EPSRC Engineering and Physical Sciences Research Council

ORCHID

Non-intrusive Load Monitoring using Prior Models of General Appliance Types

Oliver Parson, Siddhartha Ghosh, Mark Weal, Alex Rogers

Agents, Interaction, and Complexity Research Group

School of Electronics and Computer Science

University of Southampton

Scenario

Smart meters are being deployed by many countries on national scales (all houses in the UK by 2020)

In home displays have access to low resolution data from smart meters (5 second intervals in UK)

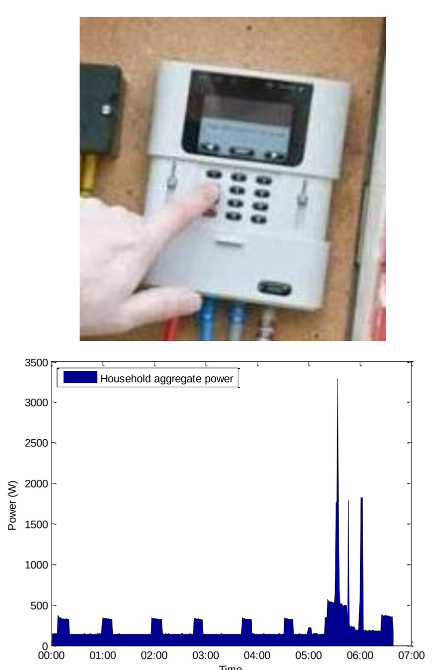

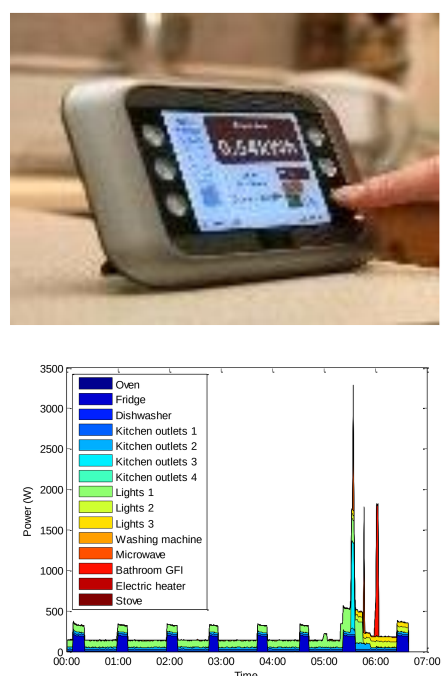

Aim: Provide disaggregated feedback of the energy usage of individual appliances to the household occupants to empower them to optimise their energy use

Existing Approach: FHMM

Existing approaches have modelled households using factorial hidden Markov models

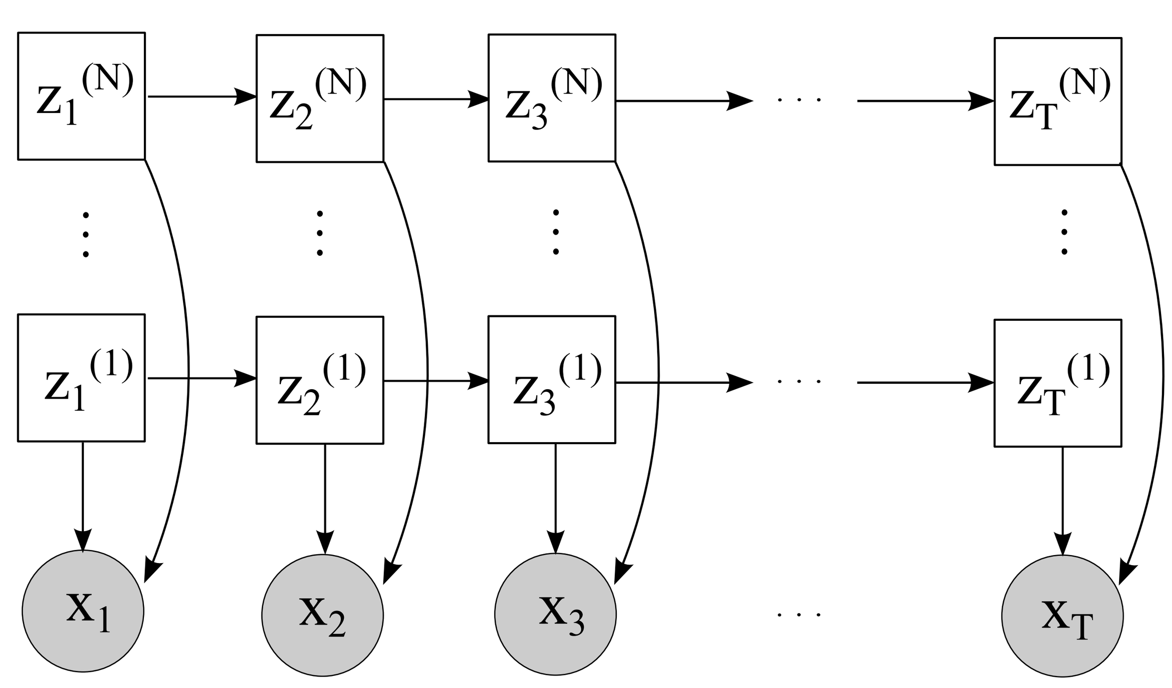

$z_t^{(n)}$ – hidden state of appliance n in time slice t

x_t – observed sum of emissions in time slice t

However, training approaches perform poorly when sub-metered data is unavailable

Prior Models of Appliances

We construct general models of appliances from sub-metering many appliances of the same type

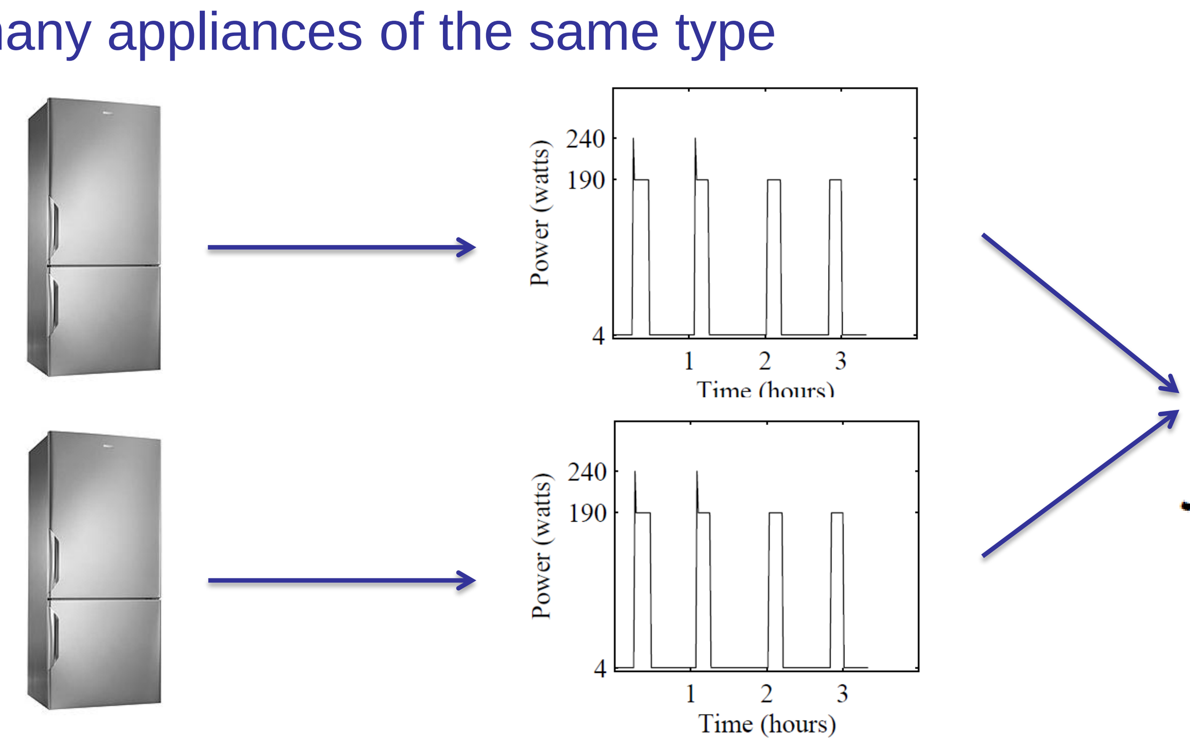

$$A_{i,j}$$

$$\mathcal{N}(\mu_{z_t}, \sigma^2_{z_t})$$

We learn a single set of HMM parameters that will generalise well using large amounts of sub-metered appliances

- Transition matrix calculated from frequency of on/off events
- Gaussian distribution calculated from measured power demand

Tuning Prior Models using Aggregate Data

1. Prior Appliance Models
   - We tune prior appliance models to specific appliance instances using only aggregate data
   - We do this by identifying and exploiting periods during which only one appliance is changing state

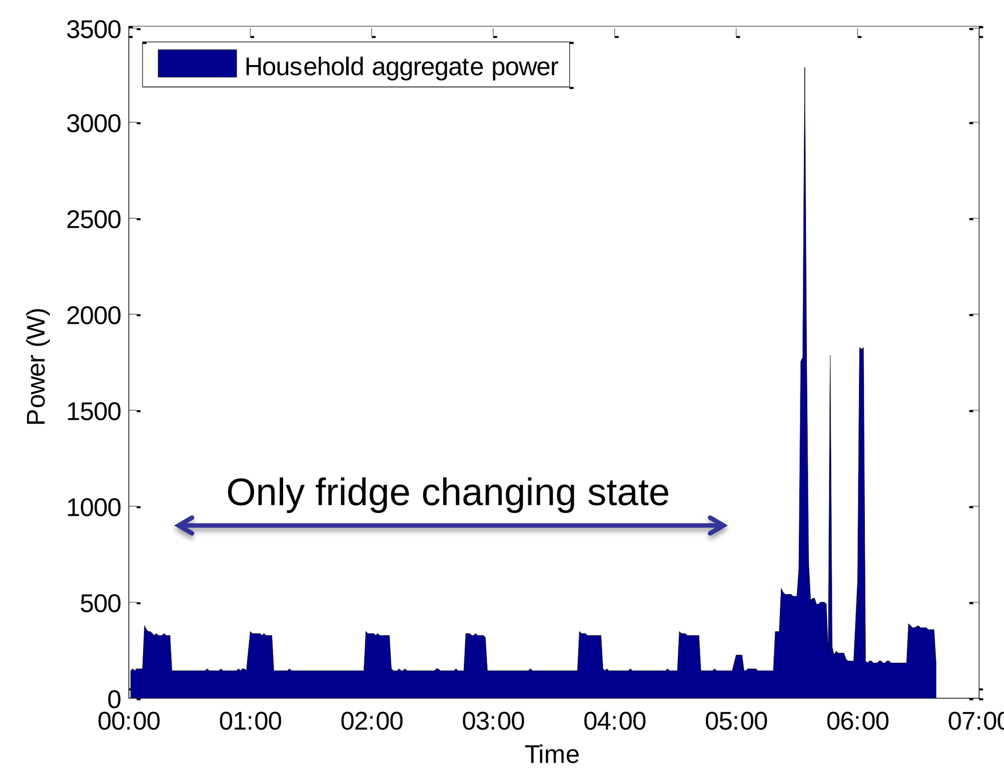

2. Specific Appliance Models
   - We use such specific appliance models as the input parameters to a FHMM
   - We apply variants of exact inference (Viterbi) and approximate inference algorithms (sampling) to determine appliance usage

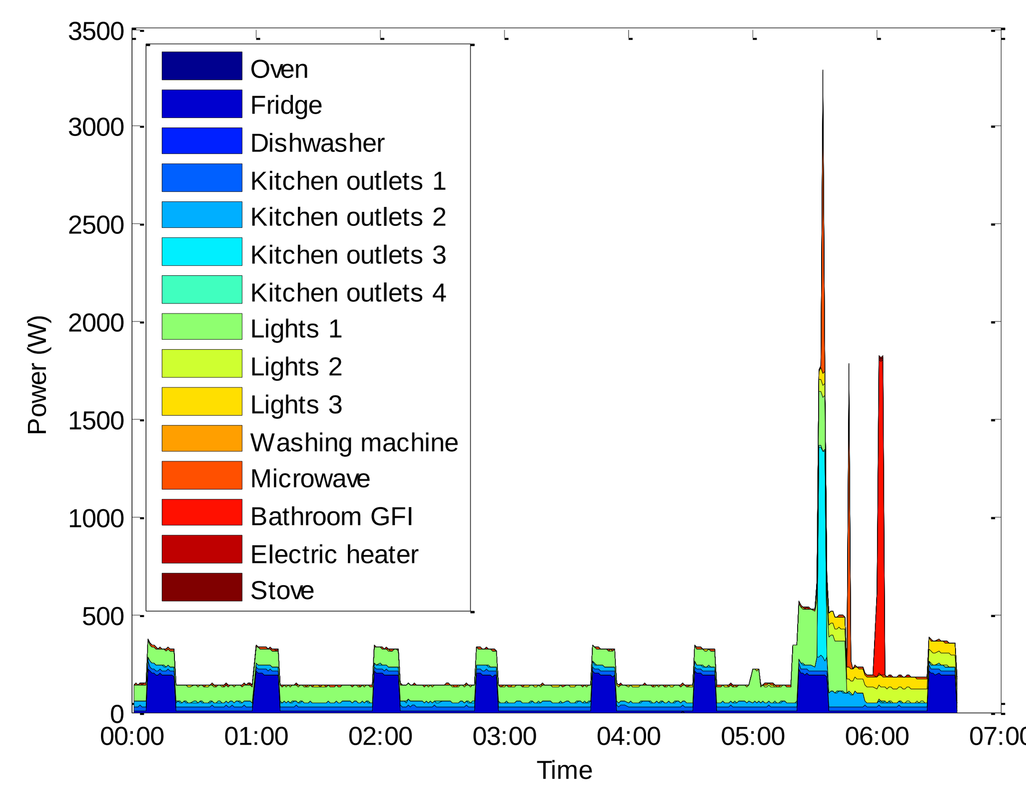

Evaluation

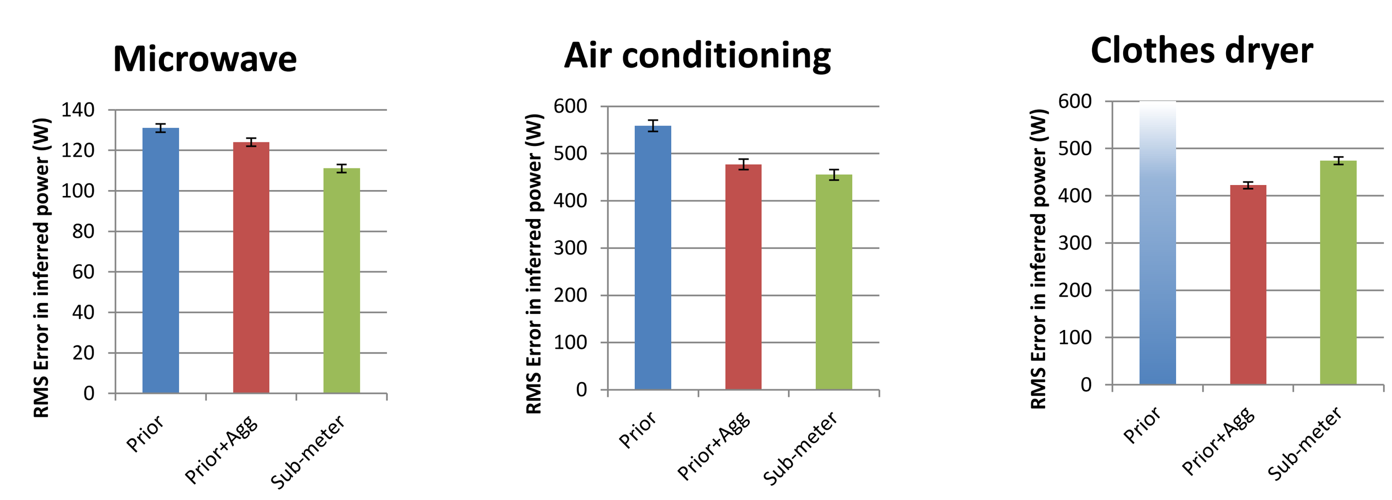

- Our approach which uses aggregate data to train the model outperforms a variant in which the prior is not trained

We deployed our approach to real homes in Southampton:

- We applied our approach to a number of real houses in Southampton
- This allowed us to quantify the savings of recommendations
- In this example, the household can save £39 or 131 $kgCO_2$ at the press of a button

BAE SYSTEMS

ACFR AUSTRALIAN CENTRE FOR FIELD ROBOTICS

SECURE

UNIVERSITY OF Southampton