Exploratory Workshop Scheme

Scientific Review Group for the Bio-Medical Sciences

Scientific Review Group for Life, Earth and Environmental Sciences

ESF Exploratory Workshop on

Noise in Decision Making: Theory Meets Experiment

Sant Fruitós de Bages (Catalonia, Spain), 28-31 May 2013

Convened by:
Albert Compte, Gustavo Deco, Jaime de la Rocha, Alex Roxin and Klaus Wimmer

---

SCIENTIFIC REPORT

---

1. Executive summary

The workshop “Noise in Decision Making: Theory Meets Experiment” took place in the medieval monastery of Món Sant Benet (http://www.monstbenet.com/en), in Sant Fruitós de Bages, 65 km away from Barcelona. The workshop was held over 2.5 days, from May 28 to May 31, 2013 (starting in the evening of May 28 with a visit of the monastery and an informal reception). See the workshop website at http://www.crm.cat/2013/DecisionMaking/ for more details and a group picture.

The final number of participants was 30 (24 invited speakers, 1 journal editor, 5 conveners) and 1 ESF representative. Workshop conveners came from three different institutions in Barcelona, and the remaining participants came from all over Europe (19 participants) and from the United States (6 participants). Co-sponsoring from the Swartz foundation was available to cover the expenses for US-based participants. Administrative and organizational support was available from the CRM (Center for Mathematical Research, http://www.crm.cat).

The entire event took place in the facilities of Món Sant Benet. All participants stayed at the hotel Món (http://www.hotelmonstbenet.com/), a two-minute walk away from the monastery where the workshop took place. Additional interaction between the participants was facilitated by two cultural events: a visit of the monastery and a visit of the Alicia foundation (a culinary research center), both on-site. Participants were able to mingle during the breaks and also during short walks to the restaurants where we had lunch and dinner. Lunches took place in a separate hall with exclusive use for the workshop participants. The somewhat remote location (compared to downtown Barcelona) permitted close interaction between the participants during the whole day, in a quiet setting without distraction. Discussions continued until later in the night in the hotel lobby. The general atmosphere was informal, open to discussions and extremely friendly.

Scientific objectives and agenda of the meeting

The workshop aimed at exploring the role of “noise” in decision making. The purpose of the workshop was to bring together experimental and theoretical neuroscientists in order to come up with a combined approach towards elucidating the origin of variability in brain activity and behavior. Behavior seems inherently stochastic: even for identical repetitions of a simple perceptual decision task, our responses are variable. What do we know about the mechanisms that generate this “noisy” behavior? Brain activity is also highly irregular and seemingly stochastic. A classical theoretical study proposed a simple circuitry in which variable behavior can emerge from stochastic neural activity (Shadlen et al., *Journal of Neuroscience* **16**:1486, 1996). This model has served as a fundamental framework to design, analyze and interpret experimental data over the years. Recent theoretical and experimental results are prompting, however, for a revision of the model postulates.

In order to focus on the role of neuronal variability, the mechanisms which generate it, and its impact on decisions and behavior, we asked the workshop participants to specifically address the following questions:

1. What is the origin of variability in brain activity (contributions from external inputs, synaptic and neural mechanisms, micro-circuit dynamics, global brain states, etc.)?
2. What is the role, if any, of neuronal variability in decision making / perception / behavior?

The goals of the workshop were:

- Discuss the most recent findings from a variety of experimental and theoretical approaches.
- Create a space for researchers with different backgrounds to engage in a dialogue and forge new relationships.
- Identify common interests, as well as new formulas for collaborations (such as sharing of experimental data and theoretical models) and opportunities for further collaborative actions at the European and international levels in order to foster research projects that truly integrate theory and experiment.
- To try to reformulate the standard framework laid out more than fifteen years ago.
- Provide the workshop talks to a broader audience as "online talks" that are freely available on the internet.

The workshop talks (25 minutes + 5 minutes for questions) were organized in 5 coherent sessions. Discussions continued through the coffee breaks and during lunch time. There was an additional discussion session at the end of each day (1 hour) for extensive debates of topics that arose throughout the day. Topics for these discussions were collected on a flip-chart during the day. The workshop closed with a discussion session which aimed at defining a new framework for the role of noise in decision making.

Overall conclusions

The workshop constituted a highly productive series of presentations and discussions. All speakers made an effort to share their insights regarding the two proposed questions for the workshop. Experts on the neural basis of perception presented the latest findings in neural or brain activity recorded during sensory discrimination tasks and other attention-demanding sensory-guided behavior, while modelers presented the state-of-the-art in modeling stochastic network dynamics or in normative approaches to probabilistic codes. The round table discussion on the second day served to initiate an active dialogue between researchers that laid out the current conceptual framework of perceptual decision making and the ones that challenged this "classical" view. The overall consensus was that a new model of perceptual decision making must include additional factors such as extrinsic variability (for example in the stimulus), trial-to-trial variability of top-down feedback signals, and expectation biases. The importance of stimulus and reward history in experimental tasks was highlighted, together with the importance of the strategy that a subject uses to perform the task. All these effects potentially contribute to neural variability but they are very different in nature from irreducible intrinsic noise sources. For example, in rodent experiments there is evidence that learning continues even in highly trained animals without ever reaching a steady-state performance, thus leading to stochastic behavior.

A main conclusion of the workshop was that it turned out that the time is not yet ripe to converge onto a new unified model of perceptual decision making. Especially the origin and the causal role of neural co-variability (i.e. noise correlations) were hotly debated. Participants reached almost opposite conclusions, with some claiming a crucial, causal role for noise correlations on decision making and behavior, while others considered them almost negligible and merely reflecting ongoing computations. The fact that we could not agree on a new coherent framework highlights even more the importance of bringing scientists with different expertise and seemingly opposing points of view together. Participants were enthusiastic about inspiring comments they received related to their own work. Overall, we concluded that the workshop solved as a starting point that will lead to a continuous and intensive dialogue on the nature of “Noise in Decision Making”.

Outcome

Most of the participants agreed on being filmed and their talks are provided online on the internet: https://www.youtube.com/user/CRMatematica.
The workshop organizers are currently writing a perspective article summarizing the emerging view on “Noise in Decision Making”. We also plan to organize follow-up meetings.

Feedback from the workshop participants

We asked the workshop participants for feedback using an anonymous online evaluation questionnaire. We received 17 responses (from 25 participants in total, excluding the conveners). A detailed summary of the workshop evaluation can be found at https://docs.google.com/forms/d/1R6T5jbPWTckaSPb5txI39N8LSbOnAxsJOB4VqwxRLcs/viewanalytics. The overall scores were very high, confirming the impression that this was a highly successful meeting. For example the aspects related with the venue (rooms, food, location) received an average score of 4.7 (on a scale from 0 to 5), the talks a 4.5.
Filming the talks and providing them as “online talks” on the internet was seen controversial: 65% of the participants are not sure if this is a good idea. The main reasons are that they do not want to post unpublished results or preliminary data.
Selection of participants was conceived as very good (4.7), focused (1.6, where 0 is very focused and 5 very scattered) and merit-based (4.2). The overall score of the workshop was 4.5 (very good), and 14 out of 17 participants would like to participate in a follow-up workshop in a few years. The urgent need for ongoing discussion on this topic and the willingness of the participants to continue a close interaction are a major accomplishment of the workshop. As a concrete example, several participants were enthusiastic about starting new collaborations bringing together experimental and theoretical approaches.

2. Scientific content of the event

Session 1: Neuronal variability and computation

The first two talks in this session addressed the experimental evidence for a particular type of computation occurring in the brain of monkeys during a simple perceptual decision-making task. This computation is an integration of noisy sensory evidence in support of the two possible decisions to be made. Both speakers sought indirect evidence for this type of computation by looking at measures of variability in the number of spikes generated by neurons in an area of the parietal cortex involved in the task. Put simply, they took advantage of the model prediction of the level of variability (or noise) to look for something similar in the actual data.
The second two talks focused on computational models in which input signals to a neuronal population should be faithfully reconstructed by looking at the neuronal output. In this case one finds that an optimal solution yields neurons which fire in a seemingly irregular fashion although the only noise in the system comes directly from the input. This prompted the speakers to suggest that noise should be primarily extrinsic to neuronal circuits.

1. **Anne Churchland** (***Insights about neural computation from analyses of spike-count variance***) talked about behavioral data from simple perceptual decision making tasks with monkeys are commonly fit by so-called drift-diffusion or noisy evidence accumulation models. These models describe an integration of noisy evidence to a bound, at which point a decision is made. The speaker in this session explained how she has sought for evidence of such a process by looking at spiking data from neurons. The main gist is that such a noisy evidence accumulation model provides tight constraints on the variability of the spike-count from neurons across trials. Specifically, in such a process, the variance in the spike-count across trials should increase linearly. The spiking data from monkeys shows a similar increase, ruling out certain alternative models and building support for an evidence-accumulation process.

2. **Mike Shadlen** (***Firing rate autocorrelation as a signature of noisy evidence accumulation***) gave a talk which was a continuation of the previous one by Churchland. Shadlen described another measure of variability in the neuronal spike counts which can be used to build further support for an underlying process of noisy evidence accumulation. Specifically, while the previous speaker focused on variability in the spike-count of one neuron during the same epoch of time across trials, this speaker looked at variability between different epochs of time. He constructed matrices of the covariances in the spike-count of a single neuron across trials and at different epochs of time. Again, the drift-diffusion or noisy evidence accumulation model predicts a definite form for this covariance matrix. Namely, the matrix has a banded structure in which the covariance drops off away from the diagonal. The experimentally constructed matrices were in qualitative agreement with this type of structure.

3. **Christian Machens** (***Some new insights on tuning in neural populations***) asked the question: "Is there a simple computational principle which can explain the shape of neuronal tuning curves seen in experiment?" The computational principle he explores is that of optimal linear decoding. The central idea is that a population of neurons receives an particular input,

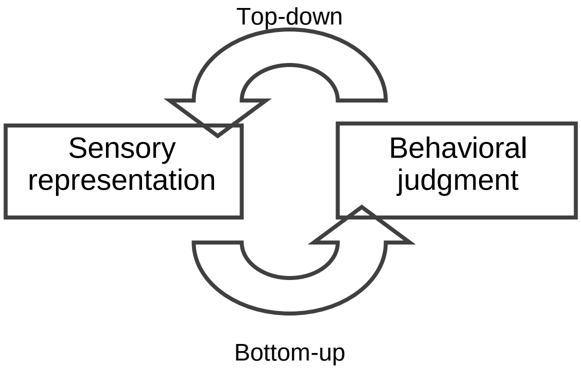

Figure 1. Network scheme used in many talks illustrating the so-called bottom-up connections from sensory areas to decision-related areas, and the top-down connections going the opposite direction.

generating some activity in each neuron. Now the task is to read out the activity of the neurons in such a way as to recover the initial input signal. If the readout is linear then a simple minimization procedure gives the solution, which is to generate a wide array of neuronal tuning curves. He cited examples from the oculo-motor system, the cricket motion-sensing system and neurons in primary visual cortex.

4. **Sophie Denève** (***Learning optimal spike-based representation using predictive coding***) gave an extension of the computational techniques introduced in the previous talk by Machens. Denève showed how the linear decoding framework could be used when neurons are not merely modeled as continuous rate variables, but rather also taking into account the fact that they generate spikes. Once again the task of the computational population of neurons is to track a particular input. In order to do this optimally, it turns out that the timing of the spikes generated by each neuron while seemingly highly irregular is actually very precise. In this case the irregularity of the spiking activity is actually dictated by the irregularity of the input itself since the model has no intrinsic source of noise.

Session 2: Impact of variability on behavior / decision making

1. Néstor Parga (*Noise correlations in decision-making tasks*) showed how noise correlations between frontal lobe neurons are related to detection performance in a vibrotactile detection task. His results suggest that behavioral outcomes are crucially affected by the state of cortical networks before stimulus onset times. Additionally, he introduced a generalization of choice probability, which often use to quantify the relationship of single neurons and behavior, to pairs of neurons.

2. Bruce Cumming (*Measuring interneuronal correlations to understand choice Probability*) introduced two interpretations of the relationship between firing rates of sensory neurons and behavioral choice: in the bottom-up interpretation variability of sensory neurons causes perceptual choices, whereas in the top-down interpretation it reflects the perceptual choices. He then went on to explore the structure of noise correlations predicted by the bottom-up vs. top-down interpretation. Multi-electrode recordings from primary visual cortex of monkeys performing a coarse direction discrimination task showed a correlation structure that could largely be explained with the top-down interpretation.

3. Klaus Wimmer (***Stimulus fluctuations together with top-down feedback can account for the dynamics of choice probabilities***) presented a computational model of perceptual decision making that took into account bottom-up and top-down contributions to choice probability and noise correlations. He suggested that choice probability can be decomposed into an early, bottom-up component that reflects a causal influence of stimulus fluctuations (or sensory variability) on behavior and later, top-down component that is caused by top-down inputs. Both components are linked by the decision dynamics such that they naturally produce sustained choice probability as found across many experiments. The model made several predictions which were verified in a classical data set.

4. Matthias Bethge (*Measuring the contribution of individual neurons to collective decisions*) presented an exact analytical solution for choice probability arising in a feed-forward network. Choice probabilities of sensory neurons are uniquely defined by the correlation matrix and the read-out weights. Findings such as that correlations within pools should be larger than correlations across pools that arose from simulations studies can now be understood and derived in a rigorous framework. As a useful application, Bethge showed how his formula can be inverted such that choice probabilities and pair-wise correlations between neurons can be used to infer the structure of read-out weights.

5. Satu Palva (*Multiscale neuronal dynamics underlying inter-individual variability in behavioral performance*) took a completer view on variability, focusing on fluctuations in human performance on sub-second to minute time-scales (obeying a scale-free dynamics). She showed that evoked MEG / EEG activity in a stimulus detection task was weak and proposed that performance is determined by intrinsic brain dynamics. A neural correlate of the perceptual hit-rat were infra-slow fluctuations in EEG that were also related to inter-subject differences in performance. On a shorter time-scale, intrinsic oscillations were predictive of subject’s performance in a visual working memory task.

6. Catherine Tallon-Baudry (*Some specificity of conscious decisions*) measured MEG activity in subjects performing a visual stimulus detection task and found that pre-stimulus fluctuations in the gamma range seems to reflect a decision bias, whereas fluctuations in the alpha range correlated with an attention-driven baseline shift. In a further experiment she used TMS to stimulate the FEF with different frequencies and found that 50 Hz pulses specifically lead to a decision bias. Finally, she investigated the potential impact on bodily signals on decision making. Specifically, she showed that fluctuations in neural responses to heartbeats before stimulus onset partly predicted stimulus detection.

Session 3: Expectation, priors and learning

In this very focused session speakers presented data showing the dependence of decision on variables other than the current stimulus evidence (e.g. trial-history and on statistics of the stimulus timing. The last three speakers presented models which could describe certain properties of the data (e.g. the time-course of choice probabilities, the time-course of psychophysical kernels, etc) or which could explain the behavioral variability found in experiments.

1. Hendrike Nienborg (*Past decisions partially account for present variability in sensory neurons and behavior*) presented a new analysis on old data consisting on unit recordings from monkey V2 performing a disparity discrimination task (Nienborg & Cumming, 2009). In her previous work published in 2009, she convincingly showed that part of the correlation between neural variability in V2 and behavioral choices (i.e. choice probability) was due to top-down signals coming from higher visual or non-vial areas. The question she was asking this time is whether the nature of that top-down signal is (1) a pre-decision signal meaning is that is originated even before the stimulus is presented (e.g. an expectation signal) or (2) a post-decision signal sent to sensory areas once a the decision had been formed via top-down connections. In other words whether the causality of the inputs shown in Figure 1 was (1) first the activation of top-down inputs and then bottom-up or (2) first bottom-up and then top-down. She address this question by first quantifying the dependence of behavior and neural activity on previous trial history. She used a Generalized Linear Model (GLM) to describe the behavioral data based on three terms (i) the last two stimuli, (ii) past rewards and (iii) past choices. She first found that the two monkeys had a tendency to switch independently of past rewards. She argued that this was due to monkeys picking on the fact the sequence of stimuli was not entirely random but only "pseudorandom" meaning that there was a slight increased probability for the stimuli to switch. She then fitted the same type of GLM model to predict the neural responses and found that past choices and rewards could predict neural responses equally well as behavioral choices. She then asked, "Is the influence of past rewards and choices in choice behavior mediated by the influence in neural responses (i.e. expectation signal)?". By means of clever comparisons between different ways to fit the GLM model she concluded that this was not the case. Finally she asked the contrary "Can the influence of past choices/rewards on neural responses occur "via" the influence in behavioral choice?". She concluded from the analysis that the past history influence on neuronal variability could be fully explained by its influence on choice, suggesting that the nature of the top-down signal was post-decision.

2. Laura Busse (*Influence of expectations on decision making*) presented recent published work (Katzner et al, J. of Vision 2012). In that study, she had investigated one of the key features of active perception, which is the ability to predict critical sensory events. She showed that subjects can implicitly learn statistical regularities in the timing of events and use them to improve behavioral performance. She used a signal detection approach to investigate whether such improvements in performance result from changes of perceptual sensitivity or rather from adjustments of a response criterion. In a regular sequence of b riefly presented stimuli, human observers performed a noise-limited motion detection task by monitoring the stimulus stream for the appearance of a designated target direction. She manipulated target predictability through the hazard rate, which specifies the likelihood that a target is about to occur, given it has not occurred so far. Analyses of response accuracy revealed that improvements in performance could be accounted for by adjustments of the response criterion; a growing hazard rate was paralleled by an increasing tendency to report the presence of a target. In contrast, the hazard rate did not affect perceptual sensitivity. Consistent with previous research, she also found that reaction time decreases as the hazard rate grows. A simple rise-to-threshold model could well describe this decrease and attribute predictability effects to threshold adjustments rather than changes in information supply. She concluded that, even under conditions of full attention and constant perceptual sensitivity, behavioral performance can be optimized by dynamically adjusting the response criterion to meet ongoing changes in the likelihood of a target.

3. Encarni Marcos (***Neural variability in premotor cortex is modulated by trial history and influences motor decision***) talked about recent published work (Marcos et al, Neuron 2013). Marcos analyzed and modeled data from dorsal premotor cortex from monkeys performing a countermanding task where subjects sometimes have to cancel a planned movement. Marcos found that the behavioral response of the monkey depended on the history of previous trials. In particular the reaction time decreased when the previous trial was a GO trial. The reaction time was even lower when the last two trials had been Go trials. She then tested whether a signature of a known trial history effect might be evident in the neural activity of premotor neurons. She quantified the trial-to-trial variability of the spiking activity of single neurons the Variance of the Conditional Expectation (VarCE), a method to isolate the spike variance solely due to variability in the firing rate (Churchland et al 2010). Her main finding was that trials that were just after a “stop” trial were highly variable (i.e. large VarCE): some had a much higher-than-average firing rate, and some a much lower-than-average firing rate. In contrast, trials preceded by a GO trial, or by two GO trials, had lower and lower VarCE, respectively. Encarni’s dataset is particularly intriguing because the two conditions she compared had nearly identical firing rate means. By examining VarCE, she was able to uncover a neural mechanism that would have been invisible using traditional analyses. She finally presented a model of an attractor neural network which could somehow capture the behavior observed in the data by means of an external signal to the network which varied depending on the previous history. Marcos results highly the importance of history dependence on movement planning and how this dependence affects the way in which this commands are encoded in the neural activity by means of changes in its variability.

4. Ralf Haefner (*A normative explanation for task-dependent correlations and choice probabilities based on the neural sampling hypothesis*) presented unpublished work from a recent model which implements the sampling hypothesis and can account for many of the observations on neural activity described in two alternative forced choice tasks (2AFT). He first presented the increasingly popular view describing perception as a probability inference problem. Thus, perception is not simply about estimating what see when we look at an image for instance, but about estimating the full probability distribution of the things we could be seeing. Having the complete probability distribution can lead to estimating the most likely stimulus but also to more sophisticated computations such as estimating the most rewarding choice, etc. He proposed that the brain acts as a Generative model to estimate the conditional probability distribution which allow us to perform inferences about unobserved variables given observed ones (e.g. by observing that the grass is wet, we infer the probability that it rained or the probability that the sprinkler was recently turned on). Computing these conditional probabilities exactly is, for most realistic situations, a computationally intractable task. Inference sampling is a method used in machine learning to compute these conditional probabilities. Haefner proposed that the brain computes these probabilities performing neural sampling of the probability distributions (Fiser et al 2010). The neural sampling hypothesis is realizable in natural conditions (continuous time, asynchronous firing, …) and yields certain implications: Poisson variability, tuning curves.

Haefner presented a model of the visual system implementing a generative model carrying out the neural sampling hypothesis and obtained a number of results consistent with the data : sustained choice probabilities, non-monotonic psychophysical kernels and a correlation matrix which depends on the task as observed in the experiments.

5. Zach Mainen (***Contribution of online learning to variability in perceptual decisions***) presented a new analysis performed on old data from his already classic two choice odor discrimination task (Uchida & Mainen, Nat. Neurosci 2003). With this new analysis he showed that there was a dependence of the current choice on previous trial history. This dependence was larger when the present stimulus was closer to the discrimination threshold and negligible when it was a stimulus easy to discriminate (far from threshold). He used the standard accumulation to boundary model to conclude that the dependence was due to a constant readjustment of the discrimination boundary. In other words, animals kept learning the task once they had reached a steady state performance such that the discrimination rule was always being updated based on the recent trials. Mainen convincingly showed that animals might not always be performing the task we think they perform but, aspects of the task which we may find trivial (e.g. the fact that the task is always the same and the rules do not change) are not clear to the animal.

6. Xiao-Jing Wang (*Categorization learning depends on choice-correlated variability of mixed-selective neurons*) presented a new model to account for a categorization task performed in monkeys (Assad & Freedman, Nature 2006). His model was a hierarchical network composed of a sensory circuit (MT), an evidence accumulator circuit (LIP) and a category circuit (PFC). Connections between circuits were feedforward except between the PFC and LIP which were also top-down. He used a plasticity leaning rule which modified the weights between LIP and PFC (both feed-forward and feedback) and showed that feedback connections were necessary to achieve full performance in the task. He also showed that the model could account for the relation between neuronal sensitivity and choice probability found in many discrimination tasks.

7. Alexandre Pouget (*Not noisy, just wrong: the role of suboptimal inference in behavioral variability*) presented ongoing work in collaboration with Peter Latham and work from a recent paper (Beck et al, Neuron 2012). In the first part he argued that people keep discussing about the size of noise correlations whereas what really determined their impact on coding is their relation with the structure of tuning curves. In particular he talked about he calls f * f' correlations, which are correlations in the "direction where most impair encoding" and argued that this could be extremely small and yet have a tremendous impact in bounding the fisher information of a population of cells. He said that this correlations could be "masked" but large correlations of the kind which have no impact because tuning curves are generally heterogeneous (Ecker et al, J. Neurosci. 2012) implying that by measuring correlations alone was almost impossible to determine their impact on coding. The only way to determine their impact would be to record from hundreds of neurons simultaneously and try to decode the stimulus. If decoding increases monotonically with the number of neurons used, then correlations had no impact, whereas is decoding performance decreases with the number of neurons, then correlations had a deleterious impact bounding the discrimination error. In the second part of the talk he proposed that behavioral variability is not due to intrinsic sources generating large amplitude noise but to an amplification of small noise present already at the receptors level due to sub-optimal inference. In particular Pouget argued that the brain might not always be doing an optimal job inferring the correct stimulus but the use of suboptimal filters in decoding the stimulus might lead to a large amplification of small fluctuations and to a variable response.

Session 4: Network dynamics and brain state

The first two speakers discussed large-scale brain activity. In particular they discussed how the anatomical and functional connectivity of a large number of brain areas can be used to build a dynamical model, which in turn can be compared to experimental data from fMRI. The role of noise in such models was discussed by the second speaker.
There was some discussion regarding the complexity of the models used. The first speaker, for example, discussed results from a model incorporating anatomical connectivity between 50.000 brain areas, with a complex dynamical system at each node. This prompted a question about the usefulness of a model with such a high-dimensional parameter space.

1. Petra Ritter (***Brain states: models, maps, and behavior***) gave an overview of a project called the "Virtual Brain" project, which makes use of anatomically data of human brain connectivity collected via Diffusion Tensor Imaging (DTI). The "Virtual Brain" project allows the user to simulate dynamics on the network extracted via DTI by inserting a dynamical system at each node. The resulting dynamics can be compared to data collected from fMRI of the same human subjects.

2. Gustavo Deco (***The importance of being balanced***) discussed a similar model to that of the first speaker, although with many fewer brain areas. Deco was particularly focused on reproducing the slow fluctuations seen in the so-called "default networks" of the human brain: various clusters of brain areas which show synchronized activity in fMRI studies.
He argued that one can reproduce such "default network" dynamics by using the network of functional connectivity extracted from fMRI studies, coupled with a network of spiking neurons at each node. Importantly, the "default network" activity seen in the model reflects synchronization between different nodes of the fluctuations in the activity at each node about a stationary state of activity.

3. Alfonso Renart (*Competitive dynamics during spontaneous activity in cortical circuits*) presented follow up work on his recent article about the Asynchronous State in cortical circuits (Renart et al, Science 2010). He first showed that the distribution of correlation coefficients of pairs of cortical neurons (from auditory and somatosensory cortices) is centered at zero and exhibits large positive and negative tails. The time scale of these correlations was also relatively slow (~50-100 ms). He presented a new method to organize the matrix of correlation coefficients termed the Similarity matrix, defined as the correlation between correlation coefficients. This matrix showed two clear groups of cells showing strong similarity between pairs of neurons of the same group and strong but negative similarity between pairs across groups. Renart showed that this type of structure resemble the competition between two large groups of cells. This structure is not observed in randomly connected balanced networks that can only account for the small mean of correlations. To account for the competitive structure Renart proposed an extension from the balanced network model into a network with two excitatory populations which competed via strong lateral inhibition. This network, commonly used in decision-making networks, exhibits a Pitchfork bifurcation. Below the bifurcation, the stochastic fluctuations make it exhibit the a very similar correlation structure as that found in the data.

4. Ruben Moreno-Bote (*Poisson-like spiking and contrast invariant sampling with probabilistic synapses*) addressed the question of what are the mechanisms that make cortical activity so variable**.** This is naturally a relatively old question and it was what motivated what we called a *balanced network.* Originally a balanced network was proposed to explain how a large network where neurons receive a large number of inputs can produce spike trains with large stochasticity. Moreno-Bote argued that this is the case at low firing rates, but the spike count Fano factor in a balanced network decreases as neurons are driven stronger and fire at higher firing rates. He argued that experimental data from evoked cortical responses shows large Fano factors for a very broad range of firing rates and concluded that a balanced network alone cannot reproduced this feature. He proposed that synaptic unreliability can compensate the decrease in Fano factor observed at high firing rates and generate networks that show Fano factors close to one for all rates. He ended suggesting that synaptic stochasticity could be functionally useful for perceptual statistical sampling, the process by which neuronal circuits sample the various interpretations of sensory stimulus with the corresponding probability.

5. Matteo Carandini (*Cortical state and response variability in primary visual cortex*) argued that response variability is mostly of cortical origin. He first presented a model showing that cortical cells are expected to amplify noise because of the nonlinear threshold effect due to their f-I curve. He next presented experimental data from anesthetized cats which demonstrated that noise depends on cortical state. In fact, variability was mostly due to global events (UP- and DOWN-states). The degree by which individual neurons followed the global population activity was heterogeneous (neurons were either “soloists” or “choristers”). Taking this into account, the measured pair-wise noise correlations between neurons could be predicted with high accuracy. Finally, he presented experimental data showing that choristers were more visually responsive and more connected than soloists.

Session 5: Attention and memory

1. Tatiana Pasternak (***Trial-to-trial variability of cortical neurons reveals the nature of their engagement in sensory decision making*)** reported on the dynamics of neuronal variability in prefrontal cortex during a memory-guided motion discrimination task. The task had a strongly predictable timing structure, which may emphasize the variability dynamics, as measured by the Fano Factor averaged over the population of cells. Fano Factors decreased phasically at each sensory stimulation and sustained over the delay period, especially for narrow spiking neurons, putative inhibitory cells. There was an additional dissociation between cells with rising activity during the delay, which reduced their variability more pronouncedly than cells with non-rising delay activity. This data suggests a functional specialization in prefrontal cortical populations for time-dependent control, which may operate through the reduction of variability in subnetworks of the prefrontal circuit.

2. Alexander Thiele *(Attention induced variance and noise correlation reduction in macaque V1 is mediated by NMDA receptors)* presented neurophysiological data from recordings in awake behaving monkeys that dissected the contribution of different synaptic receptors to neuronal variability during selective attention tasks. In a series of experiments with iontophoretic application of synaptic blockers in primary visual cortex while the monkeys attended to visual stimuli inside or outside the recorded neuron's receptive field, he reported that muscarinic acetylcholine receptors were implicated in attentional modulation of firing rates, but not in attentional modulations of neuronal variability (as measured by the Fano Factor). Conversely, iontophoretically applied blockers of NMDA receptors did not affect firing rate modulations by attention but did reduce the change in neuronal variability associated with attentional selection. This dissociation between synaptic receptor and attentional modulation suggests that attention operates independently modifying different aspects of the sensory code: rate tuning and spiking variability.

3. Douglas Ruff (Attention can adaptively increase or decrease interneuronal correlations in V4) presented data from his work in the laboratory of Marlene Cohen investigating how attention modulates population activity in macaque visual cortex. Ruff designed a new experiment to investigate if attention always modulates firing rates and neuronal correlations following an inverse relationship or else if firing rates and neuronal correlations can be flexibly modulated by attention to adapt to the best coding strategy. To this end he trained one monkey in a contrast discrimination task and they recorded simultaneously from multiple neurons responding to two different Gabor patches. He investigated how attention modulated the relationship between noise correlations and signal correlations across pairs of neurons. Ruff found that attention reduced correlations for pairs of neurons with similar selectivity, as previously reported, but it increased correlations for pairs of neurons with dissimilar selectivity. This supports the initial hypothesis that attention flexibly modulates population activity in a way that is best for information coding.

4. Georgia Gregoriou *(Attention and interneuronal correlations)* presented the analysis of neuronal variability and co-variability in cortical areas V4 and FEF in an attention task. Simultaneous recordings across multiple areas permitted the evaluation of how attention modulated correlations within and across cortical areas. Attention reduced Fano Factors and neuronal correlations in both areas immediately following stimulus onset, but more prominently in FEF. During stimulus presentation, correlations between activity in areas FEF and V4 were weaker in the attended compared to non-attended condition. Interestingly, the presentation of the attentional cue resulted in increased noise correlations within and across areas. This data shows that activity correlation across areas is also flexibly adjusted in attention tasks.

The **final discussion** session of the workshop, chaired by Jaime de la Rocha, revolved around trying to specify the elements of a new theoretical framework to understand the role of activity variability and co-variability in brain computations and behavior, based on the content of the workshop presentations. An online document containing a slide with an initial scheme (Figure 2) that included some possible elements, such as networks, feed-forward and feedback connectivity, stimulus variability, etc. was distributed among participants and they were requested to emphasize mechanisms or fill in the missing elements. Subsequently, each participant presented their view briefly while projecting their slide. Very different proposals were formulated, with some advocating for noise being just a result of insufficient experimental knowledge (see Figure 3), while others attributed a causal role for neuronal noise in determining behavior. Several elaborated on the role of top-down feed-back connections on perception (see e.g. Figure 4) and on the dependence of perceptual decisions on the recent past history (see e.g. Figure 5). There was general agreement in the fact that new experiments manipulating sequential stimulus regularities, internal brain states, and other possible sources of systemic trial-to-trial variability should be carried out to clarify the role of neuronal noise in decision making.

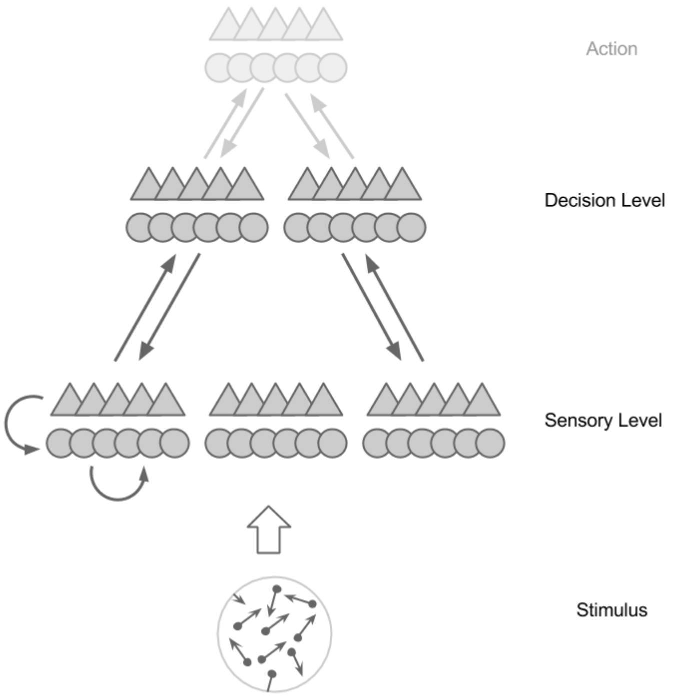

Figure 2. Model template given to the participants to add their contribution to the workshop and to give an opinion on what the future research should focus.

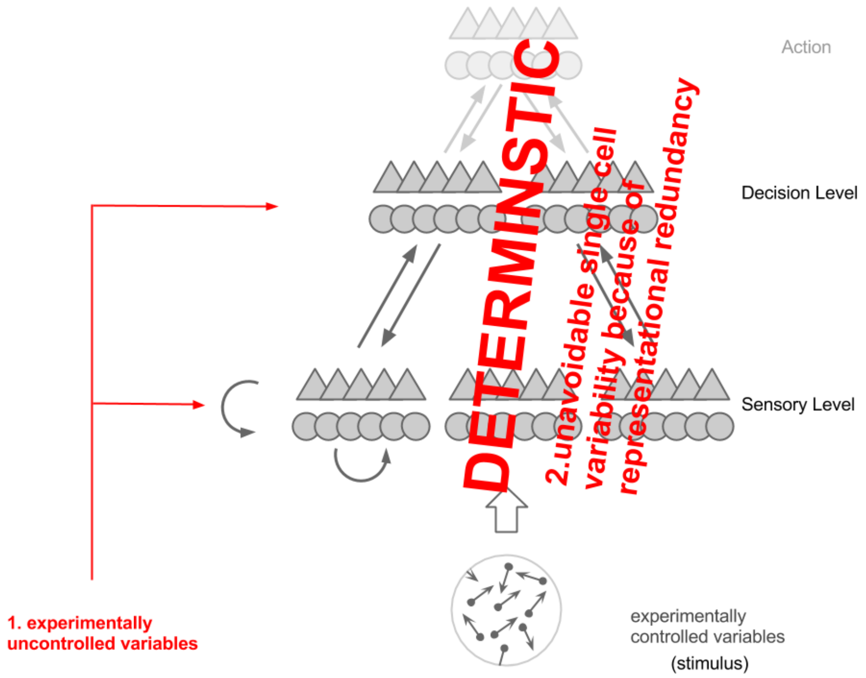

Figure 3. Slide made by one of the participants showing his view on neuronal variability.

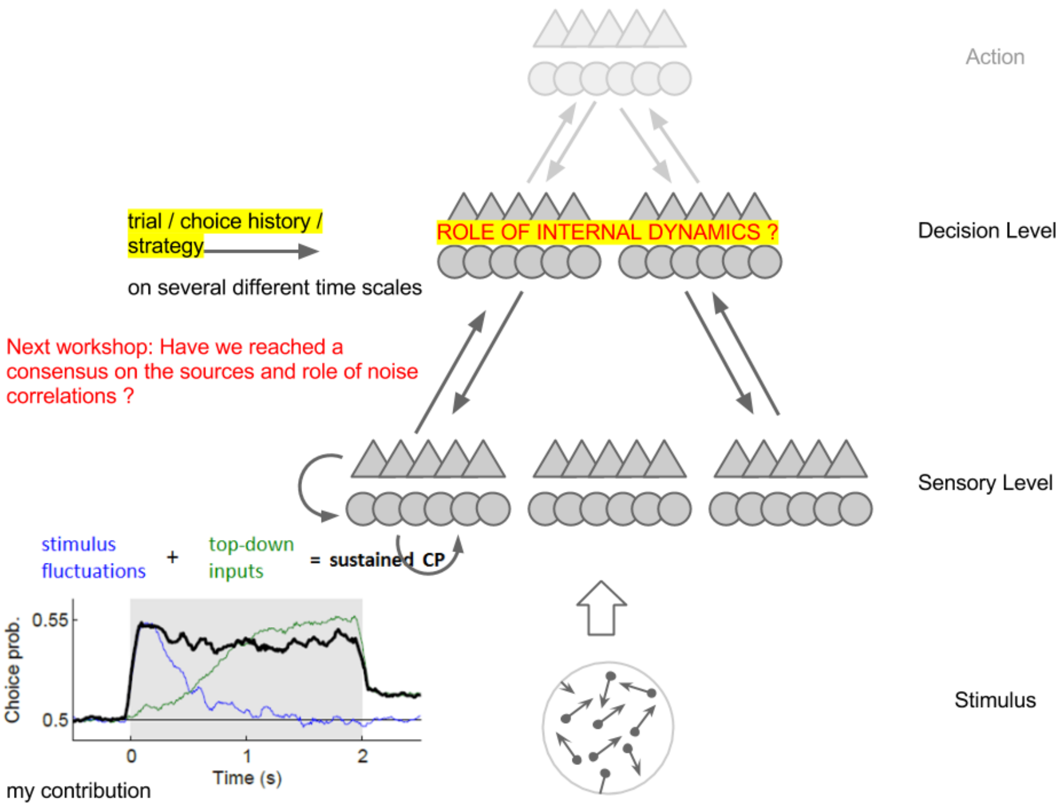

Figure 4. Slide drawn by another particpant with his contribution and his view on future steps.

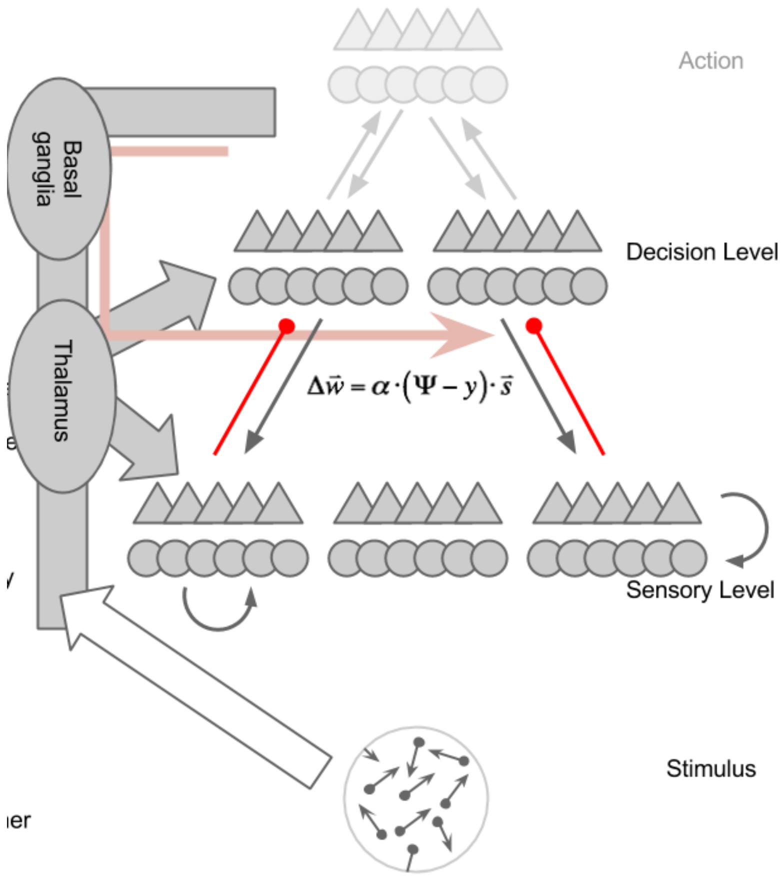

Figure 5. Slide drawn by yet another participant with his view of the model.

3. Assessment of the results, contribution to the future direction of the field, outcome

What was learnt from the workshop and new objectives

On general terms, the workshop aimed at gaining a general view on the current research on the neural mechanisms of behavioral variability. On more specific terms, one objective of the workshop was to see if the experts gathered in this meeting could converge on a new model of how neuronal and behavioral variability are related. To facilitate this convergence, the organizers set beforehand two core questions to guide discussions during the meeting and asked the participants to design their contributions according to them (see Executive Summary above). Participants were selected with a criterion to represent the diversity of approaches used in Neuroscience to address a focused topic such as the neural origin of behavioral variability. The presentations of the workshop were of the highest quality and very much focused on the core questions from a variety of perspectives, thus covering extensively the current views on the topic in Neuroscience. The resulting general view, however, emphasized the multiple factors involved in relating neuronal and behavioral variability, their insufficient experimental and theoretical characterization, and thus the difficulty of converging at this point on a new extended neural model of behavioral variability.

<u>A multiplicity of factors</u>

While there was general agreement that at least part of the neural variability observed in cortical neurons would be due to computations carried out in downstream areas being fed back onto sensory cortices via feedback connections, the relative contributions of various such top-down signals (attention, expectation/anticipation, decision, etc.), of variability being intrinsically generated in the circuit (competitive dynamics, non-optimal inference, learning, ...), or of variability being inherited from sensory input irregularity was unclear. It appeared like some critical experiments to identify the major elements contributing to neural variability in sensory cortices were still pending. Likewise, the specific mechanisms for the generation of neuronal correlations in neuronal networks were still not sufficiently understood at the theoretical level. The general feeling was that the topic of neural variability and its relation to behavioral variability, far from being an old problem largely resolved in classic studies, is a contemporary, open problem that deserves a lot of attention from both experimentalists and theoreticians. The workshop was very useful in setting the various perspectives side-by-side, and in identifying new critical objectives to advance in specifying the neural mechanisms of behavioral variability at the various research levels represented in the meeting.

<u>Convergence towards a new model</u>

The general view gained from the talks was further distilled in two general discussion sessions to see if we could converge to a new conceptual model. The conclusion was that many of the possible elements that were to be considered in this new model were still insufficiently characterized at the experimental and theoretical levels. Radically different models could still be seen as supported by current data. Workshop participants agreed that achieving a new consensus around a revised conceptual model would be an important objective for future research, but this did not seem possible at this point given our current knowledge.

<u>Experiment-theory</u>

The workshop sought to bring together computational and experimental researchers to foster exchanges across disciplines and cross-fertilization. This objective was attained, since this specific interaction was mentioned as particularly inspiring: some of the experimenters declared that the theoretical work presented at the meeting was conditioning their own approaches to these scientific issues, and some bilateral collaborations between computational and experimental labs were initiated at the meeting.

Multidisciplinarity
An effort was put to gather scientists that addressed the issue of behavioral variability from different perspectives, from single-cell perspectives, to whole-brain analyses or behavioral studies. A diversity of theoretical approaches were also present, with neural dynamics perspectives side-by-side with probabilistic models and normative generative approaches. This multidisciplinarity around a precisely focused topic was valued as a strength of the workshop that set it apart from other scientific events, and many participants found this "*focused diversity*" highly inspiring.

Contribution to future direction of the field
The workshop is expected to have an impact on the future direction of the field, given the fact that the participants in the event were a significant fraction of the most active scientists in this field of research worldwide. On the one hand, the reduced dimension of the meeting promoted closer interactions and more focused discussions, a main objective of ESF Exploratory Workshops. On other hand, this same factor could limit the workshop's impact on the long run. Indeed, relevant scientists interested in attending the workshop had to be turned down prior to the event based on the ESF rules for Exploratory Workshops. In order to mitigate the effects of the reduced dimension of the meeting, we inscribed the event in the European Month of the Brain initiative (http://ec.europa.eu/research/conferences/2013/brain-month/index_en.cfm?pg=home) to gain visibility, and we designed a strategy to reach to a broader audience by videotaping workshop talks and making them available on-line (see Follow-up actions below).

Follow-up actions
Several follow-up actions were decided at the meeting or are intended to be initiated by the organizers:

- We ran an anonymous survey among workshop participants to evaluate the event post-hoc. Seventeen participants filled out the on-line form remotely. The survey reveals that the location and logistics of the meeting were very highly valued (>70% gave the highest mark to these aspects). The cultural and social events were positively evaluated (> 70% gave marks ≥3 in a scale 0-5). The talks were perceived as high quality (>90% gave marks ≥4 in a scale 0-5), and the selection of participants was praised (>90% gave marks ≥4 in a scale 0-5). On the other hand, while 30% of participants felt that the program was balanced, a significant fraction (~50%) considered that it was too busy for various possible reasons (too short for the number of talks, too many participants, too many non-scientific events, insufficient discussion,...). Overall, the workshop was scored ≥4 in a scale 0-5 by all survey participants. The few criticisms received revolved around the necessity of more structured discussions.
- In order to share with a broader public the general view on the various perspectives regarding the relation between neural and behavioral variability, most of the talks presented at the workshop were videotaped and have been made available publicly at https://www.youtube.com/user/CRMatematica. This link has also been publicized through the Connectionists distribution list (http://grey.colorado.edu/Connectionists) that spreads this information over a broad community of computational scientists. The opportunity of disseminating the content of workshops through video-talks was also queried in the on-line survey that we ran post-hoc. A majority of responses were uncertain as to whether this was a good idea (65%), while only a minority (29%)

considered it a good idea. The reasons were mostly that speakers do not want to post publicly unpublished results or preliminary, yet unconfirmed, data.

- During the meeting we also discussed the opportunity of organizing a new workshop on this topic to evaluate periodically the advance towards the objective of a new revised model for behavioral and neural variability. We asked about this anonymously in our on-line survey and the result was that 85% of participants would be interested in participating in a follow-up workshop. We are therefore committed to help gather again experts around this topic in a few years.
- The organizers intend to write a perspective article on the topics discussed in the meeting and the conclusions achieved in a high-impact journal. This is an on-going project at this point.

4. Final programme

PROGRAMME

Tuesday, 28 May 2013

Afternoon	*Arrival*
18.30-20.00	*Visit of the monastery*
20.00-	*Reception and dinner (at the Món restaurant)*

Wednesday, 29 May 2013

09.00-09:40	**Welcome Session**
09.00-09.10	**Presentation of the European Science Foundation (ESF)** **Giovanni Pacini** (ESF Science Review Group for the Bio-Medical Sciences)
09.10-09.40	**"Why are we here?"** **Albert Compte** (IDIBAPS, Barcelona, Spain)
09.40-12.00	**SESSION 1: Neuronal variability and computation** **Session chair: Alex Roxin**
09.40-10.10	**Insights about neural computation from analyses of spike count variance** **Anne Churchland** (Cold Spring Harbor Lab., United States)
10.10-10.40	**Firing rate autocorrelation as a signature of noisy evidence accumulation** **Michael Shadlen** (Columbia Univ., New York, United States)
10.40-11.00	*Coffee / Tea Break*
11.00-11.30	**Some new insights on tuning in neural populations** **Christian Machens** (Champalimaud Found., Lisbon, Portugal)
11.30-12.00	**Learning optimal spike-based representations using predictive coding** **Sophie Denève** (ENS, Paris, France)
12.00-13.30	*Visit of the Alicia Foundation*
13.30-15.00	*Lunch*
15.00-19:00	**SESSION 2: Impact of variability on behavior / decision making** **Session chair: Petra Ritter**
15.00-15.30	**Noise correlations in decision-making tasks** **Néstor Parga** (Univ. Autónoma de Madrid, Spain)
15.30-16.00	**Measuring interneuronal correlations to understand choice Probability** **Bruce Cumming** (NIH, Bethesda, United States)
16.00-16.30	**Stimulus fluctuations together with top-down feedback can account for the dynamics of choice probabilities** **Klaus Wimmer** (IDIBAPS, Barcelona, Spain)
16.30-17.00	*Coffee / tea break*
17.00-17.30	**Measuring the contribution of individual neurons to collective decisions** **Matthias Bethge** (Univ. Tübingen, Germany)

17.30-18.00	**Multiscale neuronal dynamics underlying inter-individual variability in behavioral performance** **Satu Palva** (Univ. Helsinki, Finland)
18.00-18.30	**Some specificity of conscious decisions** **Catherine Tallon-Baudry** (CNRS, Paris, France)
18.30-19.00	**Discussion**
19.30	*Dinner*

Thursday, 30 May 2013

09.00-13:00	**SESSION 3: Expectation, priors and learning** **Session chair: Tatiana Pasternak**
9.00-9.30	**Influence of expectations on decision making** **Laura Busse** (Univ. Tübingen, Germany)
9.30-10.00	**Past decisions partially account for present variability in sensory neurons and behavior** **Hendrikje Nienborg** (Univ. Tübingen, Germany)
10.00-10.30	**Neural variability in premotor cortex is modulated by trial history and influences motor decision** **Encarni Marcos** (UPF, Barcelona, Spain)
10.30-11.00	*Coffee / tea break*
11.00-11.30	**A normative explanation for task-dependent correlations and choice probabilities based on the neural sampling hypothesis** **Ralf Haefner** (Central European Univ., Budapest, Hungary)
11.30-12.00	**Contribution of online learning to variability in perceptual decisions** **Zach Mainen** (Champalimaud Found., Lisbon, Portugal)
12.00-12.30	**Categorization learning depends on choice-correlated variability of mixed-selective neurons** **Xiao-Jing Wang** (NYU, New York, United States)
12.30-13.00	**Not noisy, just wrong** **Alexandre Pouget** (Univ. de Genève, Switzerland)
13.00-14.30	*Lunch*
14.30-18.30	**SESSION 4: Network dynamics and brain state** **Session chair: Sophie Denève**
14.30-15.00	**Brain States: models, maps and behavior** **Petra Ritter** (Charité, Berlin, Germany)
15.00-15.30	**The importance of being balanced** **Gustavo Deco** (UPF, Barcelona, Spain)
15.30-16.00	*Coffee / Tea Break*
16.00-16.30	**Competitive dynamics during spontaneous activity in cortical circuits** **Alfonso Renart** (Champalimaud Found., Lisbon, Portugal)
16.30-17.00	**Poisson-like spiking and contrast invariant sampling with probabilistic synapses** **Rubén Moreno-Bote** (Found. Sant Joan de Déu, Barcelona, Spain)
17.00-17.30	**Cortical state and response variability in primary visual cortex** **Matteo Carandini** (UCL, London, United Kingdom)
17.30-18.30	**Discussion on follow-up activities/networking/collaboration**

18.30-	*Banquet dinner*

Friday, 31 May 2013

09.00-11.00	**SESSION 5: Attention and memory** **Session chair: Bruce Cumming**
09.00-9.30	**Trial-to-trial variability of cortical neurons reveals the nature of their engagement in sensory decision making** **Tatiana Pasternak** (Univ. Rochester, NY, United States)
9.30-10.00	**Attention induced variance and noise correlation reduction in macaque V1 is mediated by NMDA receptors** **Alex Thiele** (Univ. Newcastle, United Kingdom)
10.00-10.30	**Attention can adaptively increase or decrease interneuronal correlations in V4** **Douglas Ruff** (Pittsburgh Univ., United States)
10.30-11.00	**Attention and interneuronal correlations** **Georgia Gregoriou** (Univ. of Crete, Greece)
11.00-11.10	*Coffee / Tea Break*
11.10-12:30	**Closing Session**
11.10-12.30	**Summary remarks and discussion on follow-up activities/networking/collaboration** **Jaime de la Rocha** (IDIBAPS, Barcelona, Spain)
12.30-14.00	*Lunch*
14.00	*End of Workshop and departure*

5. Final list of participants

1. Matthias BETHGE, Centre for Integrative Neuroscience, Tübingen, Germany
2. Laura BUSSE, Centre for Integrative Neuroscience, Tübingen, Germany
3. Matteo CARANDINI, University College London, United Kingdom
4. Anne CHURCHLAND, Cold Spring Harbor Laboratory, NY, United States
5. Albert COMPTE, IDIBAPS, Barcelona, Spain
6. Bruce CUMMING, National Eye Institute, Bethesda, MD, United States
7. Jaime DE LA ROCHA, IDIBAPS, Barcelona, Spain
8. Gustavo DECO, Universitat Pompeu Fabra, Barcelona, Spain
9. Sophie DENÈVE, Group for Neural Theory, Paris, France
10. Georgia GREGORIOU, University of Crete, Heraklion, Greece
11. Ralf HAEFNER, Central European University, Budapest, Hungary
12. Christian MACHENS, Champalimaud Neuroscience Programme, Lisbon, Portugal
13. Zach MAINEN, Champalimaud Neuroscience Programme, Lisbon, Portugal
14. Encarni MARCOS, Universitat Pompeu Fabra, Barcelona, Spain
15. Rubén MORENO-BOTE, Fundació Sant Joan de Déu, Barcelona, Spain
16. Charvy NARAIN, Nature Publishing Group, London, United Kingdom
17. Hendrikje NIENBORG, Centre for Integrative Neuroscience, Tübingen, Germany
18. Satu PALVA, University of Helsinki, Finland
19. Néstor PARGA, Universidad Autónoma de Madrid, Spain
20. Tatiana PASTERNAK, University of Rochester, NY, United State
21. Alexandre POUGET, University of Geneva, Switzerland
22. Alfonso RENART, Champalimaud Neuroscience Programme, Lisbon, Portugal
23. Petra RITTER, Charité , Berlin, Germany
24. Alex ROXIN, Centre de Recerca Matemàtica, Barcelona, Spain
25. Douglas RUFF, University of Pittsburgh, PA, United States
26. Michael SHADLEN, Columbia University, New York, NY, United States
27. Catherine TALLON-BAUDRY, Ecole Normale Supérieure, Paris, France
28. Alex THIELE, University of Newcastle, United Kingdo
29. Xiao-Jing WANG, New York University, NY, United States
30. Klaus WIMMER, IDIBAPS, Barcelona, Spain

6. Statistical information on participants

Country representation

Finland: 1

France: 2

Germany: 4

Greece: 1

Hungary: 1

Portugal: 3

Spain: 8

United Kingdom: 3

United States: 6

Switzerland: 1

Gender repartition

Male: 19

Female: 11