian error. Which algorithm to take? There are, among others, AID, CART, CTree, C4.5, GUIDE, M5 – all having different properties in different

settings. There is a lack of guidance as to which algorithm to select. One might narrow the possibilities down by looking at certain additional, desirable properties like unbiasedness in split variable selection but that still leaves one with a number of possibilities to consider. People looking to solve their problems with tree algorithms might easily be intimidated by the large number of possibilities and if different tree algorithms give different answers. Perhaps this contributed to many people inventing a new algorithm for their specific problem rather than work through different properties of existing algorithms and benchmarking them against each other on their data, which in turn perpetuates the fragmentation problem.

The fact that tree methods can be easily adapted to new situations can backfire. First, this sometimes leads to new tree algorithms or changes to old ones that appear *ad hoc*, which was already noted by Murthy [1998]. While experimentation is a necessary part of algorithm development, we should nevertheless take care to propose and use well-motivated, well-founded, methodologically sound procedures in the end. Second, modifications or improvements of old algorithms are often considered to be entirely new algorithms. They may be named differently and are often only viewed “as a whole” rather than emphasizing which steps in the tree induction are similar and which are different (e.g., the loss function might change but the split variable selection is the same or vice versa). Hence, fragmentation is increased more than necessary and common properties are obscured. This seems to tie in with a third bad effect: many authors who propose or apply tree algorithms either are not aware of – or choose to ignore – similar work in that area. It happens that even recent papers do not refer to work done from 2000 onwards, therefore ignoring more than a decade of active development that may be highly relevant. On the one hand this leads to reinventing the wheel, loss of time and resources and again more fragmentation of the literature. On the other hand, this also makes it hard for new algorithms to be noticed in an evergrowing, dense wood – an unfortunate situation both for developers and users alike. Perhaps all these points explain why, while there has been numerous new developments and improvements in tree modelling since the seminal work on C4.5 and CART in the 1980s, both remain the most popular tree algorithms for classification and regression [Wu et al., 2008].

2.3 The Ugly

As algorithmic models, classification and regression trees are very closely tied to their respective implementations. With few exceptions, tree models usually *are* their specific implementation. Hence they can be characterized by the specific combination of the generic computational steps that are adopted and by the way they are turned into a specific implementation – with both aspects being highly intertwined. While this algorithm-implementation dualism is rather natural for algorithmic models, there is also an ugly side to it: The potential lack of free (or any) implementations for new algorithms. This potential lack of access to the core of the actual tree model makes understanding, using, assessing, and extending it much more difficult.

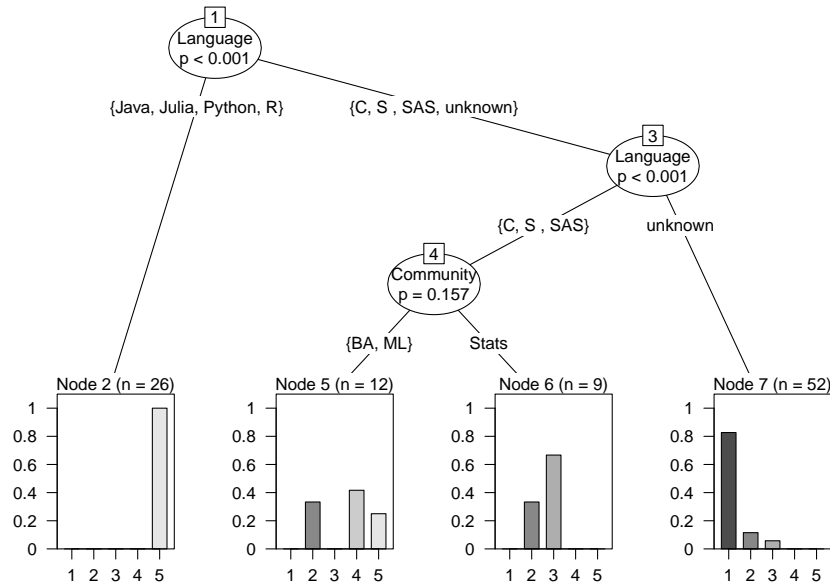


Figure 1: A classification tree of tree algorithms fitted with CTree [Hothorn et al., 2006]. The target variable is the group of implementation as defined in Section 2.3. Predictor variables were the publication “Year”, the programming “Language”, whether it allows fitting of a classification (“CT”), regression (“RT”), or model tree (“MT”), author names (“FirstAuthor”, “SecondAuthor”, “LastAuthor”), “Community” the algorithm is aimed at (machine learning vs. business analytics vs. statistics) and “Journal” venue.

To illustrate this empirically, we consider 99 algorithms – including the 83 mentioned by Loh plus 16 very recent or less known ones. For 43 of these we were not able to find an implementation¹ (see also Figure 1). More specifically, we placed all 99 algorithms into one of five broad classes pertaining to the availability of (free) software:

1. *Algorithms without an existing implementation:* These are algorithms for which a theoretical description was published but no implementation seems to exist. Often these are old algorithms or algorithms who were developed for a specific problem. Examples are AID or SUPPORT.
2. *Algorithms with a closed-source, for-profit implementation:* These are implementations of particular algorithms that are sold by a company. The code and specific implementation is kept a proprietary secret. Typical examples are M5, CHAID, or CART.

¹We searched with Google for all permutations of author names and algorithm names combined with the words “software” and “implementation”, as well as on the main author’s homepages.

3. *Algorithms with a closed-source, free-of-charge implementation:* These are implementations of particular algorithms that can be obtained free of charge, usually in an executable binary format. The code and specific implementation, however, is still kept a proprietary secret. Examples include GUIDE, CTMBR, HTL.
4. *Algorithms with an open-source, free-of-charge implementation:* These are implementations of particular algorithms that can be obtained free of charge and whose source code is open. However, these implementations either restrict or not explicitly allow copying, adaptation and distribution of the source code. Examples currently include SECRET and C4.5.
5. *Algorithms with a free and open-source implementation:* These algorithms have implementations that follow the ideas of free software (FLOSS; free, libre open-source software, see Free Software Foundation, Inc., 2013). They are open-source and give the user extensive rights with respect to copying, modification and distribution. Examples are most algorithms developed for FLOSS software packages like Weka or R, including re-implementations of closed-source algorithms, e.g., RPart, M5', LMT, and CTree but also C5.

We find that most algorithms belong to group 1 (43), followed by group 5 (29). The group of open-source algorithms/implementations (classes 4 and 5) only comprises 34% of all algorithms. To take a closer look at how this availability of (free) software depends on other characteristics of the algorithms, we naturally employ a classification tree (Figures 1, built by CTree). We see that an implementation in R, Python, Java (primarily in the packages Weka, KNIME, RapidMiner) and Julia is predictive of belonging to group 5, whereas other languages are either predictive for group 1 (if we do not know the language), of groups 2 and 3 respectively if suggested to the statistics community and for groups 4 and 5 for implementations directed at the machine learning and business analytics communities. At any rate, it shows that unfortunately FLOSS is far from being the standard for tree modeling software.

Since trees are inherently algorithmic, we view the implementation as an integral part of each algorithm. In our opinion restrictions to viewing, modifying, and sharing tree implementations is one of the main reasons for the bad sides of fragmentation discussed above. Depending on which group an algorithm belongs to this has different implications: For example, for algorithms belonging to groups 1 through 4, this leads to the need for authors proposing improvements to existing algorithms to implement the improved algorithm from scratch. Often this adapted implementation is then again not FLOSS and the problem perpetuates. Not being able to rely on existing code also applies to using an algorithm on different platforms. Another example is that the restriction of distribution and modification effectively prohibits to change the specific implementation (say, with regards to adapt it to parallel computing) or to improve it (say, with regards to speed). This also restricts the possibilities of combining the algorithms with other methods to form a pipeline of methods and distribute

the bundle. For algorithms from groups 1 through 3, a further consequence is that specific steps that may not be well documented can be hard to reproduce [as was the case with M5, Quinlan, 1993, that prompted M5', Wang and Witten, 1997, as a “a rational reconstruction”].

We strongly believe that these and other implications of a lack of free implementations led to a lot of synergy potential having been lost over the years. Partly because conceptual similarities and differences of various tree algorithms were not obvious enough and partly because the lack of reusable computational tools slowed down the pace development. Furthermore, there seems to be some confusion among practitioners as to which algorithms perform well (or even best) for their particular problems which often leads to suboptimal algorithms being used. This view appears to be shared increasingly by other researchers [e.g., Vukićević et al., 2012].

3 A Possible Remedy

We have a suggestion as to what we think will improve or even solve the problems discussed in Sections 2.2 and 2.3 while retaining the advantages mentioned in Section 2.1: (Academic) publications of tree algorithms should be accompanied by free implementations (in the FLOSS sense). This means opening the source code of past and future implementations, giving users permission to modify, adapt, and distribute it with an appropriate free software license and making the code/implementations easily publicly available, preferably with a low adoption threshold (for example in a popular language such as Python, C, Java, or R).

Specifically, we think the following six steps should be undertaken to reduce the bad and ugly aspects to a minimum while at the same time retaining the good:

- Every newly suggested algorithm or larger improvement should come with a FLOSS implementation.
- The source code of currently existing implementations should be opened and they should be licensed with a FLOSS license.
- For algorithms for which there is no up-to-date or even existing software any trademark should be relinquished and the original source code should be made freely available or re-implemented as FLOSS (as has already happened with, e.g., J4.8 or RPart).
- All implementations should be made available on a public repository or archive, the author's homepages or as freely accessible supplementary material to articles.
- When a software is used, extended, modified, etc., the software (and not only the underlying algorithm) should be referred to and cited.
- When reviewing or editing papers or algorithms we should point out the above and demand this as a new standard.

FLOSS software licenses should be used so that the copyright holders grant the rights to inspect, modify, and (re)distribute the software. Most FLOSS licenses preserve the original copyright in such modified/extended versions and in an academic context citation of the software (and not only the underlying algorithm) is appropriate.

Then, improvements to algorithms do not need to be their own algorithm but can be suggested or submitted as patches or adaptations. Code building blocks (e.g., for split point selection, or predictions, or tree visualizations) can be reused and recombined or be made computationally more efficient. Assessment and comparison of algorithms is facilitated, both for evaluating newly suggested methods or for choosing a particular model in practice. Hence, not only users of FLOSS implementations will profit but also the authors because their work is easier to understand, use, and ultimately cite.

Finally, the steps above can also reduce the fragmentation and possibly achieve a certain degree of standardization through developing and reusing computational “tree toolkits”. Some effort in this vein has been made already. For example, there are two R packages providing standardized frameworks: `partykit` [Hothorn and Zeileis, 2014] for representing, summarizing, and plotting of various tree models from different free software sources, and `caret` [Kuhn, 2008] for training, tuning, and benchmarking of various tree algorithms (among many other methods). Other efforts of providing such standardization exist as well [Vukićević et al., 2012].

4 Conclusion

In this discussion, we follow up on Loh’s review paper and take a closer look at the fragmented field of tree models. While indicative of a creative, active, and diverse community of researchers, this fragmentation also leads to undesirable side effects making it hard to understand, assess, use and compare tree algorithms. Hence, a common language for describing tree models, both conceptually and (perhaps more importantly) computationally is crucial to reduce the fragmentation. In particular, making free, libre and open-source software (FLOSS) should be an integral part of the communication about classification and regression trees. We argue that this can alleviate many of the problems caused by or following from the fragmentation while at the same time retaining the good that comes from having a bright and vibrant community.

Especially in light of open research areas that Loh mentions in his conclusions, FLOSS implementations are an effective means for reducing fragmentation in the future and tackling open hard problems in tree algorithm research faster than it was possible before. The good news is that the tree community has already started to move into this direction and an increase in FLOSS implementations can already be observed. With free platforms for statistical computing such as R or Python as well as initiatives like the “Foundation for Open Access Statistics” (FOAS, <http://www.foastat.org/>), the conditions are now better than ever before. We should use this momentum. Rather than not seeing the

wood for the trees, the whole community can grow a healthy, open, and light forest of trees within which we can all walk with intimate familiarity.

Acknowledgements

We want to thank Wei-Yin Loh for his stimulating article and Probal Chaudhuri for inviting us to join this discussion.

References

- Free Software Foundation, Inc. (2013). What is free software? Version 2013-06-18 05:16:52.
- Hothorn, T., K. Hornik, and A. Zeileis (2006). Unbiased recursive partitioning: A conditional inference framework. *Journal of Computational and Graphical statistics* 15(3), 651–674.
- Hothorn, T. and A. Zeileis (2014, March). partykit: A modular toolkit for recursive partytioning in R. Working Paper 2014-10, Working Papers in Economics and Statistics, Research Platform Empirical and Experimental Economics, Universität Innsbruck.
- Kuhn, M. (2008). Building predictive models in R using the caret package. *Journal of Statistical Software* 28(5), 1–26.
- Loh, W.-Y. (2014). 50 years of classification and regression trees. *International Statistical Review* 82(3), 329–348.
- Murthy, S. K. (1998). Automatic construction of decision trees from data: A multi-disciplinary survey. *Data Mining and Knowledge Discovery* 2(4), 345–389.
- Quinlan, J. R. (1993). *C4.5: Programs for Machine Learning*. San Mateo, California: Morgan Kaufmann Publishers.
- Vukićević, M., M. Jovanović, B. Delibašić, S. Iščamović, and M. Suknović (2012). Reusable component-based architecture for decision tree algorithm design. *International Journal on Artificial Intelligence Tools* 21(05).
- Wang, Y. and I. Witten (1997). Induction of model trees for predicting continuous classes. In *Proceedings of the Posters of the European Conference on Machine Learning*, Prague, Czech Republic. University of Economics, Faculty of Informatics and Statistics.
- Wolpert, D. H. (1996). The lack of a priori distinctions between learning algorithms. *Neural Computation* 8(7), 1341–1390.

Wu, X., V. Kumar, R. Quinlan, J. Ghosh, Q. Yang, H. Motoda, G. McLachlan, A. Ng, B. Liu, Y. Philip, Z.-H. Zhou, M. Steinbach, D. Hand, and D. Steinberg (2008). Top 10 algorithms in data mining. *Knowledge and Information Systems* 14(1), 1–37.