SOBP Computational Psychiatry Satellite 2018

Reinforcement Learning

May 9th, 8am
room Nassau
Hilton Midtown
New York City

Please provide feedback here:
https://tnusurvey.ethz.ch/index.php/472246

Organizers
Michael Browning, DPhil MBBS, University of Oxford
Quentin Huys, MD PhD, University of Zürich & ETH Zürich
Martin Paulus, MD PhD, Laureate Institute for Brain Research
Contact: tcpw@quentinhuys.com

Program

8 - 8.50	Coffee & breakfast	
8.50	Opening Remarks	Michael Browning
9 - 10.45	**Reinforcement Learning - theory**	Yael Niv
10.45 - 11.15	Coffee break	
11.15 - 11.45	**Data in search of model**	Contributed talks
	Andreas Frick	
	Merage Ghane	
	Verena Ly	
	Laurel Morris	
	Brooke Staveland	
	Robert Whelan	
11.45 - noon	**Model in search of data**	Contributed talks
	Manish Saggar	
	Shan Siddiqi	
noon - 1.30	Lunch break	
1.30 - 3.15	**Model fitting**	Michael Frank
3.15 - 3.45	Coffee break	
3.45 - 5.30	**RL applications to mental health**	Quentin Huys
5.30 - 6	Discussion	
6	Closing Remarks	Martin Paulus

Speakers

Michael Frank, PhD, Brown University
Andreas Frick, PhD, Stockholm University
Merage Ghane, MA, Virginia Tech/National Institute of Mental Health
Quentin Huys, MD PhD, University of Zürich & ETH Zürich
Verena Ly, PhD, Leiden University
Laurel Morris, Cambridge University
Yael Niv, PhD, Princeton University
Manish Saggar, PhD, Stanford University
Shan Siddiqi, MD, Harvard Medical School, MGH/McLean
Brooke Staveland, Stanford University
Robert Whelan, PhD, Trinity College Dublin

Participants

Rany Abend, PhD, *NIMH*
Brendan Adkinson, *Yale University School of Medicine*
Ben Andrew, *Austen Riggs Center*
Tali Ball, PhD, *Stanford University*
Daniel Barron, MD, *Yale University*
David Benrimoh, MD, CM, *UCL/McGill*
Isabel Berwian, *University of Zurich and ETH Zurich*
Caroline Bévalot, *Hospices Civiles de Lyon, France*
Venkat Bhat, MD, *University of Toronto*
Vanessa Brown, *University of Pittsburgh*
Mingbo Cai, PhD, *Princeton University*
Pablo Carillo, *Hospices Civiles de Lyon, France*
Linda L Carpenter, MD, *Brown University/Butler Hospital*
Henry Chase, *University of Pittsburgh*
Alexandra Cohen, PhD, *New York University*
Jessica Cooper, PhD, *Emory University*
Zhuoya Cui, *Virginia Tech*
Marilyn Cyr, PhD, PsyD, *New York State Psychiatric Institute*
Daniel Davies, PhD, *University of Cambridge*
Lorenz Deserno, PhD, *MPI for Humanan Cognitive and Brain Sciences Leipzig*
Stefan Ehrlich, MD, *Technische Universität Dresden*
Angela Fang, PhD, *Harvard Medical School*
Morgan Flynn, PhD, *Yale University School of Medicine*
Greg Fonzo, PhD, *Stanford University School of Medicine*
Anna Frey, *University of Reading*
Andreas Frick, *Stockholm University*
Bowen Fun, PhD, *Caltech*
Ken Garcia, MD,
Merage Ghane, MA, *Virginia Tech/NIMH*
Reza Hosseini Ghomi, MD, MSE, *University of Washington*
Jessica Gilbert, PhD, *NIMH*
Jimmie Gmaz, *Dartmouth College*
Jackie Gollan, PhD, *Northwestern University*
Jerry Guintivano, PhD, *University of North Carolina at Chapel Hill*
Christopher Hammond, MD PhD, *Johns Hopkins University*
Jessica Hua, MA, *University of Missouri*
Noboru Hiroi, PhD, *Albert Einstein College of Medicine*
Olof Hjorth, MSc, *Uppsala University Sweden*
Philipp Homan, *The Feinstein Institute for Medical Research*
Matthew Hudgens-Haney, PhD, *UT Southwestern Medical Center*
Eugene Young-Chul Jung,
Claire Kaplan, MS, *University of Maryland College Park*
Dr. Teresa Katthagen, *Charité-Universitätsmedizin Berlin*
Alfred Kaye, *Yale University*

Hanna Keren, PhD, *NIMH*
Donald Klein, MD DSc, *NYU*
Lisa Klevebrant, *Uppsala University*
Antonija Kolobaric, *Yale University School of Medicine*
Anna Konova, PhD, *New York University*
Stephanie Korenic, *Maryland Psychiatric Research Center*
Soyoung Kwon, PhD, *Caltech*
Jakub Kraus, *Karolinska Institute*
Poornima Kumar, PhD, *McLean/Harvard Medical School*
Ellen Leibenluft, MD, *NIMH*
John Erik Leikauf, MD, *Stanford University School of Medicine*
Dr. Julia Linke, *NIMH*
Falk Lohoff, MD, *NIAAA*
Verena Ly, PhD, *Leiden University*
Anahit Mkrtchian, *UCL/NIMH*
Andra Mihali, *New York University*
Laurel Morris, PhD, *Mount Sinai School of Medicine*
Flora Moujaes, *Yale University*
Dahlia Mukherjee, PhD, *Penn State College of Medicine*
Dr. Akshay Nair, *UCL - HD Max Planck Centre*
Kate Nussbaum, *New York University*
Georgia O'Callaghan, PhD, *NIMH*
Katie O'Connell, *Georgetown University*
David Pagliaccio, *New York State Psychiatric Institute*
Madeleine Payne, *University College London*
Maria Picó Pérez, *Bellvitge University Hospital-IDIBELL*
Daniel Pine, MD PhD, *NIMH*
Alexandra Pike, *University of Oxford*
Teodor T. Postolache, MD, *University of Maryland School of Medicine*
Hillary Raab, *New York University*
Angela Radulescu, *Princeton University*
Laura Rai, *Trinity College Dublin*
Vijay Ramchandani, PhD, *NIAAA/NIH*
Jenna Reinen, PhD, *Yale University / IBM Research*
Gail Rosenbaum, PhD, *New York University*
Manish Saggar, *Stanford University*
Betty Jo Salmeron, MD, *NIDA-IRP-Neuroimaging Research Branch*
Florian Schlagenhauf, MD, *Charite and Max Planck Institute Leipzig*
Liane Schmidt, *ICM INSERM U1127*
Maria Sekutowicz, MSc, *Charité Campus Mitte*
Mohamed Sherif, MD, MSc, *VA Connecticut / Yale University*
Shan Siddiqi, MD, *Harvard Medical School*
Brooke Staveland, *Stanford University*
Joanna Steinglass, MD, *Columbia University/NYSPI*
Douglas Stelle, MD PhD, *University of Dundee*
Samia Tasmim, *CAMH*
Michael Treadway, PhD, *Emory University*
Blair Uniacke, MD, *Columbia University*

Elisa van der Plas, *Donders Institute*
Benjamin Wade, PhD, *University of California, Los Angeles*
Rob Whelan, PhD, *Trinity College Dublin*
Alexis Whitton, PhD, *Harvard Medical School*
Stephanie Winkelbeiner, *Feinstein Institute for Medical Research*

Data in search of model - abstracts

Dopamine release and activity in the amygdala during fear conditioning

Andreas Frick[1], Fredrik Ahs[2], Johannes Bjrkstrand[3], Mark Lubberink[4]and Mats Fredrikson[5]
[1]Stockholm University, Uppsala University; [2]Uppsala University, Karolinska Institutet; [3]Uppsala University, Copenhagen University Hospital, Lund University; [4]Uppsala University; [5]Uppsala University, Karolinska Institutet;

The authors declare no conflict of interest.

Data Healthy volunteers were scanned with simultaneous bolus+infusion [11C]raclopride positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) to measure dopamine release and neural activity during a 20 minute differential fear conditioning paradigm, pairing one cue (CS+) with an aversive electrical shock while another cue (CS-) was never paired with shock. Skin conductance was recorded during the fear conditioning. Thus, the data consists of PET measure of changes in dopamine concentration (5 minute frames), fMRI measure of neural activity (3 s TR), and skin conductance responses during fear conditioning. Skin conductance responses (SCRs) to CS+ subtracted from CS- is used as the fear learning measure. Data collection is ongoing. We have so far scanned ten volunteers (meanSD age: 25.15.9 years; 7 women) and plan to scan 8 more.

Question Functional neuroimaging has revealed a critical role for amygdala in acquisition of fear memories, and animal studies have indicated that amygdala dopamine signaling is necessary for fear conditioning. However, little is known regarding the relationship between dopamine signaling, amygdala activity and fear learning in humans. Thus, we are interested in testing these relationships in our dataset. We are aware that this is a bit vague, but are happy to discuss more specific hypothesis during the meeting.

Optimal integration of perceptual and reward uncertainty in decision-making and the role of general task structure

M. Ghane[1], S.A. Japee[2], J.A. Richey[3]and L.G. Ungerleider[4]

[1]Virginia Tech/National Institute of Mental Health; [2]National Institute of Mental Health; [3]Virginia Tech; [4]National Institute of Mental Health;

The authors declare no conflict of interest.

Data: To study how individuals make choices under simultaneous perceptual and reward uncertainty (PU;RU), we systematically varied uncertainty in both domains. On each trial, participants viewed two targets (face/house) and two distractors (cars) and were instructed to maximize reward by making speeded choices. We manipulated PU (5 conditions) by altering the phase coherence of target images. Manipulations were negatively correlated between face and house targets (CohFace + CohHouse = 81%). Distractors were fixed at the lowest coherence level.

We manipulated RU (5 conditions) by varying the reward probability ratio for receiving 8 (otherwise 2) between face/house targets from to 4/1. Reward probabilities were also negatively correlated between targets. Distractors were associated with no reward and all conditions were explicitly instructed. To test the impact of task design on behavior and eventually neural response, we designed two task versions. Version one fixed RU within run (varied randomly between runs), while PU was randomly varied across trials. In version two, PU was fixed within run (varied randomly between runs) while RU varied randomly across trials.

Participants completed 30 trials x 25 conditions x 2 tasks (75 trials per run, 750 trials per task; 1500 trials total). Each run began with a self-paced cue stating the run condition. Trials started with a 1000ms fixation, followed by a 700ms cue stating the trial condition, 300ms stimulus, 1300ms mask, and 800ms feedback (identity of the stimulus selected and reward outcome of single trial). Participants had 1600ms from stimulus presentation to end of mask to make choice.

Question: How do perceptual and reward uncertainty interact and integrate to alter subjective value? Specifically, is the impact of PU and RU on behavior (and eventually associated decision-relevant neural substrates) unique to each domain, or generally reflective of the overall uncertainty and method of uncertainty manipulation (stable/dynamic) independent of domain?

The Biasing Effect of Perceived Control on Instrumental Behavior

V. Ly[1], P. Piray[2], M. Jepma[3], M. Rose[4], K. Eleftheriadou[5]and M. Delgado[6]

[1]Department of Clinical Psychology; Leiden Institute for Brain and Cognition, Leiden University, Leiden, The Netherlands; [2]Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA; [3]Department of Psychology, University of Amsterdam, Amsterdam, The Netherlands; [4]Department of Clinical Psychology, Leiden University, Leiden, The Netherlands; [5]Department of Clinical Psychology, Leiden University, Leiden, The Netherlands; [6]Department of Psychology, Rutgers University, Newark, New Jersey 07102, USA;

The authors declare no conflict of interest.

Data: Healthy participants (N=31) performed a two-phased controllability task. Phase1: Participants learned keypresses to prevent shocks for Qcontrollable R cues. Events for Quncontrollable R cues were yoked, except that random keypresses were required. Phase2: Participants learned instrumental go/nogo-responses in a probabilistic learning task based on monetary feedback for cues that were previously associated with control/no control. Controllability and optimal response were manipulated independently enabling the assessment of perceived control biases of instrumental action. Thus, there are 4 trialtypes: 1. control-go, 2. control-nogo, 3. uncontrollable-go, 4. uncontrollable-nogo; 2 cues per trialtype, 45 trials per cue, resulting in a total of 360 trials.

Question: Controllability has been theorized to support decision making by modulating control allocation and instrumental responding. A notable feature of perceived control is that it can be inherently rewarding. Indeed, the appetitive value associated with perceived control has been suggested to bias control-seeking behaviors. An outstanding question is whether the affective properties of perceived control influences instrumental responding. We tested whether perceived control would non-selectively promote instrumental performance; or whether the impact of perceived control is action-specific due to Pavlovian valence-action (appetitive-go) coupling. Model-free analyses of choice data demonstrate increased instrumental performance for controllable (versus uncontrollable) cues, but this effect was not shown to be action-specific. However, when controlled for initial action bias, participants demonstrated increased performance for control-go specifically. With a model-based approach, I hope to better control for Qbaseline R action-bias, and test for the action-specificity of control biases of instrumental learning (control potentiates go; control potentiates learning; control potentiates go-learning).

How to model novel constructs of volition and self-agency;

L. S. Morris and J. W. Murrough

Department of Psychiatry, Icahn School of Medicine at Mount Sinai

The authors declare no conflict of interest.

While motivational deficits are a core feature of several psychiatric disorders, the construct of volition (self-driven motivation) and its differentiation from externally-generated motivation has been under explored. Classic measures of volition with effort valuation tasks have been developed and validated (see Research Domain Criteria), but they entail choices between externally-generated effort expenditure for reward and do not address potential levels of internally-generated motivation (volition).

Similarly, deficits in self-agency (subjective awareness of self as distinct from other) have been suggested in psychiatric disorders and neural representations of the self-other distinction are negatively associated with schizophrenia and depression severity. However, while deficits in self-agency are noted in psychiatry, current diagnostic and therapeutic models are not informed by its empirical operationalization.

Thus, we have developed a novel battery of cognitive tasks designed to capture internally-driven motivation (volition) and valence-dependent self-agency in humans. The internal-external motivation task (IMT) captures internally-driven motivation as a free-choice of preferred effort expenditure for varying rewards. The self-agency task (SAT) captures valence-dependant agency attribution. Both tasks require no learning.

Data: For the IMT, task performance data includes effort expenditure*reward curves for two conditions of internally-driven and externally-driven motivation (the latter similar to standard effort discounting tasks) for 8 levels of monetary reward and 13 levels of effort. For the SAT, data includes trial-by-trial ratings of self versus other following positive, negative or neutral feedback based on a simple reaction time task.

Question: We aim to dissociate internally-driven from externally-driven motivation and characterise the properties of both (i.e. Utility function fitting). We also aim to demonstrate the effect of valence on agency attribution in a model that accounts for task performance.

High-Resolution, Multimodal Data in MDD and Healthy Control Participants

B. R. Staveland and L. M. Williams

Stanford University

The authors declare no conflict of interest.

Background: In fMRI data, the blood oxygenation dependent (BOLD) signal is a noninvasive proxy for hemodynamic responses to neuronal activity. Therefore, a detailed characterization of the spatial and temporal properties of the BOLD signal is fundamental for accurately inferring the underlying neuronal activity. While the temporal properties, via the balloon models, and the spatial properties, via hemodynamic point spread functions, of the BOLD signal have been characterized by existing physiologically based models, it is commonly agreed that the adjoined spatiotemporal properties are relatively poorly understood.

Data: High-spatial and high-temporal resolution multimodal data were acquired from 204 unmedicated patients with MDD during cognitive control, continuous performance, and emotional processing task paradigms. Additionally, complimentary data from 69 healthy control subjects were also acquired.

Model: We are interested in developing a synthetic, biophysically-detailed model to characterize which task-evoked circuits function within which frequency bands. Ideally, this model should formalize how EEG synchrony in a single subject relates to their connectivity in fMRI.

Relevance: Such a model has clear benefits for many areas of neuroscience, particularly those concerned with a detailed understanding of wide-spread, whole-brain connectivity. As the data were also collected in a clinical sample, we hope this model will elucidate individual differences in depression pathology.

Reinforcement learning in adult ADHD, addiction, and working memory deficits

R. Whelan[1], M. Bennett[2], Z. Cao[3]and L. Rai[4]

[1]Trinity College Dublin; [2]Trinity College Dublin; [3]University College Dublin; [4]Trinity College Dublin;

The authors declare no conflict of interest.

Data: I will (very) briefly describe several datasets in search of models. First, Probabilistic Selection Task (PST) data from adults aged 55+ (final sample size=250), including a subsample with poor working memory performance (¡15th percentile on a standardized measure). n=60 completed the PST under concurrent high density electroencephalography (EEG), with spontaneous eye blink measurements also recorded. Second, behavioural PST data from ¿120 participants in a smoking study (current smokers, ex-smokers, e-cigarette users and never smokers). Third, data from a task involving probabilistic reward and punishment, including contingency reversal phases: 31 adults with ADHD and 31 matched controls. Fourth, a Pavlovian threat-conditioning protocol that was presented to 45 adults with ADHD +/- comorbid anxiety and 45 matched controls. On each trial, participants indicated their threat-expectancy on a scale, and high density EEG was recorded throughout. Fifth, the monetary incentive delay task (MID) was administered under high density EEG (and MRI in some cases) to a range of participants, including current smokers, acutely abstinent smokers (n 150), individuals with adult ADHD (n 50) and first degree relatives of those with ADHD (n 40), non-ADHD controls (n 50) and young adult alcohol users (many with heavy use; n 90). All ADHD participants completed a medication washout period, and genetic data will be available on most participants.

Question: In general, we are searching for models that will yield insights into between-group differences. For Datasets 1-3, we are interested in drift diffusion models; for Dataset 4, in experience-weighted attraction (EWA) models; and temporal difference models for MID tasks.

Model in search of data - abstracts

Modeling brain dynamics to ground diagnostic nosology in biological features;

Manish Saggar

Stanford University

The author declares no conflict of interest.

Model: The diagnostic nosology as currently used in psychiatry is built entirely upon symptoms and not biology. Thus, although the diagnosis made using DSM-V is highly reliable, it is not necessarily (biologically) valid. With the advent of modern noninvasive neuroimaging modalities, sophisticated methods have been developed to examine both structural and functional activity/connectivity of the brain for characterizing different psychiatric disorders. Nevertheless, several issues remain in developing neuroimaging based diagnostic nosology that is disorder-specific, person-centric, and grounded in biology. Here we propose to use a combined modeling approach of Topological Data Analysis (TDA) with Biophysical Network Modeling (BNM) as a "lens" towards (a) characterizing/stratifying psychiatric illness and (b) generating biologically grounded mechanistic insights regarding how neural processes interact during ongoing cognition to give rise to different dynamical landscapes in patient populations. We have recently shown that using TDA we can reveal brain's overall dynamical organization without arbitrarily averaging neuroimaging data across space or time at the single participant level (Saggar et al. 2018 Nature Communications). By combining our TDA approach with BNM we aim to also reveal the mechanisms underlying individual differences in brain dynamics.

Question: How to ground diagnostic nosology of psychiatric disorders into biological features? To start with we are interested in application of this approach to psychiatric populations where the neural dynamics are putatively on the opposite ends of the spectrum, for example, Major Depressive Disorder (MDD) and Attention Deficiency and Hyperactivity Disorder (ADHD).

Mapping of individualized brain network architecture to refine neuropsychiatric biotyping;

Shan H. Siddiqi[1], Carl D. Hacker[2], Sridhar Kandala[3], Maurizio Corbetta[4]and David L. Brody[5]

[1]Harvard Medical School, MGH/McLean; [2]Washington University School of Medicine in St. Louis; [3]Washington University School of Medicine in St. Louis; [4]Washington University School of Medicine in St. Louis; [5]Uniformed Services University of the Health Sciences

Conflicts of interest

S. Siddiqi: Scientific advisor for SigNeuro LLC. The other authors declare no conflicts of interest.

Model: The recent advent of individualized resting-state network (RSN) mapping enables fMRI-based identification of individual differences by using machine learning to classify brain regions into subject-specific network maps. Functional connectivity between these maps can provide robust individualized biomarkers of disease. This allows construction of a general linear model (GLM) to detect associations between unique clinical phenotypes and associated connectivity profile.

This approach was initially applied to a model of neuropsychiatric disturbance associated with traumatic brain injury (TBI), which is commonly associated with major depression and post-traumatic stress disorder (PTSD). TBI-associated depression is widely believed to be clinically distinct from primary major depressive disorder (MDD) and PTSD, but clear physiologic distinctions remain elusive. Resting-state fMRI data were obtained from 91 subjects across four datasets encompassing healthy controls (n=31), patients with TBI-depression (n=16), non-depressed TBI patients (n=19), and patients with MDD (n=27). After controlling for the effect of co-morbidities and datasets, individualized network architecture revealed that comorbid PTSD was associated with relative hyperconnectivity within ventral attention/salience network. TBI, TBI-depression, and MDD were distinguished by connectivity between dorsal attention network, default mode network, and subgenual anterior cingulate cortex. Interestingly, TBI-depression and MDD showed connectivity changes in opposite directions. All of these effects were statistically significant (pi0.05) using individualized network maps, but not with group-based network atlases.

Prediction: This model can be extended to any complex neurobehavioral syndrome to discover and validate discrete biotypes. This requires amalgamation of cross-diagnostic datasets with high-quality resting-state fMRI data and clear clinical phenotyping with well-characterized co-morbidities.