

Simulations of signal flow in a functional model of the cortical column


Rüdiger Kupper, Andreas Knoblauch, Marc-Oliver Gewaltig, Ursula Körner, Edgar Körner

2007

Preprint:

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

Provided for non-commercial research and educational use only.
Not for reproduction or distribution or commercial use.



This article was originally published in a journal published by Elsevier, and the attached copy is provided by Elsevier for the author's benefit and for the benefit of the author's institution, for non-commercial research and educational use including without limitation use in instruction at your institution, sending it to specific colleagues that you know, and providing a copy to your institution's administrator.

All other uses, reproduction and distribution, including without limitation commercial reprints, selling or licensing copies or access, or posting on open internet sites, your personal or institution's website or repository, are prohibited. For exceptions, permission may be sought for such use through Elsevier's permissions site at:

<http://www.elsevier.com/locate/permissionusematerial>

Simulations of signal flow in a functional model of the cortical column

Rüdiger Kupper^{*}, Andreas Knoblauch, Marc-Oliver Gewaltig, Ursula Körner, Edgar Körner

Honda Research Institute Europe GmbH, Carl-Legien-Str. 30, D-63073 Offenbach/Main, Germany

Available online 10 November 2006

Abstract

We describe the simulation of a layered cortex model based on the cortical column as a generic *local processor*. It simulates the signal flow in the layers I–IV of a set of model columns across three hierarchical cortical areas. It demonstrates the fast formation of an initial stimulus hypothesis, and its subsequent refinement by inter-columnar communication. In this prototype simulation, we implement word recognition from a string of characters. The three cortical areas represent letters, syllables, and words, used as a metaphor for visual stimuli. Focusing on the intra- and inter-columnar dynamics, we show how the different processing subsystems interact in order to switch off expected signals and accomplish symbolic recognition of words, and how representations for new words can be constructed based on old representations (*self-reference*).

© 2006 Elsevier B.V. All rights reserved.

Keywords: Cortex; Column; Generic circuit; Predictive coding; (re-)Cognition; COREtext model

1. Introduction

How does the brain, most notably the visual system, manage to process and ultimately “understand” the immense amount of data, that is picked up by our sensors in each second of everyday life? What strategies, what neural algorithms does it use to interpret the sensory input in terms of what it “knows”, and how does it decide when to learn and memorize new content?

In [6] we have put forward a hypothesis of computation in neocortical architecture, that bridges the gap between processing of signals at the single-neuron level, and the processing of cognitive symbols at the level of knowledge representation: this model proposes the *cortical column* as the basic, generic building block of the cortex. At each point on the cortical surface, the column vertically connects neurons across the six cortical layers, with a distinct circuitry. The column is a module in a two-fold sense: it is a *sub-unit* in the architectural sense [3], and it is a *sub-process* in the algorithmic sense. Our model gives a detailed functional interpretation of the six-layered columnar cortical architecture (Fig. 1). It hypothesizes three intercommunicating processing systems in the columns at each stage of the

cortical hierarchy: in the “A-system” (middle cortical layers IV and lower III), the first wave of spikes traveling upwards in the cortical hierarchy can activate a coarse initial “local hypothesis” on the contents present in the stimulus. In the “B-system” (superficial layers II and upper III), this initial hypothesis is refined by slower processes. Finally, the “C-system” (deep layers V and VI) represents the local interpretation of the input signals that results from the local integration of bottom-up and top-down signals. Subsequently, input signals that match the local prediction are suppressed, and only differences between predicted and actual signals can reach the next higher cortical level [7]. Thus, stimulus content is effectively expressed in terms of previously acquired knowledge (*self-reference*).

In this article, we describe the COREtext 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.

2. Model

The COREtext model consists of three areas, “IT”, “V2”, and “V1”. Each area is composed of three layers, each of which is a linear array of h neural subsystems

^{*}Corresponding author. Tel.: +49 69 89011725; fax: +49 69 89011749.
E-mail address: ruediger.kupper@honda-ri.de (R. Kupper).




Fig. 1. Layered model of a cortical column as proposed in [6]. Three different subsystems at different vertical locations (layers) are intertwined within each cortical column.

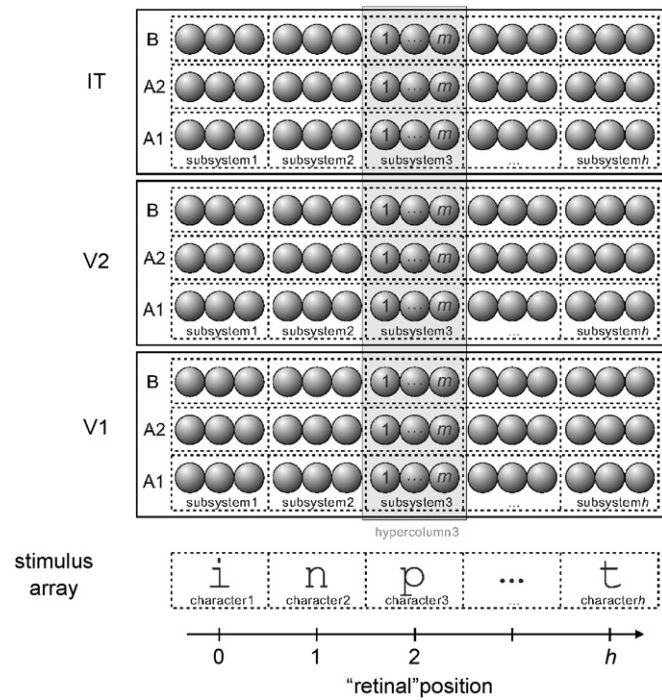


Fig. 2. Hierarchy of areas, layers and subsystems. The subsystems are “retinotopically” aligned with the stimulus array.

(Fig. 2). A neuronal subsystem is an aggregation of m neurons. The number m is referred to as the number of minicolumns in the area. The number h is referred to as the number of hypercolumns in the area. The COREtext model uses a graded-response neuron model. In each time-step, a neuron accumulates its stimulation s and its modulation m . The resulting activation a is given by $a = \Theta(s \cdot [1 + m])$. Here, Θ is a piecewise linear sigmoidal transfer function. Θ is zero below a firing threshold of 0.3, and then rises linearly up to a maximal activation value of $a = 3$, $\Theta(1) = 1$. The neurons in each subsystem participate in a “winners-take-all” competition, and all but the maximal activations in a subsystem are set to 0.

Synaptic connections between the neurons in the COREtext model are specified separately between pairs of layers. Weights between the neurons of given source and target layers take binary values, multiplied by a constant factor w : $(w_{ijkl}) = w \cdot (b_{ijkl})$, $b_{ijkl} \in \{0, 1\}$, with i the source subsystem, j the source neuron in the subsystem, k the target subsystem, and l the target neuron in the subsystem. When subsystem and neuron indexes are not of explicit interest, we can shorten this notation to $(w_{ij}) = w \cdot (b_{ij})$. A set of synaptic connections can either be driving or modulating. Inhibitory connections are realized by negative driving weights.

The COREtext model uses a synchronous update of neuronal states. In each simulation time-step, the activity of all neurons is propagated to the postsynaptic neurons, according to the connection matrix. From the resulting activation values, the “winners-take-all” competition is computed inside the neural subsystems. The result is stored as the new activation of the neurons. This update cycle repeats in every simulation time-step.

For conceptual simplicity, we choose an abstracted stimulus environment that shows a clear hierarchical structure of how complex objects are composed from smaller parts. We implement word recognition from a string of characters. Stimuli have the form of strings of lower-case characters, arranged on a one-dimensional grid. In our setup of the COREtext model, all layers have the same number h of subsystems (hypercolumns), and are aligned with the input array (Fig. 2). We choose a number of three COREtext areas that form a hierarchy, similar to sensory processing pathways in the neocortex. The areas are labeled “V1”, “V2” and “IT” for convenience, where these labels are metaphors taken from the domain of visual processing. Fig. 3 (right) relates features processed in the visual domain to the features processed in the COREtext

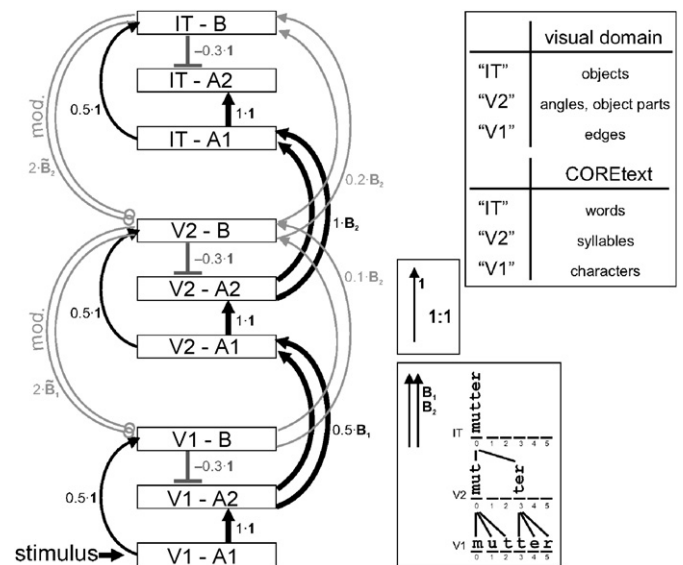


Fig. 3. Left: connectivity scheme between the COREtext layers. Right: example for the construction of a cognitive object (word) from its constituents (syllables and letters) in the local alphabets.

model. We derived the set of features represented in the three areas from the analysis of three pages of classical German literature [4]. We decomposed the text into syllables according to usual German spelling rules. After removal of duplicates the text analysis yielded a set of 386 different words, a set of 418 syllables, and a set of 30 characters. We call these sets of features the *local alphabets* L_{IT} , L_{V2} , and L_{V1} . The sizes of the local alphabets determine the number m of minicolumns in the three areas. Each hypercolumn in “IT”, “V2”, and “V1” represents the full local alphabet. I.e., each subsystem in “V2” is an aggregation of $m_{V2} = 418$ neurons, with each neuron $n_{i,s,i}$ representing the i th element of L_{V2} .

Knowledge about the composition of cognitive objects (words) is generated from German spelling rules. From the local alphabets we derive two binary knowledge matrices, \mathbf{B}_1 and \mathbf{B}_2 . We initialize the knowledge matrices in a process that imitates supervised one-shot-learning. It is similar to training all words in L_{IT} by presenting them at position 0 in the input array, while giving their correct decompositions into syllables and characters. For each word in L_{IT} we derive its constituting syllables and characters, as well as their starting position in the word. We store these decompositions in the binary knowledge matrices \mathbf{B}_1 and \mathbf{B}_2 . For example, the word “mutter” decomposes in the way shown in Fig. 3 (right).

We designed the connectivity scheme between the nine COREtext layers to reflect the main functional projections that are known to exist between the layers of cortical areas [8,1,3,2]. We aimed to create a reduced model of the functional connectivity that is at the core of the cortical columnar processing. The connectivity scheme is shown in Fig. 3, left panel. Layers of the same COREtext area are linked by 1:1-connections (single arrows). These are connections between neurons in the same minicolumn. (E.g., a neuron representing the syllable “mut” in V2–A1 projects to the neuron representing the *same* syllable and at the *same* position in V2–A2.) Stimulation enters the system at layer V1–A1. This layer is linked to the input array (Fig. 2) by 1:1-connections.

Layers of different COREtext areas are linked by connection matrices proportional to the two knowledge matrices \mathbf{B}_1 and \mathbf{B}_2 (double arrows). These are connections that span mini- and hypercolumns, and that implement the knowledge about the hierarchical composition of cognitive objects (words) in the network. Backward arrows ending in circles indicate modulatory connections. The modulatory backward connections between the B-systems are reverse to the forward connections. I.e., a neuron in IT–B projects back to exactly the same neurons in V2–B that it receives from. The matrices \mathbf{B}_1 and \mathbf{B}_2 denote these reverse connections. The exact connection values including scaling factors are shown beside the arrows in Fig. 3.

The connection scheme can be categorized into several pathways. The strong $A1 \rightarrow A2 \Rightarrow A1$ forward pathway reflects the strong driving synapses that are found between layers IV and III of cortical areas. The weaker $A1 \rightarrow B \Rightarrow$

B pathway and the top-down modulatory $B \rightarrow B$ pathway reflect the reciprocal connections between the neurons in upper layers II and III of cortical areas. The inhibitory $B \rightarrow A2$ pathway reflects local inhibition via inter-neurons in upper and middle layers of a column.

3. Results: response modes

Simulation runs consisted of a constant stimulation by a string of lower-case characters in the input array, and subsequent iteration of the update cycle. The signal flow in the COREtext system, based on the distinct connectivity and the knowledge imprinted in the synaptic connections, takes different response modes.

Fig. 4 shows a stimulation example using a string of characters that is identical to a known word. The word “mutter” was part of the analyzed text, and knowledge about its decomposition into syllables (*mut*, *ter*) and characters is imprinted into the connections between the COREtext areas. The stimulus activates the corresponding single-character-detectors at each hypercolumn in “V1”. From here, a fast wave of activation spreads via the $A1 \rightarrow A2 \Rightarrow A1$ path and the $A1 \rightarrow B \Rightarrow B$ path (cf. Fig. 3). Syllable-detectors in “V2” and word-detectors in “IT” get active according to the forward connection matrices. Starting from these first activations in the B-systems, activity propagates backwards via two pathways. Along the modulatory top-down $B \rightarrow B$ pathway, an active word detector in “IT” supports activations in all detectors in “V2” that are compatible with the word’s decomposition into syllables. Similarly, each active syllable detector supports activations in all compatible detectors in “V1”. At the same time, active neurons in the B-systems inhibit neurons in the A2-systems in the same minicolumn of the same area. (E.g., an active detector for syllable *mut* in V2–B inhibits the detector for the same syllable *mut* and at the same position in V2–A2.)

The state shown in Fig. 4 is the static pattern of activations in response to the stimulus “mutter”, that is

```

t=10
-----
IT-A1:  0:      IT-A2:  0:      IT-B:   0: mut-ter(1.3)
IT-A1:  1:      IT-A2:  1:      IT-B:   1:
IT-A1:  2:      IT-A2:  2:      IT-B:   2:
IT-A1:  3:      IT-A2:  3:      IT-B:   3:
IT-A1:  4:      IT-A2:  4:      IT-B:   4:
IT-A1:  5:      IT-A2:  5:      IT-B:   5:

V2-A1:  0:      V2-A2:  0:      V2-B:   0: mut(3.0)
V2-A1:  1:      V2-A2:  1:      V2-B:   1:
V2-A1:  2:      V2-A2:  2:      V2-B:   2: sten(0.4) ste(0.4)
V2-A1:  3:      V2-A2:  3:      V2-B:   3: ter(3.0)
V2-A1:  4:      V2-A2:  4:      V2-B:   4: er(0.4)
V2-A1:  5:      V2-A2:  5:      V2-B:   5:

V1-A1:  0: m(1.0)  V1-A2:  0:      V1-B:   0: m(3.0)
V1-A1:  1: u(1.0)  V1-A2:  1:      V1-B:   1: u(3.0)
V1-A1:  2: t(1.0)  V1-A2:  2:      V1-B:   2: t(3.0)
V1-A1:  3: t(1.0)  V1-A2:  3:      V1-B:   3: t(3.0)
V1-A1:  4: e(1.0)  V1-A2:  4:      V1-B:   4: e(3.0)
V1-A1:  5: r(1.0)  V1-A2:  5:      V1-B:   5: r(3.0)

```

Fig. 4. Stimulation example using a known word “mutter”. Neural activations are given in parentheses behind the corresponding symbols. Activations of 0 are not shown.

reached after 10 update cycles. In IT–B, the correct symbol *mutter* is active. In V2–B and V1–B compatible constituting syllables and characters are active. The activations in the B-systems support each other via the forward $B \Rightarrow B$ and backward $B \rightarrow B$ pathways. At the same time, via the inhibitory $B \dashv A2$ pathway, they “switch off” the according signals in the A2-systems of the same minicolumn. After 10 update cycles, activations in the A2-systems have completely vanished, and activations in the B-systems are mutually supportive.

Fig. 5 shows a stimulation example using a string of characters that is similar, but not identical to a known word. The stimulus “vatrxr” is a distorted version of the words “vater” and “vaters” that were part of the analyzed text. The state shown in Fig. 5 is the static pattern of activations reached after eight update cycles. The bottom-up/top-down dynamics activates the symbols *vater* and *vaters* in IT–B, showing that the system was able to compensate for the distortions. In V2–B and V1–B compatible constituting syllables and characters are active. The activations in the B-systems support each other via the forward $B \Rightarrow B$ and backward $B \rightarrow B$ pathways. Note that the two distorted characters *x* in V1–B are only weakly active, since they *do not* receive modulatory support from any syllable in V2–B. This means, they can exert only weak inhibition on their V1–A2 counterparts via the inhibitory $B \dashv A2$ pathway. After eight update cycles, all activations in the A2-systems have vanished, except for the two distorted characters in V1–A2.

Fig. 6 shows a stimulation example using the stimulus “rüdiger” that is largely dissimilar to all words that were part of the analyzed text. A variety of symbols in “IT” and “V2” get weakly active, and via the top-down modulatory $B \rightarrow B$ pathway support all compatible syllables and characters in V2–B and V1–B. At the same time, they inhibit the according signals in the A2-systems of the same minicolumn. This changed distribution of activations in the A2 and B systems changes the activations in the B-systems via the forward pathways. The mutual dependence of the

```

t=8
-----
IT-A1: 0:
IT-A1: 1:
IT-A1: 2:
IT-A1: 3:
IT-A1: 4:
IT-A1: 5:

IT-A2: 0:
IT-A2: 1:
IT-A2: 2:
IT-A2: 3:
IT-A2: 4:
IT-A2: 5:

IT-B: 0: va-ter(1.5) va-ters
IT-B: 1:
IT-B: 2:
IT-B: 3:
IT-B: 4:
IT-B: 5:

V2-A1: 0:
V2-A1: 1:
V2-A1: 2:
V2-A1: 3:
V2-A1: 4:
V2-A1: 5:

V2-A2: 0:
V2-A2: 1:
V2-A2: 2:
V2-A2: 3:
V2-A2: 4:
V2-A2: 5:

V2-B: 0: va(3.0) hat(3.0) ta
V2-B: 1:
V2-B: 2: ter(3.0) ters(3.0)
V2-B: 3:
V2-B: 4:
V2-B: 5:

V1-A1: 0: v(1.0)
V1-A1: 1: a(1.0)
V1-A1: 2: t(1.0)
V1-A1: 3: x(1.0)
V1-A1: 4: r(1.0)
V1-A1: 5: x(1.0)

V1-A2: 0:
V1-A2: 1:
V1-A2: 2:
V1-A2: 3: x(0.9)
V1-A2: 4:
V1-A2: 5: x(0.9)

V1-B: 0: v(3.0)
V1-B: 1: a(3.0)
V1-B: 2: t(3.0)
V1-B: 3: x(0.3)
V1-B: 4: r(3.0)
V1-B: 5: x(0.3)

```

Fig. 5. Stimulation example using a distorted word “vatrxr”. The best matching symbols have been found in IT–B, and the two distorted characters remain active in V1–A2, indicating parts of the stimulus that could not be explained from knowledge.

```

t=15
-----
IT-A1: 0: red-ner(0.5) rühr-t
IT-A1: 1:
IT-A1: 2:
IT-A1: 3:
IT-A1: 4:
IT-A1: 5:
IT-A1: 6:

IT-A2: 0: red-ner(0.4) rühr-t
IT-A2: 1:
IT-A2: 2:
IT-A2: 3:
IT-A2: 4:
IT-A2: 5:
IT-A2: 6:

IT-B: 0: red-ner(0.1) rühr-t
IT-B: 1:
IT-B: 2:
IT-B: 3:
IT-B: 4:
IT-B: 5:
IT-B: 6:

V2-A1: 0: red(0.8) rühr(0.8)
V2-A1: 1:
V2-A1: 2: hig(0.4) nigs(0.4)
V2-A1: 3: cher(0.6)
V2-A1: 4: der(0.6) ler(0.6)
V2-A1: 5: brin(0.2) trä(0.2)
V2-A1: 6: ri(0.2)

V2-A2: 0: red(0.5) rühr(0.5)
V2-A2: 1:
V2-A2: 2: hig(0.8) nigs(0.8)
V2-A2: 3: cher(0.8)
V2-A2: 4: der(0.8) ler(0.8) g
V2-A2: 5:
V2-A2: 6:

V2-B: 0: red(0.4) rühr(0.4)
V2-B: 1:
V2-B: 2: hig(0.3) nigs(0.3)
V2-B: 3: cher(0.4)
V2-B: 4: der(0.4) ler(0.4)
V2-B: 5:
V2-B: 6:

V1-A1: 0: r(1.0)
V1-A1: 1: ü(1.0)
V1-A1: 2: d(1.0)
V1-A1: 3: i(1.0)
V1-A1: 4: g(1.0)
V1-A1: 5: e(1.0)
V1-A1: 6: r(1.0)

V1-A2: 0: r(0.5)
V1-A2: 1: ü(0.7)
V1-A2: 2: d(0.7)
V1-A2: 3: i(0.8)
V1-A2: 4: g(0.8)
V1-A2: 5: e(0.7)
V1-A2: 6: r(0.7)

V1-B: 0: r(1.1)
V1-B: 1: ü(0.7)
V1-B: 2: d(0.7)
V1-B: 3: i(0.3)
V1-B: 4: g(0.3)
V1-B: 5: e(0.3)
V1-B: 6: r(0.3)

```

Fig. 6. Stimulation example using an unknown word “rüdiger”. Activations in the A- and B-systems of all areas keep oscillating through a set of alternatives.

two systems leads to a constantly changing pattern of weak activations. The state shown in Fig. 6 is the pattern of activations reached after 15 update cycles. This state is not stable, but activations in the A- and B-systems of all areas keep changing between reappearing patterns. As a consequence, the response to the unknown stimulus “rüdiger” is not a stable active symbol in “IT”, but a whole set of alternating symbols.

3.1. Randomized update

In order to rule out that the third response mode (iteration of a set of symbols) is an epiphenomenon caused by the synchronous update of all model neurons, we implemented a randomized update scheme. In this scheme, only a fixed fraction of randomly chosen neurons is updated in a time-step, while all other neurons keep their activations. We could confirm, that the phenomenon reproduces independently of the fraction of updated neurons. Area IT iterates the same set of symbols in all cases. Typical set sizes are 2–10, depending on the input string. The size of the iterated set, and thus the frequency of occurrence of the individual symbols, depends on the number of known symbols that overlap with the input string. (For example, the iterated words tend to start with the same letter as the input string, making the iterated sets larger for frequent German starting letters.)

The large invariance with respect to the fraction of updated neurons shows that symbol iteration in response to unknown stimuli is a robust propagation phenomenon in our network. It is rooted in the mutual dependence of feed-forward excitation and inhibitory feedback, which cannot be congruent for an unknown stimulus: the winning symbols in the B-systems support the inhibition of their constituting parts via the top-down $B \rightarrow B \dashv A2$ pathway (Fig. 3). However, an unknown stimulus activates a set of parts that is different from the constituting parts of any one known symbol at the next hierarchic level (otherwise, it would be known). The remaining, non-inhibited parts will consequently cause another symbol to win the competition at the next hierarchic level, causing an iteration of symbols that overlap with the stimulus.

4. Discussion

After a stimulus is applied, a fast wave of bottom-up activation spreads via the forward pathways. Symbols in the upper cortical areas get active after only a few monosynaptic propagations of activity. In this part of the activity spread, the whole system acts as a multilayered perceptron. This fast forward activation of a first hypothesis about the stimulus content is compatible with findings on the speed of processing of categoric information in the human visual system [9]. After the formation of a fast initial hypothesis, it is consolidated with the evident stimulation via the inter-areal top-down modulatory pathway. All stimulus parts that could be confirmed get “switched off” in the A2 systems (middle cortical layers), indicating that the active symbols in upper cortical areas correctly predict these parts of the stimulus [7]. Finally, the B-systems maintain a self-consistent explanation of the stimulus from “pure knowledge”.

In the case of a stimulus with variations or distortions (second example), a self-consistent explanation can also be established. The “switching off” of activity in the A2-systems must, however, leave residuals, since parts of the stimulus establish a bottom-up evidence that cannot be confirmed by top-down consolidation. The B-systems represent an abstracted or corrected version of the stimulus, as would be expected from “pure knowledge”. Still, information on the unexpected details is not lost: residuals in the A2-system clearly identify the parts of the stimulus that cannot be explained from knowledge. This residual activity in the A2-system can be used in several ways. First of all, its pure existence is an indication, that the recognized symbols do not entirely represent the stimulus. Second, the residual activity is specific in the position (the hypercolumn) it appears in. It indicates the exact position of the unexplained parts in the stimulus, and can thus guide a motor action, e.g., a saccade, to gather additional information on the yet unexplained parts. We show simulations of this function in a companion article (this issue) [5]. Third, the residual activity is specific in the exact symbol from the local alphabet (the minicolumn) it appears in. The residual activity in the A2-system thus fulfills the necessary prerequisites to enable incremental learning in the cortical hierarchy: it indicates *when* to learn, *what* to learn, and also *where* to learn it.

Our third example showed that an “unknown” stimulus cannot be represented across the hierarchy of areas in a self-consistent way. Instead, at all hierarchical levels, the system keeps “associating” possible symbols that are locally compatible both with the momentary bottom-up-stream of signals, and the momentary top-down-stream of hypothesized symbols. This state of activation is clearly different from the case of stable activation with residuals. It indicates that the knowledge *does not suffice* to explain the stimulus. It is important that the neural system has a means to indicate this conflict, instead of converging into some stable, but necessarily inappropriate state of activation that

would ultimately deceive the individual into taking wrong actions and drawing wrong conclusions. Still, this type of activation is more than a pure “error-state”: activations tend to converge towards sets of repeating symbols. The exact sets of alternating symbols are determined by the interaction between the evident signal, and the system’s knowledge imprinted in the synapses. Thus, they are specific to the stimulus: the system starts to paraphrase the stimulus in its own terms. The human drive to make sense from everything could actually be rooted in the core circuitry of our brains. It is also an ideal basis for the formation of new stimulus concepts. The repeating set of cognitive symbols itself can be learned as a description of the new stimulus. We think, that the hippocampus is the instance on top of the cortical hierarchy that performs this transformation [6].

5. Conclusion

Realistic stimuli occurring in a rich environment will *almost never* match a known symbol. Moreover, stimuli that are “beyond knowledge”, (be it at a lower or at a higher level of cognitive symbols), will occur frequently. It is thus a lesson to be learned from the COREtext model, that a stable pattern of activation will *almost never* be reached, and that this is not a limitation, but a meaningful state of the system. Iterating patterns of activation is the rule, not the exception in this setup of intercommunicating subsystems. The conceptually reduced connectivity of the COREtext model was derived from a thorough review of physiological findings over many years. We see this as a strong hint, that the typical, neurophysiologically observed cortical oscillations represent a natural mode of processing in the cortex, and are inherent from the unique kind of information flow inside and between cortical layers and columns. Constant iteration of *compatible alternatives* at all levels of detail may simply be the key to cortical processing, and ultimately, understanding.

Acknowledgment

We thank the anonymous reviewers for comments, which helped to clarify important aspects in the description and evaluation of our model. The comments led to a substantial improvement of our article.

References

- [1] A.P. Bannister, Inter- and intra-laminar connections of pyramidal cells in the neocortex, *Neurosci. Res.* 53 (2) (2005) 95–103.
- [2] E.M. Callaway, Feedforward, feedback and inhibitory connections in primate visual cortex, *Neural Networks* 17 (5–6) (2004) 625–632.
- [3] R.J. Douglas, K.A. Martin, Neuronal circuits of the neocortex, *Annu. Rev. Neurosci.* 27 (2004) 419–451.
- [4] H. Hesse, Siddhartha: eine indische Dichtung, Project Gutenberg Etext, 12th ed., vol. 2499, Project Gutenberg, (<http://www.gutenberg.org/etext/2499>), 2001, pp. 418–514.

- [5] A. Knoblauch, R. Kupper, M.-O. Gewaltig, U. Körner, E. Körner, A cell assembly based model for the cortical microcircuitry, *Neurocomputing* (2006), this issue, doi:10.1016/j.neucom.2006.10.092.
- [6] E. Körner, M.-O. Gewaltig, U. Körner, A. Richter, T. Rodemann, A model of computation in neocortical architecture, *Neural Networks* 12 (7–8) (1999) 989–1005.
- [7] R.P.N. Rao, D.H. Ballard, Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects, *Nat. Neurosci.* 2 (1) (1999) 79–87.
- [8] A.M. Thomson, A.P. Bannister, Interlaminar connections in the neocortex, *Cereb. Cortex* 13 (1) (2003) 5–14.
- [9] S. Thorpe, D. Fize, C. Marlot, Speed of processing in the human visual system, *Nature* 381 (6) (1996) 520–522.



Rüdiger Kupper is a senior scientist at the Honda Research Institute Europe in Offenbach/Main, Germany. He has been a member of this institute since its foundation in 2003. Rüdiger received his Ph.D. in physics from the University of Marburg, Germany, for a neural simulation model of distance-invariant object recognition. His work at the Honda Research Institute is dedicated to the discovery and computer-simulation of the fundamental processing strategies and control architecture of the human cortex. His vision is a neural theory of intelligence.



Andreas Knoblauch is currently a senior scientist at the Honda Research Institute Europe in Offenbach. He received his Ph.D. from the University of Ulm in computer science and neural information processing (2003), and holds a master degree from the same university in computer science and electrical engineering (2000). From 2003 to 2005 he worked as a postdoc at the Neural Information Processing Department at the University of Ulm and the Cognition and Brain Sciences Unit at the MRC Cambridge, UK. His research interests include modeling of cortical circuits, learning in associative networks, information retrieval, and the design of simulation software.



Marc-Oliver Gewaltig was born in 1967 in Germany. He received his M.Sc. in physics from the Ruhr-University of Bochum in 1994. After associations with the Ruhr-University Bochum and the Albert-Ludwigs University Freiburg, he received his Ph.D. in 1999. In 1998, he started working for Honda R&D Europe (Deutschland) GmbH. Since 2003, he is a principal scientist at the Honda Research Institute Europe GmbH in Offenbach/Germany. His research interests are the neuronal principles of information processing in the primate brain as well as methods for simulation of large cortical networks.



Ursula Körner is a principal scientist at the Honda Research Institute Europe in Offenbach, Germany. She received a degree in neurophysiology at the Moscow State University and obtained a Ph.D. in physiology at the Humboldt State University in Berlin. Ursula Körner's research interests concentrate on principles of information processing and learning in cortical structures.



Edgar Körner is the president of the Honda Research Institute Europe GmbH. He studied electrical engineering, control engineering, and biomedical cybernetics at the Ilmenau Institute of Technology, Germany where he was appointed professor in 1988. In 1992, he joined Honda R&D in Wako near Tokyo as a chief scientist for brain-like computation research. In 1997 he moved to Honda R&D Europe, to head the Future Technology Research division, focussing on computational neuroscience, evolutionary technology, and cognitive robotics. Since 2003, Prof. Dr. Körner is president of the Honda Research Institute Europe GmbH. His research focus is on brain-like artificial neural systems and self organization of knowledge representation for autonomous robots.