

# Novel design-based inference for adaptive experiments: enhancing efficiency while preserving regret optimalities

Nina Deliu

Using reinforcement learning and bandit algorithms to conduct adaptive randomised experiments can minimise regret, and thus maximise participants benefits. However, it poses major challenges for statistical inference (e.g., biased estimators, inflated type-I error and reduced power). Recent attempts to address these challenges typically impose restrictions on the exploitative nature of the bandit algorithm—trading off regret optimalities—and require large sample sizes to ensure asymptotic guarantees. In this talk I discuss a novel exact hypothesis test statistic, uniquely based on the allocation probabilities observed by the algorithm or design, that overcomes such issues, preserving regret optimalities. I illustrate the application of such *Allocation Probability Test to Thompson Sampling*, presenting its theoretical properties and illustrating its finite-sample performances in both simulated and in a real-world experiments.