

This is a repository copy of *Rise of the machines? Machine learning approaches and mental health; opportunities and challenges*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/130913/>

Version: Accepted Version

Article:

Tiffin, Paul Alexander orcid.org/0000-0003-1770-5034 and Paton, Lewis William orcid.org/0000-0002-3328-5634 (2018) Rise of the machines? Machine learning approaches and mental health; opportunities and challenges. *British Journal of Psychiatry*. pp. 509-510.

<https://doi.org/10.1192/bjp.2018.105>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Rise of the machines? Machine learning approaches and mental health; opportunities and challenges

Paul A. Tiffin and Lewis W. Paton

Abstract

Machine learning methods are being increasingly applied to physical healthcare. In this article we describe some of the potential benefits, challenges and limitations of this approach in a mental health context. We provide a number of examples where machine learning could add value beyond conventional statistical modelling.

Editorial

Stephen Hawking is quoted as stating that *'the development of full artificial intelligence could spell the end of the human race'* and Elon Musk that *'artificial intelligence is a fundamental risk for human civilisation'*. However, if such risks can be averted it is clear that there is a huge potential for artificial intelligence to enhance the lives of humanity. Machine learning is the basis of artificial intelligence and there has been a growing interest in the application of the approach to address diagnostic and prediction issues in healthcare. Machine learning occurs when a system is able to learn from new information in order to complete a specified task. Learning in this sense can be 'supervised' or 'unsupervised'. In supervised learning the machine is shown a series of examples so it learns how to link predictors ('features') to a target outcome, e.g. converting handwritten characters to postcodes. In unsupervised tasks the system learns how to cluster or categorise variables in the absence of training data.

In contrast to conventional statistical modelling, which usually facilitates explanation, machine learning tends to be focused solely on prediction. Thus, in medicine, the technique has been used primarily to generate automated approaches to both diagnosis and treatment. There have been a number of well-publicised 'proof of concepts' in physical healthcare, such as those provided by DeepMind (1), formerly part of Google. DeepMind have been successful in generating machine learning algorithms that can predict impending mortality in cardiac patients and renal failure.

Unlike physical health, machine learning approaches to mental health have mainly been applied in research contexts rather than being implemented in real-world settings. For example, machine learning algorithms have been evaluated in relation to the identification of patients affected by Obsessive Compulsive Disorder via functional brain images (2), the detection of markers of depression using Instagram photos (3) and, combined with automated speech analysis, the prediction of psychosis onset in clinically at-risk individuals (4). Artificial intelligence is also being used to power 'Tess', a mental health support messaging service (5).

Differing machine learning approaches vary in relation to their mathematical underpinnings. Nevertheless they often share the same relative strengths and limitations of conventional statistical methods. However, machine learning is firmly focused on optimising performance on a prediction task. As such, learning algorithms will pursue this task with all the relentlessness of a T-101 Terminator pursuing Sarah Connor through a Los Angeles police station. Unlike traditional approaches to statistics, where linear (or log linear) relationships between predictors variables and an outcome are often assumed, machine learning takes a more flexible approach. As such linear or curvilinear relationships between the input variables (features) can be postulated, sometimes in numerous permutations, in order to maximise the predictive performance of an algorithm. Moreover, 'ensemble' techniques use many (sometimes thousands) of different models simultaneously in order to link features and outcomes most effectively under differing conditions. Thus machine learning is able to harness the 'brute force' of modern computing power to bear on prediction tasks.

Inevitably, these approaches bring with them the risk of 'overfitting' - that is, deriving a well-fitting model to the training data that has little or no predictive validity in a separate dataset. This is generally a consequence of the model being fitted to noise rather than the signal of interest. Machine learning experts have thus spent considerable time and effort in developing measures that counter this risk.

Machine learning models are sometimes described as 'black boxes' in that they are usually not interpretable. This means little is learned about the relationship between the predictors and the outcome of interest. However, this also means that models cannot be easily accommodated to changing trends. For example, models based on 'web scraping' approaches that harvest images or text from the internet in order to identify individuals who may be at risk of depression or suicide may be invalidated by changing trends in the use of language, or even camera settings. This may not be immediately apparent to end-users of the algorithms.

It is clear that machine learning has the potential to address important diagnostic and clinical decision-making issues in mental health. However, in order to realise this routinely available data of acceptable quality must exist- machine learning is no 'magic bullet' for 'dirty data'. Three potential sources of such data are worth mentioning. Firstly the implementation of electronic records in health services provides both coded and free-text that can be exploited by machine learning. An example of this has been the CRIS and D-CRIS systems (6), which have enabled a variety of research projects, including the use of natural language processing in order to describe symptom profiles in severe mental illness (7). The possibility of obtaining automated diagnoses that are at least as accurate as those given by experienced clinicians could offer cost-savings as resources become ever more straitened. Nevertheless, there will inevitably be human factors and public and professional perceptions of technology that may affect its implementation. Certainly, the experience to date in the use of mathematically driven clinician decision support tools is that they are frequently overridden by practitioners, who may feel that they have more to lose than the machine if the diagnosis is wrong, even by some very small chance (8). However, if automated systems effectively communicate their own degree of uncertainty in relation to a prediction then practitioners can take this into account when making their final clinical decision. Thus in practice, as with the use of autopilots in aeroplanes, machine learning is likely to complement rather than replace human clinicians in the near future.

The roll-out of improving access to psychotherapies (IAPT) services in England has also led to large, routinely arising datasets which include repeated brief symptom measures. It is possible that, combined with other social demographic background data as well as symptom profile data, patients who are unlikely to respond to brief interventions could be identified early on and more intense therapies offered. This is in keeping with the 'stratified /personalised medicine movement' where management is tailored to individuals or subgroups of patients. Moreover, such psychosocial data could be combined with biological markers (including genetic or immunological profile) or neuroimaging to improve the predictive accuracy of models. Furthermore, there is the possibility that previous existing data from large-scale trials could be reanalysed in order to understand which patients are most likely to benefit from certain interventions, as well as answering questions related to prognosis. Ethically, it would be important to provide alternative treatments for those identified as being highly likely to benefit from a particular therapeutic approach. Moreover, in practice, the 'negative predictive value' of algorithms will not be 100%. Consequently, it may be that such automated predictions would be best employed to guide early cessation of a treatment approach that was not appearing to be effective, rather than prevent access to such a therapy altogether.

A second important source of routinely arising data is from the 'digital footprint' left by internet-based activity. Machine learning can detect when an individual's online footprint, generated from apps, social media and other personal monitoring devices may indicate an impending risk of illness or self-harm or even violence (9). The advent of affordable wearable technology, combined with machine learning techniques, could lead to an era of automated risk monitoring, although ethical considerations must be taken into account.

In theory, there could come a time when machine learning approaches can deliver aspects of mental health care remotely and independently of clinical input. Such an achievement has the potential to increase the reach of mental health care, particularly to those individuals reticent to access mental health services currently. There have been previous attempts at 'remote care', such as computerised cognitive behavioural therapy and telemedicine. Such approaches have had some success in delivering care remotely, hinting at the potential for machine learning approaches. Nevertheless, machine learning approaches must learn the lessons of other remote care attempts, such as higher drop-out rates compared to treatment as usual (10).

To conclude, machine learning should not be considered a magic solution to predictive problems in mental health care, but as one tool in the quantitative box of approaches that may be useful in both research and, eventually, practice. Like a novel class of psychotropic drug, machine learning will undoubtedly eventually find its correct place within mental health services.

Author details

Paul A Tiffin, MD

Department of Health Sciences, Seebohm Rowntree Building, University of York, Heslington, YO10 5DD

paul.tiffin@york.ac.uk

Lewis W Paton, PhD

Department of Health Sciences, Seebohm Rowntree Building, University of York, Heslington, YO10 5DD

lewis.paton@york.ac.uk

Biographical details

PAT is a Reader (Associate Professor) in Psychometric Epidemiology at the University of York, and an NIHR Career Development Fellow. His main research interest is the development and application of statistical and psychometric modelling approaches.

LWP is a Research Fellow (Assistant Professor) at the University of York. He is a statistician interested in developing novel methodological approaches.

Declarations

Both authors have nothing to declare.

This paper presents independent research part-funded by the National Institute for Health Research (NIHR). The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health.

1. DeepMind. [cited 2017 20th November]. Available from: <https://deepmind.com/>.
2. Takagi Y, Sakai Y, Lisi G, Yahata N, Abe Y, Nishida S, Nakamae T, Morimoto J, Kawato M, Narumoto J, Tanaka SC. A neural marker of obsessive-compulsive disorder from whole-brain functional connectivity. *Scientific Reports*. 2017;7(1):7538.
3. Reece AG, Danforth CM. Instagram photos reveal predictive markers of depression. *EPJ Data Science*. 2017;6(1):15.
4. Bedi G, Carrillo F, Cecchi GA, Slezak DF, Sigman M, Mota NB, Ribeiro S, Javitt DC, Copelli M, Corcoran CM. Automated analysis of free speech predicts psychosis onset in high-risk youths. *npj Schizophrenia*. 2015;1:15030.
5. X2AI. Affordable, on-demand, and quality mental healthcare for everyone using psychological artificial intelligence [cited 2017 20th November]. Available from: <https://x2.ai/>.
6. South London and Maudsley NHS Foundation Trust. D-CRIS [cited 2018 26th April]. Available from: <https://www.maudsleybrc.nihr.ac.uk/facilities/clinical-record-interactive-search-cris/>
7. Jackson RG, Patel R, Jayatilleke N, Kolliakou A, Ball M, Gorrell G, et al. Natural language processing to extract symptoms of severe mental illness from clinical text: the Clinical Record Interactive Search Comprehensive Data Extraction (CRIS-CODE) project. *BMJ Open*. 2017;7(1):e012012.
8. Roshanov PS, Fernandes N, Wilczynski JM, Hemens BJ, You JJ, Handler SM, et al. Features of effective computerised clinical decision support systems: meta-regression of 162 randomised trials. *BMJ*. 2013;346:f657.
9. Walsh CG, Ribeiro JD, Franklin JC. Predicting risk of suicide attempts over time through machine learning. *Clinical Psychological Science*. 2017;5(3):457-69.
10. Waller R, Gilbody S. Barriers to the uptake of computerized cognitive behavioural therapy: a systematic review of the quantitative and qualitative evidence. *Psychological Medicine*. 2009;39(5):705-12.