CNRS - INPT - UPS - UT1 - UTM

Dutch Poetry Generation using Encoder-Decoder Networks

Tim Van de Cruys
CNRS & IRIT, France

10 February 2017

Introduction

- Automatic poetry generation is a challenging task
- Both linguistic and literary aspects need to be taken into account
- System with multiple components
  - Syntax — Recurrent neural network language model
  - Semantics — Latent topic model (NMF)
  - Discourse — Recurrent encoder-decoder network
  - Literary constraints — Prior on network's probability distribution

Recurrent neural network language model

- Gated recurrent neural network (GRU)
  - input layer: current word at time t, represented as word embedding
  - hidden layer: layer with recurrent connection
  - output layer: probability distribution over vocabulary words
  - reset and update gate: control flow of information
- Hidden layer maintains representation of sentence history
- Gates allow fine-grained control of what information the network remembers

Graphical representation

RNNLM

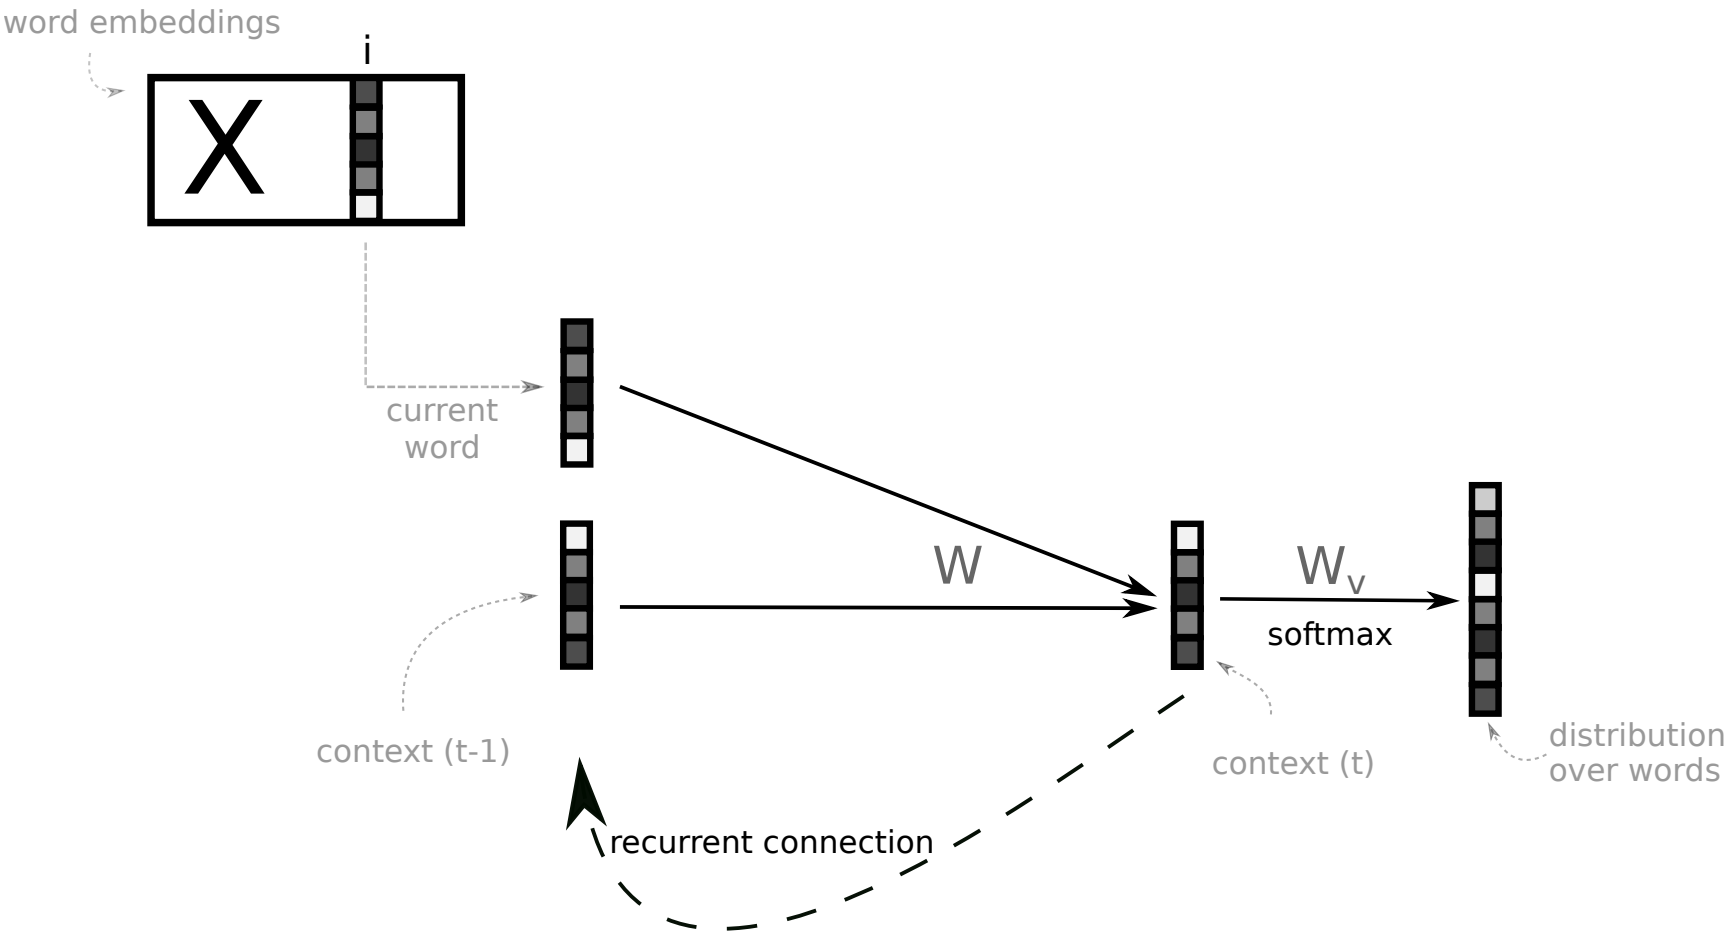

Graphical representation

RNNLM with GRU

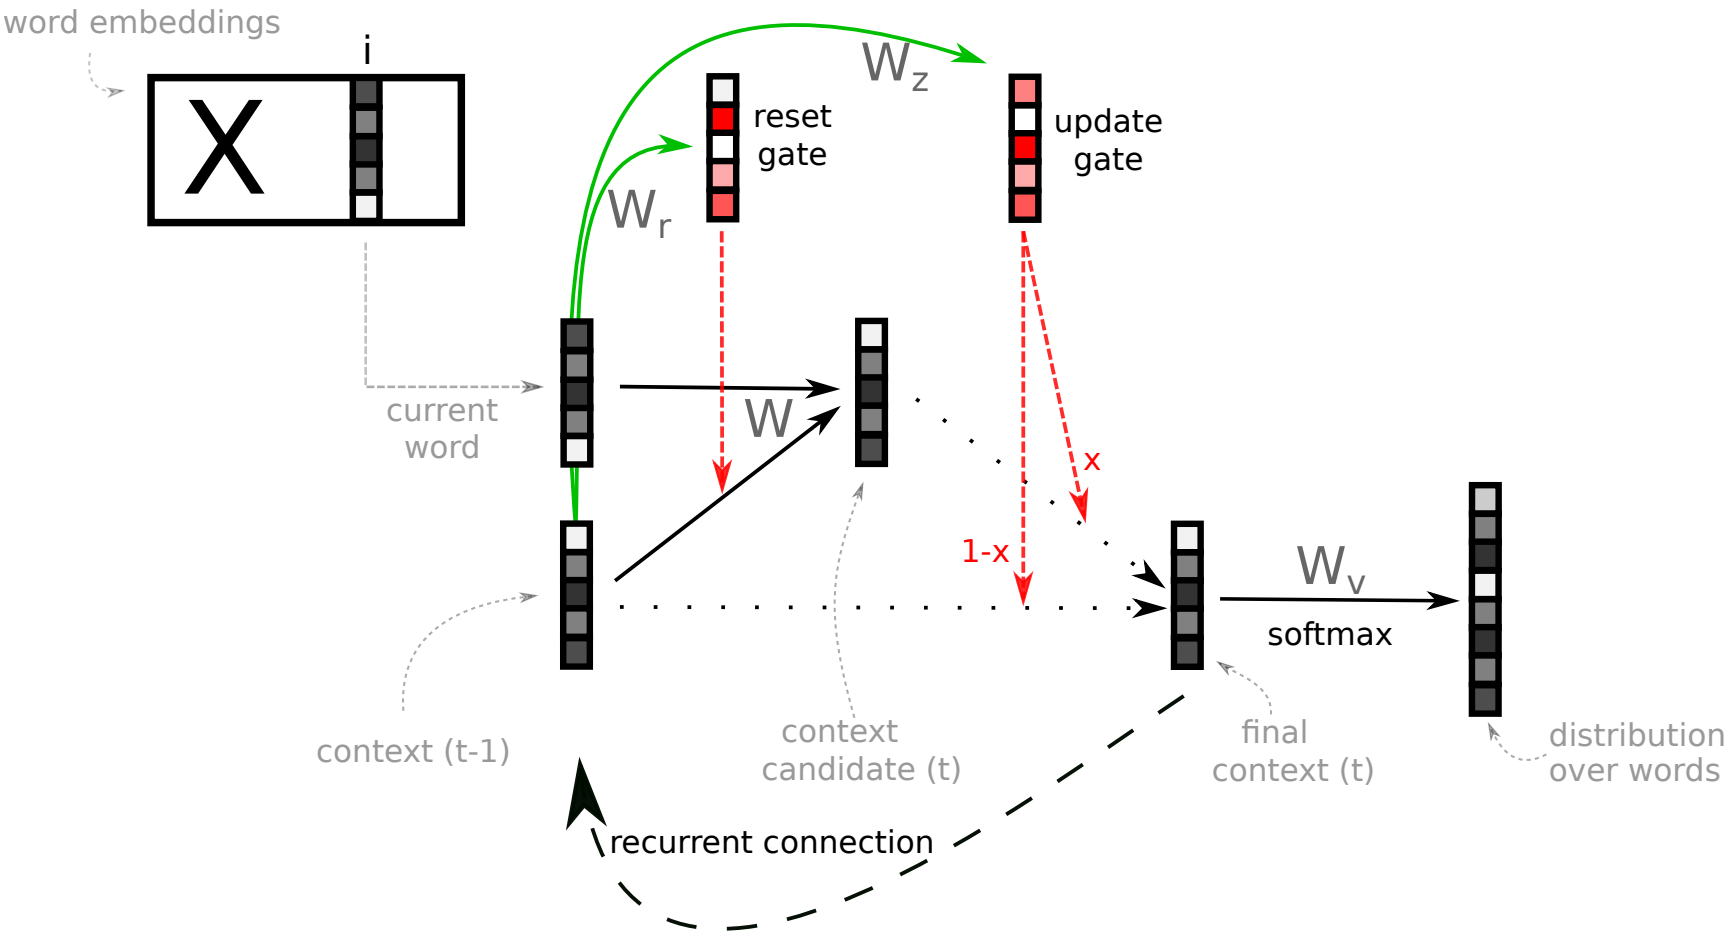

Literary constraints

Rhyme

- Phonetic syllable endings extracted from CELEX database and Dutch Wiktionary

```
onbescheiden    ('sx', 'E+d@')
rusttijden      ('t', 'E+d@')
herdershond     ('h', 'Ont')
landbouwgrond   ('Gr', 'Ont')
```

- Neural network is trained with reversed sentence order
- Last word is generated first
- First time step: rhyme enforced as prior on recurrent network's probability distribution

Graphical representation

Rhyme prior

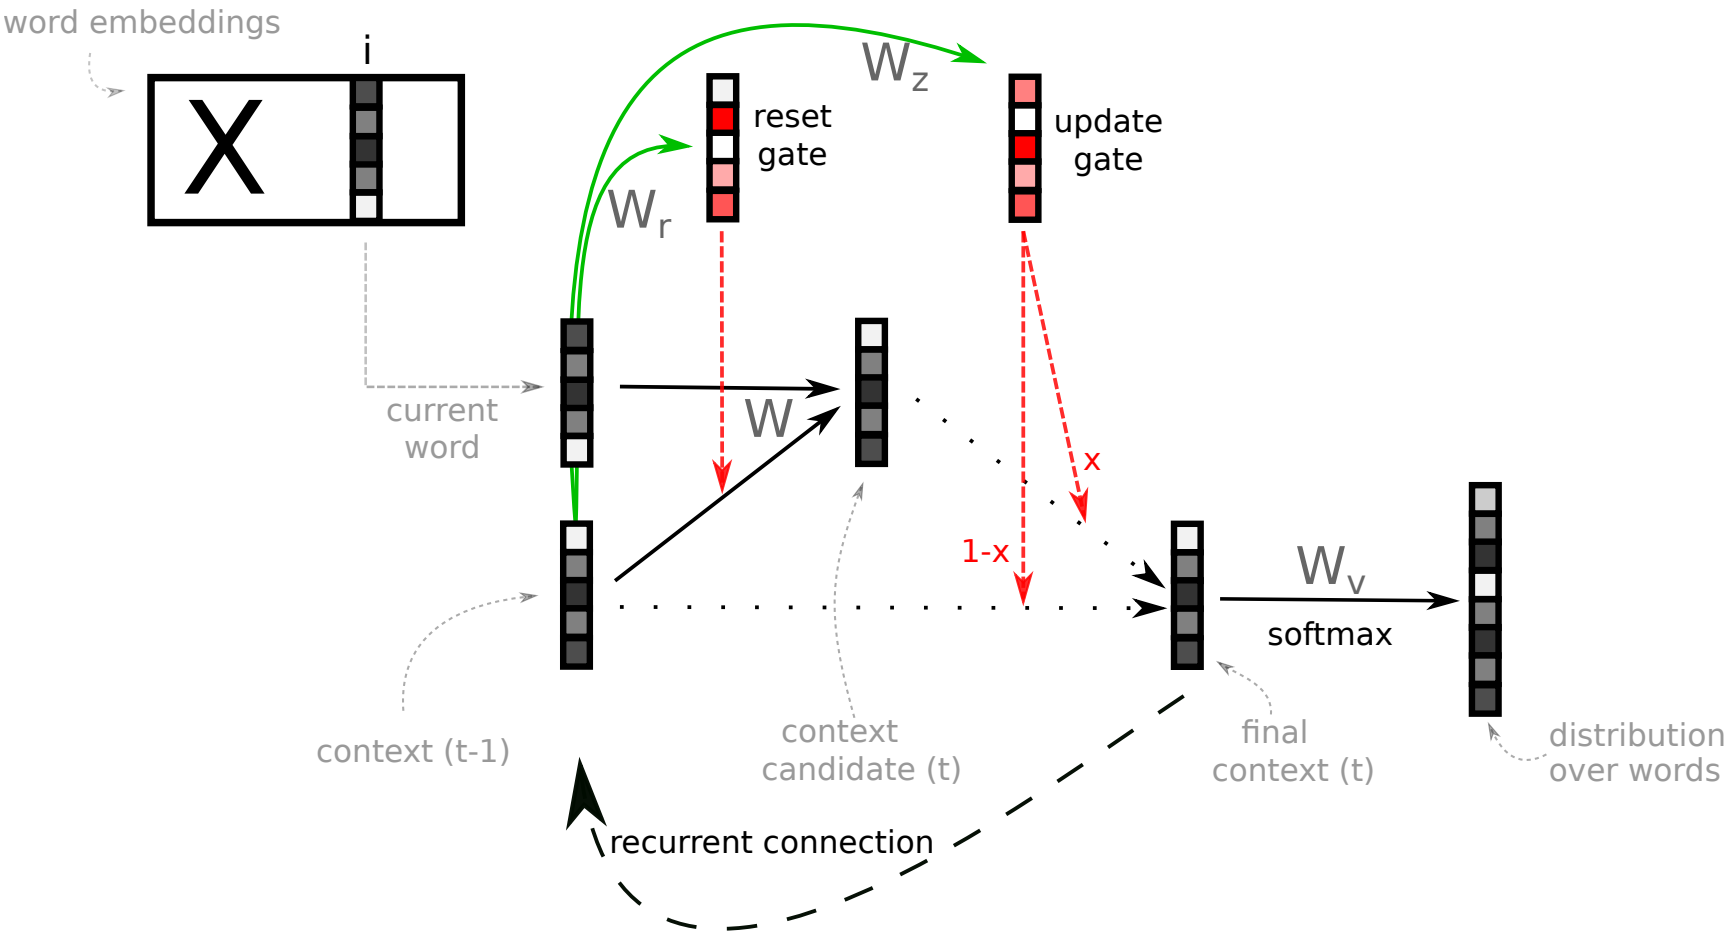

Graphical representation

Rhyme prior

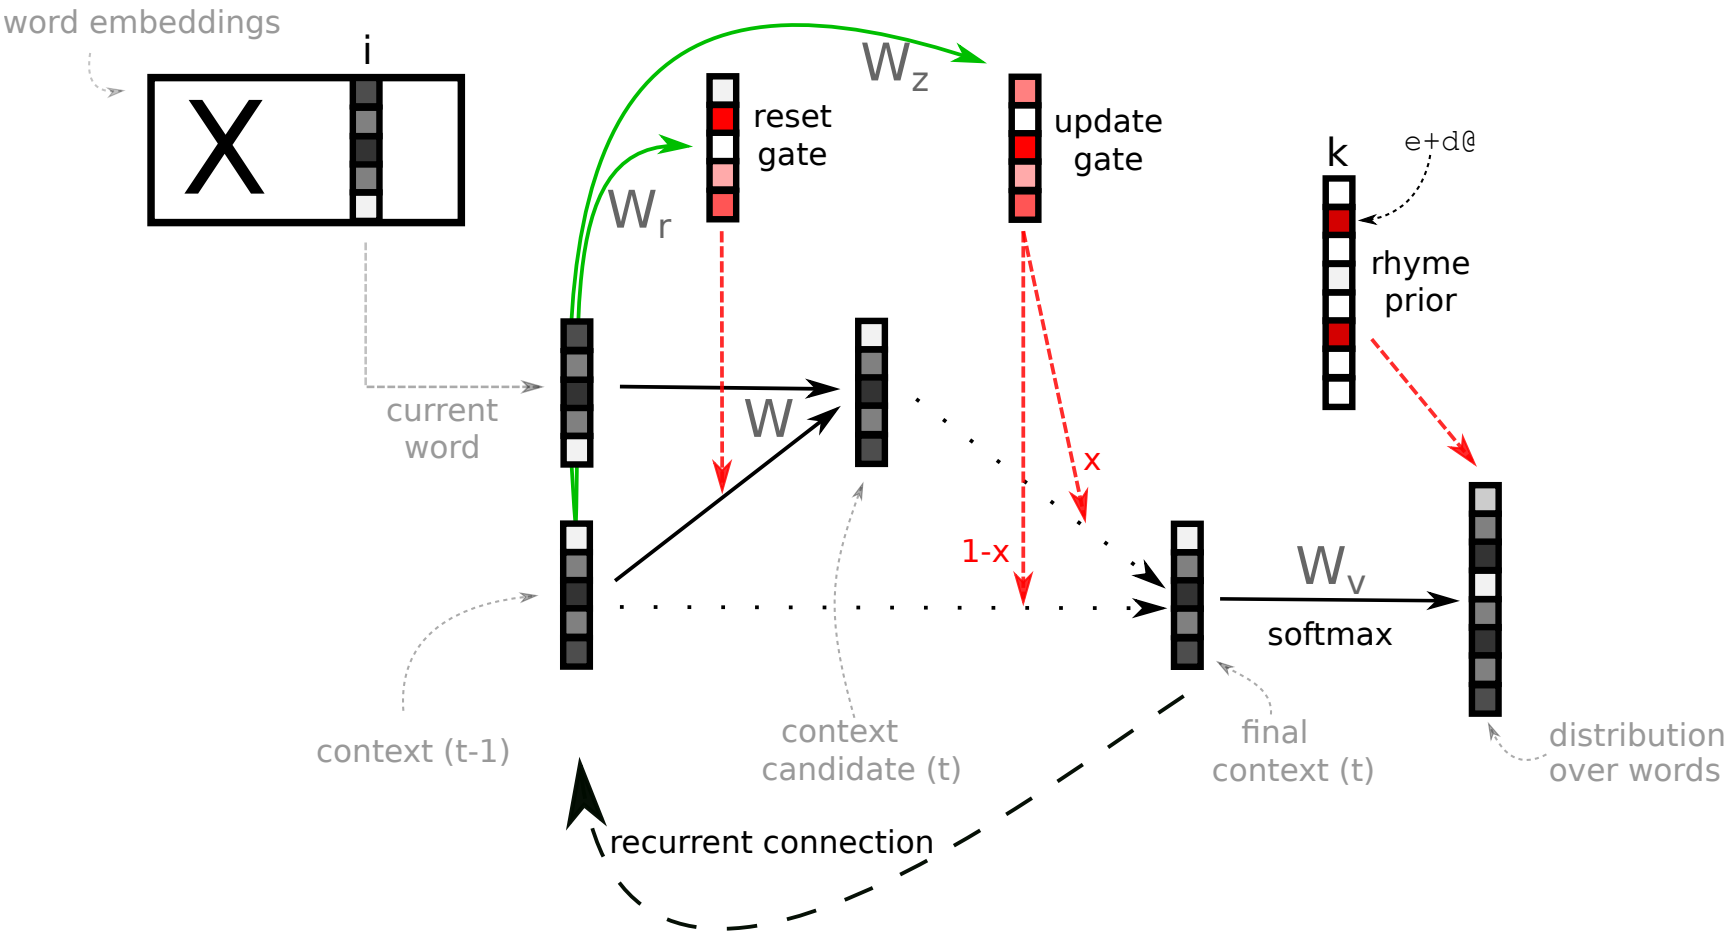

Latent topic model

Non-negative matrix factorization

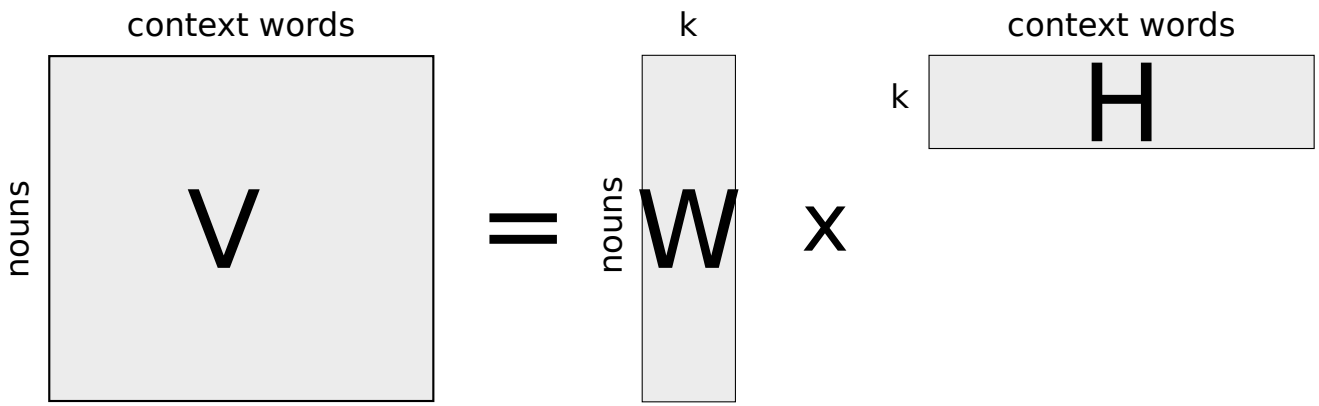

Example

dim 13	dim 20	dim 21	dim 40
slotfase	boter	tranen	barmhartigheid
ruststand	saus	verdriet	almachtige
penalty	pasta	woede	verlosser
gescoord	tomaten	angstig	aanbidden
treffer	olijfolie	schaamte	ongerechtigheid
thuisclub	gedroogde	wanhoop	heiland
thuisploeg	peper	gevoeld	heiligheid
balbezit	kaas	schrik	geopenbaard
scorebord	paprika	begreep	zaligheid
kansloos	salade	wanhopig	dienaar

Graphical representation

Semantic prior

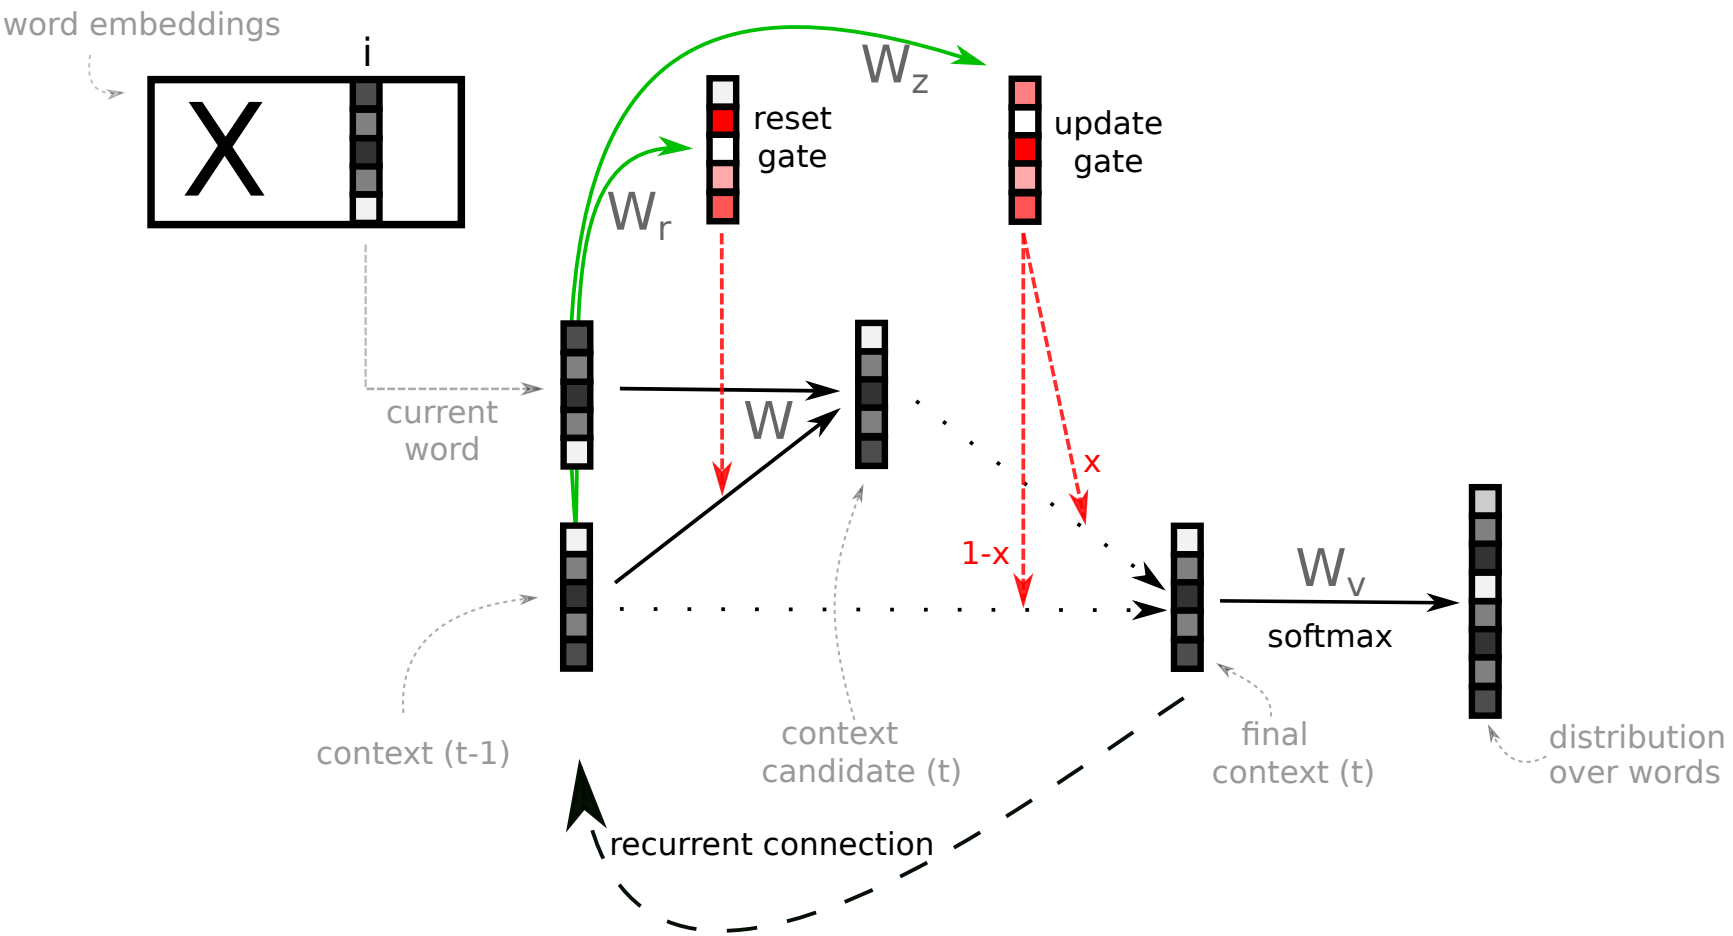

Graphical representation

Semantic prior

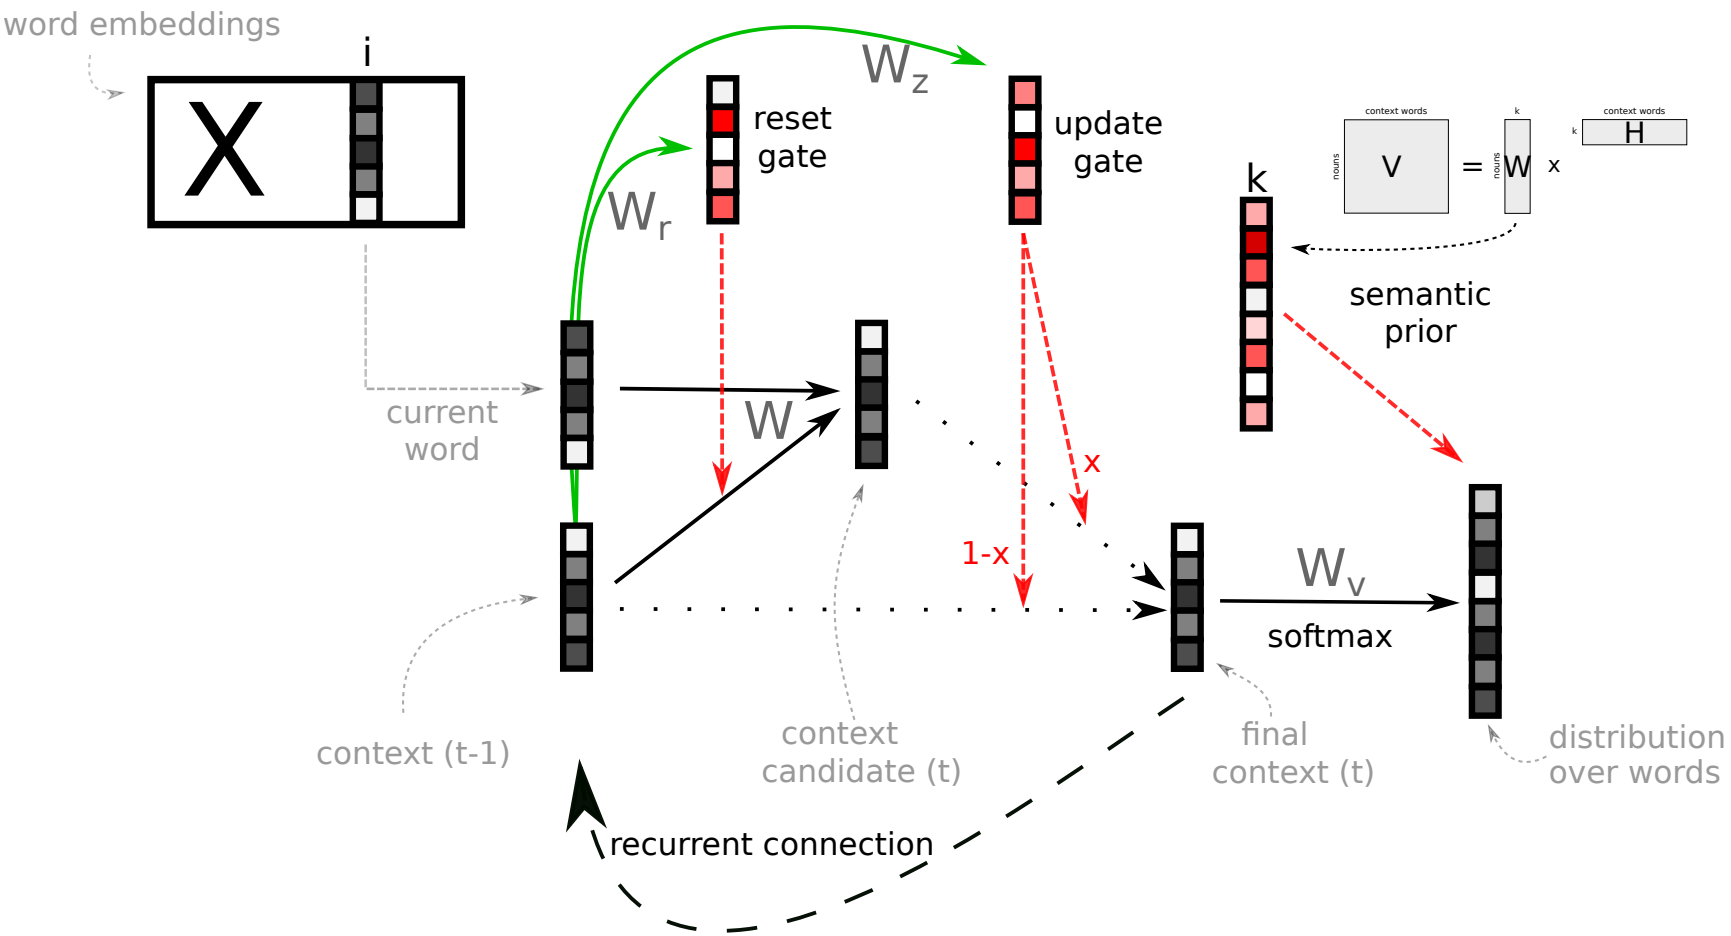

Discourse

Encoder-decoder

- Recurrent neural network outputs well-formed lines, but no coherence between successive lines
- Train recurrent encoder-decoder to predict the next sentence from previous sentence
- $\sim$ *skip thought* model [Kiros et al. 2015], but only next sentence
- Similar structure as RNNLM, but next sentence in its entirety predicted from last state of previous sentence
- Goal: model inter-sentential structure (discourse markers, coherence)

Graphical representation

Encoder-decoder

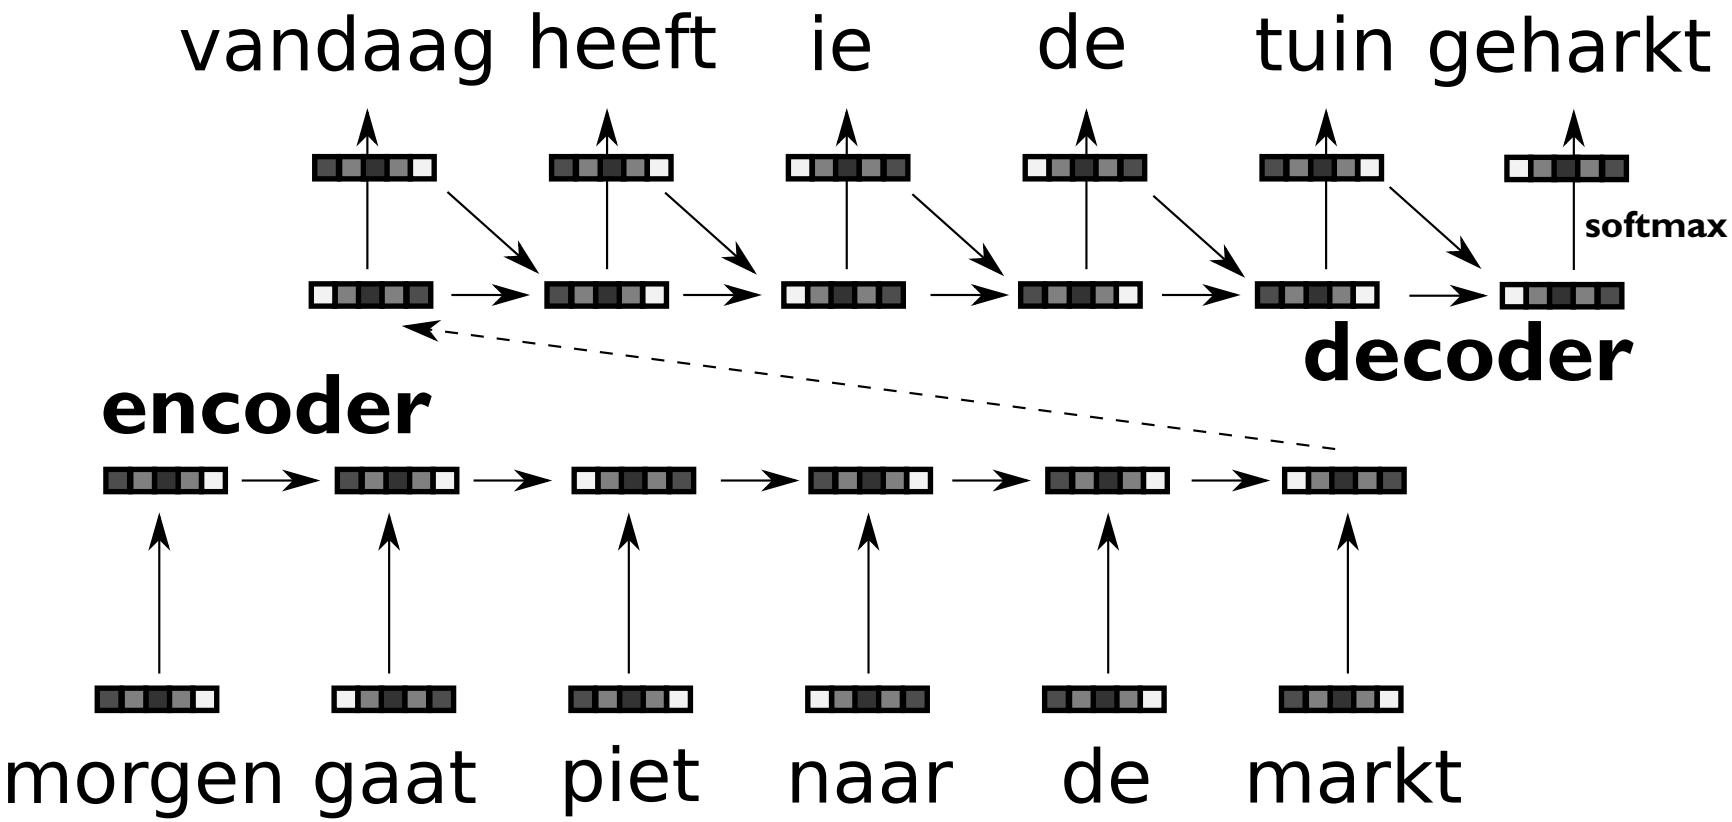

Optimization framework

- RNNLM (with rhyme and semantic priors) generates a number of possible verse continuations
- Each line is given a score using encoder-decoder (log-probability)
- Line with best score is included in poem
- (also possible to directly generate lines using encoder-decoder, but joint approach yields more variation)

Implementation details

- Neural models trained on 150 million words of web text (NLCOW)
- Vocabulary: 10K, embedding size: 256, hidden layer size: 512
- Batch stochastic gradient descent with Nesterov momentum
- Additional optimization criteria:
  - standard *n-gram* score (order 3, Kneser-Ney smoothing)
  - NMF score
  - line length
- NMF & n-gram model trained on 1.5 billion words of web text

Example (1)

```
verhit de boter in de pan met heerlijk frisse lucht
voor de rest is gewoon de smaak te raden
als ik teveel suiker eet, maak ik een vlucht
het was een hele grote vis in het kwade

vooral de smaak van het drinken was veel te groot
we hebben de rest van de dag met elkaar gesproken
de rest van de dag gaat hij een beetje dood
de rest van de dag begint de kip te koken
```

Example (2)

```
boven de felle zon scheen de zon en de zon scheen
soms waren de woorden van een mens zo boos
de zon scheen de zon schijnt de zon lijkt er een
de zon doet af en toe een hele grote doos

wat een droom, wat een heerlijk gevoel zeg
voor de huid is het licht bruin van kleur
aan het licht van licht valt de stilte weg
wanneer je sterft, dan sta je voor een deur
```

Conclusion

- Gated recurrent neural network generates well-formed lines
- Poetic and semantic constraints may be incorporated as priors on probability distribution
- Encoder-decoder network allows to model basic discourse coherence between subsequent lines
- But for real poetry, more discourse, pragmatics, and intelligent creativity is needed
- Future work
  - Further development of encoder-decoder structure (hierarchical representation)
  - Incorporate poetic optimization in training objective
  - Human evaluation


Kyunghyun Cho, Bart van Merrienboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734. 2014.


Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Skip-thought vectors. In *Advances in neural information processing systems*, pp. 3294–3302. 2015.


Daniel D. Lee and H. Sebastian Seung. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems*, pp. 556–562. 2001.


Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocky, and Sanjeev Khudanpur. Recurrent neural network based language model. In *INTERSPEECH 2010, 11th Annual Conference of the International Speech Communication Association*, pp. 1045–1048. 2010.


Tomas Mikolov, Stefan Kombrink, Lukas Burget, Jan Cernocky, and Sanjeev Khudanpur. Extensions of recurrent neural network language model. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5528–5531. 2011.


Tim Van de Cruys, Thierry Poibeau, and Anna Korhonen. Latent vector weighting for word meaning in context. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 1012–1022. 2011.

Anthology

http://www.timvandecruys.be/media/guido.pdf