

Computer Vision Applications

COMP 388-002/488-002 Computer Science Topics

Daniel Moreira
Fall 2022



LOYOLA
UNIVERSITY CHICAGO

Final Exam

COMP 388-002/488-002 Computer Science Topics

Daniel Moreira
Fall 2022



LOYOLA
UNIVERSITY CHICAGO

Question 1



What is computer vision? How does it relate to computer science, artificial intelligence, machine learning, pattern recognition, deep learning, media forensics, and biology?

Suggested comments. Why do we study computer vision? What is not computer vision? Is computer vision already a reality? Have you used or benefited from computer vision before?

Question 2

What is the semantic gap in computer vision applications? Why does it exist? How has it been addressed by the computer vision community?

Suggested comments. What are the differences between classical and recent computer vision techniques with respect to addressing the semantic gap?

Question 3

What were the recent advances in computer science and engineering that led to the resurgence of neural networks and their successful application to computer vision? Why was this resurgence coined deep learning?

Suggested comments. What are the pros and cons of deep learning? What should we as a society worry the most about it?

Question 4

What is the difference between training and inference time in deep learning solutions? What is happening to the architecture and to the weights of a network during training time? What are back-propagation and the loss function? When are they used?

Suggested comments. Which one is the most computationally intensive, training or inference time? What is the importance of annotated data for both training and inference times? What is the impact of wrongly annotated or cherry picked data?

Question 5

What are local feature image detectors and descriptors? What are they useful for? Are they still used nowadays? How are they related to deep learning?

Suggested comments. Give examples of detectors and descriptors. Give examples of applications that benefit from local features.

Question 6

What is image retrieval? How does it differ from image classification? How is image retrieval useful to media forensics and provenance analysis?

Suggested comments. What are the limitations of typical image retrieval in the face of provenance analysis? Give examples of methods that help to overcome these limitations.

Question 7

What are AlexNet, GoogLeNet, VGGNet, and ResNet? What computer vision tasks are they useful for right off the shelf (i.e., without architectural modifications)?

Suggested comments. What has made the development of these solutions possible? How does ImageNet fit in this realm?

Question 8

What are the characteristics that make the ResNet architecture unique?

Suggested comments. Why do these characteristics work? What problem do they address?

Question 9

Is it possible to use VGGNet for image retrieval? If no, please explain why. If yes, please explain how.

Suggested comments. Do you agree with the given answer? Please justify your answer.

Question 10

What are the differences between image classification, object detection, and image segmentation? How are these tasks ranked in terms of difficulty? Please justify your answer.


Suggested comments. Do you agree with the given answer? Please justify your answer. Give examples of architectures designed for each task.

Question 11

What is U-Net useful for? What are the characteristics that make the U-Net architecture unique?

Suggested comments. Why do these characteristics work? What problem do they address?

Question 12



The face detector in this example is not working for the man. Why do you think this is happening? How would you fix this problem?

Suggested comments. What should we care about when collecting and annotating training data?

Question 13

What is transfer learning in the realm of deep learning? How does one perform it? Does transfer learning impact the architecture or the weights of a neural network? Please give an example of transfer learning.


Suggested comments. What are the pros and cons of transfer learning?

Question 14

Are neural networks discriminative or generative machine learning techniques? Please justify your answer. If they are discriminative, is it possible to make them generative (or vice-versa)? If no, please explain why. If yes, please explain how.

Suggested comments. Do you agree with the given answer? Please justify your answer. Give examples of architectures that support your answer.

Question 15



The screenshot shows the top of a news article. The header is dark green with the Harrisburg University of Science and Technology logo. The article title is 'Facial recognition software paper not being published'. The breadcrumb trail is 'HOME \ UNCATEGORIZED'. The main text begins with the title and a paragraph explaining that a paper titled 'A Deep Neural Network Model to Predict Criminality Using Image Processing' was removed from the website at the request of the faculty involved in the research. The editor of the publication where it was scheduled to appear has decided to not publish the work.

If the proposed network got a high classification accuracy, why do you think this work was not published?

Suggested comments. What might the network have learned? Should we pursue such a solution? Please justify your answer.

Question 16

What is SVM? How does it compare with deep learning?

Suggested comments. Give examples of either replacement or combined use of SVM and deep learning. Is SVM still used nowadays? If no, please explain why. If yes, please cite examples.