



Rapid online learning and robust recall in a neuromorphic olfactory circuit

Nabil Imam¹ and Thomas A. Cleland²

We present a neural algorithm for the rapid online learning and identification of odourant samples under noise, based on the architecture of the mammalian olfactory bulb and implemented on the Intel Loihi neuromorphic system. As with biological olfaction, the spike timing-based algorithm utilizes distributed, event-driven computations and rapid (one shot) online learning. Spike timing-dependent plasticity rules operate iteratively over sequential gamma-frequency packets to construct odour representations from the activity of chemosensor arrays mounted in a wind tunnel. Learned odourants then are reliably identified despite strong destructive interference. Noise resistance is further enhanced by neuromodulation and contextual priming. Lifelong learning capabilities are enabled by adult neurogenesis. The algorithm is applicable to any signal identification problem in which high-dimensional signals are embedded in unknown backgrounds.

Spike timing-based mechanisms of coding and computation operating within plastic neural circuits present a central problem of interest to both neuroscience and neuromorphic computing. We have found that a coordinated set of these mechanisms, hypothesized for the neural circuitry of the external plexiform layer (EPL) of the mammalian main olfactory bulb (MOB), exhibits rapid learning of arbitrary high-dimensional neural representations and robust memory recall despite occlusion by random sources of destructive interference. On the basis of these mechanisms, we derived a neural algorithm for the learning of odourant signals and their robust identification under noise, and instantiated it in the Intel Loihi neuromorphic system¹ (Supplementary Methods). The algorithm operates over a network of excitatory and inhibitory units that embed feedforward and recurrent feedback circuit motifs. Information in the network is represented by sparse patterns of spike timing² measured against an underlying network rhythm. Learning is based on local spike timing-dependent plasticity rules, and memory is retrieved over sequential gamma-breadth packets of spiking activity. The network can be effectively trained using one-shot learning, and innately supports online learning; that is, additional training on new stimuli does not disrupt prior learning.

While both biological and artificial olfaction systems recognize chemical analytes based on activity patterns across arrays of weakly specific chemosensors^{3,4}, mammalian olfaction demonstrates levels of performance in signal restoration and identification currently unmatched by artificial systems. Indeed, the underlying identification problem is deceptively difficult. Natural odours comprise mixtures of many different odourant molecules⁵; moreover, under natural conditions, different odours from many separate sources intermingle freely and, when sampled together, chemically occlude one another in competition for primary chemosensor binding sites^{6–9}. This occlusion substantially disrupts the primary sensory activation patterns that provide the basis for odour recognition. Moreover, the patterns of potential occlusion are unrelated to the input statistics of the odours of interest, and hence unpredictable. This presents an extraordinary signal restoration challenge that has been recognized as one of the central problems in neuromorphic olfaction^{3,10,11}. By designing a neuromorphic algorithm based on computational principles extracted from the biological system, and

implementing it on a compact, field-deployable hardware platform, we sought to dramatically improve the performance and capabilities of artificial chemosensory systems deployed into uncontrolled environments.

This biological system exhibits several critical properties and mechanisms that we used to address the problem. Primary sensory representations of odour stimuli at steady state constitute intrinsically high-dimensional feature vectors, the dimensionality of which is defined by the number of receptor types (columns) expressed by the olfactory system¹²; this number ranges from the hundreds to over 1,000 in different mammalian species. Each of these receptor types induces spiking in a corresponding group of principal neurons (mitral cells; MCs). Mechanistically, fast coherent oscillations in the gamma band (around 30–80 Hz), which are intrinsic to MOB circuitry^{13–15}, phase-restrict the timing of these MC action potentials^{13,16}. This property discretizes spiking output into gamma-breadth packets, here enabling a robust within-packet phase precedence code^{17,18} that disambiguates phase-leading from phase-lagging spikes within each gamma cycle. Recurrent activity loops in OB circuitry evince control systems architecture, implementing gain control in the superficial layers^{19–21} and enabling associative attractor dynamics in the deeper network²². Odour learning in the biological system is localized and rapid, and depends substantially on plastic synapses within the MOB^{23–29}, here instantiated as spike timing-dependent plasticity rules. The neuromodulatory tuning of MOB circuit properties^{30–32} here is leveraged as an optimization trajectory rather than a fixed state variable. Adult neurogenesis in the MOB, known to be required for odour learning and memory^{26,27,33,34}, here provides indefinite capacity for lifelong learning through the permanent differentiation and replacement of plastic interneurons.

Our algorithm is derived from these computational properties of the EPL neural circuit in the biological MOB. We train and test the algorithm using data from the Vergara et al. dataset⁴, acquired from an array of 72 chemosensors mounted across a wind tunnel, and show that it rapidly learns odour representations and robustly identifies learned odours under high levels of destructive interference, as well as in the presence of natural variance arising from odourant plume dynamics. The destructive interference model, impulse noise, models the effects of intermixed, simultaneously

¹Neuromorphic Computing Laboratory, Intel Corporation, San Francisco, CA, USA. ²Computational Physiology Laboratory, Department of Psychology, Cornell University, Ithaca, NY, USA. e-mail: nabil.imam@intel.com; tac29@cornell.edu

sampled background odourants that effectively randomize the activation levels of a substantial fraction of the primary sensors on which odour recognition depends. The algorithm exhibits online learning and generalizes broadly beyond experience; accordingly, it can be trained on relatively clean diagnostic samples using one- or few-shot learning and then deployed into environments containing unknown contaminants and other sources of interference.

Network architecture and plasticity

The structure of the model network (Fig. 1a) was based on the circuitry and computational properties of the mammalian MOB, optimized for efficient implementation as a spiking network on the Loihi chip (Fig. 1b). In particular, we instantiated some core principles of MOB computation that we have hypothesized for the biological system^{12,14,15,17,35}, including (1) the dynamically acquired, learning-dependent topology of the lateral inhibitory network of the EPL, (2) gamma-discretized spike timing-based computation in the EPL, (3) the widespread excitation of granule cell interneurons (GCs) by MCs and the local, spike-based inhibition of MCs by GCs, (4) the manifestation of GC inhibition of MCs as delays in MC spike times on the gamma timescale, (5) the permanent differentiation of GCs by the process of odour learning and the consequent need for replacement by adult neurogenesis, and (6) the deployment of neuromodulation as a dynamic optimization trajectory rather than a stationary state.

Like the mammalian MOB, the neuromorphic EPL network is implicitly columnar (Fig. 1a). Each column comprises a single MC principal neuron as well as up to 50 inhibitory GC interneurons, coupled by moderately sparse intercolumnar excitatory synapses (connection probability 0.2) and local (intracolumnar) inhibitory synapses (Supplementary Methods). We activated the MCs of a 72-column EPL network using the ‘Gas sensor arrays in open sampling settings’ dataset published by Vergara and colleagues⁴ (see ‘Data availability’). Samples were drawn from an array of 72 metal oxide gas sensor elements spatially distributed across the full 1.2 m breadth of a wind tunnel⁴ (Fig. 1c). From the 180s datastreams comprising each odourant presentation in this dataset, sensor array responses were sampled (‘sniffed’) from discrete points in time and presented to the EPL network for training or testing. That is, individual odour samples (‘sniffs’) comprised discrete feature vectors in which the pattern of amplitudes across vector elements reflected odour quality, as well as concentration-based variance owing to plume dynamics in the wind tunnel.

The biological EPL network is intrinsically oscillogenic in the gamma band (30–80 Hz)^{15,36}, and MC action potentials are statistically phase-constrained with respect to these local oscillations^{13,16}. In our algorithm, MC spikes were constrained in time by an ongoing network oscillation with alternating permissive and inhibitory epochs reflecting the periodic inhibition of the OB gamma cycle (Fig. 1d). Sensory integration and MC spiking were enabled only during permissive epochs, whereas inhibitory epochs reset and held the activation of all MCs at zero. Therefore, in the absence of learning, and given stationary sensor input, the temporal patterning of spikes evoked by a given odour directly reflected sensor activation levels—stronger excitation evoked correspondingly earlier spikes—and was repeated across successive gamma cycles (Fig. 1d). Different odours evoked correspondingly different spatiotemporal spike patterns across the MC population, thereby generating a hybrid channel/phase code, or precedence code, on the gamma timescale.

Critically, this dynamical architecture enables multiple iterative cycles of processing for each sample by taking advantage of the differences in timescale between sampling (4–8 Hz in rodent sniffing, 100 Hz in the Vergara et al. dataset) and processing (30–80 Hz gamma oscillations in the rodent OB, 100 kHz in the Loihi chip). In the present instantiation of the algorithm, five gamma

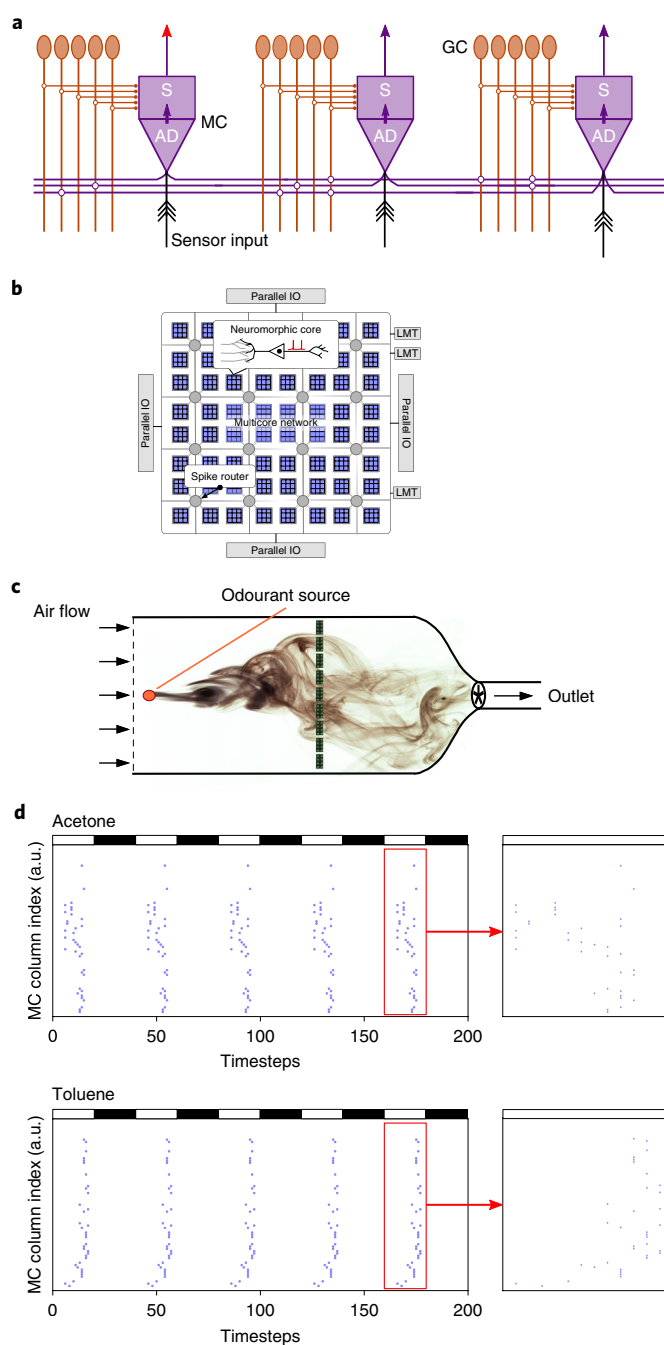


Fig. 1 | Model structure and signal encoding. **a**, Architecture of the neuromorphic model. Sensor input is delivered to the apical dendrite (AD) of each MC, which in turn excites its corresponding soma (S). MC activity propagates via lateral dendrites (purple) to excite the dendrites of GCs (orange). The distribution of excitatory connections (open circles) is sparse and independent of spatial proximity. In contrast, GC inhibitory connections are local. **b**, Architecture of the Intel Loihi neuromorphic chip¹. Neuromorphic cores (blue squares) operate in parallel and communicate through a mesh of spike routers (grey circles). LMT, Lakemont x86 processors; IO, input/output. **c**, Illustration of odourant delivery to a 72-element chemosensor array within a wind tunnel⁴. **d**, Presentation of acetone (top) or toluene (bottom) to the chemosensor array resulted in characteristic patterns of spiking activity across the 72 MCs (ordinate). Stronger sensor activation led to correspondingly earlier MC spikes within each gamma cycle. Inhibitory epochs are denoted by black bars; permissive epochs are denoted by white bars. The fifth gamma cycle is expanded in time (rightmost panels) to illustrate the distribution of MC spike times.

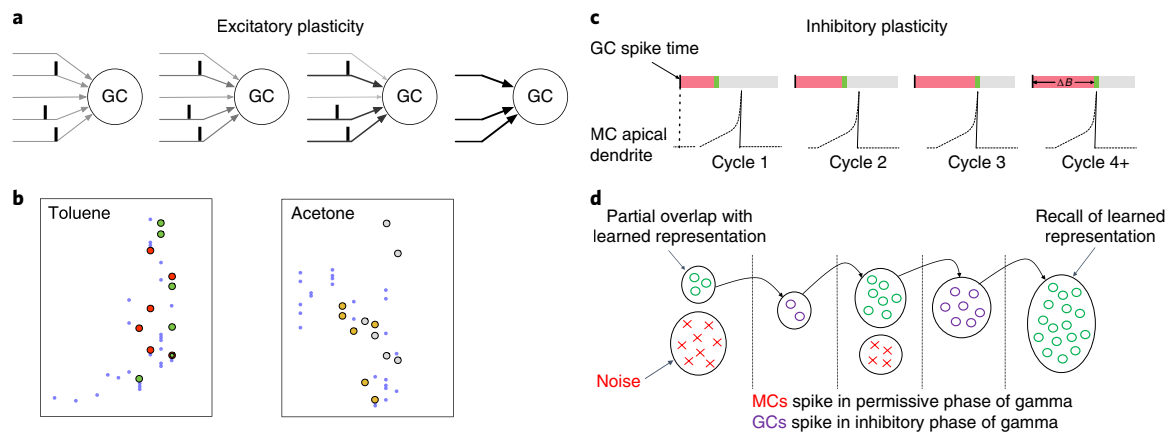


Fig. 2 | Plasticity rules. **a**, During training, coincident MC spikes converging onto a given GC activated that GC, and developed strong excitatory synaptic weights thereon, whereas other inputs to that GC were weakened and ultimately eliminated. **b**, Excitatory plasticity rendered GCs selective to higher-order features of odour representations. After training on toluene (left) or acetone (right), a number of GCs became responsive to specific combinations of activated MCs. Spike times highlighted with green spots denote those MC spikes that activated a specific GC in a network trained on toluene. Red spots denote a second such GC, responsive to toluene via a different set of activated MCs. Yellow and grey spots denote analogous MC spike populations that activate two GCs responsive to acetone in the same network. **c**, Illustration of the inhibitory plasticity rule. During training, the duration of spike-mediated GC inhibition onto its co-columnar MC (red bar) increased until the release of this inhibition (green) coincided with spike initiation in the MC apical dendrite. The learned inhibitory weight corresponded to a blocking period ΔB during which spike propagation in the MC soma was suppressed. **d**, Illustration of the iterative denoising of an occluded test sample. Partially correct representations in MCs evoke responses in some of the correct GCs, which deliver inhibition that draws MC ensemble activity iteratively closer to the learned representation. Three permissive epochs interspersed with two inhibitory epochs are depicted.

cycles, each requiring 0.4 ms to execute (Methods), were embedded within each odour presentation (“sniff”) for both training and testing. After learning, GC feedback inhibition on each successive gamma cycle iteratively modified MC spike timing and hence altered the precedence code. Network output thus was interpreted as an evolving series of representations, in which each discrete representation comprised a population of spikes, with each spike defined by the identity of the active MC and the spike latency within the corresponding gamma cycle. These representations then were classified based on their similarities to each of the representations known by the network.

Each gamma-constrained array of MC action potentials, in addition to serving as network output, also drove its complement of postsynaptic GCs across the network. During learning, the synaptic weights between MCs and GCs were systematically modified by experience. GCs were modelled as single-compartment neurons that accumulated excitatory synaptic inputs from their widely distributed presynaptic MCs. On reaching threshold, they generated spike events that inhibited their co-columnar postsynaptic MC in the subsequent gamma cycle.

GC spiking also initiated excitatory synaptic plasticity. Specifically, GCs learned to respond to higher-order stimulus features by becoming selective for specific combinations of MC spiking activity. To do this, we implemented a heterosynaptic spike timing-dependent plasticity rule that learned these input combinations in terms of a spike phase precedence coding metric on the gamma timescale¹⁷. Specifically, the weights of synapses mediating presynaptic spikes that immediately preceded a postsynaptic spike were strengthened, and the weights of all other incoming synapses, including those in which the presynaptic MC spiked at other times or not at all, were weakened (Fig. 2a). Accordingly, spiking GCs ultimately learned a fixed dependency on the synchronous firing of a set of k MC inputs, with inputs from other MCs decaying to zero (effectively a ‘ k winners take all’ learning rule). Consequently, at the end of the training period, the response to each trained odourant evoked a distributed ensemble of GCs tuned to a diversity of stimulus-specific higher-order correlation patterns (Fig. 2b).

Spikes evoked by GC interneurons delivered synaptic inhibition onto the MC of their local column. As proposed for the biological system, the weights of GC-mediated inhibitory inputs regulated the timing of MC spikes within the permissive phase of the gamma cycle, with stronger weights imposing greater MC spike time delays within each gamma cycle^{14,35,37}. Specifically, a GC spike blocked the generation of a spike on its follower MC for a period of time corresponding to the inhibitory synaptic weight. During odour learning, the durations of GC spike-evoked inhibitory windows were iteratively modified until the release of inhibition on the MC soma coincided with a threshold crossing in the MC apical dendrite resulting from integrated sensory input (Fig. 2c). During testing, the end of the GC inhibitory window permitted the MC to fire, and evoked a rebound spike in the MC even in the absence of sufficient apical dendritic input. Synaptic inputs from multiple local GCs onto a common MC were independent of one another, enabling a diverse range of higher-order GC receptive fields to independently affect the MC. During testing, occluded inputs activated some fraction of GCs, which then modified their postsynaptic MC spike times such that the representation in the next gamma cycle was closer to a learned odourant, hence activating an increased fraction of its corresponding GCs. This process continued iteratively until the learned representation was recalled (Fig. 2d and Supplementary Notes).

Odour learning and memory

We first trained the 72-column network on the odourant toluene in one shot (that is, one sniff, enabling learning over five gamma cycles), and then, with plasticity disabled, tested the response of the trained network to presentations of toluene contaminated with destructive interference. To generate this interference, we entirely replaced a proportion P of the sensory inputs with random values (impulse noise, $P=0.6$ unless otherwise indicated) to represent strong and unpredictable receptor occlusion through simultaneous activation or inhibition by other ambient odourants. The occluded inputs remained consistent over the five gamma cycles of a sniff. In a naïve network, the presentation of occluded toluene yielded an essentially stationary and uninformative representation

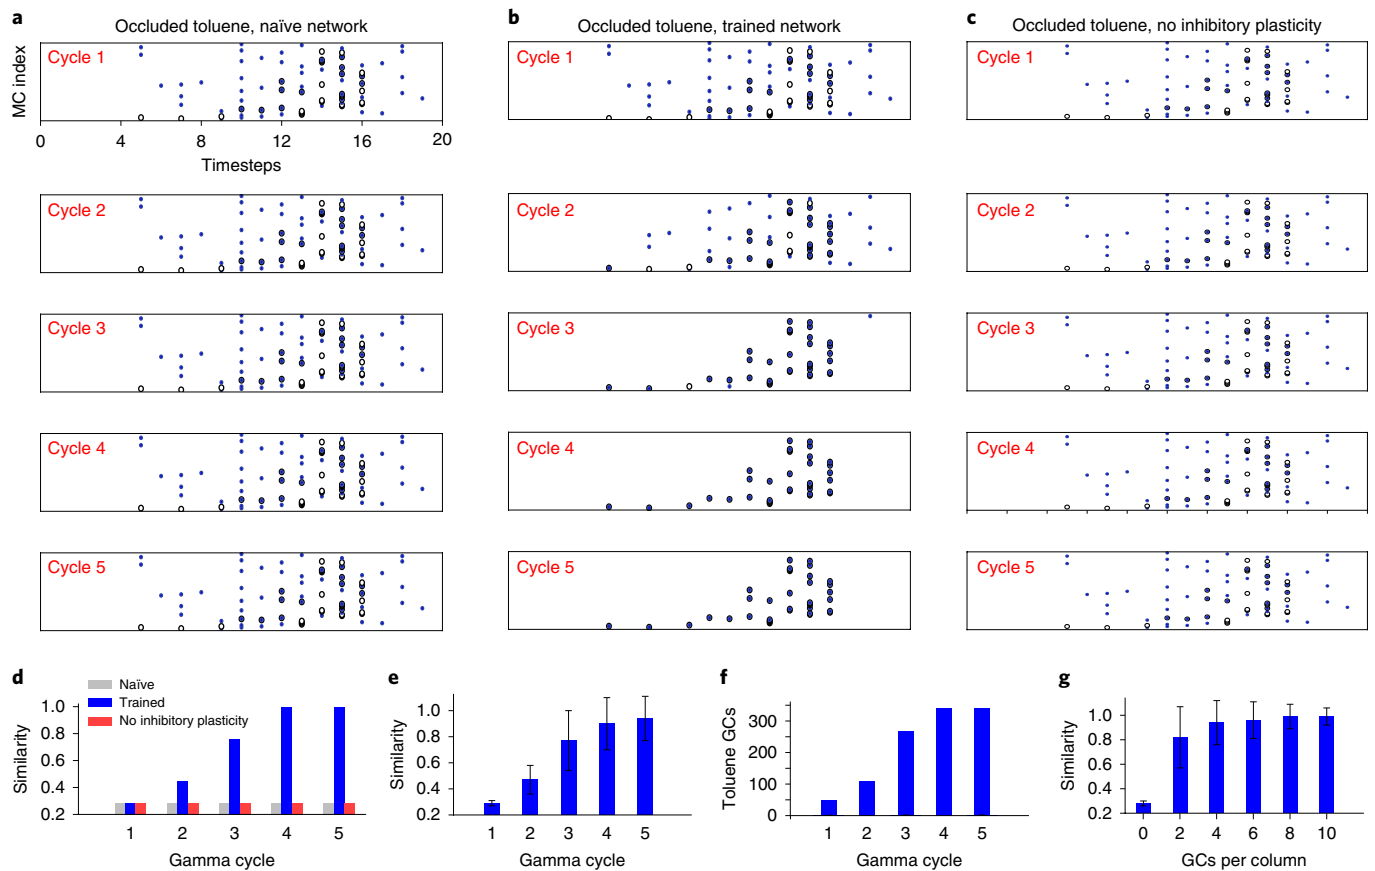


Fig. 3 | Odour-evoked MC activity patterns are attracted to learned representations. **a**, Presentation of an occluded instance of toluene to an untrained network. Blue dot rasters denote spike times evoked by occluded toluene (impulse noise $P=0.6$). The untrained network does not update the response to occluded toluene over the five gamma cycles depicted. For comparison, open circle rasters denote the spike times evoked by non-occluded toluene. **b**, Presentation of the same occluded instance of toluene to a plastic network trained on (non-occluded) toluene. The activity profile evoked by the occluded sample was attracted to the learned toluene representation over successive gamma cycles. **c**, Presentation of the same occluded instance of toluene to a network trained on non-occluded toluene with excitatory, but not inhibitory, plasticity enabled. The omission of inhibitory plasticity rendered the network unable to denoise MC representations during testing. **d**, The Jaccard similarity⁴⁴ between the response to occluded toluene and the learned representation of toluene systematically increased over five gamma cycles in the trained network (**b**), but not in the untrained network (**a**) or the network with inhibitory plasticity disabled (**c**). **e**, The Jaccard similarity increased reliably over five gamma cycles when averaged over 100 independently generated instances of occluded toluene (impulse noise $P=0.6$). **f**, During learning, the number of GCs tuned to toluene increased over the five successive gamma cycles of training. **g**, Mean Jaccard similarity in the fifth gamma cycle as a function of the number of undifferentiated GCs per column. Mean similarity is averaged across 100 occluded instances of toluene (impulse noise $P=0.6$). Five GCs per column were utilized for all other simulations described herein. Error bars denote standard deviation.

(Fig. 3a,d). However, in the trained network, the spiking activity generated by occluded toluene was attracted over the five gamma cycles towards the previously learned toluene representation, enabling clear identification of the occluded unknown (Fig. 3b,d,e). In contrast, if inhibitory plasticity (Fig. 2c,d) was suppressed during training, the trained EPL network was unable to denoise the MC representation (Fig. 3c,d).

As hypothesized for the biological OB, odour learning in the network induces the permanent differentiation of GCs (Fig. 3f) that thereby become selective for higher-order feature combinations that are relatively diagnostic of the learned odour^{38–40} (Fig. 2b). We tested whether increased allocations of GCs, enabling each MC to be inhibited by a broader selection of feature combinations, would improve odour learning and identification under noise. We found that increasing the number of undifferentiated GCs per column improved the robustness of signal restoration, increasing the similarity of the occluded signal to the learned representation after five gamma cycles (Fig. 3g). Nevertheless, we limited our simulations to five GCs per trained odourant and five gamma

cycles per sniff to avoid ceiling effects and thereby better reveal the variables of greatest interest.

This learning algorithm irreversibly consumes GCs. Each odour memory is associated with a distributed population of differentiated GCs tuned to its complex diagnostic features. Fully differentiated, mature GCs do not undergo further plasticity and hence are protected from catastrophic interference^{41,42} (Supplementary Notes). The learning of successively presented new odourants, however, would be increasingly handicapped by the declining pool of undifferentiated GCs (Fig. 3g). The competition among distinct new odourants can be substantially reduced by sparser initial MC → GC connectivity and higher numbers of GCs, among other parameters³⁸; however, genuine lifelong learning in such a system requires a steady source of undifferentiated GCs. Exactly this resource is provided to the mammalian olfactory system by constitutive adult neurogenesis. We propose that the critical role of adult neurogenesis in odour learning^{26,27,33,34,43} be interpreted in this light.

We then trained the 72-column network sequentially with all ten odourants in the dataset⁴, using a one-shot training procedure for

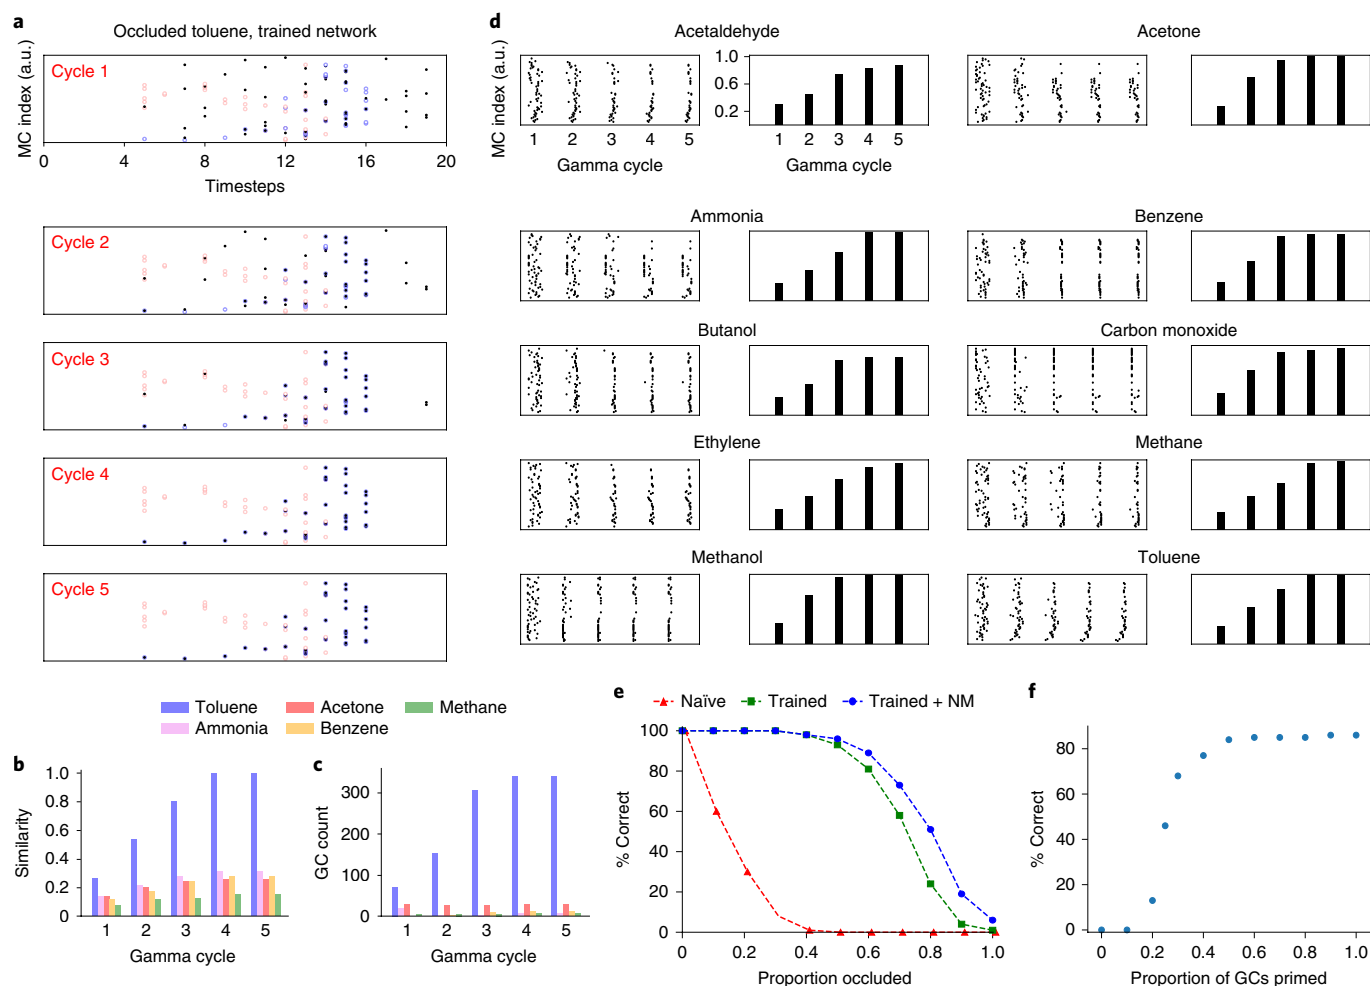


Fig. 4 | Multi-odour learning. **a**, Spike raster plot depicting attractor dynamics after training the network on all ten odourants. The representation generated by a sample of occluded toluene ($P=0.6$; black dots) was progressively drawn towards the learned representation of toluene (open blue circles) and away from the learned representations of acetone (open red circles) and the other eight odourants (not shown). **b**, The Jaccard similarity to toluene that was evoked by the occluded-toluene stimulus increased over five successive gamma cycles until the stimulus was classified as toluene (similarity >0.8). For clarity, only five odourants are depicted. **c**, The number of toluene-tuned GCs activated by the occluded-toluene stimulus progressively increased over five gamma cycles as the MC spiking activity pattern was attracted towards the learned toluene representation. GCs tuned to the other nine odourants were negligibly recruited by the evolving stimulus representation. **d**, Network activity evoked by presentation of occluded instances of each of the ten learned odourants following one-shot learning. Left: spike raster plots over five gamma cycles (200 timesteps). Right: Jaccard similarity between the activity pattern generated by each occluded odourant stimulus and the learned representation of the corresponding odourant. The same network can reliably recognize all ten odourants from substantially occluded examples ($P=0.6$). **e**, Mean classification performance across all ten odourants under increasing levels of sensory occlusion (100 impulse noise instantiations per odourant per noise level). The abscissa denotes the level of impulse noise—that is, the proportion of MC inputs for which the sensory activation level was replaced with a random value. The red curve is the proportion of correct classifications by an untrained network. The green curve is the proportion of correct classifications by a network trained on all ten odourants. The blue curve is the proportion of correct classifications by a trained network with the aid of a neuromodulation-dependent dynamic state trajectory. **f**, Effects of GC priming on classification performance under extreme occlusion. One hundred independently generated samples of occluded toluene with impulse noise $P=0.9$ were presented to the fully trained network. The putative effects of priming arising from piriform cortical projections onto bulbar GCs were modelled by lowering the spike thresholds of a fraction of toluene-tuned GCs. As the fraction of toluene-tuned GCs so activated was increased, classification performance increased from near zero to over 80% correct.

each odour. In each case, the network was trained on one odour first, followed by a second odour, then by a third, until all ten odours had been learned. Similar results were obtained irrespective of the order in which the ten odourants were trained. A set of new, undifferentiated GCs was added to the network after each odour was learned, reflecting the effects of adult neurogenesis. Critically, subsequent odour training did not disrupt the memories of previously learned odours; that is, the EPL network supports robust online learning, and is resistant to catastrophic forgetting. This capacity for online learning is essential for memory formation under

natural conditions, as well as for continuous device operation in the field; in either case, new signals of potential significance may be encountered at unpredictable times, and must be incorporated non-destructively into an existing knowledge base.

We then tested the algorithm's capacity to recognize and classify odourant samples that were strongly occluded by impulse noise, reflecting the effects of any number of independent odourous contaminants that could mask the odour of interest in uncontrolled environments. Following training on all ten odourants, sensor-evoked activity patterns generated by strongly occluded odour

stimuli (impulse noise $P=0.6$) were attracted specifically towards the learned representation of the corresponding odour (Fig. 4a–c). Notably, the same network was able to rapidly identify occluded instances of all ten odours within five gamma cycles (Fig. 4d). An odour was considered identified when the spatiotemporal pattern of its evoked spiking activity exceeded a Jaccard similarity⁴⁴ of 0.75 to one of the network's learned representations. Performance on this dataset under standard conditions was strong up to sample occlusion levels of $P=0.6$, after which increased occlusion began to gradually impair classification performance (Fig. 4e).

Neuromodulators such as acetylcholine and noradrenaline generate powerful effects on stimulus representations and plasticity in multiple sensory systems including olfaction. Traditionally, they are treated as state variables that may sharpen representations or gate learning, or bias a network towards one source of input or another^{45–47}. We instead modelled neuromodulatory effects as a dynamic search trajectory. Specifically, as the neuromodulator is released in response to active olfactory investigation (sampling), the local concentration around effector neurons and synapses rises over the course of successive sniffs, potentially enabling the most effective of the transient neuromodulatory states along that trajectory to govern the outcome of the stimulus identification process. We implemented a gradual reduction in the mean GC spiking threshold over the course of five sniffs of a corrupted odour signal, reflecting a concomitant increase in neuromodulator efficacy, and used the greatest of the five similarity values measured in the last gamma cycle within each sniff to classify the test odourant. Importantly, under very high noise conditions, each of the five 'neuromodulatory' states performed best for some of the test odours and noise instantiations, indicating that a trajectory across a range of neuromodulatory states could yield superior classification performance compared with any single state. Indeed, this strategy yielded a substantial improvement in classification performance at very high levels of impulse noise, approximately doubling classification performance at $P=0.8$ (Fig. 4e).

In the biological system, OB activity patterns resembling those evoked by a specific odour can be evoked by contextual priming that is predictive of the arrival of that odour⁴⁸. We implemented this as a priming effect exerted by ascending piriform cortical neurons that synaptically excite GCs in the OB⁴⁹, the mapping between which can be learned dynamically⁵⁰. Specifically, we presented the network with odour samples at an extreme level of destructive interference ($P=0.9$) that largely precluded correct classification under default conditions (Fig. 4e). When fractions of the population of GCs normally activated by the presented odour were primed by lowering their spike thresholds, classification performance improved dramatically, to a degree corresponding to the fraction of primed GCs (Fig. 4f). That is, even a weak prior expectation of an incoming odour stimulus suffices to draw an extremely occluded odour signal out of the noise and into the attractor.

In addition to occlusion by competing odourants, odour samples can vary based on the dynamics of their plumes (Fig. 1c), which evolve over time. We therefore tested the algorithm's ability to recognize and classify samples of each odourant that were drawn from the wind tunnel at different points in time (Fig. 5a,b). Specifically, in this paradigm, repeated samples of the same odourant differed from one another based on evolving odour plume dynamics, whereas samples of different odourants differed from one another both in plume dynamics and in the distribution of analyte sensitivities across the sensor array. Following one-shot training on all ten odours as described above, the spiking activity generated by odourant test samples was attracted over the five gamma cycles towards the corresponding learned representation. Notably, plume dynamics alone constituted a relatively minor source of variance compared with impulse noise (Fig. 5c).

We then tested the network on samples incorporating both plume dynamics and impulse noise ($P=0.4$). Following one-shot training on all ten odours, we sampled each odour across widely dispersed points in time, and contaminated each sniff with an independent instantiation of impulse noise (Fig. 5d,e). Spiking activity was again attracted over the five gamma cycles of each sniff towards the correct learned representation (Fig. 5f–h). Classification performance across levels of impulse noise from $P=0.0$ to $P=1.0$ (Fig. 5i) indicated that the addition of plume-based variability moderately reduced network performance (compare to Fig. 4e, green curve). Network performance was not affected by the introduction of noise correlations over time (Supplementary Fig. 1).

Classification performance

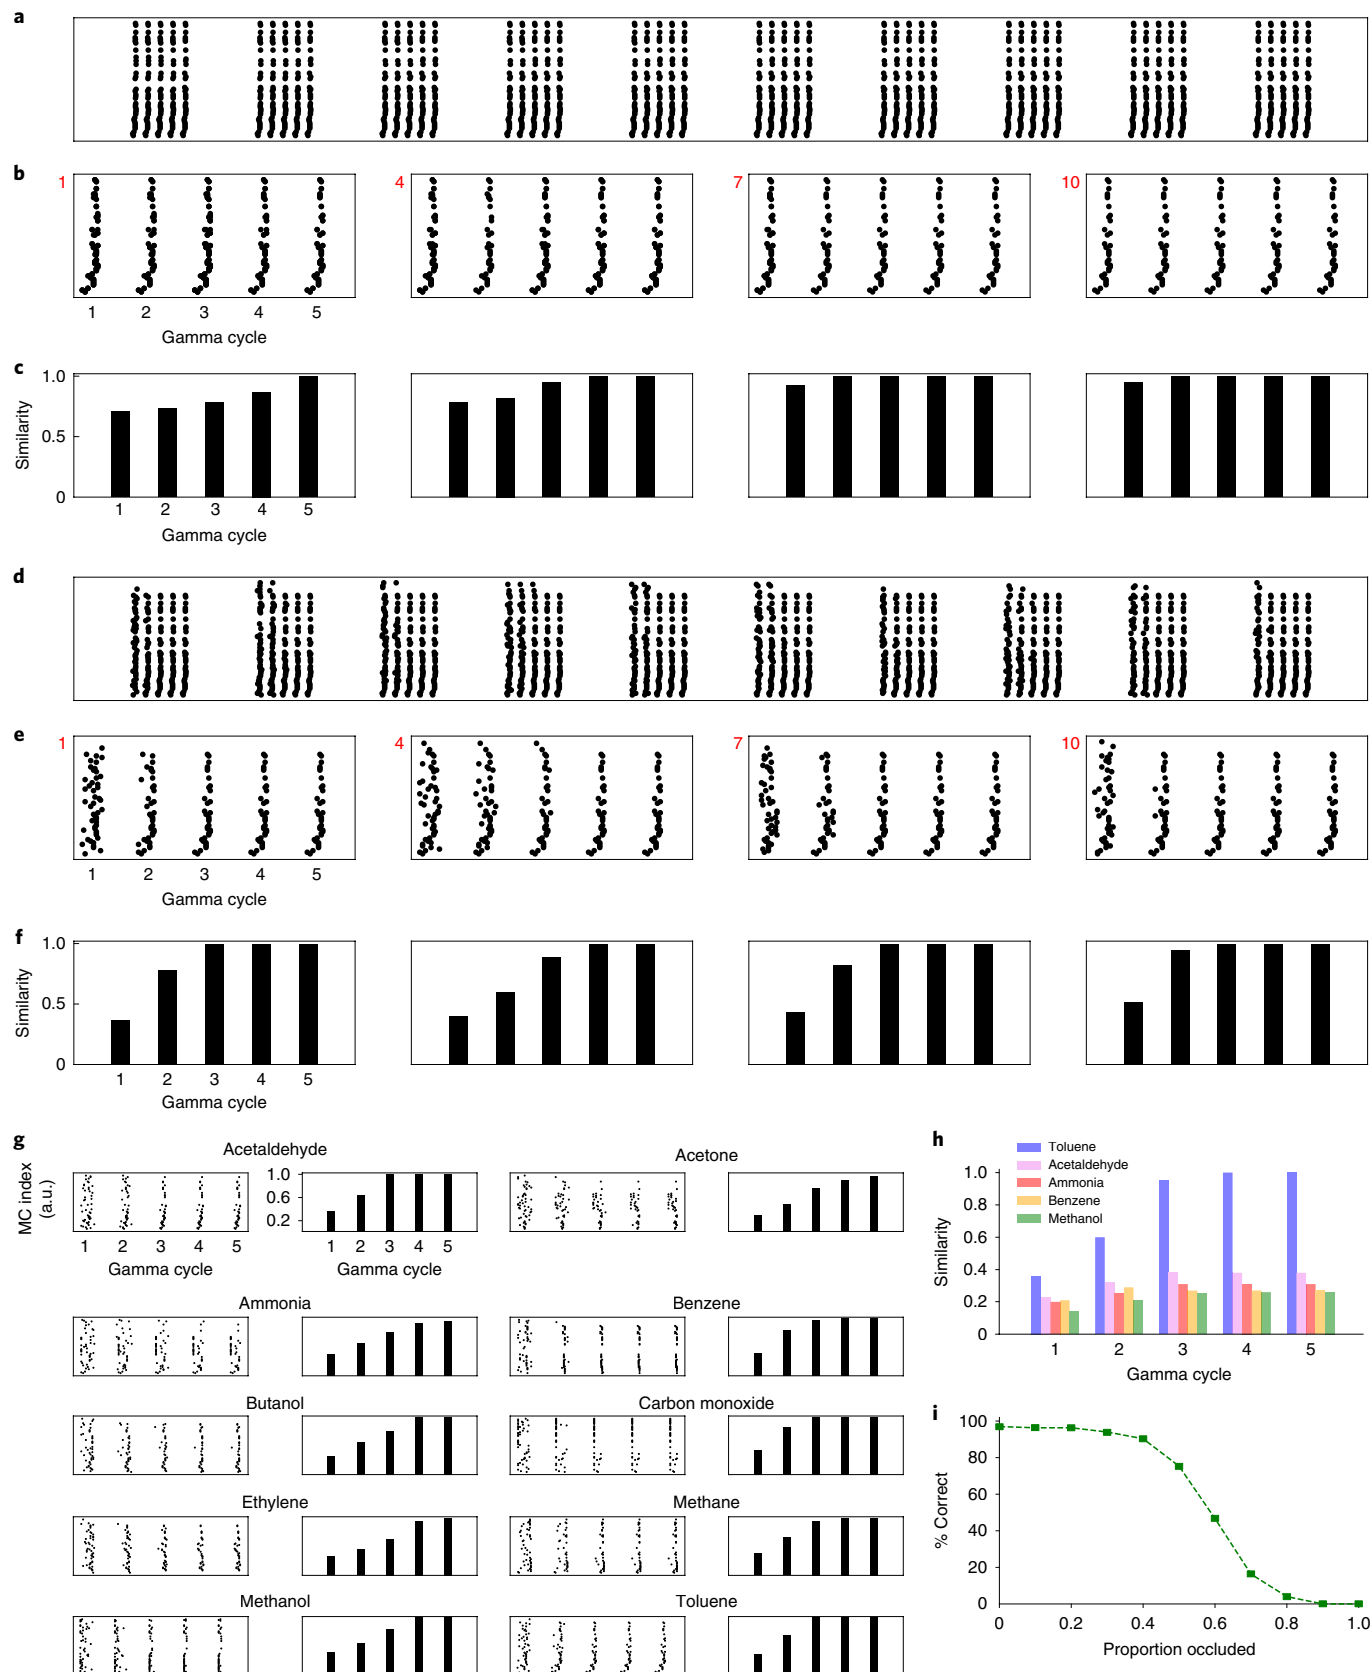
To evaluate the performance of the EPL model, we compared its classification performance to the performance of multiple conventional signal processing techniques: a median filter (MF), a total variation filter (TVF; both commonly used as impulse noise reduction filters⁵¹), principal component analysis (PCA; a standard preprocessor used in machine olfaction³) and a seven-layer deep autoencoder (DAE; Methods). Specifically, following training, we averaged the classification performance of each method across 100 different occluded presentations of each odour, with the occlusion level for each sample randomly and uniformly selected from the range $P=[0.2, 0.8]$, for a total of 1,000 test samples. Incorrect classifications and failures to classify both were scored as failures.

The neuromorphic EPL substantially outperformed MF, TVF and PCA. To model 'one sample' learning on the DAE for comparison with one-shot learning on the EPL network, we trained a DAE with one sample from each of the ten odourants over 1,000 training epochs per odourant, with the odourants intercalated in presentation. The EPL network substantially outperformed the DAE under these conditions, in which the training set contained no information about the distribution of error that would arise during testing owing to impulse noise (Fig. 6a). To improve DAE performance, we then

Fig. 5 | Odour learning with plume dynamics. **a**, Ten sniffs of toluene drawn from randomly selected timepoints within the dataset illustrate sampling variance arising from plume dynamics. Ordinate denotes MC index, ordered according to sensor locations across the wind tunnel. **b**, Higher-resolution depictions of sniffs 1, 4, 7 and 10 from **a**. **c**, Jaccard similarities between the learned representation of toluene and the activity patterns generated by plume-varying toluene stimuli across the five gamma cycles of each of the four sniffs depicted in **b**. **d**, Ten sniffs of toluene drawn from randomly selected timepoints within the dataset and also occluded with impulse noise ($P=0.4$). **e**, Higher-resolution depictions of sniffs 1, 4, 7, and 10 from **d**. **f**, Jaccard similarities between the learned representation of toluene and the activity patterns generated by plume-varying, occluded toluene stimuli across the five gamma cycles of each of the four sniffs depicted in **e**. **g**, Network activity evoked by presentation of plume-varying and occluded instances of each of the ten learned odours following one-shot learning. Left: spike raster plots over five gamma cycles. Right: Jaccard similarities between the activity pattern generated by each occluded odourant stimulus and the learned representation of the corresponding odourant. The same network reliably recognized all ten odourants from plume-varying and occluded examples. **h**, The Jaccard similarity to toluene that was evoked by the occluded, plume-varying toluene stimulus increased over five successive gamma cycles until the stimulus was classified as toluene (similarity >0.8). For clarity, only five odourants are depicted. **i**, Mean classification performance across all ten odourants, with plume dynamics, under increasing levels of sensory occlusion (100 impulse noise instantiations per odourant per noise level). The abscissa denotes the level of impulse noise. The green curve shows the proportion of correct classifications by a network trained on all ten odourants.

trained it with 500 to 7,000 samples of each of the ten odourants, with each sample independently occluded by impulse noise randomly and uniformly selected from the range $P=[0.2, 0.8]$. Under this training regime, the deep network required 3,000 samples per

odourant, including the attendant information regarding the distribution of testing variance, to achieve the classification performance that the EPL model achieved with one sample per odourant. With further training, DAE performance exceeded that of the EPL



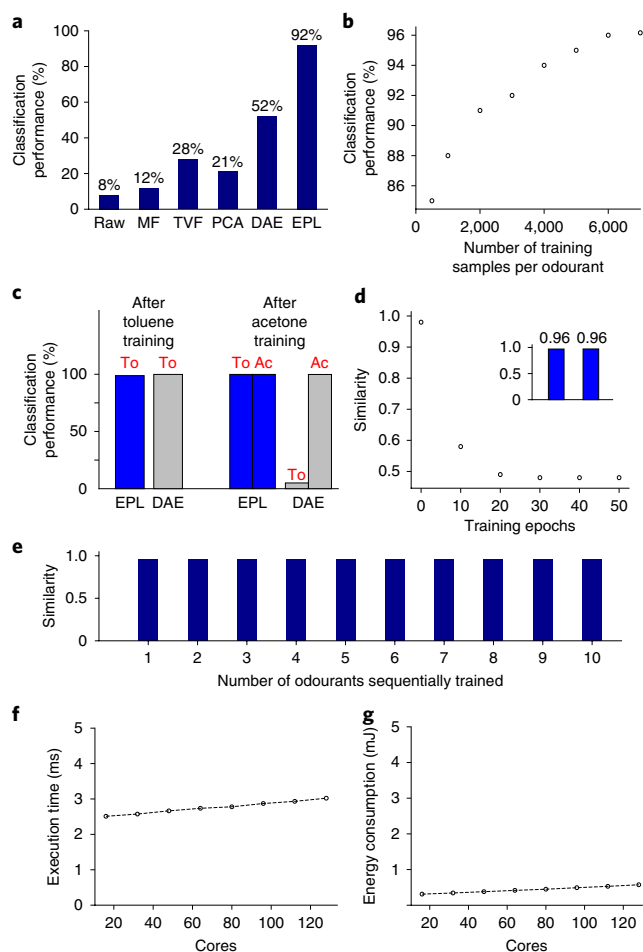


Fig. 6 | Performance evaluation. **a**, Classification performance of the EPL network compared with four other signal processing techniques. Raw, classification of unprocessed sensor signals; MF, median filter; TVF, total variation filter; PCA, principal component analysis; DAE, a seven-layer deep autoencoder; EPL, the neuromorphic EPL model. Each of the 10 odourants was presented with 100 independent instantiations of impulse noise, yielding 1,000 total test samples. **b**, The performance of the DAE improved when it was explicitly trained to map a variety of occluded instances of each odour to a common representation. To achieve performance superior to the one-shot-trained EPL network, the DAE required 3,000 occluded training samples per odourant. **c**, Online learning. After training naïve EPL and DAE networks with toluene (To), both recognized toluene with 100% accuracy. After then training the same network with acetone (Ac), the DAE learned to recognize acetone with 100% accuracy, but was no longer able to recognize toluene (catastrophic forgetting). In contrast, the EPL network retained the ability to recognize toluene after subsequent training on acetone. **d**, Gradual loss of the toluene representation in the DAE during subsequent training with acetone. The ordinate denotes the similarity of the toluene-evoked activity pattern to the original toluene representation as a function of the number of training epochs for acetone. Values are the means of 100 test samples. Inset: similarity between the toluene-evoked activity pattern and the original toluene representation in the EPL network before training with acetone (left) and after the completion of acetone training (right). **e**, Similarity between the toluene-evoked activity pattern and the original toluene representation as the EPL network is sequentially trained on all ten odourants of the dataset. Values are the means of 100 test samples. **f**, The execution time to solution is not notably affected as the EPL network size is expanded, reflecting the fine granularity of parallelism of the Loihi architecture. In the present implementation, one Loihi core corresponds to one OB column. **g**, The total energy consumed increases only modestly as the EPL network size is expanded.

network (Fig. 6b). We then tested the online learning capacities of the two networks, in which the presentations of different odourants during training were sequential rather than uniformly interspersed. After training both networks to recognize toluene using the methods of Fig. 6b, both the EPL and the DAE exhibited high classification performance. However, after subsequent training to recognize acetone, the DAE lost its ability to recognize toluene, whereas the EPL network recognized both odours with high fidelity (Fig. 6c,d). Susceptibility to catastrophic forgetting is a well-established limitation of deep networks, though some customized networks recently have shown improvements in their online (continual) learning capabilities that reflect some of the strategies of the EPL network, such as the selective reduction of plasticity in well-trained network elements⁴².

These results indicate that the EPL network ultimately serves a different purpose than techniques that require intensive training with explicit models of expected variance to achieve optimal performance. The EPL network is competitive with these algorithms overall, but excels at rapid, online learning with the capacity to generalize beyond experience in novel environments with unpredictable sources of variance. In contrast, the DAE evaluated here performs best when it is trained to convergence on data drawn from the distribution of expected variance; under these conditions, its performance exceeds that of the present EPL network. EPL network instantiations are thereby likely to be favoured in embedded systems intended for deployment in the wild, where rapid training, energy efficiency, robustness to unpredictable variance and the ability to update training with new exemplars are at a premium.

Summary

Neuromorphic computing shows great promise, but currently suffers from a paucity of useful algorithms. Seeking inspiration from the circuit-level organization of biological neural systems, with their radically different computational strategies, provides a key opportunity to develop algorithmic approaches that might not otherwise be considered. We demonstrate that a simplified network model, based on the architecture and dynamics of the mammalian olfactory bulb¹⁴ (Supplementary Discussion) and instantiated in the Loihi neuromorphic system¹, can support rapid online learning and powerful signal restoration of odour inputs that are strongly occluded by contaminants. These results evince powerful computational features of the early olfactory network that, together with mechanistic models and experimental data, present a coherent general framework for understanding mammalian olfaction as well as improving the performance of artificial chemosensory systems. Moreover, this framework is equally applicable to other steady-state signal identification problems in which higher-dimensional patterns without meaningful lower-dimensional internal structure are embedded in highly interfering backgrounds.

Methods

Dataset and odourant sampling. Sensory input to the model was generated from the 'Gas sensor arrays in open sampling settings' dataset published by Vergara and colleagues⁴ and available from the UCI Machine Learning Repository (see 'Data availability'). The dataset comprises the responses of 72 metal oxide-based chemical sensors distributed across a wind tunnel. There are six different sensor mounting locations in the tunnel, three different settings of the tunnel's wind speed and three different settings of the sensor array's heater voltage. In our present study, we consider the recordings made at sensor location 'L4' (near the mid-point of the tunnel), with the wind speed set to 0.21 m s⁻¹ and the heater voltage set to 500 V. The tunnel itself was 1.2 m wide by 0.4 m tall by 2.5 m long, with the sensors deployed in nine modules, each with eight different sensors, distributed across the full 1.2 m width of the tunnel at a location 1.18 m from the inlet (Fig. 1c). The nine modules were identical to one another. To maintain the generality of the algorithm rather than optimize it for this particular dataset, we here sampled the 72 sensors naïvely, without in any way cross-referencing inputs from the nominally identical sensors replicated across the nine modules, or attempting to mitigate the plume-based variance across these sensors. The turbulent plume shown in Fig. 1c is illustrative only; distribution maps of local concentrations in the plume, along with

full details of the wind tunnel configuration, are provided in the publication first presenting the dataset⁴.

Ten different odourants were delivered in the gas phase to the sensor array: acetone, acetaldehyde, ammonia, butanol, ethylene, methane, methanol, carbon monoxide, benzene and toluene⁴. For every tunnel configuration, each of these odourants was presented 10–20 times, and each presentation lasted for 180 s. In our present study, we consider one of these 180 s plumes (chosen at random) for each odourant.

We discretized each sensor's range of possible responses into 16 levels of activation, corresponding to 16 time bins of the permissive epoch of each gamma cycle. The discretized sensor values were composed into a 72-dimensional sensor activity vector, which then was sparsened by setting the smallest 50% of the values to zero. Accordingly, each odourant sample ('sniff') presented to the EPL network comprised a discrete 72-element sensor vector drawn from a single point in time and presented as steady state. The training set underpinning one-shot learning was based on single-timepoint samples drawn from the 90 s timepoint in each of the 180-s-long odourant presentations. Test sets for the impulse-noise-only studies (Figs. 3 and 4) comprised these same timepoints, each altered by 100 different instantiations of impulse noise. For the plume-variance studies (Fig. 5), test samples for each odourant were drawn from different timepoints in the corresponding plume (specifically, across the range 30–180 s after odourant presentation, at 5 s intervals) and were presented to the network both with and without added impulse noise.

The OB EPL model therefore was instantiated with 72 columns, such that each column received afferent excitation proportional to the activation level of one sensor. Because we here present the network in its simplest form, we treated the 72 columns as independent inputs, without crafting the algorithm to combine the responses of duplicate sensor types, to weight the centrally located sensors more strongly, or to perform any other dataset-specific modifications that might improve performance. Each model OB column comprised one principal neuron (MC) and initially five GC interneurons that were presynaptic to that MC (for a total of 360 GCs across all columns), though the number of GCs per column rose as high as 50 in the most highly trained models described herein. MCs projected axons globally across all columns and formed excitatory synapses onto GCs with a uniform probability of 0.2 (20%). Each GC, in turn, synaptically inhibited the MC within its column with a probability of unity (100%). GCs did not inhibit MCs from other columns, though this constraint can be relaxed without affecting overall network function. To reflect the mapping of the algorithm to the physical layout of the Loihi chip, we consider an MC and its co-columnar GCs to be spatially local to one another. However, there is no computational basis for the physical locations of neurons in the model; an OB column is simply 'an MC plus those inputs that can affect its activity'. MC and GC model implementation details are presented in the Supplementary Methods.

Intrinsic gamma and theta dynamics. In the biological system, the profile of spike times across MCs is proposed to reflect a phase precedence code with respect to the emergent gamma-band field potential oscillations generated in the olfactory system. Spike timing-based coding metrics are known to offer considerable speed and efficiency advantages^{52–56}; however, they require computational infrastructure in the brain to realize these benefits. Fast oscillations in the local field potential are indicative of broad activity coherence across a synaptically coordinated ensemble of neurons, and thereby serve as temporal reference frames within which spike times in these neurons can be regulated and decoded. Accordingly, these reference frames are essential components of the biological system's computational capacities.

In the OB, gamma oscillations emerge from interactions of the subthreshold oscillations of MCs with the network dynamics of the EPL (PRING dynamics¹⁴). For present purposes, the importance of these oscillations was twofold: (1) MC spike phases with respect to the gamma-band oscillations serve as the model's most informative output, and (2) by considering each oscillation as embedding a distinct, interpretable representation, repeated oscillations enable the network to iteratively approach a learned state based on stationary sensory input. Notably, in vivo, piriform cortical pyramidal neurons are selectively activated by convergent, synchronous MC spikes⁵⁷, and established neural learning rules are in principle capable of reading such a coincidence-based metric⁵⁸. Because MC spike times can be altered on the gamma timescale by synaptic inhibition from GCs, and their spike times in turn alter the responsiveness of GCs, these lateral inhibitory interactions can iteratively modify the information exported from the OB. In the neuromorphic EPL, each MC periodically switched between two states to establish the basic gamma oscillatory cycle. These two states were an active state in which the MC integrated sensory input and generated spikes (permissive epoch) and an inactive state in which the excitation level of the MC was held at zero, preventing sensory integration and spike generation (inhibitory epoch; Fig. 1d). The effects of the plastic lateral inhibitory weights from GCs were applied on top of this temporal framework (Supplementary Methods). The correspondence with real time is arbitrary and hence is measured in timesteps (ts) directly; that said, as Loihi operates at about 100 kHz, each timestep corresponds to about 10 μ s. In the present implementation, the permissive epoch comprised 16 ts and the inhibitory epoch 24 ts, for a total of 40 ts per gamma cycle. Notably, the duration of the permissive

epoch directly corresponds to the number of discrete levels of sensory input that can be encoded by our spike timing-based metric; it can be expanded arbitrarily at the cost of greater time and energy expenditures.

A second, slower, sampling cycle was used to regulate odour sampling. This cycle is analogous to theta-band oscillations in the OB, which are driven primarily by respiratory sampling (sniffing) behaviours but also by coupling with other brain structures during certain behavioural epochs. Each sampling cycle ('sniff') consisted of a single sample and steady-state presentation of sensory input across five gamma cycles of network activity. The number of gamma cycles per sampling cycle can be arbitrarily determined to regulate how much sequential, iterative processing is applied to each sensory sample, but was held at five for all experiments herein.

Importantly, these differences between the slower sampling timescale and the faster processing timescale can be leveraged to implement 'continuous' online sampling, in which each sample can be processed using multiple computational iterations before digitizing the next sample. In the present implementation, for example, the Vergara et al. dataset sampled odourants at 100 Hz—one sample every 10 ms. On Loihi, operating at 100 kHz, the 200 ts (five gamma cycles) used for the processing of a single sniff require a total of around 2 ms. As this is five times faster than the sampling rate of the sensors, there would be no update to sensor state during the time required for five cycles of processing. Examples of this algorithm operating in 'continuous' mode are presented in Supplementary Fig. 1.

Testing procedures. After training, we tested the network's performance on recognizing learned odourants in the presence of destructive interference from unpredictable sources of olfactory occlusion (impulse noise), alone or in combination with variance arising from sampling plume dynamics at different timepoints. All testing was performed with network plasticity disabled.

The responses of primary olfactory receptors to a given odourant of interest can be radically altered by the concomitant presence of competing background odourants that strongly activate or block some of the same receptors as the odourant of interest, greatly disrupting the ratiometric activation pattern across receptors on which odour recognition depends. We modelled this occlusion as destructive impulse noise. Specifically, an occluded test sample was generated by choosing a fraction P of the 72 elements of a sensor activity vector and replacing them each with random values drawn uniformly from the sensors' operating range (integer values from 0 to 15). When multiple occluded test samples were generated to measure average performance, both the identities of the occluded elements and the random values to which they were set were redrawn from their respective distributions.

Odour plume dynamics comprise a second source of stimulus variance encountered under natural conditions. To test network performance across this variance, we drew test samples from different timepoints within the odour plumes. Specifically we drew 30 samples per plume at 5 s intervals between 30 s and 180 s within the 180 s datastreams. After one-shot training with a single sample, we tested network performance on the other samples, with and without the addition of impulse noise (Fig. 5).

While the present study focuses on one-shot learning, the network can also be configured for few-shot learning, in which it gradually adapts to the underlying statistics of training samples. In this configuration, the network learns robust representations even when the training samples themselves are corrupted by impulse noise. We illustrate this effect in Supplementary Fig. 2.

Sample classification. The pattern of MC spikes in each successive gamma cycle was recorded as a set of spikes, with each spike defined by the identity of the active MC and the spike latency with respect to the onset of that permissive epoch. Accordingly, five successive sets of spikes were recorded for each sample 'sniff'. When an impulse noise-occluded sample was presented to the network, the similarities were computed between each of the five representations evoked by the unknown and each of the network's learned odour representations. In descriptive figures (but not for comparisons with other methods), the similarity between two representations was measured with the Jaccard index, defined as the number of spikes in the intersection of two representations, divided by the number of spikes in their union⁴⁴. Specifically, the permissive epoch of a gamma cycle included 16 discrete timesteps in which MCs could spike; these 16 bins were used for Jaccard calculations. Test samples were classified as one of the network's known odourants if the similarity exceeded a threshold of 0.75 in the fifth (final) gamma cycle. If similarities to multiple learned odourants crossed the threshold, the odourant exhibiting the greatest similarity value across the five gamma cycles was picked as the classification result. If none of the similarity values crossed the threshold within five gamma cycles, the odourant was classified as unknown. This combination of nearest-neighbour classification and thresholding enabled the network to present 'none of the above' as a legitimate outcome. Summary figures each consist of averages across 100 independent instantiations of impulse noise, and/or averages across 30 different test samples drawn from different timepoints in the datastream (without or with added impulse noise), for each odour in the training set.

Benchmarks. We first compared the classification performance of the EPL network to three conventional signal processing techniques: MF, TVF and PCA (Fig. 6a)^{5,51}.

The MF and TVF are filters commonly used in signal processing for reducing impulse noise, while PCA is a standard preprocessor used in machine olfaction applications³. The MF used a window size of 5, and was implemented with the Python signal processing library `scipy.signal`. The TVF used a regularization parameter equal to 0.5, and was implemented using the Python image processing library `scikit-image`. PCA was implemented using the Python machine learning library `scikit-learn`; data were projected onto the top five components.

Corrupted input signals also can be denoised by training an autoencoder, a modern rendition of autoassociative networks^{59,60}. We therefore compared the performance of the EPL network to a seven-layer deep autoencoder constructed using the Python deep learning library `Keras`. The seven layers consisted of an input layer of 72 units, followed by five hidden layers of 720 units each and an output layer of 72 units. This resulted in a network of 3,744 units, identical to the number in the EPL model when trained with ten odours. The network was fully connected between layers, and the activity of each unit in the hidden layers was L1 regularized. The network was trained with iterative gradient descent until convergence using the `Adadelta` optimizer with a mean absolute error loss function. Its training set consisted of 7,000 examples per odourant class. For the same training set, the performance of this seven-layer autoencoder exceeded that of shallower networks (six-, five-, four- and three-layer networks were tested).

For direct comparison, the outputs of all of these methods, including that of the EPL network, were presented to the same nearest-neighbour classifier for sample classification according to a Manhattan distance metric. Specifically, for each of the techniques, the output was read as a 72-dimensional vector and normalized such that their elements summed up to a value of unity. (In the case of the EPL network, the spiking output in each gamma cycle was read out as a 72-dimensional rank-order vector and normalized so that the elements summed to unity). The similarity between any two such vectors was measured as $1/(1+d)$, where d is the Manhattan distance between the two vectors. Classification performance was measured by computing this similarity between the output of training data samples and those of test data samples. A test data sample was classified according to the identity of the training data sample to which it was most similar, provided that this similarity value exceeded a threshold of 0.75 (thresholding enabled the inclusion of a 'none of the above' outcome).

We trained the DAE in three different ways for comparison with EPL network performance. First, the DAE was trained using the same ten non-occluded odour samples that were used to train the EPL model. These ten samples underwent 1,000 training epochs to ensure training convergence. This method assesses DAE performance on 'one sample' learning, for comparison with the one-sample/one-trial learning of the EPL network (Fig. 6a). Second, we trained the DAE on multiple impulse noise-occluded samples, to maximize its performance. Specifically, we trained the DAE on 500 to 7,000 training samples, where each sample comprised an independently occluded instance of each of the ten odourants. Each training set was presented for 25 training epochs to ensure convergence. The occlusion levels for each training sample were drawn from the same distribution as the test samples, being randomly and uniformly selected from the range $P = [0.2, 0.8]$. With this procedure, we show that the DAE requires 3,000 training samples per odourant to achieve the classification performance that the EPL model achieved with 1 training sample per odourant (Fig. 6c); that is, the EPL model is 3,000 times more data efficient than the DAE. Third, we trained the DAE and EPL models first on one odourant (toluene) and then, subsequently, on a second odourant (acetone) to compare the models' sequential online learning capabilities. After training on toluene, the DAE classified test presentations of toluene with high fidelity (Fig. 6c, left). However, over the course of acetone training, the similarity between test samples of toluene and the learned representation of toluene progressively declined (Fig. 6d), to the point that the DAE network became unable to correctly classify toluene (Fig. 6c, right; see ref.⁴² for strategies to improve the sequential learning performance of deep networks). In contrast, training the EPL network with acetone exhibited no interference with the pre-existing toluene representation (Fig. 6d, inset). The similarity between test samples of toluene and the learned representation of toluene was not affected as the EPL learned all of the ten odourants in sequence (Fig. 6e).

Data availability

The Vergara et al. gas sensor dataset⁴ is freely available from the UCI Machine Learning database (<http://archive.ics.uci.edu/ml/datasets/gas+sensor+arrays+in+open+sampling+settings>).

Code availability

A software version of the model reproducing the primary results of Figs. 3–5 is freely available from the ModelDB public archive (<http://modeldb.yale.edu/261864>). The Loihi source code is freely available from Github (<https://github.com/intel-nrc-ecosystem/models/tree/master/official/epl>).

Received: 15 June 2019; Accepted: 7 February 2020;

Published online: 16 March 2020

References

- Davies, M. et al. Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro* **38**, 82–99 (2018).
- Rinberg, D., Koulakov, A. & Gelperin, A. Sparse odor coding in awake behaving mice. *J. Neurosci.* **26**, 8857–8865 (2006).
- Marco, S. & Gutierrez-Galvez, A. Signal and data processing for machine olfaction and chemical sensing: a review. *IEEE Sens. J.* **12**, 3189–3214 (2012).
- Vergara, A. et al. On the performance of gas sensor arrays in open sampling systems using inhibitory support vector machines. *Sens. Actuat. B* **185**, 462–477 (2013).
- Goff, S. A. & Klee, H. J. Plant volatile compounds: sensory cues for health and nutritional value? *Science* **311**, 815–819 (2006).
- Araneda, R. C., Kini, A. D. & Firestein, S. The molecular receptive range of an odorant receptor. *Nat. Neurosci.* **3**, 1248–1255 (2000).
- Araneda, R. C., Peterlin, Z., Zhang, X., Chesler, A. & Firestein, S. A pharmacological profile of the aldehyde receptor repertoire in rat olfactory epithelium. *J. Physiol.* **555**, 743–756 (2004).
- Rospars, J. P. Interactions of odorants with olfactory receptors and other preprocessing mechanisms: how complex and difficult to predict? *Chem. Senses* **38**, 283–287 (2013).
- Rospars, J. P., Lansky, P., Chaput, M. & Duchamp-Viret, P. Competitive and noncompetitive odorant interactions in the early neural coding of odorant mixtures. *J. Neurosci.* **28**, 2659–2666 (2008).
- Persaud, K. C., Marco, S. & Gutierrez-Galvez, A. *Neuromorphic Olfaction* (CRC Press, 2013).
- Raman, B., Stopfer, M. & Semancik, S. Mimicking biological design and computing principles in artificial olfaction. *ACS Chem. Neurosci.* **2**, 487–499 (2011).
- Cleland, T. A. Construction of odor representations by olfactory bulb microcircuits. *Prog. Brain Res.* **208**, 177–203 (2014).
- Bathellier, B., Lagier, S., Faure, P. & Lledo, P. M. Circuit properties generating gamma oscillations in a network model of the olfactory bulb. *J. Neurophysiol.* **95**, 2678–2691 (2006).
- Li, G. & Cleland, T. A. A coupled-oscillator model of olfactory bulb gamma oscillations. *PLoS Comput. Biol.* **13**, e1005760 (2017).
- Peace, S. T. et al. Coherent olfactory bulb gamma oscillations arise from coupling independent cylindrical oscillators. Preprint at <https://doi.org/10.1101/213827> (2018).
- Kashiwadani, H., Sasaki, Y. F., Uchida, N. & Mori, K. Synchronized oscillatory discharges of mitral/tufted cells with different molecular receptive ranges in the rabbit olfactory bulb. *J. Neurophysiol.* **82**, 1786–1792 (1999).
- Linster, C. & Cleland, T. A. Decorrelation of odor representations via spike timing-dependent plasticity. *Front. Comput. Neurosci.* **4**, 157 (2010).
- Masquelier, T., Hugues, E., Deco, G. & Thorpe, S. J. Oscillations, phase-of-firing coding, and spike timing-dependent plasticity: an efficient learning scheme. *J. Neurosci.* **29**, 13484–13493 (2009).
- Banerjee, A. et al. An interglomerular circuit gates glomerular output and implements gain control in the mouse olfactory bulb. *Neuron* **87**, 193–207 (2015).
- Cleland, T. A., Johnson, B. A., Leon, M. & Linster, C. Relational representation in the olfactory system. *Proc. Natl Acad. Sci. USA* **104**, 1953–1958 (2007).
- Cleland, T. A. et al. Sequential mechanisms underlying concentration invariance in biological olfaction. *Front. Neuroeng.* **4**, 21 (2012).
- Miller, P. Itinerancy between attractor states in neural systems. *Curr. Opin. Neurobiol.* **40**, 14–22 (2016).
- Tong, M. T., Kim, T. P. & Cleland, T. A. Kinase activity in the olfactory bulb is required for odor memory consolidation. *Learn. Mem.* **25**, 198–205 (2018).
- Chu, M. W., Li, W. L. & Komiyama, T. Lack of pattern separation in sensory inputs to the olfactory bulb during perceptual learning. *eNeuro* **4**, 5 (2017).
- Doucette, W. & Restrepo, D. Profound context-dependent plasticity of mitral cell responses in olfactory bulb. *PLoS Biol.* **6**, e258 (2008).
- Sultan, S. et al. Learning-dependent neurogenesis in the olfactory bulb determines long-term olfactory memory. *FASEB J.* **24**, 2355–2363 (2010).
- Moreno, M. M. et al. Olfactory perceptual learning requires adult neurogenesis. *Proc. Natl Acad. Sci. USA* **106**, 17980–17985 (2009).
- Gao, Y. & Strowbridge, B. W. Long-term plasticity of excitatory inputs to granule cells in the rat olfactory bulb. *Nat. Neurosci.* **12**, 731–733 (2009).
- Lepousez, G. et al. Olfactory learning promotes input-specific synaptic plasticity in adult-born neurons. *Proc. Natl Acad. Sci. USA* **111**, 13984–13989 (2014).
- de Almeida, L., Idiart, M. & Linster, C. A model of cholinergic modulation in olfactory bulb and piriform cortex. *J. Neurophysiol.* **109**, 1360–1377 (2013).
- Devore, S. & Linster, C. Noradrenergic and cholinergic modulation of olfactory bulb sensory processing. *Front. Behav. Neurosci.* **6**, 52 (2012).
- Li, G., Linster, C. & Cleland, T. A. Functional differentiation of cholinergic and noradrenergic modulation in a biophysical model of olfactory bulb granule cells. *J. Neurophysiol.* **114**, 3177–3200 (2015).

33. Kermen, F., Sultan, S., Sacquet, J., Mandairon, N. & Didier, A. Consolidation of an olfactory memory trace in the olfactory bulb is required for learning-induced survival of adult-born neurons and long-term memory. *PLoS ONE* **5**, e12118 (2010).
34. Lepousez, G., Valley, M. T. & Lledo, P. M. The impact of adult neurogenesis on olfactory bulb circuits and computations. *Annu. Rev. Physiol.* **75**, 339–363 (2013).
35. McIntyre, A. B. & Cleland, T. A. Biophysical constraints on lateral inhibition in the olfactory bulb. *J. Neurophysiol.* **115**, 2937–2949 (2016).
36. Lagier, S. et al. GABAergic inhibition at dendrodendritic synapses tunes gamma oscillations in the olfactory bulb. *Proc. Natl Acad. Sci. USA* **104**, 7259–7264 (2007).
37. McTavish, T. S., Migliore, M., Shepherd, G. M. & Hines, M. L. Mitral cell spike synchrony modulated by dendrodendritic synapse location. *Front. Comput. Neurosci.* **6**, 3 (2012).
38. Borthakur, A. & Cleland, T. A. A neuromorphic transfer learning algorithm for orthogonalizing highly overlapping sensor array responses. In *ISOCS/IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)* 1–3 (IEEE, 2017).
39. Borthakur, A. & Cleland, T. A. Signal conditioning for learning in the wild. In *Proc. Neuro-inspired Computational Elements Workshop 8* (Association for Computing Machinery, 2019).
40. Borthakur, A. & Cleland, T. A. A spike time-dependent online learning algorithm derived from biological olfaction. *Front. Neurosci.* **13**, 656 (2019).
41. French, R. M. Catastrophic forgetting in connectionist networks. *Trends Cogn. Sci.* **3**, 128–135 (1999).
42. Kirkpatrick, J. et al. Overcoming catastrophic forgetting in neural networks. *Proc. Natl Acad. Sci. USA* **114**, 3521–3526 (2017).
43. Chow, S. F., Wick, S. D. & Riecke, H. Neurogenesis drives stimulus decorrelation in a model of the olfactory bulb. *PLoS Comput. Biol.* **8**, e1002398 (2012).
44. Levandowsky, M. & Winter, D. Distance between sets. *Nature* **234**, 34–35 (1971).
45. Moreno, M. M. et al. Action of the noradrenergic system on adult-born cells is required for olfactory learning in mice. *J. Neurosci.* **32**, 3748–3758 (2012).
46. Hasselmo, M. E. & Giocomo, L. M. Cholinergic modulation of cortical function. *J. Mol. Neurosci.* **30**, 133–135 (2006).
47. Mandairon, N. et al. Cholinergic modulation in the olfactory bulb influences spontaneous olfactory discrimination in adult rats. *Eur. J. Neurosci.* **24**, 3234–3244 (2006).
48. Mandairon, N. et al. Context-driven activation of odor representations in the absence of olfactory stimuli in the olfactory bulb and piriform cortex. *Front. Behav. Neurosci.* **8**, 138 (2014).
49. Strowbridge, B. W. Role of cortical feedback in regulating inhibitory microcircuits. *Ann. N. Y. Acad. Sci.* **1170**, 270–274 (2009).
50. Adams, W., Graham, J. N., Han, X. & Riecke, H. Top-down inputs drive neuronal network rewiring and context-enhanced sensory processing in olfaction. *PLoS Comput. Biol.* **15**, e1006611 (2019).
51. Huang, Y.-M., Ng, M. K. & Wen, Y.-W. Fast image restoration methods for impulse and Gaussian noises removal. *IEEE Sig. Proc. Lett.* **16**, 457–460 (2009).
52. Thorpe, S., Delorme, A. & Van Rullen, R. Spike-based strategies for rapid processing. *Neural Netw.* **14**, 715–725 (2001).
53. Hopfield, J. J. Pattern recognition computation using action potential timing for stimulus representation. *Nature* **376**, 33–36 (1995).
54. Izhikevich, E. M. Polychronization: computation with spikes. *Neural Comput.* **18**, 245–282 (2006).
55. Maass, W. in *Models of Neural Networks IV. Physics of Neural Networks* (eds van Hemmen, J. L. et al.) 373–402 (Springer, 2002).
56. Maass, W. & Markram, H. On the computational power of circuits of spiking neurons. *J. Comput. Syst. Sci.* **69**, 593–616 (2004).
57. Luna, V. M. & Schoppa, N. E. GABAergic circuits control input-spike coupling in the piriform cortex. *J. Neurosci.* **28**, 8851–8859 (2008).
58. Feldman, D. E. The spike-timing dependence of plasticity. *Neuron* **75**, 556–571 (2012).
59. Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y. & Manzagol, P.-A. Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion. *J. Mach. Learn. Res.* **11**, 3371–3408 (2010).
60. Xie, J., Xu, L. & Chen, E. Image denoising and inpainting with deep neural networks. In *Proc. 25th International Conference on Neural Information Processing Systems* (eds Pereira, F. et al.) 341–349 (NeurIPS, 2012).

Acknowledgements

This work was funded by the National Institute on Deafness and Other Communication Disorders (R01 DC014701 and R01 DC014367 to TAC).

Author contributions

T.A.C. conceived the neural algorithm. N.I. instantiated the algorithm and enhanced the learning rules. T.A.C. and N.I. wrote the manuscript.

Competing interests

The authors declare competing interests as follows. The underlying platform-independent algorithm is the subject of a Cornell University patent application on which the authors are listed as inventors. N.I. is currently employed by Intel Labs, developers of the Loihi neuromorphic system. T.A.C. is a member of the Intel Neuromorphic Research Community and has received research funding from Intel for related work.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s42256-020-0159-4>.

Correspondence and requests for materials should be addressed to N.I. or T.A.C.

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 2020