The problem of maximizing a concave function $f(x)$ in a simplex S can be solved approximately by a simple greedy algorithm. That algorithm can find a point x_k, in k steps, so that $f(x_k) \geq f(x^*) - O(1/k)$, where $f(x^*)$ is the maximum value of f in S. This algorithm and analysis were known before, and related to problems of statistics and machine learning, such as boosting, training support vector machines, regression, and density mixture estimation. In other work, coming from computational geometry, the existence of ε-coresets was shown for the minimum enclosing ball problem, by means of a simple greedy algorithm. These two algorithms, and some other similar greedy algorithms, are all special cases of the classical Frank-Wolfe algorithm. I'll tie these results together, review some stronger convergence results, and generalize or strengthen some coreset bounds.