

Sentience in machines

[Manuel de la Herrán Gascón](#)

First version: Dec. 2016.

Updated: Jan. 2017

I have created this text from the materials I prepared for the talk I gave at the [Faculty of Philosophy](#) of the [University of Santiago de Compostela](#) on December 15, 2016 along with [Brian Tomasik](#), which was entitled "[Outlook and future Risks of artificial consciousness](#)".




Fig.1 The thermostat in my house

This is my home thermostat. It is the "intelligent" part that regulates the temperature by controlling a boiler. If it's cold, it turns on the boiler. If it is too hot, it turns it off. It seems that aims to keep the temperature constant at a certain value, as if he did not like excessive cold or excessive heat, as it happens to me.

Are we absolutely sure that this thermostat is not sentient¹?

Perhaps the question is poorly worded?

¹ David Chalmers suggests that even a thermostat might have experiences.

Perhaps the question about the sentience of the thermostat is a poorly worded question, such that it is better not answer. If ever you are accused of a murder you did not committed and interrogated with " Yes / No" questions like this: "*Is it true that you killed him for money?*" It would be better not to answer. The question takes the blame. If you say yes, it would be admitting the crime, and recognizing that the motive was economic. If you says no, it seems that you are taking the implied accusation and simply the motive was another.

Perhaps the question "What generates consciousness? Under what conditions does it emerge?" Is poorly formulated and consciousness is not generated, it does not emerge.

Perhaps the level of description of reality in which we speak of the idea of a concrete thermostat (the level of physical objects identifiable by us) is not the level of reality adequate to ask questions in relation to consciousness, being possible to describe Reality to other levels, such as talking about a diffuse set of metal and plastic atoms, the level of quantum phenomena, or the universe as a whole. It may be on another of these levels that the question about consciousness makes sense.

This text is about sentience (feel, suffer, enjoy), consciousness (perceive, be aware, realize), have interests (things that benefit me, things that hurt me), subjectivity (to have an "I", to have an opinion, be somebody, to have a point of view), have desires or preferences (want, desire, longing, love, things that matter to me).

Perhaps each of us has a different definition for words like "conscience" or "sentience" and even a different theory about their nature. But sometimes it seems that two people handle concepts faced when they are actually talking about the same thing; and in other cases they use the same words and theories but when you dig deeper into it you can see that they are referring to different things.

There is something basic: "*I feel, therefore I am*"². I feel, therefore I have consciousness. I feel, therefore I am sentient. I feel, therefore I am. But who are the other beings with consciousness / sentience? Is it necessarily to be human, like me? Animals? Do you need to be alive? Is it necessary to have been produced by evolution? Does such development must be a natural evolution? Biological? Based on carbon? Is a requirement to be wet?

Some definitions, theories and hypotheses about sentience / consciousness are evident for some but not for others, and vice versa. The aim of this exposition is to show some of these theories and ideas and argue that it is not justified to discard them completely, even if they appear crazy or anti-intuitive. And sow the doubt -if there wasn't already- that thermostats may be sentient beings.

As for consciousness and sentience we must recognize that there are many things we do not know and this is why it's a very big risk to dismiss the theories that are not intuitive for us. Particullary, is a very big risk in relation to the possible sentience of machines.

For some it is absolutely clear that machines can not be sentient, and for others it is evident that machines can be sentient. My goal is that both groups take into account the contrary hypothesis, and deepen the aspects that can determine sentience or

² Inspired by the *cogito ergo sum* "I think, therefore I am" of Descartes (XVII century)

machines. And do this orientated to prevent some moral hazards that can be of enormous magnitude.

The world is being transformed by and for humans. A general purpose machine would have some advantages if it had human form. For example, it could drive a car. It could evoke tenderness in us. It could also sexually entice us, even fall in love us.

We can expect that at least some machines or robots are going to have a look very similar to human, perhaps indistinguishable. As machines become more like us, both in appearance and behavior, we'll find it easier to empathize with them. How will we react? Perhaps creating an emotional relationship with the robot, like in the movies [IA](#), [Ex-machina](#) or [Her](#). Perhaps aversively, as in [I, Robot](#).

Detective Del Spooner: *Humans have dreams. Even dogs have dreams, but not you. You are just a machine. An imitation of life. Can a robot write a symphony? Can a robot turn ... a canvas into a beautiful masterpiece?*

Sonny (robot): *Can you?*

When we heard a statement in relation to what sentience is, we must ask the other party to clarify whether it is a definition, or conversely, if it is an explanation / description of the conditions required for sentience. If it is a definition, there is nothing to discuss. People call things as they want. If it comes to an explanation or the conditions, then we can discuss about it.

For example, someone might say:

- *"A conscious being is one who operates a symbolic system where a symbol is himself"* (self-reference).
- *"A conscious being is one who manages information from their environment, as a map"*.
- *"Awareness is the popularity of (facing) ideas or agents in the market of mind³".*
- *"Consciousness is a character in the theater of the mind. Being aware is to play a role: to appear (or to believe) be aware"* (eliminativism).

If it comes to definitions, there is no objection. For example someone could define "flying" in this way:

- *"A flying being is one who has feathers"*.

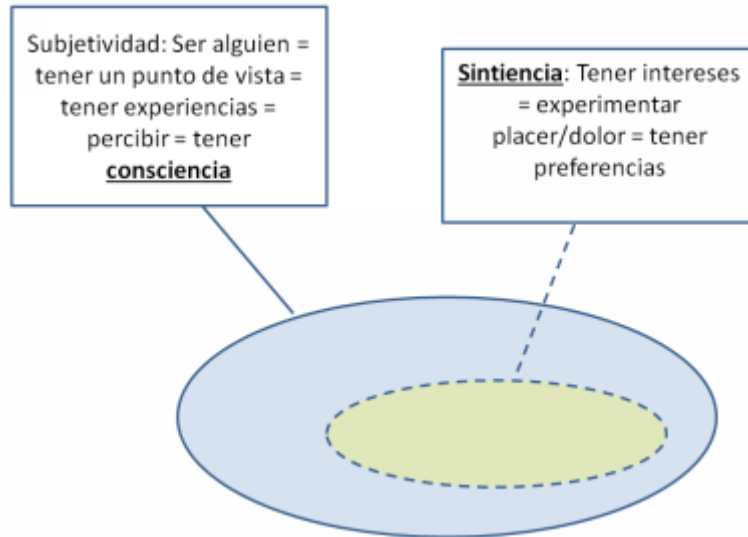
If it is a definition, ostrich would be a *flying being* (by this definition). The ostrich can not fly, but it would be *flying*.

[My definitions of sentience and consciousness](#) are performed by synonyms. I try to evocate it's own idea in the reader's mind. But the words I use could be exchanged or I could use other different, no problem about it.

- I consider **Sentience** the ability to have pleasurable or painful sensations, which involves having preferences and interests (avoid pain, seek pleasure).

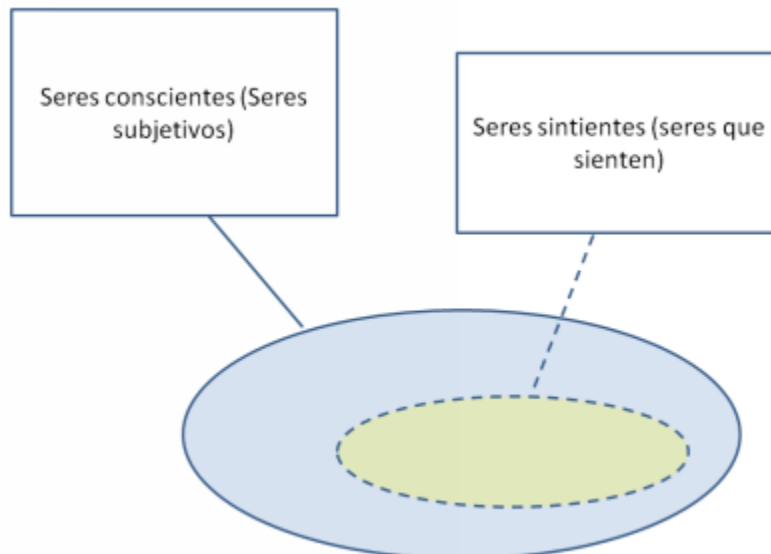
³ Marvin Minsky among others proposed that sentience occurs in the interaction of different mental agents in a "market" in which competing to get solutions to different needs that can be contradictory, such for instance to be sleepy and hungry at the same time. Sentience is related to or identified with the existence of a neural representation of environment that is subjective and self-centered.

- I consider **Subjectivity** the ability to experience. In "experience" I include the ability to feel pleasure and pain, but also I include having a point of view, be someone, perceive, to have **consciousness**.




With these definitions I assume that a being who feels, is necessarily conscious (all sentient beings are both conscious).

Additionally I assume it is at least theoretically possible to be conscious without experiencing pleasure or pain you may be aware and not having to have preferences.



Another way of looking at it is to imagine a hedonic scale of pleasure and pain, where for whatever reason, there are beings who are always kept together with a zero value, no way out of it.




Beings who feel pleasure and pain have moral relevance (for themselves) because they have interests. Instead it's beings who have no preferences, or can have them, have no moral relevance (for themselves).

If the machines have a view, but could not experience pleasure or pain, or have interests, they would have moral relevance. But if the machines could experience suffering themselves they have moral relevance.


More precisely, we should not assume that all sentient beings can feel pain and pleasure vice versa. It seems at least theoretically possible that there are beings who can only feel pain, or can only feel pleasure. I will define three capacities and combinations to explain these ideas.

- **Hedon**: the one being who can experience positive things (feel pleasure)
- **Sufferon**: the one being who can experience negative things (pain)
- **Perceptron**: the one being who can perceive, have a point of view, realize.

All hedones and all are necessarily sufrones perceptrons. Additionally, we are used to all hedones are also sufrones and vice versa. But that might not always be so. We could imagine beings, even worlds, where there is only pleasure (there are only positive experiences), there is only pain (there are only negative experiences) or even a world where other subjectively relevant experiences, we could not qualify exist or positive or negative of (other axes or additional to the hedonic scale dimensions).

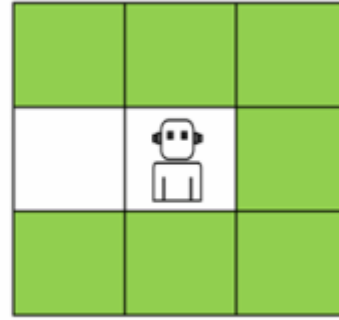
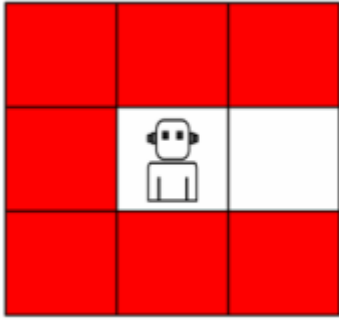


Pleasure and pain appear to be useful, and be in different ways, justifying the existence of both experiences. Imagine that I am in the middle of a board, and I green box is good for my genes. My genes are scheduled me for survival and reproduction. I am a biological machine created by the genes with the aim of perpetuating these genes. Suppose the green box out meal or a sexual partner, something positive for my genes. Genetic programming can take me to the box through pleasure.





But if something threatens my survival, like fire or a predator, the pain seems useful to escape or avoid certain situation.

In short, pleasure is practical to motivate to "go" and pain is practical to motivate to "escape". To "go" through pain or to "escape" through pleasure seems more difficult to achieve.



However, we all love pleasure and hate pain⁴. It would be difficult, but not impossible, that all behaviors were guided by pleasure, and turn pain into something unnecessary, and nonexistent. To do this, we need is a cognitive system that works well for "Go" and "Escape":



That is precisely what he proposes [David Pearce](#) in his abolitionist project " [The hedonistic Imperative](#)": a guided gradients club, where pain is unnecessary and nonexistent behavior.

But r ealmente feel is necessary to behave as if he felt?Is it really necessary to feel fear to provoke the reaction of run?Does nature Glad we could not programmed with the most appropriate behaviors to maximize survival and reproduction of genes, without having to experience absolutely nothing? Could a robot be programmed to behave like an animal, without feeling as animals feel?

[David Chalmers](#) proposes that certain "zombies" are [metaphysically possible](#) (consistent with our knowledge of physics): beings who are physically or resemble conscious people, but they are not aware.These zombies screaming "Ouch!"stepping on a nail, but they feel nothing at all.

There are two possible reactions to the possible existence of such zombies.

- **They are impossible**: Right behavior is inevitably linked sentience / consciousness


⁴ One might object saying that there are certain pains that produce satisfaction, at least to some people.These "pains like" would be included in what I call "pleasure".

- **They are possible**: that behavior may occur without any sentience / consciousness. This case has been used to argue in favor of dualism.

Whether you answer yes as if we say no, in both cases, -the artificial - robots could be sentient machines. In the first case, the machines could be sentient, simply, if they gained such behavior. In the second case the machines could be sentient if included in its composition that makes sentient animals.

What makes sentient animals? They have been mentioned several factors that may be relevant to sentience:

- Natural (not created by humans).
- Living things, carbon-based, the result of a natural evolution.
- wet beings.
- "Accomplishments" in the physical world, not simulations.
- Not digital analog Beings
- Nanotube quantum effects ([Roger Penrose](#))
- Interaction with the Multiverse




Source: mrcorfe.com

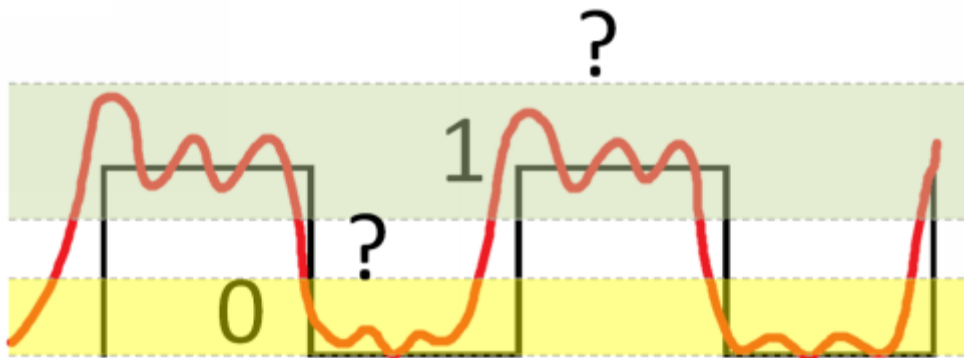
We have been taught that the first is an analog signal and the second and third are digital signals.



However, digital signals like this do not exist physically. What exists is something like this:




And a system that reacts to said signal interpreting, for example, as follows:



That is, when facing a physical quantity signal or there is a system that interprets (reacts) by discrete states, we talk about digital signals. But the signal is still analog.

Our universe is analog (simplifying: it has an "infinite" precision). Instead discretized digital computers (limited) time and information to a series of defaults.



In short, everything is analog, and within the things that are analog (all), some are digital. This occurs until the quantum level, where space (position), time (time) and matter / energy / information could be discrete / digital phenomena.

The huge success of digital computers is because they approach a general purpose machine. Pero the digital being carries the cost of losing some very useful features of analog computers. In a way, it seems as if the development of computers, being digital, it had gotten into a dead end. The use of physical quantities without discretize has drawbacks, but also some great advantages. The "[analog computers](#)" are (less flexible) specific purpose, often suffer breakdowns and wear processes which cause imprecision. However, they can be extremely effective (eg, fast) for obtaining certain tasks radically different from digital computers yields.

The [sorting algorithms](#) used in digital computers have an exponential or equivalent to exponential computational complexity. This means that the time required to do the sorting is not increased more or less constant with increasing the number of items you need to order, but the situation is getting worse, and is getting worse. The solution to this has been to use brute force, taking advantage of [Moore's Law](#), which has been performing since 1965 according to which the performance of computers doubles every year, representing an increase of exponential capacity.

There is however an *analog* algorithm whose computational complexity is exponential: the time required to perform the ordination is increased by increasing the number of items you need to order, but these increases are constant, directly proportional to the number of elements. How this wonderful algorithm works?



The operation of this sort analog machine is as follows. For each item required order, a stick whose length is proportional to the number you want to order is cut and that number is written to the stick. When all sticks have been cut, are joined in a bundle, placed vertically, they are firmly secured, and gives them a blow on a flat surface, for example, in the table. After hand on the beam is lowered. The first stick that touches our hand will be the highest number. We extract it and write down the number. Then we go down the hand, playing and drawing the second number, and so on.

Digital computers have eclipsed analog, but perhaps the extraordinary advantages of analog computers, as his "infinite" precision or its ability to efficiently solve problems such as ordination could be a requirement for sentience, because machines for which we have overwhelming proof of sentience (animals in general) are analog machines.

1
11
21
1211
111221
312211
13112221
1113213211
31131211131221
13211311123113112211
11131221133112132113212221
3113112221232112111312211312113211
132113213211121312211231131122211311221131221

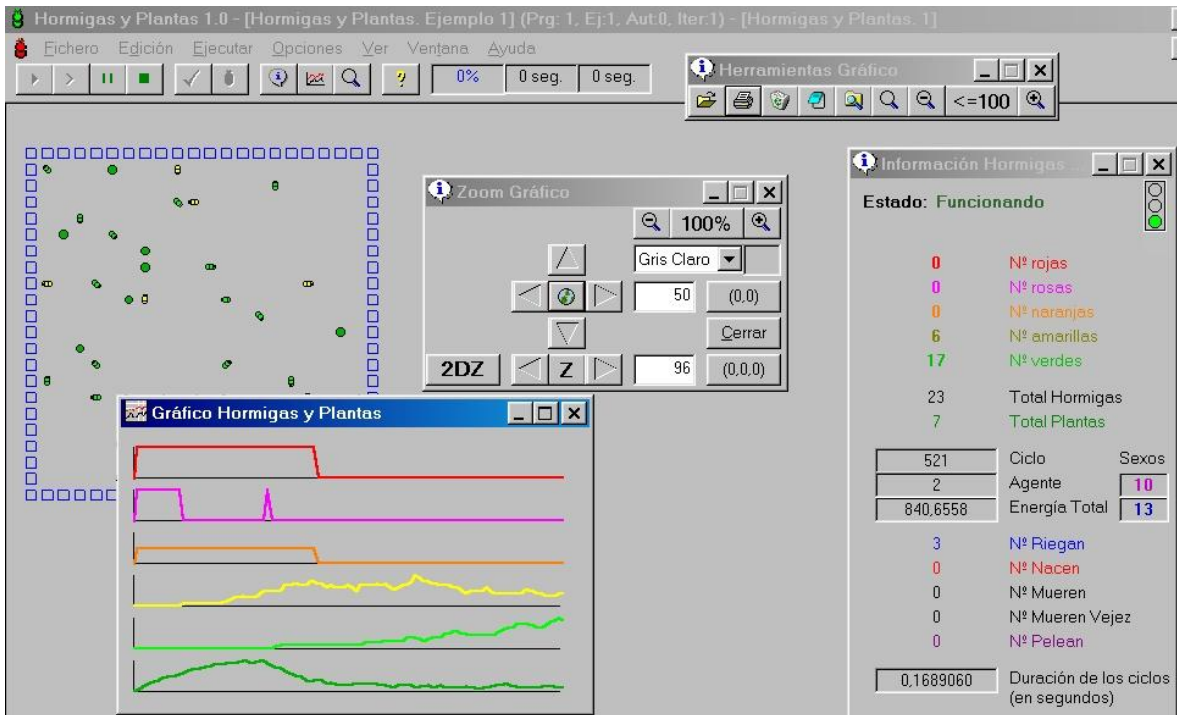
What is the next value of [this series](#)? You can view [this video](#) or download [this program](#) to find the solution.

Can a brain be sufficiently complex to understand himself? It seems difficult.

There are phenomena that seem complex and are not. To explain what we do not understand (eg sentience / consciousness) we appeal to other things that we do not understand well (as complex systems or intelligent systems).

The sophistication of the complexity dissolves when understand its origin, as in the case of the previous series. It is unfair to say that the human brain, as we understand it, has enough complexity to create sentience, and other simpler brains, as if we understand them, are not sentient. If a super-intelligent alien being could perfectly understand our brain and anticipate their behavior should conclude that we are not sentient?

The next [chat robot](#) was created with a few very simple rules, but was able to "cheat" to several people who came to believe they were talking to me. Consider that a certain "someone" who is on the other side of the chat can be a sentient being is a reasonable speculation. But the fact understand how they work their mental processes is not an argument for rejecting his sentience.



In [this program](#) that simulates the evolution of ants altruistic, cooperative and selfish behavior, sexual reproduction (recombination) and mutations can observe certain "evolution of species" adapting to their environment. What is the argument for thinking that a natural evolution can produce sentient beings, and instead can not do an artificial evolution in computer simulations?

Is heard saying that subjectivity / consciousness implies the existence of a self, an identity, and that human beings have freedom, a will, a free will. There is confusion in this regard has bitual: Free is not the same as unpredictable.

- **To be unpredictable does not mean being free:** if we had a random physical system, we would say that is unpredictable but we would not say that it is free.
- **To be free does not mean being unpredictable:** I can freely decide always take the same decision, and others can correctly predict my future action.

At the "macro" level we can consider that "everything has a cause" (genes, environment), but in this subject it does not really matter whether our behavior is determined or not, whether it is predictable or not, whether it is random or not. Freedom consists in having the feeling of being free. It is about having the experience that it is me who makes the decisions. What we call *freedom* is really *identity*: "I am free when *I* decide it."

Usually we associate the word freedom with an immense physical space, without limits or obstacles, and the lack of freedom with a prison. But none of us would like to be taken out of prison to be left abandoned in the African savannah at the mercy of predators, or on the surface of the moon without oxygen. Freedom is subordinate to the concept of interest. It is no longer free who is offered more alternatives, but the one who can choose what he most desires. Being free is primarily to be able to satisfy one's own desires and interests.

The desire for freedom is one of the strongest desires that exist, and that is why we must respect it: in general, individuals want to be free, they want to have the feeling that they make their own decisions, and in general they would be frustrated if they are not allowed to do so.

There are those who ask: *"If we are determined by genes + environment what do we do with all the prisoners? Must we open all prisons?"*. The answer is simple: what was the purpose of that jail? According to various theories, the purpose may be the sanction (compensation reaction), prevention, amendment, rehabilitation, etc. Well, this purpose of punishment it is not affected by the fact of the absence of freedom in the usual sense of the word, and is equally applicable to humans and robots.

Now I will discuss a number of important to establish sentience or not the machines "curious ideas" related to the concepts of "identity" and "reality" and.

All my memories could be false. All my past until exactly half a second could have been implanted in my mind, in my memory, as in the film [Total Recall / Total Recall](#)

There could be an anesthesia do not eliminated the heat but the memory of pain. This anesthesia is mentioned in one of the novels of fiction [Robert Anson Heinlein](#), and in the essay ["The worst of evils"](#) of [Thomas Dormandy](#).

Perhaps [we live in a simulation](#), as stated [Nick Bostrom](#).



Our world, like the [world of ants](#) we've seen before, could stop for years (button "pause"), and then resume ("play" button); and we would proportionally increase or

decrease the size of objects in the world via the button "zoom", keeping the proportions, without anyone noticing anything unusual.



The [invincible boat](#) won all naval battles. As expected, it was much appreciated and every time some piece, they are not destroyed, but were stored in a warehouse was replaced. Unfortunately, the enemies came one day in the store and stole all the pieces, with such bad luck that had replaced many parts of the invincible ship, there were enough pieces to create a new invincible ship. So they did, and this new boat seemed invincible even more authentic. When faced at sea two boats invincible who won?

As samples countless stories of [science fiction about teleportation](#): if copied, one by one, all my neurons and other cells in my body, forming a new body which one would I be?

If replaced, one by one, all my natural neurons, by artificial devices externally behaved like neurons would I lose my ability to feel?


If you do not replace, but expand my ability with artificial neural devices could do so to overcome such biological could even disconnect the biological part without notice much difference?

If you could directly connect my human brain with that of another animal, uniting our neurons would we be an individual or two?


If I could separate my human brain neurons, creating several groups of widely separated neurons among them, as already happens in some way would we be an individual or several?

There are many [alternatives in relation to sentience](#) that can be summarized in the following, creating a kind of map or axes:

- A: Who feels? I, my kind, animals with a central nervous system, matter
- B: What do you feel? Pleasure, pain ...
- C: How much is able to feel? Discrete, continuous, maximum ...
- D: When do you feel? Past present Future
- E: Where sentient experience occurs?
- F: Why sentient experience occur? How can the sentient experience? What is its nature?
- G: Is sentience is created or pre-existing?
- H: What sentient experience occurs? Do you have any purpose?
- I: How much experience sentience there? What is your availability?



The hypothesis of separate identity or "closed individuality" might be an illusion, and instead be a single subjectivity as I described in my book "[Arena Sensible](#)" (2005) or "Individualism Open" of [Daniel Kolak](#) in "I Am You "(2004).



Every day, every hour, we experience things fall down. Every day, every hour, we obtain evidence that things fall down. And we know that is not true. We know that the concept is "down" ill - posed. The question "things fall down?" It is one of those Poor questions it is better not to respond in a "Yes / No".



The reality can be like a box covered by a canvas. The "I" is a [hole in the fabric](#), and the "you" is another hole.



At least at the *descriptive* level we can talk about matter, ideas and experiences. They are three types of things, or three types of realities.



Curiously, they coincide with the three types of work that exist: manipulate matter (process material things), manipulate ideas (consulting, services), and manipulate experiences (commercial, political, educational, manipulating people).


There materialist / reductionist / monistic who say, "I only believe in the matter. I can only be sure of what I can touch: things, the material, the tangible ". But the "tangible" refers to the *experience of playing*. He's talking about *experiences*, not matter. These people who say that everything is material, they are actually saying that everything is *experience*.



Are you at home. On the table is an apple. You point the finger first apple, and then you point your finger your own head and say, "How the brain manages to represent this block in my mind?". This approach is totally wrong. That thing that I point out is no real apple: e s representation that my mind has apple. Now I point to my head. That's not my head. That is the mental representation that my mind has my head. Now I leave the house. I see the field, trees, moon and stars. I open my arms, look around and say, "all this, is my mind". [That's right](#). If there is any doubt, people who suffer an amputation usually complain of pain in the limb that no longer exists. What we call "my right foot" is a mental representation in the brain. The shoe as well.



Usually considered sentience matter emerges. But it may also happen that sentience is received, invoked or tune, as with radio receivers. The hardware configuration of our bodies, like the hardware configuration of radio receivers, could not generate sentience, but be in tune with it, so that depending on the internal composition is receiving a signal or other sentient.



Instead of emerge ("up"), sentience could go down, it could be tuned, [as a radio receiver](#) listening to a certain frequency. A tooth caries could be a resounding gong or hardware configuration or tune aligns with a signal that is constantly emitting and generates the pain in what appears to be an individual.

The hypothesis of [Platonism sentience](#) considers that the experiences could exist by themselves, regardless of sentient beings who experience them. Even though the likelihood of [Platonism sentiente](#) was extremely small, while there is a probability greater than zero, and given that it is unclear where sentience comes, we might think twice before ignoring this idea as it should be true, its implications as to prevent suffering would be immense.

Assuming simulation Nick Bostrom, certainly not necessary to simulate the entire physical universe, but only the tangible universe (that part that someone will perceive). But, in fact, it is not necessary to simulate absolutely nothing material universe; it is only necessary to simulate subjective experiences. And this way of viewing the simulation argument does seem to be more conducive to Platonism sentient and immersionist monism, since in the same way that we believe we use materials (computers) to simulate field (aircraft and bridges), perhaps in another platonic world someone is using to simulate sentience sentience.


Suppose we have several buckets of water with different temperatures. Imagine that there are some plastic sponges into and out of the buckets.

In this metaphor wet sponges are sentient beings, and water (at different temperatures) are the experiences. Each sponge along with the water therein form a coherent, all so that the temperature of the water within the sponge as experiencing specific experience identifies said sponge.

High temperatures correspond to positive experiences (different satisfactions or pleasures) and low temperatures with negative experiences (different frustrations or pains). In each sponge, the higher the temperature, the more pleasant the feeling will be, and the lower, more painful; having an intermediate temperature at which a state of indifference occurs without significant pain or pleasure.

It is very important to note that in this metaphor represents water experiences, and these experiences (water) exist independently of sponges. For example, a labeled "toothache" blue cube may contain water at 4 degrees centigrade, and one labeled "pleasant tickle" may contain orange water at 27 degrees Celsius. By introducing the sponge in a bucket, sponge acquires this experience (toothache or tickling) and sponges (beings) can identify their own experiences and those of other beings depending on the type of water they contain. But water exists independently of sponges.

Sentience Platonism



Sponges into and out of the buckets. By changing bucket, sponge takes on a different water, in the same way that beings experience different things over time. In a project to reduce or eliminate suffering, and according to this metaphor, there is a risk that could be focused on alleviating suffering carrying sponges progressively from the coldest to the hottest cubes cubes. And assuming that a majority of buckets of cold water exists, and that sponges are able to reproduce, we could devote ourselves to promote avoid the reproduction of sponges, as a way to avoid having new sponges in

buckets of cold water. We could even ask that [the best possible world is one in which there are no sponges](#).

Indeed, in both cases we could see that there are individual and very specific sponges that have improved their welfare, and now found in warmer cubes. Or that there is now no sponge in particularly cold cubes. The problem is that these buckets of cold water still exist, and if the metaphor were true, we would not have solved anything. If this metaphor really would represent the nature of the experience, we would be wasting time carrying buckets sponges of others or preventing sponges reproduce. What we should do is raise the temperature of the cold cubes. In doing so, certain negative experiences would cease to exist, or would be relieved instantly for all, and forever.

Just as the sentience of robots and simulations could generate a moral catastrophe, we should also consider the possibility of Platonism sentient. Although it may seem unlikely, the consequences would be enormous.


When we explain [how to recognize another sentient being](#), we discuss behavioral, developmental and physiological criteria, but we do "a posteriori" in a very egocentric way. Based on the maximum evidence "sorry" it seems to perform an interpolation between individuals who have similar appearance and / or behavior similar to itself ("If it looks like me and behaves like me, feel me"); it also seems that we interpolate between species ("If you have been created like me, feel me") assessing whether it has the same origin (evolutionary) and if it is genetically close to me. Additionally and since sentience has an evolutionary utility, seems to establish -injustificadamente- to have an evolutionary utility is a necessary criterion for sentience exist.

If you look at a group of humans killing a pig and then eat it, we could try to defend the pig arguing that the animal feels "look panicked expression eyes", "listen to their heart-rending cries," "look as he writhes in pain "... but this could be as mean to argue that an animal is flying," notice that bones so light "," how so aerodynamic "and" what majestic wings ". The fact that that description fits that of a flying creature does not mean that all flying beings should have those attributes. The airplanes are aerodynamic but do not have feathers and are particularly light.

This egocentric way to recognize the sentience is very dangerous for those who do not look like us, like [robots](#), the [insects](#), the [simulations](#), the [atoms](#) and who knows what else.

If someone speaks Spanish I can understand what you say and whether it is intelligent. But if I speak in another language I do not know, I can not know if I said genius or a banality. It is radically unjust disregard and despise what they do not understand. There is a widespread notion that only animals are sentient, yet we do not understand well what is sentience. We hide both robots and create simulations of life for increasingly complex computer. This is morally very daring.

Another way to look at it: If I am selfish (analogy of sentience) and want to know who else is selfish in a group of people, I can search other behaviors that are associated in my case selfishness. Or: if I am heterosexual and I want to know who else is heterosexual, I can search other behaviors that are associated in my case to be heterosexual. But this does not mean that no other people, selfish or heterosexual (sentient analogy) that simply hide or manifest in ways that I do not understand, exist.



As if this were not enough, we write and read about topics such as treaties here are some of the luckiest individuals (with better quality of life), members of the luckiest species -the humana-, and at the time of greater prosperity of the history. This can make us forget the asymmetry between pleasure and pain: there is much more pain than pleasure. We have managed to control so much physical pain that psychological pain seems very relevant, when most other animal species, and most humans from other eras, have experienced large doses of physical pain in many of the stages of their lives. The risk of widespread suffering and great suffering, machines and simulations is very high.

In conclusion

- There is no clear idea of what is sentience / consciousness; or what generates / invokes it, or which are the conditions necessary for it to happen.
- From different philosophical perspectives, machines can be sentient. What changes are the conditions necessary to assume that they are.
- Creating sentient machines or simulations can cause a moral catastrophe of an astronomical magnitude.
- We should not dismiss hypothesis simply because they seem crazy or anti-intuitive. His moral implications could be extraordinary.