Statement of Research Interests — Joelle Pineau

My research is motivated by the desire to build intelligent autonomous systems meant for human interaction. My objectives are two-fold. First, I aim to develop broadly applicable probabilistic representations and algorithms that can address the problem of planning and control under uncertainty. Second, I am committed to designing and implementing real-world intelligent systems that operate based on these techniques.

Intelligent machines are starting to permeate human-centered environments such as offices, homes, and hospitals, where they will be called upon to assist human users. The scientific challenge is then to build systems which are sufficiently robust to handle the uncertain and dynamic nature of the task domain, as well as reason appropriately about partial, missing, or conflicting information. Traditional approaches to building autonomous systems are inadequate to deal with this kind of uncertainty.

The Partially Observable Markov Decision Process (POMDP) is widely accepted as a general framework for decision-making under uncertainty. Unlike traditional planners which seek an action-selection strategy over *world states*, the POMDP seeks an action-selection strategy over *information states*. This allows us to capture uncertainty, but is often computationally prohibitive. In my thesis work, I have proposed two complementary algorithmic approaches to improve scalability. One is based on a selective sampling of the information state which provides a bounded approximation to the optimal strategy. The second is based on a structured decomposition of the problem, providing an efficient divide-and-conquer solution within the POMDP framework. These techniques improved state-of-the-art POMDP planning by an order of magnitude; whereas earlier approaches solved problems with fewer than 100 dimensions, we can now address problems in excess of 1000 dimensions.

Throughout my doctoral work, I have been involved in the Nursebot project, with the goal of developing mobile robotic assistants for elderly people with mild physical and cognitive disabilities. My major contribution with this project was to field a robot which incorporated full probabilistic decision-making. Field tests at a nursing facility showcased the robot guiding a user, offering information, and providing cognitive reminders.

One of the major challenges for interactive intelligent systems is the ability to conduct natural (unconstrained) dialogue with a variety of users. This problem is particularly well-suited for probabilistic decision-making. It requires the generation of appropriate actions, in response to a signal that is corrupted by poor articulation, ambient noise, and ambiguous utterances. Using POMDP-based approaches, we can automatically learn dialogue strategies which are robust to uncertainty. I believe this could have a tremendous impact on the acceptance of intelligent systems in human-centered environments. The long term goal is to extend the decision-theoretic approach to a wider array of such problems.

Research on probabilistic decision-making is only beginning. Intersecting the fields of machine learning, optimization and control, it provides numerous scientific challenges, both theoretical and applied. I believe my early contributions put me in an excellent position to tackle these challenges. Ultimately, the impact of this work will be measured through the seamless integration of intelligent systems in everyday workplaces and living environments.