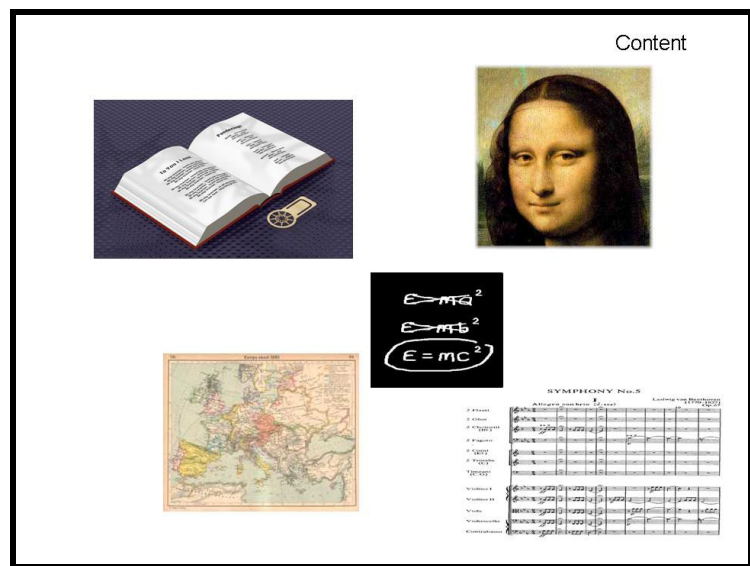


Content, Cognition & Robotics

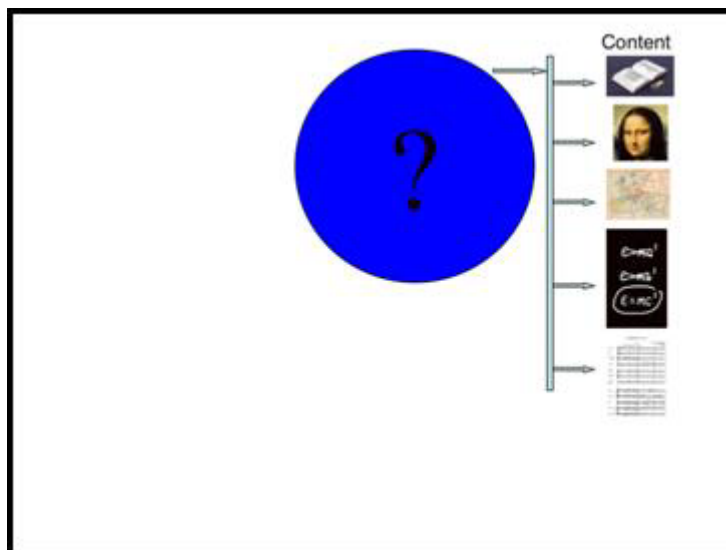
by Hans-Georg Stork

What do Content and Cognition, "Digital Content" and "Cognitive Systems", have to do with each other?

A lot, I claim, providing we agree on a suitable understanding of the concepts at issue. First, "Content". Let this term denote the external and communicable representations we people form of our perceptions and thoughts. Of course there are other, equally possible (or plausible) meanings of that term which we do not consider here. But we do note that the particular ways we humans represent our thoughts and perceptions distinguish us from about any other kind of living entity on this planet. We are content producing animals. Here are a few examples:

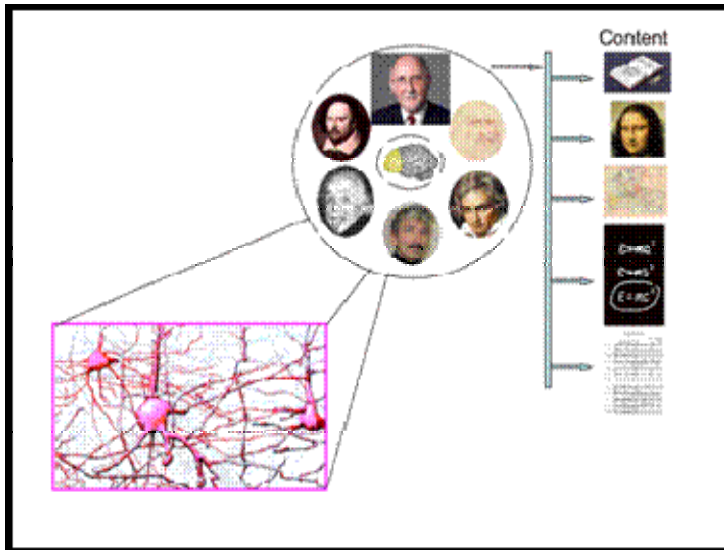


These external representations do not come out of the blue:

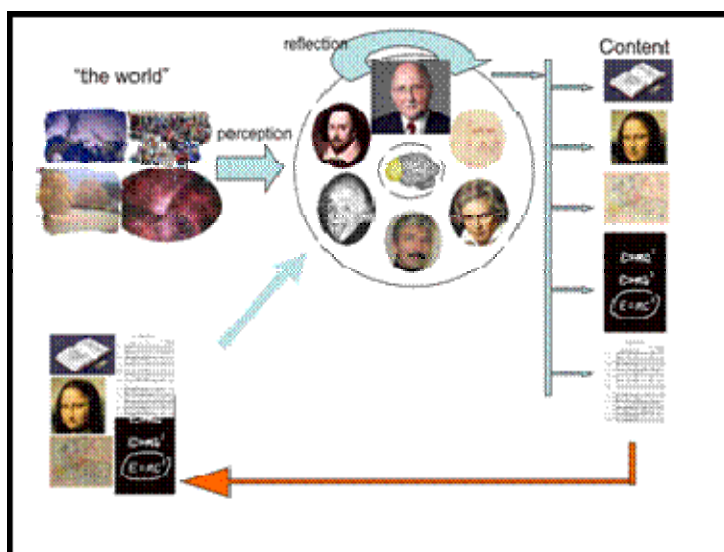
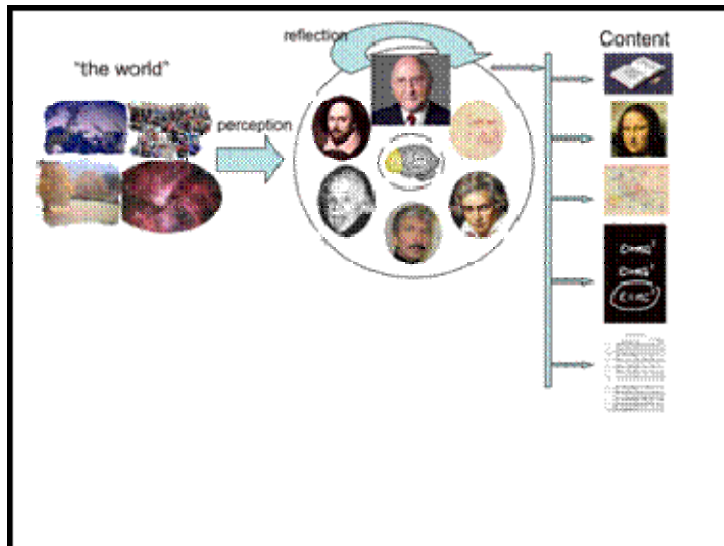


What's behind them are no more and no less than the matter and structure of our bodies, evolved over millions of years and grown and developed in their respective lifetimes. And

within our bodies it is, above all, our central organ of control, consisting of some 100 billion neurons and about 10000 times as many synapses, the human brain:

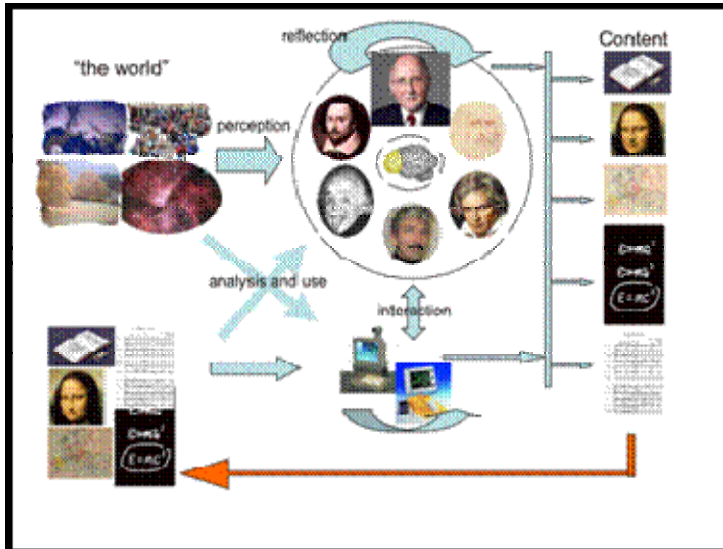


Thus "Content" is the result of our experiences and interactions, reflected and mediated by an embodied brain.




We use the term "Cognition" to denote the totality of physico-mental activities underlying our interactions with the world. The creation of "Content" is, as a matter of fact, one of the key ways we humans shape the world. Indeed, the "Content" we create becomes part and parcel of our worlds as everyone can realise when he or she reads a book, listens to the radio, watches TV, surfs the Internet or attends a concert.

So "Content" itself becomes an object of perception and reflection, of "Cognition" that is. In the long run this can be quite strenuous: beware of the "information overload"! Fortunately, for some time now, we have been supported, at least potentially, in our endeavour to understand and process not only the world we directly access through our senses but also the artificial worlds we ourselves have created. For instance like this:




This takes us to "Digital Content" and "Cognitive Systems"; systems, that should make our own "Cognition" more effective.




Such systems, tuned to the analysis of "Digital Content", could for instance be asked to classify a given image and to deliver formal and informal descriptions, of what can be seen in it, its content.

More generally, apart from enabling the generation of "Content", our cognitive capabilities are prerequisite for all sorts of actions and activities that result, deliberately or not, in changes of our living conditions, globally or locally.




But it is quite fitting here to note that we are not the only species in the animal kingdom capable of such feats.




However, we humans, unlike the very most of the other animals, consciously and purposefully produce machines and other technical systems to aid us in performing a multitude of hands-on tasks in the world of our five senses. In many ways these systems could be much better than they are today. It would be desirable if they could, of their own accord ("autonomously", so to speak) relieve us of many a cognitive burden that comes with carrying out such tasks. (Examples abound, in surgery, in driving a car, in manufacturing all sorts of goods, etc.) Instead, these machines, whatever they do, they must do it upon our command and under our tight supervision.

It is an age-old dream of mankind to build autonomous machines. One of its possible realisations has been known for some time - decades, in fact - under the term "robot", inter alia. Their most popular representatives even appear, what else to expect, in human-like shapes. Frankenstein and Golem ante portas ... ?




And here is an entire menagerie ... in 2050 their successors are supposed to defeat the best (human!) football team of the world.



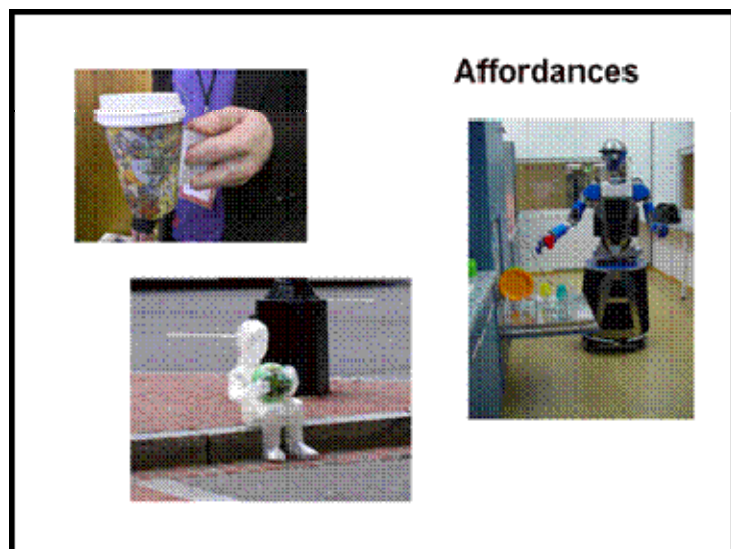
And there are even more daring visions!

But there are many robots that have not been created in man's own image:



None of these machines and systems is poised to solve chess conundrums or any other classical AI ("Artificial Intelligence") problem; they are not supposed to keep track of budgets or stock. But they have to keep their wits about themselves, for instance when it is called for to recognise a certain object if viewed from a different angle or under different lighting conditions. There are other systems that ought to understand speech and identify a speaker even in the presence of strong background noise. All of them should understand to a greater or lesser degree, aspects and features of their environment (spoken language may be part of it). They should act in compliance with what the objects afford upon which they act.

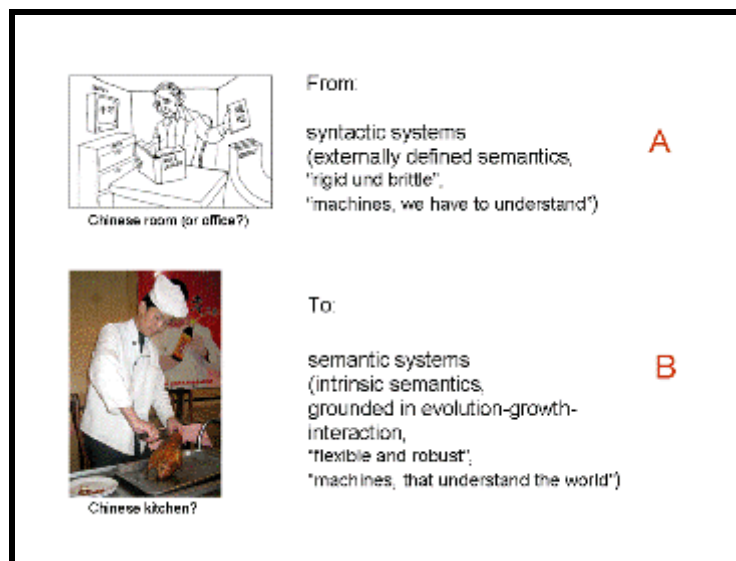
They should for instance "know" what the handle of a mug is good for, a dish washer, the curb of the pavement along a busy street, et cetera; they should know what they can do with them ("*affordances*").



Machines and systems that in this sense are "cognitive" need not necessarily be at the height of our intellectual capacities or, in the same way that we are, be conscious of what they are doing. What they are supposed to do is well within the purview of many animals, living entities that is, who have much less grey matter at their disposal than we humans have. And we do not even know exactly how animals accomplish their many remarkable feats.

What is the EU Programme "Cognitive Systems & Robotics" about, what are its aims?

For quite some time now many clever people who do research in relevant fields have been fairly sure that new approaches are being called for in order to engineer machines and systems of the kind we just described. These approaches should go way beyond what has fondly been nicknamed "Good Old Fashioned Artificial Intelligence" (GOFAI). John Searle, an American philosopher, invented the "Chinese Room" gedankenexperiment (thought experiment) that is often used somewhat teasingly to characterise GOFAI. The person in that room does not know a word of Chinese; yet he gives perfectly sensible answers to questions put to him (in writing) in classical Mandarin, based on a book of syntactic rules, compiled by people who do master Mandarin. But he has himself no clue what all the characters appearing in those rules mean.



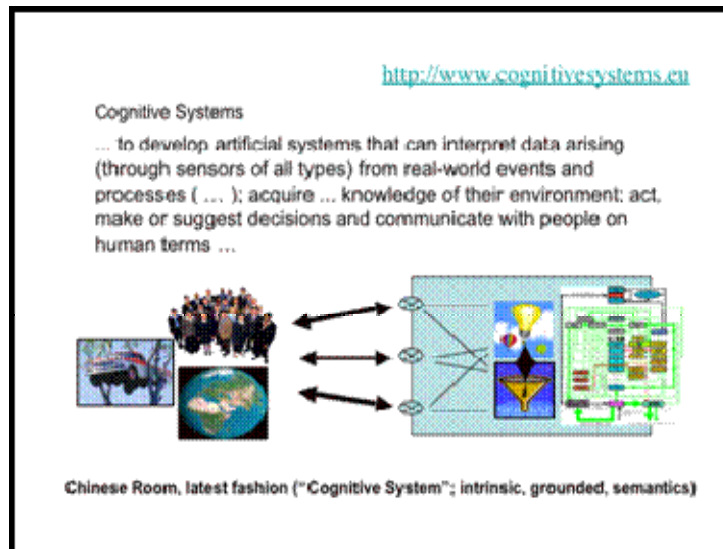
The "Chinese Room" metaphor has triggered controversies that even today, some 30 years later, have not subsided.

Some say (e.g., John Haugeland, another US philosopher): *"Don't worry, if you take care of the syntax the semantics will take care of itself"*.

Others say: *"Abstract syntax with externally defined semantics will at best give rise to virtual language games and not be able to keep up with the uncertainties of a world that is in permanent flux. Semantic intelligence does not necessarily result from syntactic intelligence."*

This view stipulates some sort of grounding of abstract symbols and systems in real-world contexts, and to generate richer and more dynamic representations of what is going on in the world and what that means.

This view appears to gain itself more and more ground. It is in fact the ground our "Cognitive Systems and Robotics" programme is largely based upon.



One of the first versions of our programme sets the following broad objectives:

... to develop artificial systems that can interpret data arising (through sensors of all types) from real-world events and processes (...); acquire ... situated knowledge of their environment; act, make or suggest decisions and communicate with people on human terms and assist them in tackling complex tasks ...

This has up until now not greatly changed. By some sort of "stepwise refinement" at best, these objectives have been and are being broken down further and made more concrete. A reader interested in more details may consult relevant chapters of the website <http://www.cognitivesystems.eu>.

There has been another constant since the programme was launched (in 2003): the emphasis on "strengthening the scientific foundations of designing and implementing artificial cognitive systems". This includes not only the pertinent formal sciences (e.g., informatics, mathematics, etc.), engineering disciplines and classical natural sciences but also at least parts of the neuro- and behavioural sciences (relating to humans and animals). Developing and advancing relevant methods of machine learning (or "learning and adaptation in and by artificial systems") is one of the key issues in this endeavour.

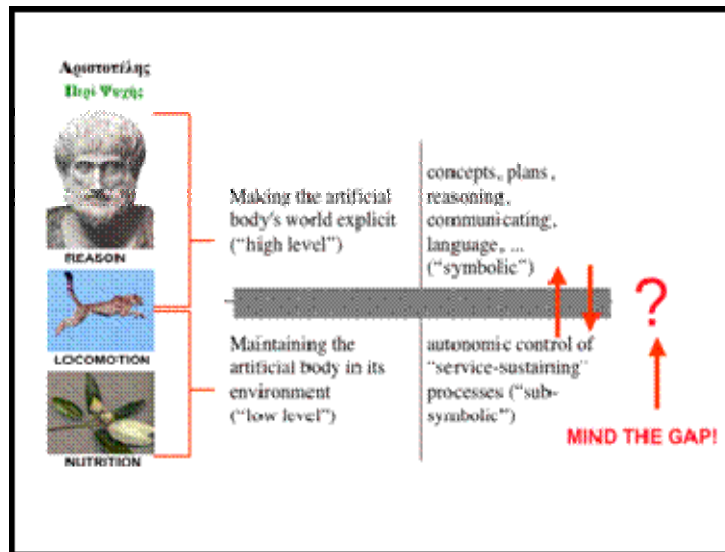
What has so far come of this programme?

Expressed in numerical terms, quite a lot. As the programme (in 2010) approaches the end of its eighth year there have been more than 100 projects on its past and present "payrolls", receiving more than half-a-billion Euros of funding (<http://cordis.europa.eu/ist/cognition/projects.htm>).

More than twenty will be added to the list next year. A brief presentation like this one therefore only allows a rather cursory assessment.

One can, for example, try to structure the set of all projects according to various criteria. According to specific "cognitive competencies" for instance, a given project is aiming to implement.

As the next slide shows, already old Aristotle had ideas on that which to this day are pretty much up to date.



As a matter of fact, many of the above projects can be more or less precisely allocated to one of the two levels described on the right hand side. Among these and apart from these there are those that more specifically focus on the transition of the "anima sensitiva" from the "sub-symbolic" to the "symbolic" domain. This yields the following very rough classification:

<http://cordis.europa.eu/ist/cognition/projects.htm>

A classification in terms of environments and pertinent issues addressed

focus on Type of environment	low-level competencies	high-level competencies	minding the gap	L E A R N I N G
1 real				
2 virtual				
3 hybrid				

L E A R N I N G

The row labels indicate different types of environments (or "umwelts") in which machines and systems act, in line with their respective competencies and capabilities. We distinguish between say, everyday worlds (the worlds of our five senses), "digital worlds" (e.g., content networks, webs), and all sorts of hybrids. Learning, as already pointed out, is a must everywhere.

These projects implement our programme. Within their limited timeframes they are supposed to add to our knowledge and prowess - if not also to our insights. In the best case they may even bring about results that can be used, in the short or medium term, to create new marketable products or to enhance existing ones. But this is not easy. Most benefits will be

reaped in the longer term. Too much groundwork still has to be laid. Scientific research clearly has to dominate but must not happen in sterile lab environments or closed studies only. Proposers are therefore called to "define their projects through hard research problems and ambitious but relevant and realistic scenarios where (even approximate or partial) solutions of these problems can lead to measurable progress".

This takes us to "applications", to services and production, for instance.

Cui bono ?

This question may spawn long-winded answers and provoke splendid arguments. I shall confine myself to the same, rather personal and perhaps somewhat proselytising answer that I gave while sitting on a panel concerned with "Robot Ethics", at the "Human-Robotics Interfaces" conference in 2008 (HRI'08).



One of the questions posed was: *"Should we continue and fund research that will produce autonomous agents that could be used and/or abused in ways that will increase suffering of humans?"*

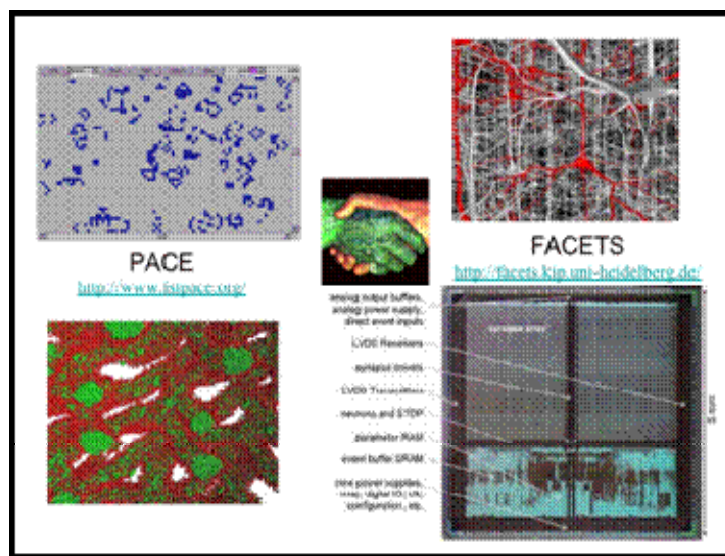
My response: *"We should fund research the expected results of which enable us to create better living conditions for everyone on this planet, and research that helps us to better understand ourselves and the world we live in. (In fact, the latter is a prerequisite for the former.) We – concerned citizens (and scientists and engineers in particular) - should make every effort to ensure that research results can not be abused to increase the suffering of people. Past experience shows that this is difficult. Open access to knowledge helps but is not sufficient. We should not give up though."*

What next ?

We can only guess. Umberto Maturana and Francisco Varela, two highly esteemed Chilean neurophysiologists and cognitive scientists, once said: "Life is cognition". This equation may go a bit too far. The somewhat weaker statement "no life without cognition" (assuming a sufficiently generous interpretation of the term "cognition") is probably less contentious. But then we can still ask: *"Things to which we ascribe cognitive capabilities, do they not in one way or another have to be alive?"*

It is therefore not surprising that biomimetics is gaining currency in artificial cognitive systems. Biometrics has a long-standing tradition, older still than the cybernetics of the late forties of the previous century. And at least since John von Neumann's cellular automata and Conway's "Game of Life" computer scientists too have been claiming their stake in a discipline that under the name of "Artificial Life (AL)" has made a remarkable career.

By now not only simulation in software but also emulation in dedicated hardware are on AL research agendas, as for example in the (FP6) EU projects PACE (Programmable Artificial Cell Evolution, <http://www.istpace.org/>) and FACETS (Fast Analog Computing with Emergent Transient States, <http://facets.kip.uni-heidelberg.de/>). PACE attended to life's elementary building blocks whereas FACETS, with neuromorphic hardware, attempted to recreate processes occurring in the living brain.



In a way both projects were committed to an "embodiment" paradigm which has become very popular among the proponents of "new AI". It is in fact a corner stone of the "new AI" and postulates the inseparability of "intelligence and corporeality" and, in a final analysis, the indivisible oneness of "Mind and Matter", in contrast to an abstract functionalist view. More succinctly:

“Mind Matters Matter!”

So it may turn out that in order to make substantial progress in the design and implementation of artificial cognitive systems we may indeed have to reinvent life itself. This is risky. Even more so as in doing so we are entering the realm of nano-objects. Of course, this is not the only risk. Ethical issues abound. We should at least be aware of an old fairy tale - one of my favourites - recounted by the Grimm Brothers, of: The Fisher and his Wife.



The enchanted flounder the fisher had one day encountered in the placid sea and returned to its element, thankfully gave in to the wife's every demand, promoted her from the filthy shack where she used to live, all the way to the papal throne, and made her richer and richer, and more and more powerful. Then the woman wanted to be God. Upon hearing this the flounder told the fisherman: "Go home. She is sitting in her filthy shack again." And the sea was roaring like hell.

There surely are various interpretations possible. A moral of the story could be: let us be mindful of what we wish when we believe we should be able to have it or make it. Not only in the world of fairy tales but in real life as well.

Luxembourg, September 2010