Using Machine Learning to triage patients

By Russell Holmes, BSG Principal Consultant

Artificial intelligence (AI) and machine learning (ML) are evolving at a rapid rate, and as they become more sophisticated, the breadth and depth of their application potential expands exponentially. In the healthcare sector, that potential could result, not only, in cost-savings for patients and medical aid providers, but more importantly in fewer lives lost by saving doctors precious moments in the diagnosis and treatment of illness.

Typically, when illness strikes, a GP either makes a diagnosis based on their experience, or refers the case to a specialist. Specialist consults generally involve a series of physical and pathological tests to determine the cause and severity, and then a treatment plan can be defined.

These diagnostic tests are costly, time-consuming and comprise of hundreds of data points. Interpreting and understanding these results requires extensive training and experience. In many cases, the true cause of the illness is hidden among comorbidities or external factors, and navigating the barrage of red-herrings requires years of experience and extensive insight. Once a diagnosis and cause are clear, determining the severity, likelihood of recovery and treatment options all require further analysis.

> ***Interpretation of diagnositic testing requires a level of context and experience AI and ML cannot match – yet***

Much of this expertise is down to insight gained through years of specialised experience and an in-depth understanding of the field. For this reason, AI and ML cannot fully replace doctors – not yet anyway – but they can be used to triage patients into high-level risk categories, based on treatment of previous cases, saving doctors time and effort.

ML assessments are based purely on the maths of previous cases and there's not much room for conjecture and opinion. This might result in 'exceptions to the rules' being incorrectly categorised, but by and large, the stats will bear out. This approach could be used to support specialists in recommending pathological testing prior to a full consult, or enable medical aid schemes to verify the need for certain diagnostic testing beforehand.

The benefit of categorising patients in such a way is twofold:

1. **Patients can be prioritised according to risk**, resulting in less time spent identifying primary causes, and more time planning specialised treatment programmes
2. **Long-term financial implications can be predicted** (to a certain degree of accuracy), helping medical aid schemes and patients alike when planning and forecasting the associated financial burden

Can we trust the model?

It's human nature to mistrust machinery, but throughout history machines have consistently saved us time and energy, and this is no different. Each AI and ML model is supported by an accuracy percentage, which is continually measured against a minimum threshold. If the threshold is not upheld, the model will shut down. There is of course a margin for error, and specialists should in no way defer to the model's diagnosis without validating it. But, doctors can use the model to support or validate their assumptions, or to enable early diagnostic testing according to case presentation and symptoms.

How does it work?

Medical practitioners, hospital groups and medical aid schemes generally have access to a wealth of information about their patients and members. By leveraging this information, a predictive diagnostic model can be built.

1. **Filtering by diagnosis:** By first filtering patients and cases by diagnosis, using the ICD (International Classification of Diseases) codes, the size of the data source is reduced, providing focus to your ML model (i.e. what diagnosis are you solving for?)
2. **Combine demographics and pathology**: Patient demographics can have an influence on the model's function – some illnesses are more common in men than women, or across racial or geographic lines. By using demographic information, linked to pathology, you are able to further enrich the model
3. **With history:** This brings in other factors – such as financial records, medical aid claims and medical history.
4. **Batches similar patients together:** Based on the above filters and data combinations, the model batches similar patients together using an unsupervised machine learning algorithm, effectively creating patient personas for the diagnosis you are solving for.

It's important, once you have a model, to then spend time with the specialist, evaluating the high-level persona(s) based on their specialised experience and understanding of the diagnosis. This will allow you to incorporate things like risk rate, mortality and survival rates, comorbidity correlations, common CPT (current procedural terminology) codes, etc.

Up to now, most of what you'll have done will be based on previous patients and cases – using what you know of them from their records and from their doctors' experience of their cases to tag and label the data. This labelled data can then be used to train a supervised machine learning model to recognise and categorise new patients. Each new patient analysed by the model is assigned a probability of accuracy score, based on how closely they match the patient persona. This probability score can be used to predict various factors, such as the accuracy of the risk, mortality likelihood, and comorbidities.

Has it been tested?

A recent project, completed at a renal care facility, created a model that triaged patients with 92% accuracy. The model was built on patient demographics, pathology, and physical and mental assessments. Despite operating with 'unclean' physical and mental health data, the model was still able to achieve high accuracy levels. The project, which was completed in a month and run by four consultants from BSG's Data and Analytics capability, quickly proved its validity as a diagnostic tool. Although not yet fully integrated into daily decision-making, the uses of the model – and others like it – are clear, and the fear that work like this is slow and costly has been firmly put to rest.

Our frontline healthcare workers have been stretched to their limits through this pandemic, and have dealt with a massive volume of cases (both COVID and non-COVID-related). If AI and ML can assist with even the most basic risk assessment, it would go a long way in helping healthcare workers prioritise their patients based on genuine need and risk, reducing the financial burden of unnecessary diagnostic testing and ultimately saving lives.

If you'd like to find out more about the work BSG has done in this area, read the project case study or get in touch.

About BSG

As a homegrown South African Consulting and Technology company, BSG is uniquely positioned to deliver solutions tailored to the South African context.

We have more than 20 years' experience across the banking, specialised financial services, insurance, healthcare, telecommunications, and oil and gas sectors. By employing a multi-skilled approach, BSG effectively leverages our clients' data to create solutions that improve the experiences of their customers and solve enterprise-scale challenges.

We understand the dynamics of Business and Technology, which allows us to create flow between supply and demand, bridging the gap between business and IT. We work with our clients to drive out success, transforming their operational platforms and creating the customer experiences they need.

Visit us at **www.bsg.co.za**

Get in touch

Jurie Schoeman
BSG Chief Executive Officer
jurie.schoeman@bsg.co.za
+27 (0)83 30 27169
+27 (0)11 215 6666