Ethan M. Meyers, Ph.D.

Curriculum vitae, updated April 2012

Contact Information

McGovern Institute for Brain Research and
Department of Brain and Cognitive Sciences
Mass. Institute of Technology, Building 46-5155D
77 Mass Ave, Cambridge, MA, 02139
Phone: (617) 447-7814
emeyers@mit.edu
web.mit.edu/emeyers/www

Education

2004-2010	Ph.D. in Computational Neuroscience, Mass. Institute of Technology Advisor: Tomaso Poggio
1998-2002	B.A. in Computer Science, Oberlin College

Current Position

2010-present	Postdoctoral Fellow, McGovern Institute for Brain Research, MIT

Past Positions

2002-2004	Technical Assistant, in the laboratory of Pawan Sinha, MIT
2002-winter	Research Intern, Dept of Biomedical Engineering, UT Memphis

Honors, Awards and Fellowships

2009-2010	Hubert Schoemaker Graduate Student Fellowship
2006-2009	National Defense Science and Engineering Graduate Fellowship
2005	Sigma Xi Scientific Research Society
2002	Phi Beta Kappa Honor Society
2002	High Honors in Computer Science, Oberlin College

Teaching Experience

Teaching Assistant, 9.011 Systems Neuroscience Core (Profs. Earl Miller, Matt Wilson)
Teaching Assistant, 9.660 Computational Cognitive Science (Prof. Josh Tenenbaum)
Teaching Assistant, 9.63 Laboratory in Cognitive Science (Prof. Aude Oliva)
Teaching Assistant, 9.00 Introduction to Psychology (Prof. Jeremy Wolfe)

Student Supervision

Ami Patel (master of engineering, MIT, 2013)
Jiahao Liang (undergraduate, MIT, class of 2014)
Mia Borzello (research assistant MGH, undergraduate UC Berkeley, class of 2011)
Mark Rodgers (undergraduate, UC Berkeley, class of 2013)

Publications

Journal articles

Meyers, E., Qi, X.L., Constantinidis C. Incorporation of new information into prefrontal cortical activity after learning working memory tasks. Proceedings of the National Academy of Sciences, 109:4651-4656-8855, 2012.

Zhang Y.,*, **Meyers, E.***, Bichot, N., Serre, T., Poggio, T., and Desimone, R. Object decoding with attention in inferior temporal cortex. Proceedings of the National Academy of Sciences, 108:8850-8855, 2011. ***These authors contributed equally**.

Ostrovsky, Y., **Meyers, E.**, Ganesh, S., Mathur, U., Sinha P. Visual Parsing After Recovery From Blindness. Psychological Science, 20:1467-1491, 2009.

Meyers, E., Freedman, D., Kreiman, G., Miller, E., Poggio T. Dynamic Population Coding of Category Information in Inferior Temporal and Prefrontal Cortex. Journal of Neurophysiology, 100:1407-1419, 2008.

Meyers, E., Wolf, L. Using Biologically Inspired Visual Features for Face Processing. International Journal of Computer Vision, Vol 76, No. 1, 93-104, 2008.

Cox, D., **Meyers, E.**, and Sinha, P. Contextually Evoked Object-specific Responses in Human Visual Cortex, Science, Vol. 303, No. 5667, 115-117, 2004.

Refereed Conference Papers

Wolf, L., Bileschi, S., and **Meyers, E.** Perception Strategies in Hierarchical Vision Systems IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2006.

Ezzat, T., **Meyers, E.**, Glass, J., and Poggio, T. Morphing Spectral Envelopes Using Audio Flow, Interspeech/Eurospeech, Lisbon, Portugal, September 2005.

Technical Reports

Isik, L., **Meyers, E.**, Leibo, J., Poggio, T. Preliminary MEG decoding results. MIT-CSAIL-TR-2012-010, CBCL-307, Massachusetts Institute of Technology, Cambridge, MA, April 20, 2012

Meyers, E., Embark, H., Freiwald, W., Serre, T., Kreiman, G., and Poggio T. Examining high level neural representations of cluttered scenes. MIT-CSAIL-TR-2010-034 / CBCL-289, Massachusetts Institute of Technology, Cambridge, MA, July 29, 2010

Book Chapter

Meyers, E., and Kreiman, G. "Tutorial on Pattern Classification in Cell Recording". In: Visual population codes. Kreigeskorte, N., and Kreiman, G. (eds.), 2011, MIT Press.

Conference abstracts

Meyers, E., Qi, X.L., Constantinidis C. The incorporation of new information into prefrontal cortical activity after learning new tasks. Poster: Cosyne, Salt Lake City, Utah, 2012.

Meyers, E., Qi, X.L., Constantinidis C. Comparing the information content of PFC before and after training in a working memory task. Poster: Society for Neuroscience, Washington, DC, 2011.

Meyers, E., Zhang, Y., Bichot, N., Serre, T., Poggio, T., Desimone R. The representation of objects in inferior temporal cortex with and without attention. Talk: Society for Neuroscience, San Deigo, CA, 2010.

Meyers, E., Zhang, Y., Bichot, N., Poggio, T., Desimone, R. Attention's influence on object representations in the inferior temporal cortex. Invited talk: The 7th International Conference on Cognitive Science, Beijing, China, 2010.

Meyers, E., Freedman, D., Miller, E., Kreiman, G., Poggio, T. The Coding of Abstract Category Information in Macaque Inferior Temporal and Prefrontal Cortex. Invited talk: The 7th International Conference on Cognitive Science, Beijing, China, 2010.

Meyers, E., Zhang, Y., Bichot, N., Chikkerur, S., Serre, T., Poggio, T., Desimone R. Decoding multiple objects from populations of macaque IT neurons with and without spatial attention. Poster: Cosyne, Salt Lake City, Utah, 2010.

Meyers, E., Freedman, D., Kreiman, G., Miller, E., Poggio T. Decoding dynamic patterns of neural activity using a 'biologically plausible' fixed set of weights. Poster: Cosyne, Salt Lake City, Utah, 2009.

Meyers, E., Embark, H., Freiwald, W., Serre, T., Kreiman, G., Poggio, T. Neural representations of cluttered scenes in macaque ventral visual cortex. Poster: Society for Neuroscience, Washington, DC, 2008.

Meyers, E., Hung, C., Freedman, D., Miller, E., Kreiman G. Decoding of ITC Cell Activity Matches Human Visual Similarity Judgments. Talk: Vision Sciences Society. Sarasota, Florida, May 2007.

Meyers, E., Freedman, D., Miller, E., Kreiman, G., Poggio, T. Reading Out Visual Information from Populations of Neurons in Inferior Temporal and Prefrontal Cortex. Poster: Cosyne, Salt Lake City, Utah, 2007.

Meyers, E., Ostrovsky, Y., & Sinha, P. Visual de-fragmentation via high spatial frequencies. Poster: Vision Sciences Society. Sarasota, Florida, May 2005.

Meyers, E., Cox, D. D., & Sinha, P. Neural responses to contextually defined faces. Talk: Vision Sciences Society. Sarasota, Florida May 2003.

Ethan Meyers - Research statement

My research interests consists of developing and applying novel machine learning data analysis methods that can give deeper insight into how information is coded in neural activity, with the ultimate goal of understanding the neural algorithms that underlie complex behaviors.

I also have a strong interest in fostering the growth of the field of neuroinformatics, which will lead to a much faster pace of discovery as the field of neuroscience enters the era of 'big data'.

Introduction

One of the primary goals of computational systems neuroscience is to understand neural processing in terms that can be translated into computer algorithms. Such an algorithmic description of neural processing could guide the creation of intelligent robotic systems, and could help in the development of treatments for psychiatric and neurological disorders that are based on a principled understanding of how the brain functions. The building blocks for developing an algorithmic understanding of neural processing is to first gain *reliable estimates* of what information is represented in particular brain regions at particular points in time, and how this information is coded in neural activity, from which algorithmic descriptions of processing can be pieced together. While current methods used by neuroscientists to analyze data can give insight into these questions, many commonly used methods can lead to incorrect conclusions (Nieuwenhuis et al., Nat. Neuro., 2011) or do an inadequate job of extracting information from data that is collected, leaving many potentially significant properties of neural systems unnoticed (e.g., Meyers et al., J. Neurophys. 2008; Meyers et al., PNAS 2012). By examining how the brain codes information in more detail, my research aims to give a clearer understanding of the algorithms that underlie intelligent behavior.

Prior work

My background is in computer science, and while I have previously done work in neuroimaging (Cox et al., Science, 2004), computer vision (Meyers et al., Int. J. Comput. Vision, 2008; Wolf et al., CVPR 2006), and studying visual deficits (Ostrovsky et al., Psychol. Sci., 2009), one of the most fruitful areas of my past research has been developing new data analysis methods that give more insight into the content and coding of information in higher level cortical regions (Meyers et al., J. Neurophys. 2008; Zhang, Meyers, et al., PNAS, 2011, Meyers et al., PNAS 2012). Below, I describe some important discoveries I have made using novel neural decoding methods to illustrate the power of the tools that I develop.

One of the most interesting neuroscience findings I have made using my data analysis methods, was to show that information stored in working memory in frontal, parietal and temporal cortices is coded by a dynamic sequence of neural activity across a population of neurons (in a way that is similar to how a synfire chain operates). I was able to make this discovery by training a classifier with data from one point in time in a trial and then testing the classifier with data at a different point in time; by observing that high decoding accuracy was only obtained when the classifier was trained and tested at the same point in time, it became apparent that different patterns of neural activity were coding the same information at different points of time in the trial (Meyers et al., J. Neurophys., 2008; see Figure 1A below). This finding upended the a long held belief that the basis of working memory was due to neurons maintaining high firing rates levels over long periods of time, and has be subsequently confirmed by several other studies. I have also shown that all the information that is available in a large population of neurons is often contained in a very small subset of highly selective neurons (Meyers et al., PNAS, 2012). This finding is not only interesting in its own right, but it also has significant implications for how one should analyze data since if only a small percentage of neurons are involved in particular computations, then averaging a selectivity measure over all neurons in a population might greatly underestimate the amount of information present.

Another line of questions that my data analyses methods are able to address concerns assessing whether neural representations are invariant to particular stimulus transformations. An important aspect of intelligent behavior is the ability to treat stimuli that have very *different physical characteristics* in a similar manner when they have *similar behavioral relevance*. For example, humans have the ability to categorize visual objects as belonging to a particular object class despite differences in its color, lighting, pose, size, etc.. Thus in order to evaluate whether a brain region contains information in an invariant manner, I have developed a decoding method where a classifier is trained on one set of conditions and is evaluated on a new but related set of conditions to see if the neural representation can generalize across the chosen transformations. Using this method, I have verified that neural representations in the inferior temporal cortex (IT) are highly position invariant (Meyers and Krieman, 2011) and that training a monkey to group stimuli as belonging to the same category influences neural representations in both IT and the PFC causing arbitrary images to be grouped together based on their behavioral significance (Meyers et al., J. Neurophys., 2008). Also, through using this method we were able to clearly see that neural representations in anterior medial face selective patches (AM) are selective for individuals regardless of the pose of their heads, while the more posterior face selective patches are not (see Fig 1B). By applying this technique in the future it should be possible to trace how information is processed in an intelligent way through particular pathways in the brain (for example, the ventral visual pathway), which will give insight into how the brain constructs neural representations of abstract categories that are useful for behavior.

To promote the application of these decoding methods, I have developed a Matlab toolbox that allows researchers to easily run the decoding analyses described above on their own dataset. The toolbox is based on a modular architecture, allowing a user to interchange pieces of the toolbox to test new algorithms or to adapt the already existing algorithms to work with new types of data (such as MEG data, local field potentials, etc.). One of my research aims is to continue to develop these decoding methods and to integrate them with other data analysis methods (such as unsupervised learning methods like factor analysis) in order to gain even more insight about neural processing.

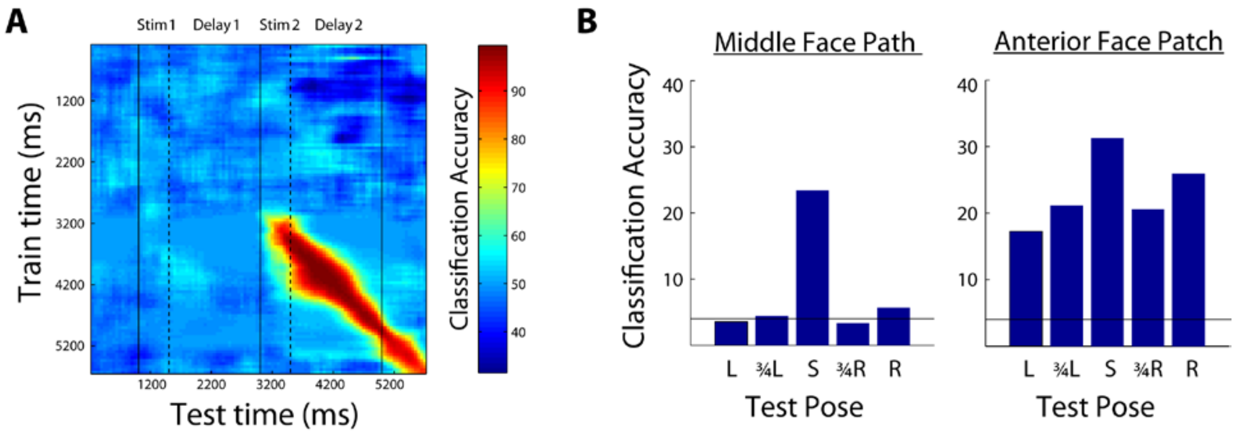

Figure 1: Results illustrating how the population decoding methods I have developed can be used to answer questions about neural coding and abstract neural representations. A, Results showing that 'match/non-match' information in a delayed match-to-sample-task is contained in a dynamic population code. We trained a classifier with data from one point in time (y-axis) and tested the classifier with data from a different trial that was taken either from the same point in time or a different point in time (x-axis). The strong diagonal band in the figure shows that high decoding accuracies were only present when the classifier was trained and tested using data from the same point in time relative to the start of the trial, indicating that even when the stimulus was constant, different neurons contained the same information at different points in the trial (from Meyers et al., PNAS, 2012). **B**, Results showing that neural populations in the anterior medial face patch (AM) represents an individual in a way that is invariant to head direction, while the middle face patches (ML/MF) do not represent faces in a pose-invariant way. These results were obtained by training the classifier to discriminate between 25 individuals using a frontal head view (S), and then evaluating whether the classifier can discriminate between individuals at the same frontal head orientation (S), and whether the classifier can generalize to other head orientations including 3/4 profile views (3/4L, 3/4R), and full profile head views (L, R). The results show that ML/MF cannot generalize to other views, since results generalizing to other views are at the 4% chance level (indicated by the black horizontal line), while AM can generalize to other views (data from Freiwald and Taso, 2010).

Research plan

My research goal is understand how populations of neurons code information with the ultimate aim of elucidating the neural algorithms that underlie intelligent behavior. To do this, I will develop and apply new data analysis methods and share these computational tools with the larger neuroscience community.

1. Creating new data analysis techniques. The major thrust of my research is to develop new computational methods that can give deeper insight into how populations of neurons represent and process information. As mentioned above, in my previous work I have used neural population decoding in novel ways to gain insight into how populations of neurons processes information. Neural population decoding is considered a "supervised learning" method because it assumes that the researchers already know what variables are of interest, and the method seeks to assess how much information is present about these known variables. However, it is likely that in many experiments neural activity is influenced by variables that are not considered a priori by a researcher. Thus one direction I am pursuing is to use "unsupervised learning" methods to examine previously unconsidered dimensions that might be present in the data. Currently I am mentoring a student on a project who is applying a commonly used unsupervised learning method called 'factor analysis' to assess whether this method can find structure in neural responses from macaque V1 and V4. Our results show that our algorithms are able to find a small subset of dimensions that capture most of the variance in the neural data, and while some of these dimensions corresponded to known variables (such as visual properties of stimuli that were shown), there also appears to be consistent variation unrelated to variables we had previously considered (Fig. 2A). By developing computational tools that can find meaning in these unaccounted for dimensions, we hope to find new functions of particular brain regions that have not been previously considered.

I am also developing computational methods that can explore information contained in the coordinated activity of populations neurons. To address this issue I have adapted some of my neural decoding methods to assess the information content in local field potential (LFPs) and in the coherence between LFPs from different brain regions. Additionally, I am exploring how different firing rate patterns, such as the burst and tonic modes seen in thalamic neurons, contribute to information processing.

Finally, I plan to adapt the computational tools I have developed to work with new types of data that have different statistical properties. For example, I am currently working with a graduate student to apply population decoding analyses to MEG signals. Our results show that we are able to decode a large amount of information from single trials of these highly noisy MEG signals, which opens up a range of possible new MEG experiments (see Fig. 2B below; Isik, CBCL memo, 2012).

2. Collaborating with experimentalists a make new discoveries. I also plan to continue to work closely with electrophysiologists to make new neuroscience discoveries. Currently I have ongoing collaborations with several electrophysiology labs including the Desimone lab at MIT, the Constantinidis lab at Wake Forrest, and the Freiwald/Tsao labs at Rockefeller/CalTech. My collaboration with the Desimone lab is particularly exciting because I am helping them analyze simultaneous neural recordings from the inferior temporal cortex, V4, and the pulvinar nucleus of the thalamus; these data allow us to explore a range of novel data analyses methods to examine how attention affects information processing and interactions between different brain regions. I also have a strong interest in forging new collaborations with faculty members at Brown. In particular, I would enjoy collaborating with David Sheinberg, and Michael Paradiso to examine neural processing of visual information, and I am interested in collaborating with Christopher Moore to examine how information is coded in oscillatory activity.

3. Developing the field of neuroinformatics and putting computational tools in experimentalists hands. A significant barrier to progress in neuroscience is that the methods used from one paper to another are often very different and thus when conflicting findings arise in the literature it is often difficult to ascertain whether the differences are "real" or whether they are merely a result of using different methods. To help alleviate this problem, I plan to work with electrophysiologists to promote the development of a few standard neuroscience datasets which will allow researchers to compare the efficacy of different data analysis methods and to determine which methods give the most reliable and insightful results. Such standard databases have had a dramatic effect on progress in the field of computer vision, and most current research in computer vision would not be possible without them. Additionally, enabling computational neuroscientists to clearly demonstrate how their data analysis methods give new insight into neural processing should help the larger neuroscience community realize the value of new computational methods and should help foster the growth of neuroinformatics research, much like how bioinformatics research has grown over the past decade.

Additionally, I intend to share the computational tools that I develop so that electrophysiologists can use them to gain deeper insights from new data that they collect. Many of the algorithms I develop require an in-depth understanding of machine learning, statistics and computer programming, and it is unrealistic to assume that merely describing these algorithms in research papers will have a substantial impact on the field if it requires that these algorithms need to be reimplemented. Currently I am developing a Matlab toolbox that will allow researchers to easily apply decoding methods their data. I have also developed a few standard data formats which should allow researchers to allow a range of computational tools to their data, thus helping to increase the pace of discovery in neuroscience.

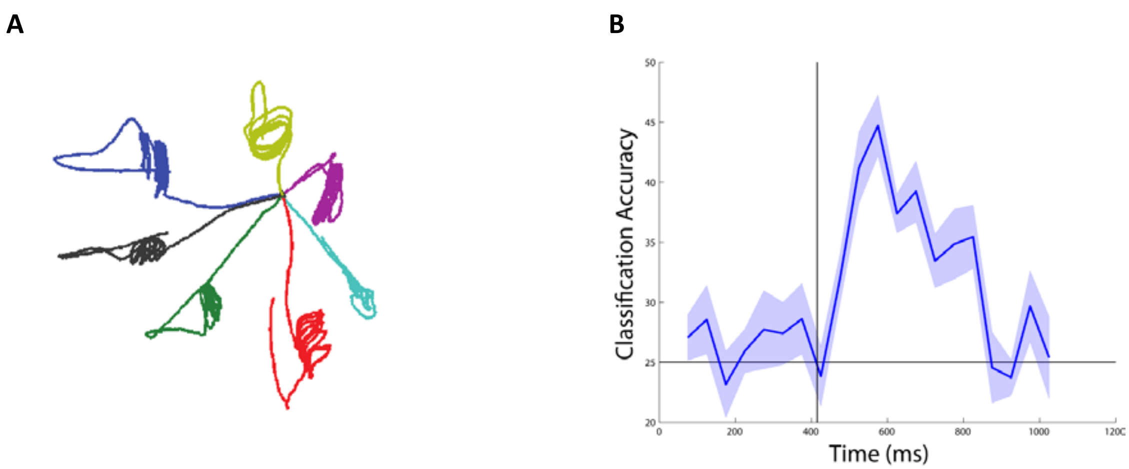

Figure 2. Factor analysis and MEG decoding results. **A,** Neural trajectories found by applying factor analysis to V1 responses to drifting gratings. The results revealed that each orientation (indicated by a different color) was well separated in two dimensions. Analyzing other dimensions (not shown) could help reveal new factors that modulate V1 responses (data provided by Adam Kohn). **B,** Decoding MEG responses to 4 visual stimuli revealed that it is possible to extract information from single MEG trials. Stimulus onset time is 400 ms (this work was done jointly with Leyla Isik).

Ethan Meyers - Teaching Statement

One of the main reasons I am pursuing a career in academia is to have opportunities to teach, mentor students and to create new educational material. I believe that there is a strong symbiotic relationship between teaching and research, where research enables one's teaching to stay on the cutting edge, and where teaching helps to ensure one's research can be easily explained and keeps one focused on broader questions in the field. I also love learning new concepts myself, and I have found from working with students that my excitement for neuroscience and statistics is contagious.

Due to my background in computer science, neuroscience and statistics, there is a broad range of classes that I would enjoy teaching. In graduate school I was a teaching assistant for an undergraduate introductory psychology course and for the graduate introductory neuroscience course, and combined with my research in neuroscience, I am well prepared to teach introductory and advanced systems neurons courses (NEUR1030 and NEUR2050). Similarly, with my extensive research experience in computational neuroscience and my experience as a teaching assistant for a computational cognitive science course, I would enjoy teaching computational neuroscience (NEUR1680). I would also be happy teaching a variety of introductory courses including NEUR 0010 and APMA 0410, and I have a strong interest in continuing to learn new material in order to improve the quality of my classes and potentially expand the range of what I can teach. Finally, I would like to introduce new courses including an advanced course on the analysis of neural data, and an introductory applied statistics course that features examples from baseball.

My teaching style adapts the appropriate approach for different types of classes. For an introductory neuroscience class where the concepts are more straightforward, I believe a simple lecture format is most appropriate, and I would be sure to introduce the historical development of the field so that students have the contextual knowledge to think critically about the questions that the field is (or should be) focused on. For other introductory courses that are more mathematical in nature (and where there might be more discrepancy in different students' abilities), I would try to leverage the students who are more adept at understanding the concepts to help explain the concepts to students who are having a harder time, because I have found that students who have recently learned new material are often very good at explaining concepts to their peers. Thus, I would try to include more group-based learning projects in these classes. Finally, for more advanced courses, I would focus on assigning class projects that explore current research themes in order to develop problem solving and creative thinking skills. My experience as a teaching assistant in a cognitive science laboratory class, and my experience mentoring undergraduate and master's students on neuroscience research projects, have been incredibly useful in teaching me how to develop projects that help students to learn how to think critically, ask and answer interesting questions, and solve challenging problems.