7. Agent-Based Modelling (and Multi-Agent Systems)

Modelling Social Interaction in Information systems

http://davidhales.com/msiis

David Hales, University of Szeged

dave@davidhales.com

What is an agent?

- So obvious that we rarely think about it!
- In philosophy and sociology agency (and intentionality) are major concepts
- An agent is a thing that has the ability to take independent actions to pursue its own goals (it is autonomous and has intentions)
- In sociology “structure and agency” is a big debate. Which is more important in social systems (similar to “nature / nurture” debate in biology)?
- It is an abstraction that allows us to make sense of the world
- Can you imagine any story that would make sense without agency?
- Many things *could* be viewed as an agent but in general it only helps when the behaviour of the thing is complex enough to imply it has autonomy and intentions
- An individual, a group or machine can be viewed as an agent

What is an agent?

- Humans see agents everywhere!
- Consider a child who ascribes agency to a toy
- Think about conspiracy theories (secret agents!)
- Many religions ascribe agency to natural forces: the sun, the sky, the sea
- We (generally) view each other as agents not objects or forces
- We sometimes ascribe agency to institutions or collectives
- No simple definition of what an “agent” is
- Agents in general are considered to be “autonomous” meaning they control their actions rather than being controlled by some other agent

Computational agents

- What is a computational agent?
- It's a computer program (but so is anything we do in computers!)
- Specifically, it's a program that takes inputs (percepts) from an environment and produces outputs (actions) that change that environment
- Not under the "control" of some other agent
- In general the environment is shared with other agents
- The actions of one agent may effect the percepts of other agents

Aside: The evolution of agency

- Sometimes concepts of “freewill” are mentioned in the context of agency
- This is huge and complex area which gets highly philosophical (so we avoid it)
- Yet it is interesting to note that the famous philosopher Daniel Dennett attempted to show how agency might emerge from evolution
- He says the concept of agency is a way of helping us make sense (predict) the behaviour of things
- Put simply, he believes that if something acts “as if” it has intentions and freedom to act then you can productively view it as an agent

Book: Dennett, D. C., (1987) The Intentional Stance, MIT Press, (Cambridge)

Agent-based modelling

- ABM uses the agency abstraction to model some phenomena (often social phenomena)
- We identify the agents, the environment, the percepts and actions
- We put some rules in the agents (that embody goals) and run it on a computer
- We look at the macro structures that emerge through interactions
- In general we don't explicitly specify the agent interactions but rather allow them to emerge in a simulation

Agent based modelling

- Almost all the simulations we have looked at so far can be considered ABM's
- Schelling's model is sometimes considered one of the first ABM's - Schelling didn't use those words however
- If someone says they are using an ABM you can ask:
  - What are the agents?
  - What percepts (inputs) and actions (outputs) do they have?
  - What is the environment?
  - How do agents interact with each other?
  - What rules do the agents use to make decisions?

Types of agents

- One way to classify different ABM is to consider how agents are modelled
- Simple reactive agents
  - Simple mapping from precepts to actions
- Adaptive / learning agents
  - Individual learning (induction / machine learning)
  - Collective or social learning (evolution)
- Logical reasoning / cognitive agents
  - Store some representation of the world (beliefs)
  - Store some goals on how to change the world (desires)
  - Deliberate about, select and execute a plan (intentions)
  - For example the BDI approach (Guest lecture Mario Paulucci will mention this – I think)

Rationality and Bounded Rationality

- Agents can be programmed to act rationally in known contexts
- For example, utility maximisation leads to a Nash Equilibrium for games with known payoff matrix
- However, symbolically, cognitive agents may use logic to deduce the best next action given their current beliefs about the world
- When we use simple rules (heuristics) describe the agents as "boundedly rational" meaning they don't behave in a deductive rational way but rather just do something simple to attempt to achieve their goals
- Consider the agents in the Schelling model. They just move randomly to try improve their situation relative to their threshold value. They do not explicitly try to optimise
- The bounded rationality and "satisficing" concept was introduced by the great Herbert Simon

Multi-Agent Systems

- If we use agents as a software abstraction to solve a practical problem this is called Multi-agent systems (MAS)
- If we use agents as an abstraction to model some phenomena this is called agent-based modelling (ABM)
- Sometimes people model MAS with ABM
- Sometimes people use MAS to implement ABM
- MAS grew out of an area called “distributed artificial intelligence” (DAI)
- The idea of DAI was to work out how to take a big problem and let independent intelligent agents work together to solve it

book: Wooldridge, Michael (2002). An Introduction to MultiAgent Systems. John Wiley & Sons.

Multi-Agent Systems

- From a software point of view, MAS apply the agency abstraction
- This can be compared with the object abstraction that characterised object-orientated programming
- Specific languages have been created such as “AgentSpeak” and “Agent0” which uses logic and the BDI framework
- Specific inter-agent communication languages have been specified such as “KQML”
- However, in general, the agency abstraction has not really taken off (since formulated in the 90’s) I believe
- Why it is not popular is a matter of debate

Video of Wooldridge introducing “Agent0” language: http://youtu.be/oOWR6wvY64E?list=UU21E9Khep4mv5LEcpYRG7tQ

Agent Based Models

- It is easy to implement ABM using an object orientated language
- We specify an object for an agent
- Insatiate a set of agents
- Loop through each agent firing its rules
- Until we decide to stop
- In NetLogo one loop through each agent is called a “Step”

Artificial Societies

- When researchers produce ABM that are not related directly to the real world (more on this later!)…
- These are sometimes called “artificial societies” because they are exploring societies “in general” rather than specific phenomena
- Axelrod's tournaments could be called “artificial societies” to explore cooperation theory
- See online free open access journal: JASSS
- http://jasss.soc.surrey.ac.uk/JASSS.html

Note: if you are looking for a paper to present during the assessment then it is worth browsing this journal to see if there is anything that looks interesting to you

What is a model?

- We think we know what a model is
- It is a representation of some other thing that allows us to understand / predict / describe that other thing
- Traditionally we talk about the Model and the Target
- The Target is the reality that we aim to capture in the model
- We judge the quality of the model by validating it against the Target (i.e. the real world)
- One way to do this is to compare the output of the model with “data” or more recently “big data”

What is a model?

- Some argue that a model that does not predict anything or can not be validated with data is meaningless
- At best a toy to play with at worst a kind of mental delusion
- This is true if you have a very practical job to do – such as for example build a rocket or a bridge
- You want to know if the bridge will stand or the rocket will fly
- You build a model based on current theory, experiment with it and then produce a design
- You then test the thing you have produced to see if it works and make measurements
- You refine your model based on these experiments
- That's engineering and science! (or is it?)
- Perhaps it is a little more complex than this

What is a model?

- On the other hand, “exploratory models” are not necessarily about prediction but understanding
- These a two different things
- For example imagine we had a “black box” which was a model of you. It could predict what you would do with 100% accuracy
- Would this oracle be useful to you? In some sense you are that model yourself
- Put another way, without being able to see into the black box would this give you any greater understanding?

Silly aside: Schopenhauer (my view)

- In his great work the philosopher Arthur Schopenhauer starts with the sentence: “The world is my representation” (English translation)
- What I think he means is that we can never know the world of reality only some representation (or model) of reality
- Since we never have direct access to reality we can only really compare representations (or models)
- In this view “Data” is just a very simple model
- Probably best not to confuse models with data however.. Back to reality…

book: Schopenhauer, “the world as will and idea”, first published 1884, English translation (1909). http://www.gutenberg.org/files/38427/38427-h/38427-h.html

Individual-Based Modelling

- Sometimes people talk of “individual-based models” IBM
- In general IBM and ABM are the same thing but..
- IBM tends to be used when what is being modelled does not include explicit agency
- IBM just means we are modelling individuals that may or may not be “agents”
- Often used in ecological modelling where you’re modelling plants at the individual level
- Sometimes IBM use probabilities for actions (derived from data) rather than rules for actions

Some ABM languages / platforms

- NetLogo (https://ccl.northwestern.edu/netlogo/) – based on logo but runs on java
- Swarm (swarm.org) – emerged out of SFI in 1990's. Originally C# libraries. Not sure of current status.
- Mason (http://cs.gmu.edu/~eclab/projects/mason/) - java libraries for ABM
- Repast (http://repast.sourceforge.net/) - multi language platform
- AnyLogic (http://www.anylogic.com/) – big commercial simulation package that supports ABM in addition to other simulation methods. Can compile down to Java

Examples of ABM's

- Schelling's segregation model (we have seen)
- Axelrod's cooperation tournaments (we have seen) and many subsequent models
- A famous "artificial society" ABM model is called The Sugarscape (using simple rules)
- A less famous ASoc model using cognitive agents is called "EOS" (Evolution of Society)
- These are all exploratory models

The Sugarscape

- Entire book about it published in 1996
- Basic methodology is to “grow in the computer” complex social behaviours in order to understand them
- They are interested in the simplest rules that can lead to a phenomena of interest
- Creates a simple artificial world and experiments with different rules to explore phenomena such as:
  - Evolution
  - Simple trade
  - Formation of cultural groups
  - Social networks and trust relationships

Epstein, Joshua M.; Axtell, Robert (1996). Growing artificial societies: social science from the bottom up. Brookings Institution Press

Sugarscape

- Some simple parts of the model are included in NetLogo models library
- Essentially the model comprises:
  - 2D grid in which are placed resources called “sugar” in non-random arrangements (sugar peaks)
  - Agents move about on the grid and consume sugar to get energy (sugar grows back)
  - If agents run out of energy they die
  - Agents reproduce, form networks, trade, fight etc.

Sugarscape

- The basic methodology in the book:
  - start with simple behaviour (e.g. moving randomly to find sugar)
  - Examine the emergent behaviour
  - Progressively add more complex rules to see how they change things
  - The rules are added to explore different kinds of social phenomena
- Consequently the Sugarscape is not a single model but rather a set of models

Sugarscape

- We will only consider a few results from the model in overview
- If you wish to understand the detail you will need to look at the book
- We will look at some general results from:
  - Chapter 3 - Sex, Culture and Conflict: The Emergence of History
  - Chapter 4 – Sugar and Spice: Trade Comes to the Sugarscape

Sugarscape – chapter 3

- Wish to "grow" a "caricature" of history. **Start with** "a social story":

"In the beginning there is a small population of agents, randomly distributed both in space and with respect to their genetic characteristics. Over time spatial agglomeration into two groups occurs as each agent – guided by the primal sugar drive – migrates to one of the two sugar peaks. There, in the midst of plenty, the pioneer agents interact sexually, producing children, who in turn beget children, and so on. All the while processes of cultural evolution are operating within each group producing culturally distinct "tribes" of agents on the two mountains. Ultimately, as population pressures mount from overexploitation of the sugar resources, each tribe spreads down into the central sugar lowlands between the two mountains. When the two tribes ultimately collide, processes of assimilation occur and feed back on the reproductive and cultural activities of the tribes, yielding complex social evolutions"

And state: Our goal, as always, is to grow this history "from the bottom up." Can the entire social history – along with all sorts of variants – be made to emerge from the interaction of agents operating under simple local rules?

Sugarscape – evolution and inheritance

- Simple rules for reproduction and death creating variable sized populations
- Genetic traits for “vision” and “metabolism”
- Various experiments on the evolution of different “population regimes”
- Explore how “inheritance” of sugar wealth changes the society
- Reduces the effectiveness of evolution and increases Gini equality measure (as would be expected)

Sugarscape – cultural processes

- Agents store “cultural tags” which is a fixed length bit string
- Each bit represents the presence or absence of some culturally learnable trait
- Agents interact locally by copying randomly chosen bits between neighbours
- Cultural groups: Agents are defined as “blue” when the string has more 0’s than 1’s or “red” otherwise
- They explore various “cultural dynamics” of how groups and tags form and spread
- They note the overtime all agents tend to converge to the same tag string – hence become part of the same cultural group

Sugarscape - combat

- They consider results of a combat rule based on the idea that one agent may kill another and take it's resources:
  - If it is a member of another tribe and,
  - If the attacking agent has more wealth (sugar) than the victim and,
  - If there are no other wealthy agents from the other tribe near by (that could retaliate)

Sugarscape - combat

- They find that if cultural transmission is turned off (with equal red and blue fixed tribe membership) then one tribe often dominates
- Although this is not the case when agents only get rewarded by a fixed amount for winning a battle (rather than full wealth of killed agent)
- Neither is it the case if cultural transmission is turned on – due to a wealthy victor becoming assimilated into the other tribe

Chapter 4 – trade

- add another resource (called spice)
- Agents need sugar *and* spice to survive
- can trade them directly between neighbours (barter pairs) – there is no central auction or clearing system
- Agents use a rather complex trade rule based on neoclassical microeconomic theory (to be discussed in a future lecture)
- Suffice to say, agents attempt to:
  - maximise their individual welfare through calculating how valuable sugar is to spice
  - Then bargain a price which is only accepted if this results in a Pareto improvement

Sugarscape - trade

- They show that in such a decentralised market the predictions of neoclassical theory do not entirely hold
- Essentially it does not converge to an efficient equilibrium promoting optimal agent welfare
- However trade:
  - Increases carrying capacity (i.e. supports larger population of agents)
  - Increases inequality (higher Gini)

Readings and Questions

- Readings
  - Gilbert et al (2005) Chapters 8 and 9.
  - Epstein, Joshua M.; Axtell, Robert (1996). Growing artificial societies: social science from the bottom up. Brookings Institution Press
  - Hales et al (2013) Agency in complex information systems – Future research directions. Unpublished draft part of an EU consultation report (see Annex 1).
- Questions
  - We have focused on "exploratory" ABM (or artificial societies). Can you find any examples of an "applied ABM"?
  - Can you think of a way ideas from the Sugarscape could be useful in engineering or computer science?
  - "I don't need models, the world IS my model!" Do you agree with this? If not why not?
  - Can you find any examples of other ABM platforms / languages?