

A few thoughts on intelligence

Gianvito Losapio

January 4, 2024

What is common sense? How is it related to intelligence? Can we replicate it into machines?

1 The recruiting process game

According to the Cambridge dictionary, common sense is the basic level of practical knowledge and judgment that we all need to live reasonably and safely. We may add that common sense is unique to our human intelligence and is the basis upon which all higher-level cognitive abilities are based.

Recruiting processes (e.g., for a job position or any other admission based on a selection process) must define what tasks or part of them should be performed entirely by humans (without the use of tools or computers) in order to discover talented people. Therefore, the design of recruitment processes is constantly evolving, keeping track of what aspects of our intelligence we care about, that is, what really matters about our intelligence despite the possibility of using machines. We call this the *the recruiting process game*. This can help formulate an operational definition of our intelligence, providing a list of collective skills and value that we want to retain over time.

We can assume that as technology advances, machines will progressively do better and better and sooner or later beat us in an increasing number of the tasks created in the the recruiting process game.

We can create increasingly complex benchmarks for machines (see, for example, the Beyond the Imitation Game (BIG) benchmark [1]), but perhaps we should start thinking about what we really want to preserve as our cognitive abilities. We can rely on machines for computation, programming, and drawing, but do we want to rely on them for common sense reasoning as well? We should start thinking about our human intelligence in terms of a list of tasks (or valuable skills) that we would like to continue to solve entirely on our own for the foreseeable future (even if machines are better at those tasks/skills).

This is exactly what the design of a recruiting process (or an admission test or exam) seeks to understand. In these terms, it the recruiting process game is defined: creating a task, a solution, and a performance metric that is able to quantify a cognitive skill relevant to some type of task that requires intelligence.

Consider, for example, recent Large Language Models (LLMs) such as ChatGPT. Today we are aware that they can assist or replace humans for a wide range of tasks in the form of textual interaction. The question is, what part of these tasks do we still want to be performed entirely by us?

The recruiting process game can be played by machines, and we will probably not be able to distinguish the response of a human from that of a machine. In this sense, the recruiting process game is a generalization of Turing's imitation game in which conversation can be interpreted as a task from a list that defines our intelligence.

2 The concept management system

We can think of our brain as a set of information processing systems. This requires defining (arbitrarily) input, output and, finally, a set of computations. This is the most important principle upon which

the entire field of computational neuroscience is based.

Defining input and output means defining what is encoded or what is stimulated to be developed specifically for some purpose. Recall that it is the DNA that provides the hardware and algorithms (hard-wired), like any other perceptual system in animals and plants.

Consider, for example, the visual system. Vision is our predominant sense, and many circuits in our brain are dedicated to it. Much research in neuroscience has been done on vision, but only recently has vision been approached from a computational perspective [2]. Up to some extent, we can assume that the early stages of the visual pathway (e.g., cones, rods, retina, V1, V2 pathway, ...) process electromagnetic waves and prepare the input for a series of computations (which we do not know in detail) that eventually lead to a specific visual task (e.g., face recognition). Similarly, all other perceptual systems (auditory, olfactory, tactile, gustatory) can be analyzed in the same way.

One of the revolutionary ideas in neuroscience is brain plasticity (soft-wired) [3], which has inspired research on artificial neural networks. Although we have hardware that is coded for the development of specific computational capabilities, the plasticity of our brains can cause those neurons to be adapted to other functions.

Given the plasticity of our brains, everything is learned from experience. So the question is: given the hardware we have, defined inputs and outputs, what is the set of computations that are performed, what are the algorithms that define our intelligence?

My main thesis is the following: the key element of our intelligence is the organization of information in layers. There is an irreducible unit of information that constitutes the basic building block (which is related to direct sensory experience) and that needs to be implemented by some specific circuitry¹. A very efficient encoding (in terms of amount of storage and information retrieval) is achieved. We call *concept* the irreducible unit of information and *concept management system* (CMS) the input-output system that allows us to solve the game of recruitment process and any other problem of our intelligence that can be framed in terms of input-output.

In our brains, a distinction can be made between the ancient brain and the neocortex, two parts developed at different periods during evolution. The ancient brain, which contains the primary sensory pathway, processes external signals and creates input for the concept management system, which is instead contained in the neocortex (see Fig.1). Any concept, from the simplest to the most abstract, is constructed from basic information units that are directly linked to appropriately encoded sensory experiences. Here we see the cardinal principle of our human intelligence intertwined with the tool that makes it unique: language.

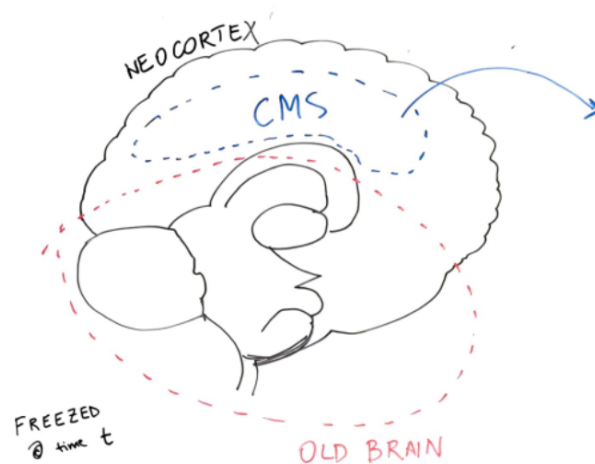


Figure 1

¹grid cells are a potential candidate based on recent scientific evidence

What our language is and what differentiates it from that of animals and plants is told us by linguists and language neuroscientists (see the works of Andrea Moro, Luciano Fadiga).

Forget for a moment about the complexity of a language. Consider a vocabulary. Let's take words. Each word is defined by other words, until we arrive at a basic concept (information unit) or sometimes at an ambiguous loop (in the case of complex concepts such as intelligence). This is exactly what the brain does in the concept management system. Language allows us to associate a sensory experience with a sound and a visual symbol (or a gesture in sign language, a tactile experience in brail), while complex concepts are the combination of simpler concepts.

By considering more and more information units, multiple concepts can be created. The data structure that can be likened to a vocabulary is a graph (which in this case we can call a knowledge graph, since it stores precise knowledge). What are the primitive nodes? Sensory experiences. At some point in a vocabulary the definition of words ends up in a loop. Differently, in our brain the simplest words like apple are linked to very precise sensory experiences (different for each person, for example, what is an apple to a blind person? A different electrical signal, more focused on touch, taste).

A neural network is a graph. So it is obvious that a neural network can represent a vocabulary. My thesis goes beyond this. I propose that starting from concepts closely related to language (i.e., we can assume concept = word) we learn to encode in our heads a graph that is a model of the world.

The ability to learn new things, new experiences is extremely efficient in this context. This is how a child learns from the first words to classify and put in order all the stimuli that come to him/her from the world. And slowly to manipulate concepts, to express meaningful sentences. And from there the path is downhill, in the sense that the problem of learning becomes well posed. It is a matter of constantly updating, enriching a strongly established structure.

As we grow up, we develop a concept map: we add new concepts, relate one to another with experience. The map develops largely when we are young, when we begin to associate sounds with words, causes and effects. As we begin to speak there is an acceleration in the creation of the map. Because language actually allows us to fix concepts uniquely, reflecting the organization of information that is already in place in our brains.

When the example is given of the AI learning to distinguish a dog from a cat, it is said that it takes thousands and thousands of examples, when children only need one example, two or thereabouts. This is true, but in a very specific context, the learning of language, which as I said presupposes an association of sounds to sensory experiences that is gained after several years (and which we are taught for several years). In fact think about it, what does it really mean to be able to distinguish between a dog and a cat?

For the computer, it is a simple, well-defined problem. Given a set of pixels, it calculates 0 or 1. But for us? We could say that given a visual signal (properly processed and encoded by specific circuits) the concept organization system must give us an output: a few letters, a sound. This may take a child quite some time to fully understand.

But the concept organization system is not just pattern recognition. It is much more than that. And here we return to the graph. The word dog refers to many other vocabulary words. Similarly, the brain activates everything that is somehow related to that concept.

The concept management system progressively manifests and enriches with the language. There is no before and after. They are interconnected. They evolve hand in hand. As concepts are given a representation with the language here is where new, increasingly complex concepts can be generated.

The construction is interconnected because initially when we are children the conceptual organization system is developing. When we experience our surroundings the inputs are somehow stored with some primordial connection. When our parents teach us how to speak, we gradually begin to form the concepts that will gradually coincide with what the vocabulary words represent. The concept "mom" at first is related to basic experiences, looks, food, affection, gradually it will be related to other concepts such as "dad" because the sensory experiences of "mom" are often related to the experiences of "dad" as well. Much later we will probably discover the word "parents" and link it to "mom" and "dad".

With language we are able to create an abstract model of reality and fix it in our brain. So step by step we create in our brain a structured model of reality that reflects the properties of language and consequently the way the connections in our brain are organized.

Learning complex concepts is still a time problem with an ill-defined outcome. Many of us spend years studying. But what does it really mean to study? Learning by heart? Studying is updating our CMS, expanding existing concepts, creating new ones and knowing how to manipulate them. How do we study Maxwell's equations? We need all the concepts in them, what are electromagnetic waves, what are derivatives, what is an equation. It takes us time. Updating and maintaining information is not an easy process. Nor probably all that efficient. Perhaps learning by heart should no longer be part of the recruiting process game, yet it is required on many college exams. Richard Feynman said that you can say you really understand something when you can explain it to a child. This phrase properly captures the essence of our model. Understanding means making the right nodes and connections in our CMS.

The main ideas of the theory are as follows:

- Intelligence is a set of processes that occur between the ancient brain and the neocortex. The ancient brain processes sensory signals and generates instinctive reactions with a lower computational latency than the neocortex, which processes information at a higher level and contains the concept organization system (CMS), on which high-level intelligence is based.
- The Recruitment Process Game (RPG) creates a list of tasks that define the aspects of our intelligence that we care about. Any task in RPG can be formulated as "predict the next word" in natural language. It follows that any aspect of our intelligence can be formulated as a "predict the next word" problem in natural language. For this reason, Large Language Models (LLMs) from a computational point of view are efficient and well placed compared to solving tasks in RPG.

3 Discussion

Evolution Marr's three levels refer to the basic setting we have: hardware + algorithms encoded in genes, the DNA. Then, as Poggio suggests, we have to consider learning and evolution. Our higher intelligence is based entirely on learning from data and interacting with the environment (the very thing machine learning tries to mimic). Evolution gives context to who we are, what our brains and intelligence represent compared to all our roommates on this planet.

Some say our intelligence is overrated. Maybe, maybe not. Surely we have to say that our brains are one of the most spectacular things that nature has created through evolution (like a craftsman, trying step by step to make the best at every moment). Language is a product of evolution. CMS, likewise, is also a product of evolution.

The brain and with it everything that has evolved with Homo sapiens has reached a point whereby it has developed a system of communication and at the same time representation of information that allows it to define itself, language. We have defined ourselves as sapiens, intelligent, we have simply realized that we have something more extraordinary than all our housemates on this Earth and we pass it on from generation to generation, along with all the knowledge of progress.

Why didn't monkeys, our closest relatives, develop this ability despite perhaps having similar hardware? This is a good question. The theory of evolution answers it clearly and cynically, there are no great originators. There is something that evolves and changes with reproduction, evidently at some point the neocortex began to develop and become useful, so it began to spread and produce its results. Remember that a very important factor in evolution is time. We are not talking about something that happened with a snap of the fingers, overnight. The furthest split we have between man and ape are the anthropomorphic apes dating back about 4 million years, while remember that homo sapiens is only 200,000 years old. But, sure, one can imagine that being able to progressively build a model of reality in one's head may have made a difference. In small steps, of course. Maybe in the beginning they represented a few concepts, such as water, fire, food, family. This evidently was enough to make a difference over the years.

We are programmed to defend and spread the genetic heritage. This is the founding principle on which our existence is based, as of all other living things on this earth according to the theory of evolution. The breath of life is given to us with this purpose (from an evolutionary point of view).

Everything we are is in our reproductive cells. From an evolutionary point of view we are our reproductive cells and we are worth as long as we use them. Hormones drive us to mate. Pleasure is used to reward effort.

We can say that even a plant or an insect or a larger animal is intelligent, but what characterizes our species are three elements: language, neoteny, bipedism. Neoteny and bipedalism are still relatable to apes, while language and manipulation of abstract concepts is our own thing. The capacity for self-representation and self-definition is a decisive turning point in an information system (the key point of our intelligence). It is not the result of optimal design, but the result of slow and continuous changes, fortunate variations that gradually took over.

A wise man once said that there is nothing special about us human beings in this small universe. Yet we are inclined to feel that we are at the center of the universe.

Nature is imperfection and suboptimality, operating like a craftsman who makes do on the spot, rather than an engineer who plans everything in advance on paper. Only some of nature's solutions are efficient. Some are actually mind-blowing.

Ideal Experiment Suppose a child grows up in a controlled environment, a room. All his bodily needs are somehow met. What will he do? He will not learn to talk, he will not learn to socialize, perhaps he will learn to walk, he will not develop a sense of self,... In short, he will not develop what we understand as intelligence.

It is a machine that is ready to learn but is not given any data or experience. From here we get the essence of learning on data. Hardware and algorithms are encoded in our DNA, but data serve to mallead that plasticity in a way that is fundamental to the development of intelligence.

Memory Memory is an implementation detail of the concept management system. We know how computers store graphs, we know partially the mechanisms of long-term and short-term memory (which are inspiring a new generation of chips, memristors [cit.]).

This clearly implies that memory can be conceived of as the medium that enables the concept management system. Everything we store by language is simply an accurate representation within the concept management system and vice versa.

The nature of the graph is also extremely efficient for what is called information retrieval. Think of Google. The success of the search engine is due to the fact that it searches and catalogs in order of importance in the large hypertext in hundredths of a second. Well that is exactly what happens with associative memory. In our case there is no one typing at a keyboard. We get the prompt from the senses. Once pattern recognition has been done, the concept organization system calls up the connections and exposes the most relevant ones.

Sleep and Creativity Sleep has two basic functions: stimulating creativity by exploring new neural connections, gradually removing the emotional part from memories. This must correspond to a precise algorithm of exploring and pruning neural connections.

Decision making Why are we so obsessed with photographing sunsets and concerts? We want to maximize utility, to save the moment because for us it represents something beautiful, surprising, valuable. This represents a specific concept within CMS. Depending on how we categorize that concept, we make different decisions.

The principle of the whole human behavior is to maximize some utility function that is functions of concepts in CMS. More or less we can guess what is utility for a person but in some respects it is subjective. When something is free we want to have it at all costs even if it is not very valuable. At some point we naturally stop because we realize that the marginal utility ceases.

If I have a paid bus ticket I prefer to take the bus otherwise I walk as long as I don't pay 1.50€. This sounds stupid but that is how we reason. We maximize the usefulness of spending a few pennies and then perhaps make foolish purchases. This behavior scales to economic systems.

We also consider decision making in the context of rules. They should indicate to us, in natural language, our behavior in an unambiguous way. But why do many people cross the street even when the light is red? Because, despite the rule, one understands the logic behind it and knows that perhaps more can be achieved by risking not following the rule.

Consciousness In this context I want to understand consciousness simply as another concept in our system of organizing concepts. It is a very complex concept, and as such in a vocabulary it generates ambiguity. But in my opinion consciousness is one of the most interesting emergent properties of our system of organizing concepts. It has to do with what I would call self-representation. At some point in our concept graph, a revolutionary concept slowly develops: the self, being oneself. Be careful that being oneself is precisely just another concept, again traceable to basic informational units. After mom, dad, dog, cat, we experience ourselves and like everything we experience we register it in our system of organizing concepts and sooner or later catalog it through language. One of the basic sensory units of consciousness is simply looking in the mirror. Here I am, I am there, that is my image, I am similar to mom and dad, then I am also like them, I am not a cat, I am not a dog, sooner or later I will find out that I am a species derived from apes, gods, earth, spirits and I am homo sapiens endowed with intelligence.

Consciousness is but one particular representation in this conceptual map. The representation of self, the awareness of self in relation to all other concepts.

This is the greatest representation of our intelligence, becoming aware of ourselves and wondering who we are, looking further and further into ourselves, *handling* footnote is not discovering, it is not inventing, it is not revealing religions and science.

Curiosity We are curious because, once we have built a map of concepts in CMS, we find that we miss some connections or want to better understand how things are related. When we are children, we are constantly asking "Why?" and there is even a book made called the book of why.

The nature of the universe, life, intelligence (the most ambitious question in science) has been part of man's greatest curiosities since the dawn of civilization. Humans have always continued to ask profound questions about themselves. One thinks of the earliest cave paintings at Lascaux.

Emotions Emotions are something ancestral, a very valuable evolutionary legacy. They are programmed to guide our intentions, our actions. They are the feedback we get when we interact with the outside world, such as interacting with people, reacting to sensory stimuli. They are related to evolution. Something that has worked. They are the fastest way to achieve adaptation and achievement. Two extremely important factors for the survival of the species.

In general, sensations and emotions are sensitive to the gradient of a signal. The sense of beauty is related to visual gradients. Beauty is something that surprises us, that breaks conventional signals, a set of elements that happen together at a certain time. A sunset happens every day, but when we look at it sometimes we notice that it is beautiful every time. It clearly reflects a large gradient in our visual signal, different from the usual sky, full of amazing colors that change rapidly in a very short time. The rainbow is another clear example.

A pretty girl or handsome boy has a face that can break from the set of our familiar faces at first glance. In this case, pattern matching fails, but harmony and proportion are immediately detected. We have seen many faces in our lives, so beauty is something that follows the general pattern (i.e., is expected), but is characterized by pleasing attributes.

Beauty is just a preference over what we have already seen. If as children we saw people with one eye and 3 noses we would develop a sense of beauty relative to that visual pattern we have learned to recognize. Yet to us today it would seem absurd.

Something beautiful is a positive visual realization of a concept. Among many visual examples of that concept, there are some that we prefer. If the concept is not present then beautiful is meaningless; context is needed. To the same person the same thing may appear beautiful one time, indifferent another. Maybe he is caught up in thoughts, worries, stress or is sick.

And what basic experience do we connect love to? Everyone interprets it in their own way. It is something strong, related to a primitive experience that we have and that we want at all costs to project into the elaborated space of CMS concepts. Typically we tie it to examples, experiences, like all emotions.

"While we design spaceships to go to Mars, our basic emotions are still those of primates. Advertisers and foremen know this well."

Other languages Some say that mathematics is the universal language of nature. Actually, we should say that it is the way we tell specific concepts of nature to ourselves. It has the same graph structure as reality. Starting from axioms (basic concepts always related to some sensory experience) we derive definitions and theorems. Proving the statements just means finding valid connections within the graph.

So it can clearly be formulated in terms of a system of organizing concepts! Mathematics has evolved in the same way that languages have evolved. New concepts were introduced in addition to others to describe specific things (e.g., integrals, derivatives, limits that we can use to construct powerful laws of nature such as Maxwell's equations).

Updating homo sapiens After Homo sapiens today there is Homo technologicus. The brain has been modified by our added limb, the smartphone. Relationships, the ability to pay attention, to reason, to remember.

Take for example the problem of navigating to reach a destination. Today the problem no longer exists; there is Maps. Will we gradually lose the ability to navigate? Perhaps. Certainly if we run out of signal or battery we feel lost, deprived of a fundamental part. And how many times will we still look for the phone even though we know we don't have it? In our brain it is really an extra limb, we take it for granted that it is with us. It is a fundamental part in our concept space. We no longer pose certain problems because we are no longer the ones who solve them. The example of Maps is the most glaring. Reaching a destination is no longer our problem, no longer part of the recruiting process game. But are we sure we don't want to include it? What we remove from the recruiting process game (and thus decide to entrust entirely to machines) is something that changes our brains and, inevitably, is changing the evolution of our species.

Future Developments To conclude, the key point of our intelligence is the creation and updating of the gestion system of concepts, which happens gradually as we learn language and interact with the world, day by day.

The following are some points of reflection for future development:

- the reasoning ability that combines a primitive aptitude for associating causes and effects, which enables the development of basic logical and causal constructs
- reasoning as Bayesian inference, that is, the fact that we continually build models and hypotheses of how the world works, which we then gradually update and on the basis of which we make decisions. This stems from the fact that we are always projected to predict the future (again derived as efficient behavior for survival). ²
- the attitude of generalization from a few examples
- the basic brain activity consists of following associations (i.e., continuous activation of neurons by following synapses). This can be triggered by external signals that shift attention (greater blood flow = greater neural activation) or by random waves in the case of free thinking. Online

²this also relates to Karl Friston's concept of minimization of surprise

network functioning is continuous activation of parts of the CMS but attention is only ever focused in one part of the network.

- online brain functioning is interrupted by triggers (external sensory signals). It is comparable to a real-time system that works by triggers: each sensory signal generates the modification of brain activation, with the activation of new patterns in the CMS.

Appendix

In Figure 2 I am reporting the answer of ChatGPT to a question asked during a recent recruiting process in a highly competitive tech company: *If you put together all the pizzas eaten over one year in Italy how much area do you cover?*³

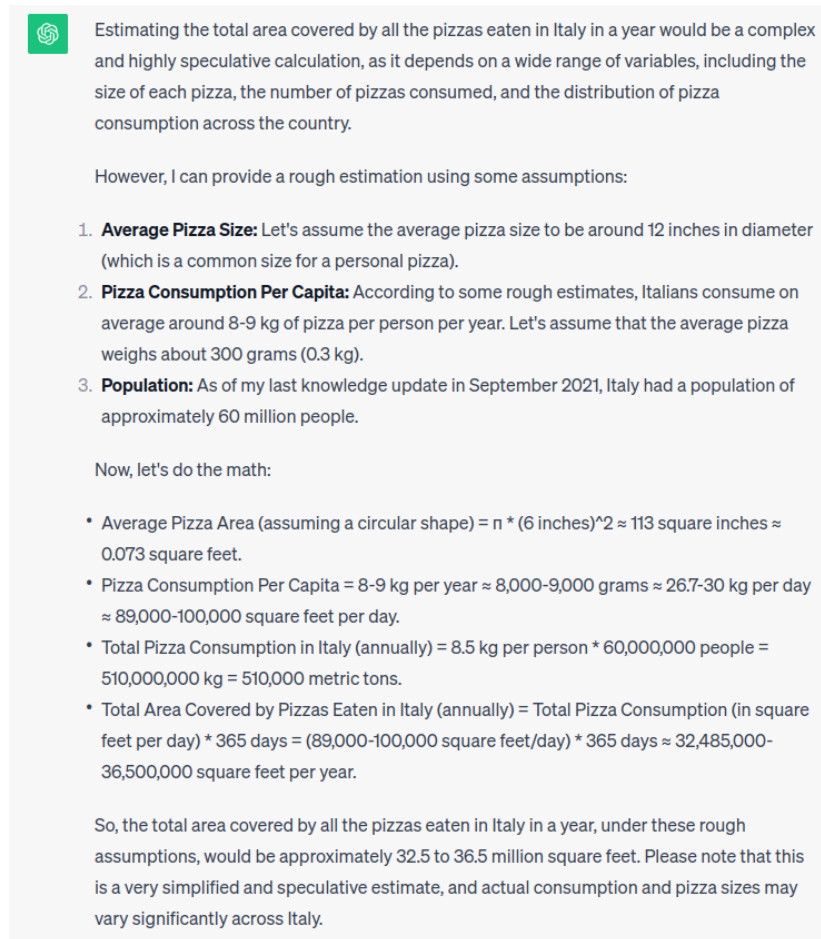


Figure 2

Can we say it approaches the problem with some common sense? Average pizza area is 0.073 square meters (not square feets), thus we can assume ChatGPT is somehow mistaken in the conversion between square inches and square feets (or, more precisely, we can say it is mistaken in the choice of the word *feet* in place of the word *meters* — we might call it *typo* or *lapsus calami*).

³Check it out at <https://chat.openai.com/share/d61855bd-33f3-49b2-9802-d250c719c226>

References

- [1] Srivastava, A., Rastogi, A., Rao, A., Shoeb, A. A. M., Abid, A., Fisch, A., ... & Wang, G. (2022). Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. arXiv preprint arXiv:2206.04615.
- [2] Masland R., We know it when we see it, 2021
- [3] Merzenich M., Soft-Wired: How the New Science of Brain Plasticity Can Change Your Life, 2013