

# PeTra: A Sparsely Supervised Memory Model for People Tracking

Shubham Toshniwal

Allyson Ettinger

Kevin Gimpel

Karen Livescu



[github.com/shtoshni92/petra](https://github.com/shtoshni92/petra)

## People Tracking

Amelia Shepherd, M.D. is a fictional character on Grey's Anatomy, portrayed by Caterina Scorsone. In her debut appearance, Amelia visited her former sister-in-law.

## People Tracking

*Amelia Shepherd*<sub>1</sub>, M.D. is a fictional character on Grey's Anatomy, portrayed by *Caterina Scorsone*<sub>2</sub>. In *her*<sub>1</sub> debut appearance, *Amelia*<sub>1</sub> visited *her*<sub>1</sub> former sister-in-law.

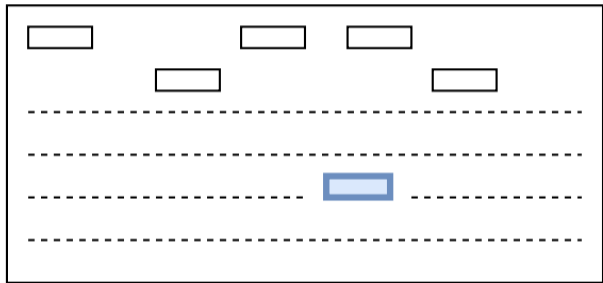
## People Tracking

**Amelia Shepherd**<sub>1</sub>, M.D. is a fictional character on Grey's Anatomy, portrayed by **Caterina Scorsone**<sub>2</sub>. In **her**<sub>1</sub> debut appearance, **Amelia**<sub>1</sub> visited **her**<sub>1</sub> former sister-in-law.

*Coreference Resolution for people references*

# Coreference Resolution

# Mention Ranking Models



Amelia Shepherd

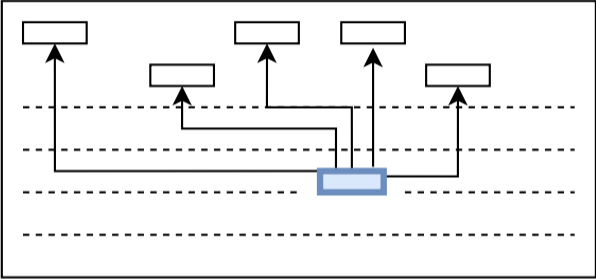
Caterina Scorsone

her

Amelia

her

# Mention Ranking Models





Amelia Shepherd   Caterina Scorsone   her   Amelia

her

Arrows point from the bottom her to the Amelia Shepherd, Caterina Scorsone, and her mentions. Arrows also point from the her mention to the Amelia mention.

# Mention Ranking Models



Impractical for long documents!



# Entity-Mention Models


Amelia Shepherd

Caterina Scorsone

⋮


Entities

# Entity-Mention Models



Entities

# Entity-Mention Models



Entities

Number of entities can be quite large!

Fixed-size Memory

PeTra

# Ideal Run




IG = Ignore, OW = Overwrite, CR = Coreference

# Ideal Run




IG = Ignore, OW = Overwrite, CR = Coreference

# Ideal Run



IG = Ignore, OW = Overwrite, CR = Coreference


# Ideal Run



IG = Ignore, OW = Overwrite, CR = Coreference

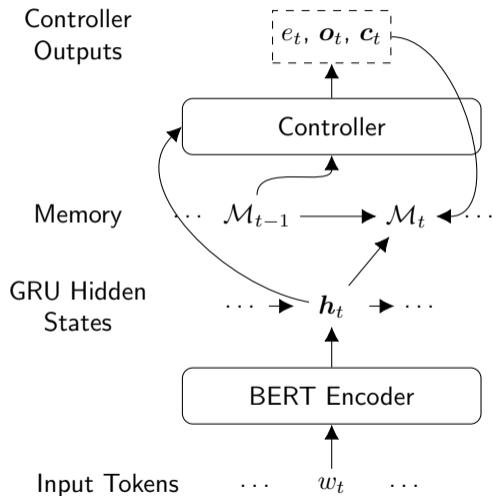


# Ideal Run



IG = Ignore, OW = Overwrite, CR = Coreference

# Model Details



## Key Outputs:

- $e_t$  Entity Probability
- $o_t$  Overwrite Prob. Vector
- $c_t$  Coreference Prob. Vector

## Sparse Supervision

*Amelia Shepherd*<sub>1</sub>, M.D. is a fictional character on Grey's Anatomy, portrayed by *Caterina Scorsone*<sub>2</sub>. In *her*<sub>1</sub> debut appearance, Amelia visited her former sister-in-law.

## Auxiliary Loss

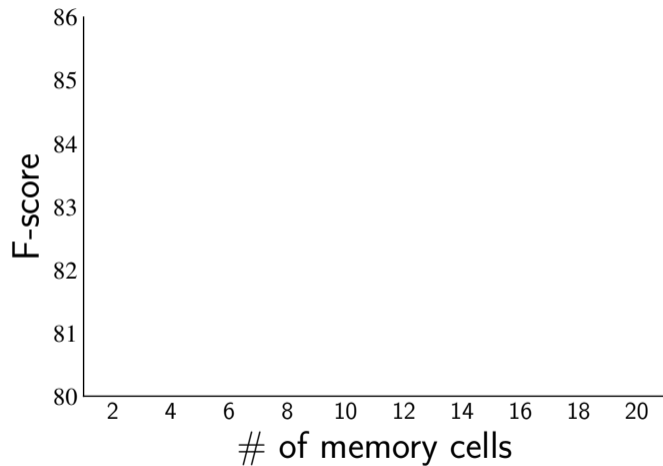
*Amelia Shepherd*<sub>1</sub>, M.D. is a fictional character on Grey's Anatomy, portrayed by *Caterina Scorsone*<sub>2</sub>. In *her*<sub>1</sub> debut appearance, Amelia visited her former sister-in-law.

Most tokens don't refer to people!

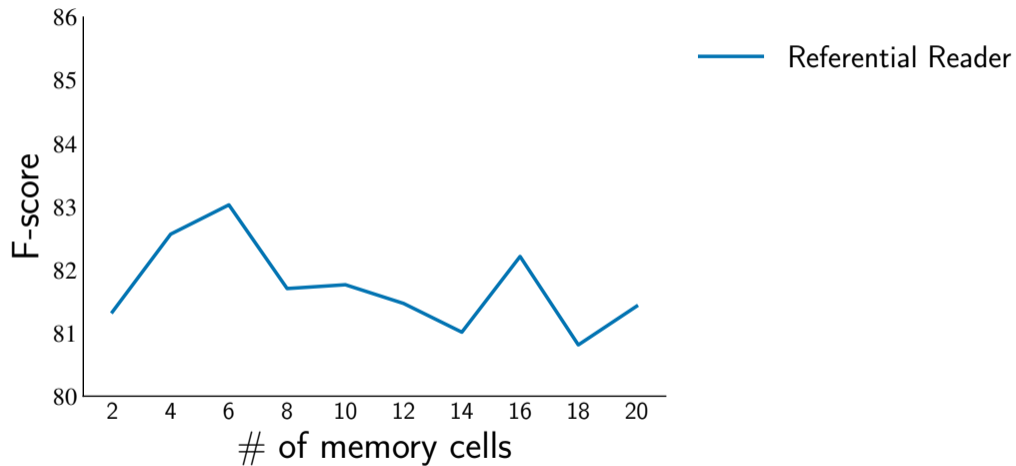
Penalize high entity probability outside labeled spans.

# Results

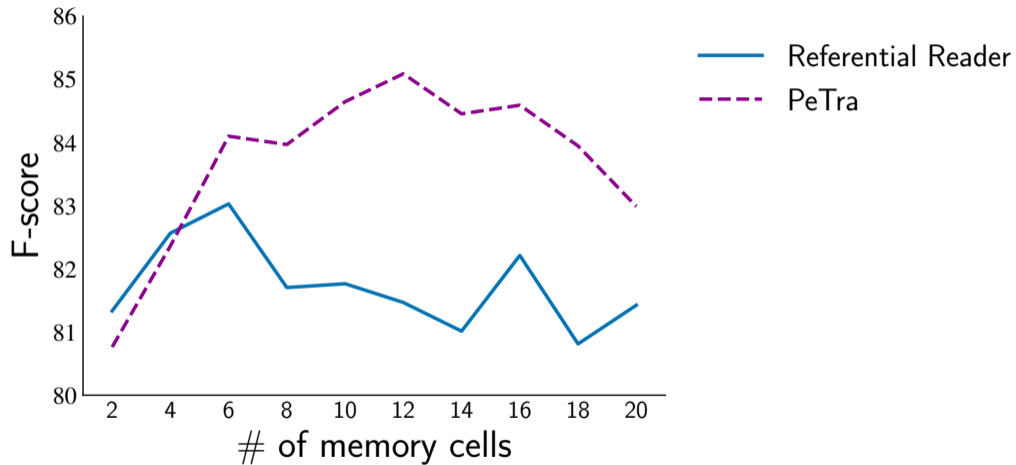
# GAP Results



## GAP Results

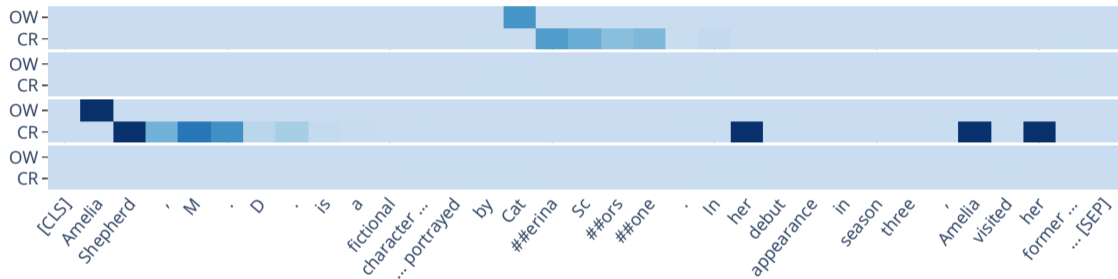


## GAP Results






# Memory Log Visualization



PeTra can resolve references beyond pronouns!

# Interpretability



# Tracking People

Are the models **actually tracking people?**

## Tracking People

Are the models **actually tracking people**?

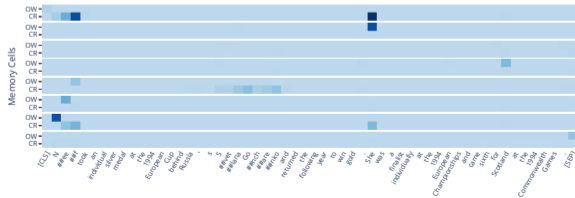
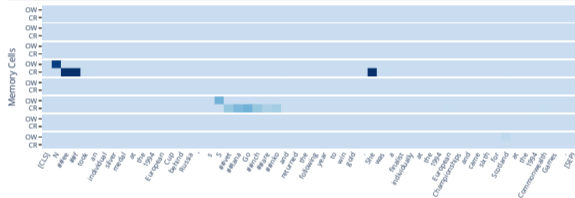
What does it mean for memory model to “track people”?

Detect all references to people.

One-to-one correspondence between people and memory cells.

# Human Evaluation

Neef took an individual silver medal at the 1994 European Cup behind Russia's Svetlana Goncharenko and returned the following year to win gold. She was a finalist individually at the 1994 European Championships and came sixth for Scotland at the 1994 Commonwealth Games.



# Human Evaluation

Neef took an individual silver medal at the 1994 European Cup behind Russia's Svetlana Goncharenko and returned the following year to win gold. She was a finalist individually at the 1994 European Championships and came sixth for Scotland at the 1994 Commonwealth Games.



PeTra  
74%





Ref. Reader  
8%

Neutral  
18%

## Conclusion

- Propose a memory model for tracking people with sparse pronoun resolution annotations.
- Outperforms previous approach on both the end task and interpretability measures.
- Future work: Extend this work to longer documents.

## References

-  Fei Liu, Luke Zettlemoyer, and Jacob Eisenstein.  
The Referential Reader: A recurrent entity network for anaphora resolution.  
*ACL*, 2019.
-  Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge.  
Mind the GAP: A Balanced Corpus of Gendered Ambiguous Pronouns.  
*TACL*, 2018.
-  Mikael Henaff, Jason Weston, Arthur Szlam, Antoine Bordes, and Yann LeCun.  
Tracking the world state with recurrent entity networks.  
*ICLR*, 2017.
-  Alex Graves et al.  
Hybrid computing using a neural network with dynamic external memory.  
*Nature*, 2016.