Humanitarian Machine Learning

STEPHEN J. BALDWIN

BONNER LEADER PROGRAM

COLLEGE of CHARLESTON
BONNER LEADER PROGRAM

Introduction

Machine learning is a subfield of computer science that gives computers the ability to learn without being explicitly programmed. Instead, computers are fed data and they learn by algorithmically identifying patterns in that data. Given the availability of data routinely collected by humanitarian organizations, *Can humanitarian organizations improve their impact by employing machine learning techniques?* The answer to such a question is vastly complex, touching on fields of computer science, mathematics, ethics, sociology, and psychology. This research attempts to scratch the surface by exploring the theoretical limitations and implications of humanitarian use of machine learning.

Current Landscape

Machine learning has countless commercial applications in modern technology and is employed frequently in healthcare, finance, social tech, and marketing. Many people experience a benefit from machine learning without even realizing it. For instance, credit card companies often utilize machine learning methodologies in order to detect fraudulent transactions. It is estimated that humans produce 2.5 quintillion bytes of data per day and 90% of the data in the world today has been created in the last two years.[1] It is not feasible for people to consume such quantities of data without the help of a computer. Therefore, it is likely that data science fields like machine learning will continue to grow. Machine learning itself can be applied virtually anywhere data is collected with the assumption that there are patterns to be found in that data.

Humanitarian efforts have not been immune to this development. GiveDirectly, a nonprofit that gives money to some of the poorest households in rural Kenya and Uganda via mobile phone transfers, attempted to identify impoverished villages with satellite imagery in 2013. Their metric for poverty was the proportion of thatched roofs to metal roofs, and although their algorithm was not accurate enough to be used, later similar efforts have been more promising.[2] In August 2016, researchers in Stanford's Sustainability and Artificial Intelligence Lab found that they could predict poverty with 75% accuracy using nighttime lights in satellite images as a rough proxy for economic wealth.[3]

Disaster relief has also benefitted from machine learning. Artificial Intelligence for Digital Response (AIDR) is an algorithm that was developed at the Qatar Computing Research Institute. This algorithm consumes microblog data and helps to identify and classify crisis locations in real time so that responders can react appropriately.

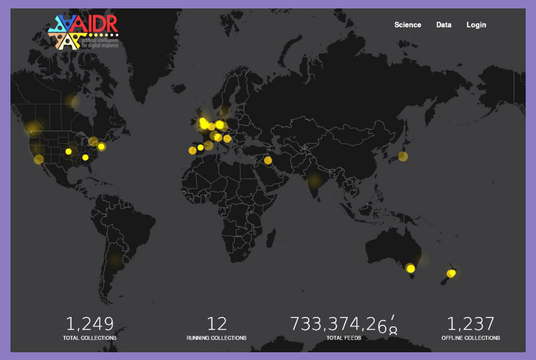

Figure 1. AIDR web application displaying current crisis areas

Perhaps no other field has adopted machine learning quite like healthcare. Machine learning has been used for medical diagnosis, bioinformatics, and DNA sequencing to name a few. In some cases, machine learning has performed better than traditional medical practices. For instance, researchers have been able to achieve a positive predictive value of 81% for autism on infants 6-12 months old using machine learning, whereas previous methods could only achieve such accuracy at 24 months. They did this by using machine learning to identify abnormal brain development patterns.[4]

Machine Learning

Machine learning is the study of algorithms that learn from, and make predictions on, data. During a training phase, such algorithms are given training data to generalize from. Training data can usually only represent a small fraction of the variability of all possible inputs, therefore generalization is key to machine learning. Applications where training data is provided alongside targets or solutions are known as supervised learning problems. These problems often involve the classification of data into a finite number of discrete categories. Where the output is one or more continuous variables, this is considered to be a regression problem. A researcher could use supervised learning to classify handwritten digits into discrete categories. This would be done by first training the algorithm with known images so that it could respond reasonably to unknown images.

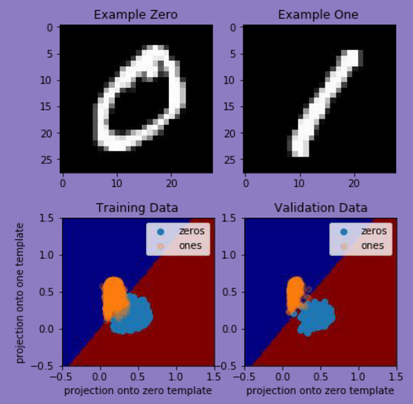

Figure 2. Supervised learning classifying handwritten 1s and 0s.

Other machine learning problems employ an unsupervised approach. They are trained by finding patterns in data with no defined target values. Such problems find solutions by clustering data rather than finding discrete categories for each datum. Finally, reinforcement learning is used to find suitable actions to take in a given situation. Such algorithms find the optimal action to take by calculating the rewards and costs for each potential action.

Case Study

The case study for this research analyzes the implications of using machine learning for loan approvals. Although this is not directly a humanitarian concern, it is a real application of machine learning with the potential to have deep social impacts. Traditional methods of risk assessment for loans utilize FICO scores which tend to hurt younger borrowers and foreign borrowers by targeting them with higher-interest loans.[5] Machine learning could conceivably overstep these biases and therefore make loans more accessible to more people with fewer defaults.

However, machine learning is merely an extension of human culture and can therefore show some of the same biases as humans if trained with biased data. Further, machine learning can create its own bias by finding correlations in data where there should be none. An example found in this research was an algorithm that linked people who pay inside rather than at the pump with being smokers and smokers were more highly correlated with a lack of creditworthiness.[5] It seems that, although machine learning has the potential to reduce traditional bias on loan risk assessment, its use should be knowledgeably pursued with proper accountability.

Conclusions

Because of the growing ubiquity of machine learning in modern day society, humanitarian efforts are not likely to go untouched. Some organizations have embraced this change. Machine learning, when used properly, has been highly effective in many domains. However, algorithms should be employed with care, especially where people are involved. It has been shown that machine learning can reflect some of the same biases people have when trained with biased data. Although beyond the scope of this research, some of this bias could be reduced using principles of proper data acquisition. Further, machine learning is only used for analysis and prediction. For this science to have a positive impact, it must be paired with passionate conscientious citizens that are willing to turn the insights of machine learning into action.

Bibliography

1. "Bringing big data to the enterprise," What is big data?, 17-Mar-2017. [Online]. Available: https://www-01.ibm.com/software/data/bigdata/what-is-big-data.html. [Accessed: 10-Apr-2017].
2. "Using Satellite Imagery to Find Villages In Need," DataKind. [Online]. Available: http://www.datakind.org/projects/using-the-simple-to-be-radical. [Accessed: 10-Apr-2017].
3. J. Neal, M. Burke, M. Xie *et al.*, "Combining Satellite Imagery and Machine Learning to Predict Poverty," in *Science,* vol. 353 AAAS, Gale Cengage, 2016.
4. M. Imran, P. Meier, C. Castillo *et al.*, "Enabling Digital Health by Automatic Classification of Short Messages," *ACM International Conference Proceeding Series*, 2016.
5. H.C. Hazlett, H. Gu, B.C. Munsell *et al*., "Early brain development in infants at high risk for autism spectrum disorder," in *Nature* no. 7641, 2017.
6. C. Lane, "Will Using Artificial Intelligence To Make Loans Trade One Kind Of Bias For Another?," NPR, 31-Mar-2017. [Online]. Available: http://www.npr.org/sections/alltechconsidered/2017/03/31/521946210/will-using-artificial-intelligence-to-make-loans-trade-one-kind-of-bias-for-anot. [Accessed: 10-Apr-2017].

CONTACT

Stephen J. Baldwin
sjbaldwin94@gmail.com