Error Bounds for Domain Adaptation

Speaker: Werner Zellinger, RICAM

Abstract: A core assumption of most machine learning systems is that future application data is drawn from the same distribution as the training data. However, this assumption is often violated: Medical diagnostic systems are not restricted to physical human variations present in the training data, and application machines often follow setups different from the training machines. Unsupervised domain adaptation uses unlabeled data from the application distribution to improve the learning on labeled data from the training distribution. In this talk, we provide first error bounds for the main two methods in this field, importance weighted least squares and domain-invariant representation learning, using methods from inverse problems.