DOI: 10.47743/ejes-2024-0202 • December 2024 • VOLUME 15, ISSUE 2

AI & Regional Science: past, present, and future

Richard L. Church ✉
University of California, Santa Barbara, USA

Abstract
Artificial Intelligence, or AI for short, has captured the attention of almost everyone in the year of 2023, when several companies introduced software that can be used by the general public and applied to many different areas of interest, including voice, video, text, and large data analysis. For example, it is possible to create a novel using AI large language models or create or ask an AI app to generate a drawing of a given subject. Overall, AI has the potential to disrupt many industries, social media, and even government. This paper presents an introduction to AI with respect to Regional Science. It discusses some of the promises and pitfalls of AI, including its use in education. AI techniques are not new to many in Regional Science as it has already been used in a number of application areas and some of these applications are pointed out in this paper. Thus, it covers the past, present, and future for the field of Regional Science.

Keywords: artificial intelligence, AI, regional science

Introduction

Perhaps the decade of 2020-2030 will be called the decade of disruption, including disruptive elements of international trade, new novel pandemics like Covid, increasing international conflicts, increasing international migration, and greater impacts of climate change, not to mention the emergence of Artificial Intelligence (AI) as a major tool. No one would deny that climate issues associated with a warming environment will usher in more variable weather systems, possible droughts of longer duration, major heat stresses in cities like Mumbai, India, and higher sea-levels threatening cities and infrastructure around the world. Add to this, one can think of AI as an emerging disruptor in industry, Hollywood, government, education, and research. The promises of AI as a major aid in education and research are great, but not without possible downsides which will be explained later in this paper.

The term Artificial Intelligence was coined in 1956 during a 2-month summer research workshop held at Dartmouth College. It has been attributed to John McCarthy who was an assistant professor of Mathematics at Dartmouth. This workshop involved 11 scientists and was organized to discuss the basic elements needed in "thinking machines". The seed for this workshop undoubtedly was planted

✉ Professor, Department of Geography, University of California, Santa Barbara, USA; e-mail: rickchurch@ucsb.edu.

by Lady Lovelace and Alan Turing. Turing (1950) was the first to propose to test whether a machine can think. Lady Lovelace stated in 1842 that Babbage's "Analytical Engine has no pretensions to originate anything. It can do whatever we know how to order it to perform" (Hartree, 1949; Lovelace, 1842). Turing states that Hartree followed this quote and stated: "This does not imply that it may not be possible to construct electronic equipment which will think for itself" (Turing, 1950, p. 450). This statement seems to have inspired Turing to consider the question: "Can machines Think?" Turing's intriguing discussion on this question is a philosophical piece on how we might test such a possibility. This test involves a person (the interrogator) who asks a series of questions of a human and a machine (both hidden from the interrogator in another room) and, based upon the responses, the person attempts to distinguish which is the machine and which is the human. Turing proposed that the communication between the interrogator and machine and human would be based upon text, so that distinguishing characteristics of speech could not be used to discern which was the machine and which was the human.

Turing's discussion clearly indicates that he had given such a task a considerable amount of thought. Depending upon the area of questioning, I suspect that we have already reached the point where the machine answers may well be more nuanced and correct than the individual's. Also, the response from a chatbot may be very fast, whereas the human may take some time to answer. For instance, if one asked one of the chatbots to give a description of water (a task that I recently did with Microsoft's 'copilot' which is based upon ChatGPT, an AI program from the company Open AI), it would be more complete and nuanced than the one I would expect the average human would give. Would we then conclude that the answer was "too good" to be human? That is, we could distinguish between the machine and the human because "humans aren't that good". But upon asking the same chatbot (Microsoft's copilot), "what is the colour of your hair?", the answer was "I apologize, but I don't have hair". That is, if we want to distinguish between a human and a machine, we could ask about human issues, like the color of their eyes or where they were born. Outside of these distinguishing questions, to make the task less discernible between machine and humans, we might be approaching the point where we have to assign the machine human characteristics (like eye color), slow the response times, and sometimes say, I am not sure, or I'll have to think about that. That is, we will have to add a few human-like elements to the response. Turing suggested asking the man and machine "can you write a sonnet about the Forth Bridge in Scotland?" I did that for Microsoft's copilot and it produced a very compelling sonnet about the Forth Bridge. I immediately thought that the sonnet was quite good, in fact, maybe a copy of something already written by someone. I then asked copilot if this was original and copilot answered with "I apologize for the oversight in my initial response. The sonnet I provided is *newly generated* and dedicated to the *Forth Bridge* in Scotland." I conclude that many of the elements that were raised by Turing have been met, that is, we have machines that can think.

We already have computers that beat world class chess champions and are competitive at playing the game GO. New forms of chatbots now converse with prospective buyers in helping them on internet retail sites, look up the account balances on credit cards, and help filling out income tax forms. Whether communication is via text of voice, the answers and conversation seem real and personable. Obviously, much progress needs to be made, but one has to acknowledge only a few years ago such programs did not exist except in a crude form.

The objective of this paper is threefold. First, it is important to explain why AI has quickly come into use with new applications that can easily be used by the general public. Second, what are the prospects of AI and its use in Regional Science? Third, how will AI's use impact industry and government, leading to changes in employment and other possible disruptions.

AI as a science has its beginning roots in the early 1950s. One of the first truly thinking machine developments occurred in 1958 with the development of the Perceptron, a simple neural network designed to perform binary classifications (Rosenblatt, 1958). The features of Perceptron were supposed to form the design of a computer rather than code that represented the process, although the earlier versions were represented only in computer code. Perceptron was the very early version of what has now been expanded into a very large multilayered (even multi-module) neural network of millions of parameters that has formed the basis for such programs called large language models. For the next 20 years there were a number of new developments in the field of AI. But the 1980s experienced what has been called the AI winter, a time when enthusiasm and funding in the field had cooled. That was followed by the development of expert systems in the 1990s. Such systems were designed to help understand, model and support decisions in specific problem areas, like environmental management and transportation planning. In the last 20 years we have experienced substantial developments in a number of research areas such as robotics, self-driving cars, chatbots, and visual recognition systems, just to name a few. But given the fact that many of the rudimentary models for AI have been in the literature for quite some time, just why has AI become so prominent today as compared to just 3 or 4 years ago? In the next section, we explore why it has emerged as the new frontier. This is followed by a short discussion of Large Language Models. Section 4 provides a brief overview of approaches found in AI, beyond large language models. Section 5 details the major focus areas of AI as well as a few selective applications in industry and research. That is followed by a short discussion of where AI fits within the field of regional science. Section 6 discusses potential problems with AI. Finally, we make a few final thoughts in section 7.

1. AI: why now?

Even if one does not know much about AI, the average person on the street thinks of AI as an emerging new technology which has been recently used in industries, like robotics in auto production, search engines in browsers like google chrome, and navigation systems like Tesla's autopilot. Although it is a seemingly new term, as explained in the introduction, AI has its roots in research going back nearly 70 years. It is only logical to ask, if many techniques in AI have been developed over the last 70 years, why has it emerged as a technology that will solve societies' major problems now? The answer to this is fourfold. First, the cost for computing has dropped substantially over the last several decades, but even more recently with the development of graphical processing chips.

Table 1 presents data associated with Intel CPUs that have been sold over the period of about four decades, starting with 1978. You will note that the first of the series of 8086 processors were able to process 33 million instructions per second. This may sound relatively fast, but you can observe in Table 1 that there was a steady increase in computational speeds and that, by 1991, the speed of the 80486 chipset was 33 times faster than the initial INTEL 8086. To put this in perspective, in the late 1960s, IBM manufactured a handful of supercomputers (360 model 91) that cost more than $4 million and had a computational speed of 16.6 MIPS. Thus, the desktop computer of 1994 (i.e. 80585) had a computer processing speed that was 10 times faster than the supercomputer of 24 years earlier. Computational speed is not the only consideration. Cost is also a major consideration. Whereas the 360/91 needed a large, air-conditioned room with peripherals that cost nearly $8 million, a complete desktop computer with an 80586 chipset could easily be purchased for about $2,500.00 in 1994. But progress did not stop there. By 2011, Intel was selling the core i7 CPU that contained 6 cores and processed 177,000 MIPS, a speed and capability that made the supercomputer of the 1960s ancient history. A competitor to INTEL, Advanced Micro Devices, introduced in 2020 a CPU with 64 cores and a processing capability of 2,356,230 MIPS.

Another major consideration is the introduction of graphical processing units (GPUs), initially introduced to provide support for 3-D rendering in computer graphics by emerging companies like Silicon Graphics. One of the major players in today's GPU marketplace is NVDIA. The processing units on a NVDIA GPU are called CUDA cores, short for a "compute unified device architecture". These CUDA cores are designed to support parallel computing to a significant degree using a specialized instruction set. Whereas a GPU from NVDIA might contain thousands of CUDA cores today's Intel CPU might have 100 cores or less. For example, an NVIDIA GeForce RTX 3090 GPU contains more than 10,000 CUDA cores (cost $1500); in comparison, a 13th generation INTEL i9-13900K contains 24 cores ($600 cost). Admittedly, the INTEL CPU is more flexible with a wider instruction set than an NVDIA GPU, but the immense amount of parallel computing possible on the GPU makes some computing considerably more efficient and faster, especially in the support of graphics needs for gaming and 3D rendering. One way to compare shear processing power between the CPU and GPU is by analysing the processing of floating-point operations. An NVIDIA GeForce RTX 3090 GPU is rated at 35.5

TFLOPS (trillion floating point operations per second, whereas an INTEL i9 13900K is 1.15 TFLOPS. This clearly demonstrates why GPUs have helped to make AI applications more commonplace.

Table 1. Computing speeds of Intel CPUs over time

CPU-name	MIPS (millions of instructions per second)	Date of introduction
8086	.33	1978
80286	1.33	1982
80386	4.3	1989
80486	8.7	1989
80486	11.1	1991
80586-Pentium	188.	1994
80686-Pentium II	541.	1996
80686-Pentium III	2,540.	1999
Core i7 3630QM (2.5Ghz)	113,093	2012
Core i7 3960X (3.3Ghz)-6 core	177,730	2011

Note: Recent computer chip processing has continued to be faster, taking advantage of multiple cores and threads. Intel is not the only company that makes CPUs. The speed of Advanced Micro Devices (AMD) Ryzen Threadripper 3990X with 64 cores that was introduced in 2020 is rated at 2,356,230 MIPS.
Source: all data in this table was sourced from various internet websites

Second, the cost of computer storage has dropped significantly over the last few decades, making the storage and retrieval of data in support of spatial and large data analytics relatively inexpensive. Table 2 presents the trend of computer disk drives storage and costs over 67 years. In 1956, the cost of a disk drive was $9200 per megabyte (3.75 Mbyte drive). Today, the cost per megabyte is now $0.0000125 where the drive has a capacity of 8 teraBytes. Note that the data in table 2 does not include computer memory costs as well as solid state drives. These two technologies have also expanded over time along with a similar trend in decreased cost be Mbyte. Thus, just as computing power has increased over time, so too has the capacity to store information. In addition, the costs of both storage and computing have dropped substantially over time. One can say that, in many ways, we have entered the era of cheap computing and storage.

Third is the fact that a large number of data sources and an immense amount of data is available for analysis. This includes data sources such as cell phone tracking data, traffic flow counts, video of traffic at intersections, Google's Streetview, parcel data for virtually every piece of land in the US. Even computer aided dispatching systems data from emergency services like fire departments are now available from many jurisdictions where private information has been scrubbed. FEMA produces flood maps and radar tracking of weather is now available in real time. The internet is also a vast source of data, not to mention the fact that some data has been made available at a price that characterizes neighborhoods and their shopping patterns and consumer profiles. Even Microsoft has made available a database of 129,591,852 computer generated building footprints derived by using computer vision algorithms on satellite imagery. Oftentimes, such data may not be exactly what one would collect for a specific purpose, but it can be used in innovative ways to tease out information and apply regional/spatial analysis.

Table 2. Cost and storage trends of computer disk storage

Year	US$/MB	Manufacturer	Size	Mbytes
1956	9200.	IBM	24”	3.75
1959	7600.	IBM	24”	7.5
1960	3600.	IBM	24”	10.
1966	1047	IBM	14”	29.
1970	259.	IBM	14”	100.
1974	185.	IBM	14”	200.
1985	31.	*	**	140.
1990	5.28	Miniscribe	**	340.
1995	0.267	Seagate	5.25” (SCSI)	9000.
2000	0.00407	Maxtor	3.5”(IDE)	30700.
2005	0.000406	Western Digital	3.5	320000
2010	0.0000550	Seagate	3.5”(SATA-2)	2000000.
2015	0.0000300	Seagate	3.5”(SATA-3)	3000000.
2020	0.0000194	Seagate	3.5”(SATA-3)	8000000.
2023	0.0000125	Western Digital	3.5”(SATA-3)	8000000.

Note: all data is sourced from: https://jcmit.net/diskprice.htm (McCallum, 2023)
* denotes manufacturer is not given; ** size of disk is not given; some of the disk drives give technology used: SCSI - small computer system interface, IDE – integrated drive electronics, SATA – serial advanced technology attachment
Source: author’s representation

Finally, there has been a continued development of software systems including languages that readily support analytical models. These include scripting software like python, statistical packages like R, and specialized software products like BARON, CPLEX, GUROBI, Geoda, and ArcGis. Add to this the fact that you can ask a program like ChatGPT to write a computer code for a specific task means that we have entered a new era for computer programming as well. Legacy computer codes are still being used, involving computer languages like COBOL, at banks, telephone exchanges, industrial plants, to name a few. Such codes have been costly to maintain and replace. It is quite possible we will soon enter a new era, where the replacement of these legacy codes will be inexpensive using AI tools. It also means that in the future, employment in the computer programming/coding sector may decline, where higher level pseudo-code instructions will be coded automatically without significant programmer attention.

2. The new era of AI has now emerged with Large Language Models

In the 1980s, AI progress had cooled and most of the funding had dried up. There were new developments in the field, but they were hampered by limited computational resources and available/retrievable digitized data. Now, the trend has reversed, with a significant growth in AI with the reasons for this being discussed in the last section. One of the key emerging areas that has caught the attention in the public arena is the maturation of the Large Language Models and their use in chat-bots. A good example of this is the program by Open AI, called ChatGPT (Generative Preformed Transformer). It involves a massive multi-layered neural network (actually a combination of modules, each with a defined general purpose and based on a neural network, with a general knowledge workspace connecting the modules) with over 1.75 trillion weighted arcs. According to Microsoft's copilot that is based upon ChatGPT, the training database is on the order of 570 gigabytes or approximately the equivalent of 300 billion words. The essence of this is that Chat is trained to be a good guesser (my terminology here). It is given a sentence with a missing word. Chat then guesses what the word is and then it scores itself on how well it did. This training involves many, many sentences and recalculations of arc weights until it gets many of the missing words correct or close to right. It is this basis for which Chat can then be used to construct text in communicating with the user, along with the immense knowledge built into the neural network (Naveed et al., 2023). But that does not mean it is infallible. For example, last November (2023), I queried ChatGPT3.5 about who Kingsley Haynes and Peter Nijkamp were. I chose these two individuals as they were the co-chairs of the session that this paper was presented in. ChatGPT answered, noting that among other things they both held PhD degrees from The Johns Hopkins University. Although Kingsley Haynes received his PhD from Hopkins, Peter Nijkamp received his PhD degree from Erasmus University. Several months later I asked the same question and Chat correctly stated that Peter Nijkamp holds a degree from Erasmus University, but then stated that Kingsley Haynes had been at the University of Pennsylvania. So, the second time, months later, the answers were still only half correct. It should also be noted that Chat will not necessarily know the answer, but it will make an intelligent guess. That is, it holds a list of possible answers with associated probabilities of being correct. It makes a selection from this list (usually the answer with the highest associated probability of being correct). It should also be mentioned that it will usually give a reference or two as to the source of information, but the user must recognize that its sources for a particular field may be limited. Most users do not understand that ChatGPT's answers can be tuned by the user. One of the key parameters in controlling the answers from ChatGPT is to set the temperature (a form of setting the temperature in a process called simulated annealing). A high temperature (say .75) makes the system more diverse and creative in its answers and it could select a "word" from the list of possible answers with a lower probability of being correct. A low temperature value (say .15) "produces a more focused and deterministic response" (ChatGPT's terminology), that is, it picks the answer with the computed highest probability of being correct. ChatGPT also provides a way in which the choice of a given word may be limited among a reduced set (Top-p) of potential answers. This reduced set is similar to a structured semi-greedy choice in a semi-greedy heuristic, a process long used in optimization. Most users do not even bother to test out various settings but suffice it to say that the answers are not necessarily deterministic and not always right.

3. Becoming familiar with AI in Regional science

To the uninitiated, AI may seem to be a mystery, but in most respects, many people in Regional Science have used techniques that can be classified as AI methods. Clustering and classifying data into categories has a long history in Regional Science research. One of the historical techniques in data clustering is called K-means. Given an *n*-dimensional space and a number of data points, each mapped within the *n*-dimensional space, can the data be clustered into *k*-sets, where each set or cluster is represented by a cluster center and each data point is assigned to the closest cluster center (using Euclidean distance)? This problem was defined by MacQueen (1967). MacQueen defined a very efficient implementation to cluster data into k-sets whereby the memory required by the computer program was very small and could be programmed on a computer where a large data set could be stored and read directly from a data tape, never having to be stored in memory. It was considered to be one of the best clustering techniques at the time. Unfortunately, MacQueen's method suffered from the fact that the starting set of cluster centers was predefined as the first *k*-points in a data set. Because it was a heuristic and had a predefined starting point, its performance could not be guaranteed in generating the best k-clusters. Numerous refinements have been made over time to improve results (as it is a heuristic), and it still remains a mainstay in data analysis and can be found in statistical programs such as R (The R-project, 2024). When asking Microsoft's copilot about K-means, the response was "K-means is not specific to AI; it is a general technique used for clustering data. However, it is commonly employed in AI-related tasks due to its simplicity and effectiveness." A cursory review of the textbook, Machine Learning (Flack, 2012) includes many techniques that can be found in use in the regional science field. In fact, such techniques as K-means, ANOVA, data visualization, principal components analysis, and many other techniques are part of the tool kit of AI.

To cite a personal example, in 2008, I was working with Professor Zeng from China University of Geosciences in Wuhan, China. His particular interest was in finding shortest paths in real road networks. We embarked on a review of techniques and began to explore some new innovative ways of solving very large problems to optimality. At the same time, he was developing applications based upon the fastest techniques found in the computer science literature, based upon such techniques as 2Q and Dykstra' algorithm with the latest forms of handing priority queues (Zeng & Church, 2009). We found that our experimental algorithm was faster than anything developed to date. But our analysis was based upon the state-of-the-art work in the computer science literature. Basically, the computer science literature had developed independently of what existed in the artificial intelligence literature. I was aware of a technique A* (called A-star) that had appeared 20 years before in the AI literature (Dechter & Pearl, 1988; Pearl, 1988) that I thought was a possible candidate as the fastest existing process that had been developed, but quickly found no reference to it having ever been tested to those in the operations research and computer science literature, including 2Q and Dykstra. We concluded that we could not declare our new process the fastest without also testing A*. That turned out to be not a simple task, due to the fact that the simple structures for priority queues (the list of the best candidates for subsequent closest nodes in a shortest path) needed to be expanded to handle A*. Prof Zeng embarked on that task and then tested the latest, fastest versions of known codes against our version of A*. The test proved that A* outperformed what was considered to be the best existing algorithms by a sizable amount when applied to large real road networks (Zeng & Church, 2009). This is an example of finding research in the AI field that can have important implications in Regional Science and Geographical Analysis. Another example is the development of Random Forests (Breiman, 2001). Random Forests are based upon an ensemble of predictive decision tree models. It involves combining the predictions of multiple models to make a final decision of classification. For the case of Random Forests, each predictive model is based upon an instantiation of a decision tree. The final prediction is based upon a majority vote or average of all of the tree predictions. This technique has been used in fields such as remote sensing (Belgiu & Drăguţ, 2016), ecology in habitat prediction (Cutler et al., 2007), water resources (Tyralis et al., 2019) and bioinformatics (Qi, 2012).

In the field of spatial optimization, which includes location modelling, districting, and network design, one can find the application of a variety of optimization techniques, both optimal and heuristic procedures. This includes such techniques as genetic algorithms, simulated annealing, swarm smarts like ant colony optimization, vertex substitution and greedy approaches, TABU search, path-relinking, and threshold accepting (Dueck & Scheuer, 1990). Virtually all of these techniques can be found in use in different AI applications. In fact, I hazard to guess that everyone in the field of Regional Science can find a technique that they have used in their work will have also been employed somewhere in an AI application. Why? The simple fact is that any useful modelling approach will eventually be used in one or more AI applications. For example, genetic algorithms have been used in machine learning programs (De Jong, 1988). Any approach that can be used to search for or optimize a set of rules (a combinatorial problem) will by its very nature involve some type of optimization technique that has been developed or applied in Regional Science. The real power of AI is based upon a combination of several techniques, something that we could not easily do a number of years ago but which we can now consider due to the immense capabilities of computer speed and storage.

4. Focus areas for AI

There are a number of substantial research and application areas in AI. They include (but are not limited to):

1. Machine learning: systems that can learn from data and make decisions or predictions;
2. Natural language models: deals with understanding and generating natural language, such as text and speech (e.g. ChatGPT described in section 3);
3. Computer vision: processing visual information from images and video (face recognition, generating a pathway within an image for car navigation, etc.);
4. Robotics: creating and controlling machines that can interact with the physical world;
5. Expert systems: emulates the knowledge and reasoning of human experts in a specific domain;
6. Fuzzy logic: attempts to handle uncertainty and imprecision.

There have been significant developments in each of these fields, or in applications that involve a combination of several of these categories. For example, developing a self-driving car is based upon computer vision, fuzzy logic, robotics and expert systems. What is important to understand is that each of these areas can have a significant impact on society and our economy. For example, WM (Waste Management Inc.) is a very large company dealing with solid waste collection and disposal in the US. It operates a number of Material Recovery facilities (MRF) involving sorting a waste stream and recycling different materials, including plastics, glass, aluminium, etc. The classic system involves human sorters, a labor-intensive process. In their new facilities, WM is now integrating optical sorters and robotic arms with visual comprehension to pick waste items from a conveyor belt. In one new test facility in Houston, TX WM found that the labour cost per ton of waste handled was reduced by 33%, throughput capacity was increased by 75%, recyclables in the residue stream was reduced by 50%, and the recycle value was increased by 30% per ton of waste handled. The value to WM of introducing robotics is substantial. The fact is that flipping burgers, communicating with customers online with chatbots, and providing better traffic management on roads will all be handled differently in the future, with substantial impacts to employment, regional economies, and quality of life.

In the research field, neural networks, random forests and other techniques originally developed within the field of AI will continue to have a large impact on science, in general, including regional science. One example of this is literature-based discovery (LBD). LBD represents new methods to analyse the existing scientific literature, with the focus of making new discoveries. For example, Lawrence Berkeley Labs used LBD to analyse abstracts of papers to gain “chemical intuition” and suggest materials for a specific purpose and found all ten suggestions identified by their software displayed the needed properties but were not known for that property in the literature (Could AI transform science, 2023). In medicine, one example involves a Quantitative Structure Activity Relationships (QSARS) model that found a compound used in toothpaste “can inhibit an essential mechanism in malaria causing parasites” (Could AI transform science, 2023). Some researchers are now using AI to suggest hypotheses given data, with the idea that new ideas might be generated. Whatever the case, it is the dawn of a new era in which AI will become a part of our research paradigm and our educational processes.

The future of AI is a topic of much discussion and speculation. Tewari (2022) has listed a number of areas in which AI will have a great impact. Among these is that AI and machine learning will transform the scientific method. This includes Literature Based Discovery and Quantitative Structure Activity Relationships techniques. We may well use AI techniques to structure or propose hypotheses or we may ask AI to look for flaws in a given paper. Some of the most promising possibilities of AI in Regional Science and geography include:

1. Spatial machine learning: this involves the use of machine learning algorithms applied to spatial data, e.g. satellite imagery, to identify patterns and make predictions;
2. Geospatial analytics: AI can be used to automate information extraction from geospatial data, including car cameras, traffic cameras, cameras on crowds at a demonstration, and even animal movements, etc.

Some now call this new field GeoAI which involves the use of AI methods such as machine learning with spatial data to glean meaningful information. Without a doubt, Regional Scientists will begin to use new AI tools for specific purposes such as transportation flow management as well as contribute to expanding existing AI methods. Overall, Regional Science has much to contribute to understanding the impacts of AI, especially in increased inequality among regions, employment categories, and socio-economic status.

In more classical terms, one can think of the main economic effects of AI fitting into two main categories: direct and indirect effects (Chen et al., 2016). It is expected that direct GDP growth will result from the investments that are occurring in industries directly associated with the development of new AI tools, the production of new parallel processors, the development of new data centers, and in companies that are expanding access to data. Indirect economic effects will result from changes in GDP resulting from those industries and governmental agencies that use these services and software. Whereas employment in companies that are the primary drivers in the new era of AI is expected to rise significantly, employment in those indirect sectors may actually decline as the application of AI may lead to enhanced productivity and the need for fewer employees. For example, it is estimated that healthcare benefits from the application of AI might not just enhance productivity of health care professions, but lead to better outcomes by improved diagnostics based on AI, finding new drug formulations, and enhanced health monitoring.

5. What are the possible problems with AI?

The age-old mantra in computer science is "Garbage in, garbage out". This is true today as AI programs like large language models are controlled by the data that is used to train them. Without understanding how data sets used in AI are collected and cleaned, there should always be doubt that the system will generate errors. The real issue is when the user does not know the difference between a "good answer" or a "flawed answer" or blindly believes exactly what a chatbot might write. We may well create people who are artificially stupid, that is, those who would blindly accept a fake news report, a fake photo, or doctored data.

Forbes magazine has conducted a survey of students (Westfall, 2023). They found that that 89% of those in the survey had used ChatGPT in doing their homework and 48% had admitted to using ChatGPT for other tasks, such as writing an essay and taking a take home test. If this trend cannot be controlled, will current education systems be circumvented by students finding answers without understanding? Should ChatGPT be banned in schools?

I asked ChatGPT what the problems of AI were. The response was quite good. It included the following issues that were all sourced from Thomas (2024):

1. Job loss due to AI automation;
2. Lack of transparency and explainability:

> AI and deep learning models can be difficult to understand, even for those who work directly with the technology. This leads to a lack of transparency for how and why AI comes to its conclusions, creating a lack of explanation for what data AI algorithms use or why they may make biased or unsafe decisions (Thomas, 2024, section 1. Lack of AI Transparency and Explainability, para. 1).

3. Algorithmic bias caused by bad data. A related issue is: "who cleans the data, and can such cleaning introduce bias?"
4. Privacy violations: AI systems can be used to collect and analyse vast amounts of personal data;

5. Socioeconomic inequality: the benefits of AI may not be distributed equally across society;
6. Market volatility: the use of AI in financial markets can lead to increased volatility;
7. Weapons automatization: autonomous weapons that operate without human oversight.

This list is sobering. We can probably expect greater inequalities due to access issues, increased surveillance and loss of privacy, nefarious uses including generating fake news, and weapons that are self-navigated and decide on targets without human intervention. Without a doubt, there is a clear need for legislation that will help protect people, data, and AI manipulation. This includes an enormous need to prevent the theft of personal identities, protecting consumers and preventing manipulation, and preventing the dissemination of fake news as well as cyber-attacks designed to disrupt services and destroy infrastructure (Acemoğlu, 2022). The impacts of damaging political discourse and the potential to empower corporations and governments against workers requires significant review and new regulations and policies (Acemoğlu, 2022).

To the classical list of impacts given above, I think it is important to consider the environmental impacts of the emerging industry of AI. The biggest of the environmental impacts comes from the increased energy needs required in running new AI models as well as the energy and water needs required in the production of new parallel processing chips. Whereas the development of many classical computer systems, like laptops and servers, have followed a declining energy need per unit of computing, the energy needed in AI applications tends to be expanding exponentially. A recent Forbes article noted that the power consumption increased by 300% in moving to the current generation of GPU chips (Kindig, 2024). It is not that the new generation of GPU chips are not energy efficient, but that the energy demand associated with the computing needs in supporting such programs as LLM is enormous compared to classical computer tasks. Kindig (2024) states that Morgan Stanley has estimated that "global data center power use will triple this year, from ~15 TWh in 2023 to ~46 TWh in 2024". Just in the US alone, it is estimated that data centers in consume between 4.6% and 9.1% of US electricity by 2030 (Electric Power Research Institute [EPRI], 2024). Much of this demand is met by an increase in the use of natural gas power plants and slower retirements of coal-fired power plants in the US. Altogether, such power needs are increasing global carbon emissions.

Concluding comments

Techniques in AI have been developed over the last 70 years. In some areas, the classic Turing test for machine intelligence seems to have been met, but only in specific areas and special applications. Without knowing it, many in Regional Science are already familiar with specialized AI techniques, like Random Forests and A-star. Other techniques, like neural networks, have been used in a number of regional science problems, like trip distribution modeling (Mozolin et al., 2000), forecasting regional employment (Patuelli et al., 2006), willingness-to-pay to reduce road noise (Bravo-Moncayo et al., 2007), and predicting demand in large-scale bike sharing networks (Lin et al., 2018). Techniques like genetic algorithms, simulated annealing, and classic clustering routines, including k-means that have been used in solving Regional Science problems, are also considered to be part of the AI framework. Perhaps the development of large language models has had the greatest impact in the last couple of years. Thus, AI is not new for us. It is just getting better and more accessible.

Students in academic programs have used ChatGPT and other AI systems to craft essays, write papers, and help in answering take-home exams. Plagiarism has been a problem in the past with students lifting information from the internet and using it verbatim in their papers, but now the problem appears to be on steroids. Faculty, researchers, and even publishers need to find ways to deal with this trend. In fact, all academic programs should consider how AI can be integrated into their curriculum with the goal of using AI for the inherent educational possibilities and, at the same time, find ways to prevent plagiarism and encourage students to think beyond what an AI chatbot might tell them. Perhaps AI systems should be required to provide metadata with generated text that could easily be checked and tracked.

Finally, AI will have a great impact on the business models of industry and government. Companies like WM will find ways in which productivity will increase while the needed labour will decrease. Grocery store shelves may be redesigned so that robots can easily restock them. Even checkout stations may be reduced to an automatic conveyor belt that identifies products/prices and leave the last task of a person to bag the groceries. Regional economies will have to adjust with changes in labour demand and training. Altogether, individual companies may be exploring how their organization can be more productive with fewer workers, but the combined impact of these changes is the subject for Regional Science researchers.

References

Acemoğlu, D. (2022). Harms of AI. In J. B. Bullock, Y.-C. Chen, J. Himmelreich, V. M. Hudson, A. Korinek, M. M. Young & B. Zhang (Eds.), *The Oxford Handbook of AI Governance* (pp. 660-706). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780197579329.013.65

Belgiu, M., & Drăguţ, L. (2016). Random forest in remote sensing: A review of applications and future directions. *ISPRS journal of photogrammetry and remote sensing*, *114*, 24-31. https://doi.org/10.1016/j.isprsjprs.2016.01.011

Bravo-Moncayo, L., Naranjo, J. L., García, I. P., & Mosquera, R. (2017). Neural based contingent valuation of road traffic noise. *Transportation Research Part D: Transport and Environment*, *50*, 26-39. https://doi.org/10.1016/j.trd.2016.10.020

Breiman, L. (2001). Random forests. *Machine learning*, *45*, 5-32. https://doi.org/10.1023/A:1010933404324

Chen, N., Christensen, L., Gallagher, K., Mate, R., & Rafert, G. (2016). Global economic impacts associated with artificial intelligence. *Analysis Group*, *1*.

Could AI transform science itself? (2023, September 16). *The Economist*. https://www.economist.com/science-and-technology/2023/09/13/could-ai-transform-science-itself

Cutler, D. R., Edwards Jr, T. C., Beard, K. H., Cutler, A., Hess, K. T., Gibson, J., & Lawler, J. J. (2007). Random forests for classification in ecology. *Ecology*, *88*(11), 2783-2792. https://doi.org/10.1890/07-0539.1

Dechter, R., & Pearl, J. (1988). The Optimality of A*. In L. Kanal & V. Kumar (Eds.), *Search in Artificial Intelligence* (pp. 166–199). Springer New York. https://doi.org/10.1007/978-1-4613-8788-6_5

De Jong, K. (1988). Learning with genetic algorithms: An overview. *Machine Learning*, *3*(2), 121–138. https://doi.org/10.1007/BF00113894

Dueck, G., & Scheuer, T. (1990). Threshold accepting: A general purpose optimization algorithm appearing superior to simulated annealing. *Journal of computational physics*, *90*(1), 161-175. https://doi.org/10.1016/0021-9991(90)90201-B

Electric Power Research Institute. (2024). *Powering intelligence: analyzing artificial intelligence and data center energy consumption* [Brochure]. https://www.epri.com/research/products/3002028905

Flack, P. (2012). *Machine Learning: the art and science of algorithms that make sense of data*. Cambridge University Press.

Hartree, D. R. (1949). *Calculating instruments and machines*. Cambridge University Press.

Kindig, B. (2024, June 20). AI power consumption: rapidly becoming mission-critical. *Forbes*. https://www.forbes.com/sites/bethkindig/2024/06/20/ai-power-consumption-rapidly-becoming-mission-critical/

Lin, L., He, Z., & Peeta, S. (2018). Predicting station-level hourly demand in a large-scale bike-sharing network: A graph convolutional neural network approach. *Transportation Research Part C: Emerging Technologies*, *97*, 258-276. https://doi.org/10.1016/j.trc.2018.10.011

Lovelace, A. (1842). Translator's notes to an article on Babbage's Analytical Engiro. In R. Taylor (Ed.), *Scientific Memoirs, 3* (pp.691-731). London. https://archive.org/details/scientificmemoir03memo

MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In L. M. Le Cam & J. Neyman (Eds.), *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, *Volume 1*, 14 (pp. 281-297).

McCallum, J. C. (2023). *Price and Performance Changes of Computer Technology with Time*. Jcmit. https://jcmit.net/diskprice.htm

Mozolin, M., Thill, J. C., & Usery, E. L. (2000). Trip distribution forecasting with multilayer perceptron neural networks: A critical evaluation. *Transportation Research Part B: Methodological*, *34*(1), 53-73. https://doi.org/10.1016/S0191-2615(99)00014-4

Naveed, H., Khan, A. U., Qiu, S., Saqib, M., Anwar, S., Usman, M., Akhtar, N., Barnes, N., & Mian, A. (2023). A comprehensive overview of large language models. *arXiv preprint arXiv:2307.06435*. https://doi.org/10.48550/arXiv.2307.06435

Patuelli, R., Reggiani, A., Nijkamp, P., & Blien, U. (2006). New neural network methods for forecasting regional employment: an analysis of German labour markets. *Spatial Economic Analysis*, *1*(1), 7-30. https://doi.org/10.1080/17421770600661568

Pearl, J. (1988). On the discovery and generation of certain heuristics. In *Readings from the AI Magazine* (pp. 58–66). American Association for Artificial Intelligence.

Qi, Y. (2012). Random Forest for Bioinformatics. In C. Zhang & Y. Ma (Eds.), *Ensemble Machine Learning: Methods and applications* (pp. 307-323). Springer. https://doi.org/10.1007/978-1-4419-9326-7_11

Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, *65*(6), 386-408. https://doi.org/10.1037/h0042519

Tewari, G. (2022, May 6). *The Future Of AI: 5 Things To Expect In The Next 10 Years. Forbes*. https://www.forbes.com/sites/forbesbusinesscouncil/2022/05/05/the-future-of-ai-5-things-to-expect-in-the-next-10-years/?sh=7e68aabe7422

The R-project for statistical computing. (n.d.). R-project. https://www.r-project.org/

Thomas, M. (2024). *12 risks and dangers of artificial intelligence*. Builtin. https://builtin.com/artificial-intelligence/risks-of-artificial-intelligence

Turing, A. (1950). Computing machinery and intelligence. *Mind*, *59*(236), 433–460. https://www.jstor.org/stable/2251299?origin=JSTOR-pdf

Tyralis, H., Papacharalampous, G., & Langousis, A. (2019). A brief review of random forests for water scientists and practitioners and their recent history in water resources. *Water*, *11*(5), 910. https://doi.org/10.3390/w11050910

Westfall, C. (2023, January 28). Educators Battle Plagiarism As 89% Of Students Admit To Using OpenAI's ChatGPT For Homework. *Forbes.* https://www.forbes.com/sites/chriswestfall/2023/01/28/educators-battle-plagiarism-as-89-of-students-admit-to-using-open-ais-chatgpt-for-homework/?sh=2378b29b750d

Zeng, W., & Church, R. L. (2009). Finding shortest paths on real road networks: the case for A*. *International journal of geographical information science*, *23*(4), 531-543. https://doi.org/10.1080/13658810801949850