The Impact of Machine Intelligence in Tourism Analysis and Management

In the popular imagination, Artificial Intelligence (AI) is envisioned as the proliferation of self-aware anthropomorphic robots at some point in the distant future; however, machine intelligence has already been deployed in varying degrees to a wide range of relatively mundane tasks, not in competition with humans but rather as a tool. Few humans as yet interact with robots or ride in self-controlled aerial drones, but the daily lives of virtually everyone alive today are being increasingly affected by AI systems that can recognize speech or images, manipulate prodigious databases, or analyze patterns of behavior.

Since the 1960s, AI has been used to mine for profits, streamline business environments, and cultivate customer buying behavior. Presently, AI is present in almost all forms of enterprise, from banks, health care providers, and insurance companies to manufacturers, utility companies, and website marketers. In the corporate workplace, AI is widely used for complex problem solving and decision support in disciplines ranging from financial management to forecasting. It is a characteristic of modern economies that data increases exponentially literally day by day. Analyzing this growing body of information and formulating quick and accurate decisions requires complex cognitive abilities available only to machine intelligence.

For consumers residing in developed countries, from the moment the alarm clock wakens them each morning to the moment they turn out the light on the nightstand each night, AI already affects them more than any scientific or cultural phenomenon has influenced people at any previous time in history—with the possible exception of nuclear weapons. As the future unfolds, the impact of AI on the personal life, work, leisure time, and health of every individual in the workplace will continue to multiply. The phenomenon will occur without a sudden revelation, yet inexorably machine intelligence will become indispensable to the existence of virtually everyone on the planet. The event will simply happen, and most people will accept the consequences as a normal, and perhaps necessary, aspect of their everyday lives.

No economic process will be more profoundly affected by the impact of machine intelligence than tourism analysis and management decision-making. Globally, tourism is the second largest economic influence, eclipsed only by the weapons trade. The socio-economic impact of the tourism economy touches not just those workers employed in tourism-related enterprises such as airlines, lodging establishments, and food and beverage operations, but also ancillary groups in fields such as transportation, manufacturing, and retail services. Virtually every government in the world derives a portion of its gross national product from tax revenues paid not only by

tourism enterprises but by their employees and related service entities which benefit from the redistribution of tourism revenues at various levels of the local economy.

Technology advancements have historically obsoleted workers in some fields, and the trend inexorably will increase with the introduction of more machines that are more intelligent. New entrants into the workforce will require training to program, maintain, and work with intelligent machines. Nevertheless, the fear of technology-related job loss concentrated at the lower end of the income scale is very real. As AI development and deployment accelerate in the future, not only white-collar workers but also executives at all levels of management will be affected by the proliferation of intelligent machines.

One reason for the increasingly prevalent adoption of Artificial Intelligence is the analytical power made possible by artificial neural networks (ANNs) using fuzzy logic. In some applications, such as fraud detection, AI has already become the technology of preference. In addition, the use of neural networks has become an established methodology for pattern recognition, particularly of images, data streams and complex data sources. Such networks have emerged as the platform for the majority of data-mining tools currently in use.

Besides fraud detection, machine intelligence is used to analyze and manage the buying habits of consumers, and manage production and delivery schedules based on fluctuating customer demand. Eliminating downtime due to unpredictable scheduling changes increases productivity.

The Logical Distinction

What most distinctively separates Homo Sapiens from other species is our innate capacity for logic. The human brain is structured in such a way that it can analyze and solve quandaries by thinking problems through. Basically, logic consists of recognition of cause-effect relationships in nature: "If this happens, then that will certainly happen." The end result of logical analysis is a decision of some sort.

On an elementary level, such analysis might take the form of something as simple as "If my connecting flight is cancelled, I will miss my onward flight." Your decision might be to book a flight with a different carrier enabling you to board the onward flight to your destination." This type of instant "if/then" shortcut is an example of heuristic reasoning.

On a more sophisticated level, a physician may detect that a patient's EKG has a particular pattern characterized by a jump between two parts called the S wave and the T wave. Upon testing, the patient's blood is found to contain elevated levels of a substance known as SGOT. The

physician, therefore, concludes the patient has suffered a myocardial infarction--a heart attack. This scenario is also an example of heuristic reasoning.

All humans exercise heuristic reasoning in their daily lives, and all logic can be reduced to such "if/then" rules, known formally as shallow reasoning. Virtually all human expertise can be reduced to shallow reasoning, using heuristic shortcuts, thus bypassing what might otherwise be a tedious and time-consuming exercise. For example, it is enough to know that if a patient's EKG result depicts an elevated S-T segment, and his blood test results indicate a high level of SGOT, he has experienced a heart attack. Why? These two findings rarely occur in concourse unless a myocardial infarction has occurred. Such shallow reasoning is the basis for the science of heuristics.

Other animals besides humans might seem to exercise logic at times; for example, when a cheetah evaluates a herd of wildebeest before deciding how and which one to attack, or when the lead goose of a flock in flight traces a transoceanic route to their mating grounds using landmarks as a guide. In such animals, actions that we humans might perceive as decisions are made on the basis of stimulus-response relationships. For every stimulus, there is an exact and predictable response, from the simplest particles that make up all matter to the most complex organisms.

In an Artificial Intelligence program, the logical rules that are required to arrive at an unambiguous solution are expressed in an algorithm. As an example, the following is a simplified rules-based algorithm for a program that predicts the effect of oceanic temperature on the stock prices of resort holdings in the tropics.

```
oceanic_temperature(rises) > cyclone_activity(increases).
oceanic_temperature(flat) > cyclone_activity(flat).
oceanic_temperature(declines) > cyclone_activity(flat).
cyclone_activity(increases) > tourist_visitation(decrease).
cyclone_activity(_) > tourist_visitation(static).
tourist_visitation(decrease) > stock_prices(decline).
tourist_visitation(increase) > stock_prices(rise).
tourist_visitation(flat) > stock_prices(flat).
```

The object "oceanic_temperature" may have one of three values: rise, flat, or decline. The text on the left of the ">" symbol is the predicate, and the part on the right, the conclusion. According to this algorithm, when oceanic temperatures rise, the risk of cyclone activity increases, causing the number of tourist visitations to decrease and stock prices to decline. If oceanic temperatures

remain flat or decline, cyclone activity remains flat, producing no effect on the stock prices of resort holdings.

Only the rules that result in a successful conclusion are processed. If a rule fails, the program backtracks and processes the next rule, continuing in this fashion until an unambiguous solution is determined. An unambiguous--or deterministic--solution is an essential quality of a rules-based algorithm. For example, it is insufficient to state that "stock prices will be affected." The program must be able to state definitively that stock prices will rise, decline, or remain flat.

Using historical data and certainty factors, an AI program can even more definitively predict the amount of increase or decline:

Oceanic temperature rises 0.5% > Stock prices will increase 38% (Certainty factor: 55)
Oceanic temperature rises 0.5% > Stock prices will increase 24% (Certainty factor: 27)
Oceanic temperature rises 1.0% > Stock prices will increase 52% (Certainty factor: 34)
Oceanic temperature rises 1.0% > Stock prices will increase 38% (Certainty factor: 24)

In this example, the program predicts a 55% certainty that stock prices will increase by 38% if oceanic temperatures rise by one half of one percent, and a 27% certainty that stock prices will increase by 24%.

Of course, the more historical information that is available, the greater the accuracy of the predictions. Even so, history does not always repeat itself, especially economic history. Certainty factors used in AI programming range from simple averages to values derived from complex trend analysis.

Using this method of solving problems based on human rules of logic, machine thinking mimics the processes that a human brain uses to assimilate and analyze information, form strategies, and decide on a course of action. An AI program tests each premise that is presented to it, to determine if the premise is true or false. If it is true, the premise becomes a fact and leads to the testing of other premises. If the premise is false, the program leaps backward and begins testing an alternative premise. Logical reiteration in the pursuit of deterministic goals constitutes the essence of machine thinking.

The Ascent of AI in Tourism

The basis for all free-market enterprise consists of three inter-connected influences: cost, competition, and demand. Commerce originated in ancient times on the basis of products that

fulfilled identifiable needs, such as tools, food, or apparel. Today, even products that seem to fulfill only psychic or imaginary needs have a basis in human behavior that can be quantified. For example, "moon rocks," which arguably serve no practical purpose, fulfill a psychic need to experience the pleasure of possessing--or, perhaps, merely the act of buying--such a product.

Pricing is as important as a product that fulfills identifiable needs. From the user's standpoint, pricing determines not only how much a product costs, but also its value--the quality in relation to the price. From a marketer's perspective, pricing determines how much money will be made from sales and, in the final analysis, whether or not the business will be successful. In every transaction, users also pay other costs besides the monetary price. Product costs fall into the following categories:

1. Monetary costs
2. Location costs
3. Time costs
4. Sensory costs
5. Psychic costs

The monetary cost is based not only upon the purchase price but also upon additional charges such as taxes and freight. However, tourists also have to pay location costs such as air and land transportation. The time spent traveling represents another type of cost. In addition, some tourists may experience a sensory cost in adjusting to unfamiliar surroundings, and a psychic cost in the form of stress, discomfort, or fear.

To some extent, all these types of costs must be taken into account when prices are set. An AI pricing strategy consists of the planning processes that are used to set a monetary price. Pricing strategies may be as simple as adding a fixed percentage or amount to the cost, or as complicated as a yield management policy.

The seller pays a cost for every product that is sold to a user. To produce a profit, the price received from the user must be higher than the cost paid by the seller. For any business to be profitable, the total revenues received from sales must exceed the total manufacturing costs plus the operating costs of the company. Thus, prices must be high enough to compensate for products that are not sold. Yet, if prices are too high, customers will be discouraged from buying. If prices are too low, the business sacrifices profits that it might otherwise have realized.

Forecasting--the analytics devoted to estimating future events--is fundamental to all aspects of management. The goals of forecasting are to reduce uncertainty and to provide benchmarks for

monitoring actual performance. AI forecasting methods are used to predict product demand, inventory levels, worker requirements, surplus rates, and raw materials.

The classic business forecasting dilemma is the scenario known as the Baker's Dozen. Assume a baker bakes 12 loaves of bread and, by the end of that day, eight loaves are sold, leaving a surplus rate of 25 percent (four loaves). The following day, the baker bakes eight loaves, and at closing time all the loaves have been sold. However, the possibility exists that if more loaves had been baked, some of the extra loaves might have been sold as well. The problem, from a forecasting perspective, is to determine the optimal number of loaves required to maximize sales while minimizing surplus, as cost, competition, and demand fluctuate. This scenario is the basis for virtually all quantitative methods used to compute airline seat quotas and hotel room availability.

With the rise of online purchasing and social media, AI began to be applied to collecting, processing, and analyzing information on a massive scale, giving rise to the pop-science terms Big Data, analytics, and data mining.

Mining for Data

Big Data, in a general sense, refers to a very large and rapidly growing volume of information, beyond the ability of conventional computer technology to process efficiently. In a typical enterprise environment, the information results from internet usage, social media, and computing task such as word processing or accounting. Analytics and data mining are the basic tools of harnessing, harvesting, and interpreting the information.

In logic, analytics refers to the logical science of analysis. In data mining, analytics involves the entire process of collecting, analyzing, and drawing conclusions from a growing body of data. From a business perspective, analytics enables management to view, at any point in time, how customers or users utilize a product or service, as well as demographic and behavioral traits, such as age, gender, buying habits, personal interests, and so forth.

Data mining gives companies a precise view of how particular segments of the customer base react to a product or service, and proposes changes consistent with those findings. In addition to exploring customers' buying patterns, analytics enables management to react much more quickly to marketplace changes. The results enable management to refresh marketing models on the fly, based on each new incoming piece of customer information, to create a more targeted offering.

Data mining techniques are largely identical to those used by computer viruses. In fact, programs that infect computers with programs designed to provide feedback about the internet

browsing habits of users are classified as malware in most antivirus databases, as are programs designed to deliver targeted advertisements. Like a virus, a data mining program secretly intrudes on a user's computer without permission. The program seeks out specific locations on the computer, such as email folders and places where downloads and preferences are stored, and then harvests pertinent data about the user. The program's instructions are in the form of "scripts" that waken automatically when a user accesses the web page to which they are attached.

Data mining scripts are now present on most e-commerce web pages, as well as the sites of many governments and institutions. The information they extract from users' computers is used to create custom product offerings, compile demographic data, and guard against fraud, but also to generate email lists for target marketing ("spam").

Most merchandising websites track customer purchases so that the most relevant products can be displayed each time the user accesses the site. Consumers' buying habits are often sold to other marketers for use in email campaigns or pop-up ads.

Merely reaping information and storing it in a big database are, by themselves, of minimal use. To be meaningful--and, of utmost importance to businesses, also profitable--the mined data must be sorted into identifiable patterns and relationships, or behaviors. Once these behaviors are defined, they then need to be validated. The Data Mining process thus occurs in three separate stages: exploration, pattern definition, and validation.

From a marketing perspective, the resulting information can be used to predict future behaviors, such as purchasing similar or related products. The data can also reveal personal traits, such as age, gender, occupation, hobbies, entertainment preferences, and health status. For example, analytics would predict that a computer user who browses for women's hair coloring and anti-wrinkle cream is most likely a female 40 years or older. Based on statistics about 40-year-old females, the program might predict that the user would also be interested in a certain style of apparel. Similarly, a 67-year-old traveller browsing for airline flights to a tropical destination might be considered a prime candidate for a cruise offering.

Left to its own devices, a data mining system may develop subtle patterns that might otherwise evade normal expectations. A classic case is a supermarket that prominently juxtaposed displays of beer and diapers, creating a seemingly humorous correlation between alcohol consumption and urinary urgency. The displays were actually based on predictive modeling derived from data mining, showing that males who have infants and shop on Saturday tend to buy both types of products. The model was validated by a surge in sales of both items.

The focus of Data Mining systems is on defining behaviors that are useful in generating reliable predictions. Web tracking—identifying and correlating websites visited by a single

user—has become a common instrument for mining data about internet users. In an experiment performed by researchers at Princeton University, Google Analytics web tracking scripts were found on more than 70% of the websites they studied. DoubleClick, an ad-serving system also developed by Google, was found on half of the sites. The researchers used software that was also developed at Princeton, called OpenWPM, to robotically survey one million sites, logging any tracking technology that was detected.

When embedded on a large number of sites, web tracking scripts can compile detailed profiles of individual users navigating the internet. Among the documents leaked by Edward Snowden was a memo indicating that the U.S. National Security Agency tapped into Google's web tracking system to identify potential suspects for surveillance.

In theory, data mining techniques are based on predictive analysis using behavioral modeling and statistics. However, methods such as Google's ad-serving technique actually have built-in biases. Advertisers choose the types of users to target. To streamline the validation process, businesses often rely on heuristic shortcuts; for example, a male who browses female swimsuit photos might be presented, without any statistical basis, an ad for an erectile dysfunction aid.

Predictably, the practice of data mining has led to controversy over privacy and ethical concerns.

Conclusion

With increasing development and deployment of Artificial Intelligence and data mining technology, the focus of tourism analysis is necessarily shifting from promotion-driven to behavior-driven. Whereas traditional tourism analysis has relied primarily on motivating buying behavior, machine intelligence will enable tourism analysts and providers to target specific consumers based on data mining and validation, using personal and demographic traits of best-odds buyers. However, Machine intelligence will also produce inexorable job losses in food and beverage service, transportation, and receptionist positions, laying the seeds for potential social problems and unrest.

Objectively, the impact of machine intelligence will have both positive and negative outcomes on host societies. Arguably, everyone alive today will be impacted to some degree, for better or worse, by the inevitable advent and rapid spread of Artificial Intelligence.

Dennis L. Foster, Ed. D.