Received:
09 November 2018

Revised:
29 January 2019

Accepted:
01 February 2019

<https://doi.org/10.1259/bjr.20180958>

Cite this article as:

Thirkettle M, Thyoka M, Gopalan P, Fernandes N, Stafford T, Offiah AC. Internet-based measurement of visual assessment skill of trainee radiologists: developing a sensitive tool. *Br J Radiol* 2019; **92**: 20180958.

FULL PAPER

Internet-based measurement of visual assessment skill of trainee radiologists: developing a sensitive tool

¹MARTIN THIRKETTLE, Ba, PhD, ²MANDELA THYOKA, ²PADMINI GOPALAN, ³NADIAH FERNANDES, MBBS, BSc, ⁴TOM STAFFORD, BSc, PhD and ⁵AMAKA C OFFIAH, BSc, MBBS, MRCP, FRCR, PhD, FHEA

¹Department of Psychology, Sociology & Politics, Sheffield Hallam University, Sheffield, United Kingdom

²Department of Radiology, Sheffield Teaching Hospitals, Sheffield, United Kingdom

³Imperial College London School of Medicine, London, United Kingdom

⁴Department of Psychology, University of Sheffield, Sheffield, United Kingdom

⁵Department of Radiology, Sheffield Children's NHS Foundation Trust, Sheffield, United Kingdom

Address correspondence to: Dr Martin Thirkettle

E-mail: m.thirkettle@shu.ac.uk

Objective: Expert radiologists exhibit high levels of visual diagnostic accuracy from review of radiological images, doing so after accumulating years of training and experience. To train new radiologists, learning interventions must focus on the development of these skills. By developing a web-based measure of image assessment, a key part of visual diagnosis, we aimed to capture differences in the performance of expert, trainee and non-radiologists.

Methods: 12 consultant paediatric radiologists, 12 radiology registrars, and 39 medical students were recruited to the study. All participants completed a two-part, online task requiring them to visually assess 30 images (25 containing an abnormality) drawn from a library of 150 paediatric skeletal radiographs assessed prior to the study. Participants first identified whether an image contained an abnormality, and then clicked within the image to mark its location. Performance measures of

identification accuracy, localisation precision, and task time were collected.

Results: Despite the difficulties of web-based testing, large differences in performance, both in terms of the accuracy of abnormality identification and in the precision of abnormality localisation were found between groups, with consultant radiologists the most accurate both at identifying images containing abnormalities ($p < 0.001$) and at localising abnormalities on the images ($p < 0.001$).

Conclusions: Our data demonstrate that an online measurement of radiological skill is sufficiently sensitive to detect group level changes in performance consistent with the development of expertise.

Advances in knowledge: The developed tool will allow future studies assessing the impact of different training strategies on cognitive performance and diagnostic accuracy.

INTRODUCTION

The accurate interpretation of radiological images in order to reach a correct diagnosis is at the heart of the expertise of a radiologist.¹ Because radiographic images are part of the "first-line" of diagnosis for traumatic medical conditions, the identification and localisation of abnormalities is a highly valuable area of clinical expertise. Accordingly, understanding the development of this expertise has attracted significant interest not only from radiologists striving to improve performance in the field,²⁻⁵ but also from psychologists, for whom radiology acts as an excellent, real-world assay of visual cognition and expertise more generally.⁶⁻¹⁰

Previous research has divided the visual expertise of the radiologist into two constituent parts; visual search

expertise and cognitive or analytical skills.^{11,12} The first step involves perceptual interrogation of the medical image, noting any abnormalities. The second step analyses what has been noted within the clinical context of the patient's presentation. While there is mixed support for performance differences in the first of these two steps when radiologists and non-radiologists are compared in tasks using non-clinical images,^{13,14} trainee radiologists must acquire both abilities to develop expertise in diagnostic radiology.¹⁵

Previous studies in visual search expertise have demonstrated that expert radiologists spend less time scrutinising each image than novices or trainees^{12,14} and visually explore images in a different manner to trainees,^{7,14} suggesting that a more profound, strategic change rather than the simple accumulation of knowledge is the foundation of expert

radiological skill. Experts develop a robust memory structure that forms the basis of an extensive knowledge base and devise analytical strategies to assist them in correlating the clinical information and image data, allowing for superior information processing. Indeed, experts demonstrate recall superior to that of novices, especially when time is limited.¹⁶

The present study describes the development of a web-based behavioural measure of visual assessment using a library of paediatric skeletal radiographs. We aimed to assess the feasibility of the online and pragmatic interpretation of radiographs and to determine whether the collection of radiographs was of sufficient quality and sensitivity to allow future longitudinal visual tracking experiments. Our hypothesis, in line with previous literature,^{1,2,16,17} is that the consultant radiologists would perform more accurately across all measures, and do so while spending less time assessing a given image than the radiology registrars, who in turn would be better and faster than the medical students.

METHODS AND MATERIALS

Study design

This was a web-based study using a bank of paediatric radiographs predominantly of fractures, but also including normal variants and congenital abnormalities. A computer-based task was developed to quantify the ability of radiologists of varying experience. The dedicated library of 150 skeletal radiographs was selected from 3000 radiographs obtained from Sheffield Children's Hospital NHS Foundation Trust on children presenting to the Emergency Department over a 6-year period (2008–2013), following trauma. The images were assessed separately by a consultant paediatric radiologist and a radiology trainee, each of whom made their assessment with access to the radiology reports—only radiographs in which there was no discrepancy between initial reporter, the consultant paediatric radiologist and the trainee were included (where available, follow-up radiographs were also assessed to help with the decision-making). For each image, the veridical location of any abnormalities as documented on the picture archiving and communications system (PACS) was recorded for comparison against participant responses and (for the purposes of computing diagnostic accuracy), was taken as the gold-standard. The two assessors, through agreement, subjectively graded each image into one of three categories of difficulty from the perspective of the second-year radiology trainee: easy, intermediate and difficult. 16 normal radiographs were included. All identifying details were removed from the images. Such a large library of images was curated in order that in future the same participant could revisit the task multiple times and be faced with a different sample of pre-assessed images, avoiding the possibility in a longitudinal study that specific image familiarity could explain any improvement in performance.

To balance the desire to replicate clinical practice as far as possible within the task against the need to quantify performance as accurately as possible across the various cognitive demands of radiology, the task was split into two stages for each image—identification of the presence of an abnormality, and, if an abnormality was detected, its localisation.

Delivery of the survey, image display and collection of participants' demographic information were all managed by the on-line survey platform Qualtrics, augmented with custom Javascript for abnormality localisation measures.

Participant selection and recruitment

12 consultant paediatric radiologists (referred to from this point on as consultant radiologists or "CR"), all members of the European Society of Paediatric Radiology Child Abuse Task-force, responded to an open invitation to participate in the study. Consultants reported having between 6 and 31 years of radiology experience (mean 15, SD 6.7) and all were practicing within the UK or EU. Concurrently, 12 radiology registrars from across the 5 years of the South Yorkshire Radiology Training Scheme (referred to from this point on as trainee radiologists or "TR"), and 39 medical students from The University of Sheffield Medical School's 5-year degree, who had not received radiology training (referred to from this point on as "MS"), were recruited. Power analysis was conducted to ensure that sufficient participants had been recruited, for testing at an $\alpha = 0.05$ level of significance. All participants were recruited through email, ensuring that only those invited to participate could access the test web page. The reading environment, computer screen and time(s) of reading, were all left to the reader's discretion. All aspects of the study were conducted in accordance with the ethical standards as laid down in the 1964 Declaration of Helsinki. The study was approved by the University of Sheffield Psychology Department Research Ethics board. Participants were exposed to a full ethics and consent statement and provided explicit consent before completing the task. Participants were given the opportunity to withdraw from the study at any time by contacting the research team. The consultants were remunerated for their time, while the registrars and medical students were entered into separate draws to win a £50 book voucher.

Procedure

After consenting to the study, participants completed a short demographic section including years of experience. Via on-screen instruction, they were then briefed on the experimental task and encouraged to ensure they were not using a particularly small screen, before completing a practice set of image responses. The instructions were repeated at this point, and the participant then began the testing session proper.

All participants completed responses to the same 30 images, 10 each of those previously ranked as easy, intermediate and difficult. 5 of these 30 images have been assessed as not containing an abnormality, and these acted as "target absent" images in the test, included to detect "false positive" responses from the participants. This gave our test image set an abnormality prevalence of 83.3%, far greater than that found in a clinical setting, but not uncommon in psychological studies.¹⁰ All images were resized prior to data collection to maintain a constant image height (600 pixels) to suit web viewing. The order of image presentation was randomised across participants to avoid any order effects. Date, time, computer, screen resolution and other related information was recorded.

Participants were asked to identify both pathology and normal variants. First, participants were asked to rank how likely they thought there was an abnormality on the given radiograph on a 6-point Likert scale (definitely yes, probably yes, possibly yes, possibly no, probably no, definitely no). The time taken for this was labelled the “decision time”. If participants clicked any of the first three options indicating they thought an abnormality likely, then they were asked to indicate location by clicking the point(s) on the radiograph where they believed there was abnormality/abnormalities. Once satisfied, they clicked the submit button to move onto the next radiograph (the time taken to click the submit button was labelled the “localisation time”). Some images in the bank contained two projections of the same area (e.g. anteroposterior and lateral knee); in these cases, participants were instructed to identify and click to locate abnormalities visible on either or both projections.

Data were collected on the accuracy of the participant’s identification of an abnormality—as compared to the reference answer for each image—the participants’ decision time, localisation time and accuracy of the participants’ localisation of each abnormality. Localisation error was calculated as the distance in pixels of the participant’s click from the reference location of the abnormality. A mixed ANOVA was used to test participant performance. Post hoc testing was used to further interrogate significant differences—significance was defined as ($p < 0.05$). Statistical analysis was performed using SPSS v. 20.

RESULTS

Results of two observers were excluded for low task engagement (1 MS) and excessively long response times (1 TR). Therefore, results presented are for 12 CR, 11 TR and 38 MS.

Identification task accuracy and inaccuracies

Participants’ responses (Figure 1) were categorised as positive or negative with respect to an abnormality and the corresponding sensitivity (proportion of images containing abnormalities

Figure 1. (A) DP and lateral L wrist. (B) DP right hand. Example radiographs from the library as presented in the experimental task showing reference locations for abnormalities (red), marks placed by consultants (blue), trainees (yellow), and medical students (green). Figure 1A was graded as easy, and Figure 1b as intermediate.

Figure 2. Average d' for each participant group for abnormality identification. Error bars show bootstrapped 95% confidence intervals.

correctly identified) and specificity were calculated for each group of observers and for each image difficulty level.

D' (“ d prime”) for each group was then calculated to combine specificity and sensitivity,¹⁸ collapsing across image difficulty. Figure 2 shows the d' across the three groups. This combined measure was significantly affected by group $F(2,58) = 29.698$, $p < 0.001$, $\eta_p^2 = 0.506$ and post hoc testing showed that only the sensitivity results of CR and TR were not significantly different ($p = 0.27$), all other post hoc comparisons were significant at the $p = 0.001$ level.

An unpaired t -test showed that the sensitivity of MS was not significantly different from zero ($t(37) = 1.90$, $p = 0.065$, $M_{diff} = 0.16$, 95% CI = -0.01 – 0.33) suggesting a discriminative ability on the identification task which was indistinguishable from guessing.

There were significant group differences ($p < 0.001$) in true and false positive and negative rates as summarised in Table 1.

Localisation error

There was no significant effect of image difficulty on localisation error ($F(2,116) = 1.705$, $p = 0.153$, $\eta_p^2 = 0.056$). Therefore, localisation error performance was analysed after collapsing across difficulty levels. CR and TR were far more accurate than MS in

Table 1. Group mean (and standard deviations) of performance for responses to the abnormality identification task

	Medical students	Trainee radiologists	Consultant radiologists
True positive responses	14.24 (2.81)	17.00 (1.95)	19.83 (2.65)
True negative responses	2.42 (1.13)	3.45 (1.13)	3.41 (0.90)
False positive responses	2.55 (1.10)	1.55 (1.13)	1.58 (0.90)
False negative responses	10.713 (2.72)	7.91 (1.97)	5.16 (2.66)

Figure 3. Group differences in localisation error (A) and time spent per image (B). Medical students were far less precise in their localisation of abnormalities compared to the reference location for each image (A), while responding far quicker than the trainee or consultant radiologists (B). Error bars show bootstrapped 95% confidence intervals.

locating abnormalities, clicking on the image far closer to the reference location (mean location error CR = 46.26 (SD 18.8), TR = 43.27 (SD 22.99), MS = 97.98 (SD 35.38) pixels, $F(2,58) = 21.185$, $p < 0.001$, $\eta_p^2 = 0.422$). Post hoc testing showed that MS were significantly less accurate than CR ($p < 0.001$) and TR ($p < 0.001$), while CR and TR were not significantly different from each other ($p = 1.0$) – Figure 3A.

Task time

On average, MS spent 15.8 s (SD 8.5 s) completing the identification and localisation tasks for each image making them faster than both TR ($M = 20.7$ s, SD = 8.1 s) and CR ($M = 36.9$ s, SD = 18.8 s). There was a significant effect of group on total time spent per image ($F(2, 58) = 16.383$, $p < 0.001$, $\eta_p^2 = 0.361$), which post hoc testing confirmed was driven by the speed of the CR, who were significantly slower than both TR ($p = 0.003$) and MS ($p < 0.001$). TR and MS were not significantly faster or slower than each other ($p = 0.6$) – Figure 3B.

As Figure 3B shows, splitting the total time per image into the two component tasks—identification and localisation—there was a significant difference between groups in time taken on the first part of the task, deciding if an image contained an abnormality or not ($F(2,58) = 19.75$, $p < 0.001$, $\eta_p^2 = 0.405$), but no difference between groups on time taken to localise the abnormality on the image ($F(2,58) = 2.27$, $p = 0.122$, $\eta_p^2 = 0.073$).

DISCUSSION

We present data from consultant radiologists, trainee radiologists and medical students to demonstrate that online testing is sensitive enough to meaningfully capture observer differences in diagnostic accuracy from radiographs, despite the near total lack of control over the conditions within which the task was performed and the hardware used by each participant to access the study. One important aspect of the hardware used for the task is screen resolution, which varied significantly between participants, and it is reasonable to suggest this may have impacted performance.

However, at the group level, differences between the resolutions used do not predict task performance. On average, consultants used higher screen resolution than both the trainees and medical students, but there was no difference between trainee radiologists and medical students. This means that it is difficult to attribute differences in task performance entirely to screen resolution, either in terms of detection accuracy—where consultants did not outperform trainees despite the higher resolution of their screen, or localisation error—where trainees significantly outperformed the medical students on similar resolution hardware. Moreover, we are not able to comment whatsoever on the ambient lighting or broader testing environment each participant chose for their participation, which potentially may also have affected their performance to some degree. While caution should be taken, particularly in cross-sectional designs, when interpreting data from internet-based studies where so much is left uncontrolled, the results of this study show that overall accuracy followed the expected pattern based on the level of radiological expertise with the caveat that our task was, unusually, also completed more slowly by the most expert participants.

It has previously been shown that the diagnostic accuracy of relatively senior radiologists for the detection of the subtle fractures of child abuse was low and that there was no correlation between years of experience and expertise, in spite of many experts being experienced.¹⁹ A further study demonstrated that UK radiologists perceive they would benefit from improved training in this field.²⁰ However, to design and assess any such training, it is imperative to fully understand what makes one individual more “expert” than another and this study is the first step towards developing that understanding.

Recent research has focussed on classification of the errors made by paediatric radiologists in visual assessment of radiographs.²¹ Our approach is to develop an easily deliverable, sensitive, and repeatable measure of the skill being acquired, by collecting image ratings from a wide set of experts and trainees and

comparing behavioural markers of performance across levels of training and experience. This approach allows the development of rapid assessment and quantification of the underlying skill differences between the expert and trainees. In turn, results from such tasks can be used to drive improvements in training interventions. Such improvements to the design and delivery of teaching during the training of trainee radiologists would potentially allow the faster development of the skills involved in the visual assessment of radiographs, leading to better diagnosis of trauma and disease and allow expert levels of performance to be attained earlier in a trainee's career, with resulting improvements in clinical performance and patient care.

Consultant radiologists, experts in their field, outperformed their intermediary and novice counterparts. Medical students achieved the poorest accuracy score of the three groups. The d' measure is a common and reliable statistic within psychological studies of visual search tasks.^{18,22} A higher d' value suggests a higher sensitivity for the task, with both lower false positive and higher true positive responses required for higher d' values. A d' of zero reflects chance performance and would indicate the participant to be wholly insensitive to the task and just guessing. In this study, MS performed so poorly as to be indistinguishable from chance. The MS had not received radiology training, so we are not surprised by this result, which does however support the use of the tool for future studies. Further support for the tool is found in disaggregated analysis, not reported here, which showed that both trainee and consultant radiologists' abnormality identification performance reduced as image difficulty rating increased.

One unusual result was the finding that consultant radiologists performed this task significantly slower than novice participants. This contradicts previous research findings¹² and runs counter to perceived wisdom on expertise development; as skill develops, both speed and accuracy improve.^{1,7} However, rather than propose our results genuinely suggest a revision to this position, it seems more likely that the open nature of our task left the participants free to perform the task at different levels of meticulousness. The near chance level of accurate performance in the novice group supports a view of these participants clicking through the task without the level of diligence shown by the consultants, resulting in far faster task performance. Future studies will need to address task engagement to ensure that participants at all levels of expertise engage with the task sufficiently to provide a reliable measure of their ability. That said, the current results support the potential of web-based assessment protocols, particularly either within a structured training program, where tutors can use the measures described here to evaluate trainees' performance or as a self-assessment program.

The next step in this project is to use the library of validated radiographs in longitudinal studies of cohorts as they complete their training, and to add eye tracking experiments to examine changes in participants' search strategies with increasing experience. This has been done in other contexts²³ and we will adapt these published methods to our needs in paediatric radiology.

CONCLUSION

The present study demonstrates that consultant radiologists performed far better than trainee radiologists or medical students in correctly detecting the presence of an abnormality on paediatric musculoskeletal radiographs. We show that a web-based delivery of the experimental task is sensitive enough to detect between group differences in performance. Previous work has reported similar findings to these under laboratory conditions, and our results add to this literature. The shift to web-based testing, and the task design—which attempted to resemble clinical practice—may explain the discrepancies between the current results and those from previous studies, both in terms of variation in task performance within groups and in the pattern of performance between levels of expertise.

Future studies will refine the testing platform and provide insight into the development of expert visual diagnostic abilities by radiology trainees through both additional physiological methodologies (e.g. visual tracking) and longitudinal studies of trainee cohorts.

ACKNOWLEDGEMENTS

Consultant Radiologists involved: Drs Marina Easty, Anna Fohr, Kate Kingston, Jeannette Kraft, Caren Landes, Rebecca Linke, Ima Moorty, Samantha Negus, Maria Raissaki, Rui Santos, Martin Stenzel, Rick van Rijn

COMPETING INTERESTS

None

FUNDING

This work was supported by The Sheffield Children's Hospital Charity. The authors certify that there is no actual or potential conflict of interest in relation to this article.

ETHICS APPROVAL

The study received local R&D and University of Sheffield ethical approval

DISCLOSURE

The authors have no conflicts of interest to declare

PATIENT CONSENT

Not applicable

REFERENCES

1. Norman GR, Coblenz CL, Brooks LR, Babcock CJ. Expertise in visual diagnosis. *Academic Medicine* 1992; **67**: S78–83. doi: <https://doi.org/10.1097/00001888-199210000-00045>

2. Wood G, Knapp KM, Rock B, Cousens C, Roobottom C, Wilson MR. Visual expertise in detecting and diagnosing skeletal fractures. *Skeletal Radiol* 2013; **42**: 165–72. doi: <https://doi.org/10.1007/s00256-012-1503-5>
3. Wood BP. Visual expertise. *Radiology* 1999; **211**: 1–3. doi: <https://doi.org/10.1148/radiology.211.1.r99ap431>
4. Potchen EJ. Measuring observer performance in chest radiology: some experiences. *J Am Coll Radiol* 2006; **3**: 423–32. doi: <https://doi.org/10.1016/j.jacr.2006.02.020>
5. Kok EM, van Geel K, van Merriënboer JJG, Robben SGF. What we do and do not know about teaching medical image interpretation. *Front Psychol* 2017; **8**: 309. doi: <https://doi.org/10.3389/fpsyg.2017.00309>
6. Evans KK, Cohen MA, Tambouret R, Horowitz T, Kreindel E, Wolfe JM. Does visual expertise improve visual recognition memory? *Atten Percept Psychophys* 2011; **73**: 30–5. doi: <https://doi.org/10.3758/s13414-010-0022-5>
7. Bourne LE, Kole JA, Healy AF. Expertise: defined, described, explained. *Front Psychol* 2014; **5**(MAR): 186. doi: <https://doi.org/10.3389/fpsyg.2014.00186>
8. Nakashima R, Kobayashi K, Maeda E, Yoshikawa T, Yokosawa K. Visual search of experts in medical image reading: the effect of training, target prevalence, and expert knowledge. *Front Psychol* 2013; **4**: 166. doi: <https://doi.org/10.3389/fpsyg.2013.00166>
9. Drew T, Võ ML-H, Wolfe JM. The invisible gorilla strikes again: sustained inattention blindness in expert observers. *Psychol Sci* 2013; **24**: 1848–53. doi: <https://doi.org/10.1177/0956797613479386>
10. Horowitz TS. Prevalence in visual search: from the clinic to the lab and back again. *Japanese Psychological Research* 2017; **59**: 65–108. doi: <https://doi.org/10.1111/jpr.12153>
11. Berbaum KS, Smoker WR, Smith WL. Measurement and prediction of diagnostic performance during radiology training. *AJR Am J Roentgenol* 1985; **145**: 1305–11. doi: <https://doi.org/10.2214/ajr.145.6.1305>
12. Krupinski EA, Graham AR, Weinstein RS. Characterizing the development of visual search expertise in pathology residents viewing whole slide images. *Hum Pathol* 2013; **44**: 357–64. doi: <https://doi.org/10.1016/j.humpath.2012.05.024>
13. Nodine CF, Krupinski EA, skill P. Perceptual skill, radiology expertise, and visual test performance with NINA and WALDO. *Academic Radiology* 1998; **5**: 603–12. doi: [https://doi.org/10.1016/S1076-6332\(98\)80295-X](https://doi.org/10.1016/S1076-6332(98)80295-X)
14. Schuster D, Rivera J, Sellers BC, Fiore SM, Jentsch F. Perceptual training for visual search. *Ergonomics* 2013; **56**: 1101–15. doi: <https://doi.org/10.1080/00140139.2013.790481>
15. Donovan T, Manning DJ. The radiology task: Bayesian theory and perception. *Br J Radiol* 2007; **80**: 389–91. doi: <https://doi.org/10.1259/bjr/98148548>
16. Drew T, Evans K, Võ ML-H, Jacobson FL, Wolfe JM. Informatics in radiology: what can you see in a single glance and how might this guide visual search in medical images? *Radiographics* 2013; **33**: 263–74. doi: <https://doi.org/10.1148/rg.331125023>
17. Reingold EM, Sheridan H. Eye movements and visual expertise in chess and medicine. In: Oxford handbook on eye movements [Internet. *Oxford University Press* 2011;: 528–50.
18. Stanislaw H, Todorov N. Calculation of signal detection theory measures. *Behav Res Methods, Instruments, Comput* [Internet. *Springer-Verlag* 1999; **31**: 137–49. cited 2018 Jul 5.
19. Offiah AC, Moon L, Hall CM, Todd-Pokropek A. Diagnostic accuracy of fracture detection in suspected non-accidental injury: the effect of edge enhancement and digital display on observer performance. *Clin Radiol* 2006; **61**: 163–73. doi: <https://doi.org/10.1016/j.crad.2005.09.004>
20. Leung RS, Nwachuckwu C, Pervaiz A, Wallace C, Landes C, Offiah AC. Are UK radiologists satisfied with the training and support received in suspected child abuse? *Clin Radiol* 2009; **64**: 690–8. doi: <https://doi.org/10.1016/j.crad.2009.02.012>
21. Taylor GA. Perceptual errors in pediatric radiology. *Diagnosis* 2017; **4**: 141–7. doi: <https://doi.org/10.1515/dx-2017-0001>
22. Green DM, Swets JA. *Signal detection theory and psychophysics*. Oxford, UK: John Wiley & Sons, Ltd; 1966. pp. 479.
23. Manning D, Ethell S, Donovan T, Crawford T. How do radiologists do it? The influence of experience and training on searching for chest nodules. *Radiography* 2006; **12**: 134–42. doi: <https://doi.org/10.1016/j.radi.2005.02.003>