

Machine Learning 2023

Final Exam

27 March 2023, 12:15–14:30

The questions start on the next page. **Please do not open the exam booklet until the scheduled starting time.**

Feel free to write on the exam booklet, and take it with you afterwards. The exam will be made available on Canvas later today, together with the correct answers. You can copy your final answers onto the booklet, so you can check how you did.

There are 40 questions, worth 1 point each.

To pass the examination, your total points for the exam plus those for the quizzes should be around 50–55. The exact pass mark will be decided *after* the exam results have been analysed.

Rules:

- You are allowed to use a calculator or graphical calculator.
- You are *not* allowed to use your phone or smartphone.
- The exam is closed-book.
- You are allowed to use the formula sheet provided through Canvas.
- The formula sheet should not have any writing on it, except in the cheat sheet box. Anything inside the box should be written by hand.

Recall questions

1. How are *hyperparameters* defined?
 - A They are parameters of a learning algorithm that are set before we start training on the data.
 - B They are parameters of which the values are found by the training algorithm itself.
 - C They are parameters that make the models train faster.
 - D They are parameters of which the values are set after training the model.
2. In principal component analysis, which is **true**?
 - A The first principal component is the direction in which the variance is the smallest.
 - B The first principal component is the direction that maximizes the reconstruction loss.
 - C The first principal component is the direction that minimizes the reconstruction loss.
 - D The first principal component is the direction in which the bias is the smallest.
3. What is a valid reason to prefer *gradient descent* over *random search*?
 - A My model is easily differentiable.
 - B I need to be sure that I find the global minimum.
 - C My loss function is not smooth.
 - D There is some computation between the output of my model and my loss function, which I do not control.
4. If we have a feature that is categorical, but our model requires numeric features, we can turn the categoric feature into one or more numeric features either by **integer coding**, or by **one-hot coding**. Which is **true**?
 - A The benefit of integer coding is that it assigns each value a separate feature.
 - B Integer coding results in more features than one-hot coding.
 - C The feature(s) introduced by integer coding are are valued between 0 and 1.
 - D The feature(s) introduced by one-hot coding are valued 0 or 1.
5. Two classes that are not linearly separable may become linearly separable if we add a new feature that is derived from one or more of the other features. Why?

- A This approach removes outliers.
 - B A linear decision boundary in the new space may represent a nonlinear boundary in the old one.
 - C Adding a new feature derived from the existing ones normalizes the data.
 - D This approach makes a non-differentiable loss function differentiable.
6. In the Expectation-Maximization algorithm for Gaussian mixture models, we train a model consisting of several *components*, and we compute various *responsibilities*. Which is **true**?
- A Each component is a normal distribution. Its responsibility for a point is the normalized probability density it assigns the point.
 - B Each component is a normal distribution. Its responsibility is the integer number of points it generates.
 - C Each component is one of the features of the data. Its responsibility for a point is the normalized probability density it assigns the point.
 - D Each component is one of the features of the data. Its responsibility is the integer number of points it generates.
7. Lise is given an imbalanced dataset for a binary classification task. Which one of the following options is useful solution to this problem?
- A She can oversample her majority class by sampling with replacement.
 - B She can use accuracy to measure the performance of her classifier.
 - C She can use SMOTE to augment her training data.
 - D She can augment the feature set by adding the multiplication of some features.
8. What is **true** about *logistic regression*?
- A Logistic regression is a linear, discriminative classifier with a cross-entropy loss function.
 - B Logistic regression finds the maximum margin hyperplane when applied to data with well-separated classes.
 - C Logistic regression is identical to the least-square linear regression.
 - D Thanks to the log-loss function, logistic regression can learn nonlinear decision boundaries.
9. Neural networks usually contain *activation functions*. What is their purpose?
- A They are used to compute a local approximation of the gradient.
 - B They are applied after a linear transformation, so that the network can learn nonlinear functions.
 - C They control the magnitude of the the step taken during an iteration of gradient descent.
 - D They function as a regularizer, to combat overfitting.

10. I am training a generator network to generate faces. I take a random sample, compare it to a randomly chosen image from the data, and backpropagate the error.

When training is finished, all samples from the network look like the average over all faces in the dataset.

What name do we have for this phenomenon?

- A Multiple testing
 - B Overfitting
 - C Dropout
 - D Mode collapse
11. Which is **true**?
- A A maximum likelihood objective for least-squares regression does not provide a smooth loss surface.
 - B The least-squares loss for linear regression can be derived from a maximum likelihood objective.
 - C Linear regression can be performed with a maximum likelihood objective but the results will be different from the least-squares loss.
 - D The cross-entropy loss is derived from assuming a normal distribution on the residuals.
12. What is an important difference between regular recurrent neural networks (RNNs) and LSTM neural networks?
- A LSTMs have a *forget gate*, allowing them to ignore parts of the cell state right away.
 - B RNNs have a *forget gate*, allowing them to ignore parts of the cell state right away.
 - C LSTMs have a vanishing gradient problem, RNNs don't.
 - D RNNs can be turned into variational autoencoders, LSTMs can't.
13. Which of the following is an ensemble method?
- A Random forest
 - B Gradient boosting
 - C AdaBoost
 - D All of the above
14. William has a dataset of many recipes and many ingredients. He doesn't know anything about the recipes except which ingredients occur in each, and he doesn't know anything about the ingredients except in which recipes they occur.
- He'd like to predict new ingredient/recipe pairs for ingredients that could be added to existing recipes.

Which is **true**?

- A He could model the recipes as instances with their ingredients as a single categorical feature, and solve the problem with a decision tree.
- B He could model the ingredients as instances and their recipes as a single categorical feature, and solve the problem with a decision tree.
- C He could model this as a matrix decomposition problem.
- D This problem requires a sequence-to-sequence model.

Combination questions

15. The mean squared error (MSE) loss function $\sum_i (y_i - t_i)^2$ and the mean absolute error (MAE) loss function $\frac{1}{n} \sum_i |y_i - t_i|$ (where the sum is over n instances, y_i is the model output for instance i and t_i is the training label) are two popular loss functions. Which one of the following options could be a reason for favoring the MAE over the MSE?
- A In MAE, the negative and positive differences do not cancel out in the sum.
 - B The mean of the error minimizes the MAE.
 - C The MAE is more robust to outliers than the MSE.
 - D There is no advantage in using the MAE over the MSE.
16. We are choosing a new basis for our data. We decide to use an *orthonormal* basis. What is the advantage of having an orthonormal basis?
- A It ensures that the basis vectors are equal to the principal components.
 - B It ensures that the basis vectors are orthogonal to the principal components.
 - C It ensures that the inverse of the basis matrix is equal to its transpose.
 - D It ensures that the data is automatically whitened in the new basis.
17. Which of these statements about principal component analysis is **false**?
- A The first principal component provides the direction of greatest variance.
 - B It is a supervised method.
 - C It can be used for dimensionality reduction.
 - D It can be used for data preprocessing.
18. In a support vector machine, which is **not** an advantage of using *soft margins*?
- A It allows some points to fall inside the margin.
 - B It allows some points to be misclassified.
 - C It allows applying gradient descent for optimization, where a hard margin doesn't.
 - D It finds better decision boundaries even for linearly-separable classes

19. What is **not true** about the naive Bayes classifier?
- A The naive Bayes classifier can be applied to multi-class classification.
 - B The naive Bayes classifier assumes that the features are independent.
 - C To avoid zero probabilities in the naive Bayes classifier, one can use the Laplace smoothing technique.
 - D The naive Bayes classifier is a probabilistic generative classifier.
20. One way to think of a convolution layer is as a fully connected layer, with some extra constraints. Which is **not** one of these constraints?
- A Some of the weights are forced to have the same value.
 - B Some of the connections are removed.
 - C The L2 norm of the weights is limited to a maximum value.
 - D All of the above are part of a convolution layer.
21. Which answer contains methods that can *all* be used as sequence-to-sequence layers?
- A Convolutions, RNN, LSTM
 - B Gradient boosting, LSTM, Deep Q-Learning
 - C Convolutions, Word2vec, Gradient boosting
 - D RNN, Deep Q-Learning, Word2vec
22. How are the Gaussian Mixture Model (GMM) and the Mixture Density Network (MDN) related?
- A The MDN is an alternative to the GMM that can also describe complex distributions, but that doesn't use Gaussians.
 - B The GMM is an alternative to the MDN that can also describe complex distributions, but that isn't a mixture model.
 - C The MDN is a neural network which uses the GMM as its output distribution.
 - D The GMM is a neural network which uses the MDN as its output distribution.
23. I have a dataset of politicians in the European parliament and which laws they voted for and against. The record is incomplete, but I have some votes for every law and for every politician.
- I would like to predict, for a new law, which politicians will vote for and which will vote against. I plan to model this as a recommender system using *matrix factorization*.
- Which is true?

- A This is not a good model, because there are too many classes and not enough instances.
- B This is not a good model, because there are not enough classes, and too many instances.
- C I would have to deal with the cold start problem, because for the new law I don't have any voting information.
- D I would have to deal with the cold start problem, because the voting record is incomplete.
24. Which is a good principle to decide what feature to use for the next split in a decision tree?
- A After the split we want the probability of the majority class to be as low as possible.
- B After the split we want the probability of the majority class to be as high as possible.
- C After the split we want the entropy of the class distribution to be as low as possible.
- D After the split we want the entropy of the class distribution to be as high as possible.
25. What is the difference between *inductive* and *transductive* learning?
- A In transductive learning the model is allowed to see the labels of the test data during training.
- B In inductive learning the model is allowed to see the labels of the test data during training.
- C In transductive learning, the model is allowed to see the features of the test data during training.
- D In inductive learning, the model is allowed to see the features of the test data during training.

Application questions

We have the following training set:

	x_1	x_2	label
a	1	3	Pos
b	2	2	Pos
c	3	1	Neg
d	6	4	Neg
e	5	5	Pos
f	4	6	Pos
g	7	7	Neg

For the following questions, it helps to draw the data and the classification boundary in feature space.

We use a linear classifier defined by

$$c(x_1, x_2) = \begin{cases} \text{Pos} & \text{if } -x_1 + 0 \cdot x_2 > -2 \\ \text{Neg} & \text{otherwise.} \end{cases}$$

26. If we turn c into a *ranking* classifier, how does it rank the points, from most **Negative** to most **Positive**?

- A c b a d e f g
- B g f e d a b c
- C g d e f c b a
- D a b c f e d g

27. How many ranking errors does the classifier make?

- A 3
- B 4
- C 5
- D 6

28. If we draw a coverage matrix (as done in the slides), what proportion of the cells will be red?

- A $\frac{5}{12}$ B $\frac{1}{2}$ C $\frac{5}{16}$ D $\frac{3}{8}$

We will use the backpropagation algorithm to find the derivative of the function

$$f(x) = \frac{\sin(x^2)}{x^3}$$

with respect to x .

First, we break the function up into modules:

$$a = x^2$$

$$b = \dots$$

$$c = \dots$$

$$f = \dots$$

29. What should be in the place of the dots?

- A $b = \sin x$, $c = x^2$
- B $b = \sin x$, $c = x^3$
- C $b = \sin a$, $c = x^2$
- D $b = \sin a$, $c = x^3$

30. In terms of the *local* derivatives. Which is the correct expression for the derivative $\frac{\partial f}{\partial x}$?

- A $\frac{2x \cos a}{c} - \frac{3bx}{c^2}$
- B $\frac{2x \cos a}{c} + \frac{3bx}{c^2}$
- C $\frac{2x \cos a}{c} - \frac{3bx}{c}$
- D $\frac{2x \cos a}{c} + \frac{3bx}{c}$

We have trained a linear support vector machine, and found the following weights and bias:

$$\mathbf{w} = \begin{pmatrix} 1.5 \\ -2 \end{pmatrix}, \quad b = -1. \quad (1)$$

31. Which of the following points are *support vectors* for the trained model?

- | | |
|-------------------------------|----------------------------|
| A $y_1 = 1, x_1 = (2, 1.5)$ | $y_2 = 1, x_2 = (0, 1)$ |
| B $y_1 = 1, x_1 = (1, -0.25)$ | $y_2 = -1, x_2 = (0, 1)$ |
| C $y_1 = -1, x_1 = (2, 1.5)$ | $y_2 = 1, x_2 = (0, -1)$ |
| D $y_1 = -1, x_1 = (2, 0.5)$ | $y_2 = 1, x_2 = (1, 0.75)$ |

32. Consider two points: $x_1 = (1, 0)$ and $x_2 = (-0.5, -1)$. How are these two points classified?

- A Both **positive**.
- B x_1 is **positive**, x_2 is **negative**.
- C x_1 is **negative**, x_2 is **positive**.
- D Both **negative**.

We are given a dataset of email, labeled **spam** and **ham**. The total number of words in the **spam** and **ham** datasets is 50 000 and 200 000, respectively.

On this dataset, we want to train a *first-order Markov model* as a basis for a Bayes classifier to detect spam.

The following table shows the frequency of several bigrams and uni-grams in the dataset.

uni- and bigrams	frequency	
	spam	ham
you	5 000	5 000
won	3 000	1 000
a	9 000	9 000
prize	2 000	1 000
you won	300	100
won a	300	100
a prize	300	100

We have two *priors* on how likely an email is to be spam. **Prior 1** follows the dataset statistics, which says that an email is as likely to be **spam** as **ham**. **Prior 2**, based on a user survey, says that the probability that a given email is **spam** is 1%.

Consider the following email:

email: you won a prize

33. How is the email classified?
- A With both priors as **ham**.
 - B With both priors as **spam**.
 - C Under prior 1 as **spam**, under prior 2 as **ham**.
 - D Under prior 1 as **ham**, under prior 2 as **spam**.
34. Under prior 1, what is the probability that the email is **spam**?
- A $\frac{1}{28}$ B $\frac{27}{28}$ C $\frac{1}{45}$ D $\frac{27}{45}$
35. Under prior 2, what is the *probability ratio* of **ham** over **spam** for the second email?
- A $\frac{11}{3}$ B $\frac{3}{11}$ C $\frac{27}{45}$ D $\frac{45}{27}$

Consider the following function:

$$\mathbf{y} = \mathbf{x} \oslash \mathbf{v} + \mathbf{b}$$

Here \mathbf{x} , \mathbf{y} , \mathbf{v} and \mathbf{b} are vectors and \oslash represents element-wise division.

This function is part of a larger computation graph, resulting in a scalar loss l . We want to implement this computation in an automatic differentiation (AD) system, as discussed in the lectures.

36. Work out the scalar derivative of y_i over x_j . Which is the correct solution?

- A** $\frac{\partial y_i}{\partial x_j} = \begin{cases} 1/v_j & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$
- B** $\frac{\partial y_i}{\partial x_j} = \begin{cases} 1/v_j & \text{if } i \neq j \\ 0 & \text{otherwise} \end{cases}$
- C** $\frac{\partial y_i}{\partial x_j} = \begin{cases} 1 + b_j & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$
- D** $\frac{\partial y_i}{\partial x_j} = \begin{cases} 1 + b_j & \text{if } i \neq j \\ 0 & \text{otherwise} \end{cases}$

The AD system will give us a vector y^∇ , such that $y_i^\nabla = \frac{\partial l}{\partial y_i}$. Using y^∇ , we want to efficiently compute a vector x^∇ such that $x_i^\nabla = \frac{\partial l}{\partial x_i}$.

37. Which operation computes this vector for us?

- A** $x^\nabla = y^\nabla \oslash v + b$
- B** $x^\nabla = y^\nabla \oslash v$
- C** $x^\nabla = y^{\nabla T} v + b$
- D** $x^\nabla = y^{\nabla T} v$

Consider the following task. The aim is to predict the class y from the binary features x_1, x_2, x_3 and x_4 .

x_1	x_2	x_3	x_4	y
A	A	B	B	Yes
A	B	A	A	No
A	A	A	A	Yes
A	A	A	B	Yes
B	B	B	B	No
B	B	A	B	No
A	B	A	A	Yes
A	A	B	A	Yes
B	A	A	A	No
B	A	B	B	No
B	B	B	B	Yes
B	B	B	A	No


38. In standard decision tree learning (as explained in the lectures), without pruning. Which would be the first feature chosen for a split?

- A** x_1 **B** x_2 **C** x_3 **D** x_4

39. If we *remove* that feature from the data, which would be chosen instead?

- A** x_1 **B** x_2 **C** x_3 **D** x_4

Consider the following (partial) decision tree:



Note that this is just an arbitrary tree, not necessarily corresponding to the answers in the previous questions.

We can place either x_1 or x_2 on the open node (indicated by the question mark).

40. At this point, what is the information gain for x_2 ?

- A -1
- B 0
- C 0.5
- D 1

Thank you for your effort. Please check that you've put **your name and student number** on the answer sheet.