


Your Coursework

(task for up to 3 student teams, coursework, 100%)

Re-produce a Published Research Paper



Forming your Teams:

Register your team of up to 3 people (i.e. one, two or three students) online at:

<https://doodle.com/poll/tbiab8e52md7fka9>

Post registration, teams can split but cannot merge, to avoid any copying of code or ideas.

Each member of the team should submit an exact copy of the final submission on Blackboard by the deadline. The report (see below) should note the full names and usernames of all members of the team.

It is up to each team to decide their best strategy to tackle this coursework, i.e. whether to divide the tasks below, or to work together on all tasks. Contributions of team members need not be explicitly stated.

However, by submitting a group coursework, you are implicitly acknowledging that all members of the team contributed equally. If this is not the case, you should email unit director with details of any issues encountered during the coursework.

Task Brief:

This assignment gives you the opportunity to appreciate the work required in replicating published research from a publicly available dataset and manuscript. It allows you to reflect on the experience of reproducing published results and potentially outperforming on your replication.

Gathering all the knowledge you acquired from the lectures and labs, read the paper below carefully and replicate the required results (Note: you are not required to re-produce all the paper's results). Feel free to take any pieces of code from the labs as a baseline, but the rest of the code should be originally yours.

The Paper:

J Pan, E Sayrol, X Giro-i-Nieto, K McGuinness and N O'Connor. Shallow and Deep Convolutional Networks for Saliency Prediction. In IEEE CVF Computer Vision and Pattern Recognition (CVPR), 2016.

https://openaccess.thecvf.com/content_cvpr_2016/papers/Pan_Shallow_and_Deep_CVPR_2016_paper.pdf

Note that our choice for paper is based on its simplicity and similarity to your labs, rather than its superior performance or exceptional novelty.

Please read the following information carefully before attempting the replication.

1) Architecture:

For this coursework, you will only be asked to implement and replicate the **shallow** architecture (Fig 2 in the paper). You should not replicate the deep architecture, even for the extension of your work.

2) Selected Dataset:

Out of the 3 datasets used in evaluating results in this paper (see Table 2), we only ask you to replicate results for the SALICON datasets (first line in the table). We have prepared this dataset for you already, as well as the evaluation metric calculations: [see **Dataset and Useful Code**]

3) Required Results:

In replicating the results, we expect you to provide code and your results for **only**:

1. SALICON dataset, on the shallow network and the published validation set, using three evaluation metrics (labelled in yellow below from Table 5).

iSUN (validation)	CC	AUC Shuffled	AUC Borji
Shallow Convnet	0.59	0.64	0.79
Deep Convnet	0.53	0.63	0.80
SALICON (val.)			
Shallow Convnet	0.58	0.67	0.83
Deep Convnet	0.61	0.73	0.86
SALICON (test)			
Shallow Convnet	0.60	0.67	0.83
Deep Convnet	0.62	0.72	0.86

Table 5. Comparison of our shallow and deep convnets.

2. Fig 3: You should be able to visualise all filters learnt by the *first layer* in your shallow convolutional network.
3. Fig 5: You should be able to choose your own examples, and visualise the image, ground-truth and predicted saliency maps.

4) Other details:

There are a few implementation details that were left unspecified by the authors. We specify these below to resolve ambiguities:

- The batch size should be set to 128
- For weight decay, we've used a value of 0.0005
- Although the authors use momentum decay, there is no easy way of implementing this in PyTorch, hence you should not try to implement any momentum decay.

- For data augmentation, we mirrored all training images and saliency maps horizontally as a pre-processing step, so you do not need to implement it in your training script.
- The authors mention: "For validation purposes, we split the training data into 80% for training and the rest for periodic validation". Instead of doing that, you should use the whole training set for training and the whole validation set for evaluating your model.
- You should not use ReLU before the maxout layer. We noted this improved results.
- For your convolutional layers, you should use padding=2 for the first convolutional layer, and padding=1 for the second and third convolutional layers.

5) Our replicated results:

Replicating papers rarely produces exact results as those reported in the published papers. It is highly advisable to publish one's code with the paper, however this is very infrequently adopted by researchers.

We have replicated the paper's results in PyTorch for you. We provide the corresponding table that we could re-produce, using the data files available to you. These are the results you are attempting to reproduce.

SALICON (val)	CC	AUC Shuffled	AUC Borji
Pan et al [2016]	0.58	0.67	0.83
Our replication	0.65	0.55	0.71

As our reimplementations produce worse results, this is the replication we expect you to reproduce.

Dataset and Helpful Code:

You can find resources we've prepared for you for this project at:

https://uob-my.sharepoint.com/:f/g/personal/qc19291_bristol_ac_uk/EiSx7UQ_S45As-13F0cmzcBy2HxhKudJodlpRmQHbUDQQ?e=xgSTKS

Final Submission:

1. An original code, **based on PyTorch** (other software engines won't be accepted – we won't accept Keras or Tensorflow), replicating the published paper. You can use your lab code from any or all group members. **We aim to run your code on BC4, so ensure it compiles and runs.**
2. A report in the IEEE conference format (https://www.ieee.org/conferences_events/conferences/publishing/templates.html) of up to 5 pages including references, submitted in PDF format. The report should include the following sections:
 - A. **Title and Team members** (names and usernames)
 - B. **Introduction:** Definition of the problem addressed by the paper Pan et al (in your own words)
 - C. **Related Work:** A summary of **more recent** published papers (i.e. after Pan et al was published in 2016) attempting to address the same problem (up to 3 works).
 - D. **Dataset:** A description of the dataset used, training/test split size, labels and file formats.

- E. **Input:** Explain what Saliency maps are. Give 1-2 examples visually from your data, by plotting these as images, and showing the ground-truth for these. Do not use the figure from the paper. We have provided code to help you visualise both the ground truth and the predicted saliency maps.
- F. **Shallow Architecture (Pan et al):** Describe through a table the architecture and all its details.
- G. **Implementation Details:** Summary of the steps you have undertaken to replicate the results, train the data and obtain the results, including any decisions you needed to make along the way. Do not include any pieces of code, but you can include pseudo-codes if needed.
- H. **Replicating Quantitative Results:** You need to present your results for table 5 as above.
 - I. **Training curves:** Include your training/test loss curves for your models, and comment on any overfitting in your training. The curves here should correspond to the same run as those in the reported table (Section H). These curves could be directly retrieved from Tensorboard. Also in this section, include your visualisation of the filters learnt in the first convolutional layer (your replication of Fig 3 in the paper).
 - J. **Qualitative Results:** This section should include sample success and failure cases based on your algorithm, similar to Fig 5 in the paper. In presenting these examples, you can plot/display the input image, ground-truth and predicted saliency maps. Particularly: (a) find 1 good example, where your prediction works. (b) find 2 problematic examples where your prediction can be criticised. These should represent two different modes of failure. Hint: Note that from our replication, the model has a strong centre bias.
 - K. **[65+] Improvements:** Using the same shallow architecture, propose, implement and test one improvement you made to your results). **Note:** if you describe multiple improvements, we will give you the lower mark (rather than the higher one), so choose the one you believe in. Cover any implementation details required to understand and replicate your modifications. Report your improved results in tabular format for all metrics. Do not include any pieces of code, but you can include pseudo-codes if needed. **Note:** Your improvement should be made using the same dataset, train/test split and evaluation metrics used earlier. Improvements can include changes to architecture, hyper-parameters, data augmentation or learning algorithm. Your choice should be justified theoretically and experimentally.
 - L. **Conclusion and Future Work:** Summarise what your report contains in terms of content and achievements. Suggest future work that might extend, generalise or improve the results in your report.

Marking Guideline.

Note: Code and report will be checked for plagiarism. Proven plagiarism will result in a 0 grade on this coursework for the whole team.

50-54

To pass this assignment, you must produce original complete (compiles and runs on BC4 using batch-mode command and PyTorch) code that replicates the results in the paper. You should produce a report with sections A-F correct and satisfactory. A partially-complete and correct attempt to address sections G, H, I and L is included (i.e. excluding J and K). Any errors or misses do not significantly affect a “replication of results” effort. Replication results (Section H) are within 0.2 difference on all three metrics.

55-64

In addition to the above, sections F, G, H, I and J would be complete, correct and reflective of your understanding of the code and the implementation. All sections (except K) are completed to an acceptable standard. Reported results are within 0.1 difference on at least one metric.

65-70

In addition to the above, a satisfactory attempt to provide improvements (K) on the published results have been achieved, correctly described, with improvements to the results. Marginal improvements will be accepted.

70-75

In addition to the above, the presentation given was to a very good standard with almost no areas of weakness. The proposed improvement is far from random and has been carefully thought of in light of the problem and misclassification errors. Section J should include interesting (rather than random) success and failure cases, with explanations of failure cases. The report’s organisation and structure should be very good.

75-80

In addition to the above, the report should be submit-able to a B-class peer review conference, i.e. it shows excellent understanding, correct and complete showcasing of the approach. Statements are concise, and any jargon out of implementation details is avoided. The chosen related work reflects state of the art on this problem. Extensive evidence of analysis, creativity & originality in concise content presentation should be shown. Code is commented, and could be easily understood and re-used by a reader.

80-100

In addition to the above, the produced code and report are exemplary, and could be given as an example for an attempt to replicate this published work. Improvements in results are beyond marginal.

General Guidelines (Department Regulations):

Deadline

The deadline for submission of all optional unit assignments is 13:00 on Friday 11th of December. Students should submit all required materials to the “Assessment, submission and feedback” section of Blackboard - it is essential that this is done on the Blackboard page related to the “With Coursework” variant of the unit.

Time commitment

The expectation is that students will spend 3 full working weeks on their two assignments. The effort spent on the assignment for each unit should be approximately equal, being roughly equivalent to 1.5 working weeks each.

Academic Offences

Academic offences (including submission of work that is not your own, falsification of data/evidence or the use of materials without appropriate referencing) are all taken very seriously by the University. Suspected offences will be dealt with in accordance with the University’s policies and procedures. If an academic offence is suspected in your work, you will be asked to attend an interview with senior members of the school, where you will be given the opportunity to defend your work. The plagiarism panel are able to apply a range of penalties, depending the severity of the offence. These include: requirement to resubmit work, capping of grades and the award of no mark for an element of assessment.

Extenuating circumstances

If the completion of your assignment has been significantly disrupted by serious health conditions, personal problems, periods of quarantine, or other similar issues, you may be able to apply for consideration of extenuating circumstances (in accordance with the normal university policy and processes). Students should apply for consideration of extenuating circumstances as soon as possible when the problem occurs, using the following online form:

<https://apps.powerapps.com/play/3172b943-0956-4b88-bf3d-3f37871d1170?tenantId=b2e47f30-cd7d-4a4e-a5da-b18cf1a4151b>

You should note however that extensions are not possible for optional unit assignments. If your application for extenuating circumstances is successful, it is most likely that you will be required to retake the assessment of the unit at the next available opportunity.

Implications of UK “travel window”

The UK Government will be instigating a “travel window” to allow students return home from University once UK national restrictions have been lifted. This will take place between the 3rd and 9th of December and as such, will occur during the optional unit assignment period. Students whose work is significantly impacted by extended periods of travel and quarantine during this period should apply for extenuating circumstances. As discussed previously, extensions are not possible for optional unit assignments. If your application for extenuating circumstances is successful, it is most likely that you will be required to retake the assessment of the unit at the next available opportunity.

