DEPARTMENT OF
PSYCHOLOGICAL AND
BRAIN SCIENCES

INDIANA UNIVERSITY
Bloomington

William Alexander, Ph.D.
wialexan@indiana.edu
(812) 856-1846
April 26, 2012

David Sheinberg, Chair
Faculty Search Committee
Department of Neuroscience
Brown University
Providence, RI

Dear Professor Sheinberg:

I write to apply for the position of Assistant Professor in the area of Computational Neuroscience. This position was posted to the Science Careers employment web site. I am currently a postdoctoral research associate working with Joshua Brown in the Department of Psychological & Brain Sciences at Indiana University, Bloomington. I am eager to apply for this position as I believe that my research background and interests correspond well with the desired characteristics outlined in the description of the position.

The job posting specifically solicits applications from individuals using computational approaches to investigate neural information processing. In my graduate work with Olaf Sporns, I investigated computational models of neuromodulatory systems in the context of ongoing autonomous behavior. For my dissertation, I demonstrated how signals generated by a reinforcement model can be exploited to selectively attend to behaviorally relevant features of a complex visual scene. In my postdoctoral work with Joshua Brown, I have expanded my interests to include computational modeling of brain regions implicated in decision making and cognitive control. I have recently published a new model describing the role of anterior cingulate cortex (ACC) in cognitive control (Alexander & Brown, 2011, *Nature Neuroscience*). The *predicted response-outcome* (PRO) model comprehensively accounts for an unprecedented array of ACC activity observed in single-unit, ERP, and fMRI studies within a unified framework, and suggests a reinterpretation of ACC as a region involved in predicting and detecting the unexpected non-occurrence of predicted events.

The description of the position additionally emphasizes that the successful candidate would be able to collaborate across multiple disciplines. On the strength of my work modeling ACC, I am currently involved in a multi-university collaboration under the IARPA Integrated Cognitive-Neuroscience Architectures for Understanding Sensemaking (ICArUS) program (over $500,000 total costs; PI Joshua Brown). The project brings together individuals with expertise in computer science, neuroscience, mathematical psychology, as well as computational neuroscience in order to develop neurally plausible computational models of brain regions involved in high-level cognitive behaviors. Due to my contribution to this project, the lead investigator has assured me of continued funding should I secure an independent tenure-track position.

Finally, my work has direct relevance to understanding possible neural correlates of psychiatric disorders, as specified in the job posting. In work currently in preparation for publication, conducted in collaboration with clinical psychologists at Indiana University, we tested predictions of the PRO model regarding ACC activity in substance dependent individuals. In this study, we find that increased aversion to risk better accounts for ACC activity in substance dependence individuals than alternative hypotheses such as increased risk seeking or attention to cues that predict reward. The results of this combined fMRI/computational modeling study suggest that the PRO model may be exploited to investigate possible mechanisms of additional behavioral and psychological disorders.

I am particularly attracted to this position for several reasons. First, I believe my research interests are an exceptional fit with the work being conducted both in the Department of Neuroscience and the Brown Institute for Brain Science. In particular, several members of the faculty share my interests in computational modeling of the brain at the level of single-unit activity as well as at higher cognitive levels. Additionally, as my postdoctoral work has included a significant amount of training in neuroimaging methods, I would be pleased to have the opportunity to initiate and contribute to research that makes use of the equipment maintained at the Brown University MRI Research Facility. Finally, my participation in the IARPA ICArUS program has given me the opportunity to collaborate with several academic and industry scientists in the greater New England area, including Michael Frank at Brown University, and I would greatly enjoy the opportunity to continue these interactions on a more frequent basis.

Please find attached the requested application materials. Letters of recommendation will arrive separately. I look forward to hearing from you.

Sincerely,
William Alexander

Attached:
Curriculum vitae
Statement of Research Interests
Statement of Teaching Philosophy
Selected Reprints

Curriculum Vitae
William H. Alexander
Department of Psychological and Brain Sciences
Indiana University
Bloomington, IN 47405
(812)856-1846
wialexan@indiana.edu

EDUCATION

Indiana University, Bloomington
Degree: Ph.D. **2006**
Major: Cognitive Psychology
2nd Major: Cognitive Science
Dissertation: "A Real-Time Model of Attention"

University of Nevada, Reno
Degree: B.A. **1999**
Major: Psychology
Minor: Philosophy

RESEARCH EXPERIENCE

Indiana University, Bloomington
Postdoctoral Researcher – Cognitive Control Lab **2007-Present**
Director: Joshua W. Brown
Research Topics: Executive Control, Decision Making, Reinforcement Learning, Computational Modeling, fMRI

Okinawa Institute of Science and Technology
Researcher – Neural Computation Unit **2006-2007**
Director: Kenji Doya
Research Topics: Attention, Computational Modeling, Neuromodulation

Indiana University, Bloomington
Graduate Student – Computational Cognitive Neuroscience Lab **2001-2006**
Director: Olaf Sporns
Research Topics: Neural Bases of Reinforcement, Embodied Cognition, Robotics

FUNDING SOURCES

Title	**Role**	**Dates**
Computational Models of Risk (AFOSR FA9550-07 1-0454). PI: Joshua Brown	Co-Investigator	**6/1/07-11/30/09**
Neural Mechanisms of Risky Behavior Avoidance NIH/NIDA R01 DA026457. PI: Joshua Brown	Co-Investigator	**8/1/2009-7/31/2011**
Integrated Cognitive Architectures for Understanding Sensemaking (subcontract). PI: Joshua Brown	Co-Investigator	**10/1/2010-2/28/2014**

TEACHING EXPERIENCE

Indiana University, Bloomington

Assistant Instructor – Advanced Statistical Analysis* **Fall, 2005**

Assistant Instructor – Statistical Techniques **Spring, 2005**

Assistant Instructor –Advanced Statistical Analysis* **Fall, 2004**

Instructor – Methods of Experimental Psychology **Spring, 2004**

* *denotes graduate-level courses*

PEER-REVIEWED JOURNAL ARTICLES

Alexander, W.H. and Brown, J.W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience.*14(10), 1338-1344.

Alexander, W.H. and Brown, J.W. (2010). Computational models of response-outcome prediction as a basis for cognitive control. *Topics in Cognitive Science* 2(4), 658-677.

Alexander, W.H. and Brown, J.W. (2010). Hyperbolically discounted temporal difference learning. *Neural Computation* 22(6), 1511-27.

Alexander, W.H. and Brown, J.W. (2010). Competition between learned reward and error outcome predictions in anterior cingulate cortex. *Neuroimage,* 49(5), 3210-3218.

Alexander, W.H. (2007). Shifting Attention Using a Temporal Difference Prediction Error and High-Dimensional Input. *Adaptive Behavior,* 15, 121-133

Alexander, W.H. and Sporns, O. (2003). An Embodied Model of Learning, Plasticity, and Reward. *Adaptive Behavior.* Vol 10(3-4), Sum 2002, pp. 143-159

Sporns, O., and **Alexander, W.H.** (2002). Neuromodulation and plasticity in an autonomous robot. *Neural Networks.* Vol 15(4-6), Jun-Jul 2002, pp. 761-774.

MANUSCRIPTS IN PREPARATION

Alexander, W.H., Fukunaga, R. and Brown, J.W. (In preparation). Aversion to risk underlies mPFC activity in substance abusers.

Alexander, W.H. and Brown, J.W. (In preparation). Decision-making under uncertainty as survival maximization.

BOOK CHAPTERS & REFEREED CONFERENCE PAPERS

Brown, J.W. and **Alexander, W.H.** (2011). Computational Neuroscience Models: Error monitoring, conflict resolution, and decision making. In: *Perception-reason-action cycle: Models, algorithms and systems*. pp. 169-186, Cutsurdis, V., Hussain, A, & Taylor, J.G. (Eds),Springer: New York, NY.

Alexander, W.H. and Sporns, O. (2006). Temporal difference learning with learned attention shifts. *Proceedings of the Fifth International Conference on Development and Learning.* Bloomington, IN.

Alexander, W.H. and Sporns, O. (2004). Interactions of environment, behavior, and synaptic patterns in a neuro-robotic model. In: *Animals to Animats 8: Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior*,pp. 13-22, Schaal, S., Ijspeert, A., Billard, A., Vijayakumar, S., Hallam, J., and Meyer, J-A. (Editors). MIT Press: Cambridge, MA.

Sporns, O. and **Alexander, W.H.** (2003). Neuromodulation in a learning robot: Interactions between neural plasticity and behavior. *Proceedings of IJCNN 2003*, 2789-2794.

Alexander, W.H. and Sporns, O (2002). Timed delivery of reward signals in an autonomous robot. In: *Animals to Animats 7: Proceedings of the Seventh International Conference on the Simulation of Adaptive Behavior,* pp. 195-204, Hallam, B., Floreano, D., Hallam, J., Hayes, G. and Meyer, J-A. (Editors), MIT Press: Cambridge, MA.

Sporns, O., and **Alexander, W.H.** (2002). Dopamine, reward conditioning, and robot behavior. In: *Proceedings of the 2nd International Conference on Development and Learnin*g, pp. 265-270, IEEE Computer Society, Los Alamitos, CA.

POSTERS/PRESENTATIONS

Alexander, W.H. and Brown, J.W. (2010). Medial prefrontal cortex predicts the outcomes of actions. Nanosymposium talk at the annual meeting of the Society for Neuroscience. San Diego, CA.

Alexander, W.H. and Brown, J.W. (2010). Discounting time and probability by reward perception. Poster at the Society for Neuroeconomics annual conference. Evanston, IL.

Alexander, W.H. and Brown, J.W. (2010). A common mechanism for time and probability discounting. Poster at the Air Force Office of Scientific Research Cognition & Decision Joint Program Review. Arlington, VA.

Alexander, W.H. and Brown, J.W. (2008). A computational neural model of learned response-outcome predictions by anterior cingulate cortex. Poster at the annual meeting of the Society for Neuroscience. Washington, D.C.

Alexander, W.H. and Brown, J.W. (2008). Error likelihood effects in anterior cingulate cortex modulated by average reward and reinforcement learning. Poster at the annual conference for the Cognitive Neuroscience Society. San Francisco, CA.

Alexander, W.H. (2004). Mutual influences of environment and behavior on the development of a neural model. Invited talk at the workshop for Neurorobotic Models in Neuroscience and Neuroinformatics. Los Angeles, CA, July 17, 2004.

Alexander, W.H. and Sporns, O. (2003). Environmental influence on behavior and development of an autonomous robot. Poster at the Annual meeting of the Society for Neuroscience, New Orleans, LA.

Sporns, O., Bulwinkle, D., Chadderdon, G., and **Alexander, W.H.** (2003). Neuro-robotic models of learning and addiction. Poster at NIH Symposium (Biomedical Information Science and Technology Initiative) Digital Biology, The Emerging Paradigm. Bethesda, MD.

Malkoc, G., **Alexander., W.H.**, and Webster, M.A. (2001). Color and Adaptation in Perceptual Grouping. Poster at the 1st Annual Meeting of the Vision Sciences Society, Sarasota, FL.

Amberg, M.D., Yamashita, J.A., Merica, B.L., **Alexander., W.H.**, and Wallace, W.P. (2001). Words with overlapping phonemes in early positions facilitate correct recall. Poster at the Annual Convention of the Western Psychological Association. Tucson, AZ.

AWARDS

Cognitive Science Summer Research Fellowship – Indiana University, 2005

Outstanding Paper – International Conference on Development and Learning (co-author), 2002

Summer Research Incentive Fellowship – Indiana University, 2002

Faculty Commendation – Indiana University, 2001-2002

National Science Foundation Graduate Research Fellowship – Honorable Mention, 2002

Cognitive Science Supplemental Fellowship – Indiana University, 2001

PROFESSIONAL SERVICE

Ad-hoc Reviewer: *Adaptive Behavior*; *Cognitive, Affective, & Behavioral Neuroscience; Cortex; Neuropsychologia*; *Topics in Cognitive Science*

PROFESSIONAL MEMBERSHIPS

Society for Neuroscience

Cognitive Neuroscience Society

Society for Neuroeconomics

Statement of Research Interests

William Alexander
Department of Psychological and Brain Sciences
Indiana University, Bloomington

In the course of goal-directed behavior, an individual must rapidly integrate and process a great deal of information in order to generate appropriate behaviors. This information may include the nature of a desired goal, how likely that goal is to be realized, or potential actions which may be required in order to achieve that goal. My research seeks to understand the computational processes involved in goal-directed behavior and decision making, and the neural substrates on which these processes are implemented. My work incorporates two complementary approaches: computational and mathematical modeling of behavior and neural activity based on empirical observation coupled with model-driven experimentation to test model predictions.

Predicting the consequences of actions

Medial prefrontal cortex, especially anterior cingulate cortex (ACC), is generally recognized as a key locus of cognitive control and executive function, and thus is implicated in a variety of behavioral functions and dysfunctions. However, despite the amount of research devoted to this region, the role of ACC is still a subject of considerable debate (reviewed in Alexander & Brown, 2010). On one hand, imaging and ERP studies of humans suggest that ACC is involved in error processing and resolving behavioral conflict. In contrast, single-unit neurophysiology studies of monkey indicate that ACC predicts and detects rewarding events.

I have recently developed a new computational model of ACC function, the *prediction of responses and outcomes* (PRO) model (Alexander & Brown, 2011) . The PRO model learns to predict the possible outcomes of an action based on environmental cues, and signals when a predicted outcome fails to occur. Intuitively, this signal reflects the "surprise" experienced when an expected event fails to occur. Using only this notion of *surprising non-occurrence*, the PRO model accounts for an unprecedented array of observed effects from imaging, EEG, single-unit recordings, and behavioral studies which have not previously been unified in a single model. These effects include error detection, prediction of error likelihood, conflict effects, single-unit prediction of reward and reward detection, tracking of environmental volatility, and effects of multiple responses.

The PRO model makes a number of specific and testable predictions regarding the neural substrates of cognitive control, and I am currently in the initial stages of testing these predictions using fMRI and behavioral methods. In work currently in preparation, I show that ACC activity in substance-dependent individuals engaged in a cognitive control task is best explained as increased aversion to risk relative to non-substance dependent subjects, in agreement with a prediction of the PRO model. I am currently piloting an additional fMRI study to test another such prediction, namely that surprising absence of a painful stimulus should elicit increased activity in ACC in the same region typically associated with processing error and behavioral conflict. If correct, this finding would suggest a more general role for ACC as a region involved in signaling deviations from predictions on both cognitive and affective levels.

While a number of effects in ACC are captured by the PRO model, ACC is only a single component underlying cognitive control and decision making. The formulation of the PRO model suggests a general algorithm which may be implemented by a distributed network of brain regions. If ACC signals surprising non-occurrences of predicted events, additional brain regions may be involved in, for example, maintaining predictions of future events or reporting the affective valence of an event. A longer term goal of my research, therefore, is to elucidate the interactions of brain regions underlying cognitive control of behavior. To this end, I am currently involved in a multi-university collaborative effort under the IARPA ICArUS (Integrated Cognitive Neuroscience Architectures for Understanding Sensemaking) program. The goal of the ICArUS program is to develop a model of the brain based on observed neuroanatomy that is capable of engaging in high-level cognitive behaviors related to interpretation of geospatial intelligence data (sensemaking).

An integrated model of reward discounting

In a second line of work, I am investigating models of temporal and probability discounting. In general, prospective gains that are delayed or uncertain are less valuable than gains that are available immediately or certain. While real-world decisions typically involve combinations of delay and uncertainty, studies of discounting behavior frequently focus on only one of these domains in isolation. As a result of this segregation, models of decision making may fail to adequately account for a range of observed behaviors.

In Alexander & Brown (2010) I demonstrated that a reinforcement learning model based on temporal difference learning, Hyperbolically Discounted Temporal Difference Learning, captures choice behavior by rats for sequences of rewards which minimize risk. The model is of interest from a purely theoretical perspective in that it is an exact recursive definition of the hyperbolic model of discounting (Mazur, 1987), a feat which was previously considered impossible by reinforcement learning theorists (Dayan & Niv, 2008). Practically, these findings rule out rational accounts of decision making that suggest that temporal discounting solves the problem of maximizing average reward. In work currently in preparation, I show that a well-known model of probability discounting (Rachlin, Raineri et al., 1991) can be mathematically derived based on the assumption that decision makers attempt to maximize the probability of survival in an uncertain environment.

Together, these results suggest a common cognitive mechanism underlying discounting of uncertain and delayed rewards, namely that reward magnitude mediates perception of risk. I am currently developing this intuition into an integrated model of rewarding discounting that I will use to investigate discounting of rewards which are both delayed and uncertain, an area of work that has received relatively little attention. Ultimately, my goal is to relate the abstract cognitive mechanisms suggested by my modeling work with the neural systems underlying discounting behavior.

One such system may involve the neuromodulator serotonin. Recent evidence from single-unit recordings suggests that the activity of serotonergic neurons in the dorsal raphe nucleus of monkeys may be involved in reporting reward magnitude (Nakamura, Matsumoto et al., 2008), and dysfunction of the serotonergic system is implicated in impulsive choice behavior. Preliminary results show that a model of discounting incorporating a serotonin-like signal that indicates the size of an expected reward and is used to mediate risk perception can yield patterns of choices for delayed and probabilistic rewards consistent with observed discounting behavior.

Alexander, W. H. and J. W. Brown (2010). "Computational models of response-outcome predictions as a basis for cognitive control." Topics in Cognitive Science **2**(4): 658-677.

Alexander, W. H. and J. W. Brown (2010). "Hyperbolically discounted temporal difference learning." Neural Comput **22**(6): 1511-27.

Alexander, W. H. and J. W. Brown (2011). "Medial prefrontal cortex function as learning to prediction action outcomes." Nature Neuroscience **14**: 1338-1344.

Dayan, P. and Y. Niv (2008). "Reinforcement learning: the good, the bad and the ugly." Curr Opin Neurobiol **18**(2): 185-96.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. The effect of delay and of intervening events on reinforcement value. Quantitative analyses of behavior. M. L. Commons, J. E. Mazur, J. A. Nevin and H. Rachlin. Hillsdale, NJ, England, Lawrence Erlbaum Associates, Inc. **5:** 55-73.

Nakamura, K., M. Matsumoto, et al. (2008). "Reward-dependent modulation of neuronal activity in the primate dorsal raphe nucleus." J Neurosci **28**(20): 5331-43.

Rachlin, H., A. Raineri, et al. (1991). "Subjective probability and delay." J Exp Anal Behav **55**(2): 233-44.

Statement of Teaching Philosophy

William Alexander
Department of Psychological and Brain Sciences,
Indiana University, Bloomington

My philosophy of teaching has been highly influenced by my experiences as an instructor of an undergraduate research methods laboratory section, an assistant instructor for graduate level statistics courses, and as a postdoctoral researcher helping to mentor graduate students in fMRI analyses. In all cases, the student was required to assimilate and apply new material in order to work with problems in psychological and neuroscientific research. I believe this skill set is critical for student development in all areas of science.

Thus, my goal as an instructor in psychology and neuroscience is to prepare my students to engage in meaningful work in their field. I believe there are three essential components to teaching which allow a student to leave a classroom ready to make the most use of material covered in class. First, the instructor must impart a baseline level of information necessary for their work. Second, the instructor should help students to engage critically with the ideas presented in the class. Finally, the instructor should be able to connect the abstract material presented in class with its potential applications in the relevant field.

Considering the large body of knowledge that must be quickly mastered by a student, it is inevitable that many classes will involve a lecture component. In order to facilitate the acquisition of basic knowledge needed by a student in order to progress in their chosen area, material presented in this format should be made accessible by placing it within the broader context of ongoing work in the subject. From my experience both as a student and in front of a class, I have found that providing specific examples of how a series of facts are applicable outside the classroom greatly enhances a student's ability to retain that information, especially when the student is asked to identify additional instances in which the material might be applied. In developing future courses, I look forward to identifying opportunities to promote increased student interaction in traditionally lecture-based classes.

Of course, merely mastering a set of facts does not necessarily permit one to engage in meaningful work outside the classroom. In order for a student to be able to apply their knowledge beyond a set of homework problems, it is necessary for them to be able to think critically about the ideas presented to them in class. As an assistant instructor for graduate level statistics, I would often field questions about *how* to do a particular problem. A technically correct answer would be to simply rehash the steps involved in computing, say, a t-statistic. A more satisfying approach from my perspective both as an instructor and a student is to query a student's understanding of a problem in order to help them understand *why* a particular analysis is called for. By engaging the student in a dialogue, I hope to encourage them to think in more general terms about how information can be used in their own work. This approach has worked well in my previous teaching experiences. Indeed, as one student wrote in an evaluation: "He explains materials so that I am able to do research in this field."

Perhaps one of the most difficult transitions a student faces is taking the skills they have developed in the classroom and applying them in their own work. By necessity, the material covered in the classroom will not have the level of complexity an individual is likely to face when they begin their own work. In my own experience, I have found that learning new skills and concepts is facilitated by applying them to the kinds of questions that a student might face in their future work. As an assistant instructor for graduate statistics, I encouraged students to use data from their research in order to gain experience in statistical analyses using real data. Furthermore, I have found it useful to have my students read scholarly articles which illustrate how the material covered in class may be applied by a practicing scientist, and to think critically about such reports in order to develop the student's ability to participate in scientific dialogue. In the words of one student: "I now look at reported findings a lot differently, and don't believe everything I hear right away."

My doctoral and postdoctoral work has equipped me to teach a wide range of classes. My field of research, cognitive neuroscience, is highly interdisciplinary, requiring familiarity with principles of psychology, computer science, neuroscience, cognitive science, as well as statistics and research methods. As noted above, I have extensive experience as an assistant instructor in advanced statistics, an area which is necessary for a wide range of fields. Furthermore, because of my research background, I am capable of leading classes in introductory and upper level psychology courses, including cognitive and biological psychology. I have a broad knowledge in the area of machine learning, a growing subfield of computer science, with direct applications to research in neuroscience. Finally, my ongoing research in neuroscience is highly relevant to courses which include material in the area of cognitive neuroscience and the neural bases of behavior.

In considering my teaching experiences thus far, I have found that students are better able to learn the skills required for success outside the classroom when they are asked to approach a subject as a practitioner in the field rather than as a student taking a class. In developing future courses, I will emphasize the student's role as a scientist-in-training in two ways. First, in order to encourage students to generalize the specific material discussed in class to a broader context, I will include current examples of relevant research in the field in the course material, and assign pertinent readings which apply an idea or concept to a concrete research question. Second, to promote critical evaluation of material, I will give students the opportunity to teach one another by including group activities which encourage them to discuss course material.