PASTE: A Tagging-Free Decoding Framework Using Pointer Networks for Aspect Sentiment Triplet Extraction

Rajdeep Mukherjee, Tapas Nayak, Yash Butala, Sourangshu Bhattacharya, Pawan Goyal

CNeRG Lab, Department of Computer Science and Engineering, Indian Institute of Technology Kharagpur

EMNLP 2021 Online & in the Dominican Republic
Conference on Empirical Methods in Natural Language Processing

Aspect Sentiment Triplet Extraction

Given a sentence, ASTE deals with extracting ***opinion triplets***, consisting of an opinion target or aspect, its associated sentiment, and the corresponding opinion term/span explaining the rationale behind the sentiment.

The weather was *gloomy* , but the food was tasty .

Triplet 1: (weather, gloomy, NEGATIVE)
Triplet 2: (food, tasty, POSITIVE)
} Triplets with Different Sentiments

The *film* was *good* , but *could have been better* .

Triplet 1: (film, good, POSITIVE)
Triplet 2: (film, could have been better, NEGATIVE)
} Aspect-Overlapped Triplets

Limitations of Prior Works

- Early efforts break the interaction between aspects and opinions by extracting them in isolation using separate *BIEOS*-based sequence taggers.
- Some of the later methods rely on word-pair sentiment dependencies, and hence cannot guarantee span-level sentiment consistency.
- One of the recent methods extends the BIEOS tags to propose a novel position-aware tagging scheme. None of their model variants can, however detect both aspect- as well as opinion-overlapped triplets.

Our Goals

- Offer a new perspective to solve the task - investigate the utility of a tagging-free scheme.
- Jointly extract aspect and opinion spans by modeling their interdependence ***during*** the extraction process.
- Present a truly end-to-end solution while effectively capturing the inter-relatedness between all three elements of an opinion triplet.
- Can we use a generative framework for the task?

PASTE Architecture

PASTE uses an encoder-decoder architecture with a Pointer Network-based decoding framework that generates an entire opinion triplet at each time step. Different from prior works, our solution is ***tagging-free*** and end-to-end.

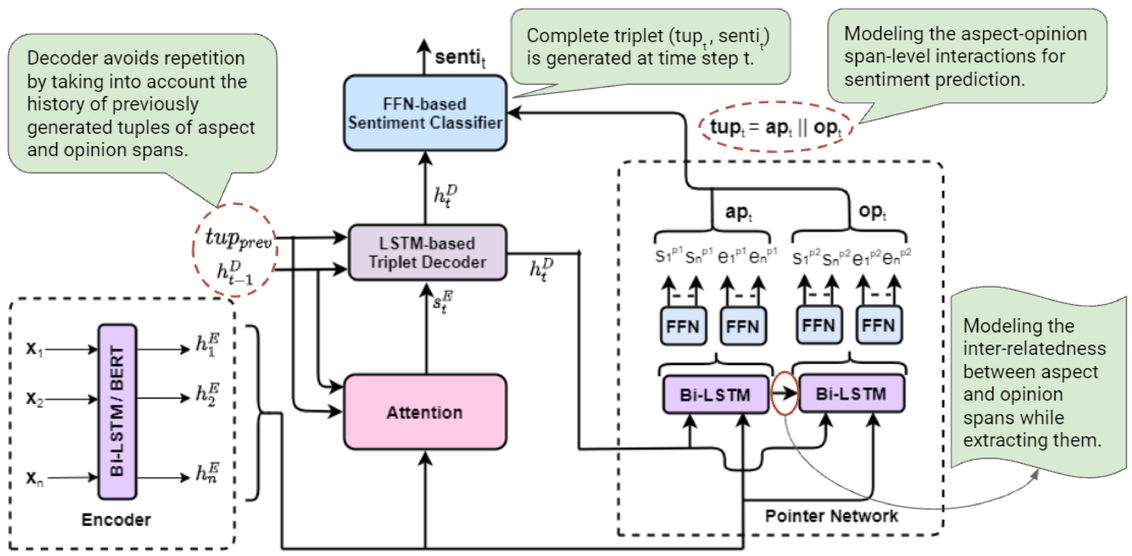

Experiments and Results

Model	Laptop			14Rest			15Rest			16Rest		
	P.	R.	F_1	P.	R.	F_1	P.	R.	F_1	P.	R.	F_1
What-How-Why	0.374	0.504	0.429	0.432	0.637	0.515	0.481	0.575	0.523	0.470	0.642	0.542
OTE-MTL	0.492	0.405	0.451	0.630	0.551	0.587	0.579	0.427	0.489	0.603	0.534	0.565
JET^{o}	0.560	0.354	0.433	0.615	0.551	0.581	0.644	0.443	0.525	0.709	0.570	**0.632**
GTS-BiLSTM	0.597	0.348	0.439	0.686	0.528	0.597	0.654	0.443	0.528	0.686	0.515	0.588
PASTE-AF	0.537	0.486	**0.510**	0.624	0.618	0.621	0.548	0.534	**0.541**	0.622	0.628	0.625
PASTE-OF	0.521	0.481	0.500	0.634	0.619	**0.626**	0.548	0.526	0.537	0.623	0.636	0.629

Figure: Comparative results on the *Laptop* (SemEval 14Lap) and *Restaurant* (14Rest, 15Rest, 16Rest) datasets from ASTE-Data-V2.

Dataset: Experiments are conducted on the **ASTE-Data-V2** (proposed by the authors of *JET*). It contains sentences from two domains: laptop and restaurant, with 27.68% of all sentences containing aspect/opinion overlapped triplets.

Baselines:

- *What-How-Why*: A 2-stage pipeline approach. Aspect-sentiment pairs and opinion spans are extracted separately.
- *OTE-MTL*: A multi-task framework. Two separate sequence taggers detect the aspect and opinion spans in isolation.
- *JET*: A novel position-aware tagging scheme. JET^{o}/JET^{t} cannot handle opinion/aspect-overlapped triplets.
- *GTS*: A novel grid tagging scheme. Models word-level interactions; can't guarantee span-level sentiment consistency.

Performance on Challenging Data Points

Model	Laptop			
	Single	*Multi*	*MultiPol*	*Overlap*
JET^{o}	0.453	0.406	0.219	0.363
OTE-MTL	0.485	0.277	0.172	0.380
GTS-BiLSTM	0.418	0.452	**0.237**	0.403
PASTE-AF	**0.506**	**0.512**	0.216	0.507
PASTE-OF	0.495	0.502	0.205	**0.511**

Figure: Comparison of F1 scores on different splits of *Laptop* dataset.

- *Single*: Sentences with a single opinion triplet.
- *Multi*: Sentences with a multiple triplets.
- *MultiPol*: Sentences containing at least two triplets with different sentiment polarities.
- *Overlap*: Sentences with overlapping triplets.

Additional results are available in the paper.

Key Highlights

- We formulate ASTE as a structured prediction problem.
- We propose a *tagging-free* position-based scheme to uniformly represent an opinion triplet, irrespective of varying lengths of aspect and opinion spans.
- The *Pointer Networks* exploit the aspect-opinion interdependence *during* the span detection process.
- The *Triplet Decoder* models the aspect-opinion span-level interactions for sentiment prediction.
- *PASTE* decodes a complete triplet at each time step, thereby making our solution truly *end-to-end*.
- Model Variants: *PASTE-AF* and *PASTE-OF* based on whether the first pointer network is used to detect the aspect or opinion span respectively.
- We obtain **15.6% recall gains** (averaged across PASTE-AF and PASTE-OF) over the respective strongest baselines on all the datasets.
- *PASTE* performs better while handling sentences with multiple, and aspect/opinion-overlapped triplets.

For Further Information

Preprint: `https://arxiv.org/pdf/2110.04794.pdf`
Github: `https://github.com/rajdeep345/PASTE`