

0

What programming language do we use in this course?

What is the most negative Python integer?

What gives the largest result?

```
"""list(range(-1,8,2))[-2]"""
```

```
"""[[i,i+1] for i in [1,2]]"""
```

```
"""[x for x in b if x % 2 == 0]""" constructs an array with ...
```

What produces [1,2,3,...10]?

```
"""d = {'a': 'b', 'b': 'd', 'c': 'a', 'd': 'd'}
```

```
d[d['a']] = d[d['c']]
```

```
print(d['a'],d['b'])"""
```

```
"""d = {[1,2]: True, [2,3]: False}"""
```

```
"""*a,b = [1,2,3]"""
```

```
"""def f(x): x[1] = 0
```

```
y = [1,2]
```

```
f(y)
```

```
print(y)"""
```

Why did the creator of Python call the language Python?

Who is the creator of Python?

What is `not` in the Zen of Python?

What is the output of this function?

1

From what name is the term 'Algorithm' derived?

What notation is commonly used to denote the asymptotic complexity of a function?

Which running time is the fastest?

What is the relationship between space and time complexity?

Pruning and look-ahead are the main ingredients of what type of algorithm?

What does this algorithm compute? / What is the output?

```
"""  
def f(items): # items is a list  
    if len(items) == 1:  
        return items[0]  
    first, *rest = items  
    temp = f(rest)  
    if first < temp:  
        return temp  
    else:  
        return first  
"""
```

The space complexity of a dynamic programming solution is typically smaller than that of a greedy approach.

Which technique helps DP to avoid duplicate computation of intermediate results?

Given a set of non-negative integers, and a value x , determine if there exists a subset with a sum equal to x .

Which is the correct recursive solution?

What is the meaning of the values in the DP array for this dynamic programming solution for the subset sum? $DP[i][j]$ is true if ...

```
"""  
def subset_sum_dp(items, x):  
    DP = [[True for i in range(len(items))] for j in range(x+1)]  
    for i in range(1, x+1):  
        DP[i][0] = items[0] == i  
        for j in range(1, len(items)):  
            if items[j] > i:  
                DP[i][j] = DP[i][j-1]  
            else:  
                DP[i][j] = DP[i][j-1] or DP[i - items[j]][j-1]  
    return DP[x][len(items)-1]  
"""
```

Will linear relaxation will find the optimal solution?

What is the max flow of this graph?

How do we find a starting node x for a linear program?

We have an LP with n variables and m constraints. In a single step, the simplex algorithm checks roughly how many points as possible next node?

What is the maximum achievable value of this LP?

.....

max $2x + 3y + z$

subject to:

$x \leq 3$

$y \leq 5$

$z \leq 2$

$y + 3x \leq 12$

$y - x \leq 4$

$y + x \leq 7$

$x, y, z \geq 0$

.....

2

Who discovered NP-completeness?

A problem in NP ...

Assume some problem in NP admits a polynomial randomized algorithm. This means that

What is the most difficult problem in NP?

$(A \leq B)$ means that ...

Computing the size of a Minimum Vertex Cover is

Consider the following problem:

"Given a graph G and two nodes s and t , is there a path of length 2022 from s to t not repeating nodes?"

This problem is:

The Minimum Circuit Size Problem (MCSP) for boolean circuits with one output is

Assume all you have are AND, OR and NOT gates with at most 2 inputs. Computing the Parity of n bits

Consider an $N \times M$ grid graph (each node is linked to its at most 4 neighbors).

For which of sizes 4×4 , 4×5 , 5×4 and 5×5 does the graph have a hamiltonian cycle?

We presented a simple 2-approximation for Vertex Cover. We also showed that Independent Set \leq Vertex Cover.

Therefore, for Independent Set we can achieve an approximation ratio of

Vertex Cover and Independent Set are both NP-hard. Consider a combination:

"Given a connected graph G and a number k , is there a vertex cover of size k that is also an independent set?"

This problem is

Consider the following approximation for BinPacking:

"Go through the items in order, always keeping a single open bin.

When the new item does not fit in the current bin, close the current bin and open a new bin."

What is the approximation ratio of this algorithm?

Consider the following algorithm for 2D Euclidean TSP: traverse the n points according to their x-coordinate (left-most towards right, and then back again to left-most). What is the approximation ratio?

How many satisfying assignments does this boolean formula have: $(x_1 \text{ or } x_2)$ and $(x_3 \text{ or } x_4)$?

3

Which person is mentioned in the cryptography history book?

Can you decrypt this message: $\text{eqorwvcvkqpcnvjkpmkpi}$?

What is a prerequisite for secure encryption against a computationally unbounded adversary?

Finding a preimage of an ideal cryptographic one-way function with binary output of length λ takes how many evaluations in expectation?

Finding with high probability a collision for an ideal collision-resistant hash function with binary output of length λ takes how many evaluations?

Given $p=1009$, $g=11$. What is the discrete logarithm of 1?

Given collision resistant hash functions h_1 and h_2 . Which of the following is also collision resistant?

How can we share a testament among 3 children with perfect security?

How can Alice and Bob compute their mean salary without revealing the salary to each other?

What problem do digital signatures and PKI solve?

Digital signatures offer which property that message authentication codes do not?

TLS often uses RSA (similar to El-Gamal) for authentication but Diffie-Hellman for key exchange. What could be a reason for not using RSA for both?

Diffie-Hellman is vulnerable to man-in-the-middle attacks. Let Alice and Bob send ... in addition to (g^a) and (g^b) respectively. Fill the blank to fix the vulnerability.

A new secure channel was established via Diffie-Hellman (c2) over a channel already secured by a pre-shared key (c1). What is the adversary capable of after the pre-shared key is leaked?

Given $p=29$, $g=2$. What is the discrete logarithm of 7?

4

What was a keyword in this lecture?

In which format is data stored in a relational database?

What does SQL stand for?

Which tree is NOT a data structure for dictionaries?

Which keys can be used to represent the relationships between the tables?

If several elements compete for the same bucket in a hash table, we have a ...

Which of the following is the correct order of keywords for SQL statements?

The HAVING clause is like the WHERE clause but...

We have two tables T1 and T2 with sizes (number of rows) s_1 and s_2 respectively.

What is the minimum and maximum size of table T3 that is an inner join of tables T1 and T2?

Consider the database from the last video (table keys are underlined):

```
movie(<u>id</u>, title, d_id, year)
director(<u>id</u>, name)
actor(<u>id</u>, name)
acting(<u>a_id</u>, <u>m_id</u>, character)</pre>
```

What does the following query return?

```
SELECT * FROM director LIMIT(5);
```

Again the same database:

```
movie(<u>id</u>, title, d_id, year)
director(<u>id</u>, name)
actor(<u>id</u>, name)
acting(<u>a_id</u>, <u>m_id</u>, character)</pre>
```

What does the following query return?

```
SELECT actor.name FROM actor
INNER JOIN acting ON actor.id= acting.a_id
INNER JOIN movie ON movie.id= acting.m_id
WHERE movie.name = "War of the Worlds";
```

Again the same database:

```
movie(<u>id</u>, title, d_id, year)
director(<u>id</u>, name)
actor(<u>id</u>, name)
acting(<u>a_id</u>, <u>m_id</u>, character)</pre>
```

What does the following query return?

```
SELECT movie.title, COUNT(*) AS actors
FROM movie
JOIN acting ON movie.id= acting.m_id
GROUP BY movie.title, movie.id
ORDER BY actors DESC
LIMIT(1);
```

Again the same database:

```
movie(<u>id</u>, title, d_id, year)
director(<u>id</u>, name)
actor(<u>id</u>, name)
acting(<u>a_id</u>, <u>m_id</u>, character)</pre>
```

Which query returns the names and IDs of all actors in movies directed by Steven Spielberg?

Given the following input (4233, 4129, 5871, 1459, 7483, 5385, 9386, 1049) and the hash function $x \bmod 9$, how many collisions are there?

New database (see slide). How many rows does the output have?

5

What is the best approach to feature modeling in machine learning?

Which regularization method optimizes the L2 norm?

Overfitting happens when:

We want to approximate the function $f(x) = x^2 + 3x + 4$. What model yields the lowest variance (assume that we sample with some noise)?

What is the F1 score for the predictions

$(\hat{f}(x_1)=0, \hat{f}(x_2)=0, \hat{f}(x_3)=1, \hat{f}(x_4)=1)$

and the ground truth labels

$(y_1 = 0, y_2 = 1, y_3 = 1, y_4 = 1)$

How can we best estimate the risk (expected loss) of a model?

What is a sufficient condition for gradient descent to find the global optimum (we assume that the loss function is differentiable)?

What is the output of the softmax function applied to the following vector:
(1, 2, 3, 4)?

What is the purpose of pruning a decision tree?

Which statement is FALSE?

Consider the following datapoints that are contained in a node v of a decision

tree: $(D_v = \{x_1 = (0,0)^T, x_2 = (0,1)^T, x_3 = (1,0)^T, x_4 = (1,1)^T\})$

with ground truth labels: $(y_1 = 0, y_2 = y_3 = y_4 = 1)$ (for a classification task).

Which split is a valid split of v that minimizes the CART loss?

What is One versus Rest (OvR)?

Why do we usually optimize the log likelihood in practice instead of just the likelihood?

Why do we use the softmax function in regression tasks?

How many features do we have for polynomial regression with degree

$(m=2)$ and 3 variables (x_1, x_2) and (x_3) as input?

6

Which person is real?

What is an activation value of a node in a neural network?

What is a perceptron?

The universal approximation theorem states: Any continuous function can be ... by a feedforward neural network with a single hidden layer

What is backpropagation?

Why is the phenomenon of vanishing gradients a problem when training neural networks?

ReLU is a piecewise linear function and therefore a neural network with ReLU activation ...

Which of the following methods can be used as a regularization technique for neural networks?

Why do we use regularization in neural networks?

Recurrent neural networks are cyclic and therefore ...

Which of the following makes use of the weight sharing technique?

Which question in this quiz were generated by a neural network?

How will the marked weights change after one learning step?

Which binary function does the following network compute?

Given a 50x50 pixel RGB image and a CNN layer with 64 filters of size 5x5. How many learnable parameters are in this layer? Padding is used to keep the output shape the same, and each filter has a bias term.

7

Who does not belong in this Chapter (see slide)?

A Turing machine (TM) can ...

If a TM runs out of tape ...

What is the halting problem?

The halting problem is ...

The mortality problem asks if a computer program can determine whether ...

The mortality problem is ...

Which of the following decidability relationship statements is true?

Which of the following is a sufficient condition for undecidability of a problem P?

In which case can you conclude that a PCP does not have a solution?

PCP is ...

In which state will the TM (see slide) be when being started on input "01011"? (Tape pointer starts on leftmost 0 of input, what is the state after reading the right most 1 of input?)

For which of the following inputs will the TM above NOT find itself in the halting state H?

Which of the following characterizes all inputs for which the TM does NOT halt?

How many 1s are on the tape after running this TM (see slide)? If the answer is undefined (e.g., the TM does not halt) then answer 0.