

Reply to: Limitations in odour recognition and generalization in a neuromorphic olfactory circuit

Received: 21 February 2024

Roy Moyal¹, Nabil Imam² & Thomas A. Cleland¹ 

Accepted: 6 November 2024

Published online: 16 December 2024

REPLYING TO N. Dennler et al. *Nature Machine Intelligence*

 Check for updates

<https://doi.org/10.1038/s42256-024-00952-1> (2024)

In their Comment, Dennler et al.¹ submit that they have discovered limitations affecting some of the conclusions drawn in our 2020 paper, 'Rapid online learning and robust recall in a neuromorphic olfactory circuit'². Specifically, they assert (1) that the public dataset we used suffers from sensor drift and a non-randomized measurement protocol, (2) that our neuromorphic external plexiform layer (EPL) network is limited in its ability to generalize over repeated presentations of an odourant, and (3) that our EPL network results can be performance matched by using a more computationally efficient distance measure. Although they are correct in their description of the limitations of that public dataset³, they do not acknowledge in their first two assertions how our utilization of those data sidestepped these limitations. Their third claim arises from flaws in the method used to generate their distance measure. We respond below to each of these three claims in turn.

Metal oxide sensor drift

We utilized samples of odour responses drawn from an array of 72 metal oxide chemosensors spatially dispersed across a wind tunnel⁴. The intended sources of variance in this public dataset include plume dynamics and wide, largely unpredictable variations in odourant concentration at specific sensor sites. Moreover, as noted by Dennler et al., there also are unintended sources of variance embedded in that dataset. To wit, the rate of drift in the response profiles of these metal oxide sensors generated cumulative changes in sensor responses over the course of data acquisition that are comparable to or greater than the odourant-specific differences in sensor responses on which odourant recognition depends. The fact that these data were acquired in sequential, odourant-specific batches over the course of 8 months ensured that these drift-based changes in sensor response profiles would be conflated with the odourant-based differences in sensor responses. Dennler et al. illustrate this with an experiment (their Fig. 1b and Supplementary Fig. 1b) in which they identify odourants based solely on the sensor drift that had accumulated by the day on which a particular odourant was tested.

That said, this problem is not relevant to our findings. We explicitly set aside the problem of metal oxide sensor drift by randomly selecting single samples from each of the ten odourant responses and occluding them with our own noise models designed to mimic interference by environmental background odourants². While the rate of baseline drift and decay is an existential threat to the utility of metal oxide-based chemosensor devices, the problem ultimately is associated with these particular materials⁵ and can be foreseeably resolved by using different chemosensor technologies^{6–10}. Accordingly, we chose to sidestep this metal oxide-associated problem and instead study the broader problem of identifying odour sources of interest in the presence of unpredictable competitive interference, predicated on reasonably stable sensor responses.

Generalization

Dennler et al. acknowledge that the EPL network convincingly restores input patterns corrupted by impulse noise (Figs. 3–5 in ref. 2). Their broad conclusion that the model does not generalize is incorrect; signal restoration is a form of generalization and comprises the central message of the paper. Generalization to different samples within bounds also is shown in Fig. 5 of ref. 2. What they do correctly point out is that a learned odourant representation does not generalize to a separate presentation of the same odourant delivered into the wind tunnel. This requires explanation.

The same odourant, sampled after separate deliveries to these 72 spatially dispersed sensors, would be encountered by each individual sensor at substantially different concentrations according to the random dispersion patterns of each plume. It is well established that concentration effects are more powerful than odourant quality-based differences in their ability to drive metal oxide sensor responses. Accordingly, their test equates to the addition of high levels of random concentration-based noise to every sensor—high enough to disrupt all odourant-diagnostic pattern information. It is no surprise that Dennler et al. failed to identify the test odourant under these conditions.

¹Computational Physiology Lab, Department of Psychology, Cornell University, Ithaca, NY, USA. ²College of Computing, Georgia Institute of Technology, Atlanta, GA, USA.  e-mail: tac29@cornell.edu

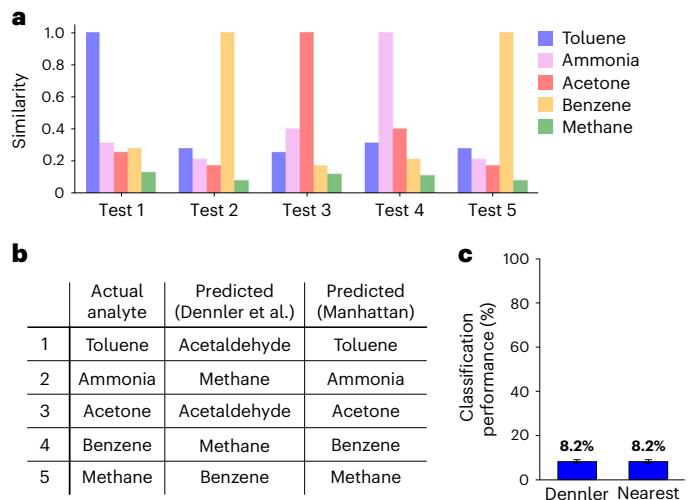


Fig. 1 | Benchmarking performance of the Dennler et al. method. **a**, When randomly generated 72-element vectors were presented to the algorithm of Dennler et al., it classified each random vector as a known odourant with Jaccard similarity 1.0. One hundred independently generated random vectors were presented, all of which were classified as knowns (five of these are depicted). The nearest-neighbour algorithm from ref. 2 classified each of the input vectors as none of the above, whereas the EPL algorithm correctly classified 99% as none of the above. **b**, We generated ten test samples by adding a small amount of noise (+1) to each of the elements of the ten training odourants. The algorithm of Dennler et al. failed to identify any of the test samples correctly (five are depicted). When we modified the algorithm to employ a Manhattan distance-based similarity metric, all of the test samples were classified accurately. **c**, We then further modified the algorithm of Dennler et al. to also include a classification threshold of $\theta = 0.75$ to exclude false positives. When we ran the same benchmark utilized in ref. 2 (100 test samples of each odourant, totalling 1,000 samples per simulation, impulse noise occlusion levels randomly selected from the range [0.2–0.8], here reporting the mean of 100 simulations performed with different random seeds), we obtained results identical to the nearest-neighbour algorithm from ref. 2 (Fig. 6a, 'Raw'), confirming their equivalence. Error bars depict standard deviations.

While it is desirable for a sensor device deployed in the wild to generalize across diverse, statistically ill-behaved odourant presentations, this is not what the EPL network is designed for. This regularization problem is addressed in part by our glomerular layer preprocessor computations^{11–14}, which specifically incorporate circuitry to provide concentration tolerance and other stabilizing effects. These transformations were excluded from ref. 2 to enable us to focus on the capabilities of EPL transformations. Clearly, also, given the goal of odourant recognition, one certainly would not disperse the elements of a multisensor array across a wide area, which randomizes per-sensor analyte concentrations and hence greatly increases the amount of uncorrelated variance without returning any benefit.

Distance measure performance

Dennler et al. submit that our findings can be “effectively addressed by using... a simple hash table,” by which they mean a measure of overlap using Jaccard similarity. This claim rests on two errors of analysis. First, the Jaccard similarity coefficients that they report in their Fig. 2 are misleading: their method, based on their own code, provides much lower certainty than that figure indicates (discussed below). Second, because Dennler et al. do not employ a similarity threshold θ to determine successful classification, they exclude the possibility of ‘none of the above’ (we used $\theta = 0.75$ in our paper, as noted therein). This forced-choice strategy leads inexorably to false-positive results, which they interpret as successful classifications. We explain these errors in more detail below.

The method of Dennler et al. (code in their Supplementary Algorithm 1) is comparable to the nearest-neighbour algorithm that we use in our paper (Fig. 6a, ‘Raw’, in ref. 2), except that they employ an inappropriate measure of similarity and do not include a threshold for classification. In brief, when a 72-element test odourant is presented for classification, it is compared against each of the (10) 72-element odourant templates. The algorithm of Dennler et al. then selects the template with the greatest number of exactly matching values to the test odourant and calculates the Jaccard coefficients between that selected template and each of the odourant templates in turn. The outcome of this strategy is that the best match is always reported as a perfect Jaccard similarity of 1.0 (that is, that of the selected template compared with itself; the actual similarity between the test odourant signature and that of any template is never reported). Each of the 5 bars corresponding to different odourant templates in the panels of their Fig. 2, then, depicts the median of ten results comprising some number of samples classified as that odourant (each contributing a Jaccard coefficient of 1.0) and some number classified as a different odourant (each contributing the Jaccard coefficient computed between the template of the odourant corresponding to the bar in question and the template of the odourant that the sample was classified as). Among other deficiencies, this practice does not reflect actual classification performance; to wit, if across ten samples of toluene, six are identified as toluene and four as some other odour, a median similarity of 1.0 to toluene is reported, even though four of the ten trials return inaccurate classifications. This strategy is unrelated to the method of our EPL attractor network, in which a test odourant signature is iteratively drawn towards the template(s) of its class, and its progressively increasing similarity to that template is directly reported.

A revealing consequence of the forced-choice strategy in Dennler et al. is that even randomly generated vectors, unrelated to any of the learned templates, are classified positively as knowns if they have even a single element that matches that of any template. To test this, we passed 100 randomly generated 72-element vectors to the algorithm of Dennler et al.; all were misclassified as knowns with perfect Jaccard similarities of 1.0 to the template of their class (Fig. 1a). An effective remedy for this error, as we presented in ref. 2, is to define a classification threshold θ such that test odourants are identified as class members only if their similarity to the template exceeds that threshold, otherwise being classified as none of the above. Accordingly, the nearest-neighbour algorithm that we used in ref. 2 (Fig. 6a, ‘Raw’) classified all 100 random vectors as none of the above, whereas our EPL model classified one as known and 99 as none of the above.

A converse problem with the algorithm of Dennler et al. is that a test odourant may be very closely related to a learned template, but have no precisely identical elements, in which case it is scored as having zero similarity to that learned template. Small amounts of noise can therefore lead the method of Dennler et al. to yield effectively random outcomes. To illustrate this, we generated ten test odourants by adding a minimal quantity of noise to each odourant template (+1 to each element). The Dennler et al. algorithm did not correctly classify any of these samples, whereas the Manhattan distance method that we used for benchmarking (Fig. 6 in ref. 2) accurately classified them all (Fig. 1b). We then modified Dennler et al.’s code to use a Manhattan distance-based similarity metric with a detection threshold of $\theta = 0.75$, and found that it then shows the same performance (8.2%) as the nearest-neighbour algorithm that we utilized in ref. 2 (Fig. 1c; compare with raw results from Fig. 6a in ref. 2). The performance benchmarks claimed by Dennler et al., then, arise from an artefactual effect of a similarity metric handcrafted for the specific test samples presented, potentiated by the lack of a threshold requirement and the reporting of all classification decisions as having similarity coefficients of unity.

Finally, Dennler et al. report a comparison between the runtime of their code on a central processing unit (CPU) and that of our EPL network on the Intel Loihi neuromorphic platform. We do not address

that specific comparison here, as their code does not actually solve the problem presented. However, a k -nearest-neighbours search algorithm using spike encoding, similar to the problem we solved, was evaluated on a multichip Loihi system; it outperformed a brute-force CPU algorithm by a factor of 685 in energy-delay product^{15,16}. Other problem instances using neural networks deployed on Loihi hardware also have shown orders of magnitude improvements in runtime and energy efficiency compared with CPU computation¹⁶.

Conclusion

Chemical sensing using metal oxide sensors is a challenging goal; it is considered an accomplishment for metal oxide sensor arrays to recognize a natural signal drawn from the same odorous location on the same or next day even without the addition of any disruptive noise¹⁷. Rather than wrestle with the limitations of metal oxide sensors, we seek to improve the back-end intelligence used for the identification of chemical odourants in the presence of competitive/occluding noise and other such challenges, in anticipation of real-world deployment using less problematic chemosensor technologies.

Data availability

The gas sensor data referred to herein¹³ are available from <http://archive.ics.uci.edu/dataset/251/gas+sensor+arrays+in+open+sampling+settings>.

Code availability

The code to generate the panels in Fig. 1 is available from <https://doi.org/10.6084/m9.figshare.25263661> (ref. 18).

References

1. Dennler, N., van Schaik, A. & Schmuks, M. Limitations in odour recognition and generalization in a neuromorphic olfactory circuit. *Nat. Mach. Intell.* **5**, 1038–1045 (2023).
2. Imam, N. & Cleland, T. A. Rapid online learning and robust recall in a neuromorphic olfactory circuit. *Nat. Mach. Intell.* **2**, 181–191 (2020).
3. Dennler, N., Rastogi, S., Fonollosa, J., van Schaik, A. & Schmuks, M. Drift in a popular metal oxide sensor dataset reveals limitations for gas classification benchmarks. *Sens. Actuators B* **361**, 131668 (2022).
4. Vergara, A. et al. On the performance of gas sensor arrays in open sampling systems using inhibitory support vector machines. *Sens. Actuators B* **185**, 462–477 (2013).
5. Müller, G. & Sberveglieri, G. Origin of baseline drift in metal oxide gas sensors: effects of bulk equilibration. *Chemosensors* **10**, 171 (2022).
6. Laplatine, L. et al. Silicon photonic olfactory sensor based on an array of 64 biofunctionalized Mach–Zehnder interferometers. *Opt. Express* **30**, 33955–33968 (2022).
7. Mulla, M. Y., Torsi, L. & Manoli, K. Electronic biosensors based on EGOFETs. *Methods Enzymol.* **642**, 403–433 (2020).
8. Bozdogan, A., Hageneder, S. & Dostalek, J. Plasmonic biosensors relying on biomolecular conformational changes: case of odorant binding proteins. *Methods Enzymol.* **642**, 469–493 (2020).
9. Aspermair, P., Mishyn, V., Szunerits, S. & Knoll, W. Electronic biosensors based on graphene FETs. *Methods Enzymol.* **642**, 371–401 (2020).
10. Cali, K., Tuccori, E. & Persaud, K. C. Gravimetric biosensors. *Methods Enzymol.* **642**, 435–468 (2020).
11. Borthakur, A. & Cleland, T. A. Signal conditioning for learning in the wild. In *Proc. 7th Annual Neuro-inspired Computational Elements Workshop 8* (Association for Computing Machinery, 2019); <https://doi.org/10.1145/3320288.3320293>
12. Moyal, R., Mama, K. R., Einhorn, M., Borthakur, A. & Cleland, T. A. Heterogeneous quantization regularizes spiking neural network activity. Preprint at <https://doi.org/10.48550/arXiv.2409.18396> (2024).
13. Imam, N. et al. Implementation of olfactory bulb glomerular-layer computations in a digital neurosynaptic core. *Front. Neurosci.* **6**, 83 (2012).
14. Borthakur, A. & Cleland, T. A. A spike time-dependent online learning algorithm derived from biological olfaction. *Front. Neurosci.* **13**, 656 (2019).
15. Frady, E. P. et al. Neuromorphic nearest neighbor search using Intel's Pohoiki Springs. In *Proc. 2020 Annual Neuro-Inspired Computational Elements Workshop 23* (Association for Computing Machinery, 2020); <https://doi.org/10.1145/3381755.3398695>
16. Davies, M. et al. Advancing neuromorphic computing with Loihi: a survey of results and outlook. *Proc. IEEE* **109**, 911–934 (2021).
17. Drix, D., Dennler, N. & Schmuks, M. Rapid recognition of olfactory scenes with a portable MOx sensor system using hotplate modulation. In *2022 IEEE International Symposium on Olfaction and Electronic Nose 1–4* (IEEE, 2022).
18. Cleland, T. A., Moyal, R. & Imam, N. Code to generate Figure 1. *Figshare* <https://doi.org/10.6084/m9.figshare.25263661> (2024).

Acknowledgements

This work was supported by NSF grants CBET-2123862 to T.A.C. and EFMA-2223811 to T.A.C. and N.I. R.M. was supported by the Eric and Wendy Schmidt AI in Science Postdoctoral Fellowship, a Schmidt Futures Program.

Author contributions

All authors contributed to the critical analysis of the Comment by Dennler et al. on our work, and to the text of the paper. R.M. conducted simulations. T.A.C. compiled and revised the paper.

Competing interests

T.A.C. and N.I. are listed as inventors on US patent US20220198245A1 and European patent EP3953868A4, both pending, on neuromorphic methods for rapid online learning and signal restoration.

Additional information

Correspondence and requests for materials should be addressed to Thomas A. Cleland.

Peer review information *Nature Machine Intelligence* thanks Yang Chai for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2024