

ICVSS 2015 – Examination

1. Invariance and distinctiveness may be an insufficient characterisation of good representation. What fundamental property of representations is not accounted for by these two concepts?

- A) speed
- B) compactness
- C) ability to tackle multiple tasks**

2. Inverting most representations is an ambiguous operation. This is:

- A) interesting, as we can infer representation properties from the ambiguity.**
- B) problematic, because the representation is discarding information contained in the image.
- C) irrelevant, as most representations are meant for discrimination, not for image generation.

3. Kernels are:

- A) Alternative to hand-crafted features
- B) An alternative approach to the construction of representations**
- C) An alternative approach to representations, as they allow using a similarity function instead of a feature vector in classification

4. What are the relation between wavelets and deep networks ?

- A) Both are computed with cascades of filtering and subsamplings**
- B) Both are linear operators
- C) Both are non-linear operators

5. Why can one remove some rectifier non-linearities within a deep network ?

- A) Because they are usually useless
- B) Because they do not modify the output of a low-pass filter which outputs positive coefficients**
- C) Because rectifiers are similar to absolute values

6. Why are multiscale transforms useful to analyze complex image scenes?

- A) Because they can compress data
- B) Because they provide stable representations relatively to deformations**
- C) Because they are sensible to frequencies

7. The chain rule is used for obtaining the back-propagation algorithm. Which of these is false?

- A) The chain rule for $y=f(g(x_1, x_2))$ decomposes as a sum over the terms respectively associated with x_1 and x_2 (consider the case where both f and g return a scalar, and x_1 and x_2 are scalar).**

B) The back-propagation algorithm computing time grows as $O(M)$ if M is the number of arcs in the computation flow graph (number of connections between neurons), assuming each arc requires $O(1)$ computation.

C) The chain rule for computing dy/dx from $y=f(g(x))$ can be generalized to the case where $g(\cdot)$ outputs a vector by replacing the dot product between $f'(g(x))$ and $g'(x)$ by the matrix product of the corresponding Jacobians.

8. Consider a convolutional layer with a $K \times K$ input kernel, an $N_i \times N_i$ dimensional input with F_i input maps and F_j output maps, and no 0-padding. Which of these is false?

A) The number of free parameters associated with the convolution is $F_i \times F_j \times K \times K$.

B) The size of the output image is $M \times M$, with $M = N - K$

C) The number of multiply-adds for backprop through the convolution is $N \times N \times K \times K \times F_i \times F_j$

Hint: try it out in 1-D with a small value of K .

9. Which of these statements is probably false?

A) Gradient-based optimization of neural networks tends to get stuck in poor local minima.

B) A distributed representation like in an MLP hidden layer can capture exponentially more distinct regions in input space than a non-distributed one like in clustering or nearest-neighbour algorithms.

C) Deep learning can yield better generalization when the priors or preferences associated with depth and distributed representations fit well with the actual data generating distribution.

10. Consider a deep neural network A with L layers and N parameters and a less deep neural network B with $L' < L$ layers and M parameters. Which of these is false?

A) if $N=M$, then the set of functions reachable by A is a superset of the set of functions reachable by B.

B) even if M is 10 times larger than N , there are functions reachable by A that are not reachable by B

C) with $N=M$, there are functions reachable by B that are not reachable by A

11. Name two pioneers of variational methods:

A) Galilei and Newton

B) Schrödinger and Heisenberg

C) Lagrange and Euler

12. What dimension does the group of 3D similarity transformations have:

A) 6

B) 7

C) 8

13. The total variation is

A) convex but non-smooth

- B) convex and smooth
- C) non-convex

14. What is the formula for computing probabilities to sample from in stochastic pooling?

- A) $p_i = \exp(a_i) / \sum_j(\exp(a_j))$
- B) $p_i = a_i / \sum_j(a_j)$**
- C) $p_i = a_i / \sqrt{\sum_j(a_j * a_j)}$

15. What components is the cost function of a deconvolutional network composed of?

- A) L1 sparsity and reconstruction of feature maps by convolving feature maps with learned filters.**
- B) L0 sparsity and reconstruction of feature maps by convolving image with learned filters.
- C) L1 sparsity and reconstruction of image by convolving feature maps with learned filters.

16. What improvements to the first layer convolutions were made to Alexnet (Alex Krizhevsky's CNN model from Imagenet 2012) due to deconvolution network visualizations?

- A) smaller stride and larger convolution kernels.
- B) larger stride and larger convolution kernels.
- C) smaller stride and smaller convolution kernels.**

17. Subspace clustering refers to the problem of (choose the best answer):

- A) Clustering data according to their membership to multiple groups
- B) Clustering data according to their membership to multiple subspaces**
- C) Clustering data according to their membership to multiple manifolds

18. Sparse subspace clustering is based on the principle of (choose the best answer):

- A) Expressing each data point as a linear combination of all data points
- B) Expressing each data point as a sparse linear combination of all data points
- C) Expressing each data point as a sparse linear combination of all other data points**

19. In sparse subspace clustering, the number of nonzero coefficients in the representation of a data point corresponds to:

- A) The dimension of the subspace containing that point**
- B) The number of subspaces
- C) The diameter of the subspace containing that point

20. What is one major difference between cells in areas V1,V2,V3,and V4 and cells in inferotemporal (IT) cortex?

- A) Cells in the IT cortex have a retinotopic map, whereas V1-V4 have very large receptive fields.
- B) Cells in V1-V4 are mostly sensitive to oriented edges, while cells in IT together form a retinotopic map.
- C) Cells in V1-V4 have localized receptive fields, together forming a retinotopic map of space, while cells in IT cortex have very large receptive fields.**

21. What is one major transformation occurring along the face patch system?

- A) Building a representation of faces that is invariant to view, illumination, and identity.
- B) Achieving illumination invariance by oriented edge detection.
- C) Construction of a view-invariant representation of individual identity.**

22. According to one computer vision algorithm, pooling light-invariant local contrast features provide a powerful way to detect faces. Describe how neural data from the macaque brain supports this idea.

- A) Cells showing similar responses as Haar wavelets indicate that the brain architecture is similar to a Viola-Jones detector.
- B) Cell recordings show that HOG type features are extracted, leading to a lighting-invariant encoding of face geometry.
- C) Cells are completely consistent in their contrast preferences to each other, and to predictions from computational experiments revealing which features are lighting-invariant.**

23. The main advantage of features (versus direct methods) for structure from motion is that they enable to match:

- A) more efficiently
- B) more accurately
- C) over a wider baseline**

24. A Wulff shape allows to implement

- A) an arbitrary anisotropic surface penalizations
- B) a convex shape prior**
- C) convex surface penalizations, yielding a nonconvex shape prior

25. A high-order ray potential for a pixel-ray containing N voxels and considering a total of C semantic classes (including free space) has how many different groups of states that need to be handled separately?

- A) C^N
- B) $N*(C-1)+1$**
- C) $N*C$

26. How does one have to customize a DLP projector to be used in a system to avoid illuminating raindrops?

- A) Increase power
- B) Decrease latency**
- C) Increase brightness

27. How did the structured light system work that tolerates bright illumination?

- A) Synchronizing light projection and detection**
- B) Increasing light output and using a high-framerate camera
- C) Filtering outdoor light by active polarization filters