

THEMATIC SESSION: Discrete Mathematics

Is it easier to count communities than find them?

Fiona Skerman (Uppsala University)

Random graph models with community structure have been extensively studied. For both the problems of detecting and recovering community structure, an interesting landscape of statistical and computational phase transitions has emerged. A natural unanswered question is: might it be possible to infer properties of the community structure (for instance, the number and sizes of communities) even in situations where actually finding those communities is believed to be computationally hard? We show the answer is no. In particular, we consider certain hypothesis testing problems between models with different community structures, and we show in the low-degree polynomial framework that testing between two options is as hard as finding the communities. Our methods give the first computational lower bounds for testing between two different "planted" distributions, whereas previous results have considered testing between a planted distribution and an i.i.d. "null" distribution. Joint work with Cynthia Rush, Alex Wein and Dana Yang.