CS5740: Natural Language Processing

Recurrent Neural Networks

Instructor: Yoav Artzi

Text Classification

- Consider the simple example of sentiment classification

  Goal: classify sentiment
  Example A: How can you not see this movie? → 👍
  Example B: You should not see this movie? → 👎

- The order of words is important

- Transformers handle it be just keeping everything in "memory" and constantly looking back to it

- But: what about maintaining a state and discard the past?

  - Why is this a good thing?

Finite State Machines

- Simple, classical way of representing state
- Current state: saves necessary past information
- Example: email address parsing

Deterministic Finite State Machines

- S – states
- Σ – vocabulary
- $s_0 \in S$ – start state
- $R: S \times \Sigma \to S$ – transition function

- What does it do?
  - Maps input $w_1, \dots, w_n$ to states $s_1, \dots, s_n$
  - For all $i \in \{1, \dots, n\}$
  $$s_i = R(s_{i-1}, w_i)$$
- Try to think on how we can use it for POS tagging and language modeling

Types of State Machines

- Acceptor
  - Compute final state s_n and make a decision based on it: $y = O(s_n)$
- Transducers
  - Apply function $y_i = O(s_i)$ to produce output for each intermediate state
- Encoders
  - Compute final state s_n, and use it in another model

Recurrent Neural Networks

- Motivation:
  - Neural network model, but with state
    - Ideally: that doesn't grow as history grows
  - How can we borrow ideas from FSMs?
- RNNs are FSMs …
  - … with a twist
  - No longer finite in the same sense

RNN

- $S = \mathbb{R}^{d_{hid}}$ - hidden state space

- $\Sigma = \mathbb{R}^{d_{in}}$ - input state space

- $\boldsymbol{s}_0 \in S$ - initial state vector

- $R : \mathbb{R}^{d_{in}} \times \mathbb{R}^{d_{hid}} \to \mathbb{R}^{d_{hid}}$ - transition function

- Simple definition of R:
$$R_{Elman}(\boldsymbol{s}, \boldsymbol{x}) = \tanh([\boldsymbol{x}, \boldsymbol{s}]\boldsymbol{W} + \boldsymbol{b})$$

* Notation: vectors and matrices are bold   Elman (1990)

RNN

- Map from dense sequence to dense representation
  - $x_1, \ldots, x_n \rightarrow s_1, \ldots, s_n$
  - For all $i \in \{1, \ldots, n\}$
  $$s_i = R(s_{i-1}, x_i)$$
  - R is parameterized, and parameters are shared between all steps
  - Example:
  $$s_4 = R(s_3, x_4) = \cdots = R(R(R(R(s_0, x_1), x_2), x_3), x_4)$$

RNNs

- Hidden states $\boldsymbol{s}_i$ can be used in different ways
- Similar to finite state machines
 - Acceptor
  - Transducer
  - Encoder
- Output function maps vectors to symbols:

$$O: \mathbb{R}^{d_{hid}} \rightarrow \mathbb{R}^{d_{out}}$$
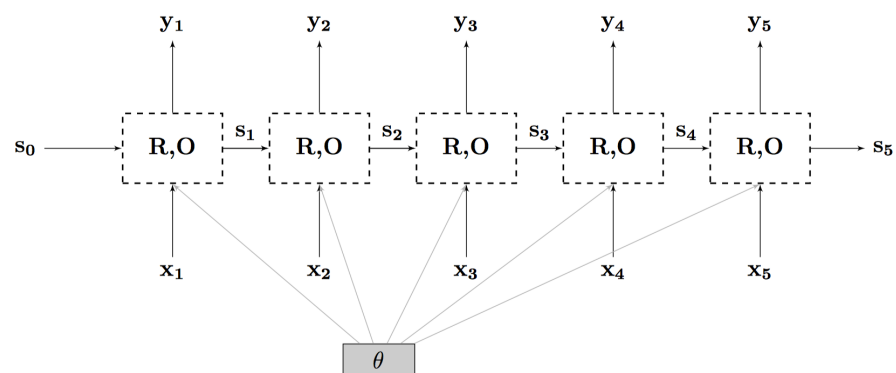
- For example: single layer + softmax

$$O(\boldsymbol{s}_i) = \text{softmax}(\boldsymbol{s}_i \boldsymbol{W} + \boldsymbol{b})$$

Visual Representation

Recursive Representation

Unrolled Representation

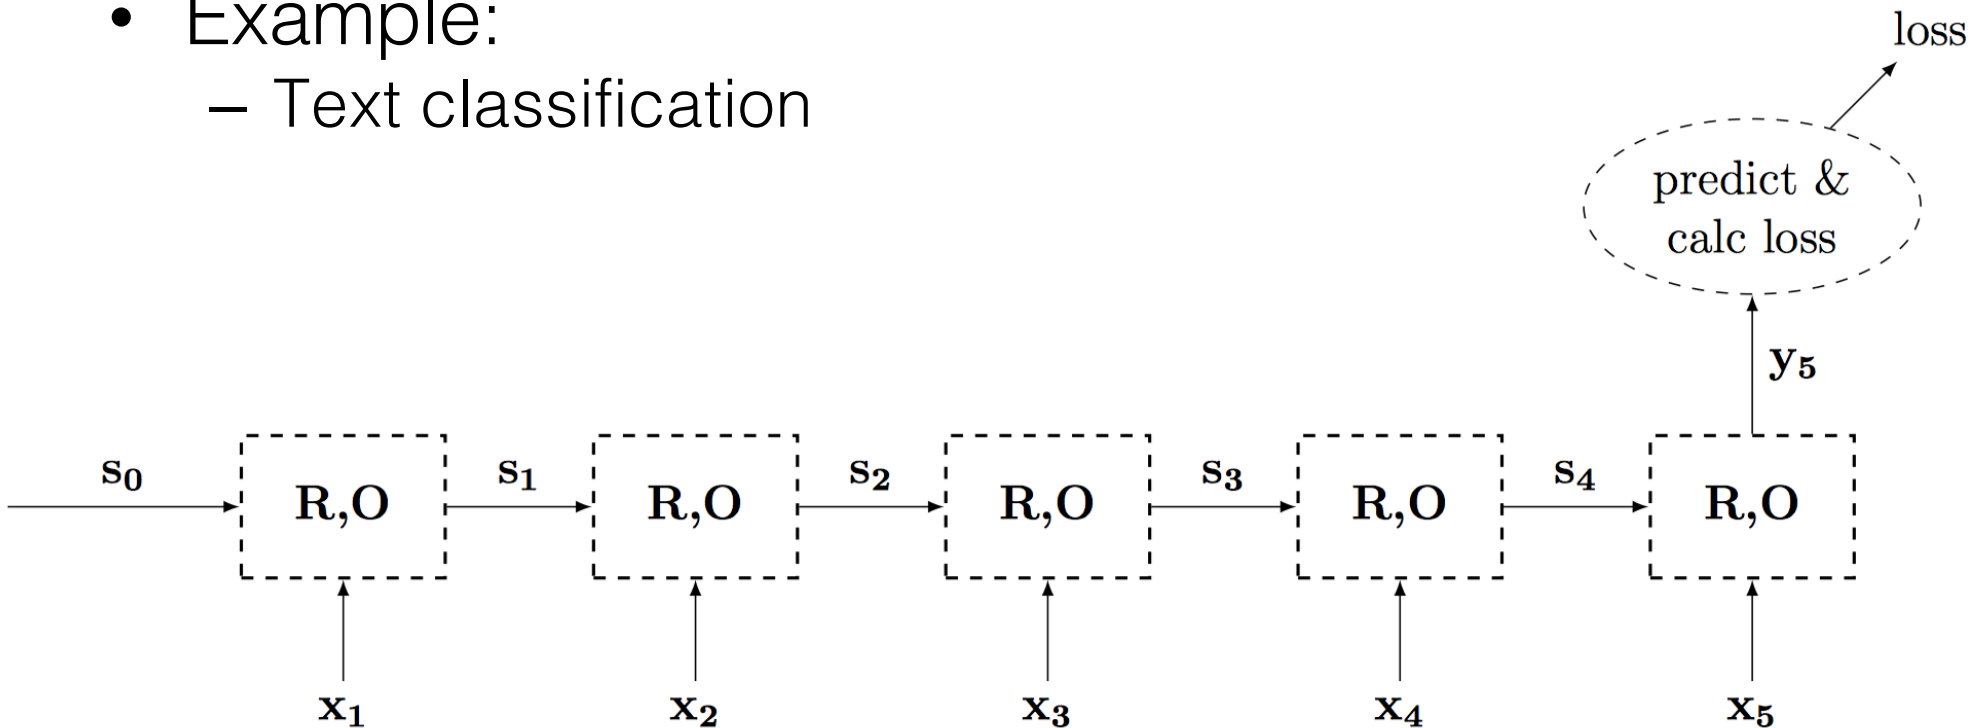# Visual Representation

Training

- RNNs are trained with SGD and Backprop
- Define loss over outputs
  - Depends on supervision and task
- Backpropagation through time (BPTT)
  - Use unrolled representation
  - Run forward propagation
  - Run backward propagation
  - Update all weights
- Weights are shared between time steps
  - Sum the contributions of each time step to the gradient
- Inefficient
  - Batch helps, common but tricky to implement with variable-size models (good helper methods in PyTorch, non-issue with auto batching in DyNet)

Types of State Machines

- Acceptor
  - Compute final state s_n and make a decision based on it: $y = O(s_n)$
- Transducers
  - Apply function $y_i = O(s_i)$ to produce output for each intermediate state
- Encoders
 - Compute final state s_n, and use it in another model

RNN: Acceptor Architecture

- Only care about the output from the last hidden state
- Train: supervised, loss on prediction
- Example:
  - Text classification

RNN: Transducer Architecture
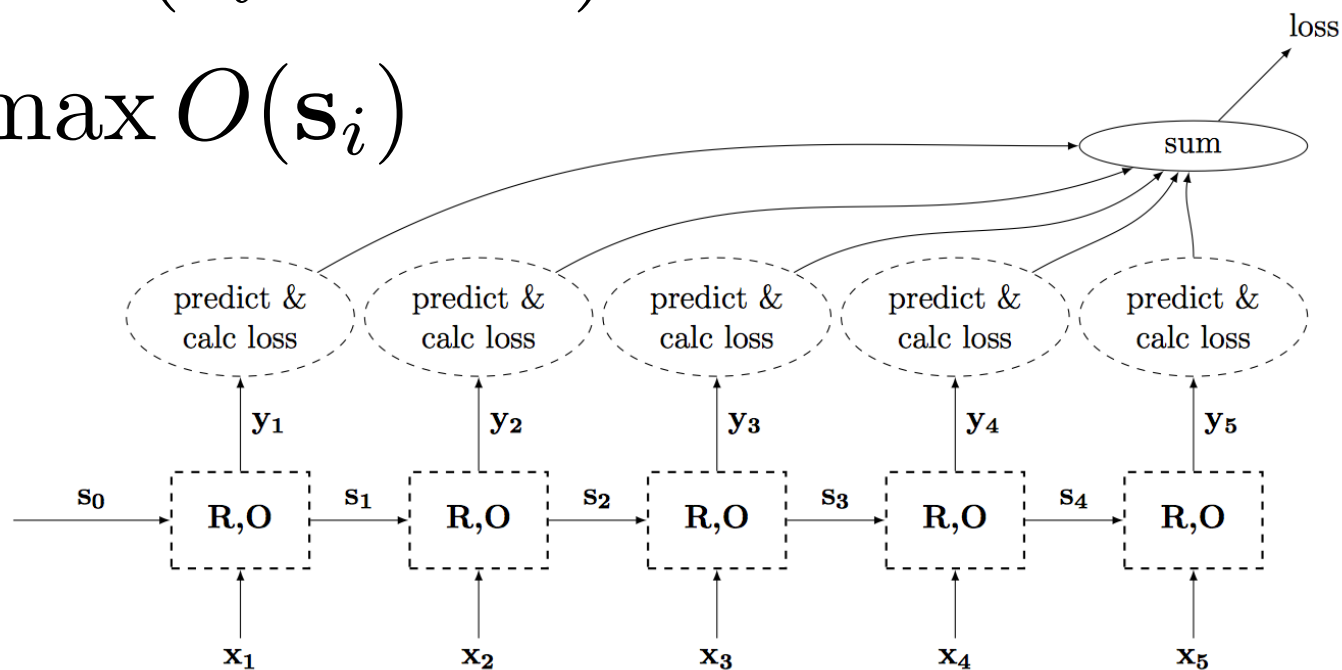
- Predict output for every time step

RNN: Transducer Architecture

$$X = \mathbf{x}_1, \ldots, \mathbf{x}_n$$

$$\mathbf{s}_i = R(\mathbf{s}_{i-1}, \mathbf{x}_i), i = 1, \ldots, n$$

$$O(\mathbf{s}_i) = \mathrm{softmax}(\mathbf{s}_i \mathbf{W} + \mathbf{b})$$

$$\hat{y}_i = \arg\max O(\mathbf{s}_i)$$

Language Modeling

- Input: $X = x_1, \ldots, x_n$
- Goal: compute $p(X)$
- Bi-gram decomposition:

$$p(X) = \prod_{i=1}^{n} p(x_i \mid x_{i-1})$$

- RNNs are non-Markovian models:

$$p(X) = \prod_{i=1}^{n} p(x_i \mid x_1, \ldots, x_{i-1})$$

Language Modeling

- Input: $X = x_1, \ldots, x_n$
- Goal: compute $p(X)$
- Model:

$$p(X) = \prod_{i=1}^{n} p(x_i \mid x_1, \ldots, x_{i-1})$$

$$p(x_i \mid x_1, \ldots, x_{i-1}) = O(\boldsymbol{s}_i) = O(R(\boldsymbol{s}_i, \boldsymbol{x}_{i-1}))$$

$$O(\boldsymbol{s}_i) = \text{softmax}(s_i \boldsymbol{W} + \boldsymbol{b})$$

- Predict next token $\hat{y}_i$ as we go:

$$\hat{y}_i = \text{argmax} O(\boldsymbol{s}_i)$$

RNN: Transducer Architecture

- Predict output for every time step
- Examples:
  - Language modeling
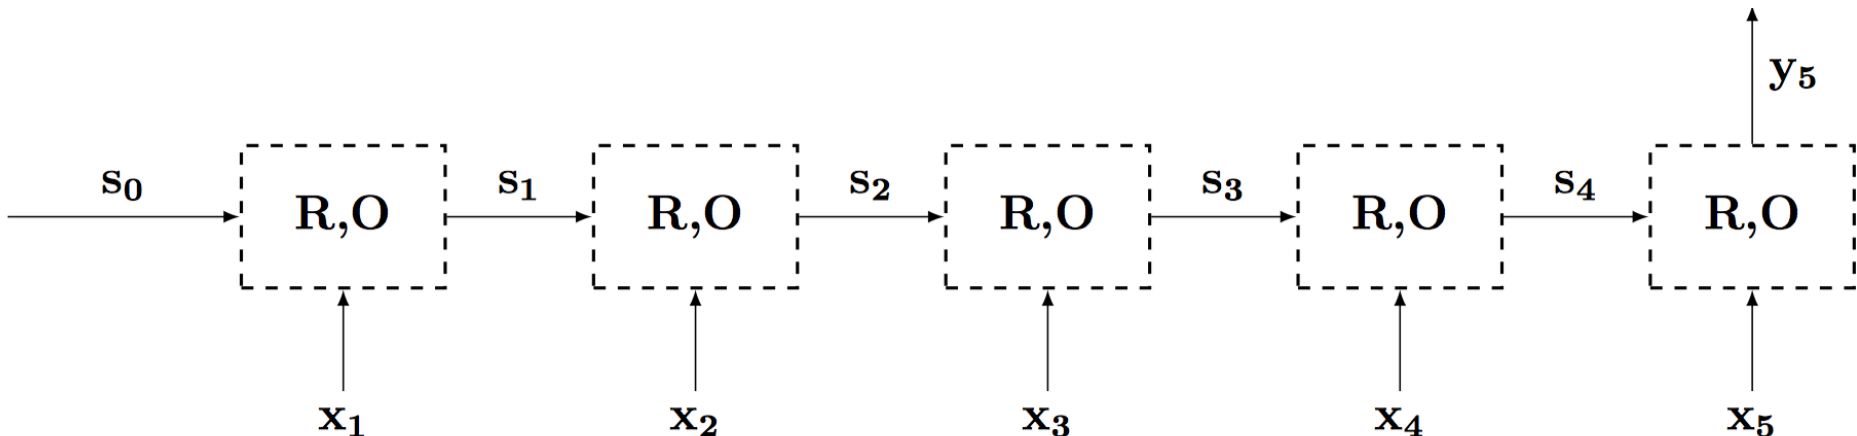 - NER
  - Extractive QA

RNN: Encoder Architecture

- Similar to acceptor
- Difference: last state is used as input to another model and not for prediction
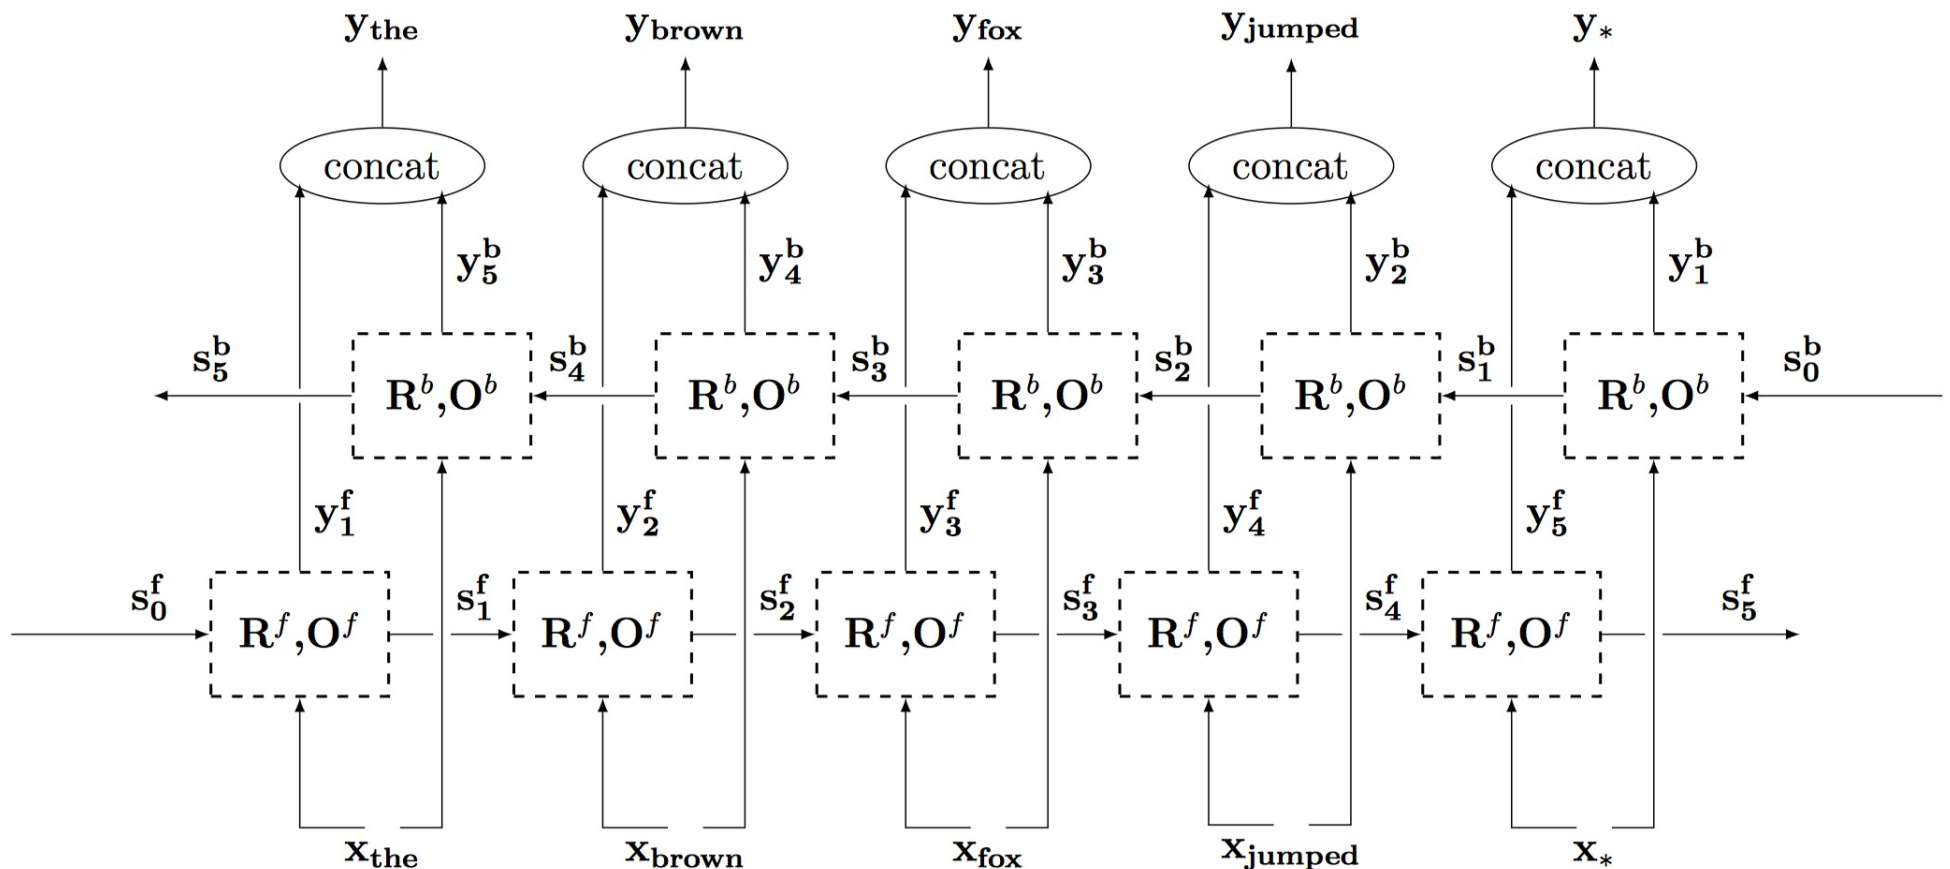
$$O(s_i) = s_i \rightarrow y_n = s_n$$

- Example:
  - Sentence embedding (like words, but sentences)
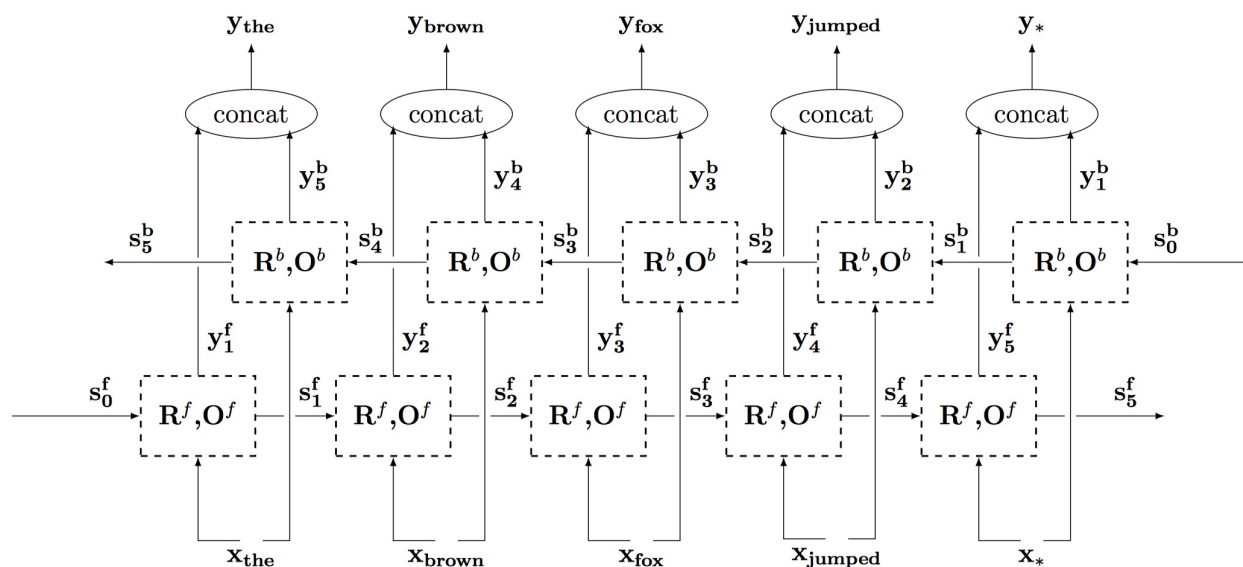
Bidirectional RNNs

- RNN decisions are based on historical data only
  - How can we account for future input?
- When is it relevant? Feasible?
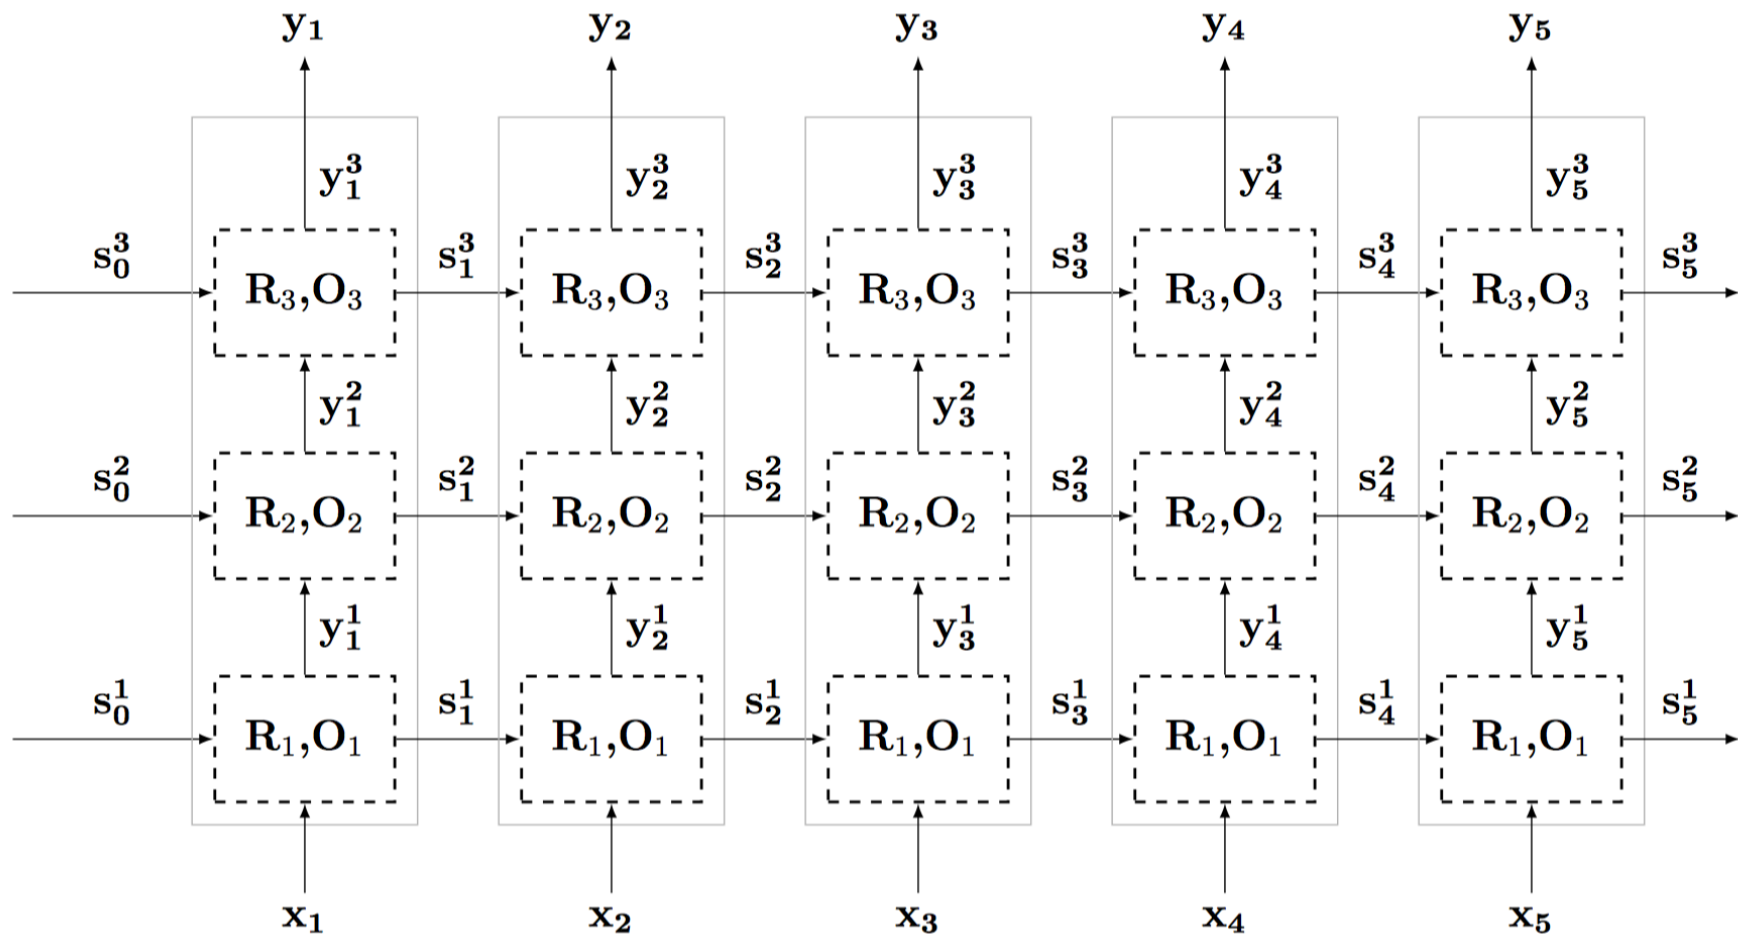
Bidirectional RNNs

- RNN decisions are based on historical data only
  - How can we account for future input?
- When is it relevant? Feasible?
    - When all the input is available. Not for real-time input.
- Probabilistic model, for example for language modeling:

$$p(X) = \prod_{i=1}^{n} p(x_i \mid x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$$
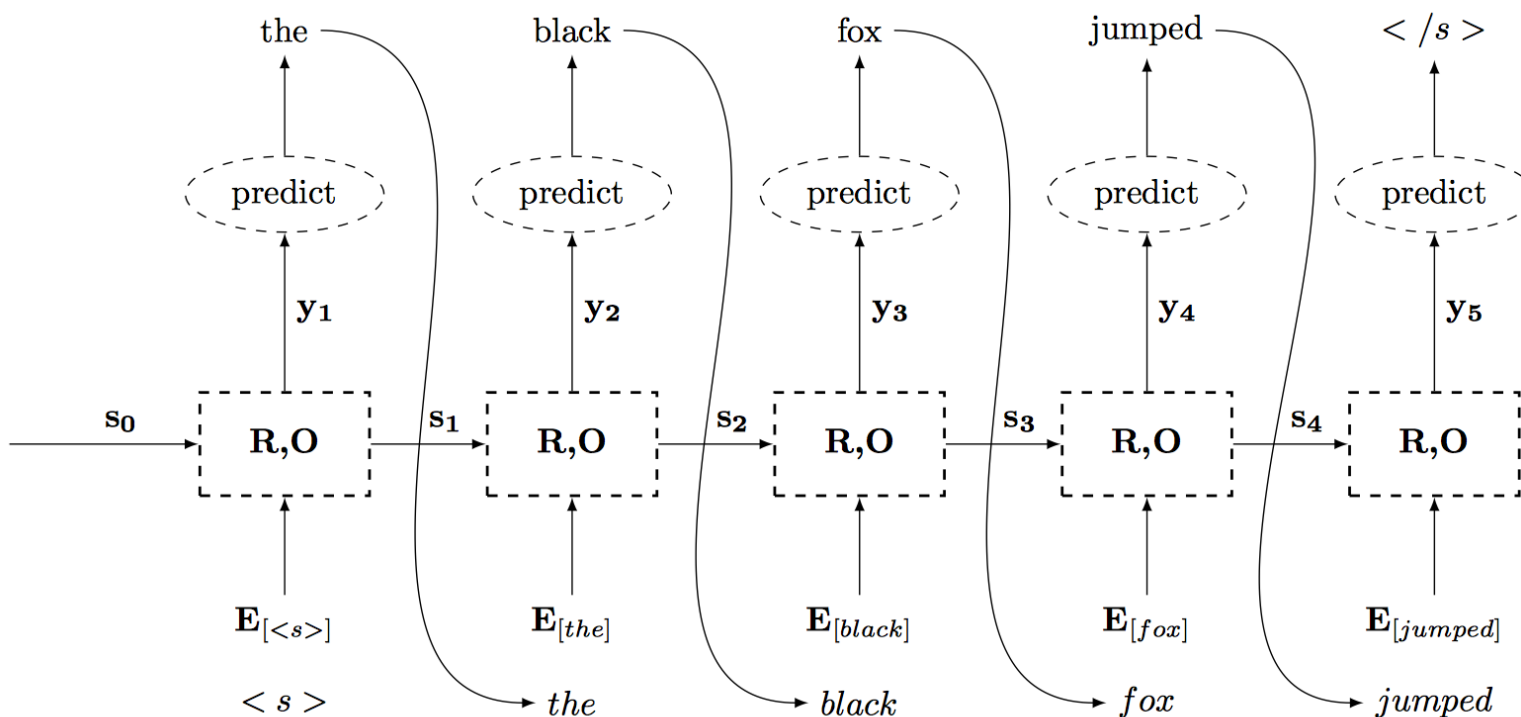
Deep RNNs

- Can also make RNNs deeper (vertically) to increase model capacity

RNN: Generator

- Special case of the transducer architecture
- Generation conditioned on s_0
- Probabilistic model:

$$p(X \mid s_0) = \prod_{i=1}^{n} p(x_i \mid x_1, \ldots, x_{i-1}, s_0)$$
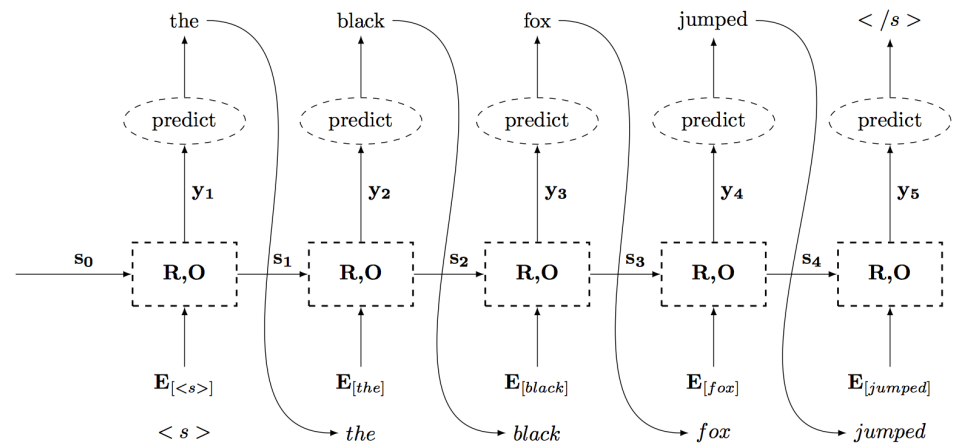
RNN: Generator

- Stop when generating the STOP token
- During learning (usually): force predicting the annotated token and compute loss

$$\mathbf{s}_j = R(\mathbf{s}_{j-1}, E(\hat{\mathbf{t}}_{j-1}))$$

$$O(\mathbf{s}_j) = \text{softmax}(\mathbf{s}_j \mathbf{W} + \mathbf{b})$$

$$\hat{\mathbf{t}}_j = \arg\max O(\mathbf{s}_j)$$

Example: Caption Generation

- Given: image I
- Goal: generate caption
- Set $s_0 = \mathrm{CNN}(I)$
- Model:

$$p(X \mid I) = \prod_{i=1}^{n} p(x_i \mid x_1, \ldots, x_{i-1}, I)$$

Examples from Karpathy
and Fei-Fei 2015

"little girl is eating piece of cake."

"baseball player is throwing ball in game."

"woman is holding bunch of bananas."

"a young boy is holding a baseball bat."

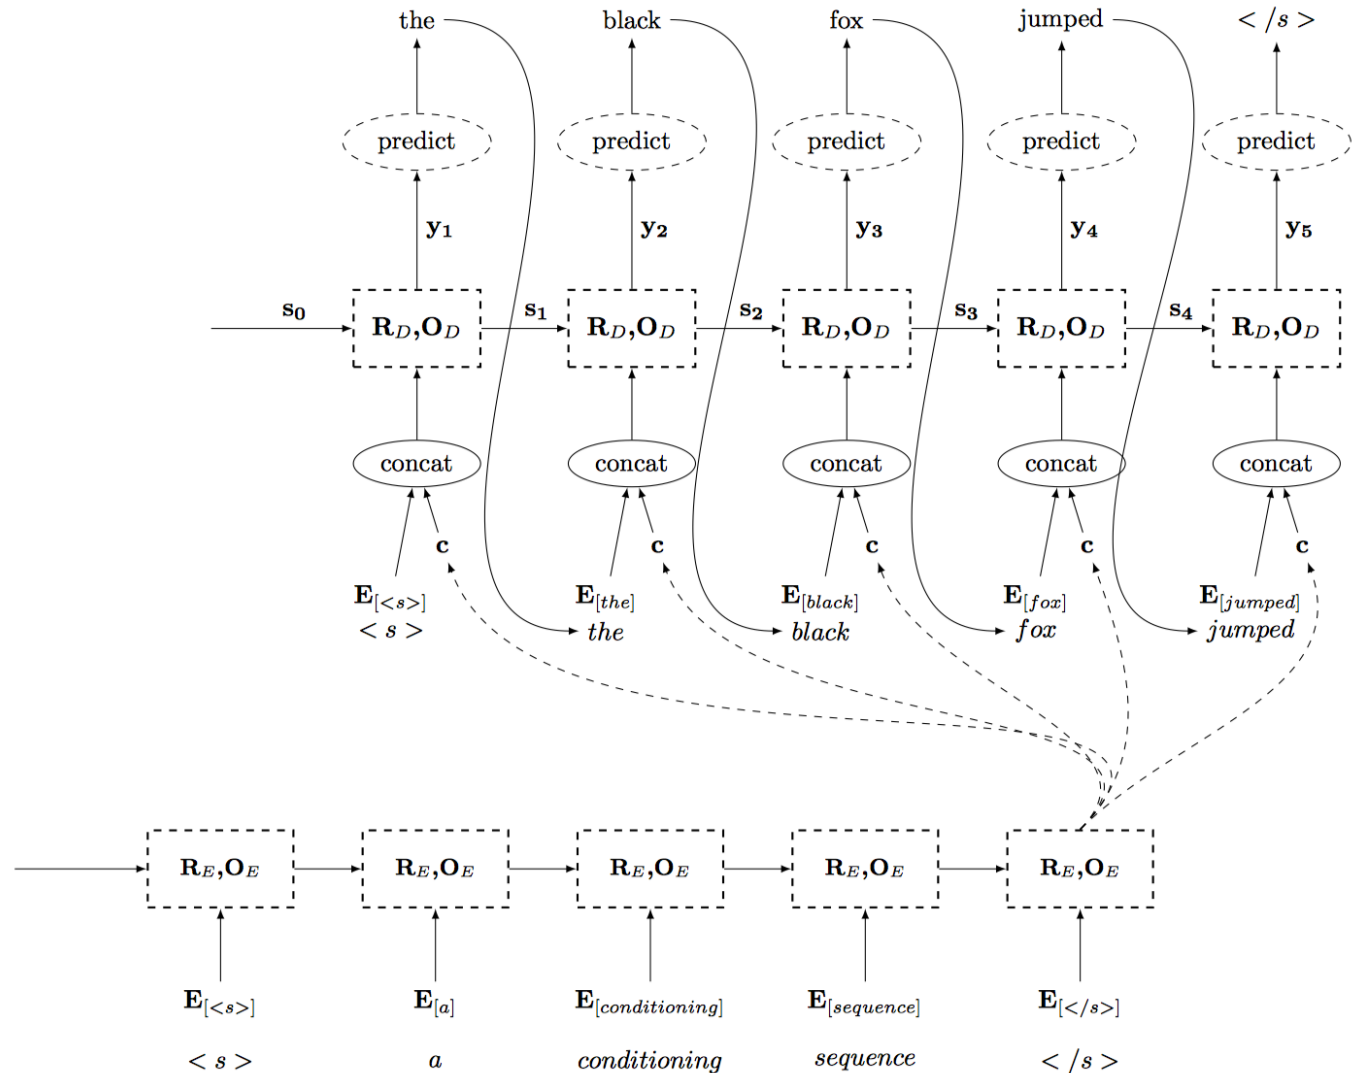"a cat is sitting on a couch with a remote control."

"a woman holding a teddy bear in front of a mirror."

Sequence-to-Sequence

- Connect encoder and generator
- Many alternatives:
  - Set generator s_0^d to encoder output s_n^e
  - Concatenate s_n^e with each step input during generation
- Examples:
  - Machine translation
  - Chatbots
  - Dialog systems
- Can also generate other sequences – not only natural language!
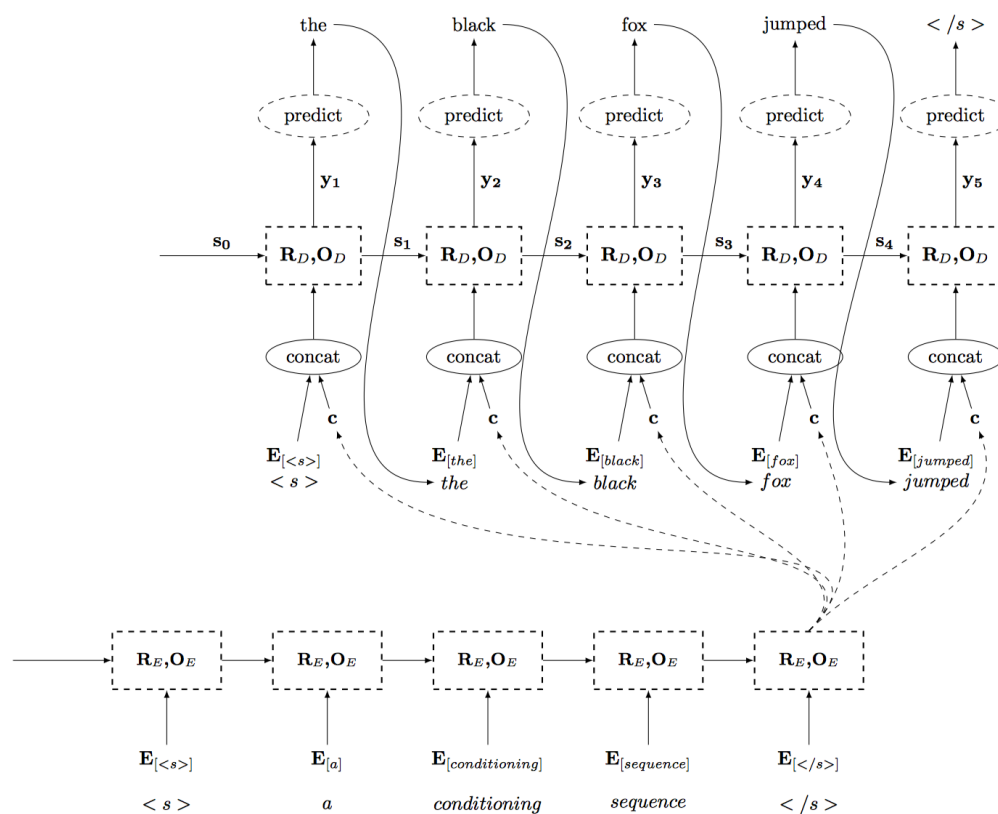
Sequence-to-Sequence

$$X = x_1, \ldots, x_n$$

$$\mathbf{s}_i^E = R_E(\mathbf{s}_{i-1}^E, \mathbf{E}_{[x_i]}), i = 1, \ldots, n$$

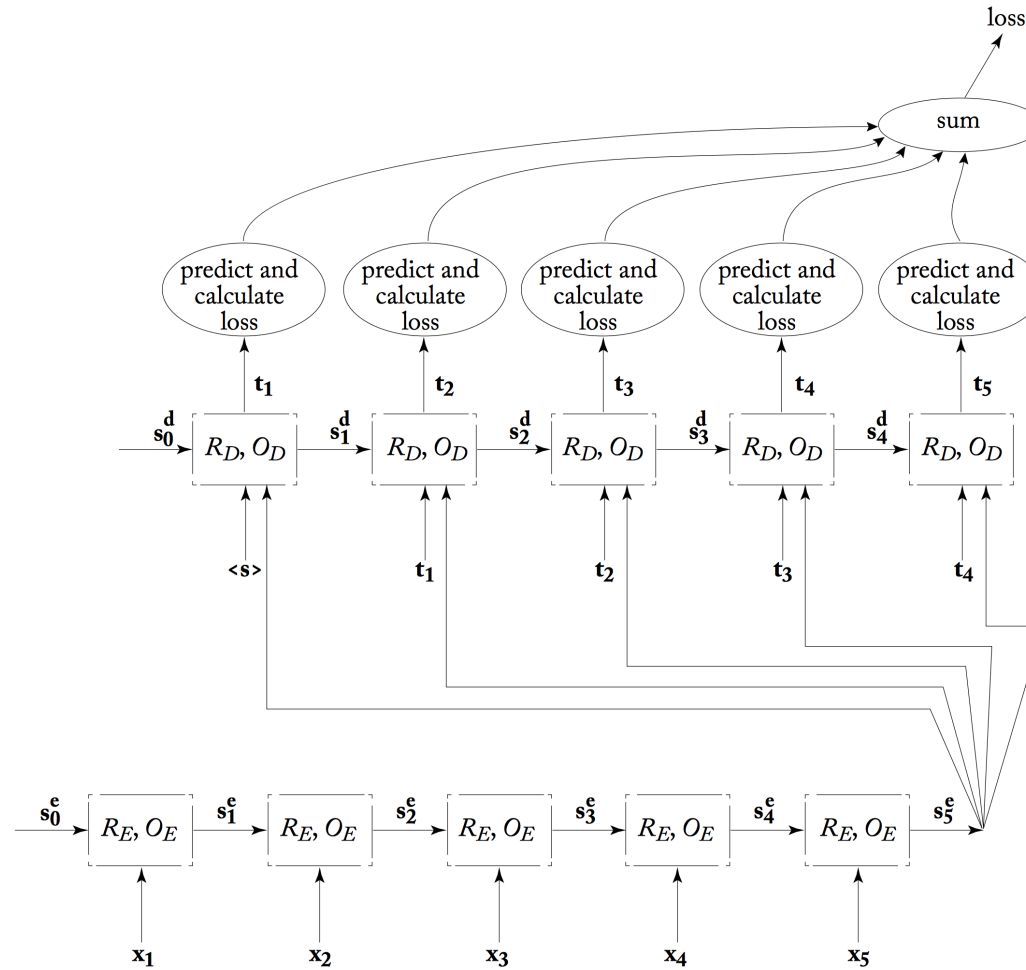$$\mathbf{c} = O_E(\mathbf{s}_n^E)$$

$$\mathbf{s}_j^D = R_D(\mathbf{s}_{j-1}^D, [\mathbf{E}_{[\hat{t}_{j-1}]}); \mathbf{c}])$$

$$O_D(\mathbf{s}_j^D) = \text{softmax}(\mathbf{s}_j^D \mathbf{W} + \mathbf{b})$$

$$\hat{t}_j = \arg\max O_D(\mathbf{s}_j^D)$$

Sequence-to-Sequence Training Graph

Long-range Interactions

- Promise: Learn long-range interactions of language from data
- Example:

  How can you not see this movie?

  You should not see this movie.

- Sometimes: requires "remembering" early state
  – Key signal here is at s_1, but gradient is at s_n

Long-term Gradients

- Gradient go through (many) multiplications
- OK at end layers → close to the loss
- But: issue with early layers
- For example, derivative of $\tanh$

$$\frac{d}{dx}\tanh x = 1 - \tanh^2 x$$

  – Large activation → gradient disappears (vanish)
- In other activation functions, values can become larger and larger (explode)

Exploding Gradients

- Common when there is no saturation in activation (e.g., ReLu) and we get exponential blowup because of product rule in backprop
- Result: reasonable short-term gradient, but bad long-term ones
- Common heuristic:
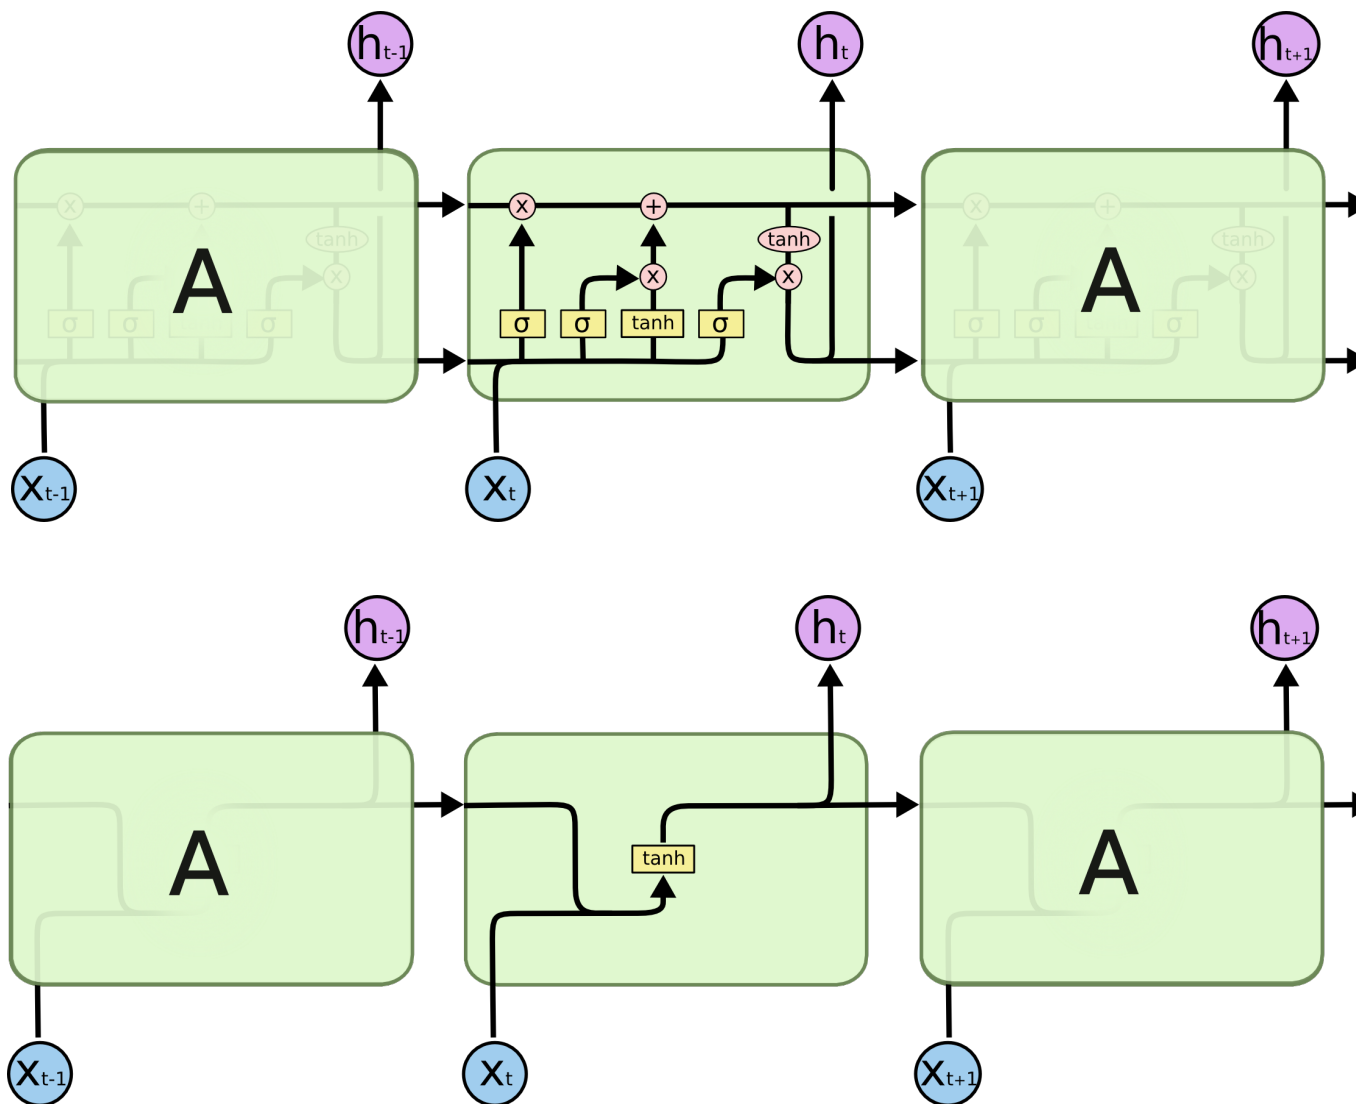 - Gradient clipping: bounding all gradients by maximum value

ReLU

$R(z) = max(0, \ z)$

Vanishing Gradients

- Occurs when multiplying small values
  - For example: when $\tanh$ saturates
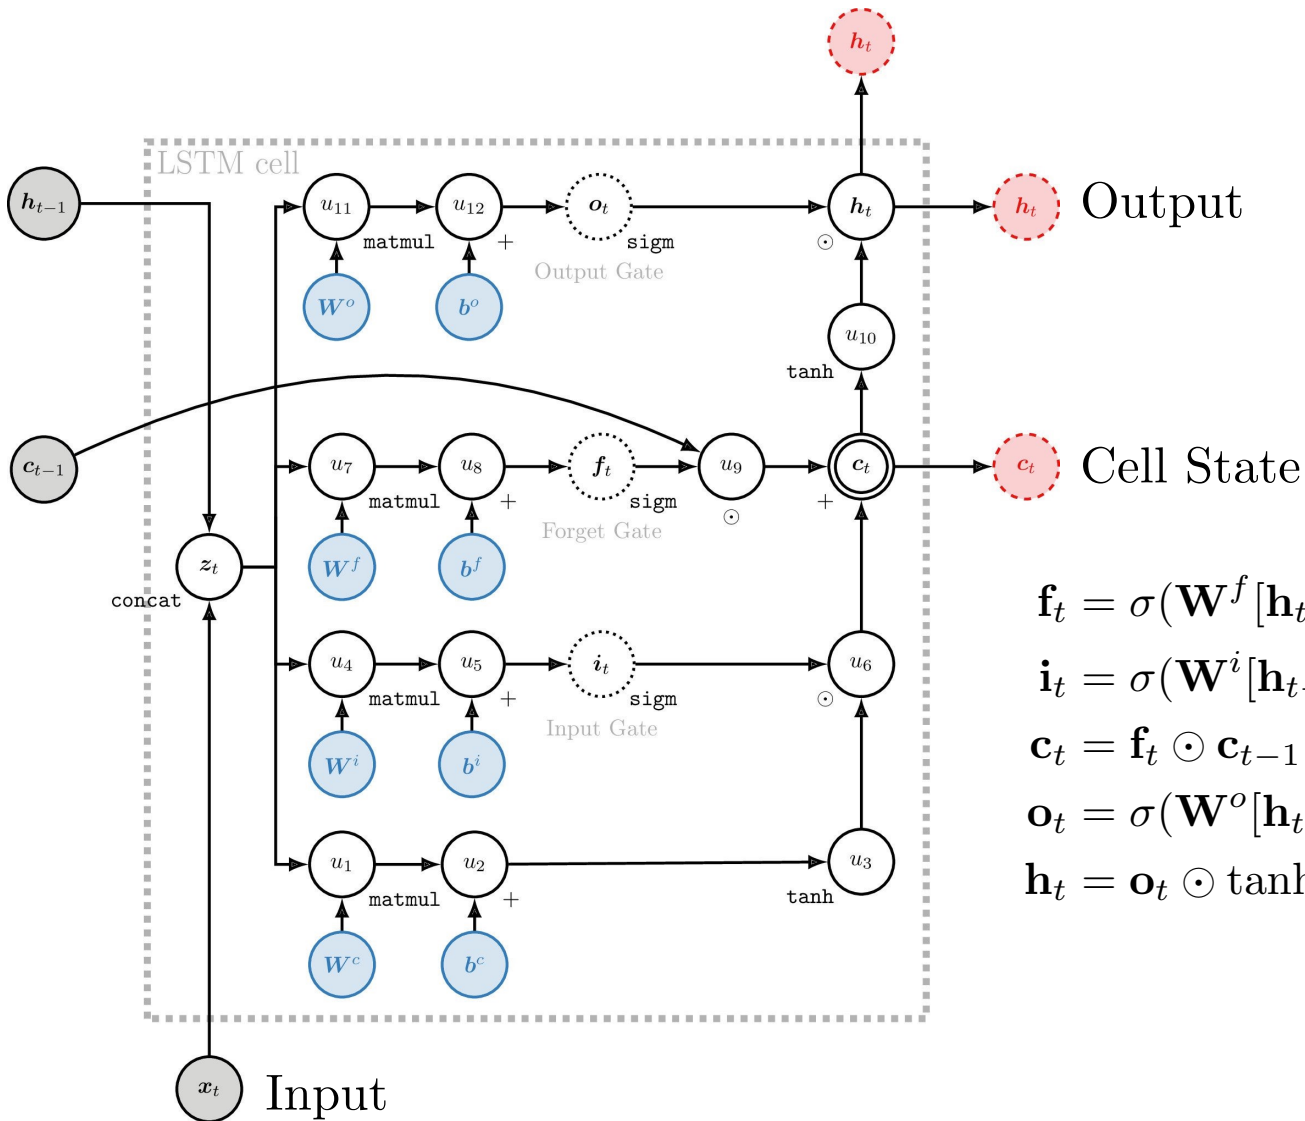- Mainly affects long-term gradients
- Solving this is more complex

Long Short-term Memory (LSTM)

Hochreiter and Schmidhuber (1997)

LSTM vs. Elman RNN

LSTM

$$\mathbf{f}_t = \sigma(\mathbf{W}^f[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}^f)$$

$$\mathbf{i}_t = \sigma(\mathbf{W}^i[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}^f)$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tanh(\mathbf{W}^c[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}^c)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}^o[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}^o)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t)$$
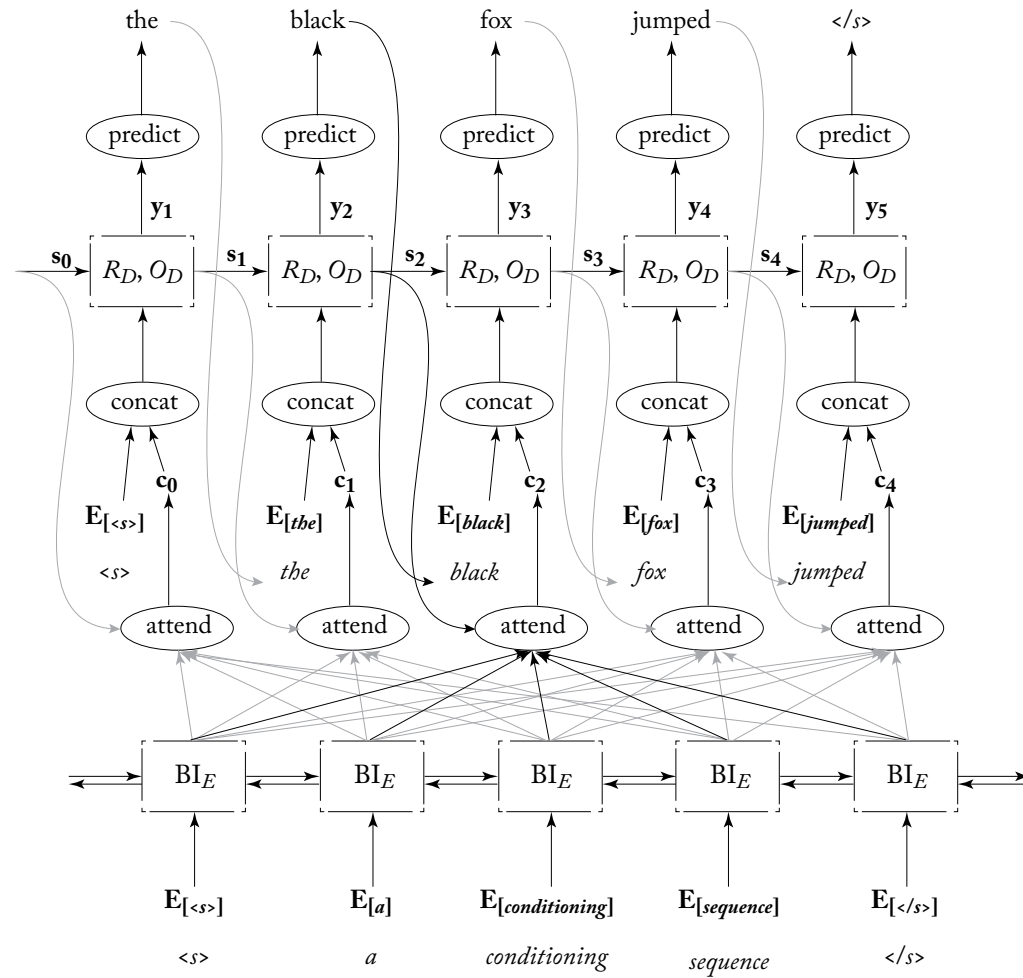
Image by Tim Rocktäschel

Attention

- In seq-to-seq models, a single vector connects encoding and decoding
  - Any concern?
  - All the input string information must be encoded into a fixed-length vector
  - The decoder must recover all this information from a fixed-length vector
- Attention relaxes the assumption that a single vector must be used to encode the input sentence regardless of length

Attention

- Encode input sentence as a sequence of vectors (already doing this)
- At each step, pick vector to use
- But: discrete choice is not differentiable
  - Make the choice soft

Attention

Attention

$$X = x_1, \ldots, x_n$$

$$\mathbf{s}_i^E = R_E(\mathbf{s}_{i-1}^E, \mathbf{E}_{[x_i]}), i = 1, \ldots, n$$

$$\bar{\mathbf{c}}_i = O_E(\mathbf{s}_i^E)$$

$$\bar{\alpha}_i^j = \mathbf{s}_{j-1}^D \cdot \bar{\mathbf{c}}_i \quad \text{Attention function}$$
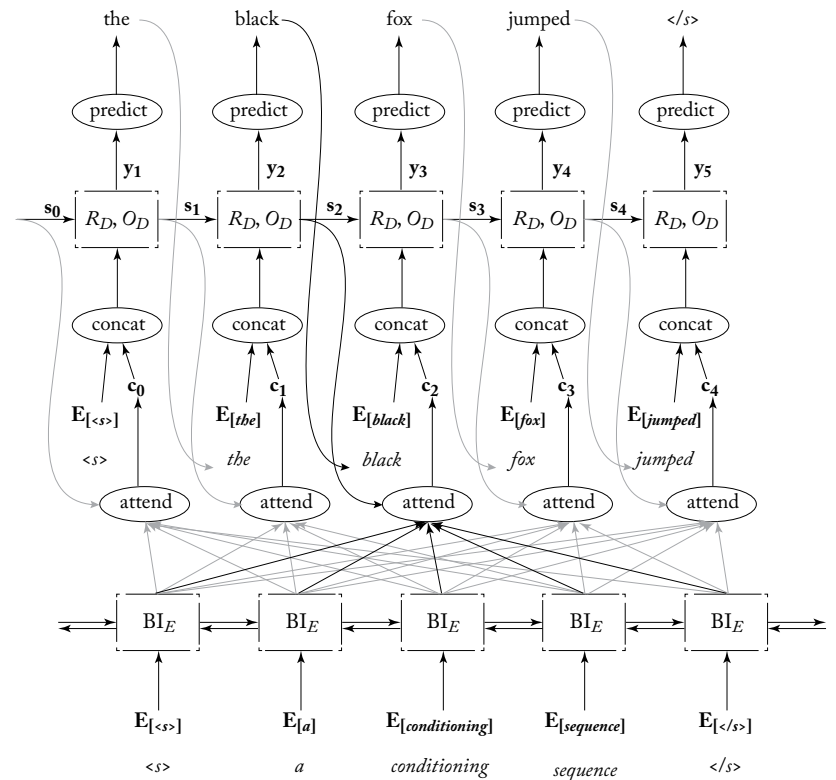
$$\alpha^j = \mathrm{softmax}(\bar{\alpha}_1^j, \ldots, \bar{\alpha}_n^j)$$

Attend

$$\mathbf{c}_j = \sum_{i=1}^n \alpha_i^j \bar{\mathbf{c}}_i$$

$$\mathbf{s}_j^D = R(\mathbf{s}_{j-1}^D, [\mathbf{E}_{[\hat{t}_{j-1}]}; \mathbf{c}_j])$$

$$O_D(\mathbf{s}_j^D) = \mathrm{softmax}(\mathbf{s}_j^D \mathbf{W} + \mathbf{b})$$

$$\hat{t}_j = \arg\max O_D(\mathbf{s}_j^D)$$

Attention

- Many variants of attention function
  - Dot product (previous slide)
  - MLP
  - Bi-linear transformation
- Various ways to combine context vector into decoder computation
- See Luong et al. 2015

Batching

- Goal: use all GPU cores
- Why is it hard?

Example 1

Example 2

Batching

- Goal: use all GPU cores
- Why is it hard?

Example 1

Example 2

- Dependencies between time steps, so can't batch the computations of an example
- Sequences have different lengths, so how to batch across examples?

Batching

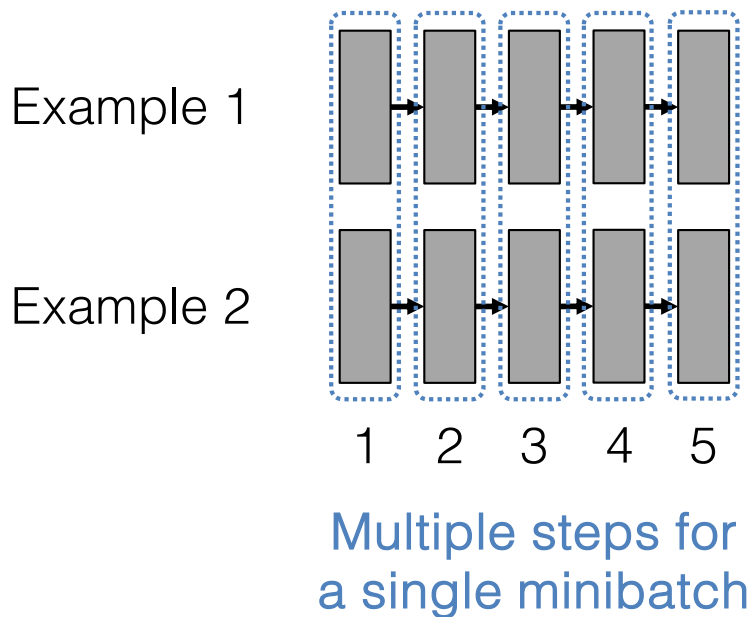- Goal: use all GPU cores
- Why is it hard?
  - Dependencies between time steps, so can't batch the computations of an example
  - Sequences have different lengths, so how to batch across examples?
- Solutions
  - Only batch examples of the same length – can get you some of the way
  - Batch examples regardless of length – get more complex

Batching

- Let's assume all examples are of the same length
- So, it's easy!

Example 1

Example 2

1   2   3   4   5

Multiple steps for a single minibatch

At each step:
- State and input are packed into tensors across examples
- Given as input to recurrence function
- Which broadcasts the computation

Batching

- What can we do if not all examples are the same length?
- Make them the same length
- Cost? Adding time steps wastes compute
- Careful: must ignore any added results

Batching

- Pad by adding input steps
- Mask to ignore any new values (e.g., when computing loss or doing predictions)
  - For example: in transducers, loss values for padding elements should be multiplied by zero
- Should still group by length to reduce cost

Example 1

Example 2

Padding