

Complex object representation in a computer simulation of the neocortex

Sven Schrader, Rüdiger Kupper, Marc-Oliver Gewaltig, Ursula Körner, Edgar Körner

2008

Preprint:

This is an accepted article published in 5th HRI Global Workshop. The final authenticated version is available online at: [https://doi.org/\[DOI not available\]](https://doi.org/[DOI not available])

Complex object representation in a computer simulation of the neocortex

Sven Schrader, Rüdiger Kupper, Marc-Oliver Gewaltig, Ursula Körner, Edgar Körner

Abstract—In vision, an object representation is constructed through a cascade of areas, each processing more and more elaborate features of the stimulus image. We hypothesize that the necessary signal flow involves a feed-forward path to compose the object representation as well as a feedback path to refine the representation and to correct errors. Using word recognition as a metaphor for visual stimuli, we have implemented this architecture in a spiking neural network. Given a string of characters, the network quickly generates an initial stimulus hypothesis after only 35 ms. Already the first wave of action potentials traveling upwards the hierarchy contains the correct local symbols (syllables, word) of the presented stimulus primitives. In most cases, the activation of the correct word precedes all other candidates with millisecond precision and is most pronounced. Thus, our model codes the correctness of a response in its relative latency. In a subsequent refinement stage, where high-level activity *modulates* lower stages, this activation dominance is propagated back, influencing its own afferent activity to exhibit a unique decision.

I. INTRODUCTION

Despite the fact that we are provided with increasing detail about neocortical connectivity (1; 2; 3; 4; 5), the principles of higher brain functions are still poorly understood. Structures that have been suggested as computational units are often based on Hebb’s widely accepted idea of *neuronal assemblies* (6) where information is processed by sets of neurons rendering functional, rather than anatomical, relations. Theoretical models, which are not subject to experimental constraints, have become important tools to investigate such principles. They are often simplified to allow analytic treatment (7; 8; 9; 10). In particular, *random* connection schemes often mimic the seemingly unspecific nature of connections in the cortex (8; 11; 12). However, the flow of information is bound to follow the structured connections within and between brain areas. To address this issue, we studied the dynamics of a large network of integrate-and-fire neurons (13) whose architecture mimics the most prominent features of cortical connectivity, its areal division and its horizontal and vertical organization. In this model, we study the integration of sensory stimuli and the stages of visual processing in the ventral stream, the so-called “what”-pathway. We investigated the potential of our

network to recognize and neurally represent (14) complex objects. Using the hierarchical composition of words, we reduce the complexity of the visual modality to a one-dimensional stimulus domain. The goal of our study is to demonstrate how the hierarchy of the network (reflecting the hierarchy of the modality) creates *compositionality*, a unique high-level representation based on the application of basic *primitives*. At this juncture, local feedback onto lower stages is of particular functional importance. Large parts of lateral connections between areas are modulatory (15). Using modulatory feedback, we can show that initial activations refine themselves iteratively by influencing their own input.

A. Columns and layers

Besides the division into several layers, large parts of the brain are characterized by a vertical organization, prompting the notion of the cortical column (16; 17; 18). Connections may span layers vertically but are laterally confined to a particular volume. Lateral connections between columns can occur but are most pronounced in layers III and V (19). The smallest of such units is the *minicolumn* (20; 21). Following earlier work (22), we hypothesize that neocortical computation is carried out by elementary modules, the minicolumns, which are repeated over the cortical surface and are general ‘processing units’ indifferent to the task at hand. Within this framework, the set of minicolumns that share the same receptive field is defined as one *macrocolumn*. We propose that sensory integration, e.g. along the visual pathway, results from the massively parallel action of those units, which per se feature a stereotypical sequence of processing: Across its layers, a macrocolumn splits recognition into subsequent stages according to a fast “local hypothesis”, shaped in layers IV and lower III which we call “A-system”, (see 22, for details) and a slower refinement phase in layers II and upper III, the “B-system”. Finally, activity in layers V and VI, the “C-System”, is the neural representation of an object. Downstream areas or efferent instances (e.g. motor output) receive this activity for further processing or generation of behavior. Extending a study using analog neurons (23), the present paper seeks to elucidate how the above architecture forms dynamic representations of composite entities with biologically realistic neurons. We describe the *neoCOREtext* model, a layered cortical model that demonstrates the formation of a fast initial stimulus hypothesis in the columns and its subsequent refinement by inter-columnar communication. We simulate the signal flow in the A- and B- systems of a set of model columns across three hierarchical cortical areas.

S. Schrader is with the Honda Research Institute Europe, email:sven.schrader@honda-ri.de

R.Kupper is with the Honda Research Institute Europe, email:ruediger.kupper@honda-ri.de

Marc-Oliver Gewaltig is with the Honda Research Institute Europe, email:marc-oliver.gewaltig@honda-ri.de

U. Körner is with the Honda Research Institute Europe, email:ursula.koerner@honda-ri.de

E. Körner is with the Honda Research Institute Europe, email:edgar.koerner@honda-ri.de

B. Latency code

Usually, post-synaptic potentials are small as compared to the distance to firing threshold (24; 25). Successful firing can therefore be reached only by a high input rate, where the neuron responds with a large variability, or by a coherent input of many neurons, which leads to a more precise firing behavior. Whether such *temporal code* is functionally used in the brain is still under debate (26; 27; 28). Spikes occurring with millisecond precision, however, point into the direction that information can indeed be found in the precise timing of action potentials (29; 30; 31; 32). Our network suggests how time can be employed as a computational variable. Although no effort has been made to explicitly incorporate or extract information using spike patterns, the network reaches a state where spike timing plays a critical role. Responses appear in the form of narrow spike volleys whose occurrence times encode its correctness. A perfect match will yield the fastest response, while incomplete matches yield successively more delayed responses. Thus, from the receiver’s point of view, the *relative latency* of a signalled local decision from its input neurons encodes directly the probability of being true (truth value). Relative latencies are a fundamental property of our model since it is the mechanism that separates correct decisions from wrong ones. We show how this separation can be further improved using several feed-forward and feedback pathways.

II. METHODS

We simulated our model with the simulation tool NEST (33), using a newly developed Python interface (34). Simulations were run on a machine with two AMD Dual Core Opteron 280 processors at 2.4 Ghz. A typical simulation took about 1.5 seconds per millisecond simulated time. Offline data analysis and display programs were implemented in Python, using the scientific packages numpy and pylab.

III. MODEL

A. Network Architecture

Without loss of generality, we have chosen a simplified stimulus environment in which complex objects are composed from smaller parts. We realize the hierarchical construction of words from syllables and letters and implement a word recognition task. The neoCOREtext model consists of three areas, “V1”, “V2” and “IT”, representing consecutive steps in visual processing. Each area is composed of three layers in which macrocolumns are arranged retinotopically. The network scheme is shown in Figure 1. Within one macrocolumn, the three layers A1, A2 and B perform the guessing-refinement stages. The minicolumns represent the local symbols at each retinal position. This is done by populations of 65 neurons. The connections are explained in Figure 2. Gray boxes symbolize the excitatory (rectangles, $N=50$) and inhibitory (squares, $N=15$) neural sub-pools. The number of connections is shown in Table III in the Appendix.

In this framework, we treat *letters* as primitives (such as edges in vision) to be processed in V1, syllables as

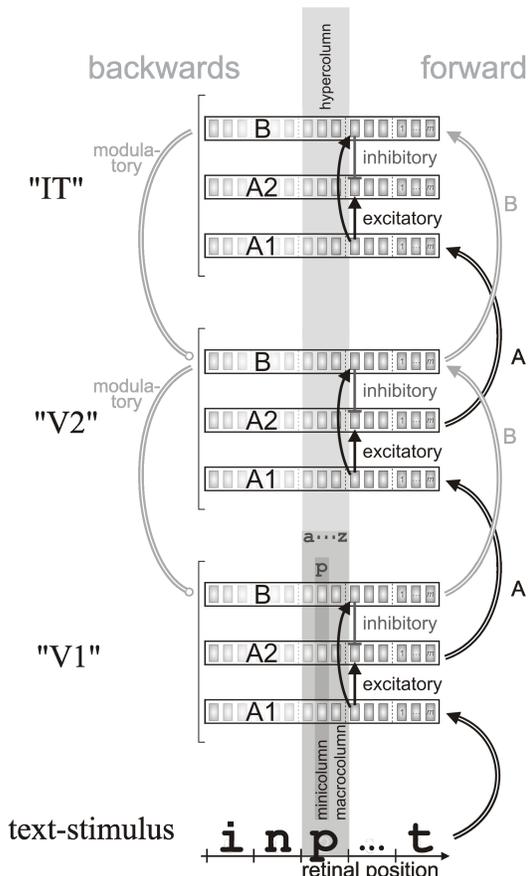


Fig. 1. The neoCOREtext model. Shown from bottom to top are the different levels of network architecture in the system. The three **areas** constitute the subsequently higher order in word processing, letters (V1), syllables (V2) and words (IT). Their feed-forward (double arrows) and feedback (round-headed double arrows) connections manifest learned knowledge about word compositionality (see Figure 2) The feedforward connections are divided into two pathways, one which connects A-systems (‘A-pathway’) and one between B-systems (‘B-pathway’) Within each area, the three cortical **layers** form the A-B system of local symbol recognition. Their connectivity include feed-forward excitation and feedback inhibition. The retinal arrangement is carried out by the **macrocolumns**, where any local symbol (words, syllables, letters) can be represented at any position of the array in the **minicolumn**. Within each layer, a gray box denotes a small neural pool representing a single local symbol, m denotes the quantity of the local alphabets. Arrows are further explained in Figure 2.

intermediate features in V2 and words as high-level features in IT. The stimulus enters the network in V1. Its task is to evolve a dynamic state which activates the macrocolumn in layer B of area IT, representing the stimulus word. Figure 2e demonstrates how inter-area connections are constructed according to the word composition into syllables and letters. Local symbols converge onto the next higher instance. 500 English words are stored by connecting the underlying syllables (531 in total) and letters (26) in a feed-forward manner. Modulatory feedback follows the opposite direction, shaping *divergent* connectivity to lower areas. The retinal width in our model is six symbols, Figures 1 and 2e show a width of five, for clarity.

The input to the network consists of two noise sources

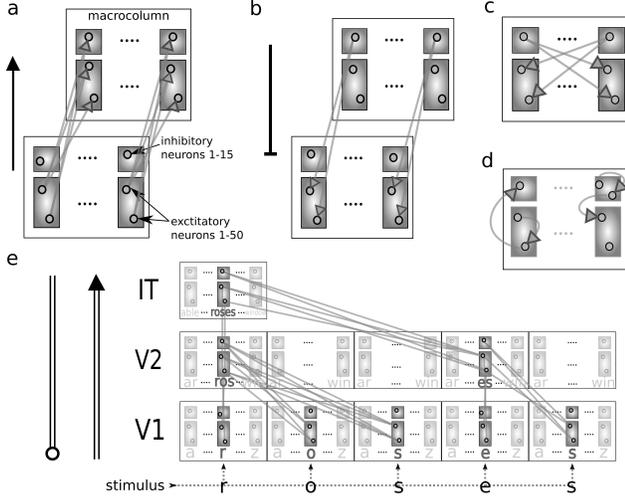


Fig. 2. Connections in the neoCOREtext system. (a) feedforward connections. Each neuron from the excitatory source population project to all neurons of the source pool. (b) inhibitory feedback. Interneurons project back to the excitatory target population. (c) lateral inhibition in IT-B. Each symbols inhibits all other symbols. (d) recurrent connectivity is implemented within and across excitatory and interneurons, respectively. (e) Connectivity scheme constituting word knowledge. Neural pools that represent symbols at particular positions on the retina are connected according to the word composition into syllables and letters. Shown is the connectivity for the word *roses*. Two minicolumns are connected in a divergent-convergent manner (gray lines), from the excitatory pool of one level to all neurons in the next level. The synapses project uni-directionally to higher areas (double-lined arrow) and back to lower areas with modulatory synapses (round-headed arrow). All words are stored in this way, multiple connections are excluded. In addition, every symbol *inhibits* six neighboring symbols from the next level in a center-surround fashion (three to the left and three to the right), leading to a local feed-forward inhibition (not shown).

emitting spikes at a constant rate. In both cases, spike intervals are randomly drawn such that the number of spikes per unit time is Poisson distributed (Poisson process). The stimulus is applied to all neurons in layer A1 of area V1 representing the letters at their positions. For instance, if the word *roses* is stimulated, the stimulus is applied to the A1-minicolumns for *r* at the first position, the one for *o* at the second, and so on (see Figure 2 e, dashed arrows). In addition to the letter stimuli, the entire network receives a constant noise mimicking the spiking activity of the surrounding network (see Appendix for details).

Altogether, the network consists of 18 macrocolumns (three areas with 6 retinal positions each). Depending on the number of local symbols, each macrocolumn comprises between 5,000 and 100,000 neurons. Each of the 3,842 minicolumns has 196 neurons (minicolumns beyond retina position 1 in IT are spared, since they are not required). Thus, the entire network comprises about 750,000 neurons and 60,000,000 synaptic connections.

B. Types of inhibition

In each minicolumn, we have implemented a fraction of inhibitory interneurons comprising 30% of all neurons. Inhibition can be reached directly, or by exciting the interneurons of the target population. The latter indirect inhibition occurs

across areas because they are not connected via inhibitory fibers (5). Feedback (B→A2), lateral and recurrent inhibition are realized directly (Figure 2 b, c and d). In addition to the feedforward connections across areas (Figure 2 e), every symbol *inhibits* six neighboring symbols from the next level in a center-surround fashion, leading to a local feed-forward inhibition. For simplicity, the features are ordered alphabetically rather than by similarity. Due to memory limitations, this lateral inhibition is currently only implemented in Layer B of area IT.

C. Neuron Model

The dynamics of each neuron is described by the integrate-and-fire model (13)

$$\tau_m \frac{dV(t)}{dt} = -V(t) + V_0 + R_m(E(t) + I(t)), \quad (1)$$

where τ_m denotes the membrane time constant, R_m the resistance and E and I the sums of all incoming excitatory and inhibitory synaptic currents, respectively. When V reaches a fixed threshold V_θ , a spike is emitted and the membrane potential is reset to the resting potential V_0 for the time of the refractory period t_{ref} . Each incoming spike elicits synaptic currents E and I which follow an α -function $\alpha(t) = J \frac{t}{\tau_{syn}} \exp(-\frac{t}{\tau_{syn}})$ (35) after a fixed delay δ ,

$$E, I(t) = \sum_j \sum_{i, t_i^j \leq t} \alpha(t - t_i^j - \delta), \quad (2)$$

where t_i^j represents the time of occurrence of the i th spike from the j th presynaptic excitatory (inhibitory) neuron and J the amplitude of the alpha-function (Greek letters in Table III). Modulatory action potentials to a lower B system (round headed arrows in Figures 1 and 2 e), enhance the synaptic weight amplitude J by a factor that decays exponentially to one, $J(t) = J_0 \times (1 + \mu(t))$, where

$$\mu(t) = f_{mod} \sum_j \sum_{i, t_i^j \leq t} \exp\left(\frac{-t - t_i^j - \delta}{\tau_{mod}}\right), \quad (3)$$

and t_i^j are the spike times from the next higher B system. In all other layers, synaptic weights are constant, $J(t) = J_0$. Further details and parameters are given in the appendix.

IV. RESULTS

A. Initial hypothesis

We first consider the early (<50 ms) stimulus response of the network. Figure 3 shows the responses of areas V1 and IT after stimulation of the letters *g,a,r,d,e,n*. The symbols in area V1 respond with a higher activity until stimulation is switched of at 150 ms (Figure 3 b). Across layers, the network activity becomes more structured. The phasic input creates pulse-like patterns that feed onto the next area. The main output layer of area V1 (A2) emits almost only pulses at around 30 ms.

This trend continues in area IT, where symbol activity is entirely composed of pulses (Figure 3 c). Spike volleys

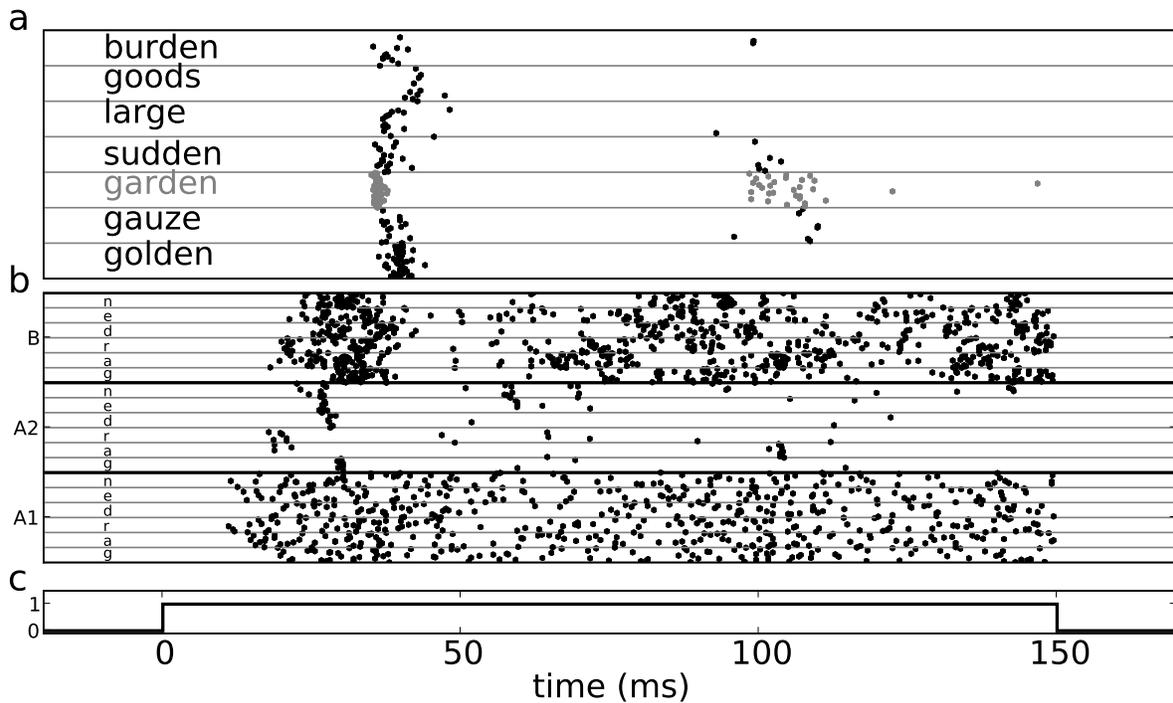


Fig. 3. Recognition of a word. (a) the stimulus to V1 starts at 0 ms and lasts for 150 ms. Background noise starts at -100 ms. (b) rasterplot of the spike activities in area V1 during stimulation of the letters forming the word *garden*. Each row (separated by horizontal lines) stands for the neurons representing the highlighted symbol. Shown are only symbols with strong responses, other symbols are almost silent.

appear at around 35 ms for the word that was stimulated, *garden*, along with words that share a high overlap. We refer to this set of symbol responses as *initial candidate list*. It comprises about one percent of the entire set of 500 words. Furthermore, Figure 3 indicates that the stimulus word is activated ahead of the other candidates and its activity is both strongest (in terms of number of spikes) and temporally most confined. Thus the correct word dominates the other candidates in terms of speed, strength and precision.

To assess how the network generates its initial hypothesis on average, we stimulated with 104 polysyllabic words (their letters in V1). For each stimulus, the correct word and a small list of candidates is activated in IT-B at around 35 ms. As in the example of Figure 3, the initial candidates comprise about 1% of all learned words. The responses are summarized in Figure 4. Response pulses are shown in the two-dimensional space, defined by their times of occurrence and their temporal spread. Generally, the correct words appear earlier and are stronger and more precise than their alternatives. The average lead of the correct word is 1.8 ms. In 46 cases, the correct word fails to lead other candidates, but still is the strongest symbol, except in a single trial. In 18% of the trials, the correct response is the third to appear or later.

B. Symbol refinement

If the stimulus persists for a longer period, response pulses reoccur, representing a second set of word-candidates at around 100 ms. Thus, area IT responds with waves of pulses that are 60 ms apart. Without feedback, the second

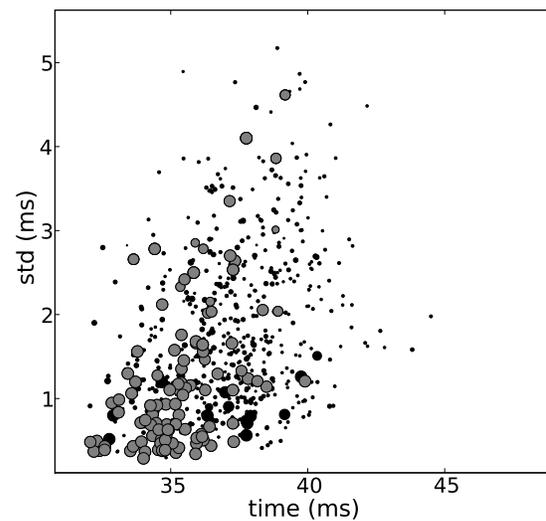


Fig. 4. Symbol activity in IT-B. Shown are the occurrence-times (abscissae) and precisions (ordinate) of word-representing pulses for 104 recognition trials. Each point denotes a different word occurrence. The strengths of the pulses (number of spikes) is denoted in the point diameter. Parameters are estimated by fitting a Gaussian rate profile on the individual pulses. The temporal spread corresponds to the standard deviation. Correct words are emphasized in gray, false-positives are black. Time is shown relative to stimulus onset.

wave contained approximately the same words as before (not shown). But with feedback (inhibitory and modulatory, see Section III-A), the second set is reduced. As indicated in Figure 3 a, the second wave represents mainly the correct symbol. Figure 5 shows four examples of the recognition dynamics.

In many cases, the dynamics reduces the activity until only the correct word remains. As in Section IV-A we investigate the dynamics of the second wave for the same 104 stimuli. As shown in Figure 6, the ratio of correct versus incorrect symbols is greatly increased (from 0.38 to 0.86, if false-positives are present). 43 of the words are found correctly and uniquely in the second wave. However, in 10 cases the correct word does not appear in the second wave. The rest of the trials contain the correct word in the second wave along with weaker false-positives that share a great overlap with the correct word (cf. Figure 5 b).

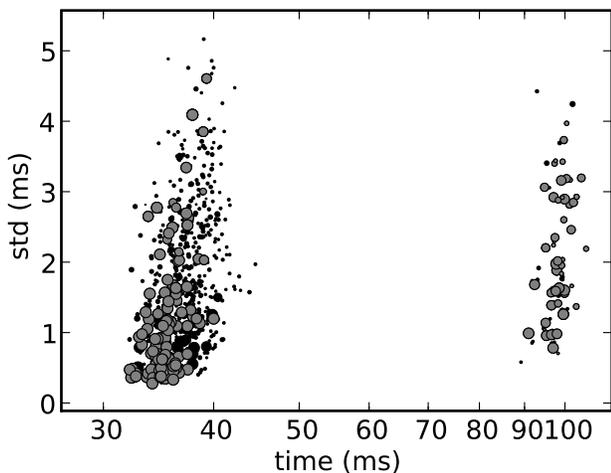


Fig. 6. Word refinement in IT-B. Shown are two waves of word-pulses. The first wave corresponds to Figure 4. About 60 ms after the first wave, a second wave appears with almost only correct words (gray).

C. Symbol shift

To ascertain how the correct word ‘overrules’ the wrong alternatives, we investigate the effect of modulatory feedback in isolation. To this end, we have switched off the B→A2 inhibition (cf Figure 1). Thus, the candidates can reoccur, and the network can not refine its hypothesis. Comparison of the two symbol waves (Figure 7) reveals that the correct words, which already appear at the onset of the first wave, reoccur faster than the wrong candidates and are shifted further ahead from the distribution. The average lead is enhanced by 2.7 ms. Thus, modulatory feedback alone appears to increase the temporal segregation of correct versus incorrect words. Yet due to the missing feedback inhibition, the relative amount of false-positives during the second wave is larger.

In summary, word recognition takes place in two main stages. In the first wave at about 35 ms, responses comprise

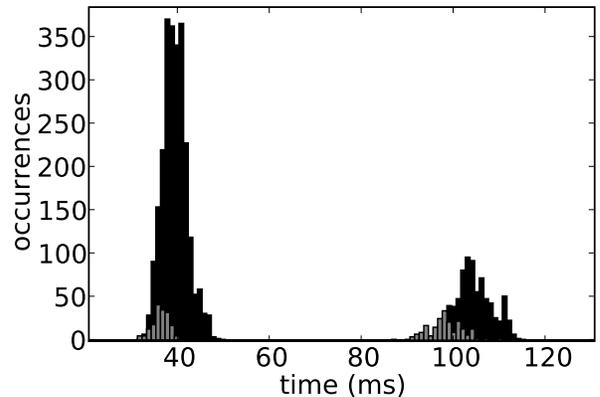


Fig. 7. Shifting the correct symbols. The four histograms count the occurrence times of words without feedback inhibition (B to A2). In the early phase (around 40 ms), the correct words (gray) are located at the early onset of the false-positive distribution (black). During the second wave (around 100 ms), correct symbols appear earlier than before. The lead of correct words is increased by 2.7 ms on average.

about one percent of all possible words and represent a set of close candidates that come into question. From 104 trials, 55% of the earliest and 99% of the strongest responses correspond to the correct word. In the second wave (100 ms), 90% of the correct words reappear with the strongest response. Modulatory feedback isolates the correct words by shifting them further ahead.

V. DISCUSSION

In this paper we have presented the neoCOREtext model, a network that is able to shape successively higher representations of complex objects. Using three hierarchy stages that refer to steps in visual processing, entities are dynamically ‘composed’ from a set of primitives. For conceptual simplicity, we have abstracted the sensory modality to the perception of words given a string of characters. The underlying neuron model is an integrate-and-fire neuron with action potential generation. Our network features horizontal (layered) and vertical (columnar) organization. We interpret those levels of *anatomical* modularization as a division into *functional* building blocks: Across **layers and areas**, a fast initial hypothesis is formed via the A-A path which is refined by feedback processes between the B-systems. Within **columns**, local symbols compete by lateral inhibition and are subject to feedback inhibition.

We showed in Section IV-A that higher areas respond in pulse-like activity, even if the stimuli are continuously applied. This sharpening is caused by the converging knowledge-connectivity (Figure 2). For instance, a syllable that is correctly excited by its letters (usually more than one), receives a stronger input than unrelated syllables. Firing threshold is quickly reached and the spikes of the underlying population are more coherent. Consequently, the refractory periods are temporally aligned, rendering the neural pool excitable again at the same time. This leads to two predictions.

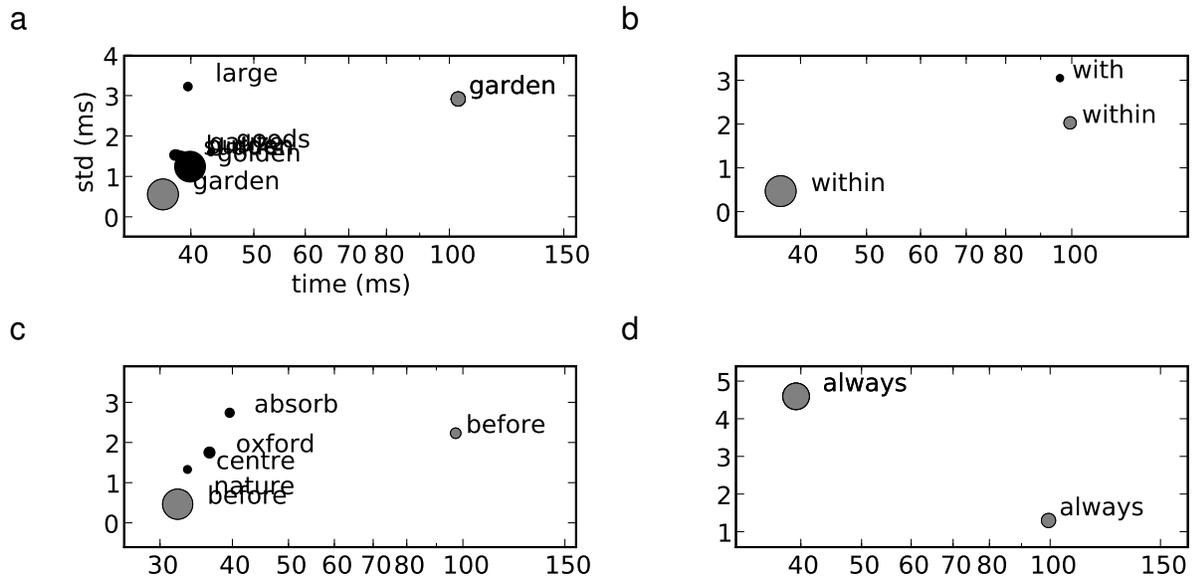


Fig. 5. Word recognition examples. Shown are four examples of word recognition tasks. Pulses are symbolized as in Figure 4. (a) An example for distinct recognition in the presence of many candidates (*garden*). (b) False-positives can occur in the second wave (*within*). (c) The earliest word occurrence during the first wave for all 104 trials (Figure 4), *before*. (d) Example where only the correct word is elicited (*always*). This example corresponds to the point with the highest temporal spread (std) in the first wave (Figure 4)

First, wrong symbols should not be active at all. Figure 5 shows that only words with a certain overlap are activated during the first pulse, most of the words do not appear. Second, symbols that are closer to the correct word should appear earlier, because their input is higher and threshold is reached earlier. Figure 5 a shows that the correct word, *garden*, precedes the similar word *golden*, which differs by two letters. This is confirmed in Figure 4 where, on average, the correct word precedes the rest of the candidates by a few milliseconds. This small lead allows fast lateral inhibition to suppress the wrong alternatives.

Our network demonstrates that precise spike volleys can arise from phasic inputs and, more importantly, it suggests *how* information could be temporally coded. A response gains its truth value by its relative latency. The earlier a symbol appears, the more likely it is correct. Therefore, information can be extracted as soon as possible. This is already the case for the first wave of candidates (Figure 4), when feedback has not yet arrived. The correct word can be read out in almost all cases by a threshold or a simple winner-takes-all circuit. This suggests how sensory information can be processed in situations when there is no time to refine the hypothesis. For example, it has been shown that decisions can already be made within short periods that exclude feedback paths (36).

During the second wave, the set of word candidates is reduced by 73% (Figure 6) and often consist of only the correct word. We conclude that refinement works as follows. First responses are propagated via both feedforward paths, (“A” and “B” in Figure 1). They elicit a quick response with a coarse hypothesis on the stimulus. Symbols that are

recognized in the B-systems suppress the activity in layer A2 with feedback inhibition (B-A2, see Figure 1), weakening the contribution of the A-pathway. We state that this suppression is critical in cortical processing - symbols that are already recognized are taken out from the initial pathway, allowing other symbols to be resolved. On the other hand, the impact of correct symbols to the next area has been up-regulated by the B-system with modulatory feedback. This results in an elevated activity along the B-pathway. Thus, by means of modulation, a B-system ‘gives credibility’ to its strongest sources, predicting the most certain symbols. This *Predictive Coding* scheme has a direct impact on strengths and relative latencies of subsequent responses. Correct responses occur stronger and earlier than before (Figure 7). Although this shift is only in the millisecond range (2.7 ms on average), a fast lateral inhibition uses this time lead and silences the wrong symbols. Patterns that are already part of the established interpretation of the input are kept at the B-system, while its removal from the A-system enables a further refinement of the interpretation by adding more details of the input description.

We have shown that the responses of the word-area IT are correct for the majority of the words. However, in 10 trials (9.6%), the correct word failed to reoccur in the second wave of activity. These words did not produce a weaker stimulus signal (e.g. they were not shorter than others). Obviously, more reliable symbol repetition can be achieved by strengthening the feed-forward path or by attenuating the feedback. For instance, all correct words reoccurred in the case without feedback inhibition (Figure 7). On the other hand, this leads to an increase of word candidates and

complicates the separation into correct and wrong symbols. Reliability could rather be increased by reaching more precise responses during the *first* wave. In fact, the words that disappeared in the second wave failed to have the first or most precise response in the initial wave and were also part of a bigger candidate set (up to 11 candidates in the first wave, 7% of all words on average, as opposed to 1.2%). Therefore, the weights of the feedforward path should be adjusted to elicit equally sized candidate lists. This can be reached by using weight normalization which takes word similarity into account.

Some hundred thousand neurons seem to be a high amount to code for 500 words. Our network size denotes an upper bound for two reasons. First, we have realized all local symbols with no regard to their retinal positions. For instance, the syllable *ing* is also implemented at the first retinal position although it only occurs at the end of words. Obviously, symbols that are never used are not learned or would not endure in a real-world situation. Second, nearby neurons often fail to have identical receptive fields (37) indicating that there is a high overlap between neural populations. This receptive overlap, however, requires a similarity measure for symbols. In our model, we have neglected such distance measures in order to demonstrate the network’s basic recognition capabilities.

Having equipped the network with word-knowledge in a supervised fashion, the question remains as to how complex objects are *learned*. A possible learning signal could emanate from layer A2, where unresolved symbols sustain a higher activity. Thus, ‘ignorance’ of a symbol is translated into an error signal that can mediate between symbols by means of synaptic plasticity.

APPENDIX

Parameters of the integrate-and fire neuron are given in Table I. The individual neurons receive their synaptic

parameter	value	unit
R_m	200	$M\Omega$
V_0	0.0	mV
V_θ	20.0	mV
τ_m	20	ms
t_{ref}	2.0	ms
τ_{syn}	0.5	ms
τ_{mod}	50.	ms
f_{mod}	0.001	1

TABLE I

PARAMETERS OF THE INTEGRATE-AND FIRE NEURON AS IN EQU. (1) OF SECTION III-C

currents as shown in Table II. Capital indices denote the areas and layers. If no area is given, the equation holds for all three layers. A superscript “—” denotes the next lower area, otherwise terms refer to the same area. N denotes the quantity of the connections, Greek letters the weights and delays (see Table III).

source	target	weight/quantity	delay (ms)	remark
noise	all	7.1 kHz		Poisson noise
stimulus	V1-A1	2.0 kHz		Poisson noise
A1	A2	$N_\alpha = 50,$ $\alpha = 0.3 \text{ pA}$	$\delta = 1.0$	
A1	B	$N_\alpha = 50,$ $\alpha = 0.3 \text{ pA}$	$\delta = 1.0$	
B	A2	$N_\eta = 50,$ $\eta = -0.6 \text{ pA}$	$\delta = 1.0$	
V1-A2	V2-A1	$N_\zeta = 30,$ $\zeta = 0.1 \text{ pA}$	$\delta = 1.0$	
V1-B	V2-B	$N_\zeta = 30,$ $\zeta = 0.1 \text{ pA}$	$\delta = 1.0$	
V2-A2	IT-A1	$N_\gamma = 30,$ $\gamma = 0.1 \text{ p}$	$\delta = 1.0$	
V2-B	IT-B	$N_\zeta = 30,$ $\zeta = 0.12 \text{ pA}$	$\delta = 1.0$	
IT-B	V2-B	$f_{mod} = 0.001$	$\delta = 1.0$	
V2-B	V1-B	$f_{mod} = 0.001$	$\delta = 1.0$	
IT-B	IT-B	$N_\lambda = 10,$ $\lambda = -0.1 \text{ pA}$	$\delta_2 = 0.5$	lateral inhibition
all	all	$\rho = 0.1 \text{ pA},$ $N_{ee} = N_{ei} = 5,$ $N_{ie} = N_{ii} = 1.5$	$\delta = 1.0$	recurrent connections within one minicolumn

TABLE III

PARAMETERS OF SYNAPTIC CONNECTIVITY.

ACKNOWLEDGMENT

The authors would like to thank Jochen Eppler for help on the simulation environment and for comments on the manuscript. We thank Andreas Knoblauch for comments on the manuscript.

REFERENCES

- [1] V. Braitenberg and A. Schüz, *Anatomy of the Cortex: Statistics and Geometry*. Berlin, Heidelberg, New York: Springer-Verlag, 1991.
- [2] B. Hellwig, “A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex,” *Biological Cybernetics*, vol. 2, no. 82, pp. 111–121, Feb. 2000.
- [3] R. J. Douglas and K. A. C. Martin, “Neuronal circuits of the neocortex,” *Annual Review of Neuroscience*, vol. 27, pp. 419–451, 2004.
- [4] S. Song, S. Per, M. Reigl, S. Nelson, and D. Chklovskii, “Highly nonrandom features of synaptic connectivity in local cortical circuits,” *Public Library of Science, Biology*, vol. 3, no. 3, pp. 0507–0519, 2005.
- [5] A. M. Thomson and C. Lamy, “Functional maps of neocortical local circuitry,” *Frontiers in Neuroscience*, vol. 1, pp. 19–42, 2007.
- [6] D. O. Hebb, *The organization of behavior: A neuropsychological theory*. New York: John Wiley & Sons, 1949.
- [7] D. J. Amit and N. Brunel, “Model of global spontaneous activity and local structured activity during delay

$E_{V1-A1}(t) =$	$P + S(t)$	$+ \rho (N_{ee} E_{A1}(t - \delta) + N_{ie} I_{A1}(t - \delta))$	
$I_{V1-A1}(t) =$	$P + S(t)$	$+ \rho (N_{ei} E_{A1}(t - \delta) + N_{ii} I_{A1}(t - \delta))$	
$E_{A2}(t) =$	$P + N_{\alpha} \alpha E_{A1}(t - \delta)$	$+ \rho (N_{ee} E_{A2}(t - \delta) + N_{ie} I_{A2}(t - \delta))$	$+ N_{\eta} \eta I_B(t - \delta)$
$I_{A2}(t) =$	$P + N_{\alpha} \alpha E_{A1}(t - \delta)$	$+ \rho (N_{ei} E_{A2}(t - \delta) + N_{ii} I_{A2}(t - \delta))$	
$E_{V1-B}(t) = \mu(t)[$	$P + N_{\alpha} \alpha E_{A1}(t - \delta)$	$+ \rho (N_{ee} E_B(t - \delta) + N_{ie} I_B(t - \delta))]$	
$I_{V1-B}(t) = \mu(t)[$	$P + N_{\alpha} \alpha E_{A1}(t - \delta)$	$+ \rho (N_{ei} E_B(t - \delta) + N_{ii} I_B(t - \delta))]$	
$E_{V2-A1}(t) =$	$P + N_{\gamma} \gamma E_{A2}^-(t - \delta)$	$+ \rho (N_{ee} E_{A1}(t - \delta) + N_{ie} I_{A1}(t - \delta))$	
$I_{V2-A1}(t) =$	$P + N_{\gamma} \gamma E_{A2}^-(t - \delta)$	$+ \rho (N_{ei} E_{A1}(t - \delta) + N_{ii} I_{A1}(t - \delta))$	
$E_{IT-A1}(t) =$	$P + N_{\gamma} \gamma E_{A2}^-(t - \delta)$	$+ \rho (N_{ee} E_{A1}(t - \delta) + N_{ie} I_{A1}(t - \delta))$	
$I_{IT-A1}(t) =$	$P + N_{\gamma} \gamma E_{A2}^-(t - \delta)$	$+ \rho (N_{ei} E_{A1}(t - \delta) + N_{ii} I_{A1}(t - \delta))$	
$E_{V2-B}(t) = \mu(t)[$	$P + N_{\alpha} \alpha E_{A1}(t - \delta)$	$+ \rho (N_{ee} E_B(t - \delta) + N_{ie} I_B(t - \delta))$	$+ N_{\zeta} \zeta E_{B-}]$
$I_{V2-B}(t) = \mu(t)[$	$P + N_{\alpha} \alpha E_{A1}(t - \delta)$	$+ \rho (N_{ei} E_B(t - \delta) + N_{ii} I_B(t - \delta))$	$+ N_{\zeta} \zeta E_{B-}]$
$E_{IT-B}(t) =$	$P + N_{\alpha} \alpha E_{A1}(t - \delta)$	$+ \rho (N_{ee} E_B(t - \delta) + N_{ie} I_B(t - \delta))$	$+ N_{\zeta} \zeta E_{B-} + N_{\lambda} \lambda I_B(t - \delta_2)$
$I_{IT-B}(t) =$	$P + N_{\alpha} \alpha E_{A1}(t - \delta)$	$+ \rho (N_{ei} E_B(t - \delta) + N_{ii} I_B(t - \delta))$	$+ N_{\zeta} \zeta E_{B-} + N_{\lambda} \lambda I_B(t - \delta_2)$

TABLE II
EQUATIONS FOR THE INDIVIDUAL SYNAPTIC CURRENTS.

- periods in the cerebral cortex,” *Cerebral Cortex*, vol. 7, pp. 237–252, 1997.
- [8] N. Brunel, “Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons,” *Journal of Computational Neuroscience*, vol. 8, no. 3, pp. 183–208, 2000.
- [9] D. Hansel and G. Mato, “Asynchronous states and the emergence of synchrony in large networks of interacting excitatory and inhibitory neurons,” *Neural Computation*, vol. 15, pp. 1–56, Jan 2003.
- [10] E. Boustani and A. Destexhe, “A master equation formalism for macroscopic modeling of asynchronous irregular activity states,” *Neural Computation*, p. in press, 2008.
- [11] A. Kumar, S. Schrader, A. Aertsen, and S. Rotter, “The high-conductance state of cortical networks,” *Neural Computation*, vol. 20, no. 1, pp. 1–43, 2008.
- [12] T. P. Vogels, K. Rajan, and L. F. Abbott, “Neural network dynamics,” *Annual Review of Neuroscience*, vol. 28, pp. 357–376, 2005.
- [13] H. C. Tuckwell, *Introduction to Theoretical Neurobiology*. Cambridge: Cambridge University Press, 1988, vol. 1, ch. 3, The Lapique model of the nerve cell, pp. 85–123.
- [14] H. B. Barlow, “Single units and sensation: a neuron doctrine for perceptual psychology?” *Perception*, vol. 1, no. 4, pp. 371–394, 1972.
- [15] Y. Yoshimura, H. Sato, K. Imamura, and Y. Watanabe, “Properties of horizontal and vertical inputs to pyramidal cells in the superficial layers of the cat visual cortex,” *J Neurosci*, vol. 20, no. 5, pp. 1931–1940, Mar 2000.
- [16] V. B. Mountcastle, “The columnar organization of the neocortex,” *Brain*, vol. 120 (Pt 4), pp. 701–722, Apr 1997.
- [17] M. Hübener, D. Shoham, A. Grinvald, and T. Bonhoeffer, “Spatial relationships among three columnar systems in cat area 17,” *Journal of Neuroscience*, vol. 17, no. 23, pp. 9270–9284, 1997.
- [18] K. H. Britten, “Clustering of response selectivity in the medial superior temporal area of extrastriate cortex in the macaque monkey,” *Vis Neurosci*, vol. 15, no. 3, pp. 553–558, May-Jun 1998.
- [19] J. S. Lund, A. Angelucci, and P. C. Bressloff, “Anatomical substrates for functional columns in macaque monkey primary visual cortex,” *Cerebral Cortex*, vol. 12, pp. 15–24, Jan 2003.
- [20] M. R. del Rio and J. DeFelipe, “Double bouquet cell axons in the human temporal neocortex: relationship to bundles of myelinated axons and colocalization of calretinin and calbindin d-28k immunoreactivities,” *J Chem Neuroanat*, vol. 13, no. 4, pp. 243–251, Oct 1997.
- [21] D. P. Buxhoeveden and M. F. Casanova, “The minicolumn hypothesis in neuroscience,” *Brain*, vol. 125, no. Pt 5, pp. 935–951, May 2002.
- [22] E. Körner, M.-O. Gewaltig, U. Körner, A. Richter, and T. Rodemann, “A model of computation in neocortical architecture,” *Neur. Netw.*, vol. 12, no. 7–8, pp. 989–1005, 1999.
- [23] R. Kupper, A. Knoblauch, G. Marc-Oliver, U. Koerner, and E. Koerner, “Simulations of signal flow in a functional model of the cortical column,” *Neurocomputing*, vol. 70, no. 10–12, pp. 1711–1716, November 2007.
- [24] A. Destexhe, M. Rudolph, and D. Pare, “The high-conductance state of neocortical neurons in vivo,” *Nat Rev Neurosci*, vol. 4, pp. 739–751, 2003.

- [25] M. Matsumura, D. Chen, T. Sawaguchi, K. Kubota, and E. E. Fetz, "Synaptic interactions between primate precentral cortex neurons revealed by spike-triggered averaging of intracellular membrane potentials *in vivo*," *Journal of Neuroscience*, vol. 16, no. 23, pp. 7757–7767, December, 1 1996.
- [26] S. N. Baker and R. Lemon, "Precise spatiotemporal repeating patterns in monkey primary and supplementary motor areas occur at chance level," *Journal of Neurophysiology*, vol. 84, pp. 1770–1780, 2000.
- [27] M. Abeles, "Time is precious," *Science*, vol. 304, pp. 523–524, 2004.
- [28] R. VanRullen, R. Guyonneau, and S. J. Thorpe, "Spike times make sense," *Trends in Neurosciences*, vol. 28, no. 1, pp. 1–4, 2005.
- [29] M. Abeles, E. Vaadia, H. Bergman, Y. Prut, I. Haalman, and H. Slovin, "Dynamics of neuronal interactions in the frontal cortex of behaving monkeys," *Concepts in Neuroscience*, vol. 4, no. 2, pp. 131–158, 1993.
- [30] A. Riehle, S. Grün, M. Diesmann, and A. Aertsen, "Spike synchronization and rate modulation differentially involved in motor cortical function," *Science*, vol. 278, no. 5345, pp. 1950–1953, December 1997.
- [31] Y. Prut, E. Vaadia, H. Bergman, I. Haalman, S. Hamutal, and M. Abeles, "Spatiotemporal structure of cortical activity: Properties and behavioral relevance," *Journal of Neurophysiology*, vol. 79, no. 6, pp. 2857–2874, 1998.
- [32] R. Johansson and I. Birznieks, "First spikes in ensembles of human tactile afferents code complex spatial fingertip events." *Nature Neuroscience*, vol. 2, no. 7, pp. 170–177, Feb. 2004.
- [33] M.-O. Gewaltig and M. Diesmann, "NEST (Neural Simulation Tool)," *Scholarpedia*, vol. 2, no. 4, p. 1430, 2007.
- [34] J. M. Eppler, M. Helias, E. Muller, M. Diesmann, and M.-O. Gewaltig, "Pynest: A convenient interface to the nest simulator," *Frontiers in Neuroinformatics*, 2008, submitted.
- [35] J. J. B. Jack, S. J. Redman, and K. Wong, "The components of synaptic potentials evoked in cat spinal motoneurons by impulses in single group Ia afferents," *Journal of Physiology (London)*, vol. 321, pp. 65–96, 1981.
- [36] S. Thorpe, D. Fize, and C. Marlot, "Speed of processing in the human visual system," *Nature*, vol. 381, pp. 520–522, 1996.
- [37] G. C. DeAngelis, G. M. Ghose, I. Ohzawa, and R. D. Freeman, "Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons," *J Neurosci*, vol. 19, no. 10, pp. 4046–4064, May 1999.