gcplyr: an R package for microbial growth curve data analysis

Michael Blazanin[1]

[1]Department of Ecology and Evolutionary Biology, Yale University, New Haven, CT 06511, USA.

Correspondence to: Mike Blazanin, Department of Ecology and Evolutionary Biology, Yale University, 165 Prospect St, New Haven, CT, 06520, USA; mike.blazanin@yale.edu

Abstract

Characterization of microbial growth is of both fundamental and applied interest. Modern platforms can automate collection of high-throughput microbial growth curves, necessitating the development of computational tools to handle and analyze these data to produce insights. To address this need, here I present a newly-developed R package: gcplyr. gcplyr can flexibly import growth curve data in common tabular formats, and reshapes it under a tidy framework that is flexible and extendable, enabling users to design custom analyses or plot data with popular visualization packages. gcplyr can also incorporate metadata and generate or import experimental designs to merge with data. Finally, gcplyr carries out model-free (non-parametric) analyses. These analyses do not require mathematical assumptions about microbial growth dynamics, and gcplyr is able to extract a broad range of important traits, including initial density, lag time, growth rate and doubling time, carrying capacity, diauxie, area under the curve, extinction time, and more. gcplyr makes scripted analyses of growth curve data in R straightforward, streamlines common data wrangling and analysis steps, and easily integrates with common visualization and statistical analyses.

Availability

gcplyr is available from the central CRAN repository (https://CRAN.R-project.org/package=gcplyr), or from GitHub (https://github.com/mikeblazanin/gcplyr).

Introduction

Characterization of microbial population growth dynamics has been of fundamental and applied interest since nearly the dawn of microbiology [1]. Today, growth curves are a ubiquitous technique to study microbial growth. Indeed, modern automated platforms, including plate readers, can collect high-throughput growth data over time on hundreds of samples simultaneously. Yet, this data-generation capacity has outpaced the development of computational tools to handle and analyze microbial growth data, presenting new challenges.

First and foremost, data are rarely output in the ideal format for analysis, visualization, or publication. Reorganizing data manually can be tedious and fraught with the potential for introduction of errors. Moreover, since output files vary between different plate readers, scripted reorganization may require tailored code for each output format. Despite this, only a few existing software tools provide utilities for streamlined data wrangling and reorganization (Table S1) [2–8].

Once data are reorganized, scientists face the challenge of converting raw data into quantitative microbial traits. Many existing computational tools use parametric analyses of microbial growth curve data (Table S1) [3, 4, 6, 8–18]. Parametric analyses fit a mathematical model of microbial population growth to observed data, then extract the fitted parameter values to quantify traits. This approach can be powerful, but it has drawbacks for many users. First, users must explicitly or implicitly choose a model of growth, making error-prone mathematical assumptions about the form the growth data should take. Moreover, for accurate results users must laboriously verify that their data meet the assumptions of the model and that the fitting process converged on an appropriate solution. Finally, some dynamics do not fit any known models of microbial growth, leaving researchers with few to no options to analyze such growth data.

Given these drawbacks, in recent years some groups have developed software to analyze microbial growth data with non-parametric or model-free approaches (Table S1) [3, 6–8, 14, 15, 17–26]. These analyses make no specific mathematical assumptions about the form of the growth data, instead extracting parameters of interest directly from the data itself or from non-parametrically smoothed transformations of the data. However, many of these tools can only quantify a few traits of interest non-parametrically.

Here, I present my newly-released R package, gcplyr. gcplyr is built as a versatile tool to wrangle plate reader data and carry out model-free analyses. gcplyr is used in R, a popular scripting language for scientific data analysis and visualization. With gcplyr, users can import growth curve data from a wide range of plate reader formats and carry out model-free analyses to extract a number of traits, including lag time, growth rate and doubling time, carrying capacity, area under the curve, and more. Notably, the model-free analyses implemented by gcplyr enable the characterization of dynamics that do not fit traditional mathematical models of microbial growth, and therefore were difficult or impossible to analyze using the parametric methods of many existing tools. Examples of these dynamics include diauxic shifts [27] and growth in the presence of antagonists [e.g. bacteriophages [28]]. Additionally, gcplyr provides a framework for data organization that is flexible and extendable, making it easy for users to run custom analyses or integrate gcplyr workflows with other R packages.

Implementation

gcplyr is an open-source R package [29] that can be installed with R's built-in install.packages function from the central CRAN repository (https://CRAN.R-project.org/package=gcplyr), or from GitHub (https://github.com/mikeblazanin/gcplyr). gcplyr has been written to minimize the number of external dependencies needed for installation.

Usage

gcplyr is usable with a basic working knowledge of the R coding language. A comprehensive manual and tutorials are included with the gcplyr installation and available online (https://mikeblazanin.github.io/gcplyr/). Within R, gcplyr functions make it easy to import data and metadata, merge them with experimental design information, smooth and calculate derivatives as necessary, and analyze curves to produce summary growth curve metrics (Fig 1). gcplyr was designed to easily integrate growth curve analyses with other common R tasks, including plotting with ggplot2 [30], data transformation with tidyverse [31], and statistical analysis.

Inputs

Growth curve data is typically output by plate reader software as block-shaped (also called 'matrix') or wide-shaped (also called 'column-wise') (Fig 1). gcplyr can import any tabular file containing block-shaped or wide-shaped data, with user-set arguments enabling functionality regardless of where the data is located in the input file. This is an especially useful advance for block-shaped data, where each timepoint is saved separately, and which only one recently-released tool can parse [5]. Additionally, user-set arguments can designate metadata located in input files to be incorporated into the resulting data.

Reshaping data

After reading input files into R, gcplyr extends the functionality of tidyr [32] to reshape data into a 'tidy' format (also known as 'long' format) [33] ideal for subsequent visualization and analysis steps. Tidy data have all observations in a single column, with each unique datapoint with its own row, and additional columns specifying the timepoint, well, and any other information. Tidy data is the best layout for most analyses [33], is consistent with requirements of data repositories like Dryad, and is the expected input for a number of popular R packages [31]. By transforming data into tidy data, gcplyr makes it easy for users to also visualize their data using ggplot2 [30], manipulate their data using dplyr [34], or apply any of the other tidyverse packages to their data [31].

Incorporating experimental designs

gcplyr can import experimental design information from user-supplied spreadsheet files, or users can use gcplyr functions to directly generate experimental designs in R. Once imported or generated, these designs can be exported to files or easily merged with imported growth curve data.

Smoothing data

Unlike fitting-based approaches, the model-free analyses that gcplyr implements can be sensitive to experimental noise in growth curve data. To smooth noise in raw density data, gcplyr implements several well-established smoothing functions, including moving average, moving median, loess [29, 35–37], and

GAM [38–41]. Smoothing functions are tuned by user-set smoothness parameters. The package documentation discusses best practices for setting smoothness parameters and smoothing data.

Calculating derivatives

Model-free calculation of many growth curve metrics depends on identifying features in the derivatives of density data. gcplyr can calculate both the population rate of growth (derivative) and cellular growth rate (per-capita derivative), using the slope of density over time. However, such derivatives can be sensitive to experimental noise, even if density data has been smoothed. To address this, gcplyr calculates derivatives by fitting a linear regression on a rolling window of multiple points [7], with user-set parameters defining the width of the window. The package documentation discusses best practices for calculating derivatives and setting their smoothing parameters.

Characterizing microbial growth

Finally, gcplyr quantifies traits of interest by computing metrics from density and derivatives data. Many of these metrics are quantified by identifying peaks or valleys in the density or derivatives (e.g. maximum growth rate, doubling time, diauxie, carrying capacity), using a local extrema-finding algorithm implemented in gcplyr. Other metrics are quantified by identifying threshold-crossing events in the density or derivatives. Finally, some metrics are quantified using purpose-built functions that implement previously-established methods (e.g. lag time [42, 43] or area under the curve).

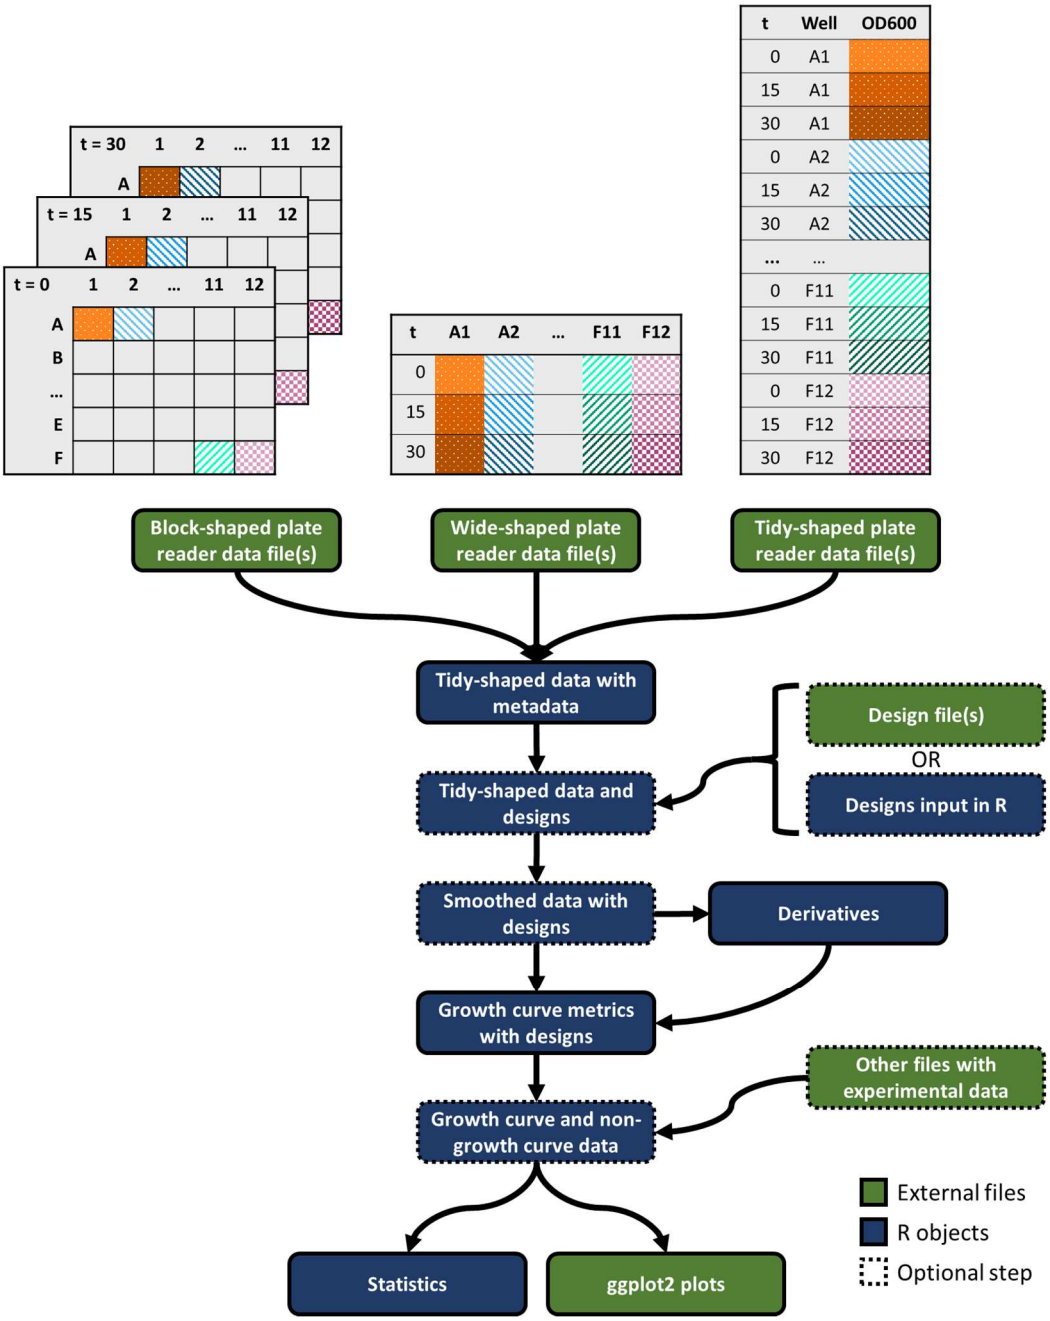

Fig 1. Workflow to use gcplyr to analyze microbial growth curve data. Growth curve data are commonly found in one of three layouts: block-shaped, wide-shaped, or tidy-shaped. Block-shaped data are organized to match the physical layout of the multi-well plate. Wide-shaped data contain one column for each well from a plate. Tidy data contain all the observations in a single column. gcplyr functions import data files regardless of the data layout, reshape them into tidy-shaped data, then merge them with experimental design information. Data and designs can then be smoothed and have derivatives calculated to extract growth curve metrics. Metrics can be easily merged with other non-growth curve experimental data before statistical analyses and visualization.

Discussion

High-throughput microbial growth data has necessitated computational tools for data wrangling and analysis. Here I introduced gcplyr, a new R package built specifically to address this need.

First, gcplyr wrangles and reshapes growth curve data. gcplyr recognizes that there are common layouts for growth curve data (Fig 1), and it flexibly imports such files. During data reorganization, gcplyr employs a tidy data framework, providing a number of benefits over other data organization approaches:

1. It enables users to integrate as many pieces of experimental design information as desired.
2. It enables users to leverage the general-use functions in gcplyr and other packages to generate custom analyses to identify unique features in their data.
3. It enables users to easily integrate their growth curve analyses with existing visualization (e.g. ggplot2) and statistics packages in R.
4. It enables users to merge growth curve data analyses with other sources of data.

gcplyr is the first package designed for scriptable wrangling of growth curve data in the R programming ecosystem (Table S1). This provides high customizability for users, and enables them to easily incorporate gcplyr with other scripted analyses. In contrast, many existing growth curve analysis tools have little to no data wrangling capability (Table S1), requiring users to reformat files manually. Some packages do provide data reshaping via non-R user interfaces, including bletl via Python scripting [6], and Parsley [5], QurvE [3], and AUDIT [4] via graphical user interfaces.

gcplyr is also the first package in R that is designed for scriptable non-parametric analyses for a wide array of microbial growth traits (Table S1). Many existing tools quantify traits of interest by fitting parametric mathematical models of microbial growth to growth curve data, including biogrowth [16], growthcurver [14] and AUDIT [4]. Some tools provide non-parametric analyses through non-R interfaces, including the Python packages bletl [6], phenom [24], and AMiGA [23], and the graphical user interface-based QurvE [3]. And some scriptable R packages provide non-parametric estimation of specific traits, including area under the curve by growthcurver [14], and maximum growth rate by growthrates [15]. gcplyr builds on these approaches to quantify many different traits non-parametrically, including maximum growth rate, lag time, area under the curve, carrying capacity, diauxie, and more.

gcplyr is a versatile tool designed for users who have a basic level of familiarity with programming in R. Because gcplyr is used via R scripts, it can provide customized data wrangling and growth curve analysis for each user. However, this flexibility comes with a tradeoff: it is less accessible to users who have no experience with R programming, who may be better served by software with graphical user interfaces (GUI) [3–5, 9–13, 19–22].

All growth curve analyses are limited by their input data. Most frequently, growth curves are collected by automated spectrophotometers, where optical density is used as a proxy for microbial cell density [although see [44] for an alternate proxy for cell density]. However, optical density has limited sensitivity to detect low concentrations of microbial cells, and typically saturates before maximum concentrations of cells are reached [45]. These limitations constrain any analysis of optical density-based growth curves, but may especially affect non-parametric analyses like those of gcplyr.

Additionally, non-parametric analyses have limitations relative to model-fitting approaches. First, input data can contain measurement error, which can more strongly affect non-parametric analyses. To address this, gcplyr includes several smoothing methods, although these are not a panacea for noise. Additionally, non-parametric approaches often require more fine-grained temporal resolution for precise quantification of microbial growth traits. Finally, the traits quantified by non-parametric approaches can sometimes differ from those estimated by model fitting. For instance, non-parametric analyses quantify the maximum realized cellular growth rate, while parametric analysis quantify the intrinsic growth rate [46]. These limitations highlight the relative strengths and weaknesses of the non-parametric approach taken by gcplyr. Parametric methods can be best when density dynamics closely follow known mathematical models of microbial growth and users can validate that model fitting converged on appropriate solutions, while non-parametric analyses enable quantification of traits outside those conditions.

Ultimately, growth curves are a widespread experimental approach in the microbial sciences. However, until now, tools capable of wrangling and doing model-free analysis of growth curve data were limited. By enabling these steps, gcplyr streamlines high-quality model-free analysis for a wide range of applied and fundamental research on microbial growth.

Acknowledgements

Thanks to Paul Turner for support during this project, to Jyot Antani, Alita Burmeister, Noah Houpt, Albert Vill, and Jordan Lewis for feedback on drafts of this paper, and to all the gcplyr alpha and beta testers.

References

1. Monod J. The Growth of Bacterial Cultures. Annual Review of Microbiology. 1949;3:371–94.

2. Hughes SM. plater: Read, Tidy, and Display Data from Microtiter Plates. The Journal of Open Source Software. 2016;1.

3. Wirth NT, Funk J. QurvE: Robust and User-Friendly Analysis of Growth and Fluorescence Curves. 2023.

4. Coutin NPJ, Giaever G, Nislow C. Interactively AUDIT Your Growth Curves with a Suite of R Packages. G3. 2020;10:933–43.

5. Csibra E, Stan G-B. Parsley: a web app for parsing data from plate readers. Bioinformatics. 2023;:btad733.

6. Osthege M, Tenhaef N, Zyla R, Müller C, Hemmerich J, Wiechert W, et al. bletl - A Python package for integrating BioLector microcultivation devices in the Design-Build-Test-Learn cycle. Engineering in Life Sciences. 2022;22:242–59.

7. Hall BG, Acar H, Nandipati A, Barlow M. Growth rates made easy. Molecular Biology and Evolution. 2014;31:232–8.

8. Vaas LAI, Sikorski J, Hofner B, Fiebig A, Buddruhs N, Klenk HP, et al. Opm: An R package for analysing OmniLog® phenotype microarray data. Bioinformatics. 2013;29:1823–4.

9. Bukhman YV, DiPiazza NW, Piotrowski J, Shao J, Halstead AGW, Bui MD, et al. Modeling Microbial Growth Curves with GCAT. Bioenergy Research. 2015;8:1022–30.

10. Huang L. IPMP 2013 — A comprehensive data analysis tool for predictive microbiology. International Journal of Food Microbiology. 2014;171:100–7.

11. Liu Y, Wang X, Liu B, Yuan S, Qin X, Dong Q. Microrisk Lab: An Online Freeware for Predictive Microbiology. Foodborne Pathogens and Disease. 2021;18:607–15.

12. Veríssimo A, Paixão L, Neves AR, Vinga S. BGFit: Management and automated fitting of biological growth curves. BMC Bioinformatics. 2013;14.

13. Vervier K, Browne HP, Lawley TD. CarboLogR: a Shiny/R application for statistical analysis of bacterial utilisation of carbon sources. 2019;:695676.

14. Sprouffske K, Wagner A. Growthcurver: An R package for obtaining interpretable metrics from microbial growth curves. BMC Bioinformatics. 2016;17:17–20.

15. Petzoldt T. growthrates: Estimate Growth Rates from Experimental Data. 2022.

16. Garre A, Koomen J, Besten H den, Zwietering M. biogrowth: Modelling of Population Growth. 2022.

17. Cuevas DA, Edwards RA. PMAnalyzer: A new web interface for bacterial growth curve analysis. Bioinformatics. 2017;33:1905–6.

18. Cuevas DA, Garza D, Sanchez SE, Rostron J, Henry CS, Vonstein V, et al. Elucidating genomic gaps using phenotypic profiles. F1000Research. 2016;3:1–28.

19. Swain PS, Stevenson K, Leary A, Montano-Gutierrez LF, Clark IBN, Vogel J, et al. Inferring time derivatives including cell growth rates using Gaussian processes. Nat Commun. 2016;7:13766.

20. Fernandez-Ricaud L, Kourtchenko O, Zackrisson M, Warringer J, Blomberg A. PRECOG: A tool for automated extraction and visualization of fitness components in microbial growth phenomics. BMC Bioinformatics. 2016;17:249.

21. Jung PP, Christian N, Kay DP, Skupin A, Linster CL. Protocols and Programs for High-Throughput Growth and Aging Phenotyping in Yeast. PLoS ONE. 2015;10:e0119807.

22. Olsen B, Murakami CJ, Kaeberlein M. YODA: Software to facilitate high-throughput analysis of chronological life span, growth rate, and survival in budding yeast. BMC Bioinformatics. 2010;11:141.

23. Midani FS, Collins J, Britton RA. AMiGA: Software for Automated Analysis of Microbial Growth Assays. mSystems. 2021;6:e00508-21.

24. Tonner PD, Darnell CL, Bushell FML, Lund PA, Schmid AK, Schmidler SC. A Bayesian non-parametric mixed-effects model of microbial growth curves. PLOS Computational Biology. 2020;16:e1008366.

25. Tonner PD, Darnell CL, Engelhardt BE, Schmid AK. Detecting differential growth of microbial populations with Gaussian process regression. Genome Res. 2017;27:320–33.

26. Hemmerich J, Wiechert W, Oldiges M. Automated growth rate determination in high-throughput microbioreactor systems. BMC Research Notes. 2017;10:617.

27. Blazanin M, Moore JP, Olsen S, Travisano M. Fight not flight: parasites drive the bacterial evolution of resistance, not avoidance. 2023;:2023.04.29.538831.

28. Blazanin M, Vasen E, Jolis CV, An W, Turner P. Theoretical validation of growth curves for quantifying phage-bacteria interactions. bioRxiv. 2023. https://doi.org/10.1101/2023.06.29.546975.

29. R Core Team. R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing; 2022.

30. Wickham H. ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag New York; 2016.

31. Wickham H, Averick M, Bryan J, Chang W, McGowan LD, François R, et al. Welcome to the tidyverse. Journal of Open Source Software. 2019;4:1686.

32. Wickham H, Vaughan D, Girlich M. tidyr: Tidy Messy Data. 2023.

33. Wickham H. Tidy Data. Journal of Statistical Software. 10.

34. Wickham H, François R, Henry L, Müller K. dplyr: A Grammar of Data Manipulation. 2023.

35. Savitzky A, Golay MJ. Smoothing and differentiation of data by simplified least squares procedures. Analytical chemistry. 1964;36:1627–39.

36. Cleveland WS, Devlin SJ. Locally weighted regression: an approach to regression analysis by local fitting. Journal of the American statistical association. 1988;83:596–610.

37. Cleveland WS. Robust locally weighted regression and smoothing scatterplots. Journal of the American statistical association. 1979;74:829–36.

38. Trevor Hastie, Robert Tibshirani. Generalized Additive Models. Statistical Science. 1986;1:297–310.

39. Wood SN. Thin-plate regression splines. Journal of the Royal Statistical Society (B). 2003;65:95–114.

40. Wood SN. Generalized Additive Models: An Introduction with R. 2nd edition. Chapman and Hall/CRC; 2017.

41. Wood SN. Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. Journal of the Royal Statistical Society (B). 2011;73:3–36.

42. Swinnen IAM, Bernaerts K, Dens EJJ, Geeraerd AH, Van Impe JF. Predictive modelling of the microbial lag phase: a review. International Journal of Food Microbiology. 2004;94:137–59.

43. Yates GT, Smotzer T. On the lag phase and initial decline of microbial growth curves. Journal of Theoretical Biology. 2007;244:511–7.

44. Mackie AM, Hassan KA, Paulsen IT, Tetu SG. Biolog Phenotype MicroArrays for Phenotypic Characterization of Microbial Cells. In: Paulsen IT, Holmes AJ, editors. Environmental Microbiology: Methods and Protocols. Totowa, NJ: Humana Press; 2014. p. 123–30.

45. Biesta-Peters EG, Reij MW, Joosten H, Gorris LGM, Zwietering MH. Comparison of Two Optical-Density-Based Methods and a Plate Count Method for Estimation of Growth Parameters of Bacillus cereus. Applied and Environmental Microbiology. 2010;76:1399–405.

46. Ghenu A-H, Marrec L, Bank C. Challenges and pitfalls of inferring microbial growth rates from lab cultures. Frontiers in Ecology and Evolution. 2024;11.

Table S1. gcplyr and other available microbial growth curve data wrangling and analysis tools.

Software	Used via	Param. analysis	Non-param. analysis	Scriptable	Imports	Data reshaping	Incorporate designs	Plotting	Smoothing	Citation
gcplyr	R package	No	Yes	Yes	Block, wide, tidy	Yes, into wide or tidy	Yes	With ggplot2	Yes	This paper
growthcurver		Yes	Minimally	Yes	Wide	No	No	Built-in	No	[14]
growthrates			Minimally	Yes	Tidy	No	No	With ggplot2	Yes	[15]
biogrowth			No	Yes	Wide	No	No	Built-in	No	[16]
opm			Minimally	Yes	Specialized formats	Yes, into opm-specific class	Yes	Built-in	Yes	[8]
plater		No	No	Yes	Block	Yes, into tidy	No	No	No	[2]
QurvE	GUI or R package	Yes	Yes	Yes	Wide, specialized formats, tidy	Yes, from specialized formats into wide	Yes	Built-in	Yes	[3]
AUDIT	GUI	Yes	No	Yes, in part	Tidy, specialized formats	Yes, from specialized formats into tidy	Yes	Built-in, with ggplot2	Yes	[4]
Parsley		No	No	No	Block, wide	Yes, into tidy	Yes	No	No	[5]
bletl	Python package	Minimally	Yes	Yes	Specialized formats	Yes	Yes	Built-in	Yes	[6]
AMiGA	Python package or command line	No	Yes	Some	Wide	No	Yes	Built-in	Yes	[23]
fitderiv	GUI or Python package	No	Yes	Some	Wide	No	No	Built-in	Yes	[19]
phenom	Python code	No	Yes	Yes	Wide	No	Yes	No	Yes	[24]
B-GREAT			Yes	Yes	Wide	No	Yes	No	Yes	[25]
PMAnalyzer	bash	Yes	Minimally	No	Wide	No	No	Built-in	No	[17, 18]
GrowthRates	Command line	No	Yes	No	Wide, specialized formats	Yes	No	No	No	[7]
-----	MATLAB code	No	Yes	Yes	One well at a time	No	No	No	No	[26]
GCAT	GUI	Yes	No	No	Wide	No	Yes	Built-in	No	[9]
PRECOG		No	Yes	No	Wide	No	No	Built-in	Yes	[20]
IPMP 2013		Yes	No	No	Wide	No	No	Built-in	No	[10]
GATHODE		No	Yes	No	Wide	No	Minimally	Built-in	Yes	[21]
Microrisk Lab		Yes	No	No	Wide	No	No	Built-in	No	[11]
CarboLogR		Yes	No	No	Specialized formats	No	Yes	Built-in	No	[13]
YODA	Webpage GUI	No	Yes	No	Wide	No	No	No	No	[22]
BGFit		Yes	No	No	Wide	No	No	Built-in	No	[12]