1135-68-662 **Lorena Mejia Domenzain*** (lorenamejiadomenzain@outlook.com), Porfirio Diaz #163 casa 6, 10200 Ciudad de Mexico, Mexico, **Natalia Dibbern** (ndibbern@lion.lmu.edu), 1 LMU Drive MSB, Los Angeles, CA 7559, and **Harjasleen Malvai** (jasleenmalvai@gmail.com). *Taming Information Leaks in Machine Learning.*

A supervised machine learning algorithm takes a dataset along with known labels as input and outputs a learned function which is used to make predictions about new data points. However, using machine learning could compromise privacy at various stages: private information about the training data could be inferred from the output of the algorithm, or the learned function applied to a new data point could reveal private information about the data point (which we call feature vector). Preserving privacy while providing utility poses a set of interesting problems, the first of which is finding precise, achievable definitions for the concepts of privacy and utility in the context of machine learning algorithms. Semantic security, which implies full privacy, is impossible to achieve in general. We present a weaker definition of privacy, differential privacy, which quantifies the level of privacy attained as a measure of the probability of having similar outcomes from close inputs. We extend differential privacy to quantify feature vector privacy, and show methods of achieving privacy for training data and feature vectors under these definitions. We also pick measures of utility and then present trade-offs between utility and privacy for various machine learning algorithms. (Received September 12, 2017)