

Post-doctoral position in multi-agent learning, bandit algorithms, and computer games

Context:

ANR EXPLORA Project (<http://sites.google.com/site/anexplora/>),
multi-agent learning (MAL), UCT (UCB for Trees), multi-armed bandit problem (MAB).

Objectives:

- (1) Experimental evaluation of MAL algorithms on matrix or stochastic games,
- (2) Many-armed bandit problem, continuous action games such as Voronoi game (VG),
- (3) Other attractive objectives according to the above context and the profiles of candidates.

Location:

LIPADE (<http://lipade.math-info.univ-paris5.fr/>),
UFR mathématiques et informatique (<http://www.math-info.univ-paris5.fr/>),
Université Paris Descartes (<http://www.parisdescartes.fr/>),
45 rue des saints-pères 75006 Paris France.

Contact:

Bruno Bouzy (e-mail: bruno.bouzy@parisdescartes.fr, <http://www.mi.parisdescartes.fr/~bouzy>)

Required skills:

Programming (C, Java, C++),
machine learning, game theory, computer games, Monte-Carlo, tree search,
fluent English language, PhD in computer science.

Length: 12 months.

Dates: Start: between September 1st and December 1st 2010,
End: before November 30th 2011.

Salary: 23,000 euros/year ~ 1,920 euros/month.

To apply: send one pdf containing:

a motivation letter with a statement of interest, a CV, 2 recommendation letters.

References:

Y. Shoham, R. Powers, T. Grenager, If multi-agent learning is the answer, what is the question? AI Journal, 171, 7, May, pp. 365-377, (2007).

B. Bouzy, M. Métivier, Multi-agent Learning Experiments on Repeated Matrix Games, accepted at ICML-10 (2010).

L. Kocsis, Cs. Szepesvári, Bandit Based Monte-Carlo Planning, ECML, Springer-Verlag, LNCS/LNAI 4212, September 18-22, pp. 282-293, (2006).

The Voronoi Game page: <http://www.voronoi-game.com/>