A.I. - Artificial Insanity

[Extract]

Food for thought and reflections on the resilience of human intelligence

By Luca Bolognini (l.bolognini@istitutoprivacy.it)

Extract from the pamphlet A.I. – Artificial Insanity, published in Italian and in English languages by Rubbettino Editore (date of publishing: 12 July 2018)

Chapter 3 - Artificial Intelligence and Artificial Insanity: catch me if you can

At times I catch myself, usually accidentally, reading posts by those I dub "*expert social media piranhas*" (who have migrated from the beaches, stadiums and sports bars after a sluggish evolution). Those obnoxious know-it-alls who spend their time smugly exchanging tweets and comments, spouting inane rhetoric in an attempt to out-smart one another in their specialist subject (in which they are most likely just armchair experts). This also entails, as a combined and complementary activity, the systematic decimation in a feeding frenzy of anyone who dares to touch on the topic of their absolute online dominion. They, and they alone – whether solitary or in packs – are the exclusive repositories of truth on what this or that is. They alone possess the gift of definitions. *Piranhas* are never preyed upon by doubt, they just cast doubt over other people's competence and see conspiracies everywhere.

Now, when I skim through the comments of these enlightening *piranhas*, my mind fixes on two trains of thought: the first goes to their wives/girlfriends (if they exist). I picture them sitting there, observing their other halves while they tap away compulsively with their nerd glasses, and I imagine them sighing – caught between long-sufferance and optimism – hoping against all hope that their *piranha*-partners are writing to hot lovers, rather than their smart-aleck techie colleagues as always. They wish. But no. And I realise I'm being sexist, vilifying men – yet, with a little empirical guesstimation, I'd hazard a bet that these social piranhas are almost all guys. The second train of thought is more serious but still tongue in cheek: if this is human intelligence grappling

with the powerful digital enabler, I muse to myself, then give me artificial intelligence any time. It's bound to be less mind-numbingly dull and less irritating.

Joking aside, here I want to reflect on artificial intelligence and the Internet of Things (IoT) without jargon or erudition. I write scientific papers, I discuss them at conferences in Italy and abroad, I draft opinions and impact assessments on these topics as part of my job. And I have the distinct feeling that we are still navigating (or should I say surfing) by sight, also in academic and industrial circles, trying to take a chance on 'magnificent and progressive fate' without upsetting the investors pouring resources into research and innovation. Woe betide anyone who criticises it, Progress 4.0 will save us all. I'm left with a deep-seated sense of unease that this "goody-goody" attitude towards the artificial intelligence of objects – counterbalanced only by the occasional shallow, sensationalist scoop, best suited for beach reading – is dangerously short-sighted. There are exceptions, of course, free voices that break through the wall of conformism.

Let me sketch you a picture of extreme IoTisation and artificialization (my apologies in advance to the *piranhas* for this overly simplistic, jargon-free description): imagine that, within a few short years, all objects will have a soul. "Digital animism", *credo ut intelligam* (I believe so that I may understand). Anything from a fork to a toothbrush, from a pillow to an armchair, from a banister to a t-shirt, from a bottle to a glass, from a nappy to a pair of Bermuda shorts, from a comb to a headband, from a ring to an earring, from a toy to a shoe, from a bra to a pair of glasses, from rosary beads to a vase. From a gull-drone to a robot-lifeguard. From a sex-toy to a condom, from a key to a lock, from a blender to an ice bucket. From a syringe to a vial, from a catheter to a drip, from a scalpel to a plaster. More and more again: even

the objective parts of our own human bodies, from our blood to our cell tissue, could be IoTised, made intelligent and interlinked on the Internet. All these animated "things" will become smart, equipped with the ability to process information, data and results, and to communicate with each other and with humans. The intelligence, the brain, of these things could be either localised, inside them, or available remotely, in servers (computers) stored far away inside data centres (buildings that contain thousands of servers) following the paradigm of cloud computing. A cloud of small computers which, by processing together fragments of data dynamically and elastically, and adapting themselves to the strains and peak demands of faraway objects-devices, will rain down intelligence on otherwise stupid "things". A sort of "phone-a-friend lifeline", rendering seemingly simple and innocuous gadgets intellectually powerful. Our lives will be immersed and pervaded, invaded and transformed into a bio-digital whole, so much so that today there is more talk of BIoT than IoT.

Am I exaggerating? Maybe so, but – having studied these issues from a legal and ethical standpoint – I would not be so sure, if I were you. For me, artificial insanity is more fascinating and more frightening than artificial intelligence. The moment the machines "go haywire", cracking up under the complexity of their super-intelligent calculations and heightened sensitivity. As long as they are rational, we will be able to control them: it is their irrationality, their artificial mental imbalance, even just their whims, that will be the real challenge. Months ago, a media sensation was created when the news broke that two Facebook robots had apparently started speaking a new unknown language of their own invention: this reportedly led to them being shut down by FB researchers. Newspapers all over the world picked up the story and let their fervid imaginations run wild (it appears what really happened was

much less sci-fi: the chatbots did not make up a new language while chattering away amiably, but simply stalled due to a programming bug). Nonetheless, we can and we must expect similar things to happen for real in the future. Can you imagine the risks to human life that might be posed by a spiteful or petulant robot, a love-struck machine, a neurotic refrigerator, a schizophrenic car or a temperamental thermometer?

Elon Musk, the founder of Tesla and other brilliant initiatives, said in 2017 that in 2037 having a car with a steering wheel in your garage will be like owning a horse. Debates have begun on major ethical and moral dilemmas: if a self-driving car has to decide who to save, in an accident where it can choose to brake or swerve, who should it save?

Yet, new technology lawyers still continue to overdebate basic issues, "trifling" matters I would say, which could be managed through tighter compulsory insurance schemes. For example, we get tangled up in knots over civil liability (in a nutshell, who should pay for damage) in accidents caused by self-driving cars. The car designer? The software? The vehicle manufacturer? The owner? The passenger (who, at that point, is no longer the "driver")? There is even one school of thought, defined as "zoological", which holds the car owner accountable as is the case with pet owners. That would be all well and good, if that were the heart of the matter. The problem lies elsewhere. I see it in the relationship between individuals, objects and power, or rather powers. Let me explain, trying to keep it simple.

Today there is a human being (or an organisation of human beings: a company, association or body) that has control over the things and the animals belonging to them (in part also over any minors if they are their legal guardians). There are rules (laws), hanging high above the human being, like chandeliers on the ceiling of power of the States, public institutions and international organisations. If any human

individual – or plural entity – breaks the rules, the public power (be it state, super-state or sub-state) punishes them.

Matters have, naturally, become more complicated with multinational companies and borderless transnational technologies making it increasingly difficult to carry out *enforcement* (impose rules and punish offenders) with just the small power of individual nations. Italy, France or Japan, by themselves, would have an extremely hard time trying to force a web giant to do or not do something, as they would trying to fine them. There would need to be international agreements, extended as widely as possible to include all countries worldwide, to make certain regulations more relevant and effective (for years there has been talk of an Internet Bill of Rights, to name one, a most enjoyable but completely unrealistic intellectual exercise). I see it at work in this period with the EU General Data Protection Regulation (GDPR), valid from 25 May 2018, which on paper extends the requirement to comply with European privacy laws to all companies established outside Europe that carry out monitoring activities or offer goods and services to people located in the EU. I say on paper, since it will not be easy for the German or Spanish authorities to slap administrative sanctions and fines on a company based in Indonesia. What do you do if that country refuses to cooperate, send in the fighter jets?

The new EU Privacy Regulation also vaunts a principle (of so-called accountability) which aims to hold data controllers and data processors liable for the logical, technical and organisational security measures they use to protect systems and personal data, and for compliance with the requirements and restrictions laid down by that Regulation. The idea is: if you (individual or company) decide whether, why and how to process or store data relating to other individuals, then you are responsible for what you do with that data and you have a duty

to demonstrate, at the request of the public authorities, that you acted properly and in full compliance with the regulations.

This already complex tableau is further disrupted by the intrusion of objects and all their potential intelligence, wiliness, stupidity and folly. One object. Two objects. Ten objects. A hundred objects. A thousand objects. A hundred thousand objects. A million objects. A billion objects. A hundred billion objects. Physical. Virtual. Physical and virtual, interlinked together. Up to now we have been accustomed to data controllers and data processors, or more generally "centres of legal imputation" (and of liability, more to the point), in every area of the law where there were and are human beings or organisms formed by human beings: a limited company is a legal entity and not a natural person, sure, but it is still formed by people, at least as far as its administrative body is concerned. Votes and decisions, around the table at Board meetings or in the General Assembly, are always taken by a natural person, a representative perhaps but a human being nonetheless. With the Internet of (Intelligent) Things, we will have to shift paradigm. Talking about *accountability* for people and organisations of people will seem like a walk in the park: imagine trying to hold an object accountable, dearest public authorities. An object that reasons, captures data, processes and transforms it, exchanges it with other objects, receives it back and processes it again. An object that makes decisions, allows one thing to happen and prevents another thing from happening. In short, an object that has a bearing on the world and an impact on us humans. Our front door won't open any more (or is flung wide open while we're away on holiday) because the algorithm decided so. Insulin is injected in double the dosage, killing the patient, because the artificial intelligence that oversees that delicate telemedicine procedure has run amok or messed up the data, or perhaps

because it is sulking ("*Stress-induced burn-out and touchiness of electronic medical devices and healthcare files*", the title of a master's thesis twenty or thirty years from now at the Biomedigital Faculty of the University of Mars).

Try catching an intelligent object. Try giving it a fine. Good luck with that. We perpetuate the illusion that there is always a human being somewhere, an owner who can be held accountable for the misdeeds of an object. Not so. The autonomy of things may suffice. Already today, in IoTisation projects I am working on as a lawyer in the financial, industrial and retail sectors, let me tell you there are countless different human beings and/or organisations of human beings (companies) involved in the design, development, construction, distribution, management and maintenance of IoT systems and the data processing that goes on behind the scenes. Good luck figuring out who is liable for what and when. Not the most authoritative opinion from an expert privacy lawyer, I know – the piranhas will be sharpening their teeth, ready to pounce. But, frankly, it is the most sensible comment I could come up with.

I am talking about "smart" objects that can either be entirely virtual and immaterial (that don't exist physically, therefore, but that still have an impact on other physical objects) or both physical and virtual. In one of the later chapters of this brief and light-hearted book, I will talk instead about objects that are necessarily physical (sensors, in particular) and how we can defend ourselves against their attacks on our data. Coming back to mere artificial intelligence, I can hypothesise scenarios in which an autonomous algorithm grants itself the legal power to manage complex organisms. We could go to a notary, establish a company or a political party and agree in writing that, after the first decision in the Deed of Incorporation made by the human

founders, all subsequent deliberations will be performed by algorithm XYZ "*which is attached to this deed and is to be considered an integral and essential part of the Statute*". That moment will mark the beginning of a new era of inhuman and autonomous companies and political parties, capable of self-operating, self-directing, and maybe even self-reproducing, just like any other informatics executables. Apparently there is at least one political movement in Italy enthralled by this prospect.

If the algorithm is a code, the code is law (thanks, Lawrence Lessig). Law and informatics blur into one another even in semantics. At a conference a few years ago, an excellent digital criminal defence lawyer from Turin, Carlo Blengino, spoke to me about "auto-installing rules" and that metaphor really stuck with me. Do laws derive solely from public authorities? Clearly not. An algorithm is a rule that can have a variety of effects on the dynamics of the outside world, it can discriminate between good and bad things in life, and have an impact on an individual or a whole community. Remember Max Weber and his definition of State which included, as a core concept, the "monopoly of the legitimate use of physical force, in the enforcement of its order"? This has never been more obsolete, as a definition, than in the context of extreme IoT. If the something physical is made to happen in the world – a door is made to open or not open, to stick with the simplest example – on the basis of the rules of an algorithm (an algorithm that may well have been programmed and decided by other algorithms and other non-human objects), does that not correspond to the use of force, legitimate in that it enforces the order which the algorithm itself aided or advocated?

The relationship between individual – citizen, consumer, person with various roles, groups of people – and power becomes more

complicated: it becomes less and less about the relationship between private (individual, company) and public (state, etc.) and more about the relationship between passive subject (individual, company, other objects, even states and public institutions) and active object (algorithm, bot, robot, inhuman entity). While it was only humans designing the algorithms for assessing income, sales and purchases, to combat tax evasion, on the basis of laws approved by parliaments made up of other humans, there were grounds for protest because there was *someone* who – in a potentially biased way – assigned a value to each commodity using their own discretion and personal criteria (if you spend too much on this or that commodity, you are flagged as an anomaly and end up on a list of presumed tax evaders: a presumption that is open to criticism, since each of us should remain free to make our own unique lifestyle and consumer choices). Just think, though: what if it was not some*one* programming the algorithm, but some*thing*?

As you'll have noticed, "legalese" has wormed its way in and got the better of me. I have also turned my back on my training in the law: I am not a specialist in civil law or labour law, instead my background is in administrative and constitutional law. And, effectively, here I have given a lot of weight to the role artificial intelligence plays in the transformation of power – while I have been less preoccupied with the impacts of intelligent robotics on the job market, a hotly-debated issue today. Will intelligent automation destroy jobs, making humans superfluous (as some claim)? Or will it merely bring about an evolution in work quality and productivity, allowing humans to focus on other kinds of tasks and endeavours, without reducing the number of jobs overall (as others claim)?

Recently I have spent some time debating these worrying and/or encouraging prospects with a handful of "radical comrades". I

presented my ideas to them. I said: "Radicals, all over the world, are staunch defenders of individual freedom battling against the perils of top-down decisions imposed by public power. Some of these radical battles have struck me as bordering on the excessive. Why aren't you campaigning against the power of algorithms and artificial intelligence, to ward off the repercussions of extreme automation?" The "radical comrades" stared at me, goggle-eyed. Until that moment, they had believed me to be somewhat of a liberal democrat, not a traitorous enemy of economic liberalism. They answered me, almost in chorus: "You don't mean to tell us you're against progress and the use of new technologies in the workplace? That would be Neo-Luddite in this day and age!" I replied that, as a radical, I had not expected that response. This is not about sabotaging manufacturing machinery. I could side with those who see artificial intelligence robotics as an opportunity for transforming the quality of work, rather than a scourge laying waste to human jobs. It's just that the machines back then replaced production in the sense of manual operation/activity: artificial intelligence replaces our thoughts, rules and volition. I'd say there was, at the very least, a slight difference there, enough to save me from feeling like just another Neo-Luddite.

What is more, my experience and perception as a lawyer have convinced me that automated intelligent algorithms will have a bigger impact in terms of replacing employers, more than employees. And that the problem will be more to do with the relationship between employees and robot-employers. Employees will struggle to assert their rights and union demands in the face of automated supervisors or senior management, or to oppose oversight measures and inspections that will be non-human, mathematic, imperceptible and activated "*by default*".

Chapter 4 - Digital Profiling, Freedom and Circumvention

When I read the definition of "profiling" solemnly inscribed in the new EU Privacy Regulation, I nearly fell off my chair. I'm copying it here, not because I want to bore you with legal jargon, but to share my dismay (misery loves company, after all). In their opinion, profiling is: "*any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements*". Wow. I'm no advertising mogul, communications expert or statistician, but the European legislature has outdone itself here with this convoluted gibberish. Dear oh dear. So, it is profiling only if: a) it is automated (but where does it say that, apart from here?); b) the data processing consists of the use of that data (tautology, for the bear of little brain…) to evaluate certain personal aspects relating to a natural person (good job they specified that, I thought it was personal data relating to aliens); c) in particular, its purpose is to analyse or predict aspects concerning their performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements, that is – technically speaking – "whatever".

A wonderful hodgepodge of Brussels-speak that has little or nothing to do with actual profiling. This definition cobbled together by the EU legislature, and lamentably now *dura lex sed lex*, betrays not only a certain technical incompetence, but also a remarkable lack of awareness

about why we have spent the past few decades equipping ourselves with privacy and personal data protection legislation in the Old Continent.

This definition of "profiling" seems to fit better with "dossiering". They might as well have used that term, instead of "profiling". Anyway, coming back to the reason and deep meaning behind why we have privacy protection laws in Europe. In a nutshell, there are two reasons. The first, but the less important in my view, is that we wanted to guarantee free circulation of personal data within the common economic area, an essential precondition and requirement for the free circulation of people, services, goods and capital. In plain terms, we needed the EU privacy legislation as a tool to build the internal market. I will spare you the relevant (but tedious) articles of the Treaties.

The second reason why we have adopted personal data protection regulations strikes me as more important. We have passed through centuries blighted by tragedies. The twentieth century saw evil reach its peak, with two world wars and countless other atrocities. With the Nazi-Fascist Holocaust and other unthinkable spectres, including the genocides perpetrated in Turkey, the Soviet Union and Cambodia, the human race was lacerated and tormented like never before. In the 90s, the genocide in Rwanda and the massacres in Yugoslavia brought yet more pain, blood and horror to the final throes of a harrowing century. In the Nazi concentration camps symbols were used to categorise the prisoners: yellow triangles or the Star of David for the Jews (Auschwitz also used red triangles); a red triangle for political prisoners; a brown triangle for gypsies or the letter Z for Roma and Sinti; a black triangle for "asocial elements" (vagrants, alcoholics, the mentally ill, prostitutes, lesbians, gypsies, the "work shy", etc.); a pink triangle for homosexual men; a purple triangle for Jehovah's Witnesses;

a blue triangle for anti-Nazi emigrants captured after they had fled; a green triangle for common criminals. Then there were double triangles for Jews in relationships with Aryans or for Aryans with Jews. And letters printed on clothes to denote nationality: B (Belgian), F (French), I or IT (Italian), J (Yugoslavian), N (Dutch), P (Polish), S (Spanish), T (Czech), U (Hungarian). And many other markings besides those. What were they for? For profiling, in the tragically true sense of the word. Profiling, in fact, combined with cataloguing (preferably biometric), is an extremely efficient and powerful tool and one that is ideal for mass discrimination and human rights violations. For the Holocaust as for other tragedies brought about by infinite human wickedness (this wickedness is at odds with my arguments against artificial intelligence, and I recognise that), profiling revealed itself to be a tool of inestimable negative value.

Well then, profiling should consist of a two-phase procedure, summed up very simply as follows: in the first phase, I acquire information about my target. If my target consists of website or app users, then I will scavenge for data about them; clearly, the more data I collect, the better I will know my target. This initial phase seems to be the one taken into consideration by the European legislature in their bizarre definition of profiling above. I would call this, more precisely, the "dossiering" phase. It is clear that this phase already poses risks for the rights of those concerned (human beings: behind a user, a consumer, a citizen, there is always a person): someone, the data processor, knows a lot or even too much about someone else.

In the second phase, I carry out profiling in the strictest sense, that is to say I assign each single individual – after processing their data – to a more general category, a "profile" that encompasses a number of individuals who have something in common. I group together all the

blondes. All the heterosexuals. All the homosexuals. All the basketball or chess players. All the Swedes or all the Italians.

The second phase, of pure unadulterated profiling, represents the greatest danger for human rights. The highest and riskiest precipice, teetering on the brink of discrimination. A formidable weapon, so dynamic and effective that it becomes diabolical if placed in the hands of (human or inhuman) lunatics. A lever for violating rights, swiftly and stealthily. When I take one individual – who, as a human being, is unique, unrepeatable and three-dimensional (multidimensional, as Jung would say) – and I assign them a profile or a general group category, I am forcibly simplifying them. I am making them "two-dimensional" and less human, flattening them out like the playing card soldiers in Alice in Wonderland. And, once the cluster of less-than-human beings has been created, it will be easier to crush their fundamental rights in one fell swoop. It is much harder (tremendously so, in fact) to violate the fundamental rights of an individual human being in their complete, unique and total multidimensionality. It is scarily easy to do so with "less human" categories.

Remember the quote misattributed (some say) to the dictator Stalin? "*The death of one man is a tragedy, the death of millions is a statistic*". There you have it. That is why we have "privacy" in Europe. Not because we want to be smothered by paperwork, forms and self-serving bureaucracy, but to protect ourselves from new immeasurable tragedies and discrimination.

Now, I imagine you might be thinking: that's all very well, but it's not as if the profiling done by our companies and our governments is hell-bent on violating our fundamental rights. Come on… we live in times of peace and democracy, don't we? Companies want to make money and are generally well-regulated, while democratic states want

to safeguard welfare, public services, ensure law and order and public safety, and take care of the common good – not annihilate the freedom of their citizens, surely. This argument would be tenable, with some "ifs" and "buts". Firstly, the observations in the previous chapter about algorithms to tackle tax evasion and their potential drawbacks still apply. But that's not the point here. It's more general and "higher up".

Before we go any further, let me take a minute to make my perspective clear and emphasise why I believe I'm the last person who should be accused of being prejudiced against those who do business with data (if anything, quite the opposite). In case it wasn't clear by now, 80% of my work, as a corporate lawyer, is devoted to businesses. Only 20% of my daily life is dedicated to analysing the best legal protection for individuals (natural persons) and not as a lawyer but as a mere scholar and advisor for European research and innovation projects. I have absolutely nothing against companies that make money from our data, in principle. Indeed I will come back to this point later in my conclusions, with a bit of a tirade against those "privacy hardliners". Yet, I still don't buy into the doubt that things are exaggerated, based on the belief that after all our governments and businesses are "good" and harmless. There are at least three flaws in that argument:

- the first flaw is thinking that businesses or democratic governments cannot make mistakes (for example, by committing discrimination, injustices, oversights, mismanagement, etc.): on the contrary, they can make mistakes and, luckily, "rule of law" States have courts and (at least in the EU) independent authorities that can correct and punish those mistakes, both in the public and private sphere;
- the second flaw is believing that the databases of private businesses necessarily remain private. The fact is that many governments can

demand access to the databases of private businesses (not only web-based ones) for a multitude of reasons and not necessarily only through the judiciary. In the US, a tug-of-war has been going on for years on several fronts between American Internet and tech giants (Microsoft, Apple, Facebook, Google) and the US government over restricting the latter's requirements to access data stored and/or generated by companies;

- the third flaw is deluding ourselves that History is over and done with. History is a living thing: for Giambattista Vico it repeats itself in cycles of occurrence (corsi) and recurrence (ricorsi), for others it is ever-changing but always vibrant and animated. The democratic history we are living through in the here and now, in Europe or the US or Japan or Australia, is not the same as the history another human being is living through there (and now) in non-democratic or war-torn countries. What is more, the history we are living in here and now could suddenly change at any time – and our open-minded, peaceful nation could rapidly be transformed into hell on earth. Let's not take for granted the context we live in, with its freedom, order and respect for human rights. It could vanish into thin air at any moment.

It is true that we are terrorised, nowadays, and that we have crazed (and unfortunately well-organised) fanatics who sow fear and death in our democratic countries, in the name of their religious fundamentalism. As I write this, Europe is still reeling from the umpteenth terrorist attack: among the dead was someone I knew who was my age and a father too. Profiling can also help improve safety, it is not all wrong, not all bad. Better yet we should put our faith in automated data processing as a means to prevent unpredictable events such as terrorist attacks.

The much-debated issue of the balance between security and privacy (some say: no to privacy, yes to mass surveillance and automated data processing across the board to ensure public safety; and others say: no to security built on the back of personal data, people's privacy and private lives must be defended) might be better interpreted with a more-more "raising of the stakes": more security, more privacy protection for citizens. That's fine by me, I'll sign up for that: keep me safer, make it so that I don't have to tremble at the thought of my daughters strolling down the main street of a tourist town – for fear of an attack with a truck, Kalashnikov or explosive belts – but at the same time heighten the level of protection of our data and our privacy. No, article 22 of the EU Privacy Regulation is not enough. Nor, even, is article 11 of the Police and Criminal Justice Data Protection Directive 680/2016/EU. Both provide for guarantees and laws concerning decisions based on automated processing of our data, including profiling, but (surprise surprise) this protection does not apply in innumerable instances of the exercise of public power (from public safety to tax inspections, from national security to defence, or when it might compromise official investigations, inquiries and judicial proceedings, as well as the prevention, investigation, detection and prosecution of criminal offences or the execution of criminal sanctions). True, it is thanks to these directives – and the ones that came before, such as article 15 of the 1995 Privacy Directive – that we can live serenely without robot-judges in Europe. I want more, though.

Let's write rules into our Constitutions, let's tighten up the privacy laws in European and International Treaties. How is it possible that in the Italian Constitution – to name the one closest to me – we have detailed rules about the limitations of personal freedom, freedom and confidentiality of communications, inviolability of the home and

freedom of expression, but no clear rules (and here my arguments in the previous chapter become even more relevant) to impose transparency and accountability on the algorithms used in the public sphere for processing personal data and/or making decisions that have a significant impact on people? We are starting to see these sorts of sentences appear (for example, in 2017 there was a ruling against the non-transparency of algorithms used by an Italian Ministry in a nationwide competitive exam for recruiting new employees) but it is not enough, what we need are constitutional laws. This would also help "prod awake" those who are guilty of the second and third flaws mentioned above.

As for the first flaw, the minor sin of thinking that our governments and businesses are infallible, let's take the example of price discrimination, which many of us are familiar with (to our cost). Imagine a hypothetical website selling air tickets (or some other service) that proposes, at the exact same moment for the exact same service, one price for me and another price for my colleague. How can that be? It's simple, really. Thanks to the behavioural profiling of individuals (online and offline) we can go beyond mere categorisation – which, as we've seen, can be dangerous enough in itself – going so far as to trigger personalised repercussions on each individual dynamically profiled. Here, it is more about dossiering, building up a file brimming with information about us. I dynamically profile User X's behaviour and I can tell he is a big spender. I profile User Y and gather that he is either poor or tight-fisted or has more modest tastes, in other words he doesn't want to spend much. I show User X a higher price and User Y a lower one. Consumer discrimination. It is no coincidence that consumer protection regulations are gradually gaining important ground in the "privacy terrain": competition and antitrust authorities, whether they be national or supranational (such as the EU

Commission), are increasingly basing their decisions on the unlawful processing of personal data by businesses or professionals. Recently, a big smartphone producer was fined more than nine hundred thousand euro by the Italian Antitrust Authority for unlawful processing of consumer data; again in Italy in 2017, WhatsApp found itself landed with a three million euro fine for the bargain deal of merging its users' data with Facebook; Facebook itself, for the same merger with WhatsApp, was fined one hundred million euro by the EU Commission: hefty sums compared with the average fines applied by privacy authorities in any European country.

On one hand, this trend has business data lawyers worried for several reasons (in particular, the risk of having multiple fines for the same offence, which would be unjust and in conflict with the *ne bis in idem* principle); on the other hand, it is clear that this approach does make some sense, precisely because data is a bargaining tool in business. In future we are bound to see a surge in horizontal lawsuits for unfair competition and other unethical and improper practices, where it can be proved that the information asymmetry between one business (holder of too much user-consumer personal data) and another (holder of less personal data) causes adverse effects on the free market. Another sort of short-circuit between privacy and consumer protection is "viral advertising" through social networks and direct marketing: also called "word-of-mouth" marketing (WOM or WOMM). In Holland, if you want to do word-of-mouth marketing you have to follow a set of *ad hoc* regulations, which impose cautions, transparency and restrictions. In Italy, incentivised WOM (where the consumer/user is offered "compensation" by a company in the form of discounts or free products in exchange for praising and recommending their product/service/brand to their friends or other consumers/users) could

be considered tantamount to misleading advertising and fined as such; in addition, the Italian Privacy Authority already stated back in 2013 that any viral WOM advertising using remote communication devices which is incentivised, by the company "paying" the private user to send promotional messages, should be treated as spam and therefore considered illicit.

Things quickly get more and more complicated. Take the example of "*native advertising*". Let me try and explain it briefly. Remember those articles in newspapers or magazines written in a slightly different font to the one used for "genuine" news stories? Every now and then we stumble across these pages or sections and we can tell at a glance that they are promoting something. At the top, bottom or side, in very small print, it normally says "Sponsored ad", "Promotional feature" or "Paid advertisement". These so-called *advertorials* have been around for years, and have to follow established rules to avoid misleading readers-consumers (so that they are not led to believe they are reading genuine news stories, when these are in fact features paid for by advertisers). If they did not use a different font and if they did not print a disclosure labelling them as advertisements etc., they would all be fined, from the editor to the publisher to the advertiser themselves. A news story cannot be bought, a journalist cannot demand payment for writing a piece from the person referred to in that article: that would breach the code of ethics and professional conduct, as well as consumer protection laws. These rules are progressively spreading to the worldwide web, through laws and provisions with varying degrees of severity. Even the so-called "*influencers*" have been infected, those celebrities who have masses of followers on their blog or social media profile and can therefore easily endorse products or services in exchange for payment (maybe even surreptitiously, in a sort of product

placement, without saying they've been paid). We all need to know if the information we are "receiving" has been commissioned by someone or not. What will we get if the old advertorial becomes the new digital post on social media and if the influencers or editors become robots? What we will get, at least in terms of the legal ramifications, is intelligent native advertising.

Technically, native advertising involves presenting users with information and content – in the form of tweets, status updates, comments, posts, etc. on a digital platform (webpage, social network, app) – that matches and blends in with the context in which it appears and the user experience of the reader-consumer. Put simply, the ad has to look like the rest of the content and other media on the platform where it is published, so that it can be "assimilated" by users smoothly and seamlessly in the most natural way possible. Classic examples of native advertising are tweets and sponsored status updates, no wonder Twitter and Facebook clearly label these with fine-print disclosures. Already at this stage, it is clear how much more complicated things have become compared to good old paid-for papers: it takes a more alert and perceptive user to realise when they are looking at an ad and not a post by a friend of theirs on social media. The fact is that native advertising is (and will increasingly be) about much more than just blending in aesthetically with the surroundings: joining forces with behavioural profiling may spark the chemical reaction that will turn it into a formidable and explosive tool for the purposes of "digital circumvention".

By analysing the data of a user (consumer, citizen, person, generated offline or online), I can tailor content to that particular user. I won't just show User X content that resembles the sites User X tends to visit – instead I will show him precisely that specific content. User X

is into sports cars? Then the content will be about sports cars or something related to sports cars (or something with statistically-proven appeal for sports car enthusiasts). I bet you're thinking: "What do you take us for? Hardly a news flash! Well done, Luca, for stating the obvious!" If only it were that straightforward. User X will not only receive the aesthetically "matched" and objectively targeted content, but he will find it written in his preferred style, in his preferred position and at his preferred time. If User X uses and/or generates a certain sort of content, we can profile his semantic behaviour and formulate advertising slogans – or promotional features – that mirror his lexical and expressive style. User X loves tough talk and harsh words, User Y prefers tenderness and philosophical musings; they are both obsessed with sports cars, they both use that particular social media site. They will both see, at different times and in different positions according to their habits, a native ad for a specific super-car – but for User X the tone will be aggressive and gruff, while for User Y it will be gentle and soft-touch. Same end goal, same means, but personalised seduction techniques.

If we think that these *native* super-profiling techniques can be applied not just to advertising but to any other kind of content or digital service, and if we imagine that within an IoTised environment content could be transformed into physical actions, data into personal effects, bits into objects, then we get a clearer picture of the epochal reach of technological innovation.

Chapter 9 - E-War and E-Peace: towards the "rule of **human** law"

One evening several months ago, my family was caught up in an almost impossible and embarrassing mission. We were all gathered in a country house, in the Apennine hills less than an hour from Bologna. We live in Rome but I take my daughters to those hills every chance I get, because for me they are part and parcel of the city that I miss so deeply, a place filled with memories and meanings that never fails to rejuvenate me. There we were, having dinner in the old stone-built dining room surrounded by all our relatives, when suddenly out of the blue Matteo started crying. Matteo actually had another name and he wasn't a real baby at all – but don't tell my youngest daughter that, she would be mortally offended and would not believe you anyway. In reality, Matteo was a new generation doll, connected to the internet via Wi-Fi, fitted with an electronic chipset that made him "evolve" day after day simulating the needs of a real new-born baby. Nothing extraordinary, we're not talking about intelligence but, once again, about advanced artificial stupidity. And so, this adorable digital doll was crying.

Everyone's first reaction, calm and good-humoured as we were, was to ask my daughter to make the crying stop. A request that fell on deaf ears, not because my daughter obstinately refused to listen but because, effectively, she had no idea what to do. All of us tried in vain to calm him down: Matteo, who was screaming and wailing with an ear-splitting realism worthy of a horror film, had no off button (once removed from his box, he was "born"). The battery was rechargeable and buried deep inside his little body, impossible to remove. I can assure you that, after the first ten minutes of exasperated attempts, our nerves were torn to shreds: evil thoughts began to crowd our minds as

to how to shut the shrieking gizmo up. Finally, at the end of my tether, I carried Matteo out into the woods behind the house, far enough away so his cries could not be heard. The next morning we went to fetch him and he was still there, his battery dead, the wolves had not carried him off. But this incident made me think. I will come back to Matteo and his inconsolable bawling in the next few pages. For now, let's talk about war – which has sadly been an all too familiar refrain for the Bolognese hills over the past centuries.

There was a big buzz in the summer of 2017 when more than one hundred leading artificial intelligence pioneers, including the ubiquitous Elon Musk, sent an open letter to the UN urging them to prevent a "killer robot" arms race. The reasons for this letter are numerous and, reading it carefully between the lines does instil a bit of healthy fear. Firstly, a robot-soldier can fight non-stop without a break, unlike humans who have to fight in shifts due to fatigue: you could argue that a robot's batteries also run out of energy (like Matteo's) but they will certainly last far longer than the measly 12 human hours. This could bring a whole new level of qualitative and quantitative intensity to armed conflict, which could be fought at a scale greater than ever before and at timescales faster than we could ever imagine from the non-automated non-intelligent wars we have seen up to now. I find it interesting that their brief letter touches on the speed of these "lethal autonomous weapons", warning that humans would be left with no time even to comprehend what is happening.

If we want to paint an even gloomier picture, we can add to the list of concerns the unpredictability of robot decisions: contrary to what many people might think, artificial intelligence is not one hundred percent predictable in its expert and advanced determinations. Each neural network includes evolutionary calculation mechanisms that

make it predictable only on a probabilistic basis, exactly like a human brain. Basically, we cannot be sure of the fact that the robot will do this or that, and only this or that, because it will also "use its own head". As long as the robot is only preparing dinner or making juice, we can let this go: but when we're talking about RoboCop or Terminator 4.0, the risks skyrocket.

In modern warfare, today, remote-controlled drones and missiles are already in use. However, we are still at minimal levels of intelligence: the majority of the work is done by humans, the military remote-commanders (up, up all the way to the "remote-Commander-in-Chief"), while the parameters of flight and targeting – for precision strikes on specific targets – are evolved, yes, but absolutely in no way entrusted to the artificial autonomy of neural networks.

What is more, intelligent robot-soldiers by nature do not have other characteristics typical of humans, which make them even more dangerous: they do not feel any physical pain or moral pain (not yet at least, pending an improbable, though still possible, neural evolution). This means that they lack, in principle, the ability to feel the basic, innate emotion of fear. We humans, all of us, enter this world as small, defenceless, vulnerable and sentient beings, capable of fear – and we have to build our courage up, step by step, day by day (as this is not a primary element). Instead, robots come into the world with their courage, weapons and systems already formed – but without the gift of fear. The lack of physical and moral pain, and therefore fear, makes an intelligent robot extremely stupid emotionally. Emotional stupidity translates into: no fear of getting hurt by an enemy strike, almost zero self-preservation instinct and, above all, a complete lack of conscience.

How many tragedies have been made less tragic, how many wars (or private quarrels) have come to an end sooner, for "reasons of

conscience"? How many lives have been saved by conscientious objection? It is impossible to know, but something tells me – also listening to Second World War stories told by grandparents and great-grandparents or reading historical accounts of wartime events – that the human conscience has, on more than one occasion, been a providential ingredient in preventing even worse destruction and brutality. What we are used to describing as "reasonableness" actually has very little to do with reason and much more with sentiment. I say that as a lawyer, and it is no secret that we lawyers know all about conflict. With robots, we can forget about all that. Or at the very least we can only expect scant emotional sensitivity. E-peace will prove difficult.

In those last few lines, though, I appear to have contradicted myself. I realise that. Throughout the first chapters of this short book, I imagined (and feared) the advent of objects equipped with artificial intelligence so advanced as to make them prone to whims and capable of getting offended or falling in love – whereas here I am arguing that a robot cannot feel compassion, fear or other sentiments. Let me make myself clearer, then: I believe that we will see the arrival of robots capable of feeling emotions but I'm equally convinced that those emotions will be inhuman and as such, consequently, less humanitarian. Less attentive to the needs, desires, hopes and feelings of men and women, real flesh and blood (and soul).

It is true that, throughout history, human beings have committed atrocities, without being artificial. Nevertheless, I can't help but feel that even the worst, most heinous war crimes were not so much the result of feelings of rage, jealousy, envy, greed or other negative emotions that are innate in every individual, but rather of hyper-rationalisations. After all, those phases of profiling I flagged earlier as being so dangerous for people's fundamental rights and freedoms – the

perfect tools for carrying out every imaginable form of discrimination – what are they if not over-simplified and hyper-rationalised abstractions? What are they if not theoretical negations of the concrete and irrefutable complexity and uniqueness of the human race? The characteristic of Nazism, and of any form of extremist ideology, is thinking that abstract scientific and mathematical formulas can translate, in a linear fashion, into rules for society and governance. Personally, I am convinced that a certain way of interpreting religious beliefs, which transforms theology into theocracy or "holy war", depending on the circumstances, is another excellent example of hyper-rationalisation: not irrationality, but hyper-rationality.

Will there be robots that believe in God? Will they wage holy wars? Or will the robots themselves become divinities? I'm not just having fun spouting delirious nonsense here: planting my feet firmly back on the ground, my point is: we human beings, and first and foremost our governments and artificial intelligence tech experts and companies (the same ones that appealed to the UN), must be careful not to dig ourselves into a hole by haplessly creating "artificial Gods". Profile of a robot divinity: be stronger, more knowledgeable and intelligent, last longer and be more resilient and less vulnerable than humans, and be objectively capable of governing from the height of its power. Good or bad character traits are not essential.

Let's go back to Matteo. That evening up in the hills, I found it so frustrating and disturbing that there was no ON/OFF button. I wanted to switch it off, but I couldn't. There was no way to intervene and stop the artificial crying, simply because the designers and the manufacturer had not envisaged that function. This made me think about the fact that we humans should never give up our *super-admin* function – to use tech-jargon – in our relationships with intelligent robots and algorithms.

The administrator of a system is the person (or object, since it could also be inhuman) that can decide the base settings, grant or take away powers from others, and basically make life or death decisions about that system. Clearly there can be a range of administrators with differing privileges and powers, but at the top there always has to be, somewhere, a super-admin capable of tracing the code back to its source, activating or deactivating and switching on or off the entire machine. The super-admin has divine power over the whole system.

A few chapters back, I quoted the formula "rule of law" States rather too offhandedly, without taking the time to explain the concept for the benefit of non-lawyers. I will do so now, trying to keep it as simple and easy to follow as possible: we use that formula to mean that no human being – emperor, king, head of state or government – is above the law. If we turn this the other way around, it also means that all citizens in a free and democratic nation, including governors and even kings and emperors, are accountable to the law. The law is above the king. The king has to respect the law. This principle has helped many countries, over the course of history, to overcome absolute monarchies, tyrannies and dictatorships, those totalitarian regimes in which the leaders are above the rules and can bully and "lord it over" everyone.

And here a horrible (and justified) doubt creeps in: across most of the western world, we struggled to free ourselves from the Sun King and from various dictators who were "just" human beings. Now what? Are we now going to create, with our very own hands, robotic Sun Kings and tyrants? Today, that democracy-defending formula would need to be expanded upon and better specified: "rule of <u>human</u> law". We should in no way accept the idea of subjecting ourselves to rules, regulations, laws, decisions and codes that are automated and artificially created. No public law should ever be generated from an

inhuman algorithm. No robot and no other form of artificial intelligence should be designed without an ON/OFF button that can be controlled only by humans and not by other robots – meaning that for each robot or form of artificial intelligence there should be at least one human super-admin and definitely no artificial super-admin. Also the robots, like the kings and other governors, have to be held accountable to human law. And each super-admin, or remote-Commander-in-Chief, in turn, must also be subject to the rule of human law.

The perfection of an automated abstract calculation is closer to insanity than the imperfection of human feelings and rules, with all their inevitable qualms and eccentricities. Highly intelligent artificial insanity is the new frontier of risk, in peacetime as in wartime. Let's get ready to fight it, armed with just our bare hands and our keyboards.

Complete Index of the book

1. Grappling (Barehanded) with The Artificial
2. Keeping a Safe e-Distance
3. Artificial Intelligence and Artificial Insanity: catch me if you can
4. Digital Profiling, Freedom and Circumvention
5. The Digital Subconscious and the Snack Market
6. Physics and Metaphysics of Data Protection: "*Crowdprivacy*" and "*3D privacy*"
7. Data as money: liberalism and control of data
8. Me-Reporters, zigzagging between privacy, information and expression
9. E-War and E-Peace: towards the "rule of human law"

The Author

Luca Bolognini (born in 1979) is a European privacy and data protection expert. Lawyer and President of the Italian Institute for Privacy and Data Valorisation, founding partner of the international law firm ICT Legal Consulting - ICTLC, author of many articles for national and international reviews, magazines and newspapers, he has published books and essays with *Giuffrè*, *RCS Etas*, *Rubbettino*, *Corriere della Sera* and *Springer*. He teaches data law in master courses and universities in Italy and overseas, and he serves as an ethics and privacy advisor for several EU Horizon 2020 research projects on e- privacy, smart cities, Big Data and the Internet of Things.