Is AI Really What You Think It Is?

Each time I fail to avoid exposure to wild stories, predictions, verbal expressions of individual anxiety (or whatever you want to call them) regarding Artificial Intelligence (AI), I cannot help but wonder whether some (hopefully not all) of my contemporaries actually do perceive AI as a "thing", some item one employs to bridge a gap or fill a crack, you know, "just put it right here and Bob's your uncle". I loathe being the one to break the news, but that is not going to happen anytime soon … or ever.

Neither "a Fairly New Technology" Nor "only one Click Away"

The other day, I heard a presumed expert on the matter state in an interview that "AI is a fairly new technology and we have to adapt to the idea that it will change the world of labour". Well, not quite but yes, of course. (Common decency demands for this person to remain unnamed here, because referencing them as a source would mean adding insult to injury.)

Granted, compared to the "hand axe" or the "controlled use of fire" (both technologies, having arguably changed the Stone Age world of labour for good), AI is "fairly new". Yet in this day and age, when the majority of work force (of the foreseeable future) cannot remember (let alone imagine) a world turning without permanent internet connectivity, anyone who still has to "adapt to the idea" of AI's probability of influencing the world of labour (and beyond) appears to have missed a few stages of more recent human (and technological) development.

It wasn't before 2006 that AI really took off (due to technical limitations "in the early days"), but even so "Artie" (as a practical implementation) has already come of age (at the time of writing). As proof of concept, however, "he" has already reached pre–retirement.

In 1983, when this writer was still an unruly teenager (and many a precious reader may not have even been born yet), one Stanisław Lem, a Polish philosopher, novelist, and essayist, published an essay titled *Waffensysteme des 21. Jahrhunderts oder die verkehrte Evolution,* and by doing so broke about every rule in the book. (Why it was first published in German rather than Polish or English? Well, read it and consider the time and place it was conceived, the answer will inevitably follow.) Little

did I know at the time that Lem had been writing about all sorts of "artificial intelligence" for more than two decades already (e.g., *The Invincible*, 1963/64 or *Eden*, 1959). If you, precious reader, just happened to think "yeah well, science fiction …", you may want to read one or more of the titles mentioned above and then think again. His descriptions (Dare I call them "predictions"?) of what the greater public know as "robotics" and "nanotechnology" today were so accurate and detailed that he made poor old Michel de Nostredame (arguably better known to some by the Latinised mononym "Nostradamus") look like a delusional junkie.

Every once in a while, I wonder what would have happened, if he had the book published in the early 2000s rather than back in the day, when the internet was still in diapers. He probably would have had to book a cosy room in a South American embassy well before the day of its (probably legally challenged) date of release. Yet, as he used to quip, his books were never read. Those who did read them, failed to understand, and those who did understand, forgot.

So what we are discussing here is by no means a "revolution of the world of labour" or a "change of game" that's going to hit us anytime soon, but rather an evolution that's been in progress for considerably longer than any of us can possibly remember. In other words, we (that would be you, the reader and I, the writer) have already been living — both in theory and practice — with artificial intelligence our entire lives; and so have our parents (both the individuals and the generation). The only difference between them and us is, most of them didn't waste a single thought on it, and if they did, boldly dismissed it as "sci–fi crap", while we become gradually aware of this reality … and the various intricacies that come with it.

Divide and Conquer

I cannot possibly prove it, yet still I'm rather certain that to try and break down complex tasks, challenges, or problems we face into smaller, more digestible ones is one of the oldest strategies we have known and applied. Assigning these subtasks each and all to the party best suited to tackle it successfully is only a matter of logic.

It is by instinct that we try to survive (as individuals, societies, or species), but to actually accomplish this goal takes conscious decisions and intellectual effort. Trying to "survive better" takes

even more sophisticated, applicable means. To this end, we need to see and recognise patterns (perceive the "bigger picture", if you will), both within and around us. The faster and more efficiently we acquire the necessary skills to arrive at this stage, the more successful we will be.

So what was a small band of hunters and gatherers (to name a random example) to do in order to get by rather comfortably? Right! They assigned individual tasks to those members of their "society" with the highest probability of tackling them properly. They learned to collaborate, in other words.

The faster and physically more robust hunted and killed their prey, using weapons manufactured by their fellows particularly skilled making those tools. Those who took more interest in the subtleties of life, and had a mind to deviate from the narrow path of caution every now and again, took charge of home and hearth. They specialised, in other words, for they had figured out at some point that assigning certain tasks to the wrong party (or none in particular) could easily lead to disaster.

A good chin–wag while guarding the fire or a slight deviation in preparing your food, is fine and dandy (and if all goes well, might even lead to a fully developed language and a greater variety in your diet), but allowing for either (idle chatter or unforeseen deviations in your modus operandi) while lying in wait for a wild animal is a promising approach for the hunting party to quickly become prey, one and all. So everything in moderation and its own sweet time … and handled by the party most capable of getting the job done.

Fast–forward a day or two to the moment AI takes the stage (unfortunately, it is also the moment when matters turn a tad more complicated than they used to be). What better instant for our "natural intelligence" to eventually kick in and apply the strategy discussed above. Let's break down the problem and tackle the pieces one at a time, lest confusion should prevail and complexity overwhelm us, shall we?

All Those "Funny" Terms …

The first logical step is to make sure we all have a solid grasp of the technical terms that are going to be used in the process. Knowing the subject at hand (and myself) fairly well, I venture the guess

there will be plenty. Hence, I'd be a darn fool neglecting to make immediately sure we all are (quite literally) on the same page at all times. Here's a quick overview of what's going to be discussed on the following pages (terms and phrases in question will be printed *in italics*):

Intelligence (in humans or, more generally, in any life form) is commonly considered the ability to make proper sense and use of available information. The word's Latin origin, *intelligere*, leaves no room for interpretation here: to understand, comprehend, or perceive. (Even "Leo", the AI–driven assistant implemented in the *Brave* browser, knows that. A little more on him and his "little friends" in a bit …) I do actually prefer "information" to "knowledge" (at least — but not exclusively — in this particular context), as "knowledge" (the term) conveys more often than not a quite questionable connotation. One may be intelligent, but "know" close to nothing (i.e., is poorly or not educated), while another one may be "quite learned" (has plenty of "knowledge" or is "knowledgeable"), but not able to put this resource to good use. Bear with me, my reason for insisting on this distinction will be gradually unveiled, as soon as we begin discussing individual models of artificial intelligence.

Artificial Intelligence is a different animal altogether. It is undoubtedly artificial (in that it is not the result of natural causes), but intelligent it is not — at least not by the definition mentioned above. The reason for that is rather obvious, actually: it lacks creativity altogether. Creativity means to create "something from (almost) nothing" (sorry, I couldn't think of a more striking metaphor). It means to create something original; something that hasn't existed (or been known in that form) before. So, while the creation of (any model of) AI is in and of itself a manifestation of creativity (even though it is basically just a simulation of individual aspects of natural intelligence), its output is not. AI–driven tools may be able to generate or reproduce results much faster (and if all goes well even better) than you or I, but they are not able to create or produce anything they haven't been introduced to before; they won't make "something out of (almost) nothing", without having first received detailed training in a particular field of action (not to mention an obscene amount of pertinent data). They may also gather and analyse data faster than any of us ever could, but they are not able to comprehend the implications of their findings, without being thoroughly trained to do just that — and how. Generally speaking, the way it is commonly used, AI is merely an umbrella term to

describe (any and all) aspects of approach to garner useful information from a large set of available information and distil the result into a digestible sample by way of computational power — the one term to describe them all, if you will. The three individual aspects most commonly known to the greater public (or more precisely, the aspects the greater public is most familiar with) are the following (which is why I shall focus on these three in order to keep things simple):

Generative AI is an umbrella term for any model designed to "generate" output (of any kind; e.g., text, graphics, music, etc.), derived from available data. "ChatGPT" is probably all the information the average person today needs to (believe to) know what we are talking about, but there are others, too.

Predictive AI is the umbrella term for models designed to derive useful information from a larger set of available information in order to make decisions based on the results. The machine's take on *quantitative* analysis. Obviously, the potential range of application for these instances is quite open and diverse: political elections, sports events, economic developments, etc.; basically, anything that relies upon statistical data.

Discriminative AI is the umbrella term for models designed to distinguish between sets of information. The machine's take on *qualitative* analysis. The most obvious example (and arguably the one the precious reader will be most familiar with) is a "CAPTCHA" (which stands for "Completely Automated Public Turing test to tell Computers and Humans Apart"; Source: Cambridge Dictionary, since I keep forgetting the "Public" part and then wonder what "P" might stand for).

Now, what's left? The "real catchy" terms, of course. Those terms that could easily keep impressionable people awake at night. Like:

Large Language Models (LLMs) are basically neural networks (see Deep Learning below) trained to simulate natural speech. How does that even work? Well, quite similar to learning French in the Foreign Legion, I would imagine: You get a word (or a string of words or even vast sets of content) thrown at you (that would be your input) and you are supposed to process it correctly (and fast). How well (or poorly) you did is determined by your instructor (who is the provider of both the input and algorithm). Depending on the way your instructor communicates orders (the algorithm), the

support of your comrades (the nodes in the network), and the situation you are in (the hardware infrastructure), it may take you several months (or many more months, if the computational power is less advanced) to get anywhere with your language skills. Yet once you are fluent, you definitely know the language inside out (which does not necessarily mean you will know what you are talking about). At the end of the day, you just try to predict the next word (or string of words), based on the input you received and the context, delivered by the network and the algorithm. In technical terms, an LLM is a predictive AI–model supported by a neural network. Whether it functions as a useful "conversation partner" or simply as a "talking head" depends mainly on the respective input.

Artificial Narrow Intelligence (ANI) is what you might (already) see at work in various industries (and you most certainly will see more of it in the foreseeable future). Whenever one needs a repetitive, rather specialised work process executed (almost) without flaw or glitch, an ANI is the go–to model. Why so? Simply because the "human factor" has been taken out of the equation. A machine never gets tired of performing the same (simple) task (or answering the same dull questions) over and over again, for hours, days, or months on end. Machines don't get bored, they don't lose focus, or are concerned about their "work–life balance". On a sunny day, they are every entrepreneurs wet dream. If, however, things happen to go sideways, they may easily become every entrepreneurs nightmare. I trust the precious reader is able to connect the dots without my assistance …

If you happened to wonder why I kept using the singular ("performing the same task" instead of "tasks"), that's because it takes a different model for each individual task. That's one of the downsides of this concept. Their range of operation is rather limited; they may not (all) be restricted to one single task per node only, but the operations one model may accomplish is quite narrow (hence, the name). It is artificial because it is not prone to natural flaws and deviations, that is both its advantage and disadvantage. It may be more useful (if all goes well) than any of us, but it doesn't "know" what it is doing. Unfortunately, that also means it does not realise mistakes (unless someone took it upon them to train it to accomplish exactly this task). If it happens to produce faulty results (be that items or answers), it will do so until someone (with natural intelligence or

common sense) puts an end to it (so much for its “intelligence”). This, at long last, brings us to the “premier class” of artificial intelligence …

Artificial General Intelligence (AGI) is — How shall I put it: “sorry to say” or “happy to report”? — the wet dream of every movie producer who’s run out of plausible plot ideas. The way it has been portrayed in several (I’d even say, way too many) cinematic works over the decades, your favourite humanoid has the same probability of “independent breathing” as Frankenstein’s monster: none. It is undoubtedly an interesting idea, but lacks any scientific basis — not to mention its practical applicability (for quite a while to come … or ever). Building a small number of humanoids and teaching them some tricks is one thing (and costly enough), but developing the infrastructure to make them flawlessly function and seamlessly interact among themselves (and with humans) independently (which would also include their maintenance) poses a decidedly more profound challenge.

Let’s not get carried away by the sight of a small host of “bots” playing soccer (or similar ventures occasionally making the news) — seemingly without human intervention. That is cute, I agree. However, it is a far cry from AGI (humanoid or not) taking charge and running the world as we know it. How could these obstacles be overcome? Well, Machine Learning (ML) would be one option (at least to some extent), Deep Learning would be another.

It’s been repeatedly said that AI is comparable to the intelligence of the average six–year–old. While this notion may not sound too promising when looking to develop intelligent tools that may save (or cost) human life, this supposed cognitive limitation could actually be an advantage — not to mention that there is no serious way of telling (at the time of writing) whether or not the above assumption is actually true.

Machine Learning is — outrageously simplified — the machine equivalent of an unspoilt toddler attending a premium kindergarten. Imagine a highly–trained teacher instructing an extremely focussed child (with the cognitive abilities of one twice its own age) how to tell one pattern from the other. Obviously, just about every six–year–old would be a “rock star” in kindergarten.

That the aforementioned CAPTCHAs still seem to work so poorly is not the method's fault, it is more often than not the quality of the samples used to train the model that causes the many false positives and negatives.

The "machine" does not actually "see" the image, but simply takes what patterns of information are available, scanning them layer by layer — brightness of pixels, alignment or clustering of pixels, etc. — depending on the computational power of the Graphics Processing Unit (GPU), one layer at a time or several layers simultaneously, and tries to predict what is to be seen in the image. The, by now, quite famous "too many fingers" in AI–generated portraits of people are testimony to that misinterpretation of information. The most promising method to mitigate that flaw would be deep learning.

Deep Learning is probably the methodology closest to human acquisition of knowledge. Imagine machine learning as a means to learn to negotiate a path that is full of y–junctions. Eventually, "the machine" will learn to determine which of the two available options it has to take to travel successfully (or with a higher probability of success, really) to the end. In deep learning, on the other hand, it may take advantage of an artificial neural network. The trainer weighs the possible options and so tells the machine, which is promising a better chance to get along, depending on the given situation, like a vast network of streets, with some that allow you to travel faster (on a sunny day) and others that make for safer travelling (when the weather's fierce). What we call "intuition" is more often than not simply "experience" that leads us to make decisions we don't seem to think about (because we calculate the odds so fast, we don't even realise it). In deep learning, we provide the machine with (a simulation of) experience, that allows it to make situational decisions it would otherwise rule out for contradicting pure logic. In that sense, Lem was spot on, calling neural networks "swarm intuition" rather than "swarm intelligence". Again, this is a bold, plain–language simplification of the matter.

What's It to You, or Me, or Us?

Well, as is the case with many a problem, you ought to know what you are looking at to get what you are looking for. In other words, if AI is to you "just about anything a computer does faster than

you do", then AI is not going to be of much use to you. And you need to understand that "AI" is not merely some "code monkey" that somehow fell off the tree into your waiting lap for you to blindly delegate tasks to you'd rather not do yourself. Fancy an example? Here goes:

Take your mobile phone. Yes, that one. It doesn't matter what make or model it happens to be. Now type a text message (but don't send it, if you have a thread of dignity left). What you (probably) just experienced was AI, if on a rather basic level. Whether the text correction or completion feature kicked in as soon as you typed a letter or two, both are varieties of *predictive* AI. The face you are pulling and the suppressed curse you uttered are testimony to the quality of training that particular AI model has received until its latest update. It won't get any better than this, unless someone is prepared to contribute hours upon hours to fine–tune the training or rewrite the algorithm from scratch. Both approaches would be quite costly and bothersome, but, let's be honest, the potential benefit would not justify the effort.

Would you really want this tool to have its way and send the message, without curating it yourself? Rather not. The quite evident consequence is, the longer your message, the more bothersome it gets to curate it. Yet your goal is to save time typing those messages (without making a fool of yourself). Hence, the tool you ought to be looking for would be a *generative* AI model. You'd still have to proofread the result, but even the first shot would be way closer to what you were aiming at. Turning off "autocorrect" and "autocomplete" and using a generative AI model to write your messages would save time (in the long run). Now, there you have it: Since we are aware of what we are looking at, we instantly know what we are looking for, and employ the tool best suited for our needs. Yet (and I wouldn't mind, if you marked my words) you should never be content with whatever was coughed up by a model that was trained with material you have never checked yourself, based on an algorithm you have not written yourself, and offer the result to the world at large as your own brainchild, for you are certain to embarrass yourself beyond intelligible words (sooner or later).

So, Theoretically We Could …, Couldn't We?

Of course, theoretically we could automate just about everything, save boatloads of time and money and the world along the way. What's going to get in our way, though, is the real world (the one beyond bits and bytes). Let's take the above example (the mobile phone): it's a fairly simple model, close enough to everyone's (presumed) reality and thus just perfect for our thought experiment.

What would it take to make "autocorrect" and "autocomplete" useful features (work as expected for every user, regardless of make or model or carrier, in other words)? Well, first of all, a paradigm shift in politics, economics, and society — and also in certain areas of science. Yes, unfortunately, we'd really have to go to such lengths to actually get anywhere with this rather simple task. Yet that single step would instantly have a lot more people gasping for air than the random utterly moronic text message.

Next, we'd not only have to make the source code of all models publicly available, but also all messages sent via individual apps (anonymised, of course). These would have to be constantly reviewed and compared with each other, because only then would the individual models have a fair chance of keeping up with the ways we all tend to use and develop our languages on a daily basis. To eventually make proper use of all the harvested information, the models would have to be tied to a (sophisticated) LLM. The implications of all these steps are obvious to the naked eye. Each of our telephones would suddenly make all those three–letter organisations that are constantly suspected and accused of trying to breach our personal rights look like bloody amateurs.

Yes, we would have to go that far to actually make it happen — we would literally have to go "all in" — or all individual efforts taken in the process would be in vain. Would it really be that complicated? No. Truth be told, I boldly simplified the matter; if "complicated" were a mathematical expression, the proper term to describe the process would be "complicated^2".

To get this one feature on your mobile phone to work decently, would require us to utterly remodel the way we communicate and interact (with each other and with technology). One doesn't

have to be capable of seeing through walls to confidently predict that none of this is going to happen anytime soon — at least not for the sake of avoiding random typos.

I hear you. Big talk, no proof. Um, not so fast. Let's think it through, step by step, shall we? I trust we are in agreement by now that "autocomplete" and "autocorrect" are (rather simple) predictive AI models. To make more accurate predictions (or, in the case of "autocorrect", identify incorrect words more effectively), it would have to be trained more thoroughly (i.e., learn more words and know how these are commonly used, in unity and division, and in various contexts).

That this training is not going to be conducted by individual users is a rather cheap prediction. None of these features was invented yesterday, but I have yet to meet a single person who is content with the way they work. Consequently, the training would have to be initiated (and conducted) by developers of individual apps. This is not going to happen, either! Simply because it is not economical to go to such lengths for a handful of users (provided that globally there are nearly five billion mobile phones in operation at the time of writing, even a million customers would have to be considered "a handful"). So to make that venture feasible, all developers (of applications) would have to go "open source" (make source codes publicly available). That in turn would not "get the nod" from (most) manufacturers of mentioned phones (for any number of implausible but also not so implausible reasons). Yet even if (economically competing) manufacturers agreed to join forces, carriers (companies providing connectivity) would balk (for pretty much the same reasons). That would leave us with two possible scenarios: Individual users take to customise the instances they use by manually creating "whitelists" (words and phrases to be used) and "blacklists" (words and phrases not to be used) — and repeat that tedious process with every new phone they use — or developers put in a lot more effort and provide extended dictionaries. As mentioned already, neither is going to happen or it would be a common custom already.

And that brings us back to square one, the paradigm shift. How is any model supposed to ever learn to work the way we expect it to, if we constantly do our best to keep pertinent information from it? Without access to a reasonable LLM (in that context, it would have to be a neural network trained with texts that real people, speaking any particular language, tend to send to other real

people, using that very same language, in everyday situations), these features have the odds of a snowball in hell to ever work effectively. And since not every user is prepared (or able) to develop their own, fully customised model, it would have to be an open source model that derives information from a huge pool of text messages in order to determine the probable applicability of its own predictions (and return that favour, of course). You may not see it from where you sit, but the personal right advocates huddling back there in their legally fortified corner are already up in arms, a button on their waistcoat reading "privacy or death".

Pretty much the same goes for applied algorithms and methodology of training. Yet in their case, it is not so much privacy concerns that motivate the mob, but competitive advantages (or rather the fear of losing them).

So yes, theoretically, we could, but whether we should or will is a different matter altogether — one that will probably take us a fair number of decades to figure out and settle. After all, it is such ventures, like, shall we call it, the "Chinese model", that give "AI" and "social transparency" a bad name. It should go without saying that transparency must never be used to oppress (or blackmail) others, but, hey, we are only humans, logic is not our strongest suit.

But what, in more practical terms, really keeps us from taking (universal) advantage of AI, is the depletion of necessary resources. Of course, AI could also help us find ways to save resources, but thus far only at the cost of spending considerably more resources than we would have to without it. So for the time being, the best we can hope for is a zero–sum game of sorts. In other words, some few will profit, a considerably larger number will not see any gain (or difference), a vast number will suffer, and if we manage to keep things from going utterly sideways, the total sum of damages will remain the same. The imminent peril brought about by AI is most certainly not the vision that "machines will try to take over" or even "enslave humans" and "abolish humanity" (in order to replace it with … what? … machinity?), but that what little natural intelligence may be left in some of us is increasingly considered superfluous (or outdated) by a majority and gradually retired.

How many People Will Lose Their Job Because of AI?

After hours of carefully polishing my crystal ball, I am prepared to confidently predict that the number of people losing their job "because of AI" will be approximately as high as the number of professional drivers, accountants, secretaries, or sales persons who lost their job after the introduction of the automobile, calculator, typewriter, or cash register.

Not trying to call any one expert out, but who came up with that ludicrous fear mongering? What's the general purpose (or your individual satisfaction) of having innocent people on edge? While the actual reason for losing their job is probably utterly irrelevant to those who suffered the termination of their contract, it is rather safe to say that they have not been (are not and will not be) "let go because of AI".

The people actually at risk of losing their job to AI (yet still not "because of AI") are executive officers, especially those who think AI might be the remedy to rectify their own bad business decisions. The point of introducing AI into the business world is to improve future business, not to cover or compensate for questionable past business.

What's with Leo and "His Little Friends"?

I earlier promised to talk about Leo, the AI–driven assistant implemented in *Brave* (and, of course, also about his contenders), and shall not fail to deliver. The other day, I happened to see an advert about Opera's new (well, I'm not even sure it is all that new, it merely was presented to me as if it was) browser model and it was impossible to not recognise how emphatically it was promoted as the "new pink", what with its alleged AI skills front and centre. I may be a cynical bastard at times — no wait, poor choice of words: I definitely am a cynical bastard at times — but, pray tell, when exactly has "our product is (at long last) just as modern as several others" become a (unique) selling point? Did anyone really think the hits returned by a search engine query have been hand–selected by a host of busy but gloomy researchers in the basement of the *Google* factory? Well, yeah, some of the results seem to have come about that way, but I guess you know what I mean … Seriously, what did you think years' worth of heated global debates concerning "Google's algorithm" have been about?

Granted, a search engine is not a browser and, also granted, not every browser conducts its own search queries. Yet hasn't it become obvious by now that everyone and their cousin is dabbling with AI? And how difficult is it to realise that everyone considers their own child a prodigy? I'm not certain about them all, but *Brave* for one has introduced "Leo" a while ago (and I cannot remember them making such a fuss about it). In fact, it took me a while to even realise that the first entry was actually just the gist of what the search results had to offer in response to my query — concocted by Leo. Tell you what: it didn't hurt me in the least. Quite to the contrary, it more often than not was just what I really needed. It wasn't before I recognised a certain phrase ("Based on the search results, …") being frequently used that I realised there was the love child of predictive and generative AI at play. For that's what Leo really is: a model that tries to predict from earlier searches (not mine but all its users') what content might be most useful (to me) and then generates a summary of its findings. And if I really feel like it, I can have a happy chat with "him", without having to register or log into my account. So yeah, Opera, welcome to the present.

What Is AI, Then?

So if you think about "AI", see to free yourself of the mental picture of robots of any kind — be that actual robots (machines that assemble products), autonomous cars, electronic receptionists, and what have you — or software that appears capable of providing all the right answers in no time, but rather consider it a concept. All those robots and software are merely individual "proof of concept". What exactly do they prove, then, you ask? Well, basically, that the strategy discussed earlier ("divide and conquer") works — if properly applied and flawlessly executed — effectively, also in fields of operation where humans are by nature rather ill equipped or even perfectly inept.

It was Elena Schirenc who, with her invaluable insight, advice, and dedication, delivered you from quite some words and complexity and me from a number of embarrassing errors. Thank you.

This essay is dedicated to the memory of Nina Pankl (2006–2024), who introduced me to artists' general difficulty of drawing hands and the AI's "too many fingers" issue. May you rest in peace and may "your hands" always be impeccable.