



Neural-inspired sensors enable sparse, efficient classification of spatiotemporal data

Thomas L. Mohren^{a,b}, Thomas L. Daniel^b, Steven L. Brunton^a, and Bingni W. Brunton^{b,1}

^aDepartment of Mechanical Engineering, University of Washington, Seattle, WA 98195; and ^bDepartment of Biology, University of Washington, Seattle, WA, 98195

Edited by John G. Hildebrand, University of Arizona, Tucson, AZ, and approved August 16, 2018 (received for review May 23, 2018)

Sparse sensor placement is a central challenge in the efficient characterization of complex systems when the cost of acquiring and processing data is high. Leading sparse sensing methods typically exploit either spatial or temporal correlations, but rarely both. This work introduces a sparse sensor optimization that is designed to leverage the rich spatiotemporal coherence exhibited by many systems. Our approach is inspired by the remarkable performance of flying insects, which use a few embedded strain-sensitive neurons to achieve rapid and robust flight control despite large gust disturbances. Specifically, we identify neural-inspired sensors at a few key locations on a flapping wing that are able to detect body rotation. This task is particularly challenging as the rotational twisting mode is three orders of magnitude smaller than the flapping modes. We show that nonlinear filtering in time, built to mimic strain-sensitive neurons, is essential to detect rotation, whereas instantaneous measurements fail. Optimized sparse sensor placement results in efficient classification with approximately 10 sensors, achieving the same accuracy and noise robustness as full measurements consisting of hundreds of sensors. Sparse sensing with neural-inspired encoding establishes an alternative paradigm in hyperefficient, embodied sensing of spatiotemporal data and sheds light on principles of biological sensing for agile flight control.

sparse sensing | neural encoding | sensory arrays | sparse optimization | insect flight control

In both living systems and modern technology, there is a tension between gathering vast and increasing quantities of heterogeneous data (e.g., the internet of things) and acquiring targeted data gathered by specialized sensors (1, 2). Large numbers of sensors would provide extensive information about the system and its environment but may, in turn, command high energetic costs. Indeed, big data demand synthesis and significant processing, often to identify which few features of the data are meaningful, particularly when the crucial information is obscured by large, nonrelevant signals or noise. In contrast, each specialized sensor can extract features tailored to the signal, but unanticipated features in the data may be lost. The tradeoff between flexibility and efficiency relies in part on the relative difficulty of acquiring, transforming, and performing complex computations on the data. In addition, local computations alleviate expensive data transfers and may reduce the latency of a decision. Here, we focus on understanding and designing systems with sparse and efficient sensing strategies that leverage both correlations in space and dynamics in time.

Recent advances in sparse sensing rely on the observation that many signals in nature exhibit relatively simple, low-dimensional patterns, so that signal reconstruction or classification can be achieved with a small subset of all possible sensors. In particular, compressed sensing theory states that if the information of a signal x is sparse in a transformed basis Ψ , then the signal may be reconstructed from relatively few incoherent measurements (3–6). The number of measurements may be further reduced by taking two additional perspectives. First, if we do not use a universal transform basis (e.g., Fourier, wavelets, etc.) but instead learn Ψ from training data, sensor selection may be tai-

lored to a specific task (7). Second, when only classification is required, reconstruction can be circumvented and the number of measurements needed is orders of magnitude smaller still (8). Here we use the sparse sensor placement optimization for classification (SSPOC) (9) approach to identify the locations of a few key strain sensors tailored to inform body rotation.

We turn to flight control in insects as inspiration for a sensing strategy by which temporal and spatial information is combined. Flying insects are remarkably adept at making rapid and robust corrections to stabilize their body orientation in response to gusts. This robust flight control relies on multimodal integration of visual and mechanical information; vision is crucial for flight—indeed, insects rarely fly without it—yet the slow timescale of visual processing cannot support the rapid maneuvers observed in free flight (10–12). Insects accomplish this task using mere tens to hundreds of neurons acting as strain sensors distributed on their bodies (13–16), despite the complexity of the surrounding fluid dynamics (17). Efficient, distributed sensing and computing have also been explored in nature-inspired engineering (18); some examples include insights gained from flying insects (19–21), birds and bats (22, 23), and fish (24–27). In particular, flying insects sense mechanical deflections using neurons associated with mechanosensory structures known as campaniform sensilla on their wings (28, 29) and halteres (30–32). Halteres are structures derived from wings and function as gyroscopes;

Significance

Winged insects perform remarkable aerial feats in uncertain, complex fluid environments. This ability is enabled by sensation of mechanical forces to inform rapid corrections in body orientation. Curiously, mechanoreceptor neurons do not faithfully report forces; instead, they are activated by specific time histories of forcing. We find that, far from being a bug, neural encoding by biological sensors is a feature that acts as built-in temporal filtering superbly matched to detect body rotation. Indeed, this encoding further enables surprisingly efficient detection using only a small handful of neurons at key locations. Nature suggests smart data as an alternative strategy to big data, and neural-inspired sensors establish a paradigm in hyperefficient sensing of complex systems.

Author contributions: T.L.D. and B.W.B. designed research; T.L.M. performed research; T.L.M., T.L.D., S.L.B., and B.W.B. analyzed data; and T.L.M., T.L.D., S.L.B., and B.W.B. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Data deposition: All of the code for modeling and classification used in this study has been deposited on GitHub and is available at <https://github.com/tlmohren/Mohren.WingSparseSensors>.

See Commentary on page 10545.

¹To whom correspondence should be addressed. Email: bbrunton@uw.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1808909115/-DCSupplemental.

Published online September 13, 2018.

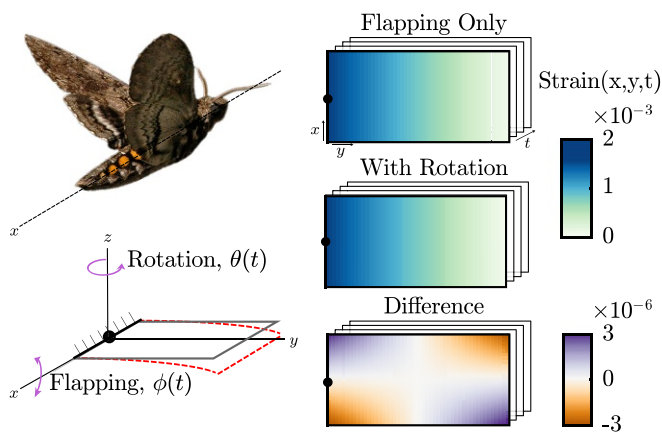


Fig. 1. A simulated flapping wing model with and without externally induced body rotation differs by a twisting mode three orders of magnitude smaller than the dominant flapping mode. From the flapping wing simulation, we obtain span-wise normal strain over a dense grid on the wing as a function of space (x, y) and time t . Hawkmoth image courtesy of Armin J. Hinterwirth (photographer).

however, most insects lack halteres, and it is believed that wings serve a similar function. In this work, we focus on campaniform sensilla on wings, where they have been shown to mediate flight posture control (29) and encode mechanical stimulus features (33). Even so, they do not resemble typical engineered sensors,

as they do not directly report physical measurement quantities. Instead, mechanosensitive neurons encode physical strain by a transformation that may be summarized as a temporal filter followed by a nonlinear activation function (32, 34); this encoding has been well characterized in animal experiments (33).

In this paper, we combine SSPOC with mechanical modeling to show that neural-inspired encoding of mechanical strain experienced by a flapping wing is crucial for reliable, efficient classification of spatiotemporal data associated with body rotation. Body rotation leads to Coriolis acceleration, which induces spatiotemporal twisting modes on the wing (35). Classifying between flapping wings with and without body rotation is challenging in part because these twisting modes are three orders of magnitude smaller than the flapping modes (Fig. 1). In contrast, we show that merely 10 neural-inspired sensors placed at key locations can achieve similar classification accuracy as a dense grid of sensors distributed over the entire wings. We find that this performance is robust to large, noisy disturbances added to the biomechanical wing model. Further, the experimentally derived nonlinear encoder is not unique; instead, exploration of filter function space reveals a large plateau of similar encoders that perform comparably well at this classification task. Analyzing the locations of these few, key neural-inspired sensors offers mechanistic clues to how biology senses in this hyperefficient regime.

Neural-Inspired Sparse Sensors

Here we take a reverse-engineering perspective to ask, What is the smallest number of strain-sensitive neurons required

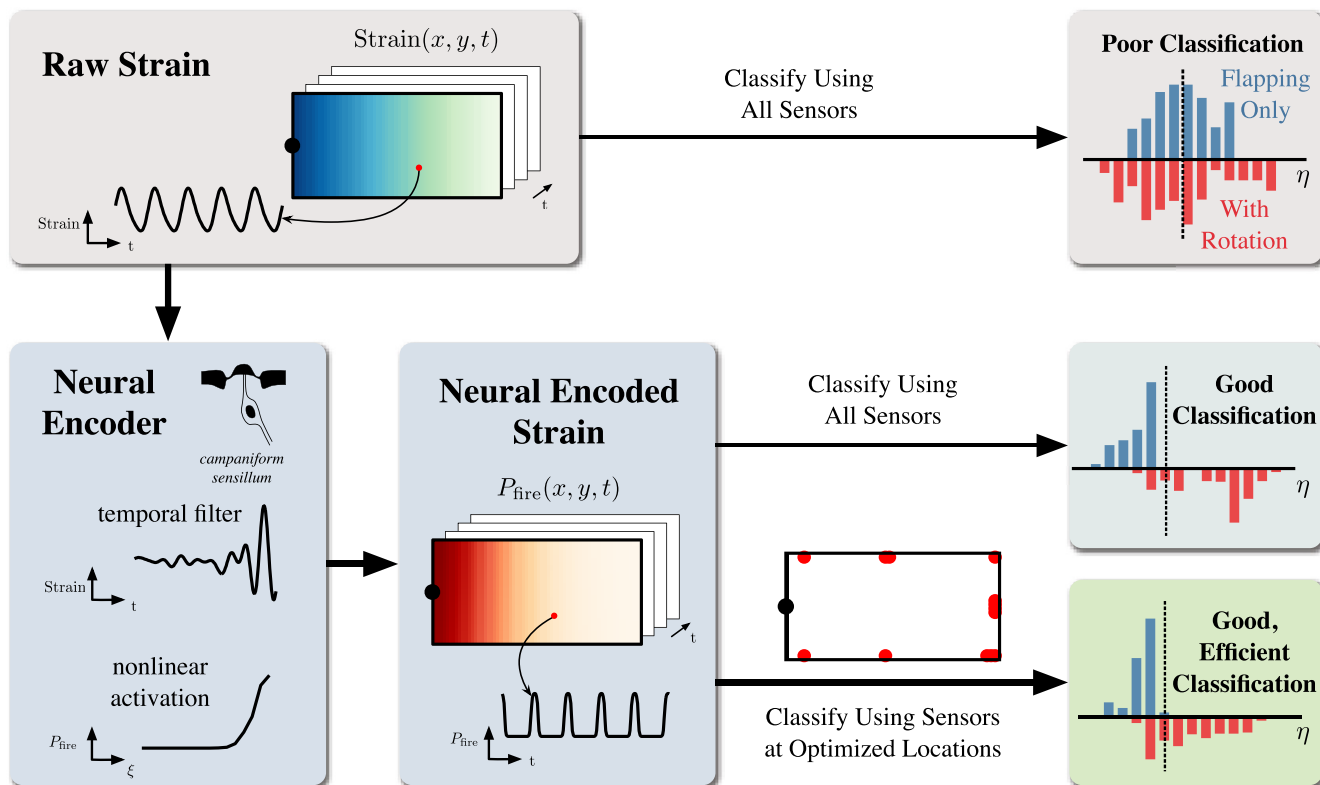


Fig. 2. A schematic of classifying body rotation using sparse neural-inspired strain sensors placed on a flapping wing. Raw span-wise normal strain is obtained from the structural simulation in two conditions, flapping only and flapping with rotation (35). The raw strains in these two conditions are not linearly separable, leading to poor classification even using all of the sensor locations. Alternatively, raw strain is encoded by a neural-inspired filter and transformed into the probability of a mechanosensory neuron firing an action potential (33, 36). The neural encoder is approximated by experimental recordings of campaniform sensilla and summarized as a temporal spike-triggered average (STA) filter followed by a nonlinear activation function, transforming raw strain into probability of the neuron firing an action potential P_{fire} . We define P_{fire} to be a neural-encoded strain. The neural-encoded strain separates well with a linear classifier; further, this performance can be achieved remarkably efficiently using approximately 10 sensors at key locations (9).

to inform body rotation, and where should they be placed? Answering these questions requires an integrated approach, combining tools from biomechanical simulations, neurophysiology, and sparse optimization. This analysis will demonstrate the role of neural-inspired nonlinear filtering in time and how it enables a dramatic reduction in the number of required sensors through sparse optimization in space. All of the code for modeling and classification is openly available and can be found at github.com/tlmohren/Mohren_WingSparseSensors, and details of our approach are found in *SI Appendix*.

First, we simulate a flapping wing using a Euler–Lagrange model with parameters based on a hawkmoth (35). The flapping wing produces spatiotemporal strain fields sampled at a dense grid on the wing. We consider two conditions, given by flapping without body rotation and with body rotation of 10 rad/s (Fig. 1). Through a Coriolis force, wing flapping combined with body rotation in an orthogonal axis activates a very small twisting mode in the strain field, and detecting this rotation is a significant challenge. We use the simulation data to train a supervised machine-learning classifier to distinguish between flapping with and without rotation. Random perturbations are added to the flapping and rotational velocities, and the classification accuracy is assessed on validation data from simulations that were not used in training.

To gauge the role of neural encoding in this task, we compare the performance of classifiers trained using either raw strain from the structural model or neural-encoded strain (Fig. 2). The encoding performed by single mechanosensory neurons on the insect wing is approximated by two functions, both of which are derived directly from neurophysiological experiments (33). In short, extracellular recordings of nerve action potentials were made at the wing hinge while mechanical stimuli were delivered to the wing tip through a motor. Analysis of the mechanical features leading to action potentials were summarized in a temporal spike-triggered average (STA) filter followed by a nonlinear activation function (36). We define the neural-encoded strain data as the probability of a mechanosensory neuron firing an action potential.

Next, we solve for the locations of a small subset of sensors among the dense grid of sensors on the wing that are sufficient to support classification. Sensor locations are selected by exploiting the inherent sparsity in the training data. Our approach uses sparsity-promoting regression and is an extension of SSPOC (9). Starting with the truncated basis Ψ and the discriminant vector between the two categories \mathbf{w} , we solve for a sparse vector $\mathbf{s} \in \mathbb{R}^n$ that achieves the discrimination $\Psi^T \mathbf{s} = \mathbf{w}$. Here, \mathbf{s} has the same shape as the full-state discriminant vector $\Psi \mathbf{w}$ but contains mostly zeros. In particular, we use an elastic net penalty to formulate the sparse optimization problem (37–39),

$$\mathbf{s} = \arg \min_{\mathbf{s}'} \left\| \mathbf{w} - \Psi^T \mathbf{s}' \right\|_2 + \alpha \left\| \mathbf{s}' \right\|_1 + (1 - \alpha) \left\| \mathbf{s}' \right\|_2, \quad [1]$$

where $\|\cdot\|_2$ is the ℓ_2 norm, $\|\cdot\|_1$ is the ℓ_1 norm, and α is a hyperparameter of the optimization (here we use $\alpha = 0.9$ to emphasize the sparsity-promoting term). The few nonzero elements of \mathbf{s} correspond to desired sensor locations; these few sensors are able to closely match the performance of full-state classification.

Results

Our primary result is that classification of flapping with and without body rotation requires neural-inspired encoding of strain data. In addition, only a few neural-inspired sensors are needed for classification, showing remarkable robustness to large-magnitude disturbances. We further characterize how well a family of neural-inspired encoders, including the one derived directly from experimental recordings, are able to perform this classification.

Neural-Inspired Encoders Are Essential. The raw strain data reveal that body rotation orthogonal to the axis of flapping introduces a torsional mode in the flapping wing orthogonal to the axis of flapping (*SI Appendix*, Fig. S2) (35). Although a signature of the rotation is measurable by strain, the magnitude of this torsional mode is three orders of magnitude smaller than that of the flapping modes (Fig. 1). In addition, raw strain data of flapping only vs. flapping with rotation are not linearly separable (Fig. 3, black diamond). The data under both conditions are overlapping sinusoidal time series at the wing flap frequency; therefore, it is not possible to construct any linear hyperplane that separates them.

In contrast, neural-encoded strain (Fig. 2) enables a linear classifier to detect body rotation, achieving accuracy on validation data of 90% (Fig. 3, red diamond). In particular, a STA temporal filter selects a short time history of raw strain that matches the activation of strain-sensitive wing mechanoreceptors, and a nonlinear activation function transforms the raw strain into a probability of firing an action potential, which we define to be the neural-encoded strain. To highlight the importance of the neural encoding, we extensively tested other nonlinear classifiers on the raw unfiltered strain data (*SI Appendix*, Table S1), and none of them consistently approach the accuracy obtained with even a simple linear classifier on neural-encoded strain data.

A Few Key Neural-Inspired Sensors Are Required. Importantly, very few of the neural-inspired sensors in the simulation are required

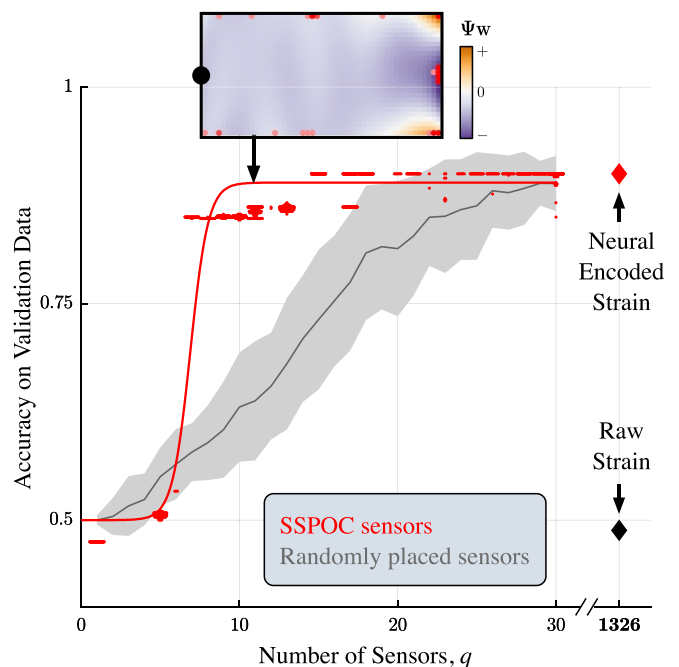


Fig. 3. Classification using about 10 neural-encoded sensors placed at key locations on the wing achieves accuracy comparable with that of classification using all sensors. Flapping wing structural simulations were computed with moderate disturbance amplitudes ($[\dot{\phi}^*, \dot{\theta}^*] = [0.31, 0.1]$ rad/s). The classification accuracies shown are validated on an epoch of the simulation separate from what was used for training. Sparse sensors are learned from training data from trials with random disturbances using SSPOC (red, each dot is an individual trial) and compared with randomly placed sensors (gray, mean and SD). The red line is a sigmoidal fit to the SSPOC sensors accuracy. *Inset* shows a probability distribution of SSPOC sensor locations on the wing for $q = 11$ sensors, averaged over 100 training sets with random instances of noisy disturbances. The opacity of the red dots is proportional to the likeliness of sensor solutions at that location; most sensors are found at the periphery of the wing. *Inset* background shows the full-state discriminant vector $\Psi \mathbf{w}$.

for classification, achieving accuracy approaching what is possible with all sensors. This performance is made possible by exploiting the inherent low rank structure of the data, which is evident in the singular value spectrum of neural-encoded strain (*SI Appendix*, Fig. S10). Although the raw strain data are even lower rank than the neural-encoded strain, flapping with and without body rotation remains not linearly separable (*SI Appendix*, Fig. S11).

Indeed, ~ 25 randomly placed sensors perform just as well on average as using all 1,326 sensors (Fig. 3, gray curve). It is possible to further reduce the number of sensors by selecting optimized locations, and ~ 10 SSPOC sensors achieve comparable performance (Fig. 3, red dots). The relationship between the number of SSPOC sensors q and validated accuracy follows a sigmoidal shape (Fig. 3, red curve).

The optimized sensor locations are shown in Fig. 3, *Inset* for $q = 11$ sensors, and they are distributed at distinct locations at the periphery of the wing. These locations include the far edge of the wing away from the body, where the full-state discriminant vector $\Psi\mathbf{w}$ has large amplitude.

Classification Is Robust to Disturbances. The few key sensors discovered by the SSPOC optimization reliably classify body rotation even when the magnitude of disturbances is large. Fig. 4 shows the validated classifier accuracies for increasing disturbances in both the flapping ϕ and rotational θ axes. Smaller disturbances support classification with fewer sensors. Even so, the asymptotic full-state accuracy is approached for all rotational disturbances less than 10 rad/s, at which the disturbances equal the steady rotation velocity. The performance of sparse sensors is characterized for finer resolutions of disturbances in *SI Appendix*, Fig. S12, and the probability distribution of sensors at each disturbance is shown in *SI Appendix*, Fig. S13.

Interestingly, when the mean classification degrades for larger disturbances, the distribution of accuracy at a given number of sensors q becomes bimodal. In other words, sparse sensor optimization on some sets of training data achieves accuracy that approaches the asymptotic full-state accuracy, but other ran-

dom instances lead to poor classification. Comparing the sensor location distributions for the good classification vs. the poor classification cases, we see that sensors at the far edge of the wing away from the body are crucial for classification (*SI Appendix*, Fig. S14).

Variations on the Theme of Experimentally Derived Neural Encoders.

So far, we have used a parameterized neural encoder fitted directly to electrophysiological recordings of campaniform sensilla in insects (33). Now, we explore the effects of systematic variations to the neural encoder's parameters to determine whether the experimentally derived encoder is uniquely suited to the task. The temporal filter and the nonlinear activation function both have two parameters each. We vary each pair of parameters while holding the others fixed at their experimentally derived values.

The performance achievable by this family of neural encoders is summarized by the fewest sensors required to achieve 75% classification accuracy. For each encoder, a full sweep of validated accuracy is computed with at least 10 iterations of random disturbance at each value of q . A sigmoidal fit of the relationship between q and accuracy (as in the red curve in Fig. 3) is then used to determine at what q the accuracy exceeds 75%. For some regimes in the encoder parameter space, this accuracy is never achieved for any number of sensors.

The temporal filter STA has two parameters, frequency and width. Fig. 5, *Top* row shows that the experimentally derived STA (in red boxes) is surrounded by a large plateau in parameter space with comparably STA-like functions. Further, higher-frequency filters tend to perform better, whereas the width of the filter is less crucial as long as it is not too narrow. In the limit of the narrowest STA, the temporal filter acts as an identity and does not transform the data; in other words, here the encoding is achieved by the nonlinear activation function alone. The fact that this regime of parameter space is still able to classify rotation, albeit requiring a larger number of sensors, hints at the importance of the nonlinearity. This observation is corroborated by the results of nonlinear classifiers trained on raw strain data (*SI Appendix*, Table S1). The STA acts as a temporal filter for disturbances, and without it, classification accuracy degrades for larger noise amplitudes (*SI Appendix*, Fig. S12).

Campaniform sensilla nonlinear activation functions generally have a sigmoidal shape [although variations have been observed (32, 33)]. Fig. 5, *Bottom* row shows that the half-maximum (half-max) of the sigmoidal function does not impact classification accuracy. Similarly, the precise slope of the sigmoid is not crucial, as long as it is not too sharp or too shallow. In the limit of unit slope with zero half-max (middle of leftmost column of parameter space in Fig. 5, *Center* column), the nonlinear activation function becomes linear. Without this nonlinearity, classification never achieves 75% accuracy; in other words, the nonlinear activation function dramatically improves classification.

Although the experimentally derived neural encoders are well suited to perform body rotation classification, they are not unique. In the context of this nature-inspired classification task, the observed properties of campaniform sensilla are found in a large parameter space of similar encoders, most of which can support robust and sparse classification of body rotation.

Discussion

This paper takes inspiration from nature to demonstrate how classification of subtle dynamic regimes in spatiotemporal data can be achieved with remarkably few sensors. Specifically, we explore how strain-sensitive neurons on a flapping wing can detect body rotation, an ethologically relevant task for flying insects. We show that the task can be accomplished efficiently

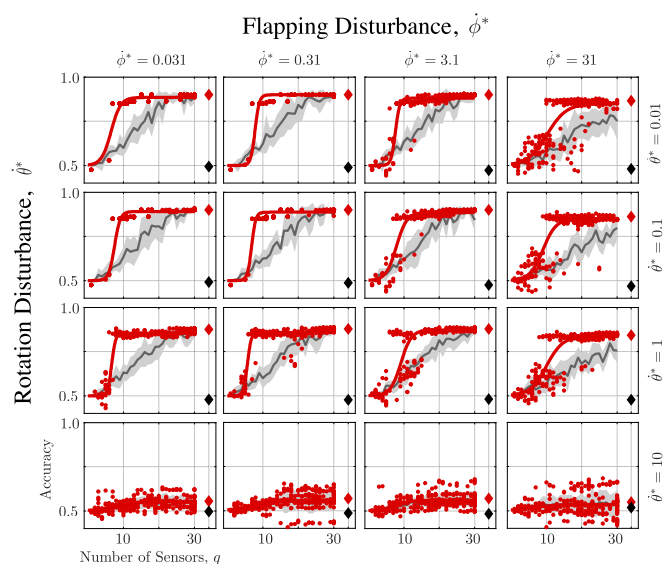


Fig. 4. Classification accuracy is robust for moderate- to large-magnitude disturbances in flapping ϕ and in rotation θ . Each panel of the 4×4 grid shows the classification accuracy for varying numbers of sensors (gray, random sensors; red, SSPOC sensors; black diamond, all sensors without encoding; red diamond, all sensors with encoding). The levels of rotation disturbance represent 0.1%, 1%, 10%, and 100% of the SD of steady flapping $\bar{\phi}$ and of the magnitude of constant rotation $\bar{\theta}$.

framework motivates development of hardware demonstrations using flexible materials (48, 49). Recent innovations in 3D printing technology have enabled manufacturing of flexible structures with embedded strain sensors (50). Some of these sensors are capacitive devices with low temporal resolution (51), and they have been limited in number and energy budget on small devices. We suggest that our neural-inspired sensing perspective may pivot both of these limitations into advantages in the design of autonomous microrobotic implementations.

- Boyd D, Crawford K (2011) *Six Provocations for Big Data* (Social Science Research Network, Rochester, NY), SSRN Scholarly Paper 1926431.
- Lazer D, Kennedy R, King G, Vespignani A (2014) The parable of google flu: Traps in big data analysis. *Science* 343:1203–1205.
- Candès EJ, Romberg J, Tao T (2006) Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans Inf Theor* 52:489–509.
- Donoho DL (2006) Compressed sensing. *IEEE Trans Inf Theor* 52:1289–1306.
- Baraniuk RG (2007) Compressive sensing. *IEEE Signal Process Mag* 24:118–120.
- Romberg J (2008) Imaging via compressive sampling. *IEEE Signal Process Mag* 25:14–20.
- Manohar K, Brunton BW, Kutz JN, Brunton SL (2018) Data-driven sparse sensor placement. *IEEE Control Syst Mag* 38:63–86.
- Proctor JL, Brunton SL, Brunton BW, Kutz JN (2014) Exploiting sparsity and equation-free architectures in complex systems. *Eur Phys J Spec Top* 223:2665–2684.
- Brunton BW, Brunton SL, Proctor JL, Kutz JN (2016) Sparse sensor placement optimization for classification. *SIAM J Appl Math* 76:2099–2122.
- Collett T, Land M (1975) Visual control of flight behaviour in the hoverfly *Syrirta pipiens* L. *J Comp Physiol* 99:1–66.
- Theobald JC, Warrant EJ, O’Carroll DC (2009) Wide-field motion tuning in nocturnal hawkmoths. *Proc R Soc Lond B Biol Sci* 277:853–860.
- Sponberg S, Dyrh JP, Hall RW, Daniel TL (2015) Luminance-dependent visual processing enables moth flight in low light. *Science* 348:1245–1248.
- Nalbach G, Hengstenberg R (1994) The halteres of the blowfly calliphora. *J Comp Physiol A Neuroethol Sensory Neural Behav Physiol* 175:695–708.
- Sherman A, Dickinson MH (2003) A comparison of visual and haltere-mediated equilibrium reflexes in the fruit fly *Drosophila melanogaster*. *J Exp Biol* 206:295–302.
- Sane SP, Dieudonné A, Willis MA, Daniel TL (2007) Antennal mechanosensors mediate flight control in moths. *Science* 315:863–866.
- Sponberg S, Krapp HG (2007) Sensory systems and flight stability: What do insects measure and why? *Adv Insect Physiol* 34:231–316.
- Bomphrey RJ, Nakata T, Phillips N, Walker SM (2017) Smart wing rotation and trailing-edge vortices enable high frequency mosquito flight. *Nature* 544:92–95.
- Brunton SL, Noack BR (2015) Closed-loop turbulence control: Progress and challenges. *Appl Mech Rev* 67:050801.
- Combes SA, Daniel TL (2001) Shape, flapping and flexion: Wing and fin design for forward flight. *J Exp Biol* 204:2073–2085.
- Faruque I, Humbert JS (2010) Dipteran insect flight dynamics. Part 1: Longitudinal motion about hover. *J Theor Biol* 264:538–552.
- Faruque I, Humbert JS (2010) Dipteran insect flight dynamics. Part 2: Lateral-directional motion about hover. *J Theor Biol* 265:306–313.
- Song A, et al. (2008) Aeromechanics of membrane wings with implications for animal flight. *AIAA J* 46:2096–2106.
- Clark RP, Smits AJ (2006) Thrust production and wake structure of a batoid-inspired oscillating fin. *J Fluid Mech* 562:415–429.
- Triantafyllou MS, Triantafyllou GS (1995) An efficient swimming machine. *Sci Am* 272:64–71.
- Allen JJ, Smits AJ (2001) Energy harvesting eel. *J Fluids Structures* 15:629–640.
- Whittlesey RW, Liska SC, Dabiri JO (2010) Fish schooling as a basis for vertical-axis wind turbine farm design. *Bioinspiration Biomimetics* 5:035005.
- Leftwich MC, Tytell ED, Cohen AH, Smits AJ (2012) Wake structures behind a swimming robotic lamprey with a passively flexible tail. *J Exp Biol* 215:416–425.
- Dombrowski UJ (1991) Untersuchungen zur funktionellen organisation des flugsystems von *manduca sexta* (l). Dissertation (Universität Köln).
- Dickerson BH, Aldworth ZN, Daniel TL (2014) Control of moth flight posture is mediated by wing mechanosensory feedback. *J Exp Biol* 217:2301–2308.
- Dickinson MH (1999) Haltere-mediated equilibrium reflexes of the fruit fly, *Drosophila melanogaster*. *Philos Trans R Soc Lond B Biol Sci* 354:903–916.
- Huston SJ, Krapp HG (2009) Nonlinear integration of visual and haltere inputs in fly neck motor neurons. *J Neurosci* 29:13097–13105.
- Fox JL, Fairhall AL, Daniel TL (2010) Encoding properties of haltere neurons enable motion feature detection in a biological gyroscope. *Proc Natl Acad Sci USA* 107:3840–3845.
- Pratt B, Deora T, Mohren T, Daniel T (2017) Neural evidence supports a dual sensory-motor role for insect wings. *Proc R Soc Lond B Biol Sci* 284:20170969.
- Fox J, Daniel T (2008) A neural basis for gyroscopic force measurement in the halteres of *holorusia*. *J Comp Physiol A* 194:887–897.
- Eberle AL, Dickerson BH, Reinhall PG, Daniel TL (2015) A new twist on gyroscopic sensing: Body rotations lead to torsion in flapping, flexing insect wings. *J R Soc Interface* 12:20141088.
- Aljadeff J, Lansdell BJ, Fairhall AL, Kleinfeld D (2016) Analysis of neuronal spike trains, deconstructed. *Neuron* 91:221–259.
- Tibshirani R (1996) Regression shrinkage and selection via the lasso. *J R Stat Soc Ser B Methodol* 73:267–288.
- Zou H, Hastie T (2005) Regularization and variable selection via the elastic net. *J R Stat Soc Ser Stat Methodol* 67:301–320.
- Clemmensen L, Hastie T, Witten D, Ersbøll B (2011) Sparse discriminant analysis. *Technometrics* 53:406–413.
- McCulloch WS, Pitts W (1943) A logical calculus of the ideas immanent in nervous activity. *Bull Math Biophys* 5:115–133.
- Rumelhart DE, Hinton GE, Williams RJ (1986) Learning representations by back-propagating errors. *Nature* 323:533–536.
- LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521:436.
- Goodfellow I, Bengio Y, Courville A, Bengio Y (2016) *Deep Learning* (MIT Press Cambridge, MA), Vol 1.
- Yamins DL, DiCarlo JJ (2016) Using goal-driven deep learning models to understand sensory cortex. *Nat Neurosci* 19:356–365.
- Barlow HB (1961) Possible principles underlying the transformations of sensory messages. *Sensory Communication*, ed Rosenblith W (MIT Press, Cambridge, MA), pp 216–234.
- Sharpee T, Rust NC, Bialek W (2004) Analyzing neural responses to natural signals: Maximally informative dimensions. *Neural Comput* 16:223–250.
- Brunton SL, Brunton BW, Proctor JL, Kaiser E, Kutz JN (2017) Chaos as an intermittently forced linear system. *Nat Commun* 8:19.
- Majidi C (2014) Soft robotics: A perspective—current trends and prospects for the future. *Soft Rob* 1:5–11.
- Dean WA, Ranganathan BN, Penskiy I, Bergbreiter S, Humbert JS (2017) Robust gust rejection on a micro-air vehicle using bio-inspired sensing. *Mechatronics and Robotics Engineering for Advanced and Intelligent Manufacturing*, eds Zhang D, Wei B (Springer International Publishing, Basel), pp 351–362.
- Muth JT, et al. (2014) Embedded 3d printing of strain sensors within highly stretchable elastomers. *Adv Mater* 26:6307–6312.
- Shin HS, Castano LM, Humbert JS, Bergbreiter S (2016) Sensing skin for detecting wing deformation with embedded soft strain sensors. *IEEE Sensors*, 10.1109/ICSENS.2016.7808417.