Table 3: Performances of `JointSem`'s different variations. `JointSem-NoAtt` is the entity-based ranking without attention. `JointSem-SpotAtt` uses the spot attention with entity-based ranking. `JointSem-EntityAtt` uses the entity attention with entity-based ranking. `JointSem-All` is the full model. Relative performances and Win/Tie/Loss are all compared with `JointSem-NoAtt`.

	ClueWeb09-B					ClueWeb12-B13				
Method	**NDCG@20**		**ERR@20**		**W/T/L**	**NDCG@20**		**ERR@20**		**W/T/L**
`JointSem-NoAtt`	0.2919	–	0.1835	–	–/–/–	0.1258	–	0.1012	–	–/–/–
`JointSem-SpotAtt`	0.3005	+2.95%	0.1882	+2.61%	83/55/56	0.1247	−0.88%	0.1010	−0.27%	31/32/37
`JointSem-EntityAtt`	0.2999	+2.74%	0.1872	+2.06%	85/50/59	0.1240	−1.48%	0.1058	+4.52%	36/34/30
`JointSem-All`	**0.3054**	+4.62%	**0.1926**	+5.00%	88/49/57	**0.1314**	+4.44%	**0.1076**	+6.31%	46/27/27

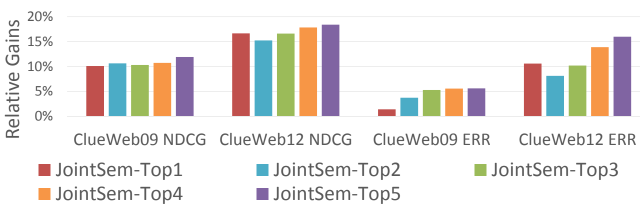

Figure 1: Relative improvements of `JointSem` with different numbers (TopK) of candidate entities per spot. The relative gains marked by the y-axis are compared with `LeToR-Qe-Dw`.

Although in most cases the Top1 entity is the right choice, in general, considering more candidate entities, especially when using all top 5, improves the ranking accuracy. Our manual examination found that improvements are often seen on ambiguous queries whose most popularly linked candidate entities are not the right choice. Without many contexts in short queries, an entity linking system tends to merely choose the most popular candidate entity. For example, the query 'bobcat' refers to bobcat the company, but bobcat the animal is the more popular choice for the entity linking system. `JointSem`'s soft-alignment avoids such over-commitment, and lets the final ranking model select the most useful one(s).

4.3 Effectiveness of Joint Modeling

This experiment studies the effectiveness of joint modeling by comparing `JointSem` and its sub-models. The results are listed in Table 3. `JointSem-NoAtt` uses the Top1 entity per spot as fixed and only includes the ranking part ($f_r(E, d)$). `JointSem-SpotAtt` includes the surface form attention part (f_s), but only the Top1 entities are included with uniform weights; it is similar to the recent attention-based ranking model with word-entity duet [6], but without document entities. `JointSem-EntityAtt` includes the soft-alignment and entity weighting (f_e), but without surface form weighting. `JointSem-All` is the full model. The relative performances and Win/Tie/Loss are all compared with `JointSem-NoAtt`.

The ranking part alone provides better or comparable performance with baselines. Adding in the spotting or the linking part individually helps on ClueWeb09 but has mixed effects on ClueWeb12. Only `JointSem-All` provides stable 5% improvements, confirming the importance of jointly modeling the linking and the utilization of entities for document ranking.

5 CONCLUSION

This work addresses the discrepancy between entity linking and entity-based ranking systems by performing the two tasks jointly. Our method, `JointSem`, spots and links entities in the query, and then uses the linked entities to rank documents. The signals from spotting and linking are incorporated as entity importance features, and the similarities between entities' texts and the document are used as ranking features. `JointSem` uses a joint learning-to-rank model that combines all three components together, and directly optimizes them towards the end-to-end ranking performance.

Experiments on two TREC Web Track datasets demonstrated the effectiveness of `JointSem`, and the influences of the two novelties: the soft-alignment includes multiple entities per spot thus is more robust to ambiguous queries; and the joint modeling stably combines the features from spotting, linking, and ranking together.

This work demonstrates that entity linking, a widely studied natural language processing task, and document ranking, a core information retrieval task, can be, and should be developed together. A future direction is to incorporate entity linking in documents with more advanced entity-based ranking systems.

6 ACKNOWLEDGMENTS

This research was supported by National Science Foundation (NSF) grant IIS-1422676, a fellowship from the Allen Institute for Artificial Intelligence, and DARPA grant FA8750-12-2-0342 under the DEFT program. Any opinions, findings, and conclusions in this paper are the authors' and do not necessarily reflect those of the sponsors.

REFERENCES

[1] David Carmel, Ming-Wei Chang, Evgeniy Gabrilovich, Bo-June (Paul) Hsu, and Kuansan Wang. 2014. ERD'14: Entity recognition and disambiguation challenge. In *Proceedings of SIGIR 2014*. ACM.
[2] Jeffrey Dalton, Laura Dietz, and James Allan. Entity query feature expansion using knowledge base links. In *Proceedings of SIGIR 2014*. ACM.
[3] Paolo Ferragina and Ugo Scaiella. 2010. Fast and accurate annotation of short texts with Wikipedia pages. *arXiv preprint arXiv:1006.3498* (2010).
[4] Xitong Liu and Hui Fang. 2015. Latent entity space: A novel retrieval approach for entity-bearing queries. *Information Retrieval Journal* 18, 6 (2015), 473–503.
[5] Hadas Raviv, Oren Kurland, and David Carmel. 2016. Document retrieval using entity-based language models. In *Proceedings of SIGIR 2016*. ACM, 65–74.
[6] Chenyan Xiong, , Jamie Callan, and Tie-Yan Liu. Word-entity duet representations for document ranking. In *Proceedings of SIGIR 2017*. ACM.
[7] Chenyan Xiong and Jamie Callan. EsdRank: Connecting query and documents through external semi-structured data. In *Proceedings of CIKM 2015*. ACM.
[8] Chenyan Xiong and Jamie Callan. 2015. Query expansion with Freebase. In *Proceedings of ICTIR 2015*. ACM.
[9] Chenyan Xiong, Russell Power, and Jamie Callan. Explicit semantic ranking for academic search via knowledge graph embedding. In *Proceedings of WWW 2017*. ACM.