



Microsoft Research AI

Deep Communicating Agents for Abstractive Summarization

Asli Celikyilmaz


Deep Learning Technology Center
Microsoft Research AI

Antoine Bosselut, Yejin Choi

Paul G. Allen School of CSE
University of Washington

Xiaodong He

JD AI Research



Generation by Multi-Agent Communication

- Improves discovery of salient concepts in long text.
- Enables multiple encoders communicate through message passing and agent attention.
- Generates more coherent summaries with increased focus
- Reduces decoder redundancy

Human: Mr Turnbull was interviewed about his childhood and his political stance. He also admitted he **planned to run for prime minister if Tony Abbott had been successfully toppled in February's leadership spill**. The words 'primed minister' were controversially also printed on the cover.

Single-Agent: Malcolm Turnbull is set to feature on the front cover of the GQ Australia in a bold move that will no doubt set senators' tongues wagging. **Posing in a suave blue suit with a pinstriped shirt and a contrasting red tie**, Mr Turnbull's confident demeanour is complimented by the bold, confronting words printed across the page: 'primed minister'.

Multi-Agent: Malcolm Turnbull was set to run for prime minister if Tony Abbott had been successfully toppled in February's leadership spill. He is set to feature on the front cover of the liberal party's newsletter.

Semantic Cohesion

Encourage sentences in the summary to be informative without repetition.

$$L_{SEM} = \sum_{q=2}^Q \cos(s'_q, s'_{q-1})$$

Self Critical REINFORCE

Enables exploring new sequences and comparing them to the best greedily decoded sequence.

$$L_{RL} = (r(\tilde{y}) - r(\hat{y})) \sum_{t=1}^N \log p(\hat{y}_t | \hat{y}_1 \dots \hat{y}_{t-1}, d)$$

Intermediate Rewards

Sentence based rewards (r) as opposed to end of summary rewards. \hat{o}_q is q^{th} generated sentence

$$r(\hat{o}_q) = r([\hat{o}_1, \dots, \hat{o}_q]) - r([\hat{o}_1, \dots, \hat{o}_{q-1}])$$

Automatic Evaluations

CNN/DailyMail Model	R-1	R-2	R-L
SummaRuNNer-extractive ¹	39.60	16.20	35.30
Pointer + Coverage ²	39.53	17.28	36.38
Controlled Summarization ³	39.75	17.29	36.54
Intra-Attention, RL ⁴	41.16	15.75	39.08
Intra-Attention, Mixed ⁴	39.87	15.82	36.90
DCA (Ours)	41.69	19.47	37.92

NYT Model	R-1	R-2	R-L
MLE ⁴	44.26	27.43	40.41
RL ⁴	47.22	30.51	43.27
Mixed ⁴	47.03	30.72	43.10
DCA (Ours)	48.08	31.19	42.33


Agents




Ablation Study

CNN/DailyMail Model	R-1	R-2	R-L
Single Agent, MLE	36.12	14.38	33.83
Single Agent, MLE + SEM	36.90	15.02	33.00
DCA, MLE+SEM, no comm	37.45	15.90	34.56
DCA, MLE+SEM	39.52	17.12	36.90
DCA, MLE+SEM, caa	41.11	18.21	36.03
DCA, MLE+SEM+RL, caa	41.69	19.47	37.92

Head-to-Head Human Evaluations

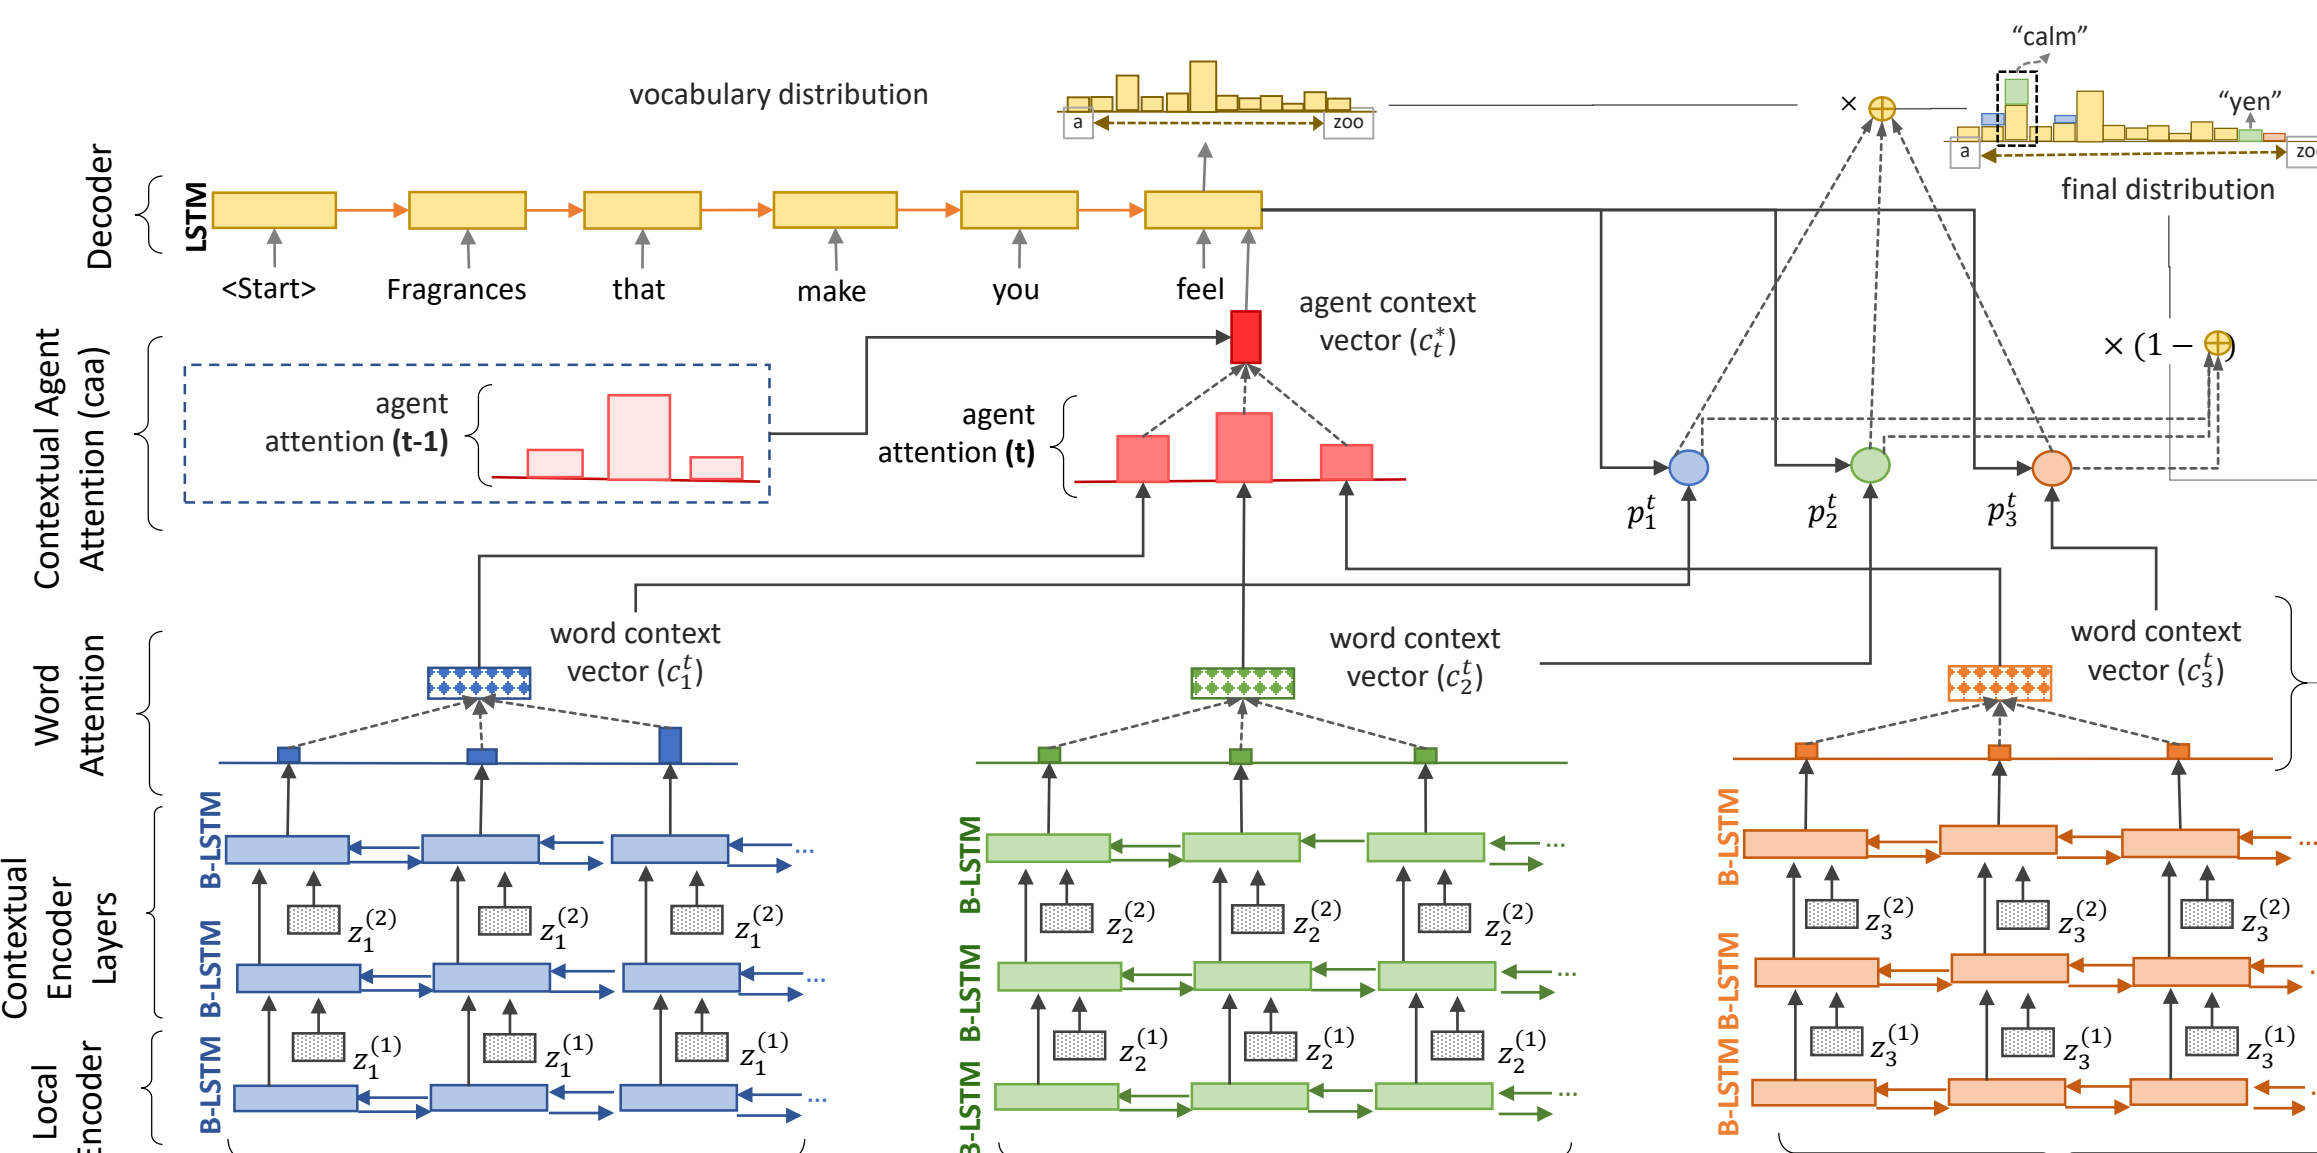


Agent Distribution and Communication



- Average ROUGE-L scores for summaries that are binned by each agent's average attention when generating the summary.
- When agents contribute equally to the summary, the ROUGE-L score increases.

Multi-Agent Encoder-Decoder with Pointer Network

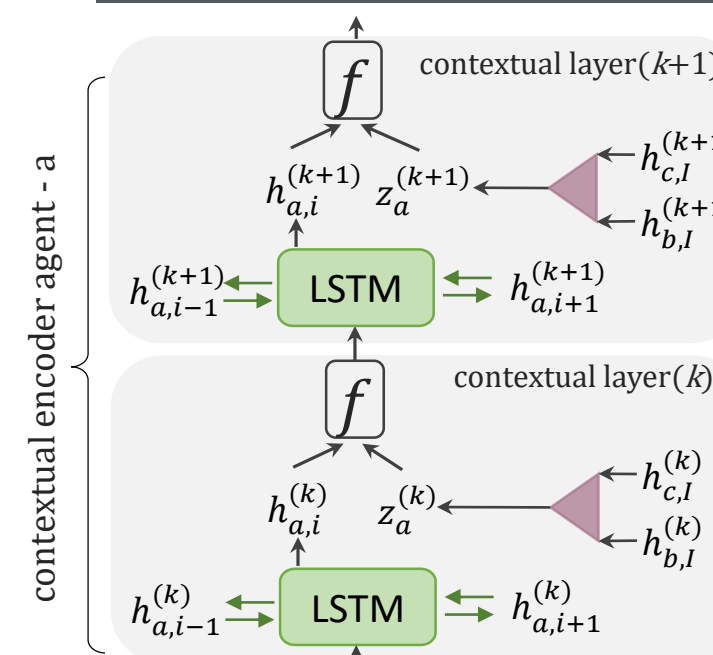


Prime Minister Tony Abbott has been thrown a lifeline by Australian voters just as his internal opponents appeared ready to dump him, with the latest opinion poll confirming a pro-government shift is under way. The Fairfax-Ipsos poll ...

Labor's primary vote dropped four points to sit at 36 per cent, while the Coalition gained four points to hit 42 per cent, according to the poll of 1,406 voters. The results confirm a huge surge in government support suggested last week by ...

Labor opposition leader Bill Shorten remains the preferred prime minister compared to Mr Abbott, but his margin has been slashed from 16 points to five since the previous poll to sit at 44-39. On Sunday, Mr. Abbott laughed off ...

Multi-Agent Encoder Message Passing



Agent b and c transmit the last hidden state output (h) of the current layer k as a message which are mean pooled:

$$z^{(k)} = \frac{1}{M-1} \sum_{m \neq a} h_{m,l}^{(k)}$$

Multi-Agent Pointer Network

$$p_a^t \in [0,1]$$

$$p_a^a('calm') = p_a^t * P('calm') + (1 - p_a^t) * att('calm')$$

$$p('calm') = \sum_a agt-attn(a) * p_a^a('calm')$$

Generation switch

Balances generation and copying

Per-agent word probability

Probability distribution over the extended vocabulary. Computed for each agent from generation probability and pointer attention for each word

Final word probability

Probability distribution over the extended vocabulary.