

RESEARCH ARTICLE

Discovering data-driven microbial growth models with symbolic regression

T. Anthony Sun¹  | Dovydas Kičiatovas¹ | Inga-Katariina Aapalampi² |
Teemu Kuosmanen¹ | Teppo Hiltunen² | Ville Mustonen¹

¹Department of Organismal and Evolutionary Biology, Department of Computer Science, University of Helsinki, Helsinki, Finland

²Department of Biology, University of Turku, Turku, Finland

Correspondence

T. Anthony Sun
Email: anthony.sun@helsinki.fi

Ville Mustonen
Email: v.mustonen@helsinki.fi

Funding information

Research Council of Finland, Grant/Award Number: 345829, 346128 and 364234

Handling Editor: James Orr

Abstract

1. Connecting mathematical models with empirically measured microbial growth has remained challenging, as numerous competing models based on different theoretical approaches can fit observations. Therefore, we develop a method to automatically propose growth models from microbial data alone. We validate this approach using an available dataset of *E. coli* grown on known resources, and study 14 species across various concentrations of a rich medium.
2. The inherently interpretable approach of symbolic regression infers explicit dynamical models directly from growth data. Using symbolic regression natively, does not favour biologically interpretable models, but we find cumulative population gain to be a more informative machine learning feature than population size.
3. Random Forest machine learning allows us to relate this finding to the approximation of a constant-rate per capita resource consumption. This suggests that the area under the growth curve (AUC) measured in routine experiments provides information on the effective resource dynamics governing microbial growth. Finally, we use theoretical insights to inform the symbolic regression algorithm and favour biologically interpretable models.
4. Overall, we found that balancing between data fit, parsimony and biological relevance favoured both the simplest, linear approximation and models based on Monod dynamics, with either one or two underlying resources. Therefore, our approach to read growth laws off of microbial batch cultures provides insights on data-driven modelling.

KEYWORDS

bacterial growth, Monod dynamics, Random Forest regression, resource-consumer model, symbolic regression

1 | INTRODUCTION

Microbial growth is a key process in the ecology and evolution of microbes. For instance, understanding it quantitatively is essential for public health in the face of rising antibiotic resistances in

human pathogens (Antimicrobial Resistance Collaborators, 2022; Wong, 2024), and for various biotechnology applications (Lässig et al., 2023; Yeoman et al., 2020). Yet, quantifying how microbial populations convert resources into biomass remains challenging, with most experiments relying on chemically complex media and

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *Methods in Ecology and Evolution* published by John Wiley & Sons Ltd on behalf of British Ecological Society.

short culturing times. So far, there has been no simple way to use batch culture data to infer how microbial populations grow without relying on strong prior modelling assumptions. We know for certain that a growing microbial population uptakes resources from the environment and converts them to new biomass, and previous research has established that growth in batch cultures involves two intrinsic nonlinearities. On the one hand, the typical lag phase observed in growth experiments (Allen & Waclaw, 2018) implies that minimal models must include a time-dependent adjustment function to account for physiological changes occurring in bacterial cells encountering a fresh culture medium (Baranyi & Roberts, 1994; Borse et al., 2024; Huang, 2011; Ram, Dellus-Gur, et al., 2019). On the other hand, the maximum per capita growth rate has a saturating response to an increase in resource concentration, demonstrated by Monod's classical work on a single resource (Monod, 1949).

A wealth of theoretical work has proposed models for multi-resource growth (Allen & Waclaw, 2018; Baranyi & Roberts, 1994; De Jong et al., 2017; Ghenu et al., 2024; Held et al., 2024; Huang, 2011; Letten, 2022; Zwietering et al., 1990). Yet, quantitative analyses of empirical bacterial growth have often circumvented any explicit account of resource dynamics by using traditional models based on logistic growth (Asnicar et al., 2023; Ram, Dellus-Gur, et al., 2019), which might misrepresent microbiological processes at work (Balsa-Canto et al., 2019; Hatton et al., 2024; Ishizawa et al., 2024; Picot et al., 2023; Ram, Obolski, et al., 2019; Turchin, 2001). Our inability to fully use the biologically more appropriate resource-consumer models does not arise from a lack of theoretical tools (Cui et al., 2024), but from the fact that model complexity often far surpasses experimental evidence. This evidence typically covers observed population data but not resource dynamics, which creates a disconnect between observable phenotypes related to macroscopic growth (e.g. yield, effective growth rate, lag time) and underlying microscopic mechanisms (e.g. maximum growth rate, resource consumption) (Miguel Trabajo et al., 2024).

Indeed, several incompatible multi-resource models may fit experimental data equally well and produce, in effect, indistinguishable simulated observations (Breiman, 2001; Ghenu et al., 2024; Held & Manhart, 2024; Peleg & Corradini, 2011; Tang & Riley, 2021). This means that, in addition to the widespread question of parameter identifiability (Balsa-Canto et al., 2019), models are themselves not generally discoverable at realistic resolutions in microbiology. Even with an improved experimental resolution, we know that no algorithm can definitively find the generative model in arbitrarily complex cases (Richardson, 1969) nor in microbiological conditions (Shumaylov et al., 2025). The problem is that the various theoretical approaches to generate hypotheses themselves cannot be criticised on the basis of data: Since traditional processes first assume a set of theoretical models, then gather empirical observations to fit the models before subsequent model selection, the initial choice of models itself is methodologically immune to data-based criticism. Thus, how the data could

allow us to identify a true, generative model of microbial growth is not generally a well-posed problem, outside the ideal case of perfect knowledge of all latent variables.

Therefore, we use a radically different approach of hypothesis generation, which is to accept working with approximations sufficient for specific use cases. This means asking which models suffice to reproduce the observed behaviour, to propose null hypotheses *allowed* by empirical observations, irrespective of our theoretical priors, at first. In that sense, a good null-model trades off among data-drivenness, goodness-of-fit, biological interpretability and parsimony. In this data-driven approach, generating models as null hypotheses should not primarily result from our assumptions but from finding patterns in the data, before model selection. Since theoretical input does not occur at the model proposal stage, it now becomes an essential part of model selection: after gathering empirical data, we should propose theoretical models in an unbiased way to fit data and, finally, filter model proposals through model selection, which now includes verifying that they make biological sense.

So far, an interpretable, yet automated, method to uncover dynamical growth models from experimental data has been missing. One interpretable alternative to black-box machine learning is symbolic regression, a method for the unbiased discovery of relationships between predictors and a dependent variable (Kronberger et al., 2025). By directly searching the space of mathematical expressions, it can fruitfully serve in hypothesis generation (Radwan et al., 2025). While computationally intensive, computing power advancements in recent decades have made symbolic regression viable and increasingly popular across a multitude of scientific disciplines, yielding impressive results in recovering and expanding known laws in physics (Course & Nair, 2023) and, in the context of microbial ecology, in reverse-engineering Lotka-Volterra and logistic-based models (Gaucel et al., 2014; Hsin et al., 2023; Martin et al., 2018; Regalado, 2021). At its basic usage, this methodology requires no assumptions or any other inputs than the data itself, but its true potential is in the ability to specify various constraints for the mathematical expressions by defining allowed mathematical operations and partial functional expressions. The latter constraint is extremely important for utilisation of the known mathematical laws and general domain knowledge, as showcased in this study. Symbolic regression is not an alternative to existing models, but an alternative to the approach of specifying candidate models *a priori*.

Here, we make minimal theoretical assumptions to better understand growth data obtained for 14 bacterial species in Reasoner's 2A (R2A) medium (Figure 1). Since our ultimate goal is to propose models in arbitrary experimental settings with tractable or intractable resources, we use a complex medium where resources are not tightly controlled to represent the general case, in addition to validating our approach using *E. coli* data obtained with known chemical resources (Held et al., 2024). The aim is to infer an effective resource dimensionality telling us how complex the medium is from the point of view of a growing population of bacteria. We demonstrate that symbolic regression can reverse-engineer explicit dynamical models for the microscopic growth process. We find cumulative population




FIGURE 1 Workflow for the production and analysis of the datasets.

gain to be a more informative machine learning feature than population size, then link this finding to the approximation of a constant-rate per capita resource consumption. We use theoretical insights to inform the symbolic regression algorithm and favour biologically interpretable models. The algorithm suggests that bacterial growth widely follows a Monod dynamics in our experiments, with either one or two effective underlying resources, even in our chemically complex medium. This is generally a slightly better approximation than a linear-consumption model, but it remarkably applies to even a complex medium.

2 | MATERIALS AND METHODS

2.1 | Experimental methods for bacterial cultures

Our experiments involved 16 species from the HAMBI collection (Helsinki Institute of Life Science HiLIFE, Biodiversity Collections Research Infrastructure HUBCRI, Microbial Domain Biological Resource Centre HAMBI) summarised in Table S1. All were able to grow in Reasoner's 2A medium and the following culture protocol yielded our data, summarised in Figure S1.

Solution M9 contained 11.28 g of M9 minimal medium salts (MP BIOMEDICALS, lot 142, 662) in 1.0 L of MILLI-Q purified water. The culture medium (100%) was prepared by dissolving then autoclaving 3.0 g of Reasoner's 2A broth powder (NEOGEN, lot UK208078/190) into 1.0 L of solution M9. We prepared a range of 20 different concentrations C of culture medium, separated by an incremental factor of 0.8, from 200% to 2.88%, by successive dilutions into autoclaved solution M9. Bacterial pre-cultures were run by inoculating each species into 6.0 mL of culture medium (100%), then incubating for 48 h in aerobic conditions at 30°C under agitation at 70 rpm and finally diluted into autoclaved solution M9 by a factor of 10^{-4} .

Monocultures were run in four batch replicates by inoculating 110 μ L of each culture medium concentration C with 20 μ L of pre-culture onto a 96-well plate, then incubating for 48 h in aerobic conditions at 30°C under agitation at 800 rpm. An AGILENT BioTEK LOGPHASE 600 device automatically recorded time and optical density at 600 nm every 10 min as a proxy to population size, yielding 16 species $\times$ 4 biological replicates $\times$ 20 concentrations = 1280 experimental growth curves, each comprising 289 time points. We considered 14 among the 16 species because two did not reach stationary phase. We derived all other observations numerically.

2.2 | Theoretical methods and models

Consider a bacterial population of size N growing in a substrate of concentration S . The variables S and N are time-dependent. To account for the lag period observed in microbiological batch cultures, minimal models of bacterial growth assume a time-dependent physiological adjustment α in their population dynamics (Baranyi & Roberts, 1994; Ghenu et al., 2024; Zwietering et al., 1990):

$$\rho_{\text{obs}} \stackrel{\text{def}}{=} \frac{1}{N} \frac{dN}{dt} = \alpha(t) \times \rho(S), \quad (1)$$

where ρ is the time-independent and potentially resource-dependent growth function. How physiological adjustment delays bacterial cells' growth dynamics is usually represented by similar mathematical forms (Baranyi & Roberts, 1994; Borse et al., 2024; Ghenu et al., 2024; Huang, 2011; Lo Grasso et al., 2023). From fits to real data, we know that just two phenomenological parameters suffice to capture the adjustment function (Borse et al., 2024; Ram, Dellus-Gur, et al., 2019), namely the initial physiological state of the cell (population) q (dimensionless) and the rate of change m (per unit time) in adjusting its physiology to the present environmental state:

$$\alpha(t) = \frac{q}{q + e^{-mt}}. \quad (2)$$

Monod (1949) showed that the maximum ρ_{max} of the empirical per capita growth rate ρ_{obs} was an increasing then saturating function of the resource concentration. The demonstration concerns only the maximum per capita growth rate $\rho_{\text{max}}(S)$ in a bacterial growth experiment (Angaroni et al., 2025; Monod, 1949). However, this type II functional response pattern (Holling, 1959a) is usually

modelled using a hyperbolic function (Holling, 1959b; Michaelis & Menten, 1913; Monod, 1949) and considered a function of time within a dynamical model (Miguel Trabajo et al., 2024):

$$\rho(S) = g_{\text{max}} g(S) \text{ with } g(S) \stackrel{\text{def}}{=} \frac{S}{K + S}, \quad (3)$$

where K is the substrate concentration inducing half-maximum response. Without further assumptions, this functional form cannot model multi-resource growth (Monod, 1949), which may be complicated by non-mutually exclusive resources (Perrin et al., 2020), colimitation (Held et al., 2024; Saito et al., 2008), and sequential preference (Kremling et al., 2018). As it is unclear how several resources ($S_0, \dots, S_s$) together impact bacterial growth, multi-resource growth models combine each resource's Monod contribution to the total growth in various ways, notably with multiplicative or additive effects (Held et al., 2024; Held & Manhart, 2024; Letten, 2022; Saito et al., 2008):

$$\rho_M(S_0, \dots, S_s) = g_{\text{max}} \prod_{x=0}^s g_x(S_x), \quad (4)$$

$$\rho_A(S_0, \dots, S_s) = g_{\text{max}} \sum_{x=0}^s g_x(S_x). \quad (5)$$

The multiplicative-effect multi-resource model combines Equations (1) and (4). It assumes that growth requires all resources, which are thus biologically not substitutable. The additive-effect multi-resource model combines Equations (1) and (5). It assumes that any resource allows for some growth, which implies substitutable resources. A common framework can unify several formulations (Tang & Riley, 2021), but at the cost of a combinatorial explosion of theoretical parameters, as the number of resources increases. Overall, experimental data provides no definitive, general support for either model over the other (Held & Manhart, 2024; Tang & Riley, 2021), so that we have no prior hypothesis in our experimental setting involving several species and a rich medium.

2.3 | Numerical methods

Each dataset consisted of a range of bacterial growth curves at different initial medium concentrations C for a given species. We used PYTHON (3.10.14) for analysis, NUMPY (1.26.4) and PANDAS (2.2.2) for numerical protocols, MATPLOTLIB (3.8.4) for visualisations, SCIKIT-LEARN (1.4.2) for Random Forest methods, PYSR (1.3.1) for symbolic regression (Cranmer, 2023). Code is available at: <https://github.com/anthony-sun/SR-Monod> (Sun & Kičiatovas, 2026). Since the experimental batch replicates were extremely similar, we averaged population sizes across the four replicates, for each time point. This is unlikely to impact the study in practice, due to a very low coefficient of variation (median interreplicate-CV: 1.3%), but greatly facilitates computation, which is a major challenge of symbolic regression. The Supporting Information details the protocol used for simulated data and expands on the symbolic regression methods.

For symbolic regression, the aim was to find explicit mathematical growth models which we could then interpret, in light of biological knowledge, as revealing latent resources. Alternative frameworks exist, such as general additive modelling (Wood, 2025), with specific constraints on how variables may combine. We split each dataset into a training set (50%) and a testing set (50%) comprising every other growth curve along the range of medium concentrations. This ratio allowed for a sufficient testing dataset keeping the time series structure of whole growth curves, and avoided artefacts which could result from an uneven spacing of the training/testing medium concentrations. In any case, it was not possible for this unusual ratio to unduly inflate the performance of the symbolic regressor, much to the contrary. For a given upper bound on the complexity of output expressions, the algorithm optimised data fits based on mean squared-error using the default method (with crossover and mutation) and parameters in Cranmer (Cranmer, 2023). More complex functional forms resulted from randomly combining feature variables and constants using: addition (+), multiplication ($\times$), division (/) and constant power ($^k, k > 0$). We imposed biology-grounded restrictions against negative constants and non-integer powers of the relative medium concentration C . To facilitate numerical computations, we normalised each dataset's cumulative population gain N_c by its maximum value, and we used the exponential negative scaled time $U = e^{-t}$ and the exponential negative cumulative population gain $F = e^{-N_c}$ as features. Finally, we used expression templates constraining the search for mathematical expressions as detailed in the results.

For Random Forest regression, the aim was to suggest which observed variables would best predict growth rate and thus better reflected a latent resource dynamics. We split each dataset into a training set (75%) and a testing set (25%) comprising every fourth medium concentration along the range of medium concentrations. Since we were not interested in evaluating the Random Forest methods per se, but to see how various feature sets enable prediction, it was not critical to avoid artefacts which could result from an uneven spacing of the training/testing medium concentrations. The predicted response variable was the per capita growth rate ρ_{obs} . Feature variables included: One population size feature, namely N or N_c , the relative initial medium concentration C , the time t . We used all default settings of the SCIKIT-LEARN package, namely 100 estimators, the coefficient of determination R^2 as a performance score, and no cap on the maximum tree depth by default, unless otherwise specified.

3 | RESULTS

3.1 | Unconstrained symbolic regression predicts growth

Our hypothesis was that growth curves in a range of medium concentrations should allow for the inference of dynamical growth models. We demonstrated it with a symbolic regression algorithm. Symbolic regression consists of optimising a functional

form predicting a response variable under constraints on formula complexity, thus providing explicit mathematical models with little theoretical bias (Cranmer, 2023). We used this tool to explore which key observables are more appropriate for a blind model search.

We found that models relying on cumulative population gain N_c systematically outperformed models relying on population size N (Figure 2, $\{C, N\}$ and $\{C, N_c\}$) in terms of fitting performance. Feature set $\{C, N_c\}$ yielded the best overall performance (Figure 2), especially with more complex models. Thus, our naïve symbolic regression algorithm showcases the potential for N_c to inform us on microbial growth resources while providing mathematically explicit models. Yet, it could not yield substantial biological insight beyond the importance of N_c , because the models that it inferred optimised only their fitting performance and did not generally reflect basic theoretical expectations from the literature.

3.2 | Cumulative population gain allows for black-box prediction of growth

Before addressing the issue of biological relevance, we asked why the cumulative population gain feature N_c yielded better results than population size N alone. We found that the observed behaviour was at least consistent with the theoretical approximation that resource consumption is a constant-per-capita-rate process. Indeed, the law of mass action assumes that a bacterial population of size N growing in a substrate of concentration S uptakes resources at a constant rate v relative to its size:

$$\frac{1}{S} \frac{dS}{dt} = -vN \Rightarrow S = CS^0 \exp \left[-v \underbrace{\int_0^t N(u) du}_{N_c} \right], \quad (6)$$

where C is the relative initial concentration of the growth medium compared to a reference, and $S^0 = S_{(C=1, t=0)}$ is the resource density in the reference concentration. This constant-rate model for per capita resource consumption readily extends to several independent resources and naturally gives a paramount role to the cumulative population gain N_c . In laboratory terms, the cumulative population gain N_c corresponds to the area under the (growth) curve (AUC) sometimes reported in microbiology.

This theoretical expectation yielded a testable implication, independent from symbolic regression: models relying on the traditional, instantaneous population size N should perform worse than models relying on the cumulative population gain N_c , but could be supplemented by providing time data t to approximate the integral in Equation (6). To test that independently from symbolic regression, we compared the performance of a black-box Random Forest regressor using either N or N_c . Our aim was not to optimise the performance of the machine learning algorithm but to compare the use of different feature sets.




FIGURE 2 Performance of the unconstrained symbolic regressor in predicting the per capita growth rate from different feature sets in experimental data, given by the coefficient of determination R^2 . Crosses show the median train score across species. The absence of highly complex models when using certain feature sets means that the symbolic regressor did not propose any highly complex model. The cumulative population gain predictor N_c outperforms the population size predictor N as well as its combination with time t .

Effectively, using N_c instead of N improved the performance of a Random Forest regressor in predicting per capita growth rate ρ_{obs} (Figure 3, $\{C, N\}$ and $\{C, N_c\}$); two-sided correlated t -test on the models with no maximal tree depth: $p = 0.0011$). This was not due to statistical smoothing because we observed the same effect in smooth simulations (Figure S2). Feature N_c effectively captured the impact of both N and t since feature sets $\{C, N, t\}$, $\{C, N_c\}$ and $\{C, N_c, t\}$ achieved a similar improvement. They also yielded simpler decision trees: While relying on N led to a median maximum tree depth of 35, relying on N_c instead reduced it to 25 (Table 1). When capping the maximum depth of all decision trees from the start, all three yielded a systematic improvement of the Random Forest regression compared to feature set $\{C, N\}$ (Figure 3).

When relying on population size N to predict the per capita growth rate ρ_{obs} , the Random Forest regressor showed systematic overfitting, with a gap between train and test scores. Models relying instead on cumulative population gain N_c reduced this overfit.

The feature sets $\{C, N, t\}$, $\{C, N_c\}$ and $\{C, N_c, t\}$ roughly yielded a similar improvement. The theoretical formulation of Equation (6), implying a map from resource S onto the unidimensional space N_c , suggested that the algorithms could use N and t together to computationally approximate N_c . In any case, N_c effectively offered a parsimonious, unidimensional approximation of the bidimensional spaces (N, t) and (N_c, t) .

Our data confirmed that, compared to population size N , cumulative population gain N_c improved the performance of a Random Forest black-box, including by reducing the overfitting. This cannot prove that resource consumption has a constant per capita rate, but showed that this was a reasonable approximation, consistent with the data in our experiments.

3.3 | Constrained symbolic regression infers biological growth model

Next, we hypothesised that a symbolic regression algorithm could infer a biologically interpretable growth function when guided by theory-based expectations. Thus, we ran three rounds of symbolic regression: one with minimal biological assumptions (unconstrained as previously), and two informed by supplementary theoretical assumptions, using the following expression templates:

$$\rho_{\text{obs}} = \frac{\hat{q}}{\hat{q} + e^{-\hat{m}t}} \times \hat{p}(C, N_c), \quad (7)$$

$$\rho_{\text{obs}} = \frac{\hat{q}}{\hat{q} + e^{-\hat{m}t}} \times \hat{p}(C \times \hat{h}(N_c)), \quad (8)$$




FIGURE 3 Performance of a black-box Random Forest regressor in predicting the per capita growth rate from different feature sets in experimental data, given by the coefficient of determination R^2 . Crosses show the median train score across species. The cumulative population gain predictor N_c outperforms the population size predictor N .

TABLE 1 Unrestricted complexity of a Random Forest regressor using 100 trees to predict the per capita growth rate: Median across 14 experimental datasets, each comprising 5780 points and corresponding to one species.

Feature set	Train score	Test score	Median maximum tree depth (SD)
$\{C, N\}$	98.1%	72.0%	35 (5)
$\{C, N, t\}$	99.9%	91.3%	27 (4)
$\{C, N_c\}$	99.9%	91.6%	25 (3)
$\{C, N_c, t\}$	99.9%	93.0%	27 (3)

where the algorithm looked for the undefined parameters $\hat{q}$ and $\hat{m}$, and for the undefined functions $\hat{p}$ and $\hat{h}$. Note that those templates effectively mean using feature set $\{C, N_c, t\}$. Template (7) assumed Equations (1) and (2) by imposing the search for adjustment parameters q and m . Template (8) inspired by Equation (6) additionally assumed that the growth's resource dependency is proportional to the medium's relative concentration C .

We focused on how symbolic regression could inform microbial growth in two use cases corresponding to different research questions. In the first use case, our approach determined the biological substitutability or non-substitutability of known resources. In the second use case, it inferred the effective dimensionality of

the resource space, as perceived by a growing bacterial population, when the exact limiting resources were not known.

3.3.1 | Inferring non-substitutability of known resources

The first use case concerns simulations and empirical data in which chemical resources are known and varied independently from one another in the growth medium. This experimental setting allows for a modelling investigation of whether resources are substitutable or not. We simulated growth curves using a dynamical Monod model with either a non-substitutable or substitutable bidimensional resource, combining respectively Equations (1–4) and (6), or Equations (1–3), (5 and 6). Using data obtained under independently varying resource conditions, the symbolic regressor was able to correctly learn the substitutability of the provided resources (Figure S4).

We repeated those observations with available batch culture data where *E. coli* was grown on glucose as a carbon source and ammonium as a nitrogen source (Held et al., 2024). Among the best-fitting, biologically interpretable models proposed by the algorithms corresponded to a multiplicative, bidimensional, dynamical Monod formulation (Figure 4). This multiplicative model biologically implied that the two inferred resources were not substitutable, in




FIGURE 4 Fitting performance of models proposed by the symbolic regressor from growth data obtained with batch cultures of *E. coli* on known resources: Glucose and ammonium, which are non-substitutable. Circle: No template; triangle pointing down: Template (7); triangle pointing up: Template (8).

agreement with the fact that glucose was the only carbon source, and ammonium the only nitrogen source. Moreover, the fact that the one-resource models proposed relied on glucose concentration rather than on nitrogen concentration aligns with the key impact of the carbon source on the growth dynamics.

This provides for a validation of our approach through simulations and experimental data. Namely, in the ideal case where resources are known, our approach is able to inform us on their biologically relevant properties, such as their substitutability.

3.3.2 | Inferring effective dimensionality of unknown resources

The second use case concerns the more general situation where resources are not tightly controlled, not precisely known or not chemically well-defined. There, the question is rather how a given microbial species perceives a given medium: What is the effective dimensionality of this medium as reflected by the growth dynamics? Experimental datasets were obtained in R2A, a rich culture medium containing yeast extract, thus yielding no expectations regarding resource limitations. The algorithms generated 1068 models compatible with the data. A data-driven approach cannot prove nor disprove any mathematical expression: in simulated data, for example, it may or may not propose the generative model (Supplementary

Information) with no noticeable difference in fitting abilities or other performance metrics. This is an expected consequence of the known lack of model discoverability (Balsa-Canto et al., 2019; Shumaylov et al., 2025; Tang & Riley, 2021). Instead, a rigorous interpretation is that proposed models are sufficient to approximate empirical observations. To avoid focusing on well-fitting models with no biological interpretability, we manually annotated models to pick those which were compatible with theoretical expectations (Supplementary information). Overall, we found few of the proposed models to be interpretable in light of the literature in microbiology (Figure 5).

Table 2 summarises the best-fitting models that remained biologically interpretable in experiments. In many species, the algorithm initially proposed a linear model, where the resource-dependent growth ρ was a linear function of a unidimensional, latent environmental resource S . This remained the best interpretable model for one of 16 experimental species in R2A medium, suggesting that experimental conditions did not cover the saturating part of Monod's formulation. In all other species indeed, we retrieved a dynamical formulation of the Monod model corresponding to Equations (1) and (6), which slightly improved fits. This also implied an effectively unidimensional resource space where the per capita growth rate was a saturating function of a single environmental resource. In four cases, the algorithm further found that the observed growth was better explained by an




FIGURE 5 Fitting performance of models proposed by the symbolic regressor for R2A data with unknown resource dimensionality. Circle: No template; triangle pointing down: Template (7); triangle pointing up: Template (8).

TABLE 2 Formula, interpretation and inferred resource dimensionality of the best-fitting, biologically interpretable growth models for ρ in Equations (1) and (3).

Model for $\rho(S)$	Interpretation	Res. Dim.	Number of experimental species
$g_{\max} S$	Linear	1	1 sp.: 2443
$g_{\max} \frac{S}{K + S}$	Monod	1	9 sp.
$g_{\max} S_1 \times \frac{S_2}{K + S_2}$	Linear $\times$ Monod	2	2 sp.: 403, 1896
$g_{\max} \frac{S_1}{K_1 + S_1} \times \frac{S_2}{K_2 + S_2}$	Monod $\times$ Monod	2	1 sp.: 2160
More complex	Linear and Monod	≥ 2	1 sp.: 2164

effectively multidimensional resource, with constant-per-capita-rate or Monod-like components.

However, the fitting improvement brought about by more complex models was often limited (Figure 5). This means that linear models are good first approximations in many cases. Modellers may improve fitting performance by considering instead Monod's model, which has a limited cost in terms of parameters and complexity. Whether the fitting improvement is worth the increased complexity will depend on the specific objective at hand, especially when considering multidimensional Monod models.

4 | DISCUSSION

4.1 | Biological implications

In this study, we used symbolic regression to learn explicit growth models from population size data, without relying on prior theory. While fits to the data were excellent, an unconstrained use of the method did not ensure biological relevance, thus emphasising the need for a biologically bespoke usage of symbolic regression. However, our approach suggested that cumulative population gain N_c provides a powerful and simple way to indirectly observe the latent resource space experienced by growing microbial populations. We confirmed this using random forest machine learning and linked it with theoretical expectations resulting from the law of mass action expressed by Equation (6).

Finally, we combined symbolic regression with theoretical insights based on the literature, namely: biological constants are positive; expressions combine through addition, multiplication, division and power; biological mechanisms involve only integer powers of the medium concentration C ; the law of mass action of Equation (6) applies, that is, N_c is an important state variable; the effective per capita growth rate follows Equation (1) as the product of an adjustment function and a resource-dependent part. We validated our approach using an available dataset in which resources were known (Held et al., 2024) before applying it to our experimental data in R2A. We found that the dynamical resource-consumer model based on Monod's formulation in Equation (3), which we never assumed, best trades off among fit to data, symbolic simplicity and biological interpretability, which extends previous non-dynamical findings (Angarone et al., 2025). Our experimental observations effectively follow the approximation, formulated by Equations (1) and (6), of a dynamical Monod growth model with latent resources consumed at a constant per capita rate. Overall, they highlight the potential of symbolic regression as a data-driven hypothesis generation method while emphasising the need to integrate existing theory and biological knowledge when interpreting its outputs.

Our approach also provides an access to the effective resource dimensionality experienced by a growing microbial population. The effective resource dimensionality represents the minimal number of effective, independent resources theoretically required to reproduce

the observed population dynamics. In most cases, the growth dynamics can be approximated as relying on a unidimensional resource (Table 2) (Monod, 1949), such as a linear combination of the concentration of one or several compounds present in the medium. Even in our chemically complex medium, the number of effective resources remained remarkably low. The exact link between data-driven models and mechanistic resource dynamics is an excellent follow-up question raising several interpretative difficulties. For instance, if a complex medium can be approximated as relying on a unidimensional resource, which we may interpret as limiting, it does not follow that a single chemical acts as the key limiting resource throughout the growth experiment. Moreover, the fact that two species each sees the same medium as a unidimensional resource space, does not imply that they see the same linear combination of chemically defined resources, although this provides a useful null-hypothesis. Additionally, consider a chemically defined medium with a single limiting resource 1 yielding model A with resource 1'. If enriched with a second chemically defined resource 2, it could be approximated by a new model B with resources 1' and 2'. There is no guarantee that model B reduces to model A when one of its resources is fixed.

The key point is that modelled effective resources are defined purely from their impact on the growth dynamics and, as is the case with all machine learning methods, reflect the data provided. Therefore, a strict interpretation of symbolic regression outputs cannot extend to mean that we reveal chemical resources. Departing from a chemical view allows us to translate population dynamics into effective ecological niches, which can be generally disconnected from chemical resources (Meszéna et al., 2006). Among multidimensional resource models, multiplicative effects suggest that latent resources are non-substitutable (Held & Manhart, 2024), whereas additive effects suggest that they are substitutable. Thus, our approach connects microbiology experiments with a rich literature on multidimensional niches (Hutchinson, 1961; MacArthur & Levins, 1967; Posfai et al., 2017) and their overlap, which determine species coexistence in communities (Badali & Zilman, 2020; Chesson, 2018; Pastore et al., 2021; Rozas Garcia et al., 2023). Future research could modulate the concentration of single chemical resources in a culture medium to allow for a high-definition view of the growth process in a complex environment and discuss the interpretation of symbolic regression outputs in light of more chemical evidence (Takano et al., 2023).

4.2 | Methodological impact of symbolic regression in microbial ecology

Despite its success in physics (Course & Nair, 2023), symbolic regression has mainly been constrained to search logistic-based models in microbiology (Gaucel et al., 2014; Hsin et al., 2023; Martin et al., 2018; Regalado, 2021), thus overlooking two limitations. First, such models fail (Balsa-Canto et al., 2019; Hatton et al., 2024; Ishizawa et al., 2024; Picot et al., 2023; Ram, Obolski, et al., 2019; Turchin, 2001) to account for key resource dynamics

mediating the interactions (Meroz et al., 2024; Miguel Trabajo et al., 2024) that explain how bacterial communities assemble (Coyte et al., 2021; Ho et al., 2024; Lee et al., 2023), grow, organise and change (Pacheco et al., 2021), through competition and resource utilisation traits (Ho et al., 2022; Litchman et al., 2015). Secondly, assuming a functional form a priori leaves the plethora (Allen & Waclaw, 2018; Baranyi & Roberts, 1994; De Jong et al., 2017; Ghenu et al., 2024; Held et al., 2024; Huang, 2011; Letten, 2022; Zwietering et al., 1990) of growth functions proposed in the literature unexploited. Rather than developing more and more complex theory, it is essential to better connect empirical data on microbial growth with available models offering easy computation and best understanding.

We demonstrated how to combine pre-existing knowledge with the potential for symbolic regression to automatically infer an effective growth function from data, with little to no prior constraints on its functional form. Since symbolic regression does not provide automatic scientific interpretation, it is useful in generating, but not testing, hypotheses (Radwan et al., 2025). Clearly, it should not be confused for a tool to recover arbitrary generative models. Instead, symbolic regression constructs models that are sufficient to reproduce observations. This is independent from generative processes due to well-known parametric identifiability and model discoverability caveats (Held & Manhart, 2024; Ram, Obolski, et al., 2019; Shumaylov et al., 2025; Tang & Riley, 2021). Notably, our proposed models involve exclusively observable variables. Only scientific interpretation in light of pre-existing knowledge expressed by Equation (6) allows for a translation into a latent resource space. Indeed, symbolic regression outcomes cannot reject alternative hypotheses, which may still receive other empirical support, for example, through knowledge about specific metabolites involved. Thus, future research should further investigate the link between effective resources inferred by symbolic regression and specific chemicals. A general rule-of-thumb is that assuming more resources than retrieved by symbolic regression is not necessary from the viewpoint of microbial growth.

The bottom line is that all biologically interpretable models proposed by symbolic regression could capture the observed behaviours and act as null hypotheses. Model proposals must be discussed based on goodness-of-fit, especially when comparing them, and in light of a specific research question at hand. Indeed, symbolic regression does not suffice to complete the model selection process, which is distributed between constraints on model search before symbolic regression, fit to data during symbolic regression, biological interpretability after symbolic regression, and balancing fits with model complexity. The specific coarse-graining level that researchers should apply to symbolic regression outputs depends on their particular use case. For several of our species indeed, the fitting improvement brought about by considering several resources is small, compared to a unidimensional resource. This means that researchers interested in a rough approximation can generally assume a unidimensional dynamical Monod model, while researchers interested in resource consumption of their species in a particular

medium may focus on more complex models. In the extreme case, fitting a linear-consumption model remains the easiest approach, due to a low number of parameters, and often yields roughly acceptable fits. However, symbolic regression sets a null-hypothesis upper bound on resource dimensionality.

5 | CONCLUSIONS

Our study demonstrates that symbolic regression offers a promising, empirically viable and scalable way to analyse routine experiments in microbiology to infer null-models for growth and effective resource dimensionalities. Compared to the traditional method of fitting a posited family of models to data, our approach does not predefine candidate models. Instead, we build a cohort of candidate models based on the data first, then assess their theoretical and biological relevance. When resources were known and varied independently, the symbolic regressor was able to inform their metabolic interplay. When resources were not known, it was able to suggest an effective number of resources perceived by each species. We found that it is reasonable to assume a constant-rate per capita resource consumption and, based on this assumption, we expect the dynamical Monod model to capture the growth behaviour of most species in batch monocultures. Microbiologists can now assess this assumption using experimentally measured growth curves, and knowledge of resource consumption dynamics could inform our predictions and control of microbial community interactions.

AUTHOR CONTRIBUTIONS

Conceptualization: T. Anthony Sun, Dovydas Kičiatovas, Teemu Kuosmanen, Teppo Hiltunen, Ville Mustonen. Data curation: T. Anthony Sun, Dovydas Kičiatovas, Inga-Katariina Aapalampi. Formal analysis: T. Anthony Sun, Dovydas Kičiatovas, Ville Mustonen. Funding acquisition: Teppo Hiltunen, Ville Mustonen. Investigation: T. Anthony Sun, Dovydas Kičiatovas, Inga-Katariina Aapalampi, Teemu Kuosmanen, Teppo Hiltunen, Ville Mustonen. Methodology: T. Anthony Sun, Dovydas Kičiatovas, Teemu Kuosmanen, Teppo Hiltunen, Ville Mustonen. Resources: Teppo Hiltunen, Ville Mustonen. Software: T. Anthony Sun, Dovydas Kičiatovas, Ville Mustonen. Supervision: Teppo Hiltunen, Ville Mustonen. Visualization: T. Anthony Sun, Dovydas Kičiatovas. Writing—original draft: T. Anthony Sun, Dovydas Kičiatovas. Writing—review and editing: T. Anthony Sun, Dovydas Kičiatovas, Inga-Katariina Aapalampi, Teemu Kuosmanen, Teppo Hiltunen, Ville Mustonen.

ACKNOWLEDGEMENTS

The authors wish to thank Frédéric Guillaume, the reviewers and associate editors, for comments on successive versions of the manuscript, as well as CSC—IT Center for Science, Finland, for computational resources. This work was in part supported by Research Council of Finland (345829, 346128 and 364234) to Ville Mustonen and Teppo Hiltunen. Open access publishing facilitated by Helsingin yliopisto, as part of the Wiley - FinELib agreement.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

PEER REVIEW

The peer review history for this article is available at <https://www.webofscience.com/api/gateway/wos/peer-review/10.1111/2041-210x.70335>.

DATA AVAILABILITY STATEMENT

Data and code produced for this study are available via <https://doi.org/10.5281/zenodo.20119884> [reference: Sun2026_Zenodo].

The study also used publicly available data accessible in the supplementary material of: <https://doi.org/10.1073/pnas.2400304121> [reference: Held2024PNAS].

ORCID

T. Anthony Sun  <https://orcid.org/0000-0002-5277-7861>

REFERENCES

- Allen, R. J., & Waclaw, B. (2018). Bacterial growth: A statistical physicist's guide. *Reports on Progress in Physics*, 82(1), 016601. <https://doi.org/10.1088/1361-6633/aae546>
- Angarani, F., Peruzzi, A., Alvarenga, E. Z., & Pinheiro, F. (2025). Translating microbial kinetics into quantitative responses and testable hypotheses using Kinbiont. *Nature Communications*, 16, 6440. <https://doi.org/10.1038/s41467-025-61592-6>
- Antimicrobial Resistance Collaborators. (2022). Global burden of bacterial antimicrobial resistance in 2019: A systematic analysis. *Lancet*, 399(10325), 629–655. [https://doi.org/10.1016/S0140-6736\(21\)02724-0](https://doi.org/10.1016/S0140-6736(21)02724-0)
- Asnicar, F., Thomas, A. M., Passerini, A., Waldron, L., & Segata, N. (2023). Machine learning for microbiologists. *Nature Reviews Microbiology*, 22(4), 191–205. <https://doi.org/10.1038/s41579-023-00984-1>
- Badali, M., & Zilman, A. (2020). Effects of niche overlap on coexistence, fixation and invasion in a population of two interacting species. *Royal Society Open Science*, 7(2), 192181. <https://doi.org/10.1098/rsos.192181>
- Balsa-Canto, E., Alonso-del-Real, J., & Querol, A. (2019). Mixed growth curve data do not suffice to fully characterize the dynamics of mixed cultures. *Proceedings of the National Academy of Sciences of the United States of America*, 117(2), 811–813. <https://doi.org/10.1073/pnas.1916774117>
- Baranyi, J., & Roberts, T. A. (1994). A dynamic approach to predicting bacterial growth in food. *International Journal of Food Microbiology*, 23(3–4), 277–294. [https://doi.org/10.1016/0168-1605\(94\)90157-0](https://doi.org/10.1016/0168-1605(94)90157-0)
- Borse, F., Kičiatovas, D., Kuosmanen, T., Vidal, M., Cabrera-Vives, G., Cairns, J., Warringer, J., & Mustonen, V. (2024). Quantifying massively parallel microbial growth with spatially mediated interactions. *PLoS Computational Biology*, 20(7), e1011585. <https://doi.org/10.1371/journal.pcbi.1011585>
- Breiman, L. (2001). Statistical modeling: The two cultures. *Statistical Science*, 16(3), 199–215. <https://doi.org/10.1214/ss/1009213725>
- Chesson, P. (2018). Updates on mechanisms of maintenance of species diversity. *Journal of Ecology*, 106(5), 1773–1794. <https://doi.org/10.1111/1365-2745.13035>
- Course, K., & Nair, P. B. (2023). State estimation of a physical system with unknown governing equations. *Nature*, 622(7982), 261–267. <https://doi.org/10.1038/s41586-023-06574-8>
- Coyte, K. Z., Rao, C., Rakoff-Nahoum, S., & Foster, K. R. (2021). Ecological rules for the assembly of microbiome communities. *PLoS Biology*, 19(2), e3001116. <https://doi.org/10.1371/journal.pbio.3001116>
- Cranmer, M. (2023). Interpretable machine learning for science with PySR and SymbolicRegression.Jl. *arXiv* <https://doi.org/10.48550/arXiv.2305.01582>
- Cui, W., Marsland, R., & Mehta, P. (2024). Les Houches lectures on community ecology: From niche theory to statistical mechanics. *arXiv* <https://doi.org/10.48550/arXiv.2403.05497>
- De Jong, H., Casagrande, S., Giordano, N., Cinquemani, E., Ropers, D., Geiselmann, J., & Gouzé, J.-L. (2017). Mathematical modelling of microbes: Metabolism, gene expression and growth. *Journal of the Royal Society Interface*, 14(136), 20170502. <https://doi.org/10.1098/rsif.2017.0502>
- Gaucel, S., Keijzer, M., Lutton, E., & Tonda, A. (2014). Learning dynamical systems using standard symbolic regression. In *Genetic programming* (pp. 25–36). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-662-44303-3_3
- Ghenu, A.-H., Marrec, L., & Bank, C. (2024). Challenges and pitfalls of inferring microbial growth rates from lab cultures. *Frontiers in Ecology and Evolution*, 11, 1313500. <https://doi.org/10.3389/fevo.2023.1313500>
- Hatton, I. A., Mazzarisi, O., Altieri, A., & Smerlak, M. (2024). Diversity begets stability: Sublinear growth and competitive coexistence across ecosystems. *Science*, 383, eadg8488. <https://doi.org/10.1126/science.adg8488>
- Held, N. A., Krishna, A., Crippa, D., Battaje, R. R., Devaux, A. J., Dragan, A., & Manhart, M. (2024). Nutrient colimitation is a quantitative, dynamic property of microbial populations. *Proceedings. National Academy of Sciences. United States of America*, 121(52), e2400304121. <https://doi.org/10.1073/pnas.2400304121>
- Held, N. A., & Manhart, M. (2024). Are microbes colimited by multiple resources? *Current Opinion in Microbiology*, 80, 102509. <https://doi.org/10.1016/j.mib.2024.102509>
- Ho, P. Y., Good, B. H., & Huang, K. C. (2022). Competition for fluctuating resources reproduces statistics of species abundance over time across wide-ranging microbiotas. *eLife*, 11, e75168.
- Ho, P.-Y., Nguyen, T. H., Sanchez, J. M., DeFelice, B. C., & Huang, K. C. (2024). Resource competition predicts assembly of gut bacterial communities in vitro. *Nature Microbiology*, 9(4), 1036–1048. <https://doi.org/10.1038/s41564-024-01625-w>
- Holling, C. S. (1959a). The components of predation as revealed by a study of small-mammal predation of the European pine sawfly. *The Canadian Entomologist*, 91(5), 293–320. <https://doi.org/10.4039/ent91293-5>
- Holling, C. S. (1959b). Some characteristics of simple types of predation and parasitism. *The Canadian Entomologist*, 91(7), 385–398. <https://doi.org/10.4039/ent91385-7>
- Hsin, J., Agarwal, S., Thorpe, A., Sentis, L., & Fridovich-Keil, D. (2023). Symbolic regression on sparse and noisy data with gaussian processes. *arXiv* <https://doi.org/10.48550/arXiv.2309.11076>
- Huang, L. (2011). A new mechanistic growth model for simultaneous determination of lag phase duration and exponential growth rate and a new Bělehrádek-type model for evaluating the effect of temperature on growth rate. *Food Microbiology*, 28(4), 770–776. <https://doi.org/10.1016/j.fm.2010.05.019>
- Hutchinson, G. E. (1961). The paradox of the plankton. *The American Naturalist*, 95(882), 137–145. <https://doi.org/10.1086/282171>
- Ishizawa, H., Tashiro, Y., Inoue, D., Ike, M., & Futamata, H. (2024). Learning beyond-pairwise interactions enables the bottom-up prediction of microbial community structure. *Proceedings of the National Academy of Sciences of the United States of America*, 121(7), e2312396121. <https://doi.org/10.1073/pnas.2312396121>
- Kremling, A., Geiselmann, J., Ropers, D., & de Jong, H. (2018). An ensemble of mathematical models showing diauxic growth behaviour. *BMC Systems Biology*, 12, 82. <https://doi.org/10.1186/s12918-018-0604-8>
- Kronberger, G., Burlacu, B., Kommenda, M., Winkler, S. M., & Affenzeller, M. (2025). *Symbolic regression*. Chapman and Hall. <https://doi.org/10.1201/9781315166407>

- Lässig, M., Mustonen, V., & Nourmohammad, A. (2023). Steering and controlling evolution—From bioengineering to fighting pathogens. *Nature Reviews Genetics*, 24(12), 851–867. <https://doi.org/10.1038/s41576-023-00623-8>
- Lee, H., Bloxham, B., & Gore, J. (2023). Resource competition can explain simplicity in microbial community assembly. *Proceedings of the National Academy of Sciences of the United States of America*, 120(35), e2212113120. <https://doi.org/10.1073/pnas.2212113120>
- Letten, A. D. (2022). Rescomp: An R package for defining, simulating and visualizing ODE models of consumer-resource interactions. *BioRxiv* <https://doi.org/10.1101/2022.01.11.475574>
- Litchman, E., Edwards, K. F., & Klausmeier, C. A. (2015). Microbial resource utilization traits and trade-offs: Implications for community structure, functioning, and biogeochemical impacts at present and in the future. *Frontiers in Microbiology*, 6, 254. <https://doi.org/10.3389/fmicb.2015.00254>
- Lo Grasso, A., Fort, A., Mahdizadeh, F. F., Magnani, A., & Mocenni, C. (2023). Generalized logistic model of bacterial growth. *Mathematical and Computer Modelling of Dynamical Systems*, 29(1), 169–185. <https://doi.org/10.1080/13873954.2023.2236681>
- MacArthur, R., & Levins, R. (1967). The limiting similarity, convergence, and divergence of coexisting species. *The American Naturalist*, 101(921), 377–385. <https://doi.org/10.1086/282505>
- Martin, B. T., Munch, S. B., & Hein, A. M. (2018). Reverse-engineering ecological theory from data. *Proceedings of the Royal Society B: Biological Sciences*, 285(1878), 20180422. <https://doi.org/10.1098/rspb.2018.0422>
- Meroz, N., Livny, T., & Friedman, J. (2024). Quantifying microbial interactions: Concepts, caveats, and applications. *Current Opinion in Microbiology*, 80, 102511. <https://doi.org/10.1016/j.mib.2024.102511>
- Meszéna, G., Gyllenberg, M., Pásztor, L., & Metz, J. A. J. (2006). Competitive exclusion and limiting similarity: A unified theory. *Theoretical Population Biology*, 69(1), 68–87. <https://doi.org/10.1016/j.tpb.2005.07.001>
- Michaelis, L., & Menten, M. L. (1913). Die Kinetik der Invertinwirkung. *Biochemische Zeitschrift*, 49, 333–369.
- Miguel Trabajo, T., Guex, I., Dubey, M., Sarton-Lohéac, E., Todorov, H., Richard, X., Mazza, C., & van der Meer, J. R. (2024). Inferring bacterial interspecific interactions from microcolony growth expansion. *microLife*, 5, uqae020. <https://doi.org/10.1093/femsm/luqae020>
- Monod, J. (1949). The growth of bacterial cultures. *Annual Review of Microbiology*, 3(1), 371–394. <https://doi.org/10.1146/annurev.mi.03.100149.002103>
- Pacheco, A. R., Osborne, M. L., & Segrè, D. (2021). Non-additive microbial community responses to environmental complexity. *Nature Communications*, 12(1), 2365. <https://doi.org/10.1038/s41467-021-22426-3>
- Pastore, A. I., Barabás, G., Bimler, M. D., Mayfield, M. M., & Miller, T. E. (2021). The evolution of niche overlap and competitive differences. *Nature Ecology & Evolution*, 5(3), 330–337. <https://doi.org/10.1038/s41559-020-01383-y>
- Peleg, M., & Corradini, M. G. (2011). Microbial growth curves: What the models tell us and what they cannot. *Critical Reviews in Food Science and Nutrition*, 51(10), 917–945. <https://doi.org/10.1080/10408398.2011.570463>
- Perrin, E., Ghini, V., Giovannini, M., Di Patti, F., Cardazzo, B., Carraro, L., Fagorzi, C., Turano, P., Fani, R., & Fondi, M. (2020). Diauxie and co-utilization of carbon sources can coexist during bacterial growth in nutritionally complex environments. *Nature Communications*, 11(1), 3135. <https://doi.org/10.1038/s41467-020-16872-8>
- Picot, A., Shibasaki, S., Meacock, O. J., & Mitri, S. (2023). Microbial interactions in theory and practice: When are measurements compatible with models? *Current Opinion in Microbiology*, 75, 102354. <https://doi.org/10.1016/j.mib.2023.102354>
- Posfai, A., Taillefumier, T., & Wingreen, N. S. (2017). Metabolic trade-offs promote diversity in a model ecosystem. *Physical Review Letters*, 118(2), 028103. <https://doi.org/10.1103/physrevlett.118.028103>
- Radwan, Y. A., Kronberger, G., & Winkler, S. (2025). A comparison of recent algorithms for symbolic regression to genetic programming. In A. Quesada-Arencibia, M. Affenzeller, & R. Moreno-Díaz (Eds.), *Computer aided systems theory—EUROCAST 2024, number 15172 in lecture notes in computer science* (pp. 157–171). Springer Nature. https://doi.org/10.1007/978-3-031-82949-9_15
- Ram, Y., Dellus-Gur, E., Bibi, M., Karkare, K., Obolski, U., Feldman, M. W., Cooper, T. F., Berman, J., & Hadany, L. (2019). Predicting microbial growth in a mixed culture from growth curve data. *Proceedings of the National Academy of Sciences*, 116(29), 14698–14707. <https://doi.org/10.1073/pnas.1902217116>
- Ram, Y., Obolski, U., Feldman, M. W., Berman, J., & Hadany, L. (2019). Reply to Balsa-Canto et al.: Growth models are applicable to growth data, not to stationary-phase data. *Proceedings of the National Academy of Sciences of the United States of America*, 117(2), 814–815. <https://doi.org/10.1073/pnas.1917758117>
- Regalado, D. (2021). Approximate analytic solution to the three species Lotka-Volterra differential equation model. *Recoletos Multidisciplinary Research Journal*, 9(2), 123–128. <https://doi.org/10.32871/rmrj.2109.02.09>
- Richardson, D. (1969). Some undecidable problems involving elementary functions of a real variable. *Journal of Symbolic Logic*, 33(4), 514–520. <https://doi.org/10.2307/2271358>
- Rozas Garcia, E., Crumpton, M. J., & Galla, T. (2023). Niche overlap and Hopfield-like interactions in generalized random Lotka-Volterra systems. *Physical Review E*, 108(3), 034120. <https://doi.org/10.1103/physreve.108.034120>
- Saito, M. A., Goepfert, T. J., & Ritt, J. T. (2008). Some thoughts on the concept of colimitation: Three definitions and the importance of bioavailability. *Limnology and Oceanography*, 53(1), 276–290. <https://doi.org/10.4319/lo.2008.53.1.0276>
- Shumaylov, Z., Zaika, P., Scholl, P., Kutyniok, G., Horesh, L., & Schönlieb, C.-B. (2025). When is a system discoverable from data? Discovery requires chaos. *arXiv* <https://doi.org/10.48550/ARXIV.2511.08860>
- Sun, T. A., & Kičiatovas, D. (2026). Symbolic regression Monod models. *Zenodo* <https://doi.org/10.5281/zenodo.20119884>
- Takano, S., Vila, J. C. C., Miyazaki, R., Sánchez, Á., & Bajić, D. (2023). The architecture of metabolic networks constrains the evolution of microbial resource hierarchies. *Molecular Biology and Evolution*, 40(9), msad187. <https://doi.org/10.1093/molbev/msad187>
- Tang, J., & Riley, W. J. (2021). Finding Liebig's law of the minimum. *Ecological Applications*, 31(8), e02458. <https://doi.org/10.1002/eap.2458>
- Turchin, P. (2001). Does population ecology have general laws? *Oikos*, 94(1), 17–26. <https://doi.org/10.1034/j.1600-0706.2001.11310.x>
- Wong, C. (2024). Antibiotic resistance is a growing threat: Is climate change making it worse? *Nature*. <https://doi.org/10.1038/d41586-023-04077-0>
- Wood, S. N. (2025). Generalized additive models. *Annual Review of Statistics and Its Application*, 12(1), 497–526. <https://doi.org/10.1146/annurev-statistics-112723-034249>
- Yeoman, K., Fahner, B., Lea-Smith, D., & Clarke, T. (2020). Microbial biotechnology and agriculture. In K. Yeoman, B. Fahner, D. Lea-Smith, & T. Clarke (Eds.), *Microbial biotechnology*. Oxford University Press. <https://doi.org/10.1093/hesc/9780198822813.003.0008>
- Zwietering, M. H., Jongenburger, I., Rombouts, F. M., & van't Riet, K. (1990). Modeling of the bacterial growth curve. *Applied and Environmental Microbiology*, 56(6), 1875–1881. <https://doi.org/10.1128/aem.56.6.1875-1881.1990>

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

Table S1. Bacterial species used in the experiments and their HAMBI code. The last two did not generally reach stationary phase and were not considered for further analyses.

Figure S1. Overview of the experimental data, showing the optical density of bacterial cultures and the per capita growth rate as a function of time.

Figure S2. Black-box Random Forest regression predicting the per-capita growth rate from different feature sets shows that the cumulative population gain predictor N outperforms the population size predictor N in simulated data (multiplicative model with five resource dimensions).

Figure S3. Performance of a Random Forest regressor, given by the coefficient of determination R^2 , using different feature sets to predict the per capita growth rate in data obtained: from experiments (Rs0); from simulations using an α -adjusted model with various levels of resource dimensionality (Ms1–Ms5 and As1–As5). Each group comprises 16 species and shows, for each, the algorithm performance on the training set (dots), its performance on the testing set (lighter shade), the importance of the population size feature (darker shade), and its standard deviation (white segment).

Figure S4. Fitting performance of models proposed by the symbolic regressor from growth data simulated using a dynamical Monod model relying on a bidimensional resource space, assuming either that there source dimensions were non-substitutable (left) or substitutable (right). White: unidimensional Monod model; blue: non-substitutable bidimensional Monod model; red: substitutable bidimensional Monod model; orange: higher-dimensional Monod model; grey: not biologically interpretable. Circle: no template; triangle pointing down: Template (7); triangle pointing up: Template (8).

How to cite this article: Sun, T. A., Kičiatovas, D., Aapalampi, I.-K., Kuosmanen, T., Hiltunen, T., & Mustonen, V. (2026). Discovering data-driven microbial growth models with symbolic regression. *Methods in Ecology and Evolution*, 17, 2135–2148. <https://doi.org/10.1111/2041-210x.70335>