Expert radiologists have improved memory performance for images with focal abnormalities.

Hayden M. Schill¹, Jeremy M. Wolfe^{2,3}, Timothy F. Brady¹

¹Department of Psychology, University of California, San Diego

²Department of Surgery, Brigham & Women's Hospital

³Departments of Ophthalmology and Radiology, Harvard Medical School

Abstract

Memory capacity depends on prior knowledge, both in working memory and in long-term memory. For example, radiologists have improved long-term memory for medical images compared to novices. Furthermore, people tend to remember abnormal or surprising items best. This is often claimed to arise primarily because such items attract additional attention at encoding. How do expertise and abnormality interact when experts are actively searching for abnormalities; e.g. radiologists looking at mammograms? In the current work, we investigate whether expert radiologists (N=32) show improved memory performance for abnormal images compared to novice participants (N=60). We consider two types of “abnormality.” A mammogram can have a focal abnormality that can be localized or it could simply be the mammogram of a woman known to have cancer (e.g. the image of the breast contralateral to the focal abnormality). Must an image have a focal abnormality for additional attentional processing to be engaged? We found that experts have better memory for mammograms than novice participants and enhanced memory for abnormal images relative to normal images. Overall, radiologists showed no memory benefit for the contralateral-abnormal images and did not discriminate them from normal images, but had enhanced memory for images with focal abnormalities. Our results suggest that focal abnormalities play an important role in enhancing memory of expert observers.

Keywords: Expertise, Radiology, Recognition Memory, Long-term Memory

Word count: Abstract 213 words, Text 6230 words (without references)

Significance Statement

Perceptual errors in reading medical images can result in fatal consequences. Thus, in the current work we are broadly interested in what it is radiologists process and remember about medical images, and how this relates to well-known cognitive psychology effects like abnormality advantages and expertise advantages in memory. In addition to asking how expertise interacts with memory for abnormal images theoretically, we also investigate a potential actionable way to boost radiologist's performance in reading medical images based on the "crowd-within" effect, where averaging performance across multiple readings of an image – even by a single radiologist – could improve performance, particularly if the images are poorly remembered (and so judgments are relatively independent each time). By measuring the crowd-within effect in the context of a study of radiologists' memory for images, we found that memory is indeed quite imperfect even after only 30 subsequent images, and that this allows for radiologists to give somewhat independent assessments of mammograms if shown them twice: i.e., averaging a given radiologists' rating of abnormality across multiple exposures to the same mammogram resulted in better accuracy than looking at either rating alone. This suggests that there is some independent information in each judgement. Our results suggest that radiologists have improved, though far from perfect memory for mammograms; that such memory is enhanced by focal abnormalities (e.g., lesions) being present in the image; and that there may be situations in which having a second look at a case will prove to be beneficial to radiologist performance.

Imagine after playing a round of chess, you are asked to recreate the board from some critical moment in the game. For most people, this task would prove very difficult. However, if you were a world-class chess player, this might be quite easy. Becoming an expert in a domain such as chess changes our memory for items in that domain of expertise (Chase & Simon, 1973; de Groot, 1946), allowing us to store more information if this information is consistent with the expectations we have formed as a result of our expertise (Gobet & Simon, 1996).

A large literature is devoted to quantifying the specific short- and long-term memory differences in experts compared to novices (e.g., Ericsson & Kintsch, 1995; Engle & Bukstel, 1978; Vincente & Wang, 1998; Gobet & Simon, 1996). Curby, Glazek, & Gauthier (2009), for instance, showed increased visual working memory capacity in experts who specialize in cars. When observers were probed on whether a presented car was different from one seen in a previous group, experts were significantly better than novices, indicating better visual working memory capacity. Experts also display better long-term memory performance: for example, expert radiologists have been shown to have better long-term memory for mammograms compared to controls (Evans, Cohen, Tambouret, Horowitz, Kreindel, & Wolfe, 2010). This effect was limited to images in their domain of expertise, as the memory benefit did not transfer to natural scenes or isolated objects.

Why do experts show this increase in memory performance for their domain of expertise? In the expertise literature, many authors claim that at least in part, memory improvement occurs because existing knowledge allow experts to know what variation to expect for an image in an expert's domain (e.g., Vincente & Wang, 1998). In particular, experts may have developed specialized processing mechanisms for their domain of expertise which take advantage of the way stimuli vary in that domain. For example, experts employ holistic processing strategies (Watson & Robins, 2014; Richler, Wong, & Gauthier, 2011). Such holistic processing strategies may involve recruiting specialized machinery in fusiform cortex which is known to be involved in holistic face

perception (Kanwisher & Yovel, 2006) and which is also activated when experts view items in their domain of expertise (Gauthier, Skudlarski, Gore, & Anderson, 2000; Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999; Bilalić, Langer, Ulrich & Grodd, 2011).

Memory benefits in experts may also be considered a manifestation of a broader phenomenon where information that is understood as meaningful is better remembered (e.g., Bartlett, 1932; Asp, Störmer, Brady, 2019). The Baker effect is an example: People remember that someone *is* a baker more so than remembering that someone's *name* is Baker because the profession Baker activates a rich set of meaningful associations that the name Baker does not (McWeeny, Young, Hay, Ellis, 1987). Recently, Brady et al. (2016) illustrated an increase in working memory capacity for meaningful real-world objects compared to simple, non-meaningful stimuli. They found that the Contralateral Delay Activity (the CDA), a neural marker of visual working memory, was greater for meaningful stimuli, as though participants could recruit additional memory resources when the stimuli were meaningful. For experts, connecting information to meaningful concepts may result in an enriched ability to understand and remember stimuli in their domain of expertise (e.g., Ericsson & Kintsch, 1995), perhaps in addition to changes in more bottom-up perceptual processing such as holistic processing.

Memory for abnormal and surprising items

Abnormal or surprising items tend to play a critical role in many areas of expertise (i.e. a cancerous mammogram, a gun in a suitcase). The literature, largely conducted on non-expert populations, suggests that abnormal items tend to be better remembered than normal or expected items (Friedman, 1979; McDaniel & Einstein, 1986; Light, Kayra-Stuart, & Hollander, 1979; Pedzek, Whetstone, Reynolds, Askari, and Dougherty, 1989; Hollingworth & Henderson, 2003). For

example, Friedman (1979) showed observers a series of expected and unexpected images and found that fixations to the unexpected image were almost twice as long as fixations to the expected images. Furthermore, during a recognition task after presentation of the images, observers generally noticed only the changes that had been made to unexpected images. Thus, they concluded that abnormal or surprising items tend to be better remembered, in large part because they attract additional attention. A decade later, Reynolds, Askari, and Dougherty (1989) had participants sit in a mock office and view consistent and inconsistent objects for 35 seconds. After eliminating the ability to guess based off of context cues, they found a robust advantage for the recall of inconsistent objects over consistent objects, as well as a greater advantage for detecting a change in inconsistent objects across trials compared to consistent objects.

The literature on memory for surprising or abnormal items has primarily focused on observers who have no special expertise in the stimuli they are remembering (Henderson & Hollingworth, 1999; Friedman, 1979; Pedzek et al., 1989). General research on the applicability of memory for abnormal items, and how it interacts with other sources of knowledge (i.e. meaning, holistic processing) that affect what information experts tend to encode, is severely limited and yet would be highly pertinent to many domains (i.e. airport baggage screeners, the army, radiologists). Thus, while experts tend to have perceptual expertise about the relevant information to encode, we do not know if the 'abnormality' memory advantage applies to expert populations.

The current experiments

To understand how expertise effects memory and the role of abnormality on memory, the current experiment asks if expert radiologists have better memory for abnormal images (i.e. a cancerous case) compared to normal images. We focus on radiologists' memory for mammograms for three reasons: First, search for signs of breast cancer involves a usefully specific perceptual

expertise. Second, as mentioned previously, it is already known that expert radiologists have better recognition memory for medical images than do novices (Evans, Cohen, Tambouret, Horowitz, Kreindel, & Wolfe, 2010). Third, mammography offers two varieties of abnormality. Of course, an image can be abnormal because of the presence of a visible, localized abnormality like a mass or suspicious calcifications. However, a mammogram could also be considered abnormal if it simply comes from a clinically positive patient. A woman with a mass in the left breast might have no localized signs of cancer in the contralateral, right breast. Yet there is evidence that radiologists have an ability to detect a “gist” of abnormality even from the contralateral breast, and that this gist allows experts to perform at above chance levels when asked to categorize images as coming from normal or abnormal patients (Evans, Haygood, Cooper, Culpan & Wolfe, 2016). In other words, radiologists do not always need to see a physical lesion to know that an image is abnormal. These studies of the gist of abnormality have involved giving radiologists only a brief (250-500 msec) glance at an image. This is sufficient for expert radiologists to gain some evidence of abnormality, but it remains unknown whether this ability to detect a non-focal sign of abnormality (i.e. the contralateral breast) impacts radiologists’ memory for normal versus abnormal images. The extent to which an image must have a focal abnormality for additional memory processing to be engaged is particularly important as it could bear on broader questions about the source of abnormality benefits on memory and whether they require focal attention (as proposed by many; e.g., Friedman, 1979).

The goal of the current study is to provide basic knowledge about the role of abnormalities in visual memory for experts. With this in mind, the questions guiding this experiment are the following: Do radiologists show improved memory performance for abnormal images in either short- or long-term memory compared to novices without medical training? If so, do experts also show enhanced memory for non-localized abnormality signals (i.e. from contralateral-abnormal images) or does any abnormality advantage in memory depend upon having a focal abnormality?

Finally, we explore whether radiologists demonstrate the “crowd-within” effect, which states that when a person responds to the same question twice, the average of those responses is more accurate than either response alone (Vul & Pashler, 2008). To anticipate our results, we find enhanced memory for mammograms in radiologists compared to novices. We find an abnormality advantage in radiologists for focal abnormalities as well as a robust, but small, “crowd-within” effect. We find no evidence that radiologists make use to a non-focal gist of abnormality either in judgment or memory.

Experiment 1: Novices

Experiment 1 was conducted using novice (non-radiologist) observers. The design, number of participants, exclusion and analysis plan for this experiment were preregistered (URL for this experiment: <http://aspredicted.org/blind.php?x=xr3843>).

In this experiment, novice participants view a series of mammograms and judge whether they are normal or abnormal and whether they have previously seen these images. We would expect both judging whether an image is normal or abnormal as well as remembering the images to be difficult, as this task is designed for expert radiologists. However, novice performance provides a useful baseline for comparing radiologist performance, allowing us to determine, for example, whether our stimulus set contains images that are memorable regardless of expertise.

Method

Participants

60 participants (23 female, mean age 38 years) were recruited for this experiment through Amazon’s Mechanical Turk, which offers monetary compensation for participation in online tasks. Mechanical Turk workers are reasonably representative of the American adult population (Berinsky,

Huber & Lenz, 2012; Buhrester, Kwang, & Gosling, 2011; Difallah, Filatova, & Ipeirotis, 2018), and provide data that is comparable to data obtained when participants are tested in experimental psychological laboratories (e.g., see Brady & Alvarez, 2011 for a comparison in a visual memory context). All participants gave informed consent, were compensated at a rate of approximately \$10/hour, were located in the United States, and had a Hit Approval Rate greater than 95%. Informed consent procedures were approved by the Institutional Review Board of the University of California, San Diego.

Stimuli & Procedure

Observers viewed images of a single breast. The stimulus set consisted of 80 abnormal cases and 40 normal (non-cancerous) cases. All images were pre-classified by a group of trained radiologists who did not participate in the study. Half of the abnormal images contained a visible abnormality (i.e. a lesion was present) and half were images of the breast contralateral to the breast with the lesion (i.e. still an abnormal case, but with no focal indication of that abnormality). Thus, the entire set consisted of 40 normal images, 40 focal-abnormality images (herein referred to as abnormal) and 40 non-focal abnormality images (herein referred to as contralateral-abnormal). Each image subtended approximately 16 x 20 degrees of visual angle at an estimated viewing distance of approximately 60 cm from the screen.

On each trial, one image was present for 3 seconds. This image was randomly chosen to be either normal, abnormal, or contralateral-abnormal. Critically, each image was also either a new image (presented for the first time in the experiment) or a repeated image from 3 trials back or 30 trials back (3-back and 30-back, respectively). Detecting 3-backs was our index of short-term memory and detecting 30-backs was our index of more long-term memory. In total, with repetitions,

there were 210 trials: 120 new images (40 per condition), plus 90 repeat images (30 per condition, split evenly between 3 back and 30 back).

After being displayed for three seconds, each image was immediately followed by two sets of response buttons, where participants needed to indicate their answers to the following questions: (1) Was the image abnormal or normal? (2) Have you seen this image before? Using a standard computer mouse, observers were told to indicate their level of confidence on a six-point rating scale ranging from confident yes/abnormal to confident no/normal (Figure 1). We collected confidence ratings instead of simple yes/no answers to allow for receiver operating characteristic analysis (ROC). There was no time constraint imposed on responding. The initiation of the next trial was contingent on answering both questions of the current trial.

Figure 1. *Method.* N=60 novice participants rated a sequence of 210 images on normal/abnormal and old/new. Images could repeat either after 3 or 30 subsequent images and be either normal, abnormal, or contralateral-abnormal.

Before the experiment began, participants were presented with instructions and several demographic questions (Gender; Age; 'Are you a radiologist?'; 'Do you have a job where you read medical images, (i.e. tech, medical physicist)?'). Instructions were written for a novice population with no medical training. For novice participants, abnormal cases were broadly defined as “images that might contain lesions, or cancer, or otherwise might be something worthy of follow-up if you were a radiologist.” Note that even if novices were sensitive to abnormal global texture or other

index of contralateral-abnormality, these instructions may lead them to think primarily about focal lesions. However, if the images looked troubling to the experts, we might expect them to give them more attention and, thus, to remember them better.

Exclusion Criterion & Analysis Plan

Our exclusion criteria and analyses were decided in advance (see preregistration). Individual trials were excluded if participants took less than 1500ms or >15,000ms to respond (based on pilot data). Participants were excluded if they took less than 15 minutes (0 excluded) or more than 1 hour to complete the study (3 excluded), were radiologists (1 excluded) or had any prior experience reading medical images (0 excluded). Participants were also excluded if they had more than 80% identical responses (e.g., picked the exact same answer on nearly every trial; 1 observer excluded) or had more than 20% of trials excluded on the basis of the reaction time criteria (1 excluded). After applying these a priori exclusion criteria, 7 participants were excluded from analysis, leaving a final sample of 53 participants.

In our analysis, we looked at the results of the two response questions separately. First, we analyzed the confidence ratings of classifying an image as abnormal or normal. Secondly, we looked at the confidence ratings representing short- and long-term memory for images. In order to do this, we conducted ROC analysis for 3-back and 30-back as a function of image type (normal/abnormal/contralateral-abnormal). We also generated ROCs for the normal/abnormal judgments. ROCs were summarized by area under the curve (AUC) and compared using t-tests. As noted, we are interested in whether, within the group of novice participants, there is a benefit in memory performance for any type of image (e.g. as judged by normal vs. abnormal AUC). Since the novices lack medical experience, any such effect would give us insight into the nature of the image set (i.e. memorability or distinctiveness), rather than any medical knowledge of the novice. Finally,

we conducted image similarity analyses to quantify how image differences might be influencing these results.

Image similarity analysis

Because normal, focally abnormal and contra-abnormal images are necessarily different image sets, it is useful to compare how distinctive each set of images is from all the other images people see in order to look at the effect this has on memory. To do so, we conducted a gabor wavelet pyramid (GWP) analysis, which computes features of the images and compares them (Kay, Naselaris, Prenger, & Gallant, 2008; Greene, Baldassano, Esteva, Beck, & Fei-Fei, 2016). To assess the level of similarity in the different image types, the GWP represents each image as the output of a bank of multi-scale Gabor filters. Prior work has shown that these features can successfully model object representation in early visual areas (Kay et al., 2008). Following the exact procedure and parameters provided by Greene et al. (2016), each image was converted to grayscale, down sampled to 128 by 128 pixels, and represented with a bank of Gabor filters at three spatial scales (3, 6 and 11 cycles per image with a luminance-only wavelet that covers the entire image), four orientations (0, 45, 90 and 135 degrees) and two quadrature phases (0 and 90 degrees). An isotropic Gaussian mask was used for each wavelet, with its size relative to spatial frequency such that each wavelet has a spatial frequency bandwidth of 1 octave and an orientation bandwidth of 41 degrees (Greene et al., 2016). This gave a set of features for each image, which we then compared to all 120 images to compute a distance / dissimilarity score by computing the dot product of each images features to each other images after subtracting the mean across images and normalizing the feature vectors to unit length.

Results (Experiment 1: Novices)

Performance on the classification task

First, we looked at how confident novices were at classifying an image as either normal or abnormal for each image type (Figure 2). We found a significant difference between normal and abnormal images ($t(52) = 4.78, p < 0.001$), but not between normal and contralateral-abnormal images ($t(52) = 1.94, p > 0.05$).

Figure 2. *Classification Task:* Overall performance in novices on labeling an image as normal or abnormal. The confidence rating scale is now plotted on the y axis. We found a significant difference in confidence in classifying normal vs. abnormal images, which seems to be driven by a few salient abnormal images. Novices are not confident in distinguishing between any image type. Error bars represent standard error or the mean.

While participants did not have training to distinguish normal from abnormal medical images, a small number of images in the set are extremely saliently abnormal (i.e. a salient white spot, which looks abnormal even to novice observers). Looking at ratings by image (Figure 3) reveals that these images are largely responsible for this difference between normal and abnormal ratings. Thus,

at least for a small subset of images, even novice observers can notice the abnormality, leading to above chance classification performance broadly.

Figure 3: *Image ratings in the classification task.* Example images and their confidence ratings for each observer in the classification task. As can be seen with the third image, most observers rated this as abnormal with high confidence. Altogether, the two or three brightly striped vertical lines in the abnormal image set indicate that those and only those images were reliably rated as abnormal by a large majority of participants.

Note that the y-axis in Figure 2 represents the confidence ratings for novices. It is clear that the novices are generally not confident in distinguishing any image type, with average responses tightly clustered near the middle of the rating scale for all conditions. Another way of visualizing this data is on an ROC curve (Figure 4), where novices fall almost on top of the dotted diagonal line indicative of chance performance, with an AUC of only 0.54, although as above, this difference from chance is highly reliable across participants ($t(52) = 4.21, p < 0.001$).

Figure 4. *ROC for normal/abnormal categorization.* Novices are very close to the diagonal line representative of chance performance, indicating that they do not perceive a strong difference between normal and abnormal images. The significant effect is driven by a few salient images (see Figure 3).

Memory for abnormal images

Figure 5 shows the ROCs for the 3-back and 30-back memory tasks. Since novices were not, for the most part, able to perceive contralateral-abnormal images as different from normal images in the classification task, we focused exclusively on memory differences between normal and abnormal images. Overall, independent of image type, novices have better short-term memory (averaged AUC of 0.70 for detecting 3-backs) than long-term memory (averaged AUC of 0.64 for detecting 30-backs). Interestingly, breaking down performance across image conditions reveals that novices show a small normality benefit: they remember normal images better than abnormal images in both the 3-back condition and the 30-back condition, with only the 30-back yielding a significant difference. We found an AUC benefit of 0.069 for normal images at 3-back ($t(52)=5.48$, $p<0.001$) compared to abnormal, and an AUC difference of 0.026 for normal images at 30-back ($t(52)=1.70$, $p = 0.096$) compared to abnormal.

Figure 5: *Novice performance on the memory task.* As noted, the dashed line indicates chance and the higher the curve represents better memory performance. Novices had stronger memory for images in the 3-back (short-term) condition than in the 30-back (long-term) condition. Novices also show a small effect of normality with memory for normal images being better than for abnormal images in both 3 and 30-back conditions.

Given the weak performance at discriminating normal from abnormal images, it is rather surprising that normality had any effect. Since it is unlikely that the advantage for normal cases has anything to do with their medical status, we examined the data for evidence of more basic effects of visual similarity. The lower performance in the abnormal conditions is largely driven by an increased false alarm rate in the contralateral-abnormal and abnormal image sets (based on classifying “new” as all images with rating > 3). This is consistent with an image similarity account in which novices would be more likely to false alarm to new images in the contralateral-abnormal and abnormal conditions simply because these images are more similar to one another than images in the normal set. In other words, if the normal images were more dissimilar to each other compared to the others,

this would explain why novices have better memory for the normal condition (i.e. it is easier to determine if an image of a dog is new if that dog is presented in a series of different animals than if it is presented in a set of similar dogs. Obviously, the similarity effects in our stimuli are much smaller.).

Similarity Matrix - Gabor Wavelet Pyramid Analysis:

We tested this image similarity hypothesis by conducting a gabor wavelet pyramid (GWP) analysis, which computes features of the images and compares them, as described in the Methods (Kay, Naselaris, Prenger, & Gallant, 2008; Greene, Baldassano, Esteva, Beck, & Fei-Fei, 2016).

We found increased dissimilarity among normal images relative to contralateral-abnormal and abnormal images – with normal images being more dissimilar, on average, to the other images. This is consistent with the small difference in memory favoring normal images being driven by image similarity differences across sets (Figure 6). Thus, the small normality benefit found in the current study is likely a result of image similarity, which can provide a useful baseline for considering memory for the same images in expert radiologists in Experiment 2.

Figure 6. *Gabor Wavelet Analysis.* Level of distinctiveness across image types produced by the gabor wavelet analysis. Normal images were found to be more distinctive (and thus more discriminable in memory) than either abnormal or contralateral-abnormal images.

Experiment 2: Radiologists

Experiment 2 was the same as Experiment 1, except conducted on radiologist observers.

Method

Participants

32 expert radiologists (14 female, average age = 49 years) were recruited during the 2018 Radiological Society of North America (RSNA) conference in Chicago, Illinois. All radiologists gave informed consent and were not compensated (beyond being entered into a lottery for a \$500 gift card). Informed consent procedures were approved by the Institutional Review Board of the University of California, San Diego.

Similar to Experiment 1, data from observers were excluded if they took less than 15 minutes or more than 1 hour to complete the study, had more than 80% identical responses or had more than 20% of trials excluded. Under these guidelines, no radiologists were excluded from analysis, leaving a final sample of 32 participants.

Stimuli & Procedure

Stimuli were the same as those described in Experiment 1. The main procedural difference was that the experiment was conducted at the RSNA conference in Chicago, IL, where the experimenters could explain the instructions in person. The experimental design was the same as described in Experiment 1. Unlike in Experiment 1, in Experiment 2, we gave more general instructions, asking for any abnormality rather than specifically asking participants to look for focal lesions or cancer: "For each image please judge whether the image is abnormal or normal, and whether you have previously seen it during the course of the experiment."

Results

In this section, we compare the performance of expert radiologists to the performance of novice participants in Experiment 1. In particular, we will investigate how non-experts compare to experts in judgments of image classification (i.e. normal vs. abnormal), and critically, whether experts show differential memory for abnormal versus normal images.

Performance on the classification task

Similar to Experiment 1, we first analyzed performance on the classification task by looking at the confidence ratings of labeling each image as either normal or abnormal: how good are radiologists at simply distinguishing abnormal from normal images? Unsurprisingly, radiologists are very good at distinguishing abnormality (Figure 7A). Radiologists were significantly more confident that an abnormal image was abnormal instead of normal ($t(31) = 13.17, p < 0.001$). There was not a significant difference between the normal and contralateral-abnormal image conditions ($t(31) = 0.43, p = 0.67$). In the original study of Evans et al. (2016), they found this signal of abnormality in the gist information (~ 250 ms). It is likely that observers utilize different decision processes at the different time scales (i.e. gist vs. 3 seconds in our experiment), which would explain the differences in our results. Recall, however, that our interest is in their memory for these images. Contralateral-abnormal images might be remembered better if their vaguely suspicious appearance caused radiologists to devote more attention to them.

Figure 7B shows the ROC curve for distinguishing focal-abnormal images from normal images in radiologists. ROCs were summarized by area under the curve (radiologists = 0.72, recall that novices = 0.54). As noted in Experiment 1, controls are close to the diagonal line indicative of chance, whereas radiologists show a typical curvilinear ROC indicative of a perceived (and significant) difference between the two conditions with an AUC well above chance ($t(31) = 19.8, p < 0.001$).

Figure 7A: Bar graph depicting confidence rating scale on the y axis. Radiologists clearly distinguished abnormal from normal images, but they did not distinguish between contralateral-abnormal and normal images. **7B:** ROC depiction of performance for labeling an abnormal image as abnormal instead of normal (ignoring contralateral-abnormal images).

Memory for abnormal images

Figure 8 shows performance by radiologists on the memory task. Radiologists have better memory for abnormal images in both memory conditions, but the advantage for abnormal images is only significant in the 30-back condition ($t(31)=-2.86$, $p<0.05$). We found an AUC advantage of 0.2 for abnormal images at 3-back, although this was not significant ($t(31)=1.62$, $p=0.12$).

Figure 8. *Radiology performance on the memory task.* Radiologists have better memory for abnormal images in both the short- and long-term memory conditions. However, only memory at long delays (30-back) was significant.

Radiologists showed no memory benefit for the contralateral-abnormal images, even at long delays ($p=0.24$). Since radiologists were not able to distinguish between contralateral-abnormal images and normal images in the classification task, this result would be expected. Overall, independent of image type, radiologists have better short-term memory (averaged AUC of .852 for detecting 3-backs) than long-term memory (averaged AUC of .752 for detecting 30-backs) for medical images. Taken together, these results suggest that, in general, VSTM is stronger than LTM and that experts have better memory for abnormal images particularly at long delays. Recall from the wavelet analysis in Experiment 1 that the normal images in our data set are less similar to each other than the abnormal images, and thus based on simple similarity, memory for normal images should be better than abnormal. The fact that abnormal image memory is better in radiologist observers suggests that the effect of expertise more than compensates for differences between the stimulus categories in image similarity. To see what the effect of abnormality is, independent of baseline image similarity differences, we can compare radiologists' memory performance to novices'

performance with the same images (in terms of d' , a linear scale, so we can subtract). Doing so reveals a significant abnormality benefit at both 3-back ($t(31)=-3.61, p=0.001$) and 30-back in expert radiologists ($t(31)=-3.69, p=0.001$), where relative to novice participants, radiologists were specifically better at remembering abnormal images.

Figure 9. Using novices as a baseline to account for image similarity, there were robust abnormal effects at both 3-back and 30-back.

Relationship between classification and memory

Figure 10 once again dichotomizes both judgments of abnormality (abnormal > 3 as ‘abnormal’) and old/new to create d' scores for each participant. It then plots overall memory performance across all conditions for individual radiologist and novice participants against how accurately they discriminate abnormal from normal images, showing a clear dichotomy in memory strength and categorization accuracy between novice and expert readers. Taken together, memory d' is correlated with classification d' ($r=0.57, p<0.0001$) but if each group is analyzed separately, this relationship is not reliable (radiologists: $r=0.29, p=0.10$, novices: $-0.10, p=0.48$).

Figure 10. Individual performance showing the difference between radiologists and novices on the classification and memory tasks.

“Crowd-within” Analysis

Finally, the structure of this experiment means that we have two classification ratings for each item in the memory set. This makes it possible for us to combine them in order to test the “crowd-within” hypothesis of Vul & Pashler (2008). They found that averaging a single person’s responses to the same question results in better performance than single responses alone. This is what one would expect if a single judgment did not incorporate all of the information people could possibly have about a question. If this is true for assessments of mammograms by expert radiologists, at least with the limited viewing time of the current study, we expect averaging a given radiologist’s rating of abnormality across multiple exposures to the same mammogram should result in better accuracy than looking at either rating alone.

In fact, we did find that averaging radiologists’ responses from the first and second time that they saw an image resulted in slightly higher performance in the 30-back condition ($d' = 0.97$;

AUC=0.745) compared to single item performance ($d' = 0.94$; AUC=0.716; $t(31)=3.46$, $p=0.002$) (Figure 11), but not in the 3-back condition ($t(31)=1.15$, $p=0.259$). Unsurprisingly, this effect was not present in novices, since their performance was very poor on both responses (Figure 12; all $p>0.10$). Thus, expert performance can be improved (albeit, rather modestly) by averaging more than one response. It remains to be seen whether this benefit would occur if radiologists were offered unlimited time to process each image.

Figure 11. *Crowd-Within Analysis:* Left, Radiologists (Exp. 2): The blue line is the ROC for distinguishing focally abnormal mammograms vs. normal mammograms when the radiologists first see the image. The red line is the average of the first time seeing it and their responses seeing it at 30-back. Right, Novices (Exp. 1): Once again, the blue line is the ROC for distinguishing focally abnormal mammograms vs. normal mammograms when novices first see the image. The red line is the average of the first time seeing it and their responses seeing it at 30-back.

General Discussion

In the current study, we examined short- and long-term memory performance in novices and radiologists for normal versus abnormal mammography images. First, across experiments, we looked at how confident and how competent novice and expert observers were at classifying medical

images as either normal or abnormal. Unsurprisingly, radiologists were much better than novices at this task. Novices did show some ability to distinguish abnormality, although this appeared to be largely the result of a few salient images. In memory for the images, we found a small normality benefit in novice participants' memory. They were slightly more likely to accurately recall whether or not they had seen one of the normal images. A gabor wavelet analysis (Greene et al. 2016) suggested that this was a function of stimulus similarity. Similarity was higher in the abnormal group of images. It is well known that stimulus discriminability strongly affects visual recognition (e.g., Bellhouse-King & Standing, 2007). This suggests that the more distinctive normal images were more memorable to novices, thus explaining why novices remembered normal images better than any other image type.

Interestingly, memory for abnormal images was better for the radiologists. This shows how expertise changes distinctiveness. For an expert, the abnormal images would have an added attribute (that mass, that calcification) that would help to distinguish the item in memory. Previous work has found mixed results when investigating memory improvements in radiologists. For example, Hardesty et al. (2005) investigated radiologists' long-term memory for medical images presented months later and found that none of the radiologists remembered cases that they had read previously. Evans et al. (2016) found mixed results when investigating whether abnormality improves memory in expert observers, including radiologists. Our results provide context to these ambiguities, as they suggest that expert radiologists do have stronger memory for abnormal images even in a long-term memory setting. However, our long delays were only on the order of minutes, not months, and so it remains unclear how such advantages would last over long durations.

It is worth noting that in the classification task, radiologists performed on average much worse than radiologists in the clinic would with unlimited viewing time ($d' = 2.5-3.0$, as in D'Orsi et al., 2013). One reason for this might be that each image in our study was only presented for 3

seconds each. For instance, Evans et al. (2013) showed radiologists only a brief glimpse of mammograms and varied timing from 250ms to 2000ms. The AUC for radiologists in their experiment for 500ms, 1000ms, and 2000ms viewing times was 0.65, 0.66, and 0.72. In our experiment with a presentation time of 3000ms, we found an AUC of 0.72. Thus, our 3000ms presentations resulted in a similar level of performance to Evans et al. (2013) 2000ms presentations, which is well below what is expected with unlimited viewing time and consistent with viewing time being the main constraint that leads to lower performance in our study.

Finally, we looked at whether radiologists show a “crowd-within” effect (Vul & Pashler, 2008). When presented with the same case twice, radiologist performance improved when averaged across those responses compared to either response alone, but only in the 30-back condition and only modestly even then. This suggests that by the time radiologists are presented with the same images 30 images later, they can give a response that is somewhat independent of their first response, and that there is information that the radiologists are not using the first time they see an image. A “crowd-within” technique might be of some clinical use, although it would be important to determine whether such benefits persist when radiologists are given unlimited time to process the images rather than only the 3 seconds used in the current study as well as whether this effect can be made larger with an even larger delay between the first and second presentation of an image (as found by Vul & Pashler, 2009).

Given the Evans et al., (2016) finding that there is a “gist of abnormality” that is present in the contralateral breast, even when no localizable abnormality is present, we were interested to know if these contralateral-abnormal images had any advantage over normal images in expert memory. We found no such evidence. In our experiment, we also found no difference in the immediate abnormality ratings of those images compared to normal. It is likely we did not find this result due to the fact that we presented images at longer encoding time (3000ms). Moreover, typical stimulus

exposure in mammogram “gist” studies has been less than a second; 500 msec is typical. It is possible that presenting images for longer encoding times might actually obscure the gist information - overwriting an initial “gist” impression with more semantic or meaningful information. It is possible that we could observe the contra-abnormal effect even at long encoding times if we explicitly directed participants to look for a more general abnormal texture or gist, since standard instructions may lead radiologists to expect that ‘abnormal’ would imply a focal lesion of some sort. By that measure, these contralateral-abnormal images would not be labeled as abnormal.

Our findings have implications for any applied field that utilizes expert intelligence in making inferential decisions. In particular, understanding how memory differs in experts is critical in situations where decisions need to be made by people who have significant expertise. Just in radiology, for instance, there may be situations in which having a second look at an image or a case will prove to be valuable and future studies should determine the potential usefulness and feasibility of integrating this strategy to the clinic. Furthermore, understanding expertise is important for many other professions where it may be informative to understand how the expert differs from the novice.

References

- Asp, I., Störmer, V. S., & Brady, T. F. (2019.). Greater visual working memory capacity for visually-matched stimuli when they are recognized as meaningful.
<https://doi.org/10.31234/osf.io/r6njf>
- Bartlett, F. C. (1932). *Remembering: An experimental and social study*. Cambridge: Cambridge University
- Bellhouse-King, M. W., & Standing, L. G. (2007). Recognition Memory for Concrete, Regular Abstract, and Diverse Abstract Pictures. *Perceptual and Motor Skills*, 104(3), 758–762.
<https://doi.org/10.2466/pms.104.3.758-762>

- Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk. *Political Analysis*, 20(3), 351–368.
<https://doi.org/10.1093/pan/mpr057>
- Bilalic, M., Langner, R., Ulrich, R., & Grodd, W. (2011). Many Faces of Expertise: Fusiform Face Area in Chess Experts and Novices. *Journal of Neuroscience*, 31(28), 10206–10214.
<https://doi.org/10.1523/JNEUROSCI.5727-10.2011>
- Brady, T. F., & Alvarez, G. A. (2011). Hierarchical Encoding in Visual Working Memory: Ensemble Statistics Bias Memory for Individual Items. *Psychological Science*, 22(3), 384–392.
<https://doi.org/10.1177/0956797610397956>
- Brady, T. F., Störmer, V. S., & Alvarez, G. A. (2016). Working memory is not fixed-capacity: More active storage capacity for real-world objects than for simple stimuli. *Proceedings of the National Academy of Sciences*, 113(27), 7459–7464. <https://doi.org/10.1073/pnas.1520027113>
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data? *Perspectives on Psychological Science*, 6(1), 3–5.
<https://doi.org/10.1177/1745691610393980>
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4(1), 55–81.
[https://doi.org/10.1016/0010-0285\(73\)90004-2](https://doi.org/10.1016/0010-0285(73)90004-2)
- Curby, K. M., Glazek, K., & Gauthier, I. (2009). A visual short-term memory advantage for objects of expertise. *Journal of Experimental Psychology: Human Perception and Performance*, 35(1), 94–107.
<https://doi.org/10.1037/a0015233>
- de Groot, A. D. (1946). *Het denken van den schaker: een experimenteel-psychologische studie*. Amsterdam, The Netherlands: Noord-Hollandsche Uitgevers Maatschappij.

- Difallah, D., Filatova, E., & Ipeirotis, P. (2018). Demographics and dynamics of Mechanical Turk workers. In *Proceedings of the Eleventh ACM international conference on web search and data mining* (pp. 135-143). ACM.
- D'Orsi, C. J., Getty, D. J., Pickett, R. M., Sechopoulos, I., Newell, M. S., Gundry, K. R., ... D'Orsi, E. M. (2013). Stereoscopic digital mammography: Improved specificity and reduced rate of recall in a prospective clinical trial. *Radiology*, 266(1), 81–88.
- Engle, R. W., & Bukstel, L. (1978). Memory processes among bridge players of differing expertise. *The American Journal of Psychology*, 673-689.
- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, 102(2), 211–245. <https://doi.org/10.1037/0033-295X.102.2.211>
- Evans, K. K., Cohen, M. A., Tambouret, R., Horowitz, T., Kreindel, E., & Wolfe, J. M. (2011). Does visual expertise improve visual recognition memory? *Attention, Perception, & Psychophysics*, 73(1), 30–35. <https://doi.org/10.3758/s13414-010-0022-5>
- Evans, K. K., Haygood, T. M., Cooper, J., Culpan, A.-M., & Wolfe, J. M. (2016). A half-second glimpse often lets radiologists identify breast cancer cases even when viewing the mammogram of the opposite breast. *Proceedings of the National Academy of Sciences*, 113(37), 10292–10297. <https://doi.org/10.1073/pnas.1606187113>
- Friedman, A. (1979). Framing pictures: The role of knowledge in automatized encoding and memory for gist. *Journal of Experimental Psychology: General*, 108(3), 316–355. <https://doi.org/10.1037/0096-3445.108.3.316>
- Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature neuroscience*, 3(2), 191.

- Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform face area increases with expertise in recognizing novel objects. *Nature neuroscience*, 2(6), 568.
- Greene, M. R., Baldassano, C., Esteva, A., Beck, D. M., & Fei-Fei, L. (2016). Visual scenes are categorized by function. *Journal of Experimental Psychology: General*, 145(1), 82–94.
<https://doi.org/10.1037/xge0000129>
- Gobet, F., & Simon, H. A. (1996). Recall of random and distorted positions: Implications for the theory of expertise. *Memory & Cognition*, 24, 493-503
- Henderson, J. M., & Hollingworth, A. (1999). HIGH-LEVEL SCENE PERCEPTION. *Annual Review of Psychology*, 50(1), 243–271. <https://doi.org/10.1146/annurev.psych.50.1.243>
- Hollingworth, A., & Henderson, J. M. (2003). Testing a conceptual locus. *Memory & Cognition*, 31(6), 930–940. <https://doi.org/10.3758/BF03196446>
- Kanwisher, N., & Yovel, G. (2006). The fusiform face area: a cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 361(1476), 2109–2128. <https://doi.org/10.1098/rstb.2006.1934>
- Kay, K. N., Naselaris, T., Prenger, R. J., & Gallant, J. L. (2008). Identifying natural images from human brain activity. *Nature*, 452(7185), 352–355. <https://doi.org/10.1038/nature06713>
- Light, L. L., Kayra-Stuart, F., & Hollander, S. (1979). Recognition memory for typical and unusual faces. *Journal of Experimental Psychology: Human Learning and Memory*, 5(3), 212.
- McDaniel, M. A., & Einstein, G. O. (1986). Bizarre imagery as an effective memory aid: The importance of distinctiveness. *Journal of experimental psychology: Learning, memory, and cognition*, 12(1), 54.
- McWeeny, K. H., Young, A. W., Hay, D. C., & Ellis, A. W. (1987). Putting names to faces. *British Journal of Psychology*, 78(2), 143–149. <https://doi.org/10.1111/j.2044-8295.1987.tb02235.x>

- Pedzek, K., Whetstone, T., Reynolds, K., Askari, N., & Dougherty, T. (1989). Memory for real-world scenes: The role of consistency with schema expectations. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 15, 587-595.
- Richler, J. J., Wong, Y. K., & Gauthier, I. (2011). Perceptual expertise as a shift from strategic interference to automatic holistic processing. *Current directions in psychological science*, 20(2), 129-134.
- Torralba, A., Fergus, R., & Freeman, W. T. (2008). 80 Million Tiny Images: A Large Data Set for Nonparametric Object and Scene Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(11), 1958–1970. <https://doi.org/10.1109/TPAMI.2008.128>
- Vincente, K. J., & Wang, J. H. (1998). An ecological theory of expertise effects in memory recall. *Psychological review*, 105(1), 33.
- Vul, E., & Pashler, H. (2008). Measuring the Crowd Within: Probabilistic Representations Within Individuals. *Psychological Science*, 19(7), 645–647. <https://doi.org/10.1111/j.1467-9280.2008.02136.x>
- Watson, T. L., & Robbins, R. A. (2014). The nature of holistic processing in face and object recognition: current opinions. *Frontiers in Psychology*, 5. <https://doi.org/10.3389/fpsyg.2014.00003>