Recurrent Neural Networks (RNNs).

(Sequence-to-sequence learning)

Time-series data

⇒ Temporal dependence

⇒" The input & output sizes vary from example to example"

Application

① Language Translation

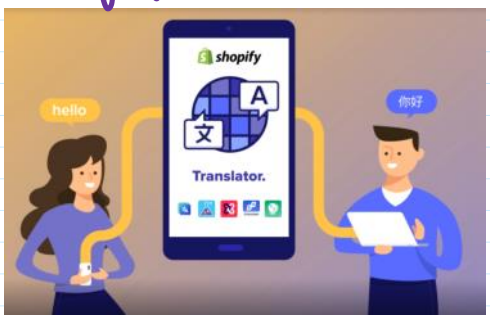

② Speech Signal to text.

③ Image Captioning

④ Visual Question Answering :-

Onehot Encoding

3 class dataset

cat	dog	bird

samples →

1. Cat → [1 0 0]
2. Cat → [1 0 0]
3. Bird → [0 0 1]
4. Dog → [0 1 0]
5. Bird → [0 0 1]

In RNNs :—

1. The weather is good today.
2. The weather is pleasant today.

The → [1 0 0 0 0 0]
weather → [0 1 0 0 0 0]
is → [0 0 1 0 0 0]
good → [0 0 0 1 0 0]
pleasant → [0 0 0 0 1 0]
today → [0 0 0 0 0 1]

Word2Vec Embeddings.

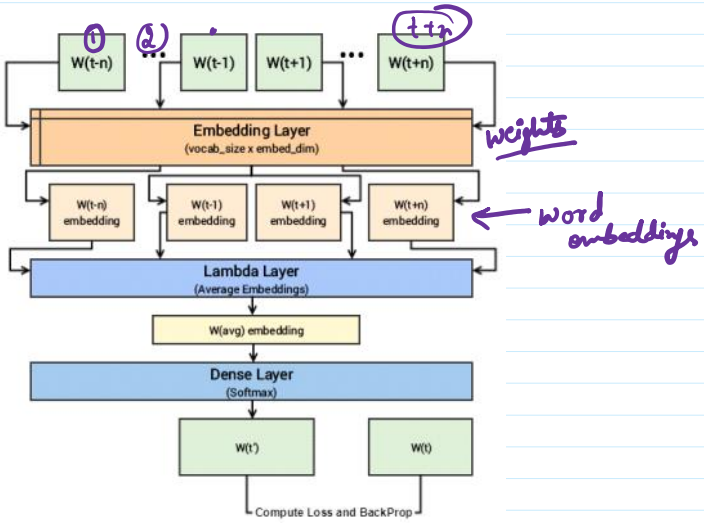

Steps :-

① Words are passed to an embedding layer (initialized weights).

② Word embeddings → Lambda Layer (Avg. embeddings)

③ Average Embedding → Predicted Word (Softmax)

④ Loss (predicted, actual).

Input → $\{x_1, x_2, \ldots, x_t, \ldots, x_{n-1}, x_n\}$

Output → $\{y_1, y_2, \ldots, y_t, \ldots, y_{n-1}, y_n\}$