Radiology reporting: can AI really be the panacea?

RAD Magazine, 46, 538, 21-22

Dr Rizwan Malik
Consultant radiologist, Royal Bolton Hospital

There is a worldwide shortage of radiologists, but a year-on-year increase in demand for imaging. According to The Royal College of Radiologists *Clinical Radiology UK Workforce Consensus 2018 Report*,[1] in the UK a tenth of radiologist jobs were unfilled in 2018 and scan outsourcing spend was nearly £140 million.

In this article we explore how the application of AI processes might help bridge the gap between staffing and demand while maintaining, and potentially even enhancing, the quality of the reports aiding radiologists, clinicians and the wider clinical pathways for the benefit of patients.

The aim for any AI application within radiology should be to assist in arriving at an outcome or diagnosis from the imaging more quickly than possible by traditional reporting methods, creating more workflow capacity while maintaining the same quality of diagnosis and service. It is this arrival at a clinically actionable outcome that we will define here as a 'report' rather than the more prosaic notion of words generated by a radiologist.

To do this effectively AI needs to encompass the four stages of the patient pathway: pre-imaging; image acquisition; image interpretation and reporting; and post-reporting. Strategies to improve reporting aren't just achieved by working on image interpretation and reporting – improvements in image acquisition and post-reporting can also be made for a better overall workflow.

The five main ways AI can assist in this pathway include:

1, expedite pertinent information to referring clinician;
2, appropriate case prioritisation for reporting;
3, reduce inefficient aspects of reporting process;
4, improve/maintain reporting standards;
5, reduce the administration burden related to reporting.

Reporting workflow needs to be made more efficient by removing or reducing repetitive tasks and a good AI solution will enable this. Appropriate case prioritisation to target the limited resource (in this case the reporting radiologist) more effectively include either highlighting cases from referring clinicians that need urgent reporting, or, in the context of screening programmes where there may be a huge number of normal scans, to tease out those cases that may be more likely to demonstrate an abnormality. Contemporaneous peer and discrepancy review can help with improving the quality of reporting, and while structured reporting templates help maintain data consistency between radiologists and departments, the current methods are somewhat cumbersome and thus finding more efficient means to impart information concisely but while preserving ease of data mining would be useful. Post-reporting administration overheads can be reduced by targeting the processes for adding cases to MDT meeting schedules as well as forwarding reports for onward management as appropriate.

A more streamlined reporting workflow will also benefit the referring clinician by providing a shortened waiting time for the imaging findings and reduction of their administrative burden by lessening the need to chase reports from the radiology department and send for onward management.

Patients will benefit by getting their diagnosis, treatment and eventual discharge quicker, and a shorter waiting period for results lessens the anxiety associated with a prolonged diagnostic process.

How can AI make improvements at the image acquisition stage?

At the image acquisition stage there are a number of ways that AI can expedite the management of patients, leading to reduced time to diagnosis and earlier treatment.

In the ER setting, critical findings for urgent care can be flagged, such as brain bleeds, strokes or pulmonary embolism from CT, highlighting patients for thrombolysis, and from x-ray diagnosis of pneumothorax.

Even when the findings are not so critical, there are a number of circumstances where earlier quantification of the reason for the imaging can greatly expedite patient care, for example, in the x-ray detection of lines and drains placement in an ICU setting: by obviating the need for human review of the x-rays to ensure a line is correctly sited (which might sometimes take many hours) we reduce the needless delay in initiating the appropriate treatment. Similar, albeit less pressing, examples include identifying fractures or flagging pathology on chest x-rays such as consolidation and effusion.

AI solutions have been deployed with demonstrable benefits in the context of tuberculosis screening. Late diagnosis hampers public health outcomes since TB is so contagious and delay increases the risk of spread. Another big challenge is patient tracking and treatment adherence, especially as TB programmes are often conducted in slums and other low income geographies; if the diagnosis arrives after weeks (when the x-ray vans return to cities and over-burdened radiologists eventually get to review the images) if there is no trackable address it can be hard to track down the patient.

AI products are also being targeted at the management of national screening programmes in two distinct ways:

1, For breast screening to use AI to act as the second reader, thus allowing radiologists to read more scans by reducing their need to perform double reporting duties.
2, In the developing lung cancer screening programmes using AI so that 'normal' scans can be quickly filtered out, flagging any potential malignancies, which in turn can help develop more effective one-stop-shop management.

Lesion change analysis can expedite review and reporting of lesions that have increased in size, while lesions with no change or shrinkage can be posted as less important.

This expedition of management means earlier reassurance for patients. Patients with 'normal' findings could be discharged quicker by a radiographer, so no more waiting around for radiologists or A&E clinicians to review normal imaging findings before discharge. Screening programme management will reduce the reporting workload by filtering appropriate cases with abnormal results for review.

Improvements to the post-processing stage

Improvement of worklist management and examination routing means the right case is automatically sent to the right radiologist to report, ideally matching the consultant's specialisms. KPI's can be automatically reviewed and assessed. Urgent findings are flagged and prioritised whereas 'normal' findings can be redirected to a different management pathway.

Improvements to diagnosis and reporting

Once the reporting stage is reached, smart hanging protocols should arrange the images for optimal viewing, including any relevant priors – both of the same and different modalities (such as abnormal chest x-rays that prompted chest CTs), and by automatically coregistering images either by anatomical start point or to match lesions on the current and previous imaging (again both with the same type of imaging such as current and prior CT as well as, for example, a lesion on the MRI scan that has been arranged to better quantify the same lesion first identified on the CT scan).

Improvements at this stage can also target non-pixel data by aggregating and providing better access to any clinical details relevant to the request or associated reports – for example for other types of investigations a radiologist might not otherwise see on the radiology systems such as reports or even the video files of a colonoscopy examination for a patient referred for a colon CT scan.

There are studies which highlight that humans can be poor at differentiating between false positives and true positives. Therefore by instead of being 'always on' and requiring reporter validation of lesions it has identified, the AI can be used as an 'on demand' analysis aid to provide a second opinion on areas of uncertainty where necessary; this will reduce the risk of false positives.

There is significant scope for AI segmentation tools in their current guise to be improved. Consider that at the moment when reviewing a full body CT, the reporter may need to load the 'chest AI' tool that will either automatically identify a lesion or more likely need some level of user interaction to measure the lesion(s), then scroll to the abdomen and unload the chest AI tool and now load the liver or colon tool and again intervene manually to measure the lesion. After this each lesion and perhaps its change from the previous scan must be dictated into the report. This involves multiple steps of needless mouse clicks, and a waste of vast computation power for this amount of manual measurement. Instead, as a standard part of the reporting workflow one would wish the solution to segment lesions across the entire scan, rather than just a specific area, and automatically measure them, with identification of any changes. These findings should then be embedded into the final report with (as appropriate) reference to RECIST criteria and perhaps with hyperlinks to the location of the lesion on the CT. This would reduce a vast amount of the repetitive, inefficient part of reporting, freeing up time and freeing up the radiologist to use their attention for more appropriate fuzzy logic engagement with the imaging.

AI should also provide image analysis support. This can range from the rather prosaic straightforward aids, such as identification of fractures, consolidation and other such pathology on chest x-rays to improve diagnostic quality and confidence, to image disease pattern analysis for example in the classification of aetiology of interstitial lung disease on high resolution chest CTs. Taking this further, we might consider the analysis of attenuation or enhancement characteristics or MR signal profiles to identify likely aetiology (such as in adrenal lesions). More advanced still would be the use of radionomics – extracting a large number of quantitative imaging features, potentially including aspects in the imaging not visible to the human eye to help classify tumours and thus either reduce the need for additional imaging, or else obviate the need for invasive biopsies.

Software should also facilitate clinical decision support by providing access to the latest pathways or protocols to ensure that any onward management conforms to the latest standards or guidelines, for example the correct imaging schedule for lung nodule follow-up in line with locally adopted guidelines.

The 2019 Digital Health Hype Cycle[2] unfortunately identified voice recognition solutions as still firmly rooted in the trough of disillusionment. However, when we consider the use of natural language processing, I do not mean in the context of the much-needed improvement in the accuracy of voice recognition software, or even to provide suggested 'next words or phrases' (something we have had access to for many years on the humble mobile phone with predictive text messaging). Instead, semantic reporting methods may help to bridge the gap and need between the "art" of reporting – individualised styles – with the desire for standardised (template-based) reporting for data consistency so desired by those wanting to mine this data. The result may be a superficial report layer, which is clinical-facing and that may be more focused (for example you might simply say "normal") whereas the deeper layers contain all the referenceable meta data (for example "normal abdomen" refers to normal liver, spleen, kidneys and so forth). Such methodology might allow data-minable structured reports to be created without the currently cumbersome methods of template-based reporting and without compromising a radiologist's individual reporting style.

We may consider the use of AI to automatically generate a summary of the salient points in the body of the report into a conclusion, or perhaps auto-insert the correct TNM staging or other relevant current guidelines from specified report descriptors. This would aid not only consistency but also quality and efficiency of the process.

If we are accepting that AI outputs can be used to guide management, then a logical progression would be to proceed to automatic report creation. This would also be of benefit in reducing the workload of secondary report readers, for example examinations forwarded for review from the national screening programmes. It could also be applied to fracture reporting, or any examinations that are read as normal to cut down on the amount of reports needing creation.

In-built peer review and QA software could compare the final report with the imaging itself and flag any discrepancies or omissions, particularly useful for radiology trainees not only to aid their training but only to expedite their engagement with service provision and thus increase the pool of reporting availability.

Improvements to the post-reporting workflow

Automation of many of the post-reporting procedures has the potential to significantly reduce the administrative burden of the radiology department. Possibilities include:

1, The referring clinician or department can be notified of critical findings or potential next steps (based on current guidelines).
2, Not only can onward examinations be suggested but they could be automatically scheduled at the correct intervals.
3, Patients with normal findings can be automatically added to the discharge pathway.

Significant sources of inefficiency and frustration currently are the MDT management processes, which are labour intensive. In many institutions, once it is deemed necessary by the radiologist to discuss a case at the MDT this must be notified by call or email to the MDT coordinator who compiles all such requests. It is not unusual for a finalised MDT list to be released to the MDT radiologist only a short time before the meeting for the cases to be prepped. This often results in cases added that have not yet been reported, or have been reported subsequent to being added and don't have cancer on them (or are otherwise not appropriate for that meeting), or might require additional investigations that have not yet been completed. It would be helpful if, instead, cases that need discussing in further detail can automatically be assigned to the next appropriate MDT meeting from the time of reporting; with clear visibility of report readiness or appropriateness this would free up radiologist time which in turn can be used for reporting.

Summary

Is AI the panacea to radiology reporting? No. However, as we have explored, there is plenty of scope for AI processes to help clinical and radiology pathways arrive at an outcome from imaging more quickly, as well as create more capacity of the same or even enhanced quality. The general principles through which such AI tools could achieve this include expediting pertinent information to the clinician, appropriate case selection/prioritisation for reporting, reducing some of the repetitive or inefficient aspects of the reporting process, improving/maintaining reporting standards and reducing the administration burden related to reporting.

Reference

1, The Royal College of Radiologists. Available at: www.rcr.ac.uk/system/files/publication/field_publication_files/clinical-radiology-uk-workforce-census-report-2018.pdf

2, The Digital Health Hype Cycle. www.healthcare.digital/single-post/2019/01/12/The-Digital-Health-Hype-Cycle-2019.