In this talk, I will introduce hinge-loss Markov random fields (HL-MRFs), a new kind of probabilistic graphical model that supports scalable collective inference from richly structured data. HL-MRFs unify three different approaches to convex inference: LP approximations for randomized algorithms, local relaxations for probabilistic graphical models, and inference in soft logic. I will show that all three lead to the same inference objective. HL-MRFs typically have richly connected yet sparse dependency structures, and I will describe an inference algorithm that exploits the fine-grained dependency structures and is much more scalable than general-purpose convex optimization approaches. Along the way, I will describe probabilistic soft logic, a declarative language for defining HL-MRFS.