TODTLER: Two-Order-Deep Transfer Learning
Supplementary Material

Jan Van Haaren
Department of Computer Science
KU Leuven, Belgium
jan.vanhaaren@cs.kuleuven.be

Andrey Kolobov
Microsoft Research
Redmond, WA, USA
akolobov@microsoft.com

Jesse Davis
Department of Computer Science
KU Leuven, Belgium
jesse.davis@cs.kuleuven.be

This supplementary material contains the AUCPRs and CLLs as well as the learning curves for all experiments. It also provides the characteristics of the three domains we consider in our evaluation.

Experimental results

To evaluate each system, we measured the area under the precision-recall curve (AUCPR) and the test set conditional log-likelihood (CLL) for the predicate of interest. Both metrics are worth looking at since CLL measures the quality of the probability estimates whereas AUCPR is insensitive to the large number of true negatives. Unlike in ROC space, random guessing in PR space does not always correspond to a value of 0.50 but is skew-dependent (Boyd et al. 2012). In our experimental setup, random guessing yields an AUCPR of 0.07 for WebKB, 0.08 for Yeast, and 0.37 for Twitter.

Tables 2 and 3 show the area under the precision-recall curve (AUCPR) and conditional log-likelihood (CLL) for predicting a protein's function in the Yeast domain. The former table shows the results for transferring from the WebKB domain, while the latter table shows the results for transferring from the Twitter domain.

Tables 4 and 5 show the area under the precision-recall curve (AUCPR) and conditional log-likelihood (CLL) for predicting a web page's class in the WebKB domain. The former table shows the results for transferring from the Yeast domain, while the latter table shows the results for transferring from the Twitter domain.

Tables 6 and 7 show the area under the precision-recall curve (AUCPR) and conditional log-likelihood (CLL) for predicting an account's type in the Twitter domain. The former table shows the results for transferring from the WebKB domain, while the latter table shows the results for transferring from the Yeast domain.

Learning curves

Figure 1 shows learning curves for predicting protein function in Yeast when transferring knowledge from WebKB. Figure 2 shows learning curves for predicting protein function in Yeast when transferring knowledge from Twitter. Figure 3 shows learning curves for predicting a web page's class in WebKB when transferring knowledge from WebKB. Figure 4 shows learning curves for predicting a web page's class in WebKB when transferring knowledge from Twitter.

Dataset characteristics

Table 1 reports the number of types, predicates, constants, true ground atoms, and possible ground atoms in all three domains.

Table 1: Dataset characteristics for all domains.

	Twitter	**WebKB**	**Yeast**
Types	3	3	7
Predicates	3	3	7
Constants	378	4,396	3,105
True ground atoms	3,142	50,432	15,015
Possible ground atoms	53,748	4,732,804	1,387,014

References

Boyd, K.; Santos Costa, V.; Davis, J.; and Page, D. 2012. Unachievable Region in Precision-Recall Space and Its Effect on Empirical Evaluation. In *Proceedings of the 29th International Conference on Machine Learning*, volume 2012, 349.

Table 2: Area under the precision-recall curve (AUCPR) and conditional log-likelihood (CLL) for predicting a protein's function in the Yeast domain when transferring knowledge from the WebKB domain. The best result is in bold.

	AUCPR				**CLL**			
WebKB → Yeast	TODTLER	DTM-10	DTM-5	LSM	TODTLER	DTM-10	DTM-5	LSM
1 training database	**0.206**	0.162	0.158	0.109	-0.573	**-0.507**	-0.518	-3.481
2 training databases	**0.214**	0.131	0.140	0.068	**-0.468**	-0.772	-1.000	-7.539
3 training databases	**0.222**	0.149	0.102	0.086	**-0.419**	-0.635	-1.303	-5.639

Table 3: Area under the precision-recall curve (AUCPR) and conditional log-likelihood (CLL) for predicting a protein's function in the Yeast domain when transferring knowledge from the Twitter domain. The best result is in bold.

	AUCPR				**CLL**			
Twitter → Yeast	TODTLER	DTM-10	DTM-5	LSM	TODTLER	DTM-10	DTM-5	LSM
1 training database	**0.209**	0.169	0.173	0.109	-0.491	-0.460	**-0.459**	-3.481
2 training databases	**0.229**	0.154	0.148	0.068	**-0.442**	-0.497	-0.772	-7.539
3 training databases	**0.214**	0.137	0.158	0.086	-0.467	-0.475	**-0.408**	-5.639

Table 4: Area under the precision-recall curve (AUCPR) and conditional log-likelihood (CLL) for predicting a web page's class in the WebKB domain when transferring knowledge from the Yeast domain. The best result is in bold.

	AUCPR				**CLL**			
Yeast → WebKB	TODTLER	DTM-10	DTM-5	LSM	TODTLER	DTM-10	DTM-5	LSM
1 training database	**0.369**	0.192	0.192	0.156	**-0.213**	-0.239	-0.239	-0.230
2 training databases	**0.405**	0.156	0.159	0.162	**-0.209**	-0.229	-0.229	-0.229
3 training databases	**0.447**	0.162	0.167	0.161	**-0.190**	-0.228	-0.227	-0.228

Table 5: Area under the precision-recall curve (AUCPR) and conditional log-likelihood (CLL) for predicting a web page's class in the WebKB domain when transferring knowledge from the Twitter domain. The best result is in bold.

	AUCPR				**CLL**			
Twitter → WebKB	TODTLER	DTM-10	DTM-5	LSM	TODTLER	DTM-10	DTM-5	LSM
1 training database	**0.369**	0.156	0.156	0.156	**-0.214**	-0.441	-0.441	-0.230
2 training databases	**0.393**	0.159	0.162	0.162	**-0.197**	-0.229	-0.228	-0.229
3 training databases	**0.434**	0.171	0.171	0.161	**-0.196**	-0.227	-0.227	-0.228

Table 6: Area under the precision-recall curve (AUCPR) and conditional log-likelihood (CLL) for predicting an account's type in the Twitter domain when transferring knowledge from the WebKB domain. The best result is in bold.

	AUCPR				**CLL**			
WebKB → Twitter	TODTLER	DTM-10	DTM-5	LSM	TODTLER	DTM-10	DTM-5	LSM
1 training database	**0.816**	0.633	0.694	0.327	**-0.338**	-1.673	-2.077	-4.892

Table 7: Area under the precision-recall curve (AUCPR) and conditional log-likelihood (CLL) for predicting an account's type in the Twitter domain when transferring knowledge from the Yeast domain. The best result is in bold.

	AUCPR				**CLL**			
Yeast → Twitter	TODTLER	DTM-10	DTM-5	LSM	TODTLER	DTM-10	DTM-5	LSM
1 training database	**0.817**	0.633	0.694	0.327	**-0.321**	-1.673	-2.077	-4.892

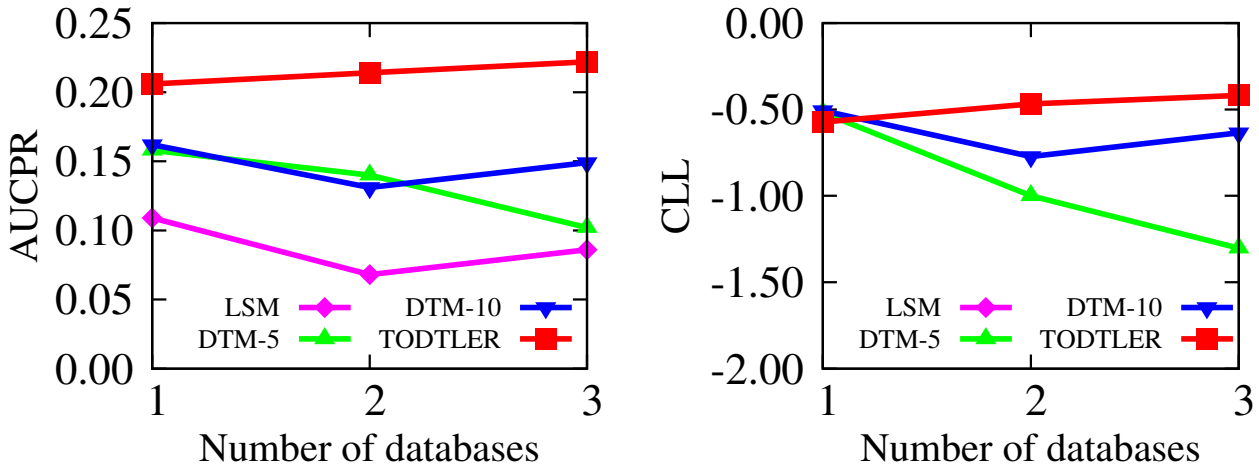

(a) AUCPR (b) CLL

Figure 1: Learning curves for predicting protein function in Yeast when transferring knowledge from WebKB.

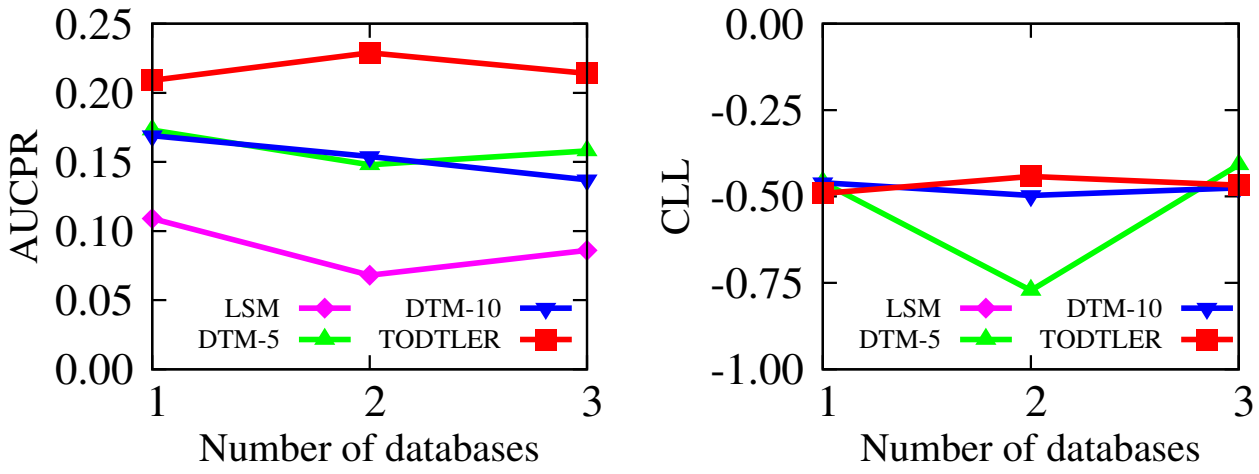

(a) AUCPR (b) CLL

Figure 2: Learning curves for predicting protein function in Yeast when transferring knowledge from Twitter.

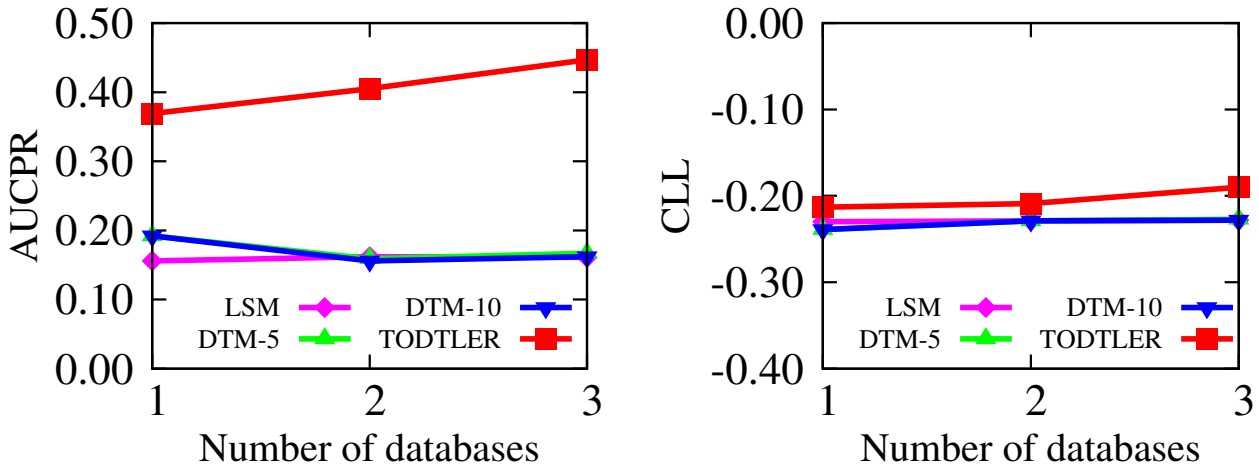

(a) AUCPR (b) CLL

Figure 3: Learning curves for predicting a web page's class in WebKB when transferring knowledge from Yeast.

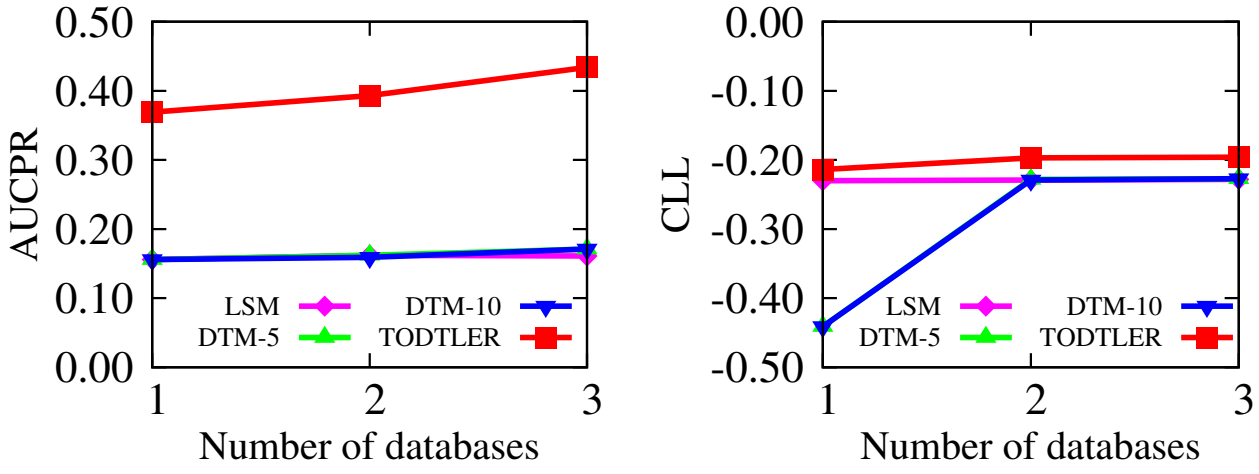

(a) AUCPR (b) CLL

Figure 4: Learning curves for predicting a web page's class in WebKB when transferring knowledge from Twitter.