THE UNIVERSITY OF AUCKLAND

SEMESTER TWO, 2019
Campus: City

COMPUTER SCIENCE

COMPSCI 725 Practice Exam
Sample Answers

(Time allowed: TWO hours)

NOTE: Do not write your name on your answer sheet!

Attachment 1 is referenced in Question 1. It presents experimental data from Mohanty (2019).

Attachment 2 lists the required readings in this course.

This is a **closed-book** examination.

Calculators are **not permitted** in this examination.

1. (15 marks, in total) The following questions refer to the experimental data from Mohanty (2019), as reproduced in Appendix 1 of this examination.

(a) (10 marks) You are a member of a digital-forensics team that is defending Mr X – who is on trial for murder. You have been asked to point out weaknesses in the prosecution's argument that Mr X's camera was used to take an incriminating photo.

Here is the prosecution's argument. Police have recovered an incriminating picture from a social media website. Police have also recovered a camera from Mr X. The police's forensics team have consulted with an independent expert, who advised them on how to compute an e-PRNU fingerprint digest by using photos taken by Mr. X's camera. The independent expert also advised them to use nine different cameras to take nine additional photos (one per camera). The expert then performed the e-PRNU matching process, as described in Mohanty (2019). The expert's report states: "It is extremely likely that Mr X's camera was used to take the incriminating photo, because the e-PRNU of this camera is an excellent match to the PRNU of the incriminating photo. Furthermore, the e-PRNU of Mr. X's camera is not a good match to the PRNU of any of the nine other photos provided."

Mr. X insists that his camera could not have been used to take the incriminating photo. Your job is to provide expert advice to Mr. X's lawyers. Is there any reasonable doubt about the validity of the police expert's report, based on your knowledge of Mohanty (2019)? Explain briefly.

- Marking rubric: 5 marks for identifying at least one of the major weaknesses, and 5 marks for an understandable and accurate discussion of this weakness.

  Model answer, from Clark: I'd characterise this report as being vague, but not invalid, when it is read by an expert in forensic testing. However when this report is read by a non-expert, it might be misinterpreted as providing strong evidence that "Mr. X's camera took the incriminating photo." This would be an invalid interpretation, for several reasons.

  Weakness #1: Police selected nine "different cameras". However these cameras may have been selected in some way which makes their fingerprints likely to be very different to the one in the incriminating image.

  Weakness #2: Police used their nine cameras to take nine images. However these images may, for various technical reasons which I won't explain here, have very little recoverable PRNU – thereby making it just about certain that the fingerprint of Mr. X's camera doesn't match any of them.

  Weakness #3: Prosecution's expert report vaguely describes a way of finding an "excellent match" to the fingerprint of Mr. X's camera. However prosecution's expert has offered no evidence that they're following any established forensic methodology when selecting cameras, extracting fingerprints, taking images, and matching fingerprints to images. It seems very likely to me that there has been no careful study of the forensic validity of the results of the e-PRNU matching method. Indeed, prosecution's expert may be relying on a single experimental finding which

has not even been replicated! So: I suggest we argue that the report has the forensic credibility of the testimony of an eyewitness with a high blood-alcohol level. This report gives us a vague and possibly very-distorted idea of the linkage (if any) between Mr. X's camera and the incriminating photo, but it cannot reveal anything more than this.

Weakness #4: Although I have no ROC curves for a forensically-sound matching method, Figure 4 of Mohanty (2019) suggests to me that there must be millions of digital cameras with fingerprints which, if they were given to prosecution's expert, would have been reported to be an "excellent match" to the incriminating photo. Unless the prosecution provide persuasive evidence that only a small number of cameras were in the vicinity of the area in which the incriminating photo was taken, Mr. X's camera is extremely unlikely to be the one which took the incriminating image. Here's my detailed reasoning: prosecution's expert is unlikely to have used a match threshold for which TPR < 0.90, for otherwise they'd have an unacceptably high probability (greater than 10%) of a false negative. At the TPR $= 0.90$ point on the ROC curve for e-PRNU in Mohanty (2019), the FPR is about 20%. If anyone were to extract a PRNU fingerprint from ten million digital cameras, I'd be confident they'd find millions of (what prosecution's expert has described as) "excellent" matches to the incriminating image. To summarise: prosecution's expert might be asked to reveal the FPR for their testing methodology, and they might also be asked to estimate the number of cameras currently in existence which would give a false-match to the incriminating image. I'd be extremely confident that their answer (or their refusal/inability to answer) would expose the forensic weakness of the e-PRNU matching process in prosecution's case against Mr. X.]

- Student answer #1, with marking comments in blue ink and square brackets.
  Regarding the case against Mr. X:

  I see three main issues which cast doubt on the certainty of the police claims.

  Firstly, the fact that there was a single image pulled from social media, it is hard to guarantee that the image was not tampered or processed [OK, this is indeed a weakness in prosecution's case, although it is not an invalidity in the expert's report. The chain of evidence on the incriminating photo links it to a social-media site, so its PRNU or visually-apparent content – or both – may have been modified prior to uploading] this is said to have little effect on the PRNU however this was analysed on the standard method not e-PRNU so there is no data to prove this. [I don't understand your argument here. Perhaps you don't recall that the fingerprint digest described in Mohanty (2019) is a subset of the PRNU fingerprint of a camera? I'd say expert is somewhat confusing, but not invalid, to be reporting on their matching of "the e-PRNU of Mr. X's camera to the PRNU of any of the nine photos".]

  Secondly, the police took the advice of an ex[p]ert on how to perform these complex procedures. I would ask what sort of experience these police have in doing this sort of technical analysis. [OK, this question may reveal some information helpful

to Mr. X. But: this is not a weakness in expert's report. Furthermore, the defense shouldn't waste time on a variety of "fishing expeditions" in the hope of discovering weaknesses, instead they should be focussing on a few lines of enquiry where they have good reason to believe they'll soon be developing a strong defense. You haven't explained why it would be helpful, to defense, to establish that the police lack expertise in e-PRNU beyond what they have already demonstrated: an ability to extract a fingerprint from Mr. X's camera that matches the incriminating image but not any of the other images.]

Thirdly, and most importantly, the expert advised the police to do the matching on 9 different cameras 1 image each, however used the fingerprint digest of many photos from the suspect camera. This could be seen as inconsistent process for each camera. [?? I'm concerned that you don't understand PRNU matching. The expert doesn't need fingerprints of other cameras to support their finding of an "excellent match" of one camera's fingerprint to the PRNU in just one of the ten images. And police can't extract a fingerprint from Mr. X's camera without taking multiple images from it.] Additionally, if the investigators are looking at completely different Types and B[r]ands of cameras e.g. Sony, [Canon], Lei[c]a / Wecam goPro etc. the[n] it is clear that they would likely have different PRNU signatures. [Well done on noticing this!] Given that there are only 10 others of completely differ[ent] sources the validity of their conclusion seems to be a stretch. [Yes, I can see that this is indeed a weakness, but you're not explaining it adequately.]

[10 marks. Your answer is something of a scattershot, in that it identifies five different issues without adequately explaining how any of these demonstrate a "weakness" in the expert's report which would raise a "reasonable doubt" about its validity. Despite this lack of responsiveness to the question asked , your scattershot did hit two of my expected targets, so I'm awarding full marks for its demonstration of excellent technical competence in applying information from Mohanty (2019) to this particular forensic application.]

- Student answer #2.

  We *first* argue about any weaknesses in the prosecution's argument:
  - *Assuming* the experiment setup and methodology from Mohanty (2019) in Attachment 1 is *reasonable* [?? what do you mean by "reasonable"?], the independent expert's advice of using nine 'different' cameras does *not* state they are of different emphModels/Brands. However in Mohanty's method, this is explicitly stated. This may have created a bias in the prosecution's argument to, say, one particular camera model. [Indeed there is a potential for bias here, well done on noticing it! However your answer is incomplete, for if defense were to argue successfully for this bias to be removed, would this be strengthening prosecution's case?]
  - Further, only a *single* photo is taken by each camera versus 40 query images done by Mohanty (2019).[This is a possibly-relevant difference in experimental

setup, so shows understanding of Mohanty (2019); but you haven't indicated why it raises any "reasonable doubt" about the expert's report. I can't see any relevance.]

We can then assess the *extent* to which Mohanty (2019)'s e-PRNU *methodology itself* (compared to the "conventional scheme") may be valid in this context. This may provide some degree of doubt in the prosecution's argument at the very least (the above details about the differing camera models is ambiguous). [Well done, see Weakness #3 in my model answer.]

[9 marks. This student has shown an ability to compare the experimental design described in Mohanty (2019) to the experimental design described in this question. They have devised and explained a clear methodology for their analysis of the situation, which would make it easier for other members of the defense team to check their findings for accuracy (aka "precision") and completeness (aka "recall"). I'd characterise their answer as a scattershot with precision = 1/3, because of its two "misses" (which were almost devoid of explanation), and its adequate identification and explanation of my model answer's Weakness #3.

Note: I'm not assessing answers on recall, because my exam question did not call for this. To demonstrate that a conclusion is invalid, I'd recommend you explain one reason carefully and understandably. However it's usually a good idea to start by listing a number of possible reasons for invalidity in a "brainstorming" session (i.e. to generate a list with high recall), then to focus your attention on the most promising item in your list (i.e. to distinguish the TP items from the FP items in your list). This student has done an excellent job of accurately identifying their third idea as being the most important one for their defense team to pursue, thereby raising the precision of their short-listed recommendation to 1/1. Well done!]

- Student answer #3.

  The report of police expert's report could be subject to a false positive. The e-PRNU generated could have error rates which led to a false positive. This is why PRNU of X's camera led to a match to PRNU of incriminating photo. [This answer correctly identifies "false positive" as a very important consideration in Mr. X's defense. The reference to "error rates" in the "e-PRNU generated" is confusing; but the student may be referring to the increased FPR caused by the approximation process of computing a fingerprint digest. I'm pretty sure this student has good understanding of e-PRNU, and they have shown some ability to apply this understanding to the forensic scenario of this question. This student would have received excellent marks if they had explained, in an understandable fashion, that the prosecution's (unexplained) decision to use e-PRNU matching has resulted in them obtaining less reliable evidence than they would have been able to obtain from a trustworthy external expert who used conventional PRNU matching. 7 marks.]
- Student answer #4.

1) Each individual camera has its own e-PRNU fingerprint, it is not safe to say if the nine different cameras are not a good match to the PRNU of the incriminating photo, then Mr. X's camera was used to take the incriminating photo. [Your reasoning is unclear, and it would strengthen the prosecution's case considerably if defense would accept your proposition that "every camera has its own e-PRNU fingerprint". Mohanty (2019) don't explain why true-positive and false-positive errors are visible in the "conventional method" (blue) ROC curve of Figure 4; however one possibility is that matching errors are unavoidable due to the lack of diversity in fingerprints of different cameras.]

Also, the nine different cameras only take one additional photos each to calculate the e-PRNU fingerprints, the error rate is not balanced and there exists chance to calculate the e-PRNU fingerprint wrongly. [No, the expert didn't compute a fingerprint for any camera. The only e-PRNU fingerprint discussed in this case is the one from Mr. X's camera. I can't guess what you mean by "error rate is not balanced" – but if you are asserting that the expert didn't select a match-threshold for which FPR = TPR, you'd have to explain how you learned this from their report, and why such a choice would cause some assertion (but which one?) in their report to be invalid. You haven't explained why an erroneous calculation of an e-PRNU fingerprint would cause Mr. X's camera to be falsely matched to the incriminating image.]

2) The incriminating photo was recovered from a social media website and there is no evidence that the photo has not been modified in terms of PRNU. [Good point! However, the possibility of a PRNU-modification attack was not discussed in Mohanty (2019), so – unless someone on your team can describe how such an attack may plausibly have occurred and then quickly collect some evidence to show that it did in fact occur, I'd say it's a very low-priority line of defense for your team to investigate.

I have formed the impression that this student has some understanding of the processes of fingerprint-calculation and fingerprint-matching, as described in Mohanty (2019), with some ability to apply this knowledge in a forensic context. 7 marks.]

- Student answer #5.
  If I were the expert that is defending Mr. X, I would provide the following advice to Mr. X's lawyers as follows:
  - give an explanation of e-PRNU. e-PRNU is a forensic method to prove that a picture is taken by which camera. It is based on the PRNU noise which each camera is different. [This is OK as a very brief overview. However I think it is likely to mislead the defense lawyers, because e-PRNU is not well-established as a forensic method. If you had informed them that e-PRNU has been proposed, very recently, by some academic researchers, and that its matching performance has been characterised by only a single unreplicated experiment on ten cameras, you'd be accurately identifying a weakness in the forensic expert's report.]

  - the doubt about the incriminating picture froma social media website [?? What is your doubt about this picture? Do you doubt that it is incriminating? Why is this doubt relevant to the validity of forensic expert's report?]
  - the fact that the accuracy of e-PRNU is not 100% correct [?? Are you referring to TPR or FPR? If e-PRNU were 99.99% accurate in both TPR and FPR, would this imply reasonable doubt about the validity of the expert's report?]
  - the fact that the picture may be taken by Mr. X's camera, but it doesn't mean Mr. X takes the picture.[Why is this consideration relevant to the validity of the expert's report?]

  [This student has written a broader report than I had asked for, but they have good instincts. It's a very good idea – whether you're writing to experts or to non-experts – to briefly explain the terminology and concepts in an introductory section. It's also a good idea to briefly discuss a few important issues that are outside the scope of your brief, but which may not have been included in your scope because the person who commissioned your report may not be aware of their importance. So: I won't mark down the student for answering some questions that I hadn't asked. The time they spent on answering these questions will reduce the time they were able to spend on answering the question I had asked. And: answering a broader question is generally a good strategy to gain some marks on an exam question that you're unable to answer. In this case: the student has shown no ability to answer the question I had asked, but by answering a broader question they have shown that they have some understanding of Mohanty (2019). 5 marks.]

- Student answer #6. As a digital forensic investigator, my advice to Mr. X layer's [sic] is that the e-PRNU is different for different pictures taken by various cameras. [This is inaccurate advice. Mohanty (2019) explains that a camera's PRNU is estimated from its images, averaged to compute a "fingerprint", then "trimmed down" to a fingerprint digest which can be matched by the e-PRNU method. This process introduces errors, making it unsafe to assume – especially when advising Mr. X's defense lawyers! – that a fingerprint digest will uniquely identify a camera.]

  Obviously the nine images do not match with Mr. X camera because they are taken from different cameras. [This is very inaccurate advice, showing no understanding of the FPR data from Figure 4.]

  If the Police thinks Mr. X is a suspect they should a[c]quire other images taken from Mr. X camera and should validate e-PRNU fingerprint with the incriminating picture. if there is a more accurate match, only then they can confirm that Mr. X has taken the incriminating picture from his camera. [This shows no understanding of the process – already used by police – to compute a fingerprint digest for Mr. X's camera. if there is a more accurate match, only then they can confirm that Mr. X has taken the incriminating picture from his camera. [This shows very weak technical understanding of how a false-positive error could be resolved. Running additional tests on Mr. X's camera won't tell you anything about the likelihood that any other camera will have a fingerprint digest that'd be an excellent match to the incriminating image. Also, as a member of the defense team, you should not be trying to give advice to the Police on the additional work they might do

at some point in the future, if they have the time and other resources required to build a stronger case.]
[4 marks – shows only a weak understanding of Mohanty (2019), and no understanding of what could be learned from a ROC curve.]

(b) (5 marks) A junior member of your forensics team has produced an alternative visualisation of the data in Figure 4 of Mohanty (2019), shown below in Figure 1. After you confirm that this plot is accurate, can you use this as part of your defense of Mr X? Explain briefly.

Figure 1: PR curve: e-PRNU with digest vs conventional scheme with digest, redrawn from Fig. 4 of Mohanty (2019).

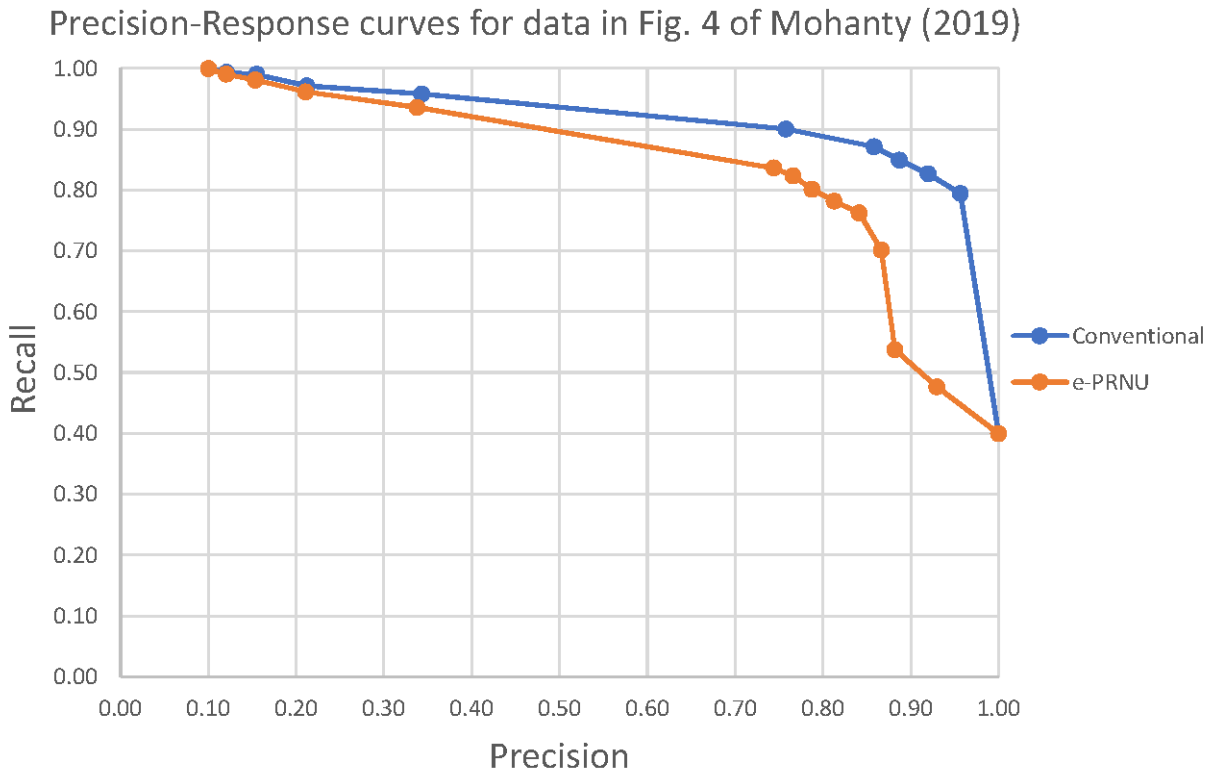

- Marking rubric: 3 marks for an accurate and relevant interpretation of Figure 1, 4 marks for showing relevant knowledge of PRNU and e-PRNU, 2 marks for constructing an argument which will help the defense team.
- [Model answer: Figure 1 tells us that the recall of the e-PRNU method is significantly different to the recall of the Conventional method, in a ten-camera match with precision of 0.7 or more. Precision is a measure of how many other cameras have fingerprints which match the incriminating image, at the (unknown) match-threshold chosen by prosecution's expert. This expert has vaguely characterised the match as "excellent", so they must have been running a high-precision test, i.e. 0.7 or higher.

Recall is defined as TP/(TP+FN) – so it's a measure of the conditional probability of a true-positive result, given that the match result is positive. I'd expect the prosecution's expert to select a match threshold for which recall is above 0.9, otherwise they'd have a chance of greater than 10% of not matching this image to the camera which actually did take the incriminating image.

From the reasoning above, I see several possibilities: 1) Figure 1 may be incorrect for the (vaguely described) matching methodology conducted by prosecution; 2) The prosecution's expert has run a test with recall 0.8 or higher, in which case precision is lower than 0.8 – giving us very weak evidence linking Mr. X's camera to the incriminating image; 3) The prosecution's expert may have taken a significant risk of a false-negative result (in which no cameras are found to match the incriminating image), by running a test with recall lower than 0.8. Case 3 is probably the worst from the defense's point of view, but even if we (later) discover this to be what the expert actually did, we can still ask the expert to estimate the precision of their matching methodology for the case of $n = 10$ cameras (perhaps 0.92), and for the case of $n = 100$ cameras. If prosecution have strong evidence that fewer than n cameras could have taken the incriminating image, then we can argue that they should supply a PR curve for n cameras, and to reveal their matching threshold on this curve, in order to evaluate the strength of their evidence.]

- Student answer #1.

Yes, this can be used as a defence because the recall rate of the e-PRNU is low which could be a significant factor in deciding the image taken is from Mr X camera or not. For a perfect match the recall rate should be higher. However in this case as we see the recall rate is low.

[0 mark. This answer is so vague that it gives me no reason to believe the student has any understanding of recall, precision, or how they apply to this case.]