Unpublished manuscript prepared for the Ernst Strüngmann Forum on
Dynamic coordination in the brain: From Neurons to Mind

Coordination: What Is It, and Why Do We Need It?

Christoph von der Malsburg
Frankfurt Institute for Advanced Studies
Ruth-Moufang-Str. 1, 60438 Frankfurt am Main, Germany
malsburg@fias.uni-frankfurt.de

What is coordination? There are many domains in which this issue is central, for instance when preparing a meal and setting the dinner table, when writing a novel or when composing music. The aesthetic feelings a piece of art arouses in our mind often reflect just the level of coordination between our sensations. Also when building and running a company or when organizing a conference, structural elements must relate to each other in a meaningful way, must fit together and form a whole for the result to be functional. Man-made structures inherit their coordination from that within their creators' brains. All insight we can glean from coordinated artifacts is therefore relevant for understanding coordination in the brain.

What is Coordination in the Brain?

Coordination means putting things together that belong together. Each term in this sentence raises questions. In the following I discuss some of them. My emphasis is more on questions than on answers. As things in our brain* are to be put together both on the slow timescale of learning and on the fast timescale of thinking, we must keep an eye on both.

What is the stuff that is to be put together?

The brain has a great variety of modalities or subsystems, corresponding to internal and external senses, to motor control, to memory, to emotions, to behavioral control and more. Minsky (1988) speaks of the Society of Mind. Superficial modalities are formed by wiring with the sensory or motor or humoral periphery and are thus defined (onto-)genetically. More centrally, there is considerable plasticity, at least in the cortical system, as shown by many cases of neurological recovery (Weiller and Chollet 1994). Different modalities focus on different reflections of phenomena and contain patterns that are similar, so that they can be mapped and piled onto each other. This pervasive phenomenon of cortical localization requires a natural measure of similarity between neural patterns, the need for which we will encounter repeatedly later on.

Each subsystem of the brain is able to create a large variety of alternate activity patterns, of which at any given time only one or a few can clearly and unambiguously be expressed. For much of the time, neural activity in a subsystem expresses an ambiguous superposition of different states, which is to be reduced in stages under the influence of signal exchange. This has to happen in a coordinated way, so that patterns are co-activated that belong together. This reduction of uncertainty is, for instance, modeled in probabilistic formulations (Pearl 1988, Bishop 2006).

What does it mean to belong together?

---

* I am adopting the attitude of Spinoza and am seeing brain and mind as two sides of the same coin. When using electrode or microscope we see the brain while psychophysics or introspection lets us see the mind.

Neural patterns belong together because they were generated by stimuli that are statistically linked in the environment (for which a common cause is often responsible), or because they successfully interact to attain goals or to playfully exercise and develop capabilities.

Thus, belonging-together is both defined in a passive, recording mode, and in an active, creative mode. The creative mode is responsible for the brain's ability to deal with new situations, and is grossly neglected in the neuroscientific literature.

How does the system learn what belongs together?

A first motive for patterns to be related is similarity. Different modalities, by contrast, speak different languages and contain neural patterns that are not similar to each other. How can the brain find out which patterns in different modalities belong together? The only means to establish *de novo* pattern association is simultaneity, significant correlation in time. Unfortunately, however, it is not useful to associate with each other all neural patterns in the brain just on the basis of simultaneity (as is implied in associative memory models, Hopfield 1982). To an overwhelming extent, pattern simultaneity out there in the world is accidental and not of lasting value. If synaptic plasticity had the form of indiscriminate stickiness the brain would soon be cluttered with all-to-all connections. It is therefore important to single out associations that are significant. The natural definition of pattern significance – recurrence – suggests the strategy (which actually is adopted by much of the artificial neural network literature, Haykin 1994) to sift through the input for those patterns that appear again and again with statistical significance. Unfortunately, this strategy fails for input fields of realistic size, as the number of patterns to keep track of is much too large and literal repetition of a given pattern is too unlikely. It is rather necessary for the system to possess a similarity measure and to have powerful prejudices concerning the nature of significant patterns. The question is to be asked what is the nature of these prejudices?

Once different modalities have already accumulated a sufficient mass of pattern associations between them, they can use general laws of composition to creatively generate novel, modality-spanning composite patterns. An important issue is the nature of these laws of composition. A typical (or maybe the typical) law may be that overlapping patterns in one modality are to be associated with overlapping patterns in another.

The Detection of Significant Patterns by Focal Attention

Why is the information content of attention limited? Is this an imperfection of the brain, or is this fact of functional significance? Focal attention powerfully (though not exclusively) restricts learning to a small subset of active neurons at any one time (for discussion of this subject see Jiménez, 2003). Key questions to ask are how this restriction is expressed and how this restriction acts back on learning. (One proposal is based on gamma rhythms, see Fell et al., 2003, and on the ensuing concentration of neural spikes into narrow temporal windows as a boost to synaptic plasticity.) This reduces the input and the memory domains to small sectors (as modelled in Jacobs et al., 1991). This restriction addresses a fundamental problem of present-day models of learning – the scaling-up to realistically large input and storage domains – by restricting system modification to within the narrow focus of attention. The reason that the informational content of attention is limited is, in this view, not imposed by an imperfect mechanism but by the necessity to preclude confusion.

On what informational basis is the focus of attention formed? Attention is very much at the heart of the problem of coordination: it has to bring together sets of patterns that belong together. If it could only unite patterns that had already been associated in the past – in passive mode, so to speak – it wouldn't help with the problem of learning. For that, it has to creatively find new associations.

External events, signaled by temporally isolated sensory signals (which are emphasized by the filter properties of our senses) can bring together specific patterns that light up simultaneously in different modalities. Patterns can also be brought together on the basis of abstract properties with which they are tagged. For instance, if the senses can be focussed on the same point in space (for which colliculus superior seems to be well equipped, at least in the mouse, see Dräger and Hubel, 1975) patterns aroused from that point can be associated. The general Gestalt laws (Koffka, 1935; Crick, 1994) can define significant patterns as figures set apart from a background. An object moving against a static background can, for instance, be made to stand out and be made the focus of attention on the basis of common motion (Spelke 1998). Thus, the Gestalt laws formulate abstract properties (common motion, color, stereo depth, spatial grouping, good form, edge continuity), which help to tag novel patterns as significant.

A statistical definition of pattern significance cannot be all, it also needs a biological definition. Important classes of patterns or events have to be genetically defined in order to tell the individual what is required for success in life. This definition must be laid down in some abstract fashion such that concrete occurrences can be recognized and singled out. Ethologists describe many such schemata, such as the facial schema with which the human infant is born, the definition of mother goose for the gosling or a red dot on the beak of the seagull to indicate to the chick the source of food (Toats, 1980). Upon the recognition of the releaser for an innate cue, attention is focussed on the recognized stimulus, which is then separated from the ground, appropriate action is induced (like an orienting reflex or grasping), and an appropriate sector of memory is singled out for modification by the stimulus. This sequence of events may be called schema-based learning.

Important questions remain regarding the technical implementation of these processes.

How is coordination evaluated?

We have a keen feeling for the level of coordination in our brain. Sometimes we feel distracted, confused, can't make up our mind, feel there's something awry, something's not quite right. Sometimes, by contrast, we feel focused, are highly concentrated, fully conscious of a situation, and then there are these precious moments when we suddenly feel we have it, it all falls in place, we shout Eureka! To a large extent, aesthetic pleasure is due to the level of coordination generated in our brain by the object of our attention. It is not conscious insight into the structure of the piece of art that we experience, but just some direct feeling of the level of coordination in our brain. This Aha! effect, this falling into a state of organization is what the Gestaltists called reorganization of the perceptual field and spoke of insight (Köhler, 1925).

What are the structures in our brain that are responsible for the evaluation of this measure of coordination, and what is the nature of this measure? Of course it is not some superior intellectual entity whose insight into the subject matter serves as judge. It must rather be some signal that can be "mechanically" generated and globally evaluated. The essence of it may be the level of non-trivial agreement of independent signals at convergence points, evaluated over the whole brain.

How is purpose defined, enforced and achieved?

In an active mode, the brain has to coordinate patterns such as to achieve its purposes. The central issue here is that purpose (like, I am hungry and am looking for food) on the one hand, and the generation of neural patterns that serve to achieve that purpose (opening the fridge or calling the pizza delivery service) on the other hand are defined on very different levels of detail, and generally in different parts of the brain. The newborn contains, presumably in mid-brain, the schematic definition of a set of fundamental goals. These form a hierarchy, the honing of which keeps us busy over much of the course of our life. Goals are activated either spontaneously or in response to some stimulus, like "danger" or "thirst", they tend to be mutually exclusive, and they come equipped with powerful mechanisms of enforcement. Complex tasks require the attainment of goals and subgoals in hierarchical fashion. There are deep questions concerning the nature, the establishment and the implementation of goals in our brain.

Behavioral patterns usually have a number of functional components. For each of these there is a range of possible role fillers (in the looking-for-food scenario, relevant roles concern possible food-stuffs, sources of food [fridge, delivery service etc.], possible modes of getting there, etc.). In a concrete situation with a concrete goal in mind, the system has to select appropriate role fillers that together interact functionally to attain the goal. It is an important question to discuss how this type of coordination is achieved by the interaction between the goal schema containing a set of role descriptions, possible role fillers that have the ability to combine appropriately and the sensory patterns describing the situation.

Important questions also concern the form in which behavioral schemata and goal descriptions are implemented neurally, the mechanisms by which these schemata are triggered and prioritized, the reward mechanisms by which the achievement of goals are evaluated, the mechanisms by which the activity of the brain is biased in the direction of goal fulfillment, and the mechanisms by which, over the course of our life, goal schemata are elaborated in richer and richer ways in terms of detailed sensory and motor patterns.

How is the environment's statistics to be captured?

The brain receives signals from the environment over many millions of fibers and influences it through many output fibers. All our brain can ever know and learn is contained in the statistics of these activity patterns. Of all possible combinations of individual neural input or output signals, a space of vast volume, only a minute sub-volume is ever realized in terms of actual signals. Exhaustive recording of global activity patterns is not possible, nor would it make sense, as no sensory-motor activity pattern ever has a chance to recur. Only by extracting significant sub-patterns, by ordering them in groups of similar patterns, and by developing schemata for their arrangement is it possible to capture the environment's statistics and cope with the ever-novel situations we are faced with. This needs a prejudice as to how to define significant patterns to be extracted, needs a general similarity measure by which these patterns are to be grouped, and needs a pre-established format for the representation of pattern arrangements. Here one is between a rock and a hard place. If the prejudices are too weak, the system is drowned in variance that cannot be captured in a realistic finite system. If, on the other hand, the system's prejudices are too narrow, the reality of the environment may be missed (The bias-variance dilemma of Geman et al., 1992). Another indication that the system's prejudices must be tuned to the environment are the no-free-lunch theorems (Wolpert and Macready, 1997), according to which any learning or optimization mechanism can be totally vitiated by an environment that doesn't fit its a priori assumptions. In one word, the brain needs powerful a priori assumptions and they must fit the actual environment! What are these a priori assumptions?

What is the nature of our environment's statistics?

This is probably the most crucial of all questions, since, as argued, the brain in its mode of operation must be tuned to the environment. In fact, the whole point is that the brain must coordinate with the environment.

Some important aspects of sensory pattern statistics are due to the media through which they are transmitted. The visual medium, for instance, is the optical radiation field captured by the eye, and the patterns that appear on our retinae are shaped by the laws of reflection and propagation of light, by geometry and motion. These transformations need to be inverted for the brain to uncover the structure of the patterns out there. The laws of transformation are to a large extent independent of the environmental patterns themselves. The first sensory stages of the brain can reduce the complexity of the input patterns tremendously by inverting these transformations, reducing large sets of patterns to invariance classes (Wiskott, 2006).

What is the regularity of the world beyond that? What are the repeating patterns? Or, rather, in what general format do they come? According to Kant's analysis we must come equipped with what he called categories, a priori structures that make it possible for us to absorb information in the first place. Among these he counted space, time and causality. We take it for granted that repeating spatial and temporal patterns and causal sequences of events play an important role in our environment. It is a wide-open question, however, how whole scenes are to be decomposed to find repeating patterns, and how to formulate the general rules of composition by which our environment generates its configurations in ever-new ways. Coordination, to relate again to our theme, is the ability to create internal scenes that capture the reality of the environment. The brain's task, then, is to extract from the environment patterns and their relations, together with a measure of likelihood for their relevance, and to thus acquire the ability to complement partial information in a given scene with additional detail familiar from the past, in order to generate a more complete description of the scene. Our challenge is to second-guess the general form – the architecture – on the basis of which this is possible.

What is the nature of structural relationships?

The patterns that we experience never repeat precisely. When recording a novel pattern it is therefore important to be able to define a halo of other patterns that are similar to it. This implies a similarity measure, or some definition of the likelihood that a sensory pattern is to be identified with a stored pattern. If properly constructed, the stored structure and the similarity measure can decide with high statistical significance whether a perceived pattern is an accidental arrangement of elements or whether it is the same pattern again. What is this similarity measure? The simplest idea of pattern recognition is template matching, where a rigid pattern, the template, is moved over an image to find an identical fit. The "motion" takes care of the invariance aspect if it includes all possible transformations (e.g., translation, scaling and rotation). Template matching has long since fallen out of favor because identical fits are never found in real images. A first step to solve the problem is to take the "template" to pieces and endow the pieces with flexible relationships so that distortion can be dealt with. A necessary next step is to replace the pieces of the model with statistical models of possible variants. (A version of this is the leading mechanism of face recognition, as described, for instance, in Wolfrum, 2008.) And then, the pieces can themselves be replaced by composite models, creating a hierarchical structure.

What has been described here for vision analogously applies to other modalities. Also motor patterns (including speech) form patterns within patterns, each being a role filler, each permitting a range of variants, the whole put together flexibly to permit continuous time warping and, of course, further nesting. Hidden Markov models (Rabiner and Juang, 1986) capture essential aspects of this. The contention is that this kind of architecture applies to all the modalities of the brain individually, and the the brain as a whole.

Out of these considerations arises a picture according to which mental objects are hierarchical graph structures, with concrete patterns put in spatial and temporal relationships with each other. Graphs are in general embedded in or linked to more abstract graphs (which, due to their abstractness, are called schemata), whose nodes refer to classes of exchangeable sub-patterns, and whose links describe permitted relations. Recognition is the process by which abstract graphs are mapped to concrete patterns homeomorphically (that is, with equivalent parts mapping to each other under preservation of relations). Coordination is the process by which concrete brain states are generated under the guidance of abstract descriptions (including formulation of goals) on the one hand and of sensory input on the other. Usually, several abstract schemata conspire to create a detailed description.

Important questions are how the repertoire of neural behavior is tuned to the construction of such hierarchical descriptions, how hierarchical descriptions are developed in the brain on the basis of experience, and how this architecture is to be described in concrete mathematical terms.

How are neural patterns put together?

The brain is endowed with an architecture on the basis of which it tends to fall into globally ordered patterns that in their structure are tuned to the world in which we live. This ability may be put in analogy to the process of crystallization, in which constituent atoms or molecules, by exerting their preferences as to the shape of the local environment, create global order out of local interactions. For crystallization to start, a seed is required – a minimal structure such that further molecules quickly find a niche into which to fall.

In the case of the brain, the constituent elements are neurons, and it is an important question how their behavioral repertoire is structured such as to favor global order. Some aspects of this repertoire are already emerging. Outgrowing processes are guided by chemical or electrical signals such as to favor ordered connectivity patterns, as exemplified by the ontogenesis of retinotopy (Goodhill, 2007), whereas intrinsic plasticity (Butko and Triesch, 2007) regulates the duty cycle of cellular activity.

If the behavioral repertoire of neurons (or of a collection of neural types) were known to any degree of precision, it should be possible to simulate on the computer the growth of ordered connectivity and activity states. So far, this venture has had some level of success, especially in modeling the ontogenesis of ordered connectivity structures. However, modeling the generation of neural states that can be interpreted as mental objects is still out of reach. So, questions are remaining regarding the repertoire of neural behaviors, and, possibly equally important, regarding the equivalent of seed structures with which to start the process of coordination.

Concluding Remarks

To coordinate stuff in our mind, even while performing routine tasks, we regularly have to put things together that have never been put together before. These associations are creative acts that are not imposed on us from outside. When, for instance, recognizing an object, we are applying a schematic description of this type of object to a concrete image, associating abstract features with concrete instantiations. When manipulating the object or describing it verbally, this enables us to directly relate to the specific character of the feature instances. Our brain certainly contains the neural pathways to connect what is to be connected, but these pathways are embedded – are drowned – in a virtual continuum of others that at present are irrelevant but are all required and useful when their time comes. Although I intended to pose questions and nothing but questions in this essay, I couldn't do so without slipping in more and more of my conviction that the task of our brain, in any given situation, is not just to activate a subset of all neurons, but to select from vast numbers of physical connections a tiny subset. The task of selecting connections is bigger than that of selecting neurons by the factor synapses outnumber neurons, about ten thousand. If we ignore this task, we are, according to this view, ignoring 99.99% of the information in our brain's state.

If indeed the brain in its rapid state changes is mainly concerned with selecting structured connectivity patterns, and if we need to study and understand synaptic dynamics in addition to neural dynamics to bridge the chasm between mind and brain, we have a problem. Experimental technology is highly developed to study neural dynamics, to a smaller extent to study static or slowly changing connectivity, can even record short-term modification of synaptic effects for individual connections, but the imaging of whole, rapidly changing, connectivity patterns is presently beyond imagination. When Ludwig Boltzmann first established statistical mechanics, he was ridiculed by his colleagues Ernst Mach, Wilhelm Ostwald and others for his atomistic ideas, and it took three decades until experiments made the reality of atoms and molecules concrete enough to convince the community. Shall the neurosciences also wait for decades for the necessary revolution? Dedicating years of effort to an experimental exploration is too risky if not supported by community convictions and therefore doesn't happen, but the community is not convinced without focussed efforts. Let's hope that this vicious circle can be broken with the help of concrete computer models of cognitive functions whose demonstrable success rests on outlandish physiological assumptions.

Acknowledgement I acknowledge support by BMBF project Bernstein, FKZ 01GQ0840.

References

Bishop, C.M. (2006) Pattern Recognition and Machine Learning. Springer.

Butko, N., & Triesch, J. (2007). Learning sensory representations with intrinsic plasticity. Neurocomputing 70, 1130–1138.

Crick, F. (1994) The Astonishing Hypothesis. London: Simon and Schuster.

Dräger UC, Hubel DH (1975) Responses to visual stimulation and relationship between visual, auditory, and somatosensory inputs in mouse superior colliculus. J Neurophysiol 38, 690–713.

J. Fell, G. Fernández, P. Klaver, C.E. Elger and P. Fries (2003) Is synchronized neuronal gamma activity relevant for selective attention?, *Brain Res. Brain Res. Rev.* **42** , 265–272.

Geman, S., Bienenstock, E. and Doursat, R(1992) Neural Networks and the Bias/Variance Dilemma. Neur. Comp. 4, 1-58.

G.~J. Goodhill, G.J. (2007) Contributions of theoretical modeling to the understanding of neural

map development.

Neuron 56, 301--311.

Haykin, S. (1994) Neural Networks. A Comprehensive Foundation. Macmillan.

J.J. Hopfield: Neural networks and physical systems with emergent collective computational abilities. PNAS 79 (1982) 2554–2558.

Jacobs, R.A., Jordan, M.I., Nowlan, S.J., Hinton, G.E. (1991) Adaptive Mixtures of Local Experts. Neur. Comp. 3, 79-87.

Jiménez, L (Ed.). (2003). Attention and implicit learning. Amsterdam: Benjamins.

Köhler W. (1925) The Mentality of Apes. New York, Harcourt Brace.

Koffka, K. (1935) Principles of Gestalt Psychology, Harcourt,New York.

Minsky, M. (1988), The society of mind. Simon and Schuster.

Pearl, J. (1988) Probabilistic Reasoning in Intelligent Systems: Networks of plausible Inference. Morgan Kaufmann.

Rabiner, L. and Juang, B. (1986) An introduction to hidden Markov models. IEEE ASSP Magazine 3, 4-16.

Spelke, E.S. (1998). Nativism, empiricism, and the origins of knowledge. Infant Behavior and Development , 21 , 181 – 200.

Toates, F (1980). Animal Behaviour – A Systems Approach. Wiley.

Weiller, C., Chollet, F. (1994). Imaging recovery of function following brain injury. Curr. Opin. Neurobiol. 4, 226–230.

Wiskott, L. (2006) **How does our visual system achieve shift and size invariance?** In *23 Problems in*

Systems Neuroscience, eds. J. L. van Hemmen and T. J. Sejnowski, publ. Oxford University Press, New York.

Wolfrum, P., Wolff, C., Lücke, J., and von der Malsburg, C. (2007). A Recurrent Dynamic Model for Correspondence-Based Face Recognition. Journal of Vision 8, 1--18.

Wolpert, D.H. And Macready, W.G. (1997) No free lunch theorems for optimization
IEEE Trans. Evol. Comp. 1, 67-82.