

# UC Irvine

## UC Irvine Previously Published Works

### Title

Efficient spiking neural network model of pattern motion selectivity in visual cortex

### Permalink

<https://escholarship.org/uc/item/5kg2s3wz>

### Journal

Neuroinformatics, 12(3)

### ISSN

1539-2791

### Authors

Beyeler, M  
Richert, M  
Dutt, ND  
[et al.](#)

### Publication Date

2014

### DOI

10.1007/s12021-014-9220-y

Peer reviewed

1 **Efficient spiking neural network model of pattern**  
2 **motion selectivity in visual cortex**

3\*Michael Beyeler<sup>1</sup>, Micah Richert<sup>2,3</sup>, Nikil D. Dutt<sup>1</sup>, and Jeffrey L. Krichmar<sup>1,2</sup>

4<sup>1</sup>Department of Computer Science, University of California Irvine, Irvine, CA,  
5USA

6<sup>2</sup>Department of Cognitive Sciences, University of California Irvine, Irvine, CA,  
7USA

8<sup>3</sup>Brain Corporation, San Diego, CA, USA

9

10

11Keywords:

12pattern motion selectivity, spiking neural network, MT, GPU, real-time,

13CARLsim

14

15

16

17Correspondence:

18Michael Beyeler

19University of California, Irvine

20Department of Computer Science

21Irvine, CA 92697

22mbeyeler@uci.edu

23Office of Jeff Krichmar: (949) 824-5888

## 24**Abstract**

25 Simulating large-scale models of biological motion perception is challenging,  
26 due to the required memory to store the network structure and the  
27 computational power needed to quickly solve the neuronal dynamics. A low-  
28 cost yet high-performance approach to simulating large-scale neural network  
29 models in real-time is to leverage the parallel processing capability of  
30 graphics processing units (GPUs). Based on this approach, we present a two-  
31 stage model of visual area MT that we believe to be the first large-scale  
32 spiking network to demonstrate pattern direction selectivity. In this model,  
33 component-direction-selective (CDS) cells in MT linearly combine inputs from  
34 V1 cells that have spatiotemporal receptive fields according to the motion  
35 energy model of Simoncelli and Heeger. Pattern-direction-selective (PDS)  
36 cells in MT are constructed by pooling over MT CDS cells with a wide range of  
37 preferred directions. Responses of our model neurons are comparable to  
38 electrophysiological results for grating and plaid stimuli as well as speed  
39 tuning. The behavioral response of the network in a motion discrimination  
40 task is in agreement with psychophysical data. Moreover, our  
41 implementation outperforms a previous implementation of the motion  
42 energy model by orders of magnitude in terms of computational speed and  
43 memory usage. The full network, which comprises 153,216 neurons and  
44 approximately 40 million synapses, processes 20 frames per second of a  
45  $40 \times 40$  input video in real-time using a single off-the-shelf GPU. To promote  
46 the use of this algorithm among neuroscientists and computer vision  
47 researchers, the source code for the simulator, the network, and analysis  
48 scripts are publicly available.

## 491. Introduction

50 Visual motion perception is a challenging problem that is critical for  
51 navigating through the environment and tracking objects. Several software  
52 packages are available to the public that deal with the neurobiologically  
53 plausible modeling of motion perception in the mammalian brain, such as  
54 spatiotemporal-energy models like the motion energy model of Simoncelli  
55 and Heeger (1998), or gradient-based models like ViSTARS (Browning et al.  
56 2009a, 2009b). However, in order for these frameworks to become practical  
57 in, for example, neuromorphic or robotics applications, they must be capable  
58 of running large-scale networks in real-time. Moreover, to take advantage of  
59 state-of-the-art neuromorphic hardware, the elements of the algorithms need  
60 to be spiking neurons (Indiveri et al. 2006; Merolla et al. 2007; Vogelstein et  
61 al. 2007; Khan et al. 2008; Srinivasa and Cruz-Albrecht 2012). Developing  
62 such a simulation environment is challenging, due to the required memory to  
63 store the network structure and the computational power needed to quickly  
64 solve the equations describing the neuronal dynamics. A low-cost yet high-  
65 performance approach to simulating large-scale spiking neural networks  
66 (SNNs) in real-time is to leverage the parallel processing capability of  
67 graphics processing units (GPUs) (Nageswaran et al. 2009; Fidjeland and  
68 Shanahan 2010; Yudanov et al. 2010; Richert et al. 2011).

69       Based on this approach, we present a two-stage model of visual area  
70 MT that we believe to be the first large-scale spiking network to demonstrate  
71 pattern direction selectivity. The model combines and extends two previous  
72 incarnations of the motion energy model (Simoncelli and Heeger 1998; Rust  
73 et al. 2006). Broadly speaking, our model integrates the V1 stage of  
74 Simoncelli and Heeger (1998) with the MT stage of Rust et al. (2006) in the  
75 spiking domain. More precisely, our model uses a bank of spatiotemporal  
76 filters (Adelson and Bergen 1985; Simoncelli and Heeger 1998) to model the  
77 receptive fields of directionally selective neurons in V1, which then project to  
78 component-direction-selective (CDS) cells in area MT. However, the local

79 motion estimates coded by the spike patterns of these neurons often vary  
80 drastically from the global pattern motion of a visual stimulus, because the  
81 local motion of a contour is intrinsically ambiguous (“aperture problem”).  
82 Therefore, in order to construct pattern-direction-selective (PDS) cells in MT  
83 that signal the global pattern motion, we implemented three design  
84 principles introduced by Rust et al. (2006): 1) spatial pooling over V1 or MT  
85 CDS cells with a wide range of preferred directions, 2) strong motion  
86 opponent suppression, and 3) a tuned normalization that may reflect center-  
87 surround interactions in MT. Whereas the implementation by Rust et  
88 al. (2006) was restricted to inputs that are mixtures of sinusoidal gratings of  
89 a fixed spatial and temporal frequency, our model can operate on any  
90 spatiotemporal image intensity.

91       The motion energy model of Simoncelli and Heeger (1998), henceforth  
92 referred to as the S&H model, is conceptually equivalent to an elaborated  
93 Reichardt detector at the end of the V1 stage (van Santen and Sperling  
94 1985), and is a specific implementation of the intersection-of-constraints  
95 (IOC) principle at the end of the MT PDS stage (Bradley and Goyal 2008). The  
96 IOC principle in turn is one possible solution to the aperture problem; that is,  
97 a velocity-space construction that finds the global pattern motion as the  
98 point in velocity-space where the constraint lines of all local velocity samples  
99 intersect. Adelson and Movshon (1982) differentiated among three methods  
100 to estimate the global pattern motion; 1) IOC principle, 2) vector average  
101 (VA), and 3) blob or feature tracking, which may be equally valid approaches  
102 to solving the aperture problem (for a recent review on the topic see Bradley  
103 and Goyal (2008)). Although the S&H model is not complete, in the sense  
104 that it does not specify the exact pattern or object velocity, the model in  
105 particular and the IOC principle in general are consistent with various  
106 experimental data.

107       In the present paper, we introduce a large-scale spiking neuron model  
108 of cortical areas critical for motion processing, which is efficient enough to  
109 run in real-time on available processors. We show that the responses of

110neurons in the network are comparable to electrophysiological results for  
111grating and plaid stimuli, as well as speed tuning. The behavioral response of  
112the network in a two-alternative forced choice (2AFC) motion discrimination  
113task (that is, a random dot motion coherence task) is in agreement with  
114psychophysical data. Moreover, our implementation outperforms a previous  
115rate-based C/Matlab implementation of the S&H model by up to a factor of  
11612 in terms of computational speed and by orders of magnitude in terms of  
117memory usage. The full network, which comprises 153,216 neurons and  
118approximately 40 million synapses, processes 20 frames per second of a  
11940 × 40 input video in real-time using a single off-the-shelf GPU.

120       The network was constructed using an open-source SNN simulator  
121(Richert et al. 2011) that provides a PyNN-like programming interface; its  
122neuron model, synapse model, and address-event representation (AER) are  
123compatible with recent neuromorphic hardware (Srinivasa and Cruz-Albrecht  
1242012). To promote the use of this algorithm among the neuroscientist and  
125computer vision research communities, the source code for the simulator,  
126the network, and analysis scripts are publicly available at  
127<http://www.socsci.uci.edu/~jkrichma/CARLsim/>.

## 1282. **Methods**

### 1292.1 **The simulator**

130The present model was developed on a simulator that was previously  
131published in Nageswaran et al. (2009) and Richert et al. (2011). The first  
132study demonstrated real-time performance for a simulation of 100,000  
133neurons on a single NVIDIA C1060 GPU. The latter added a wide range of  
134functionalities, such as equations for synaptic conductances, spike-timing-  
135dependent plasticity (STDP), and short-term plasticity (STP). The present  
136release builds on this mainly by: 1) providing the complete source code for a  
137detailed large-scale model of visual motion processing in V1 and MT, 2)  
138improving the original model to demonstrate PDS responses and speed  
139tuning, and 3) introducing source code-level optimizations that improve GPU  
140memory management and ensure code stability. Whereas the optimizations  
141should be applicable to a wide range of GPU architectures, they are not  
142directly relevant to this paper and will thus not be discussed (for more  
143information please refer to the release notes).

144The main code to run the experiments described in this paper can be found  
145in the file "examples/v1MTLIP/main\_v1MTLIP.cpp", which is part of the CARLsim 2.1  
146software package. The "examples" directory also contains a number of other  
147experiments that were part of a previous code release—for more information  
148refer to Richert et al. (2011). Matlab scripts to analyze the network output  
149and create the figures can be found in the directory "scripts/v1MTLIP/". Please  
150note that Matlab is not necessary to use the simulator, as the scripts are  
151provided mainly for analysis purposes.

#### 1522.1.1 **Setting up a simulation**

153Step-by-step instructions on how to set up, interact with, and run a  
154simulation can be found in the tutorial on our website and in our previous  
155code release (Richert et al. 2011). For the reader's convenience, we include  
156here a representative example to illustrate the ease of setting up and

157running a simulation. Listing 1 randomly connects ten Poisson spike  
158generators (gIn) firing at 50 Hz mean rate to a population of 100 excitatory  
159Izhikevich neurons (gEx), records and stores the spike times in a binary file  
160"spkEx.dat", and runs the network for a second of simulation time:

```
#include "snn.h"
CpuSNN sim("My network");

// set up network
int gIn=sim.createSpikeGeneratorGroup("input", 10, EXCITATORY_NEURON);
int gEx=sim.createGroup("excitatory", 100, EXCITATORY_NEURON);
sim.setNeuronParameters(gEx, 0.02f, 0.2f, -65.0f, 8.0f); // RS neurons
sim.connect(gIn, gEx, "random", 1.0, 1.0, 0.10f, 1, 20, SYN_FIXED);

// write spike times to file
sim.setSpikeMonitor(gEx, "spkEx.dat");

// set spike rates and run network
PoissonRate inSpikes(100);
for (int i=0; i<100; i++)
    inSpikes.rates[i] = 50.0f; // 50 Hz
sim.setSpikeRate(gIn, &inSpikes);
sim.runNetwork(1,0); // run for 1 sec and 0 msec
```

#### 161**Listing 1**

162In this example, connectivity (achieved through `CpuSNN::connect(...)`) is random  
163with an initial weight of 1.0, a maximum weight of 1.0, a 10 % (0.10)  
164connection probability, a synaptic delay uniformly distributed between 1 ms  
165and 20 ms, and static synapses (SYN\_FIXED). Note that any type of  
166connectivity profile is possible by using a callback mechanism. For a  
167description of the Izhikevich neuron model please refer to Section 2.1.3.

#### 168**2.1.2 CPU vs. GPU simulation mode**

169A major advantage of our simulator is the possibility to run a simulation  
170either on standard x86 central processing units (CPUs) or off-the-shelf NVIDIA  
171GPUs, simply by passing a constant with value CPU\_MODE or GPU\_MODE as an  
172additional function argument to `CpuSNN::runNetwork(...)`. A new feature is the  
173option to pass a “device index” to the same method, which can be used in  
174multi-GPU systems to specify on which CUDA device to establish a context.  
175For example, Listing 2 would run a built network for one second on the  
176second GPU (if such a device exists):



```
CpuSNN sim("My network");
... // build network
int run_sec = 1; int run_msec = 0; // run for 1 s and 0 ms
bool onGPU = true; // run on GPU
int ithGPU = 1; // run on 2nd device (0-indexed)
sim.runNetwork(run_sec, run_msec, onGPU?GPU_MODE:CPU_MODE, ithGPU);
```

177**Listing 2**

178 The two simulation modes allow the user to exploit the advantages of  
 179 both architectures. Whereas the CPU is more efficient for relatively small  
 180 networks, the GPU is most advantageous for network sizes of 1,000 neurons  
 181 and up (Nageswaran et al. 2009; Richert et al. 2011). It has been  
 182 demonstrated that a GPU implementation (on NVIDIA GTX-280 with 1 GB of  
 183 memory) for a simulation of 100,000 neurons and 50 million synaptic  
 184 connections can run up to 26 times faster than a CPU version (Core2 4600  
 185 @ 2.13 GHz with 4 GB of memory) of the same network (Nageswaran et al.  
 186 2009). On the other hand, the CPU mode allows for execution of extremely  
 187 large networks that would not fit within the GPU's memory.

188 It is worth noting that a simulation can be run in CPU mode even if the  
 189 code is compiled in the presence of CUDA source files. An example of this  
 190 hybrid mode is the network explained in the present work, which contains a  
 191 V1 stage purely written in CUDA. In this case the network would be allocated  
 192 on the CPU's memory, but the generation of motion energy responses would  
 193 be delegated to the GPU.

194**2.1.3 Neuron model**

195 The simulator currently supports four parameter Izhikevich point-neurons  
 196 (Izhikevich 2003). Other neuron models will follow in future releases. The  
 197 Izhikevich model aims to reduce Hodgkin-Huxley-type neuronal models to a  
 198 two-dimensional system of ordinary differential equations,

$$\frac{dv(t)}{dt} = 0.04v^2(t) + 5v(t) + 140 - u(t) + i_{syn}(t) \quad (1)$$

$$\frac{du(t)}{dt} = a(bv(t) - u(t)). \quad (2)$$

199 Here (1) describes the membrane potential  $v$  for a given external current  $i_{syn}$ ,  
 200 whereas (2) describes a recovery variable  $u$ ; the parameter  $a$  is the rate

201 constant of the recovery variable, and the parameter  $b$  describes the  
202 sensitivity of the recovery variable to the subthreshold fluctuations of the  
203 membrane potential. All parameters in (1) and (2) are dimensionless;  
204 however, the right-hand side of (1) is in a form such that the membrane  
205 potential  $v$  has mV scale and the time  $t$  has ms scale (Izhikevich 2003). The  
206 Izhikevich model is well-suited for large-scale simulations, because it is  
207 computationally inexpensive yet capable of spiking, bursting, and being  
208 either an integrator or a resonator (Izhikevich 2004, 2007).

209 In contrast to other simple models such as the leaky integrate-and-fire  
210 (LIF) neuron, the Izhikevich neuron is able to generate the upstroke of the  
211 spike itself. Thus the voltage reset occurs not at the threshold, but at the  
212 peak ( $v_{cutoff}=+30$ ), of the spike. The action potential downstroke is modeled  
213 using an instantaneous reset of the membrane potential whenever  $v$  reaches  
214 the spike cutoff, plus a stepping of the recovery variable:

$$v(v>30)=c \quad \text{and} \quad u(v>30)=u-d. \quad (3)$$

215 The inclusion of  $u$  in the model allows for the simulation of typical spike  
216 patterns observed in biological neurons. The four parameters  $a$ ,  $b$ ,  $c$ , and  $d$   
217 can be set to simulate different types of neurons. Unless otherwise specified,  
218 excitatory neurons in all our simulations were modeled as regular spiking  
219 (RS) neurons (class 1 excitable,  $a=0.02$ ,  $b=0.2$ ,  $c=-65$ ,  $d=8$ ), and all  
220 inhibitory neurons were modeled as fast spiking (FS) neurons (class 2  
221 excitable,  $a=0.1$ ,  $b=0.2$ ,  $c=-65$ ,  $d=2$ ) (Izhikevich 2003, 2004).

#### 222 2.1.4 Synapse model

223 A simulation can be run with either a current-based or a conductance-based  
224 neuron model (sometimes referred to as CUBA and COBA, respectively). All  
225 experiments in the present study were run in COBA mode.

226 In a conductance-based model, each ionic current that contributes to  
227 the total current  $i_{syn}$  (see (1)) is associated with a conductance. The simulator  
228 supports four of the most prominent synaptic conductances found in the  
229 cortex: AMPA (fast decay), NMDA (slow decay and voltage-dependent),

230 GABA<sub>a</sub> (fast decay), and GABA<sub>b</sub> (slow decay), which are modeled as dynamic  
 231 synaptic channels with zero rise time and exponential decay according to

$$\frac{dg_r(t)}{dt} = -\frac{1}{\tau_r}g_r(t) + w \sum_i \delta(t-t_i), \quad (4)$$

232 where  $\delta$  is the Dirac delta, the sum is over all presynaptic spikes arriving at  
 233 times  $t_i$ ,  $w$  is the weight of that synapse,  $\tau_r$  is its decay time constant, and  
 234 the subscript  $r$  denotes the receptor type; that is, AMPA, NMDA, GABA<sub>a</sub>, or  
 235 GABA<sub>b</sub>. Unless otherwise specified, a spike arriving at a synapse that is post-  
 236 synaptically connected to an excitatory (inhibitory) neuron increases both  
 237  $g_{AMPA}$  and  $g_{NMDA}$  ( $g_{GABA_a}$  and  $g_{GABA_b}$ ). In our simulations we set the time  
 238 constants to  $\tau_{AMPA} = 5$  ms,  $\tau_{NMDA} = 150$  ms,  $\tau_{GABA_a} = 6$  ms, and  $\tau_{GABA_b} = 150$  ms  
 239 (Dayan and Abbott 2001; Izhikevich et al. 2004). The rise time of these  
 240 conductances was modeled as instantaneous, which is a reasonable  
 241 assumption in the case of AMPA, NMDA, and GABA<sub>a</sub> (Dayan and Abbott  
 242 2001), but a simplification in the case of GABA<sub>b</sub>, which has a rise time on the  
 243 order of 10 ms (Koch 1999).

244 Then the total synaptic current  $i_{syn}$  in (1) for each neuron is given by:

$$i_{syn} = -g_{AMPA}(v-0) - g_{NMDA} \frac{\left(\frac{v+80}{60}\right)^2}{1 + \left(\frac{v+80}{60}\right)^2} (v-0) - g_{GABA_a}(v+70) - g_{GABA_b}(v+90), \quad (5)$$

245 where  $v$  is the membrane potential of the neuron, and the subscript  
 246 indicates the receptor type. This equation is equivalent to the one described  
 247 in Izhikevich et al. (2004).

## 248.2 The network

249 The network architecture is shown in Fig. 1. Grayscale videos are fed frame-  
 250 by-frame through a model of the primary visual cortex (V1), the middle  
 251 temporal area (MT), and the lateral intraparietal cortex (LIP). Bold black  
 252 arrows indicate synaptic projections. Note that inhibitory populations and

253projections are not shown for the sake of clarity. Numbers in parentheses  
254next to an element are the equations that describe the corresponding  
255neuronal response or synaptic projections, as will be explained in the  
256subsections below.

257       The V1 model consisted of a bank of spatiotemporal filters (rate-based)  
258according to the S&H model (Simoncelli and Heeger 1998), which will be  
259described in detail in Section 2.2.1. At each point in time, a  $32 \times 32$  input  
260video frame was processed by V1 cells at three different spatiotemporal  
261resolutions (labeled “3 scales” in Fig. 1). Simulated V1 simple cells computed  
262an inner product of the image contrast with one of 28 space-time oriented  
263receptive fields (third derivatives of a Gaussian), which was then half-wave  
264rectified, squared, and normalized within a large Gaussian envelope. V1  
265complex cell responses were computed as a weighted sum of simple cell  
266afferents that had the same space-time orientation, but were distributed  
267over a local spatial region. We interpreted these filter responses as mean  
268firing rates of Poisson spike trains (labeled “Hz” in the figure) as explained in  
269Section 2.2.1, which were first scaled to match the contrast sensitivity  
270function of V1 simple cells, and then used to drive Izhikevich spiking neurons  
271representing cells in area MT.

272       Area MT consisted of two distinct populations of spiking neurons  
273(explained in Section 2.2.2), the first one being selective to all local  
274component motions of a stimulus (CDS cells), and the other one responding  
275to the global pattern motion (PDS cells). MT CDS cells responded to three  
276different speeds (1.5 pixels/frame, 0.125 pixels/frame, and 9 pixels/frame)  
277illustrated as three distinct populations in the MT CDS layer of Fig. 1.  
278Divisive normalization between these populations enabled the generation of  
279speed tuning curves that are in agreement with neurophysiological  
280experiments (Rodman and Albright 1987). The three MT CDS populations  
281consisted of eight subpopulations, each of which was not only selective to a  
282particular speed but also to one of eight directions of motion, in 45 degree  
283increments. PDS cells were constructed by 1) pooling over MT CDS cells with

284a wide range of preferred directions, 2) using strong motion opponent  
285suppression, and 3) employing a tuned normalization that may reflect  
286center-surround interactions in MT (Rust et al. 2006). PDS cells were  
287selective to the same speed as their CDS afferents. For the purpose of this  
288paper we only implemented PDS cells selective to a speed of 1.5  
289pixels/frame (see MT PDS layer in Fig. 1) to be used in a motion  
290discrimination task. However, it is straightforward to implement PDS cells  
291that are selective to another speed.

292       A layer of decision neurons (see Section 2.2.3) was responsible for  
293integrating over time the direction-specific sensory information that is  
294encoded by the responses of MT PDS cells. Analogous to the MT layer, the  
295decision layer consisted of eight subpopulations, each of which received  
296projections from a subpopulation of MT PDS cells selective to one of eight  
297directions of motion. This information was then used to make a perceptual  
298decision about the presented visual stimulus, such as determining the global  
299drift direction of a field of random moving dots in a motion discrimination  
300task (presented in Section 3.3). Fig. 1 exemplifies this situation by showing a  
301snapshot of the network's response to a random dot kinematogram (RDK)  
302where dots drift to the right at a speed of 1.5 pixels/frame. The  
303subpopulation of decision neurons that is coding for rightward motion is  
304activated the strongest. The temporal integration of sensory information  
305might be performed in one of several parietal and frontal cortical regions in  
306the macaque, such as LIP, where neurons have been found whose firing rate  
307are predictive of the behavioral reaction time (RT) in a RDK task (Shadlen  
308and Newsome 2001; Roitman and Shadlen 2002).

309       The following subsections will explain the model in detail.

### 310**2.2.1 Spatiotemporal-energy model of V1**

311The first (V1) stage of the S&H model was implemented and tested in a  
312Compute Unified Device Architecture (CUDA) environment (Richert et al.  
3132011). This part of the model is equivalent to Eqs. 1-4 in Simoncelli and

314 Heeger (1998) and their subsequently released C/Matlab code, which can be  
 315 obtained from: <http://www.cns.nyu.edu/~lcv/MTmodel/>. Unless otherwise  
 316 stated, we used the same scaling factors and parameter values as in the  
 317 S&H model.

318 A visual stimulus is represented as a light intensity distribution  
 319  $I(x, y, t)$ , that is, a function of two spatial dimensions  $(x, y)$  and time  $t$ . The  
 320 stimulus was processed at three different spatiotemporal resolutions (or  
 321 scales),  $r$  (labeled “3 scales” in Fig. 1). The first scale,  $r=0$ , was equivalent to  
 322 processing at the original image (and time) resolution. The other two scales  
 323 were achieved by successively blurring the image with a Gaussian kernel.  
 324 The three stimuli  $I_r(x, y, t)$  can thus be expressed as:

$$\begin{aligned}
 I_0(x, y, t) &= I(x, y, t) \\
 I_1(x, y, t) &= \exp\left(\frac{-(x^2 + y^2 + t^2)}{2}\right) * I_0(x, y, t) \\
 I_2(x, y, t) &= \exp\left(\frac{-(x^2 + y^2 + t^2)}{2}\right) * I_1(x, y, t),
 \end{aligned} \tag{6}$$

325 where  $*$  denotes convolution. In order to circumvent the non-causality of  
 326 these convolutions (the response depends both on past and future stimulus  
 327 intensities), a time delay of four frames was introduced (see (Simoncelli and  
 328 Heeger 1998)).

329

330 **V1 simple cells.** A large body of research has found that neurons located in  
 331 V1 that project to MT are directionally selective and may be regarded as  
 332 local motion energy filters (Adelson and Bergen 1985; DeAngelis et al. 1993;  
 333 Movshon and Newsome 1996). In our network, V1 simple cells are modeled  
 334 as linear space-time-oriented filters whose receptive fields are third  
 335 derivatives of a Gaussian (Simoncelli and Heeger 1998). These filters are  
 336 very similar to a Gabor filter, but more computationally convenient as they  
 337 allow for separable convolution computations.

338 The full set of V1 linear receptive fields consisted of 28 space-time  
 339 orientations that are evenly distributed on the surface of a sphere in the

340spatiotemporal frequency domain. The  $k$ th space-time-oriented filter in the  
341V1 population can be described by a unit vector  $\hat{u}_k = (\hat{u}_{k,x}, \hat{u}_{k,y}, \hat{u}_{k,t})'$  that is  
342parallel to the filter orientation, where  $k=1, 2, \dots, 28$  and ' denotes vector  
343transposition. For more information please refer to Simoncelli and  
344Heeger (1998). An example of a spatiotemporal receptive field is illustrated  
345in Fig. 2, where the colored ovals correspond to the orientation of the  
346positive (green) and negative (red) lobes of the spatiotemporal filter. If a  
347drifting dot traces out a path (dashed line) in space ( $x$ , for now ignoring  $y$ )  
348and time ( $t$ ) that is oriented in the same way as the lobes, then the filter  
349could be activated by this motion (Fig. 2a). A dot moving in the orthogonal  
350direction would not elicit a filter response because its path intersects both  
351positive and negative lobes of the filter (as depicted in Fig. 2b).

352 First, input images were filtered with a 3D Gaussian corresponding to  
353the receptive field size of a V1 simple cell:

$$f_r(x, y, t) = \exp\left(\frac{-(x^2 + y^2 + t^2)}{2\sigma_{v1simple}^2}\right) * I_r(x, y, t) \quad (7)$$

354where  $*$  is the convolution operator,  $r$  denotes the scale, and  $\sigma_{v1simple} = 1.25$   
355pixels.

356Then the underlying linear response of a simple cell at spatial location  $(x, y)$   
357and scale  $r$  with space-time orientation  $k$  is equivalent to the third-order  
358derivative in the direction of  $\hat{u}_k$ ; that is,

$$L_{kr}(x, y, t) = \alpha_{v1lin} \sum_{T=0}^3 \left[ \sum_{Y=0}^{3-T} \left[ \frac{3!}{X!Y!T!} (\hat{u}_{k,x})^X (\hat{u}_{k,y})^Y (\hat{u}_{k,t})^T \frac{\partial^3 f_r(x, y, t)}{\partial x^X \partial y^Y \partial t^T} \right] \right] \quad (8)$$

359where  $!$  denotes the factorial,  $X=3-Y-T$ , and  $\alpha_{v1lin} = 6.6084$  is a scaling  
360factor. Note that the two sums combined yield exactly 28 summands. This  
361operation is equivalent to Eq. 2 in the original paper, and can also be  
362expressed using vector notation:

$$L_r = \alpha_{v1lin} M b_r, \quad (9)$$

363 where  $L_r$  is the set of all V1 responses at scale  $r$ , each element of  $b_r$  is one of  
 364 the separable derivatives in (8) at scale  $r$ , and each element of the  $28 \times 28$   
 365 matrix  $M$  is a number  $3!!(X!Y!T!)(\hat{u}_{k,x})^X(\hat{u}_{k,y})^Y(\hat{u}_{k,t})^T$ . Each row of  $M$  has a  
 366 different value for  $k$ , and each column of  $M$  has different values for  $X$ ,  $Y$ , and  
 367  $T$ . We will make use of this notation in Section 2.2.2, where we will explain  
 368 the construction of synaptic projections from V1 to MT.

369 At this stage of the model it is possible that filter responses  $L_{kr}$  at  
 370 positions  $(x, y)$  close to the image border have become unreasonably large.  
 371 We suppressed these edge effects by applying a scaling factor to  $L_{kr}$   
 372 whenever  $(x, y)$  was near an image border.

373 Simple cell responses were constructed by half-squaring and  
 374 normalizing the linear responses  $L_{kr}$  from (8) within a large Gaussian  
 375 envelope:

$$S_{kr}(x, y, t) = \frac{\alpha_{filt \rightarrow rate, r} \alpha_{v1rect} [L_{kr}(x, y, t)]^2}{\alpha_{v1norm} \exp\left(\frac{-(x^2+y^2)}{2\sigma_{v1norm}^2}\right) * \left(\frac{1}{28} \sum_{k=1}^{28} [L_{kr}(x, y, t)]^2\right) + \alpha_{v1semi}^2}, \quad (10)$$

376 where  $[.]$  denotes half-wave rectification, and  $*$  is the convolution operator.  
 377 The scaling factors  $\alpha_{v1rect}=1.9263$  and  $\alpha_{v1semi}=0.1$  (the semi-saturation  
 378 constant) had the same values as in the original S&H model. Instead of  
 379 having a single global normalization, our normalization occurs within a large  
 380 spatial neighborhood (Gaussian half-width  $\sigma_{v1norm}=3.35$  pixels), which is  
 381 thought to be more biologically realistic. Therefore the scaling factor  
 382  $\alpha_{v1norm}=1.0$  had to be adjusted to compensate for the implementation  
 383 difference. This was done simultaneously by setting  $\alpha_{filt \rightarrow rate, r}=15\text{Hz}$ , a  
 384 scaling factor to map the unit-less filter responses at each scale  $r$  onto more  
 385 meaningful mean firing rates, as will be explained below. In brief, we opted  
 386 to reproduce the contrast sensitivity function reported for V1 cells projecting  
 387 to MT (Movshon and Newsome 1996). Other than that, the computation in  
 388 (10) is conceptually equivalent to Eqs. 3–4 in Simoncelli and Heeger (1998).



389

390**V1 complex cells.** V1 complex cell responses were computed as local  
391weighted averages of simple cell responses,

$$C_{kr}(x, y, t) = \alpha_{v1comp} \exp\left(\frac{-(x^2 + y^2)}{2\sigma_{v1comp}^2}\right) * S_{kr}(x, y, t), \quad (11)$$

392where the half-width of the Gaussian was  $\sigma_{v1comp} = 1.6$ , and  $\alpha_{v1comp} = 0.1$  is a  
393scaling factor.

394 The responses  $C_{kr}(x, y, t)$  described in (11) served as output of the  
395CUDA implementation. These responses were interpreted as mean firing  
396rates of Poisson spike generators, following the procedure described in the  
397next subsection. V1 complex cells then projected to MT CDS cells as  
398explained in Section 2.2.2.

399

400**Converting filter responses to firing rates.** In order to find a meaningful  
401mapping from unit-less filter responses to mean firing rates, we opted to  
402reproduce the contrast sensitivity function reported for V1 cells projecting to  
403MT (Movshon and Newsome 1996), which is shown in Fig. 3. The red line is  
404the electrophysiological data adapted from Fig. 7 of Movshon and  
405Newsome (1996), whereas the blue line is our simulated data. In order to  
406arrive at this plot, we presented a drifting sinusoidal grating of varying  
407contrast to V1 simple cells coding for scale  $r=0$ , and computed their mean  
408response  $\langle S_{k0} \rangle$  from (10) over a stimulation period of one second. The drifting  
409grating had a spatial frequency of  $\omega_{spat} = 0.1205$  cycles/pixel and a temporal  
410frequency of  $\omega_{temp} = 0.1808$  cycles/frame, which is equivalent to the one used  
411in Section 3.1 for MT direction tuning. Because the grating was drifting to the  
412right, we only looked at the subpopulation of V1 simple cells that responded  
413maximally to this stimulus (which was true for  $k=24$ ). The mean firing rate of  
414neurons in this subpopulation,  $\langle S_{24,0} \rangle$ , was then averaged over all cells in the  
415subpopulation and plotted in Fig. 3 (blue curve) for  $\alpha_{v1norm} = 1.0$  and

416  $\alpha_{filt \rightarrow rate, 0} = 15 \text{ Hz}$ . Vertical bars are the standard deviation on the population  
417 average. The scaling factor  $\alpha_{v1norm}$  was gradually changed until the curvature  
418 of the blue graph approximated the curvature of the electrophysiological  
419 data. The scaling factor  $\alpha_{filt \rightarrow rate, 0}$  was then adjusted such that the simulated  
420 responses saturated at approximately 100 Hz.

421 In order to tune V1 simple cells at the other two scales, that is,  $S_{k1}$  and  
422  $S_{k2}$  from (10), we used a RDK stimulus, which is depicted as the sample input  
423 in Fig. 1 and explained in detail in Section 3.3. We chose scaling factors that  
424 would give equal response magnitudes at all three scales in response to the  
425 RDK stimulus, which resulted in  $\alpha_{filt \rightarrow rate, 1} = 17 \text{ Hz}$  and  $\alpha_{filt \rightarrow rect, 2} = 11 \text{ Hz}$ .

426 Because these filter response were transformed to mean firing rates, it  
427 was straight-forward to assign the responses  $C_{kr}(x, y, t)$  described in (11) to  
428 mean firing rates of Poisson spike generators, which served as input to the  
429 spiking neurons in area MT. The exact mapping of V1 complex onto MT CDS  
430 cells is given in (12) (see Section 2.2.2).

### 431 2.2.2 Two-stage spiking model of MT

432 The two-stage model of MT is based on the idea that CDS cells represent an  
433 earlier stage of motion processing than PDS cells (Movshon et al. 1985; M. A.  
434 Smith et al. 2005). The present model is built on this idea, making MT CDS  
435 cells similar in terms of direction and speed tuning to the model V1 complex  
436 cells used by Simoncelli and Heeger (1998). In fact, it has been shown that  
437 MT cells exhibit speed tuning characteristics similar to V1 complex cells  
438 (Priebe et al. 2006), which has led to the suggestion that speed tuning in MT  
439 might be inherited from V1. Livingstone and Conway (2007) have shown that  
440 even some V1 simple cells are speed-tuned in macaque. Whereas CDS cells  
441 give responses whose selectivity is stable and consistent from the time they  
442 are first activated, PDS cells often respond with different and broader  
443 selectivity when first activated, sometimes even resembling CDS cells, and  
444 only over a time-course on the order of 100 ms do they establish pattern

445selectivity (M. A. Smith et al. 2005). At least in anesthetized monkeys, MT is  
446believed to consist of roughly 40 % CDS cells, 25 % PDS cells, and 35 %  
447unclassified cells (Movshon et al. 1985). However, in awake animals the  
448situation might be more complicated (Pack et al. 2001).

449 All cells in MT were Izhikevich spiking neurons, whose membrane  
450potential was thus described by a pair of coupled differential equations (see  
451(1) and (2)).

452

453**Component-direction-selective cells.** CDS cells are selective to a  
454particular direction and speed of motion (an orientation in space-time). The  
455name is an indication that these cells, when presented with a plaid stimulus  
456consisting of two superimposed sine gratings, preferably respond to the  
457motion of each grating (component) rather than the global motion pattern  
458produced by the combination of the two gratings (Movshon et al. 1985).

459 MT CDS cells in our model responded preferentially to motion in one of  
460eight different directions (in 45 degree increments) and three different  
461speeds (1.5 pixels per frame, 0.125 pixels per frame, and 9 pixels per frame)  
462at any pixel location. These values can be easily adjusted by running the  
463Matlab script "scripts/v1MTLIP/projectV1toMT.m". The response properties of MT  
464CDS cells were given by 1) a set of both excitatory and inhibitory  
465interpolated weights (as explained next) coming from V1 complex cells  
466(Simoncelli and Heeger 1998), and 2) projections from an inhibitory group of  
467MT interneurons to account for response normalization.

468 Because the directional derivatives of a Gaussian are steerable  
469(Freeman and Adelson 1991), the response of an arbitrarily oriented filter  
470can be synthesized from a fixed bank of basis filters (the third derivatives of  
471a Gaussian). Thus the projection weights from V1 complex cells to MT were  
472interpolated as follows. Let  $\hat{\alpha} = (\hat{\alpha}_x, \hat{\alpha}_y, \hat{\alpha}_t)'$  be the unit vector parallel to an  
473arbitrary space-time orientation (direction and speed of motion), akin to the

474 unit vectors  $\hat{u}_k$  described in Section 2.2.1. Then we can write the third  
 475 directional derivative in direction of  $\hat{\alpha}$  analogously to (9) as:

476

$$\frac{\partial^3 f_r}{\partial \hat{\alpha}^3} = [v'(\hat{\alpha})M^{-1}]b_r \cdot w_{\hat{\alpha}} b_r, \quad (12)$$

477 where the matrix  $M$  and the vector  $b_r$  are the same as in (9), each element of  
 478 the vector  $v(\hat{\alpha})$  is a number  $6/(X!Y!T!)\hat{\alpha}_x^X \hat{\alpha}_y^Y \hat{\alpha}_t^T$  analogous to (8), and '

479 denotes vector transposition. The product  $[v'(\hat{\alpha})M^{-1}]$  thus is a set

480  $w_{\hat{\alpha}} = (w_{\hat{\alpha},1}, \dots, w_{\hat{\alpha},28})$  of interpolated weights, where the  $k$ th element of this

481 vector,  $w_{\hat{\alpha},k}$ , determined the strength of the projection from the  $k$ th V1

482 complex cell onto a MT CDS cell. The two cells were connected only if they

483 were located at the same pixel location,  $(x, y)$ . Speed tuning arose from the

484 fact that  $\hat{\alpha}$  corresponds to a specific direction and speed of motion. Thus, in

485 order to achieve MT CDS cells tuned to different speeds,  $\hat{\alpha}$  was the only

486 parameter that needed to be adjusted (refer to the Matlab script mentioned

487 above). A MT CDS cell received projections from V1 complex cells at all three

488 spatiotemporal resolutions,  $r$ . Note that it is possible to construct a network

489 with the same functionality by using only one spatiotemporal resolution,

490 which has been shown in Simoncelli and Heeger's own C/Matlab

491 implementation. Using multiple spatiotemporal resolutions, however, makes

492 the network more robust in responding to motion of different-sized objects.

493 Because the interpolated weights could assume both positive and

494 negative values, it was necessary to relay the projections with negative

495 weights to a population of inhibitory neurons. In this case (that is, if  $w_{\hat{\alpha},k} < 0$ ),

496 the weights in (12) are applied to excitatory projections from V1 complex

497 cells to the MT inhibitory population (where  $w_{\hat{\alpha},k,inh} = |w_{\hat{\alpha},k}|$ ), and the inhibitory

498 population sends one-to-one connections back to the pool of MT CDS cells.

499 Overall the interpolated weights are equivalent to the parameters  $\rho_{nm}$  in

500 Eq. 5 of Simoncelli and Heeger (1998).

501 In order to model response normalization equivalent to the one in Eq. 6  
 502 of Simoncelli and Heeger (1998), we introduced another pool of inhibitory  
 503 interneurons, which integrated the activity of all MT CDS cells within a large  
 504 Gaussian neighborhood (across direction and speed), and projected back to  
 505 all three pools of MT CDS cells with one-to-one connections. This response  
 506 normalization is important to qualitatively reproduce the speed tuning curves  
 507 (see Section 3.2).

508

509 **Pattern-direction-selective cells.** PDS cells differ from CDS cells in that  
 510 they, when presented with a plaid stimulus consisting of two superimposed  
 511 sine gratings, preferentially respond to the overall motion direction, not the  
 512 individual components (Movshon et al. 1985). Because visual stimuli typically  
 513 contain many oriented components, local motion measurements must be  
 514 appropriately combined in order to sense the true global motion of the  
 515 stimulus (aperture problem). Thus it has been suggested that PDS neurons  
 516 reflect a higher-order computation that acts on V1 or MT CDS afferents  
 517 (Movshon et al. 1985). MT PDS cells in our model received direct input from  
 518 CDS cells, and thus conserved their speed and direction preferences.

519 Pooling over MT CDS cells and opponent suppression were  
 520 implemented by pooling CDS responses across spatial position and across  
 521 direction preference, such that the strength of a projection from a CDS cell  
 522 selective to motion direction  $\theta_{CDS}$  at location  $(x_{CDS}, y_{CDS})$  to a PDS cell  
 523 selective to motion direction  $\theta_{PDS}$  at location  $(x_{PDS}, y_{PDS})$  can be expressed as:

$$w_{CDS \rightarrow PDS} = \alpha_{CDS \rightarrow PDS} \cos(\Delta\theta) \exp\left(\frac{-((\Delta x)^2 + (\Delta y)^2)}{2\sigma_{PDS, pool}^2}\right), \quad (13)$$

524 where  $\Delta\theta = \theta_{PDS} - \theta_{CDS}$ ,  $\Delta x = x_{PDS} - x_{CDS}$ ,  $\Delta y = y_{PDS} - y_{CDS}$ , the half-width of the  
 525 Gaussian neighborhood  $\sigma_{PDS, pool} = 3$  pixels, and  $\alpha_{CDS \rightarrow PDS}$  is a scaling factor. If

526 the resulting weight was negative, due to  $|\Delta\theta| > \frac{\pi}{2}$ , the projection was relayed  
 527 to a population of inhibitory interneurons. Following the reasoning of Rust et

528al. (2006), the pattern index of a MT cell can be reduced simply by  
 529sharpening the cosine tuning component in (13) (see third column of Fig. 6 in  
 530Rust et al. (2006)).

531 Tuned normalization was implemented by an inhibitory self-connection  
 532with a narrowly tuned Gaussian across direction (see second column of Fig. 6  
 533in Rust et al. (2006)). Analogous to previous projections, this was  
 534implemented by relaying the inhibitory projection to a pool of inhibitory  
 535interneurons:

$$w_{PDS \rightarrow PDS, inh} = \exp\left(\frac{-(\Delta\theta)^2}{2\sigma_{PDS, tuned, dir}^2}\right) \exp\left(\frac{-((\Delta x)^2 + (\Delta y)^2)}{2\sigma_{PDS, tuned, loc}^2}\right), \quad (14)$$

536where  $\sigma_{PDS, tuned, dir} < 45$  deg (such that only one of the eight subpopulations  
 537was activated),  $\sigma_{PDS, tuned, loc} = 2$  pixels, and the inhibitory population sent one-  
 538to-one connections back to the pool of MT PDS cells.

### 5392.2.3 Spiking layer of LIP decision neurons

540A layer of decision neurons was responsible for integrating over time the  
 541direction-specific sensory information that is encoded by the responses of MT  
 542PDS cells. This information was then used to make a perceptual decision  
 543about the presented visual stimulus, such as determining the global drift  
 544direction of a field of random moving dots in a motion discrimination task  
 545(presented in Section 3.3). A good candidate for such an integrator area in  
 546macaques might be LIP, where neurons have been found whose firing rate  
 547are predictive of the behavioral reaction time (RT) in a motion discrimination  
 548task (Shadlen and Newsome 2001; Roitman and Shadlen 2002).

549 Spiking neurons in a simulated LIP area were grouped into eight pools  
 550of 50 neurons, each pool receiving projections from exactly one of the eight  
 551pools of MT PDS cells with 10 % connection probability. As a result of this  
 552connectivity profile, each pool of decision neurons accumulated sensory  
 553evidence for a particular direction of motion, based on the response of MT  
 554PDS cells.

555 Additionally, each decision pool received inhibitory projections from  
556 other decision pools if the two preferred directions of motion were close to  
557 opposite. More precisely, a decision neuron in pool  $i$  (thus selective to  
558 direction  $\theta_i$ ) received an inhibitory projection from neurons in pool  $j$   
559 (selective to direction  $\theta_j$ ) with strength

$$w_{dec,inh \rightarrow dec} = [\cos(\theta_i - \theta_j + \pi)]^J, \quad (15)$$

560 and 10 % connection probability.

561 LIP decision neurons did not employ any internal noise.

## 562 2.2.4 Implementation details

563 In order for our implementation to be useful to researchers already working  
564 with the S&H model, we tried to stay as close to the S&H C/Matlab  
565 implementation as possible. However, there are a few minor differences  
566 worth mentioning. First, as explained in Section 2.2.1, we normalize V1  
567 simple cell responses in a large Gaussian neighborhood rather than across  
568 the whole population. Second, whereas the S&H model deals with edge  
569 effects by temporarily “padding” the input image with an invisible border, we  
570 opted for the computationally more economical alternative to simply  
571 decrease the responses of V1 simple cells located close to image borders.  
572 Third, in the S&H C/Matlab implementation there are two additional scaling  
573 factors (called  $v1Blur$  and  $v1Complex$ , with values 0.99 and 1.02, respectively)  
574 that we do not apply in order to save execution time. Fourth, our model  
575 processes input images at three different scales as described in (6), which is  
576 a feature that is not implemented in the original S&H model.

577 The most crucial mathematical operation in the V1 stage of the model  
578 is the convolution. Because the filter kernels used in our implementation are  
579 relatively small, employing the fast Fourier transform (FFT) would actually  
580 hurt performance. Instead we perform all convolution operations in the  
581 space-time domain using a custom function, which makes use of the fact  
582 that the Gaussian filter and its derivative are dimensionally separable. Future

583work could be directed towards further optimizing the convolution operation  
584in CUDA.



### 5853. Results

586 We conducted a number of experiments to ensure the accuracy and  
587 efficiency of our implementation. Here we demonstrate that the network is  
588 able to exhibit direction and speed tuning for drifting bar and plaid stimuli  
589 that are in agreement with neurophysiological recordings, and that the  
590 network qualitatively reproduces both the psychometric and chronometric  
591 function in a 2AFC motion discrimination task. Additionally, we measured  
592 both the computational performance and memory consumption of our model  
593 and compared it to the S&H C/Matlab implementation.

594 GPU simulations were run on a NVIDIA Tesla M2090 (6 GB of memory)  
595 using CUDA, and CPU simulations (including Matlab) were run on an Intel  
596 Xeon X5675 at 3.07 GHz (24 GB of RAM). The same exact network running  
597 on a single GPU produced all results; the only difference per experiment was  
598 the presented input stimulus. The full network consisted of 153,216 neurons  
599 and approximately 33 million synapses, which corresponds to a  $32 \times 32$   
600 pixels input resolution.

#### 6013.1 Direction tuning

602 We tested the ability of our model MT cells to signal the direction of motion  
603 for drifting grating and plaid stimuli. Responses were simulated for CDS cells  
604 and PDS cells in MT. The first stimulus was a drifting sinusoidal grating  
605 consisting of spatial and temporal frequency components that were  
606 preferred by MT neurons selective to a speed of 1.5 pixels per frame (that is,  
607  $\omega_{spat}=0.1205$  cycles/pixel,  $\omega_{temp}=0.1808$  cycles/frame). The second stimulus  
608 was a pair of superimposed gratings drifting in a direction orthogonal to their  
609 orientation, which together formed a coherently drifting plaid pattern. The  
610 two gratings both had the same spatial frequency  $\omega_{spat}$ , but their orientation  
611 and drift direction differed by 120 degrees. The direction of these particular  
612 patterns lay equidistant between the directions of motion of the two

613component gratings. The stimulus contrast for both grating and plaid was  
61430 %.

615        Our model was able to reproduce direction tuning curves that are in  
616agreement with single-cell electrophysiological data (Movshon et al. 1985;  
617Rodman and Albright 1989; Movshon and Newsome 1996) for V1 cells, MT  
618CDS cells, and MT PDS cells. Fig. 4 shows polar plots of direction tuning for  
619V1 neurons (Panels b and f), MT CDS cells (Panels c and g), and MT PDS cells  
620(Panels d and h), where the angle denotes motion direction and the radius is  
621the firing rate in spikes per second (compare also Fig. 9 in Simoncelli and  
622Heeger (1998) and Fig. 1 in Rust et al. (2006)). Tuning curves were obtained  
623by calculating the mean firing rate of a neuron's response to a drifting  
624grating during two seconds of stimulus presentation. These responses were  
625averaged over all neurons in the population selective to the same direction  
626of motion (black: mean neuronal response, blue: mean plus standard  
627deviation on the population average, green: mean minus standard  
628deviation). As a result of suppressing edge effects, neurons that coded for  
629locations closer than five pixels from the image border were only weakly  
630activated, and were thus excluded from the plot. The tuning curves in the  
631top row were generated in response to the sinusoidal grating drifting  
632upwards, which is illustrated in Panel a. Analogously, the tuning curves in the  
633bottom row were generated in response to the plaid stimulus drifting  
634upwards, which is illustrated in Panel e (red arrow: pattern motion direction,  
635black arrows: motion direction of the grating components). The direction  
636tuning curve for gratings is unimodal for all three neuron classes, but the  
637direction tuning curve for plaids shows two distinct lobes for V1 complex  
638cells (Panel f) and MT CDS cells (Panel g). Each lobe corresponds to one of  
639the component gratings of the plaid. Only MT PDS cells (Panel h) responded  
640to the motion of the entire plaid pattern rather than to the motions of the  
641individual component gratings.

642        In order to quantify the pattern selectivity of our model PDS cells, we  
643computed the pattern index for each CDS and PDS cell (see Fig. 5) using the

644 standard technique (Movshon et al. 1985; Movshon and Newsome 1996; M.  
645 A. Smith et al. 2005). Based on the tuning curve for the drifting grating  
646 described above, we generated two predictions for each cell's tuning curve  
647 to drifting plaids (Fig. 5a); either the cell would respond to the plaid in the  
648 same way as it responded to the grating ("pattern" prediction, black solid  
649 line), or it would respond independently to the two grating components  
650 ("component" prediction, black dashed line). We then computed the  
651 correlation ( $r_c, r_p$ ) between the cell's actual response to a plaid stimulus and  
652 the component and pattern predictions. To remove the influence of  
653 correlations between the predictions themselves, we calculated partial  
654 correlations  $R_c$  and  $R_p$  for the component and pattern predictions,  
655 respectively, using the standard formulas:

$$R_c = \frac{r_c - r_p r_{pc}}{\sqrt{(1-r_p^2)(1-r_{pc}^2)}} \quad R_p = \frac{r_p - r_c r_{pc}}{\sqrt{(1-r_c^2)(1-r_{pc}^2)}}, \quad (16)$$

656 where  $r_c$  and  $r_p$  are the simple correlations between the data and the  
657 component and pattern predictions, respectively, and  $r_{pc}$  is the simple  
658 correlation between the predictions (Movshon and Newsome 1996). Because  
659 the sampling distribution of Pearson's  $r$  is not normal, we converted the  
660 correlation measures  $R_c$  and  $R_p$  to a Fisher  $Z$ -score,

$$Z_c = \frac{0.5 \ln\left(\frac{1+R_c}{1-R_c}\right)}{\sqrt{\frac{1}{df}}} = \frac{\operatorname{atanh}(R_c)}{\sqrt{\frac{1}{df}}} \quad Z_p = \frac{\operatorname{atanh}(R_p)}{\sqrt{\frac{1}{df}}}, \quad (17)$$

661 where the numerator is the Fisher  $r$ -to- $Z$  transformation and  $df$  is the  
662 degrees of freedom, equal to the number of values in the tuning curve (in  
663 our case 24) minus three (M. A. Smith et al. 2005). The  $Z$ -scores of all CDS  
664 and PDS cells (excluding neurons coding for locations closer than five pixels  
665 from the image border) in the network are plotted in Fig. 5b. Each value of  
666  $Z_c$  and  $Z_p$  was tested for significance using a criterion of 1.28, which is  
667 equivalent to  $P=0.90$  (M. A. Smith et al. 2005). For a PDS cell (red) to be

668judged as pattern-selective, the value of  $Z_p$  had to exceed the value of  $Z_c$  by  
669a minimum of 1.28 (black solid lines). All PDS cells in Fig. 5b met this  
670criterion and, therefore, were indeed pattern-selective. Analogously, all CDS  
671cells (blue) could be judged as component-selective.

### 6723.2 Speed tuning

673We next considered the ability of our implementation to reproduce MT speed  
674tuning curves as demonstrated in Simoncelli and Heeger (1998). MT neurons  
675have been divided into three distinct classes based on their speed tuning  
676properties (Rodman and Albright 1987). The first class of neurons is  
677relatively sharply tuned for a particular speed and direction of motion  
678(“speed-tuned” or “band-pass”). This class of neurons is also strongly  
679suppressed by motion in the anti-preferred (opposite) direction; the  
680suppression is strongest when the stimulus moves in the opposite direction  
681at roughly the preferred speed. The second class of neurons prefers low  
682speeds in both the preferred and anti-preferred direction (“low-pass”). The  
683third class responds to high speed stimuli in both directions (“high-pass”).

684 Fig. 6 faithfully reproduces the speed tuning characteristics of these  
685three distinct classes (compare also Fig. 10 in Simoncelli and Heeger (1998)).  
686The stimulus consisted of a single bar drifting over the entire visual field  
687either to the right (preferred direction) or to the left (anti-preferred direction)  
688at different speeds. Each data point is the mean firing rate of a particular MT  
689CDS neuron located near the center of the visual field, averaged over the  
690time course of a specific speed and direction configuration. The relatively low  
691mean firing rates can be explained by the fact that the stimulus resides  
692outside the neuron’s receptive field for most of the time. The first neuron  
693class (Panel a, “band-pass”) preferentially responded to a bar moving at  
6941.5 pixels per frame to the right, and was strongly suppressed when the bar  
695moved at the same speed to the left. The second neuron class (Panel b, “low-  
696pass”) exhibited a preference for low speeds (0.125 pixels per frame) in both  
697directions. With increasing speed the response of the neuron to dots moving

698in the anti-preferred direction weakened. This behavior can be explained by  
699the fact that the Fourier planes corresponding to low speed motions in  
700opposite directions are both close to the  $\omega_t=0$  plane, and thus close to each  
701other (Simoncelli and Heeger 1998). Also, this class of neurons was  
702suppressed by fast stimuli moving in either direction. Similarly, the third  
703neuron class (Panel c, “high-pass”), which had a high preferred speed  
704(9 pixels per frame) in one direction, was excited by fast stimuli moving in  
705the opposite direction, but was suppressed by slow stimuli moving in either  
706direction.

### 707**3.3 Random dot kinematogram**

708In order to compare the performance of the model with behavioral data from  
7092AFC motion discrimination tasks, we developed a paradigm equivalent to  
710the RDK experiments performed with monkeys and humans (Roitman and  
711Shadlen 2002; Resulaj et al. 2009). We constructed a simple decision  
712criterion based on the race model (Shadlen and Newsome 2001; P. L. Smith  
713and Ratcliff 2004), in which eight pools of decision neurons (one for each of  
714the directions of motion, 50 neurons per pool) sum the responses of MT PDS  
715cells selective to a particular direction and speed of motion. The first decision  
716pool to emit 500 spikes (on average ten spikes per neuron) “won the race”  
717and thus signaled a choice for that direction. A correct decision was the  
718event in which the winning decision pool was selective to the actual motion  
719direction of the stimulus. The time it took the network to reach the decision  
720threshold was termed the reaction time (RT).

721 The RDK stimulus was constructed out of approximately 150 dots  
722(15 % dot density, maximum stimulus contrast) on a 32x32 input movie. An  
723example frame is shown as the input stimulus in Fig. 1. Each stimulus frame  
724was presented to the network for 50 ms. A trial consisted of 20 stimulus  
725frames of a particular motion direction and coherence level. Motion  
726coherence in the stimulus was varied between 0 and 50 %. Coherently  
727moving dots drifted in one of eight possible directions, in 45 degree

728 increments, at a speed of 1.5 pixels per frame. Note that, therefore, only MT  
729 PDS cells that were selective to this particular stimulus speed were  
730 connected to the decision layer.

731 Choice accuracy and RT as a function of task difficulty (coherence of dot  
732 motion) are shown in Fig. 7 (Panel a and b, respectively), where the thick red  
733 lines are human behavioral data extracted from a RT experiment (see Fig. 3  
734 and Table 2 in Roitman and Shadlen (2002)) and simulated data is shown in  
735 blue. Each data point (blue) is the mean outcome of 80 trials (fixed  
736 coherence level, ten repetitions per motion direction), and the vertical bars  
737 are the standard error and standard deviation for accuracy (Panel a) and RT  
738 (Panel b), respectively. As in Fig. 3 in Roitman and Shadlen (2002), we did  
739 not show RTs on error trials.

740 Our network performance is comparable to human accuracy, and it  
741 qualitatively emulates the effect of motion strength on RT. Decreasing RT for  
742 a relatively easy task (e.g., high motion coherence) is a direct consequence  
743 of the race model. Conversely, when the difficulty of a decision is high (e.g.,  
744 low coherence level), information favoring a particular response grows more  
745 slowly (P. L. Smith and Ratcliff 2004), and the probability of making an error  
746 is higher (Shadlen and Newsome 2001). The quantitative difference between  
747 behavioral and simulated RT in Fig. 7 could be eradicated by fine-tuning the  
748 excitatory weights from MT cells to the decision layer. However, such an  
749 exercise would be meaningless, because our model does not take into  
750 consideration neural areas involved in characteristics of the decision-making  
751 process that influence the length of RT, such as the time-course of LIP  
752 neuronal dynamics or the gating of saccadic eye movements (Shadlen and  
753 Newsome 2001), which have been successfully modeled in detail by others  
754 (Grossberg and Pilly 2008).

### 755 **3.4 Computational performance**

756 In order to compare our CUDA implementation of V1 (that is, the file  
757 v1colorME.cu) to the original, unmodified S&H implementation (which features

758code in both C and Matlab) we computed V1 complex cell responses (see  
759Section 2.2.1) at a single spatiotemporal scale to a drifting sinusoidal grating  
760(the same stimulus as described in Section 3.1) and recorded the model's  
761execution time. The S&H C/Matlab code was executed as  
762shModel(stim,pars,'v1Complex'), where stim was the input stimulus, and pars were  
763the default parameters (shPars). Fig. 8a shows the execution time per video  
764frame for both models. Our GPU implementation (red) was not only faster  
765(except for relatively small networks) than the S&H C/Matlab implementation  
766(blue), but it also scaled better with network size. Note that the C/Matlab  
767implementation was a single-threaded computation. The largest speedup, a  
768factor of 12, was observed for a network consisting of  $96 \times 96 = 9,216$   
769neurons. It is likely that even greater speedups could have been achieved on  
770larger networks, but these networks could not run with the S&H C/Matlab  
771implementation because they ran out of memory. Timing was performed  
772using standard commands tic and toc in Matlab, and the <ctime> function time  
773in C++/CUDA. For the S&H C/Matlab implementation, the time it took to  
774create the stimulus was not included in the time measurement. On the other  
775hand, in the CUDA implementation the stimulus had to be read from file  
776frame-by-frame and copied to the GPU card. However, we did not include the  
777time it takes to transfer the response back from the device to the host.

778        Additionally, the S&H C/Matlab implementation is memory-intensive  
779(see Fig. 8b), and execution times for networks above size  
780 $128 \times 128 = 16,384$  could not be computed because the CPU ran out of  
781memory, even though we had a relatively large amount of RAM (24 GB)  
782available. Measuring memory usage in Matlab is not straight-forward. In  
783order to demonstrate the excessive memory consumption of the S&H  
784C/Matlab implementation (see Fig. 8b) we opted to measure two metrics: the  
785size of the output argument ans to function call shModel (blue, filled circle in  
786Fig. 8b) and the maximum memory usage of the Matlab process at any point  
787in time (blue, open circle). The first was measured with native Matlab  
788command whos, and the latter was measured by running a bash script in the

789background that reported the memory usage of the process every second  
790(using linux command `ps`). The blue dashed line is the 24 GB limit of the  
791system's RAM. Note the log scale on the ordinate. Less memory was required  
792to run the process than to store the output argument, which consisted of a  
793matrix whose size was proportional to the product of the stimulus  
794dimensions and the number of frames. A straightforward way of making the  
795S&H C/Matlab implementation capable of handling large inputs would thus be  
796to break up the output argument into smaller chunks of data. On the other  
797hand, the memory usage of the GPU implementation was significantly lower  
798(red line in Fig. 8b) and scaled better with network size. We used CUDA  
799command `cuMemGetInfo` to identify the amount of allocated memory on the  
800GPU. The red dashed line is the upper limit of GPU memory available to the  
801user (roughly 5.2 GB on our card).

802 Comparing the performance between GPU simulation mode and CPU  
803simulation mode with the full network on the specific processor remains to  
804be demonstrated. Recall from Section 2.1.2 that in GPU mode all data  
805structures are allocated on the GPU, whereas in CPU mode the network  
806would be allocated on the CPU's memory, and only the generation of motion  
807energy responses (written in CUDA) would be delegated to the GPU. Hence  
808we evaluated the computational performance by running the full network in  
809both CPU and GPU mode with input images from  $16 \times 16$  pixels (38,784  
810neurons) to  $64 \times 64$  pixels (610,944 neurons). The simulation speed is given  
811as the ratio of execution time over the simulation time (see Fig. 9a) for  
812networks run in CPU mode (blue) and GPU mode (red). Note that in both  
813modes, the V1 CUDA implementation was executed (green), whose run-time  
814is part of the total simulation time (in blue and red). The GPU simulations not  
815only ran faster, but also simulation speed scaled better with network size.  
816Note that the CPU simulation was a single-threaded computation. The full  
817network at  $40 \times 40$  input resolution (239,040 neurons) ran in real-time on the  
818GPU. At  $32 \times 32$  input resolution (153,216 neurons) the simulation was 1.5  
819times faster than real-time. This result compares favorably with previous



820releases of our simulator (Nageswaran et al. 2009; Richert et al. 2011),  
821which is partly due to code-level optimizations, but mostly due to differences  
822in GPU hardware and the V1 stage of the network being spatiotemporal  
823filters instead of spiking neurons. As the network size increased, the GPU  
824simulations showed a significant speedup over the CPU (see Fig. 9b).  
825Speedup was computed as the ratio of CPU to GPU execution time. The  
826largest network we could fit on a single GPU roughly corresponded to  $64 \times 64$   
827input resolution (610,944 neurons), which ran approximately 30 times faster  
828than on the CPU. Larger networks currently do not fit on a single GPU and as  
829such must be run on the CPU, which would be more than 70 times slower  
830than real-time judging from Fig. 9a.

#### 8314. Discussion

832 We presented a large-scale spiking model of visual area MT that 1) is capable  
833 of exhibiting both component and pattern motion selectivity, 2) generates  
834 speed tuning curves that are in agreement with electrophysiological data, 3)  
835 reproduces behavioral responses from a 2AFC task, 4) outperforms a  
836 previous rate-based implementation of the motion energy model (Simoncelli  
837 and Heeger 1998) in terms of computational speed and memory usage, 5) is  
838 implemented on a publicly available SNN simulator that allows for real-time  
839 execution on off-the-shelf GPUs, and 6) is comprised of a neuron model,  
840 synapse model, and address-event representation (AER), which is compatible  
841 with recent neuromorphic hardware (Srinivasa and Cruz-Albrecht 2012).

842       The model is based on two previous models of motion processing in MT  
843 (Simoncelli and Heeger 1998; Rust et al. 2006), but differs from these  
844 models in several ways. First, our model contains the tuned normalization in  
845 the MT stage that was not present in Simoncelli and Heeger (1998) but  
846 introduced by Rust et al. (2006). Second, the implementation by Rust et  
847 al. (2006) was restricted to inputs that are mixtures of 12 sinusoidal gratings  
848 of a fixed spatial and temporal frequency, whereas our model can operate on  
849 any spatiotemporal image intensity. Third, MT PDS cells in our model sum  
850 over inputs from MT CDS cells as opposed to inputs from V1 cells, although  
851 the two approaches are conceptually equivalent. Fourth, instead of using  
852 linear summation and a static nonlinear transformation, all neuronal and  
853 synaptic dynamics in our model MT were achieved using Izhikevich spiking  
854 neurons and conductance-based synapses.

855       One could argue that the inclusion of Izhikevich spiking neurons and  
856 conductance-based synapses is unnecessary, since previous incarnations of  
857 the motion energy model did not feature these mechanisms yet were  
858 perfectly capable of reproducing speed tuning and motion selectivity.  
859 However, our approach is to be understood as a first step into modeling

860large-scale networks of visual motion processing in more biological detail,  
861with the ultimate goal of understanding how the brain solves the aperture  
862problem, among other open issues in motion perception. Integrating the  
863functionality demonstrated in previous models with more neurobiologically  
864plausible neuronal and synaptic dynamics is a necessary first step into  
865analyzing the temporal dynamics of model neurons in MT, which may 1) help  
866to explain how MT PDS cell establish their pattern selectivity not instantly but  
867over a time-course on the order of 100 ms (M. A. Smith et al. 2005) and 2)  
868enable the addition of spike-based learning rules such as STDP; both of  
869which might be harder to achieve with previous model incarnations.  
870Additionally, the introduction of the present neuron model, synapse model,  
871and address-event representation (AER) did not affect performance, yet  
872enabled the integration of the S&H model with recent neuromorphic  
873hardware (Srinivasa and Cruz-Albrecht 2012) (see also Section 4.3).

874 On the other hand, it is possible (if not likely) that some response  
875dynamics produced by the neural circuitry in the retina, the lateral  
876geniculate nucleus (LGN), and V1 may account for certain response  
877properties of neurons in MT. Thus future work could be directed towards  
878implementing the entire early visual system in the spiking domain. However,  
879for the purpose of this study we deem a rate-based preprocessor to be an  
880adequate abstraction, as the core functionality of directionally selective cells  
881in V1 seem to be well-characterized by local motion energy filters (Adelson  
882and Bergen 1985; DeAngelis et al. 1993; Movshon and Newsome 1996).

#### 883**4.1 Neurophysiological evidence and model alternatives**

884There is evidence that MT firing rates represent the velocity of moving  
885objects using the IOC principle. A psychophysical study showed that the  
886perception of moving plaids depends on conditions that specifically affect the  
887detection of individual grating velocities (Adelson and Movshon 1982). This is  
888consistent with a two-stage model in which component velocities are first  
889detected and then pooled to compute pattern velocity. Subsequent

890physiological studies broadly support such a cascade model (Perrone and  
891Thiele 2001; Rust et al. 2006; M. A. Smith et al. 2005).

892        However, other psychophysical results exist where the perceived  
893direction of plaid motion deviates significantly from the IOC direction (Ferrera  
894and Wilson 1990; Burke and Wenderoth 1993). Alternatives to the IOC  
895principle are, for example, vector average (VA) or feature tracking. VA  
896predicts that the perceived pattern motion is the vector average of the  
897component velocity vectors. Blob or feature tracking is the process of  
898locating something (a “feature”) that does not suffer from the aperture  
899problem, such as a bright spot or a T-junction, and tracking it over time  
900(Wilson et al. 1992). Ultimately, one needs to consider the interactions of the  
901motion pathway with form mechanisms (Majaj et al. 2007), and model the  
902processing of more complex stimuli (e.g., motion transparency, additional  
903self-motion, multiple moving objects) (Raudies et al. 2011; Layton et al.  
9042012). Clarifying by which rule (or combination of rules) the brain integrates  
905motion signals is still a field of ongoing research. For recent reviews on the  
906topic see (Bradley and Goyal 2008; Nishida 2011).

907        Although clear evidence for spatiotemporal frequency inseparability in  
908MT neurons has been found (Perrone and Thiele 2001), which supports the  
909idea of a motion energy model, later studies reported it to be a weak effect  
910(Priebe et al. 2003; Priebe et al. 2006). The actual proportion of neurons in  
911the primate visual system that are tuned to spatiotemporal frequency is  
912currently not known.

#### 913**4.2 Model limitations**

914Although our model is able to capture many attributes of motion selectivity  
915(e.g., direction selectivity, speed tuning, component and pattern motion), it  
916is not yet complete for the following reasons. First, it does not explicitly  
917specify the exact pattern velocity, but instead reports an activity distribution  
918over the population of MT neurons, whose firing rates are indicative of the  
919observed pattern motion. In order to estimate the speed of a target stimulus,

920it has been proposed to use a suitable population decoding mechanism that  
921operates on MT responses (Perrone 2012; Hohl et al. 2013). Second, our  
922model does not attempt to predict the temporal dynamics of MT PDS cells,  
923which often respond with broad selectivity when first activated, sometimes  
924even resembling CDS cells, and only over a time-course on the order of  
925100 ms establish their pattern motion selectivity (M. A. Smith et al. 2005). A  
926possible explanation for these temporal dynamics is given in Chey et al.  
927(1997). Third, it does not consider the visual form pathway and abstracts  
928early visual details that may be critical for operation in natural settings.  
929Fourth, the extent to which each stage in the motion energy model can be  
930mapped onto specific neuronal populations is rather limited. Tiling the  
931spatiotemporal frequency space according to the motion energy model is  
932biologically implausible, and the temporal extent of the filters is  
933unrealistically long (especially the low speed filters). However, a way to  
934combine spatiotemporal filters based on V1 neuron properties into a pattern  
935motion detector has been proposed in Perrone and Thiele (2002).

936       Another more fundamental limitation is that the S&H model (or for that  
937matter, any spatiotemporal-energy based model including the elaborated  
938Reichardt detector) can only sense so-called first-order motion, which is  
939defined as spatiotemporal variations in image intensity (first-order image  
940statistics) that give rise to a Fourier spectrum. Second-order stimuli, such as  
941the motion of a contrast modulation over a texture, are non-Fourier and thus  
942invisible to the model, yet can be readily perceived by humans (Chubb and  
943Sperling 1988). In addition, the existence of a third motion channel has been  
944suggested, which is supposed to operate through selective attention and  
945saliency maps (Lu and Sperling 1995). Also, MT has been shown to be  
946involved in color-based motion perception (Thiele et al. 2001).

947       There is also a plainly technical limitation to our model, which is  
948manifested in the amount of available GPU memory. Due to their size, large-  
949scale spiking networks have demanding memory requirements. The largest  
950network that could fit on a single NVIDIA Tesla M2090 (with 6 GB of memory)

951was comprised of 610,944 neurons and approximately 137 million synapses,  
952which corresponds to processing a  $64 \times 64$  input video. In order to run larger  
953networks on current-generation GPU cards, a change in model or (software  
954and hardware) architecture is required. One should note that this is only a  
955temporary limitation and could become obsolete as soon as with the next  
956generation of GPU cards. Another possible solution would be to employ multi-  
957GPU systems; however, more work is required to efficiently integrate our  
958SNN simulator with such a system.

### 959**4.3 Practical implications**

960The present network might be of interest to the neuroscientist and computer  
961vision research communities for the following reasons.

962       First, our implementation outperforms the S&H C/Matlab  
963implementation by orders of magnitude in terms of computational speed and  
964memory usage. Thus our CUDA implementation can be used to save  
965computation time, as well as be applied to input resolutions that the  
966C/Matlab implementation cannot handle due to memory constraints.  
967Additionally, the CUDA implementation can act as a stand-alone module that  
968could potentially be used in computer vision as an alternative to  
969computationally expensive operations such as Gabor filtering for edge  
970detection or dense optic flow computations.

971       Second, we have demonstrated that our approach is fast, efficient, and  
972scalable; although current GPU cards limit the size of the simulations due to  
973memory constraints. Nevertheless, our model processes a  $40 \times 40$  input  
974video at 20 frames per second in real-time, which corresponds to a total of  
975239,040 neurons in the simulated V1, MT, and LIP areas, at 20 frames per  
976second using a single GPU, which enables the potential use of our software in  
977real-time applications ranging from robot vision to autonomous driving.

978       Third, our implementation might be of particular interest to the  
979neuromorphic modeling community, as the present neuron model, synapse  
980model, and AER are compatible with recent neuromorphic hardware

981(Srinivasa and Cruz-Albrecht 2012). Thus our algorithm could be used as a  
982neural controller in neuromorphic and neurorobotics applications. Future  
983work could be directed toward creating an interface by which networks can  
984be automatically exported onto neuromorphic hardware.

985 Fourth, because of the modular code structure, our implementation  
986can be readily extended to include, for example, higher-order visual areas or  
987biologically plausible synaptic learning rules such as STDP. Thus our  
988implementation may facilitate the testing of hypotheses and the study of the  
989temporal dynamics that govern visual motion processes in area MT, which  
990might prove harder to study using previous (rate-based) model incarnations.

991 Lastly, the network was constructed using a SNN simulator that is  
992publicly available at <http://www.socsci.uci.edu/~jkrichma/CARLsim/>. The  
993present release features the complete source code for the simulator, the  
994network, and analysis scripts. As such it is the next step towards our goal of  
995making efficient simulations of large-scale spiking networks available to a  
996wide range of researchers, without the need of a cluster or supercomputer.

### 9975. Information Sharing Statement

998The source code for the simulator, for the network, and analysis scripts are  
999publicly available at <http://www.socsci.uci.edu/~jkrichma/CARLsim/>. This  
1000website does also feature installation instructions, source code  
1001documentation and a tutorial on how to set up, run, and interact with a  
1002simulation. In order to run the simulator in CUDA mode, the NVIDIA CUDA  
1003software developer kit must be installed (freeware, available at  
1004<https://developer.nvidia.com/cuda-downloads>).

1005**6. Acknowledgments**

1006This work was supported by the Defense Advanced Research Projects Agency  
1007(DARPA) subcontract 801888-BS. We thank Jayram M. Nageswaran for his  
1008work developing the custom spiking neural network simulator. We also thank  
1009Michael Avery, Kris Carlson, and Steve Grossberg for valuable feedback and  
1010discussion on this project.

1011



## 10127. **References**

1013Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the  
1014perception of motion. *J Opt Soc Am A*, 2(2), 284-299.

1015Adelson, E. H., & Movshon, J. A. (1982). Phenomenal coherence of moving  
1016visual patterns. *Nature*, 300(5892), 523-525.

1017Bradley, D. C., & Goyal, M. S. (2008). Velocity computation in the primate  
1018visual system. *Nature Reviews Neuroscience*, 9(9), 686-695, doi:Doi 10.1038/  
1019Nrn2472.

1020Browning, N. A., Grossberg, S., & Mingolla, E. (2009a). Cortical dynamics of  
1021navigation and steering in natural scenes: Motion-based object  
1022segmentation, heading, and obstacle avoidance. *Neural Networks*, 22(10),  
10231383-1398, doi:DOI 10.1016/j.neunet.2009.05.007.

1024Browning, N. A., Grossberg, S., & Mingolla, E. (2009b). A neural model of how  
1025the brain computes heading from optic flow in realistic scenes. *Cogn Psychol*,  
102659(4), 320-356, doi:10.1016/j.cogpsych.2009.07.002.

1027Burke, D., & Wenderoth, P. (1993). The Effect of Interactions between One-  
1028Dimensional Component Gratings on 2-Dimensional Motion Perception.  
1029*Vision Research*, 33(3), 343-350, doi:Doi 10.1016/0042-6989(93)90090-J.

1030Chey, J., Grossberg, S., & Mingolla, E. (1997). Neural dynamics of motion  
1031grouping: from aperture ambiguity to object speed and direction. *Journal of*  
1032*the Optical Society of America a-Optics Image Science and Vision*, 14(10),  
10332570-2594, doi:Doi 10.1364/Josaa.14.002570.

1034Chubb, C., & Sperling, G. (1988). Drift-Balanced Random Stimuli - a General  
1035Basis for Studying Non-Fourier Motion Perception. *Journal of the Optical*  
1036*Society of America a-Optics Image Science and Vision*, 5(11), 1986-2007,  
1037doi:Doi 10.1364/Josaa.5.001986.

1038Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience : computational*  
1039*and mathematical modeling of neural systems* (Computational  
1040neuroscience). Cambridge, Mass.: Massachusetts Institute of Technology  
1041Press.

1042DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1993). Spatiotemporal  
1043organization of simple-cell receptive fields in the cat's striate cortex. II.  
1044Linearity of temporal and spatial summation. *Journal of Neurophysiology*,  
104569(4), 1118-1135.

1046Ferrera, V. P., & Wilson, H. R. (1990). Perceived direction of moving two-  
1047dimensional patterns. *Vision Research*, 30(2), 273-287.

1048Fidjeland, A. K., & Shanahan, M. P. Accelerated simulation of spiking neural  
1049networks using GPUs. In *Neural Networks (IJCNN), The 2010 International*  
1050*Joint Conference on, 18-23 July 2010 2010* (pp. 1-8).  
1051doi:10.1109/IJCNN.2010.5596678.

1052Freeman, W. T., & Adelson, E. H. The design and use of steerable filters. In  
1053*IEEE Pattern Analysis and Machine Intelligence, 1991* (Vol. 13, pp. 891-906)

1054Grossberg, S., & Pilly, P. K. (2008). Temporal dynamics of decision-making  
1055during motion perception in the visual cortex. *Vision Research*, 48(12), 1345-  
10561373, doi:DOI 10.1016/j.visres.2008.02.019.

1057Hohl, S. S., Chaisanguanthum, K. S., & Lisberger, S. G. (2013). Sensory  
1058population decoding for visually guided movements. *Neuron*, 79(1), 167-179,  
1059doi:10.1016/j.neuron.2013.05.026.

1060Indiveri, G., Chicca, E., & Douglas, R. (2006). A VLSI array of low-power  
1061spiking neurons and bistable synapses with spike-timing dependent  
1062plasticity. *Ieee Transactions on Neural Networks*, 17(1), 211-221, doi:Doi  
106310.1109/Tnn.2005.860850.

1064Izhikevich, E. M. (2003). Simple model of spiking neurons. *Ieee Transactions*  
1065*on Neural Networks*, 14(6), 1569-1572, doi:Doi 10.1109/Tnn.2003.820440.

1066Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons?  
1067*Ieee Transactions on Neural Networks*, 15(5), 1063-1070, doi:Doi  
106810.1109/Tnn.2004.832719.

1069Izhikevich, E. M. (2007). *Dynamical systems in neuroscience : the geometry*  
1070*of excitability and bursting* (Computational neuroscience). Cambridge, Mass.:  
1071MIT Press.

1072Izhikevich, E. M., Gally, J. A., & Edelman, G. M. (2004). Spike-timing dynamics  
1073of neuronal groups. *Cereb Cortex*, *14*(8), 933-944, doi:DOI  
107410.1093/cercor/bhh053.

1075Khan, M., Lester, D., Plana, L., Rast, A., Jin, X., & Painkras, E. SpiNNaker:  
1076Mapping neural networks onto a massively-parallel chip multiprocessor. In  
1077*IEEE International Joint Conference on Neural Networks, 2008* (pp. 2849-  
10782856)

1079Koch, C. (1999). *Biophysics of computation : information processing in single*  
1080*neurons* (Computational neuroscience). New York: Oxford University Press.

1081Layton, O. W., Mingolla, E., & Browning, N. A. (2012). A motion pooling model  
1082of visually guided navigation explains human behavior in the presence of  
1083independently moving objects. *J Vis*, *12*(1), doi:10.1167/12.1.20.

1084Livingstone, M. S., & Conway, B. R. (2007). Contrast affects speed tuning,  
1085space-time slant, and receptive-field organization of simple cells in macaque  
1086V1. *Journal of Neurophysiology*, *97*(1), 849-857, doi:10.1152/jn.00762.2006.

1087Lu, Z. L., & Sperling, G. (1995). Attention-Generated Apparent Motion.  
1088*Nature*, *377*(6546), 237-239, doi:Doi 10.1038/377237a0.

1089Majaj, N. J., Carandini, M., & Movshon, J. A. (2007). Motion integration by  
1090neurons in macaque MT is local, not global. *Journal of Neuroscience*, *27*(2),  
1091366-370, doi:10.1523/JNEUROSCI.3183-06.2007.

1092Merolla, P. A., Arthur, J. V., Shi, B. E., & Boahen, K. A. (2007). Expandable  
1093networks for neuromorphic chips. *Ieee Transactions on Circuits and Systems*  
1094*I-Regular Papers*, *54*(2), 301-311, doi:Doi 10.1109/Tcsi.2006.887474.

1095Movshon, J. A., Adelson, E. H., Gizzi, M. S., & Newsome, W. T. (1985). *The*  
1096*analysis of moving visual patterns* (Pattern recognition mechanisms). New  
1097York: Springer.

1098Movshon, J. A., & Newsome, W. T. (1996). Visual response properties of  
1099striate cortical neurons projecting to area MT in macaque monkeys. *Journal*  
1100*of Neuroscience*, *16*(23), 7733-7741.

1101Nageswaran, J. M., Dutt, N., Krichmar, J. L., Nicolau, A., & Veidenbaum, A. V.  
1102(2009). A configurable simulation environment for the efficient simulation of

1103large-scale spiking neural networks on graphics processors. *Neural*  
1104*Networks*, 22(5-6), 791-800, doi:DOI 10.1016/j.neunet.2009.06.028.

1105Nishida, S. (2011). Advancement of motion psychophysics: Review 2001-  
11062010. *Journal of Vision*, 11(5), doi:Artn 11  
1107Doi 10.1167/11.5.11.

1108Pack, C. C., Berezovskii, V. K., & Born, R. T. (2001). Dynamic properties of  
1109neurons in cortical area MT in alert and anaesthetized macaque monkeys.  
1110*Nature*, 414(6866), 905-908, doi:10.1038/414905a.

1111Perrone, J. A. (2012). A neural-based code for computing image velocity from  
1112small sets of middle temporal (MT/V5) neuron inputs. *J Vis*, 12(8),  
1113doi:10.1167/12.8.1.

1114Perrone, J. A., & Thiele, A. (2001). Speed skills: measuring the visual speed  
1115analyzing properties of primate MT neurons. *Nat Neurosci*, 4(5), 526-532.

1116Perrone, J. A., & Thiele, A. (2002). A model of speed tuning in MT neurons.  
1117*Vision Research*, 42(8), 1035-1051.

1118Priebe, N. J., Cassanello, C. R., & Lisberger, S. G. (2003). The neural  
1119representation of speed in macaque area MT/V5. *Journal of Neuroscience*,  
112023(13), 5650-5661.

1121Priebe, N. J., Lisberger, S. G., & Movshon, J. A. (2006). Tuning for  
1122spatiotemporal frequency and speed in directionally selective neurons of  
1123macaque striate cortex. *Journal of Neuroscience*, 26(11), 2941-2950,  
1124doi:10.1523/JNEUROSCI.3936-05.2006.

1125Raudies, F., Mingolla, E., & Neumann, H. (2011). A model of motion  
1126transparency processing with local center-surround interactions and  
1127feedback. *Neural Comput*, 23(11), 2868-2914, doi:10.1162/NECO\_a\_00193.

1128Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of  
1129mind in decision-making. *Nature*, 461(7261), 263-U141, doi:Doi  
113010.1038/Nature08275.

1131Richert, M., Nageswaran, J. M., Dutt, N., & Krichmar, J. L. (2011). An efficient  
1132simulation environment for modeling large-scale cortical processing. *Front*  
1133*Neuroinform*, 5, 19, doi:10.3389/fninf.2011.00019.

1134Rodman, H. R., & Albright, T. D. (1987). Coding of Visual Stimulus Velocity in  
1135Area Mt of the Macaque. *Vision Research*, 27(12), 2035-2048, doi:Doi  
113610.1016/0042-6989(87)90118-0.

1137Rodman, H. R., & Albright, T. D. (1989). Single-unit analysis of pattern-  
1138motion selective properties in the middle temporal visual area (MT). *Exp*  
1139*Brain Res*, 75(1), 53-64.

1140Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral  
1141intraparietal area during a combined visual discrimination reaction time task.  
1142*Journal of Neuroscience*, 22(21), 9475-9489.

1143Rust, N. C., Mante, V., Simoncelli, E. P., & Movshon, J. A. (2006). How MT cells  
1144analyze the motion of visual patterns. *Nat Neurosci*, 9(11), 1421-1431,  
1145doi:Doi 10.1038/Nn1786.

1146Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual  
1147decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of*  
1148*Neurophysiology*, 86(4), 1916-1936.

1149Simoncelli, E. P., & Heeger, D. J. (1998). A model of neuronal responses in  
1150visual area MT. *Vision Research*, 38(5), 743-761, doi:Doi 10.1016/S0042-  
11516989(97)00183-1.

1152Smith, M. A., Majaj, N. J., & Movshon, J. A. (2005). Dynamics of motion  
1153signaling by neurons in macaque area MT. *Nat Neurosci*, 8(2), 220-228,  
1154doi:Doi 10.1038/Nn1382.

1155Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple  
1156decisions. *Trends in Neurosciences*, 27(3), 161-168, doi:DOI  
115710.1016/j.tins.2004.01.006.

1158Srinivasa, N., & Cruz-Albrecht, J. M. (2012). Neuromorphic Adaptive Plastic  
1159Scalable Electronics Analog Learning Systems. *Ieee Pulse*, 3(1), 51-56,  
1160doi:Doi 10.1109/Mpul.2011.2175639.

1161Thiele, A., Dobkins, K. R., & Albright, T. D. (2001). Neural correlates of  
1162chromatic motion perception. *Neuron*, 32(2), 351-358.

1163van Santen, J. P. H., & Sperling, G. (1985). Elaborated Reichardt Detectors.  
1164*Journal of the Optical Society of America a-Optics Image Science and Vision*,  
11652(2), 300-321.

1166Vogelstein, R. J., Mallik, U., Culurciello, E., Cauwenberghs, G., & Etienne-  
1167Cummings, R. (2007). A multichip neuromorphic system for spike-based  
1168visual information processing. *Neural Comput*, 19(9), 2281-2300, doi:DOI  
116910.1162/neco.2007.19.9.2281.

1170Wilson, H. R., Ferrera, V. P., & Yo, C. (1992). A Psychophysically Motivated  
1171Model for 2-Dimensional Motion Perception. *Visual Neuroscience*, 9(1), 79-97.

1172Yudanov, D., Shaaban, M., Melton, R., & Reznik, L. GPU-based simulation of  
1173spiking neural networks with real-time performance & high accuracy. In  
1174*Neural Networks (IJCNN), The 2010 International Joint Conference on, 18-23*  
1175*July 2010 2010* (pp. 1-8). doi:10.1109/IJCNN.2010.5596334.

1176

1177

## 11788. **Figure captions**

1179**Fig. 1** Network architecture.  $32 \times 32$  grayscale images are fed through model  
1180V1, MT, and LIP (as explained in Sections 2.2.1 – 2.2.3). Shown is a snapshot  
1181in time of the network's response to an example RDK stimulus in which 50 %  
1182of the dots drift to the right. Black bold arrows denote synaptic projections.  
1183Inhibitory projections and populations are not shown. Numbers in  
1184parentheses next to an element are the equations that describe the  
1185corresponding neuronal response or synaptic projections (see text). V1 filter  
1186responses were mapped onto mean firing rates by reproducing the contrast  
1187sensitivity function reported for V1 cells projecting to MT, as explained in  
1188Section 2.2.1

1189**Fig. 2** A drifting dot traces out a path (dashed line) in space ( $x$ , ignoring  $y$ )  
1190and time ( $t$ ). The colored ovals correspond to the orientation of the positive  
1191(green) and negative (red) lobes of a spatiotemporal filter **a** If the filter is  
1192oriented in the same way as the dot's space-time path it could be activated  
1193by this motion **b** A dot moving in the opposite direction would always contact  
1194both positive and negative lobes of the filter and therefore could never  
1195produce a strong response. Adopted from (Bradley and Goyal 2008)

1196**Fig. 3** The contrast sensitivity function of model V1 simple cells (blue) is  
1197plotted against electrophysiological data adapted from Fig. 7 of (Movshon  
1198and Newsome 1996). Each data point is a V1 mean response to a drifting  
1199grating, averaged over both one second of stimulus presentation and all  
1200neurons in the subpopulation. Vertical bars are the standard deviation on the  
1201population average

1202**Fig. 4** Polar plots of direction tuning for a sinusoidal grating **a-d** and a plaid  
1203stimulus **e-h** drifting upwards, where the angle denotes motion direction and  
1204the radius is the firing rate in spikes per second. Tuning curves were  
1205obtained by taking the mean firing rate of a neuron to a drifting grating  
1206during two seconds of stimulus presentation, averaged over all neurons in

1207the population selective to the same stimulus direction (black: mean  
1208neuronal response, blue: mean plus standard deviation on the population  
1209average, green: mean minus standard deviation). Shown are mean  
1210responses for V1 complex cells (**b** and **f**), MT CDS cells (**c** and **g**), and MT PDS  
1211cells (**d** and **h**). Only MT PDS cells **h** responded to the motion of the entire  
1212plaid pattern rather than to the motions of the individual component gratings  
1213**Fig. 5** The pattern index is computed for all MT CDS cells (blue) and all MT  
1214PDS cells (red), and plotted as a Fisher Z-score. The black solid lines are the  
1215classification region boundaries, indicating that all MT CDS cells have indeed  
1216been classified as component-selective, and all MT PDS cells have been  
1217classified as pattern-selective

1218**Fig. 6** Speed tuning curves for three different classes of MT neurons. The  
1219stimulus consisted of a single bar drifting over the entire visual field either to  
1220the right (preferred direction) or to the left (anti-preferred direction) at  
1221different speeds **a** Response of a “speed-tuned” neuron (selective to motion  
1222at 1.5 pixels per frame) **b** Response of a “low-pass” neuron (selective to  
1223motion at 0.125 pixels per frame) **c** Response of a “high-pass” neuron  
1224(selective to motion at 9 pixels per frame)

1225**Fig. 7** Random dot kinematogram. The RDK stimulus was constructed out of  
1226approximately 150 dots (15 % dot density, maximum stimulus contrast) on a  
122732x32 input movie **a** Psychometric function. The network’s accuracy  
1228increased with increasing motion strength (coherence level) **b** Chronometric  
1229function. The network’s RT decreased with increasing motion strength

1230**Fig. 8 a** Execution time of a Matlab implementation (blue) of V1 complex  
1231cells versus a CUDA implementation (red) **b** Observed memory usage for the  
1232Matlab implementation (blue) and CUDA implementation (red)

1233**Fig. 9 a** Simulation speed is given as the ratio of execution time over the  
1234simulation time for networks run in CPU mode (blue) and GPU mode (red). In  
1235both cases, the V1 CUDA implementation was executed (green), which is  
1236part of the total simulation time (in blue and red). Note the log scale on the



1237ordinate. The GPU simulations did not only run faster, but simulation speed  
1238scaled better with network size **b** Speedup is given as the ratio of CPU  
1239execution time over GPU execution time