

MGCNN:

Dataset	Reproduced RMSE	Original RMSE
Synthetic netflix (RGCNN)	0.007244	0.0053
Synthetic netflix (sRGCNN)	0.015296	0.0106
MovieLens (sRGCNN)		0.929
Flixster (sRGCNN, users+items)		1.1788
Flixster (sRGCNN, only users)	0.916766	0.9258
Douban (sRGCNN)	0.794480	0.8012
Yahoo music (sRGCNN)	22.341204	22.4149

GCMC:

Dataset	Reproduced RMSE	Original RMSE
MovieLens 100K	0.9106027/0.9033639	0.910/0.905
MovieLens 1M	0.83188426	0.832
MovieLens 10M		0.777
Flixster (only users)	0.912, 0.926, 0.910	0.941
Flixster (users+items)	0.91047454	0.917
Douban	0.735115	0.734
Yahoo music	20.883778	20.5

Notes:

- GCMC trains much faster than MGCNN (because it is a spatial graph convnet technique and does not require large matrix operations as in a spectral graph convnet)

Datasets info (from GCMC):

Dataset	Users	Items	Features	Ratings	Density	Rating levels
Flixster	3,000	3,000	Users/Items	26,173	0.0029	0.5, 1, ..., 5
Douban	3,000	3,000	Users	136,891	0.0152	1, 2, ..., 5
YahooMusic	3,000	3,000	Items	5,335	0.0006	1, 2, ..., 100
MovieLens 100K (ML-100K)	943	1,682	Users/Items	100,000	0.0630	1, 2, ..., 5
MovieLens 1M (ML-1M)	6,040	3,706	—	1,000,209	0.0447	1, 2, ..., 5
MovieLens 10M (ML-10M)	69,878	10,677	—	10,000,054	0.0134	0.5, 1, ..., 5

Table 1: Number of users, items and ratings for each of the MovieLens datasets used in our experiments. We further indicate rating density and rating levels.

Dataset	New RMSE (1 RGGCN layer)	GCMC RMSE
MovieLens 100K (without side info)	1.08	0.910
MovieLens 100K (with side info)		0.905
MovieLens 1M		0.832
MovieLens 10M		0.777
Flixster (without side info)	1.361	0.933 (one run)
Flixster (users+items)		0.941
Flixster (only users)		0.917
Douban (without side info)	0.934	0.791 (one run)
Douban		0.734
Yahoo music (without side info)	29.56	23.16 (one run)
Yahoo music		20.5

- RMSE was using the hyperparameters in Berg et al.
- Also only one RGGCN conv layer was used
- need to rerun with maybe more epochs for two RGGCN conv layers
- Side features not used because of input dimension mismatch: seems like side features only provided for subset of users/items → can't concat that with full features for all the users/items.

Dataset	New RMSE (2 RGGCN layer)		GCMC RMSE
	Train RMSE	Test RMSE	
MovieLens 100K (without side info)	0.9418 1.005 (with dropout)	1.075 1.08 (with dropout)	0.910
MovieLens 100K (with side info)			0.905
MovieLens 1M			0.832
MovieLens 10M			0.777
Flixster (without side info)	0.8012 0.866 (with dropout)	1.264 1.25 (with dropout)	0.933 (one run)
Flixster (users+items)			0.941
Flixster (only users)			0.917
Douban (without side info)	0.7228 0.769 (with dropout)	0.958 0.935 (with dropout)	0.791 (one run) 0.748 (one run, sum)
Douban			0.734
Yahoo music (without side info)	14.05 21.61 (with dropout)	30.88 25.68 (with dropout)	23.16 (one run)
Yahoo music			20.5

- Severely overfitting. For GCMC the train and test RMSE is very close (in fact val RMSE is usually lower than train RMSE). But with RGGCN that's not the case.
- 16052760 parameters compared to 3048760 using GCMC
- Results below this are with dropout

Possible innovations:

- Maybe need to treat users and items differently? run them with different weights (this is prob alr done, the way the code is)
- Learning rate decay? Use smaller learning rate and use learning rate decay like spatial graph convnets "To summarize, the type of adaptivity provided by algorithms like Adam deal with adjusting the relative learning rates of different parameters, not decreasing the learning rate over time, so it's a different type of adaptive."

Dataset	New RMSE (2 RGGCN layer), hidden (100, 75)		GCMC RMSE
	Train RMSE	Test RMSE	
MovieLens 100K (without side info)	0.972	1.05	0.910
MovieLens 100K (with side info)			0.905
MovieLens 1M			0.832
MovieLens 10M			0.777
Flixster (without side info)	0.898	1.20	0.933 (one run)
Flixster (users+items)			0.941
Flixster (only users)			0.917
Douban (without side info)	0.716	0.966 0.905 0.937 with unidirectional edges	0.791 (one run) 0.748 (one run, sum)
Douban			0.734
Yahoo music (without side info)	21.79	24.52	23.16 (one run)
Yahoo music			20.5

For 1 layer (of 2 RGGCN):

- 3059560 params for RGGCN, douban
- 3059570 params for RGGCN, flixster
- 2167449 params for RGGCN, yahoo
- 1372060 params for RGGCN, movielens 100k

PARAMETER COMPARISON

Dataset	Layers	Hidden	Output	Params

--	--	--	--	--

[illegible]

Dataset	New RMSE (2 RGGCN layer), hidden (100, 75), using SUM		GCMC RMSE
	Train RMSE	Test RMSE	
MovieLens 100K (without side info)	(gets stuck at 17th epoch) 1.24	1.27	0.910
MovieLens 100K (with side info)			0.905
MovieLens 1M			0.832
MovieLens 10M			0.777
Flixster (without side info)	(gets stuck at 6th epoch) 1.46	1.44	0.933 (one run)
Flixster (users+items)			0.917
Flixster (only users)			0.943
Douban (without side info)	(gets stuck at 4th epoch) 1.14	1.14	0.791 (one run) 0.748 (one run, sum)
Douban			0.734
Yahoo music (without side info)			23.16 (one run)
Yahoo music			20.5

- 12821160 params for sumRGGCN douban
- 25023170 params for sumRGGCN flixster
- 173,887,692 params for sumRGGCN yahoo
- 5733660 params for sumRGGCN movielens 100k