**Postdoctoral positions in neural/symbolic AI at Johns Hopkins**

Current machine learning methods such as learning in deep neural networks produce systems that perform well but whose behavior often cannot be satisfactorily understood and explained. This inadequacy compromises trust in artificial intelligence systems and is highly problematic for many practical applications. The mysterious success of opaque cognitive systems also severely limits their scientific value. A preeminent challenge posed by current artificial intelligence, in addition to performing well in practical applications, is to also explain the remarkable success of non-symbolic neural network models in cognitive domains such as language and reasoning where research since antiquity has documented in great detail the crucial role of rule-based symbolic computation.

A general approach to these problems is to integrate neural and symbolic computation. Two postdoctoral research positions are available to conduct research on such integration within a general framework, Gradient Symbolic Computation, that is currently under development by Paul Smolensky and collaborators including those at Microsoft AI & Research Group. Symbol structures are embedded in continuous activation-vector spaces via Tensor Product Representations, and grammars are embodied in connection weights within networks that converge to vectors that encode grammatical symbol structures. Application domains to be pursued include reasoning and natural-language text-generation in question-answering tasks that are particularly demanding of deep comprehension.

The term of both positions is one year, renewable for a second year given satisfactory progress. Funding for the positions has been generously provided by Microsoft Research/Redmond Washington, where related research is under active development. The successful candidates will be housed in the Cognitive Science Department at Johns Hopkins University. In addition to conducting research, they will be expected to share their expertise with colleagues in the Cognitive Science Department and the Center for Language and Speech Processing through leading seminars as well as consultation and/or collaboration with other researchers at Johns Hopkins and Microsoft Research. In addition to considerable experience with coding in AI-related areas, applicants should have strong backgrounds in at least two of the following fields: mathematical analysis of neural computation; formal theoretical linguistics; symbolic AI techniques for reasoning and language processing; large-scale neural-network system-building. Candidates will need to have in hand a PhD in a relevant field upon starting the postdoctoral position at a mutually-agreed upon date, ideally prior to Sept. 1, 2017.

Applications will be considered as they are received until the positions are filled. Applications should be sent to [GSCpostdoc@jhu.edu](mailto:GSCpostdoc@jhu.edu). An application will consist of (i) a letter explaining the relevant knowledge and skills that the candidate would bring to research in the target area as well as the training they seek to obtain from the position, (ii) a research statement, (iii) a CV with links to publications, and (iv) the names and email addresses of 3 referees who may be contacted by the search committee.

Johns Hopkins University is an Equal Opportunity, Affirmative Action employer. Minorities, women, Vietnam-era veterans, disabled veterans and individuals with disabilities are encouraged to apply.