Talk outline:

1. State the problem, without explaining the terms involved: given a normalized monotone submodular function f over a domain U and a matroid M over U, compute $\max_{S \in M} f(S)$. What we get: a $1-1/e$ approximation algorithm, which is the optimal approximation ratio.

2. State the special case of (weighted) maximum coverage.

3. Basic facts about maximum coverage: NP-hard, NP-hard even to approximate better than $1-1/e$, the greedy algorithm.

4. Properties of coverage functions: normalization, monotonicity, diminishing returns (introduce the notation $f_A(x)$). Mention that the greedy algorithm is still effective.

5. MAX-SAT as a coverage problem over a partition matroid.

6. Bad example for the greedy algorithm: $(x \vee y) \wedge x \wedge \neg x$, with weights $1, \epsilon, 1$. Approximation ratio for the greedy algorithm.

7. Continuous greedy paradigm: find a continuous relaxation, solve relaxation, lossless rounding. We want a combinatorial algorithm!

8. Greedy doesn't work. Next simplest algorithm: local search. Also doesn't work.

9. Non-oblivious local search. Outline of "ideal" algorithm (we can worry about details later).

10. How to find the best auxiliary objective function, and what we get:

$$g(A) = \sum_{B \subseteq A} f(B) \int_0^1 p^{|B|-1}(1-p)^{|A|-|B|} \frac{e^p}{e-1}\, dp.$$

11. What it amounts to in the special case of MaxCover: $\ell_0 = 0$, $\ell_1 = 1$, $\ell_{k+1} = (k+1)\ell_k - k\ell_{k-1} - \frac{1}{e-1}$.

12. Probabilistic interpretation of $g_A(x) = \sum_{B \subseteq A} f_B(x) \int_0^1 p^{|B|}(1-p)^{|A|-|B|} \frac{e^p}{e-1}\, dp$. Mention de Finetti's theorem to peak people's interests.

13. Idea of the proof: the fundamental inequality

$$\frac{e}{e-1} f(A) \geq f(B) + \sum_{i=1}^{r} [g(A) - g(A - a_i + b_i)].$$

14. Randomized vs. deterministic algorithms.