Diversity in Action
The SmartSociety perspective

Michael Rovatsos
The University of Edinburgh

University of Leeds, 9th June 2016

SmartSociety

- Hybrid and Diversity-Aware Collective Adaptive Systems: "*when people meet machines to build a smarter society*"
- 4-year €6.8M FET Integrated Project, co-ordinated by University of Trento
- Brings together AI, computer science, human factors, privacy, ethics

UNIVERSITY OF TRENTO - Italy

UNIVERSITY OF OXFORD

THE UNIVERSITY of EDINBURGH

UNIVERSITY OF Southampton

TECHNISCHE UNIVERSITÄT WIEN
Vienna University of Technology

אוניברסיטת בן-גוריון בנגב
Ben-Gurion University of the Negev

UHOPPER

imaginary

Project Context

- **Hybridity**
  - Vision of people and machines collaborating and complementing each other to tackle hard societal problems
- **Diversity**
  - Diverse populations of interacting humans and machines with different knowledge, skills, objectives, and expectations
- **Collectives**
  - How do we compose individual interactions to obtain collective action and globally coherent social computations?
- **Adaptation**
  - How can we understand and support collective adaptation in a complex socio-technical systems?

Ask & Share

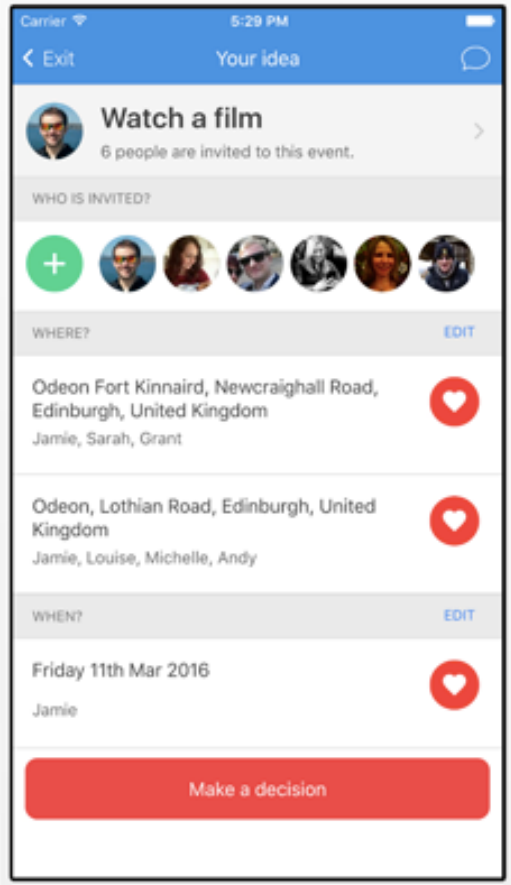
- An app for doing things together
  - **without** knowing the what/ who/how
- Combines
  - human-driven crowdsourcing
  - automated activity recommendation
- Current scenario tourism, but aim for general model of collective problem solving



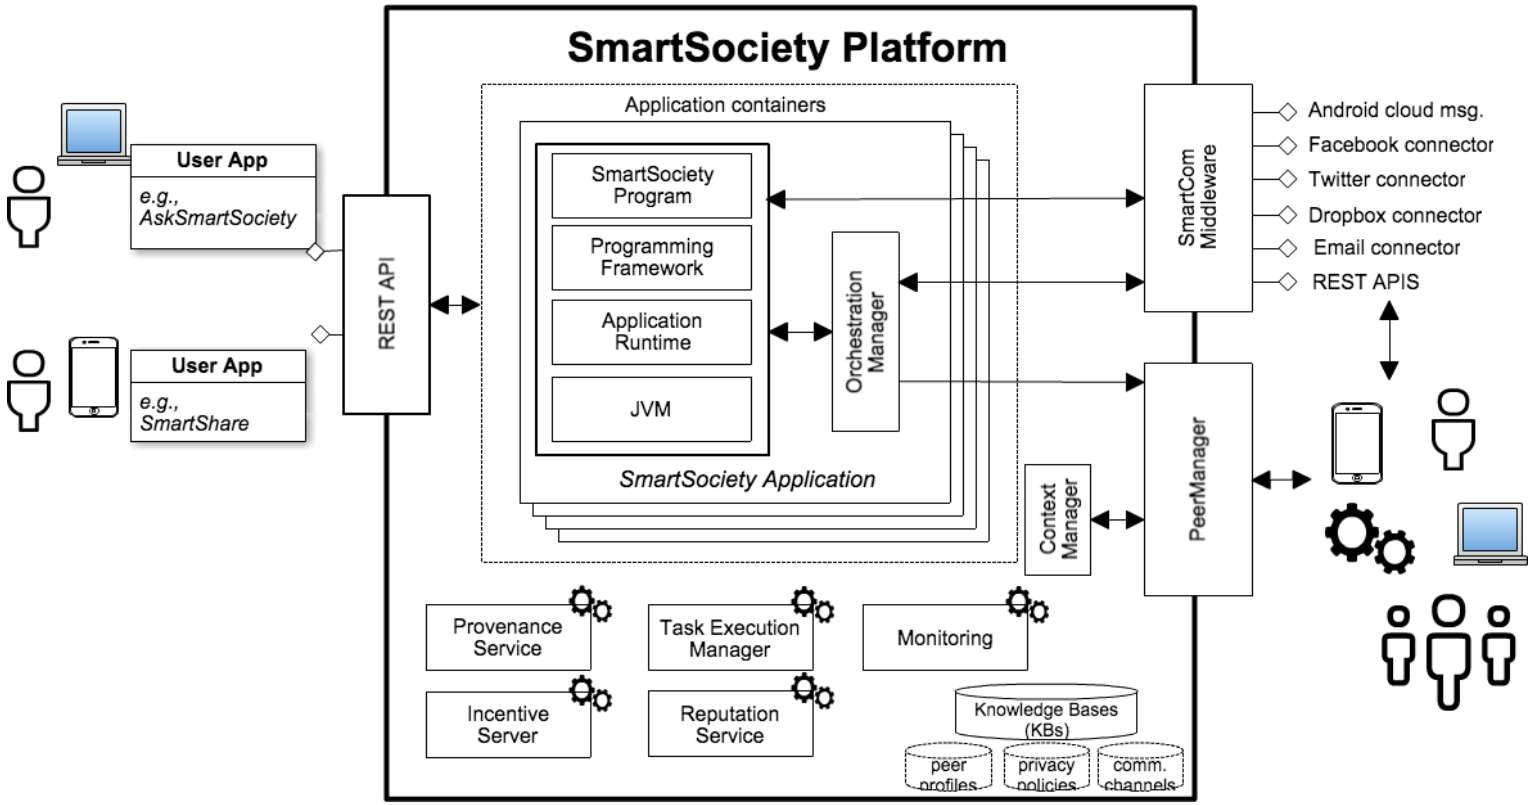
Where's the novelty?

- Conventional apps rely on
  - Static content supplied a priori
  - Search/filtering/recommendation
  - "Browse & pick" feel
- Limitations
  - Limited scalability & flexibility
  - No learning from human input
  - No explicit user/platform goals
  - No incentivisation
  - No transparency

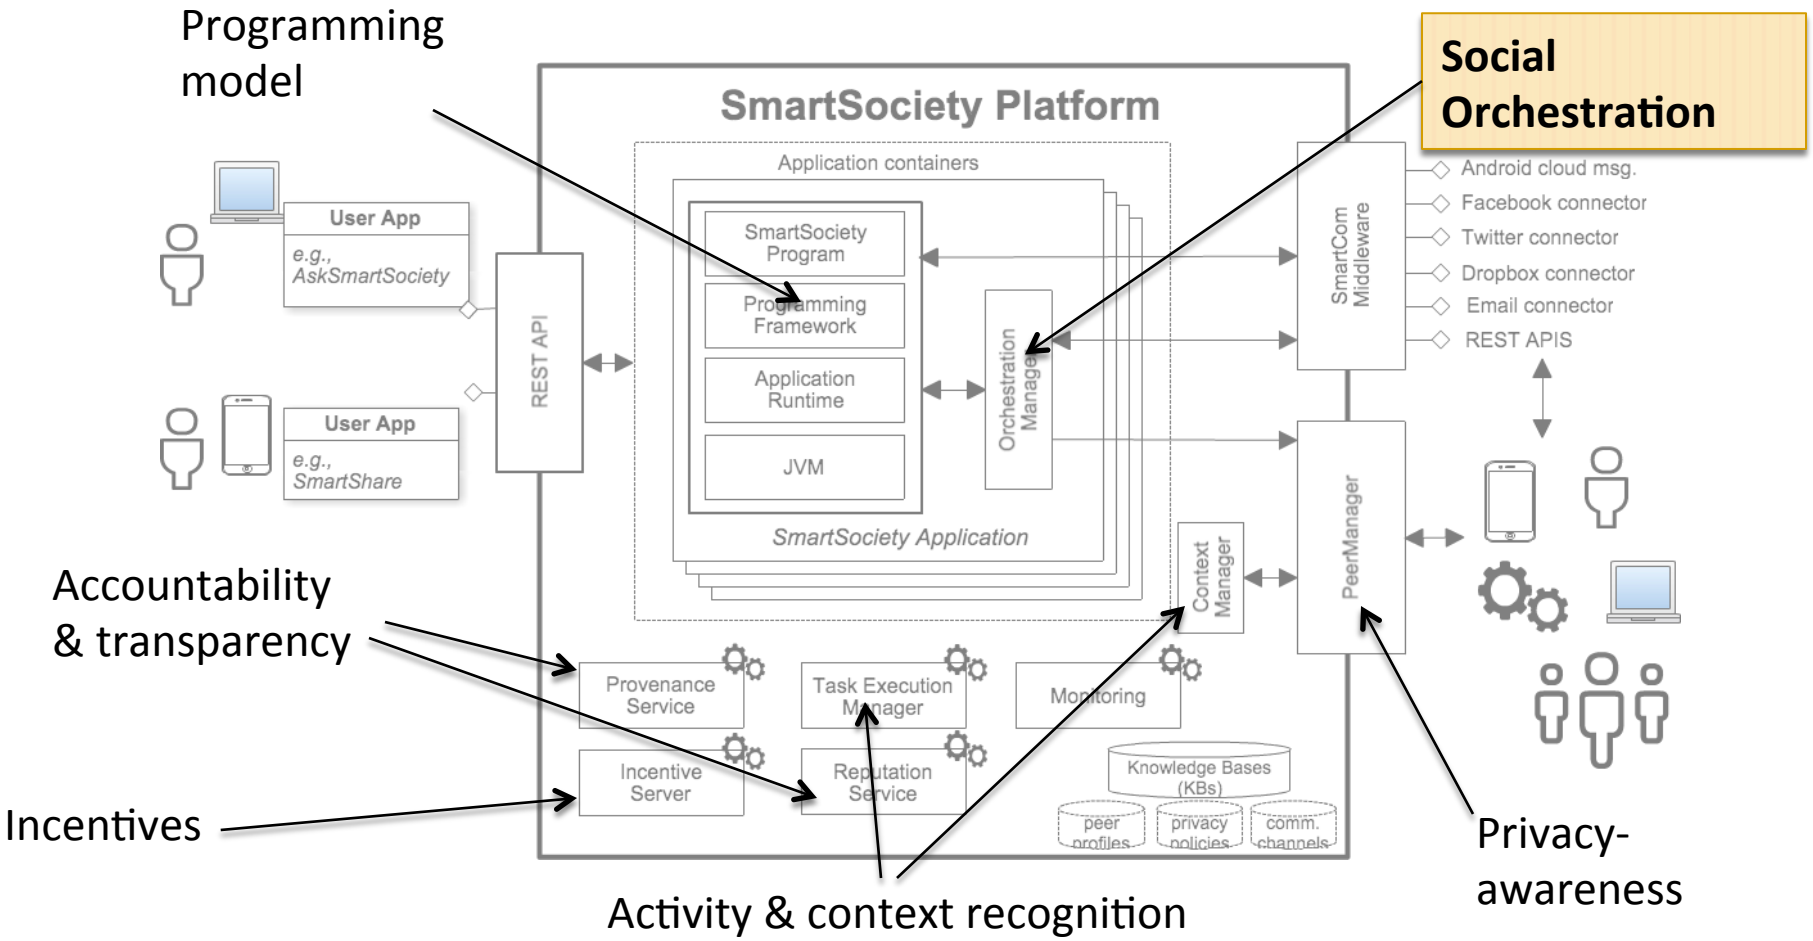
decidz



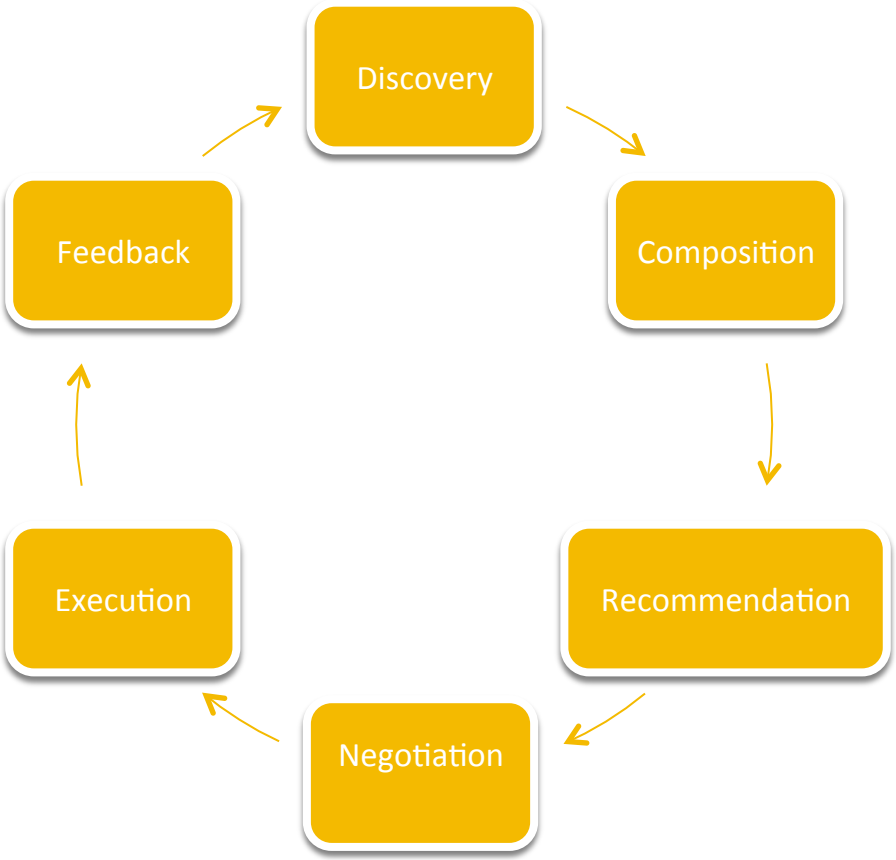
The SmartSociety Platform



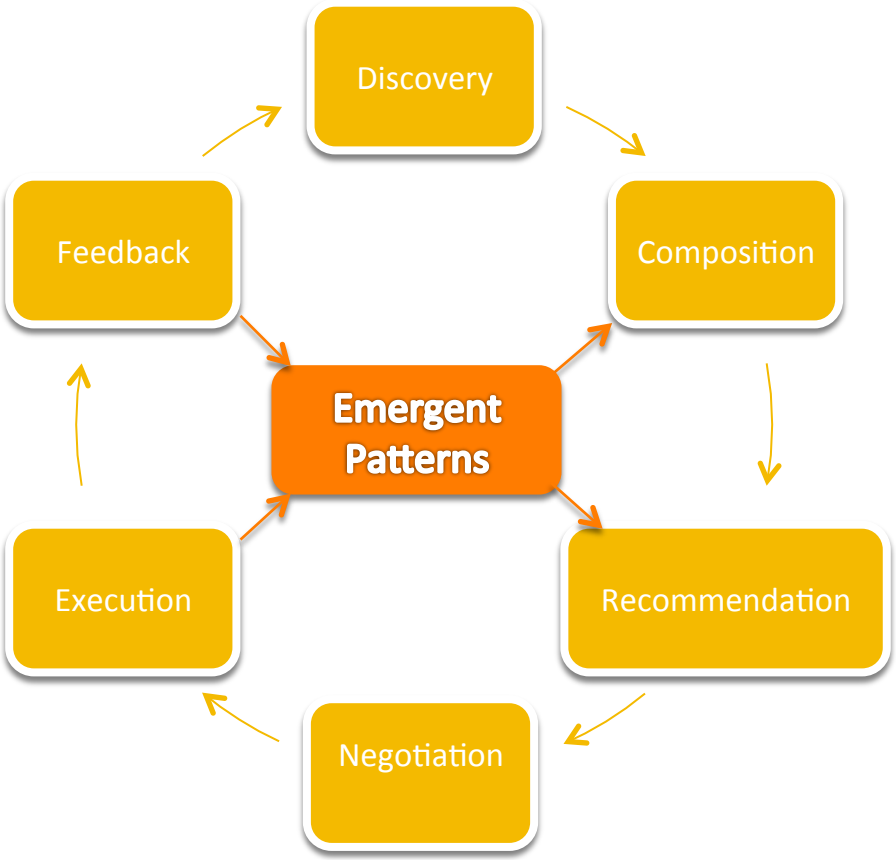
What's in the box(es)



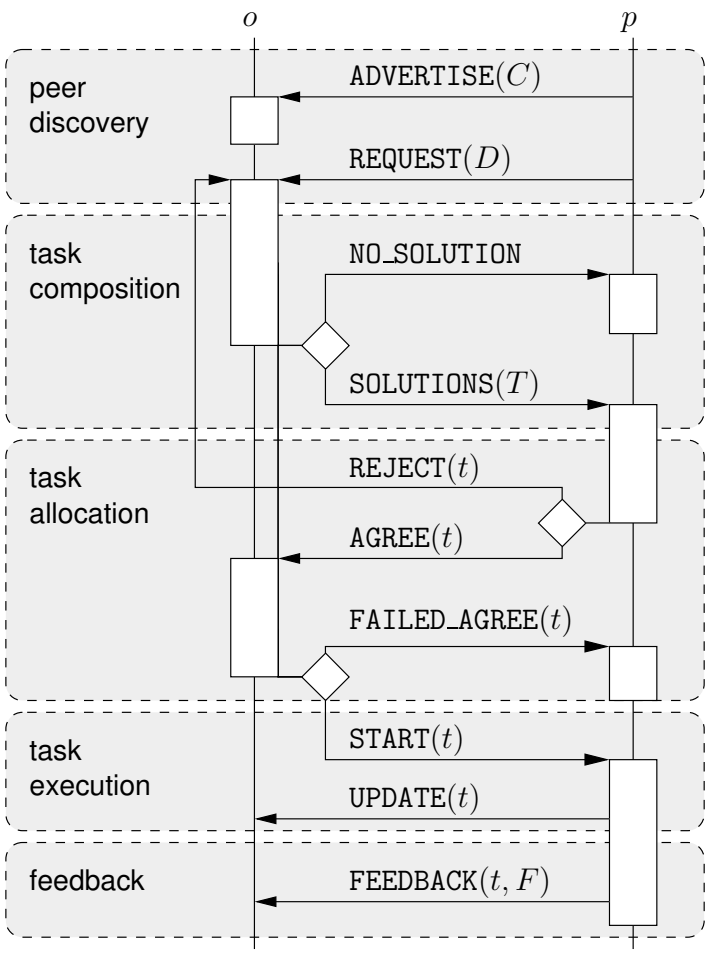
Social orchestration



Adaptive social orchestration



The “agent protocol” view



Task composition

- Combinatorial problem of allocating **groups** of users to **shared** tasks, where **task requests** come from users
- **Hard constraints** restrict the groupings and task properties that can be realised in principle
- **Soft constraints** determine which coalition structures and task features are preferred by system and/or users
- *Examples:*
  - **Ridesharing** where drivers share their cars with other passengers for a specific journey (our vanilla example)
  - Meeting scheduling, citizen science tasks, workforce shift allocation, clinical workflow management, etc etc

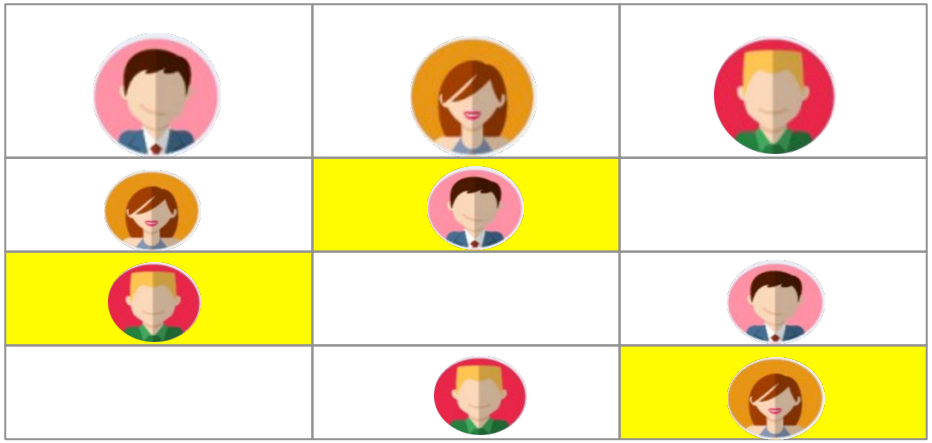
Diversity in task composition

- In traditional mechanism design, global allocations are computed given individual preferences and global criteria
  - E.g. social welfare maximisation, Pareto optimality, strategy-proofness, etc.
- Mechanisms are proposed that provably satisfy these properties, solution can therefore be **imposed** on users
- **Diversity** implies that users cannot report their preferences
  - System never captures all relevant decision variables
  - Solutions cannot be computed/considered exhaustively
  - Utility of solutions cannot be determined by users a priori

From task allocation to task recommendation

Key problems:

1. How to compute “optimal” **sets** of solutions
2. How to **influence** users’ choices



3. How to **learn** users’ preferences

User and system utility

User's utility function u_i depends on user's requirements and preferences

Global (system) utility function depends on social welfare and maximising task completion

$$U_s = \sum_{i \in I} u_i + \sum_{i \in I} \sum_{j \in J} x_{i,j}$$

Computing allocations



MIP*

$\rightarrow V^*$

Objective $max_{a \in A} U_s(a)$

Constraints Hard feasibility constraints

MIP^first

$\rightarrow a \in R$

Objective $min_{a \in A} \sum_{i \in I} \sum_{i' \in I | i' > i} |u_i(a) - u_{i'}(a)|$

Constraints MIP*

$U_s(a) \cdot h \geq V^*$

MIP^others

$\rightarrow a' \in R$

Objective $min_{a' \in A} \sum_{i \in I} |u_i(a) - u_i(a')|$

Constraints MIP^first

$a' \notin R$

Influencing users

We want to modify users' utility artificially so that their choices lead to a feasible global solution

- Explicit Approaches:
  - intervention
  - (possible) future reward
- Implicit Approaches:
  - discounts
  - **taxation**

Taxation

MIP*

$\rightarrow V^*$

Sponsored Solution

MIP$^{\text{first}}$

$\rightarrow a \in R$

MIP$^{\text{others}}$

$\rightarrow a' \in R$

Objective

$$min \sum_{i \in I} |u_i(a) - u_i(a') + \tau_i(a')| + M\Big(\sum_{i \in I} \big(u_i(a) + \epsilon - u_i(a) + \tau_i(a')\big)\Big)$$

Constraints MIP$^{\text{first}}$

$$a' \notin R$$

Noiseless and Constant Noise Models

$$u_i(a) + \epsilon \geq u_i(a') - \tau_i(a')$$

Logit Model (also goes into objective function)

$$\frac{u_i(a)}{\Big(\sum_{a'' \in R} \big(u_i(a'') - \tau_i(a'')\big) + u_i(a') - \tau_i(a')\Big)} \geq \psi$$

Learning users' preferences

- Developed a model of **coarse preferences** in combination with **user typing** to allow for fast preference elicitation
- Should allow us to optimise queries when recommending solutions using exising **active learning** methods
  - Pose queries that also maximise the expected value of information, and include this in the optimisation procedure
- **Challenge:** not suitable for collectives, only for single users
  - Basically optimisation constraints would have to include full Baysian update of probability distributions on preferences

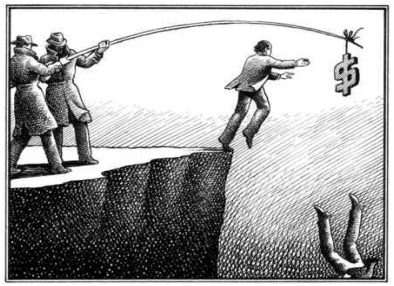
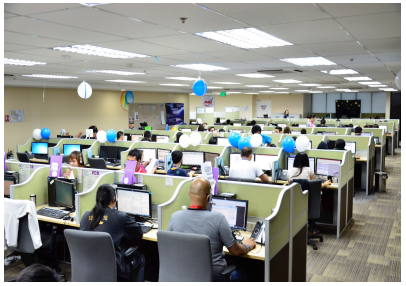
Conclusions

- Creating large-scale collective intelligence that utilises human and machine strengths requires addressing diversity at different levels
- Harnessing the power of diversity involves
  - Relaxing assumptions that otherwise allow us to provide hard guarantees
  - Tackling new algorithmic challenges (and representational ones I have not mentioned

Conclusions

- In return, we get systems that better model a broader range of human decision making
  - Increases likelihood that people will engage with future collective intelligence systems
- I have omitted many other aspects of SmartSociety where diversity also creates new challenges
  - Activity recognition, transparency, privacy, etc.

Promises and perils

Promise		*Peril*	
Man-machine collaboration		Manipulation	
Personalisation		Surveillance	
Collective intelligence		Humans as cheap labour	

What magical trick makes us intelligent?
The trick is that there is no trick.
The power of intelligence stems from our vast **diversity**, not from any single, perfect principle.

- Marvin Minsky