

Limits on the computational expressivity of non-equilibrium biophysical processes

Corresponding Author: Professor Suriyanarayanan Vaikuntanathan

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

Paper summary:

This paper explores the ability of biochemical networks modeled as Markov jump processes to perform complex pattern recognition. Specifically, it examines how input parameters encoded in the rate matrix of these networks influence their ability to classify chemical states. The authors identify fundamental thermodynamic constraints on such classification tasks. They also draw comparisons between the computations of Markov processes and transformer-like architectures, offering new perspectives on biochemical information processing.

Before addressing the main comments on the paper, I would like to draw the authors' attention to the unsatisfactory formatting of the supplementary materials. There are numerous broken references to equations and figures, which significantly hinder readability. These issues are present on pages 2, 7, 13, 16, 18, 20, and 21, as well as in the captions of Figures 1–4. I strongly recommend the authors carefully review and correct these errors to make the supplementary materials more accessible and easier to follow.

Comment 1.

In Section II, the authors introduce a simple example illustrating the challenge of pattern recognition, where the steady-state probability $\pi(F_a, F_b)$ depends on two control parameters F_a and F_b . Each parameter perturbs a single edge ($M=1$) of the graph. For successful pattern recognition, the resulting function must exhibit non-monotonicity with sign-alternating derivatives. However, in the case of $M=1$ it is impossible. The authors cite their prior work [26] to establish this limitation. While the explanation is clear and interesting, it is not an original contribution of the current paper.

Comment 2.

The authors introduce the hyperparameter M , which defines the number of edges perturbed by F_a . For M greater than 1, the system overcomes monotonic constraints, motivating the estimation of an upper bound on the number of derivative roots in Equation 4. However, this bound raises several concerns:

— **Relevance to Real Networks:** The bound is derived from the polynomial structure of the steady-state solution, but only the positive real roots are physically meaningful. For arbitrary polynomials, the actual number of such roots is often significantly lower than the theoretical bound. For instance, even with M equal to 1, the bound predicts one root, while the true value is zero. The polynomial form does not necessarily reflect the physical constraints or typical behavior of biochemical networks. For notable M , achieving R values comparable to M requires specific forms (for example, orthogonal polynomials), and it is unclear whether the Arrhenius-like parameterization supports such forms.

— **Lack of Numerical Evidence:** The authors briefly discuss the saturation of the bound in the supplementary materials but fail to include numerical results demonstrating this. The main text does not address the reliability of these bounds.

Comment 3.

The authors draw an analogy between the computations of Markov jump processes and transformer-like architectures, citing similarities such as the Boltzmann (softmax) form in Equation 2 and the parameter factorization in $\psi(\theta)$ and $\chi(F)$. However, this comparison feels superficial. The Boltzmann form is a general feature and could be equally well linked to classical techniques like kernel regression (for example, see Kernel Regression on Wikipedia).

Moreover, the analogy is discussed only briefly in a paragraph below Equation 110, and the form of Equation 110 appears to differ from the reference cited as 23 in the supplementary material.

Suggestions:

- Provide a systematic and detailed analysis of the connection to transformers, rather than relying on their popularity.
- Clarify the differences between Equation 110 and the cited reference.
- Consider comparing the Boltzmann form to other classical methods to situate the contribution more broadly.

Comment 4:

The statement that "under the linear dynamics of Markov jump processes it is necessary to break detailed balance to perform non-trivial computations" is compelling but based solely on numerical simulations. While the result is intriguing, it lacks rigorous theoretical backing.

Suggestions:

- Provide a theoretical explanation or proof for this observation.
- If theoretical support is not feasible, include a broader range of numerical examples to strengthen the claim.

Comment 5:

The presentation of Equation 8 is unclear. In the main text, it is written for any F , but the supplementary material indicates that the systematic analysis is conducted only for F equal to 0. This discrepancy suggests a potential typo in Equation 8, which might need correction.

Minor: Equation 87 appears to be incorrect, as the term m' should also be inside the exponential. For the case of $F = 0$, this issue does not affect the subsequent steps, but it should still be corrected for consistency and clarity.

Conclusion:

Overall, the paper addresses an interesting and important topic, providing new insights into the computational capabilities of biochemical networks. However, it would benefit from more rigorous analysis, clearer analogies, and additional numerical validation.

(Remarks on code availability)

Reviewer #2

(Remarks to the Author)

The authors study limitations and potentialities of pseudo-linear chemical networks to classify external chemical inputs of different kinds. The paper is mostly clear, the goal is well motivated and the methodology to address relevant issues is properly described.

The idea of mapping chemical systems to trained networks to study their ability to process information (in general terms) is a field of growing interest and this work provides a systematic study of this approach for a limited, yet quite relevant, class of chemical networks. I found the paper interesting and potentially inspiring for subsequent studies.

Below, I list some comments on this work:

i) A very recent work [1] addresses similar problems to the presented in this manuscript, focusing mostly on how chemical systems (not limited to pseudo-linear ones) can reconstruct arbitrary dynamics. A discussion linking this study to the aforementioned one might be relevant to contextualise the work and emphasise the importance of nonlinear reaction schemes. Along the same line, it would be nice to numerically show a nonlinear version of one of the model presented in the paper, without chemostatting all enzymes or adding intermediate non-linear steps. I understand that the argument proposed by the authors hold only when a spanning tree decomposition is possible, but the interest in studying nonlinear networks is increasing in the field. A proof that qualitatively the results hold when some nonlinearities appear in the reaction schemes might substantiate the results (also in the light of the results in [1]). The authors themselves mention that nonlinear networks can be expressive even in equilibrium conditions, citing two papers. Substantiating this statement within the proposed framework might be an important addition.

ii) The authors put emphasis on the role of input multiplicity to enable nonlinear classification or discrimination of multiple classes. Are there any ways to find the optimal or minimal M for encoding a given structure? It is a common situation in chemical systems that how many pathways are influenced is not known nor tunable. Knowing the optimal or minimal way to encode given features might give hints on the inner structure of chemical systems. At the same time, when the network is known, revealing crucial edges for encoding might be relevant to design targeted inputs. Could the authors elaborate on this aspect?

iii) Section IV is not very clear. The intuitive idea of counting degrees of freedom and constraint is neat, but it is hard to follow the details without taking a in-depth look at the Supplemental Material. Consider reducing it including only the intuitions or extending it explaining more technical details.

iv) The authors mention that the steady state of trained networks is dominated by specific pathways? However, even from the Supplemental Material, it is unclear whether this property has been thoroughly explored or merely observed in isolated cases. If the second option is correct, the text in the main is misleading on the generality of this statement. Also, how the property of the dominating tree are connected with nonequilibrium properties? I do not expect these pathways to have a trivial connection with nonequilibrium.

Along this line, do these pathways reflect the importance of multiple serial intermediates? In connection with the point above, is it also important where these intermediates are, not only the fact that there exists a pathway with multiple serial intermediates?

Are there any minimal motifs allowing for accurate (and sensitive) responses?

v) What happens when F_{ij} are perturbed but the underlying networks is at equilibrium but can be trained only changing barriers? I would expect worse performance since the network is not intrinsically out-of-equilibrium before the external perturbation on the asymmetric part of the rates. On this point, a recent work [2] discussed the flexibility of a chemical systems in and out of equilibrium in the case of pseudo-linear networks. Is there any connection between the authors' results on this aspect and the ability to select and amplify chemical species by increasing driving?

vi) Can the authors explain better the argument on sharpness bounds in the paragraph before the discussion section?

Another minor point. Some references to figures, articles, and formulas in the Supplemental Material appear as '?'.

Refs.

[1] Dack, et al. arXiv:2406.03456 (2024)

[2] Liang, et al. PRL 132, 228402 (2024)

(Remarks on code availability)

Reviewer #3

(Remarks to the Author)

Manuscript "Limits on the computational expressivity of non-equilibrium biophysical processes" studies non-equilibrium biomolecular systems that are able to classify high-dimensional chemical systems. The authors combine theoretical tools from statistical physics, graph theory and machine learning to derive novel and general limits to the performance of these systems.

This manuscript presents very elegant, rigorous, and important results. Limits to the performance of these systems have been studied for a long time. However, such general results are usually available only in the equilibrium case. In the non-equilibrium case, our understanding has been mostly informed by specific models (such as Hopfield's proofreading model). By providing general results in the non-equilibrium case, the manuscript substantially advance our understanding of these processes.

My only concerns on the manuscript are about its presentation style. I feel that the authors could organize and present their results in a better way, with the broad audience of Nature Communication in mind. The manuscript is overall well written, but gives several concepts and ideas for granted. I encountered this problem both with the introduction of biological examples and theoretical tools.

For example, the "glycan code" system seems to be a main motivation, being mentioned in the first sentence of the abstract and at the beginning of the introduction. However, the authors actually describe the functioning of this system only in the caption of Fig. 1. I was not familiar with this system, and I found the description in the caption too brief. It would be great if the authors could expand on this example (and possibly others), introducing them in more detail and concretely discussing the implications of their theoretical results for these specific systems.

The presentation of the technical results could be improved as well. For example, to avoid making the manuscript too technical, the authors graciously describe their main result into a Box. However, sections II and III of the manuscript are quite hard to follow without having studied the box, as they repeatedly refer to quantities that are introduced there. This fact seems to defy the scope of having a box. In general, I think the authors should make an effort to make their result sections self-contained, and understandable to a reasonably broad audience of quantitative biologists/theoretical biophysicists.

(Remarks on code availability)

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

In the revised version, the authors have made substantial changes in response to our comments, including additional numerical calculations and new results for the mutual information. We are believe that the paper is now suitable for publication in Nature Communications.

(Remarks on code availability)

Reviewer #2

(Remarks to the Author)

The authors have addressed most of my concerns, and I particularly appreciated the new section dedicated to the maximization of mutual information. Regarding nonlinearity, I understand the challenges in dealing with general nonlinear systems. However, I found the discussion in the Supplementary Information not completely satisfactory on this aspect. I would have preferred a fully numerical analysis of a nonlinear chemical network, without additions that appear somewhat ad hoc, e.g., unidirectional nonlinearities. At any rate, this might just be a matter of taste and does not affect the validity of the paper's results. Therefore, I suggest this paper for publication in Nature Communications.

(Remarks on code availability)

Reviewer #3

(Remarks to the Author)

The authors successfully addressed my concerns, as well as those raised by other reviewers. I am happy to support publication of the manuscript as it is.

(Remarks on code availability)

Reviewer #4

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

(Remarks on code availability)

Open Access This Peer Review File is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source.

The images or other third party material in this Peer Review File are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

We wish to sincerely thank the reviewers for their helpful comments on the manuscript. In response to their suggestions, we have added further discussion of key points and refined descriptions of our results. We have also included several new results in the main text and the Supplementary Material. In the updated main text and supplement, new or revised text is shown in blue. The major changes are summarized below, and a point-by-point response follows.

List of changes:

- New analysis on optimizing mutual information between network steady states and input distributions, showing that greater input multiplicity M enables more mutual information to be transmitted (including main text section [See Fig 6] and Supplementary Material section with figure)
- New section in the Supplementary Material Section X (and SI Fig 14) on how bimolecular (non-linear) reactions in the chemical dynamics can increase expressivity.
- Revised derivation of the bound on the number of turning points as a function of M
- New numerical verification on bound of number of turning points as a function of M
- Updated SI Figure 8 to show differences in tree activation clustering between trained and untrained random networks
- Revised and expanded text throughout to improve clarity and better contextualize our findings
- Removed the Box (as requested by the Editorial staff and one of the referees) and placed it into the main text (now shown in red text)
- Fixed several typos and broken references throughout

Reviewer #1 (Remarks to the Author):

Paper summary:

This paper explores the ability of biochemical networks modeled as Markov jump processes to perform complex pattern recognition. Specifically, it examines how input parameters encoded in the rate matrix of these networks influence their ability to classify chemical states. The authors identify fundamental thermodynamic constraints on such classification tasks. They also draw comparisons between the computations of Markov processes and transformer-like architectures, offering new perspectives on biochemical information processing.

Before addressing the main comments on the paper, I would like to draw the authors' attention to the unsatisfactory formatting of the supplementary materials. There are numerous broken references to equations and figures, which significantly hinder readability. These issues are present on pages 2, 7, 13, 16, 18, 20, and 21, as well as in the captions of Figures 1–4. I

strongly recommend the authors carefully review and correct these errors to make the supplementary materials more accessible and easier to follow.

Thank you for raising this issue. We have corrected the broken references in the revised manuscript and apologize for any confusion resulting from those errors in the previous draft.

Comment 1.

In Section II, the authors introduce a simple example illustrating the challenge of pattern recognition, where the steady-state probability $\pi(F_a, F_b)$ depends on two control parameters F_a and F_b . Each parameter perturbs a single edge ($M=1$) of the graph. For successful pattern recognition, the resulting function must exhibit non-monotonicity with sign-alternating derivatives. However, in the case of $M=1$ it is impossible. The authors cite their prior work [26] to establish this limitation. While the explanation is clear and interesting, it is not an original contribution of the current paper.

We thank the referee for bringing up this question. It is true that Ref. 26 contains the original derivation of the monotonicity of the response $\pi(F_a)$, which serves as a foundational result for this manuscript. As the referee notes, we do not claim this result as a novel contribution of the current work, and we cite it appropriately in the text. **Rather, this paper builds significantly upon it by interpreting the monotonicity constraint as a fundamental limitation on computational expressivity. Further, the present work identifies and explores how input multiplicity (M) can enable greater computational expressivity in such systems.**

Identifying how input multiplicity enables flexible contours, greater multi-class capacity, improved transmission of mutual information (see the new results), and potential applicability to more complex reaction dynamics (with an example in the new results in the Supplemental information and also in the summary of changes above) is new progress. We anticipate that in the future these findings can help interface with recent experiments in synthetic biology such as <https://www.science.org/doi/10.1126/science.add8468>. These results go well beyond what is in Ref. 26, which instead focuses on how dynamical control can arise in non-equilibrium systems due to structural constraints like monotonicity.

Comment 2.

The authors introduce the hyperparameter M , which defines the number of edges perturbed by F_a . For M greater than 1, the system overcomes monotonic constraints, motivating the estimation of an upper bound on the number of derivative roots in Equation 4. However, this bound raises several concerns:

— Relevance to Real Networks: The bound is derived from the polynomial structure of the

steady-state solution, but only the positive real roots are physically meaningful. For arbitrary polynomials, the actual number of such roots is often significantly lower than the theoretical bound. For instance, even with M equal to 1, the bound predicts one root, while the true value is zero. The polynomial form does not necessarily reflect the physical constraints or typical behavior of biochemical networks. For notable M , achieving R values comparable to M requires specific forms (for example, orthogonal polynomials), and it is unclear whether the Arrhenius-like parameterization supports such forms.

We appreciate the reviewer's concern regarding our previous presentation of the bound on the number of turning points as a function of M . **In response, we have refined the derivation of this bound and revised our discussion accordingly in both the main text and the supplement.** We now refer to a previous study (Ref. 39), which proves that the maximum number of positive critical points for a rational function $f(y)/g(y)$, where f and g are polynomials of degrees n and m with non-negative coefficients, is $\min(n, m)$ if $n \neq m$, and $m - 1$ if $n = m$. This result directly applies to the matrix-tree theorem expression for the steady state, which is a rational function with both numerator and denominator of degree $2M$, yielding a bound of $2M - 1$ for $M > 1$. For $M = 1$, our additional monotonicity constraint prevents the number of turning points from reaching the bound of 1 predicted by the formula. As the reviewer correctly noted, this bound pertains only to positive roots, which are the physically relevant solutions in our setting. **Importantly, we (numerically) further demonstrate that this bound can indeed be saturated by biochemical networks with Arrhenius-like parameterizations (please see answer to the next question).**

— Lack of Numerical Evidence: The authors briefly discuss the saturation of the bound in the supplementary materials but fail to include numerical results demonstrating this. The main text does not address the reliability of these bounds.

This is a helpful suggestion. In our previous draft, we stated that we observed saturation of the predicted bound up to $M = 3$ but did not provide graphical evidence. **We have now added SI Figure 3, in which we extensively sample random edge weights in a fully connected graph with six nodes for values up to $M = 4$. We plot histograms of the observed number of turning points at each node over 10,000 samples. The results verify the predicted scaling: for $M = 1$, there are no observed turning points in any sampled graph, while for $M > 1$, the number of turning points is upper bounded by $2M - 1$.** Extending the numerical evidence beyond $M = 4$ is computationally challenging because graphs that saturate the bound are statistically rare, but we have no reason to doubt that the validity of the bound holds for arbitrary M . We thank the referee again for this suggestion and hope they find the numerical evidence now presented convincing.

Comment 3.

The authors draw an analogy between the computations of Markov jump processes and transformer-like architectures, citing similarities such as the Boltzmann (softmax) form in Equation 2 and the parameter factorization in $\psi(\theta)$ and $\chi(F)$. However, this comparison feels superficial. The Boltzmann form is a general feature and could be equally well linked to classical techniques like kernel regression (for example, see Kernel Regression on Wikipedia).

Moreover, the analogy is discussed only briefly in a paragraph below Equation 110, and the form of Equation 110 appears to differ from the reference cited as 23 in the supplementary material.

Suggestions:

- Provide a systematic and detailed analysis of the connection to transformers, rather than relying on their popularity.
- Clarify the differences between Equation 110 and the cited reference.
- Consider comparing the Boltzmann form to other classical methods to situate the contribution more broadly.

We have revised our discussion of the comparison to transformers in both the main text and SI Section VII. **Specifically, we followed the reviewer's helpful suggestion to highlight the similarity between the matrix-tree theorem expression and kernel-based classification architectures such as support vector machines.** We also emphasize that the comparison to Ref. 110 is meant to be primarily qualitative. While the nonlinear feature maps $\psi(i, \theta)$ and $\chi(i, F)$ bear conceptual similarities to those in transformer architectures, we do not propose an exact mapping. Instead, we hope to encourage future work exploring the unique computational properties of classification functions that emerge naturally from Markov processes, possibly guided by ideas and metrics developed in the context of transformers.

Comment 4:

The statement that "under the linear dynamics of Markov jump processes it is necessary to break detailed balance to perform non-trivial computations" is compelling but based solely on numerical simulations. While the result is intriguing, it lacks rigorous theoretical backing.

Suggestions:

- Provide a theoretical explanation or proof for this observation.
- If theoretical support is not feasible, include a broader range of numerical examples to strengthen the claim.

We thank the referee for bringing up this important point. In our previous version, we intended the supplementary section titled “How non-equilibrium driving allows expressivity” to provide a theoretical explanation for this observation. In that section, we show that the non-equilibrium parameters F_{ij} provide the greatest flexibility for positioning the learnable feature vectors $\psi(i, \theta)$ in the high-dimensional tree space. In particular, non-zero values of F_{ij} allow these vectors to differ anisotropically across nodes i . We also include several case studies to illustrate this conceptual framework. In our revisions, we have added clearer references to the Supplementary Material for readers seeking further detail.

Comment 5:

The presentation of Equation 8 is unclear. In the main text, it is written for any F , but the supplementary material indicates that the systematic analysis is conducted only for F equal to 0. This discrepancy suggests a potential typo in Equation 8, which might need correction.

Thank you for pointing this out. In our derivation of the maximum in the supplement, we evaluate the derivative at $F=0$ but we note that the evaluation point can be shifted to any arbitrary F_0 by absorbing a corresponding factor into the definition of the ζ parameters, which are the maximization variables. Therefore, the precise location at which the derivative is evaluated is not essential for establishing the maximum. In the main text, we now clarify that Equation 8 should be evaluated at the location of the decision boundary to inform on the sharpness of the transition between class regions.

Minor: Equation 87 appears to be incorrect, as the term m' should also be inside the exponential. For the case of $F = 0$, this issue does not affect the subsequent steps, but it should still be corrected for consistency and clarity.

Thank you for this correction - the exponent should indeed include m' . We have fixed it.

Conclusion:

Overall, the paper addresses an interesting and important topic, providing new insights into the computational capabilities of biochemical networks. However, it would benefit from more rigorous analysis, clearer analogies, and additional numerical validation.

We appreciate the reviewer's close reading of our manuscript and their helpful and clear suggestions for improvements. As described above we have made many edits and performed new calculations in response to their suggestions.

Reviewer #2 (Remarks to the Author):

The authors study limitations and potentialities of pseudo-linear chemical networks to classify external chemical inputs of different kinds. The paper is mostly clear, the goal is well motivated and the methodology to address relevant issues is properly described.

The idea of mapping chemical systems to trained networks to study their ability to process information (in general terms) is a field of growing interest and this work provides a systematic study of this approach for a limited, yet quite relevant, class of chemical networks. I found the paper interesting and potentially inspiring for subsequent studies.

Below, I list some comments on this work:

i) A very recent work [1] addresses similar problems to the presented in this manuscript, focusing mostly on how chemical systems (not limited to pseudo-linear ones) can reconstruct arbitrary dynamics. A discussion linking this study to the aforementioned one might be relevant to contextualise the work and emphasise the importance of nonlinear reaction schemes. Along the same line, it would be nice to numerically show a nonlinear version of one of the model presented in the paper, without chemostatting all enzymes or adding intermediate non-linear steps. I understand that the argument proposed by the authors hold only when a spanning tree decomposition is possible, but the interest in studying nonlinear networks is increasing in the field. A proof that qualitatively the results hold when some nonlinearities appear in the reaction schemes might substantiate the results (also in the light of the results in [1]). The authors themselves mention that nonlinear networks can be expressive even in equilibrium conditions, citing two papers. Substantiating this statement within the proposed framework might be an important addition.

Thank you for these comments. The provided reference is indeed relevant in illustrating how computational power can be systematically improved in chemical reaction networks, and we have included a brief discussion of this work in the Discussion section. We have also expanded on this point by referencing a recent study that builds reservoir computers from chemical reaction networks: <https://www.nature.com/articles/s41586-024-07567-x>.

In addition, we followed the reviewer's suggestion to include a preliminary treatment of bimolecular reactions within our current computational framework (see SI Section X and SI Fig 14). As the reviewer noted, non-linear dynamics preclude direct application of the matrix-tree theorem (MTT) to solve for the system's steady state. However, we outline an approach that uses the MTT to study steady-state properties even in the presence of non-linearities, by self-consistently solving for the steady-state occupancies that appear on both sides of the MTT expression. The more complex algebraic dependence of the steady state on the input driving in these non-linear systems can be interpreted as "effectively" increasing the value of M , which controls the number of ways each input influences the steady-state response. We find that including non-linearities enables non-monotonic response functions. Our analysis further implies

that, as in linear systems, applying inputs more than once in non-linear systems increases expressivity by raising the order of the polynomials appearing in the MTT expression.

ii) The authors put emphasis on the role of input multiplicity to enable nonlinear classification or discrimination of multiple classes. Are there any ways to find the optimal or minimal M for encoding a given structure? It is a common situation in chemical systems that how many pathways are influenced is not known nor tunable. Knowing the optimal or minimal way to encode given features might give hints on the inner structure of chemical systems. At the same time, when the network is known, revealing crucial edges for encoding might be relevant to design targeted inputs. Could the authors elaborate on this aspect?

This is an interesting and important, yet difficult, question. We believe that a fully topological answer, one that rigorously characterizes how to design network structures that optimally couple input driving to output responses, lies outside the scope of the current work as it would require a complex analysis of its own.

However, to explore this issue from a different angle (without explicitly addressing structural circuit design), **we have added a new analysis in which we study expressivity without assuming a specific encoding scheme for the output.** Throughout the paper, we primarily examine a one-hot encoding scheme, where the choice of the output node, and its structural relationship to the input driving edges, remains a relatively unexplored factor influencing classification performance. In our new analysis (**see Fig 6 of new main text**), we instead consider a more general task: optimizing the mutual information between the full steady-state distribution and the input distribution, which does not require pre-selecting output nodes. Interestingly, we find that the network autonomously selects nodes that serve as effective readouts of distinct regions of the input space, essentially reconstructing a one-hot encoding scheme. This self-organized output selection highlights how the system naturally identifies nodes best coupled to the input driving. Moreover, we observe that increasing M leads to higher optimized mutual information, consistent with our previous findings under one-hot encoding and supporting the broader conclusion that expressivity increases with M . We anticipate that future work could extend this training approach to uncover deeper topological principles governing the coupling between input and output nodes.

iii) Section IV is not very clear. The intuitive idea of counting degrees of freedom and constraint is neat, but it is hard to follow the details without taking a in-depth look at the Supplemental Material. Consider reducing it including only the intuitions or extending it explaining more technical details.

We have rewritten parts of this section to improve clarity and accessibility. We acknowledge that the argument is technical and challenging to present in a condensed format. To address this, we

aimed to describe it in the main text at a more intuitive level—focusing on the scaling relationship between the number of conditions that must be satisfied for multi-class classification and the number of degrees of freedom available to satisfy them. Our goal is for this framing to convey the essential intuition behind the result, leaving the details of counting constraints and inequality conditions to the supplement.

iv) The authors mention that the steady state of trained networks is dominated by specific pathways? However, even from the Supplemental Material, it is unclear whether this property has been thoroughly explored or merely observed in isolated cases. If the second option is correct, the text in the main is misleading on the generality of this statement.

This is a helpful point. **We have repeated our analysis of low-dimensional structure in the tree activation space using a different graph, now presented in the new SI Figure 8 alongside the original results in SI Figure 7.** Unlike the fully connected graph used previously, this new graph is sparse and randomly generated. To better understand the clustering of tree activations, we have also augmented both SI Figures 7 and 8 to compare the trained and untrained cases directly. We find that input driving alone can account for much of the observed clustering in tree activations. As before, we observe that training reduces the inverse participation ratio in the sparse random graph, indicating that the network response becomes more structured and non-uniform. However, unlike in the fully connected graph, we do not find clear evidence that training enhances clustering of tree activations by input class in the sparse random case. We hypothesize that this may be due to the significantly smaller number of trees in the sparse random graph (284 vs. 1,296), which could make clustering more difficult. We have revised the main text to clarify these findings and to emphasize that these two examples serve as case studies that should be extended in future work to provide a more systematic picture.

Also, how the property of the dominating tree are connected with nonequilibrium properties? I do not expect these pathways to have a trivial connection with nonequilibrium. Along this line, do these pathways reflect the importance of multiple serial intermediates? In connection with the point above, is it also important where these intermediates are, not only the fact that there exists a pathway with multiple serial intermediates?

Are there any minimal motifs allowing for accurate (and sensitive) responses?

As with the other topology-focused question above, we believe that, while very interesting, these questions warrant a dedicated treatment in future work. Addressing them fully would require detailed analyses that goes beyond the current scope. We note, however, that our result on the effectiveness of serial, but not parallel, implementations of extra driven edges in sharpening transitions does indeed point to the importance of chemical intermediates in shaping the non-equilibrium response.

We have added text in the discussion pointing to these questions as a promising direction for follow-up studies. We have also included a reference to a potentially relevant concept from Hopfield-like associative memory models, namely the “feature-to-prototype” transitions, in which the emergence of low-dimensional collective modes is associated with improved classification accuracy. This connection may provide a useful framework for exploring how topological or structural features give rise to expressive and efficient representations.

v) What happens when F_{ij} are perturbed but the underlying networks is at equilibrium but can be trained only changing barriers? I would expect worse performance since the network is not intrinsically out-of-equilibrium before the external perturbation on the asymmetric part of the rates. On this point, a recent work [2] discussed the flexibility of a chemical systems in and out of equilibrium in the case of pseudo-linear networks. Is there any connection between the authors' results on this aspect and the ability to select and amplify chemical species by increasing driving?

When inputs are presented through changes to the F parameters but the underlying network does not break detailed balance, computations are still possible, though in a more limited capacity. We explore this scenario in SI Figure 9. While the absence of non-equilibrium driving in the network prevents the learnable $\psi(i, \theta)$ vectors from differing anisotropically across nodes i , the $\chi(i, F)$ data vectors are still positioned in distinct regions of tree space due to the non-equilibrium nature of the input. This spatial separation generally allows the $\psi(i, \theta)$ vectors to be learned such that they overlap appropriately with their corresponding $\chi(i, F)$ vectors.

The observed connection to Ref. [2] is helpful. That work and recent works by Horowitz, Esposito, Busiello, Lin, Polettini, and others have investigated bounds on the magnitude of non-equilibrium responses in terms of both equilibrium and non-equilibrium network parameters. In our previous draft, we cited several of these studies in relation to our results on the sharpness of transitions between classification regions, where we believe there is a clear conceptual link. However, in the interest of space, we have chosen not to elaborate on the connections to these works individually in the main text. We hope that future efforts will integrate our findings on expressivity, turning points, and computational architectures within Markov networks with these recent thermodynamic studies of non-equilibrium response magnitudes.

vi) Can the authors explain better the argument on sharpness bounds in the paragraph before the discussion section?

Thank you for pointing out that this section was not sufficiently clear. The prevented saturation of the predicted bound on sharpness due to network entanglement is a difficult technical point, which we discuss in greater detail in the Supplementary Material. In response, we have rewritten that paragraph in the main text to hopefully convey more clearly the central idea and

provide better guidance for readers seeking the full explanation.

Another minor point. Some references to figures, articles, and formulas in the Supplemental Material appear as '?'.
Thank you for raising this issue. We have corrected the broken citations and references in our revised draft.

Refs.

[1] Dack, et al. arXiv:2406.03456 (2024)

[2] Liang, et al. PRL 132, 228402 (2024)

Reviewer #3 (Remarks to the Author):

Manuscript "Limits on the computational expressivity of non-equilibrium biophysical processes" studies non-equilibrium biomolecular systems that are able to classify high-dimensional chemical systems. The authors combine theoretical tools from statistical physics, graph theory and machine learning to derive novel and general limits to the performance of these systems.

This manuscript presents very elegant, rigorous, and important results. Limits to the performance of these systems have been studied for a long time. However, such general results are usually available only in the equilibrium case. In the non-equilibrium case, our understanding has been mostly informed by specific models (such as Hopfield's proofreading model). By providing general results in the non-equilibrium case, the manuscript substantially advance our understanding of these processes.

My only concerns on the manuscript are about its presentation style. I feel that the authors could organize and present their results in a better way, with the broad audience of Nature Communication in mind. The manuscript is overall well written, but gives several concepts and ideas for granted. I encountered this problem both with the introduction of biological examples and theoretical tools.

For example, the "glycan code" system seems to be a main motivation, being mentioned in the first sentence of the abstract and at the beginning of the introduction. However, the authors actually describe the functioning of this system only in the caption of Fig. 1. I was not familiar with this system, and I found the description in the caption too brief. It would be great if the authors could expand on this example (and possibly others), introducing them in more detail and concretely discussing the implications of their theoretical results for these specific systems.

We are very grateful to the reviewer for their reading of our manuscript and their positive feedback, as well as for the helpful suggestion to expand the background material to make the

work more accessible to a broader audience. While space constraints prevent detailed discussions of specific systems such as the glycan code, we have added new text to better introduce the glycan system and to clarify how it motivates our study of the classification capacity of chemical reaction networks. We have also expanded the discussion section and added references throughout the main text to more explicitly connect our results to the glycan code. In particular, we highlight several studies emphasizing the high promiscuity of both glycan biosynthetic enzymes and glycan-receptor interactions—features which, based on our findings, may be important for the computational capabilities of this biochemical system.

The presentation of the technical results could be improved as well. For example, to avoid making the manuscript too technical, the authors graciously describe their main result into a Box. However, sections II and III of the manuscript are quite hard to follow without having studied the box, as they repeatedly refer to quantities that are introduced there. This fact seems to defy the scope of having a box. In general, I think the authors should make an effort to make their result sections self-contained, and understandable to a reasonably broad audience of quantitative biologists/theoretical biophysicists.

We acknowledge the difficulty in balancing the technical detail of some of our results with their intuitive explanation in the main text. We have eliminated the Box and replaced it with a section to be read after the introduction, as we agree that the remaining results cannot be well understood without that background material. We have also edited the rest of the text in line with the reviewer's suggestion to make the key results as transparent as possible.

We wish to sincerely thank the reviewers for their helpful comments on the manuscript. We are glad the paper has been accepted for publication in principle.

Reviewer #1 (Remarks to the Author):

In the revised version, the authors have made substantial changes in response to our comments, including additional numerical calculations and new results for the mutual information. We believe that the paper is now suitable for publication in Nature Communications.

We thank the reviewer for their feedback.

Reviewer #2 (Remarks to the Author):

The authors have addressed most of my concerns, and I particularly appreciated the new section dedicated to the maximization of mutual information. Regarding nonlinearity, I understand the challenges in dealing with general nonlinear systems. However, I found the discussion in the Supplementary Information not completely satisfactory on this aspect. I would have preferred a fully numerical analysis of a nonlinear chemical network, without additions that appear somewhat ad hoc, e.g., unidirectional nonlinearities. At any rate, this might just be a matter of taste and does not affect the validity of the paper's results. Therefore, I suggest this paper for publication in Nature Communications.

We thank the reviewer for their feedback. In future work we plan to explore the capacity of nonlinear reaction networks in line with the reviewer's suggestion.

Reviewer #3 (Remarks to the Author):

The authors successfully addressed my concerns, as well as those raised by other reviewers. I am happy to support publication of the manuscript as it is.

We thank the reviewer for their feedback.

Reviewer #4 (Remarks to the Author):

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

We thank the reviewer for their feedback.