Leveraging Just a Few Keywords for Fine-Grained Aspect Detection Through Weakly Supervised Co-Training

Giannis Karamanolakis, Daniel Hsu, Luis Gravano
Department of Computer Science, Columbia University
{gkaraman, djhsu, gravano}@cs.columbia.edu

COLUMBIA UNIVERSITY

Online User-Generated Reviews

amazon prime | Electronics | Televisions | EN

32-Inch 720p LED TV

Great tv for six months then out goes the sound
March 18, 2016
Size: 32-Inch | Style: TV | Verified Purchase

Great tv for six months then out goes the sound. Called samsung for help and the best they could do was send me on a 240 mile round trip to a samsung approved service center that is only open in the day time so I would have to take half a day off work twice to get it fixed. So unless you live close to a service center think twice about this tv.

Good picture and very lightweight. But it would not stay powered off.
December 31, 2015
Size: 32-Inch | Style: TV | Verified Purchase

I had this installed on the wall bracket in the exercise room and was watching a show while finishing up the installation. It is very light and seems well built cosmetically. Problem was when I used the remote to shut it off, it did go off, no audio or picture. Then 3 seconds later it came right back on again... spooky I know. So I shut it down again, it went off and right back on again. looking around to see if my wife had the smart remote from upstairs from the Samsung big screen playing games with my psyche, not to be found. So I boxed it up, sent Amazon a return request, they sent UPS the next day and I had the replacement in 2 days no charge for anything. This TV shuts off and is a welcome addition to the man cave and the exercise room. Note, this TV has a 19 volt DC power converter similar to your laptop charger. It is not a 120VAC direct power cord to the TV from the wall. The power supply sits on the floor and a cord to the TV. 5 stars for the replacement TV

A great TV for an amazing price
January 16, 2016
Size: 32-Inch | Style: TV | Verified Purchase

I purchased this to replace an older Philips LCD TV of the same size; after going on 8 years it finally died. The first time I ordered this it arrived with a cracked screen, which I'm suspecting was due to the packaging on Amazon's end (or lack there of). The replacement arrived in perfect condition though and it's extremely user-friendly and easy to set up. It only has 2 HDMI ports, but since this a second TV that I use in my bedroom, it works fine for me. I don't have cable in my bedroom and the two HDMI ports work perfectly for my Roku and DVD player.

The picture quality of this is very good and the sound is exceptionally good for a TV that is so thin. The actual TV itself only weighs maybe five or six pounds. My only complaint is that the legs of the TV can only be mounted to the very ends, making it difficult to fit on a smaller surface; it barely fit on the top of my dresser. The legs are also only about an inch high, making it impractical to put anything underneath of it, as it will block the picture. (To visualize, my tiny Roku 2 barely fits underneath it without blocking the screen). Other than that, it's great TV and exceptionally well priced.

yelp | Find Restaurants | Near Civic Center, Manhattan, NY | Log In | Sign Up

Tavern on the Green

Claimed
1377 reviews
$$$ · American (New), Cocktail Bars Edit

Write a Review | Add Photo | Share | Save

Britni R.
Houston, TX
325 friends
3 reviews
2 photos
Share review
Embed review

10/8/2019

I went to Tavern on the green before they closed during the recession and remember it very fondly. On my recent visit to New York, I went for brunch and was incredibly disappointed. It has become very touristy, the booths and chairs are rundown, glasses were not clean, the food is mediocre food and the staff are so rude.

The location is ideal, but other than that, there is no reason for going.

Useful | Funny | Cool

Rich H.
New York, NY
145 friends
236 reviews
126 photos
Elite '2019
Share review
Embed review

10/11/2019
1 check-in

This is a great venue for food, drinks, Sunday Brunch and impressing your out of town guests. It's right in Central Park so gives you a good opportunity to walk around the park either before or after. Outdoor seating on nice days is your best option. Great service, great food and great ambiance. If you go during brunch, give their Bloody Mary a try. It's one of my favorites!

Useful | Funny | Cool

Stephanie G.
New York, NY
80 friends
58 reviews
60 photos
Elite '2019
Share review
Embed review

7/30/2019
3 photos

I would actually rate this place a 3.5 stars if I could. I came here for lunch with my husband during restaurant week. The restaurant and the patio are absolutely beautiful, and it is definitely a perfect date spot. The service was friendly and efficient. The food was good, but not amazing. I ordered a watermelon appetizer and got a salmon dish as my entree. The salmon was great, and it was on top of a corn chowder/bacon sauce. Very tasty. The watermelon feta salad however, was pretty bad. It was maybe 20 pieces of watermelon cubed, with 6 pieces of feta, drizzled with balsamic vinegar. I am obsessed with watermelon feta salads, and I make them frequently at home. I'm a pretty poor cook, so the fact that my salad is 10x better than theirs says a lot.

My husband also ordered the salmon, and he got the squash flatbread for his appetizer. He was very happy with both dishes. He also ordered the wine pairing for $17, and wow. The wine was spectacular. Very delicious.

All in all, we left satisfied. The watermelon salad really left a bad taste in my mouth for the restaurant, because it's pretty shocking that they would serve something that was so bad, when other dishes were rather tasty. I definitely think it's worth a visit because it's a classic NY restaurant, but if you're just in it for the food, I think there are many other restaurants with better tasting, cheaper dishes.

Online User-Generated Reviews

amazon prime | Electronics | Televisions | EN

32-Inch 720p LED TV

TV Aspects

1. Price
2. Image Quality
3. Ease of Use
4. …

remote to shut it off, it did go off, no audio or picture. Then 3 seconds later it came right back on again... spooky I know. So I shut it down again, it went off and right back on again. looking around to see if my wife had the smart remote from upstairs from the Samsung big screen playing games with my psyche, not to be found. So I boxed it up, sent Amazon a return request, they sent UPS the next day and I had the replacement in 2 days no charge for anything. This TV shuts off and is a welcome addition to the man cave and the exercise room. Note, this TV has a 19 volt DC power converter similar to your laptop charger. It is not a 120VAC direct power cord to the TV from the wall. The power supply sits on the floor and a cord to the TV. 5 stars for the replacement TV

A great TV for an amazing price
January 16, 2016
Size: 32-Inch | Style: TV | Verified Purchase

I purchased this to replace an older Philips LCD TV of the same size; after going on 8 years it finally died. The first time I ordered this it arrived with a cracked screen, which I'm suspecting was due to the packaging on Amazon's end (or lack there of). The replacement arrived in perfect condition though and it's extremely user-friendly and easy to set up. It only has 2 HDMI ports, but since this a second TV that I use in my bedroom, it works fine for me. I don't have cable in my bedroom and the two HDMI ports work perfectly for my Roku and DVD player.

The picture quality of this is very good and the sound is exceptionally good for a TV that is so thin. The actual TV itself only weighs maybe five or six pounds. My only complaint is that the legs of the TV can only be mounted to the very ends, making it difficult to fit on a smaller surface; it barely fit on the top of my dresser. The legs are also only about an inch high, making it impractical to put anything underneath of it, as it will block the picture. (To visualize, my tiny Roku 2 barely fits underneath it without blocking the screen). Other than that, it's great TV and exceptionally well priced.

yelp | Find Restaurants | Near Civic Center, Manhattan, NY | Log In | Sign Up

Tavern on the Green Claimed
1377 reviews
$$$ · American (New), Cocktail Bars Edit

Write a Review | Add Photo | Share | Save

Restaurant Aspects

1. Food Quality
2. Ambience
3. Service
4. …

Elite '2019
Share review
Embed review

Park so gives you a good opportunity to walk around the park either before or after. Outdoor seating on nice days is your best option. Great service, great food and great ambiance. If you go during brunch, give their Bloody Mary a try. It's one of my favorites!

Useful | Funny | Cool

Stephanie G.
New York, NY
80 friends
58 reviews
60 photos
Elite '2019
Share review
Embed review

7/30/2019
3 photos

I would actually rate this place a 3.5 stars if I could. I came here for lunch with my husband during restaurant week. The restaurant and the patio are absolutely beautiful, and it is definitely a perfect date spot. The service was friendly and efficient. The food was good, but not amazing. I ordered a watermelon appetizer and got a salmon dish as my entree. The salmon was great, and it was on top of a corn chowder/bacon sauce. Very tasty. The watermelon feta salad however, was pretty bad. It was maybe 20 pieces of watermelon cubed, with 6 pieces of feta, drizzled with balsamic vinegar. I am obsessed with watermelon feta salads, and I make them frequently at home. I'm a pretty poor cook, so the fact that my salad is 10x better than theirs says a lot.

My husband also ordered the salmon, and he got the squash flatbread for his appetizer. He was very happy with both dishes. He also ordered the wine pairing for $17, and wow. The wine was spectacular. Very delicious.

All in all, we left satisfied. The watermelon salad really left a bad taste in my mouth for the restaurant, because it's pretty shocking that they would serve something that was so bad, when other dishes were rather tasty. I definitely think it's worth a visit because it's a classic NY restaurant, but if you're just in it for the food, I think there are many other restaurants with better tasting, cheaper dishes.

Fine-Grained Aspect Detection

- Segments (e.g., sentences, clauses) of a review may discuss different aspects
- **Goal:** Classify individual segments to *K* pre-defined aspect classes

Great price for an excellent LED TV

Verified Purchase

Sentence	**Aspect**
Great Tv for the price.	**Price**
Easy to setup.	**Ease of Use**
The audio is ok for the tiny speakers.	**Sound Quality**
Much better than the 20" tube tv.	**General**

Fine-Grained Aspect Detection

- Segments (e.g., sentences, clauses) of a review may discuss different aspects
- **Goal:** Classify individual segments to *K* pre-defined aspect classes

Great price for an excellent LED TV

Verified Purchase

<u>**Sentence**</u>	<u>**Aspect**</u>
Great Tv for the price.	**Price**
Easy to setup.	**Ease of Use**
The audio is ok for the tiny speakers.	**Sound Quality**
Much better than the 20" tube tv.	**General**

- **This work:** Train neural networks **without** ground truth segment labels

Fine-Grained Aspect Detection

- Segments (e.g., sentences, clauses) of a review may discuss different aspects
- **Goal:** Classify individual segments to *K* pre-defined aspect classes

Great price for an excellent LED TV

Verified Purchase

Sentence	Aspect
Great Tv for the price.	**Price**
Easy to setup.	**Ease of Use**
The audio is ok for the tiny speakers.	**Sound Quality**
Much better than the 20" tube tv.	**General**

- **This work:** Train neural networks **without** ground truth segment labels
  - **Supervised** learning: require (many) segment labels

amazon

- **>1B products**
- **>10K categories**
- **New products added every day**

Fine-Grained Aspect Detection

- Segments (e.g., sentences, clauses) of a review may discuss different aspects
- **Goal:** Classify individual segments to *K* pre-defined aspect classes

Great price for an excellent LED TV

Verified Purchase

Sentence	Aspect
Great Tv for the price.	**Price**
Easy to setup.	**Ease of Use**
The audio is ok for the tiny speakers.	**Sound Quality**
Much better than the 20" tube tv.	**General**

- **This work:** Train neural networks **without** ground truth segment labels
  - ~~**Supervised** learning: require (many) segment labels~~
  - **Unsupervised** learning (topic models): may not capture the *K* aspects of interest

Fine-Grained Aspect Detection

- Segments (e.g., sentences, clauses) of a review may discuss different aspects
- **Goal:** Classify individual segments to *K* pre-defined aspect classes

Great price for an excellent LED TV

Verified Purchase

Sentence	**Aspect**
Great Tv for the price.	**Price**
Easy to setup.	**Ease of Use**
The audio is ok for the tiny speakers.	**Sound Quality**
Much better than the 20" tube tv.	**General**

- **This work:** Train neural networks **without** ground truth segment labels

  - ~~**Supervised** learning: require (many) segment labels~~
  - **Unsupervised** learning (topic models): may not capture the *K* aspects of interest

???

Topic 1	*Price*
Topic 2	*Ease of Use*
…	…
Topic 100	*Sound Quality*

Fine-Grained Aspect Detection

- Segments (e.g., sentences, clauses) of a review may discuss different aspects
- **Goal:** Classify individual segments to *K* pre-defined aspect classes

Great price for an excellent LED TV

Verified Purchase

Sentence	Aspect
Great Tv for the price.	**Price**
Easy to setup.	**Ease of Use**
The audio is ok for the tiny speakers.	**Sound Quality**
Much better than the 20" tube tv.	**General**

- **This work:** Train neural networks **without** ground truth segment labels

  - ~~**Supervised** learning: require (many) segment labels~~
  - ~~**Unsupervised** learning (topic models): may not capture the *K* aspects of interest~~

  \+ **Weakly supervised** learning: leverage descriptive seed words as supervision

Seed Words → **Neural Network** → *Price*, *Ease of Use*, …, *Sound Quality*

~~Ground Truth Labels~~

Learning with Seed Words

- We assume a small set of indicative **seed words** per aspect:

Aspect	**Seed Words**
Price	price, value, money, worth, paid
Image	picture, color, quality, black, bright
Sound	sound, speaker, noise, loud, volume

Learning with Seed Words

- We assume a small set of indicative **seed words** per aspect:

Aspect	Seed Words
Price	price, value, money, worth, paid
Image	picture, color, quality, black, bright
Sound	sound, speaker, noise, loud, volume

- **Seed words** are easier to collect than **ground truth labels**

Humans

Manually:

Aspect	Seed Words
Price	price, value, money, worth, paid
Image	picture, color, quality, black, bright
Sound	sound, speaker, noise, loud, volume

Small Number of Labeled Segments

Automatically:

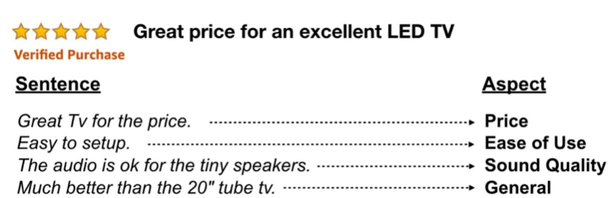

Great price for an excellent LED TV
Verified Purchase

Sentence	Aspect
Great Tv for the price.	Price
Easy to setup.	Ease of Use
The audio is ok for the tiny speakers.	Sound Quality
Much better than the 20" tube tv.	General

Aspect	Seed Words
Price	price, value, money, worth, paid
Image	picture, color, quality, black, bright
Sound	sound, speaker, noise, loud, volume

  - Spoiler: we cope with **noisy** seed words!

How to Leverage Seed Words in Neural Networks?

How to Leverage Seed Words in Neural Networks?

- **Previous approaches:** use seed words for initialization

[Lund et al., 2017]
[Angelidis and Lapata, 2018]

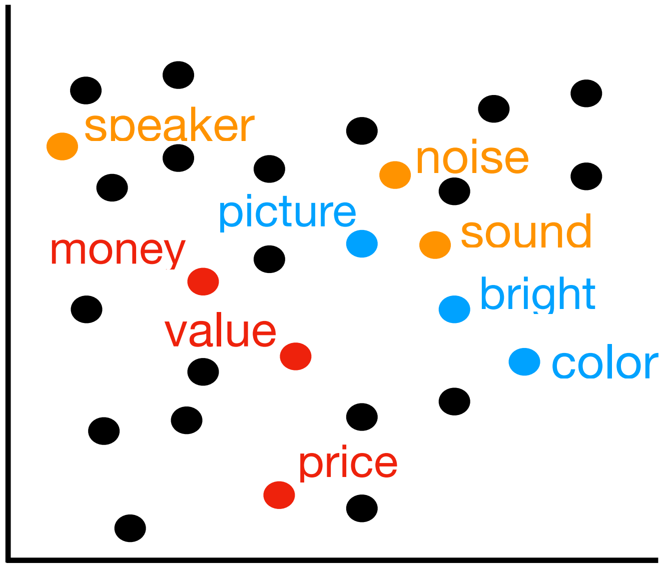

Pre-Trained Word Embeddings

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

unweighted / weighted average

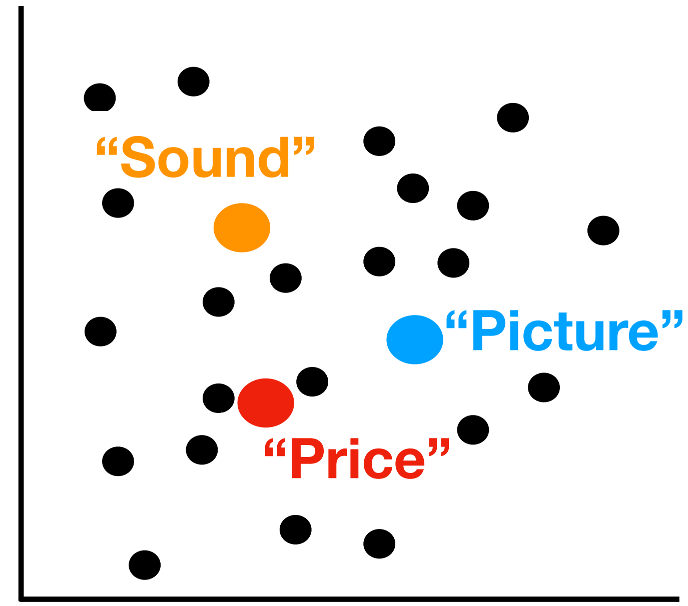

Initial Aspect Embeddings

Train Unsupervised Model

How to Leverage Seed Words in Neural Networks?

- **Previous approaches:** use seed words for initialization

[Lund et al., 2017]
[Angelidis and Lapata, 2018]

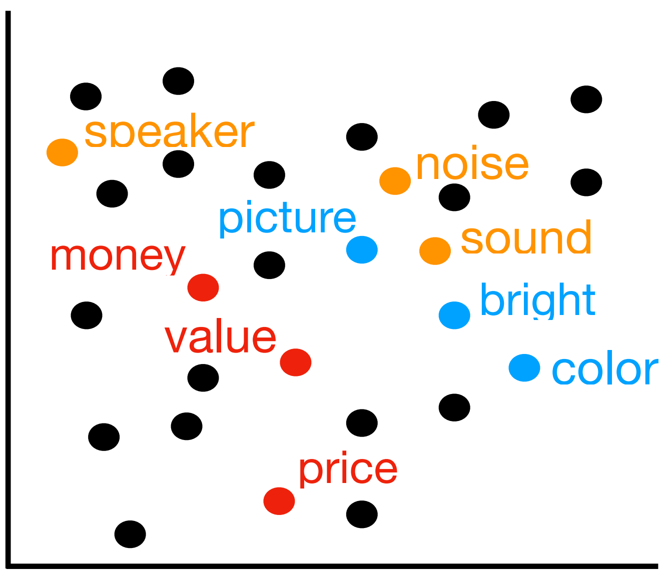

Pre-Trained Word Embeddings

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

unweighted / weighted average

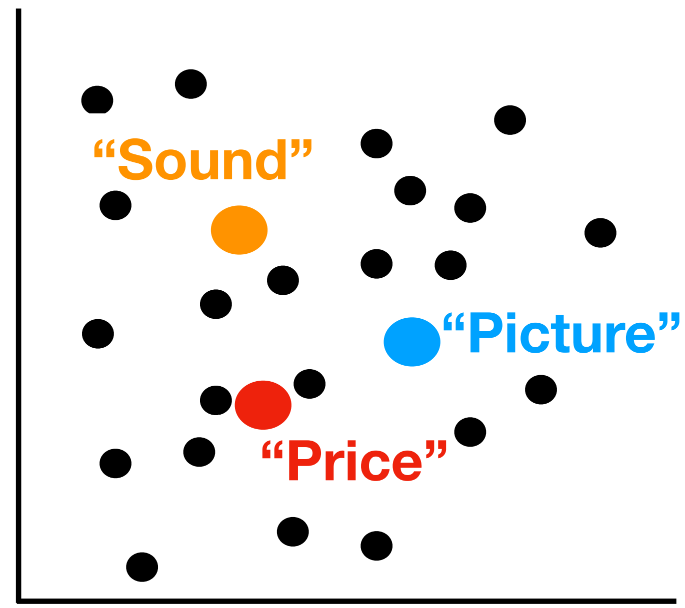

Initial Aspect Embeddings

Train Unsupervised Model

- **Limitations of previous approaches:**
    1. Individual seed words are **not** used during training
    2. Aspect embedding quality is **sensitive** to word embedding quality
    3. Aspect embedding fine-tuning risks from **diverging** from pre-defined aspects

How to Leverage Seed Words in Neural Networks?

- **Previous approaches:** use seed words for initialization

[Lund et al., 2017]
[Angelidis and Lapata, 2018]

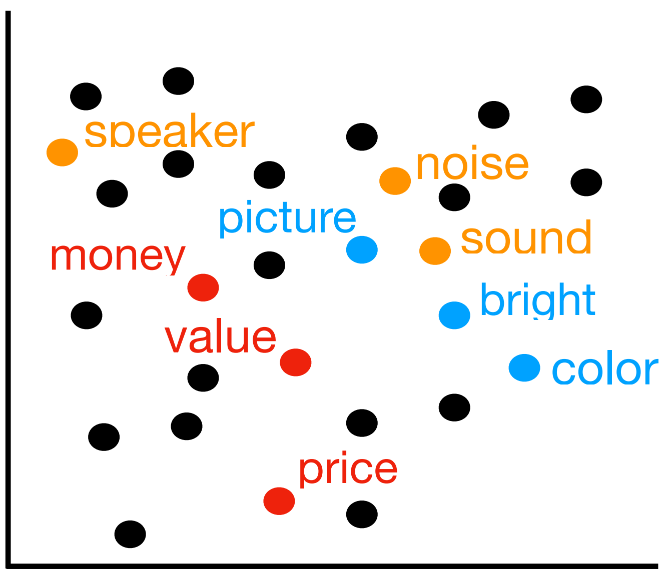

Pre-Trained Word Embeddings

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

unweighted / weighted average

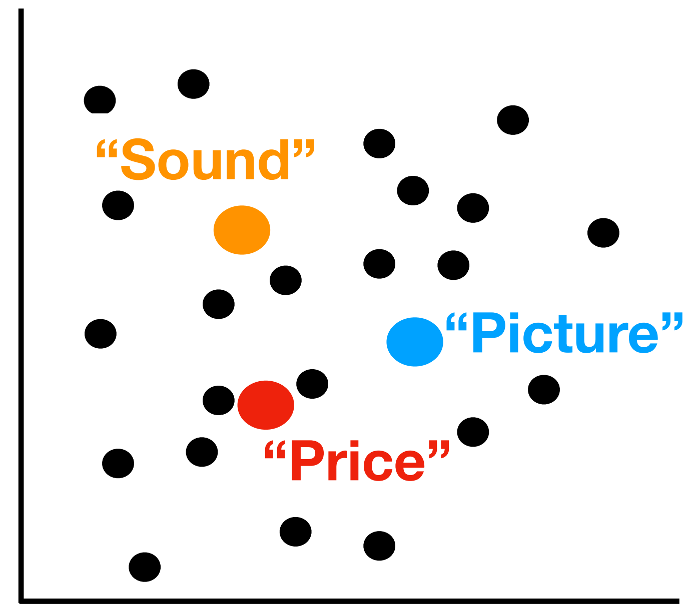

Initial Aspect Embeddings

Train Unsupervised Model

- **Limitations of previous approaches:**
  1. Individual seed words are **not** used during training
  2. Aspect embedding quality is **sensitive** to word embedding quality
  3. Aspect embedding fine-tuning risks from **diverging** from pre-defined aspects

- **Our approach:** Weakly Supervised Co-Training
  + We exploit the **predictive** power …
  + of each individual seed word …
  + during training

Outline

1. Fine-Grained Aspect Detection with Seed Words

2. Weakly Supervised Co-Training

3. Experiments

4. Conclusions and Future Work

Weakly Supervised Co-Training

- How to train neural networks using seed words?
- Our **Teacher-Student** approach:
    - + Leverage **seed words** in a bag-of-seed-words classifier (**Teacher**)
    - + Use **Teacher**'s predictions to train a neural network (**Student**)
    - + Adaptively cope with **noisy** seed words through **iterative** co-training

Aspect	**Seed Words**
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

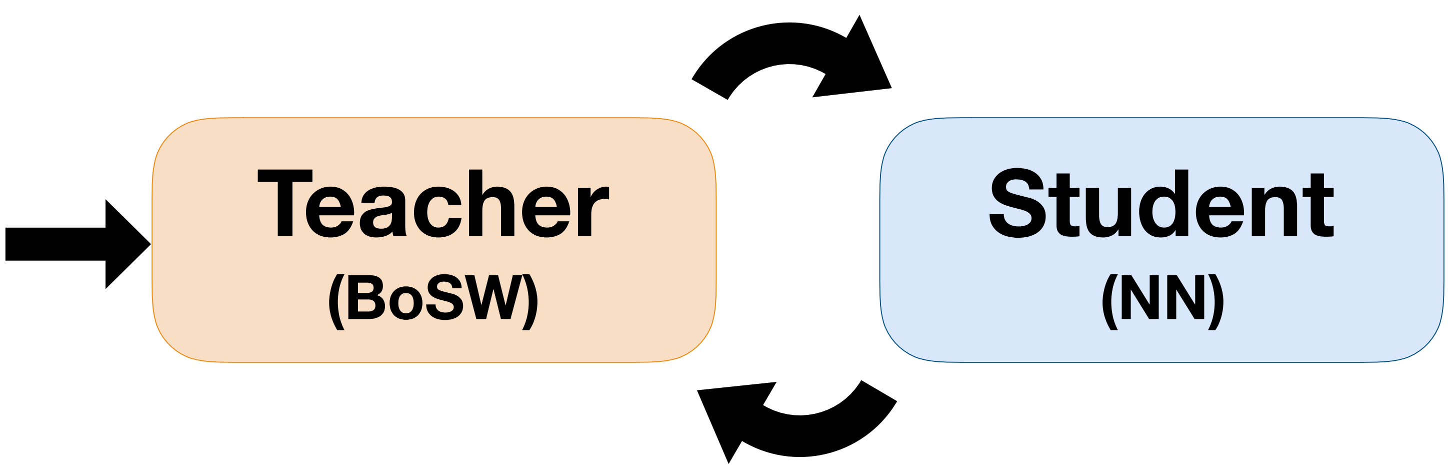

Weakly Supervised Co-Training

1. Teacher
2. Student
3. Teacher -> Student
4. Student -> Teacher

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

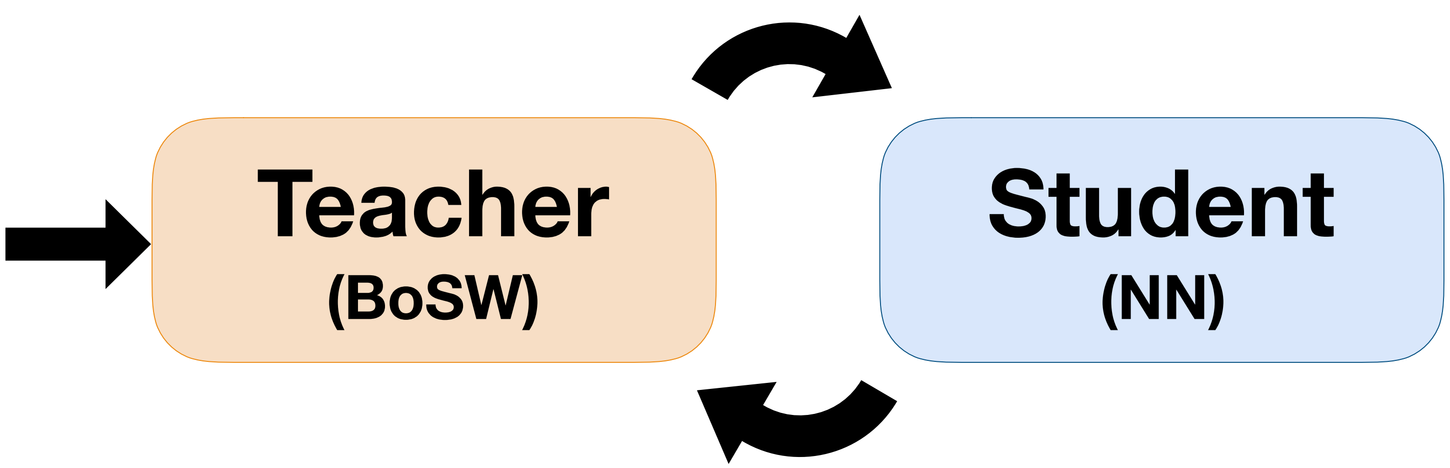

Weakly Supervised Co-Training

1. **Teacher**
2. Student
3. Teacher -> Student
4. Student -> Teacher

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

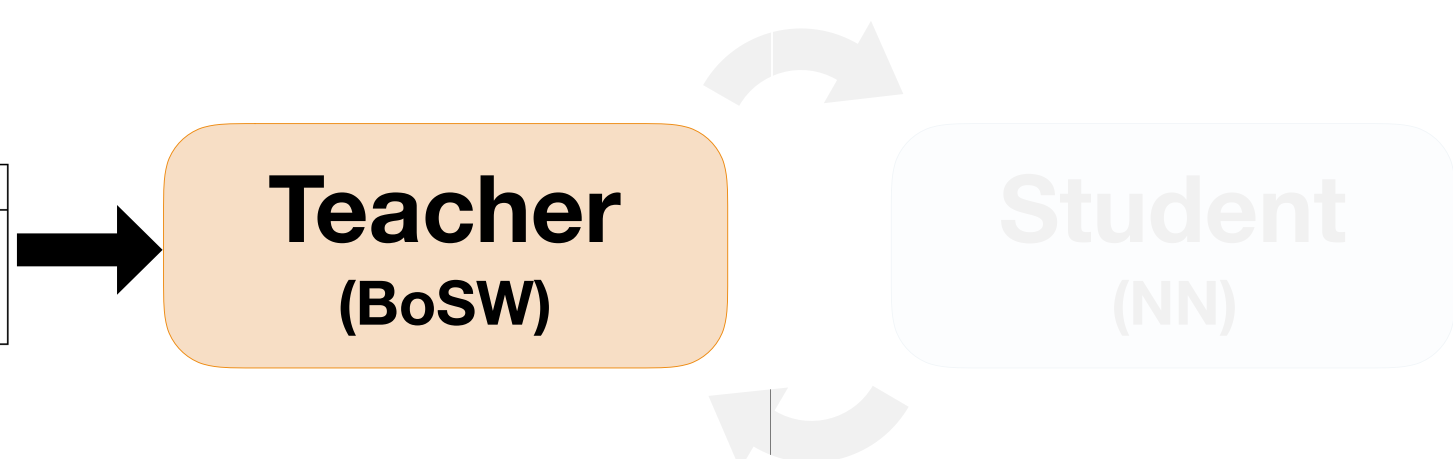

Teacher: Leverages Seed Words to Predict Aspects

- **Teacher:** a **Bag-of-Seed-Words** classifier
  - Uses **seed words** to predict aspect probabilities $\langle q^1, \ldots, q^K \rangle$

softmax
aspect probabilities

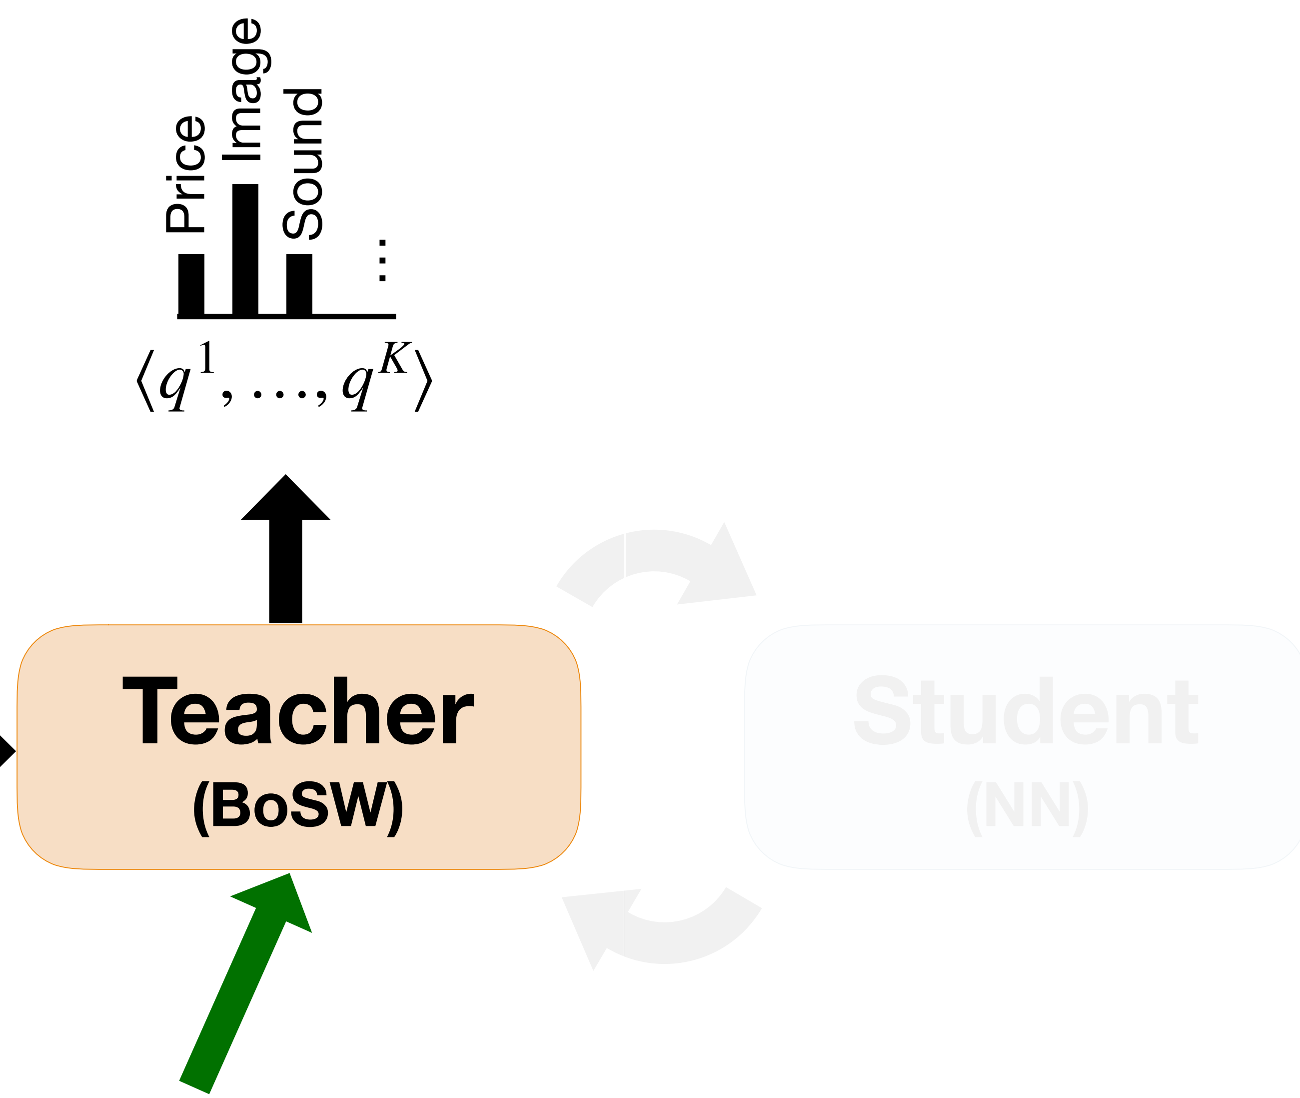

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

segment: *“The **picture** looks very pixelated when playing blu-ray movies”*

Weakly Supervised Co-Training

1. Teacher
2. **Student**
3. Teacher -> Student
4. Student -> Teacher

Student: Leverages Seed Words and Non-Seed Words

- **Student:** an embedding-based neural network
  - Uses **seed words** and **non-seed words** (context) to predict aspects

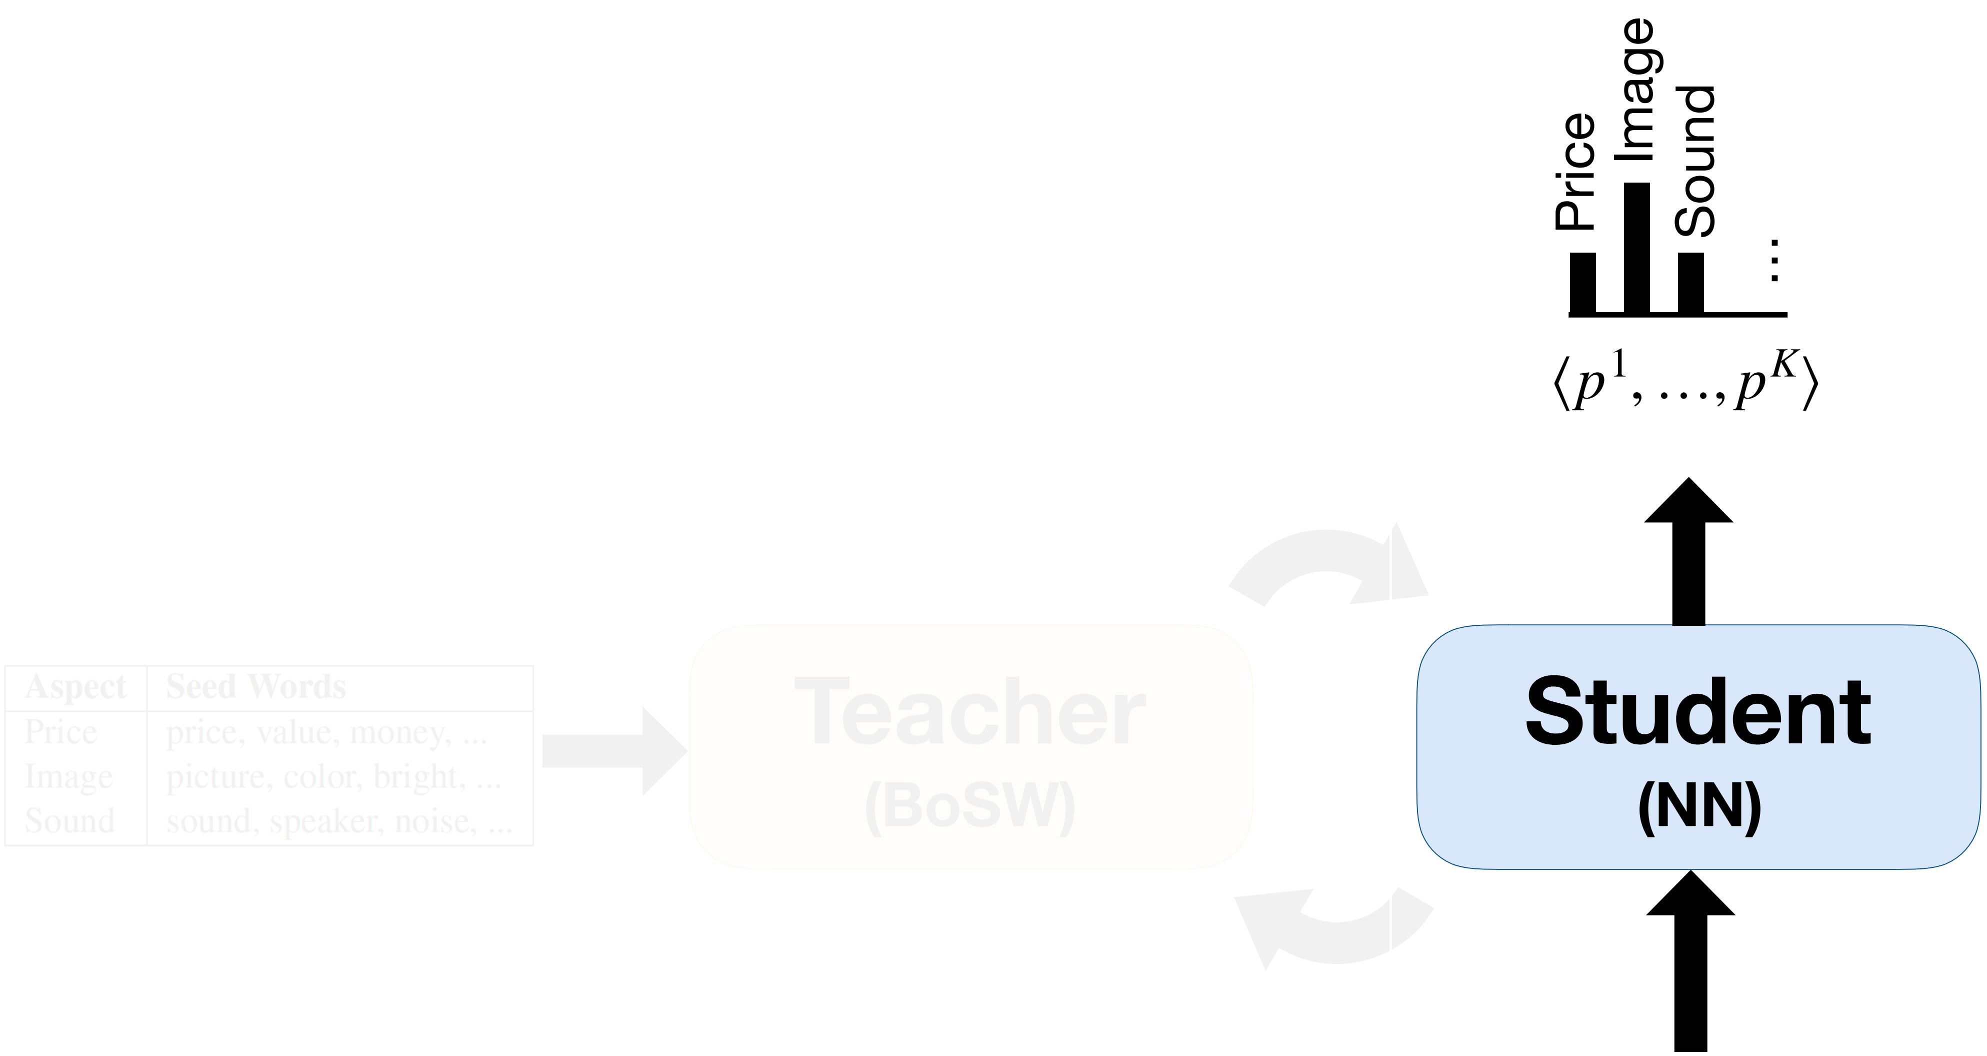

aspect probabilities

segment: *"The **picture** looks very pixelated when playing blu-ray movies"*

Student: Leverages Seed Words and Non-Seed Words

- **Student:** an embedding-based neural network
  - Uses **seed words** and **non-seed words** (context) to predict aspects

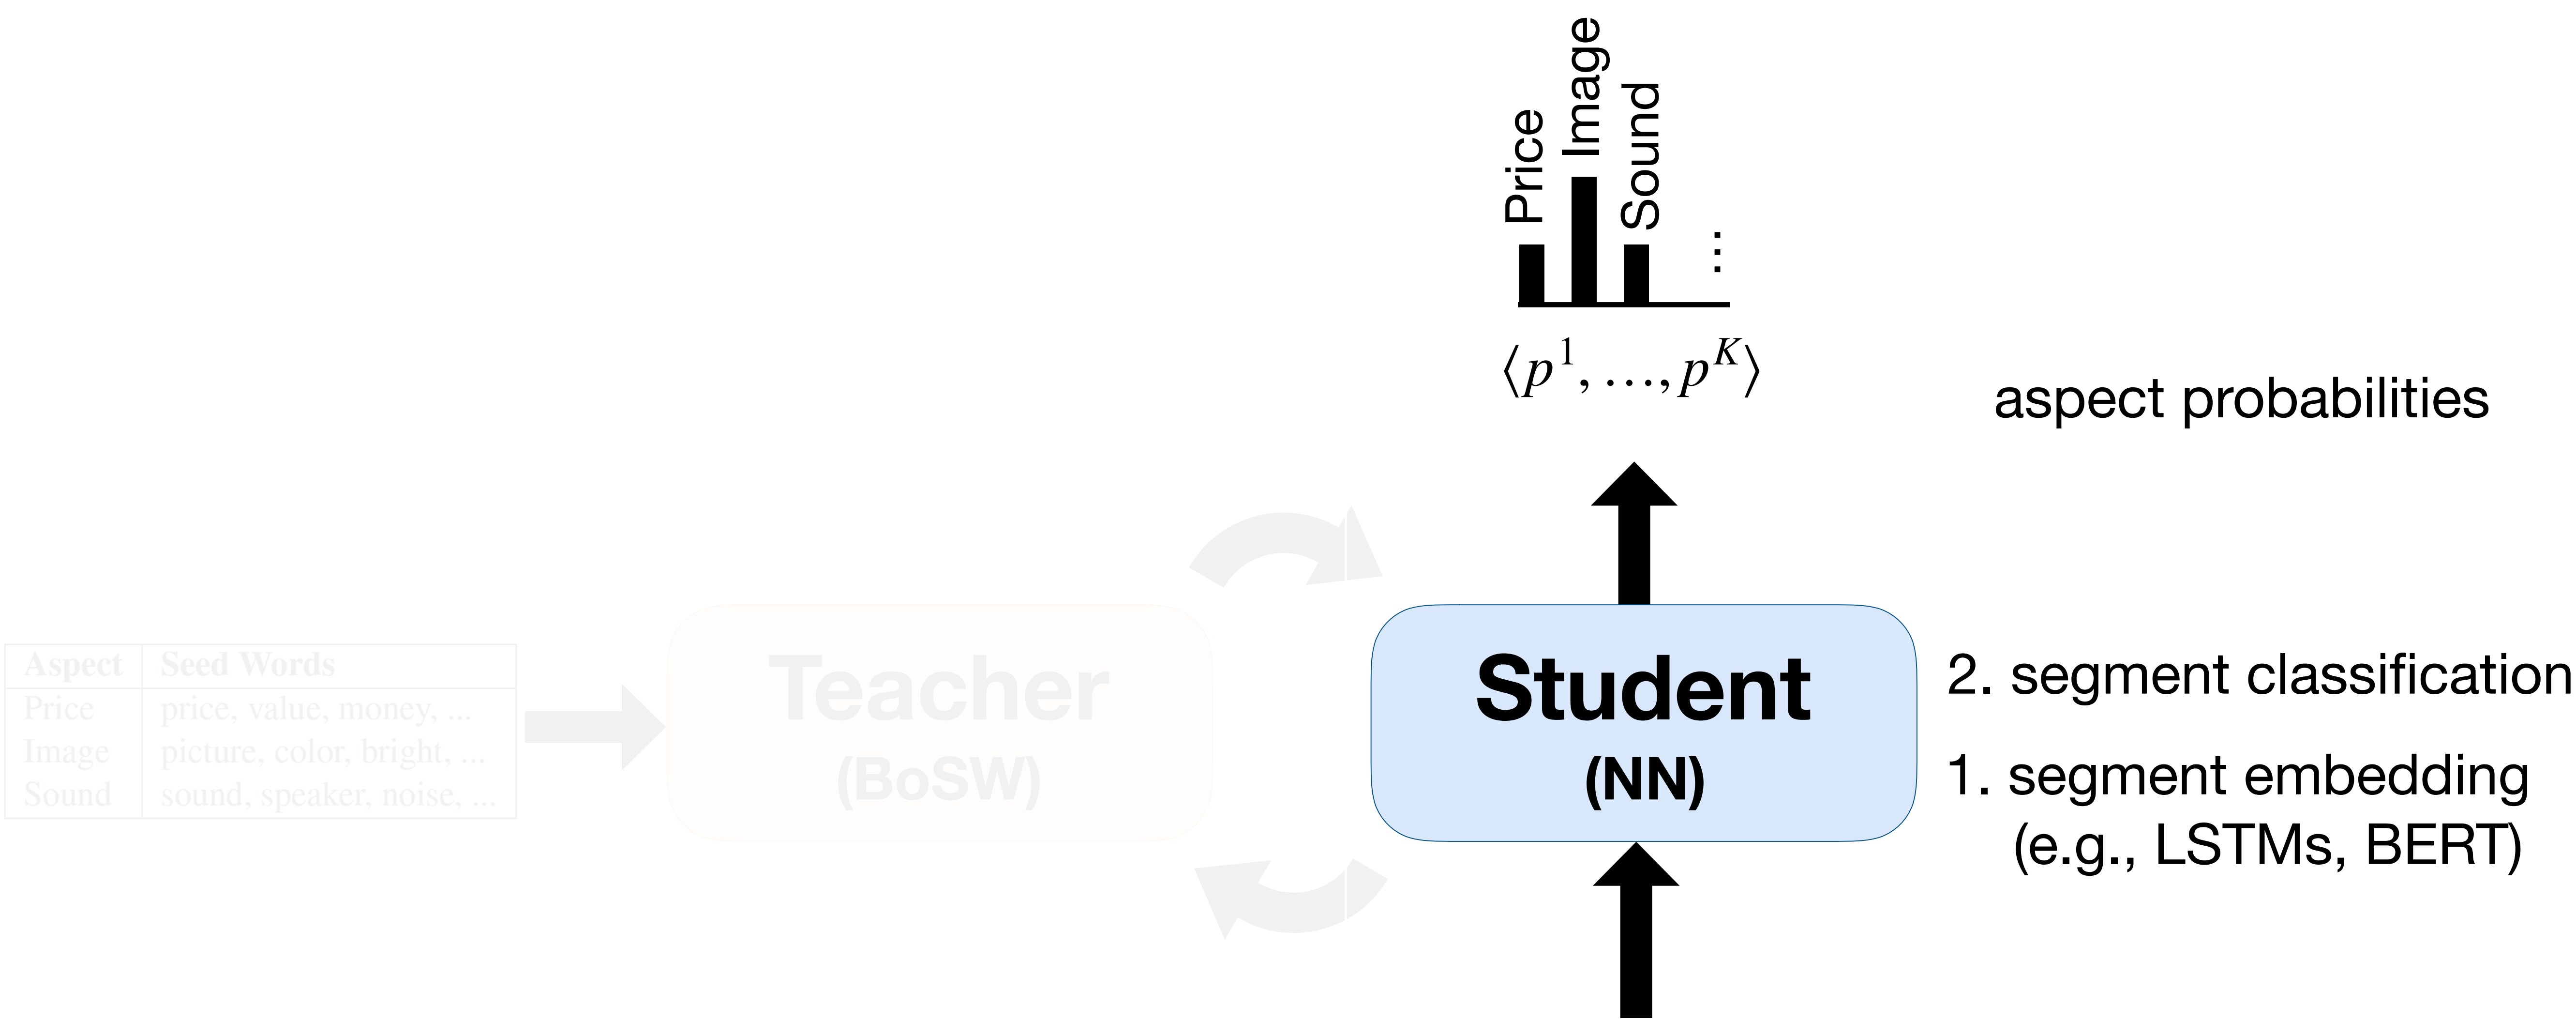

segment: *"The **picture** looks very pixelated when playing blu-ray movies"*

Weakly Supervised Co-Training

1. Teacher
2. Student
3. **Teacher -> Student**
4. Student -> Teacher

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

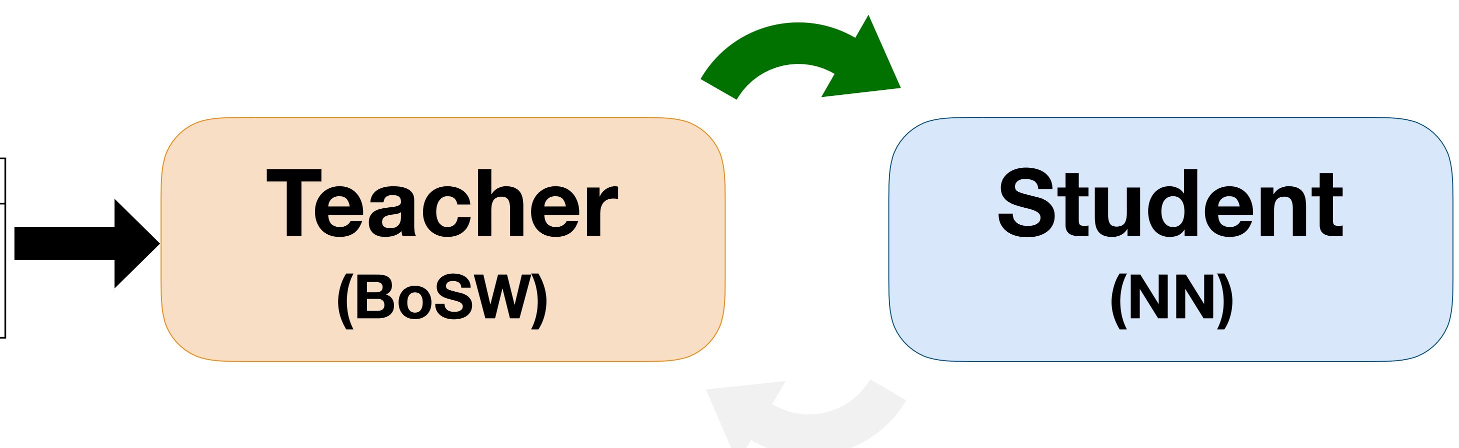

Training Student Using Teacher’s Predictions

- **Student** is trained through the “distillation” objective [Ba & Caruana, 2014; Hinton et al 2015]

$$H(q_i, p_i) = -\sum_k q_i^k \log p_i^k$$

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

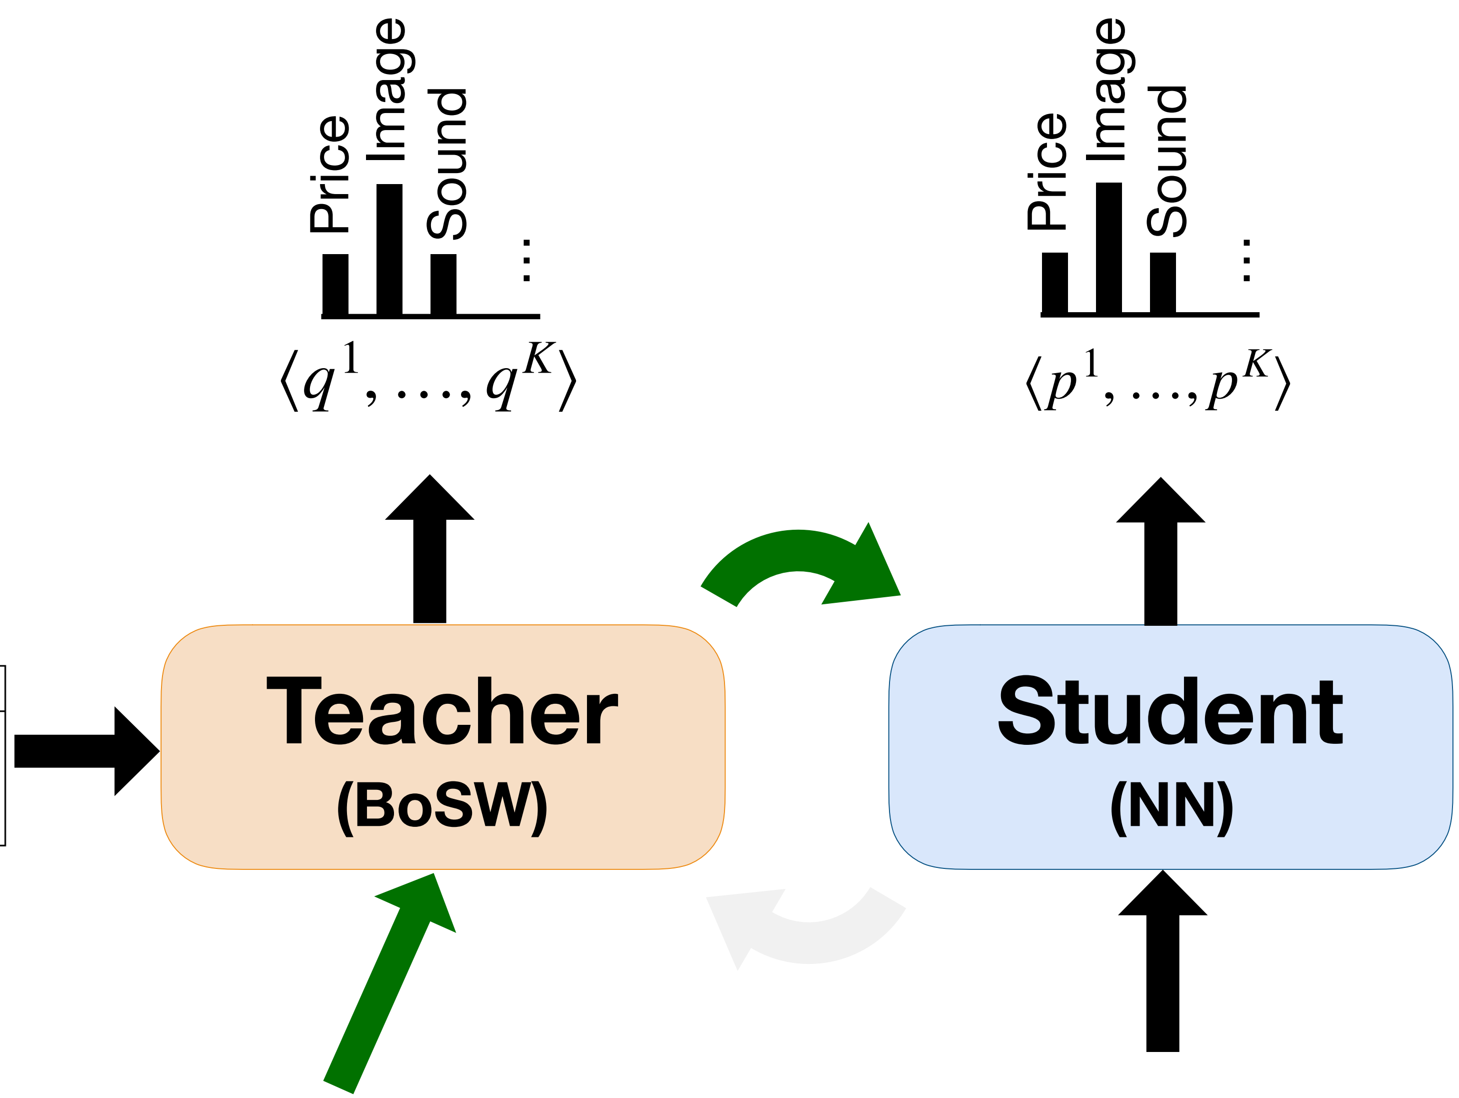

segment: *“The **picture** looks very pixelated when playing blu-ray movies”*

Training Student Using Teacher's Predictions

- **Student** is trained through the "distillation" objective [Ba & Caruana, 2014; Hinton et al 2015]
- **Student** also considers the **context** of **seed words**

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

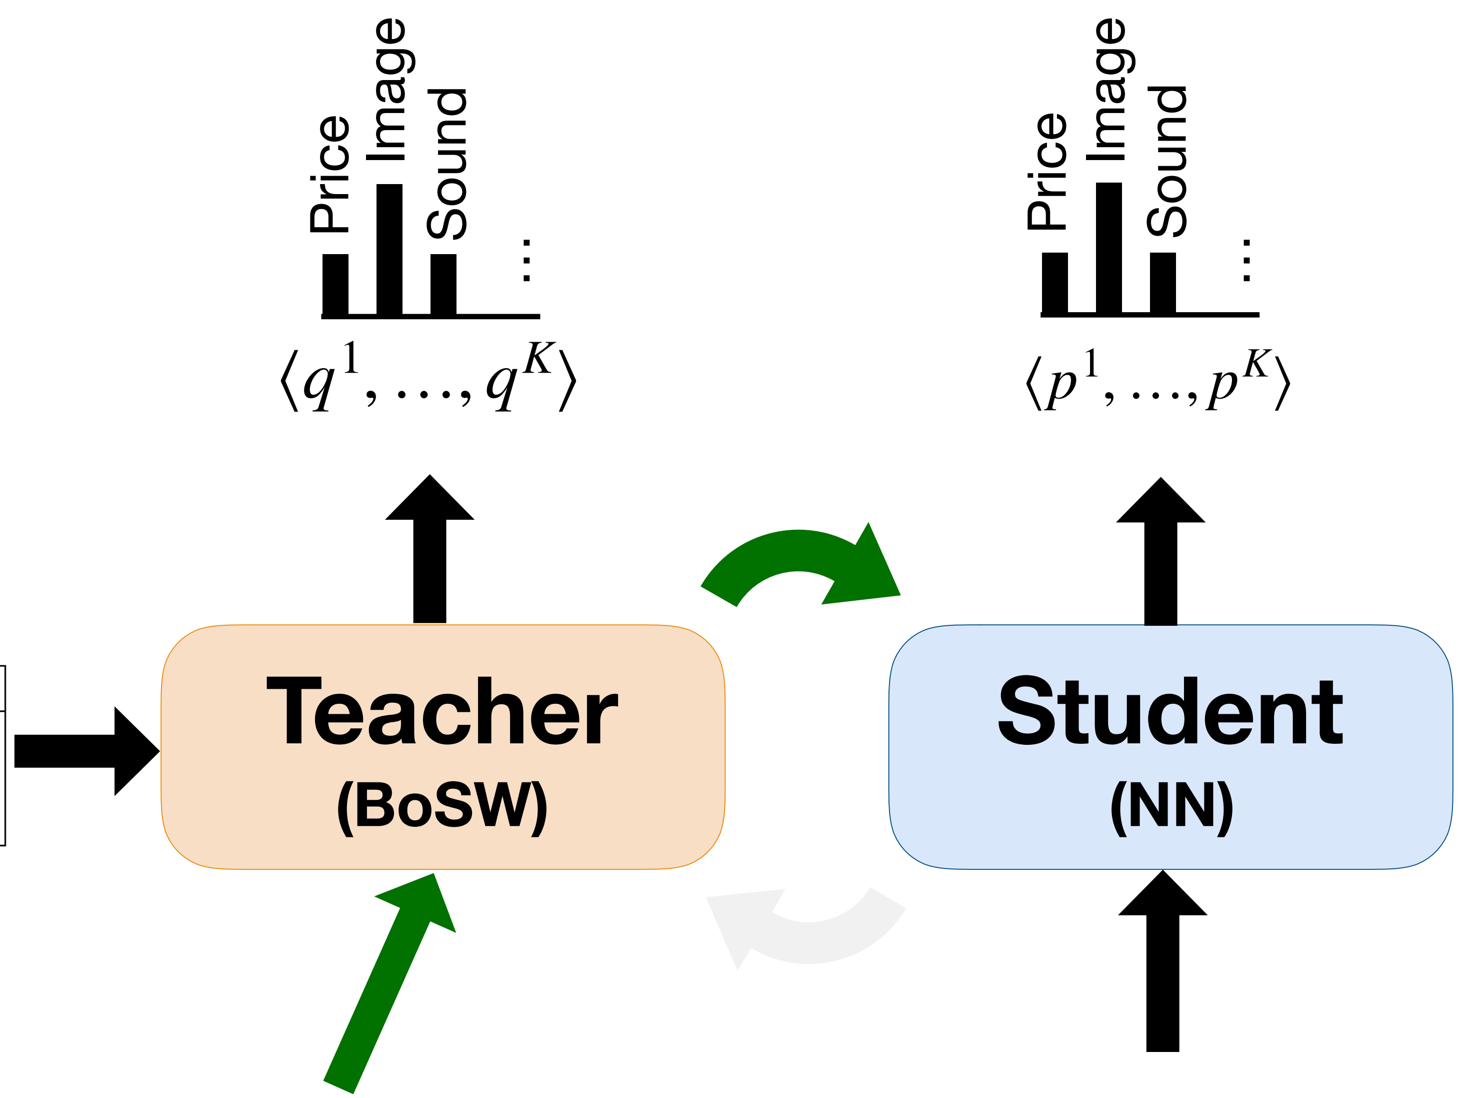

segment: *"The **picture** looks very pixelated when playing blu-ray movies"*

Training Student Using Teacher's Predictions

- **Student** is trained through the "distillation" objective [Ba & Caruana, 2014; Hinton et al 2015]
- **Student** also considers the **context** of **seed words**
- **Student** predicts aspects even if no **seed words** appear

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

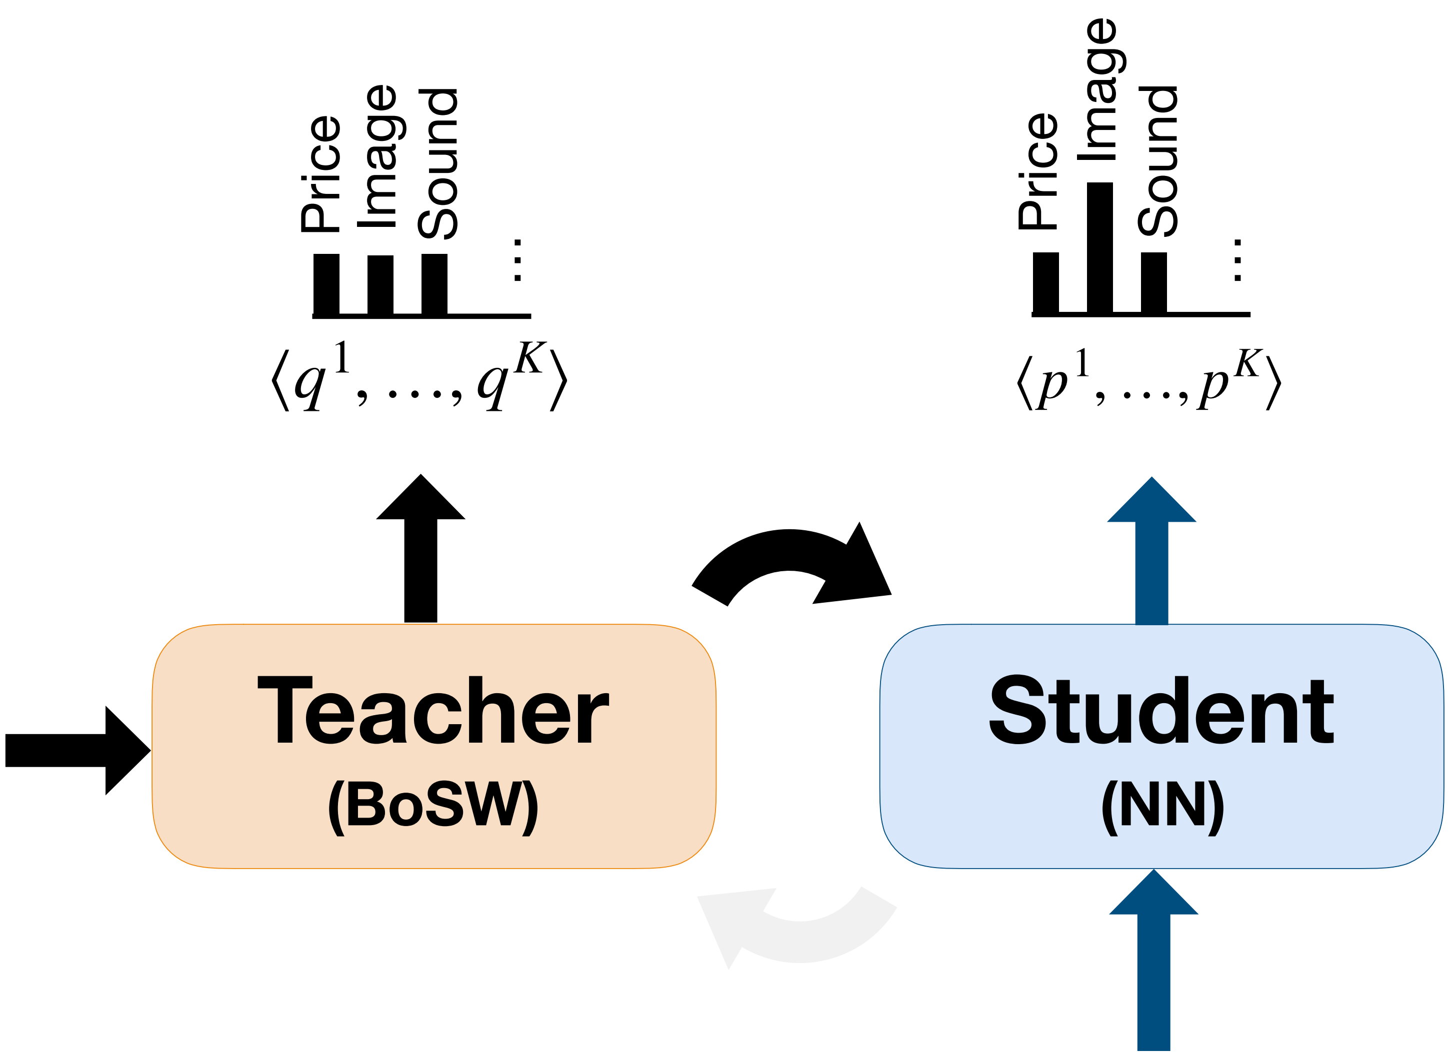

segment: *"The <UNK> looks very **pixelated** when playing **blu-ray movies**"*

Weakly Supervised Co-Training

1. Teacher
2. Student
3. Teacher -> Student
4. **Student -> Teacher**

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

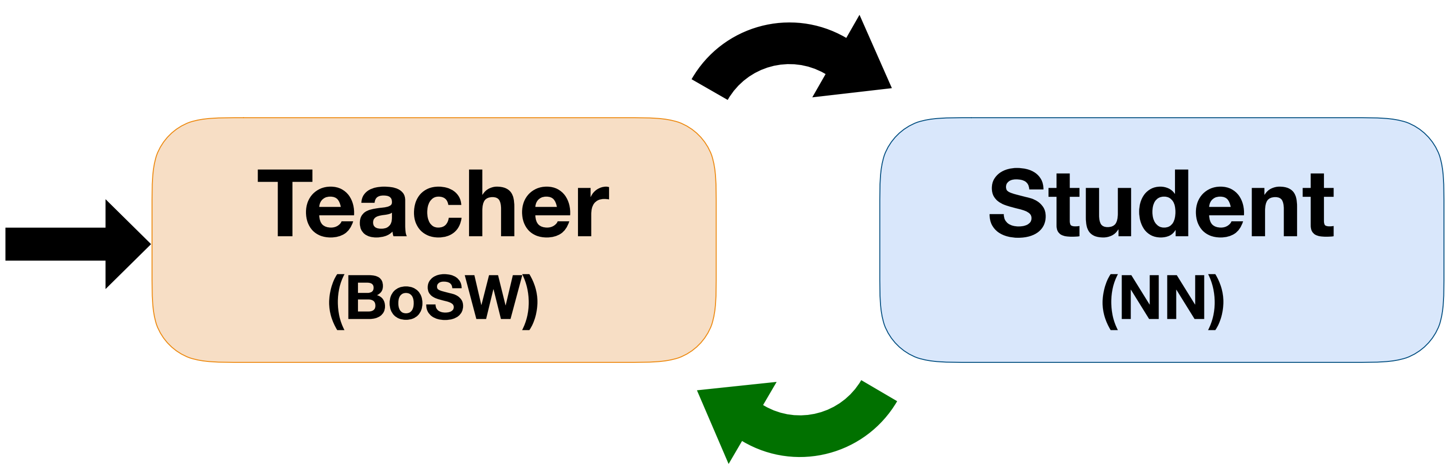

Improving Teacher Using Student's Predictions

- Seed words may be **noisy**: only weakly indicative of aspects

Qualities	Price	Image	Sound
bright	**0%**	**100%**	**0%**

"bright" should also
be indicative of **"Sound"** aspect

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

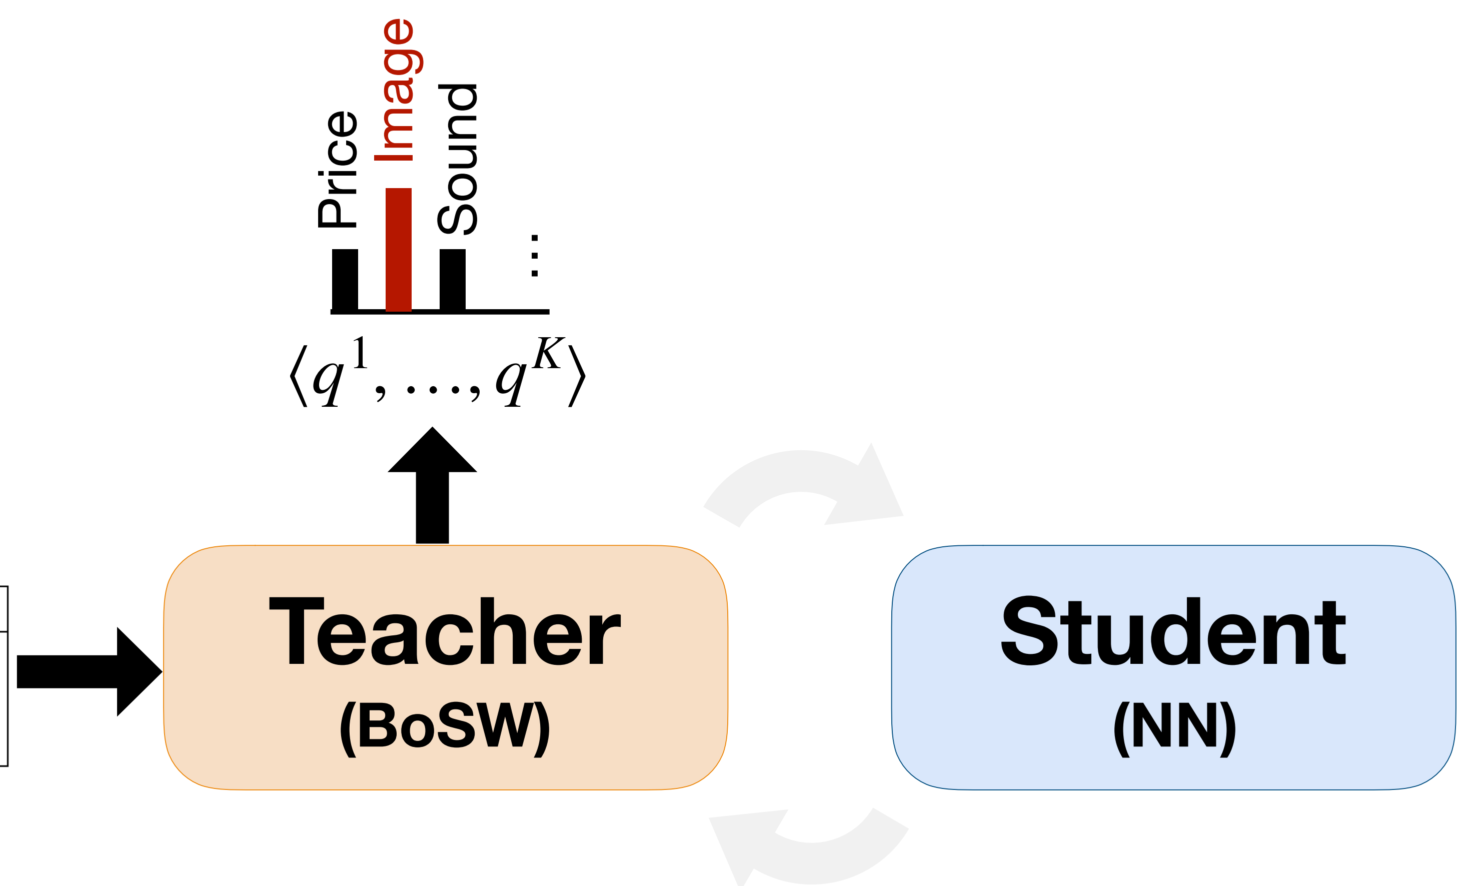

Improving Teacher Using Student’s Predictions

- Seed words may be **noisy**: only weakly indicative of aspects
- **Teacher** estimates seed word **qualities** … using **Student**’s predictions!

Qualities	**Price**	**Image**	**Sound**
bright	**0%**	**100%**	**0%**

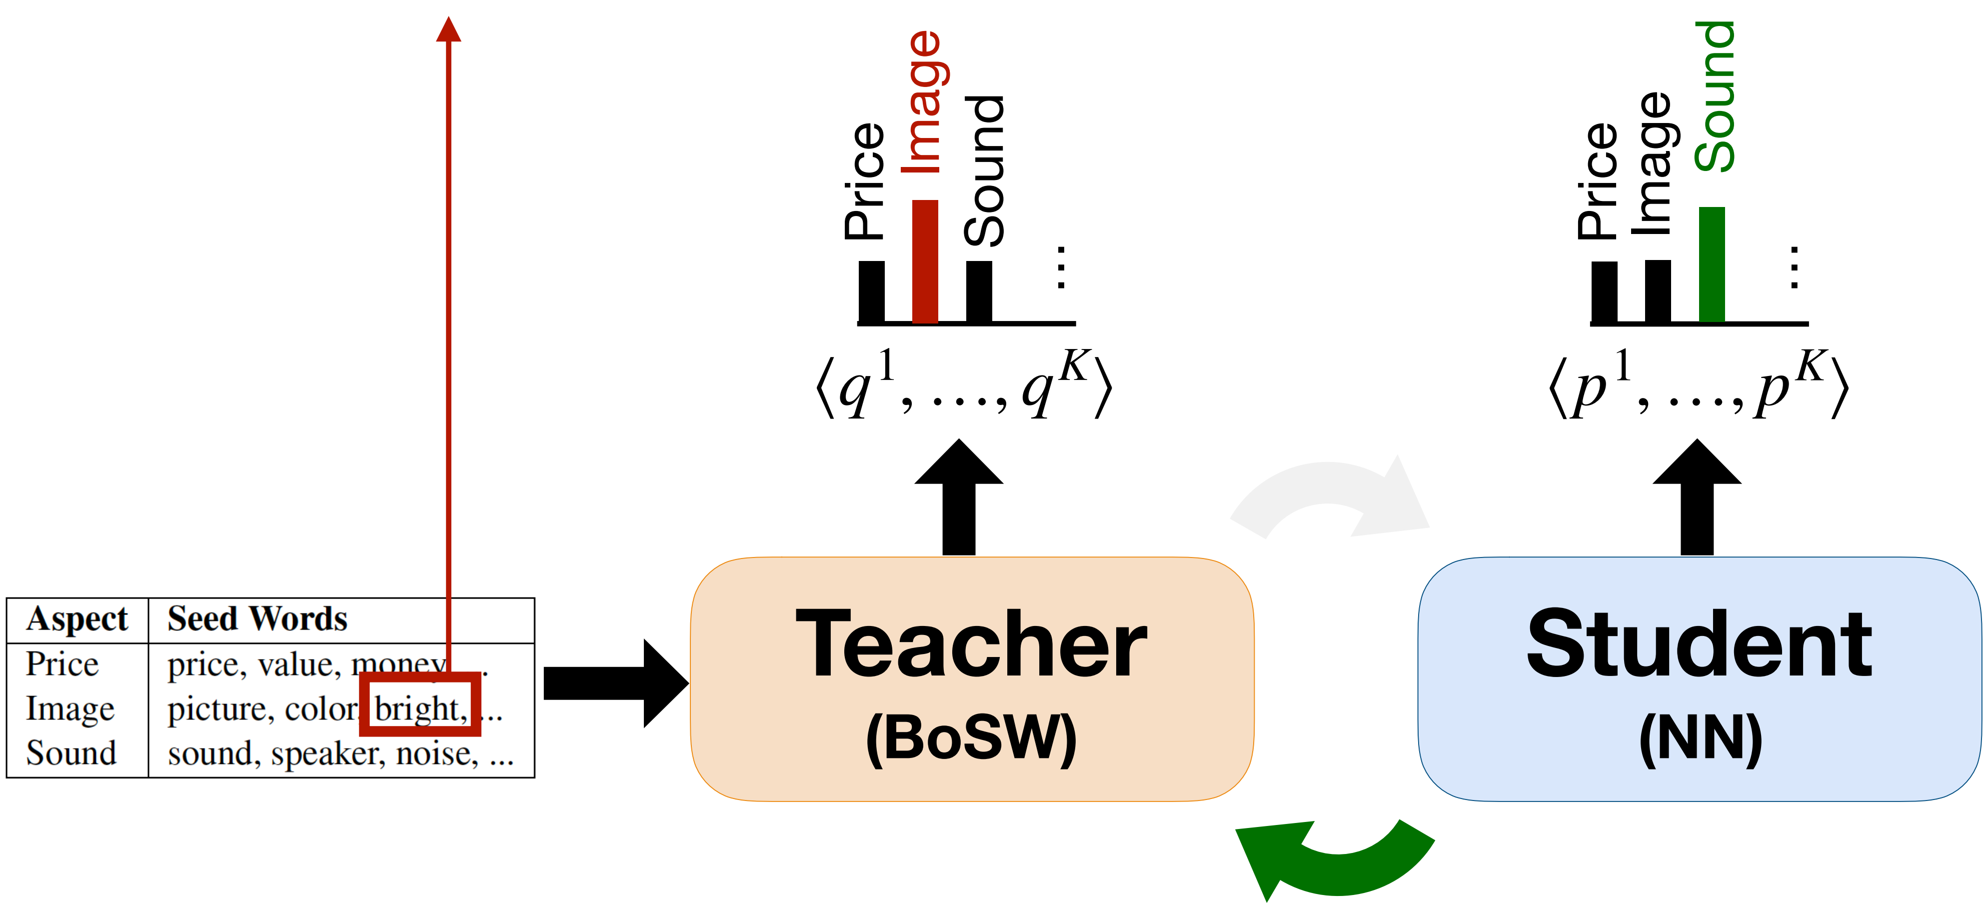

Improving Teacher Using Student's Predictions

- Seed words may be **noisy**: only weakly indicative of aspects
- **Teacher** estimates seed word **qualities** … using **Student**'s predictions!

Qualities	**Price**	**Image**	**Sound**
~~**bright**~~	~~**0%**~~	~~**100%**~~	~~**0%**~~
bright	**0%**	**70%**	**30%**

Model Bootstrapped EM (Khetan et al. 2018)

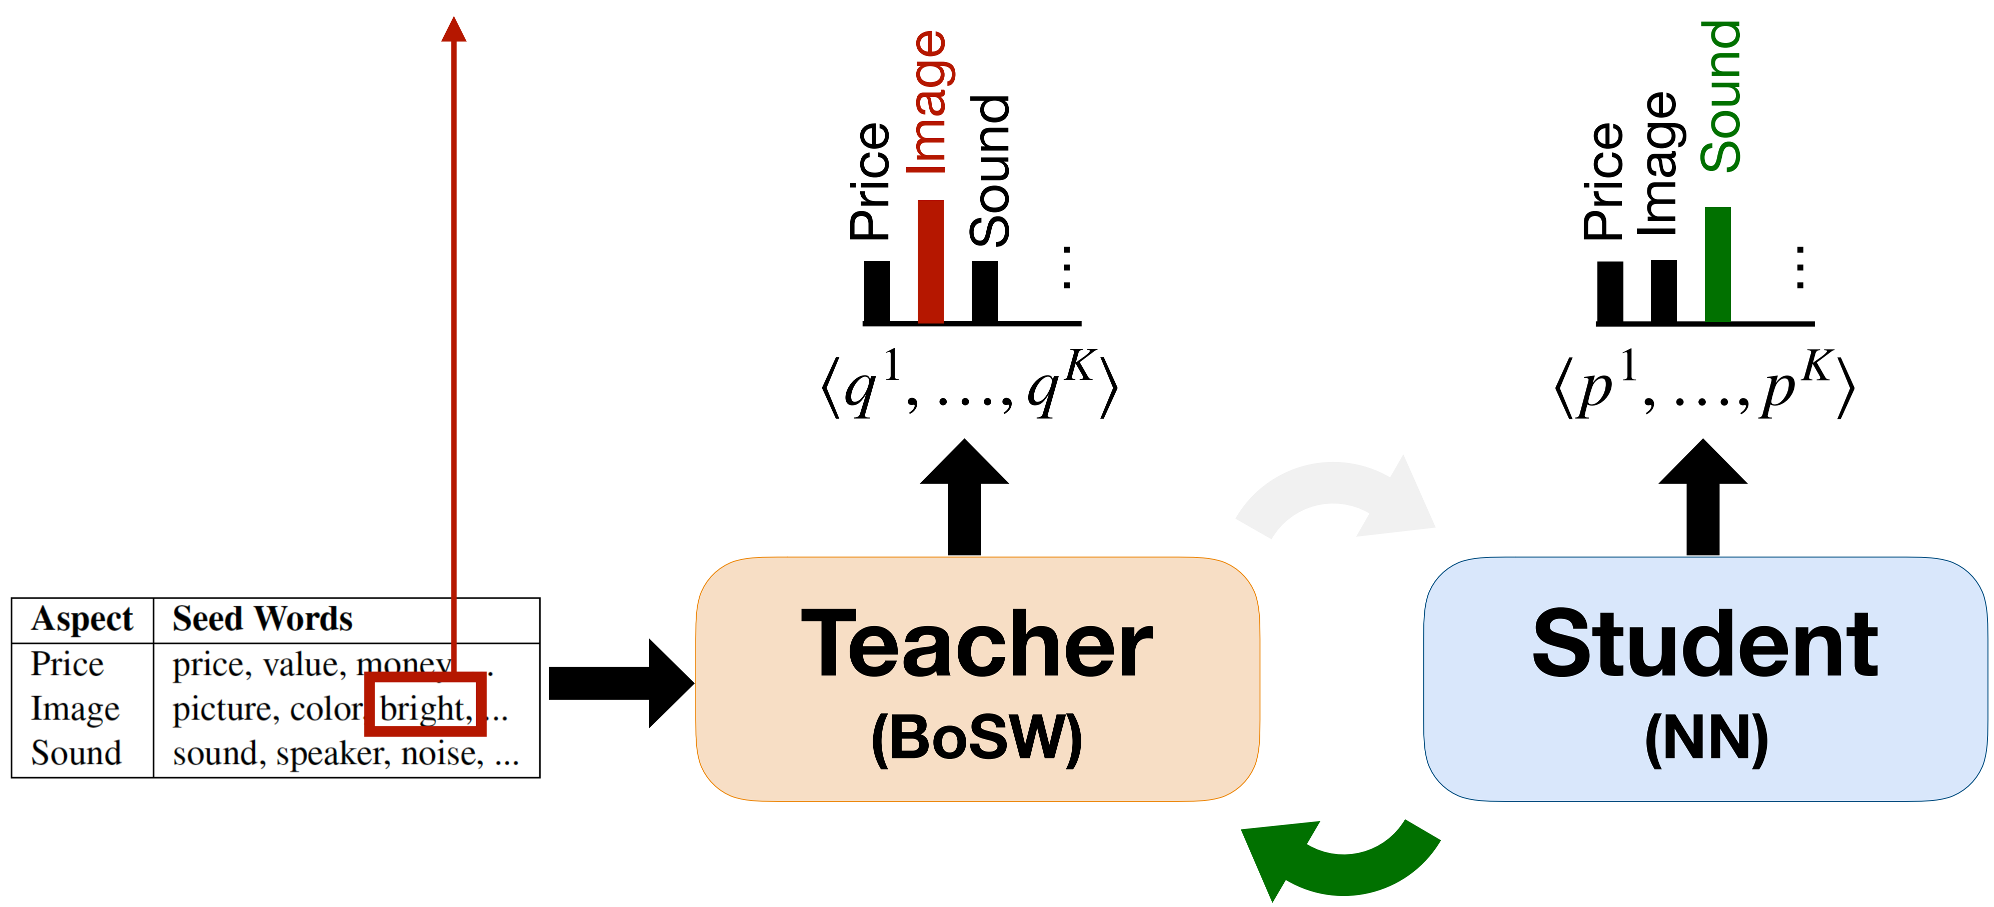

Improved Teacher Predicts Weighted Aspect Probabilities

- **Improved Teacher:**
  - Uses **seed words** and **quality estimates** to predict aspect probabilities $\langle q^1, \ldots, q^K \rangle$

Qualities	Price	Image	Sound
bright	0.0	0.7	0.3
noise	0.0	0.2	0.8
...	...	...	...

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

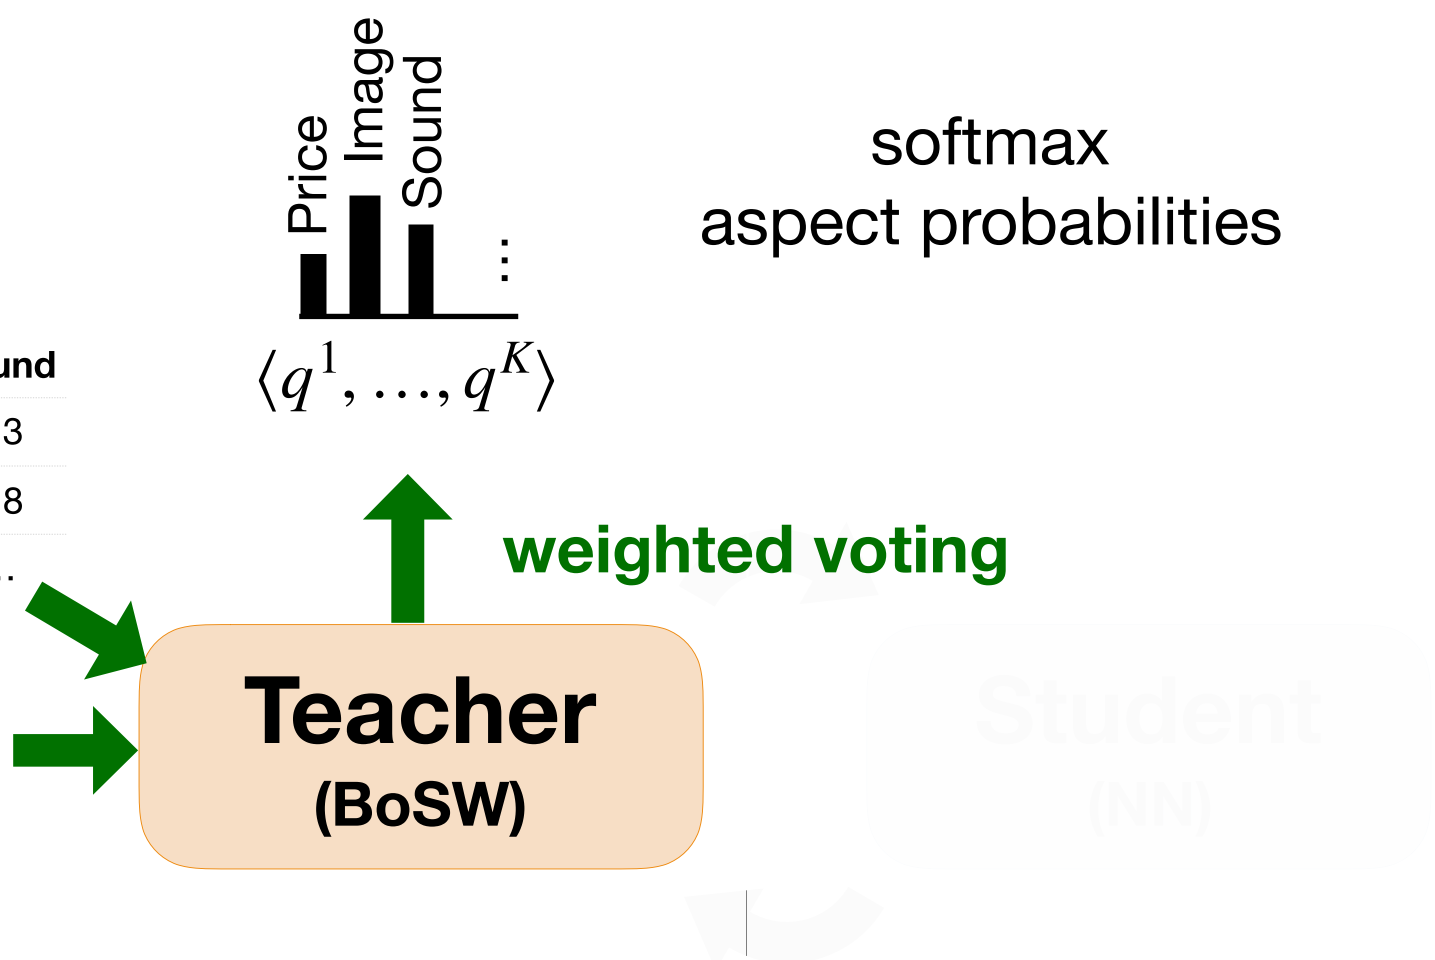

Iterative Co-Training Pipeline

- **T0:** Apply **Teacher** on unlabeled segments (assuming “perfect” seed words)
- **S0:** Train **Student** using **Teacher**’s predictions
- **T1:** Update **Teacher**’s weights (seed word qualities) using **Student**’s predictions
- **S1:** … Iterate until convergence

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

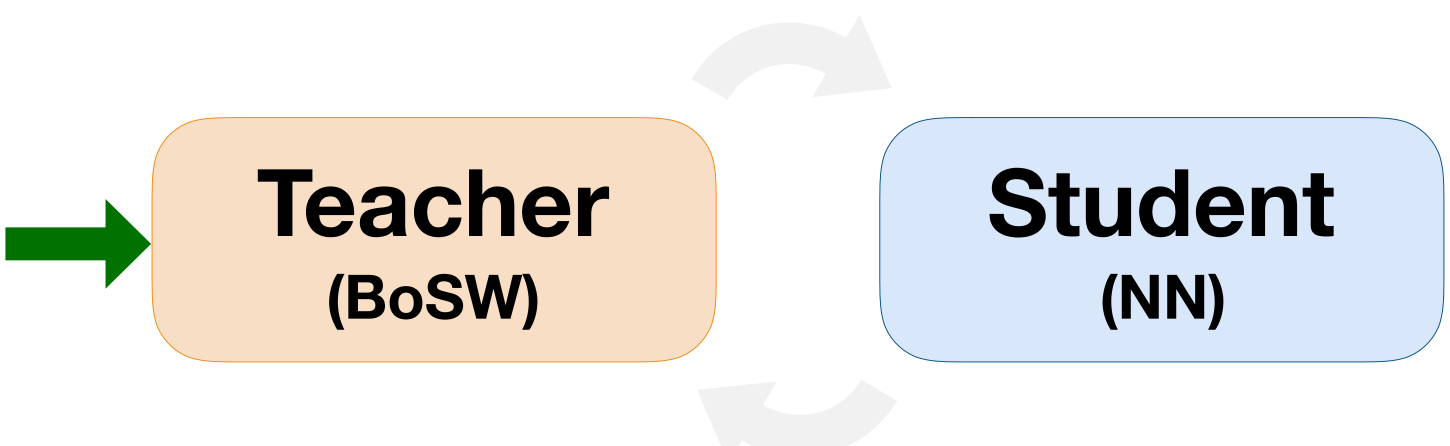

Iterative Co-Training Pipeline

- **T0:** Apply **Teacher** on unlabeled segments (assuming "perfect" seed words)
- **S0:** Train **Student** using **Teacher**'s predictions
- **T1:** Update **Teacher**'s weights (seed word qualities) using **Student**'s predictions
- **S1:** … Iterate until convergence

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

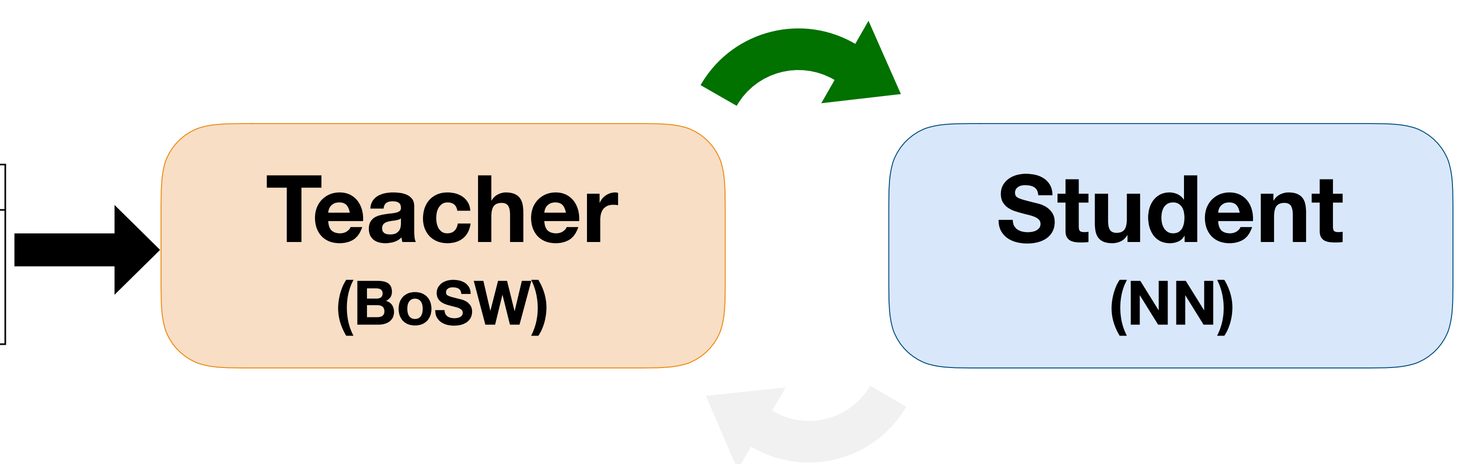

Iterative Co-Training Pipeline

- **T0:** Apply **Teacher** on unlabeled segments (assuming "perfect" seed words)
- **S0:** Train **Student** using **Teacher**'s predictions
- **T1:** Update **Teacher**'s weights (seed word qualities) using **Student**'s predictions
- **S1:** … Iterate until convergence

Qualities	Price	Image	Sound
bright	0.0	0.7	0.3
noise	0.0	0.2	0.8
…	…	…	…

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

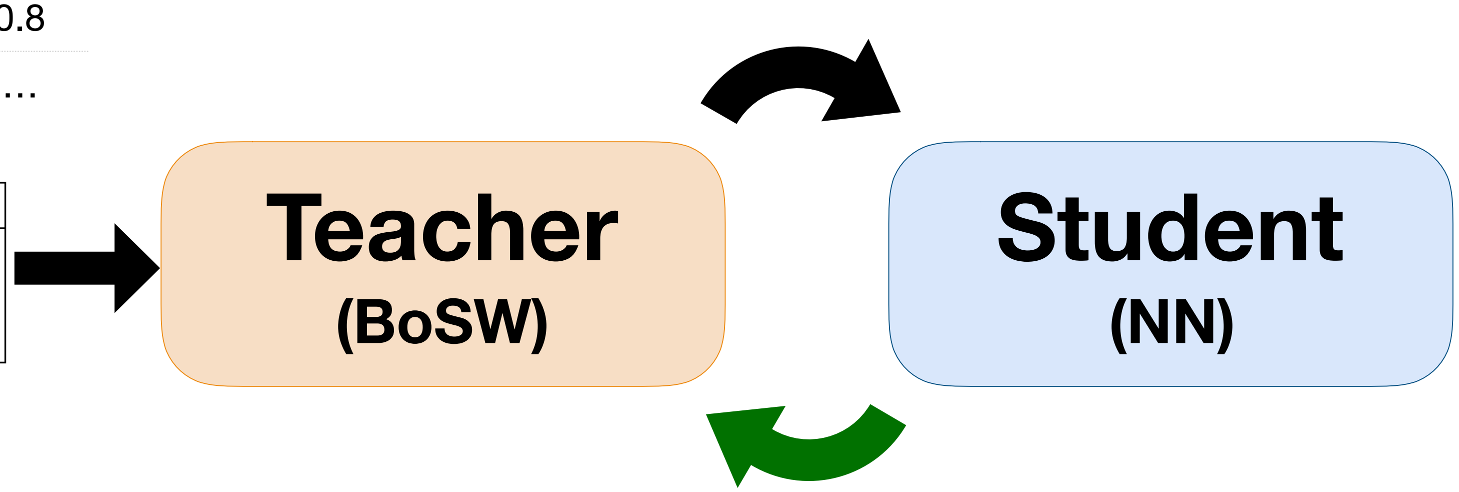

Iterative Co-Training Pipeline

- **T0:** Apply **Teacher** on unlabeled segments (assuming “perfect” seed words)
- **S0:** Train **Student** using **Teacher**’s predictions
- **T1:** Update **Teacher**’s weights (seed word qualities) using **Student**’s predictions
- **S1:** … Iterate until convergence

Qualities	Price	Image	Sound
bright	0.0	0.7	0.3
noise	0.0	0.2	0.8
…	…	…	…

Aspect	Seed Words
Price	price, value, money, ...
Image	picture, color, bright, ...
Sound	sound, speaker, noise, ...

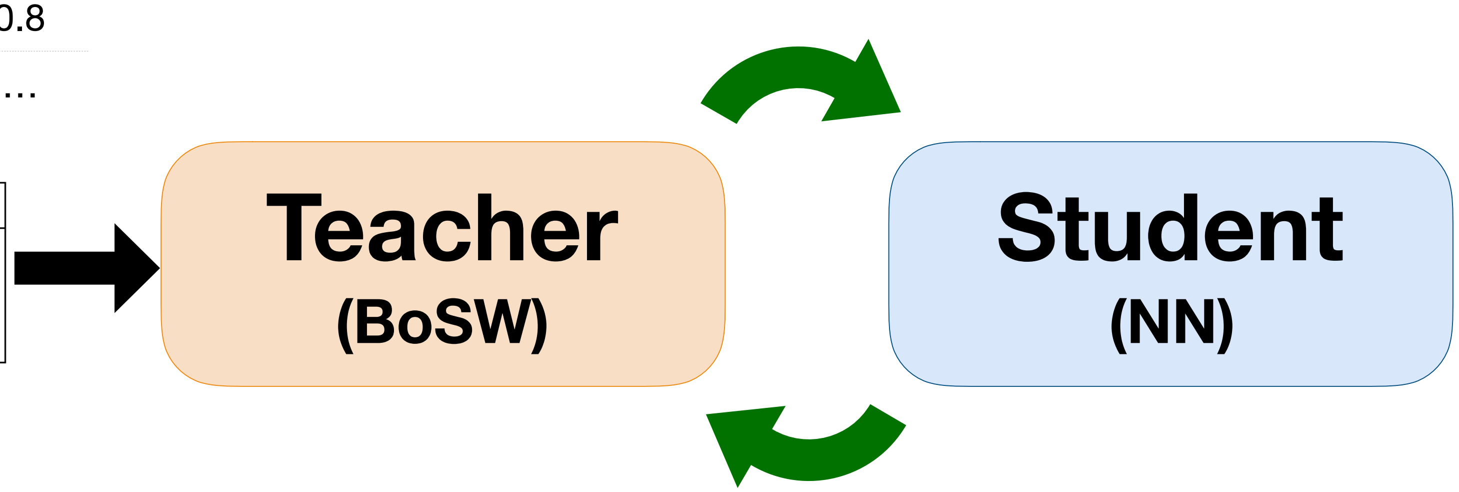

Outline

1. Fine-Grained Aspect Detection with Seed Words

2. Weakly Supervised Co-Training

3. Experiments

4. Conclusions and Future Work

Experiments: Segment-Level Aspect Detection

- **Datasets:**
  1. OPOSUM-Bags&Cases
  2. OPOSUM-Keyboards
  3. OPOSUM-Boots
  4. OPOSUM-Bluetooth Headsets
  5. OPOSUM-TVs
  6. OPOSUM-Vacuums
  7. SemEval-Restaurants-English
  8. SemEval-Restaurants-Spanish
  9. SemEval-Restaurants-French
  10. SemEval-Restaurants-Russian
  11. SemEval-Restaurants-Dutch
  12. SemEval-Restaurants-Turkish

OPOSUM

Amazon product reviews

9 aspects / domain: Quality, Looks, Price, …

SemEval-2016

Restaurant reviews

12 aspects / language: Ambience, Service, Food, …

Experiments: Segment-Level Aspect Detection

- **Model Comparison**

    1. **LDA-Anchors** (Lund et al., 2017)
    2. **ABAE:** Aspect-Based AutoEncoder (He et al., 2017)
    3. **MATE:** Multi-Seed Aspect Extractor (Angelidis & Lapata, 2018)
    4. Our **Teacher-Student** framework

Experiments: Segment-Level Aspect Detection

- **Model Comparison**

  1. **LDA-Anchors** (Lund et al., 2017)
  2. **ABAE:** Aspect-Based AutoEncoder (He et al., 2017)
  3. **MATE:** Multi-Seed Aspect Extractor (Angelidis & Lapata, 2018)
  4. Our **Teacher-Student** framework

same
seed words

Experiments: Segment-Level Aspect Detection

- **Model Comparison**

  1. **LDA-Anchors** (Lund et al., 2017)
  2. **ABAE:** Aspect-Based AutoEncoder (He et al., 2017)
  3. **MATE:** Multi-Seed Aspect Extractor (Angelidis & Lapata, 2018)
  4. Our **Teacher-Student** framework

  same seed words

- **Methodology:** experiments on 12 datasets separately
  - Training:
    - Use 1M **unlabeled** segments
    - Use 30 **seed words** per aspect (extracted by Angelidis & Lapata, 2018)
  - Evaluation:
    - Use 750 **labeled** test segments
    - Report micro-averaged F1 (average over 5 different runs)

We Leverage the Predictive Power of Seed Words

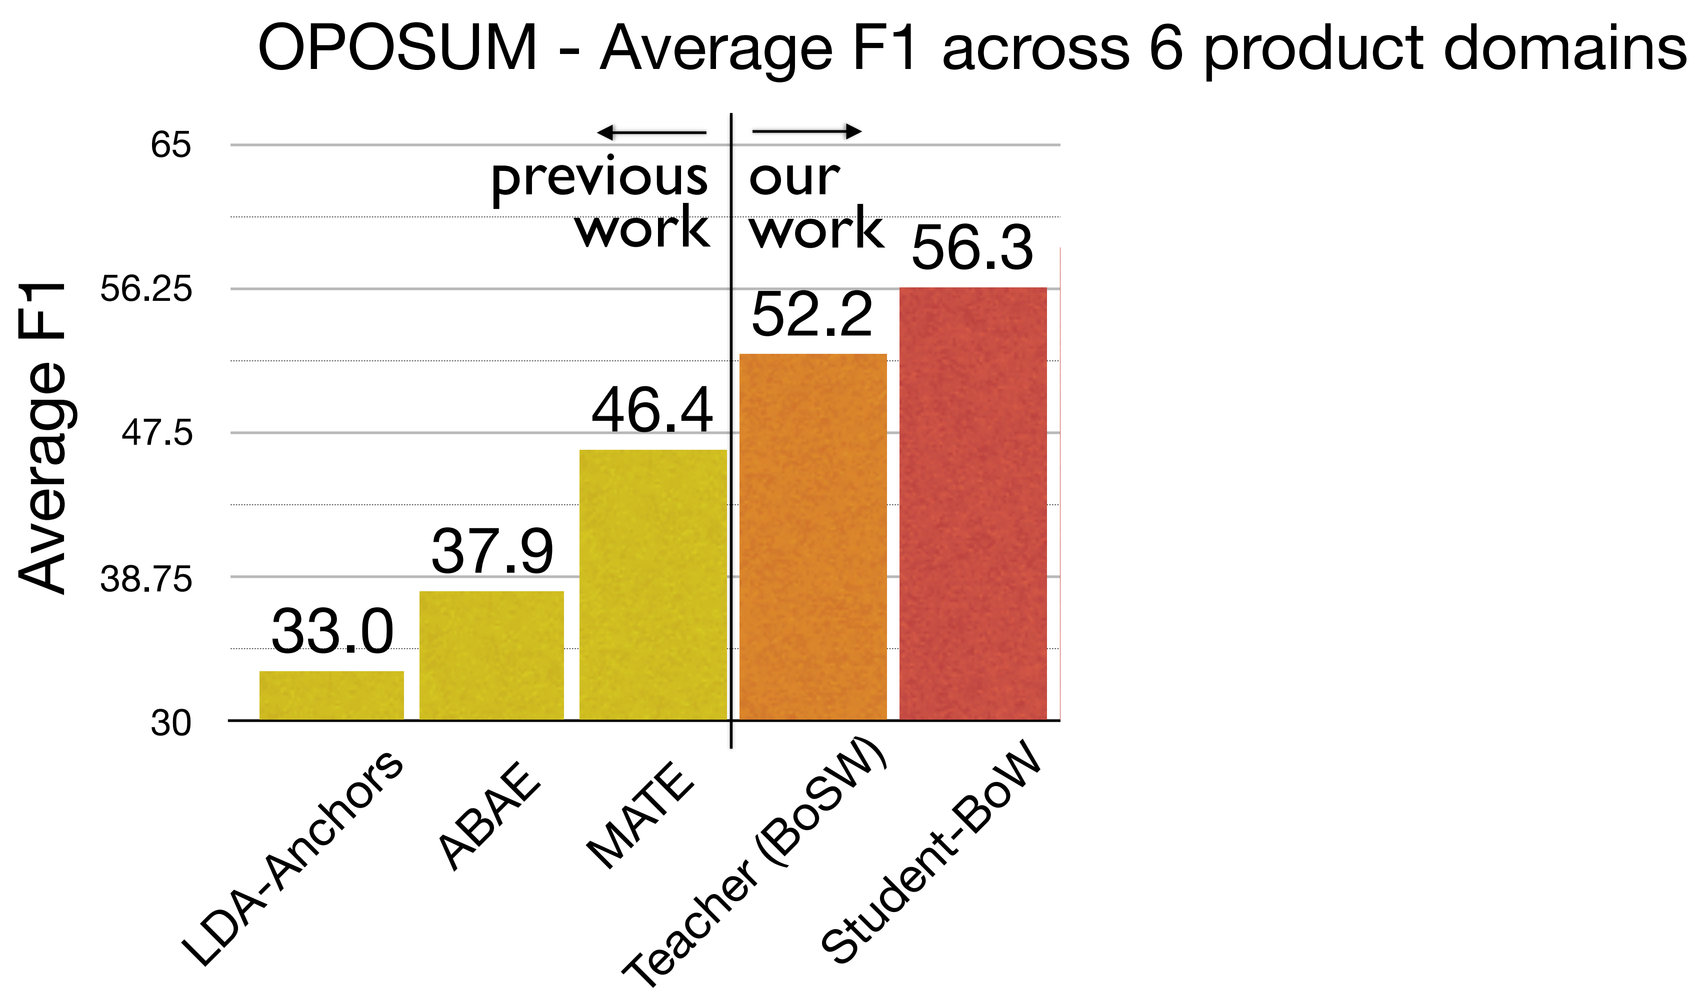

- **Teacher** and **Student-BoW** outperform previous approaches

Recent Embedding Techniques Boost Student's Performance

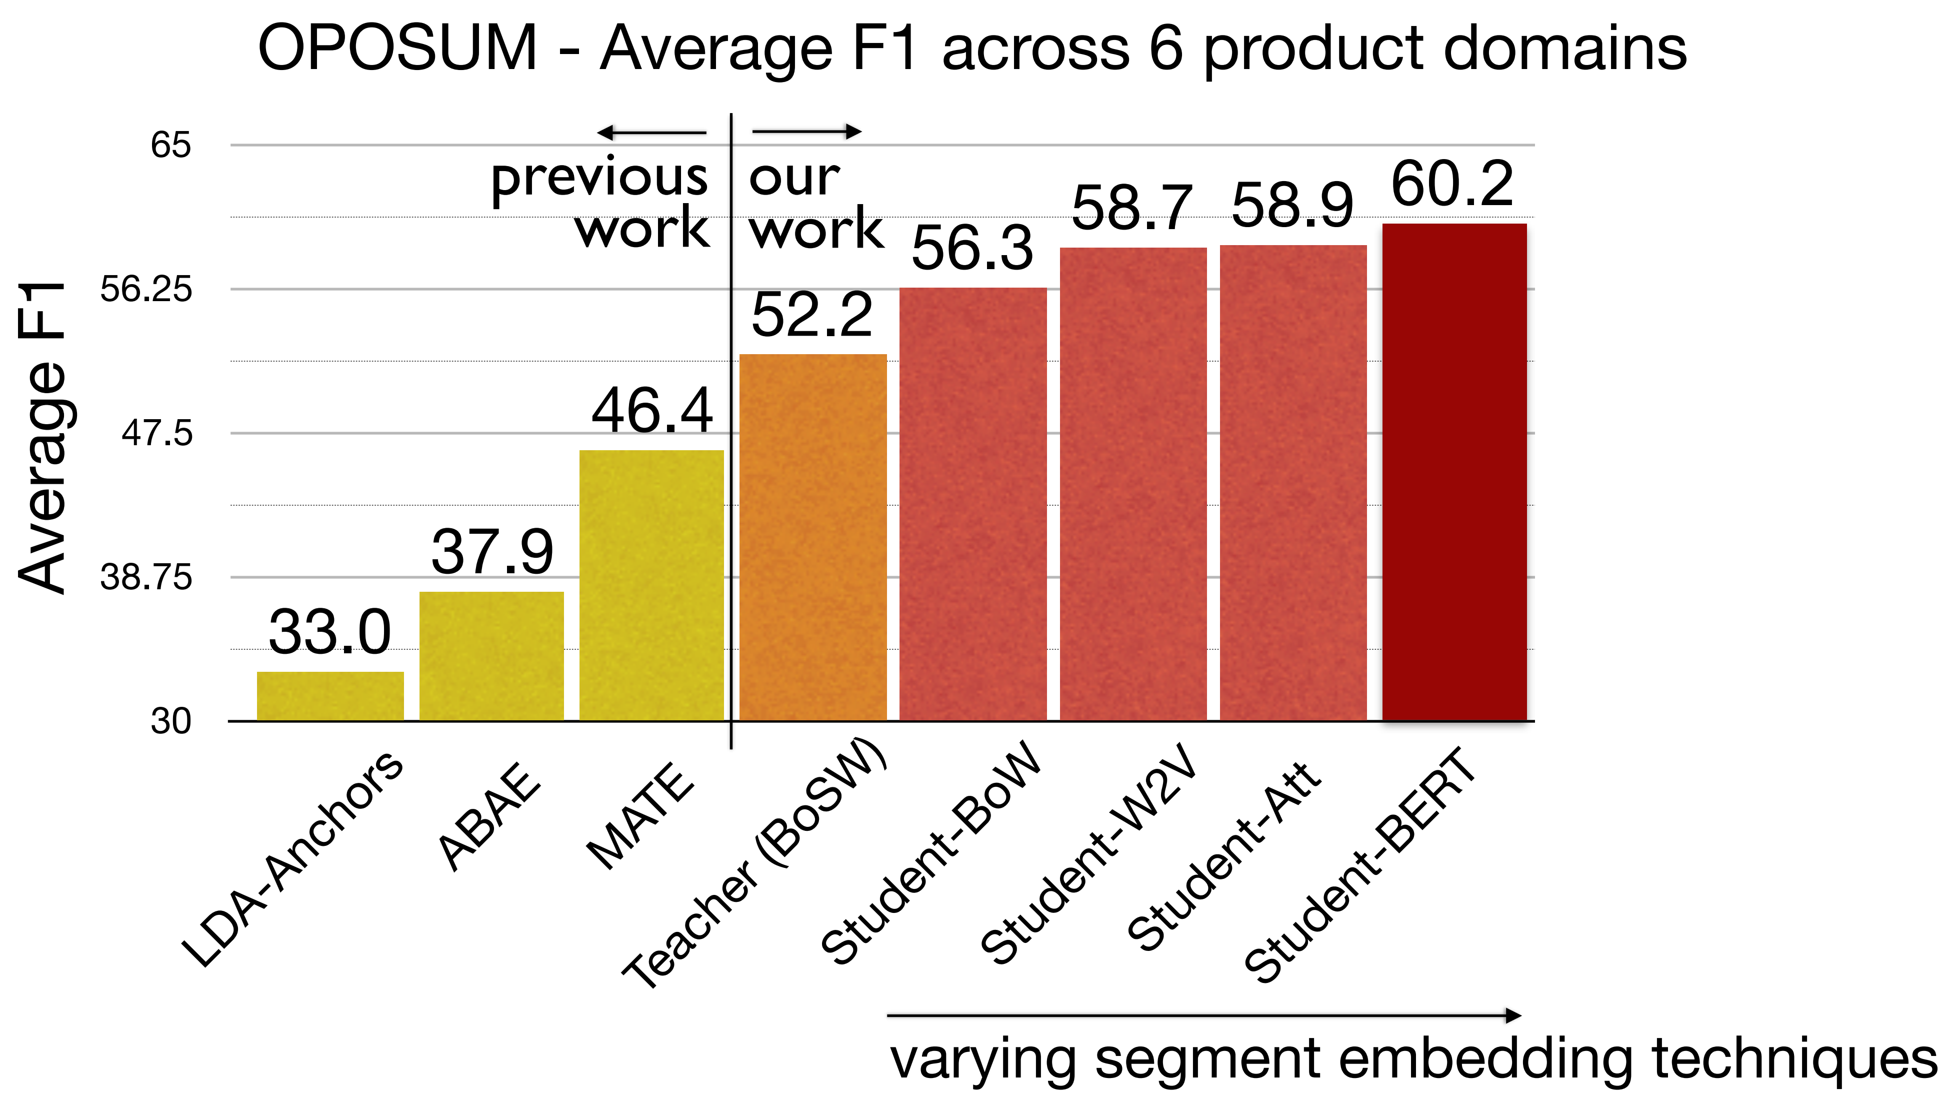

- Our framework allows using any type of embedding technique
- **Student-BERT** outperforms **MATE** by **29.7%** across 6 product domains
- **Student-*** outperform **MATE** across all domains and languages

More results in our paper.

Iterative Co-Training Leads to Better Teacher and Student

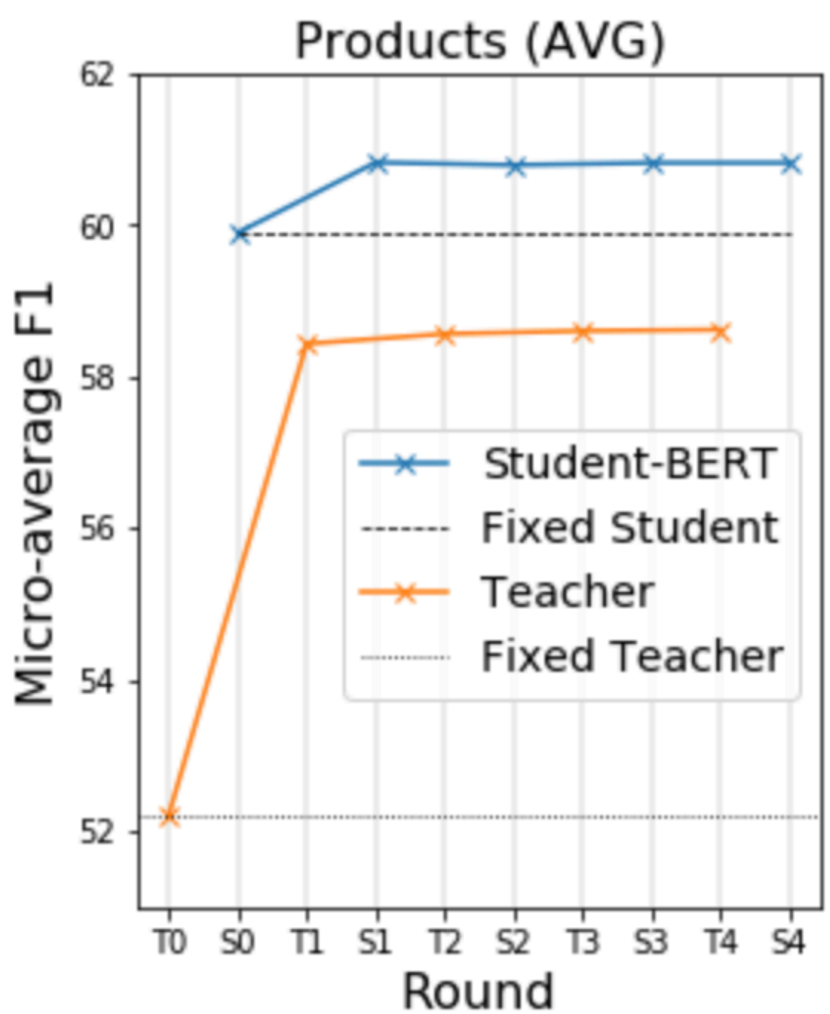

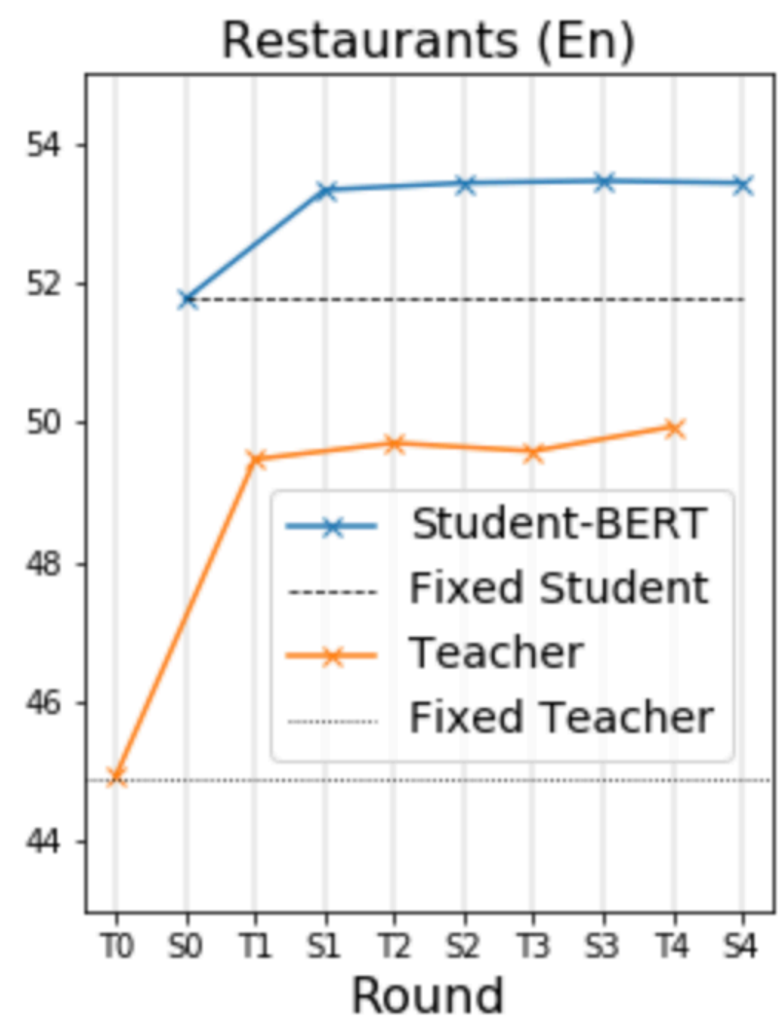

T<i>: **Teacher** update on round <i>
S<i>: **Student** update on round <i>

- Most improvement achieved after **one** round of co-training

Outline

1. Fine-Grained Aspect Detection with Seed Words

2. Weakly Supervised Co-Training

3. Experiments

4. **Conclusions and Future Work**

Summary: Leveraging Seed Words in NNs

- Previous work: use seed words for initialization of aspect embeddings
- We instead leverage the predictive power of seed words for training NNs
- Our **weakly supervised co-training** approach:
  1. Use **seed words** in a bag-of-seed-words **Teacher**
  2. Use **Teacher**’s soft predictions to train a **Student** NN (through distillation)
  3. Use iterative **Teacher-Student** co-training to cope with **noisy seed words**
- Future work:
  1. Discover **new seed words** with a human in the loop
  2. Leverage **logical rules**
  3. Apply to **more NLP tasks** where classes are indicated by seed words/rules

Thank you!

Contact
gkaraman@cs.columbia.edu
https://gkaramanolakis.github.io

Columbia University

Extra Slides

Learning with Seed Words

- We assume a small set of indicative **seed words** per aspect:

Aspect	Seed Words
Price	price, value, money, worth, paid
Image	picture, color, quality, black, bright
Sound	sound, speaker, noise, loud, volume

- **Seed words** are easier to collect than **ground truth labels**
- **Seed words** are potentially more powerful than **ground truth labels**

Traditional label y_i → **Single** segment s_i

Seed Word e.g., "picture" → **Multiple** segments containing seed word

Results - OPOSUM (product reviews)

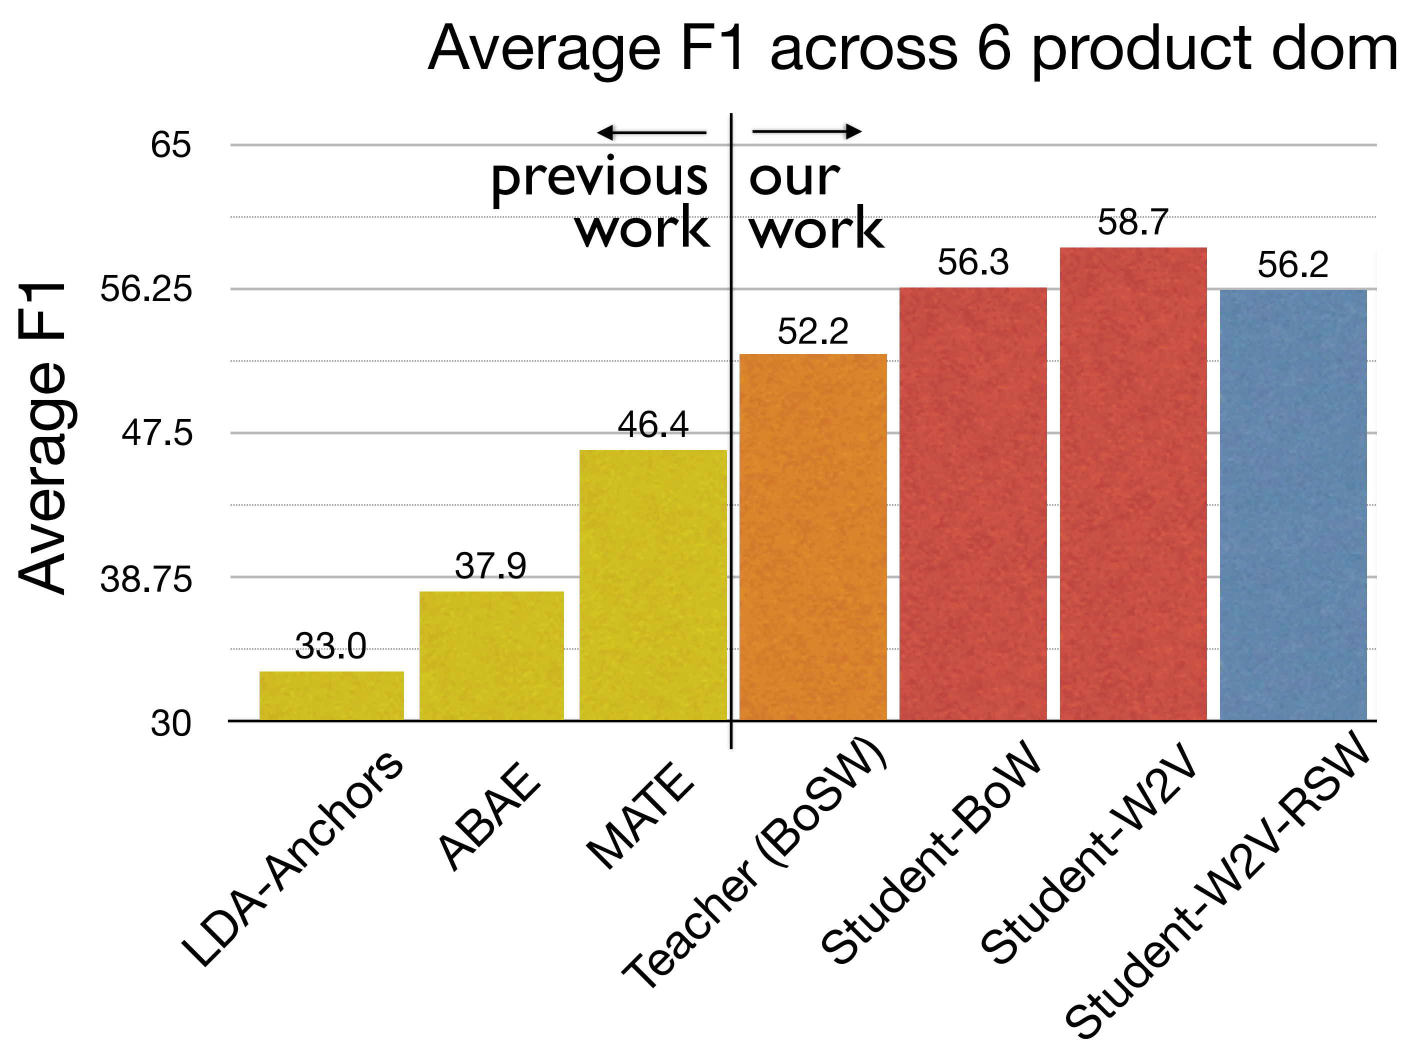

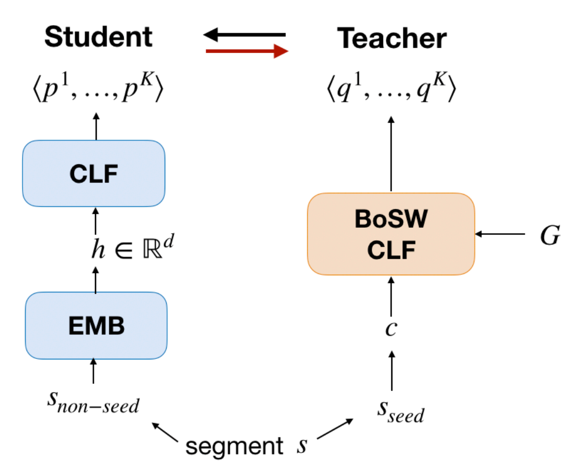

- **Student** leverages non-seed words:
    - Removing seed words (RSW) in **Student-RSW** better than **Teacher**

Results - All Datasets

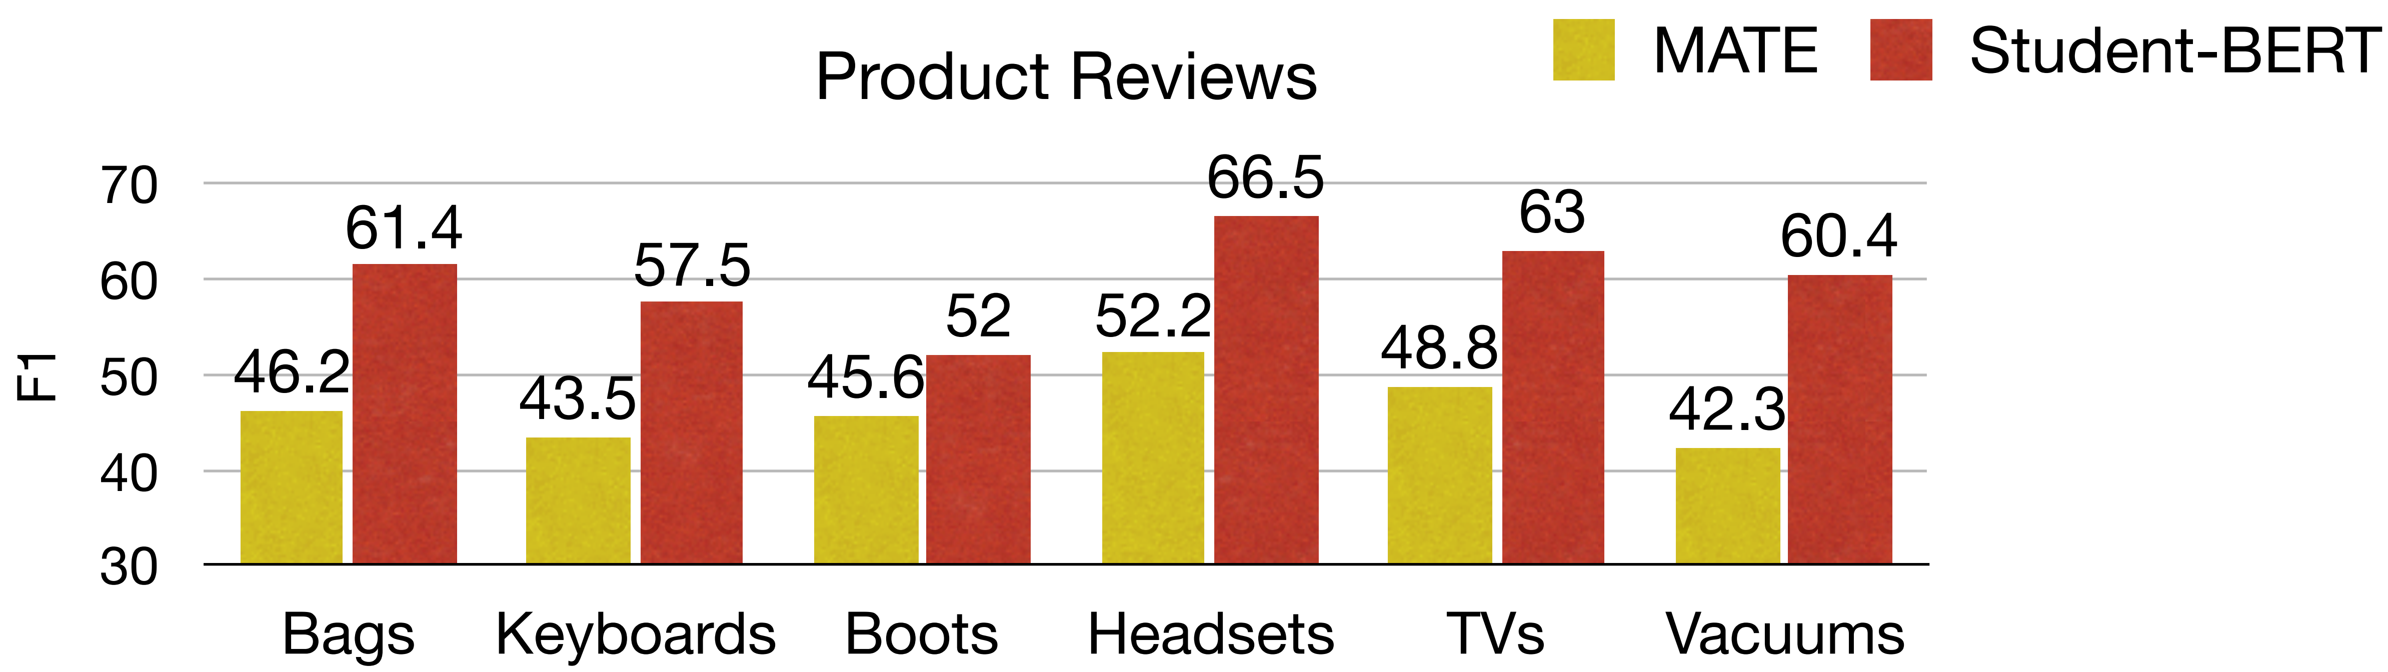

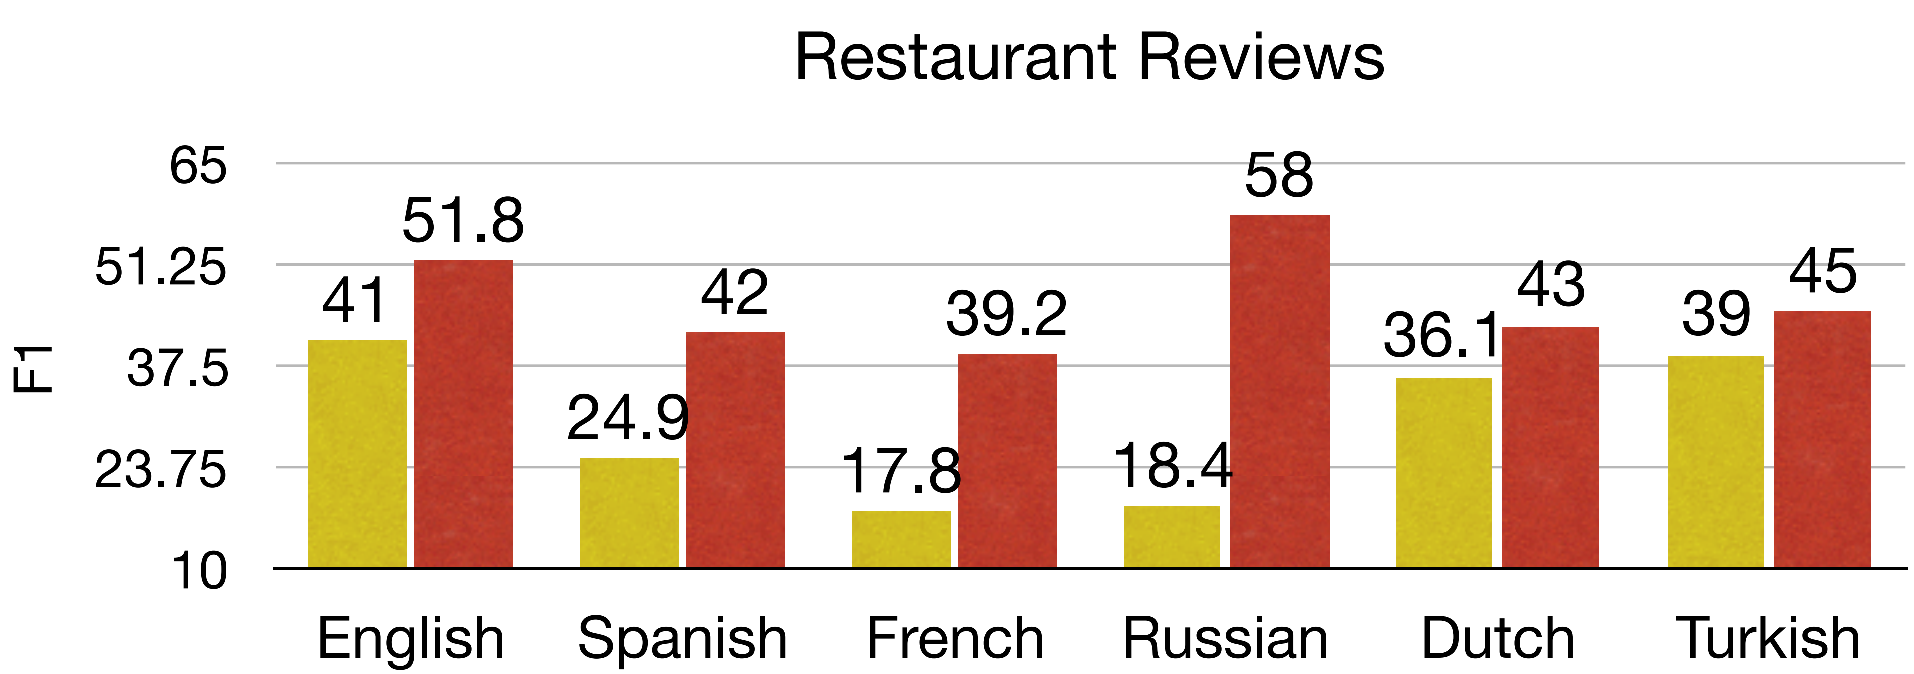

- **Student-BERT** outperforms **MATE** across all product domains and languages