Cellular computing without bioengineering

Paul Ahavi[1], Thi-Ngoc-An Hoang[1], Philippe Meyer[1], Sylvie Berthier[2,3], Federica Fiorini[2,4], Florence Castelli[2,3], Olivier Epaulard[2,5], Audrey Le Gouellec[2], Jean-Loup Faulon[1,*]

[1] University Paris Saclay, INRAE, AgroParisTech, MICALIS Institute, 78350 Jouy-en-Josas, France.
[2] University Grenoble Alpes, CNRS, CHU Grenoble, Grenoble INP, TIMC-IMAG UMR 5525, 38000 Grenoble, France
[3] University Paris Saclay, CEA, INRAE, Médicaments et Technologies pour la Santé (MTS), MetaboHUB, 91190 Gif-sur-Yvette, France
[4] Department of Drug Metabolism & Pharmacokinetics, Boehringer Ingelheim Pharma GmbH & Co. KG, Biberach, Germany.
[5] Infectious Diseases Department, Grenoble-Alpes University Hospital, 38000 Grenoble, France
[*] Corresponding author. Email Jean-Loup.Faulon@inrae.fr

Abstract:

Biological systems possess inherent computational capabilities, yet leveraging them for practical computing tasks remains underexplored. Here, we repurpose *Escherichia coli* as a living reservoir within a reservoir computing (RC) framework to solve regression and classification tasks. Using growth rate measurements across diverse media, our *E. coli*-based physical RC consistently outperforms linear models and rivals multilayer perceptrons. We extend this approach to genome-scale metabolic models of multiple bacterial strains, revealing that computational performance correlates with phenotypic diversity. As a clinically relevant application, we classify COVID-19 patient plasma samples by disease severity, achieving high accuracy using only bacterial growth curves. Our system operates without complex instrumentation, offering a low-cost, scalable, and infrastructure-independent alternative for biocomputation and diagnostics.

Performing computation with living organisms is one of the main endeavors of synthetic biology where many electronic-like circuits have been engineered in a bottom-up fashion (*1*)(*2*). Engineering and tuning such circuits is challenging and several factors such as metabolic burden (*3*), orthogonality issues (*4*) and low noise tolerance (*5*) strongly restrain their complexity.

Synthetic biology devices have largely been inspired by behaviors observed in natural systems. Indeed, information processing systems analogous to logic gates (*6*), switches (*7*), and perceptrons (*8*) have all been discovered in various organisms. This raises the question: instead of engineering bottom-up devices, could we develop a top-down approach where natural organisms are repurposed to perform complex computational tasks?

The study of physical systems' ability to solve computational tasks falls within the field of Artificial Intelligence known as reservoir computing (RC). Initially developed as an alternative to classical artificial neural networks, particularly recurrent neural networks (RNNs) (*9*), RC offers an efficient training process.

RC systems come in two broad forms: conventional and physical RC (**Fig. 1**).

In a typical conventional RC framework, input data is provided to a reservoir, the states of the reservoir nodes are read, and this information is fed into a post-readout, usually a simple linear classifier or regressor, that interprets the reservoir states to produce the final output. Unlike RNNs with trainable weights, conventional RC uses a fixed-weight model, making it faster to train.

In physical RC, the reservoir is replaced by an actual physical system or process. Such RC begins by injecting input data into the physical reservoir, where inputs are mapped into a high-dimensional, time-evolving state. These states are then passed to a post-readout layer, similar to conventional RC. Physical RC has many applications, including the seminal liquid state machine for pattern recognition in a bucket of water (*10*), chemical RC developments for classification tasks and forecasting solutions of ODE systems (*11*), and RC systems using rat cortical neurons cultured on micropatterned substrates to solve classification tasks (*12*).

The aim of the current study is to use bacteria as biocomputing reservoirs and to probe to which extent bacterial models in conventional RC systems and *E. coli* (K-12 MG1655) in a physical RC system can solve machine learning regression and classification tasks of various difficulties. We also present a practical *E. coli* physical RC application with the prediction of COVID-19 disease severity (mild vs. severe) using plasma samples.

Bacterial reservoir computing framework

In typical machine learning regression and classification tasks, training sets are composed of features and labels, and the goal is to predict the labels from the features. In the present study, bacteria act as our reservoirs. We transform the problem features into nutrients provided to the strain, observe the strain's behavior by measuring its phenotype, and feed these phenotype measurements to a readout feedforward neural network, which outputs the problem's solution (**Fig. 1**).

One challenge in this workflow is finding the most appropriate mapping between features and nutrients. There is no generic procedure in the RC literature; instead, there are various methods depending on the reservoir (e.g., using water waves to mimic sound waves for classifying voice recordings in liquid state machines (*10*)). In our case, one strategy could be an arbitrary one-to-one mapping of features to nutrients. However, this strategy might not yield the best performance. Another approach could be to find the optimal nutrients by training a prior-input layer that generates nutrient sets. While this is not customary in RC, there is no reason why the prior layer cannot be trained (*13*). To train the prior layer, one needs to backpropagate a gradient from the readout to the prior layer. This requires a gradient-compatible model that mimics the behavior of the strain, which in turn necessitates replacing the physical strain reservoir with a conventional reservoir that acts as a proxy for the strain.

Many solutions have been developed over several decades to model strain behavior. Among these, we are particularly interested in models dealing with metabolism, as we provide the strain with various metabolites (nutrients) and measure phenotypes such as growth rates. While kinetic models have been proposed to model metabolism, they are limited in scope regarding species and sizes (*14*). Instead, we propose to use genome-scale metabolic models (GEMs, *cf.* supplementary materials **M2**) that have been developed for many strains (*15*). To capture the dynamics of these models, we choose flux balance analysis (FBA) (*16*) a well-established technique in systems biology and bioengineering. FBA is a linear program making use of the Simplex algorithm (*17*), which does not allow gradient backpropagation to the prior layer. To overcome this issue, we replaced the FBA reservoir by an Artificial Metabolic Network (AMN, *cf.* **fig. S1**), a surrogate of FBA which we previously developed (*18*), and is gradient backpropagation compatible.

Throughout this paper we will be using an *E. coli* physical RC (strain K-12 MG1655) and a conventional twin based on an *E. coli* AMN model. We will also make use of conventional RC for other bacteria for which GEM reservoirs are available. All regression coefficients and accuracies

reported in the paper have been obtained using a cross-validation procedure described in supplementary materials **M10**.

Benchmarking conventional GEM-based reservoir

We first constructed a conventional RC system using an AMN model (built using GEM model iML1515), trained on growth rates acquired cultivating *E. coli* K-12 MG1655 across 280 nutrient media (*cf.* **fig. S1**, and supplementary materials **M3** and **M5**).

To assess the system's computational capacity, we evaluated it on 10 benchmark regression tasks from the OpenML platform (supplementary materials **M6** and table **S1**). For each task, input features were first mapped to nutrient combinations by a prior layer. These media were then processed through a trained AMN that computed the full set of reaction fluxes. Finally, the flux states were decoded by a feedforward readout to predict the target output readout (**Fig. 2.a**).

We developed two different methods for the prior layer. The first method (fixed prior) does not require training; instead, features are one-to-one mapped to medium metabolites using permutation importance (*cf.* supplementary materials **M8** and **table S2**). The second method (generative prior) outputs medium composition from features, with a generative model trained to produce the best medium composition for each element of the training set.

We compared the performances of these architectures with classical regressors: multi-linear regression (MLR), multilayer perceptron (MLP), and XGBoost (XGB) (*cf.* **Fig. 2.b-c**, and **table S3** for classical regressors hyperparameters).

The generative prior architecture produced higher regression accuracy than MLR and performed on par with MLP across multiple tasks (**Fig. 2.**). In contrast, the fixed prior consistently underperformed, suggesting that arbitrary feature-to-metabolite mappings limit learning capacity. We nonetheless observe that for some problems, the fixed prior architecture outperforms MLR, indicating that some learning is still performed with that prior.

To assess the ability of bacterial strain models to solve regression tasks, we tested 10 GEMs (**table S4**) on the 10 regression problems mentioned above. Because experimental phenotype data were unavailable for all strains, GEM reservoirs were trained on FBA-simulated data with the 28 media metabolites used previously (supplementary materials **M4**). As shown in **Fig. 2.**, the generative prior architecture consistently outperformed both MLR and MLP across many strains, while the

fixed prior was less effective (**fig. S3a**). Performance varied between strains, independently of the prior type.

To explore these differences, we quantified the number of distinct phenotypes each GEM could produce with the 28 media metabolites (supplementary materials **M11**). **Fig. 2.** and **fig. S3b-c** reveals a positive correlation between phenotype diversity and regression accuracy, consistent across both priors. This suggests that strains producing only a few phenotypes (e.g., growth vs. no-growth) are less suited to complex regression tasks. The trend is even stronger when considering phenotype Shannon entropy.

Solving regression and classification tasks with a physical *E. coli* reservoir

As illustrated in **Fig. 3.**, the physical *E. coli* RC uses a live bacterial strain as the reservoir. Unlike conventional reservoirs, it (1) considers only growth rates as phenotypes and (2) relies on measured (not predicted) growth data to generate readouts. To implement this, we first trained a conventional twin RC (similar to the one shown in **Fig. 2.**), matching the physical strain's behavior but outputting only growth rates. This twin was trained on 280 media compositions with associated growth rates (supplementary materials **M3**). We then selected media compositions optimized by the twin for each of the 10 regression tasks, cultured *E. coli* K-12 MG1655 accordingly, and measured growth. These values were fed back into the readout layer to compute final predictions.

Our conventional twin makes use of two different prior architectures : generative and selective (**table S2).** The generative prior, seen previously, is producing media compositions some of which might already be in our 280 media compositions set while others are novel. We note that growth rates need to be measured for these novel media compositions when feeding the generative prior output to the physical RC. The selective prior is forced to choose media compositions that are already in the 280 media compositions set. That selection is performed using a Gumbel-softmax (*19*) activation function (supplementary materials **M8**). With this prior all growth rates have already been measured and it is not necessary to carry out additional growth measurements.

We note that since our physical RC and conventional twin are considering only growth rate as phenotype is it not necessary to train a readout layer, in fact after normalizing the growth rate and the problem readout the problem readout *is* the growth rate.

Results on the 10 regression problems obtained with the conventional twin and with generative or selective priors are presented in **Fig. 3.** and **fig. S4**. As seen before, the twin conventional RC outperforms MLR and yields performances similar to MLP.

We next evaluated the physical *E. coli* reservoir on three representative regression tasks with varying complexity. Additional growth measurements were collected for generative priors, while selective priors relied on premeasured media. Compared to the conventional twin, the physical reservoir showed reduced performance under the generative prior but matched it under the selective prior (**Fig. 3c–e**). This performance gap reflects the imperfect correlation between predicted and measured growth rates ($R^2 = 0.78$ in **fig. S4**).

Turning our attention to classification problems, we chose to benchmark our *E. coli* reservoirs against both linear and non-linear classification tasks, similar to those studied in Baltussen *et al.* (*11*). We focused on eight classification tasks, generating a two-dimensional input space from 500 randomly sampled coordinates (supplementary materials **M7**). We used the conventional reservoir from **Fig. 2.** with a fixed prior and the *E. coli* physical reservoir from **Fig. 3.a** with a selective prior (*cf.* supplementary materials **M8** for the mapping between coordinates and nutrients) . The post-readout layer applied a sigmoid function to either all reaction fluxes (fixed prior) or growth rates (selective prior).

As shown in **Fig. 3.f-g**, the conventional RC system with fixed prior performed better on classification than regression tasks. For non-linear tasks, it outperformed SVM and surpassed MLP in the SINE and DOTS tasks. This aligns with our previous finding (**Fig. 2d-e**) that strain-RC performance depends on phenotype diversity, favoring binary classification where only two classes and phenotypes should be required.

The physical RC system with selective prior, validated experimentally, consistently outperformed SVM and matched or exceeded MLP on several tasks (AND, OR, LINEAR, SINE, DOTS). This demonstrates its capacity for non-linear classification, akin to chemical reservoirs (*11*), though it uses only growth rate measurements.

Classifying plasma samples with a physical *E. coli* reservoir

Classifying COVID-19 patients to predict whether they will develop a mild or severe form of the illness has been a significant research focus since the onset of the pandemic (*20*). Most methods performing classification are making use of proteomic and metabolomic profiling conducted on biological samples (*20*)(*21*)(*22*) collected from patients upon their admission to the hospital. After having been trained on these omics data, ML techniques are used to predict the disease outcome

(e.g., mild or severe according to whether ICU admission was required (*21*)(*23*)). In line with this approach, we utilized a cohort of COVID-19 patients from the CHUGA University Hospital (supplementary materials **M12**) and performed untargeted metabolomics analyses (via high resolution mass spectrometry) on plasma samples collected on 81 COVID-19 patients (50 mild, 31 severe, *cf.* supplementary materials **M13**).

a demonstrates that disease outcomes can be predicted using this metabolomics data with cross-validation accuracies greater than 0.9 using feature selection (supplementary materials **M14** and **fig. S6**). As our goal is to develop a physical reservoir for classifying the samples, we also successfully made predictions by extracting signals from the metabolomics data corresponding to metabolites expected to cross the *E. coli* membrane and contributing to its growth.

Having demonstrated that medium metabolite MS signals could be used to distinguish mild from severe samples, we tested if a conventional *E. coli* GEM reservoir could also distinguish mild vs. severe based on strain phenotypes. The architecture of that reservoir is shown in **b**. Without measured growth rates, we trained an AMN with data produced by a GEM (iML1515) coupled with FBA as our conventional reservoir. The data was generated using the classical objective of maximizing growth rate when provided with the MS signals of metabolites predicted to cross *E. coli* membrane. When using the conventional reservoir, we were unable to classify mild vs. severe with sufficient confidence when the post-readout was based solely on growth rate. However, when the post-readout was based on FBA-computed flux distributions for all reactions, we reached accuracies above 0.9 using feature selection for several classifiers (**c** and **fig. S7**). These high accuracies must be interpreted carefully, since FBA-computed fluxes show limited agreement with experimental measurements (**fig. S2c**). At best, such results supports further development of a physical reservoir.

Seeking to construct a physical *E. coli* reservoir growing on plasma, we collected samples from patients of the CHUGA COVID-19 cohort before they developed either mild or severe forms of the disease *(*supplementary materials **M12**). To this cohort, we added plasma samples from COVID-negative patients, obtained from the French Blood Establishment (Etablissement français du sang, EFS). Since *E. coli* cannot grow in the presence of antibiotics, any samples containing antibiotics were excluded. The final dataset consisted of 27 negative samples, 37 mild cases, and 36 severe cases. All these samples had to be pretreated in order to optimize bacterial growth on

plasma and therefore classification potential across sample categories (supplementary materials **M15**).

By growing *E. coli* on negative, mild, and severe plasma samples, we first experimentally confirmed our conventional reservoir prediction: such a reservoir is unable to solve the mild vs. severe classification task when relying solely on growth rate as input to the post-readout classifier. We also observed we could not classify mild vs. severe when taking lag time and maximum optical density taken at 600 nm (OD_{MAX}) as input features in addition to growth rates (**table S5**). To further increase the number of features available to the classifier, we could not use flux distributions as in a conventional reservoir. Instead, we utilized optical density measurements taken at all time points along the growth curves (
d).

Using this approach and 20-fold cross-validation, the *E. coli* physical reservoir achieved a classification accuracy of 0.84 ± 0.006 for distinguishing mild from severe samples. Furthermore, the reservoir demonstrated high accuracy (0.94 ± 0.005) in distinguishing negative samples from positive ones. These results were obtained by selecting sets of time points that maximized classification accuracies (supplementary materials **M9** and **fig. S8**). Results for other classifiers (SVM and MLP) and other folds (5 and LOO) are presented in **table S6**.

For the mild vs. severe classification task, the selected time points are all within the first 6 hours. In contrast, for the negative vs. positive classification, the selected time points extend up to 8 hours. Generally, *E. coli* biomass at 8 hours is lower in positive samples compared to negative ones, and we observe that growth is generally delayed in the first 4 hours when *E. coli* is grown on severe samples compared to mild ones. These finding are substantiated by an analysis we performed using Generalized Additive Mixed Models (GAMMs (*24*), *cf.* supplementary materials **M16**). Additionally, observed differences in OD profiles across sample categories are consistent with metabolite levels reported in the literature (*cf.* supplementary text).

Discussion

This study demonstrates the potential of leveraging conventional reservoirs constructed from bacterial strain models, alongside with physical *E. coli* reservoirs, to address both regression and classification tasks.

Our conventional bacterial RC systems outperform standard regression methods, including multi-linear regression (MLR) and multi-layer perceptron (MLP). In classification tasks, even untrained fixed priors surpassed SVMs in all cases and outperformed MLPs in several. Our physical RC implementation using an *E. coli* strain with selective priors further confirms this robustness, outperforming MLR and SVM across all tasks and matching or exceeding MLP in many regression and classification tasks (**fig. S4** and **Fig. 3.f**).

The discrepancy between regression and classification performance likely reflects task complexity, binary classification requires only two separable outputs, while regression demands broader phenotypic diversity. Supporting this, **Fig. 2.d-e** and **fig. S3** show a positive correlation between the number of internal states (phenotypes) and regression accuracy. This trend aligns with findings in reservoir computing, where increased state diversity enhances computational capacity (*25*)(*26*)(*27*).

Compared to the chemical reservoir (*11*) and biological neuronal reservoirs (mBNNs) (*12*), the *E. coli* physical reservoir offers a practical balance of classification performance, simplicity, and cost-effectiveness. The setup requires only standard microbiology equipment: an incubator and a plate reader with minimal consumables. In contrast, the chemical reservoir depends on mass spectrometry and flow control systems, while mBNNs require optogenetics, microfluidics, and live-cell imaging, all of which demand substantial infrastructure and technical expertise. Despite this simplicity, the *E. coli* reservoir delivers competitive accuracies across a range of regression and classification tasks, relying solely on growth rates as inputs to the readout.

The main key advantage of the *E. coli* reservoir compared to other published physical reservoirs lies in its applicability to clinically relevant settings. Standard ML pipelines applied to untargeted metabolomics data on plasma samples yield accuracies around 0.9 for COVID-19 disease severity classification. Using OD_{600} growth curves, our *E. coli* reservoir distinguished between mild vs. severe COVID-19 cases with a cross-validated accuracy of 0.84 with 20-fold cross-validation (0.86 with LOO), and achieved even higher accuracy (0.94) when classifying between COVID-19 positive and negative samples.

Aside from using omics analyses only few studies have been carried out to predict COVID-19 severity and these generally include human health records as features (*28*). Compared to these methods, our physical *E. coli* reservoir approach offers several advantages. Omics-based pipelines are generally expensive, time-consuming, and require specialized infrastructure and expertise. Our reservoir system achieves competitive accuracy without including patient data, operates with

minimal cost and instrumentation, and produces results within hours using only a plate reader and bacterial cultures. These properties make it a highly accessible, scalable, and infrastructure-independent alternative for severity prediction, particularly valuable in decentralized or low-resource healthcare settings.

All together this work supports the idea that living systems possess intrinsic computational capacity, echoing the concept of cellular computing supremacy proposed by Goñi-Moreno (*29*). Our findings position bacterial RC as a biomimetic and infrastructure-light alternative capable of performing meaningful and useful information processing without the need for complex genetic re-engineering.

References and Notes

1. P. E. M. Purnick, R. Weiss, The second wave of synthetic biology: from modules to systems. *Nat Rev Mol Cell Biol* **10**, 410–422 (2009).
2. Y. Xiang, N. Dalchau, B. Wang, Scaling up genetic circuit design for cellular computing: advances and prospects. *Nat Comput* **17**, 833–853 (2018).
3. O. Borkowski, F. Ceroni, G.-B. Stan, T. Ellis, Overloaded and stressed: whole-cell considerations for bacterial synthetic biology. *Current Opinion in Microbiology* **33**, 123–130 (2016).
4. A. J. Meyer, T. H. Segall-Shapiro, E. Glassey, J. Zhang, C. A. Voigt, Escherichia coli "Marionette" strains with 12 highly optimized small-molecule sensors. *Nat Chem Biol* **15**, 196–204 (2019).
5. J. M. L. Ho, C. A. Miller, S. E. Parks, J. R. Mattia, M. R. Bennett, A suppressor tRNA-mediated feedforward loop eliminates leaky gene expression in bacteria. *Nucleic Acids Research* **49**, e25–e25 (2021).
6. R. Silva-Rocha, V. De Lorenzo, Mining logic gates in prokaryotic transcriptional regulation networks. *FEBS Letters* **582**, 1237–1244 (2008).
7. A. B. Oppenheim, O. Kobiler, J. Stavans, D. L. Court, S. Adhya, Switches in Bacteriophage Lambda Development. *Annu. Rev. Genet.* **39**, 409–429 (2005).
8. B. Scheres, W. H. van der Putten, The plant perceptron connects environment to development. *Nature* **543**, 337–345 (2017).
9. G. Tanaka, T. Yamane, J. B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, A. Hirose, Recent advances in physical reservoir computing: A review. *Neural Networks* **115**, 100–123 (2019).
10. C. Fernando, S. Sojakka, "Pattern Recognition in a Bucket" in *Advances in Artificial Life*, W. Banzhaf, J. Ziegler, T. Christaller, P. Dittrich, J. T. Kim, Eds. (Springer Berlin Heidelberg, Berlin, Heidelberg, 2003; http://link.springer.com/10.1007/978-3-540-39432-7_63)vol. 2801 of *Lecture Notes in Computer Science*, pp. 588–597.
11. M. G. Baltussen, T. J. De Jong, Q. Duez, W. E. Robinson, W. T. S. Huck, Chemical reservoir computation in a self-organizing reaction network. *Nature* **631**, 549–555 (2024).
12. T. Sumi, H. Yamamoto, Y. Katori, K. Ito, S. Moriya, T. Konno, S. Sato, A. Hirano-Iwata, Biological neurons act as generalization filters in reservoir computing. *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2217008120 (2023).
13. L. Carcano, E. Plebani, D. P. Pau, M. Piastra, Pre-trainable Reservoir Computing with Recursive Neural Gas. arXiv arXiv:1807.09510 [Preprint] (2018). https://doi.org/10.48550/arXiv.1807.09510.

14. S. Srinivasan, W. R. Cluett, R. Mahadevan, Constructing kinetic models of metabolism at genome-scales: A review. *Biotechnology Journal* **10**, 1345–1359 (2015).
15. Z. A. King, J. Lu, A. Dräger, P. Miller, S. Federowicz, J. A. Lerman, A. Ebrahim, B. O. Palsson, N. E. Lewis, BiGG Models: A platform for integrating, standardizing and sharing genome-scale models. *Nucleic Acids Research* **44**, D515–D522 (2016).
16. J. D. Orth, I. Thiele, B. Ø. Palsson, What is flux balance analysis? *Nat Biotechnol* **28**, 245–248 (2010).
17. H. Karloff, "The Simplex Algorithm" in *Linear Programming*, H. Karloff, Ed. (Birkhäuser, Boston, MA, 1991; https://doi.org/10.1007/978-0-8176-4844-2_2)*Modern Birkhäuser Classics*, pp. 23–47.
18. L. Faure, B. Mollet, W. Liebermeister, J.-L. Faulon, A neural-mechanistic hybrid approach improving the predictive power of genome-scale metabolic models. *Nat Commun* **14**, 4669 (2023).
19. E. Jang, S. Gu, B. Poole, Categorical Reparameterization with Gumbel-Softmax. arXiv arXiv:1611.01144 [Preprint] (2017). https://doi.org/10.48550/arXiv.1611.01144.
20. B. Shen, X. Yi, Y. Sun, X. Bi, J. Du, C. Zhang, S. Quan, F. Zhang, R. Sun, L. Qian, W. Ge, W. Liu, S. Liang, H. Chen, Y. Zhang, J. Li, J. Xu, Z. He, B. Chen, J. Wang, H. Yan, Y. Zheng, D. Wang, J. Zhu, Z. Kong, Z. Kang, X. Liang, X. Ding, G. Ruan, N. Xiang, X. Cai, H. Gao, L. Li, S. Li, Q. Xiao, T. Lu, Y. Zhu, H. Liu, H. Chen, T. Guo, Proteomic and Metabolomic Characterization of COVID-19 Patient Sera. *Cell* **182**, 59-72.e15 (2020).
21. F.-X. Danlos, C. Grajeda-Iglesias, S. Durand, A. Sauvat, M. Roumier, D. Cantin, E. Colomba, J. Rohmer, F. Pommeret, G. Baciarello, C. Willekens, M. Vasse, F. Griscelli, J.-E. Fahrner, A.-G. Goubet, A. Dubuisson, L. Derosa, N. Nirmalathasan, D. Bredel, S. Mouraud, C. Pradon, A. Stoclin, F. Rozenberg, J. Duchemin, G. Jourdi, S. Ellouze, F. Levavasseur, L. Albigès, J.-C. Soria, F. Barlesi, E. Solary, F. André, F. Pène, F. Ackerman, L. Mouthon, L. Zitvogel, A. Marabelle, J.-M. Michot, M. Fontenay, G. Kroemer, Metabolomic analyses of COVID-19 patients unravel stage-dependent and prognostic biomarkers. *Cell Death Dis* **12**, 258 (2021).
22. X. Bi, W. Liu, X. Ding, S. Liang, Y. Zheng, X. Zhu, S. Quan, X. Yi, N. Xiang, J. Du, H. Lyu, D. Yu, C. Zhang, L. Xu, W. Ge, X. Zhan, J. He, Z. Xiong, S. Zhang, Y. Li, P. Xu, G. Zhu, D. Wang, H. Zhu, S. Chen, J. Li, H. Zhao, Y. Zhu, H. Liu, J. Xu, B. Shen, T. Guo, Proteomic and metabolomic profiling of urine uncovers immune responses in patients with COVID-19. *Cell Reports* **38**, 110271 (2022).
23. I. Roberts, M. Wright Muelas, J. M. Taylor, A. S. Davison, Y. Xu, J. M. Grixti, N. Gotts, A. Sorokin, R. Goodacre, D. B. Kell, Untargeted metabolomics of COVID-19 patient serum reveals potential prognostic markers of both severity and outcome. *Metabolomics* **18**, 6 (2022).
24. S. N. Wood, *Generalized Additive Models: An Introduction with R* (Chapman and Hall/CRC, ed. 2, 2017; https://www.taylorfrancis.com/books/9781498728348).
25. J. Dambre, D. Verstraeten, B. Schrauwen, S. Massar, Information Processing Capacity of Dynamical Systems. *Sci Rep* **2**, 514 (2012).
26. R. Martínez-Peña, J. Nokkala, G. L. Giorgi, R. Zambrini, M. C. Soriano, Information Processing Capacity of Spin-Based Quantum Reservoir Computing Systems. *Cogn Comput* **15**, 1440–1451 (2023).
27. M. Dale, J. Dewhirst, S. O'Keefe, A. Sebald, S. Stepney, M. A. Trefzer, "The Role of Structure and Complexity on Reservoir Computing Quality" in *Unconventional Computation and Natural Computation*, I. McQuillan, S. Seki, Eds. (Springer International Publishing, Cham, 2019; http://link.springer.com/10.1007/978-3-030-19311-9_6)vol. 11493 of *Lecture Notes in Computer Science*, pp. 52–64.
28. T. Ozawa, S. Chubachi, H. Namkoong, S. Nemoto, R. Ikegami, T. Asakura, H. Tanaka, H. Lee, T. Fukushima, S. Azekawa, S. Otake, K. Nakagawara, M. Watase, K. Masaki, H. Kamata, N. Harada, T. Ueda, S. Ueda, T. Ishiguro, K. Arimura, F. Saito, T. Yoshiyama, Y.

Nakano, Y. Muto, Y. Suzuki, R. Edahiro, K. Murakami, Y. Sato, Y. Okada, R. Koike, M. Ishii, N. Hasegawa, Y. Kitagawa, K. Tokunaga, A. Kimura, S. Miyano, S. Ogawa, T. Kanai, K. Fukunaga, S. Imoto, Predicting coronavirus disease 2019 severity using explainable artificial intelligence techniques. *Sci Rep* **15**, 9459 (2025).

29. Á. Goñi-Moreno, Biocomputation: Moving Beyond Turing with Living Cellular Computers. *Commun. ACM* **67**, 70–77 (2024).

Acknowledgments:

We thank Alexandra Martin (INRAE, University of Paris Saclay) for her help in acquiring *E. coli* growth curves for different media. We thank Claire Roman (Campus Innovation Paris, Air Liquide & IRIMAS, University Haute Alsace) for her help and suggestions regarding the statistical analyses.

Funding:

ANR funding agency grant ANR-21-CE45-0021-01 (P.A, A.L.G, J.L.F)

PEPR B-BEST France 2030 program and ANR funding agency grant ANR-22-PEBB-0008 / (P.M., J.L.F)

UE HORIZON BIOS program, grant 101070281 (A. H, J.L.F).

ANR funding agency grant, ANR-24-CE44-1190 project (A.L.G.)

Foundation University Grenoble Alpes and Foundation Air Liquid for BIOMARCOVID grant (A.L.G.)

Author contributions:

Conceptualization: J.L.F

Methodology: J.L.F

Investigation: P.A, P.M, T.N.A.H. , J.L.F

Funding acquisition: A.L.G, J.L.F.

Project administration: A.L.G., J.L.F.

Supervision: A.L.G, J.L.F

Writing – original draft: P.A, T.N.A.H., P.M, A.L.G, J.L.F.

Competing interests:

Authors declare that they have no competing interests.

Data and materials availability:

Will be available upon journal publication

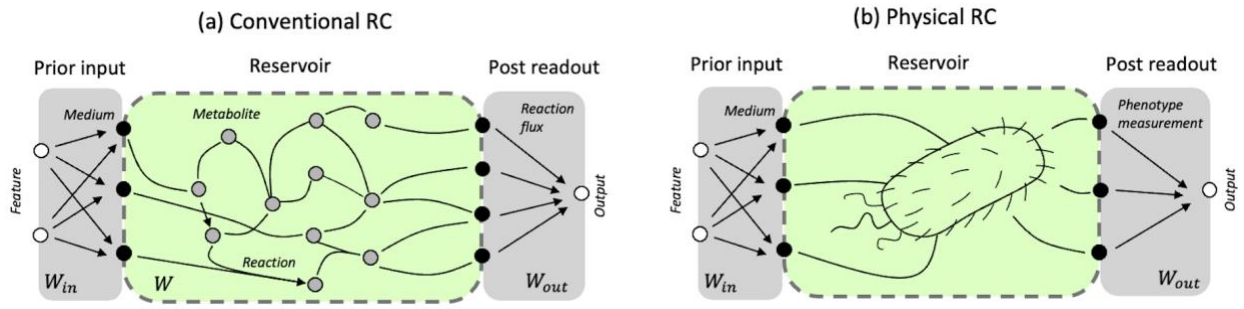

Fig. 1. Bacterial physical and conventional RC. The Figure, inspired by Tanaka et al. (9) was adapted for bacterial strain and models. (a) A conventional RC is a fixed weight (W) model (here a genome scale metabolic model). (b) In physical RC the reservoir is realized by a physical system or device (here a bacterial strain). In both cases the post readout neural network weights (W_{out}) are learned during training and the prior input weights (W_{in}) and reservoir weights are fixed. Yet, with conventional RC, the weights W_{in} can be learned through training as long as the reservoir allows gradient backpropagation.

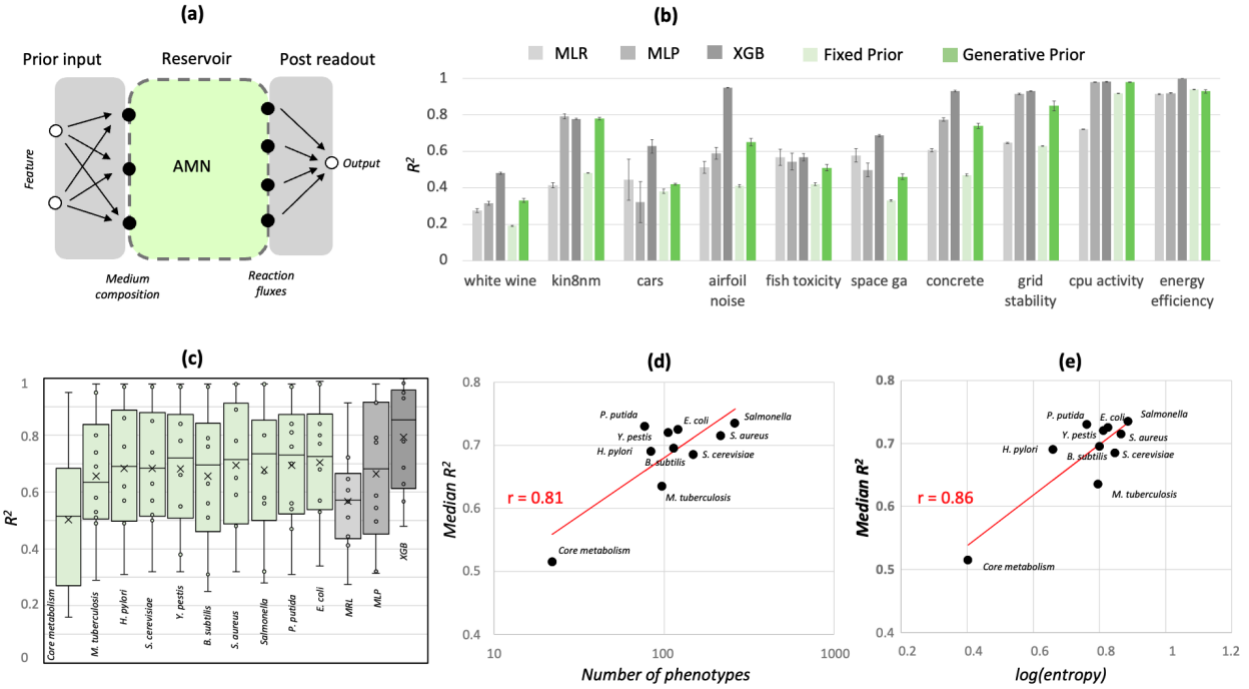

Fig. 2. Regression problem performances with conventional *E. coli* RC. (a) Schematic of the RC system architecture used for regression problems. The prior input layer maps problem features to medium composition, the reservoir (composed of a trained AMN) computes all reaction fluxes, and the post-readout layer converts reaction fluxes into problem readouts via a feedforward neural network. (b) Performance comparison of different regression models (MLR, MLP, XGB, and conventional RC with fixed and generative priors) across 10 problems from the OpenML platform (**table S1**). (c) Box plots for regression coefficients obtained for 10 bacterial conventional RC systems (**table S4**) with generative prior (**table S2**) and for the 10 regression problems of **table S1**. (d) Media regression coefficient (median R^2) vs. number of phenotypes when using generative prior. The Pearson correlation coefficient (r) between median R^2 values and the linear fit ($y = 0.19\ x + 0.31$) is 0.81. (f) Median R^2 vs. log(entropy) when using a generative prior. The Pearson correlation coefficient between median R^2 values and the linear fit ($y = 0.40\ x + 0.38$) is 0.86. Entropy calculation is given in supplementary materials **M11**. All R^2 values were acquired on 5-fold cross-validation sets for 3 repeats. Data are in the supplementary file 'Reservoir_Regression'. Results obtained with a fixed prior, the same 10 GEMs and 10 regression problems are given in **fig. S3**.

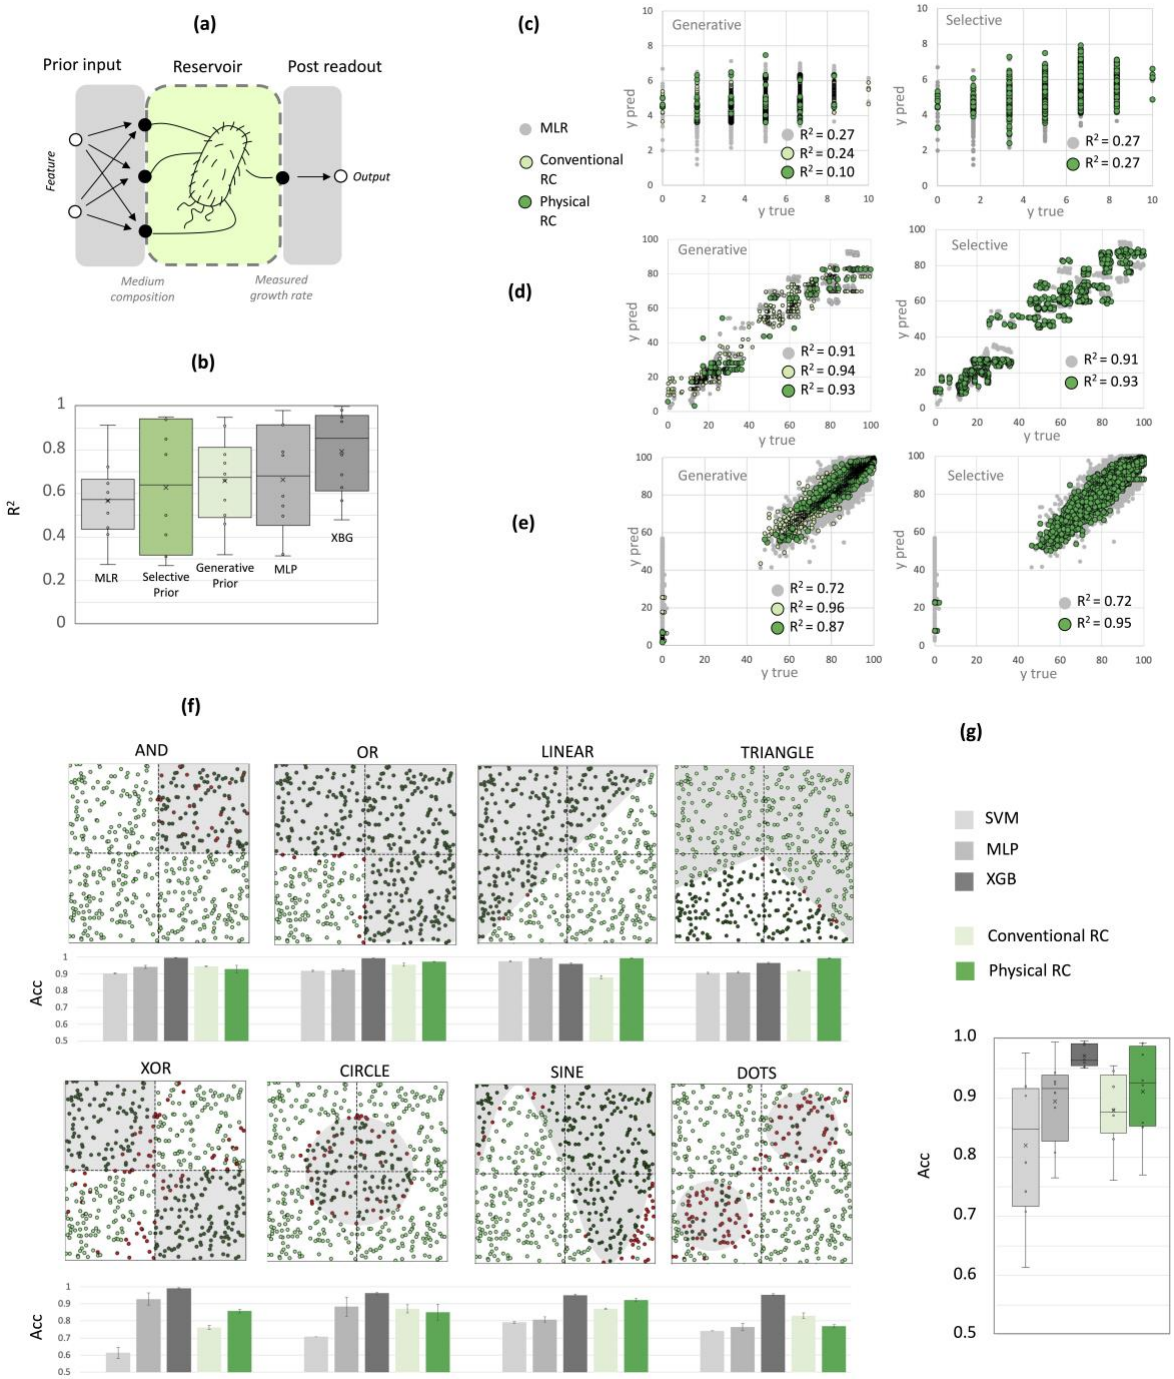

Fig. 3. Regression problem performances with physical *E. coli* RC. (a) Architecture of the physical RC system. (b) Box plot of R^2 values obtained for 10 regression problems of **Table S1** with MLR, MLP , XGB and conventional twin models with the selective and generative priors. Individual R^2 values for each of the 10 regression problems are given in **Fig. S4**. (c) Results obtained for the 'white wine' regression problem with MLR and conventional and physical RC making use of generative and selective priors. All physical RC results are derived from measured growth rates. With the generative prior 4898 values are predicted with conventional RC, 1639 of which corresponding to an experimentally measured growth rate (many values corresponding to the same growth rate). With the selective prior all values are derived from experimentally measured growth rates. (d) Results obtained for the 'energy efficiency' regression problem. With the generative prior 768 values are predicted with conventional RC, 162 of which corresponding to an experimentally measured growth rate. With the selective prior all values are derived from experimentally

measured growth rates. (e) Results obtained for the 'cpu activity' regression problem. With the generative prior 8192 values are predicted with conventional RC, 3057 of which corresponding to an experimentally measured growth rate. With the selective prior all values are derived from experimentally measured growth rates. Results presented in all panels were acquired on cross-validation sets (5-fold). Data are in supplementary file 'Reservoir_Regression'. (f) Performances for 4 linear (AND, OR, LINEAR, TRIANGLE) and 4 nonlinear (XOR, CIRCLE, SINE, DOTS) classification tasks with SVM, MLP and XGB classical classifiers along with *E. coli* conventional RC system with fixed prior and *E. coli* physical RC system with selective prior. Red dots are false positives and false negatives for the physical RC predictions. (g) Color codes and box plots of accuracies obtained for the 8 classification tasks of panel (f). Data are in supplementary file 'Reservoir_Classification'. Results using Mathews Regression coefficients (MCCs) and 26-fold cross validation as in Baltussen et al (11) are given in **fig. S5**.

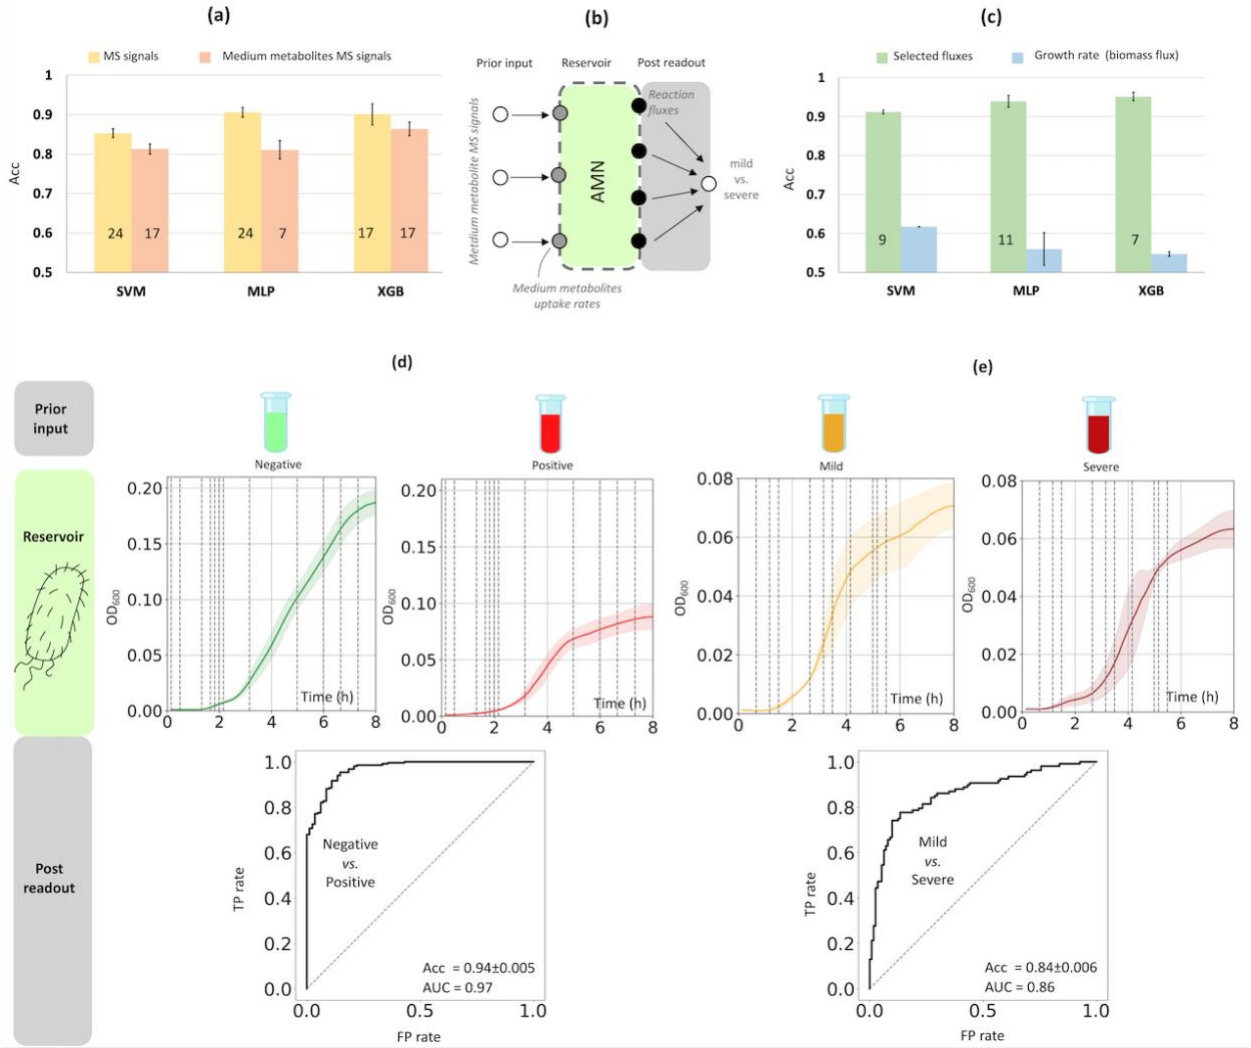

Fig. 4. Classifying COVID-19 plasma samples with metabolomics data and with *E. coli* conventional and physical reservoirs. (a) Classifying (mild vs. severe) COVID-19 with metabolomics data. MS signals: accuracies using all 624 metabolite signals detected by MS for 81 patients. Medium metabolites MS signals: accuracies using 56 MS signals for metabolites expected to cross the *E. coli* cell wall. Accuracies are calculated for 20-fold cross validation and using feature selection (supplementary materials **M9** and **fig. S6**). The numbers indicated on the bar charts are the numbers of features selected. The parameters used for the classifiers (SVM, XGB and MLP) are given in **table S3**. (b) The conventional RC takes as input medium metabolite MS signals, these are fed to a reservoir initially trained on FBA simulations with the GEM iML1515. The reservoir outputs the flux distribution over all reactions of the GEM, the distribution is fed to a post-readout performing a classification. This conventional RC is used to predict outcome (mild vs. severe) taking as input medium metabolite MS signals. (c) Classification is performed using all reaction fluxes (or just the reaction biomass flux) produced by the AMN reservoir of panel b when fed with the metabolite MS signals of the 81 COVID-19 patients. The numbers indicated on the bar charts are the numbers of features selected (supplementary materials **M9** and **fig. S7**). All data are in the supplementary files 'Covid_Data_MS_ML' and 'Reservoir_Covid'. (d) Classifying COVID-19 negative vs. positive plasma samples with *E. coli* physical reservoir. The prior input layer is a negative or positive plasma sample. The reservoir is MG1655 on which growth curves (2 technical replicates) are acquired using a plate reader recording OD_{600} at 48 time points taken from 0 to 8h (*cf.* supplementary materials **M15**). The readout layer is a XGB classifier that takes as input OD_{600} values for selected time points and prints out an accuracy for 20-fold cross validation repeated 3 times. The time points indicated by dashed vertical lines are the one selected among all time points in order to maximize accuracy

(supplementary materials **M9** and **Fig. S8**). Growth curves are shown for N1 from **fig. S9** and M10 from **Fig. S10**. (e) Classifying COVID-19 mild vs. severe plasma samples with *E. coli* physical reservoir. The set up is identical to panel d. Growth curves are shown for M9 from **fig. S10** and S8 from **fig. S11**. Growth curves for all samples are plotted in **figs. S8-S11** and OD_{600} values can be found in the supplementary file 'Reservoir_Covid'.