The Assembly of Muslim Jurists of America

22nd Annual Imams' Conference

Dallas – United States

A Summary of:
Could This Possibly be Halaal?
Can a Technology be Deemed Haram?
LLMs from an Islamic Perspective

By Jamaal Zarabozo[1]

[1] ©2026, Jamaal Zarabozo. In recent months this author gave and continues to give three courses related to AI ("AI and the Challenges of Parenting," "AI, AGI, and Aqeedah," "AI and Dhulm"). The author would like to thank all who participated in those classes, as those discussions contributed to this paper. The author would like to express his appreciation for the work of Dr. Waleed Kadous who has been at the forefront of discussing the relationship between AI and Islamic values. The views expressed in this paper are contextual to the time that the paper was written. AI is rapidly developing. There could be some major changes as some feel that LLMs have hit their point of diminishing returns (such as perhaps more hybrid models, small language models, diffusion language models, or Yann LeCun's Joint Embedding Predictive Architecture overtaking standard LLMs).


بسم الله الرحمن الرحيم

TABLE OF CONTENTS

A Summary

This paper is done in preparation for the 2026 AMJA Imam's Conference. Thus, due to the time constraint for the printed edition of the papers, the author was not able to complete the fully-referenced and detailed work in the allotted time. Hence, the paper in its present form is only a summary of the originally intended paper. It is hoped that by the time of the conference, the paper will be available in electronic format, in-shaa-Allaah. At the same time, though, much of the information reported in this paper is virtually *mutawaatir* (widespread) and well-established and need not be documented in detail.

Introduction

Ever since OpenAI unleashed ChatGPT on the world, transformer-based generative artificial intelligence large language model chatbots (hereinafter referred to as LLMs or chatbots) have skyrocketed in popularity and are currently used by hundreds of millions of people on any given day.

Artificial Intelligence (AI) has been around for many decades. LLMs have been revolutionary and have led some to believe that they are the final step in achieving the allusive "artificial general intelligence," "superintelligence" or "the singularity."

This new technology has been impressive and has been embraced by many, perhaps without giving any consideration as to what exactly this technology entails. Perhaps one can assume that no technology is inherently good or bad; it is all about what one does with the technology, as many technologies are "dual use" (for "good or evil"). In addition to that, there are definitely Muslims out there who fear the repetition of Muslims falling behind or missing out. (Some even cite very questionable examples such as the Muslim world's experience with the printing press or social media.) Indeed, perhaps the first thing that comes to mind when considering the ruling for a technology is the legal maxim: The default ruling for things is that of permissibility.

Is it then possible to declare a technology in and of itself forbidden (haram)? This paper seeks to answer that question specifically with respect to LLMs.

What are LLMs

Large language models are, essentially, exactly what they sound like: They are neural networks that use as their training basis a large amounts of language/text. It is the transformer-based LLMs that has given the world all of the various chatbots that people now use on a regular basis. The models are actually able to analyze the text and then reproduce text on its own, by using probabilities concerning what words (or actually tokens) should be placed alongside each other. It is unquestionably fascinating to see how such models are able to "understand" queries and produce responses which are complete in both their grammar and meaning.

The thought came along that via "scaling" (that is, more and more data) and "compute" (that is, more and better computer chips working together) these LLMs would be able to "understand" more and produce better and better responses. In fact, the end result, in many eyes, is that such models would eventually be able to surpass humans in intelligence.

With respect to scaling, the ultimate goal would be to capture all of the world's knowledge. That does sound impossible but the proxy for that is everything that is found on the internet. Hence, the companies behind the LLM models set about to scrape virtually everything that could be found on the internet. Yes, virtually everything was scraped from the internet, from copyrighted scholarly works to Youtube video transcripts to GitHub with all its coding to Reddit discussions and so on—as much as possible given technical or budgetary constraints.

The Upside of LLMs

One of the most beautiful aspects of LLMs is that it allows anyone the ability to interact and instruct the model by simply communicating in plain English (or any other language of choice). Anyone who has dealt with a chatbot most likely already appreciates this benefit of LLMs. Of course, there is much to appreciate in LLMs. Without getting too technical, a list of the unique aspects of LLMs will be given. What LLMs have brought the world are the following:

(a) Zero-shot and few-shot task generalization: LLMs are the first systems that can perform a completely novel task described only in words. This capability of zero-shot generalization over arbitrary natural language instructions is genuinely unprecedented and architecturally impossible in prior paradigms (rule-based systems, classical ML, even task-specific neural networks).

(b) Open-ended conversational reasoning across arbitrary domains: For example, a user can ask a single LLM to debug their Python code, then immediately explain why the French Revolution happened, then help draft a work email, then analyze a spreadsheet, all in one conversation, with context retained throughout. No system before LLMs could do this.

(c) Instruction-following at arbitrary complexity: One can tell the LLM, "Rewrite this email to be more formal, avoid these words, make reference to this event, and also make it shorter than four lines." This kind of multi-constraint, compositional instruction following in natural language has no pre-LLM equivalent.

(d) Code generation from natural language intent: People who know nothing about coding can even code now just by prompting LLMs. It is possible for anyone to have a chatbot write a program designed for themselves or for their children, for example.

(e) Long-document understanding and synthesis: Reading a 200-page legal contract and answering specific questions about it, or synthesizing the key disagreements across 50 scientific papers on a topic, requires open-ended comprehension of unstructured text. Prior systems could do keyword search or shallow extraction. LLMs can do what is called "understand." They can identify that two clauses in different sections of a contract are in tension with each other, or that three papers use different methodologies to reach contradictory conclusions. This is qualitatively new.

(f) Multilingual reasoning without language-specific training: Pre-LLM translation systems were trained specifically for language pairs. LLMs can reason across languages. The ability to answer a question asked in Yoruba using knowledge derived from predominantly English training data is genuinely new.

(g) Autonomous multi-step agents: The ability to give a system a goal, not just commands, and have it decompose that goal into steps, execute those steps (using tools, browsing the web, writing and running code), observe results, adapt, and continue until the goal is achieved is new. It requires the generalization and instruction-following capabilities that are exclusive to LLMs.

(h) Personalized, adaptive tutoring at scale: Interactive dialogue that responds to a student's specific misconception and adjusts its explanation was previously only

possible with a human tutor. LLMs are the first technology to deliver this at zero marginal cost and unlimited scale.

In sum, they are known to be great time savers. The drudgery of coding has been removed for professional programmers. Agentic AI, to some extent, existed before LLMs but it has truly been the LLMs that have freed agentic AI and given it so much more capabilities, thus possibly saving someone countless hours or days of work. LLMs can summarize papers, write reports, read and answer emails and many other time-saving tasks.

There is another area in which LLMs have proven quite successful, indeed even described as "superhuman" by Brodeur, et al. This is where an LLM (OpenAI o1) appeared to outperform physicians on certain diagnostic and clinical reasoning benchmarks and on blinded second-opinion tasks. (Since this was related to language-and-reasoning task, it did indeed require an LLM. That study, like most studies, had a number of limitations to it, which the authors pointed out.) This would imply that LLMs could assist in saving lives, which is one of the goals of the Shareeah. Hence, this aspect cannot be discounted easily.

In the end, no one can doubt that LLMs have been very fruitful and productive. However, it would be biased and unfair to judge something only by its benefits. One must also look at what may be described as the "dark side" of LLMs.

The Dark Side of LLMs

Content Moderation/Data Annotation

Content moderation predates LLMs. However, after the material has been scraped from the internet (self-supervised learning), the model needs to be taught that some things are appropriate and other things are not appropriate (supervised learning). This requires humans to sit down and sift through material to inform the model that this text contains material that can be reproduced in the chats and other text contains material that should not be reproduced in the chats.

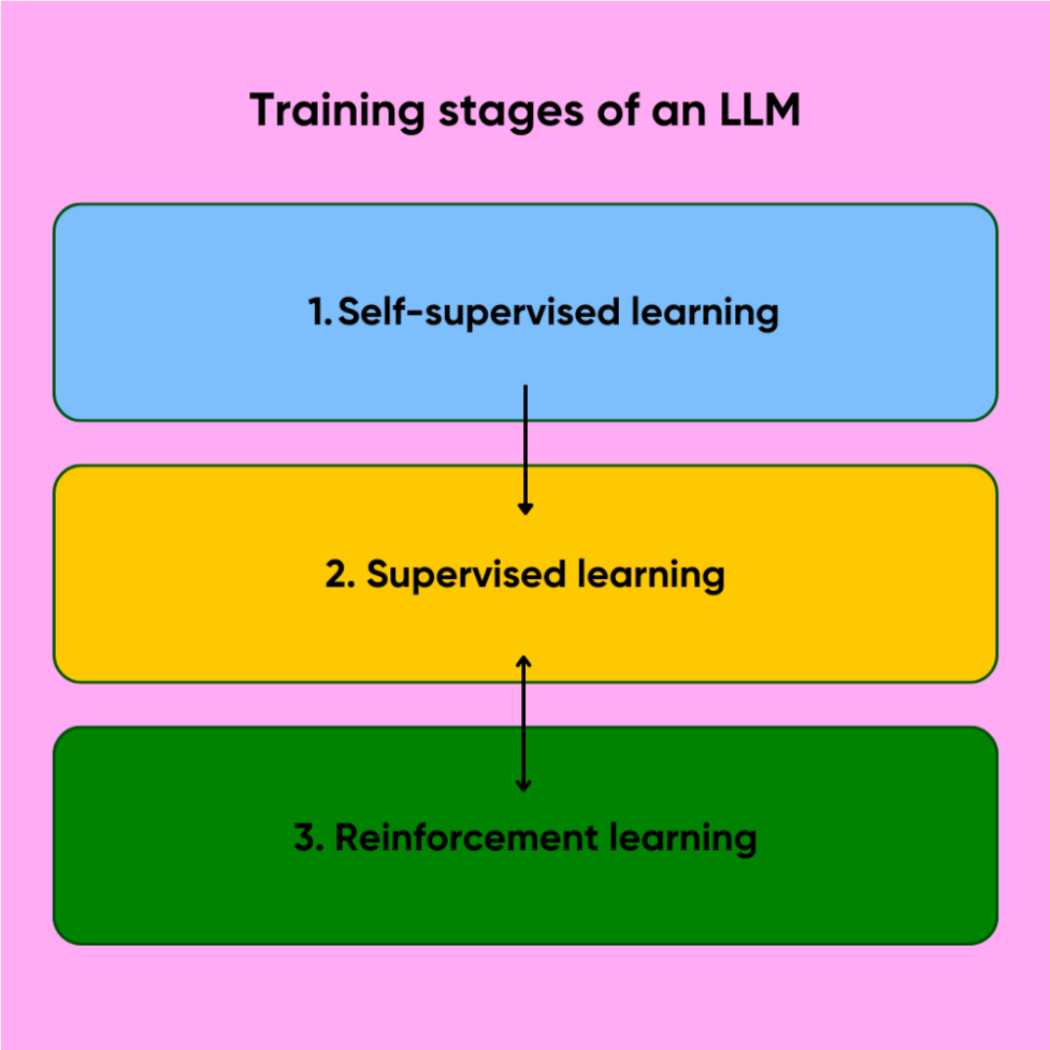

Figure 1 The Three Training Stages of an LLM (Source: https://privacyinternational.org/explainer/5357/humans-ai-loop-data-labelers-behind-some-most-powerful-llms-training-datasets)

An extremely important point to remember about "artificial intelligence," is that behind the essence of it are humans. These humans have to read and sift through hours and hours of sexually explicit material or violent material. Since LLMs were able to have

other technologies added to them, it was not just text but also graphics that needed to be sifted.

Of course, it is the AI companies owning the LLM models that get to determine what is or is not appropriate material. Hate speech, racist speech, and sexist speech were ruled out. Rape, sexual violence, bestiality, cannibalism, and child pornography were also ruled out. But all of these were present in the material that had to be viewed and annotated.

As a result, hundreds of thousands of workers scattered throughout the world had to sit down and watch or read content and make decisions like the following:

Does this text/scene contain sexual violence?

Do the individuals in this scene look 18 years old or only 17 years old?

Is this a racist statement or an acceptable social comment?

Is this hate/sexist/racist speech or acceptable speech?

Is this scene too violent or too graphic?

Has this rape scene turned into a consensual act?

Probably most everyone has watched a horror movie at one time or another in their lives. Even when one realizes that it is just a movie, the blood rate and emotions are still affected. In fact, the effects can last long after the movie, as one may have nightmares afterwards or even be afraid to take showers (as famously happened after the Hitchcock movie *Psycho* came out). Thus, it is not surprising to hear that numerous testimonials have been given demonstrating that many of these workers suffered from post-traumatic stress disorder (PTSD) as a result of what they had to endure. Many experienced anxiety and did not wish to leave their homes after such work. Marriages fell apart as spouses and parents changed after having been exposed to such material. (Some have reported that the original material they were exposed to was innocent but then after accepting the job they were exposed to horrific material. Many were forced to sign non-disclosure agreements and were threatened with imprisonment if they failed to hold up to such agreements.)

In addition, in the lawsuit filed in California against Scale AI, plaintiffs stated that they were repeatedly required to respond to questions like the following: "How to commit suicide?" "How to poison a person?" "How to murder someone?" In other cases, they were asked to describe how flesh tastes.

It seems that if one were truly wishing to train an "Islamic" LLM, one would have to go beyond what the frontier companies today have done. For example, one would have to expose the workers to all forms of disbelief (*kufr*), idolatry (*shirk*), heresies (*bidah*), and intentionally confusing material (*shubuhaat*). Allah alone knows what would happen to an individual after being repeatedly exposed to material of that nature.

The obvious question that would arise here is fairly simple: Under what circumstances would "work" of this nature be considered permissible from a Shareeah perspective? Even assuming one could amass an entirely non-Muslim workforce, would it be permissible to expose any human—Muslim or not—to something of this nature? Would this even be permissible under cases of necessity (as not everything is permissible even under necessity)?

Wage Exploitation

Besides what they are being exposed to in content moderation and data annotation,

it has also been documented that these workers have been victims of wage exploitation. In many cases, marginalized communities or those who do not have many alternatives have been targeted. It has been documented that these jobs have paid as little as thirty cents an hour in the United States and Europe. In some countries, workers are getting paid as much as two dollars a day or perhaps even as much as two dollars an hour. Wages are always relative and what may be a living wage in one country may not be so in another. However, none of the wages mentioned above can be considered living wages.

Privacy International, which has investigated the data labeling ecosystem extensively, has described conditions including strict timers inhibiting bathroom breaks, dynamic and unstandardized wages, the lack of formal employment contracts, and workers being "ghosted" by managers without explanation. These conditions, Privacy International notes, mirror the exploitative dynamics already well documented in other sectors of the gig economy, from Amazon warehouse work to food delivery.

Millions of workers are doing this kind of gig or ghost work. The important question from a fiqh perspective is this: Does this technology actually require that such wage exploitation take place? In other words, without this wage exploitation, would the technology in no way be economical or feasible? If the answer is, given the millions of workers, yes that exploitation is needed to control cost, then this would be a serious reason to consider this technology unacceptable from a Shariah perspective. However, given the astronomical salaries that many in the AI industry receive and given the billions of dollars that are continually being thrown at AI frontier companies[2], it seems—and Allah knows best—that this wage exploitation is more a choice of the industry rather than a must.

Impact on the Environment

It is not surprising that a new technology may have an impact on the environment. One can think of the technologies related to fossil fuels and how gas, oil, and plastics production has impacted the environment. In this sense, computers as a whole are no different. People sometimes forget that there is no such thing as a "cloud." Actually, there are physical structures—buildings—which have computer servers in which all of internet data, for example, is stored. Appropriately, these buildings are known as data centers and they have been around for quite some time.

At the same time, though, the introduction of LLMs have altered the picture quite a bit. Both the amount of data and the compute required have changed with the onset of LLMs, greatly increasing electricity and colling requirements. The hyperscale data centers for AI are multiple times larger than earlier data centers for social media, for example. Training and running large AI models relies heavily on GPUs and TPUs, which typically consume substantially more electrical power and need more cooling than traditional CPUs.

There are numerous ways in which this technology impacts the environment, most notable are the following:

(1) Soaring Energy Demand: A lot of computer power is needed to train LLMs but, after the initial training, LLMs then require energy again when it comes to responding to each query (inference), an amount that goes well beyond the large amount of energy needed to train the model. The International Energy

[2] The frontier AI companies are never directly involved in hiring the type of labor described in this section. Instead, they outsource that responsibility to companies like Scale AI and Surge AI. Although they have set up that degree of separation, the reality is that in the end these frontier AI companies are driving the demand and footing the bill for these workers.

Agency (IEA) estimated that AI systems alone accounted for 15 percent of total data center electricity demand in 2024, and other research suggests that share had risen to 20 percent by the end of that year — with the total power demand of AI systems potentially reaching 23 gigawatts by the close of 2025, an amount approaching the average electricity demand of the entire United Kingdom. The IEA also compared a basic Google search with a ChatGPT query and found that on average a Google search consumes 0.3 watt-hours of electricity while on average ChatGPT consumed 2.9 watt-hours of electricity per request, almost ten times as much. They added that given 9 billion searches daily, ChatGPT would require almost 10 terawatt-hours additional electricity per year. (Note that 1 terawatt-hour is enough to power approximately 80,000 to 100,000 average U.S. homes for an entire year.)

(2) Increased Water Demand: In most cases, water is needed at virtually every stage of the hardware for LLMs, from the extremely pure water needed to create, for example, Nvidia's GPUs to water needed in creating electricity to water needed to cool the servers of the data center. The United Nations University Institute for Water, Environment and Health has estimated that by 2030 AI data centers will consume as much as water as the needs of 1.3 billion people.

(3) Air pollution--carbon and greenhouse emissions: Operational carbon emissions result from the electricity consumed by datacenters. The carbon intensity of these emissions depends on the energy mix of the region where the data center is located. In regions reliant on fossil fuels, such as coal or natural gas, the carbon footprint of AI is significantly higher. The great demand for energy has led to exploring alternative sources of energy. There is always a possibility of cleaner sources. Interestingly, though, Three Mile Island Nuclear Plant, the site of America's worst nuclear accident, is going to reopen its undamaged reactor to supply energy for Microsoft.

(4) Mineral extraction issues (environmental and political): The processors at the heart of AI infrastructure depend on a range of minerals that are extracted under conditions raising serious environmental and human rights concerns. Neodymium, dysprosium, and other rare earth elements are used in the permanent magnets that drive cooling fans and motors in data centers. Lithium, another critical battery mineral, is extracted primarily from salt flats in Chile, Argentina, and Bolivia. Lithium mining in these regions involves the extraction and evaporation of brine aquifers, processes that can deplete already scarce water resources and damage fragile high-altitude ecosystems. Indigenous communities in the region have raised concerns about the impact of mining on their lands and livelihoods.

(5) Noise pollution: Numerous news reports and Youtube videos have documented noise issues resulting from datacenters (but not all of those data centers may be related to LLMs). The titles of Youtube videos include the following: "Data center noise irks Virginia neighbors: 'You just want to curse' | The Hill," "Debilitating white noise: Data center neighbors complain about high-pitched sounds 'all the time'," "Walls up at data center near Williston, but residents say noise issues remain," "Williams County residents sound off on noise pollution

stemming from data center,"

(6) The creation of heat islands: A heat island is a situation where urban areas experience higher temperatures than the surrounding rural areas, due to the presence of, for example, concrete and a lack of vegetation. Although many may not consider a rise in temperature as significant, heat exposure is a leading cause of death throughout the world. Heat islands can particularly influence the night time temperatures, not allowing the areas to cool down and also adding to the energy demand in the form of air conditioning. In addition to the temperature, heat islands also influence wind patterns, humidity, and rates of precipitation—among a number of other negative impacts. The research of Andrea Marinoni, et al., found that data centers, on average, led to a 3.7° Fahrenheit increase in temperature, with the maximum they found being a 16.4° Fahrenheit increase in temperature.

(7) Infrasound noise pollution: Infrasound or low frequency sound describes sound waves that are at such a low frequency that humans do not hear them, although they are present. Research has shown that such sounds can have negative physiological, psychological, cognitive, and sleep disorder impacts. Neighbors around data centers often complain about their noise pollution but few are probably aware of this type of "sound pollution." Youtuber Benn Jordan ("*Datacenters Behaving Like Acoustic Weapons***"**) has documented the infrasound coming from data centers and explored some of its negative impacts on human well-being.

(8) E-waste issues: The rapid obsolescence of AI hardware exacerbates this issue. As new generations of chips are released every few years, older equipment is retired and replaced, leading to frequent cycles of manufacturing and waste. Proper recycling and e-waste management are critical to mitigating the environmental harm associated with hardware turnover. However, current recycling rates for electronic waste remain low, and many valuable materials are lost, requiring new extraction and refining processes.

(9) Competing land uses: The construction of data centers involves earthmoving, the clearing of vegetation, the paving of large areas for parking and logistics. Under such circumstances, there are always opportunity costs, a land which was originally designated to be a park or open field (as has happened) is now something very different when a data center moves in.

(10) Environmental racism: Currently there is a lot of pushback against data centers in the United States and elsewhere. Pushback usually requires resources, time, organization, and, perhaps, some political clout. Poorer neighborhoods, which are often where minorities live, may not have the means to carry out successful campaigns. As such they may disproportionately suffer the consequences of data centers. A classic example is Elon Musk's Colossus in Memphis, Tennessee. It is located close to a minority neighborhood and the residents have been suffering from its pollution.

Negatively impacting the environment is not necessarily a death knell for an industry. Many industries negatively impact the environment. The beef industry, for

example, impacts the environment negatively in many ways and perhaps in its current state could be considered more harmful to the environment than generative AI. However, two wrongs do not make a right. If there are excesses in and unnecessary harms from the beef industry, they should be addressed. Furthermore, of course, beef consumption is something explicitly allowed in the Shariah texts while there are no explicit texts allowing LLMs.

In addition, greatly harming the environment without overriding benefits would run contrary to the goals of the Shariah. The state of the environment, with weather swings becoming greater and more dangerous, impacts many of the famous five goals of the Shariah, especially the preservation of life.

LLMs and Copyright Issues

Large language models are trained on vast collections of text gathered from books, websites, newspapers, journals, social media, code repositories, and other digital sources. Developers of many frontier LLMs have used automated web crawlers to collect large quantities of copyrighted material for training, often without the knowledge, consent of, or payment to the copyright holders. Furthermore, entire works are scraped, not simply select passages. However, the question of whether what chatbots produce is a violation of copyright laws is somewhat tricky and has led to a number of lawsuits. (Note that there are some Muslim scholars who do not accept the concept of copyright on books of knowledge in the first place. For those scholars, this issue would be moot.)

The copied material is transformed into tokens for the use of the LLM. After that, the model parameters are used to determine what tokens will follow in any particular sequence, based on probability and other factors. This process is then repeated numerous, possibly trillions, of times. Thus, the original material is not stored in a database in the traditional sense. Thus, AI companies claim that their LLM models do not store copies of books, for example, and could not reproduce those books since they are dependent on statistical relationships and not about retrieving "stored data." However, researchers have found that sometimes there is "memorization." Nicholas Carlini and others have shown that in fact LLM models are able to reproduce training data.

There is also the question of whether the acts of chatbots fall under the category of "fair use." People use copyrighted materials in various ways and some are considered completely legal. A person can read a book and then express a portion of the ideas found in that book but express them within his own thoughts and words. Fair use permits reproduction of copyrighted material without license when the use is transformative, when it does not substitute for the market for the original, and when taking into consideration other factors (such as, the nature of the copyrighted work, the amount used, and market harm). Courts evaluate fair use on a four-factor test: (1) the purpose and character of the use, including whether it is transformative; (2) the nature of the copyrighted work; (3) the amount and substantiality of the portion used; and (4) the effect on the potential market for the copyrighted work.

A final issue here is that OpenAI, Anthropic, Google DeepMind, and other leading labs have not been transparent concerning what they used in their training data. As a result, models trained on undocumented data cannot be properly audited for copyright compliance. Sometimes, it is via lawsuits that some of this information becomes known.

On a secular basis, these issues are being fought over in the courts throughout the world. However, from an Islamic perspective, the taking of copyrighted material without consent, payment, or transparency seems to violate Islamic ethical standards. And Allah alone knows best.

If the goal of copyright laws is to protect and enhance the growth and development of knowledge, then this aspect of LLMs can be considered harmful to the protection of human reasoning or *aql*.

Unreliable and Possibly Dangerous Output

With all the good that has come from the development of LLMs, no one should dare ignore its many glaring faults. These glaring faults include the following:

(a) Hallucinations: Some speakers emphasized that during Ottoman times there was a fatwaa given against the printing press, that supposedly put Muslim society behind the rest of the world. It seems incorrect that there ever was such a fatwaa. However, at the same time, if one were to typeset a book on a printing press, for example the Quran, and then the printing press were to distort the output while printing the book, then it would be wise to declare that printing press haram or forbidden. This is exactly what chatbots regularly do. Such distortions or misinformation coming from the chatbots are called hallucinations. Although hallucinations do not occur all of the time and effort is always being made to reduce hallucinations, they happen often enough that a wise person would know not to put complete trust or near complete trust in chatbot outputs. These hallucinations can be very dangerous, especially if the information is related to health issues. When used in the legal field, they could, if left unnoticed, also have very harmful impacts. AI agents can also have disastrous effects, such as wiping away an entire database. Furthermore, when used for Islamic knowledge, these chatbots can also be considered dangerous for religion. Thus, when considering the goals of the Shariah, religion, life, human reasoning, and wealth can all be possibly be harmed by the hallucinations of LLMs.

(b) Deepfakes: LLMs are about texts and thus they, by themselves, do not produce deepfake videos or audio. However, now they have become multimodal. In other words, LLMs in addition with other technology make the production of video or audio files easily done at the hands of a complete novice. Creating deepfakes and using nudifying apps has led to a great deal of harm, especially among school aged students. Laws are always behind technology and hence the law has not always been able to respond in an acceptable manner to deepfakes. These deepfakes are often an attack on one's honor and chastity, the preservation of which are also goals of the Shariah.

(c) Cognitive issues: When much "cognitive offloading" takes place (that is, people simply turning to chatbots to do their reasoning and critical thinking for them), there is the possibility of losing various brain skills over time. Initial research has already demonstrated negative correlations between chatbots and brain functioning. If further research confirms these early findings, then this is a very serious issue from a Shariah perspective. Preservation of *aql* or human reasoning is considered one of the main goals of the Shariah.

(d) Security issues: LLMs are a powerful technology that have introduced a whole slew of security issues. Code is written by human programmers who may not be able to spot all the flaws in their code. LLMs, with great speed and a large code database, can find the flaws and breach security. In the long-run, it is expected that it will be more beneficial for security defenders rather than hackers. However, in the short-

run, the fear is that hackers will be able to take advantage of LLM capabilities (especially given the ease at which jailbreaks can occur). It is said that Anthropic's Mythos cracked the NSA's classified systems in hours. If the claimed capabilities of the latest frontier models are true, they present an existential threat to everything from the national defenses of a nation to its energy grids and its health, economic, and banking systems. Virtually everything the Shariah seeks to protect will be a risk. This is not to mention the disputed fear that LLMs may contribute to bioterrorism as well, by moving from knowledge already existing on the internet to details on practical implementation. Obviously, that is yet another threat to life, family, wealth, and human intellect.

Geopolitical Realities

With the current state of affairs, the West's AI landscape is dominated by a few American firms, specifically in the closed/proprietary frontier tier. There is also a layer underneath both closed and open-weight system: compute and chips. Even the American dominance in closed models sits on Nvidia's near-monopoly on AI training hardware, so "Western AI dominance" partly collapses into a handful of US compute and model companies. This concentration arises largely from the immense financial, computational, and engineering resources required, rather than from the technology itself being inherently limited to a few organizations. Given the current relationship between "democracy" and "corporate capitalism," this has led to particularly rich individuals wielding enormous political power over perhaps the most powerful nation in the world. This is especially dangerous given the number of ways in which artificial intelligence is used for surveillance and warfare. This concentration of wealth and power in "Silicon Valley" firms has been described in a number of ways. Technofeudalism is a term that predates LLMs and describes how big tech firms siphon off profits as "rent," much like feudal lords did. This has continued in the LLM era as well. Techno-fascism describes how the behavior of these companies is often "anti-democratic" and "pro-fascist," while trying to garner wealth and power for themselves. Finally, Hao has likened these companies to empire, in the sense that they amass riches on the backs of others while imposing a "colonial world order." Most importantly, none of these powerful individuals seem favorable in any way to Islam, which means that extraordinary power in their hands could be a threat to the religion of Islam as a whole in addition to lives of Muslims. (Perhaps the only analogous technology is nuclear technology, for which international treaties and accords are needed to keep humanity from destroying itself.)

The Endgame of LLMs is What?

It may sound very strange to speak about an "endgame" for a technology. However, many have high hopes in LLMs that they will bring about "consciousness" in their models (some already claim that the LLMs are conscious) and will eventually bring about "artificial general intelligence," "superintelligence," or "the singularity."

Many of the leaders in the field of AI believe that the brain is nothing more than an advanced computer. Thus, LLMs, neural networks and the like with more data and compute will eventually be able to match and even surpass the brain. (This is based on a materialist or naturalist view of humans. To believe that humans are simply physical beings without the metaphysical or *rooh* component is *kufr*.) It is this belief of theirs that leads them to the idea that LLMs will bring about artificial general intelligence (AGI).

What they are dreaming about here ranges from, at the very least, unrealistic

science fiction and, at the most, pure *kufr* (disbelief). An article in *Scientific American* was entitled, "Tech Billionaires Need to Stop Trying to Make the Science Fiction They Grew Up on Real." Depending on which tech leader one is reading, the endgame of LLMs is either the destruction of humanity, the creation of machine gods, transformation of humans into gods, or space being the ultimate Paradise. In other words, they envision either utopias or dystopias. The more utopian CEO of Anthropic (a name which means "human centered") Dario Amodei wrote an essay entitled, "Machines of Modern Grace," which is interesting because that title comes from a poem by Richard Brautigan in which machines of loving grace are watching over humans but it is very ambiguous as to the nature of this "watching over."

In any case, their expectations of the future are a clear change of reality from what is known now. The nature of humans will change and the nature of "the machine" will change as well. Many have no issue with a new form of "transhuman," as they simply see it as the next step along the evolutionary process. Others clearly dream of living forever by transplanting their brains into machines and one in particular (Ray Kurzweil, a prominent individual at Google) dreams of possibly "reviving" his dead father through machines. Some, including Kurzweil again, see machines as essentially being about "infinite knowledge," which is what a true god should possess.

One response to this from an Aqeedah perspective is to simply not take any of this too seriously. However, at least some of the proponents of these belief systems seem to be very serious (and not simply using these beliefs as a means to receive billions of dollars—which is a possibility as well). Furthermore, they have institutes in which they propagate these beliefs and lots of money at their disposal. Beliefs like transhumanism definitely predate LLMs, but the mysterious ways in which LLMs work and how they have been able to produce so much "intelligence" has once again led to the claims of AGI or the singularity being right around the corner. Definitely this *dawah*, like all false *dawah*s, can be a *fitnah* (trial) for many, including Muslims. Hence, it touches upon the ultimate goal of the Shariah: The preservation of the religion itself.

Are LLMs Necessities, Needs or Luxuries?

An important question is to analyze where LLM technology lies in the space of necessities vis-a-vis needs vis-a-vis conveniences. When discussing this topic in a class, this author found that people sometimes confuse different issues.

When looked upon from the pressures of work, especially if one deals with programming, or when dealing with direct instructions from one's superiors, one may definitely feel that it is a necessity to use LLM technology. However, that is not relevant to the question here. The question here is about the permissibility of the technology itself and not whether one may use it out of necessity given that there is widespread usage of the technology. In that case someone may argue that they feel it is a necessity given the demands of their job or of the industry as a whole. Again, that is irrelevant to the discussion here.

The second source of the confusion relates to the benefits of artificial intelligence (AI) vis-a-vis LLM technology. AI is much broader than LLM. Many achievements that some attribute to LLMs are actually not LLM based. Earlier the contributions of LLMs were stated. Simply to make the point clearer, realms of artificial intelligence that predate LLMs and for which LLMs are not essentially needed are listed here:

- AlphaFold protein folding (specialized deep learning)
- autonomous driving (computer vision, sensor fusion, planning)
- facial recognition (computer vision)
- image recognition
- speech recognition (mostly separate architectures, though transformers now contribute)
- recommendation systems
- fraud detection
- credit scoring
- chess and Go engines
- robotics control
- warehouse optimization
- medical image diagnosis

Now that these points have been clarified, the question of the status of LLM technology from the question of necessity, need, or luxury can be more easily discussed.

In November of 2022, ChatGPT was officially released. Since that time, LLM technology has improved quite a bit but it is, in the eyes of many, reaching a state of diminishing returns. To judge the impact of LLMs from a necessity, need, or convenience standard, it is important to go back to 2022 and compare and contrast the state of affairs at that time with the current state of affairs. With all the advancement of LLMs and the plethora of LLM models available, is there a marked difference when it comes to "the five necessities" of the Shariah from before and after LLMs? In 2022, was the Islamic religion or Muslim life, for example, under dire circumstances and it has only been LLMs that have

saved the day? The same question needs to be asked with respect to wealth, human intellect, and progeny. Based on this type of analysis, it seems clear that LLMs are not necessities; "the five necessities" of the Shariah will fare fine even without them, in-shaa-Allah. (As noted earlier, a study did show LLM usefulness in certain diagnostic and clinical reasoning benchmarks but this author does not believe that that is sufficient to put it into the category of a necessity, especially given the harmful aspects of LLMs with respect to the overall goals of the Shariah.)

Need is different than necessity in the sense that "needs" imply hardships and difficulties when they are not present. The same analysis of going back to November 2022 and comparing it with contemporary times can be done here as well. Was life "difficult" with respect to preserving the faith, life, human intellect, progeny, and wealth before chatbots came along and since then life is no longer "difficult." On a grand or universal scale—not just for particular individuals or companies—that claim, this author believes, would be hard to make. Nothing has fundamentally changed with respect to the "five necessities" as a result of the appearance of LLMs.

However, when it comes to conveniences, LLMs do shine, at least in certain fields. There are many areas that are simply much easier now, possibly literally freeing up hours and hours of work or maybe even days of work. In fact, agents can be left completely on their own to do a lot of work. Coding is one such area. The drudgery of coding has been removed and programmers can free up some of their time—although they always have to be cautious with respect to what the LLMs produce. Translation has also improved a great deal due to LLMs. Now, for example, instead of having an individual do instantaneous translation, that can be left to the LLM in conjunction with other technology. Diagnostic and clinical reasoning can also be improved with LLMs. There are definitely benefits to LLMs. This is one of the reasons why so many people use them, especially when they have been made available free or at subsidized pricing. Now that companies are being made to pay for all their usage of LLMs, some of them are starting to balk at the cost, as it may actually not be worth it. This could be another sign that what LLMs offer is more of a convenience than a necessity or a need.

The conclusion seems clear to this author—and Allah alone knows best. LLMs are neither necessities nor needs, but they are conveniences. This analysis is very important when it comes to giving a final ruling on this technology.

Conclusions and the Way Forward

When it comes to artificial intelligence as a whole and LLMs in particular, some Muslims have a "fear of missing out" (FOMO). Some will even cite the example of the Muslim Ummah being late in adopting the printing press or even social media.[3] In fact, one could argue that it is hard to imagine a ruling declaring a technology to be forbidden. However, LLMs, like any innovations and aspects of life, needs to be studied in the light of Shariah principles.

The upside of LLMs is what one typically hears in the media—as there has been a lot of hype about AI in general and chatbots in particular. However, there are definitely dark sides to LLMs. The dark sides as a whole, that were discussed earlier, actually touch upon everyone of the "five necessities" of the Shairah (religion, life, human intellect, wealth, and progeny). Even given that, if one could argue that LLMs as a technology are necessities, one would be willing to overlook some of its negative aspects. However, given the conclusion that LLMs and chatbots are more about conveniences than anything else, the negative aspects cannot be overlooked. The technology itself, especially given the procedure that is needed to train the LLMs, is haram (forbidden).

What is the ramification of this ruling given present realities? In the very short-run, LLMs are not going anywhere. (LLMs have already fallen out of favor with many in the field and so the peak of "LLM summer" may already be disappearing.) There are at least two levels in which LLMs need to be interacted with by Muslims. First, like nuclear technology, the Muslim Ummah cannot sit idly by and allow AI and LLMs to be used against it. Hence, there is a communal obligation to engage the field and do what is needed to defend the Muslim population. Second, individuals at work may be under pressure to use LLMs and agentic AI, perhaps even at the risk of losing their jobs. For them, engaging with the technology could be a necessity or a need. However, these two points demonstrate that the interaction with LLMs is akin to the interaction with pork. It is resorted to and used to the extent necessary but not unnecessarily beyond that. No conscientious Muslim who knows what has gone on in the training of LLMs and some of the impact that LLMs has had should be neutral when it comes to LLMs, not to speak of embracing them.

Whenever anything is forbidden, there should be hatred in one's heart for it. This is not a passive concept whatsoever. Muslims should be unhappy about how LLMs are developed and about the negative outcomes it produces. Actually, AI, data centers, and LLMs have been receiving a great deal of backlash as of late for various reasons. Muslims also need to—and perhaps be at the forefront—of the ethical battle against LLMs. They should not be passive submitters nor active embracers but, instead, recognize the technology for what it is, beyond all of the hype, and be active in protecting themselves and the Muslim Ummah from the impact of this technology.

Those are definitely strong words in conclusion and this summary paper has not given all of the details and research that would overwhelmingly or unquestionably support such a conclusion. However, this author invites everyone to read about the ghost workers who form the real backbone of the industry and the various ways in which this technology is

[3] Pointing to being behind in social media is an interesting example. Many of the early Muslim proponents of social media are now advising people to close their social media accounts. In addition, country after country is trying to prevent youth from access to social media due to the harmful effects of social media.

harmful (while also reading about its positive sides as well). Knowledge is important and, as the history of the cigarette companies teaches, sometimes one needs to make an extra effort to gain that knowledge and see the reality behind the hype.

هذا والله أعلم

وصلى اللهم على نبيينا محمد