

How many steps for a lightbulb

For regulatory purposes, artificial is, hopefully, the easy bit. It can simply mean “not occurring in nature or not occurring in the same form in nature”. Here, the alternative given after the *or* allows for the possible future use of modified biological materials.

Fortunately for would-be regulators, though, the philosophical arguments might be sidestepped, at least for a while. Let’s take a step back.

Defining the terms: artificial and intelligence

From a philosophical perspective, *intelligence* is a vast minefield, especially if treated as including one or more of consciousness, thought, free will and mind. Although traceable back to at least Aristotle’s time, profound arguments on these Big Four concepts still swirl around us.



State of AI in 2015

In 2014, seeking to move matters forward, Dmitry Volkov, a Russian technology billionaire, convened a summit on board a yacht of leading philosophers, including Daniel Dennett, *Paul Churchland* and David Chalmers.

Fortunately for would-be regulators, though, the philosophical arguments might be sidestepped, at least for a while. Let’s take a step back and ask what a regulator’s immediate interest is here?

Logically, then, it is the way that the majority of AI scientists and engineers treat “intelligence” that is of most immediate concern.

In 2014, seeking to move matters forward, Dmitry Volkov, a Russian technology billionaire, convened a summit on board a yacht of leading philosophers, including Daniel Dennett, Paul Churchland and David Chalmers.

Intelligence and the AI community

Until the mid 2000s, there was a tendency in the AI community to contrast artificial intelligence with human intelligence, an action that merely passed the buck to psychologists.

In November 2007 an AI pioneer at Stanford University, addressed this issue:

The problem is that we cannot yet characterize in general what kinds of computational procedures we want to call intelligent.

John McCarthy

This informal definition signposts things that a regulator could manage, establishing and applying objective measures of ability (as defined) of an entity in one or more environments (as defined). The core focus on achievement of goals also elegantly covers other AI.

Another constraint is that AIXI lacks a “self-model” (but a recently proposed variant called

“reflective AIXI” may change that).



Smart boats and yatches are already here. What does it take to acquire one?

First, the informal definition may not be directly usable for regulatory purposes because of AIXI’s own underlying constraints. One constraint, often emphasised by Hutter, is that AIXI can only be “approximated” in a computer because of time and space limitations.


Second, for testing and certification purposes, regulators have to be able to treat intelligence as something divisible into many sub-abilities (such as movement, communication, etc.). But this may cut across any definition based on general intelligence.

Intelligence measures an agent’s ability to achieve goals in a wide range of environments.

From a consumer perspective, this is ultimately all a question of drawing the line between a system defined as displaying actual AI, as opposed to being just another programmable box.

If we can jump all the hurdles, there will be no time for quiet satisfaction. Even without the Big Four, increasingly capable and ubiquitous AI systems will have a huge effect on society over the coming decades, not least for the future of employment.

This informal definition signposts things that a regulator could manage, establishing and applying objective measures of ability (as defined) of an entity in one or more environments (as defined). The core focus on achievement of goals also elegantly covers other intelligence-related concepts such as learning, planning and problem solving.



How would you identify an Android?

The future of AI

First, the informal definition may not be directly usable for regulatory purposes because of AIXI's own underlying constraints.

- AIXI can only be “approximated” in a computer because of limitations.
- Another constraint is that AIXI lacks a “self-model”.
- Regulators have to be able to treat intelligence as something divisible.
- May cut across any definition based on general intelligence.
- Other intelligence-related concepts such as problem solving.

Second, for testing and certification purposes, regulators have to be able to treat intelligence as something divisible into many sub-abilities (such as movement, communication, etc.). But this may cut across any definition based on general intelligence.

But if the Big Four do ever (seem to) show up in AI systems, we can safely say that we'll need not just a yacht of philosophers, but an entire regatta.