When can you use AI with ABM?

Bruce Edmonds

Centre for Policy Modelling, UK.

UiT Norway University of the Arctic, Tromsø.

Umeå University, Sweden.

Purpose of paper/talk

1. To look at the different purposes for which AI entities/algorithms have been developed

2. To briefly summarize their current success at each of these

3. To critique the assumption that because an AI is good for one purpose it is good for another

4. Then look at a brainstormed list of possible uses of AI with ABM

5. Then see what is necessary in terms of quality of achievement at each AI purpose for these uses of AI with ABM

Links to paper and these slides at the end

Preliminary remarks

- There are lots of *different* ML algorithms, not just the transformers that are good at language

- I will not have time to go into the details of the different approaches here, please…

  1. read the paper (link at end of slides)

  2. come and discuss with me

- I have been talking about AI for a looooong time, writing papers about AI and social interaction, e.g.

  Edmonds, B. (2000). The Constructability of Artificial Intelligence (as defined by the Turing Test). *Journal of Logic Language and Information*, **9**:419-424. (http://cfpm.org/cpmrep53.html)

Limitations of talk

- I will be primarily talking about LLMs as they are the most used/applied AI technique but will also mention other Machine Learning approaches

- The field of AI/ML is changing very fast so some of the assessments and summaries in this talk may date quickly

- Although I brainstormed possible uses of AI with ABM, I might have missed some

- Some of the envisaged uses of AI with ABM are theoretical – I have not come across any actual instances of them yet

ABM and AI have a lot in common…

…they both lie between our simplistic understanding and the complexity of the world, *not completely understable* and…

…they are both *plagued* by their surface plausibility

…where, *at best*, some calibration is done but *only rarely* any independent validation
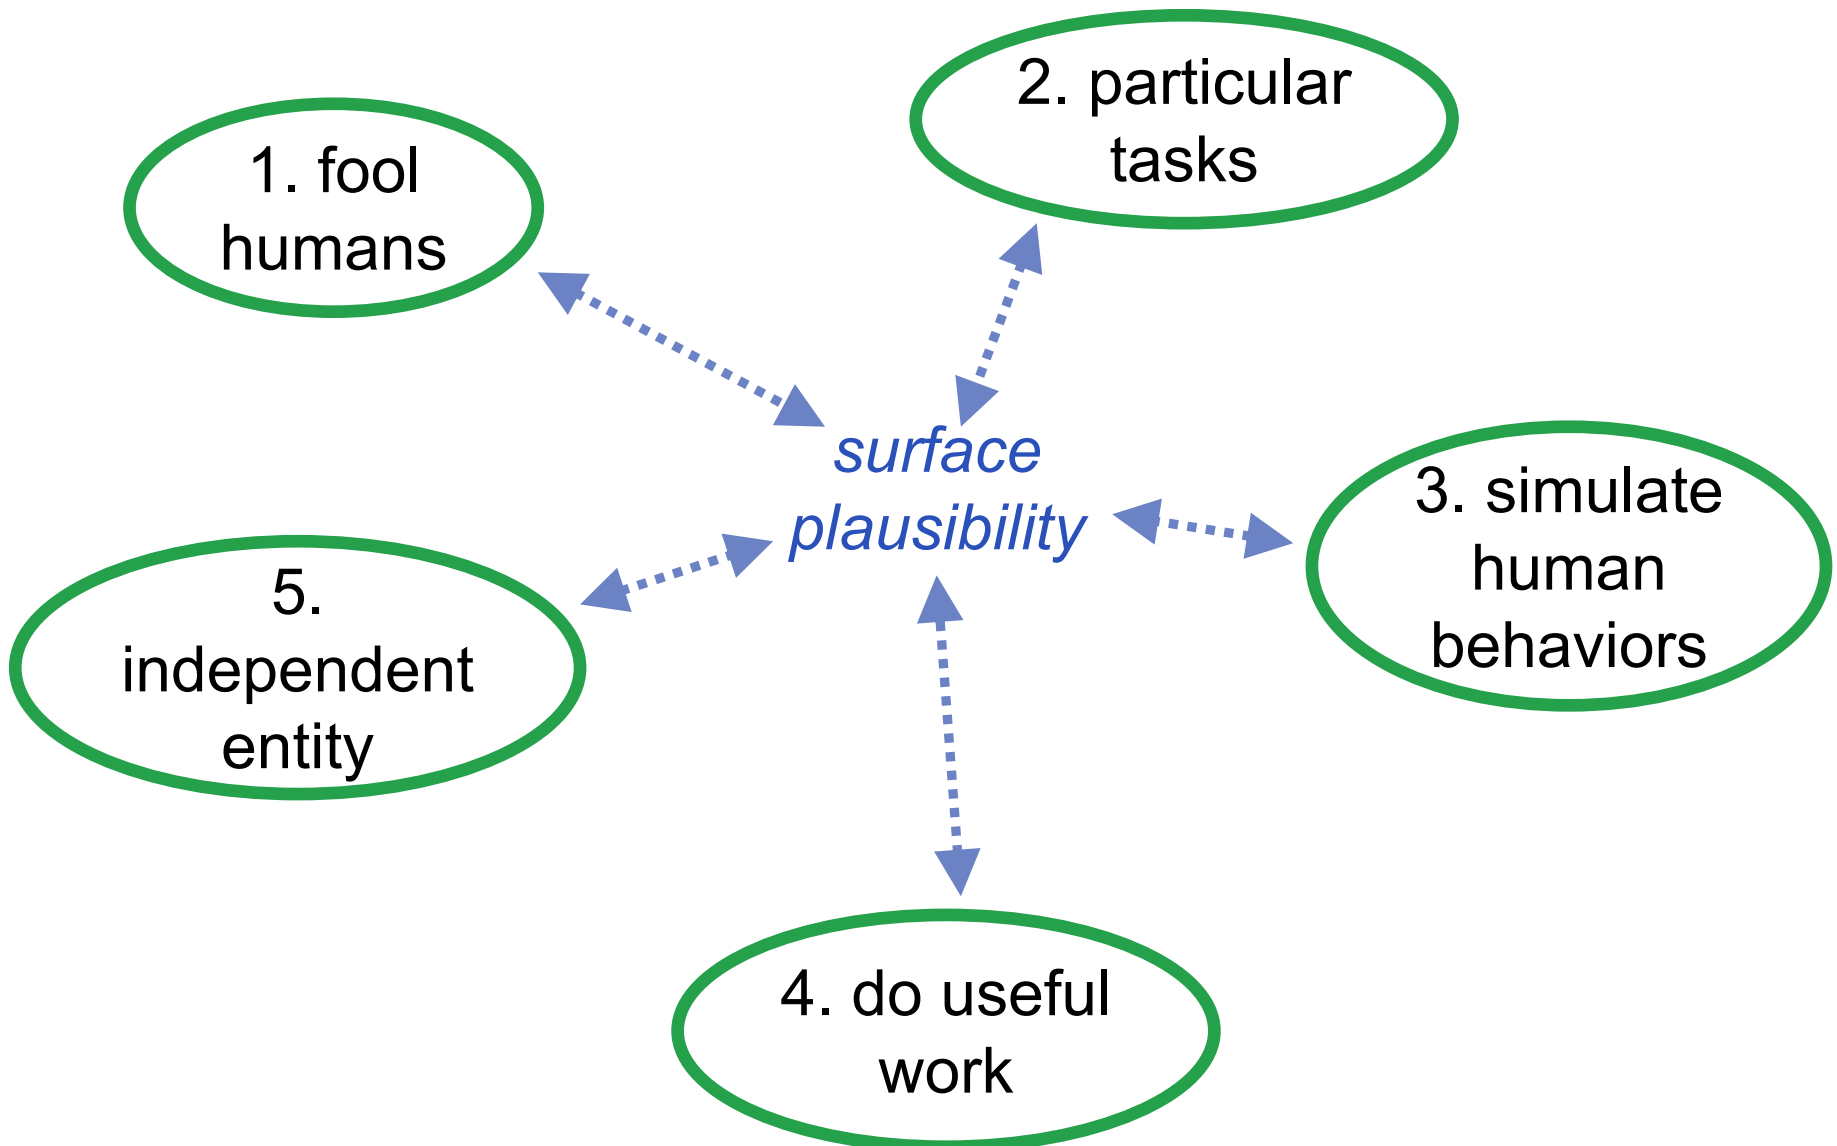
Different **purposes for AI**

1. To fool humans into thinking AIs are human

2. To do particular tasks as well as humans

3. To simulate human behaviour

4. To do useful work for humans

5. To make decisions as a successfully independent entity

Summary of progress (*July 2025*)

- AI entities can fool humans into thinking they are human but, so far, only within restricted contexts.

- AI entities are already much better than humans at a variety of specific tasks. It is expected that the range of tasks they excel at will continue to increase.

- At the moment AI entities are not very good at simulating humans. This could be an area that develops but only if they are developed and rigorously validated for this task.

- AI entities already do useful work for humans, but this is limited by their lack of embedding within human society and their "clunkiness" in AI-human cooperation.

- Making decisions for a successfully independent entity is a very hard and environment-dependent task, fortunately.

Just because an AI is good at one does not mean it is good for another!

1. fool humans

2. particular tasks

3. simulate human behaviors

4. do useful work

5. independent entity

surface plausibility

Uses of AI with ABM (part 1)

a) Prediction of unknown data

b) Explanation of known data

c) To assess the ability of LLMs to work with humans

d) To train AIs for teamwork with humans

e) To play parts in a game designed to inform an ABM

Uses of AI with ABM (part 2)

f) To illustrate how a group of interacting humans might behave

g) To program ABMs

h) To analyse complex formal data

i) To analyse qualitative data

What is necessary in terms of establishing that AI purposes succeed for uses with ABM?

	Fool humans	Particular tasks	Simulate human behaviour	Useful work	Make decisions for self
Prediction		✓	✓		?
Explanation			✓		?
Test working with humans		✓	✓	✓	✓
Train AI's for teamwork			✓	✓	✓
Play parts in experiment	✓		✓		✓
Illustrate human behaviour	✓		✓		✓
Help code ABMs		✓		✓	
Analyse formal data		✓		✓	
Analyse qualitative data		✓	✓	✓	

Conclusions (part 1)

- AI is developing quickly, so some of these assessments might date, but the general arguments should still hold.

- So far, AI systems have managed to:
  - fool humans but only in restricted circumstances,
  - do a range of tasks as well as or better than humans,
  - do some useful work for humans

- So far, they have not been very good at:
  - simulating human behaviour
  - or making decisions as an independent entity.

Conclusions (part 2)

- AI systems differ depending on the task they are developed for.

- If one wants to use AI for some purpose then one needs to match the kind of AI system to your purpose.

- Currently, LLMs do not mimic human behaviour beyond a level of vague plausibility.

- If one is going to rely on LLMs to mimic human behaviour, then this needs to be checked. Narrowing down behaviour using "prompt engineering" is insufficient.

- There are uses of LLMs to aid in simulation tasks, but they are currently not that of replacing agents that represent humans with LLMs in any simple manner.

The End!

HAPPY - O - METER

GREETINGS HUMAN!

My Email: bruce@edmonds.name

My publications: http://bruce.edmonds.name/pubs.html

These slides are at: http://cfpm.org/slides as:
AI and ABM talk - SSC2025.pdf

Centre for Policy Modelling: http://cfpm.org

cfpm.org