N. C. Saxena (2023)
Public Enterprise
Volume 27, 2023
https://doi.org/10.21571/pehyj.2023.2701.03

International Center for Promotion of Enterprises
ICPE

Using Machine Learning to improve the performance of Public Enterprises

Naresh C Saxena
Independent researcher, Canada
drncsaxena@gmail.com

Abstract

Governments all over the world are aiming to improve the performance of their public enterprises. The goal of this method is to help public enterprise CEOs transition from a bureaucrat mentality to a leader mentality. Profitable operation of public sector enterprises through cutting edge technology is included in this effort. Machine learning is a subset of Artificial Intelligence using mathematics and statistics to learn from data and, sometimes, make predictions. is able to analyse massive quantities of data, and identify patterns in that data. In doing so, it provides new sources of inspiration and innovation. Machine learning opens new possibilities to drive positive change; helping computers learn from data to create models humans would not be able to build otherwise. The challenge is to maintain trust and deliver real benefits for everyone. To be useful, the applications of machine learning will, require close interaction between IT professionals and public enterprise executives.

KEYWORDS: Machine learning, Public enterprise, Artificial intelligence

Introduction

Can machines think? The question has baffled the philosophers and laymen alike for a long time (Mays, 1952). At the dawn of the computer age, British mathematician Alan Turing, the designer of programmable computing device, began a classic article titled ‘Computing Machinery and Intelligence’ with this question (Turing, 1950). He did not answer the question and discarded it as too meaningless to deserve discussion. Turing was not committing himself to the view that to think means thinking like a human (Dennet, 2008). However, in a prescient section entitled ‘Learning Machines’, Turing (1950) had anticipated machine learning (Harnad, 2018).

Machine learning (ML), a branch of artificial intelligence, relates to the construction of systems that can learn from data thus giving computers the ability to learn without being explicitly programmed. ML lies at the intersection of computer science, engineering, statistics and often appears in other disciplines, especially business. ML techniques can be applied to many problems that need to interpret and act on data. Arthur Samuel, whose papers in the 1950's on the subject are still worth studying, defined ML as a field of study that gives computers the ability to learn without being explicitly programmed. Samuels (1959) wrote a checkers playing program, had the program play over ten thousand games against itself and work out which board positions were good and bad depending on wins and losses. Forty years later, this definition was formalised as: "A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E" (Mitchell, 1997, p. 2). In the case of the checkers learning problem, task T is playing games, performance measure P is proportion of games won and experience E is playing practice games. Thus, searching for definitions in cognitive terms favoured by philosophers, e.g., "Can Machines think?" has given way to formulating definitions in operational terms, i.e., "Can it do what we, the thinking entities, can do?"

ML and data mining have so much overlap that it is not easy to distinguish between the two disciplines. ML is more focused on prediction, based on known properties learned from the training data. Data mining focuses on the discovery of (previously) unknown properties in the data. But, ML also employs methods like "unsupervised learning" where there is no training data and data mining uses many ML methods, albeit sometimes with a slightly different goal in mind. In ML, performance is usually evaluated with respect to the ability to reproduce known knowledge, while in Data mining, the key task is the discovery of previously unknown knowledge. The research in the two fields is converging. The best known conferences in the two fields, the European Conference on Machine Learning (ECML) and European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD), were co-located in 2001. In 2008 the conferences were merged into one conference and the division into traditional ML topics and traditional data mining topics has been removed. In the 1980's the field of ML was largely empirical and ad hoc. Over a period of quarter of century, a series of models have been developed that combine technical depth and broad applicability, thus giving the field theoretical adequacy as also effective applications.

Applying Machine Learning

Last 15 years have seen spectacular applications of ML. Siri in iPhones predicts the meanings of human voices and tries to provide the desired answers. Photo album in Facebook recognises faces to be tagged in photos. LinkedIn predicts who the users want to connect with. Since 2013, the Excel program of MS Office can comb very large amounts of data to find meaningful patterns. For example, with bi-directional access to live X (formerly Twitter), it can scan millions of Twitter posts and create charts to show which product is getting the most buzz. A new version of MS Outlook employs ML to review the e-mail habits of users to see whether a user wants to read each message that comes in. Casinos are able to use face recognition to bar entry to card counters. Driverless cars are allowed in several cities in the US. Less glamorous but equally useful applications are found in the field of operations and logistics like supply chain demand forecasting. For example, forecasting the manufacturer's demand under information asymmetry – the bullwhip effect – can be tackled through advanced ML techniques (Carbonneau et al., 2008).

Big business have been placing big bets on ML applications. IBM, seeing an opportunity in data-hunting services, created a Business Analytics and Optimization Services group in April 2009. Microsoft has designed ML software that can trawl internal corporate computer systems with a view to predict which software applications are most likely to fail when seemingly unrelated programs are tweaked. It has combined even more ML with its cloud computing system, called Azure, to rent out data sets and algorithms so that businesses can build their own prediction engines. Drug companies are using ML to understand spread of diseases. The ML approach is being found useful for finding business partners and building reciprocal relationships (Mori et al., 2012). There is hardly any area of business which is untouched by ML. It is viewed as the foundation of a better and smarter future. In the academic world of computer science, enthusiasm for ML has been growing at all levels for the last quarter of a century. In 2010, the Turing Award, recognised as the Nobel Prize of Computing, went to Leslie Valiant, an innovator in this field for his transformative contributions to the theory of computation, including the theory of probably approximately correct (PAC) learning. Next year, at Udacity (a start-up launched by Stanford University) an online course on ML, taught by Peter Norvig, Google's director of research, attracted a record 160,000 students.

ML programmes may not be able to answer questions relating to meaning of life but are getting better at answering queries relating to enterprises. The programmes, however, are yet to understand nuances of languages and machine translation is improving at a very slow place. Even so, confidence in the capabilities of ML is growing. The idea of algorithms sucking up and sorting through the otherwise useless garbage of data is appealing and it is easy to be attracted to the mystique of ML. It seems like a new world where thinking need not be constrained by what may or may not be possible. ML seems to be becoming like motherhood and apple pie. Nobody is opposed to it. The reality, however, is a bit more complex. There is a huge unmet demand for applications whereas much of the output is either not being used or is unusable. Netflix paid out the prize money but never used the winning algorithm because extraneous factors came into play. Like any breakthrough technology, ML involves some forethought and discipline before being let loose in the enterprise, especially when it affects public welfare. This article looks at the cutting edge applications of ML in the public sector. The key questions to ask are: "What kinds of business problems do we want to solve?" "What questions do we want to be able to answer?" and finally, "How much would they be worth if we could answer them?"

Transforming the public sector

Machine Learning has emerged as a transformative technology in the public sector, revolutionizing the way public enterprises operate. Its applications span a wide range of domains, bringing about greater efficiency, improved decision-making, and enhanced public services. There are two interconnected areas where ML is making a significant impact in the public sector viz. predictive analytics for improved decision-making and improved public services.

Predictive analytics involves using historical and real-time data to anticipate future trends, needs, and potential issues. For example, in the realm of healthcare, ML models analyse patient records and medical data to forecast disease outbreaks, enabling health authorities to mobilise resources, disseminate preventive measures, and allocate medical staff more effectively. This not only improves response times but also saves lives and resources.

In the context of management of urban areas, ML has contributed to creation of smart cities where, among other activities, it predicts traffic congestion and public transportation demand, allowing city authorities to optimise routes, schedules, and infrastructure investments. ML is used in urban planning to optimise traffic flow, manage public transportation systems, and plan sustainable cities. This can lead to reduced congestion, lower emissions, and more liveable urban environments.

ML improves public transportation systems by predicting ridership, optimising routes, and scheduling services more efficiently. This ensures that public transit resources are allocated in a way that meets passenger demand while minimising costs. Predictive analytics in this field has improved the efficiency of public transport systems and reduced congestion, making citizens' daily commutes less inconvenient.

ML is employed in financial sectors for risk assessment and fraud detection. Algorithms can analyse financial transaction data to identify patterns associated with fraudulent activities or unusual financial behaviour. More importantly, ML algorithms can protect citizens from identity theft.

By harnessing the power of predictive analysis, public organizations can allocate resources more efficiently, enhance public services, and ultimately better serve the needs of their constituents. ML can greatly assist public sector agencies in providing better services to citizens by enhancing efficiency, improving user experiences, and optimizing resource allocation.

ML-powered chatbots and virtual assistants are increasingly being used in government websites and platforms. These AI-driven systems can answer common citizen queries, provide information, and guide users through various processes, such as applying for permits or accessing public services. They operate 24/7, improving accessibility and responsiveness.

ML analyses citizen data to offer personalized recommendations and services. This personalisation enhances the user experience and increases citizen engagement. For instance, in healthcare, ML can help agencies provide tailored health information and treatment plans. In education, it can offer personalized learning paths for students, tailor learning experiences for individual students and personalising curricula to meet their unique needs and abilities. It also assists in administrative tasks, such as student enrolment and assessment scoring.

Typically, public enterprises serve a large number of people, especially in developing countries. ML automates routine administrative tasks, such as data entry, document processing, and form filling. This reduces the administrative burden on government employees, allowing them to focus on more complex and value-added tasks. It also speeds up service delivery and reduces errors. ML processes and analyses large datasets to provide insights that inform decision-making. Public agencies can use ML to forecast demand for services, allocate resources more efficiently, and identify areas that require additional attention or investment. This leads to better planning and resource management.

Public enterprises lay great emphasis on non-discrimination and affirmative action. ML improves accessibility for citizens with disabilities. For example, Natural Language Processing (NLP) and speech recognition can help translate spoken or written text into accessible formats, making information and services available to a wider range of users.

Public enterprises have to remain cognizant of public sentiments. ML can analyse citizen feedback and sentiment from various sources, such as social media, surveys, and public comments. This helps agencies understand public sentiment and concerns, allowing them to adapt and improve services accordingly. Responding to emergencies is another function of public enterprises. ML aids in forecasting and responding to emergencies more effectively. It can predict natural disasters, manage resources during crises, and analyse social media data to identify incidents in real-time, helping first responders react promptly. ML can improve healthcare delivery by assisting in diagnostics, patient monitoring, and predictive analytics. Public health agencies can use ML to predict disease outbreaks and allocate resources accordingly, ensuring better healthcare services to the population (Woo, 2013).

ML can significantly enhance efficiency of resource allocation in the public enterprises by leveraging data-driven insights and predictive analytics to ensure that resources are distributed more efficiently and effectively. ML models can analyse historical data and patterns to predict future demand for public services or resources. For instance, in healthcare, ML can forecast patient admissions, helping hospitals staff appropriately and allocate resources like beds and medical equipment based on expected demand. ML can optimise healthcare resource allocation by predicting patient needs and disease outbreaks. This allows hospitals to allocate staff, beds, and medical supplies more effectively, improving patient care. ML is transforming patient care through predictive diagnostics, personalized treatment plans, and improved public hospital management.

Increasingly, public enterprises are using ML to analyse budget data and historical spending to make data-informed decisions about how to allocate resources across various programs and departments. This helps in ensuring that funds are distributed to areas that need them most and are likely to have the greatest impact.

ML can be a valuable tool in emergency response and disaster management. It cannot predict natural disasters but can assess the potential impact, allowing agencies to pre-position resources and personnel in high-risk areas to respond more rapidly and effectively. ML can be used to monitor and manage environmental resources, such as water and energy. Predictive models can forecast usage patterns and identify areas where conservation efforts are needed (De Lucia et al., 2020).

An illustrative list of popular applications and the ML techniques therein currently in use by public enterprises are given in table 1.

Table 1: Examples of common ML applications in Public Enterprises. (Techniques most commonly used in bold)

ML Problem	Applications		
	Public Finance	Public services	Other
Classification: identify the maximally distinguishing attributes using training data set	Risk classification (**RI – Rule Induction**)	Market segmentation (**RI – Rule Induction**)	Churn management (**CBR – Case Based Reasoning**)
Prediction: Finding probable future values or distributions of attributes	Forecasting default (**RI – Rule Induction**)	Customer reaction to promotions (**GA – Genetic Algorithms**)	Network behaviour (**NN – Neural Networks**)
Detection: Identifying causes of irregular patterns	Suspicious transactions (**NN – Neural Networks**)		Cost estimation (**NN – Neural Networks**)
Association: identifying rules governing relationships		Market basket Analysis (**VS – Visualisation**)	Similarity assessment (**ILP – Inductive Logic Programming**)

Unlearning and Relearning

Concepts in ML are better defined than in other branches of Artificial Intelligence, e.g., reasoning, expert systems, planning, natural language, robotics and vision (Hoffman, 2011). But, it does not seem to be headed for continued evolutionary progress in the standard mode. It is full of uncertainty and excitement. Some of the excitement has been due to the recent explosion of digital data from new realms – sensor signals, surveillance tapes, social network chatter, opening of public records and so on. Data overload which is seen as a problem in some disciplines is considered a boon by scholars in the field of ML as they hunt for meaningful patterns in these troves.

ML papers mostly describe a new algorithm and the research relies heavily on repositories - collections of databases, domain theories, and data generators for empirical analysis of ML algorithms. The new algorithm's behaviour is illustrated on synthetic data sets and then the paper reports results on a collection of standard data sets available in an archive. The most commonly used archive is UCI maintained by University of California, Irvine. This repository focuses on tasks like classification and regression and has made cross-domain studies in this area straightforward and commonplace leading to more than a thousand papers citing this repository. At the same time, it has had a harmful effect in the long term as the researchers are neglecting complex tasks like reasoning, problem solving and language understanding (Langley, 2011). There is no consensus within the ML community about what role UCI data set serves other than helping scholars churn out research papers. They are less useful than synthetic data, since the researchers do not control the process through which data are generated and they cannot be called real world data as they are not associated with real world users, experts or operational systems. Worse, they have depreciated the value of formulating problems and defining features. This has been a major concern for public enterprises.

Most public enterprises do not know how best to use accumulated data for improving decision-making. With the exponential growth of technology, public enterprises not only need better tools to understand the data they currently have, but also to prepare themselves for the data they will have. Data has piled up for a long period of time and public enterprise executives are afraid to get rid of it. They know that much of the data is useless but have no idea what use could be made of it in the future. Apart from normal decision making, the data could also be used in litigation and government investigation. This has led to information governance-related paralysis in the public enterprises. The cost of data storage is significant if one includes the cost of all the overheads of managing the different systems and the clogging of the networks as data moves back and forth. ML combined with an iterative workflow that leverages a small amount of human input to identify relevant information can be used to address data hoarding. On the legal side, some software applications have come in the market using which public enterprises can quickly find information that is safe for deletion and feel confident that they have made the right decision. As data accumulation continues to accelerate, more research is required in this direction. Some kind of data-management scheme has to be applied as data arrives so that wasteful volumes of data do not start accumulating and useful data do not fall through the cracks.

Experimental evaluation revolves around performance metrics. The race to improve performance metrics, however abstract, has meant that there is little interest in research on knowledge-generating mechanisms in the public enterprises. At the same time availability of large data sets meant lack of interest in using background knowledge to improve learning rate, marginalising research on learning faster from fewer experiences. A researcher claims victory when her algorithm makes an improvement in say accuracy in classification across certain datasets. The methodologies of showing improvement are open to question. What the meaning of that improvement is in a real world situation is rarely examined in the public enterprises. Moreover, mindless competition among the algorithms reveals little about the sources of power or the effects of domain characteristics. Public enterprises need to use ML diagnostically (what data to use) and if need be opportunistically (finding new data sources). Public enterprises operating in federal countries need to be careful about aggregate data from different states, because laws in different states could be different.

Pitfalls for public enterprises

The need to explicitly programme devices limits our capacity for innovation. ML allows us to build models that, after being trained with large data sets, can help predict real-world outcomes and uncover fresh insights at speed and at scale. Since the term 'machine learning' was coined in 1959, the technique has been used to build models that lie at the core of daily applications. However, the topic is not without controversy.

Many public enterprises use terms such as 'predictive analytics' or 'behavioural insights' instead of ML. This could be because of technophobia or because of the fear that when fraught with bias (the 'dark side'), ML can be dangerous. One reason is simply that probability is a factor in the algorithms. There is always a potential that they will be off, therefore because they make so many forecasts, it is possible that some of them will be inaccurate. The amount and quality of the data used to train the algorithms, the particular ML technique selected, and whether the system uses only explainable algorithms (i.e., algorithms that allow humans to explain how they arrived at their decisions) may all affect the likelihood of errors, which may prevent the system from achieving maximum accuracy.

Another reason is concept drift. The relationship between the inputs the system uses and its outputs is not stable over time or may be misspecified. When a data is trained in normal times, it would give faulty predictions in covid times creating chaos at a time when the performance of public enterprises needs to be most efficient. Concept drift is more frequent in countries with a diverse population. When the training data comes from one set of population, it may give faulty results when used on another set.

Even if the patterns the algorithm learned are stable and there is no concept drift, covariate shift can occur. For example, a certain policy or procedure may be designed using data from large urban area. With samples from certain sociodemographic groups who have underlying characteristics not commonly seen in rural areas. Such disparities may be discovered only when more errors are made while in actual use than it did during testing. Given the diversity of applications and the pace at which they are changing, it is becoming increasingly difficult to foresee what will happen in the environment that systems operate in, and no amount of data can capture all the nuances that occur in the real world.

Yet another reason ML can make inaccurate decisions has to do with the complexity of the overall systems it is embedded in. Often the data are of images. The quality of any analysis depends on how clear the images provided are, the specific algorithm used, the data that algorithm was trained with, whether the person inputting the images received appropriate instruction, and so on. With so many parameters, it is difficult to assess whether and why a mistake was made, let alone be certain about its behaviour.

While all enterprises are expected to act in accordance with law, public enterprises are also obliged to act in accordance with public policy which evolves with time. The governments as also the public views the senior management as the final decision-maker and do not apply liability to the software makers. As more black-box or autonomous systems make recommendations with much weaker involvement by the managers, this would cause serious problem in determining the liability.

Ethics is an important part of management of public enterprises. Products and services that make decisions autonomously will also need to resolve ethical dilemmas — a requirement that raises additional risks and regulatory and product development challenges. Scholars have now begun to frame these challenges as problems of responsible algorithm design (Babic et al., 2021). They include the puzzle of how to automate moral reasoning. Such designs are important for public enterprises.

Application of ML in public enterprises cannot follow the pattern of the private sector. Public enterprises lay emphasis on affirmative action. Relying on past data, ML can make a guess that certain categories of people are not reliable employees or borrowers; whereas these categories of people need to be helped as a matter of public policy. Done right, ML can improve efficiency, help prioritise resources, and gain valuable insights we might have never known. But public enterprises need to be aware of pitfalls. The problem is compounded by the multiple and possibly mutually incompatible ways to define fairness and encode it in algorithms.

Conclusion

ML algorithms have gained an aura of objectivity and infallibility. The use of these tools, however, introduces a new level of risk and complexity in policy (Osoba & Welser IV, 2017). Human-like semantic biases result from the application of standard ML to ordinary language — the same sort of language humans are exposed to every day (Caliskan et al., 2017).

ML should not be viewed as a lifeless technology by public enterprises, but rather as a living organism. The success of these systems in the real world cannot be predicted by laboratory testing alone. The use of these platforms and how users will respond to decisions made by executives should be thoroughly evaluated. Public enterprises should subject their new ML-based applications to randomized controlled trials to ensure their safety and fairness prior to rollout. They need to analyse products' decisions in the actual situations, where there are various types of users, to see whether the quality of decisions differs across them. They should consider testing them in limited situations to get a better idea of their accuracy and behaviour when various factors are at play — for instance, when users do not have equal expertise, the data from sources varies, or the environment changes. Failures in real-world settings signal the need to improve or retire algorithms.

As ML-based products and services and the environments they operate in evolve, public enterprises may find that their technologies do not perform as initially intended. It is therefore important that they set up ways to check that these technologies behave within appropriate limits. There is a significant communication gap between various branches of research on public enterprises. For example, in case of consumer choice, insights from market science are yet to be fully integrated in the management of public enterprises. It is important to note that effective resource allocation through ML requires high-quality data, well-designed algorithms, and responsible governance to ensure fairness and transparency. Additionally, the adoption of ML should be aligned with the goals and priorities of the public sector organisation to maximize its impact on resource optimization. While ML offers numerous benefits in the public sector, it also raises important ethical and privacy concerns, such as data security, bias in algorithms, and transparency. Striking a balance between reaping the advantages of ML and addressing these challenges is crucial for responsible and effective adoption in the public sector. Moreover, collaboration between the public enterprises, the academia and the civil society are essential to harness the full potential of ML in the public interest.

Acknowledgements

The author is thankful to two anonymous referees for helpful suggestions and Q. de Silva and K. Robertson for research assistance.

References

Babic, B., Cohen, I. G., Evgeniou, T., & Gerke, S. (2021). When machine learning goes off the rails a guide to managing the risks. *Harvard Business Review, 2021*(January-February), 1-10.

Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, *356*(6334), 183-186.

Carbonneau, R., Laframboise K., & Vahidov, R. (2008). Application of Machine Learning techniques for supply chain demand forecasting. *European Journal of Operational Research, 184*(3), 1140–1154.

De Lucia, C., Pazienza, P., & Bartlett, M. (2020). Does good ESG lead to better financial performances by firms? Machine learning and logistic regression models of public enterprises in Europe. *Sustainability*, *12*(13), 5317.

Denett, D. C. (2008). Can Machines think? In Levitin, D. J. (Ed.), *Foundations of Cognitive Psychology: Core Readings*. (pp. 35-54). Pearson.

Harnad, S. (2008). The Annotation Game: On Turing (1950) on Computing, Machinery, and Intelligence. In Epstein, R., G. Roberts and G. Beber, (Eds.), *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer. Evolving Consciousness*, (pp. 23-66). Springer.

Langley, P. (2011). The changing science of machine learning. Machine Learning, 82(3): 275-279.

Mays, W. (1952). Can Machines Think? *Philosophy, 27*, 148-162.

Mitchell, T. M. (1997). *Machine Learning*, McGraw Hill. Hoffman, L. (2011). A Lifelong Learner. *Communications of the ACM*, *54*(6): 127-128.

Mori J., Kajikawa, Y., Kashima, H., & Sakata, I. (2012). Machine learning approach for finding business partners and building reciprocal relationships. *Expert Systems with Applications, 39*(12), 10402–10407.

Osoba, O. A., & Welser IV, W. (2017). *An intelligence in our image: The risks of bias and errors in artificial intelligence*. Rand Corporation.

Samuel, A. (1959). Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development*, *3*(3), 210–229.

Turing, A. M. (1950). Computing Machinery and Intelligence. *Mind*, *LIX*(236): 433-460.