



# *Applied Deep Learning*



## Word Embeddings



March 22nd, 2021 <http://adl.miulab.tw>



**National  
Taiwan  
University**  
國立臺灣大學

2

# Review

## Word Representation

# Meaning Representations in Computers

- Knowledge-based representation
- Corpus-based representation
  - ✓ Atomic symbol
  - ✓ Neighbors
    - High-dimensional sparse word vector
    - Low-dimensional dense word vector
      - Method 1 – dimension reduction
      - Method 2 – direct learning

# Meaning Representations in Computers

- Knowledge-based representation
- **Corpus-based representation**
  - ✓ Atomic symbol
  - ✓ Neighbors
    - High-dimensional sparse word vector
    - Low-dimensional dense word vector
      - Method 1 – dimension reduction
      - Method 2 – direct learning

# 5 Corpus-based representation

- Atomic symbols: **one-hot** representation

car [0 0 0 0 0 0 1 0 0 ... 0]  
                                                          ↑  
                                                          car

Issues: difficult to compute the similarity (i.e. comparing “car” and “motorcycle”)

[0 0 0 0 0 0 1 0 0 ... 0] AND [0 0 1 0 0 0 0 0 0 ... 0] = 0  
                                                          ↑                                                          ↑  
                                                          car                                                          motorcycle

Idea: words with similar meanings often have similar neighbors

# Meaning Representations in Computers

- Knowledge-based representation
- **Corpus-based representation**
  - ✓ Atomic symbol
  - ✓ **Neighbors**
    - High-dimensional sparse word vector
    - Low-dimensional dense word vector
      - Method 1 – dimension reduction
      - Method 2 – direct learning

# Window-based Co-occurrence Matrix

## Example

- Window length=1
- Left or right context
- Corpus:

I love NTU.  
I love deep learning.  
I enjoy learning.

similarity > 0

Counts	I	love	enjoy	NTU	deep	learning
I	0	2	1	0	0	0
love	2	0	0	1	1	0
enjoy	1	0	0	0	0	1
NTU	0	1	0	0	0	0
deep	0	1	0	0	0	1
learning	0	0	1	0	1	0

### Issues:

- matrix size increases with vocabulary
- high dimensional
- sparsity → poor robustness

Idea: low dimensional word vector


# Meaning Representations in Computers

- Knowledge-based representation
- **Corpus-based representation**
  - ✓ Atomic symbol
  - ✓ **Neighbors**
    - High-dimensional sparse word vector
    - **Low-dimensional dense word vector**
      - Method 1 – dimension reduction
      - Method 2 – direct learning




# Low-Dimensional Dense Word Vector

- Method 1: dimension reduction on the matrix
- Singular Value Decomposition (SVD) of co-occurrence matrix  $X$



# Low-Dimensional Dense Word Vector

- Method 1: dimension reduction on the matrix
- Singular Value Decomposition (SVD) of co-occurrence matrix  $X$



## Issues:

- computationally expensive:  $O(mn^2)$  when  $n < m$  for  $n \times m$  matrix
- difficult to add new words

Idea: directly learn low-dimensional word vectors

# Word Representation

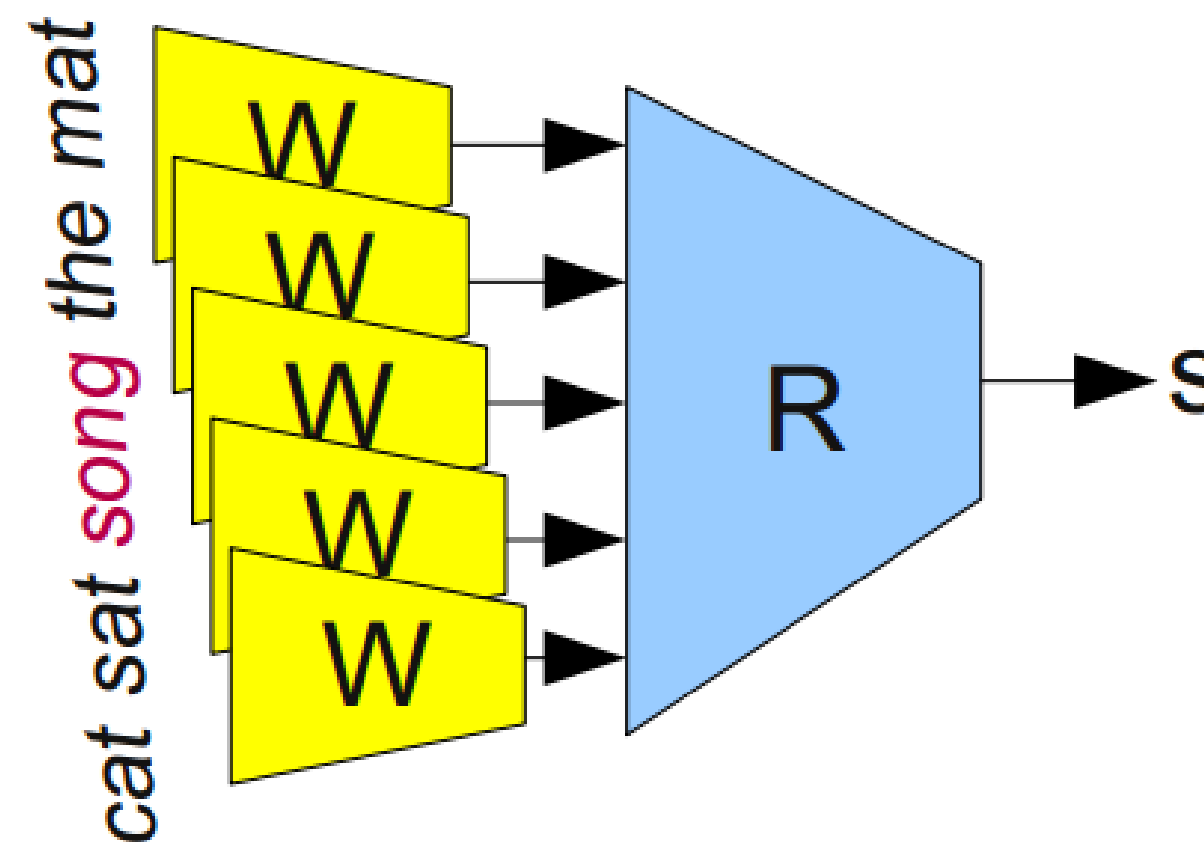
- Knowledge-based representation
- Corpus-based representation
  - ✓ Atomic symbol
  - ✓ Neighbors
    - High-dimensional sparse word vector
    - Low-dimensional dense word vector
      - Method 1 – dimension reduction
      - Method 2 – direct learning → word embedding

# Word Embedding

- Method 2: directly learn low-dimensional word vectors
  - Learning representations by back-propagation. (Rumelhart et al., 1986)
  - A neural probabilistic language model (Bengio et al., 2003)
  - NLP (almost) from Scratch (Collobert & Weston, 2008)
  - Recent and most popular models: **word2vec** (Mikolov et al. 2013) and **Glove** (Pennington et al., 2014)

# Word Embedding Benefit

- Given an unlabeled training corpus, produce a vector for each word that encodes its semantic information. These vectors are useful because:
  - semantic similarity between two words can be calculated as the cosine similarity between their corresponding word vectors
  - word vectors as powerful features for various supervised NLP tasks since the vectors contain semantic information
  - propagate any information into them via neural networks and update during training



14

# Word Embeddings

## **Word2Vec**

# Word2Vec – Skip-Gram Model

- Goal: predict surrounding words within a window of each word
- Objective function: maximize the probability of any context word given the current center word

$$w_1, w_2, \dots, \underbrace{w_{t-m}, \dots, w_{t-1}, w_t, w_{t+1}, \dots, w_{t+m}}_{\text{context window}}, \dots, w_{T-1}, w_T$$

$w_I$                        $w_O$

$$p(w_{O,1}, w_{O,2}, \dots, w_{O,C} \mid w_I) = \prod_{c=1}^C p(w_{O,c} \mid w_I)$$

$$C(\theta) = - \sum_{w_I} \sum_{c=1}^C \log p(w_{O,c} \mid w_I)$$


$$p(w_O \mid w_I) = \frac{\exp(v_{w_O}'^T v_{w_I})}{\sum_j \exp(v_{w_j}'^T v_{w_I})}$$

target word vector  
outside   target word

Benefit: faster, easily incorporate a new sentence/document or add a word to vocab

# 16 Word2Vec Skip-Gram Illustration

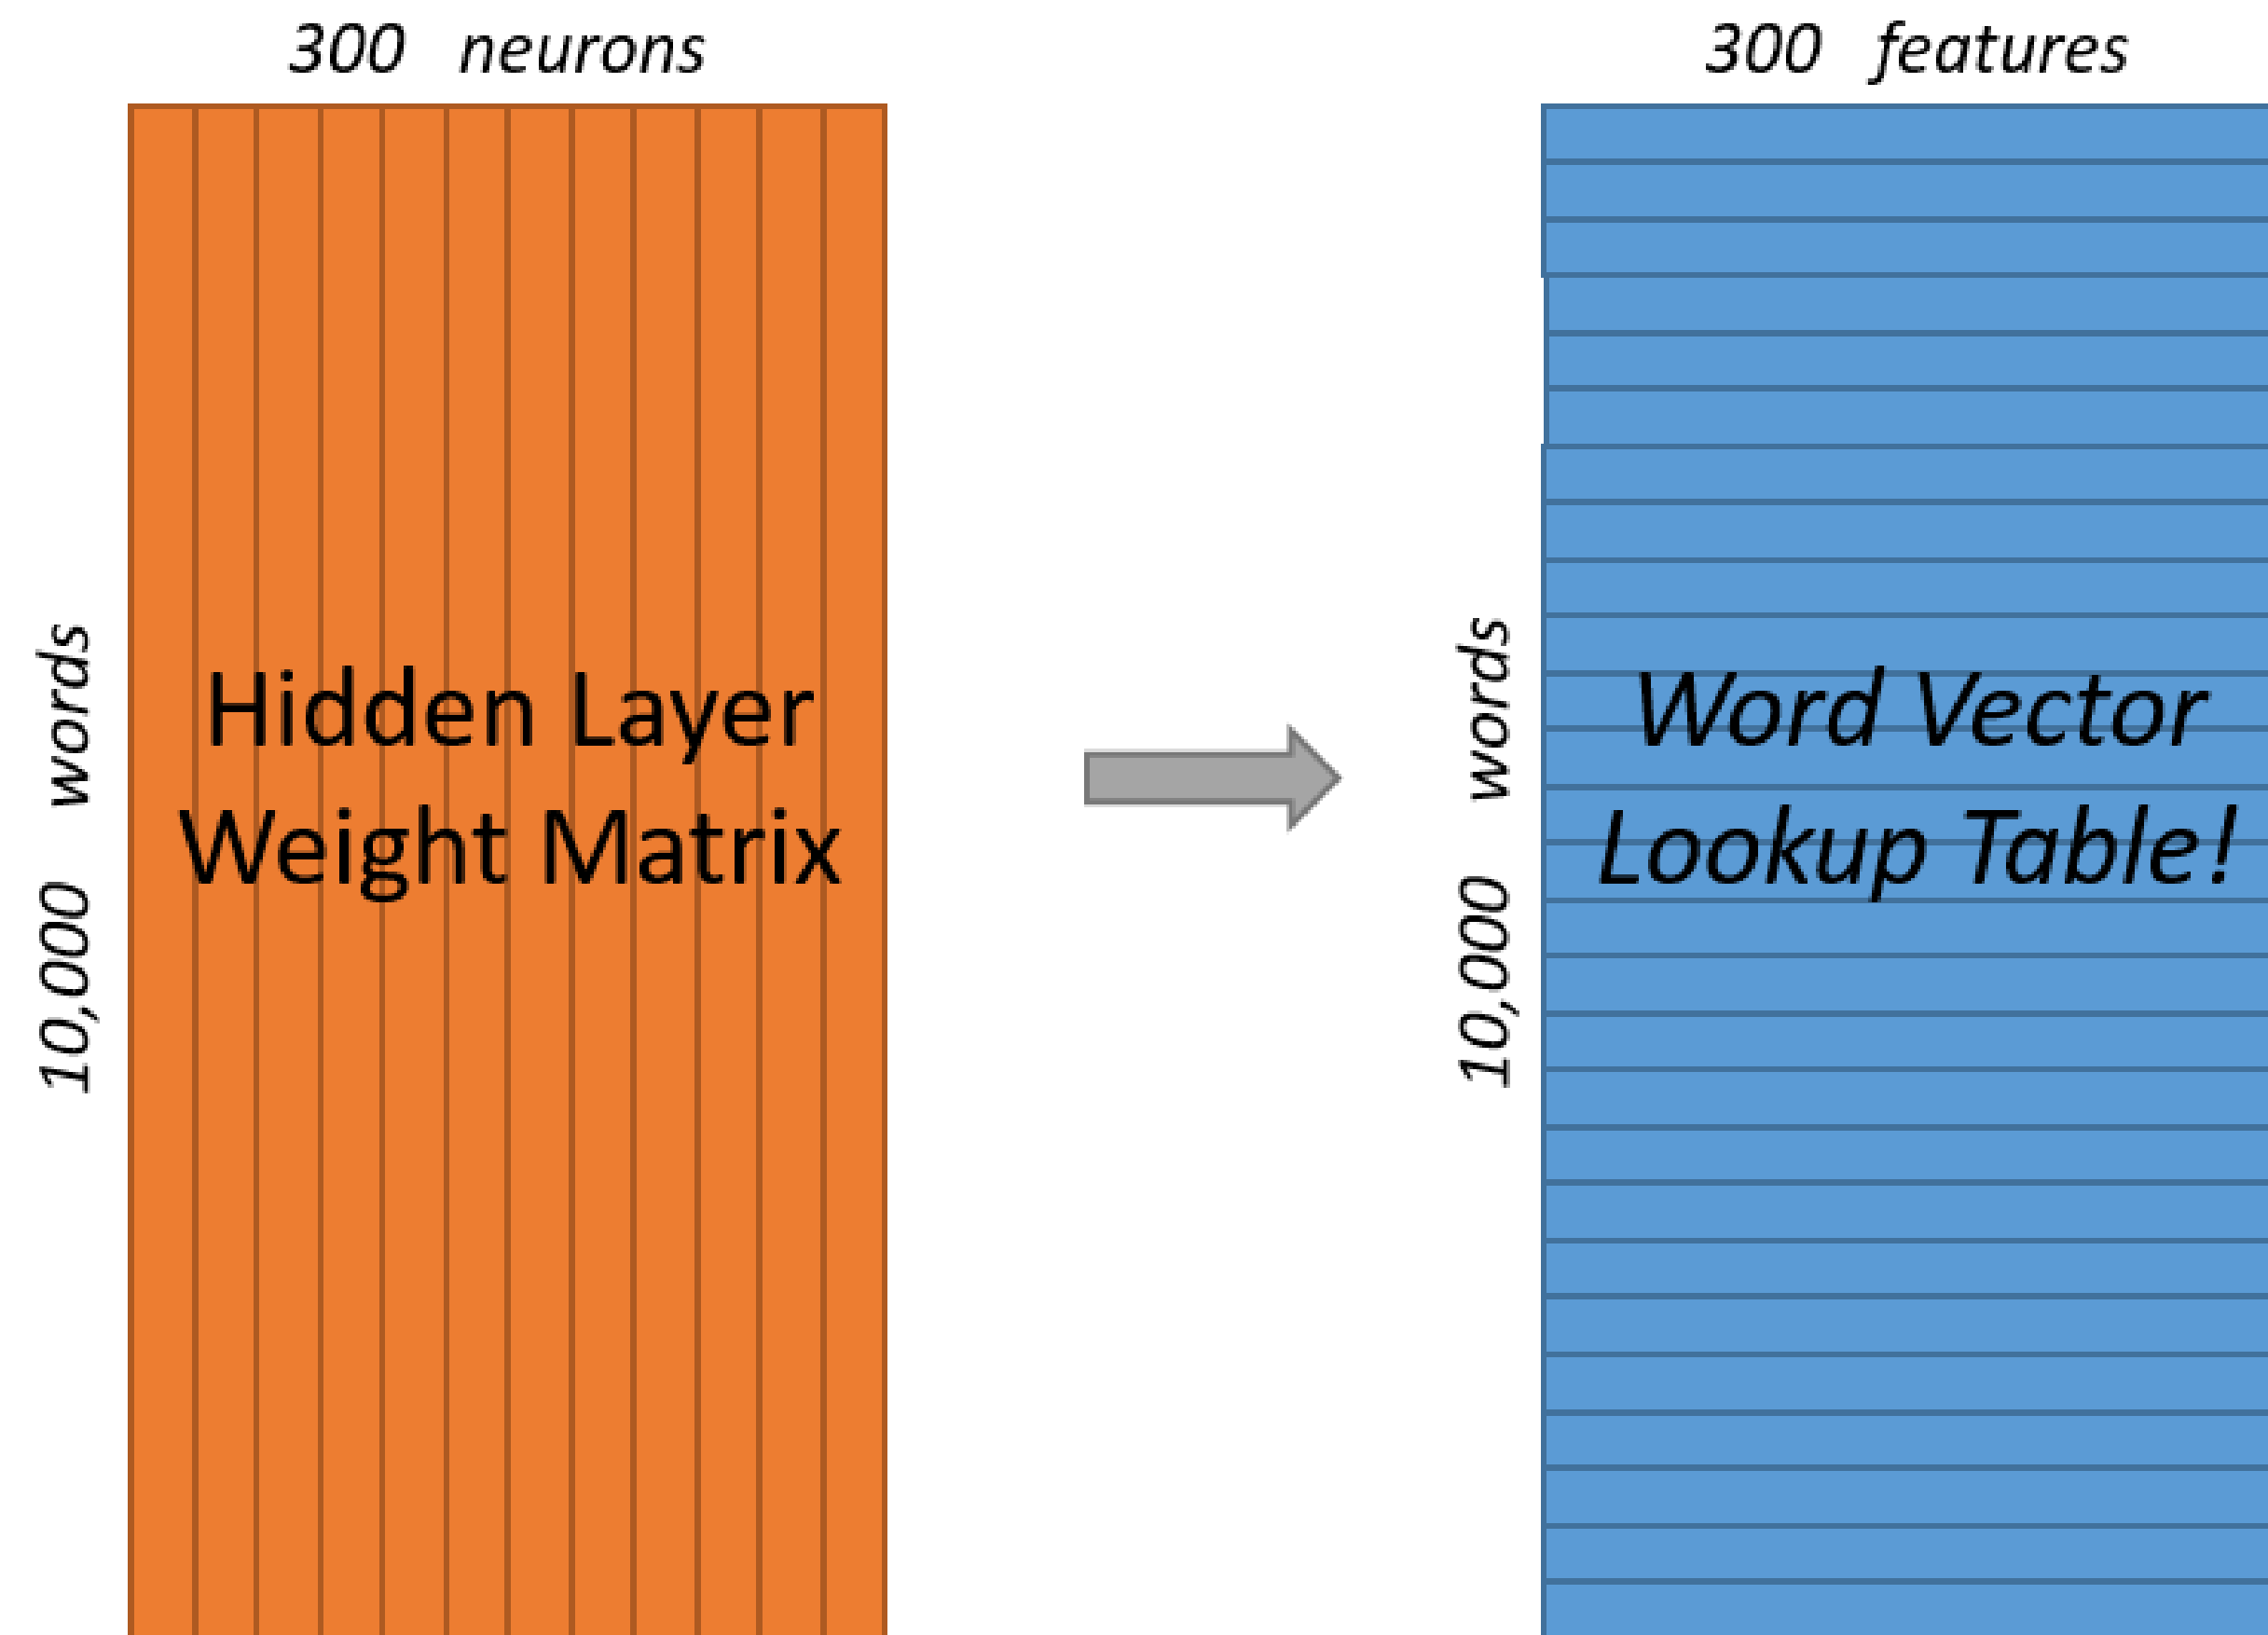
- Goal: predict surrounding words within a window of each word





# Hidden Layer Matrix $\rightarrow$ Word Embedding Matrix

$$W_{V \times N}$$


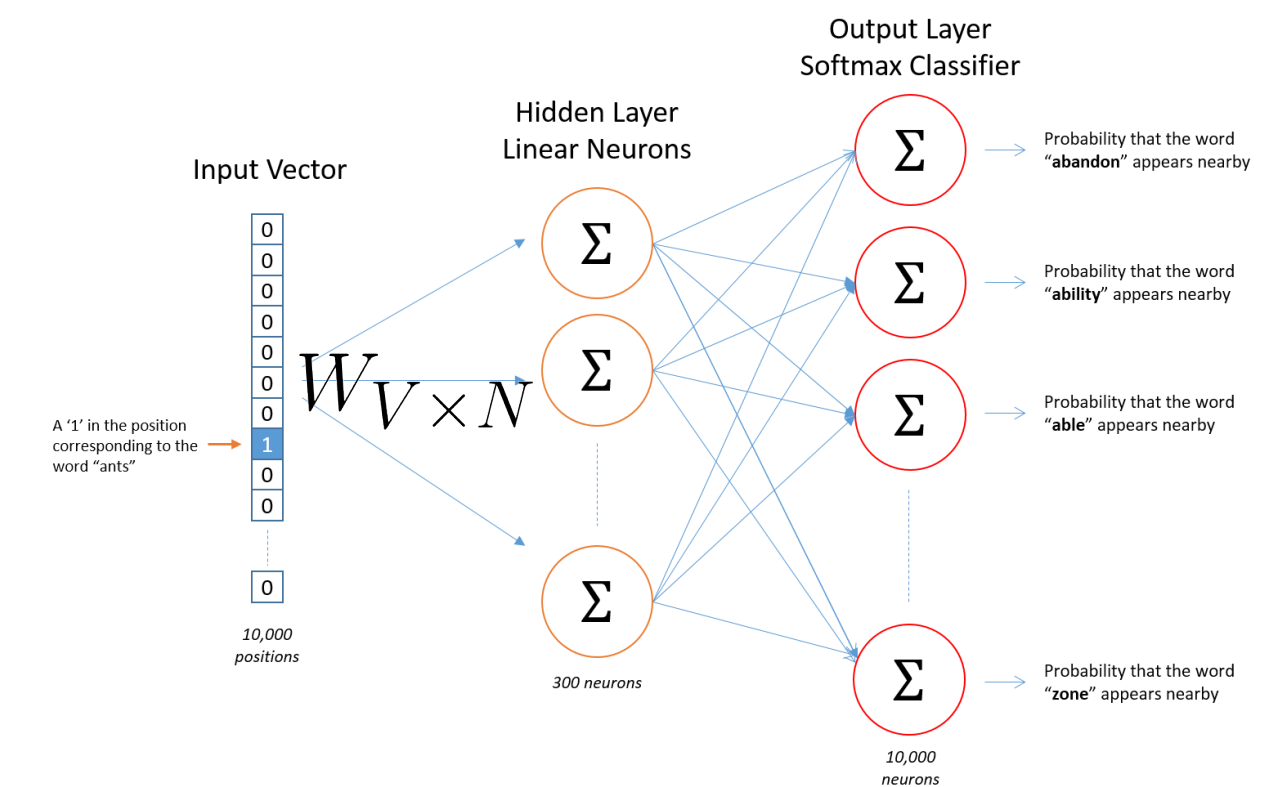


# Weight Matrix Relation

- Hidden layer weight matrix = word vector lookup

$$h = x^T W = W_{(k, \cdot)} := v w_I$$

$$[0 \ 0 \ 0 \ 1 \ 0] \times \begin{bmatrix} 17 & 24 & 1 \\ 23 & 5 & 7 \\ 4 & 6 & 13 \\ 10 & 12 & 19 \\ 11 & 18 & 25 \end{bmatrix} = [10 \ 12 \ 19]$$




Each vocabulary entry has two vectors: as a **target** word and as a **context** word

# Weight Matrix Relation

- Output layer weight matrix = weighted sum as final score

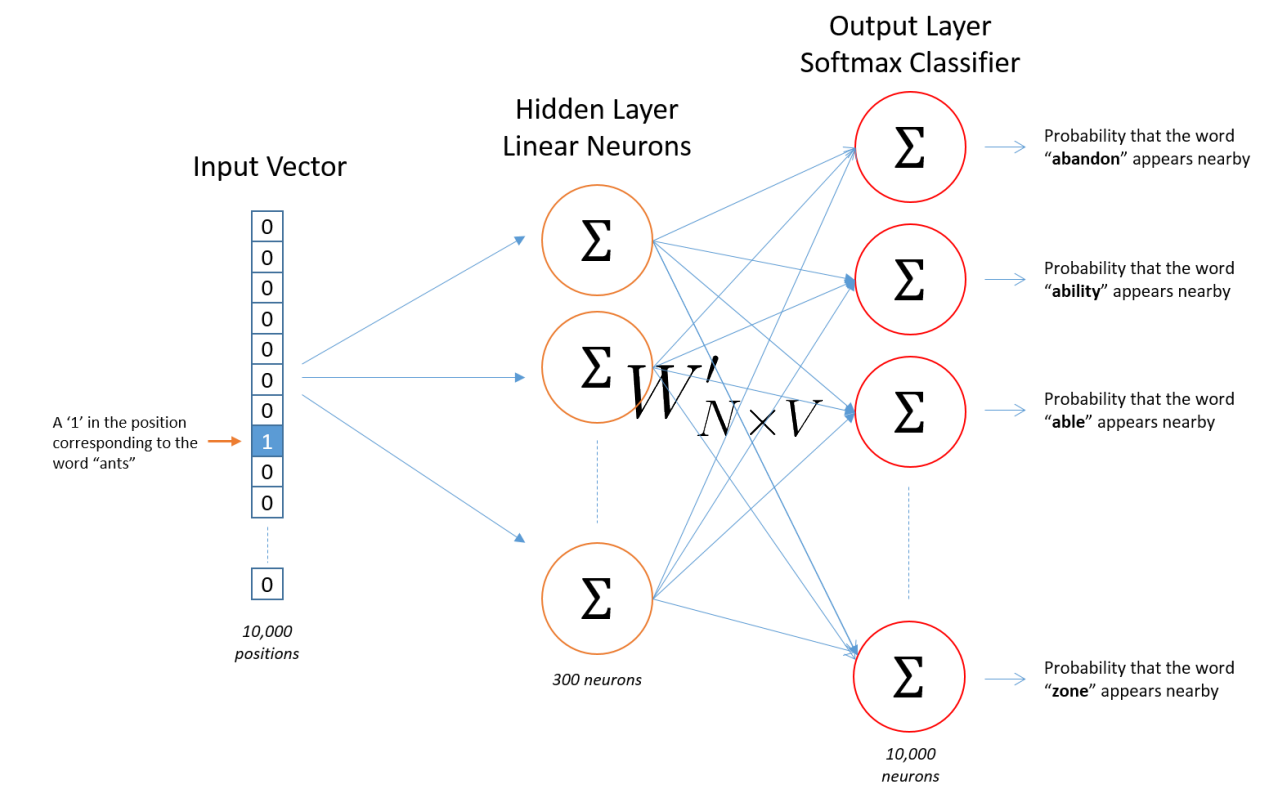
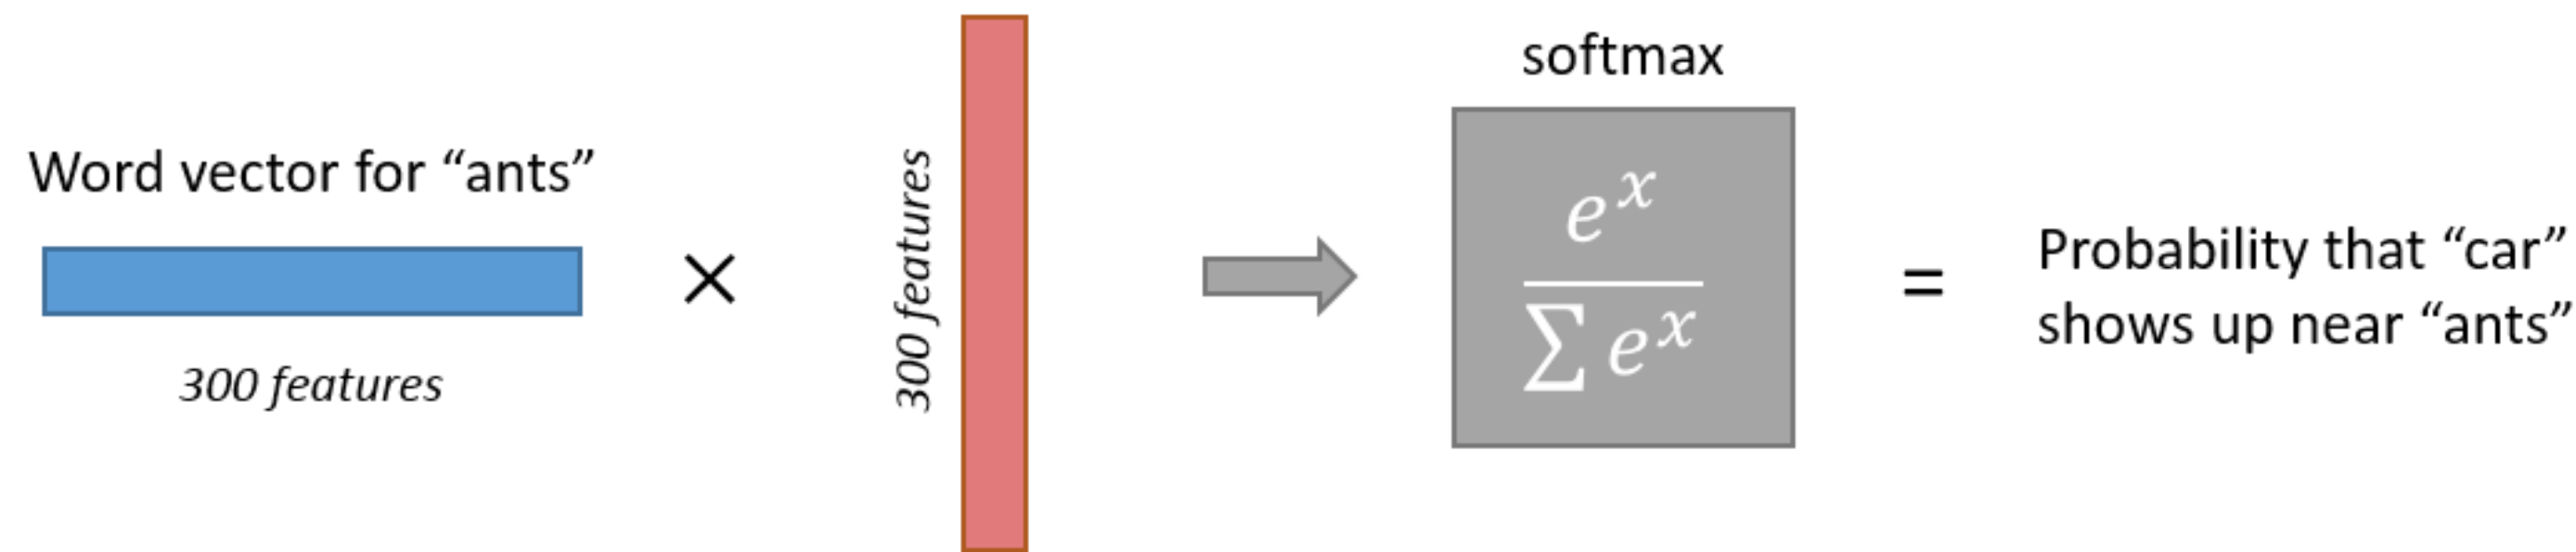
$$s_j = hv'_j$$



$$p(w_j = w_{O,c} \mid w_I) = y_{jc} = \frac{\exp(s_{jc})}{\sum_{j'=1}^V \exp(s_{j'})}$$


within the context window Output weights for "car"

softmax



Each vocabulary entry has two vectors: as a **target** word and as a **context** word

# Word2Vec Skip-Gram Illustration

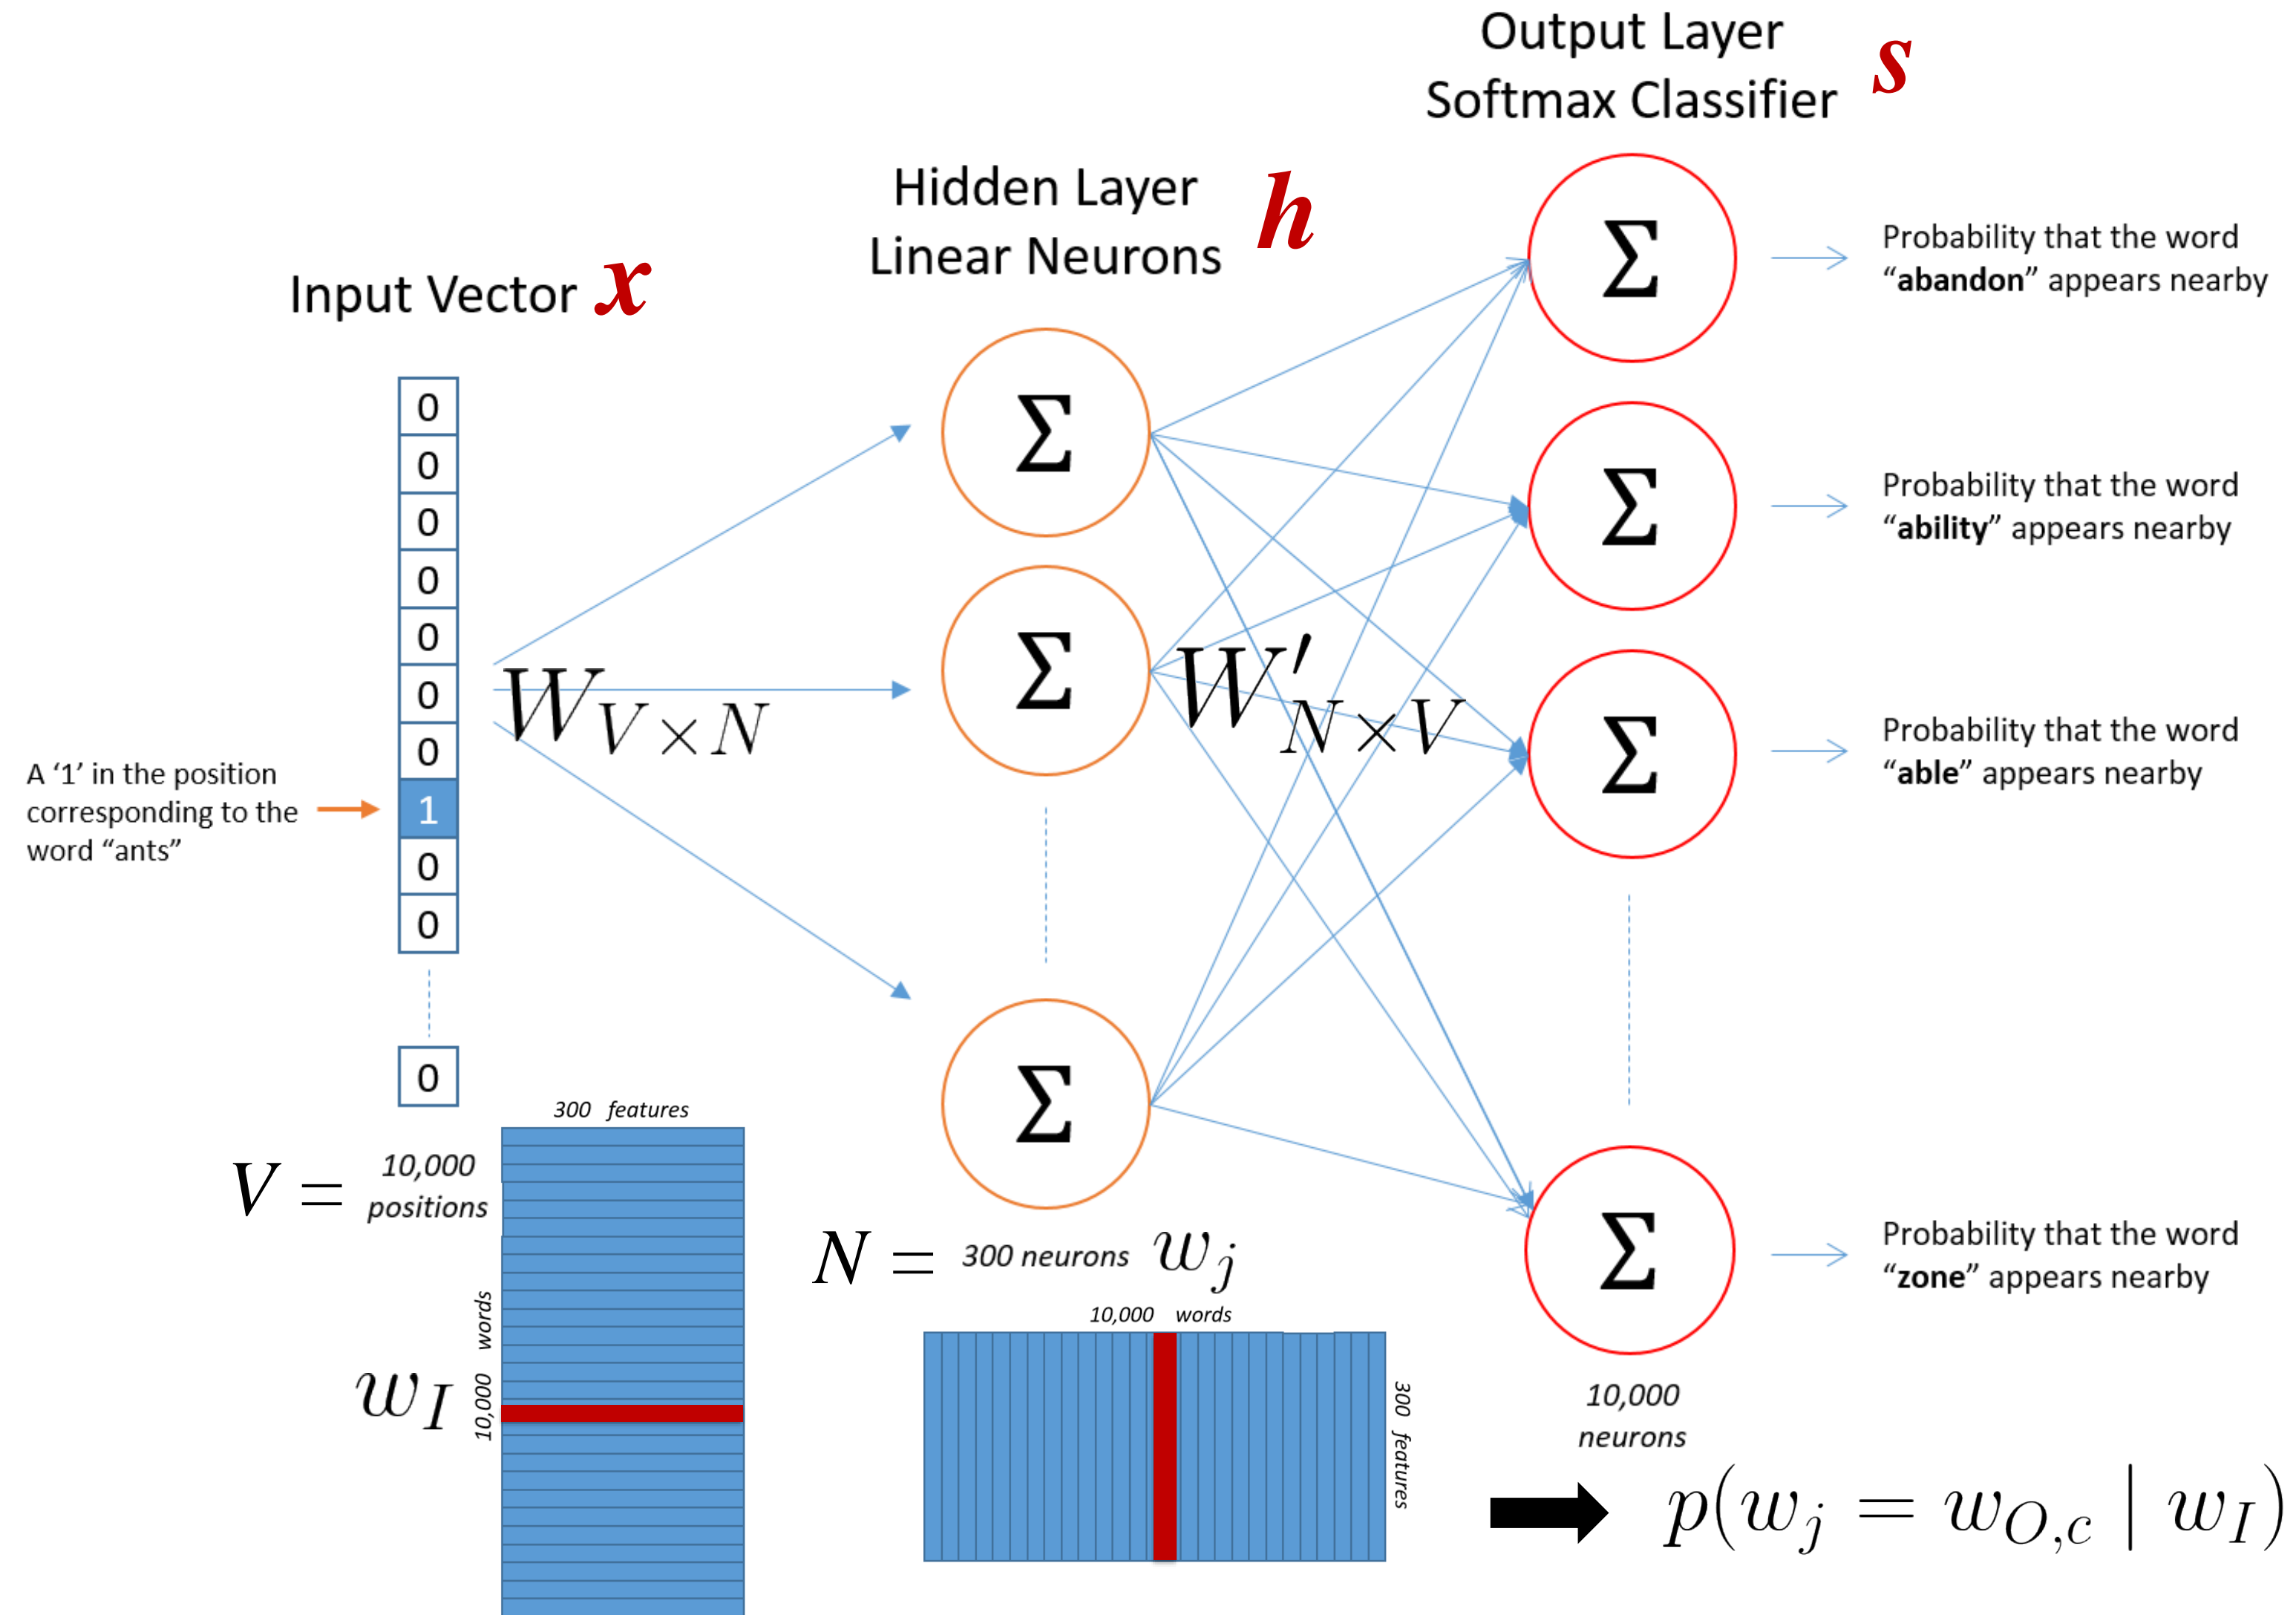


21

# Word Embeddings

## Word2Vec Training

# Word2Vec Skip-Gram Illustration



# Loss Function

- Given a target word ( $w_I$ )

$$\begin{aligned} C(\theta) &= -\log p(w_{O,1}, w_{O,2}, \dots, w_{O,C} \mid w_I) \\ &= -\log \prod_{c=1}^C \frac{\exp(s_{j_c})}{\sum_{j'=1}^V \exp(s_{j'})} \\ &= -\sum_{c=1}^C s_{j_c} + C \log \sum_{j'=1}^V \exp(s_{j'}) \end{aligned}$$

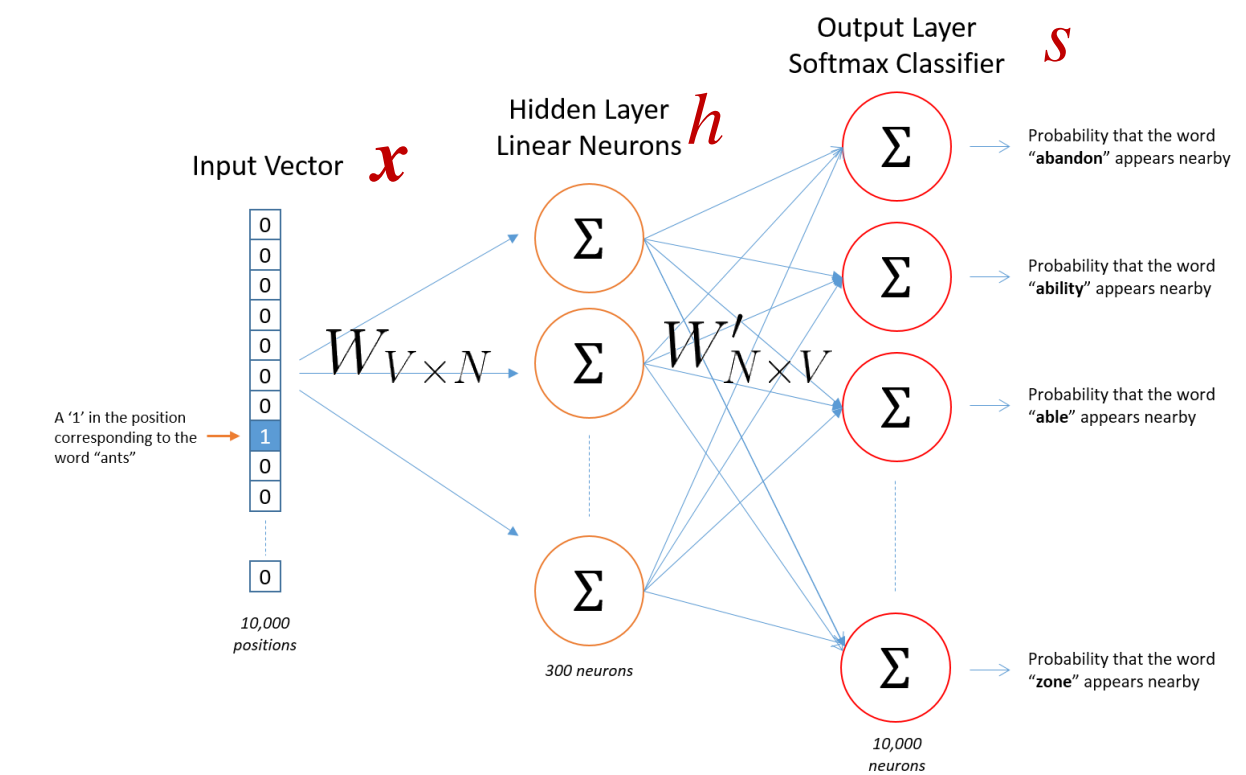
# SGD Update for $W'$

- Given a target word ( $w_I$ )

$$\frac{\partial C(\theta)}{\partial w'_{ij}} = \sum_{c=1}^C \frac{\partial C(\theta)}{\partial s_{jc}} \frac{\partial s_{jc}}{\partial w'_{ij}} = \sum_{c=1}^C (y_{jc} - t_{jc}) \cdot h_i$$

$$\frac{\partial C(\theta)}{\partial s_{jc}} = y_{jc} - \underbrace{t_{jc}}_{=1, \text{ when } w_{jc} \text{ is within the context window}} := \underbrace{e_{jc}}_{=0, \text{ otherwise}} \text{ error term}$$

$$s_j = v'_{w_j}{}^T \cdot h$$



$$w'_{ij}{}^{(t+1)} = w'_{ij}{}^{(t)} - \eta \cdot \sum_{c=1}^C (y_{jc} - t_{jc}) \cdot h_i$$



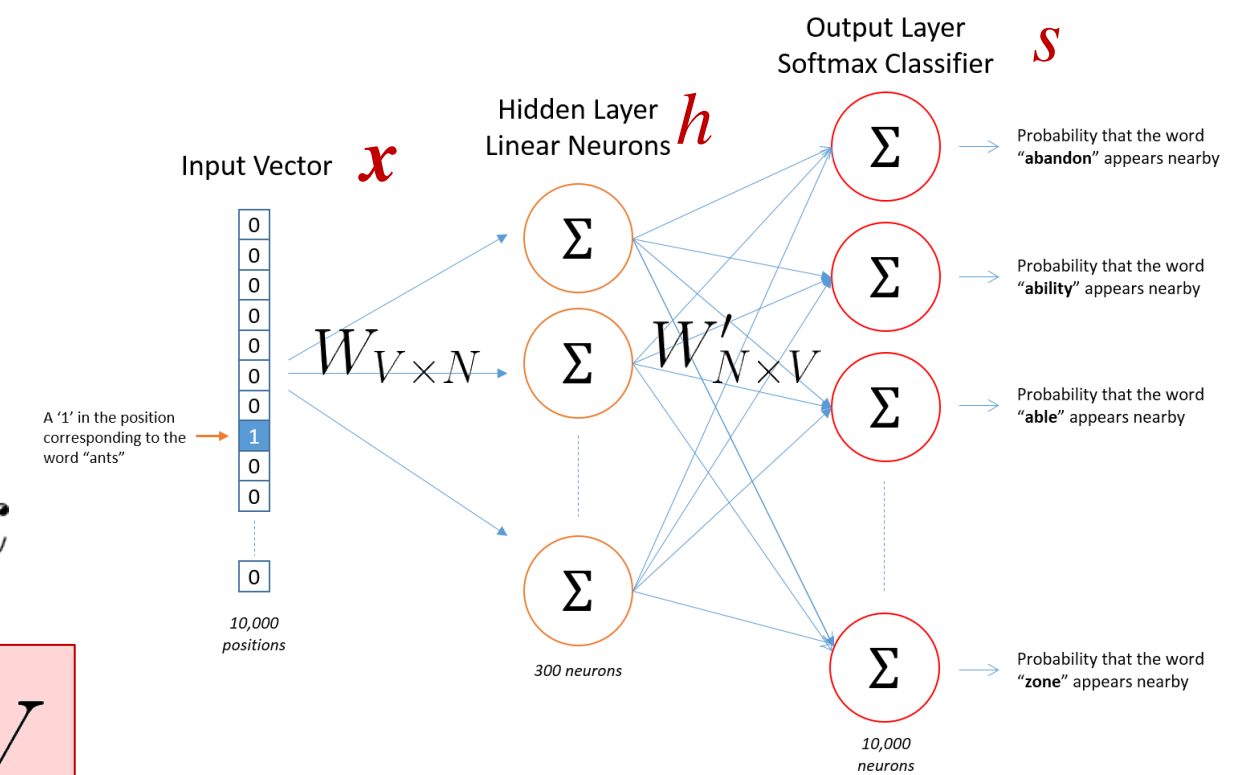
# SGD Update for $W$

$$\frac{\partial C(\theta)}{\partial w_{ki}} = \frac{\partial C(\theta)}{\partial h_i} \frac{\partial h_i}{\partial w_{ki}} = \sum_{j=1}^V \sum_{c=1}^C (y_{jc} - t_{jc}) \cdot w'_{ij} \cdot x_k$$

$$h = x^T W$$

$$\frac{\partial C(\theta)}{\partial h_i} = \sum_{j=1}^V \frac{\partial C(\theta)}{\partial s_j} \frac{\partial s_j}{\partial h_i} = \sum_{j=1}^V \sum_{c=1}^C (y_{jc} - t_{jc}) \cdot w'_{ij}$$

$$s_j = v'_{wj}{}^T \cdot h$$



$$w_{ij}^{(t+1)} = w_{ij}^{(t)} - \eta \cdot \sum_{j=1}^V \sum_{c=1}^C (y_{jc} - t_{jc}) \cdot w'_{ij} \cdot x_j$$

# SGD Update

$$w'_{ij}{}^{(t+1)} = w'_{ij}{}^{(t)} - \eta \cdot \sum_{c=1}^C (y_{jc} - t_{jc}) \cdot h_i$$

$$EI_j = \sum_{c=1}^C (y_{jc} - t_{jc})$$

$$v'_{w_j}{}^{(t+1)} = v'_{w_j}{}^{(t)} - \eta \cdot EI_j \cdot h$$

$$w_{ij}{}^{(t+1)} = w_{ij}{}^{(t)} - \eta \cdot \sum_{j=1}^V \sum_{c=1}^C (y_{jc} - t_{jc}) \cdot w'_{ij} \cdot x_j$$

$$EH_i = \sum_{j=1}^V EI_j \cdot w'_{ij} \cdot x_j$$

$$v_{w_I}{}^{(t+1)} = v_{w_I}{}^{(t)} - \eta \cdot EH^T$$

large vocabularies or large training corpora → expensive computations

limit the number of output vectors that must be updated per training instance  
→ hierarchical softmax, sampling


27

# Word Embeddings

## Negative Sampling

# Hierarchical Softmax

- Idea: compute the probability of leaf nodes using the paths



$$O(N) \rightarrow O(\log N)$$

# Negative Sampling

- Idea: only update a sample of output vectors

$$C(\theta) = -\log \sigma(v'_{w_O}{}^T v_{w_I}) + \sum_{w_j \in \mathcal{W}_{\text{neg}}} \log \sigma(v'_{w_j}{}^T v_{w_I})$$

$$v'_{w_j}{}^{(t+1)} = v'_{w_j}{}^{(t)} - \eta \cdot EI_j \cdot h$$

$$EI_j = \sigma(v'_{w_j}{}^T v_{w_I}) - t_j$$

$$v_{w_I}{}^{(t+1)} = v_{w_I}{}^{(t)} - \eta \cdot EH^T$$

$$EH = \sum_{w_j \in \{w_O\} \cup \mathcal{W}_{\text{neg}}} EI_j \cdot v'_{w_j}$$

$$w_j \in \{w_O\} \cup \mathcal{W}_{\text{neg}}$$

# Negative Sampling

## Sampling methods

- Random sampling  $w_j \in \{w_0\} \cup \mathcal{W}_{\text{neg}}$
- Distribution sampling:  $w_j$  is sampled from  $P(w)$  **What is a good  $P(w)$ ?**



Idea: less frequent words sampled more often

## Empirical setting: unigram model raised to the power of 3/4

Word	Probability to be sampled for “neg”
is	$0.9^{3/4} = 0.92$
constitution	$0.09^{3/4} = 0.16$
bombastic	$0.01^{3/4} = 0.032$


31

# Word Embeddings

## **Word2Vec Variants**

# Word2Vec Skip-Gram Visualization <https://ronxin.github.io/wevi/>

- Skip-gram training data:  
apple|drink^juice,orange|eat^apple,rice|drink^juice,juice|drink^milk,milk|drink^rice,water|drink^mil  
k,juice|orange^apple,juice|apple^drink,milk|rice^drink,drink|milk^water,drink|water^juice,drink|juic  
e^water





# wevi: word embedding visual inspector

Everything you need to know about this tool - Source code

Control Panel

Config:

```
{"hidden_size":5,"random_state":1,"learning_rate":0.2}
```

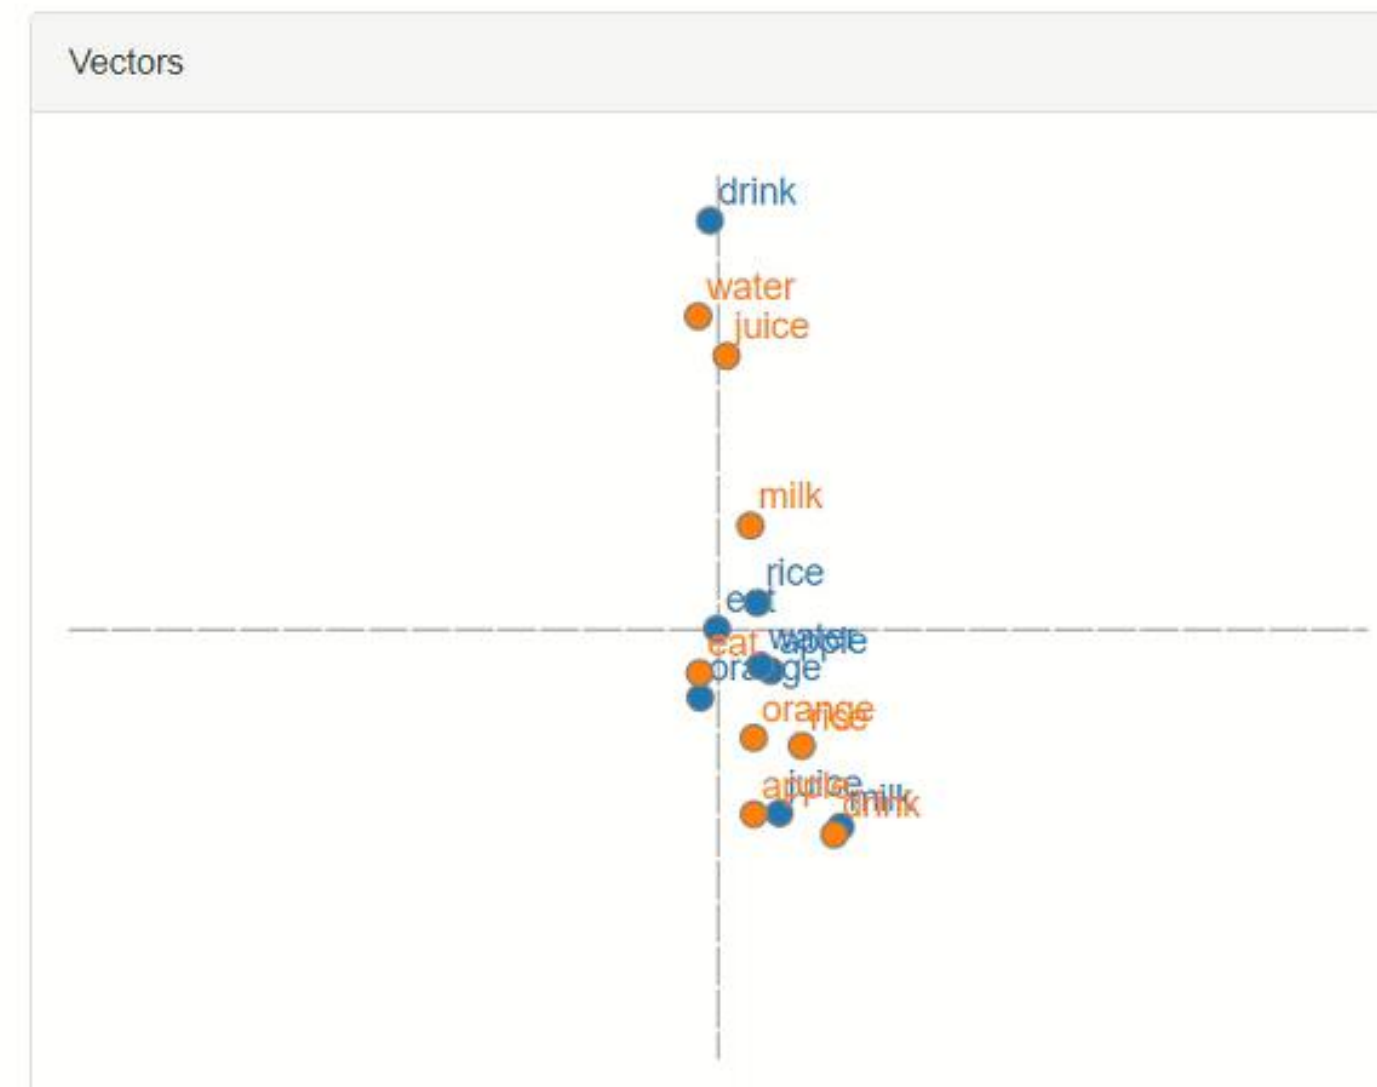
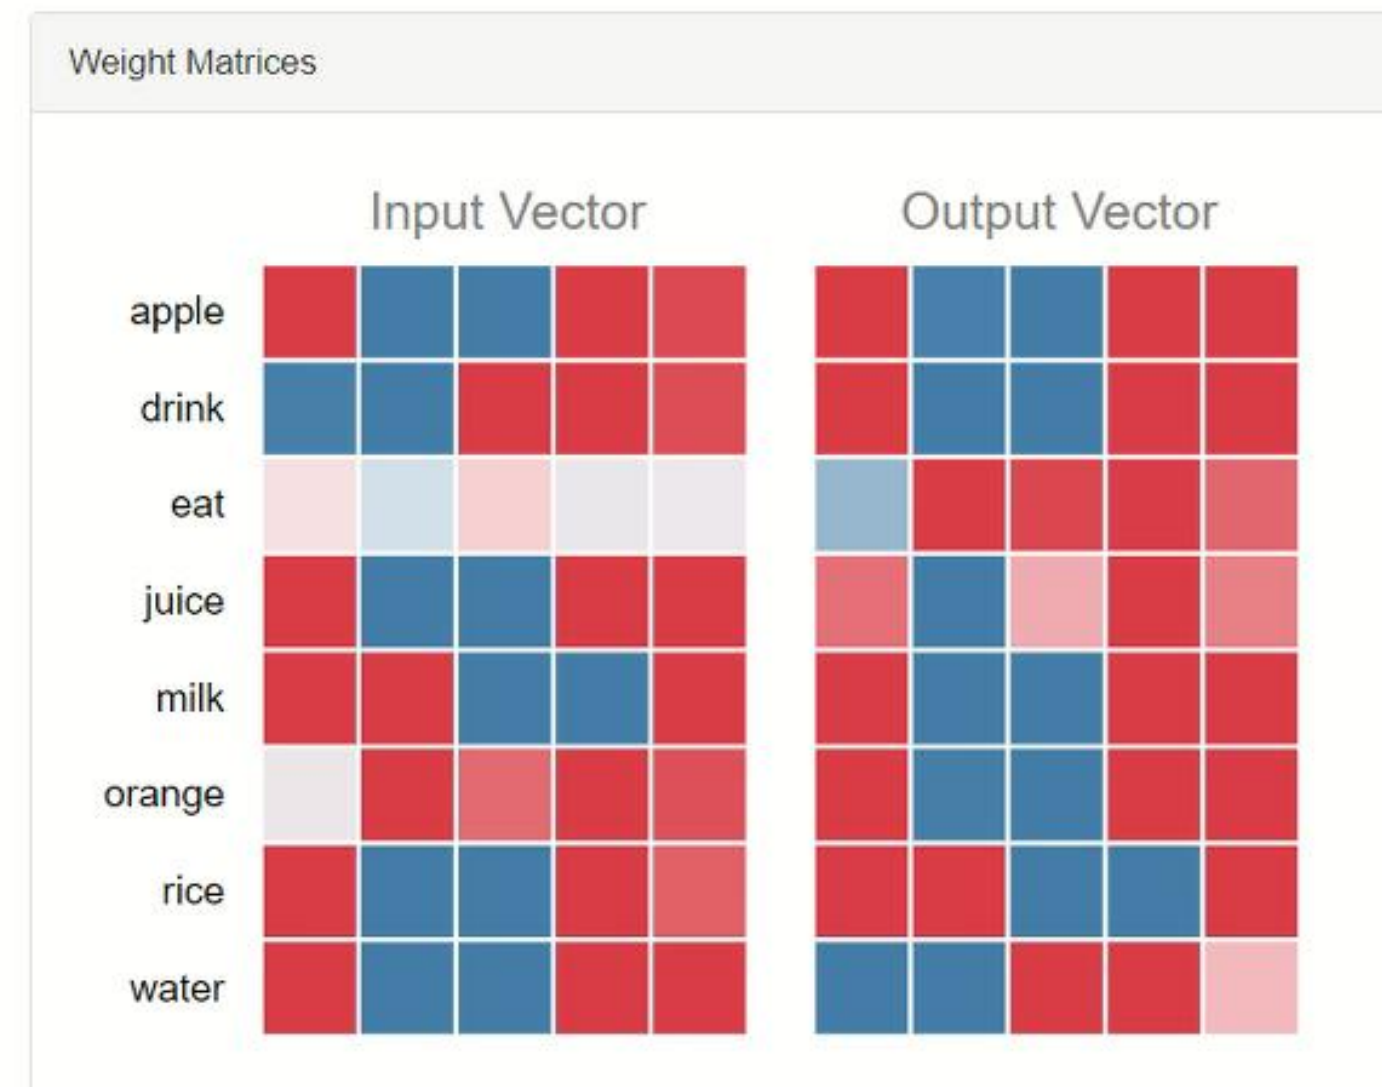
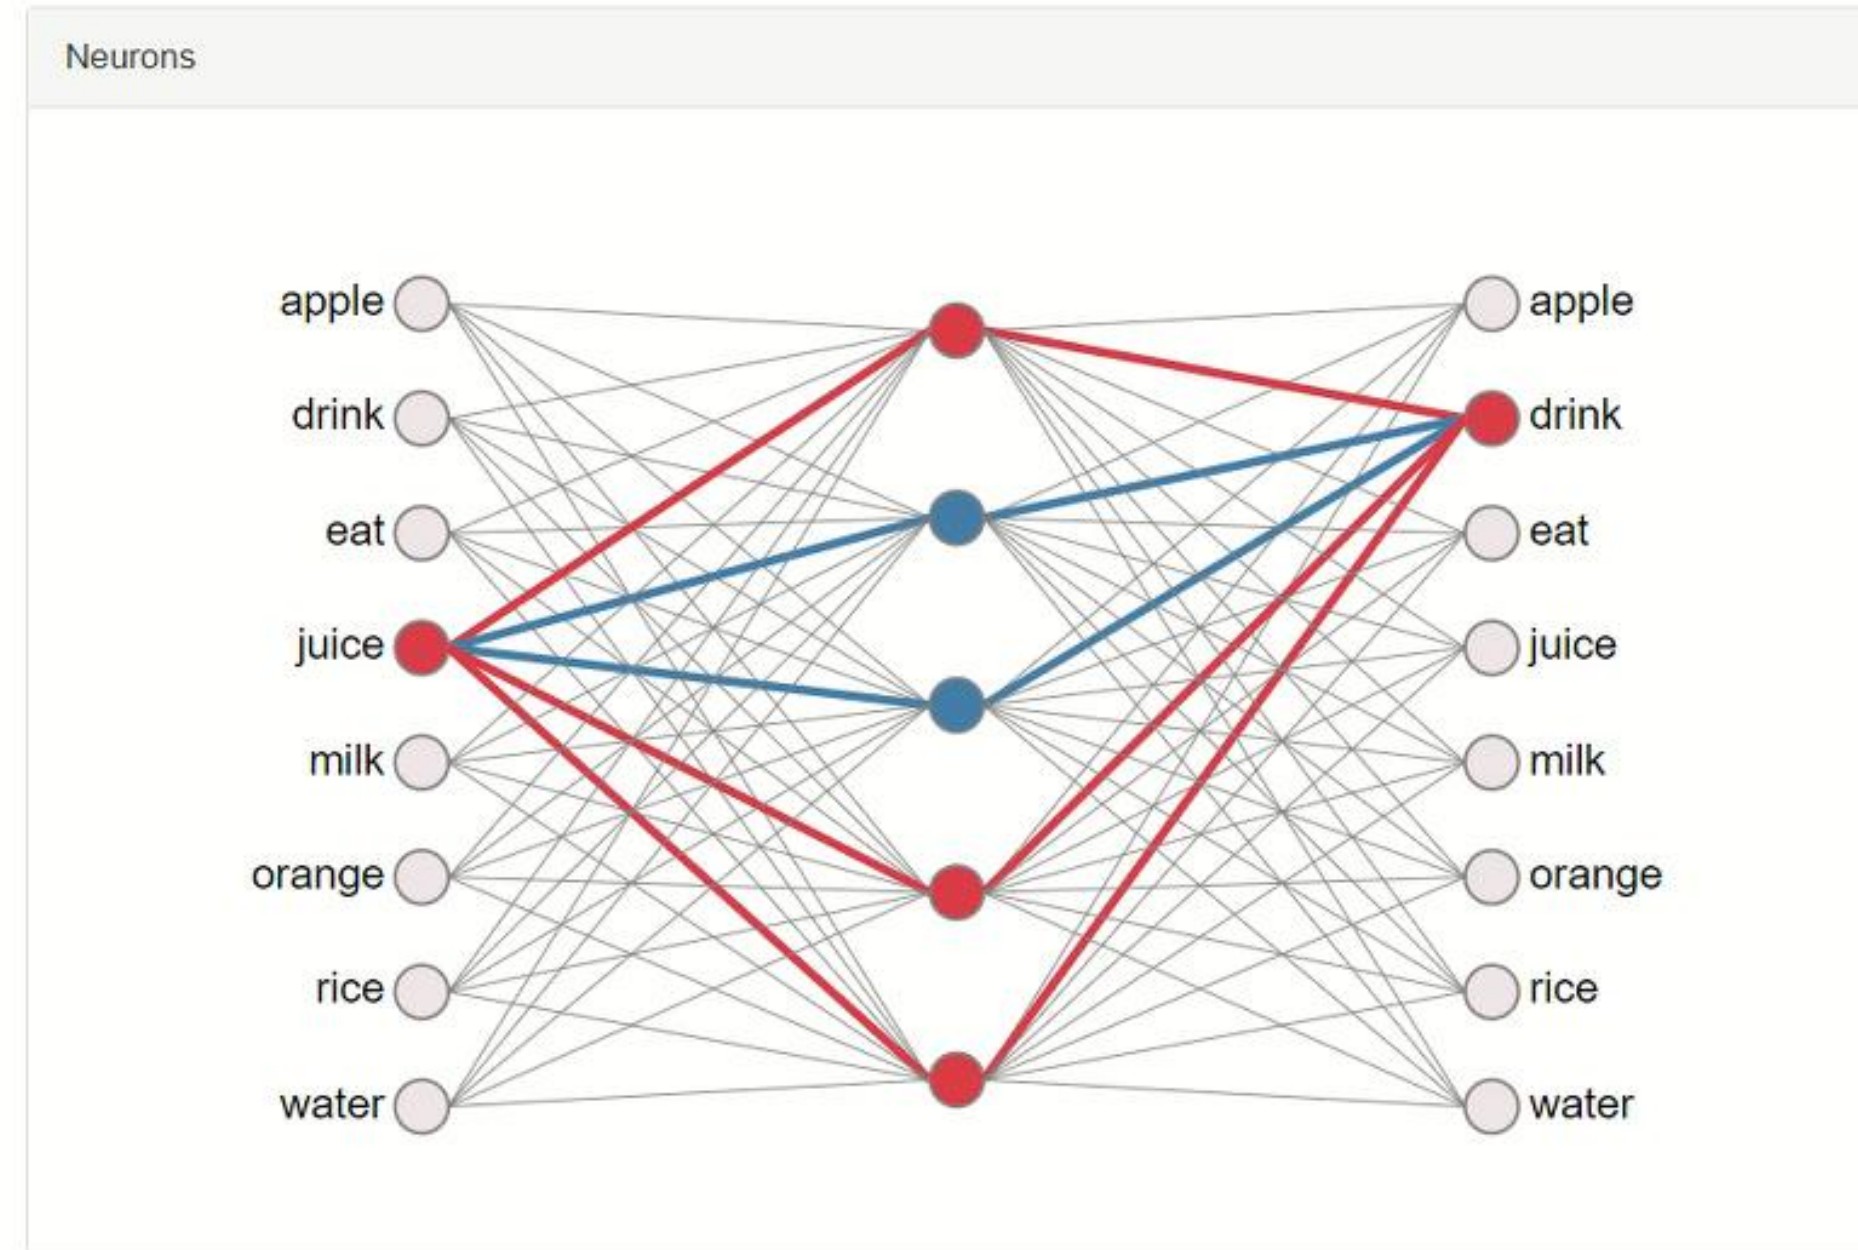
Training data (context|target):

```
apple|drink^juice,orange|eat^apple,rice|drink^juice,juice|drink^milk,milk|drink^rice,water|drink^milk,juice|orange^apple,juice|apple^drink,milk|rice^drink,drink|milk^water,drink|water^juice,drink|juice^water
```

Presets: Fruit and juice (Skip-gram)

Update and Restart    Update Learning Rate

Next   20   100   500   PCA



# Word2Vec Variants

- ◎ **Skip-gram**: predicting surrounding words given the target word (Mikolov+, 2013)

better

$$p(w_{t-m}, \dots, w_{t-1}, w_{t+1}, \dots, w_{t+m} \mid w_t)$$

- ◎ **CBOW (continuous bag-of-words)**: predicting the target word given the surrounding words (Mikolov+, 2013)

$$p(w_t \mid w_{t-m}, \dots, w_{t-1}, w_{t+1}, \dots, w_{t+m})$$

- ◎ **LM (Language modeling)**: predicting the next words given the preceding contexts (Mikolov+, 2013)

first


$$p(w_{t+1} \mid w_t)$$

Practice the derivation by yourself!!

# Word2Vec CBOW

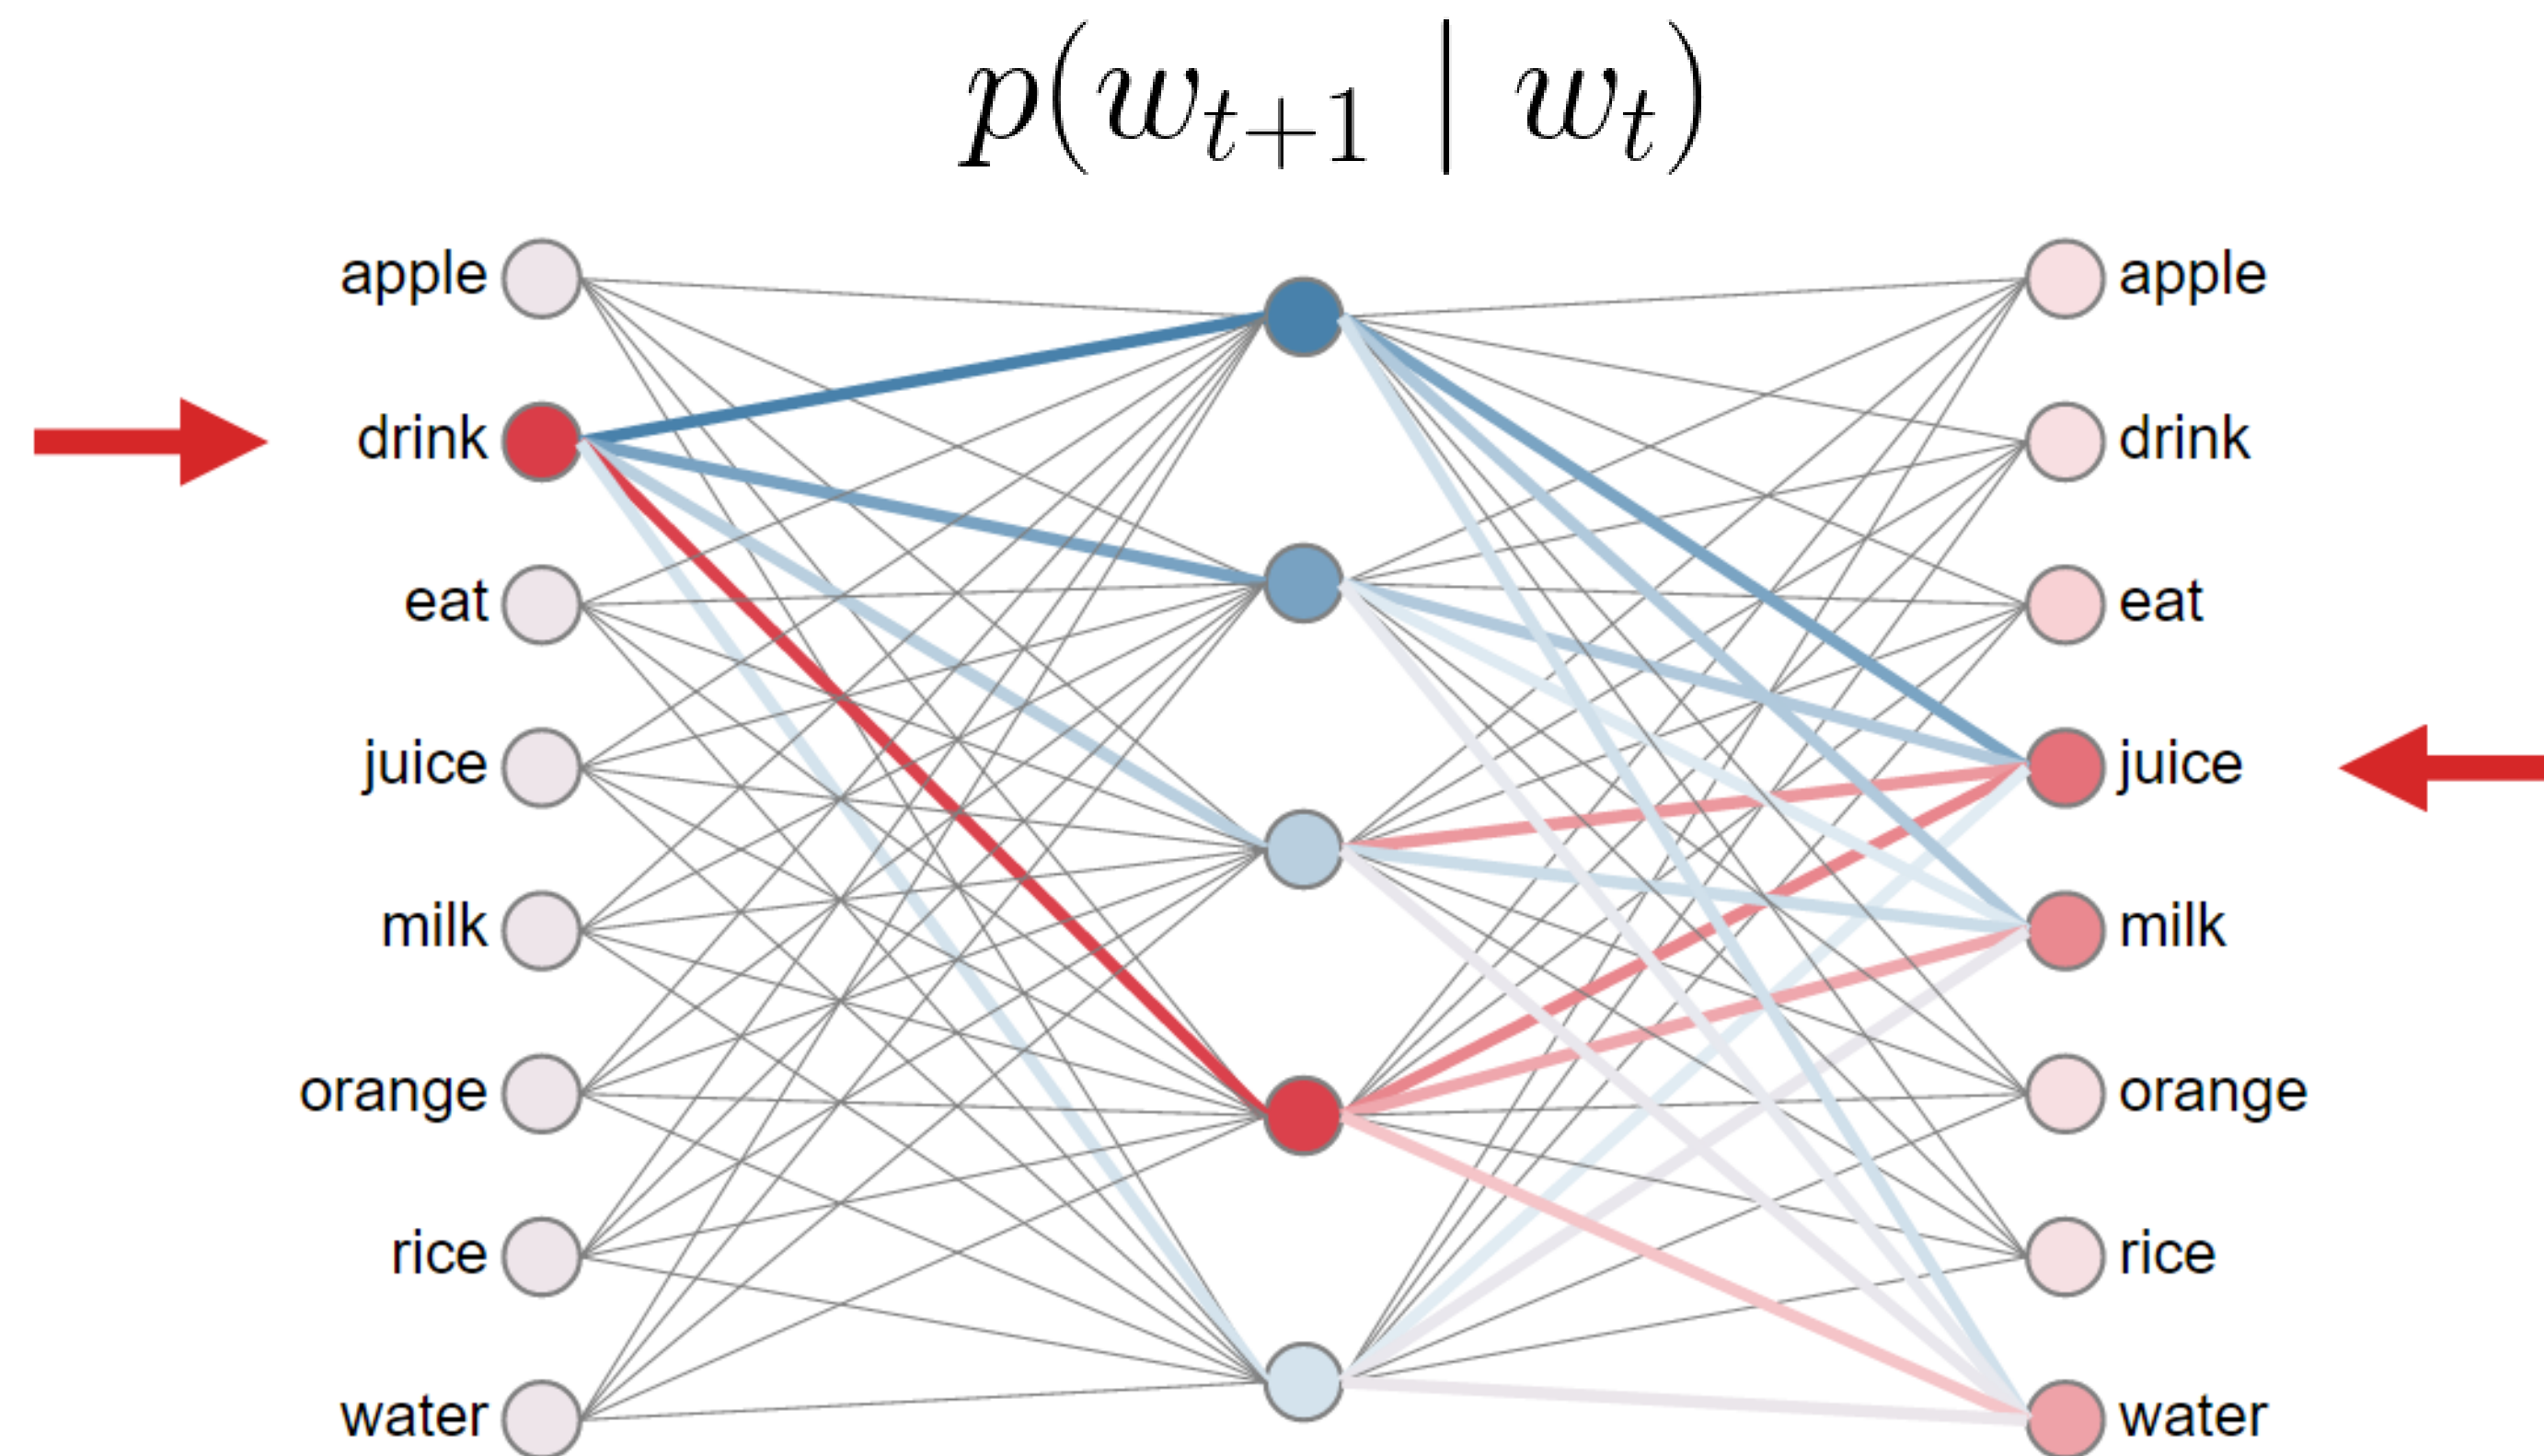
- Goal: predicting the target word given the surrounding words

$$p(w_t \mid w_{t-m}, \dots, w_{t-1}, w_{t+1}, \dots, w_{t+m})$$



# Word2Vec LM

- Goal: predicting the next words given the proceeding contexts



37

# Word Embeddings

## GloVe

# Comparison

## Count-based

- LSA, HAL (Lund & Burgess), COALS (Rohde et al), Hellinger-PCA (Lebret & Collobert)
- Pros
  - ✓ Fast training
  - ✓ Efficient usage of statistics
- Cons
  - ✓ Primarily used to capture word similarity
  - ✓ Disproportionate importance given to large counts

## Direct prediction

- NNLM, HLBL, RNN, Skipgram/CBOW (Bengio et al; Collobert & Weston; Huang et al; Mnih & Hinton; Mikolov et al; Mnih & Kavukcuoglu)
- Pros
  - ✓ Generate improved performance on other tasks
  - ✓ Capture complex patterns beyond word similarity
- Cons
  - ✓ Benefits mainly from large corpus
  - ✓ Inefficient usage of statistics

Combining the benefits from both worlds → GloVe

# GloVe

- ⊙ Idea: **ratio of co-occurrence probability** can encode meaning
- ⊙  $P_{ij}$  is the probability that word  $w_j$  appears in the context of word  $w_i$

$$P_{ij} = P(w_j | w_i) = X_{ij} / X_i$$

- ⊙ Relationship between the words  $w_i$  and  $w_j$

	$x = \text{solid}$	$x = \text{gas}$	$x = \text{water}$	$x = \text{random}$
$P(x   \text{ice})$	large	small	large	small
$P(x   \text{stream})$	small	large	large	small
$\frac{P(x   \text{ice})}{P(x   \text{stream})}$	large	small	$\sim 1$	$\sim 1$

# GloVe

- ◎ The relationship of  $w_i$  and  $w_j$  approximates the ratio of their co-occurrence probabilities with various  $w_k$

$$F(w_i, w_j, \tilde{w}_k) = \frac{P_{ik}}{P_{jk}}$$

$$F(w_i - w_j, \tilde{w}_k) = \frac{P_{ik}}{P_{jk}}$$

$$F((v_{w_i} - v_{w_j})^T v'_{\tilde{w}_k}) = \frac{P_{ik}}{P_{jk}} \quad F(\cdot) = \exp(\cdot)$$

$$v_{w_i} \cdot v'_{\tilde{w}_k} = v_{w_i}^T v'_{\tilde{w}_k} = \log P(w_k | w_i)$$



# GloVe

$$v_{w_i} \cdot v'_{\tilde{w}_j} = v_{w_i}^T v'_{\tilde{w}_j} = \log P(w_j | w_i) \quad P_{ij} = X_{ij} / X_i$$

$$= \log P_{ij} = \log(X_{ij}) - \log(X_i)$$

$$v_{w_i}^T v'_{\tilde{w}_j} + b_i + \tilde{b}_j = \log(X_{ij})$$

$$C(\theta) = \sum_{i,j=1}^V f(P_{ij}) (v_{w_i} \cdot v'_{\tilde{w}_j} - \log P_{ij})^2$$


$$C(\theta) = \sum_{i,j=1}^V f(X_{ij}) (v_{w_i}^T v'_{\tilde{w}_j} + b_i + \tilde{b}_j - \log X_{ij})^2$$

# GloVe – Weighted Least Squares Regression Model

$$C(\theta) = \sum_{i,j=1}^V f(X_{ij}) (v_{w_i}^T v'_{\tilde{w}_j} + b_i + \tilde{b}_j - \log X_{ij})^2$$

## Weighting function should obey

- $f(0) = 0$
- $f(x)$  should be non-decreasing so that *rare co-occurrences* are not overweighted
- $f(x)$  should be relatively small for large values of  $x$ , so that *frequent co-occurrences* are not overweighted



fast training, scalable, good performance even with small corpus, and small vectors

43

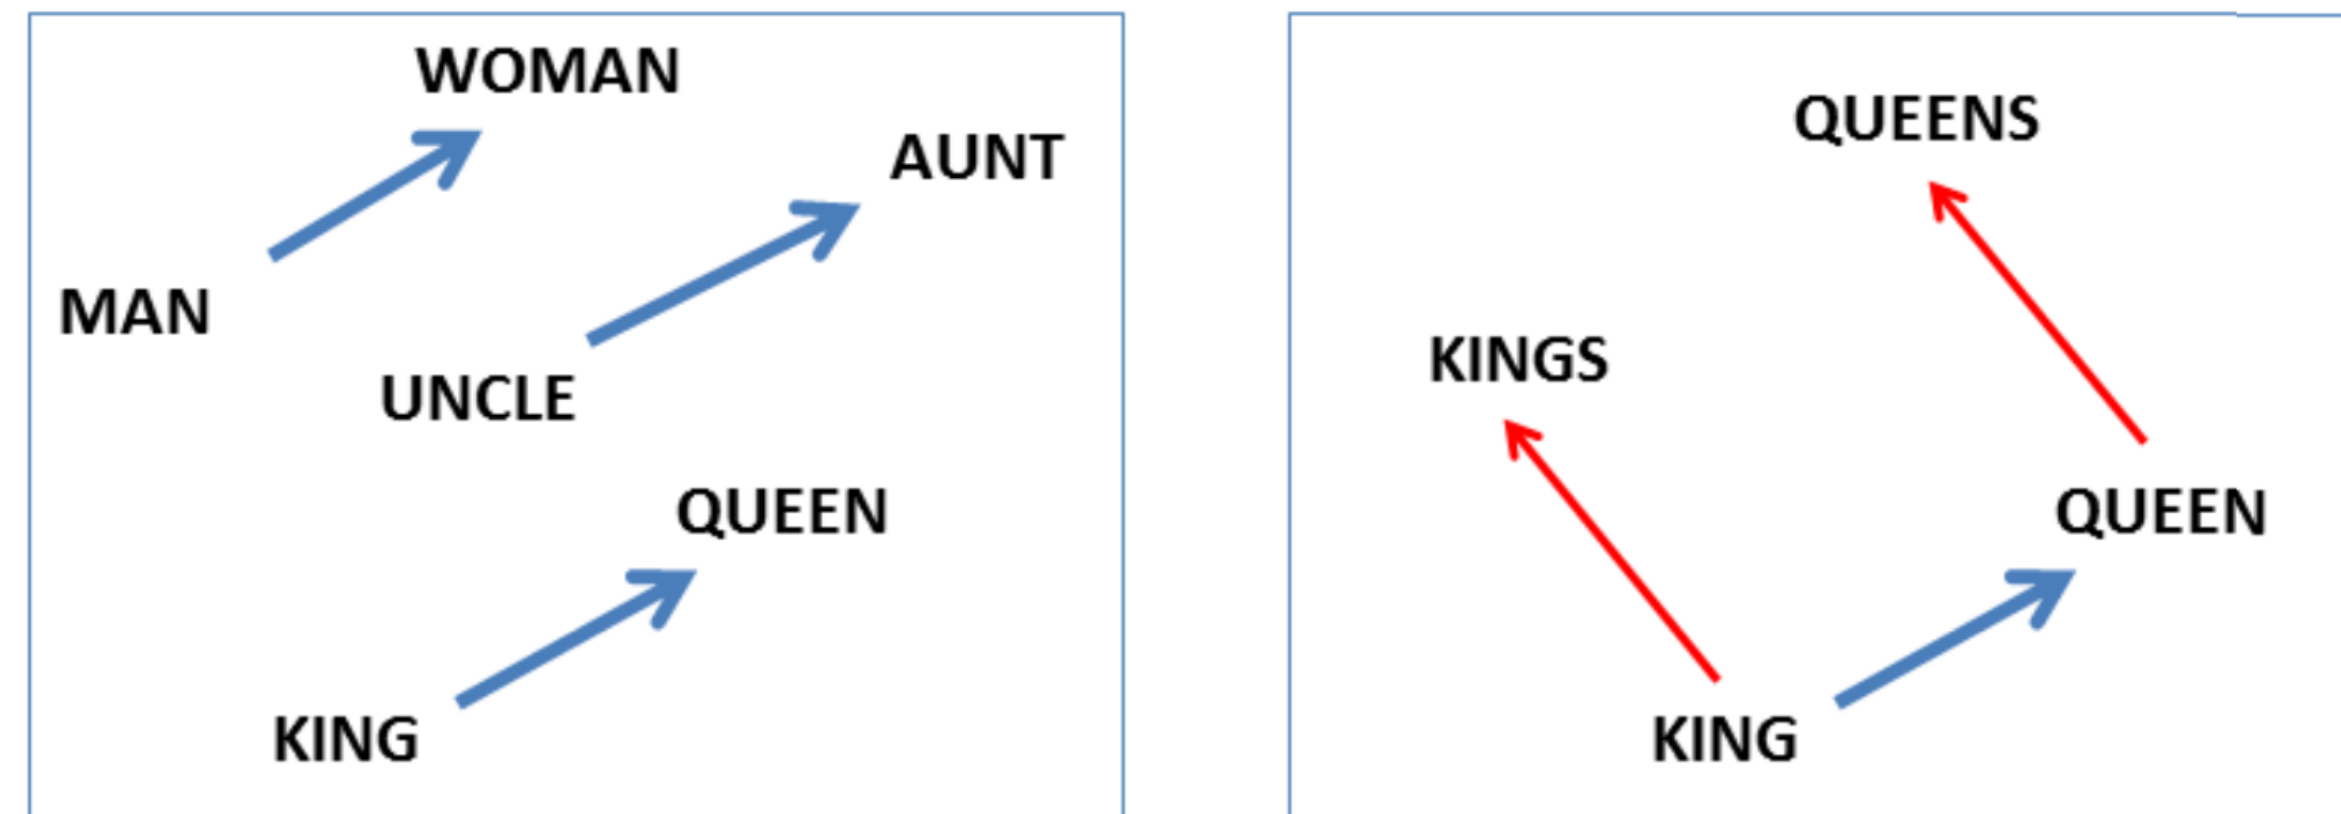
# Word Vector Evaluation

# Intrinsic Evaluation – Word Analogies

- Word linear relationship  $w_A : w_B = w_C : w_x$

$$x = \arg \max_x \frac{(v_{w_B} - v_{w_A} + v_{w_C})^T v_{w_x}}{\|v_{w_B} - v_{w_A} + v_{w_C}\|}$$

- Syntactic** and **Semantic** example questions [[link](#)]



Issue: what if the information is there but not linear

# Intrinsic Evaluation – Word Analogies

- Word linear relationship  $w_A : w_B = w_C : w_x$
- Syntactic and **Semantic** example questions [[link](#)]

## city---in---state

Chicago : Illinois = Houston : Texas  
 Chicago : Illinois = Philadelphia : Pennsylvania  
 Chicago : Illinois = Phoenix : Arizona  
 Chicago : Illinois = Dallas : Texas  
 Chicago : Illinois = Jacksonville : Florida  
 Chicago : Illinois = Indianapolis : Indiana  
 Chicago : Illinois = Austin : Texas  
 Chicago : Illinois = Detroit : Michigan  
 Chicago : Illinois = Memphis : Tennessee  
 Chicago : Illinois = Boston : Massachusetts

Issue: different cities may have same name

## capital---country

Abuja : Nigeria = Accra : Ghana  
 Abuja : Nigeria = Algiers : Algeria  
 Abuja : Nigeria = Amman : Jordan  
 Abuja : Nigeria = Ankara : Turkey  
 Abuja : Nigeria = Antananarivo : Madagascar  
 Abuja : Nigeria = Apia : Samoa  
 Abuja : Nigeria = Ashgabat : Turkmenistan  
 Abuja : Nigeria = Asmara : Eritrea  
 Abuja : Nigeria = Astana : Kazakhstan

Issue: can change with time

# Intrinsic Evaluation – Word Analogies

- Word linear relationship  $w_A : w_B = w_C : w_x$
- Syntactic** and Semantic example questions [[link](#)]

## superlative

bad : worst = big : biggest  
 bad : worst = bright : brightest  
 bad : worst = cold : coldest  
 bad : worst = cool : coolest  
 bad : worst = dark : darkest  
 bad : worst = easy : easiest  
 bad : worst = fast : fastest  
 bad : worst = good : best  
 bad : worst = great : greatest

## past tense

dancing : danced = decreasing : decreased  
 dancing : danced = describing : described  
 dancing : danced = enhancing : enhanced  
 dancing : danced = falling : fell  
 dancing : danced = feeding : fed  
 dancing : danced = flying : flew  
 dancing : danced = generating : generated  
 dancing : danced = going : went  
 dancing : danced = hiding : hid  
 dancing : danced = hiding : hit

# Intrinsic Evaluation – Word Correlation

- Comparing word correlation with human-judged scores
- Human-judged word correlation [[link](#)]

Word 1	Word 2	Human-Judged Score
tiger	cat	7.35
tiger	tiger	10.00
book	paper	7.46
computer	internet	7.58
plane	car	5.77
professor	doctor	6.62
stock	phone	1.62

Ambiguity: synonym or same word with different POSs

# Extrinsic Evaluation – Subsequent Task

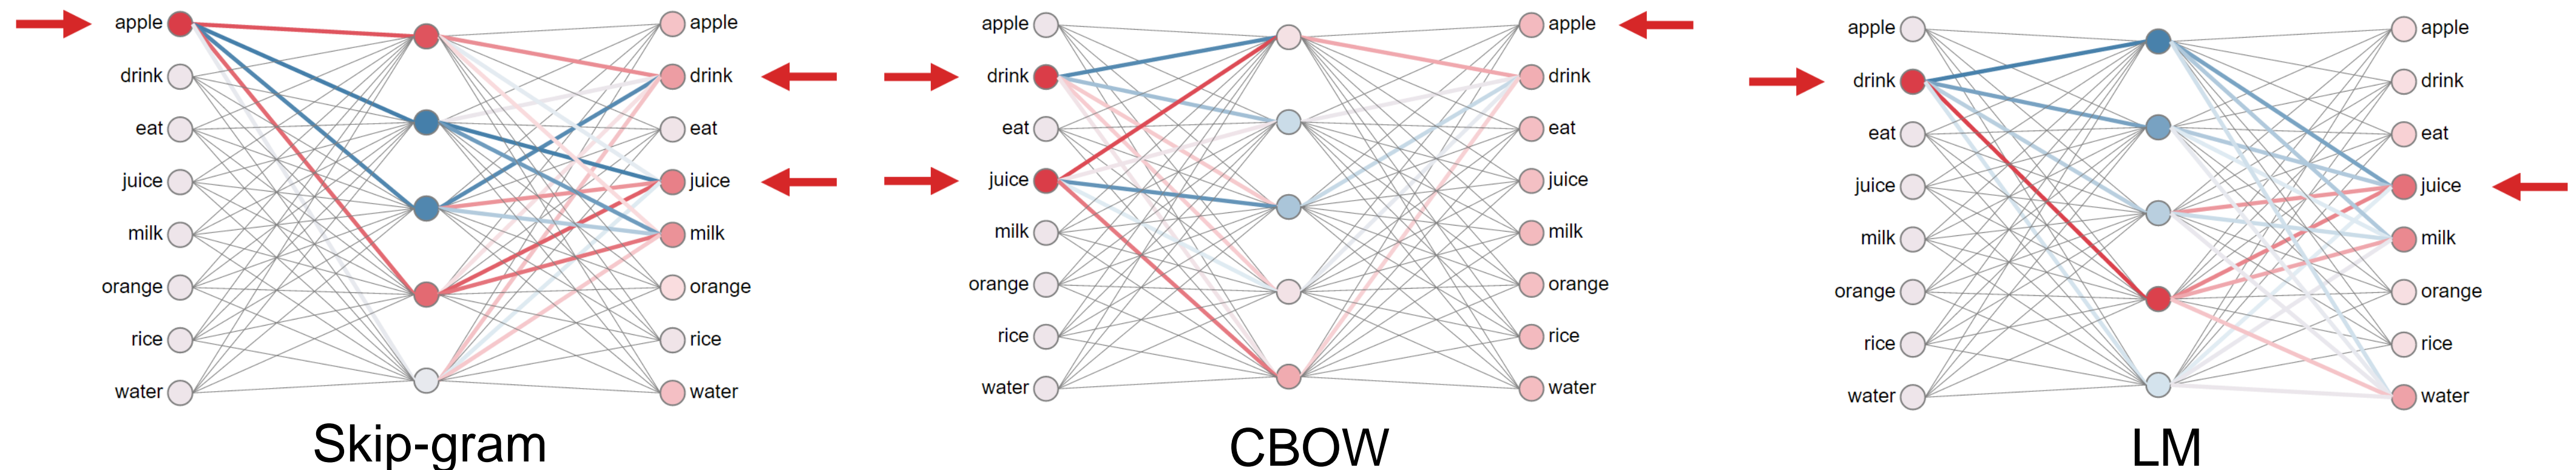
- Goal: use word vectors in neural net models built for subsequent tasks
- Benefit
  - Ability to also classify words accurately
    - Ex. countries cluster together a classifying location words should be possible with word vectors
  - Incorporate any information into them other tasks
    - Ex. project sentiment into words to find most positive/negative words in corpus



# Concluding Remarks

## Low dimensional word vector

### word2vec



### GloVe: combining count-based and direct learning

## Word vector evaluation

- Intrinsic: word analogy, word correlation
- Extrinsic: subsequent task