

Cortical Motion Perception Emerges from Dimensionality Reduction with Evolved Spike-Timing-Dependent Plasticity Rules

 Kexin Chen,¹ Michael Beyeler,^{3,4} and  Jeffrey L. Krichmar^{1,2}

¹Departments of Cognitive Sciences, ²Computer Science, University of California, Irvine, California 92697, ³Departments of Computer Science, and

⁴Psychological & Brain Sciences, University of California, Santa Barbara, California 93106

The nervous system is under tight energy constraints and must represent information efficiently. This is particularly relevant in the dorsal part of the medial superior temporal area (MSTd) in primates where neurons encode complex motion patterns to support a variety of behaviors. A sparse decomposition model based on a dimensionality reduction principle known as non-negative matrix factorization (NMF) was previously shown to account for a wide range of monkey MSTd visual response properties. This model resulted in sparse, parts-based representations that could be regarded as basis flow fields, a linear superposition of which accurately reconstructed the input stimuli. This model provided evidence that the seemingly complex response properties of MSTd may be a by-product of MSTd neurons performing dimensionality reduction on their input. However, an open question is how a neural circuit could carry out this function. In the current study, we propose a spiking neural network (SNN) model of MSTd based on evolved spike-timing-dependent plasticity and homeostatic synaptic scaling (STDP-H) learning rules. We demonstrate that the SNN model learns compressed and efficient representations of the input patterns similar to the patterns that emerge from NMF, resulting in MSTd-like receptive fields observed in monkeys. This SNN model suggests that STDP-H observed in the nervous system may be performing a similar function as NMF with sparsity constraints, which provides a test bed for mechanistic theories of how MSTd may efficiently encode complex patterns of visual motion to support robust self-motion perception.

Key words: dimensionality reduction; MSTd; optic flow; sparse coding; spiking neural network; visual motion processing

Significance Statement

The brain may use dimensionality reduction and sparse coding to efficiently represent stimuli under metabolic constraints. Neurons in monkey area MSTd respond to complex optic flow patterns resulting from self-motion. We developed a spiking neural network model that showed MSTd-like response properties can emerge from evolving spike-timing-dependent plasticity with STDP-H parameters of the connections between then middle temporal area and MSTd. Simulated MSTd neurons formed a sparse, reduced population code capable of encoding perceptual variables important for self-motion perception. This model demonstrates that complex neuronal responses observed in MSTd may emerge from efficient coding and suggests that neurobiological plasticity, like STDP-H, may contribute to reducing the dimensions of input stimuli and allowing spiking neurons to learn sparse representations.

Introduction

As an observer moves through the environment, the motion between oneself and the scene causes a change in the structure of

light, which is reflected on the retina as optic flow. Optic flow can be used to perceive the heading direction of self-motion and to guide locomotion (Warren and Hannon, 1988). Neurons in the dorsal visual pathway of the primate brain are selective for motion (Britten, 2008). The middle temporal (MT) area contains neurons that are tuned to the speed and direction of motion (Albright, 1984). Receiving primary input from MT (Raiguel et al., 1997), neurons in the dorsal subregion of the medial superior temporal (MSTd) area respond to large and complex optic flow patterns, including translation, rotation, expansion, contraction, and the intermediates of these motions (Graziano et al., 1994). MSTd was suggested to play an important role in heading perception. Physiologic studies showed that microstimulation of the

Received Feb. 18, 2022; revised May 16, 2022; accepted June 8, 2022.

Author contributions: K.C., M.B., and J.L.K. designed research; K.C. performed research; K.C. contributed unpublished reagents/analytic tools; K.C., M.B., and J.L.K. analyzed data; K.C. wrote the paper.

This work was supported by the National Science Foundation Grant IIS-1813785 and the Air Force Office of Scientific Research Grant FA9550-19-1-0306. We thank Cognitive Hardware and Software Ecosystem Community Infrastructure, supported through the National Science Foundation Grant CNS-1730158, for the computing resources.

The authors declare no competing financial interests.

Correspondence should be addressed to Kexin Chen at kexinc3@uci.edu.

<https://doi.org/10.1523/JNEUROSCI.0384-22.2022>

Copyright © 2022 the authors

MSTd area biased the perception of the monkey of heading directions based on optic flow (Britten and van Wezel, 1998; Gu et al., 2012). The causal role for self-motion perception was also established in the human functional equivalent of macaque area MST (hMST; Schmitt et al., 2020). Furthermore, heading direction can be decoded from the population activity of MSTd neurons (Hamed et al., 2003). MST was also suggested to participate in 3D velocity estimation (Mineault et al., 2012). These studies provide evidence for the role of MSTd in self-motion perception based on optic flow. However, it is not well understood how MSTd integrates input from MT to form complex selectivity to motion patterns and to perceive self-motion.

The brain is under tight metabolic constraints and uses numerous strategies to achieve efficient representations of information that allow for high performance and information transfer (Krichmar et al., 2019). Non-negative sparse coding (NSC) is an efficient population coding scheme that combines dimensionality reduction with sparsity constraints to reduce the number and activity of neurons to represent environmental features (Beyeler et al., 2019). A prior study described a computational model that can account for a wide range of MSTd visual response properties through applying a dimensionality reduction algorithm known as non-negative matrix factorization (NMF) with sparsity constraints, which implements a form of NSC to MT input activity (Beyeler et al., 2016). This MSTd model exhibited sparse, parts-based representations of the optic flow patterns that resembled the receptive fields observed in MSTd neurons. This model provided evidence that the seemingly complex response properties may emerge from MSTd neurons performing a biological equivalent of dimensionality reduction and sparse coding on their input. However, it remains to be shown how these coding schemes can be implemented in neurobiological circuits.

The current study shows how neurobiologically plausible synaptic plasticity could implement a form of NSC. Learning processes in the brain are believed to occur in the synaptic changes among neurons. Hebb's (1949) rule describes a mechanism that leads to long-term potentiation (LTP) between excitatory neurons given their consistent coactivity. Oja's (1982) learning rule adds constrained synaptic modification on top of the Hebbian model to prevent saturation of synaptic weights. In the spiking domain, the spike-timing-dependent plasticity (STDP) learning rule modifies synaptic weights according to the relative timing of spikes of the presynaptic and postsynaptic neurons (Sjöström and Gerstner, 2010). These synaptic learning rules perform statistical inferences on the input data. Oja's (2008) rule was shown to reduce the dimensionality of datasets similar to performing principal component analysis (PCA). Analogously, through mathematical derivation, STDP in combination with homeostatic scaling (STDP-H) was shown to have the same effect on the input data as NMF (Carlson et al., 2013). We propose that STDP-H could have a similar effect on MSTd receptive fields.

In the present study, we implement a spiking neural network (SNN) model of monkey MSTd based on evolved STDP-H parameters. We show that by learning to reconstruct the input stimuli with STDP-H, the network extracts representations of the input in a form that resembles the receptive fields of MSTd neurons. The spiking neurons in the network show selectivity to spiral motions and 3D heading direction similar to the selectivity observed in the brain. The present results suggest a neurobiologically plausible method for producing sparse and reduced receptive fields in the dorsal visual system and possibly other brain regions (Beyeler et al., 2019).

Materials and Methods

An SNN model of visual cortical areas MT and MSTd was created to reconstruct optic flow patterns resulting from self-movement. Figure 1 shows the overall architecture of the SNN model in which input optic flow stimulus was processed by an array of MT neurons, whose responses were then converted to Poisson spike trains and projected to a group of excitatory spiking neurons that simulated MSTd. The MSTd group projected to a group of inhibitory neurons, and the inhibitory neurons provided feedback inhibition to the MSTd group to regulate the network activity. Connection weights in the network were updated with an unsupervised learning rule that implemented STDP and STDP-H. The parameters of the learning rule were optimized through evolutionary algorithms, with an objective function measuring how well the network reconstructed the input stimuli. Neural activity was modeled with Izhikevich (2004) neurons, with regular spiking for excitatory neurons and fast spiking for inhibitory neurons.

SNN simulations were performed using CARLsim 5 software (Chou et al., 2018; Balaji et al., 2020). Parameter optimization was done with the parameter tuning interface (PTI) in CARLsim 5, which used the Evolutionary Computations in JAVA library (Luke, 2017). The evolutionary process, which included network training and validation, was performed with 1 graphics processing unit (GPU) node and 15 central processing unit (CPU) cores on the computing clusters supported by the Cognitive Hardware and Software Ecosystem Community Infrastructure (CHASE-CI; Altintas et al., 2019). We used parallel execution enabled by CARLsim and PTI to distribute computations among the GPU and CPU cores, which greatly accelerated the simulation processes. In each generation of evolutionary computation, PTI launched a total of 50 separate SNN instances that were simulated in parallel and evaluated concurrently.

Input stimuli. Input stimuli to the network were computer-generated arrays of optic flows that resembled those used in physiological experiments. To simulate the apparent motion on the retina that would be caused by an observer moving in a 3D environment, we used the model of the motion field proposed by Longuet-Higgins and Prazdny (1980), where a pinhole camera with focal length $f = 1$ cm was pointed to a frontoparallel plane located at a distance d , which simulated 3D real world points at a reference depth Z . The pinhole camera was used to project the 3D real world points, $\vec{P} = [X, Y, Z]^t$ onto a 2D image plane, $\vec{p} = [x, y]^t = f/Z[X, Y]^t$ (i.e., the retina). When the camera moved in the 3D environment, the local motion of points on the 2D image plane was described as a function of the motion of the camera and the location of the points. In our simulation, the camera moved with a translational velocity $\vec{v} = (v_x, v_y, v_z)$ and a rotational velocity $\vec{\omega} = (\omega_x, \omega_y, \omega_z)$. The optic flow component $(\dot{x}, \dot{y})$ on the 2D image plane at location (x, y) could then be computed as follows:

$$\begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} = \underbrace{\frac{1}{d} \begin{pmatrix} -f & 0 & x \\ 0 & -f & y \end{pmatrix} \begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix}}_{\text{translational flow}} + \underbrace{\frac{1}{f} \begin{pmatrix} x \cdot y & -(f^2 + x^2) & f \cdot y \\ (f^2 + y^2) & -x \cdot y & -f \cdot x \end{pmatrix} \begin{pmatrix} \omega_x \\ \omega_y \\ \omega_z \end{pmatrix}}_{\text{rotational flow}} \quad (1)$$

The first component in the right side of the equation computed the translational flow, and the second component computed the rotational flow. These two components superimposed linearly, and the rotational flow was independent of the depth d . Simulated flow fields were represented as 15×15 pixel arrays, subtending a visual angle of $90^\circ \times 90^\circ$.

For the training and validation of the network, we generated a dataset containing 1280 flow field stimuli. Each optic flow stimulus was presented to the network for 500 ms, interleaved by 500 ms of a blank stimulus. To simulate motion patterns resulting from locomotion, the linear velocity of the pinhole camera was set to 1 m/s for stimuli that included a translational flow and 0 m/s for the pure rotation motion. The angular velocity was set to 1 radian/s for stimuli that included a rotational flow, which was within the range of natural head rotation movements (Cullen, 2019), and 0 radian/s for the pure translation motion. The depth was set to 1 m from the observer. Stimuli in the training and validation dataset were sampled uniformly from a laminar motion space and a spiral motion space. The laminar motion




Figure 1. Model design. **A**, Architecture of the SNN model. Optic flow stimulus was first processed by an array of MT neurons tuned to speed and direction of motion. The MT activity was converted to Poisson spike trains as the input to the network. The MT neuron group was connected to the MSTd group with a Gaussian projection, which allowed the MSTd neurons to receive input from MT neurons that locate in different areas of the visual field and were tuned to different direction and speed. The MSTd group was reciprocally connected to a group of inhibitory neurons, which regulated the activity of the network. All connections in the network were modulated by STDP-H. The MT→MSTd connection weights and the MSTd group activity were used to reconstruct the input. A fitness function measured the network performance based on the reconstruction accuracy, and STDP-H parameters were evolved with EC to optimize the fitness function. **B**, Direction tuning curves of the simulated MT neurons. **C**, Speed tuning curves of the simulated MT neurons. The horizontal axis is plotted on a $\log_2$ scale. deg, Degrees; sec, second.

space contained unidirectional motion patterns that simulated eight directions (45° intervals in 360°) of translation. For these stimuli, the rotational component in Equation 1 was 0, and in the translational component, $v_x = \{0, \pm \sqrt{\frac{1}{2}}, \pm 1\}$, $v_y = \{0, \pm \sqrt{\frac{1}{2}}, \pm 1\}$, and $v_z = 0$.

The spiral motion space contained the following motion patterns: expansion/contraction, clockwise/counterclockwise rotation, and intermediate spiral motions, namely clockwise expanding, clockwise contracting, counterclockwise expanding, and counterclockwise contracting motions. The center of motion (COM) of these stimuli was kept near the center of the visual field, with $(v_x, v_y) \approx 0$ and $v_z = \pm 1$ for stimuli that had a translational component, and $(v_x, v_y) \approx 0$ and $v_z = \pm 1$ for stimuli that had a rotational component. To introduce a slight variability in the COM, we converted the velocities to spherical coordinates and applied a Gaussian noise with SD of 10° to the x and y velocity components before converting them back into Cartesian coordinates and using Equation 1 to compute for optic flow components. Schematic drawings of these motion patterns are shown in Figure 2.

We also generated three other datasets, which were used for testing the SNNs after training, by following the protocols used in neurophysiological studies (Graziano et al., 1994; Hamed et al., 2003; Gu et al., 2006; Takahashi et al., 2007). We used these datasets to quantify the response properties, including the Gaussian tuning in spiral space and the 3D translation and rotation heading selectivity, as well as the population encoding of heading, of our simulated MSTd neurons. These measurements allowed us to compare the selectivity of our simulated MSTd neurons to the one observed in the real brain. It is worth noting that these

three separate datasets contained stimuli not present in the training and validation dataset. We computed the fitness scores achieved by the evolved and trained models on these three datasets, which allowed us to measure the generalization ability of the model and ensure that the performance was not limited to the validation dataset.

Model design. The SNN model consisted of three neuron groups, MT, MSTd, and an inhibitory group (Fig. 1). The network model processed flow field stimuli in the following three steps: (1) Input flow fields were first processed by the MT group, (2) the MSTd group received Poisson spike trains from the MT group, and (3) the input was reconstructed with a dot product of the MT→MSTd connection weights and the MSTd neuronal activation.

The MT neuron group was organized into 3D grids with dimensions of $15 \times 15 \times 40$. The first two dimensions corresponded to the pixels of the input flow fields that subtended $90^\circ \times 90^\circ$ of the visual angle. The third dimension corresponded to the direction and speed tuning profile of the neurons. At each spatial location, there were a total of 40 MT-like model units (selective for eight directions times five speeds of motion). This organization can be seen as having multiple layers of MT neurons covering the visual field, with each layer of neurons tuned to the same direction and speed of motion but responding to different locations (Fig. 1). Neurons in this group were modeled as idealized MT neurons selective to particular combinations of speed and direction of motion. Each MT neuron had a receptive field of one pixel area of the simulated flow fields, corresponding to 3° of visual angle. In this study, the firing rate of an MT neuron R_{MT} responding to the motion flow at the location (x, y) was given by the following:




Figure 2. Schematic drawings of motion patterns the training and validation dataset was sampled from. The first two rows depict motion patterns in the laminar motion space, and the last two rows depict motion patterns in the spiral motion space.

$$R_{MT}(x, y; \theta_{pref}, \rho_{pref}) = d(x, y; \theta_{pref})s(x, y; \rho_{pref}), \quad (2)$$

where the component $d(x, y; \theta_{pref})$ computed the direction response of the neuron, and $s(x, y; \rho_{pref})$ computed the speed response of the neuron.

The direction response $d(x, y; \theta_{pref})$ was described by a circular von Mises function (Fig. 1B) as follows:

$$d(x, y; \theta_{pref}) = \exp(\sigma_{theta}(\cos(\theta(x, y) - \theta_{pref}) - 1)), \quad (3)$$

where the direction response was a function of the difference between the stimulus direction $\theta(x, y)$ and the preferred direction θ_{pref} of the neuron. The tuning bandwidth parameter σ_{theta} was set to three, corresponding to a tuning width of $\sim 90^\circ$.

The speed response $s(x, y; \rho_{pref})$ was described by a log Gaussian function (Fig. 1C), which approximated MT tuning curves observed in the macaque brain (Nover, 2005). The speed response of an MT neuron was given by the following:

$$s(x, y; \rho_{pref}) = \exp\left(-\frac{\log\left(\frac{\rho(x, y) + s_0}{\rho_{pref} + s_0}\right)^2}{2\sigma^2}\right), \quad (4)$$

where the speed response was a function of the difference between the stimulus speed $\rho(x, y)$ and the preferred speed ρ_{pref} of the neuron. The

free parameter σ determined the width of the curve and was set to 1.16. The speed offset parameter s_0 was set to 0.33, which prevented undefined logarithm values when stimulus speed approaches zero. These parameter settings were consistent with the median values across the MT population (Nover, 2005).

The MT response R_{MT} computed from Equation 2 represented the normalized firing rate of the neuron ranging from zero to one and was used to generate spike trains following a Poisson distribution with a maximum firing rate of 20 Hz. Neurons in the MT group were connected to the MSTd group following a Gaussian distribution, with which neurons that were spatially closer to each other had a higher probability of being connected and had higher initial connection weights. We defined a 2D Gaussian connectivity along the first two dimensions of the MT neuron group, and all layers of the MT neurons shared the same connectivity scheme so that each MSTd neuron had a circular receptive field covering MT neurons at different locations and tuned to different direction and speed of motion. The width of the Gaussian curve, which corresponded to the radius r of this circular receptive field, was an open parameter that was evolved through the evolutionary computation (EC) process. Activity of the MSTd group was regulated by an inhibitory neuron group, which contained the same number of neurons as the MSTd group. The MSTd neurons provided feedforward excitation to the inhibitory neuron group, and the inhibitory group provided feedback inhibition to the MSTd neurons. Both connections followed a uniform random connectivity with a 10% probability. All intergroup connections were modulated by STDP-H plasticity during training, which include the




Figure 3. Stimulus reconstruction. STDP-H performed dimensionality reduction on the input matrix $\mathbf{V}$ and decomposed it into two smaller matrices $\mathbf{W}$, the MT→MSTd weights, and $\mathbf{H}$, the MSTd activity. Each column of the input matrix represented an input instance $\vec{v}_i$, which was visualized as the original flow field. In the figure, the MT→MSTd connection weights $\mathbf{W}$ were visualized as a group of basis flow fields. The MSTd activation to this particular stimulus, $\vec{h}_i$, was represented as a column in the matrix $\mathbf{H}$, which denoted the degree of activation of the corresponding basis flow field. Darker color in this visualization corresponded to a higher level of activity. Right, The reconstructed input. The correlation score between the original and the reconstruction of this particular input instance was 0.84.

projection from MT to MSTd (MT→MSTd), from MSTd to the inhibitory group (MSTd→Inh), and from the inhibitory group back to the MSTd group (Inh→MSTd). The parameter values for STDP-H were selected through EC (see below, Optimization of the model via EC).

The resulting MSTd activity was used to reconstruct the input stimulus along with the MT→MSTd connection weights. Analogous to the sparse decomposition model proposed by Beyeler et al. (2016), the MT activity was represented as a multivariate matrix $\mathbf{V}$, where each column within the matrix represented an instance of input stimulus $\vec{v}_i$, namely, the population activity of the MT group responding to a flow field stimulus. The connection weights between the MT and MSTd groups were represented as a matrix $\mathbf{W}$, where w_{ij} was the connection weight between MT neuron i and MSTd neuron j . The activity of the MSTd neurons was represented as the coefficients in the matrix $\mathbf{H}$, where each column in the matrix $\vec{h}_i$ denoted the MSTd population activity responding to the input stimulus $\vec{v}_i$. To test the hypothesis of whether STDP-H allows the network to learn a reduced representation of the input $\mathbf{V}$ by decomposing it into two lower rank matrices $\mathbf{W}$ and $\mathbf{H}$, we took a dot product of the two matrices to reconstruct the input matrix $\mathbf{V}$. Effectively, a particular instance of input stimulus $\vec{v}_i$ was reconstructed by taking the inner product of the connection weights $\mathbf{W}$ and the MSTd activation $\vec{h}_i$ (Fig. 3).

The second column in Figure 3 shows a visualization of the connection weights between the MT and MSTd neuron groups. Each grid shows the receptive field of a MSTd neuron, determined by calculating a population vector for direction and speed of motion from $\mathbf{W}$ (Georgopoulos et al., 1982). These receptive fields acted as basis flow fields. With the MSTd neuron activity indicating the degree of activation, a linear superposition of the basis flow fields reconstructed the input flow field pattern.

STDP-H. All three intergroup connections (i.e., MT→MSTd, MSTd→Inh, and Inh→MSTd) were plastic, whose weight values were modulated by the STDP learning rule in combination with STDP-H. Homeostatic scaling was applied in a multiplicative manner, which modified synaptic weights based on the average postsynaptic firing rate R (Turrigiano et al., 1998; Turrigiano and Nelson, 2004; Carlson et al., 2013). As illustrated in Figure 4, homeostatic scaling adjusted synaptic properties (i.e., connection weight values) to keep the activity of the neurons close to the target firing rate. The total effect of STDP-H on a particular synapse w_{ij} connecting presynaptic neuron i and postsynaptic neuron j can be described as follows:

$$\frac{dw_{ij}}{dt} = \underbrace{[\alpha \cdot w_{ij}(1 - \bar{R}/R_{\text{target}})]}_{\text{homeostasis}} + \underbrace{[\beta(LTP_{ij} + LTD_{ij})]}_{\text{STDP}} \cdot K. \quad (5)$$

The α and β terms controlled the strength of STDP and homeostasis. We fixed the value of β to be one, and evolved the value of α for the




Figure 4. Illustration of the effect of homeostatic synaptic scaling on an individual neuron. As the synaptic drive (i.e., input synaptic weights) increases, the firing rate of the neuron increases and exceeds the target firing rate. Homeostatic scaling brings the activity down to the target zone by decreasing the input strength. If the synaptic drive is too low, and the activity of the neuron is below the target firing rate, homeostatic scaling raises the input strength and brings the activity of the neuron back into the target firing zone. Synaptic drive is in arbitrary units, and firing rate is normalized. Figure adapted from Turrigiano and Nelson (2004).

optimal relative strength of the two learning components. The parameter K was a term that damped oscillation in the weight updates and sped up learning, which was defined as follows:

$$K = \frac{\bar{R}}{T \cdot (1 + |1 - \bar{R}/R_{\text{target}}| \cdot \gamma)}. \quad (6)$$

Here, the parameter T was the time scale over which the firing rate of the postsynaptic neuron was averaged, and γ was a tuning factor, set to 10 and 50, respectively, in this study.

In Equation 5, the first component described the homeostatic scaling, which was a function of the ratio between the mean firing rate $\bar{R}$ and the target firing rate R_{target} of the neuron j . With homeostatic scaling, the rate of weight changes decreased if the neuron was overexcited and increased if the neuron was overly quiescent. The second component of Equation 5 described STDP, which was composed of LTP and long-term-depression (LTD). STDP strengthened or weakened the synaptic connection according to the timing of the presynaptic and postsynaptic spikes. LTP and LTD were described as decaying exponential functions, with a magnitude parameter A^+ or A^- , and a decay constant τ^+ or τ^- as follows:

$$LTP_{ij} = A^+ \exp\left(\frac{-(t_j - t_i)}{\tau^+}\right) \text{ for } t_j > t_i \quad (7)$$

$$LTD_{ij} = -A^- \exp\left(\frac{t_j - t_i}{\tau^-}\right) \text{ for } t_j \leq t_i. \quad (8)$$

Parameters A^+ , A^- , τ^+ , and τ^- were evolved through EC.

Training and validation of the model. From the training and validation dataset with 1280 flow fields, we randomly selected 160 stimuli as validation samples and the rest as training samples. In both the training and validation processes, each stimulus was presented to the network for 500 ms, followed by a blank stimulus presented for the same duration of time. During training, connection weights in the network were updated with STDP-H in an unsupervised manner. Parameters of the STDP-H learning rule were optimized through EC for each iteration of the training and validation process.

To validate whether the network had learned representations that allowed the reconstruction of unseen stimuli, we used the validation phase where we froze the connection weights and stopped STDP-H from modulating the connections. Flow fields from the validation set were presented to the network in the same way as during the training phase. Here, we recorded the activity of the MSTd neurons during the 500 ms of stimulus presentation. We reconstructed the input by multiplying $\mathbf{W}$, the connection weights between the MT and MSTd neuron groups, with $\hat{h}_i$, the corresponding MSTd neuronal response to the stimulus (see above, Model design). After obtaining a reconstruction for all flow fields from the validation set, we assembled them into a matrix $\hat{V}$, where each column contained the reconstructed MT activity pattern corresponding to an input flow field. We defined a fitness function, which accounted for both the reconstruction accuracy and network stability, as follows:

$$y = \text{corr}(V, \hat{V}) - \lambda L. \quad (9)$$

In the first term of the fitness function, we calculated the Pearson correlation between the input matrix V and the reconstructed matrix $\hat{V}$. To prevent instability in the network, we added the second term L , which described a penalty for high firing rates of the MSTd neuron population and was weighted by a scaling factor $\lambda = 0.001$. The penalty term was defined as follows:

$$L = \begin{cases} \max(\bar{R}_{MSTd}) - R_{MSTd,t} & , \text{ if } \max(\bar{R}_{MSTd}) > R_{MSTd,t} \\ 0 & , \text{ otherwise} \end{cases} \quad (10)$$

Here, $\bar{R}_{MSTd}$ was the mean firing rate of individual MSTd neurons, and $R_{MSTd,t} = 250$ Hz was the firing rate threshold. When the maximum mean firing rate of MSTd neurons exceeded the firing rate threshold, the network received a penalty score.

Optimization of the model via EC. Parameters of the STDP-H learning rule were optimized through EC. A total of 19 parameters divided into five groups were evolved, (1) STDP amplitude parameters A^+ and A^- and decay constants τ^+ and τ^- for all three intergroup connections (12 parameters), (2) target firing rates R_{MSTd} of the MSTd group and R_{Inh} of the inhibitory neuron group (2 parameters), (3) maximum connection weights for all three intergroup connections (three parameters), (4) the radius r of the Gaussian connection between MT and MSTd (one parameter), and (5) the parameter α that adjusted the relative strengths between STDP and homeostatic scaling for each of the three intergroup connections (one parameter). Table 1 shows the ranges of values for these evolved parameters.

We used an evolution strategy (ES) that took the form of $ES-(\mu, \lambda)$, in which $\mu = 5$ specified the size of the parent population, and $\lambda = 50$ specified the size of the offspring population (De Jong, 2006). In the first generation of the evolutionary process, the ES initialized the parameters for 50 individual networks. The network instances were then trained and validated with simulated flow fields in the training and validation dataset and were evaluated with the fitness function (see above, Training and

Table 1. Parameters evolved with evolutionary computation

Parameter	A^+	A^-	τ^+	τ^-	Target FR	Max Wt	Radius (r)	Alpha
Minimum	0	−4e−3	5.0	5.0	5.0	0.001	0.75*sqrt(B)	0.1
Maximum	4e−3	0	100.0	100.0	20.0	0.5	sqrt(B)	1.0

Target FR: Target firing rate; Max Wt: Maximum weight value.

validation of the model). These fitness scores were then transferred to the ES. The ES performed a binary tournament selection to select the best performing $\mu = 5$ individual networks as the parent networks for the next generation and modified their parameters through replication and mutation to produce a new population of $\lambda = 50$ network individuals. This optimization process continued for 30 generations.

Data availability. Custom code used in this study including neural network simulations and data analysis scripts are available at <https://github.com/kexinchenn/Sparse-Reduced-MSTd-SNN-Model>.

Results

The SNN model went through 30 generations of an evolutionary process that searched for optimal parameters of the STDP-H synaptic learning rule to maximize a fitness function that measured the accuracy of input reconstruction. Each evolutionary generation consisted of a training phase and a validation phase. Optic flow stimuli were generated by projecting from a pinhole camera moving at different combinations of linear and angular velocity to a frontoparallel plane. A total of 1280 computer-generated flow field stimuli that simulated different motion patterns were used in the training and validation phases. In the training phase, 1120 stimuli were randomly selected and presented to the network, and connection weights in the network were updated by STDP-H in an unsupervised manner. In the validation phase, 160 optic flow stimuli, which were not used in the training phase, were presented to the network with STDP-H disabled, and the fitness function was computed over the validation set of stimuli. After the evolutionary process, we presented several sets of novel stimuli used in physiological studies to examine the selectivity and response properties of the simulated MSTd neurons in the models.

Model performance and evolved STDP parameters

To investigate whether the SNN model was able to accurately reconstruct the input stimuli and to understand the effect of model size on the performance of the model, we experimented with different sizes of neuron populations for the MSTd group ($B = \{16, 36, 64, 100, 144\}$). All configurations shared the same network setup as described earlier (see above, Model design). The differences among different configurations were the size of the MSTd group, B , and the range of values for the Gaussian projection width, r , which was adjusted so that the MSTd neurons in all configurations covered similar portions of the visual field. To ensure reliable results from the parameter optimization process and to obtain a sizable pool of simulated MSTd neurons for subsequent analyses (a minimum of 144 MSTd neurons were required for the analysis; see below, Population encoding of heading direction), we completed 10 separate runs for the $B = 16$ configuration, so it had more variability, and five separate runs for all other configurations. In each run of the evolutionary process, a total of 1500 SNN instances were simulated (30 generations and 50 network individuals in each generation).

With a population of 50 individual networks in each generation, all five network configurations converged to optimal solutions within 30 generations (Fig. 5). For all configurations, fitness scores reached ~ 0.5 at the first generation of the evolutionary process and gradually increased over generations. At generation 30, all five configurations reached fitness scores > 0.65 (Table 2). As

shown in Table 2, the total time for completing 30 generations of optimization process and the simulation of 1500 network individuals increased with the size of the MSTd group. The evolutionary process of the smallest network with $B = 16$ MSTd neurons required 3.89 ± 0.75 d to complete, whereas the largest network with $B = 144$ MSTd neurons required 13.60 ± 2.19 d to complete. Networks with a larger MSTd group showed a better reconstruction ability and reached higher fitness scores. Among the evolved networks, configurations with $B = \{64, 100, 144\}$ MSTd neurons were able to approximate the input data with $R^2 > 0.7$, indicating an accurate reconstruction of the input stimuli. Although network configurations with more than $B = 144$ MSTd neurons may lead to an even better reconstruction accuracy, we showed in the sparseness analysis (see below, Sparse representations in simulated MSTd neurons) that the level of sparseness may not improve beyond the $B = 144$ configuration. To ensure that the reconstruction ability of the network models was not limited to the 160 samples randomly sampled from the training and validation datasets, we computed fitness scores that the evolved and trained network models achieved on the three separate datasets, which were composed of samples that were not present in the training and validation dataset (Table 3). The results showed that the models were able to generalize to other motion patterns. These datasets, which are from previous neurophysiological studies (Graziano et al., 1994; Hamed et al., 2003; Gu et al., 2006; Takahashi et al., 2007), are used in our subsequent analyses (see below, MSTd-like receptive fields and spiral selectivity, 3D translation and rotation heading selectivity, and Population encoding of heading direction).

STDP, as a biologically observed process, modifies synaptic strengths based on the time differences between presynaptic and postsynaptic spikes. The largest weight changes induced by STDP occur when the time difference between the presynaptic and postsynaptic spikes is small. As the time difference increases, the weight changes decrease exponentially. Typically, if a presynaptic neuron fires before the postsynaptic neuron, the connection is strengthened because intuitively, the presynaptic neuron caused the postsynaptic neuron to fire. If the order is reversed, the relationship between these neurons is uncorrelated, and that connection would typically be weakened. In the case of STDP for inhibitory neurons, the order might be reversed because if there is causality, the inhibition should be released (i.e., weaken the synapse). However it should be noted that many variations of timing and order have been observed experimentally.

We examined the evolved STDP parameters for the best performing network individuals with each configuration. In the evolutionary computation process, we searched the parameter space for the amplitude and time constant components of the STDP curves for LTD and LTP. Through the optimization toward an accurate reconstruction of the input spiking patterns, the STDP parameters evolved to show similar patterns across all five network model configurations as measured by the area over the LTP or LTD part of the STDP curves (Fig. 6). The MT→MSTd input connection showed stronger LTD effects than LTP. STDP for the inhibitory connections evolved the opposite pattern, the MSTd→Inh and the Inh→MSTd connections showed stronger




Figure 5. Best-so-far (BSF) fitness curve across 30 generations of evolutionary process for each network configuration. Solid lines denote the mean fitness scores of all individual runs of the network configuration. Shaded area denotes the SD of the fitness scores.

Table 2. Execution time and performance of all network configurations

Number of MSTd neurons	Execution time (days)	Fitness (BSF after 30 generations)	r^2
B = 16	3.89 ± 0.75	0.67 ± 0.02	0.45 ± 0.02
B = 36	5.28 ± 0.30	0.80 ± 0.02	0.64 ± 0.04
B = 64	7.39 ± 0.85	0.84 ± 0.01	0.71 ± 0.03
B = 100	10.40 ± 1.52	0.86 ± 0.01	0.74 ± 0.02
B = 144	13.60 ± 2.19	0.88 ± 0.02	0.77 ± 0.03

Values are mean $\pm$ SD. BSF is best so far of an EC run; r^2 is the coefficient of determination.

Table 3. Fitness scores achieved by the evolved and trained networks on the three datasets used in subsequent analyses

Number of MSTd neurons	Spiral selectivity dataset	3D translation/rotation dataset	Population encoding dataset
B = 16	0.54 ± 0.06	0.57 ± 0.05	0.56 ± 0.05
B = 36	0.74 ± 0.03	0.75 ± 0.03	0.70 ± 0.03
B = 64	0.81 ± 0.02	0.79 ± 0.01	0.72 ± 0.03
B = 100	0.81 ± 0.05	0.78 ± 0.05	0.71 ± 0.02
B = 144	0.87 ± 0.02	0.82 ± 0.01	0.76 ± 0.02

See the following sections below: Spiral selectivity dataset, 3D translation/rotation dataset, and Population encoding dataset. Values are mean $\pm$ SD.

LTP effects than LTD. Such biases agree with the hand-designed STDP learning curves used in computational studies, which demonstrated the effect of these types of curves in balancing synaptic strength and encouraging competitive learning in the synapses (Gerstner et al., 1997; Abbott and Nelson, 2000; Song et al., 2000).

These learning curves also contributed to sparseness in the learned representations. The MT→MSTd input connection was biased toward LTD, which led to a net effect of depression of all synapses and stabilized network activity by reducing the firing rate of network neurons. With a reduced overall activity in the network, correlated spikes between a particular pair of presynaptic and postsynaptic neurons are more easily differentiated, which triggered LTP consistently and strengthened the synapse. As a result, after learning through STDP, some synapses were selectively driven to be close to the maximum weight values, and other synapses close to the minimum values. On the other hand, the interconnection between the MSTd group and the inhibitory neuron group were biased toward LTP, which tended to strengthen the synapses. Stronger synapses in the feedforward MSTd→Inh connection led to a higher level of activity in the inhibitory neuron




Figure 6. Evolved STDP parameters. **A**, Evolved STDP curves visualized by plotting the changes of synaptic weight (Δw) against the time difference between the presynaptic (t_{pre}) and postsynaptic (t_{post}) spikes. The blue curves correspond to the LTP component of STDP (Eq. 7), and the yellow curves correspond to the LTD component (Eq. 8). Solid lines denote the mean values calculated from all evolved network instances (10 network instances for $B = 16$, 5 network instances for $B = \{36, 64, 100, 144\}$), and the shaded areas denote the SD. **B**, Area over the LTD or LTP component of the STDP curves.




Figure 7. Sparseness measurements of different configurations of the SNN model. Population sparseness measured how many neurons were activated by any stimulus, and lifetime sparseness measured how many stimuli any given neuron responds to. The level of sparseness increases with the increased number of MSTd neurons in the network.

group, which when combined with the strong synapses in the feedback Inh–MSTd connection, provided a powerful inhibition on the MSTd neurons. The STDP curves in all three connections encouraged a sparser activity in the MSTd neuron group. Together, it suggests that the evolutionary process in accurately reconstructing flow patterns discovered STDP curves, similar to those observed experimentally and proposed theoretically, that stabilize activity and produce sparse, reduced representations.

Sparse representations in simulated MSTd neurons

Sparse coding describes a coding scheme in which a stimulus could be represented by a small set of neurons, and each stimulus activates different subsets of neurons (Foldiak and Endres, 2008). This is in contrast to dense codes, where every stimulus activates all neurons, and local codes, where each neuron is highly selective to a particular stimulus (Arbib, 1995). Sparse codes are a favorable compromise between the dense and local codes, which reduce the overall neural activity required to represent the stimulus space (Vinje and Gallant, 2000; Beyeler et al., 2019).

Beyeler et al. (2016) showed that the MSTd units in the sparse decomposition model formed a sparse population code of the input and suggested that sparse codes may be used by MSTd to learn compact and multifaceted representations of the input. In the current study, we also investigated whether the simulated MSTd neurons in the SNN model formed sparse representations of the input activity patterns. We quantified the level of sparseness of the simulated MSTd neurons using the sparseness metrics defined in Vinje and Gallant (2000). The lifetime sparseness metric measured how selective a neuron was to different stimuli, and the population sparseness metric measured how many neurons were activated by any

given stimulus. The sparseness values range from zero to one, with one indicating very selective responses and maximum sparseness, and zero indicating a dense code. We computed both metrics for the MSTd activity over the validation trials after the models were fully evolved and trained. The same procedure was applied to all configurations of the SNN model ($B = \{16, 36, 64, 100, 144\}$).

As shown in Figure 7, network models with a larger MSTd group tended to have sparser representations of the input stimuli both in lifetime sparseness and population sparseness. Sparseness values in both metrics ranged from ~ 0.5 for the $B = 16$ configuration to ~ 0.8 for the $B = 144$ configuration. These results indicated that similar to the sparse decomposition model (Beyeler et al., 2016), the simulated MSTd neurons in all configurations of the SNN model learned a sparse population code of the input activity




Figure 8. MSTd response to spiral flow fields. **A, C, E**, Plots show the distributions of tuning with arrows spanning the spiral motion space. Each arrow represents one neuron. **B, D, F**, Histograms show the distributions of tuning by showing the percentage of simulated neurons tuned to each type of motion. **A, B**, In a population of 57 neurons recorded from the MSTd area, the tuning of MSTd neurons spanned the entire spiral space, with a large proportion of neurons tuned to expanding motions. Reprinted with permission from Graziano et al. (1994). **C–F**, Analyses of the entire population of simulated MSTd neurons obtained from five separately evolved and trained network instances with the $B = 64$ configuration. **C, D**, In a prescreened for expansion subpopulation of the simulated MSTd neurons, a large proportion of simulated neurons were tuned to expanding motions. **E, F**, In the entire simulated MSTd neuron population, 278 simulated MSTd neurons had significant tuning to spiral stimuli. The preferred spiral directions distributed evenly.

patterns. The results also indicated that the sparseness values may saturate at the $B = 144$ configuration and not improve for networks with a larger number of MSTd neurons as the fitness scores increased minimally from the $B = 100$ configuration to the $B = 144$ configuration.

MSTd-like receptive fields and spiral selectivity

We visualized the receptive fields of simulated MSTd neurons by applying population decoding on the connection weight values in the networks evolved with $B = \{16, 36, 64, 100, 144\}$ MSTd neurons (Georgopoulos et al., 1982). The 2D motion selectivity of many simulated MSTd neurons may be predicted from the visualized weight patterns. Figure 3 shows the visualized weight patterns of an evolved and trained network individual with the $B = 64$ configuration as an example. Simulated MSTd neurons in the network had large receptive fields and preferred a mixture of translational and rotational flow. The receptive fields depicted a wide range of optic flow patterns, including translation, rotation, expansion, contraction, and intermediate of these basic flow patterns. Networks with other configurations ($B = \{16, 36, 100, 144\}$) showed similar weight patterns.

To quantify the selectivity of the simulated MSTd neurons to a continuum of spiral motions, we followed the procedure described in Graziano et al. (1994), where the tuning of MSTd neurons was determined by a Gaussian curve fit (Fig. 8A,B). We generated a set of eight basic stimuli as described in Graziano et al. (1994), which contained expansion, contraction, clockwise rotation, counterclockwise rotation, and the four cardinal directions of translation (up, down, right, and left). We presented these stimuli to the fully evolved and trained networks, and obtained the response of the simulated MSTd neurons to each stimulus. We first used a criterion that the neuronal response should be $>10\%$ of the maximum response in the population for any of the stimuli to filter out simulated neurons that were not responsive

to the basic stimuli. We then fit the tuning curve of each neuron with a Gaussian function. The peak of the Gaussian function corresponded to the preferred spiral motion of the neuron. These analyses were performed on all configurations of the network with $B = \{16, 36, 64, 100, 144\}$ MSTd neurons. In Figure 8 we present the results from the networks with $B = 64$ MSTd neurons (results of the other four configurations are shown in Fig. 9). From five separately evolved and trained network individuals with the $B = 64$ configuration, we obtained the response of 320 MSTd neurons, 317 of which responded significantly to at least one stimulus. We plotted the tuning curve of each MSTd neuron as a function of their response to the spiral motion stimuli and fit the tuning curve with a Gaussian function. We measured the goodness of fit of the Gaussian function by computing the correlation r between the Gaussian function and the tuning curve. We also recorded the width σ of the Gaussian curves and compared them to the real neurons. A large proportion of simulated MSTd neurons had smooth tuning curves and good Gaussian fits; the mean r score for the Gaussian fit was 0.97 ± 0.04 SD, and 278 of 317 neurons (87.70%) had r scores >0.9 . The average tuning width for these neurons was 74.09° , with an SE of 4.44° . These measurements were comparable to the tuning widths of MST neurons in Graziano et al. (1994; average width $\sigma = 61^\circ$, SE = 5.9°).

The modeled MSTd neurons showed similar responses to spiral flow fields as those observed experimentally. We estimated the preferred spiral motion of the neuron by finding the peak of the Gaussian curve. Figure 8, E and F, show the preferred spiral motion of the simulated MSTd population. In Figure 8E, each arrow represents the preferred motion of one simulated MSTd neuron. The population of simulated MSTd neurons showed a wide range of preferred spiral motion patterns. We categorized these neurons by their tuning to the eight basic motion types (expansion, contraction, both rotations, and four intermediate directions of spiral motion). Figure 8F shows the distribution of




Figure 9. MSTd response to spiral flow fields ($B = \{16, 36, 100, 144\}$). **A, C, E, G**, Distribution of spiral tuning of the prescreened population. Similar to the $B = 64$ configuration as shown in Figure 8D, a large proportion of simulated neurons were tuned to expanding motions. **B, D, F, H**, Distribution of spiral tuning of the entire population of simulated MSTd neurons obtained in 10 separately evolved and trained network instances with the $B = 16$ configuration, and five network instances with the $B = \{36, 10, 144\}$. Similar to the $B = 64$ configuration as shown in Figure 8F, the preferred spiral directions distributed evenly.

simulated MSTd neurons tuned to different spiral motions. Unlike the MST neurons reported in Graziano et al. (1994), where the population showed a bias to expanding stimulus (Fig. 8A,B), the simulated MSTd population in the SNN model had their preferred spiral motion distributed more evenly for each motion type.

Previous studies (Gu et al., 2006; Takahashi et al., 2007; Gu et al., 2010; Beyeler et al., 2016) suggested that the predominance of expansion-tuned MST neurons may be because of selection bias in the neuron screening process. Earlier neurophysiological studies of MSTd tended to use expansion stimuli to locate visually responsive neurons in MSTd and therefore may have led to a biased population of neurons more likely to be tuned to expanding optic flows. We applied the same selection bias and sampled a subpopulation of simulated MSTd neurons. In the prescreening process, neurons that had a strong response to the expansion stimulus were more likely to be included in this subpopulation. With this simulation of selection bias, we selected 39 neurons from the pool of simulated MSTd neurons in networks with the $B = 64$ configuration. After filtering out neurons that were not significantly tuned to any of the basic motion stimuli, we had 36 simulated neurons for the subsequent analysis. In the subpopulation that was prescreened for expansion, 47% were expansion-tuned neurons, 50% spiral-tuned neurons (42% expanding spiral and 8% contracting spiral), 3% rotation-

tuned neurons, and 0% contraction-tuned neurons. The distribution of preferred spiral motion in Graziano et al. (1994) was expansion (40%), spiral [37% (expanding spiral 28%, contracting spiral 9%)], rotation (16%), and contraction (7%). Similar to the data in Graziano et al. (1994), a large proportion of simulated MSTd neurons in the SNN model were tuned to expansion (Fig. 8C,D). The prescreened subpopulation had tuning curves with good Gaussian fits, with a mean r score of 0.98 ± 0.05 SD. The average tuning width was 64.90° , with an SE of 10.82° , which was also comparable to the MST data in Graziano et al. (1994).

3D translation and rotation heading selectivity

Neurophysiological studies characterized the 3D translation and rotation heading tuning of MSTd neurons using visual stimuli that simulated the observer's movement through a 3D cloud of stars (Gu et al., 2006; Takahashi et al., 2007). These stimuli were generated following a translation protocol, which consisted of stimuli that depicted straight translational movements along 26 directions corresponding to all combinations of azimuth and elevation angles in increments of 45° , and a rotation protocol, which consisted of stimuli that depicted rotational movements along the same 26 directions, with the direction corresponding to the axis of rotation according to the right-hand rule. These 26 directions were described by 26 movement vectors spaced 45° apart in both azimuth and elevation on the sphere.

We tested whether the simulated MSTd neurons in the SNN models had similar translation and rotation tuning as experimental observations. We followed the same protocols as the neurophysiological studies and generated visual motion stimuli to test the SNN models (Gu et al., 2006; Takahashi et al., 2007). The translational stimuli were generated with a peak velocity $v_z = 0.3$ m/s, and the rotational stimuli were generated with a peak velocity $\omega_z = 20^\circ$ /s. We recorded the activity of the simulated MSTd neurons during the stimulus presentation and obtained 3D tuning curves for each MSTd neuron by plotting the neuron activity as a function of the azimuth and elevation of the stimuli. As the heading directions were sampled around the sphere, the heading tuning function was spherical. We transformed the data using the Lambert cylindrical equal-area projection to plot the 3D tuning curves on Cartesian axes (Snyder, 1987). Figure 10 shows a comparison of 3D translation and rotation tuning in an example macaque MSTd neuron and in an example simulated MSTd neuron from the SNN model ($B = 64$). The tuning curves of the neurons were illustrated as contour maps, where the x -axis represented the azimuth angle, the y -axis represented the sinusoidally transformed version of the elevation angle, and the mean firing rate of the neuron to each stimulus was plotted as a function of azimuth and elevation angles. Similar to MSTd neurons (Takahashi et al., 2007), our simulated MSTd neurons showed broad tuning curves for translational and rotational heading directions. We computed the vector sum of the response of the unit to all 26 stimuli to determine the

preferred translation and rotation directions for the simulated neurons. The experimentally observed MST neuron shown in Figure 10, *A* and *B*, was tuned to 291° azimuth and −18° elevation under the rotation condition and was tuned to 190° azimuth and −50° elevation under the translation condition (Takahashi et al., 2007). The simulated MSTd neuron also showed strong spatial tuning under both the rotation and translation condition. This particular simulated MSTd neuron shown in Figure 10, *C* and *D*, was tuned to 35° azimuth and −27° elevation for rotational heading, and 216° azimuth and −65° elevation for translational heading.

We examined the distribution of preferred directions of the simulated MSTd neurons. Figure 11 shows the distribution of translation and rotation direction preferences of MST neurons (Takahashi et al., 2007) and of all 320 simulated MSTd neurons from five individually evolved and trained SNN models with the $B = 64$ configuration (results of the other four configurations are shown in Fig. 12). Each data point in these plots represents the 3D preferred direction of a single neuron. Note that in these plots, 360° in the azimuth angle was equivalent to 0° in the azimuth angle. Similar to the distribution reported in Takahashi et al. (2007), the simulated MSTd neurons also showed significantly nonuniform distributions in azimuth and elevation for both translation and rotation heading preferences (uniformity test, $p < 0.05$). Rotational preferences showed a bimodal distribution in azimuth, clustering at 0 and 180°, corresponding to pitch-up and pitch-down rotations, respectively. Translational preferences had bimodal distributions in both azimuth and elevation, clustering at 0 and 180° in azimuth, corresponding to leftward and rightward movement directions, and $\pm 90^\circ$ in elevation, corresponding to upward and downward movement directions. We categorized the simulated neurons with their preferred motion type by finding neurons with preferred directions falling within $\pm 30^\circ$ of each of the cardinal axes of the motion type. Similar to the recordings in Takahashi et al. (2007), the majority of the simulated MSTd neurons were tuned to yaw ($\pm 90^\circ$ in elevation) or pitch (0° and 180° in azimuth) directions for rotational headings, and a smaller percentage of simulated MSTd neurons were tuned to roll motions (90° and 270° in azimuth; Table 4). Under the translation condition, the majority of the simulated MSTd neurons were tuned to lateral (0° and 180° in azimuth) or vertical ($\pm 90^\circ$ in elevation) translational motions (Table 5). To quantify the strength of heading tuning of single MSTd




Figure 10. 3D translation and rotation tuning of example MST neurons. *A, B*, Translation and rotation tuning of an example MST neuron recorded under the visual rotation and translation conditions. Reprinted with permission from Takahashi et al. (2007). *C, D*, Translation and rotation tuning of an example MSTd neuron from a fully evolved and trained SNN model with the $B = 64$ configuration. Each contour map shows the Lambert cylindrical equal-area projection of the original data (Snyder, 1987).




Figure 11. Distribution of 3D translation and rotation tuning of MSTd neurons. *A–D*, Each data point in the scatter plot represents the preferred azimuth and elevation angles of one neuron. The histograms along the sides of the scatter plots show the marginal distributions. The 2D projections of unit-length 3D preferred direction vectors are shown in the radial plots, including the front, side, and top view. Each radial line in these plots represents one neuron. *A, B*, MST neurons recorded under the visual rotation and translation conditions. Reprinted with permission from Takahashi et al. (2007). *C, D*, Simulated MSTd neurons from five separately evolved and trained SNN model with the $B = 64$ configuration.

neurons, we computed the heading tuning index (HTI) value defined in Gu et al. (2006). An HTI ranges from zero to one, with larger values indicating stronger heading tuning. Gu et al. (2006) reported that the average HTI value for their sampled MSTd neurons was 0.48 ± 0.16 . The simulated MSTd neurons in the networks with the $B = 64$ configuration showed a




Figure 12. Distribution of 3D translation and rotation tuning of simulated MSTd neurons ($B = \{16, 36, 10, 144\}$). **A, C, E, G**, The rotational heading preferences of all simulated neurons in each network configuration. **B, D, F, H**, The translational heading preferences. Each data point in the scatter plot represents the preferred azimuth and elevation angles of one simulated neuron. The histograms along the sides of the scatter plots show the marginal distributions.

Table 4. Number of neurons tuned to each motion type under the visual rotation condition

	Yaw	Pitch	Roll
Takahashi et al., 2007	36/127 (28%)	27/127 (21%)	1/127 (1%)
SNN model ($B = 64$)	67/320 (20.94%)	95/320 (29.69%)	25/320 (7.81%)

Table 5. Number of neurons tuned to each motion type under the visual translation condition

	Lateral	Fore-aft	Vertical
Takahashi et al. (2007)	57/307 (19%)	20/307 (7%)	76/307 (25%)
SNN model ($B = 64$)	75/320 (23.44%)	29/320 (9.06%)	73/320 (22.81%)

stronger tuning, averaging 0.66 ± 0.14 and 0.68 ± 0.12 for rotation and translation, respectively.

What might be the source of the nonuniform distribution of heading preferences, especially the over-representation of lateral motion (azimuth of 0 and 180°), and how might these characteristics of the MSTd neuron population contribute to self-motion perception? Gu et al. (2010) suggested that an abundance of MSTd neurons preferring lateral heading directions may

account for the observation that primates are most accurate at judging small variations in heading references directly in front of the animal and less precise for heading around an eccentric reference. Most of their recorded MSTd neurons have broad, cosine-like tuning curves, and the preference to lateral headings caused the peak discriminability of these neurons to lie near straight ahead.

In the evolved and trained models, the simulated MSTd neurons also had broad, cosine-like tuning curves (Fig. 10). We tested whether these simulated MSTd neuron populations also had maximal discriminability for heading directions near the straight ahead. We presented motion patterns that simulated eight directions of translation in the horizontal plane ($0, \pm 45, \pm 90, \pm 135, 180^\circ$, relative to straight ahead) to the models and recorded the activity of the simulated MSTd neurons. Following the experimental protocol of Gu et al. (2010), we computed the peak discriminability of individual simulated MSTd neurons and estimated the precision of heading discrimination by computing population Fisher information from the tuning curves of the simulated MSTd neurons. To obtain a smooth tuning curve of the simulated MSTd neurons, we fitted a spline function with 0.01° resolution to the coarsely sampled data (45° spacing). We then calculated the spatial derivative of the spline fit and obtained the tuning curve slope. Peak discriminability was achieved at the steepest slope of the tuning curves. For each heading direction, we computed the population Fisher information based on the derivative of the tuning curve for each simulated neuron and the variance of the response of each simulated neuron to each heading direction, as described in Gu et al. (2010). Larger values of Fisher information indicate higher precision in heading discrimination.

To make a direct comparison with the neurophysiological data, we adopted the coordinate system used in Gu et al. (2010), where 0° azimuth corresponded to the straight-ahead direction. Figure 13 compares the distribution of peak discriminability of recorded MSTd neurons and simulated MSTd neurons in the models with different network configurations. Consistent with the neurophysiologically recorded data, most simulated MSTd neurons had their peak discriminability clustered around forward (0°) and backward (180°) headings. Figure 14 shows population Fisher information computed from the heading tuning curves for recorded MSTd neurons and simulated MSTd neurons in the network models. Consistent with the data reported in Gu et al. (2010), there was a clear dependence of Fisher information on reference heading for the simulated MSTd neuron populations, with a peak for forward headings (0°) and a minimum for lateral headings ($\pm 90^\circ$). The magnitudes of Fisher information differed between the recorded and simulated MSTd neurons, and also across different network configurations, because of the different sample sizes and differences in signal-to-noise ratio. The simulated MSTd neurons diverged from the recorded neurons in that the former population showed higher Fisher information for backward headings (180°).

Population encoding of heading direction

MSTd neurons were suggested to encode perceptual variables such as the focus of expansion (FOE) in optic flow, eye position, and pursuit to support self-motion. Perceptual variables could be decoded with an average accuracy of 2 to 3° based on single-trial response of a population of 144 MSTd neurons (Hamed et al., 2003).

In the current study we focus on investigating whether the simulated MSTd neuron population in the SNN model also encoded FOE robustly. We generated a dataset with 10,000 optic flow fields that simulated various translational motions as described in a physiological study (Hamed et al., 2003). The heading condition contained flow fields that depicted randomly selected heading directions, with an azimuth between 60 and 120° and elevation between −30 and 30°. These flow fields simulated linear movements toward a back plane located at various distances in meters, $d = \{2, 4, 8, 16, 32\}$, from the observer, with the observer moving at a velocity in m/s, $v_z = \{0.5, 0.75, 1, 1.25, 1.5\}$, without any rotational movements. The goal was to decode the MSTd population activity to predict the location of FOE, which in this case, was the same as predicting the heading direction of the observer. With the fully evolved and trained network models configured with different number of MSTd neurons, ($B = \{16, 36, 64, 100, 144\}$), we obtained MSTd group activity responding to the heading stimuli. For each configuration, we randomly selected 144 units from the pool of simulated MSTd neurons collected from all separately evolved and trained networks and used optimal linear estimators with a 10-fold cross-validation procedure to find the linear weights that could decode heading from the activity of this subpopulation of simulated MSTd neurons. Table 6 shows the prediction error of FOE of the MSTd neurons and simulated MSTd neurons from the network instances with the $B = 64$ configuration. FOE as a key perceptual variable could be decoded with a slightly higher degree of precision than the experimentally observed MSTd neurons (Hamed et al., 2003). As shown in Figure 15, FOE could be decoded with a similar level of accuracy from simulated MSTd neurons in all five configurations of the SNN model (Wilcoxon's rank sum test, $p = 1$ for all pairwise comparisons).

Discussion

By evolving STDP-H parameters of SNNs, we demonstrated that a model of MSTd was able to account for many experimentally




Figure 13. A–F, Distribution of the direction of maximal discriminability for neurons recorded in macaque MSTd (A) and simulated MSTd neurons in the SNN models (B–F). Reprinted with permission from Gu et al. (2010). Both the recorded and the simulated neuron populations showed a bimodal distribution with peaks around the forward (0°) and backward (180°) headings.





Figure 14. A, B, Population Fisher information computed from neurons recorded in macaque MSTd (A) and from simulated MSTd neurons in the SNN models (B). Reprinted with permission from Gu et al. (2010). The error bands in A show 95% confidence intervals derived from a bootstrap procedure.

response properties observed in monkeys, (1) tuning to spiral motions, (2) 3D translation and rotation selectivity, and (3) encoding of perceptual variables such as heading. The simulated MSTd neurons learned compressed and efficient representations of input activity patterns with a population that was an order of

Table 6. Error in predicting FOE using the activity of a population of 144 MSTd neurons

	FOE (x, y)
Hamed et al. (2003)	$(3.62 \pm 6.78^\circ, 3.87 \pm 4.96^\circ)$
SNN model ($B = 64$)	$(0.98 \pm 0.77^\circ, 1.13 \pm 0.87^\circ)$

**Figure 15.** Error in predicting FOE (heading) using the activity of a population of 144 MSTd neurons from each configuration of the network mode I. Simulated MSTd neurons from different configurations of the network exhibited similar levels of accuracy in these two tasks. The interquartile range box represents the middle 50% of the values. The line within the box indicates the median. The red data points mark outliers that are 1.5 times bigger than the interquartile range, and the whiskers extend to the most extreme data points that are not considered outliers. deg, Degrees.

magnitude smaller than its MT neuron input group. These representations, which accurately reconstructed input stimuli, efficiently encoded complex self-motion that could be used by downstream neural areas. The present results provide a linkage from prior machine learning studies that suggested non-negative sparse coding can produce similarities to responses observed in the cortex (Hoyer, 2003; Beyeler et al., 2016, 2019) to neurobiologically plausible computations. Specifically, our results suggest that neurobiological plasticity, like STDP-H, may be contributing to dimensionality reduction and sparse coding observed in the brain.

Response properties of simulated MSTd neurons and implications for downstream processing

Simulated MSTd neurons in our SNN model exhibited a variety of response properties observed in neurophysiological studies. First, simulated MSTd neurons had Gaussian tuning spanning the entire spiral space, and many of these neurons showed a strong representation of expanding motions (Graziano et al., 1994). Second, similar to the MSTd neurons recorded in visual translational and rotational heading experiments (Gu et al., 2006; Takahashi et al., 2007), the simulated MSTd neurons showed strong tuning to 3D translation and rotation heading directions, with more neurons preferring lateral and vertical translational movements and pitch and yaw rotational movements. The simulated MSTd neuron had maximal heading discriminability clustered around forward (0°) and backward (180°) headings. Such properties may arise from the SNN models learning to reconstruct the input optic flow dataset, which sampled headings near the center of the visual field for expanding and contracting motion patterns. It's interesting to note that these biases in heading preferences and heading discriminability were also observed in macaque MSTd neurons (Gu et al., 2006; Takahashi et al., 2007; Gu et al., 2010), which may suggest that in natural scenes,

optic flow patterns with headings near the center of the visual field are more commonly experienced during forward and backward movements. In the training and validation dataset, we sampled forward and backward movements with the same probability, whereas backward locomotion is an uncommon activity in natural movements (Perrone, 1986; Perrone and Stone (1994)). This discrepancy may explain the higher precision in discriminating backward headings with the simulated MSTd neurons and the neurons recorded in the macaque MSTd. In future studies, by varying the distribution of the training data and observing how this affects the tuning curves of the MSTd neurons, we will have a better understanding of how observed response properties of MSTd neurons may be an adaptation for efficient self-motion perception. Finally, FOE as one of the key perceptual variables could be decoded with high accuracy from our simulated MSTd neurons (Hamed et al., 2003). These results suggest that the simulated MSTd neurons encoded information essential for perceiving and guiding self-motion.

These properties emerged as the model evolved unsupervised learning parameters (STDP-H) to reconstruct optic flow patterns. This led to the efficient coding of perceptual features that may be decoded by downstream neurons for motor behaviors. For example, our SNN model showed a continuum of tuning in the spiral space that has been suggested to simultaneously encode the curvature of a path and heading (Layton and Browning, 2014). This allows MSTd to perform trajectory estimation and guide self-motion along complex trajectories. Heading tuning that spanned the entire translational and rotational movement space allows for accurate heading direction perception. The sparse, reduced representations of the global motion pattern allows for efficient encoding of key locomotion parameters including the heading direction. These observations are consistent with the hypothesis that MSTd transforms stimulus information into representations more directly linked to perception and action, supported by the connections between MSTd to many cortical and subcortical areas (Boussaoud et al., 1990; Felleman and Van Essen, 1991; Wild and Treue, 2021).

A limitation of the current study is that heading perception with the simulated MSTd neurons was only studied under the locomotion condition without eye movements. Previous studies have shown that although eye rotations during smooth pursuit distort the optic flow patterns on the retina, humans and nonhuman primates are able to perceive their heading direction correctly (Royden et al., 1992; Britten and Van Wezel, 2002). It will be interesting to investigate whether the simulated MSTd neurons could predict heading in the presence of eye movements and whether retinal mechanisms are sufficient to discount distortions to the optic flow fields, as suggested in previous computational and neurophysiological studies (Royden, 1997; Perrone and Stone, 1998; Beyeler et al., 2016; Manning and Britten, 2019). The stimulus paradigm used in this study was not well equipped for such an analysis, as the stimuli did not provide meaningful distinction between head rotation and eye rotation.

Relating STDP-H to non-negative sparse coding in the brain

It has been suggested that STDP-H may lead to sparse, reduced representations in the brain (Carlson et al., 2013; Beyeler et al., 2019). Indeed, our model showed the ability to accurately reconstruct the input stimuli in a neuron population by activating a small subset of neurons (population sparseness), which were active on only a subset of stimuli (lifetime sparseness). In our SNN model, the high-dimensional visual motion input, represented by the population activity of 9000 MT neurons, projected

to a small group of MSTd neurons ($B = \{16, 36, 64, 100, 144\}$). The simulated MSTd neurons were pressured to reduce the dimensions of the input and find efficient representations that retained relevant features in the input space. In addition, sparse coding was encouraged by the synaptic competition induced by STDP at the synapse level and by feedback inhibition (Finelli et al., 2008; Chistiakova et al., 2015). Also, the STDP curves in our SNN model that modulated the MT→MSTd connections evolved to have stronger LTD than LTP, which led to a pruning effect on the connection weights by depressing uncorrelated spiking activity that led to sparser activity (Feldman, 2000, 2012). Finally, homeostatic synaptic scaling preserved the relative weight differences between synapses on the same postsynaptic neuron and stabilized STDP learning (Carlson et al., 2013).

The biological brain operates under tight metabolic constraints and is often pressured to find strategies for accurate information encoding, which also allows for efficient transfer while minimizing the energy costs (Krichmar et al., 2019). Sparse coding is one such strategy that reduces neural activity without sacrificing performance (Graham and Field, 2007). The sparse coding strategy was demonstrated to be effective in encoding sensory input and was suggested to be ubiquitous throughout the brain (Olshausen and Field, 2004). Dimensionality reduction is a strategy to transform high-dimensional and complex information into a compressed form that is easier to transfer and read out at subsequent levels of processing. The sparse coding scheme implemented in Olshausen and Field (2004) used a greater number of neurons than the dimensions of the input space to represent the input, suggesting dimensionality expansion. This may seem contradictory to the dimensionality reduction strategy. However, Hoyer (2002) combined the two strategies to form the NSC method, which decomposes multivariate data into non-negative sparse components. Hoyer (2004) further showed that explicitly incorporating sparseness constraints is important for learning non-negative representations that match the underlying elements of the data. NSC has been demonstrated to reproduce receptive fields of neurons in the primary visual cortex by decomposing the activities of ON and OFF channels in response to natural image patches (Hoyer, 2003). NMF with sparsity constraints implements a form of NSC, which when applied to MT-like input, resulted in sparse, parts-based representations of the optic flow resembling the receptive fields of experimentally observed MSTd neurons (Beyeler et al., 2016). Outside the visual cortex, NSC was suggested to be used by many other sensory areas in the brain to encode external stimuli (Beyeler et al., 2019).

The scheme of dimensionality reduction and sparse coding induced by STDP-H in our SNN model is consistent with NSC. NSC describes the input data as a superposition of a set of sparsely activated basis functions and requires the basis functions as well as the activation values to be non-negative (Hoyer, 2002). The input optic flow patterns resulting from locomotion, despite their high dimensionality encoded in the population activity of idealized MT neurons, had some statistical regularities. For example, Beyeler et al. (2016) showed that these optic flow patterns could be reconstructed with a linear combination of a set of basis flow fields. These basis flow fields captured the key features within the data and represented the data in a lower dimensional space. In the current study, we demonstrated that STDP-H allowed the spiking neurons to detect and extract these statistical regularities among the spike trains and developed representations of the input optic flow fields similar to those obtained through NSC. Together with the sparse decomposition

model proposed in Beyeler et al. (2016), the SNN model presented in the present study challenges the prevailing views about the MST area in macaque cortex, by showing that prominent neuronal responses traditionally attributed to specialized self-motion templates might instead be a by-product of neurons performing dimensionality reduction on their inputs. Our work agrees with the hypothesis that STDP-H may be a biologically plausible neural mechanism that implements NSC in the brain (Carlson et al., 2013), and to our knowledge, may be the first computational study to demonstrate the link between the STDP-H synaptic learning rules and the NSC neural encoding scheme.

Model alternatives

A number of computational models have been proposed to reproduce the motion selectivity of MSTd neurons. Examples of these models include a model trained with unsupervised Hebbian learning (Zhang et al., 1993), the multicausal model that took the form of an autoencoder (Zemel and Sejnowski, 1998), and a two-layer back-propagation network that was trained in a supervised manner to approximate tuning curves of MST neurons (Beardsley and Vaina, 1998). Alternatively, Perrone and Stone (1994) and Perrone and Stone (1998) proposed the template model, which constructed MST units that each encoded a specific combination of heading and rotation through combining the input from specifically arranged MT-like units. By modeling the responses of the MSTd neurons to a set of novel continuous optic flow stimuli, Mineault et al. (2012) showed that nonlinear integration, especially compressive integration of MT input led to the complex selectivity observed in MST neurons. Layton and Browning (2014) defined MSTd hypercolumns that had receptive fields centered at different locations of the visual field and possessed selectivities across a spiral space of motions. Mineault et al. (2021) showed that a deep neural network trained to predict self-motion parameters could account for the selectivity of neurons in the dorsal visual stream.

Here, we used a novel modeling approach of evolving synaptic learning rules in an SNN, which posed minimal assumptions on the selection of model parameters. Selectivity and receptive fields of the simulated MSTd neurons in our models were emergent properties without hand-crafted templates, which makes our modeled neurons flexible and generalizing to other complex tasks. This approach is generalizable and has been used to replicate neural recordings in the retrosplenial cortex (Rounds et al., 2018) and the medial temporal lobe (Chen et al., 2021).

In summary, we showed that MSTd-like response properties emerge from an SNN model that was trained with STDP-H and evolutionary computation, with an objective of reconstructing the input optic flow fields. Simulated MSTd neurons in this model formed a sparse population code of the input activity pattern and encoded perceptual variables important for self-motion perception. This model demonstrated that STDP-H reduced the dimensions of input stimuli and allowed spiking neurons to learn sparse and efficient representations in a similar way that NSC performs decomposition of complex input. It suggests a biologically plausible method for how these and other representations in the brain emerge.

References

- Abbott LF, Nelson SB (2000) Synaptic plasticity: taming the beast. *Nat Neurosci* 3:1178–1183.
- Albright TD (1984) Direction and orientation selectivity of neurons in visual area MT of the macaque. *J Neurophysiol* 52:1106–1130.

- Altintas I, Marcus K, Nealey I, Sellars SL, Graham J, Mishin D, Polizzi J, Craw D, DeFanti T, Smarr L (2019) Workflow-driven distributed machine learning in CHASE-CI: a cognitive hardware and software ecosystem community infrastructure. arXiv:1903.06802.
- Arbib MA (1995) The handbook of brain theory and neural networks. Cambridge, MA: MIT.
- Balaji A, Adiraju P, Kashyap HJ, Das A, Krichmar JL, Dutt ND, Catthoor F. (2020) PyCARL: a PyNN interface for hardware-software co-simulation of spiking neural network. Paper presented at the International Joint Conference on Neural Networks. Glasgow, United Kingdom, July.
- Beardsley SA, Vaina LM (1998) Computational modelling of optic flow selectivity in MSTd neurons. *Network* 9:467–493.
- Beyeler M, Dutt N, Krichmar JL (2016) 3D visual response properties of MSTd emerge from an efficient, sparse population code. *J Neurosci* 36:8399–8415.
- Beyeler M, Rounds EL, Carlson KD, Dutt N, Krichmar JL (2019) Neural correlates of sparse coding and dimensionality reduction. *Plos Comput Biol* 15:e1006908.
- Boussaoud D, Ungerleider LG, Desimone R (1990) Pathways for motion analysis: cortical connections of the medial superior temporal and fundus of the superior temporal visual areas in the macaque. *J Comp Neurol* 296:462–495.
- Britten KH (2008) Mechanisms of self-motion perception. *Annu Rev Neurosci* 31:389–410.
- Britten KH, van Wezel RJA (1998) Electrical microstimulation of cortical area MST biases heading perception in monkeys. *Nat Neurosci* 1:59–63.
- Britten KH, Van Wezel RJ (2002) Area MST and heading perception in macaque monkeys. *Cereb Cortex* 12:692–701.
- Carlson KD, Richert M, Dutt N, Krichmar JL (2013) Biologically plausible models of homeostasis and STDP: stability and learning in spiking neural networks. Paper presented at the 2013 International Joint Conference on Neural Networks, Dallas, August.
- Chen K, Johnson A, Scott E, Zou X, De Jong KA, Nitz D, Krichmar J (2021) Differential spatial representations in hippocampal CA1 and subiculum emerge in evolved spiking neural networks. Paper presented at the International Joint Conference on Neural Networks, Shenzhen, China, July.
- Chistiakova M, Bannon NM, Chen J-Y, Bazhenov M, Volgushev M (2015) Homeostatic role of heterosynaptic plasticity: models and experiments. *Front Comput Neurosci* 9:89.
- Chou TS, Kashyap HJ, Xing J, Listopad S, Rounds EL, Beyeler M, Dutt N, Krichmar JL (2018) CARLsim 4: an open source library for large scale, biologically detailed spiking neural network simulation using heterogeneous clusters. Paper presented at the International Joint Conference on Neural Networks, Rio de Janeiro, July.
- Cullen KE (2019) Vestibular processing during natural self-motion: implications for perception and action. *Nat Rev Neurosci* 20:346–363.
- De Jong KA (2006) Evolutionary computation: a unified approach. Cambridge, MA: MIT.
- Feldman DE (2000) Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron* 27:45–56.
- Feldman DE (2012) The spike-timing dependence of plasticity. *Neuron* 75:556–571.
- Felleman D, Van Essen D (1991) Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* 1:1–47.
- Finelli LA, Haney S, Bazhenov M, Stopfer M, Sejnowski TJ (2008) Synaptic learning rules and sparse coding in a model sensory system. *PLoS Comput Biol* 4:e1000062.
- Foldiak P, Endres D (2008) Sparse coding. *Scholarpedia* 3:2984.
- Georgopoulos A, Kalaska J, Caminiti R, Massey J (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2:1527–1537.
- Gerstner W, Kempter R, van Hemmen JL, Wagner H (1997) A developmental learning rule for coincidence tuning in the barn owl auditory system. In: *Computational neuroscience: trends in research* (Bower JM, ed), pp 665–669. Boston: Springer.
- Graham D, Field D (2007) Sparse coding in the neocortex. *Evolution of Nervous Systems* 3:181–187.
- Graziano MS, Andersen RA, Snowden RJ (1994) Tuning of MST neurons to spiral motions. *J Neurosci* 14:54–67.
- Gu Y, Watkins PV, Angelaki DE, DeAngelis GC (2006) Behavioral/systems/cognitive visual and nonvisual contributions to three-dimensional heading selectivity in the medial superior temporal area. *J Neurosci* 26:73–85.
- Gu Y, Fetsch CR, Adeyemo B, Deangelis GC, Angelaki DE (2010) Decoding of MSTd population activity accounts for variations in the precision of heading perception. *Neuron* 66:596–609.
- Gu Y, Deangelis GC, Angelaki DE (2012) Causal links between dorsal medial superior temporal area neurons and multisensory heading perception. *J Neurosci* 32:2299–2313.
- Hamed SB, Page W, Duffy C, Pouget A (2003) MSTd neuronal basis functions for the population encoding of heading direction. *J Neurophysiol* 90:549–558.
- Hebb DO (1949) The organization of behavior; a neuropsychological theory. New York: Wiley.
- Hoyer PO (2002) Non-negative sparse coding. Paper presented at the 12th IEEE Workshop on Neural Networks for Signal Processing, Martigny, Switzerland, September.
- Hoyer PO (2003) Modeling receptive fields with non-negative sparse coding. *Neurocomputing* 52–54:547–552.
- Hoyer PO (2004) Non-negative matrix factorization with sparseness constraints. *J Machine Learning Research* 5:1457–1469.
- Izhikevich EM (2004) Which model to use for cortical spiking neurons? *IEEE Trans Neural Netw* 15:1063–1070.
- Krichmar JL, Severa W, Khan MS, Olds JL (2019) Making bread: biomimetic strategies for artificial intelligence now and in the future. *Front Neurosci* 13:666.
- Layton OW, Browning NA (2014) A unified model of heading and path perception in primate MSTd. *PLoS Comput Biol* 10:e1003476.
- Longuet-Higgins HC, Prazdny K (1980) The interpretation of a moving retinal image. *Proc R Soc Lond B Biol Sci* 208:385–397.
- Luke S (2017) ECJ then and now. Paper presented at the Genetic and Evolutionary Computation Conference, Berlin, July.
- Manning TS, Britten KH (2019) Retinal stabilization reveals limited influence of extraretinal signals on heading tuning in the medial superior temporal area. *J Neurosci* 39:8064–8078.
- Mineault P, Bakhtiari S, Richards B, Pack C (2021) Your head is there to move you around: goal-driven models of the primate dorsal pathway. *BiorXiv*. doi:10.1101/2021.07.09.451701.
- Mineault PJ, Khawaja FA, Butts DA, Pack CC (2012) Hierarchical processing of complex motion along the primate dorsal visual pathway. *Proc Natl Acad Sci U S A* 109:5930.
- Nover H (2005) A logarithmic, scale-invariant representation of speed in macaque middle temporal area accounts for speed discrimination performance. *J Neurosci* 25:10049–10060.
- Oja E (1982) Simplified neuron model as a principal component analyzer. *J Math Biol* 15:267–273.
- Oja E (2008) Oja learning rule. *Scholarpedia* 3:3612.
- Olshausen B, Field D (2004) Sparse coding of sensory inputs. *Curr Opin Neurobiol* 14:481–487.
- Perrone JA (1986) Anisotropic responses to motion toward and away from the eye. *Percept Psychophys* 39:1–8.
- Perrone JA, Stone LS (1994) A model of self-motion estimation within primate extrastriate visual cortex. *Vision Res* 34:2917–2938.
- Perrone JA, Stone LS (1998) Emulating the visual receptive-field properties of MST neurons with a template model of heading estimation. *J Neurosci* 18:5958–5975.
- Raiguel S, Van Hulle MM, Xiao DK, Marcar VL, Lagae L, Orban GA (1997) Size and shape of receptive fields in the medial superior temporal area (MST) of the macaque. *NeuroReport* 8:2803–2808.
- Rounds EL, Alexander AS, Nitz DA, Krichmar JL (2018) Conjunctive coding in an evolved spiking model of retrosplenial cortex. *Behav Neurosci* 132:430–452.
- Royden CS (1997) Mathematical analysis of motion-opponent mechanisms used in the determination of heading and depth. *J Opt Soc Am A Opt Image Sci Vis* 14:2128–2143.
- Royden CS, Banks MS, Crowell JA (1992) The perception of heading during eye movements. *Nature* 360:583–585.
- Schmitt C, Baltaretu BR, Crawford JD, Bremmer F (2020) A causal role of area hMST for self-motion perception in humans. *Cereb Cortex Commun* 1:tgaa042.
- Sjöström J, Gerstner W (2010) Spike-timing dependent plasticity. *Scholarpedia* 5:1362.

- Snyder JP (1987) Map projections: a working manual. Washington, DC: USA Government Printing Office ix.
- Song S, Miller KD, Abbott LF (2000) Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci* 3:919–926.
- Takahashi K, Gu Y, May PJ, Newlands SD, DeAngelis GC, Angelaki DE (2007) Multimodal coding of three-dimensional rotation and translation in area MSTd: comparison of visual and vestibular selectivity. *J Neurosci* 27:9742–9756.
- Turrigiano GG, Nelson SB (2004) Homeostatic plasticity in the developing nervous system. *Nat Rev Neurosci* 5:97–107.
- Turrigiano GG, Leslie KR, Desai NS, Rutherford LC, Nelson SB (1998) Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391:892–896.
- Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287:1273–1276.
- Warren WH, Hannon DJ (1988) Direction of self-motion is perceived from optical flow. *Nature* 336:162–163.
- Wild B, Treue S (2021) Primate extrastriate cortical area MST: a gateway between sensation and cognition. *J Neurophysiol* 125:1851–1882.
- Zemel RS, Sejnowski TJ (1998) A model for encoding multiple object motions and self-motion in area MST of primate visual cortex. *J Neurosci* 18:531–547.
- Zhang K, Sereno MI, Sereno ME (1993) Emergence of position-independent detectors of sense of rotation and dilation with hebbian learning: an analysis. *Neural Comput* 5:597–612.