

AI In Human Evolution: Reducing Violence

Studies of human cultural evolution have demonstrated the unique effect technology has had in shaping the history of cultural development and in influencing the form human society has taken. Considering the possibilities at which our current understanding of AI have hinted, it is both reasonable and important to consider the way this technology might affect the future evolution of society and to consider the ways this evolution might thereby be shaped. At the risk of serious intellectual overreach, it may prove worthwhile to consider one aspect in which human culture under the influence of AI is likely to require significant evolution if the future of human society is to be protected. Due to the ability of artificial intelligence to amplify the effectiveness of any given activity, it is clear that any future in which AI is extensively implemented will require humankind to come to grips with its violent tendencies, at the individual level as well as collectively. The use of artificial intelligence as a tool in guiding humanity toward a future far less characterized by violence than is the present will in fact be required if human society is to survive to enjoy the benefits this technology promises.

AI and violence: Where's the connection?

As imposing a premise as the necessity for evolving beyond violence may seem, the very technology that requires this feat offers significant support for its success as well. AI-directed automation promises material abundance predicated on less, rather than on more, enforced human labor. This impacts society in two ways: First, material abundance offers the opportunity to relax violent posturing as well as violent conflict over material well-being. Second, the freedom from labor necessary to produce material necessities, comforts, and luxuries frees individuals to cultivate and to express themselves in other ways, specifically making way for the kind of cultural contact and respect for competing views that tend to break down the ignorance we have of each other within which violence spawns and spreads and to replace it with the vibrant marketplace of ideas required to move our conception of the human forward.

There arise immediately two objections to be addressed. As to material abundance, the question arises: What about the resources required and the pollution produced? It is true that, as technology has enabled humanity to produce more, human population has grown to require more and is now at a level that threatens to overwhelm both the planet's resources and its ability to absorb the byproducts. Technological solutions to both of these aspects have already shaped further developments, as with advances in mining for mineral inputs and continuous innovation in energy production that has significantly reduced air pollution in cities in which the air had been previously rendered dangerous. AI-assisted development and execution of future technologies will undoubtedly continue this trend, with ensuing efficiencies ratcheting back familiar arguments over the unfavorable economics of "clean tech."

As regards the argument that freedom from necessary labor will promote the greater interpersonal and cultural awareness that will precipitate a virtuous cycle of peacemaking, a full answer leads to the further development of the scope of possibilities AI enables in mediating social issues. In fact, freedom from labor is acknowledged as much a problem to be solved as an opportunity to advance. Recognizing this situation as an opportunity rather than as a problem comprises the heart of the thesis as stated above and the guide to its solution that follows.

If that's the good, what's the bad?

The dual-use problem attendant upon technological advances may ultimately only be overcome though recognizing the potential for advancement inherent in the dual-stream theory of human evolution. "Dual-use" refers to the possibility of applying any given technology in either a constructive or a destructive manner, an example being the use of nitrogenous fertilizer for increasing crop yield or for improvising explosives. "Dual-stream" refers to the conception of human evolution as having both a biological and a cultural component, the cultural component clearly outpacing the biological in rapidity. It is this rapid cultural component of human evolution that must be called upon not only to outpace, but to in some degree displace elements of biological evolution if the issue of violence is to be handled before developments in AI result in either purpose-built weapons or in dual-use technologies that threaten the future of human civilization.

If AI is to play a part in intervening in humanity's cultural evolution, where is it to be directed? how is it to be applied? what will be lost in the process? and how is such a wholesale conscious invasion of sovereign human freedom to be justified? The diversity of human cultural expression and values stand forbiddingly against any claim to an engineered "solution," just as history cries out against the very use of the term in light of horrifying examples of actions undertaken under its banner. What is the cost to the advancement of the species considering the potential to dampen the constructively competitive instinct and to limit the creative space destruction has often created if violence is to be quashed at the species level? And if violence has been an indelible part of human nature since the emergence of the species, does an effort to ameliorate and finally to eliminate its expression to the extent possible not present an affront to the meaning of experience as human beings, as individuals possessed of the concept of free will as it pertains to the expression of all natural tendencies?

That sounds complicated — where to begin....

As a path forward in light of such weighty concerns, a brief survey of current efforts aimed at containing violence is in order. As any such exploration is severely complicated in view of cultural diversity, the focus here shall rest upon the prevailing values widely expressed, if not evenly enacted, in contemporary democratic Western society. It is upon this perspective that the arguments that are to follow stand. Where, then, do we put forth the effort to limit violence? In childrearing, both parental and

institutional resources are invested in influencing the development of nonviolent tendencies. Schools are charged with enforcing non-bullying policies. Every major world religion teaches peace as a primary good. In society, physical violence committed by individuals is met with legal repercussions and the most violent individuals are removed from general circulation to be treated or incarcerated. Between sovereign states, conflict is mediated by such bodies as the United Nations in an effort to avoid the escalation to violent confrontation and warfare. Militaries are maintained by nation states at enormous cost, which in ideal circumstances dissuade violent action by other international actors.

Such a brief gloss is necessarily woefully incomplete, flawed by the aforementioned uneven application of these efforts, and arguably unworkably naive as presented, though its consideration makes the point: Society invests heavily in controlling violence. The very meaning of violence has evolved to include its nonphysical aspects, as in the case of emotional abuse. This is of great importance in itself, though for the present purposes the issue of physical violence remains the focus. Since the peaceful resolution of conflict at the individual, group, and international level is central to maintaining the social order in which humanity may flourish, it is reasonable to consider this activity to be a focus of humanity going forward. How, then, might the resources made available by artificial intelligence enable the future development of the enforcement and perennially hoped-for inculcation of nonviolence in society and within the individual?

The AI among us

Artificial intelligence is grabbing headlines as well as research dollars. Headlines currently focus on the massive investment being made in China, a recent billion-dollar training initiative announced by MIT, and the self-driving automobiles that are well on their way to becoming a reality. Already in place, however, are AI-enabled algorithms that thrive in the realm of social media. Underneath the services available and apparent to personal users, the social media corporations that have come to signify the electronically-connected world function as advertisement merchants. These companies trade in personal information collected one click at a time. This enormous amount of data is transformed into ads served up to just the right person at just the right moment — at least this is what the customers are paying for. This has an enormous effect on the average individual's purchases as well as (infamously) the average individual's voting preferences.

While surely considered a necessity by the entities that avail themselves, advertisement at its root adds no fundamental value to nor can be considered a particularly valuable source of information about a product. This enormous industry exists to persuade and this persuasion has become targeted with remarkable accuracy through the use of data management techniques — the very stuff of AI. At best, this marketplace for our attention provides a platform for building the esteem in which particular brands are held. On the darker side, the influence wielded by outside governments on recent American elections through the very same means demonstrates the potential for

flagrant abuse. In addition, an argument can be made that all such advertisement represents an assault on our attention and concentration and, via the hedonic treadmill, powers the issues that come with consumer culture. It is clear that AI-powered algorithms are already very much a part of society.

AI for the public Good: Why consider violence?

With safer and more efficient transportation, more efficient and effective education, undertaking dirty and dangerous tasks, and new medical treatments all being developed as offshoots of the application of artificial intelligence to the issues facing humanity, why is targeting violence to be considered a priority? As mentioned in the general discussion of dual-use technologies above, AI is a general-purpose amplifying technology. Just as the knife, the spear, the gun, and the bomb have amplified the ability of an individual or group to render physical violence, so may AI be applied to do damage at new scales. It is conceivable that autonomous weapons will dominate advanced military forces in the near future; the ethics specific to this issue are currently hotly debated. It is also conceivable that AI-enabled software may be employed to take down an entire national electrical grid. The violence that would ensue from a successful attack of this nature could be catastrophic. Addressing these scenarios is one motivation for choosing to address violence with the potential influence of AI.

Another threat that goes hand in hand with the development of artificial intelligence is the potential for the runaway development of an AI endowed with the power to place humanity at existential risk. Research and applications have already demonstrated that AI technology reflects the biases of the data on which it is trained or upon which it is based. It may be a tall order, but a continued effort to evolve beyond violent tendencies may prove to be the best way of reducing the chance that the AI technologies that arise within the framework of human culture turn out to themselves be violent. Using the transformative potential of artificial intelligence to do so may thus turn out to be our best hope of limiting the potential of AI to destroy human civilization, a potential danger taken very seriously by numerous AI experts and individuals focused on existential risks to the human race.

How would this work, then?

As we have seen, artificial intelligence may be applied anywhere the art and science of persuasion dominates: in childrearing, in education, in commercial transactions, in interpersonal social interactions, in entertainment, in politics... virtually everywhere human beings make contact with each other. Once the issue of the influence exposure to violence has on humanity as a whole is taken to be a species-wide issue going forward, AI can serve as the means of monitoring and filtering those effects. An obvious starting point might arise with child behavioral and mental health professionals. Just as AI is already finding its way into the detection of physical conditions, such as diagnosing cancer from radiological scans, such technology may assist specialists as they analyze the gestures, drawings, and and speech patters of children. Hidden signs of violence perpetrated on or by children might in this way be

addressed at a point before long-term effects set in and begin or continue the all-too-common pattern of a cycle of violence. As such technologies develop, they might find application in every corner of professional-child interaction, from that of educators to that of physicians. In effect, AI will provide for the ability of continuous screening, catching symptoms of violence at the earliest possible intrusion on children's lives.

In law enforcement, justice, and corrections, AI may likewise find a screening function as well as providing predictive analysis of violent behavior. Ideally, this function will integrate these realms with mental health and substance abuse services to an increasing degree, enabling a tailored approach to each case of criminal violence. Again, the technology will provide assistance to professionals, exponentially magnifying their effectiveness in terms of case loads and overall outcomes. These particular services are currently singled out as being overwhelmed to the point not only of internal strain, but of strain on society itself. The ability to quickly determine the most effective action to take — which services and professionals to involve — in any violent altercation will substantially increase the ability of individuals involved to move past the event in a healthy manner that maximally preserves the stability of their interpersonal and social roles and relationships. In addition, serious issues concerning violent behavior may be identified and the potential for tragic violent events minimized.

In tackling violence on the larger scale, artificial technology may find its first inroads in the role of arbitration in the case of internationally contested matters. It would be extremely naive to assert that a technologically-derived agreement or treaty superior to those developed by human-derived diplomatic means could be found in any given dispute. However, the ability to take into account vast amounts of data and to supply suggested resolutions at an extremely rapid pace could potentially attract disputants to the use of AI-powered technologies as an adjunct to traditional diplomacy. In addition, data provided by AI assisted research may ultimately become accepted in the venue of international diplomacy in a larger sense and help to set the agenda of such bodies as the United Nations. A universal arbiter may forever stand out of reach, though the ability of AI to deal with enormous streams of data may help bring transparency to international relations which, with the hoped for knock-on benefits of reducing violent events and addressing violent tendencies at the individual level, may bring about a species-wide awareness of the dangers and opportunity cost of a perpetual arms race.

Where's the evolution come in?

It's not at all clear that, regardless of the degree to which they may continue progress in the direction of a more peaceful world, that the above applications of artificial intelligence technologies to the problem of violence will result in the necessary cultural evolution earlier arguments have posited. This brings the discussion into an extremely contentious area, which is the degree to which violence is popularized and capitalized upon in society. While there is a great range in cultural acceptance, it is a fact that the glorification of violence is a central fixture in entertainment, from cartoons to video games to movies and series, and that the focus of a great deal of ad-driven material online features violence as its primary draw. Psychological studies abound focusing on

the individual effects of this exposure; no widely-agreed upon results have emerged. The arc of the overall argument of this piece calls another question, however.

Just how does the perpetuation of narratives, themes, images, and associated cultural artifacts affect our tendencies to violence on a species level? This is a topic that expands beyond the typical bounds of scientific studies, yet has momentous implications for the future of humanity. It may be that the individual nudges addressed above will begin a process that results in a loss of the investment of cultural identity with violence and in interest in violence as entertainment, though this is rather doubtful. The value of the threat of violence as political capital, if nothing else, argues for a long future in which violence holds our attention. The other aspect of the AI-enabled future, that of freeing individuals from mundane tasks and economic competition and enabling a new level of thinking about and planning for the future may prove to be the more decisive feature here. Just how that future is to be approached will make all the difference in whether this process does indeed drive our species toward evolving a new relationship with the concept of violence.

Whereas the technologies featured above are driven by commercial and directly utilitarian goals, the idea that violence as a subject should be reframed and that society should act in a way that is meant to turn not only our actions but also our thoughts away from violence in all its depictions and forms is a normative argument, meant to direct the potential inherent in AI technologies toward a consciously selected end. This is the crux of the argument this piece presents: The future that is being rapidly shaped by the ever increasing speed of technological advance should be discussed, debated, and to the extent possible chosen by the individuals comprising the society that they shape. It is a special property of artificial technology, as previously addressed, that it offers the space and the tools in which these vitally important processes may occur. It is also a special property of AI, as previously addressed, that it presents a challenge to the future that requires these vitally important processes to occur.

A familiar metaphor arising in philosophical debate is that of humanity in its infancy, a comment on and challenge presented to the collective ability of the human race to come to grips with its own nature and the direction that nature is to take. While the ongoing debate over nature versus nurture in determining human behavior will certainly continue, new insights into the malleability of not only behavior, but also of patterns of thought and of the establishment of values have empowered motivated forces, notably commercial and political, to make use of developing technologies in their efforts to influence individuals and groups to their own ends. This has arguably already resulted in the evolution of patterns of thought and behavior so deeply ingrained as to be indistinguishable from any arising from biological evolution. This potential must be democratized if it is to serve the good of all.

In Conclusion

The power of culture as a technology driving the evolution of the human race has itself evolved to the point at which conscious choice plays a large role in its application. So

far commercial and political interests have capitalized on this power, with minimal input from the majority of those whom it affects. Artificial intelligence promises the means of truly universal education, of allowing for the cognitive energy to consider and debate, of alleviating the natural pressures to compete, and of granting the ability to profoundly influence the direction the future is to take. Alternately, AI offers the means of mass indoctrination, of commercial enslavement, and of subverting the potential of universal human advancement. Violence has here been selected as a cornerstone issue that, if taken up with serious purpose by an informed and involved population, will serve as a model for the potential of humanity to take charge of the vast changes technological advances such as artificial intelligence promise and/or threaten to instill in the fabric of human nature itself. Taking this issue seriously will set the stage for the next, as yet unforeseen, stage in the development of the human race. As the pace of change accelerates, so must the evolution of the species. Taking charge of the technologies that will shape that evolution should and must be of collective concern if the good of all is to be preserved.