Anti-differentiating approximation algorithms for min-cuts and new relationships between PageRank, spectral, and localized flow.

David Gleich, Purdue University

We formalize and illustrate the general problem of algorithmic anti-differentiation: given an algorithmic procedure, e.g., one for which worst-case approximation guarantees are available or one that has been engineered to be practically-useful but for which a precise theoretical understanding is lacking, an algorithmic anti-derivative is a precise statement of an optimization problem that is exactly solved by that procedure. We explore this concept with a case study of finding locally-biased partitions in data graphs, demonstrating connections between min-cut objectives, a personalized version of the popular PageRank vector, and the highly effective "push" procedure for computing an approximation to personalized PageRank. For example, we show that this latter algorithm solves (exactly, but implicitly) an 1-regularized 2 -regression problem, a fact that helps to explain its excellent performance in practice.