This lecture continues our theme of succinct data structures. Building on our expertise with succinct tries, we'll continue on to the practical problem of indexing text with suffix arrays/trees, where several compact and some succinct results are known, with reasonably small (though suboptimal) slowdowns in query time. We'll cover a simple version (not the best) of the historically first such result, which is a compact suffix array that incurs a $\log^{\varepsilon}$ overhead in queries. Then we'll see how to turn any suffix array into a suffix tree with very little (succinct) additional space overhead. We'll get to use our friends (rank, select, balanced parentheses) from last lecture.

6.851 Advanced Data Structures
Spring 2012