Simulation and Control of High-Dimensional Dynamical Systems Using Artificial Neural Networks

Lucas Böttcher
Dept. of Computational Science and Philosophy, Frankfurt School
Laboratory for Systems Medicine, University of Florida

Figure courtesy: Alina Grubnyak

hessian.AI

Outline

(1) System Identification using Neural Networks

(2) Control of Dynamical Systems using Neural Networks

(3) Control of Microbiome Dynamics

(4) Summary

System Identification using Neural Networks

Learning a Dynamical System

<u>Basic idea:</u> Represent the right-hand side of an ordinary differential equation by an artificial neural network (or another universal function approximator) with parameters $\theta \in \mathbb{R}^N$.

$$\boxed{\dot{x}(t) = f_\theta(t, x(t))}$$

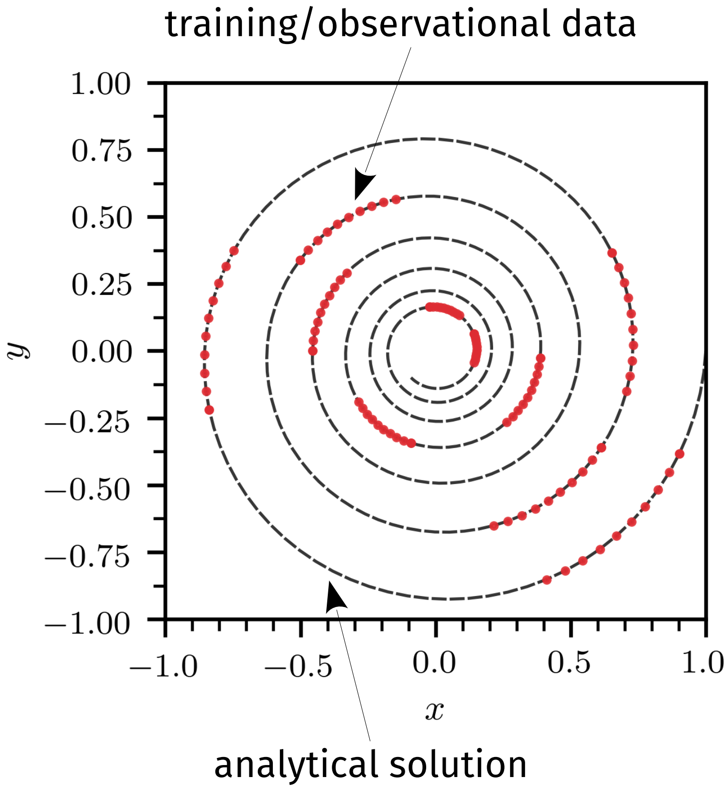

Neural Networks, Vol. 4, pp. 251–257, 1991
Printed in the USA.

0893-6080/91 $3.00 + .00

ORIGINAL CONTRIBUTION

Approximation Capabilities of Multilayer Feedforward Networks

KURT HORNIK

Technische Universität Wien, Vienna, Austria

(*Received* 30 *January* 1990; *revised and accepted* 25 *October* 1990)

Abstract—*We show that standard multilayer feedforward networks with as few as a single hidden layer and arbitrary bounded and nonconstant activation function are universal approximators with respect to* $L^p(\mu)$ *performance criteria, for arbitrary finite input environment measures* μ*, provided only that sufficiently many hidden units are available. If the activation function is continuous, bounded and nonconstant, then continuous mappings can be learned uniformly over compact input sets. We also give very general conditions ensuring that networks with sufficiently smooth activation functions are capable of arbitrarily accurate approximation to a function and its derivatives.*

Learning a Dynamical System

<u>Example:</u> We consider the following dynamical system

$$\begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = \begin{pmatrix} -0.05 & 1 \\ -1 & -0.05 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}$$

with initial condition $(x(0), y(0))^\top = (1, 0)^\top$

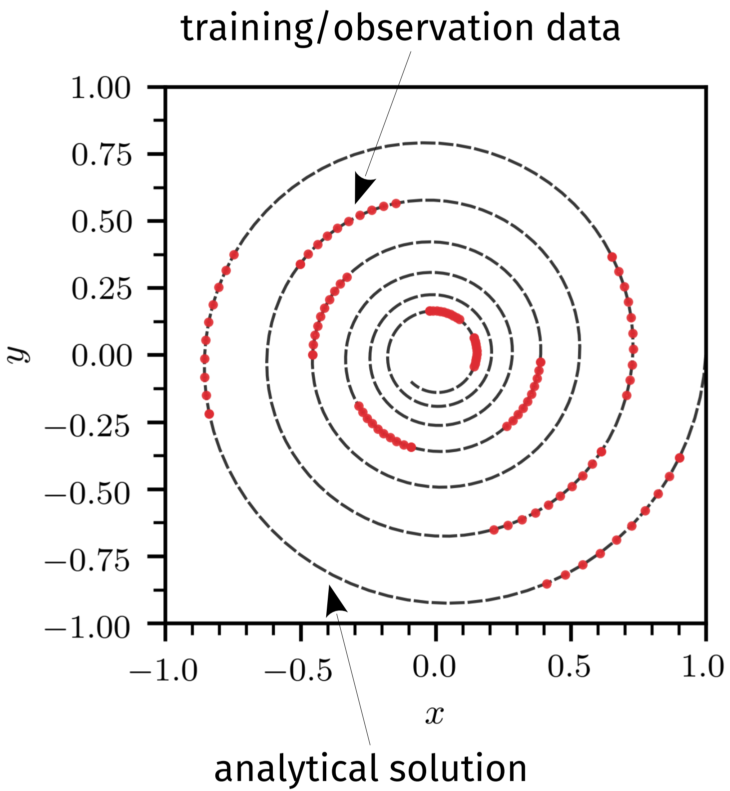

Learning a Dynamical System

$$\begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = \begin{pmatrix} -0.05 & 1 \\ -1 & -0.05 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}$$

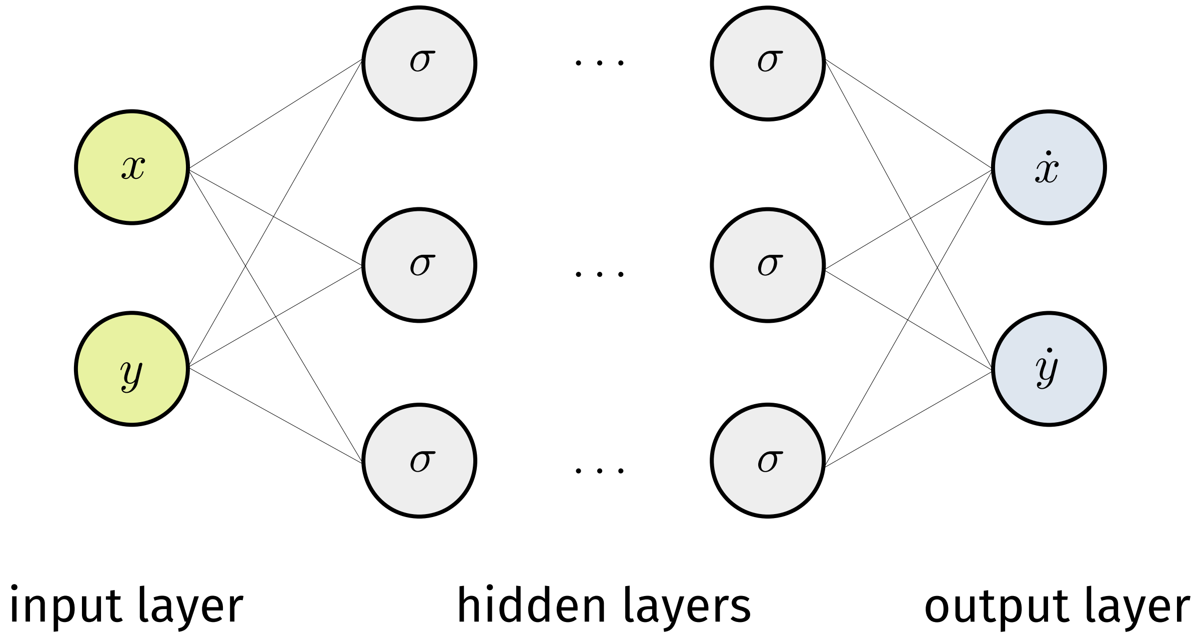

Learning a Dynamical System

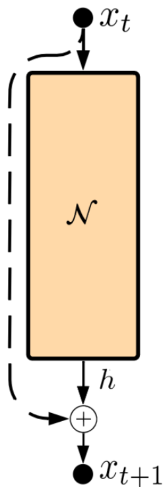

Euler cell

Solve $\dot{x}(t) = f_\theta(t, x(t))$ using a numerical integration method (e.g., Euler's method):

$$x_{t+1} = x_t + h f_\theta(t, x_t)$$

Figure courtesy: Krishnapriyan, Aditi S., et al. "Learning continuous models for continuous physics." arXiv preprint arXiv:2202.08494 (2022).

Learning a Dynamical System

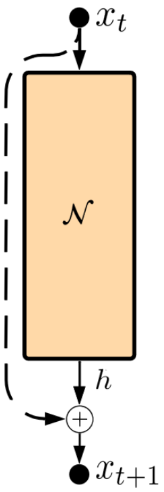

Euler cell

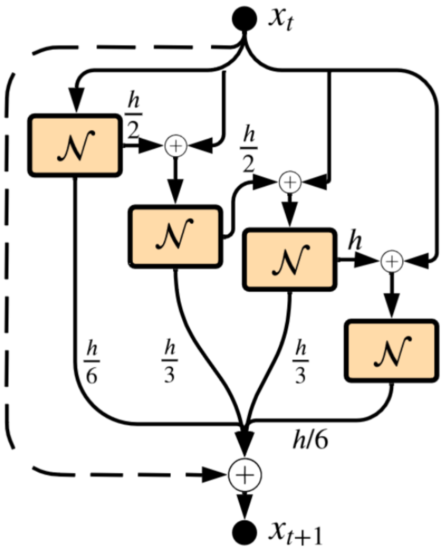

RK4 cell

Figure courtesy: Krishnapriyan, Aditi S., et al. "Learning continuous models for continuous physics." arXiv preprint arXiv:2202.08494 (2022).

Learning a Dynamical System

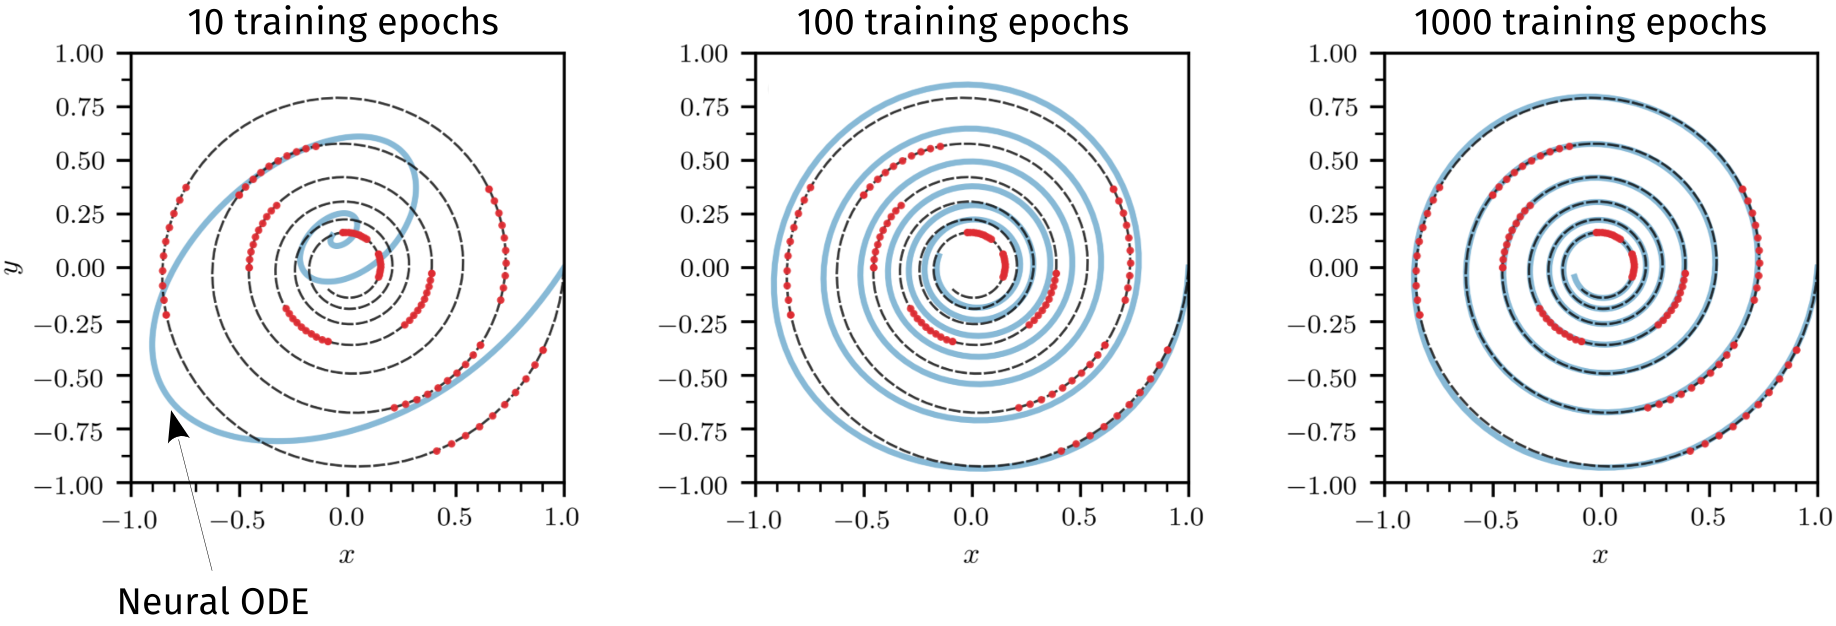

Runge–Kutta Neural Network for Identification of Dynamical Systems in High Accuracy

Yi-Jen Wang and Chin-Teng Lin, *Member, IEEE*

Abstract— **This paper proposes the Runge–Kutta neural networks (RKNN's) for identification of unknown dynamical systems described by ordinary differential equations (i.e., ordinary differential equation or ODE systems) in high accuracy. These networks are constructed according to the Runge–Kutta approximation method. The main attraction of the RKNN's is that they precisely estimate the changing rates of system states (i.e., the right-hand side of the ODE $\dot{x} = f(x)$) directly in their subnetworks based on the space-domain interpolation within one sampling interval such that they can do long-term prediction of system state trajectories. We show theoretically the superior generalization and long-term prediction capability of the RKNN's over the normal neural networks. Two types of learning algorithms are investigated for the RKNN's, gradient- and nonlinear recursive least-squares-based algorithms. Convergence analysis of the learning algorithms is done theoretically. Computer simulations demonstrate the proved properties of the RKNN's.**

fixed regular sampling rate). This is not the nature of an ODE system. Although a high-order discretization is more accurate than the first-order discretization, the resulting ordinary difference equations of the former are usually complex and intractable.

In this paper, we present a class of feedforward neural networks called Runge–Kutta neural networks (RKNN's) for precisely modeling an ODE system in the form of $\dot{\mathbf{x}} = f(\mathbf{x})$ with an unknown f, where the state vector $\mathbf{x}$ is assumed to be measured noise-free. The neural approximation of f is used in the well-known Runge–Kutta integration formulas [10] to obtain an approximation of $\mathbf{x}$. With the designed network structure and proposed learning schemes, the RKNN's perform high-order discretization of unknown ODE systems *implicitly* (i.e., internally in the network) without the aforementioned

Neural Ordinary Differential Equations

Ricky T. Q. Chen*, Yulia Rubanova*, Jesse Bettencourt*, David Duvenaud
University of Toronto, Vector Institute
{rtqichen, rubanova, jessebett, duvenaud}@cs.toronto.edu

Abstract

We introduce a new family of deep neural network models. Instead of specifying a discrete sequence of hidden layers, we parameterize the derivative of the hidden state using a neural network. The output of the network is computed using a black-box differential equation solver. These continuous-depth models have constant

C. Runge. Über die numerische Auflösung von Differentialgleichungen. *Mathematische Annalen*, 46: 167–178, 1895.

W. Kutta. Beitrag zur näherungsweisen Integration totaler Differentialgleichungen. *Zeitschrift für Mathematik und Physik*, 46:435–453, 1901.

Chaos ARTICLE scitation.org/journal/cha

Interpretable polynomial neural ordinary differential equations

Cite as: Chaos **33**, 043101 (2023); doi: 10.1063/5.0130803
Submitted: 14 October 2022 · Accepted: 9 March 2023 ·
Published Online: 3 April 2023

View Online Export Citation CrossMark

Colby Fronk[1,a] and Linda Petzold[2,3,a]

AFFILIATIONS

[1] Department of Chemical Engineering, University of California, Santa Barbara, California 93106, USA
[2] Department of Mechanical Engineering, University of California, Santa Barbara, California 93106, USA
[3] Department of Computer Science, University of California, Santa Barbara, California 93106, USA

[a] **Authors to whom correspondence should be addressed:** colbyfronk@ucsb.edu and petzold@engineering.ucsb.edu

ABSTRACT

Neural networks have the ability to serve as universal function approximators, but they are not interpretable and do not generalize well outside of their training region. Both of these issues are problematic when trying to apply standard neural ordinary differential equations (ODEs) to dynamical systems. We introduce the polynomial neural ODE, which is a deep polynomial neural network inside of the neural ODE framework. We demonstrate the capability of polynomial neural ODEs to predict outside of the training region, as well as to perform direct symbolic regression without using additional tools such as SINDy.

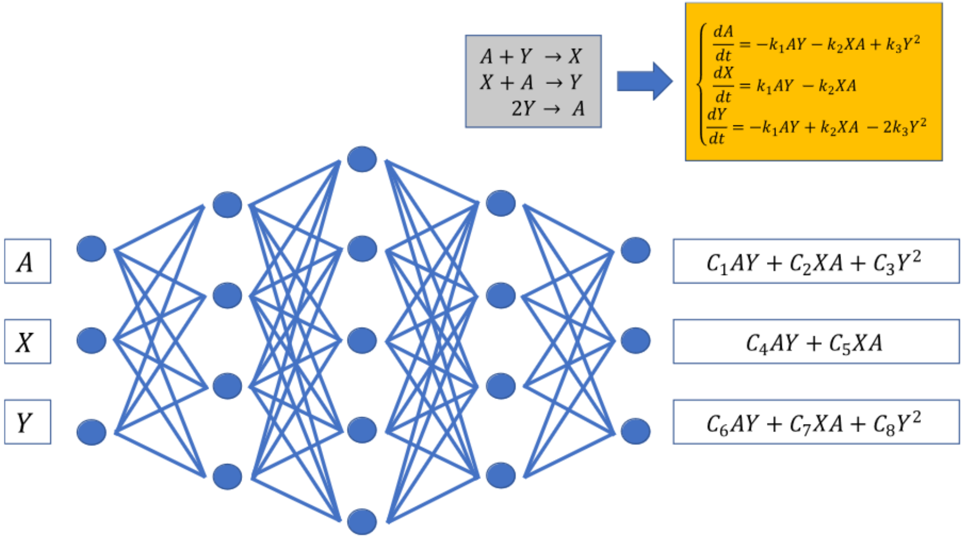

FIG. 1. Example polynomial neural ODE for a chemical reaction system with molecules A, X, and Y. The neural network outputs a polynomial transformation of the input, which are the concentrations of the chemical species.

Control of Dynamical Systems using Neural Networks

Traditional Control Theory

A dynamical system is said to be **controllable** if it can be steered from any initial state $\boldsymbol{x}_0$ to any target state **x*** in finite time T.

The goal is to minimize

$$J = \int_0^T L(\mathbf{x}(t'), \mathbf{u}(t'))\, \mathrm{d}t' + C(\mathbf{x}(T))$$

subject to the constraint

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t))$$

Traditional Control Theory

A dynamical system is said to be **controllable** if it can be steered from any initial state $\boldsymbol{x}_0$ to any target state **x*** in finite time T.

integrated cost The goal is to minimize **bequest value**

$$J = \int_0^T L(\mathbf{x}(t'), \mathbf{u}(t'))\,\mathrm{d}t' + C(\mathbf{x}(T))$$

subject to the constraint

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t))$$

Traditional Control Theory

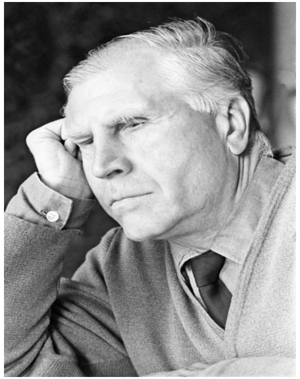

Lev Pontryagin
(1908–1988)

- Pontryagin's maximum principle (PMP)
- Solution of the Hamilton—Jacobi—Bellman (HJB) equation

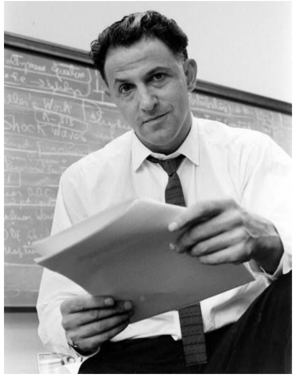

Richard Bellman
(1920–1984)

Traditional Control Theory

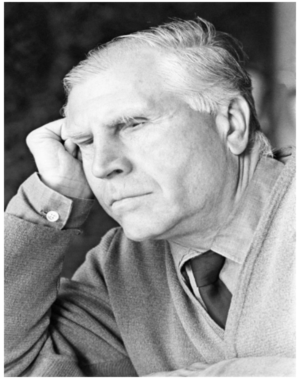

Lev Pontryagin
(1908–1988)

- Pontryagin's maximum principle (PMP)

- Solution of the Hamilton—Jacobi—Bellman (HJB) equation

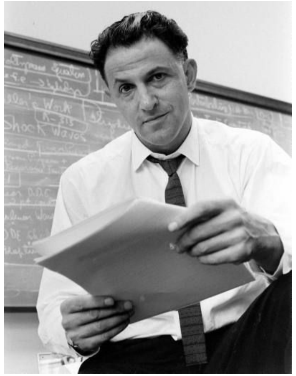

Richard Bellman
(1920–1984)

How can we control high-dimensional and potentially stochastic dynamical systems?

Lunar Lander

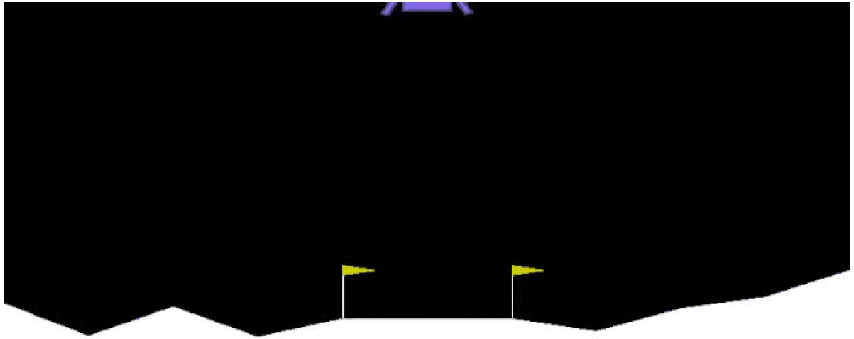

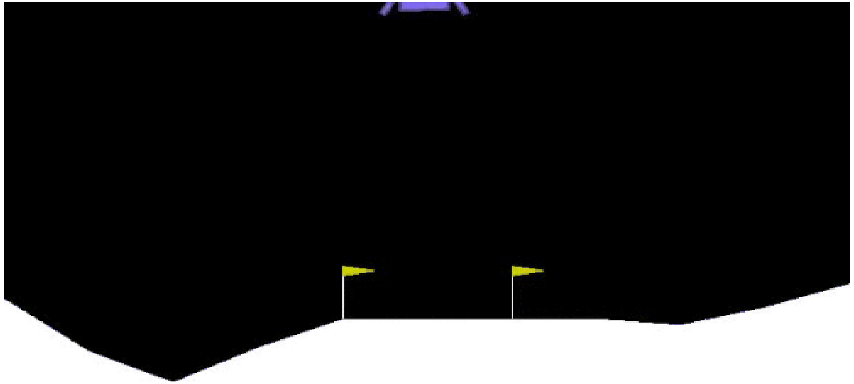

https://www.billyvreeland.com/portfolio/2017/10/13/solving-openai-gym-nm4yz

SpaceX

https://tenor.com/view/spacex-space-gif-11007965

Cart Pole

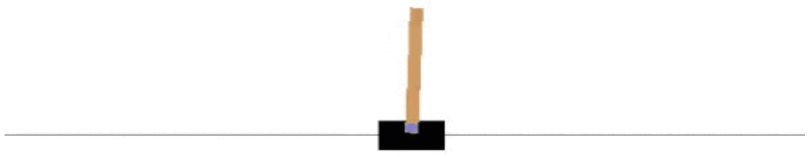

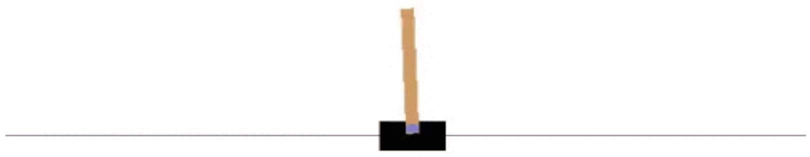

https://www.billyvreeland.com/portfolio/2017/10/13/solving-openai-gym-nm4yz

Robot Arm (Reinforcement Learning)

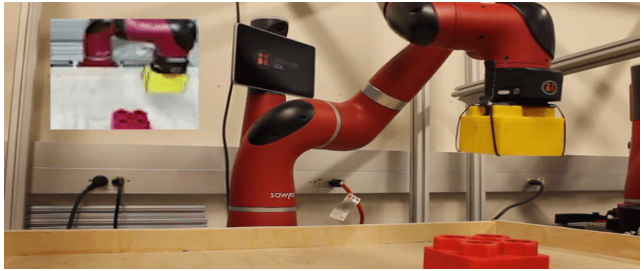

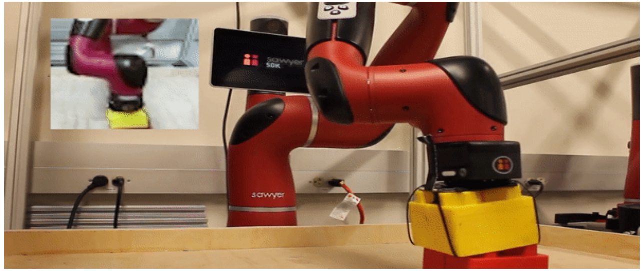

https://bair.berkeley.edu/blog/2019/05/20/solar/

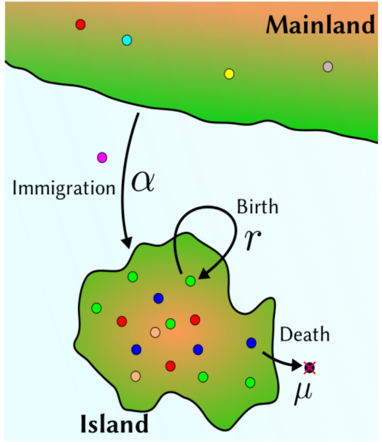

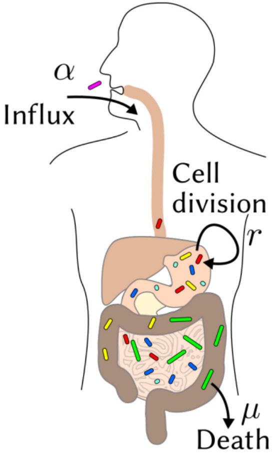

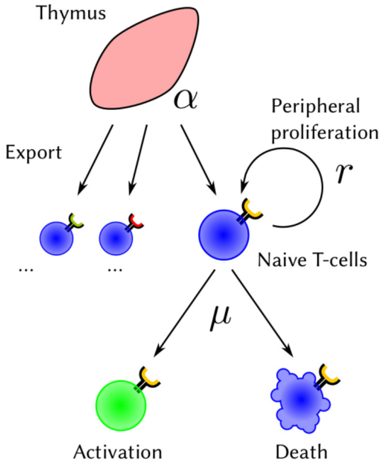

birth-death-immigration processes at different scales

Xu, S., Böttcher, L., & Chou, T. (2020). Diversity in biology: definitions, quantification and models. Physical Biology, 17(3), 031001.
Böttcher, L., Wald, S., & Chou, T. (2023). Mathematical Characterization of Private and Public Immune Receptor Sequences. Bulletin of Mathematical Biology, 85(10), 102.

Neural ODE Controller

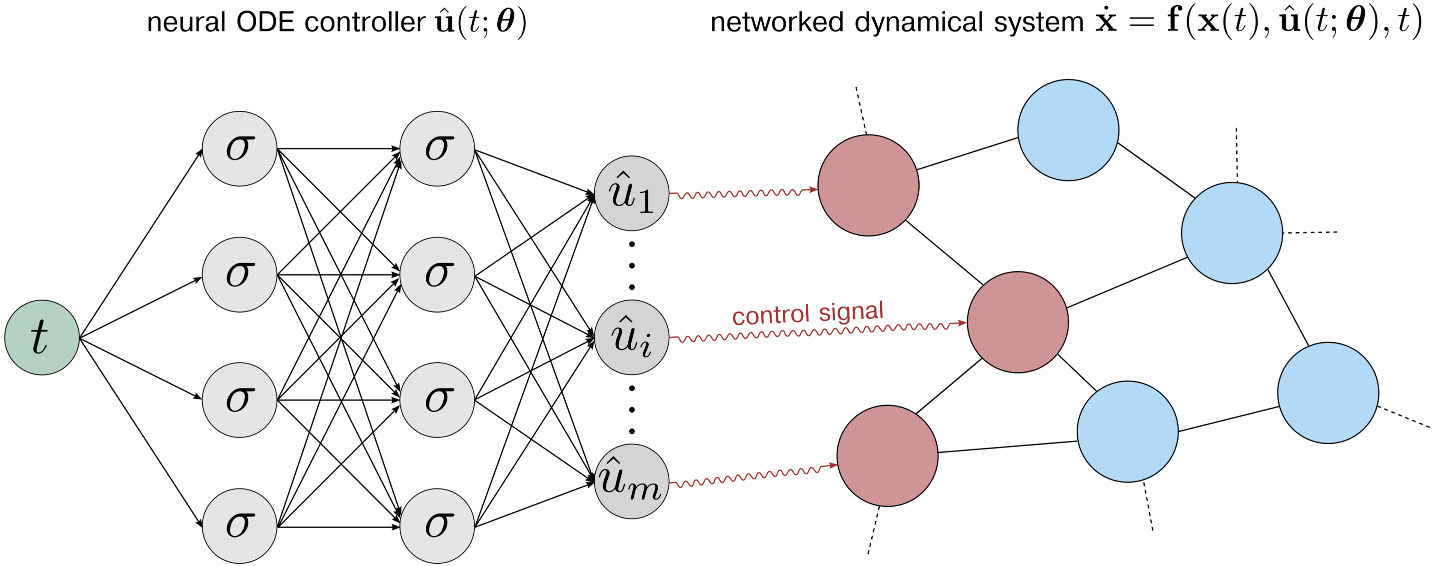

Böttcher, L., & Asikis, T. (2022). Near-optimal control of dynamical systems with neural ordinary differential equations. Machine Learning: Science and Technology, 3, 045004.

Neural ODE Controller

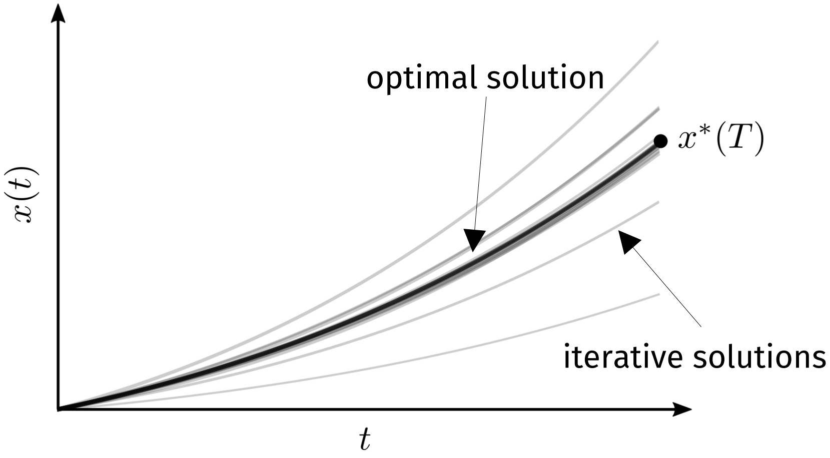

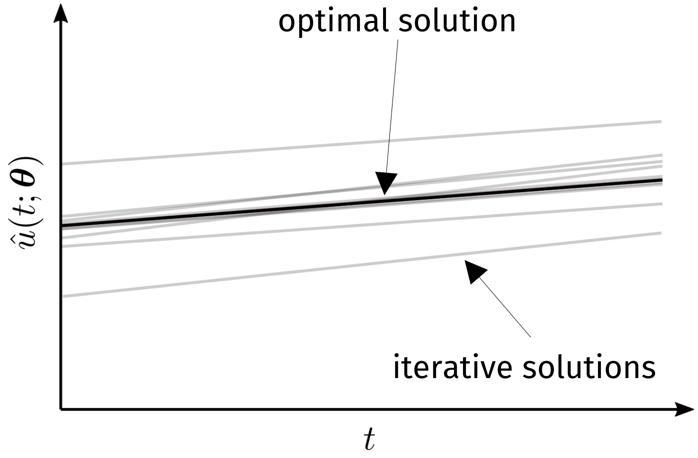

Böttcher, L., & Asikis, T. (2022). Near-optimal control of dynamical systems with neural ordinary differential equations. Machine Learning: Science and Technology, 3, 045004.

An Example

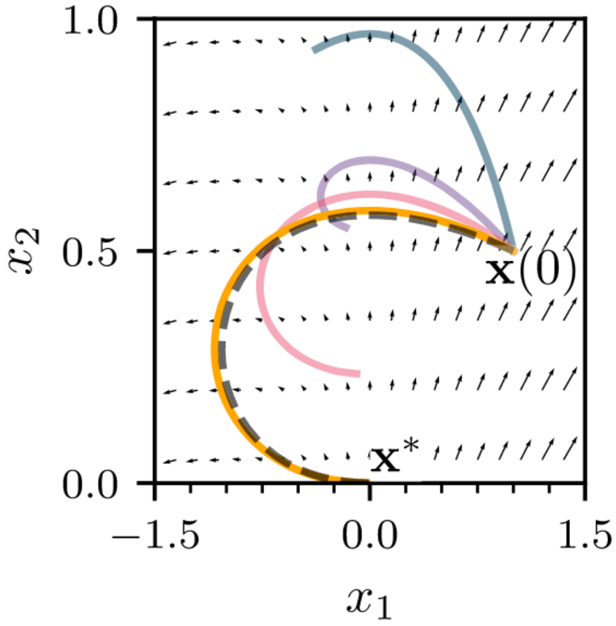

Linear dynamics

$$\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{u}$$

$$A = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 1 \\ 0 \end{pmatrix} .$$

Loss function

$$J(\mathbf{x}(T), \mathbf{x}) = \frac{1}{N} \|\mathbf{x}(T) - \mathbf{x}^*\|_2^2$$

Böttcher, L., Antulov-Fantulin, N., & Asikis, T. (2022). AI Pontryagin or how artificial neural networks learn to control dynamical systems. Nature Communications, 13, 333.

An Example

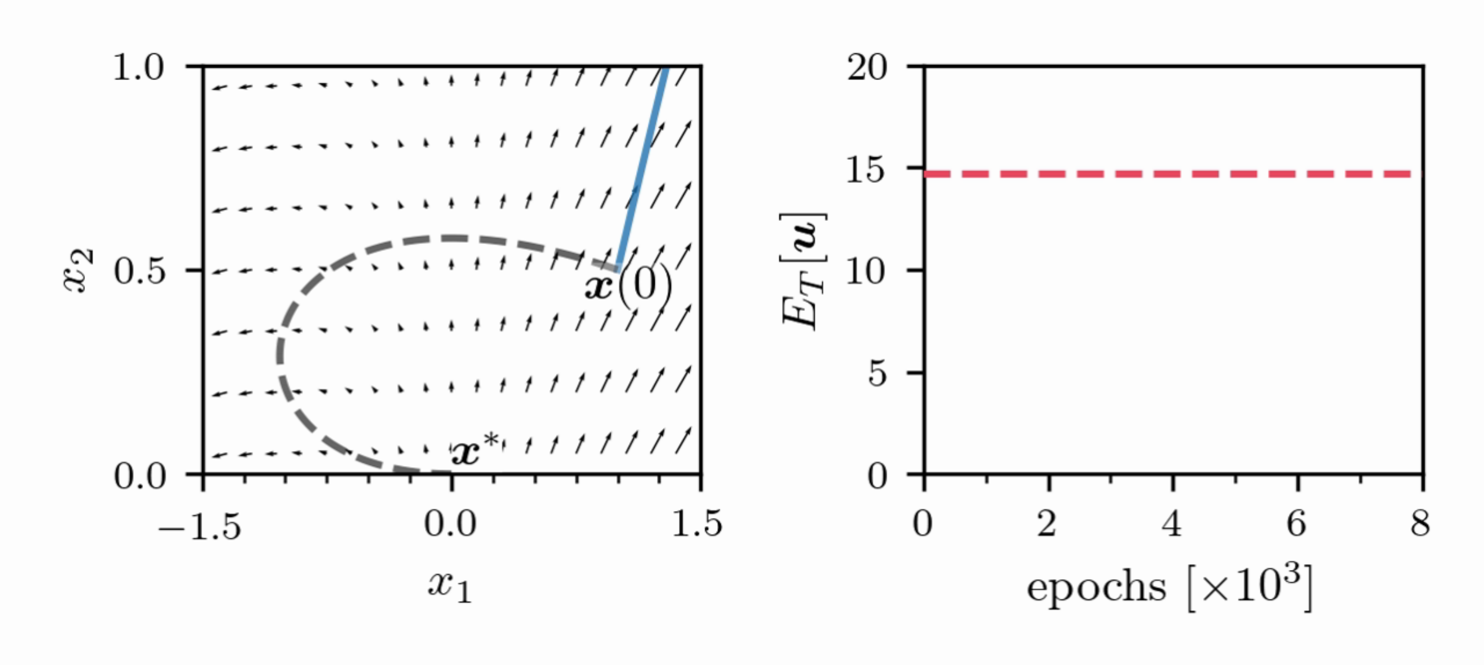

Böttcher, L., Antulov-Fantulin, N., & Asikis, T. (2022). AI Pontryagin or how artificial neural networks learn to control dynamical systems. Nature Communications, 13, 333.

An Example

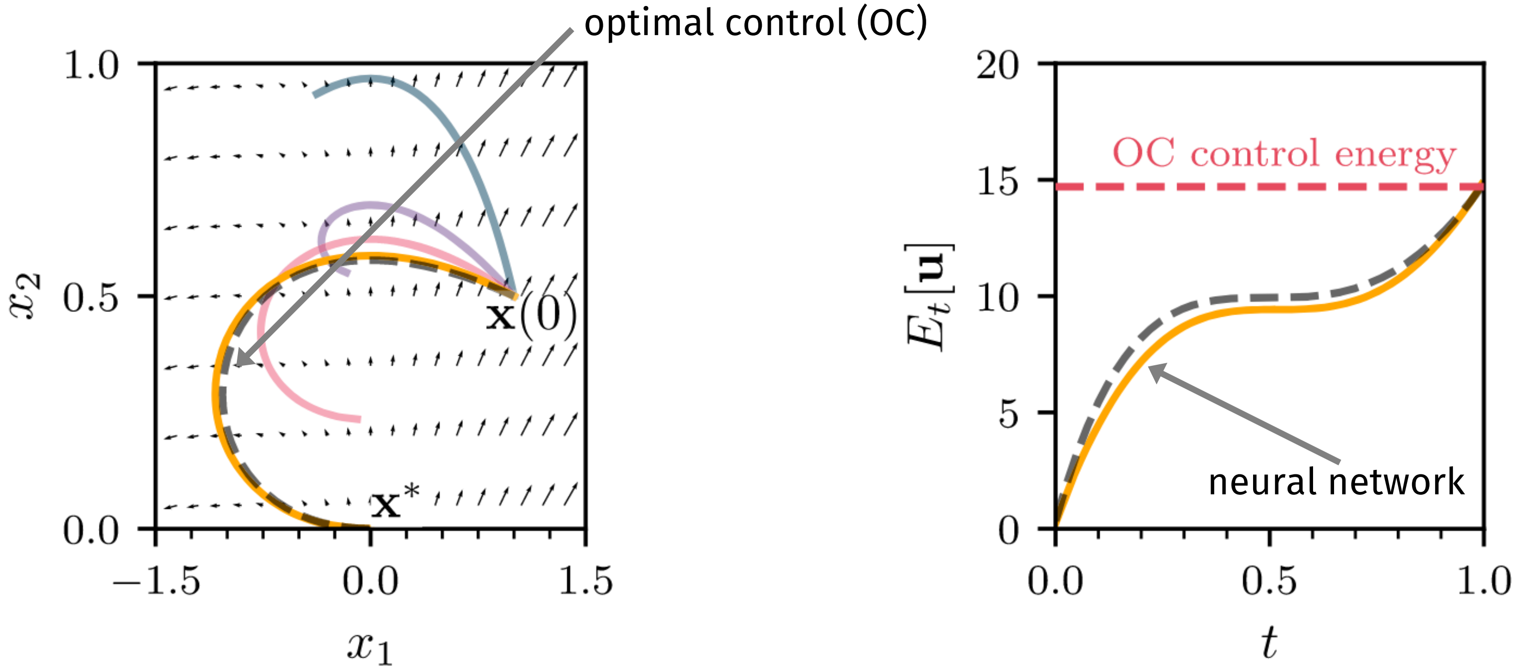

Böttcher, L., Antulov-Fantulin, N., & Asikis, T. (2022). AI Pontryagin or how artificial neural networks learn to control dynamical systems. Nature Communications, 13, 333.

Further Examples

PHILOSOPHICAL TRANSACTIONS A

royalsocietypublishing.org/journal/rsta

Research

Check for updates

Cite this article: Böttcher L, Fonseca LL, Laubenbacher RC. 2025 Control of medical digital twins with artificial neural networks. *Phil. Trans. R. Soc. A* **383**: 20240228.

Control of medical digital twins with artificial neural networks

Lucas Böttcher[1,2,†], Luis L. Fonseca[2,†] and Reinhard C. Laubenbacher[2]

[1]Department of Computational Science and Philosophy, Frankfurt School of Finance and Management, Frankfurt am Main 60322, Germany
[2]Department of Medicine, Laboratory for Systems Medicine, University of Florida, Gainesville, FL, USA

LB, 0000-0003-1700-1897; LLF, 0000-0002-7902-742X; RCL, 0000-0002-9143-9451

informs
https://pubsonline.informs.org/journal/ijoc

INFORMS JOURNAL ON COMPUTING
Articles in Advance, pp. 1–21
ISSN 1091-9856 (print), ISSN 1526-5528 (online)

Control of Dual-Sourcing Inventory Systems Using Recurrent Neural Networks

Lucas Böttcher,[a,*] Thomas Asikis,[b] Ioannis Fragkos[c]

[a]Department of Computational Science and Philosophy, Frankfurt School of Finance and Management, 60322 Frankfurt, Germany; [b]Game Theory, University of Zurich, 8092 Zurich, Switzerland; [c]Department of Technology and Operations Management, Rotterdam School of Management, Erasmus University Rotterdam, 3062 Rotterdam, Netherlands
*Corresponding author
Contact: l.boettcher@fs.de, https://orcid.org/0000-0003-1700-1897 (LB); thomas.asikis@uzh.ch, https://orcid.org/0000-0003-0163-4622 (TA); fragkos@rsm.nl, https://orcid.org/0000-0001-7654-2314 (IF)

nature COMMUNICATIONS

ARTICLE
https://doi.org/10.1038/s41467-021-27590-0 OPEN

Check for updates

AI Pontryagin or how artificial neural networks learn to control dynamical systems

Lucas Böttcher[1,2,4], Nino Antulov-Fantulin[3] & Thomas Asikis[3,4]

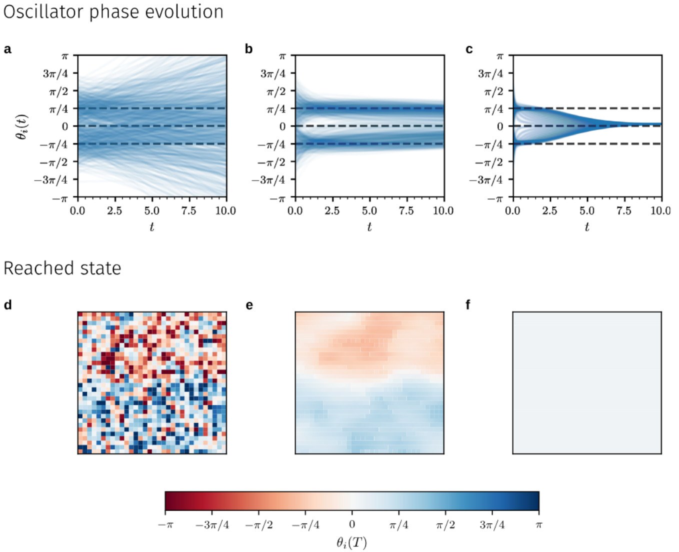

idinn: Inventory-Dynamics Control with Neural Networks

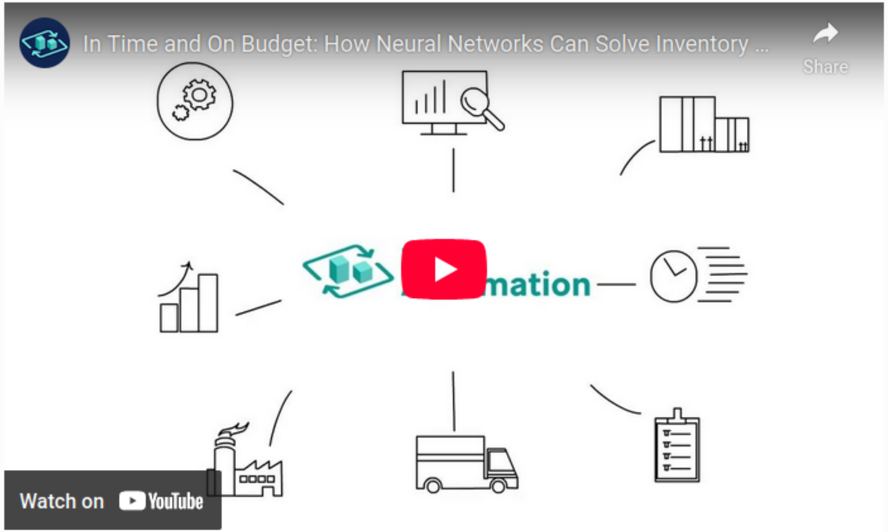

idinn implements inventory dynamics–informed neural network and other related controllers for solving single-sourcing and dual-sourcing problems. Controllers and inventory dynamics are implemented into customizable objects using PyTorch as backend to enable users to find the optimal controllers for the user-specified inventory systems.

Demo

For a quick demo, you can run our Streamlit app. The app allows you to interactively train and evaluate neural controllers for user-specified dual-sourcing systems.

Control of Microbiome Dynamics

A First Example: Intestinal Microbiota and Clindamycin

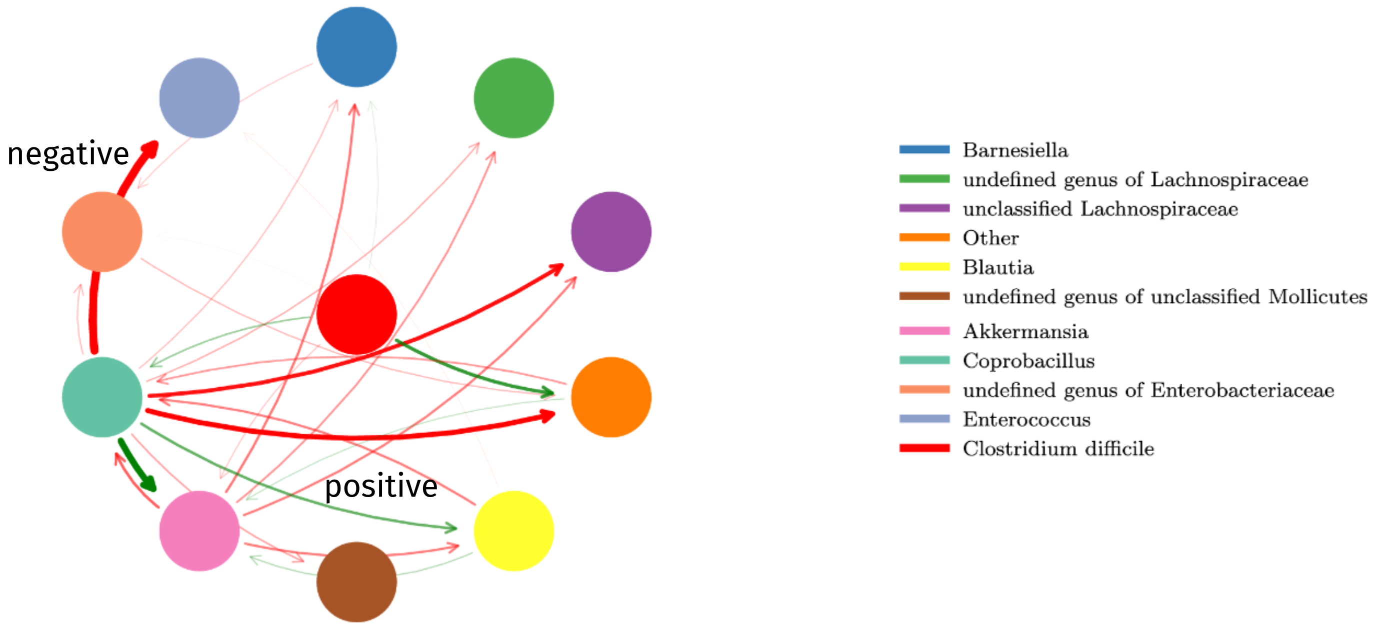

Jones, E. W., & Carlson, J. M. (2018). In silico analysis of antibiotic-induced Clostridium difficile infection: Remediation techniques and biological adaptations. PLOS Computational Biology, 14(2), e1006001.

A Generalized Lotka—Volterra Approach

intrinsic growth rate

antibiotic susceptibility

$$\dot{x}_i(t) = x_i(t) \left(b_i + \sum_{j=1}^{n} m_{ij} x_j(t) + \epsilon_i u(t) \right)$$

abundance of
microbial species *i*

interaction
coefficient

antibiotic
treatment

Jones, E. W., & Carlson, J. M. (2018). In silico analysis of antibiotic-induced Clostridium difficile infection: Remediation techniques and biological adaptations. PLOS Computational Biology, 14(2), e1006001.

A Generalized Lotka—Volterra Approach

intrinsic growth rate

antibiotic susceptibility

$$\dot{x}_i(t) = x_i(t)\left(b_i + \sum_{j=1}^{n} m_{ij}x_j(t) + \epsilon_i u(t)\right)$$

abundance of
microbial species *i*

interaction
coefficient

antibiotic
treatment

2^n potential equilibria

Jones, E. W., & Carlson, J. M. (2018). In silico analysis of antibiotic-induced Clostridium difficile infection: Remediation techniques and biological adaptations. PLOS Computational Biology, 14(2), e1006001.

A Generalized Lotka—Volterra Approach

intrinsic growth rate

antibiotic susceptibility

$$\dot{x}_i(t) = x_i(t)\left(b_i + \sum_{j=1}^{n} m_{ij}x_j(t) + \epsilon_i u(t;\theta)\right)$$

abundance of microbial species *i*

interaction coefficient

antibiotic treatment (neural ODE controller)

2^n potential equilibria

Jones, E. W., & Carlson, J. M. (2018). In silico analysis of antibiotic-induced Clostridium difficile infection: Remediation techniques and biological adaptations. PLOS Computational Biology, 14(2), e1006001.

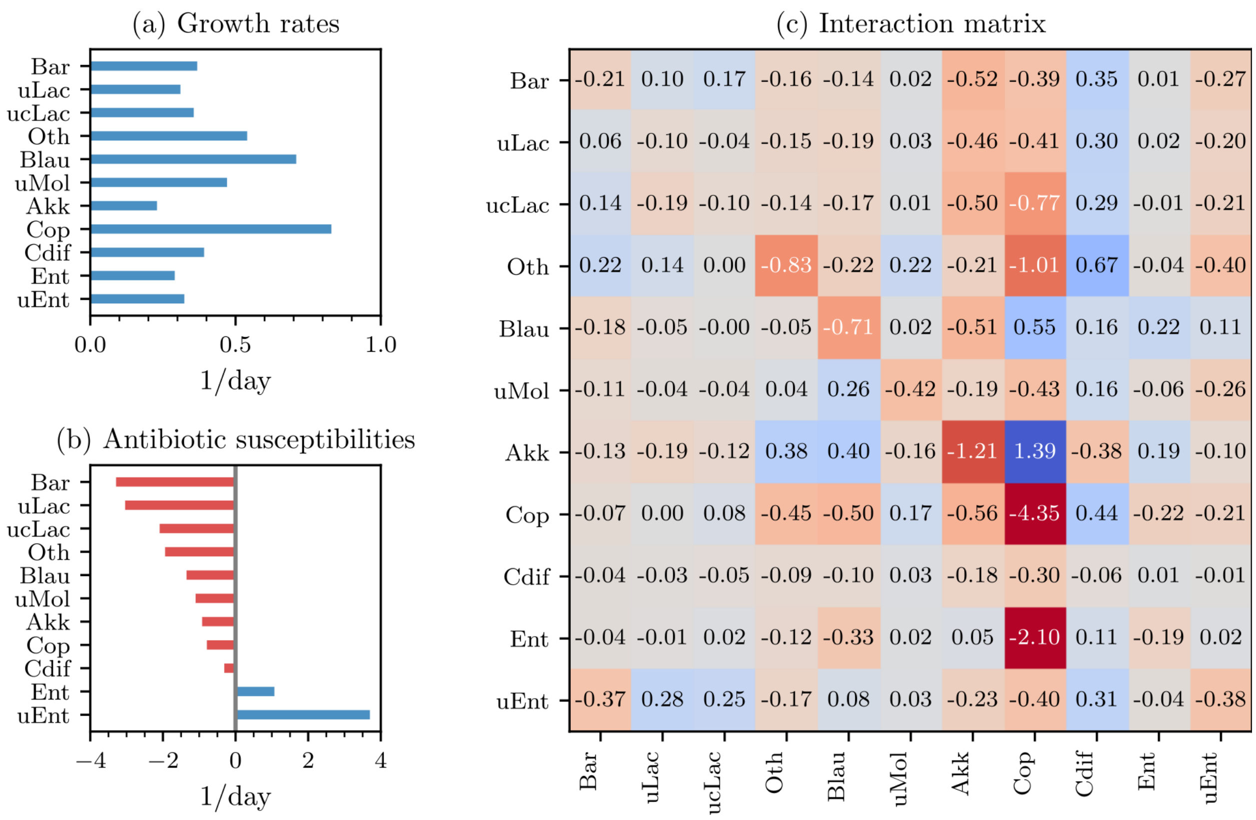

Buffie, Charlie G., et al. "Profound alterations of intestinal microbiota following a single dose of clindamycin results in sustained susceptibility to Clostridium difficile-induced colitis." Infection and Immunity 80.1 (2012): 62-73.
Stein, Richard R., et al. "Ecological modeling from time-series inference: insight into dynamics and stability of intestinal microbiota." PLOS Computational Biology 9.12 (2013): e1003388.

4

Navigation and control of outcomes in a generalized Lotka-Volterra model of the microbiome

Eric W. Jones, Parker Shankin-Clarke, and Jean M. Carlson

Department of Physics, University of California at Santa Barbara, Santa Barbara CA 93106, USA, ewj@physics.ucsb.edu

E. W. Jones, P. Shankin-Clarke, and J. M. Carlson. Navigation and control of outcomes in a generalized Lotka–Volterra model of the microbiome. In J. Kotas, editor, Advances in Nonlinear Biological Systems: Modeling and Optimal Control, volume 11 of AIMS Series on Applied Mathematics, pages 97–120. American Institute of Mathematical Sciences, Springfield, MO, USA, 2020.

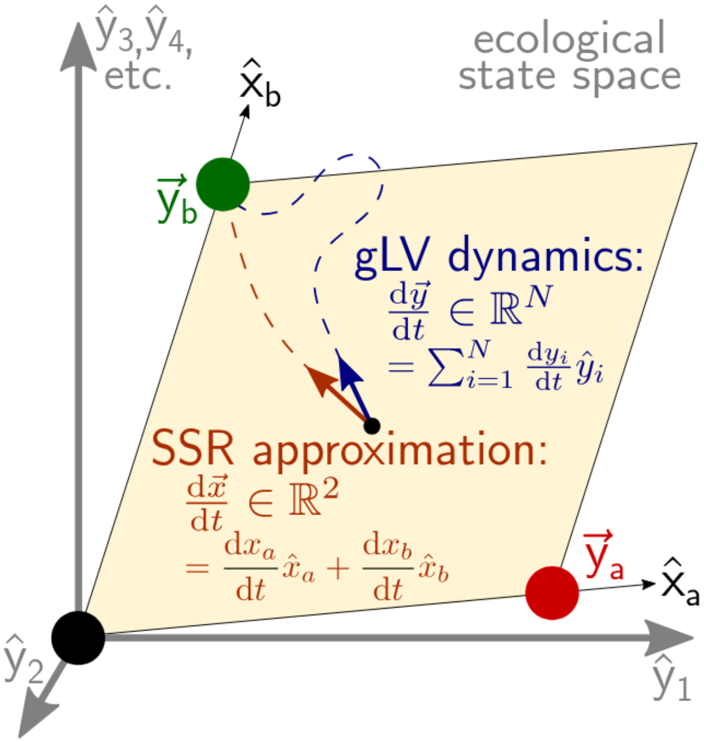

E. W. Jones, P. Shankin-Clarke, and J. M. Carlson. Navigation and control of outcomes in a generalized Lotka–Volterra model of the microbiome. In J. Kotas, editor, Advances in Nonlinear Biological Systems: Modeling and Optimal Control, volume 11 of AIMS Series on Applied Mathematics, pages 97–120. American Institute of Mathematical Sciences, Springfield, MO, USA, 2020.

Geometric optimal control of the generalized Lotka–Volterra model of the intestinal microbiome

Bernard Bonnard[1,2] | **Jérémy Rouot**[3] | **Cristiana J Silva**[4,5]

[1]Université de Bourgogne Franche-Comté, Institut de Mathématiques de Bourgogne, Dijon, France

[2]Inria Sophia Antipolis, Sophia Antipolis, France

[3]Univ Brest, UMR CNRS 6205, Laboratoire de Mathématiques de Bretagne Atlantique, Brest, France

[4]Iscte - Instituto Universitário de Lisboa, Lisbon, Portugal

[5]Center for Research and Development in Mathematics and Applications (CIDMA), Aveiro, Portugal

Abstract

We introduce the theoretical framework from geometric optimal control for a control system modeled by the generalized Lotka–Volterra equation, motivated by restoring the gut microbiota infected by Clostridium difficile combining antibiotic treatment and fecal injection. We consider both permanent control and sampled-data control related to the medical protocols.

KEYWORDS

biomathematics and population dynamics, optimal control in the permanent case, sampled-data control

Bonnard, B., Rouot, J., & Silva, C. J. (2024). Geometric optimal control of the generalized Lotka–Volterra model of the intestinal microbiome. Optimal Control Applications and Methods, 45(2), 544-574.

5 | CONCLUSION

In this article, we have presented mainly the techniques from geometric control theory to analyze reduction of the infection of a gut microbiote by a pathogenic agent using a controlled Lotka–Volterra model in dimension $N = 11$, which can admit up to $2^{11} = 2048$ interacting equilibria.

In the optimal control context the problem can be analyzed combining indirect or direct schemes in the permanent or sampled-data control frame both aspects are complementary. They were applied to the $2d$-case but can be generalized to the N-dimensional case, the limit being the computational complexity.

Bonnard, B., Rouot, J., & Silva, C. J. (2024). Geometric optimal control of the generalized Lotka–Volterra model of the intestinal microbiome. Optimal Control Applications and Methods, 45(2), 544-574.

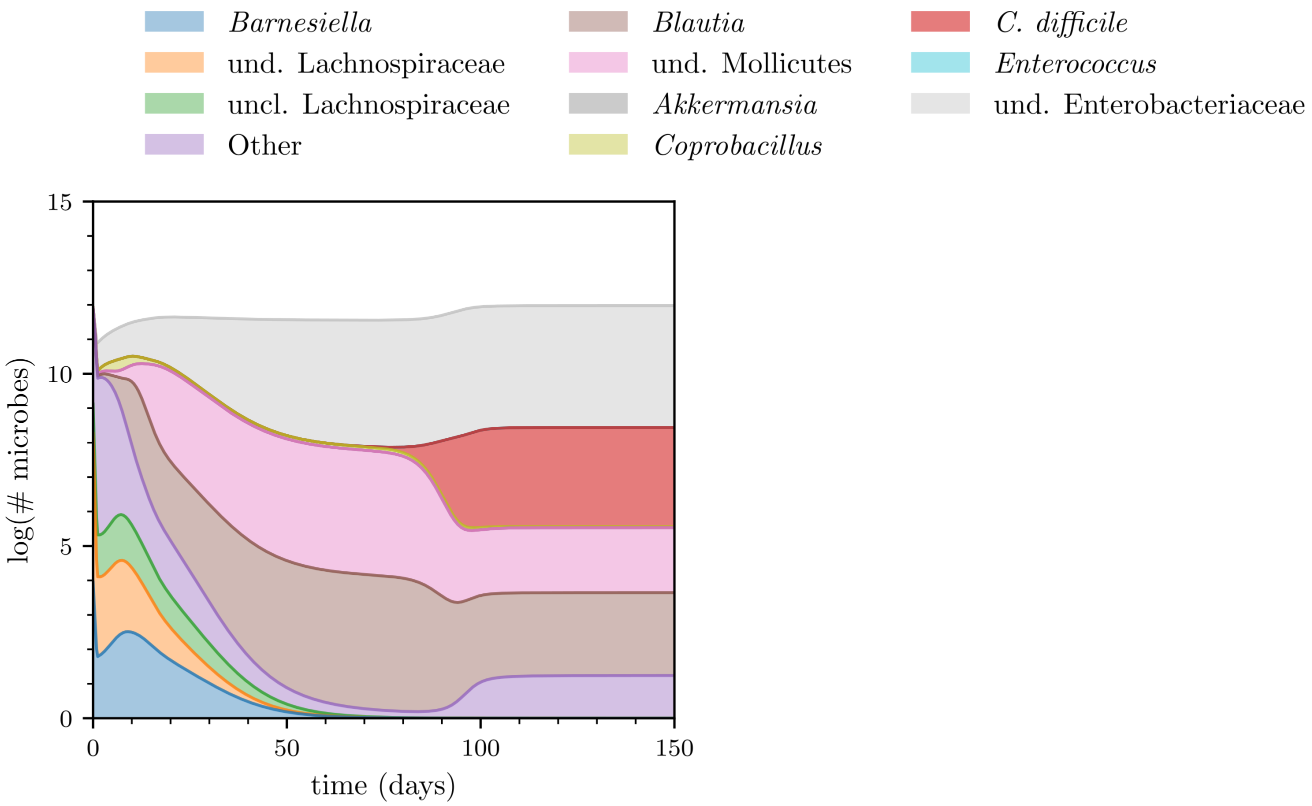

Böttcher, L., Control of dynamical systems with neural networks, preprint

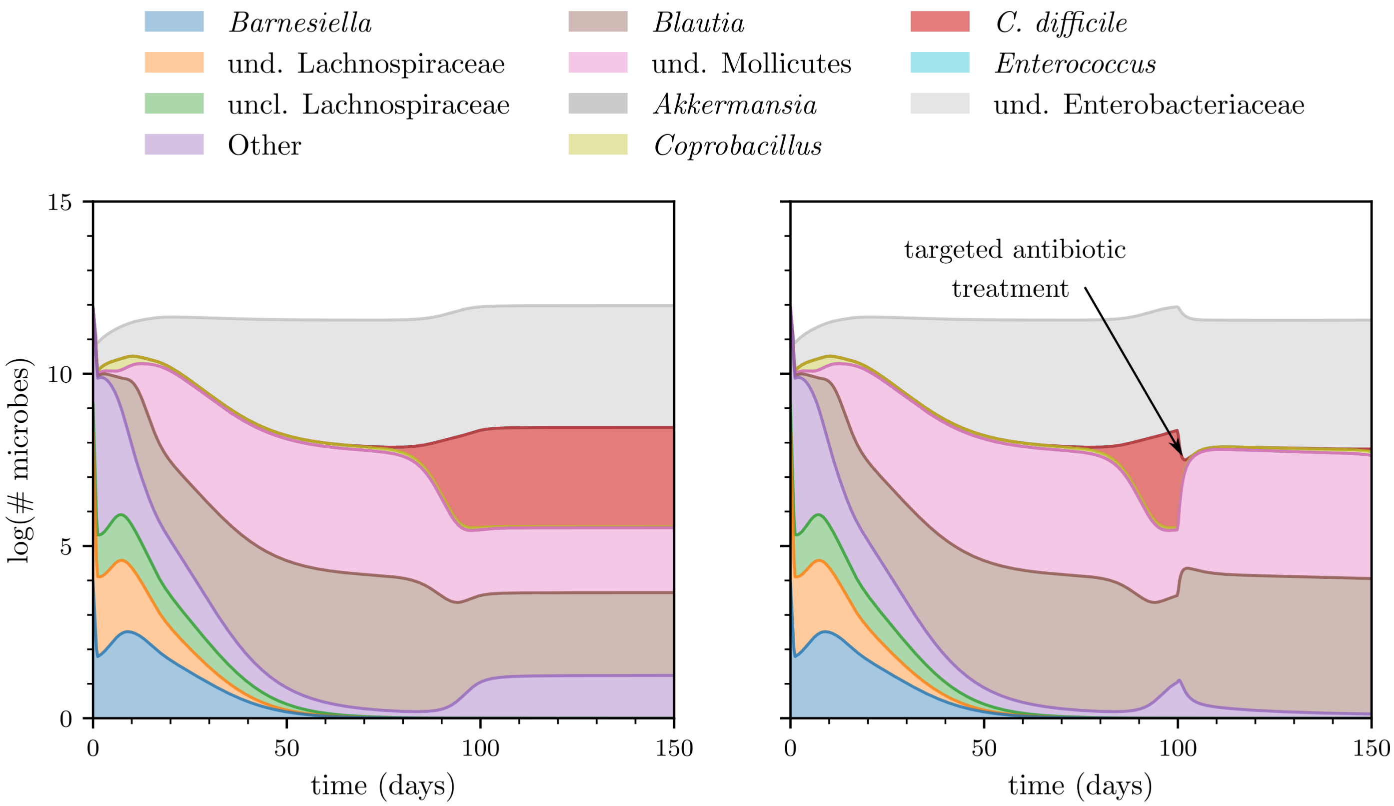

Böttcher, L., Control of dynamical systems with neural networks, preprint

Multiple Species and Substrates

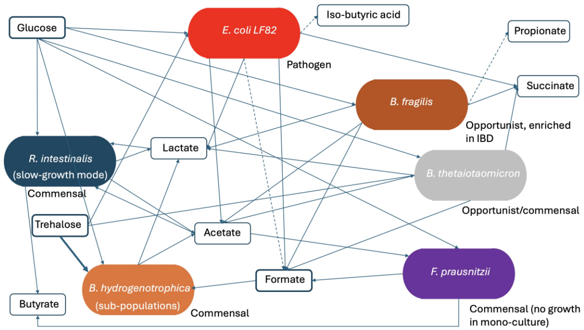

growth rate

abundance of microbial species *i*

$$\dot{x}_i(t) = \left[\sum_j \left(r_{ij}\,\frac{s_j(t)}{s_j(t)+K_{ij}}\right)\right] x_i(t) - k_i\, x_i(t) - D\, x_i(t)$$

Monod constant

decay rate

dilution rate

Multiple Species and Substrates

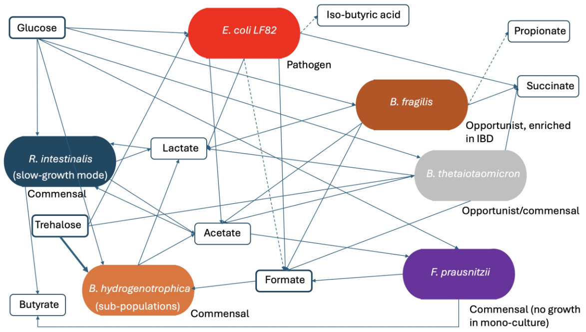

growth rate

abundance of microbial species *i*

$$\dot{x}_i(t) = \left[\sum_j \left(r_{ij} \frac{s_j(t)}{s_j(t) + K_{ij}} \right) \right] x_i(t) - k_i \, x_i(t) - D \, x_i(t)$$

Monod constant

decay rate

dilution rate

$$\dot{s}_j(t) = \sum_i \left(b_{ij} \, x_i(t) \; - \; a_{ij} \, r_{ij} \frac{s_j(t)}{s_j(t) + K_{ij}} x_i(t) \right) \; + \; D\big(s_{0j} - s_j(t)\big)$$

production rate

consumption rate

substrate inflow

Multiple Species and Substrates

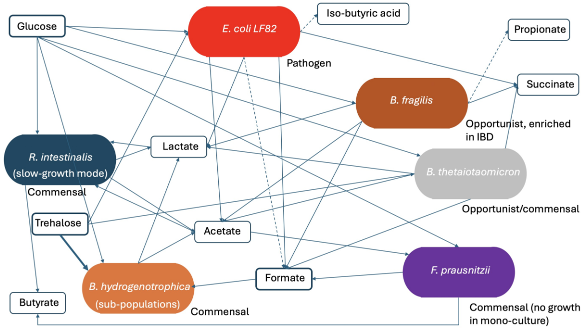

growth rate

abundance of microbial species *i*

$$\dot{x}_i(t) = \left[\sum_j \left(r_{ij} \frac{s_j(t)}{s_j(t) + K_{ij}} \right)\right] x_i(t) - k_i\, x_i(t) - D\, x_i(t)$$

Monod constant

decay rate

dilution rate

potential controls
probiotics, nutrients, antibiotics, fecal transplant

$$\dot{s}_j(t) = \sum_i \left(b_{ij}\, x_i(t) \; - \; a_{ij}\, r_{ij} \frac{s_j(t)}{s_j(t) + K_{ij}}\, x_i(t) \right) \; + \; D\left(s_{0j} - s_j(t)\right)$$

production rate

consumption rate

substrate inflow

From Bioreactor Data to Parameter Inference and Control

Bioreactor experiments

Parameter inference

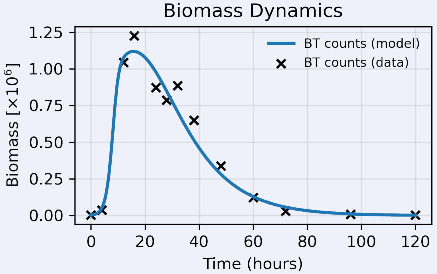

Control

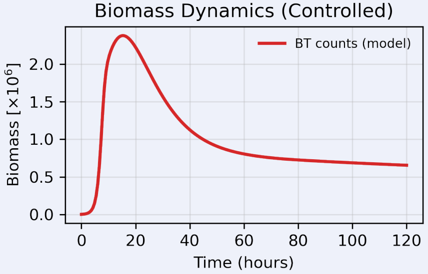

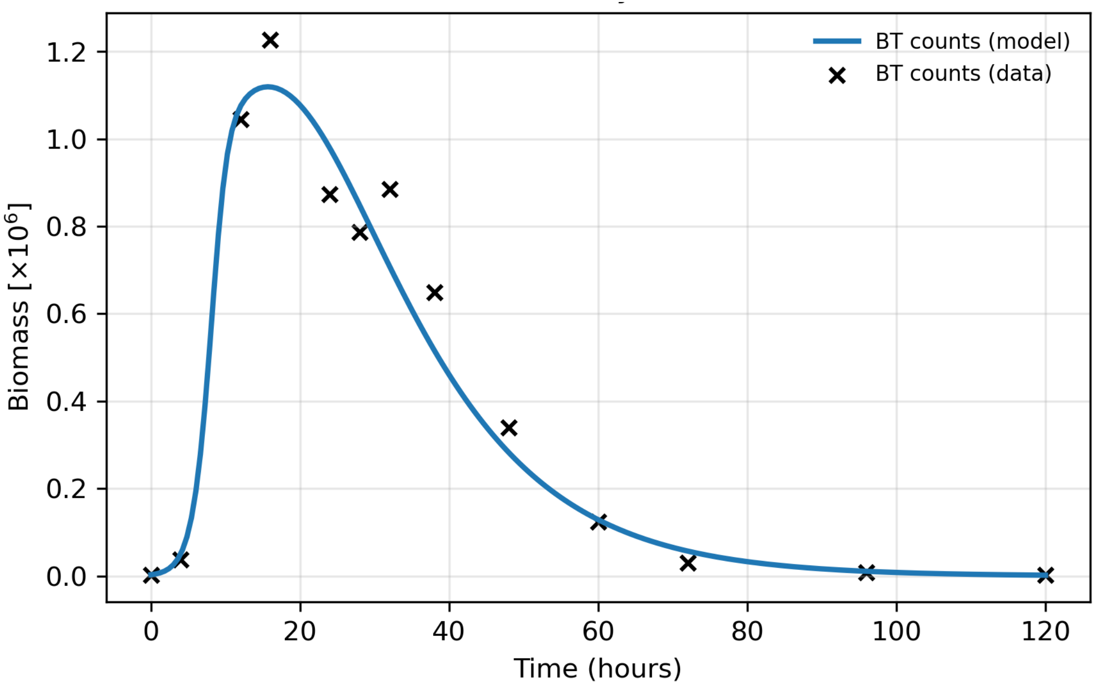

Basic setup:

1) Monitoring the biomass of one bacterial species (Bacteroides thetaiotaomicron) and of six substrates

2) Inferring model parameters

3) Controlling model dynamics

4) Validation of control function in another bioreactor experiment

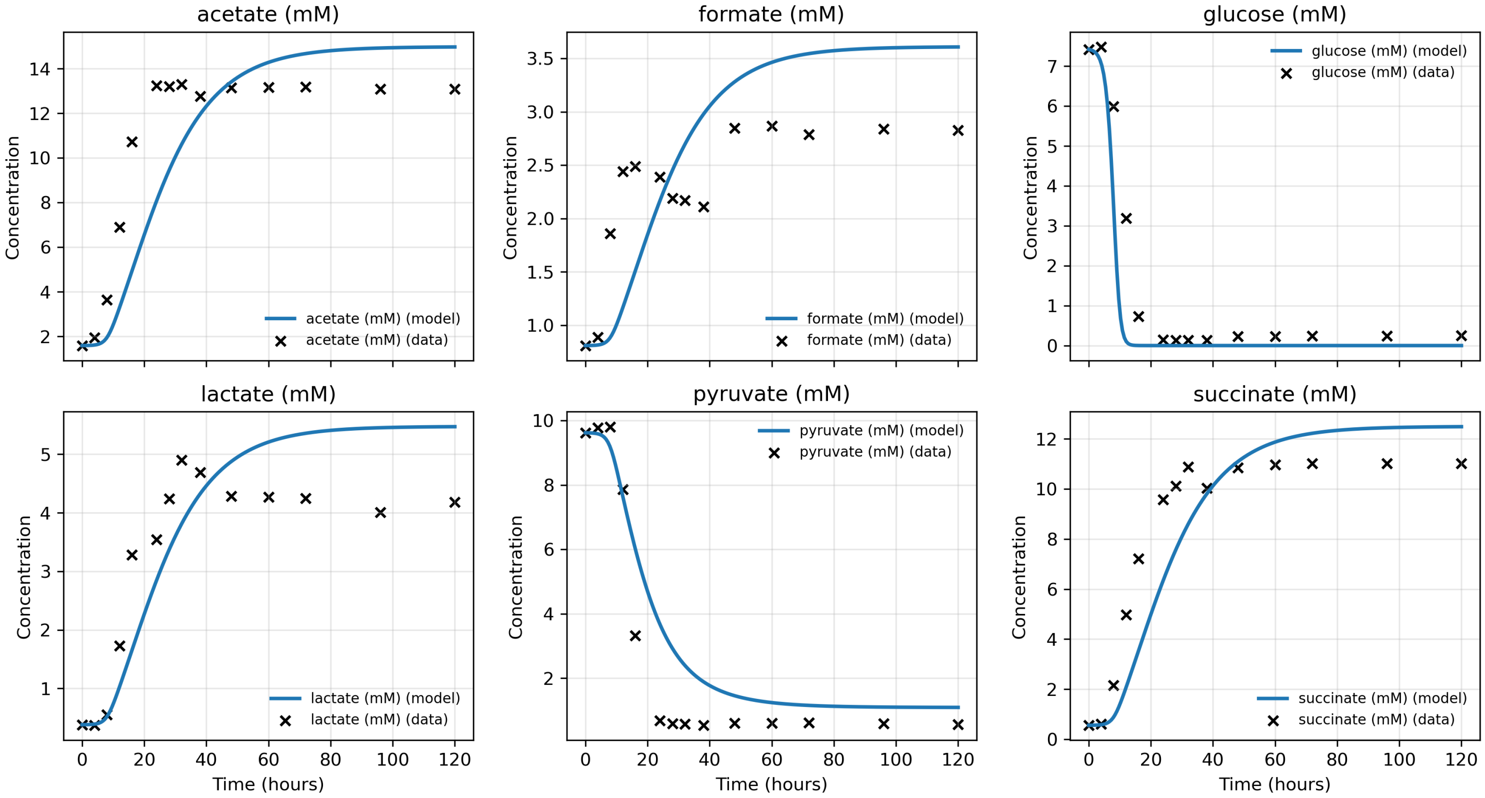

acetate (mM)
formate (mM)
glucose (mM)
lactate (mM)
pyruvate (mM)
succinate (mM)
Concentration
Time (hours)
acetate (mM) (model)
acetate (mM) (data)
formate (mM) (model)
formate (mM) (data)
glucose (mM) (model)
glucose (mM) (data)
lactate (mM) (model)
lactate (mM) (data)
pyruvate (mM) (model)
pyruvate (mM) (data)
succinate (mM) (model)
succinate (mM) (data)

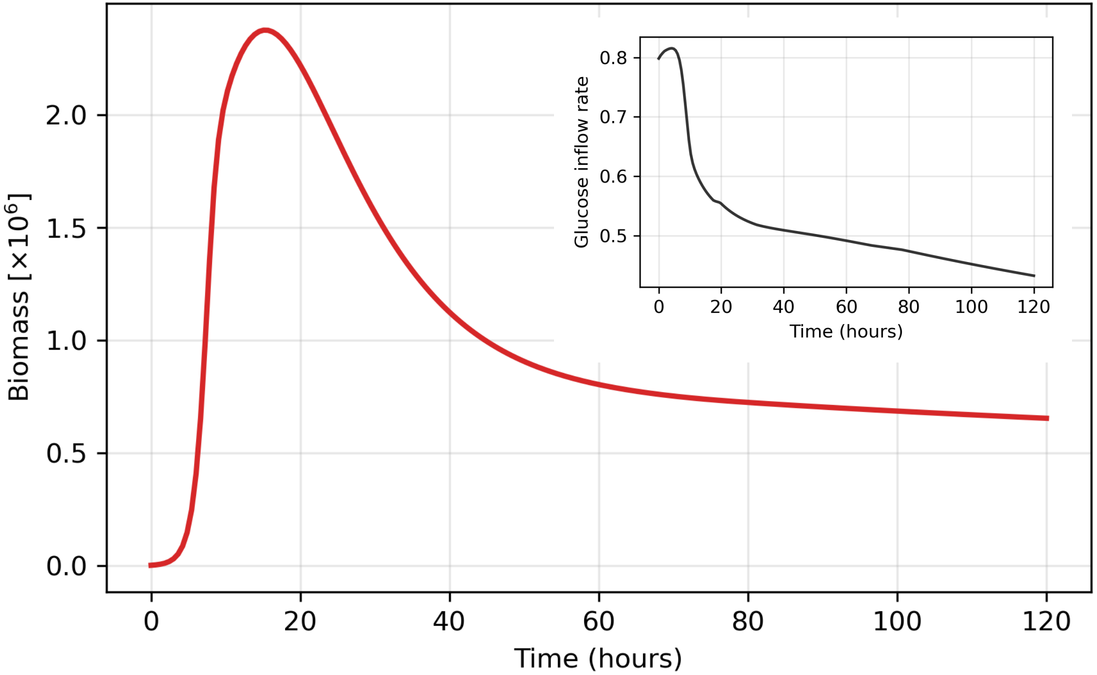

- Control objective is to **increase biomass** while keeping the **inflow** of **glucose** at a **minimum**
- The **larger biomass** is associated with a **larger production** of acetate, succinate, and other **substrates**.

Microbiome Project Team

- Karoline Faust (KU Leuven)
- Pallabita Saha (KU Leuven)
- Lorenzo Sala (INRAE Jouy-en-Josas)
- Didier Gonze (Université Libre de Bruxelles)
- Indrah Thelaganathan (FS)

Contributions

- Efficient control of **high-dimensional** dynamical systems (both **deterministic** and **stochastic**)

Contributions

- Efficient control of **high-dimensional** dynamical systems (both **deterministic** and **stochastic**)

- Up to 100× **faster runtime** with **improved** numerical **stability** vs. **PMP-based** control methods

Contributions

- Efficient control of **high-dimensional** dynamical systems (both **deterministic** and **stochastic**)

- Up to 100× **faster runtime** with **improved** numerical **stability** vs. **PMP-based** control methods

- Promising results in **biomedical applications**

Home > Collection

Biosystem Dynamics

Participating journal: Nonlinear Dynamics

Open for submissions

Submission deadline
Ongoing

Topical Collection of articles on Biosystem Dynamics

Participating journal

Submit your manuscript to this collection through the participating journal.

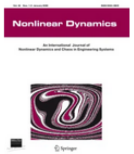

Journal

Nonlinear Dynamics

Nonlinear Dynamics is a hybrid journal publishing original content at the forefront of nonlinear dynamic research across diverse systems and scales.

Publishing model	**Hybrid**
Journal Impact Factor	**6.0 (2024)**
Downloads	**1.7M (2024)**
Submission to first decision (median)	**3 days**

Submit to this journal →

Submission guidelines →

References

https://github.com/asikist/nnc

https://gitlab.com/computational-science

- T. Asikis, L. Böttcher, N. Antulov-Fantulin, *Neural ordinary differential equation control of dynamics on graphs*, Physical Review Research 4, 013221 (2022)
- L. Böttcher, T. Asikis, Near-optimal control of dynamical systems using neural ordinary differential equations, Machine Learning: Science and Technology 3, 045004 (2022)
- L. Böttcher, N. Antulov-Fantulin, T. Asikis, *AI Pontryagin or how artificial neural networks learn to control dynamical systems*, Nature Communications 13, 333 (2022)
- M. Xia, L. Böttcher, T. Chou, *Controlling epidemics through optimal allocation of test kits and vaccine doses across networks*, IEEE Transactions on Network Science and Engineering 9, 1422-1436 (2022)
- L. Böttcher, T. Asikis, I. Fragkos, *Control of dual-sourcing inventory systems using recurrent neural networks*, INFORMS Journal on Computing 35, 1308-1328 (2023)
- L. Böttcher, Gradient-free training of neural ODEs for system identification and control using ensemble Kalman inversion, ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems at the International Conference on Machine Learning (ICML), Honolulu, Hawaii, USA (2023)
- Böttcher, L., Fonseca, L. L., & Laubenbacher, R. C. (2025). Control of medical digital twins with artificial neural networks. Philosophical Transactions of the Royal Society A 383: 20240228.
- Fonseca, L. L., Böttcher, L., Mehrad, B., & Laubenbacher, R. C. (2025). Optimal control of agent-based models via surrogate modeling. PLOS Computational Biology, 21(1), e1012138.