


Autoencoders

(30 min)

Autoencoder




Reconstruction Loss

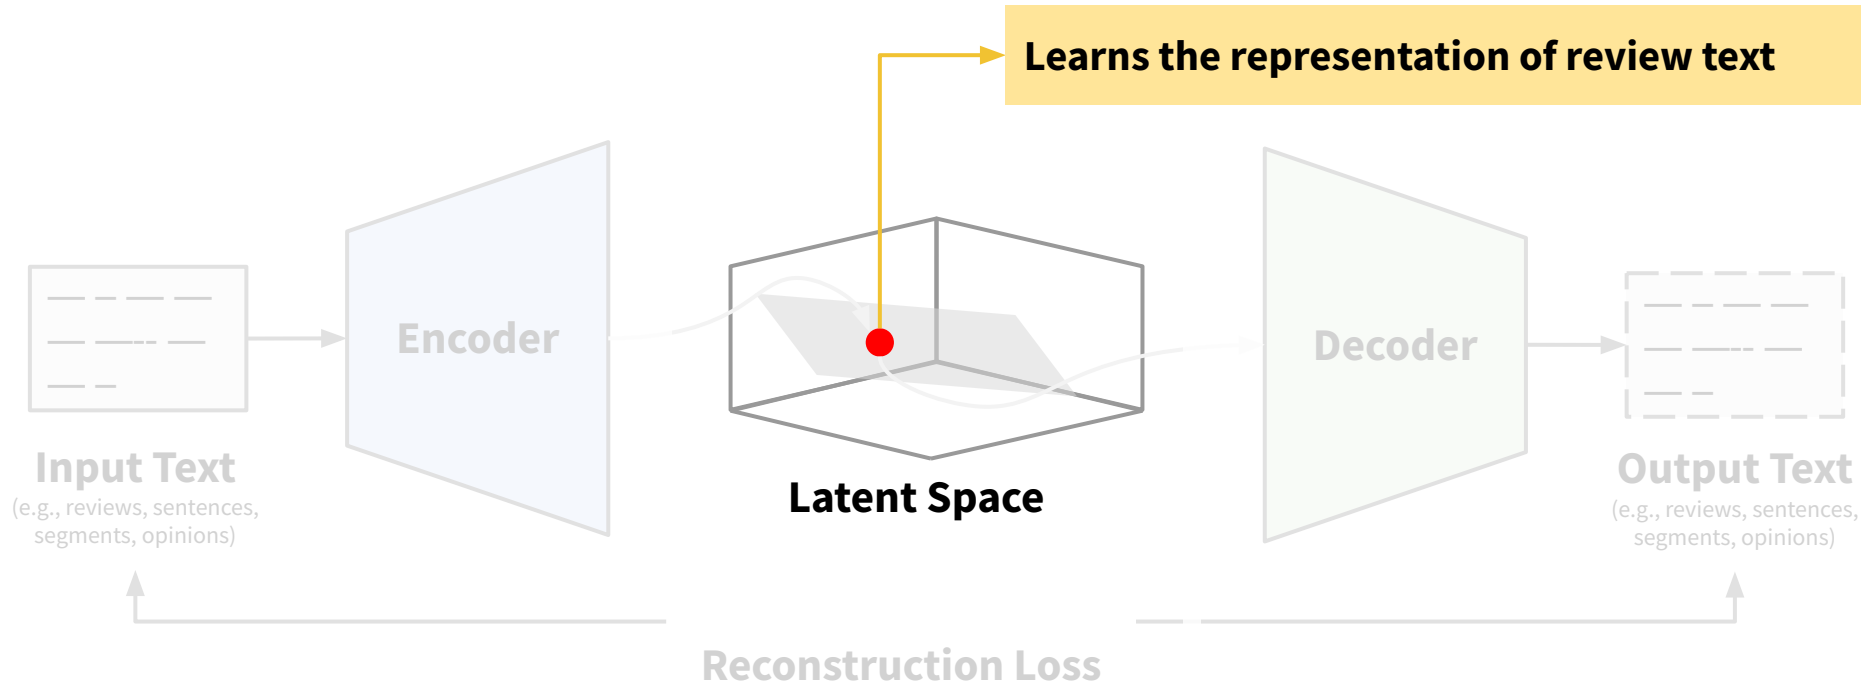
$$Loss = \mathcal{L}(x, D(E(x)))$$

Wide applications in both NLP and CV!

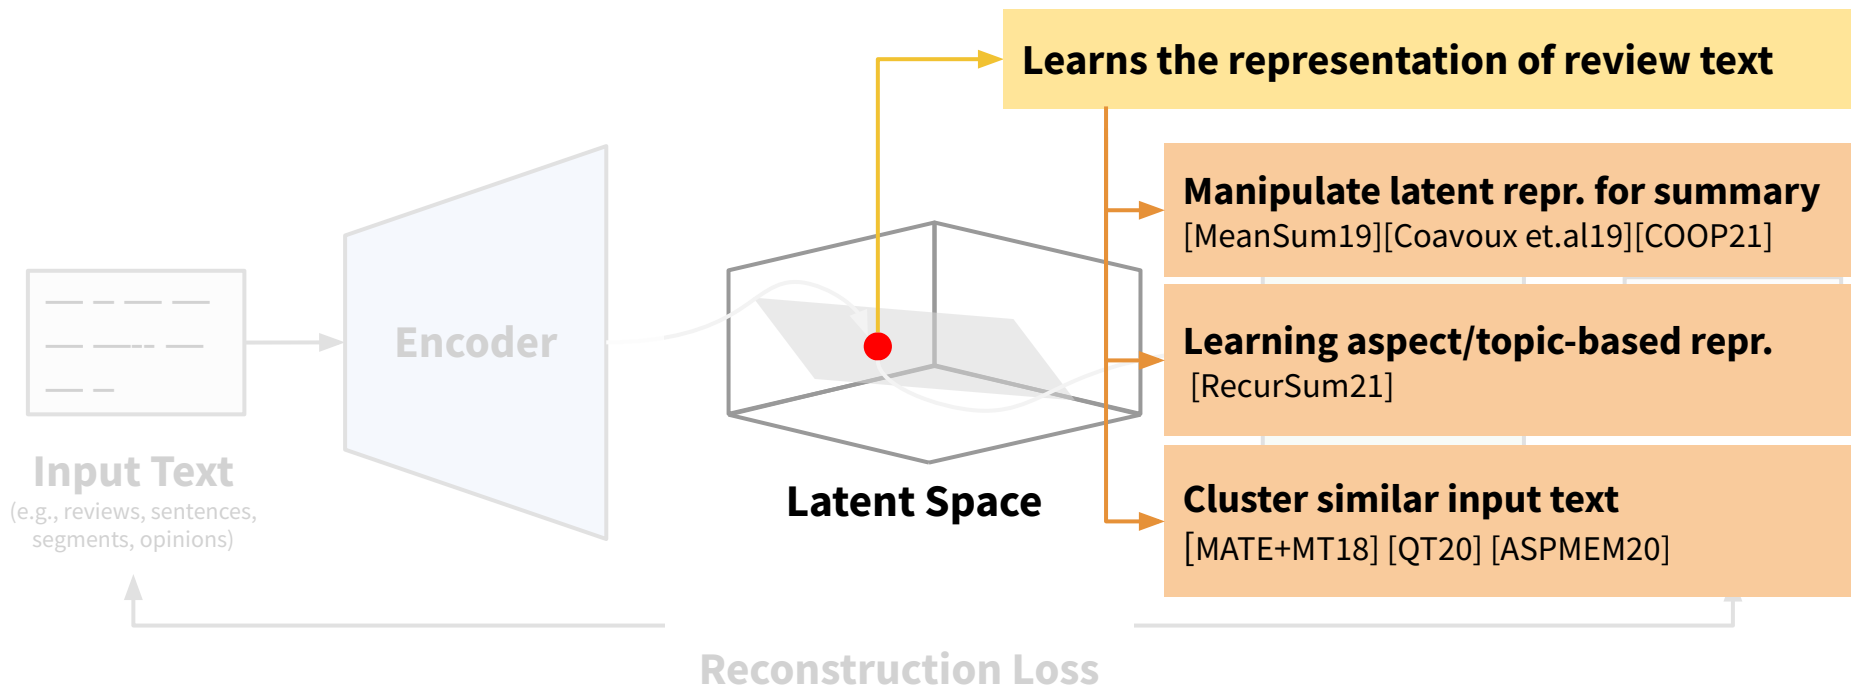
Opinion Summarization via Autoencoder



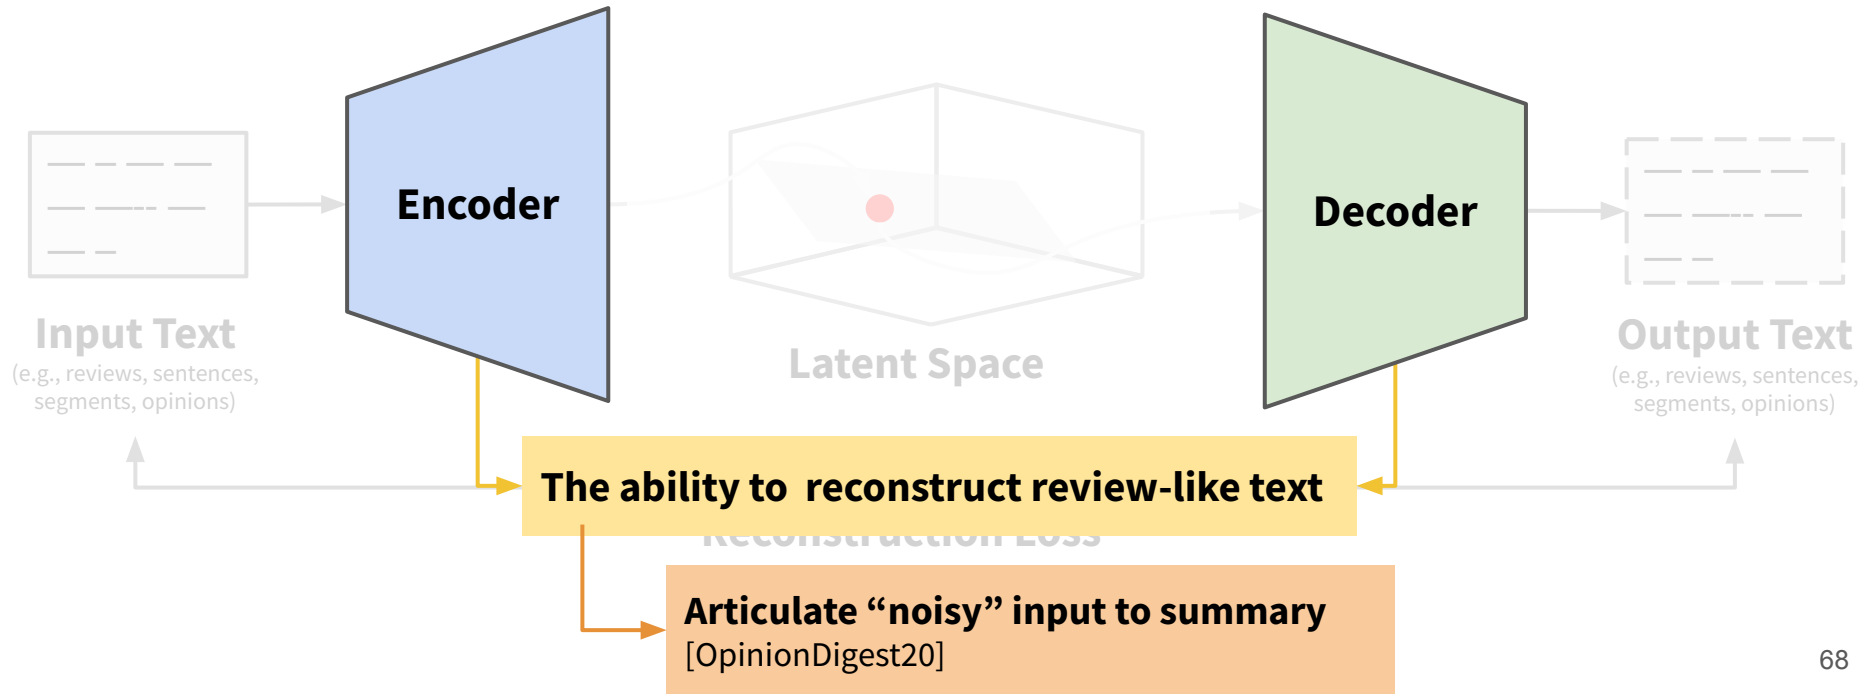
Opinion Summarization via Autoencoder



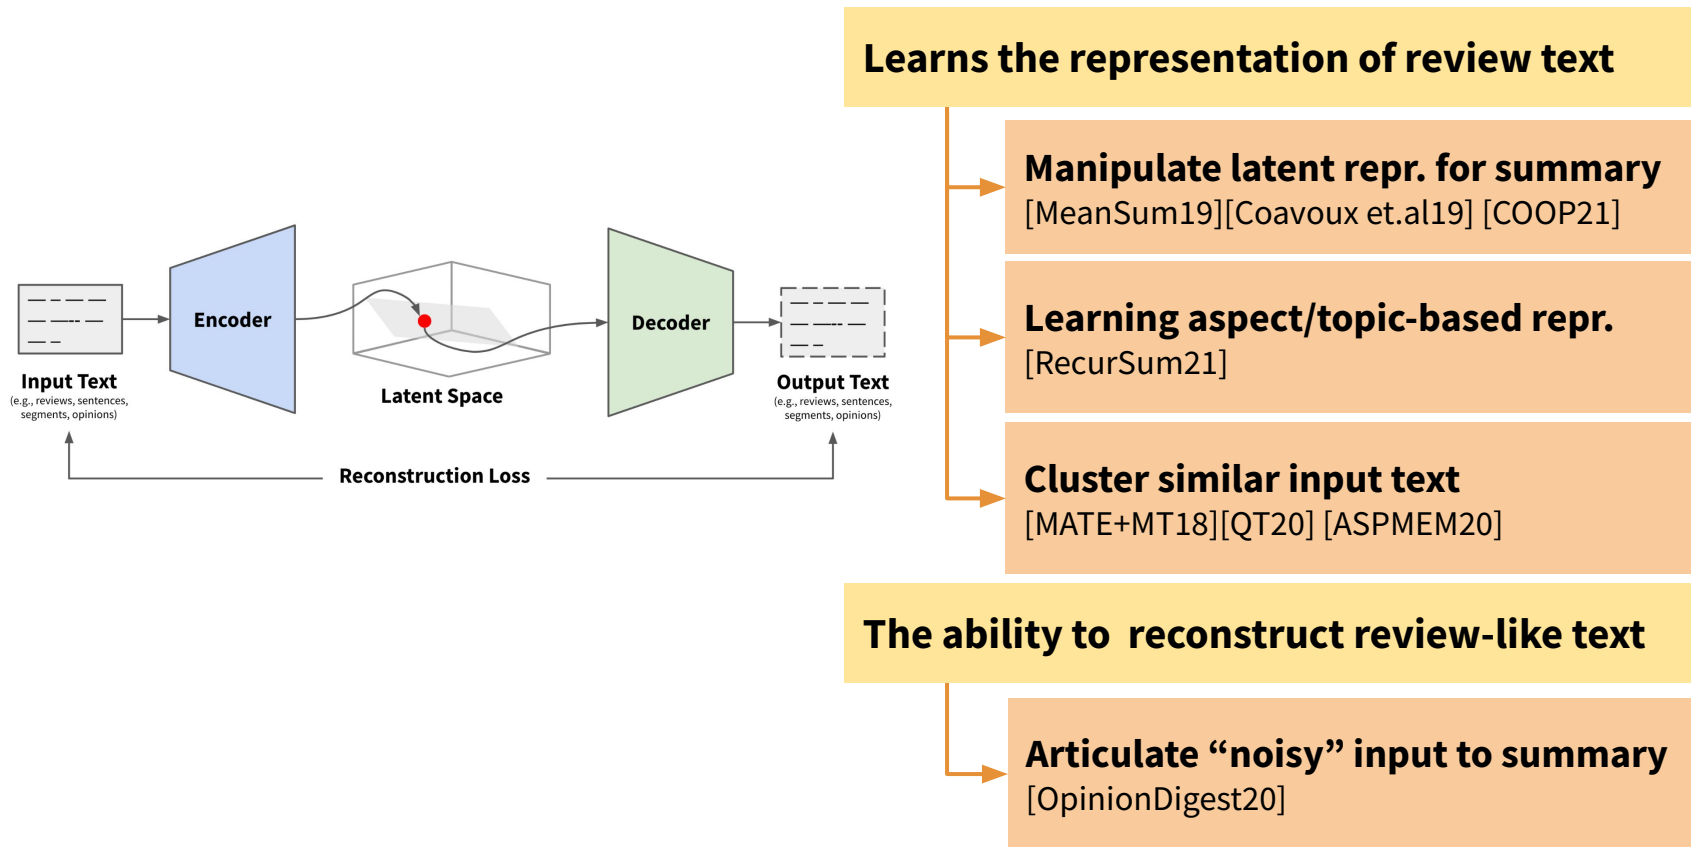
Opinion Summarization via Autoencoder



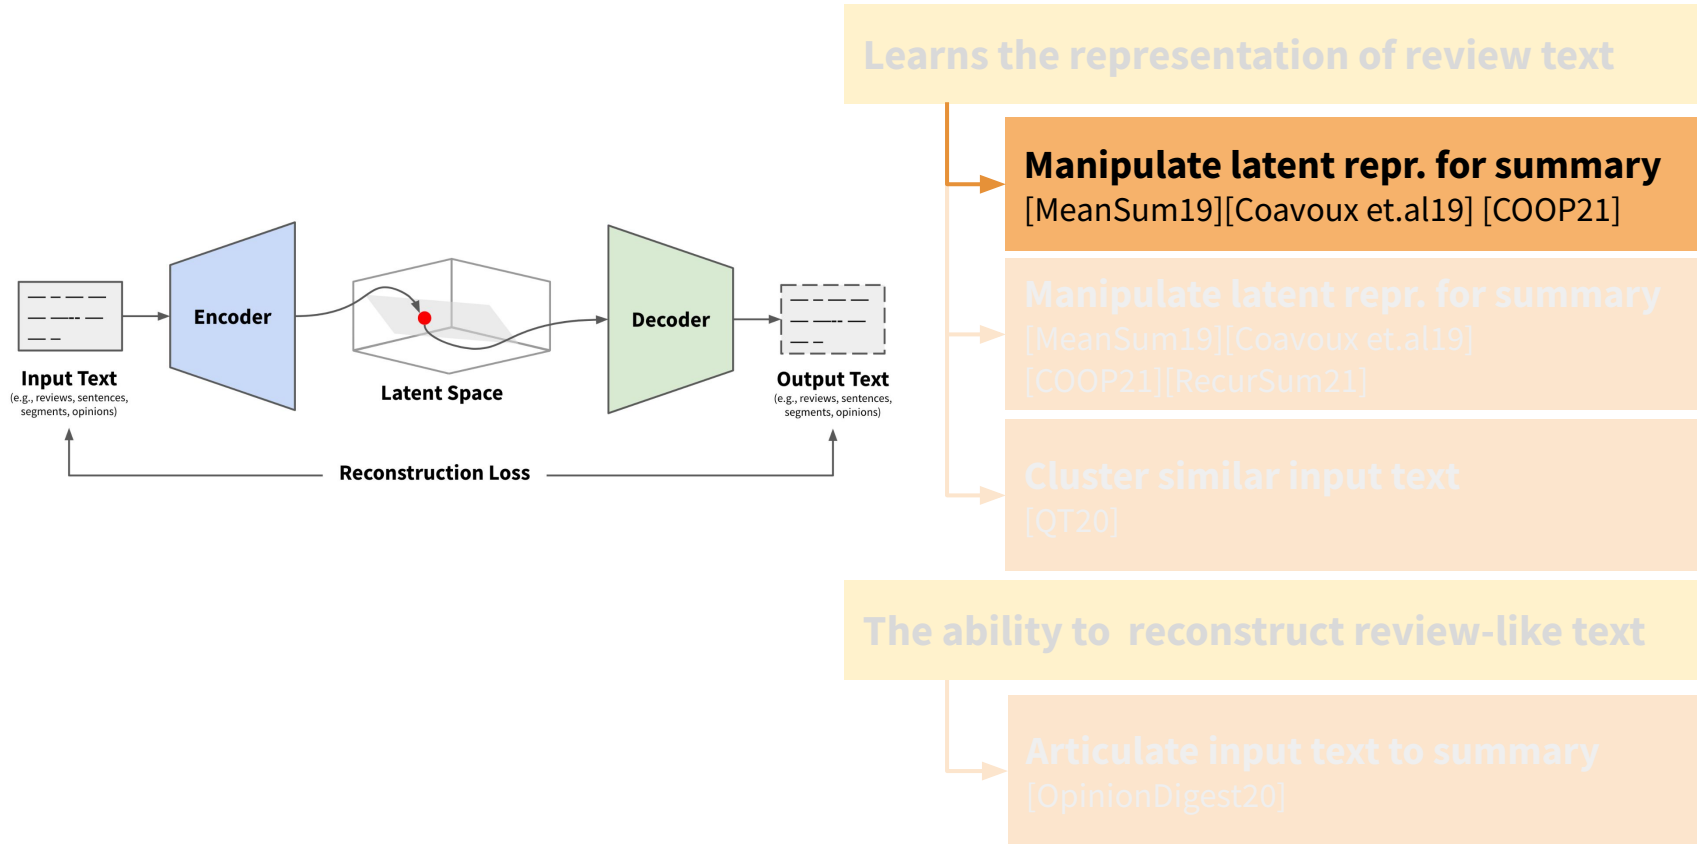
Opinion Summarization via Autoencoder



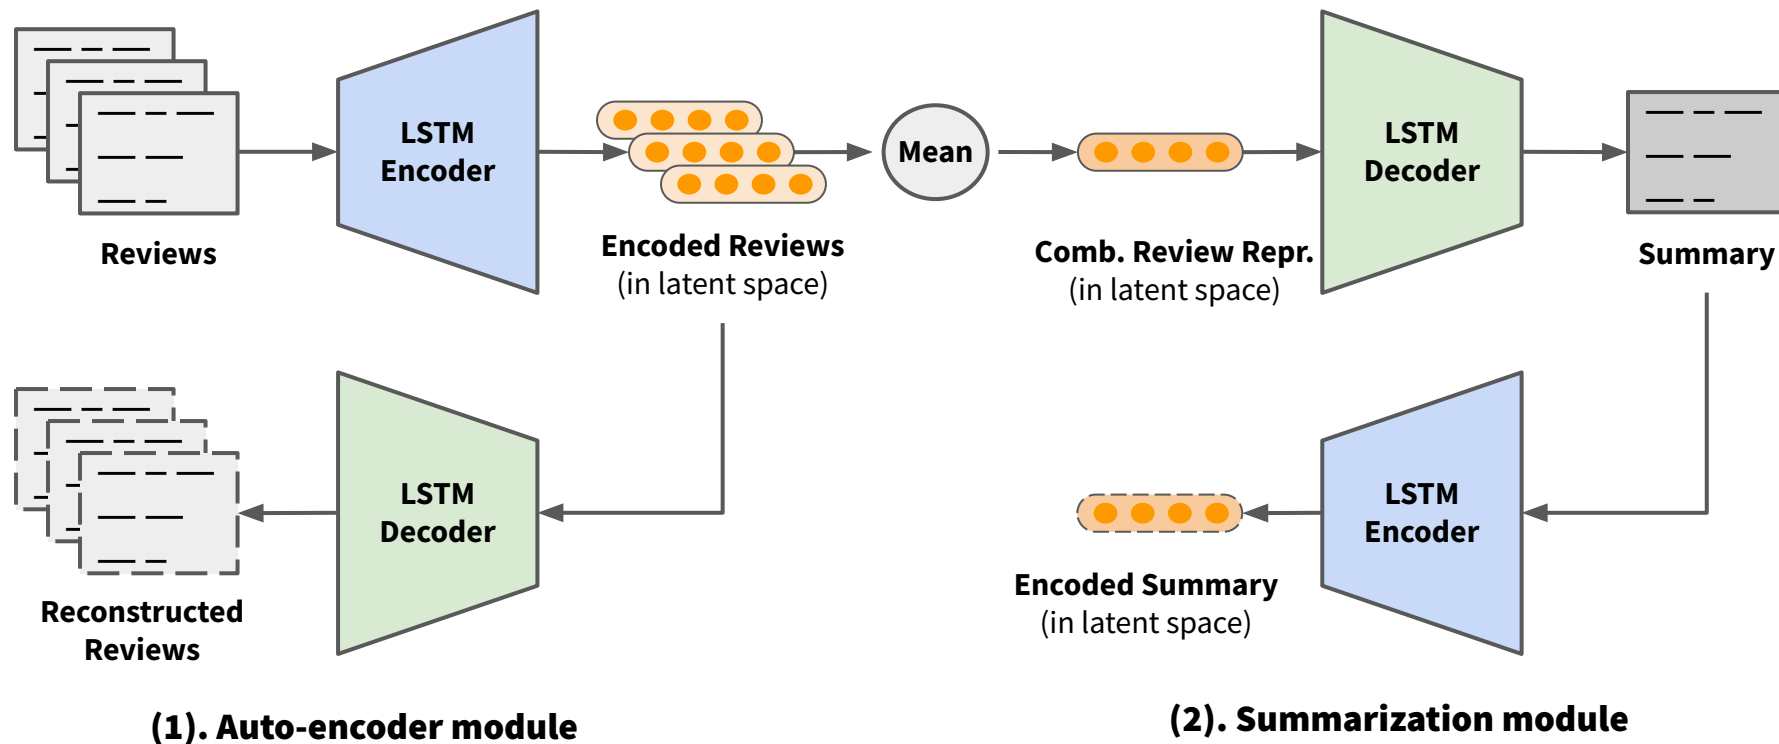
Opinion Summarization via Autoencoder



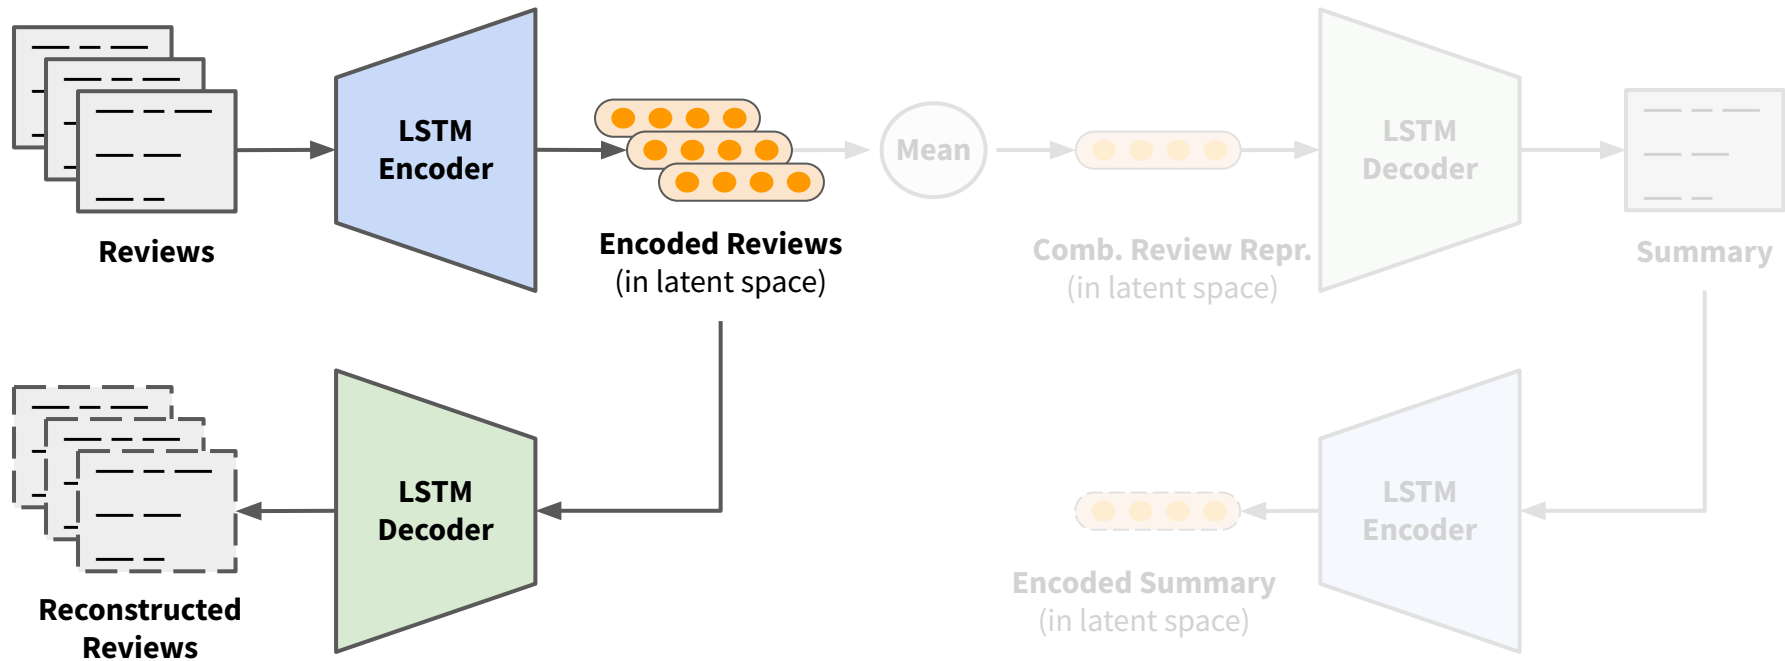
Opinion Summarization via Autoencoder



MeanSum¹



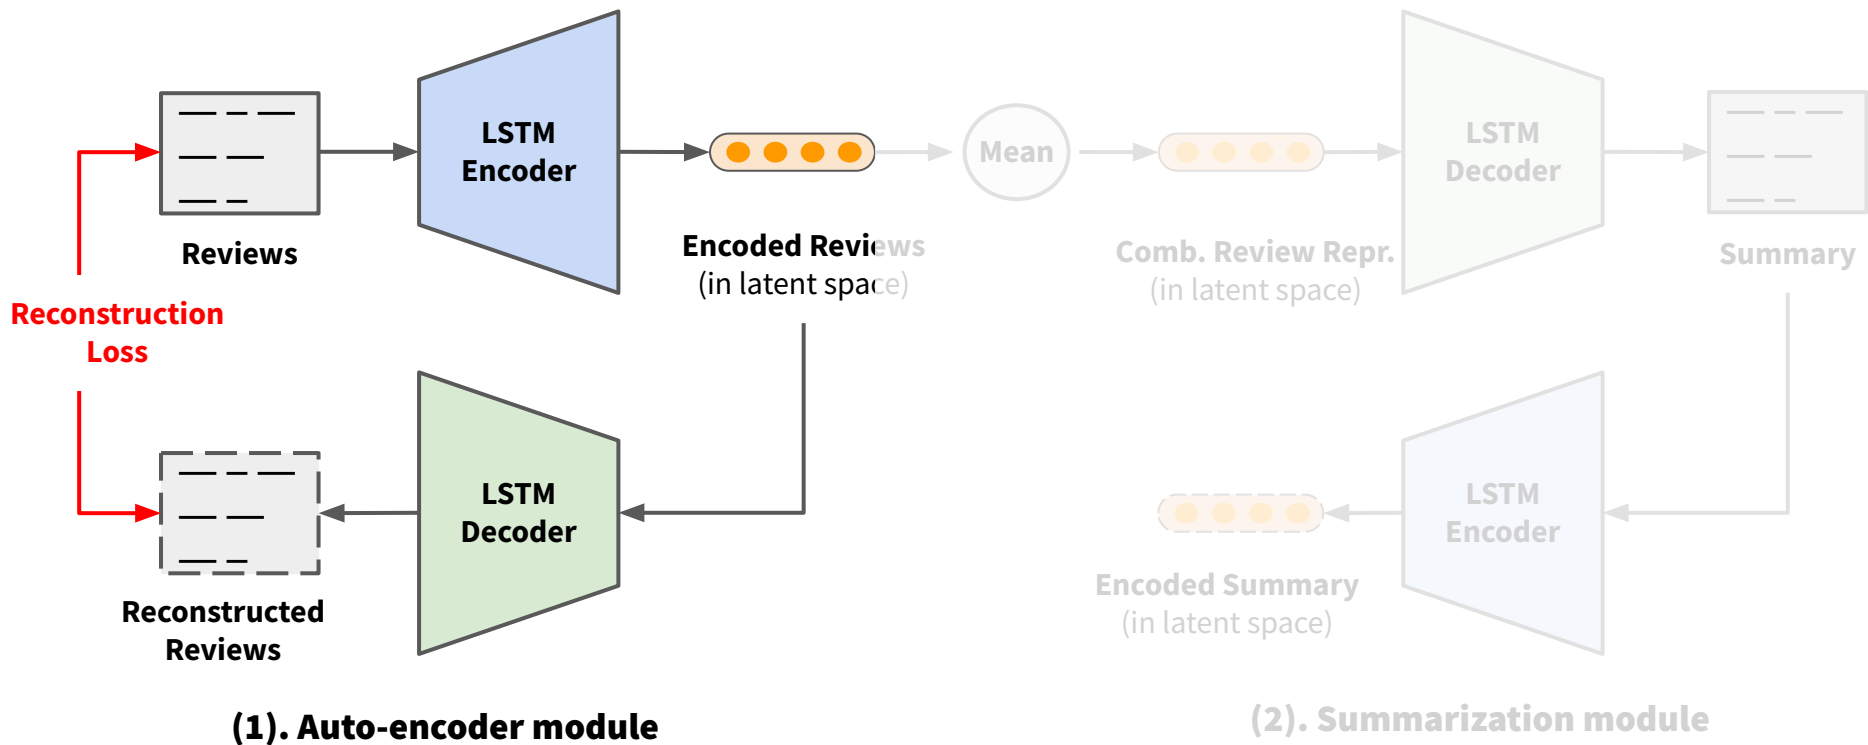
MeanSum



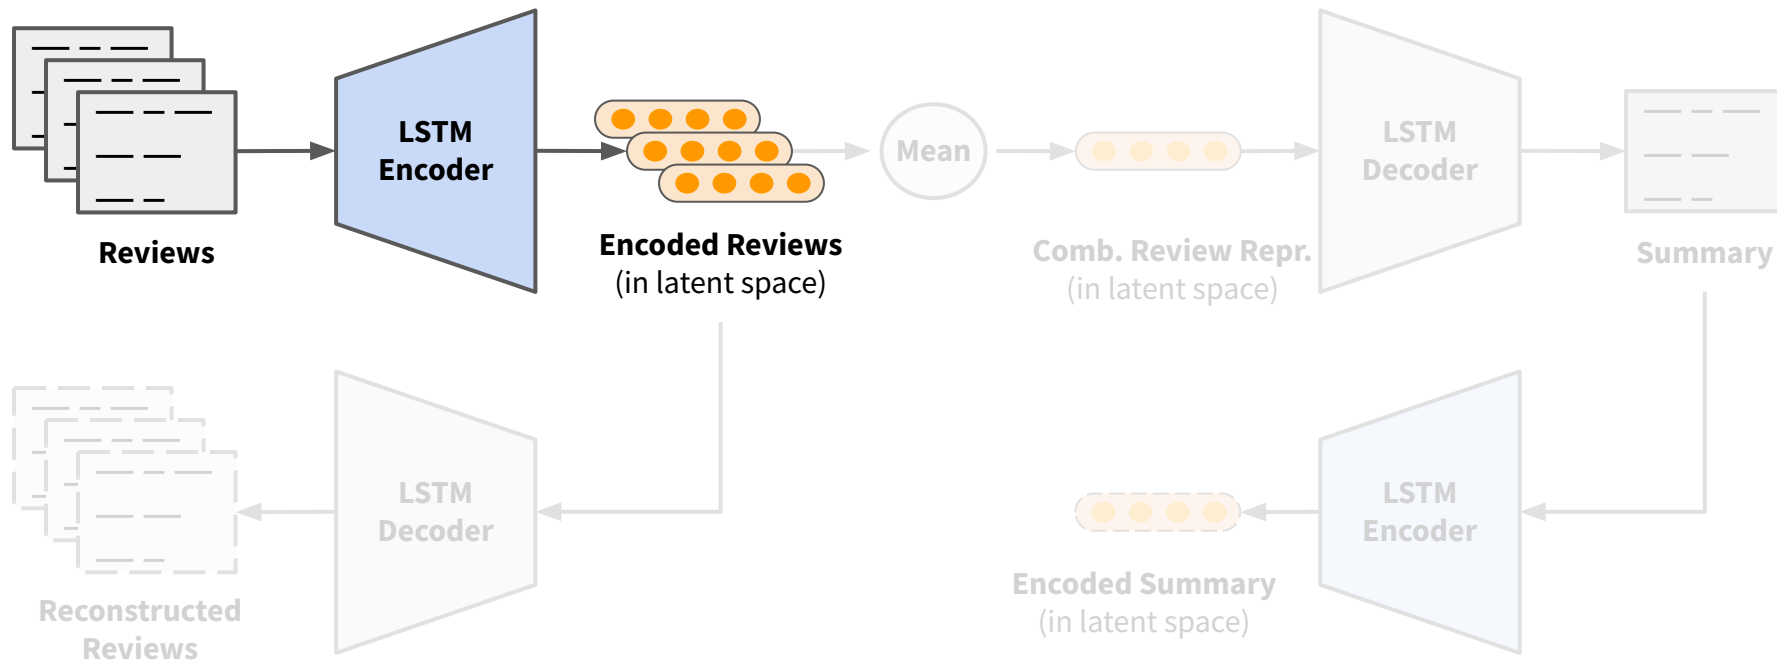
(1). Auto-encoder module

(2). Summarization module

MeanSum



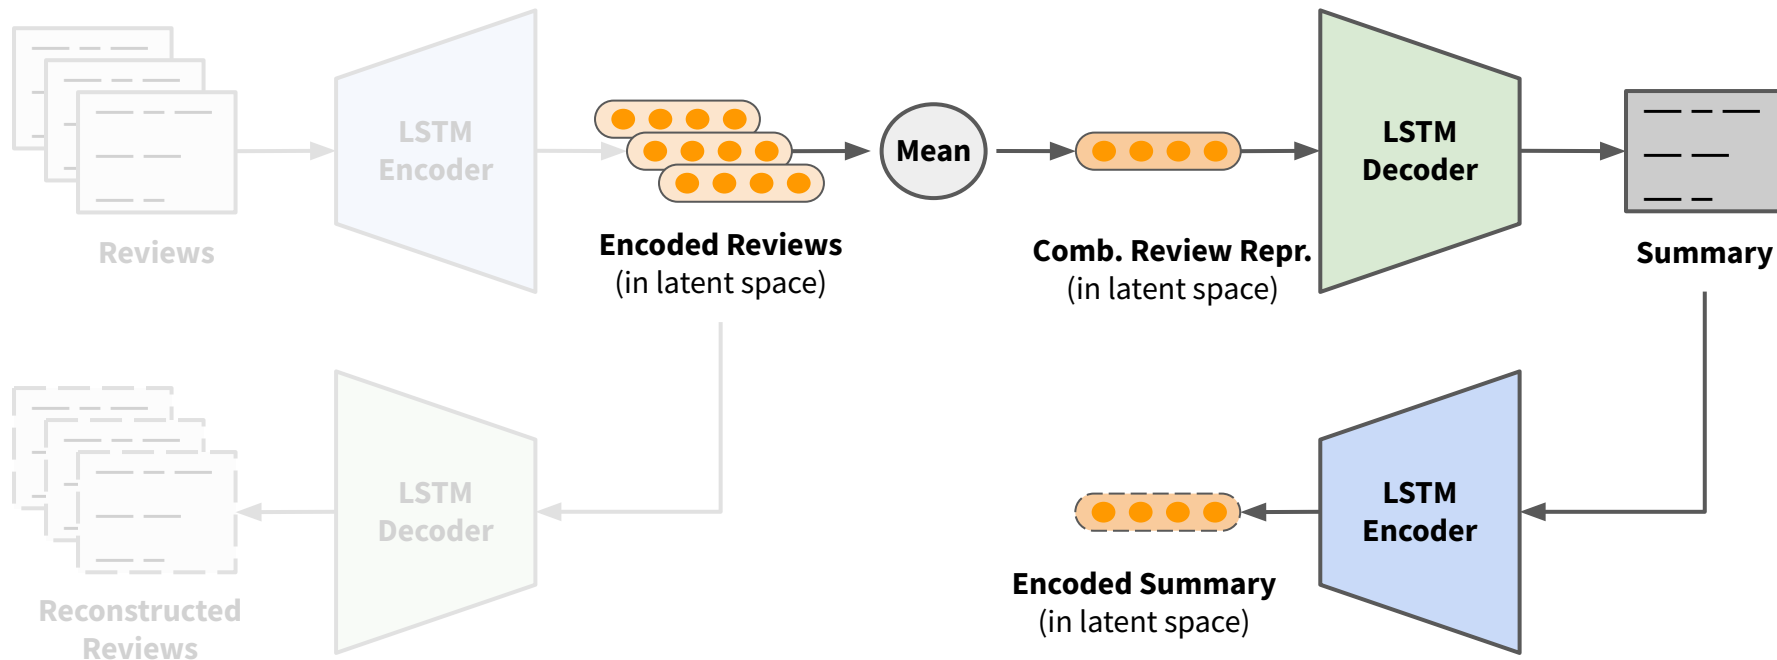
MeanSum



(1). Auto-encoder module

(2). Summarization module

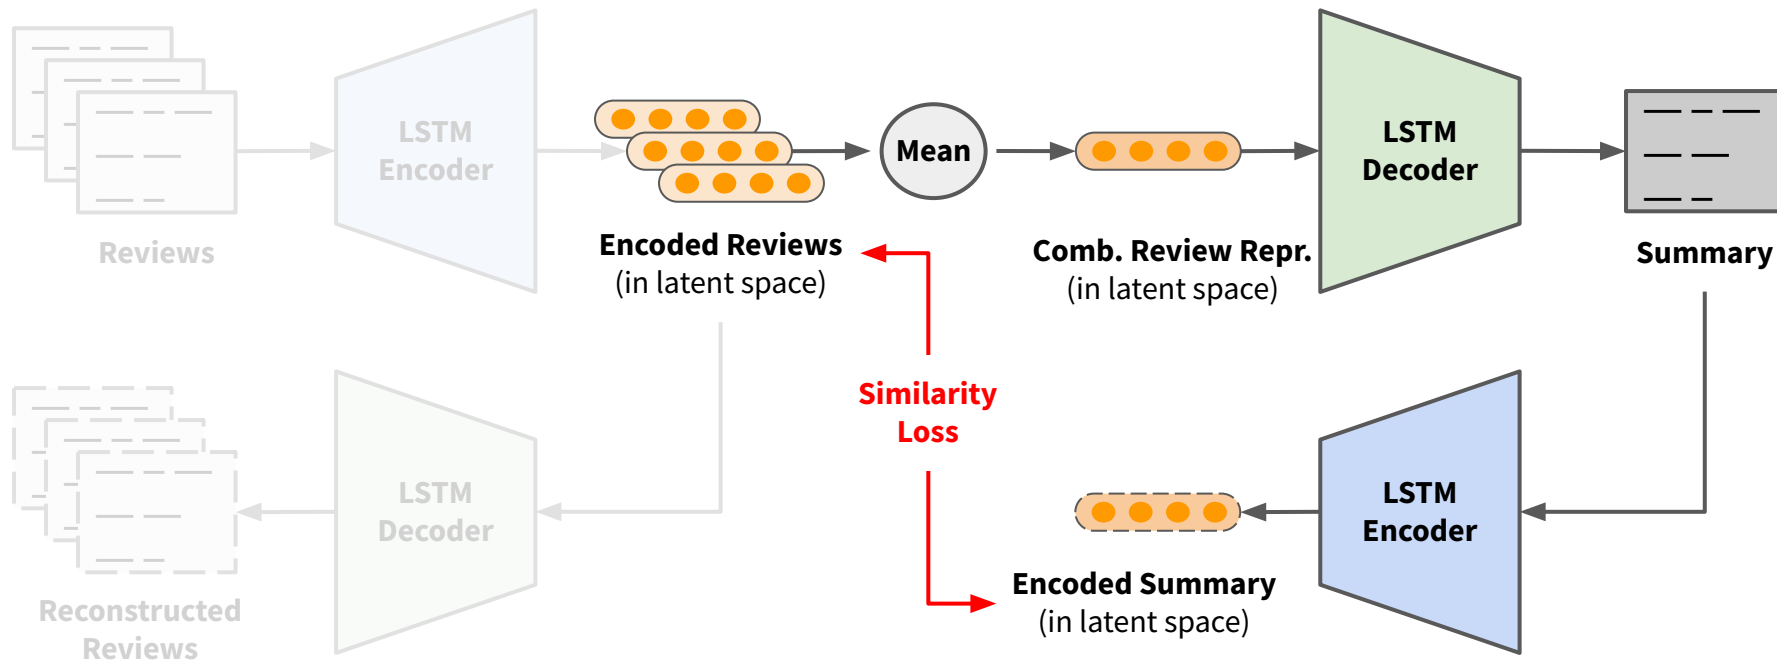
MeanSum



(1). Auto-encoder module

(2). Summarization module

MeanSum

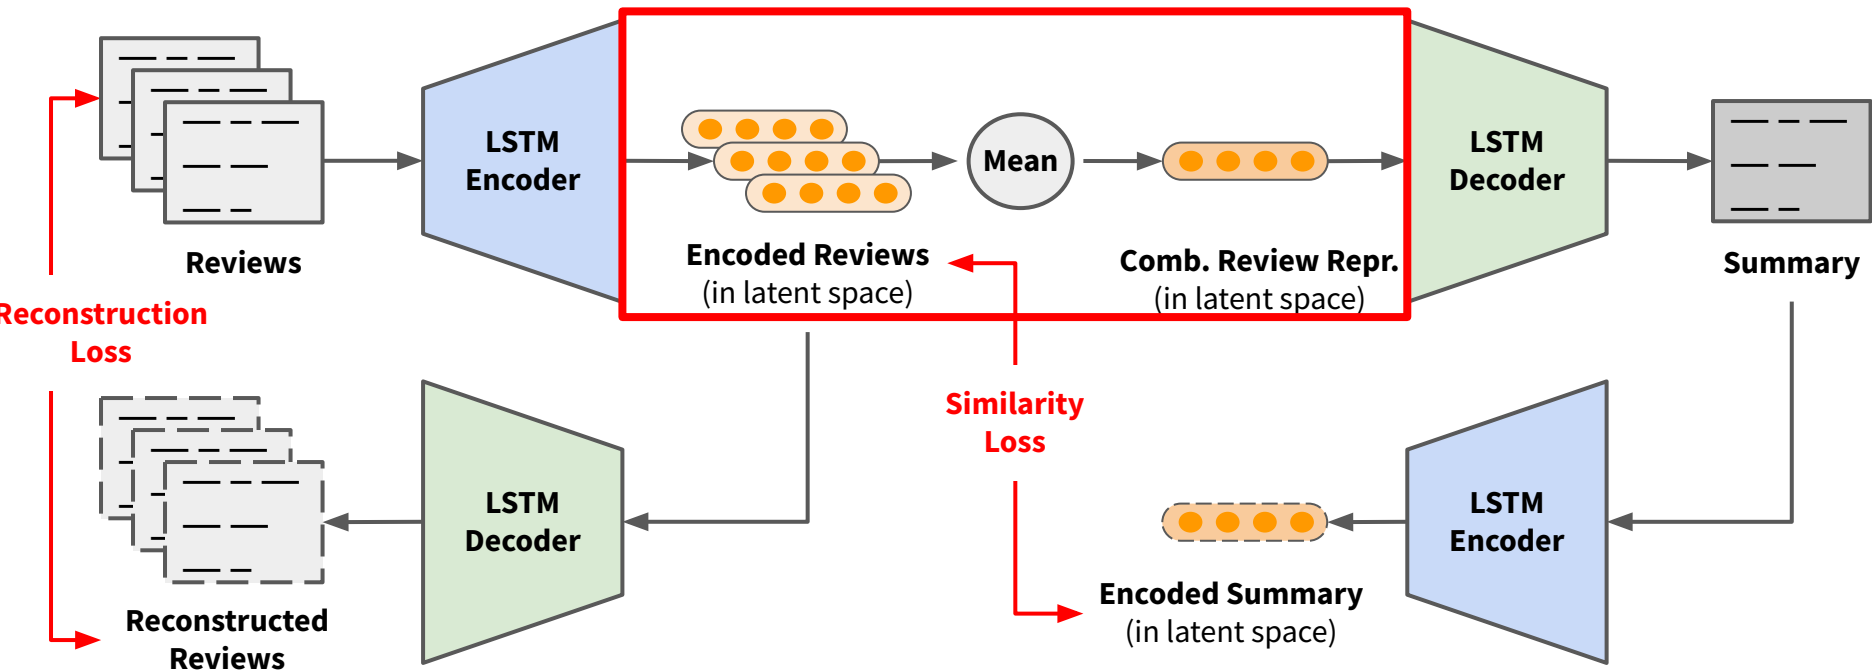


(1). Auto-encoder module

(2). Summarization module

MeanSum

$$\text{Loss} = \text{Reconstruction Loss} + \text{Similarity Loss}$$

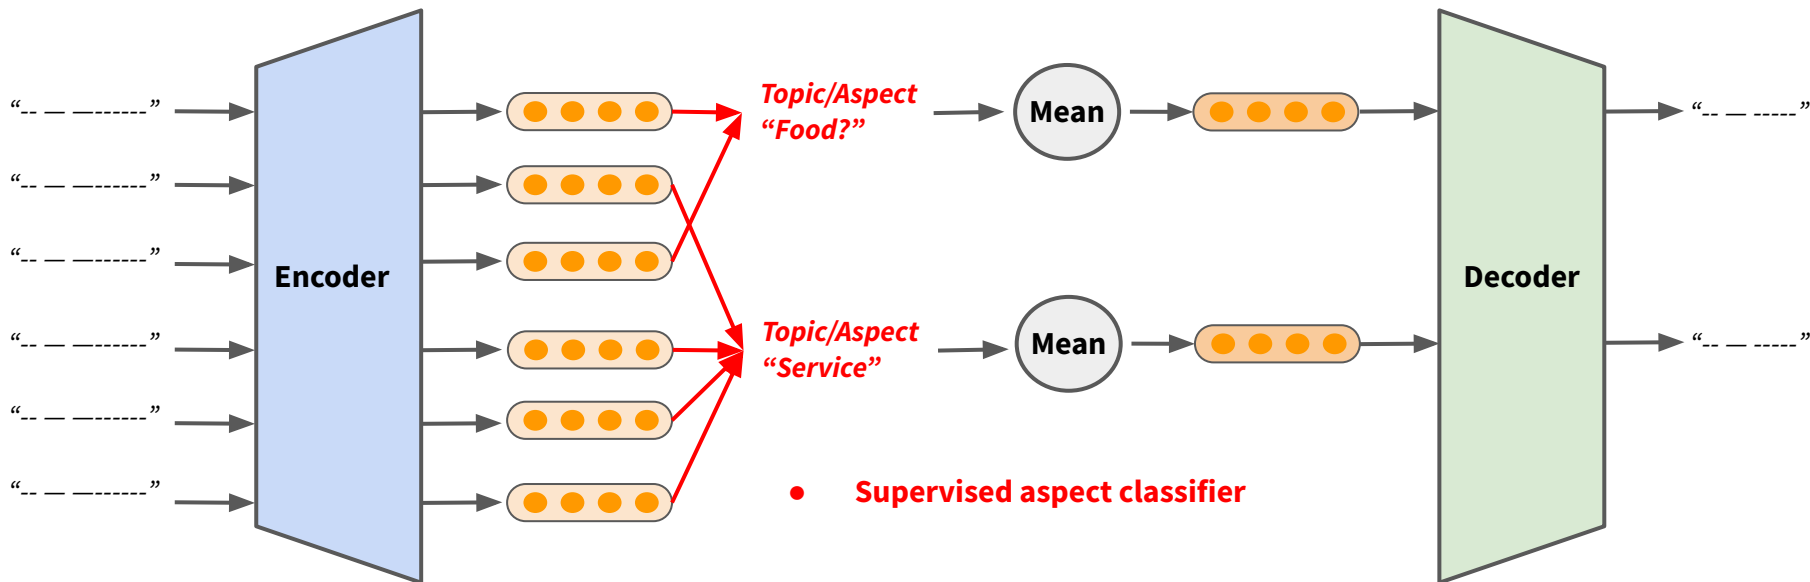


(1). Auto-encoder module

(2). Summarization module

Other works

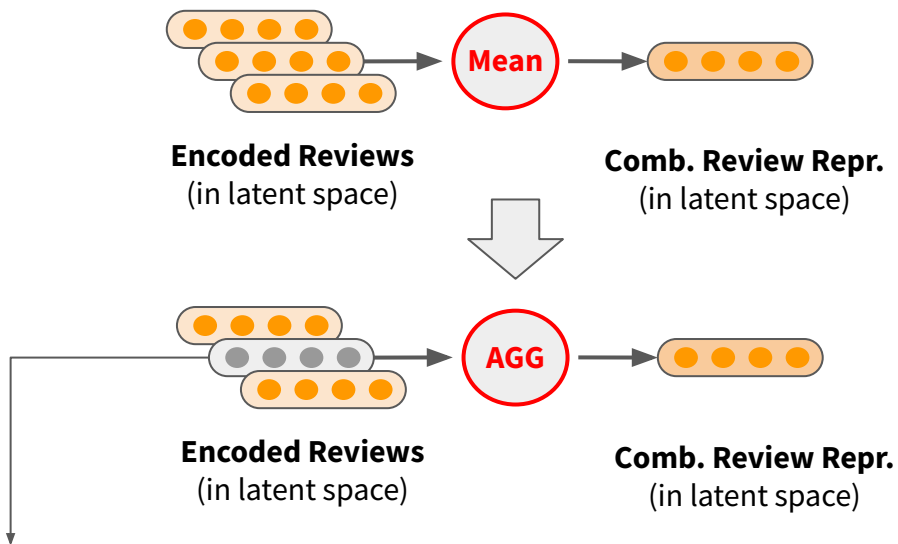
- Sentence representation learning for topic/aspect-aware summarization [Coavoux et.al19]¹



1. Coavoux, Maximin, Hady Elsahar, and Matthias Gallé. "Unsupervised aspect-based multi-document abstractive summarization." In Proceedings of the 2nd Workshop on New Frontiers in Summarization, pp. 42-47. 2019.

Other works

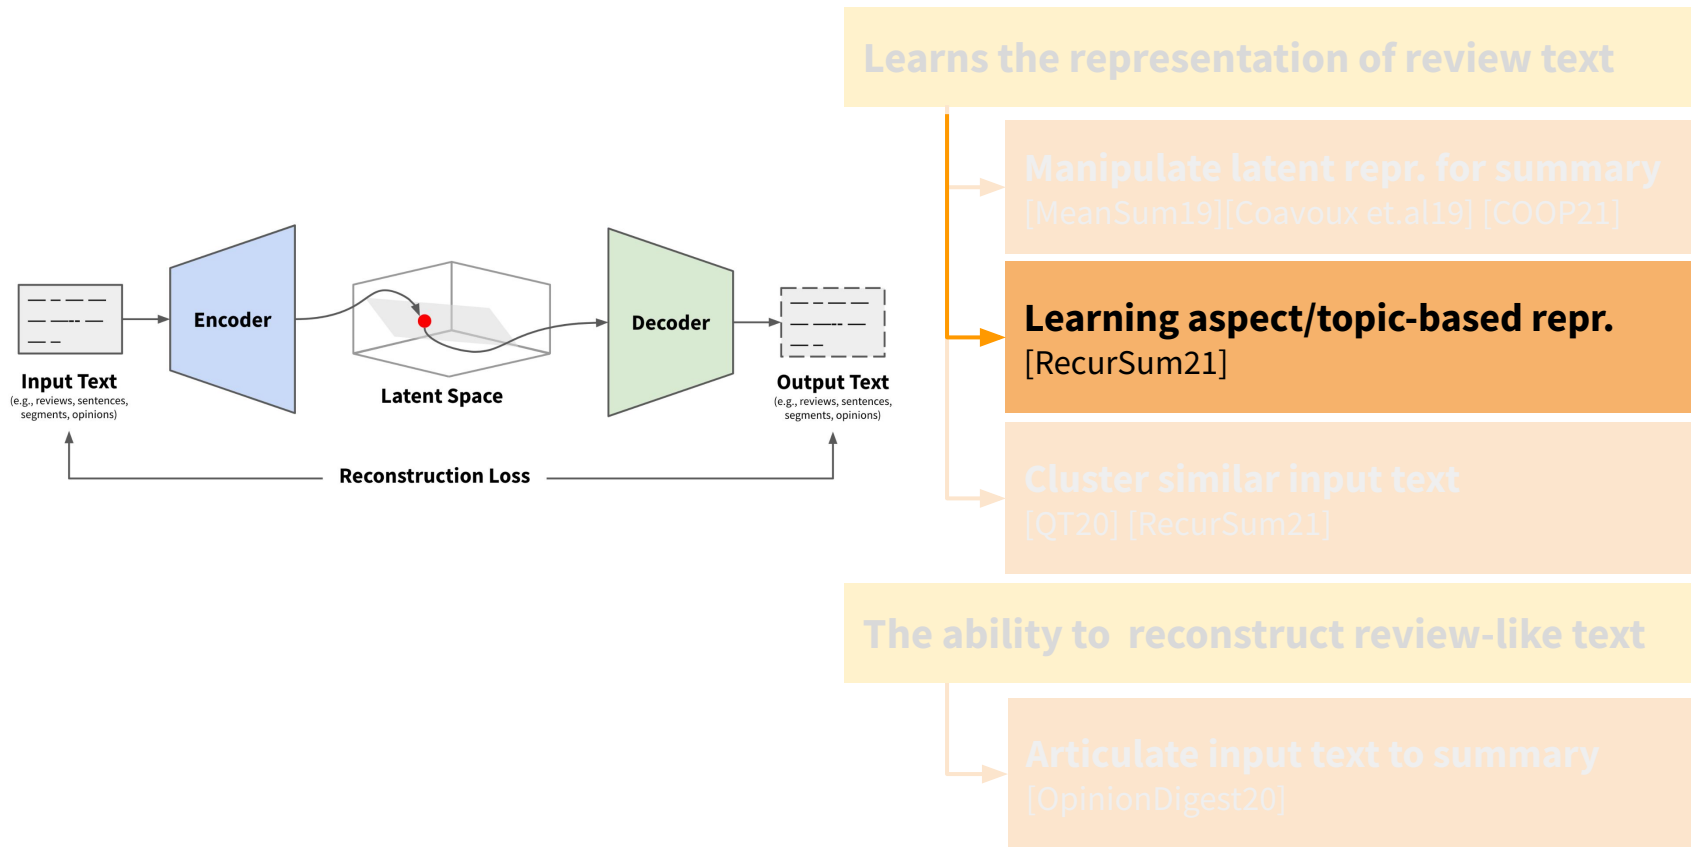
- ~~AE~~ VAE
- Representation aggregation method optimization [COOP21]¹



Select Reviews for Aggregation

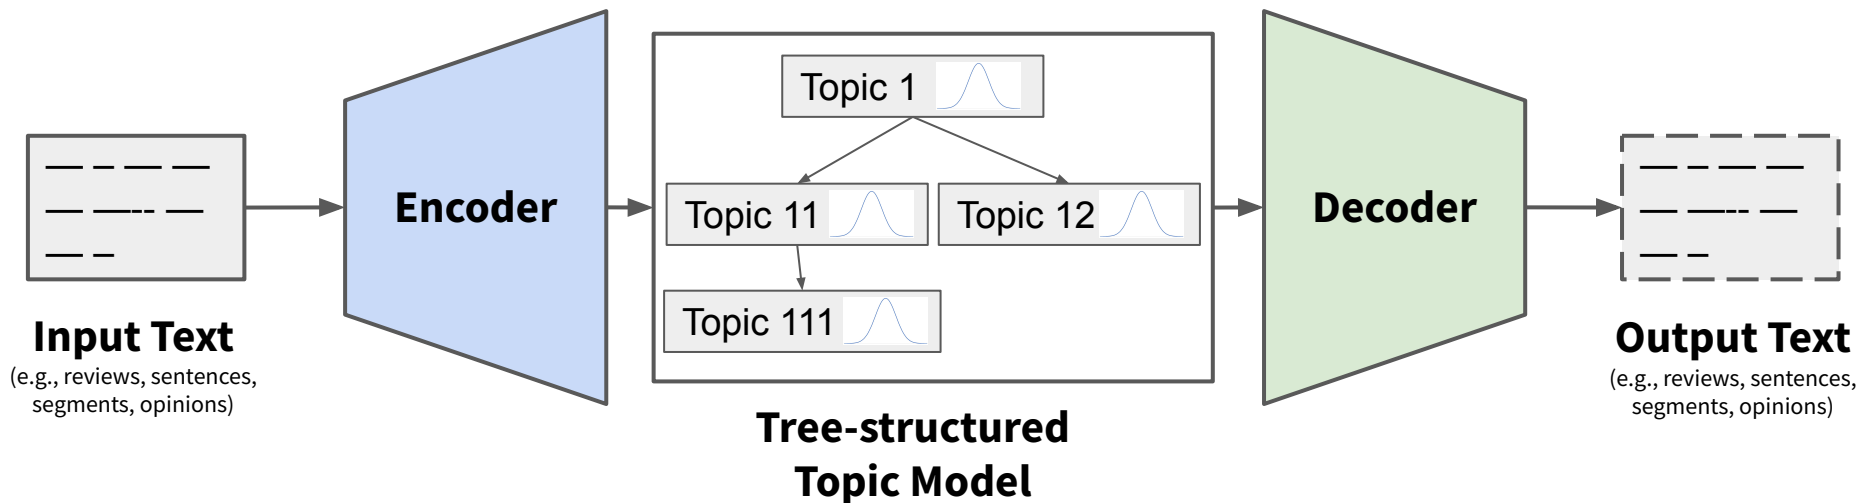
1. Iso, Hayate, Xiaolan Wang, Yoshihiko Suhara, Stefanos Angelidis, and Wang-Chiew Tan. "Convex Aggregation for Opinion Summarization." In Findings of the Association for Computational Linguistics: EMNLP 2021, pp. 3885-3903. 2021.

Opinion Summarization via Autoencoder



RecurSum¹


Use Variational-Autoencoder (VAE) to learn a recursive topic model



1. Isonuma, Masaru, Junichiro Mori, Danushka Bollegala, and Ichiro Sakata. "Unsupervised abstractive opinion summarization by generating sentences with tree-structured topic guidance." *Transactions of the Association for Computational Linguistics* 9 (2021): 945-961.
2. <https://underline.io/lecture/40739-unsupervised-abstractive-opinion-summarization-by-generating-sentences-with-tree-structured-topic-guidance>

RecurSum

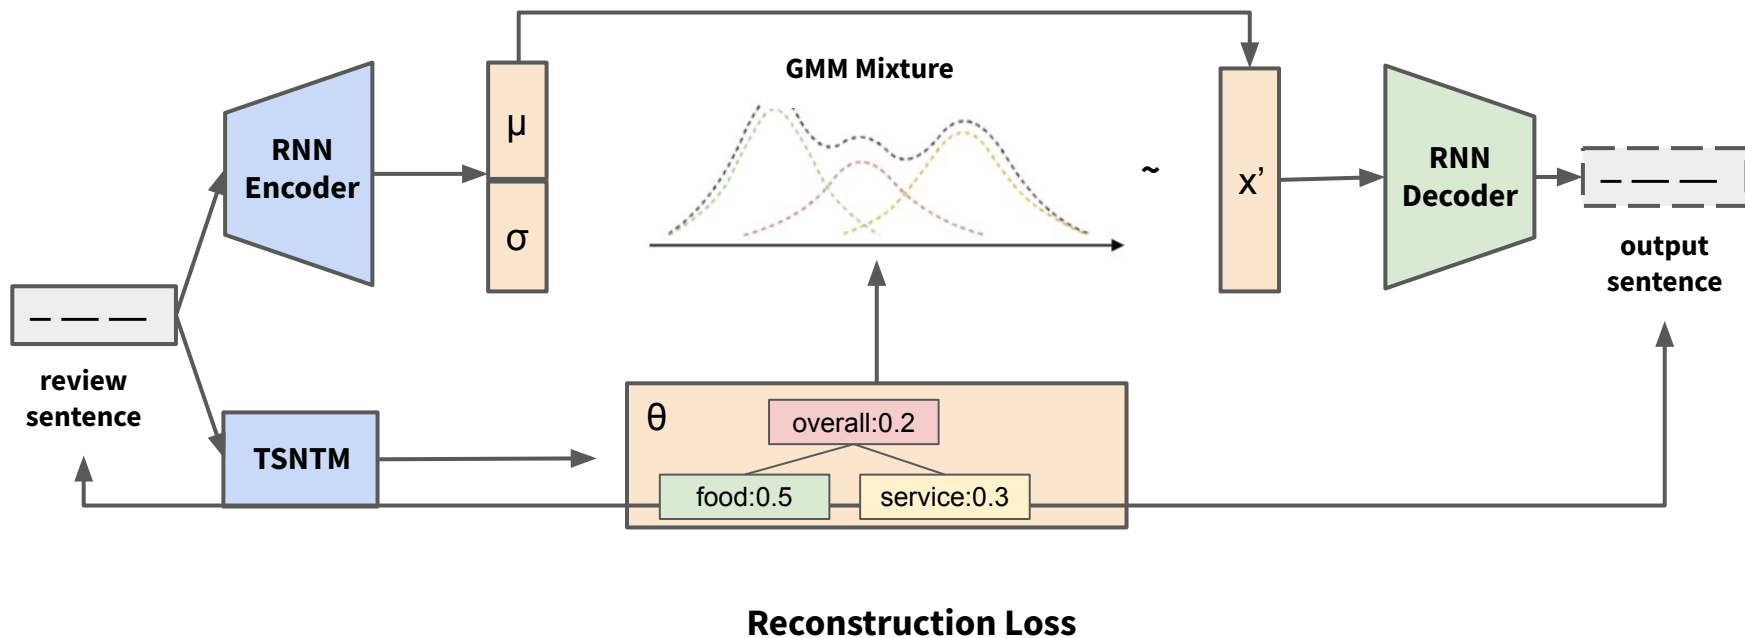
Training: learning tree-structured topic gaussians via reconstruction loss



Reconstruction Loss

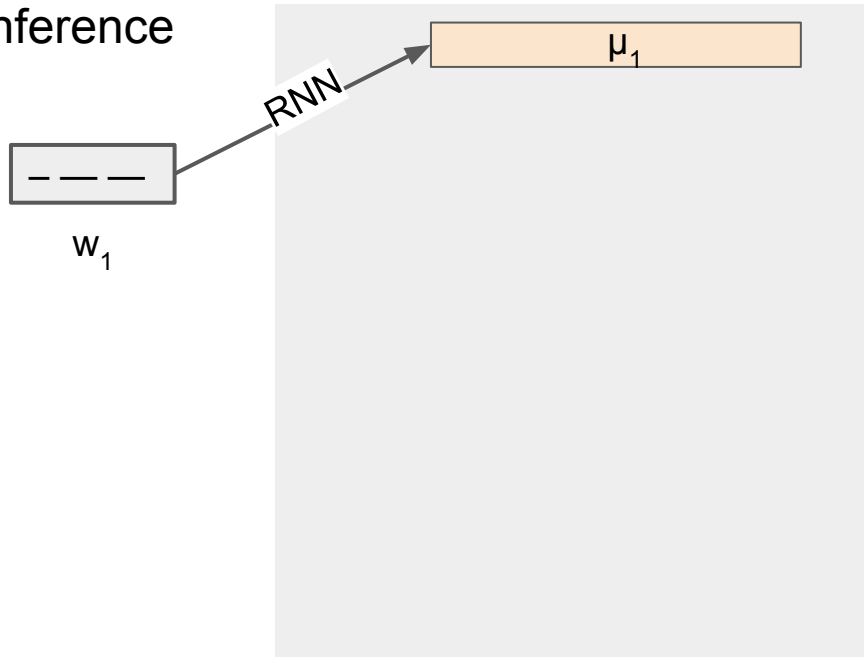
RecurSum

Training: learning tree-structured topic gaussians via reconstruction loss



RecurSum

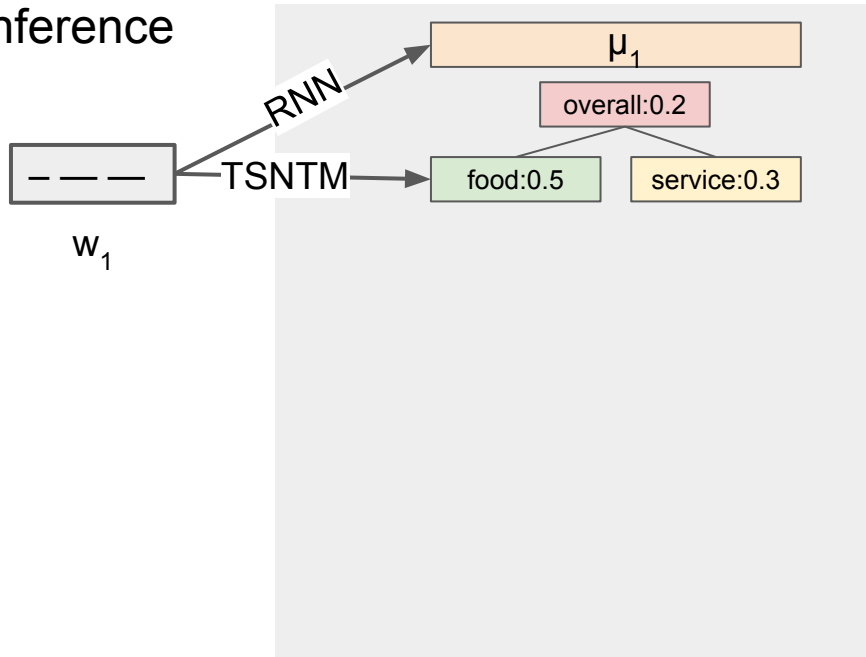
Inference



From RNN Encoder & TSNTM

RecurSum


Inference



From RNN Encoder & TSNTM

RecurSum

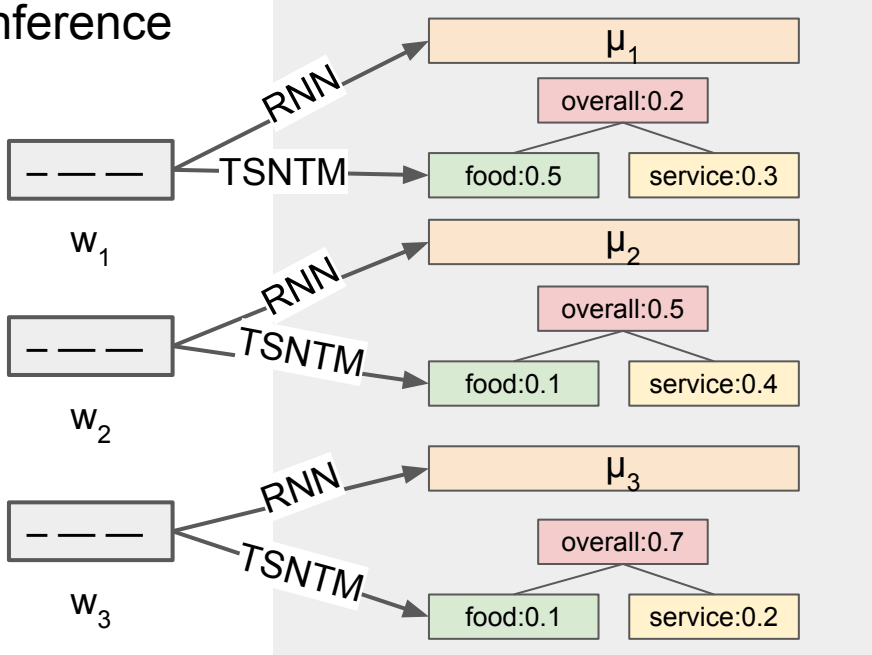
Inference




From RNN Encoder & TSNTM

RecurSum

Inference

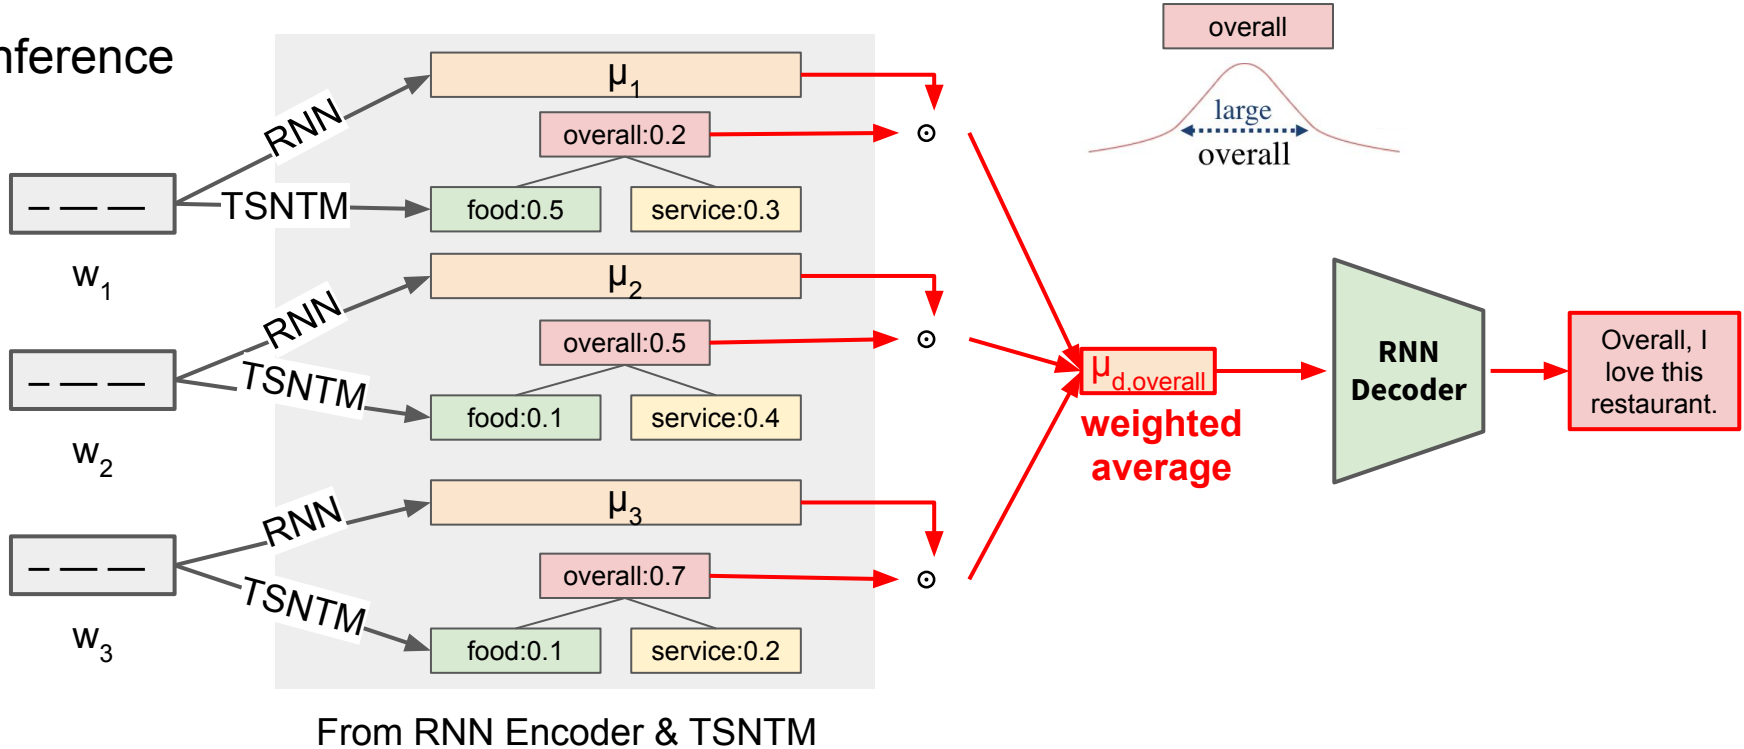


From RNN Encoder & TSNTM



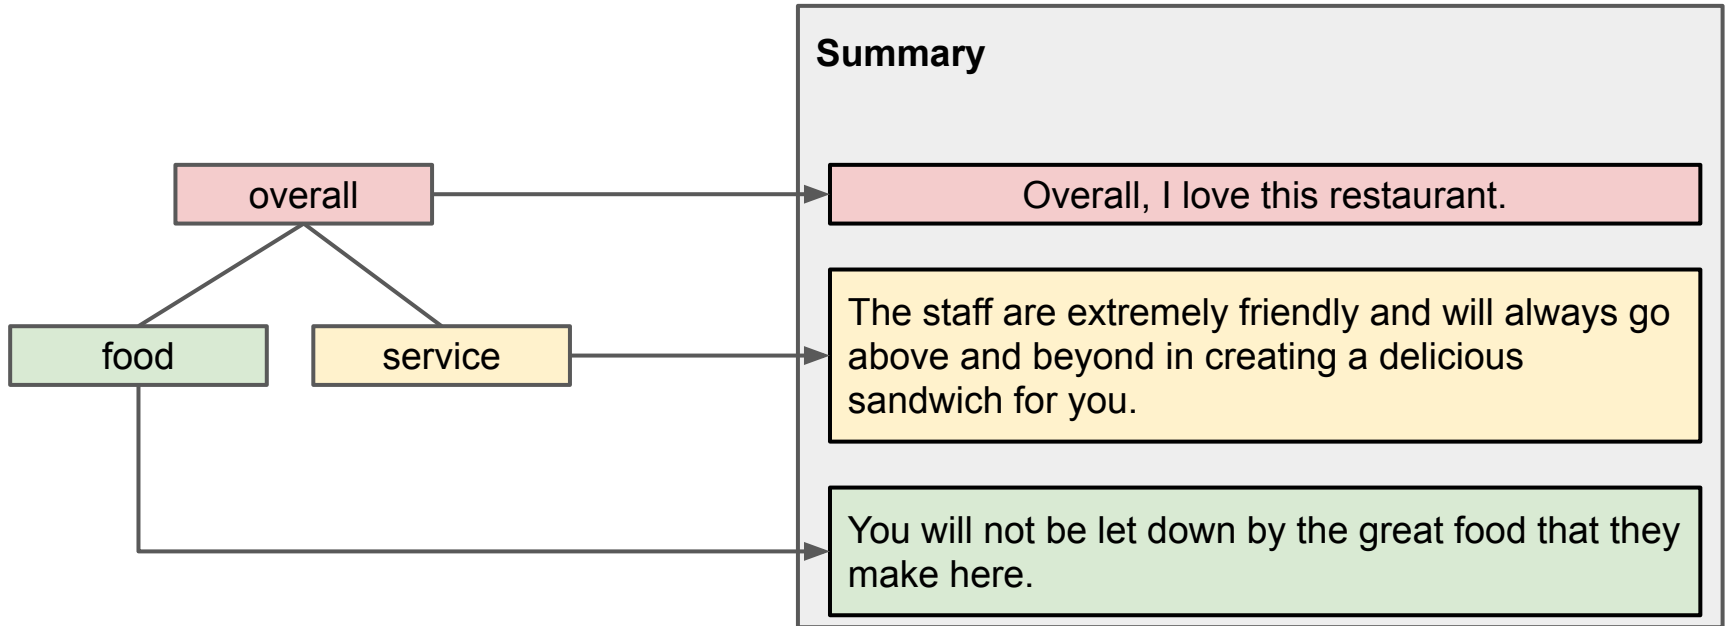
RecurSum

Inference

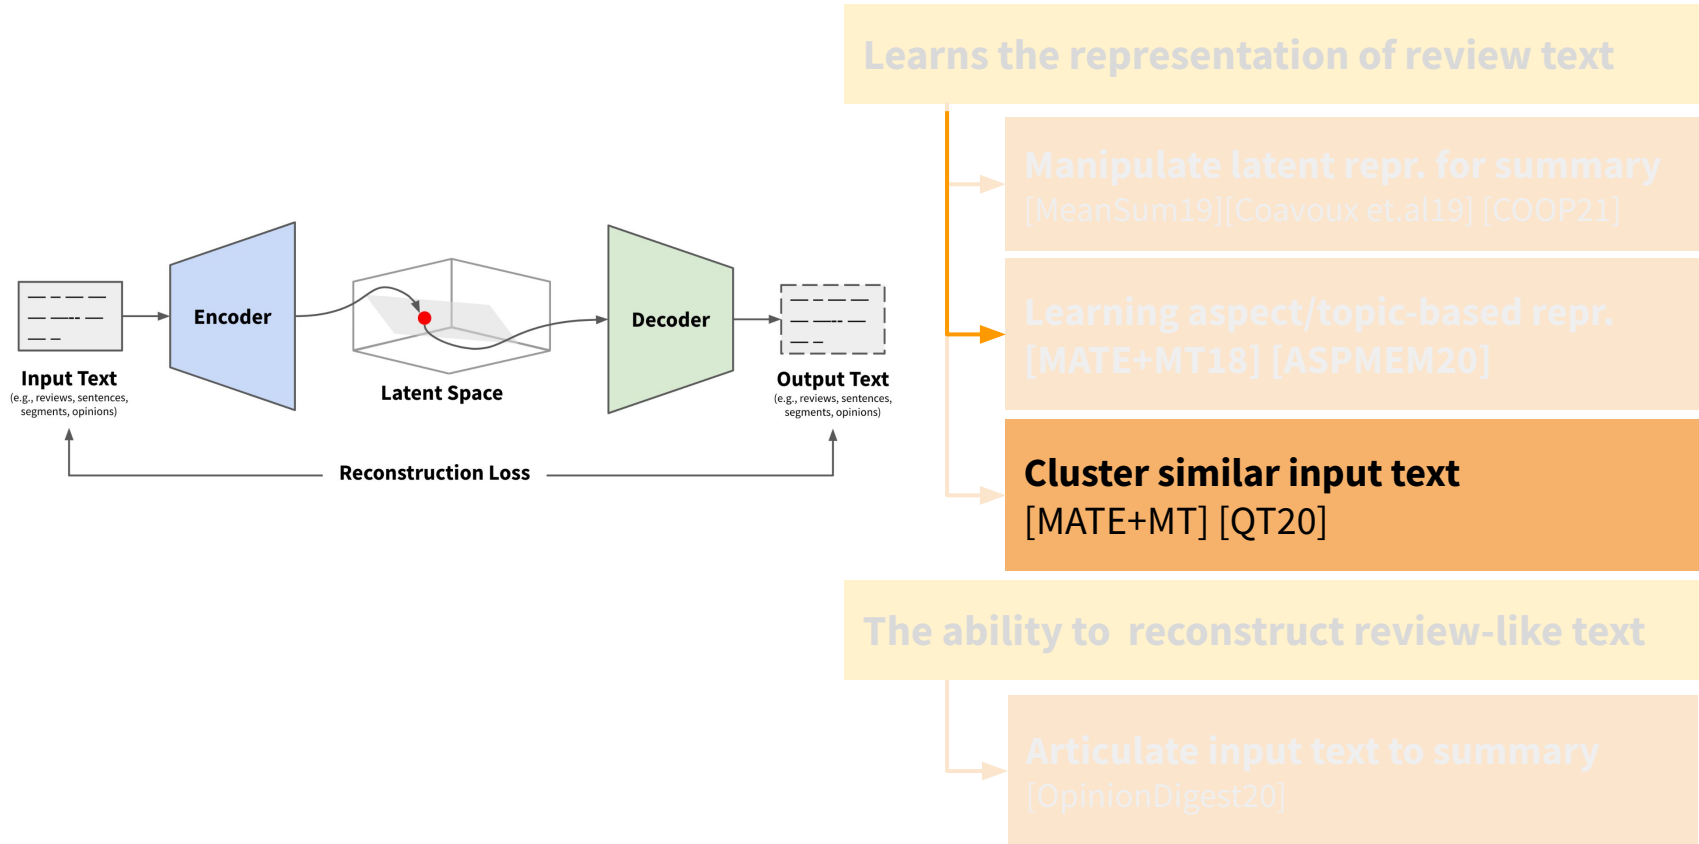


RecurSum

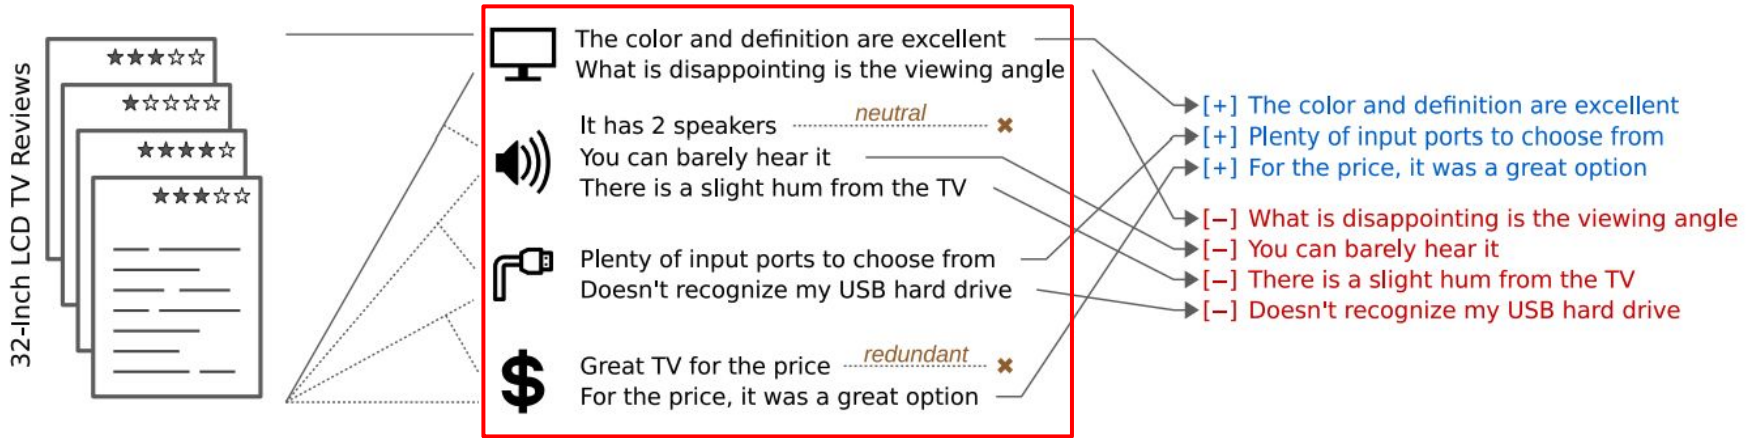
Summary Generation



Opinion Summarization via Autoencoder




MATE+MT¹



Use aspect-based autoencoder (ABAE)² to extract aspect-specific review segment

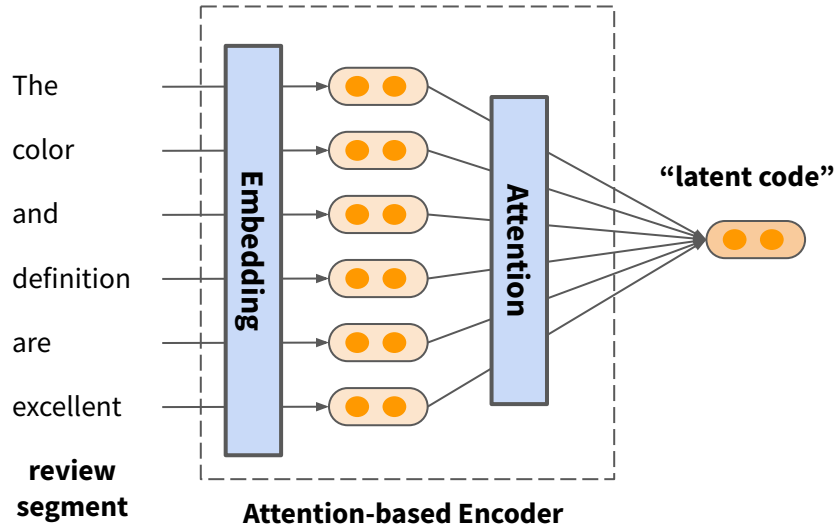
1. Angelidis, Stefanos, and Mirella Lapata. "Summarizing Opinions: Aspect Extraction Meets Sentiment Prediction and They Are Both Weakly Supervised." In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 3675-3686. 2018.
2. He, Ruidan, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. "An unsupervised neural attention model for aspect extraction." ACL, pp. 388-397. 2017.

Aspect-based autoencoder (ABAE)

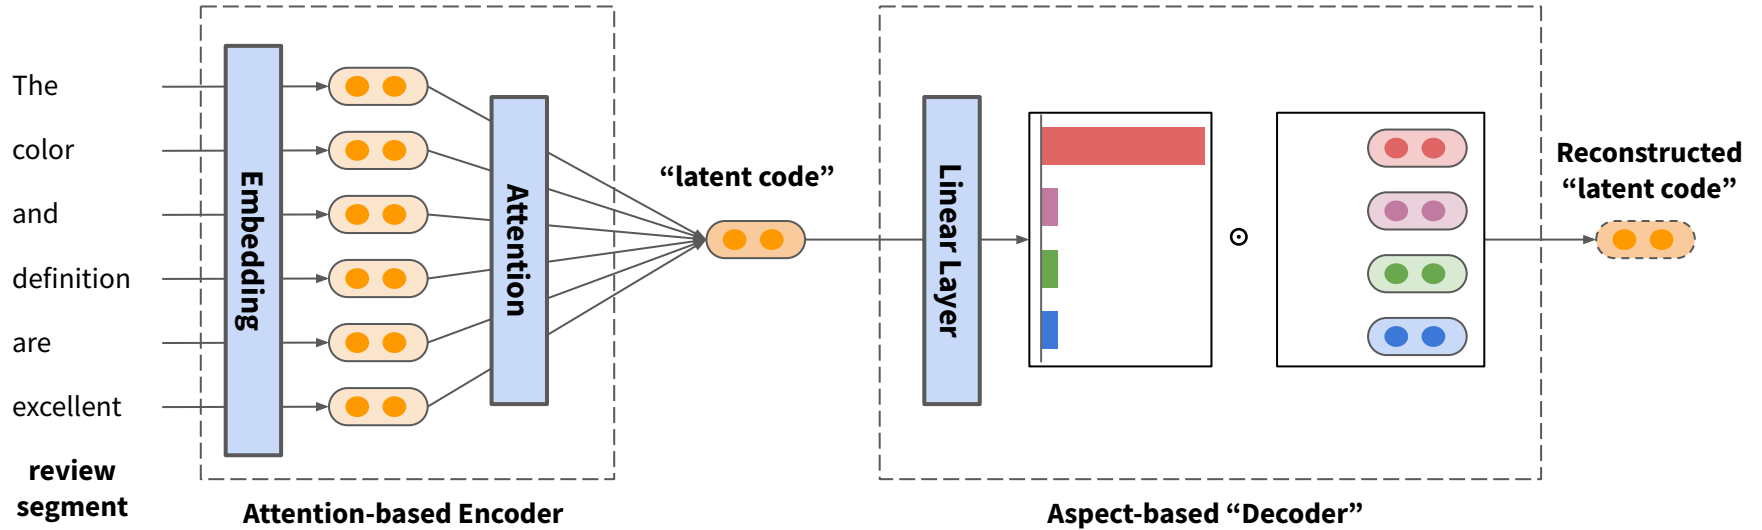


Aspect embeddings

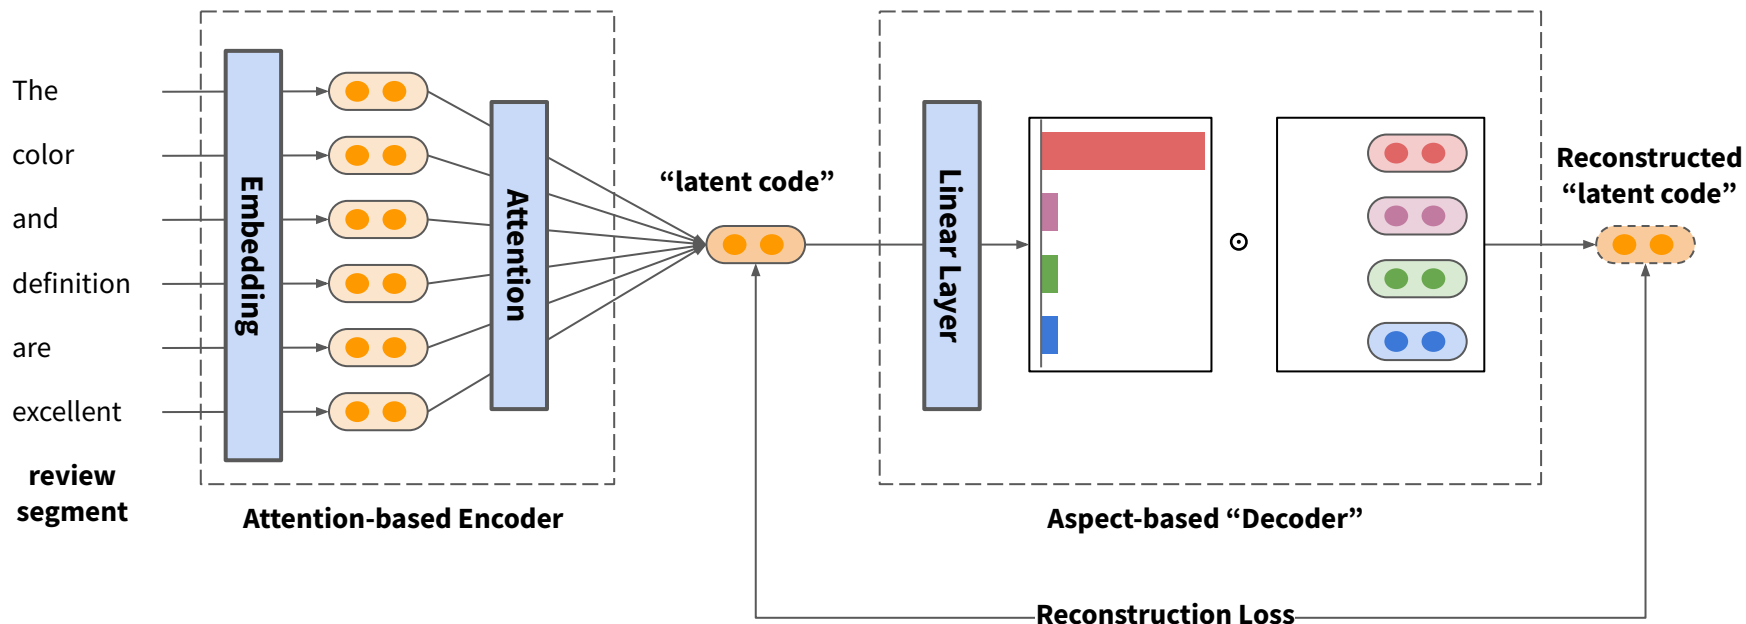
Aspect-based autoencoder (ABAE)





Aspect-based autoencoder (ABAE)

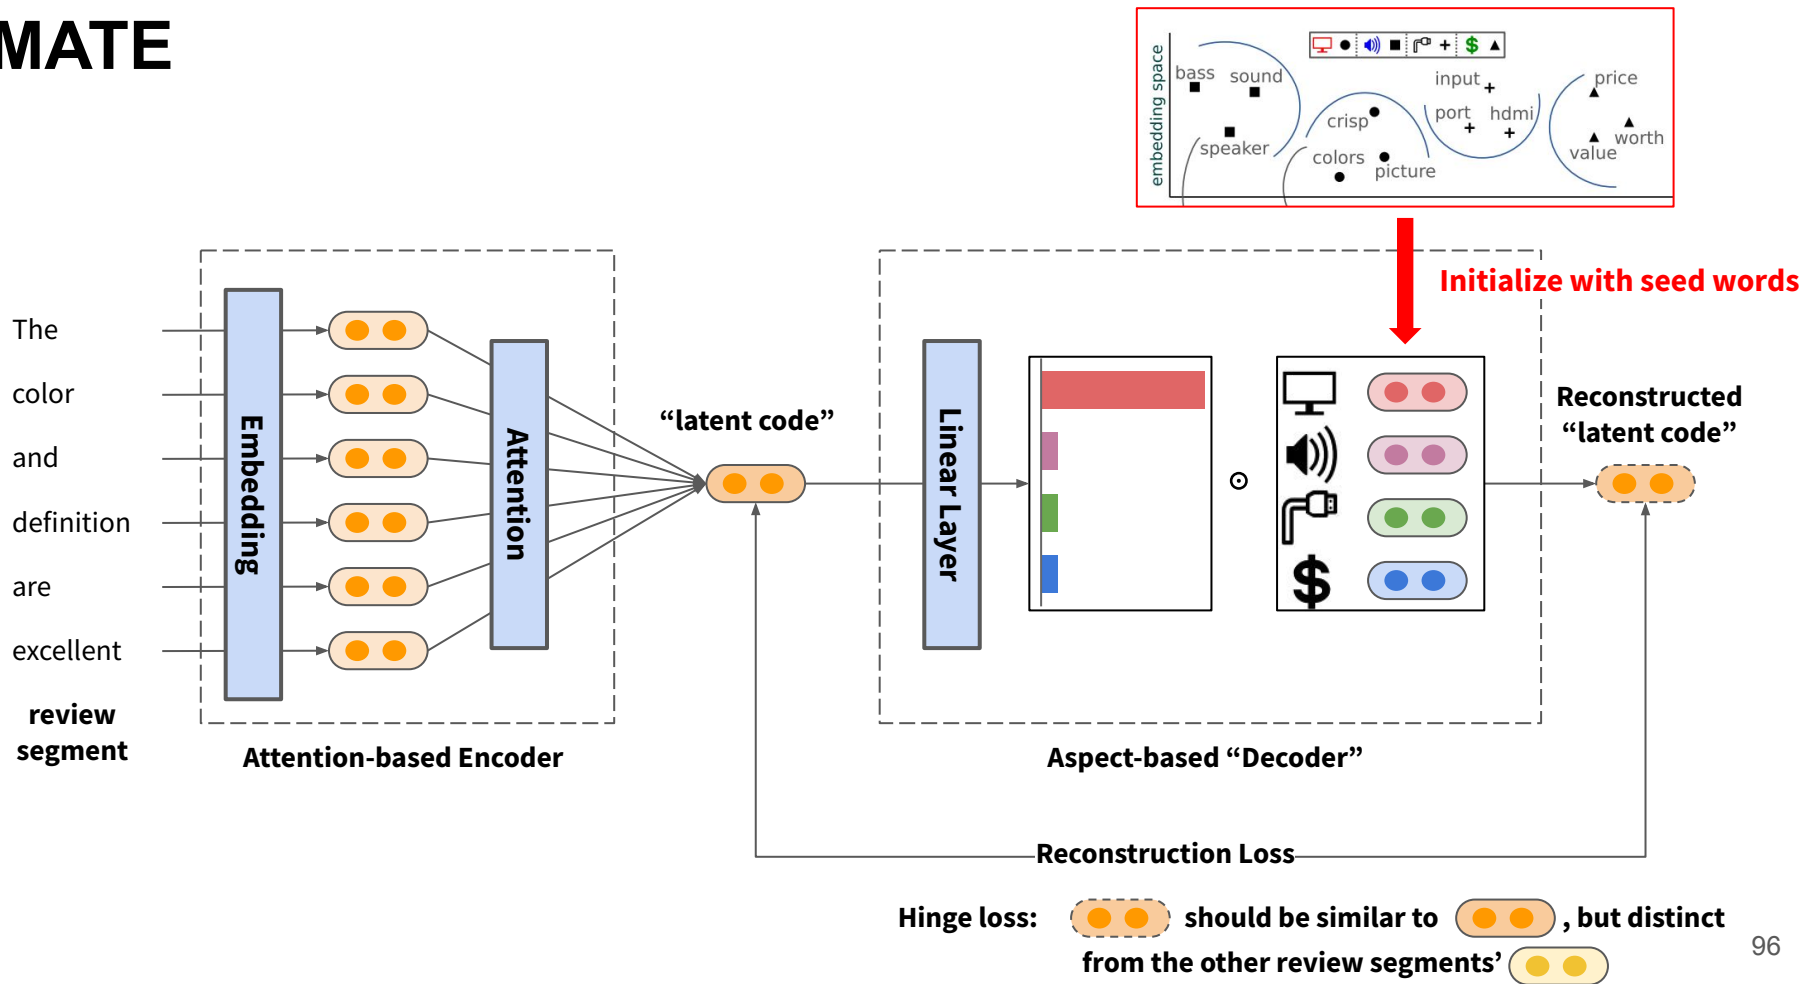


Aspect-based autoencoder (ABAE)

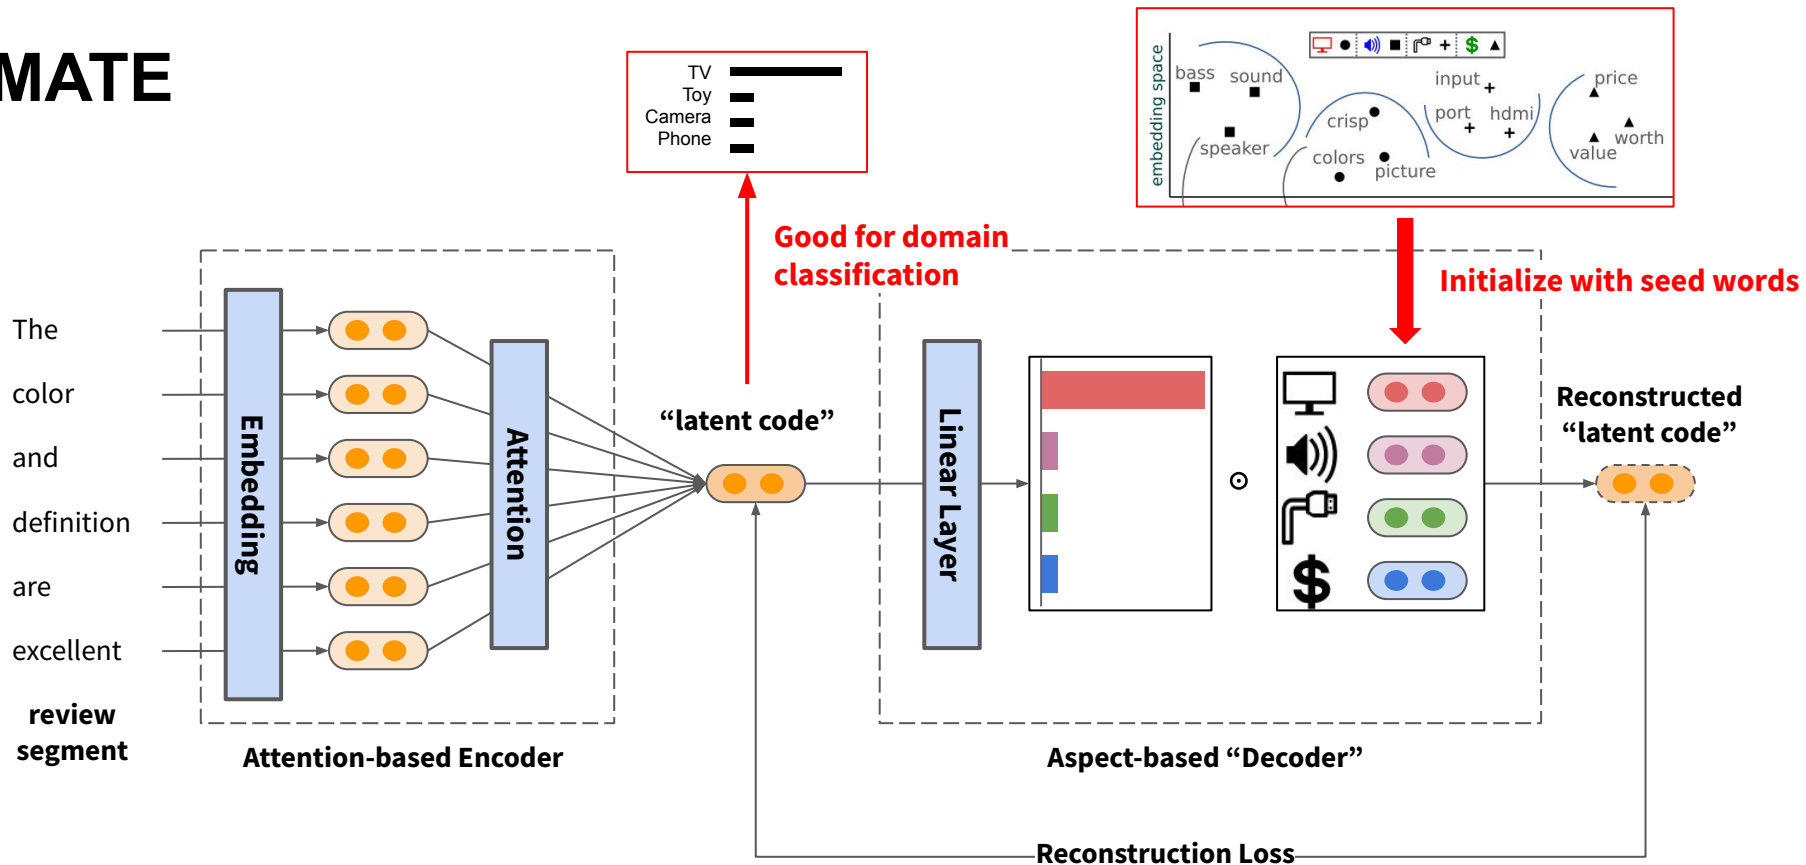


Hinge loss:  should be similar to , but distinct from the other review segments' 

MATE

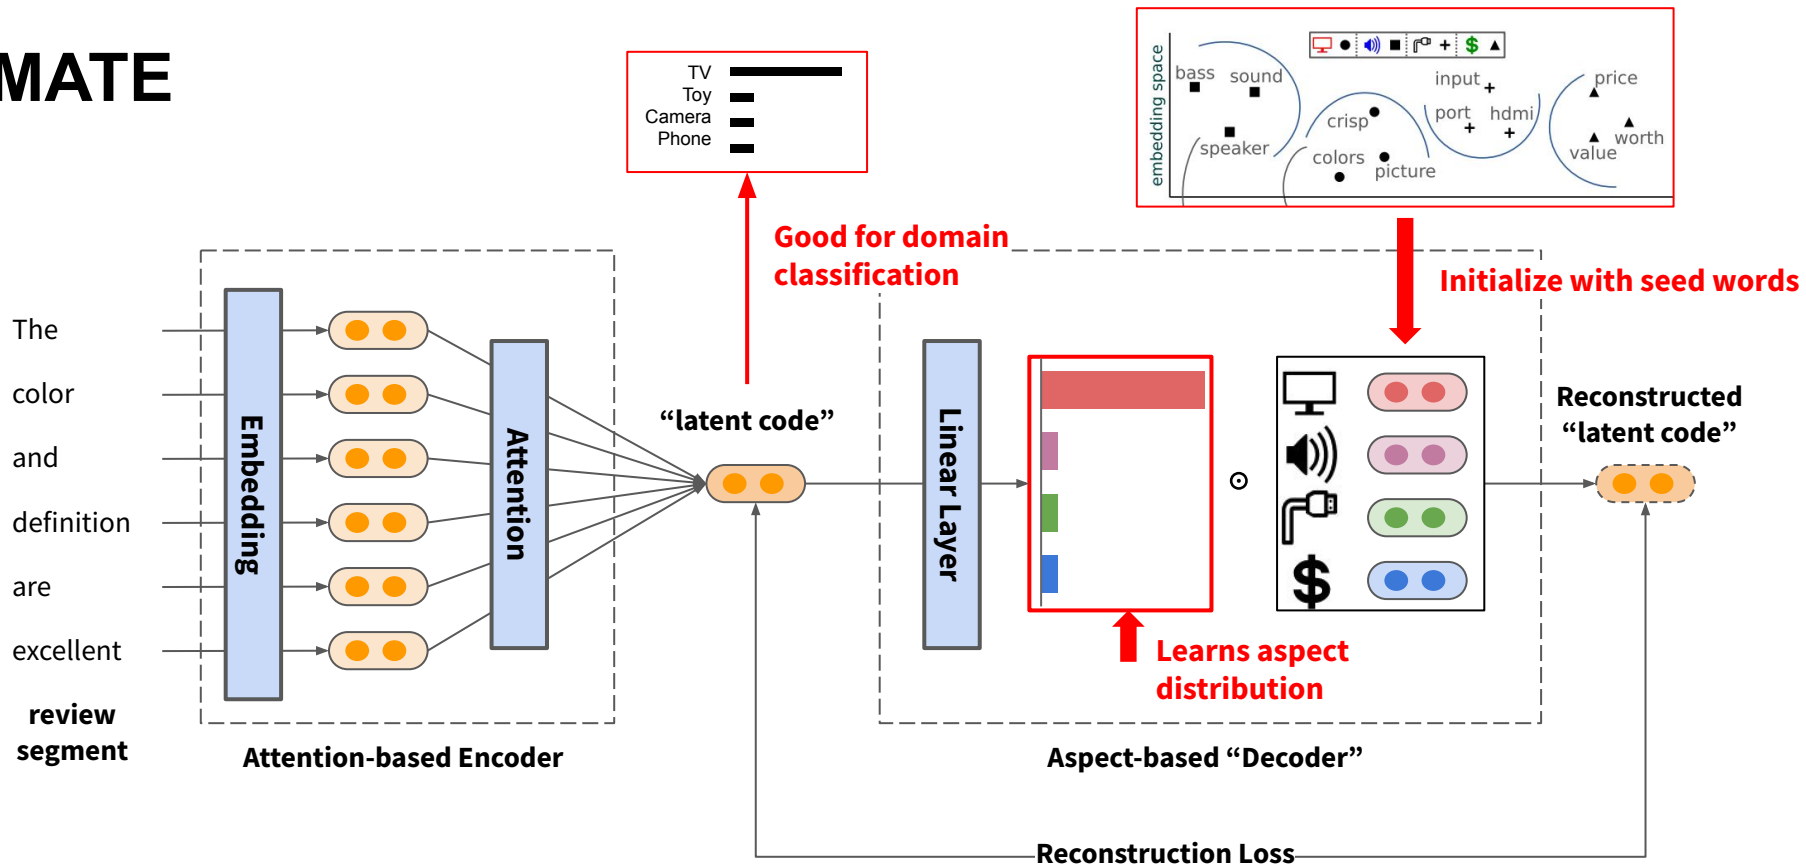


MATE



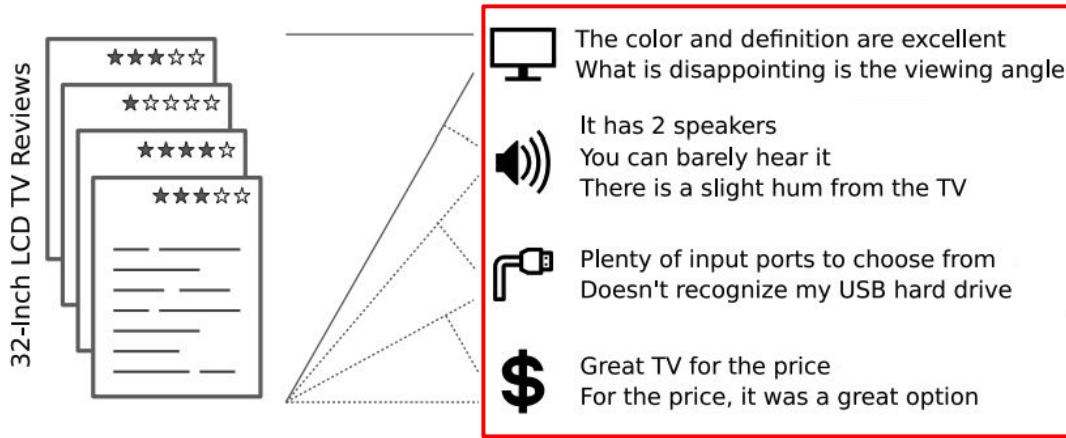
Hinge loss: should be similar to , but distinct from the other review segments'

MATE



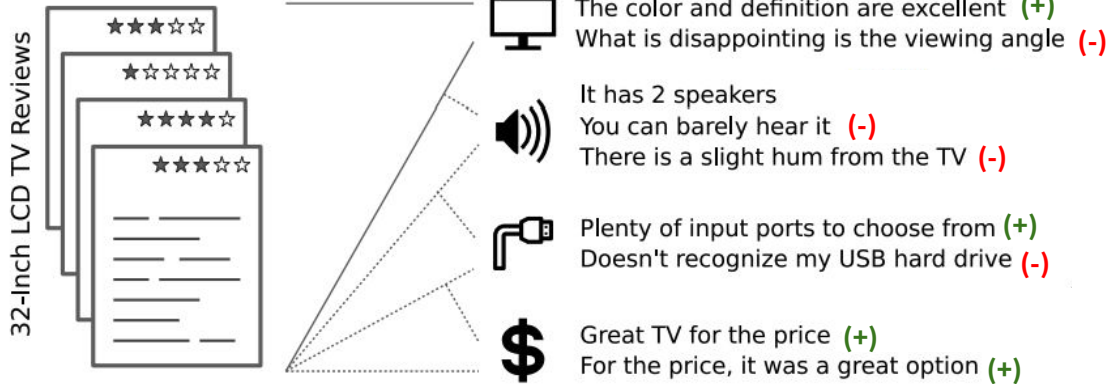
Hinge loss: should be similar to , but distinct from the other review segments'

MATE+MT



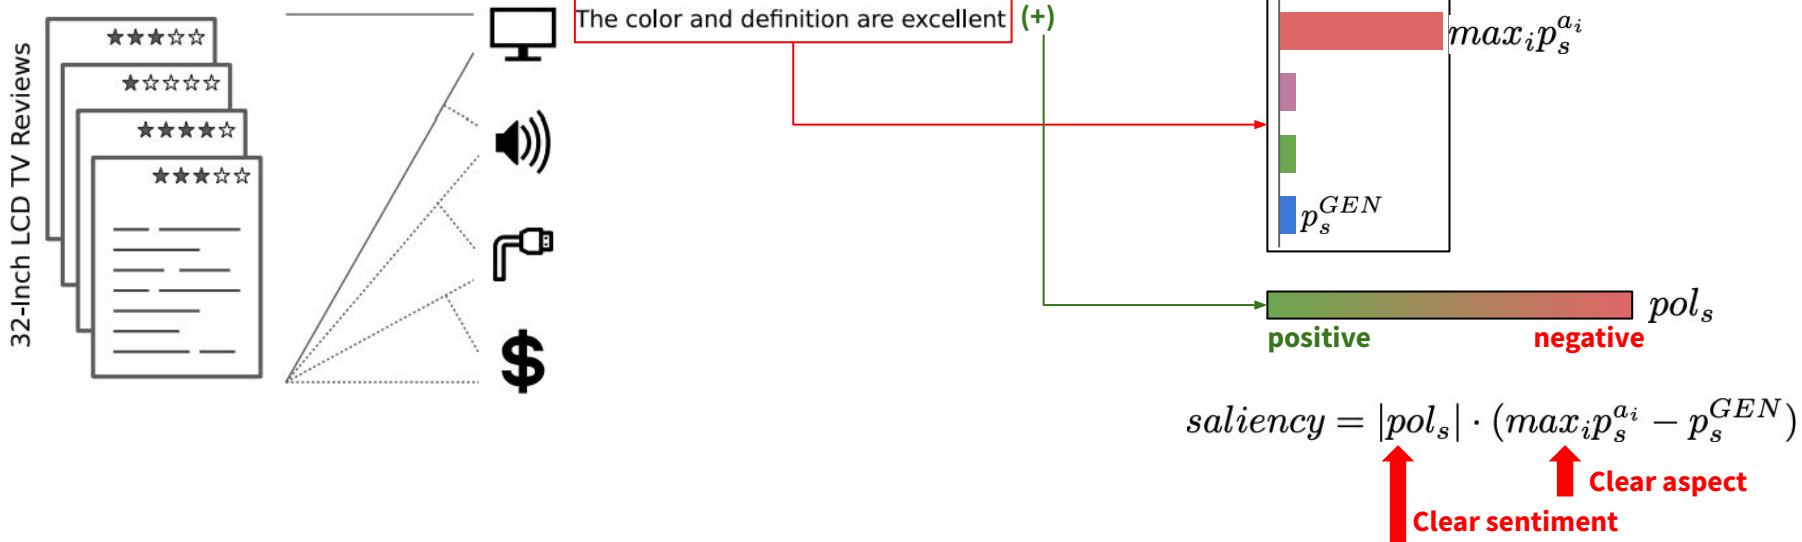
Use aspect distribution prediction to group review segments

MATE+MT




↑
Use weakly supervised classifier to provide sentiment labels

MATE+MT



QT

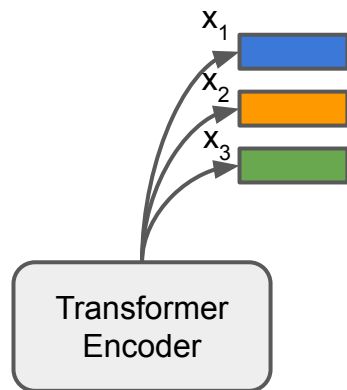
Training




1. Angelidis, Stefanos, Reinald Kim Amplayo, Yoshihiko Suhara, Xiaolan Wang, and Mirella Lapata. "Extractive opinion summarization in quantized transformer spaces." Transactions of the Association for Computational Linguistics 9 (2021): 277-293.

QT

Training




Original:
The staff was great!



Latent Space
(discrete codebook)

QT

Training




Original:
The staff was great!

Latent Space
(discrete codebook)

QT

Training




Original:
The staff was great!

Latent Space
(discrete codebook)

QT

Training

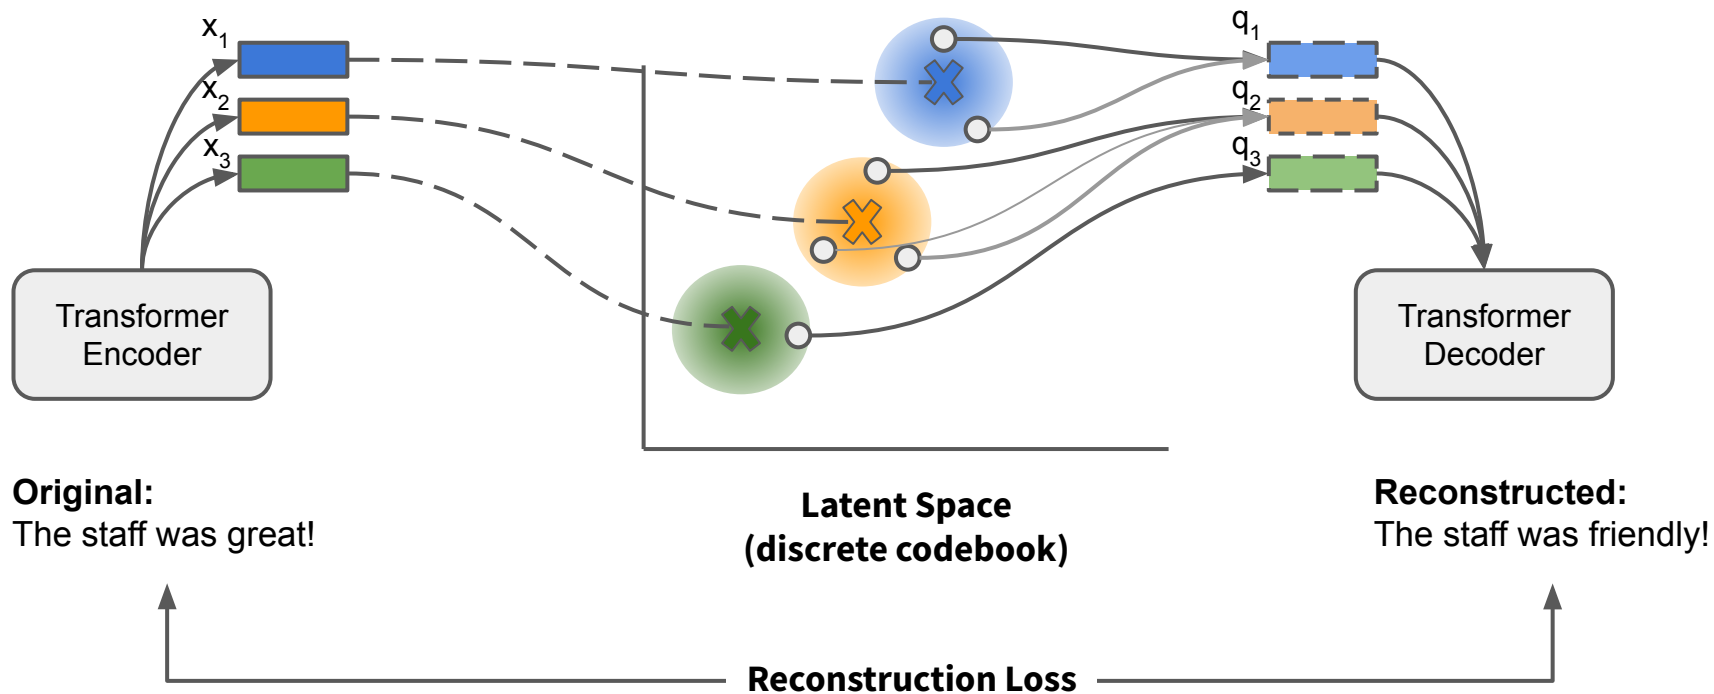


Original:
The staff was great!

**Latent Space
(discrete codebook)**

QT




Training




QT

Inference

Sentences of an entity:





Latent Space
(discrete codebook)



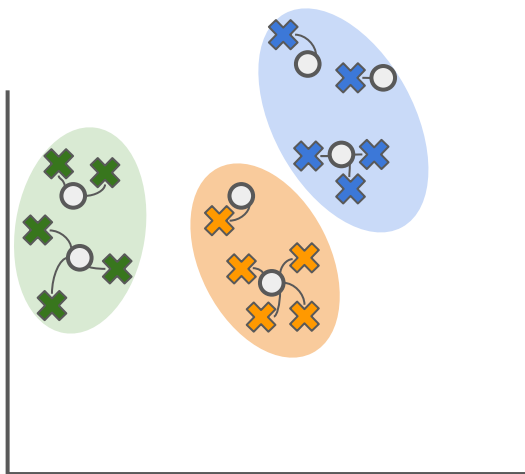
QT

Inference


Sentences of an entity:



Use seed words to find aspect-specific codes






Latent Space
(discrete codebook)



QT

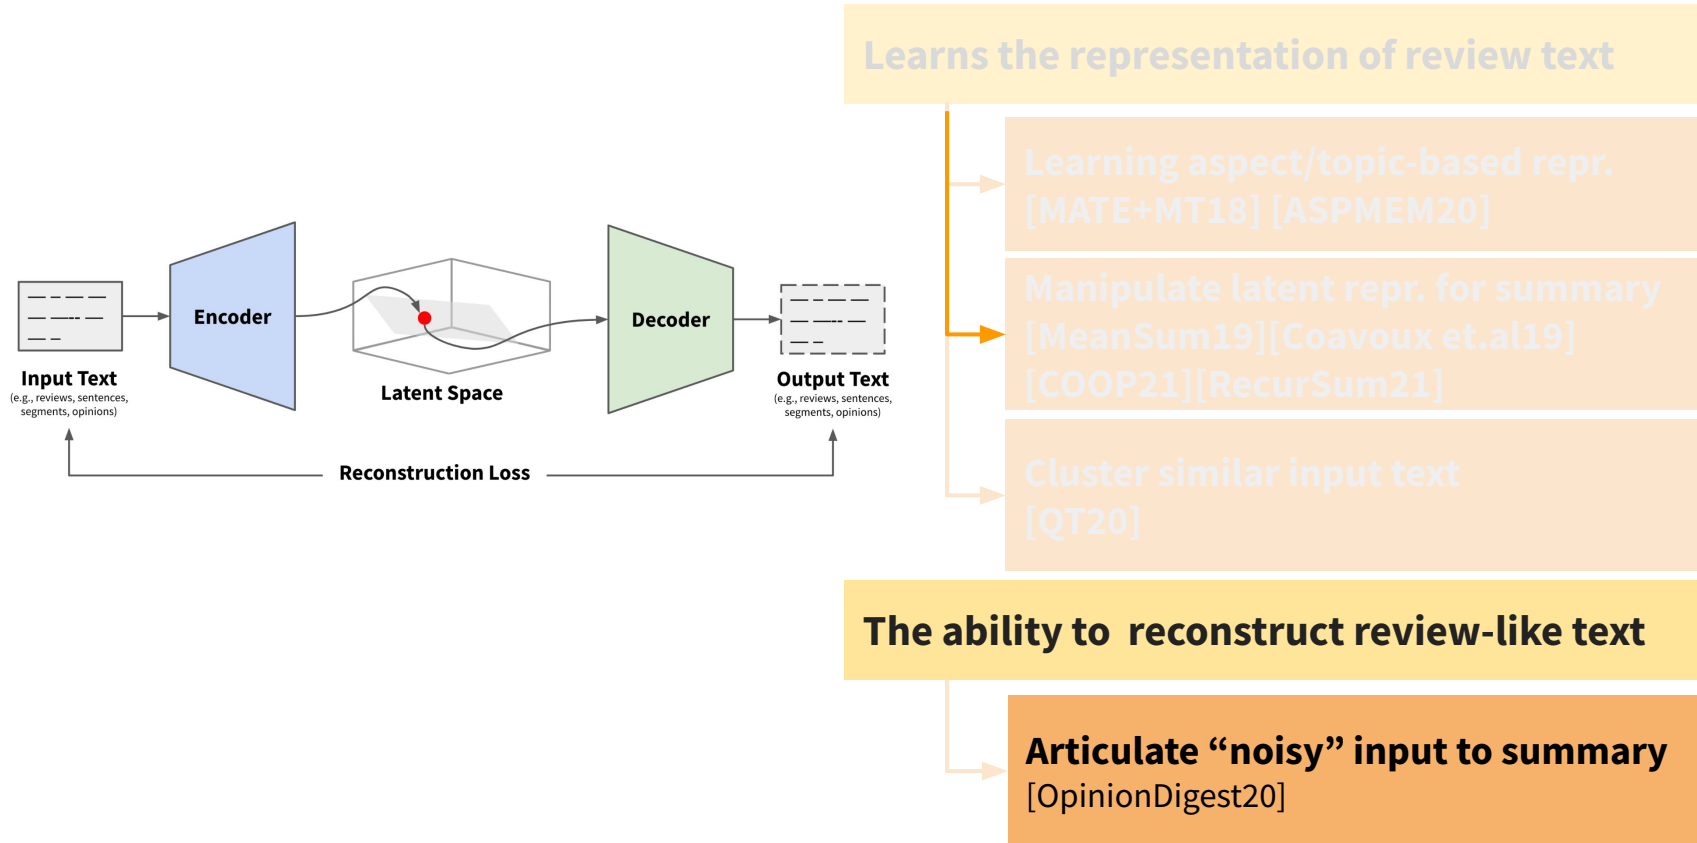
Inference

Sentences of an entity:

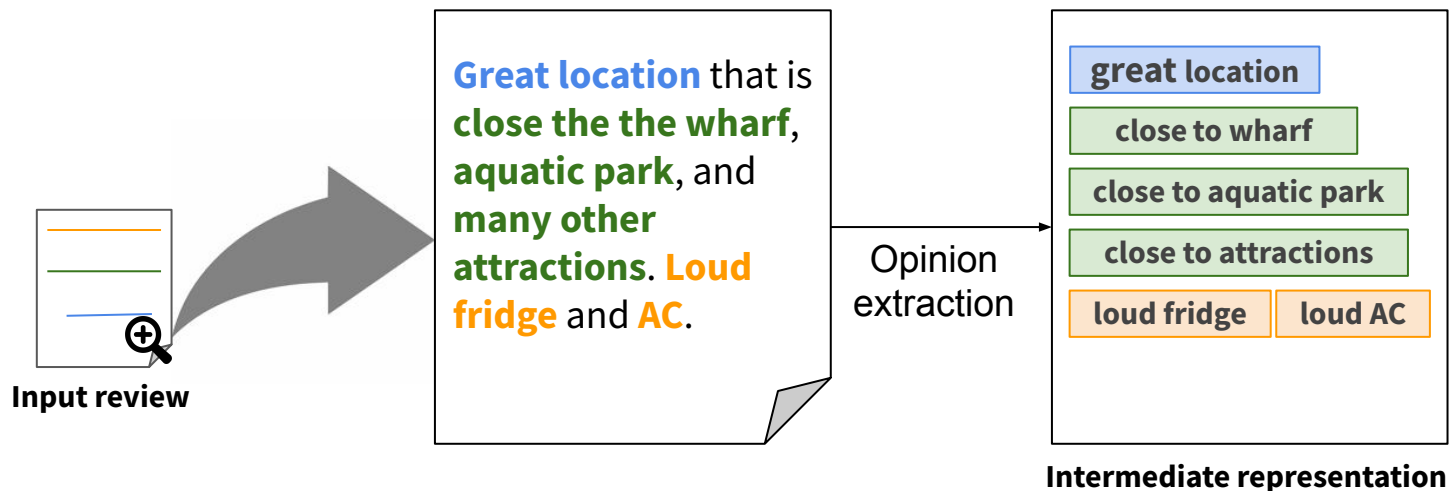


Latent Space
(discrete codebook)

Opinion Summarization via Autoencoder



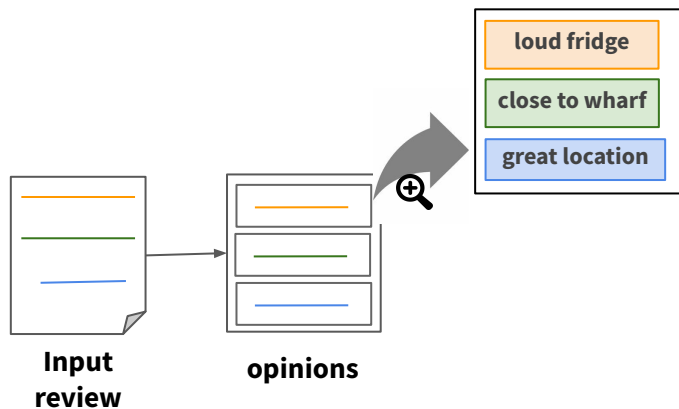
OpinionDigest¹



1. Suhara, Yoshihiko, Xiaolan Wang, Stefanos Angelidis, and Wang-Chiew Tan. "OpinionDigest: A Simple Framework for Opinion Summarization." ACL, pp. 5789-5798. 2020.


OpinionDigest

Training: Learn to reconstruct the original review from **extracted opinion phrases**



OpinionDigest

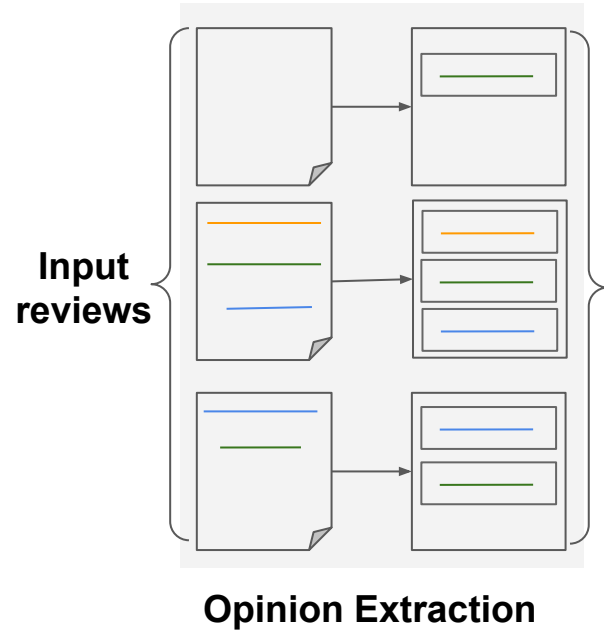
Training: Learn to reconstruct the original review from **extracted opinion phrases**



Variation of denoising autoencoder – learns to generate a review from opinion phrases

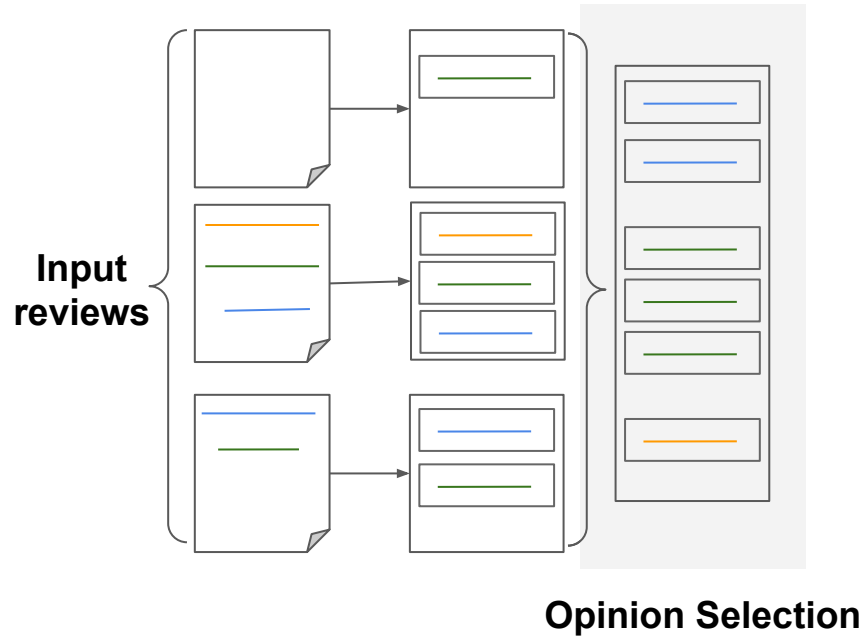
OpinionDigest

Generation (1/3): Extract opinions from input reviews



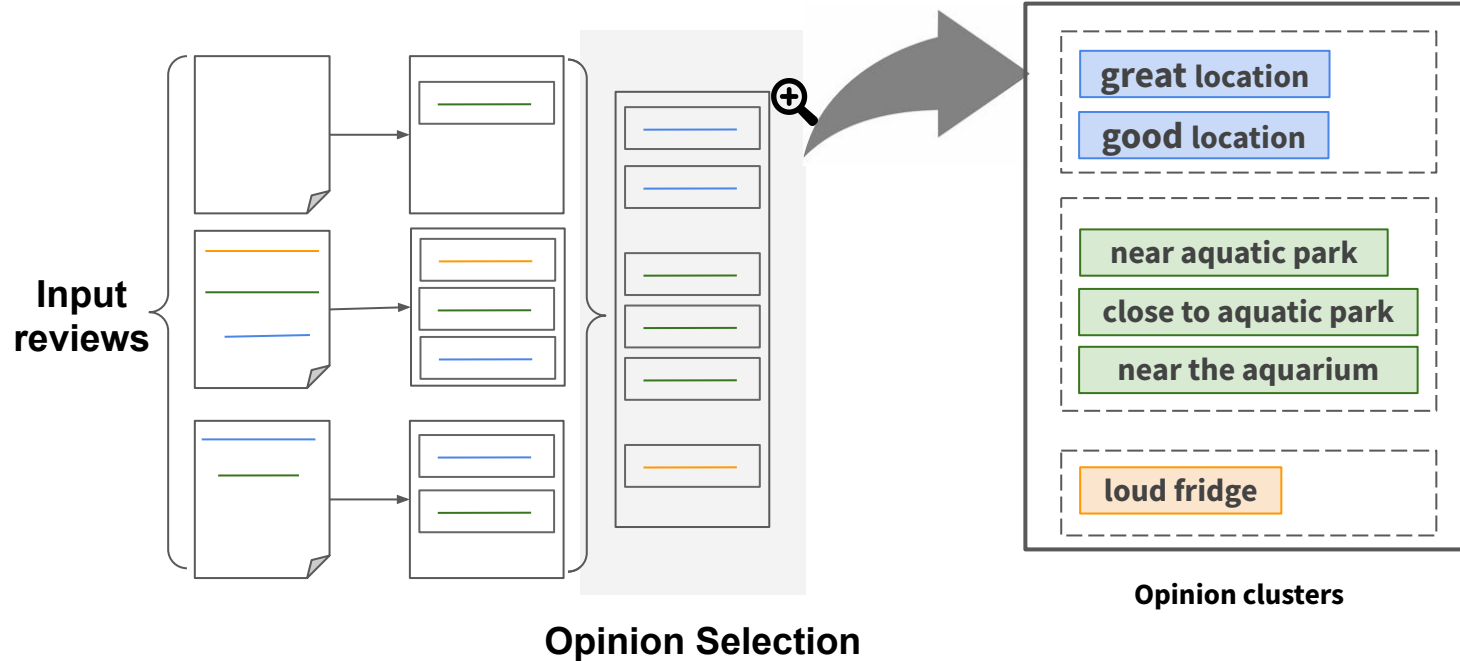
OpinionDigest

Generation (2/3): Select most popular opinions



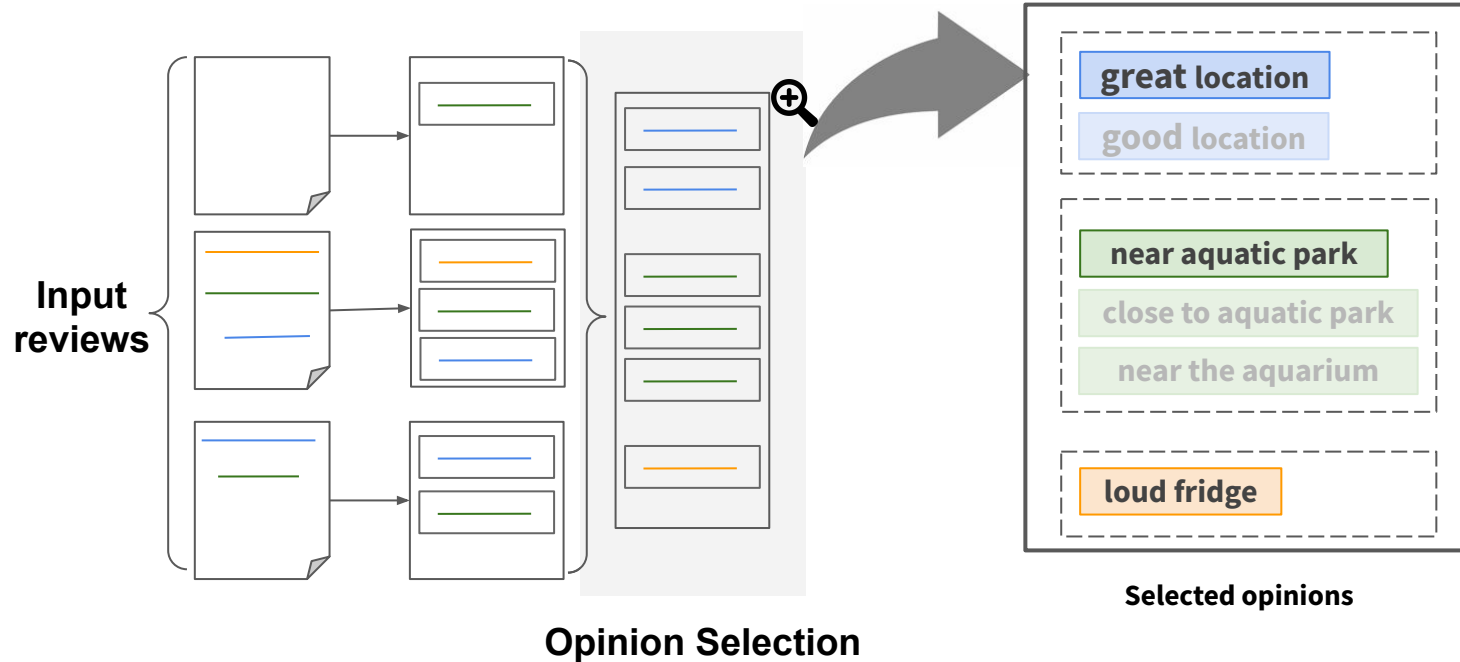
OpinionDigest

Generation (2/3): Select most popular opinions



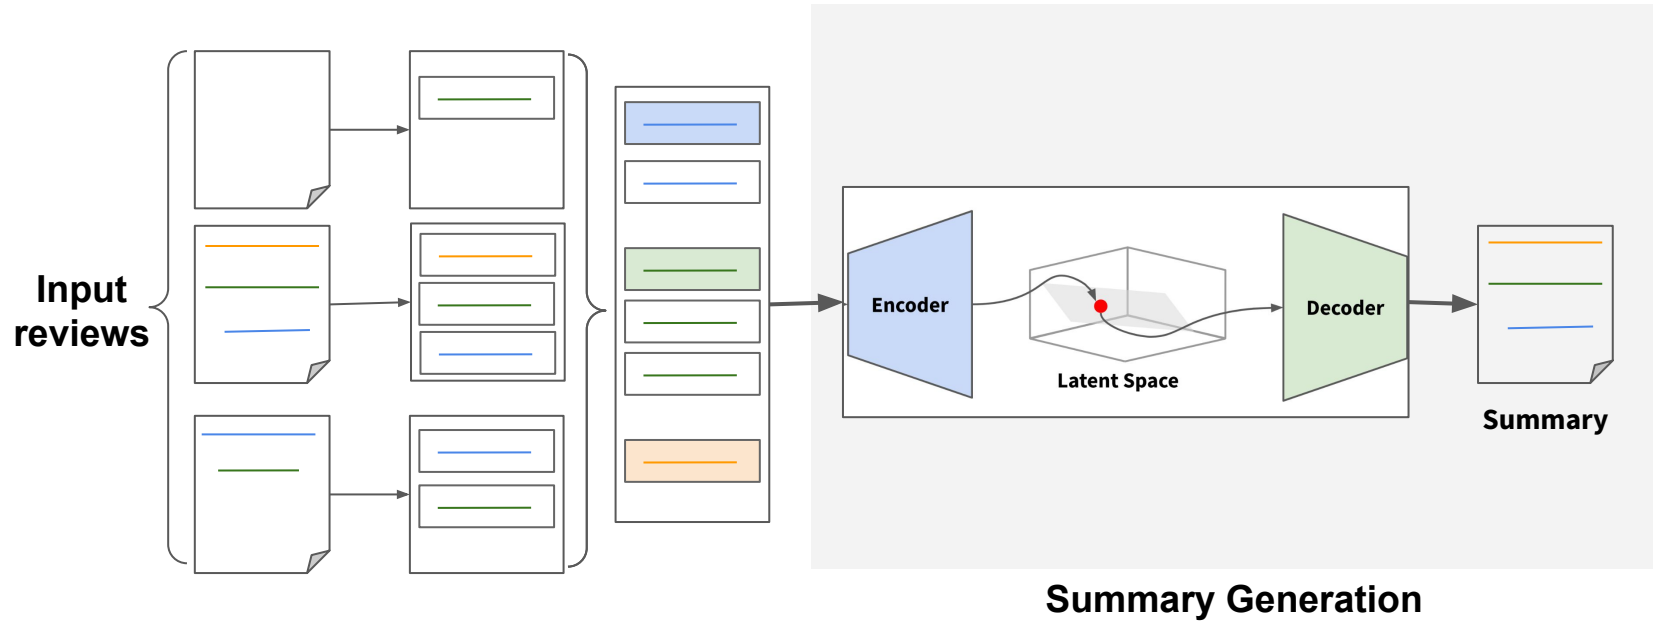
OpinionDigest

Generation (2/3): Select most popular opinions




OpinionDigest

Generation (3/3): Summary generation



Opinion Summarization via Autoencoder



Learns the representation of review text

Manipulate latent repr. for summary

[MeanSum19][Coavoux et.al19] [COOP21]

Learning aspect/topic-based repr.

[MATE+MT18][RecurSum21][ASPMEM20]

Cluster similar input text


[QT20]

The ability to reconstruct review-like text

Articulate “noisy” input to summary

[OpinionDigest20]

Autoencoder: Pros and Cons



(+) Does not require reference summaries for training
(+) Trained on individual reviews, thus can summarize large number of reviews

(-) Always generate review-like sentences as the summary due to self-supervised training
(-) Cannot summarize directly from multiple input reviews