Adrian Albert
Stanford University

Thursday, November 11, 2013
Where: 741 JMHH
When: 3:00 PM to 4:20 PM

DATA-DRIVEN MARKETING FOR SUSTAINABILITY ON THE ENERGY GRID

ABSTRACT: Utility companies in the US and Europe have deployed massive advanced sensing infrastructure (smart meters) to collect energy consumption data at fine (sub-hourly) time scales from large populations of users. Yet little understanding exists currently about how such "big data" might be used to improve operational practices and to better connect and engage with customers. A major promise of the future smart grid is the effective implementation of Demand Response (DR) strategies that dynamically match consumption on the grid with available supply in real-time. But, as utilities do not operate in a pure market context – they are heavily regulated, not-for-profit entities, and since residential consumers' price elasticity of energy use is low, typical market based strategies such as differential pricing are not effective for affecting residential demand.

The perspective I take in this work is that of data-driven energy program targeting, i.e., using smart meter readings for identifying high-potential types of users for certain demand-response and energy-efficiency programs, and reaching them with tailored marketing communications and incentives to improve their consumption behavior. Here I focus on thermally-induced consumption, which accounts for more than 40% of energy use in the US. First, I will introduce a dynamic structural model that uses hourly electricity and weather readings to characterize residential users' thermally-sensitive consumption. The model assumes that consumption is the observable outcome of latent, temperature-dependent decisions of users on whether to heat, cool, or not to use HVAC at all. Mathematically, this structure may be expressed through a non-homogeneous hidden Markov model. In effect, the model performs a coarse decomposition of a user's consumption into a base load (non-controllable), and a thermal response load (controllable). From this model I extract useful benchmarks to build profiles of individual users for use with thermal (heating or cooling) DR programs. I compare profiles generated using the model with ground truth data about HVAC operation for several real consumers, for which these end-use data is available. I then discuss the performance of the model on a large sample of residential users, and compute metrics that allow us to segment the population dynamically for the purpose of a thermally-motivated DR program. I show that a simple targeting strategy that takes into account the users' temperature-dependent thermal response may offer substantial savings over current strategies.

I then briefly describe my current work on formalizing the targeting problem when users' consumption may be structured using the model introduced above. In this setting, the utility may contact a given user to ask them to reduce their thermal usage (or control it consumption remotely) up to a limited "effort" budget. This situation introduces an opportunity cost, since marketing actions at a given point in time limit the scope of future communications. I show preliminary results on optimal marketing communication strategies in simple formulations of this problem.

Finally, I describe my past and current efforts to relate such model results with the characteristics of the actual users. Can usage patterns be used to learn about household lifestyle? To address this question, I characterize the different usage regimes using magnitude and variability (the rate of switching between regimes). Using this model, I show that certain user characteristics (appliances, lifestyle) may be inferred using features computed from the data: the presence of certain appliances is chiefly predicted by consumption magnitude, whereas lifestyle (e.g., unemployed) are best predicted by regime variability, with as much as 20% improvement over random guessing.