

Stronger contracts in machine learning

Matthew J. Holland

Osaka University

When can one guarantee non-trivial performance of a learning system with high confidence? This perfectly reasonable question is decidedly harder to answer in machine learning when compared to other engineering disciplines. Indeed, there exists a large gap between what we can guarantee in principle, and what is actually achieved in practice. In this talk, we consider a more modern definition of “robust” statistical learning systems, and explore some new avenues to achieve such robustness. Bridging the gap between theory and practice is crucial both for engineering reliable learning systems, and training people to design, use, and maintain such systems.