

# Particle methods for POMDPS

Nando de Freitas, Arnaud Doucet, Matt Hoffman,  
Ruben Martinez-Cantin , Julia Vogel

**Abstract:** I will briefly describe two ways in which particle methods can be used to solve complex POMDPS in non-Gaussian continuous state and action spaces. The first builds upon a recent paper of Marc Toussaint and Amos Storkey, where they treat the problem of planning as a structured inference task. Our extension is based on non-trivial particle smoothing methods. The second is based on our work on model-based direct policy search with response surfaces and re-planning. I will illustrate these techniques in the domains of learning in games, optimal exploration and map learning with mobile robots, and decision-theoretic dynamic scene understanding. The poster is a summary of recent joint work with Arnaud Doucet, Matt Hoffman, Ruben Martinez-Cantin and Julia Vogel.