PassQueen
a.k.a dumpsqueen.com

SAS Viya 3.5 Natural Language Processing and Computer Vision

SAS Institute A00-405

Version Demo

Total Demo Questions: 11

Total Premium Questions: 118

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Natural Language Processing	56
Topic 2, Computer Vision	51
Topic 3, Mix Questions	11
Total	118

QUESTION NO: 1 - (SIMULATION)

Complete the concept rule by adding the Boolean Operator that should be used to find concepts where the term " best " appears within five tokens of the term " picture "

Enter your answer in the field above.

ANSWER: See the explanation for the answer

Explanation:

Answer: `NEAR`

Use the `NEAR` Boolean operator between `best` and `picture`. In a concept rule, `NEAR` identifies terms that occur within five tokens of one another.

```
best NEAR picture
```

QUESTION NO: 2

Which statements are TRUE about convolutional filters in a CNN? (Choose two)

A. Filter weights are learned during training.

B. Each filter produces a feature map.

C. Filter values remain fixed and are not updated during training.

D. A filter processes only one pixel from an image.

ANSWER: A B

Explanation:

Filter weights are learned during training is correct. The training process updates convolutional filter weights to identify patterns that help minimize the model loss, such as edges, textures, shapes, or higher-level visual structures.

Each filter produces a feature map is also correct. A filter is applied across the spatial locations of its input, and the resulting activations form a feature map that indicates where the learned pattern is detected.

Filter values are not ordinarily fixed manually for the entire training process, and a convolutional filter is not restricted to processing only one pixel from an image. See the SAS Documentation portal for SAS Viya deep learning layer documentation.

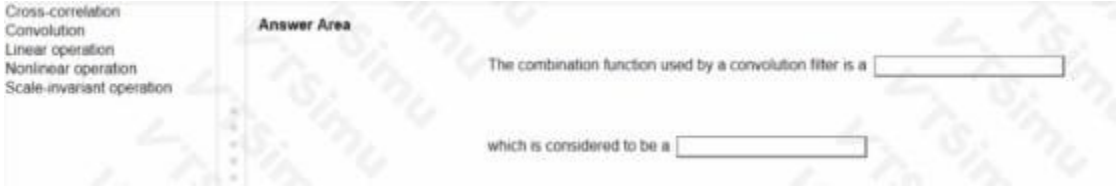
QUESTION NO: 3 - (DRAG DROP)

DRAG DROP

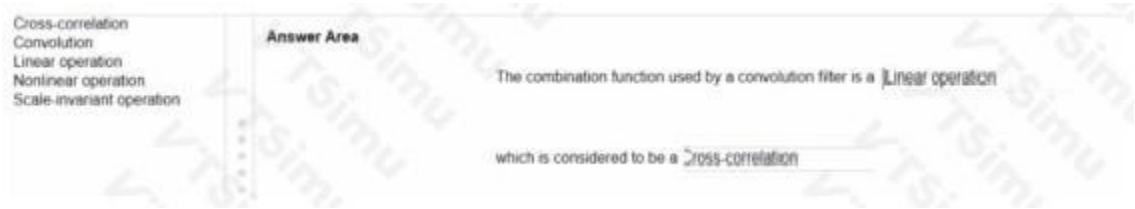
("Not*: This is an interactive item Follow the instructions to answer the question Scroll bars may appear if the windows are too small Each window can be re-sized by dragging on the

5 circles located between windows)

Complete the statement on the right with the options on the left.



ANSWER:



Explanation:

A convolution filter combines values from a local input region with the filter's learned weights. At each spatial location, the calculation forms a weighted sum of the input values, generally followed by a bias term. Because this calculation is based on multiplying input values by fixed weights and adding the resulting products, it is a linear operation with respect to the input. This local weighted-combination mechanism is what allows a model to detect learned visual or textual patterns wherever they occur in the input.

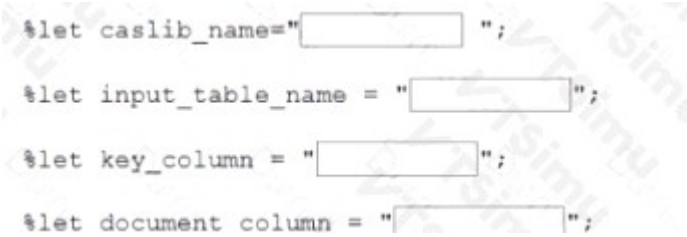
In the terminology used for convolutional neural networks, the operation performed by the filter is commonly considered a cross-correlation. The filter is slid across the input in its stored orientation, and corresponding input and filter positions are multiplied and summed. Mathematical convolution would reverse the filter before carrying out that summation. Deep-learning software and SAS documentation commonly use the familiar term convolution for the layer and filter, while the underlying sliding weighted-sum calculation is cross-correlation. This distinction does not alter the model's learning capability because the filter weights are learned during training.

Thus, completing the statement with "Linear operation" and "Cross-correlation" accurately describes both the algebraic nature of the filter calculation and the standard deep-learning interpretation of the spatial operation. Additional background on convolutional neural network layers is available in the SAS deep learning documentation and in the PyTorch Conv2d documentation, which notes that the operation is cross-correlation.

QUESTION NO: 4 - (SIMULATION)

Fill in the blanks in the code extracted from the source code download from a Category node.

Your goal is to categorize the documents in the my Score Tbl table which has been uploaded to the myCasUb library The text column Comments are identified by myKey, and the scoring output tables are to be saved in the same library as the score table



Enter your answer in the fields above.

ANSWER: See the explanation for the answer

Explanation:

Fill in the extracted Category node source-code parameters as follows:

```
%let caslib_name="myCasUb";
%let input_table_name = "my Score Tbl";
%let key_column = "myKey";
%let document_column = "Comments";
```

The CASLIB is `myCasUb`, the uploaded scoring table is `my Score Tbl`, `myKey` identifies each document, and `Comments` is the text column to categorize. Using the same CASLIB also causes the scoring output tables to be saved in `myCasUb`.

QUESTION NO: 5

Refer to the exhibit.

The pixel intensity values for the top left corner of a 224x224 grey scale image are provided in Exhibit A This image is provided as input to the convolutional filter (Exhibit B) with equal zero padding of size 2 on all sides

What is the first value in the feature map resulting from applying this filler (cross-correlation operation)?

A. 186

B. 0

C. 10 D. 97

ANSWER: A

Explanation:

186 is the first feature-map value produced by the stated cross-correlation operation. With zero padding of size 2, the filter is positioned first at the upper-left of the padded image. Consequently, locations in the filter window that extend above or to the left of the original image contribute zero. The calculation uses the pixel values that overlap the top-left image region, multiplies each of those values by the corresponding value in the displayed filter, and sums the resulting products. Because the operation is explicitly cross-correlation, the filter is used in its displayed orientation; it is not flipped horizontally or vertically before the elementwise multiplication. Applying the Exhibit B weights to the nonzero overlapping pixels from Exhibit A, while retaining zero contributions from all padded positions, yields a total of 186. This total is written to the first, upper-left location of the output feature map. SAS deep-learning documentation describes convolution layers as applying learned filters to image inputs to generate feature maps; padding controls the treatment of border positions and allows filters to be evaluated at image edges. See the SAS Deep Learning Programming Guide and the SAS documentation portal.

QUESTION NO: 6

Given an equivalent number of neurons why would you use an LSTM-based recurrent neural network over a conventional recurrent neural network?

A. Implementation advantages - easier to parallelize

B. Size - uses fewer parameters

C. Efficiency - requires fewer computations

D. Sequence memory - learns data dependencies spanning more sequence elements

ANSWER: D

Explanation:

Sequence memory - learns data dependencies spanning more sequence elements is correct. A conventional recurrent neural network passes information through a repeatedly transformed hidden state. During training with backpropagation through time, gradients can diminish or grow as they cross many time steps, making it difficult for the network to retain and learn relationships between events that are far apart in a sequence. An LSTM architecture introduces a persistent cell state and learned gates that regulate what information is written, retained, and exposed at each step. This controlled path lets useful information and gradient signals persist over longer spans, so an LSTM can model long-term temporal dependencies more reliably when the number of neurons is held equivalent. In practical NLP tasks, this is valuable when the meaning or classification of a later token depends on words, entities, or context that occurred many positions earlier. SAS describes recurrent layers as layers designed for sequential data and provides LSTM support for sequence modeling in its deep learning procedures. The underlying LSTM design is detailed in the original research paper, which explains its constant-error flow and gated memory mechanism. See the SAS Documentation portal and Long Short-Term Memory paper.

QUESTION NO: 7

An analyst is using image augmentation while training an image-classification model.

Which two practices are appropriate? (Choose two)

A. Apply label-preserving transformations to training images.

B. Keep the original class label for an augmented image.

C. Augment the test data each time the model is scored.

D. Change each augmented image to a new target class.

E. Use augmentation only after final model assessment.

ANSWER: A B

Explanation:

Image augmentation should be applied to training images so that the model sees varied, but still meaningful, versions of the training examples. Appropriate transformations can include small rotations, translations, flips, or changes in brightness when those transformations do not change the image's class. The original label must be retained because the augmented image represents the same object or class as its source image.

Validation and test data should normally remain unchanged so that their performance measures represent evaluation on a stable, held-out data set. Augmentation can increase the diversity of training examples and help reduce overfitting. See the SAS Documentation portal for SAS Viya image processing and deep learning documentation.

QUESTION NO: 8

An analyst needs to identify product serial numbers that appear in customer service notes. The serial numbers follow several business-specific formats.

Which SAS Visual Text Analytics node should the analyst use to define and apply the extraction rules?

A. Concepts

B. Topics

C. Model Comparison

D. Data Exploration

ANSWER: A

Explanation:

The Concepts node is the appropriate choice because it enables an analyst to define custom concepts that identify domain-specific entities and expressions in document text. Custom concepts can use linguistic and rule-based logic to recognize multiple valid forms of an item, such as product serial numbers, dates, addresses, or specialized business terminology.

After the concept is defined, the node can identify occurrences of the concept in the input documents and make the extraction results available to downstream processing. SAS Visual Text Analytics documentation describes concept extraction and custom concept authoring in the SAS Visual Text Analytics 3.5 User's Guide.

QUESTION NO: 9

An analyst is preparing image data for a CNN that expects identically sized input images.

Which two actions are appropriate during image preparation? (Choose two)

A. Resize images to the required input dimensions.

B. Scale or normalize pixel values consistently.

C. Assign a new class label after every resize operation.

D. Convert every image to a different width and height.

E. Remove the image identifier from the input table.

ANSWER: A B

Explanation:

Images supplied to a CNN in a batch must have compatible dimensions. Resizing images to the model input size ensures that the network can process each image consistently. Scaling or normalizing pixel values is also appropriate because it places input values on a consistent range and can improve the stability and efficiency of training.

The image data should retain the correct class label after transformations. Image-processing actions in SAS Viya can be used to load, transform, and prepare image data before it is supplied to a deep learning model. See the SAS Documentation portal for SAS Viya image processing and deep learning documentation.

QUESTION NO: 10

Which statements are TRUE about importing SAS Contextual Analysis projects into SAS

Visual Text Analytics when creating a new project? (Choose two)

A. All custom concept and some predefined concept settings are imported

B. All machine generated topics and sentiment scores are imported

C. All custom categories are imported

D. All sentiment classification and probability scores are imported

ANSWER: A C

Explanation:

When a SAS Contextual Analysis project is imported as a new project in SAS Visual Text Analytics, the reusable rule-based authoring assets are carried forward. Therefore, *All custom concept and some predefined concept settings are imported* is correct. Custom concepts contain user-authored extraction logic, and supported settings associated with predefined concepts can be retained so that the imported project preserves applicable concept-model configuration. This helps users continue developing and testing concept rules in the Visual Text Analytics interface rather than recreating the supported concept work manually.

All custom categories are imported is also correct. Categories are authored classification assets in Contextual Analysis, and their definitions are imported to preserve the project's custom categorization work. This allows the imported Visual Text Analytics project to retain the established category structure and associated custom rules for continued refinement and scoring.

The distinction is that import is intended to migrate supported project definitions and authoring assets, including custom concepts and custom categories, rather than prior generated analytical output. SAS documentation for Visual Text Analytics is available in the SAS Visual Text Analytics: User's Guide; SAS also provides product documentation through its SAS Documentation portal.

QUESTION NO: 11

Assume you are using an addLayer call with dcopout=o. 5 to add a pooling layer into a CNN model What is the outcome of the addLayer call?

A. The values from half of the output neurons of the pooling layer are randomly set to zero.

B. The output feature map size is increased 2 times on width and height.

C. The values from half of the input neurons to the pooling layer are randomly set to zero.

D. The output feature map size is reduced by half on width and height.

ANSWER: A

Explanation:

The values from half of the output neurons of the pooling layer are randomly set to zero. The parameter name in the question appears to contain typographical errors, but it clearly intends `dropout=0.5` in an `addLayer` call. In SAS Deep Learning, dropout is a regularization setting that randomly suppresses a specified proportion of a layer's output activations during training. Therefore, a dropout value of 0.5 means that each output activation produced by the pooling layer has a 50% probability of being set to zero for a training pass. The random masking changes from one training iteration to another, which helps limit co-adaptation among learned features and can improve the model's ability to generalize to previously unseen data. Dropout affects activation values; it does not by itself define the pooling window, stride, or output feature-map dimensions. Pooling-layer output dimensions are instead determined by settings such as the pooling operation, kernel dimensions, stride, and padding. For SAS documentation on deep-learning layers and their parameters, see the SAS Viya 3.5 Deep Learning Programming Guide and the SAS Documentation portal.