EPFL-Inria
INTERNATIONAL LAB

Workshop EPFL-Inria
15 et 16 février 2018, Paris

Robert West, EPFL

Title: « Privacy-Preserving Classification with Secret Vector Machines »

Abstract:
Standard supervised machine learning techniques for classification typically gather the features and labels of all data points in a matrix and vector, respectively, and then solve an optimization problem derived from this input data. For many important settings, such a setup is not acceptable, however, because it reveals all feature and label information to the learner, which is problematic if the data is sensitive, e.g., if the task is to predict user properties (such as gender, income, sexual orientation) from features such as Web-browsing histories, email text, etc. To alleviate this problem, we consider the task of training linear classification models in a decentralized client--server setting, where the server can learn an accurate model without learning any sensitive facts about individual data points (e.g., which data point is associated with which features or label). We show that the support vector machine (SVM) loss function is particularly well suited for this setting, giving rise to a model we term "secret vector machine". We demonstrate feasibility and performance by experimenting with the task of predicting user gender from tweet text and Web-browsing histories.