

Locally-biased graph algorithms are algorithms that attempt to find local or small-scale structure in a typically large data graph. In some cases, this can be accomplished by adding some sort of locality constraint and calling a traditional graph algorithm; but more interesting are locally-biased graph algorithms that compute answers by running a procedure that does not even look at most of the graph. This corresponds more closely with what practitioners from various data science domains do, but it does not correspond well with the way that algorithmic and statistical theory is typically formulated. Recent work from several research communities has focused on developing locally-biased graph algorithms that come with strong complementary algorithmic and statistical theory and that are useful in practice in downstream data science applications. We provide a review and overview of this work, highlighting commonalities between seemingly-different approaches, and highlighting promising directions for future work.