ParLearning 2012

Workshop on Parallel and Distributed Computing for Machine Learning and Inference Problems

http://researcher.watson.ibm.com/researcher/view_project.php?id=2591

Many prominent algorithms for learning and inference are notorious for their complexity. Adopting parallel and distributed computing appears as an obvious path forward, but the mileage varies depending on how amenable the algorithms are to parallel processing and the availability of rapid prototyping capabilities with low cost of entry. The workshop is one of the major meetings for bringing together researchers in **High Performance Computing** and **Artificial Intelligence** to discuss state-of-the-art algorithms, identify critical applications that benefit from parallelization, prospect research areas that require most convergence and assess the impact on broader technical landscape.

Authors are invited to submit manuscripts of original unpublished research that demonstrate a strong interplay between a parallel/distributed computing technique and a learning/inference application, such as *algorithm design* and *libraries/framework* development on **multicore/manycore architectures, GPUs, cluster, supercomputers, cloud computing** that target applications including but not limited to:

- Large scale inference using parallel TPIC models, clustering and SVM.
- Parallel natural language processing (NLP).
- Semantic inference for disambiguation of content on web or social media
- Discovering and searching for patterns in audio or video content
- On-line analytics for streaming text and multimedia content

- Comparison of various HPC infrastructures for learning
- Large scale learning applications in search engine and social networks
- Distributed machine learning tools (e.g., Mahout and IBM parallel tool)
- Real-time solutions for learning algorithms on parallel platforms

Manuscripts must be submitted by **Dec. 19, 2011**, not exceeding **10 single-spaced double-column pages using 10-point size font on 8.5x11 inch pages** (IEEE conference style), including figures, tables, and references. Accepted papers will be included in the *Proceedings of the IEEE International Symposium on Parallel & Distributed Processing, Workshops and PhD Forum (IPDPSW)*, indexed in EI and possibly in SCI.

General Co-chairs:
Sutanay Choudhury, Pacific Northwest Natl. Lab, USA
George Chin, Pacific Northwest National Lab., USA
Yinglong Xia, IBM T.J. Watson Research Center, USA

Local Chair:
Yihua Huang, Nanjing University, China

Keynote Speaker:
Haixun Wang, Microsoft Research, China

Program Co-chairs:
John Feo, Pacific Northwest National Lab, USA
Chandrika Kamath,
      Lawrence Livermore National Lab, USA
Anshul Gupta,
      IBM T.J. Watson Research Center, USA

Program Committee:
Yi Wang, Tencent Holdings Lt., China
Enhong Chen, Univ. of Sci. & Tech. of China, China
Yihua Huang, Nanjing University, China
Weizhu Chen, Microsoft Research, China
Yangqiu Song, Microsoft Research, China
Haixun Wang, Microsoft Research, China
Edmond Chow, Georgia Tech, USA
Tina Eliassi-Rad, Rutgers University, USA
Lawrence B. Holder, Washington State U., USA
Yan Liu, University of Southern California, USA
Arindam Banerjee, University of Minnesota, USA
George Chin, Pacific Northwest National Lab, USA
Oreste Villa, Pacific Northwest National Laby, USA
Sutanay Choudhury, Pacific Northwest Natl. Lab, USA
Yinglong Xia, IBM T.J. Watson Research Center, USA
Jun Wang, IBM T.J. Watson Research Center, USA
Jatin Chhugani, Intel Corp., USA
Arindam Pal, Indian Institute of Technology, India
Lexing Xie, Australian National University, Australia