

Using Posters to Recommend Anime and Mangas in a Cold-Start Scenario

Jill-Jënn Vie
Mangaki
RIKEN Center for AIP
Tokyo, Japan
jill-jenn.vie@riken.jp

Florian Yger
LMSADE
Université Paris-Dauphine
Paris, France
florian.yger@dauphine.fr

Ryan Lahfa, Basile Clement,
Kévin Cocchi, Thomas Chalumeau
Mangaki
Mitry-Mory, France
{ryan, basile, kevin, thomas}@mangaki.fr

Hisashi Kashima
Kyoto University
RIKEN Center for AIP
Kyoto, Japan
kashima@i.kyoto-u.ac.jp

Abstract—Item cold-start is a classical issue in recommender systems that affects anime and manga recommendations as well. This problem can be framed as follows: how to predict whether a user will like a manga that received few ratings from the community? Content-based techniques can alleviate this issue but require extra information, that is usually expensive to gather. In this paper, we use a deep learning technique, *Illustration2Vec*, to easily extract tag information from the manga and anime posters (e.g., sword, or ponytail). We propose *BALSE* (Blended Alternate Least Squares with Explanation), a new model for collaborative filtering, that benefits from this extra information to recommend mangas. We show, using real data from an online manga recommender system called *Mangaki*, that our model improves substantially the quality of recommendations, especially for less-known manga, and is able to provide an interpretation of the taste of the users.

Keywords—recommender system, cold-start, collaborative filtering, LASSO, tag prediction

I. INTRODUCTION

Recommender systems are useful to help users decide what to enjoy next. In the case of anime and mangas, users, easily overwhelmed by the ever-growing amount of works, end up barely scratching the surface of what Japanese animation has to offer. Collaborative filtering is a popular technique that relies on existing rating data from users on items in order to predict unseen ratings [1]. However, it is still hard to recommend items for which little information is available, e.g., items for which few or no ratings have been provided by the community. This problem has been referred to as the *item cold-start problem*.

In order to alleviate this problem, it is possible to rely on extra information about the items, such as metadata (e.g., for movies: directors, composers, release date). However, such information is not always available: new anime projects may only have a poster or a trailer, and a title. Such a poster is usually the first contact that a user has with an item and plays a large role in the user’s decision to watch it or not. Especially in the manga and anime industry, posters contain a lot of information about the characters, in order to maximize the visual appeal for the consumers. Hence, it is natural to consider posters as a source for additional metadata. In recent years, convolutional neural networks (CNNs) have established themselves as the *de-facto* method for extracting

semantic information from image content in a wide variety of tasks. We propose using a CNN for extracting meaningful tags directly from the item’s poster. Such extracted tags can help draw links between items, which can be useful when few ratings are available. In this paper, we present *BALSE*¹ (Blended Alternate Least Squares with Explanation), a new method leveraging tag information extracted from the posters for tackling the item cold-start problem and improving the recommendation performance for little-known items. We show using real data that our method provides better rating predictions than existing techniques, and gives interpretable insight about the user’s taste.

To the best of our knowledge, this is the first research work that uses tag prediction on item posters in order to improve the accuracy of a recommender system and explain to users why they are recommended little-known items.

This paper is organized as follows. We first present existing work related to this research. Then, we expose the context of collaborative filtering and item cold-start, together with a few common assumptions. We then describe our model, *BALSE*, and present some experimental results on a real dataset. We finish by discussing the results and future work.

II. RELATED WORK

Using side information in order to improve recommendations has been the core of existing research [2] and several approaches have been developed to take into account extra data about users or items, whether coming from text [3], [4], social networks [5], images or other types of data [6]–[8]. More recently, deep learning techniques have been used for this purpose. YouTube is extracting features from the videos browsed within a user history in order to improve their recommendations [9]. Researchers have also analyzed music content as extra information [10]. They managed to recover explainable latent features, corresponding to certain types of music, without any human annotation. Such side information is particularly useful in order to mitigate the cold-start problem [2], [11]–[13]. In the exact context of movies, [14] extract latent features from the posters using

¹<http://knowyourmeme.com/memes/events/balse>

CNNs and improve the recommendations using those latent features. However, those extracted features do not have semantic meaning, therefore they cannot be used to explain to the user why extra works are recommended to them.

Several approaches have tried to bridge the gap between content-based approaches and collaborative filtering [15]. The main idea behind those so-called hybrid methods is to combine different recommendation models in order to overcome their limitations and build more robust models with higher predictive power. Existing techniques can take on several names: blending or stacking [16], [17], or the general ensemble methods for machine learning estimators. These techniques use the output of different models as features for a higher-level model. This higher-level model is usually a linear model [18]. Such blended methods have played an important role in achieving top performance in challenges such as the Netflix Prize [18], [19]. The approach described in this paper builds upon these ideas, as we are presenting a blended model, but the combination we present is nonlinear. We complement a classical collaborative filtering recommender system with a fallback model that will compensate the prediction error on hard-to-predict data points, i.e. items with few ratings.

III. CONTEXT

We assume the typical setting for collaborative filtering: we have access to a $n \times m$ rating matrix R containing the ratings of n users on m items that can be either manga or anime: r_{ij} represents the rating that user i gave to item j . In practice, since users only rate the few items that they have interacted with, the rating matrix R tend to be very sparse: in the dataset considered in this paper, less than 1% of the entries are known; other popular datasets in the field [1], [20] report similar levels of sparsity. Therefore, it is challenging to infer the missing entries of R .

Another assumption is that the whole rating matrix can be explained by few latent profiles, i.e. each user rating vector can be decomposed as a combination of few latent vectors. Therefore, matrix completion is usually performed using matrix factorization: we try to devise a factorization $R \approx UV^T$ where a $n \times r$ matrix U represents the user feature vectors and a $m \times r$ matrix V represents the item feature vectors. Once this model is trained, that is, when the available entries of R match their counterparts in UV^T , computing a missing entry (i, j) of the rating matrix R is simply performed by looking at the corresponding (i, j) entry of UV^T , or, equivalently, computing the dot product $U_i^T V_j$ where U_i is the i -th row of U and V_j is the j -th row of V .

Finally, we also assume that we have access to the posters of some of the items. This is all the content we have.




Figure 1. The BALSE architecture.

IV. OUR MODEL: BALSE

We now describe BALSE (Blended Alternate Least Squares with Explanation), our model for recommending anime and mangas. The main idea is to rely on the rating matrix when possible, and on the posters when rating information barely exists. We expect a nonlinear blending of two models considering these sources of information to achieve higher performance than any of the models. BALSE is composed of several blocks:

- an *Illustration2Vec* block, which is a convolutional neural network that takes a poster as input and outputs tag predictions;
- an *ALS*² block, that performs a matrix factorization of the rating matrix for collaborative filtering using alternate least squares with λ -weighted regularization;
- a *LASSO*³ block, that performs a regularized linear regression of each row of the rating matrix, using tag predictions, in order to infer explainable user preferences;
- a *Steins gate*, that performs a blending of the outputs of *ALS* and *LASSO* models, in order to overcome their limitations and provide a final rating value.

The main architecture of our model is presented in Figure 1. Both posters and ratings are used for the predictions.

A. *Illustration2Vec*

This block extracts tag information from the posters, such as “1girl” or “weapon”. Such tags are associated with confidence weights that represent how likely a certain tag appears in a certain poster. Formally, from the poster database, we want to extract a $m \times t$ matrix T where m is the number of items and t is the number of possible tags, such that $t_{jk} \in [0, 1]$ represents how likely a tag k appears in the poster of item j . T is computed using *Illustration2Vec* [21], a VGG-16 neural network [22] that predicts a variety of tags based on illustrations, pre-trained on ImageNet and trained on manga illustrations labeled with tags from the community website Danbooru. We use the implementation provided by

²ALS stands for Alternate Least Squares.

³LASSO stands for Least Absolute Shrinkage and Selection Operator.

the authors, which is freely available. The output of the network is for each poster j , a vector $T_j = (t_{j1}, \dots, t_{jt})$ where for tag $k = 1, \dots, t$, component $t_{jk} \in [0, 1]$ represents how likely tag k describes the poster of item j . In other words, the output of Illustration2Vec is a row of matrix T . We will call such a vector a tag prediction. See Fig. 2 for an example of an output of the Illustration2Vec model.

B. LASSO

The LASSO block approximates the rating matrix R with a regularized linear regression model called LASSO [23], using the tag predictions as features for the items. We train a LASSO model for every user in the train set.

$$R \approx PT^T$$

where:

- P contains the parameters to learn, a $n \times t$ matrix of user preferences, of which the i -th row is denoted as P_i (likewise, R_i denotes the i -th row of R);
- T is the given $m \times t$ matrix of tag predictions for each item.

LASSO comes with an α parameter which induces a L1 regularization term to prevent overfitting, and to provide explanation of user preferences as we will show later. Therefore, for every user i of the train set, we estimate the parameters P_i that minimize:

$$\frac{1}{2\mathcal{N}_i} \|R_i - P_i T^T\|_2^2 + \alpha \|P_i\|_1$$

where \mathcal{N}_i is the number of items rated by user i .

The output of the LASSO block is a rating prediction for each pair (i, j) :

$$\hat{r}_{ij}^{LASSO} = \tau(P_i^T T_j).$$

where $\tau : x \mapsto \max(\min(x, 2), -2)$ is a function that shrinks its input to values between -2 and 2. Such a function prevents the regressor from providing invalid predictions that are outside the range of rating values.

C. ALS

The ALS block performs matrix factorization of the $n \times m$ sparse rating matrix R , in order to provide an estimate \hat{r}_{ij}^{ALS} for the missing entries (i, j) . Thus, we learn the parameters of the following factorization:

$$R \approx UV^T$$

where:

- U is the $n \times r$ matrix of user latent vectors;
- V is the $m \times r$ matrix of item latent vectors.

In order to avoid overfitting, we regularize the parameters to estimate. Therefore, as we want to minimize the squared error, the loss function to minimize has the following form:

$$\sum_{i,j|r_{ij} \neq 0} (r_{ij} - U_i^T V_j)^2 + \lambda (\|U_i\|_2^2 + \|V_j\|_2^2)$$

where U_i for every $i = 1, \dots, n$ are the rows of U and V_j for every $j = 1, \dots, m$ are the rows of V , and λ is a regularization parameter. This estimation is made by using alternate least squares with weighted λ -regularization (ALS-WR) [24].

Once the parameters have been learned, the prediction for rating of user i on item j is:

$$\hat{r}_{ij}^{ALS} = U_i^T V_j.$$

D. Steins Gate

At this step, we have predictions from two different blocks: ALS trained on the ratings and LASSO trained on the tag predictions of the posters. We want to improve the predictive power of the overall model, thus we learn a rule that would automatically choose the best model according to the number of ratings of the item considered. Formally, we want to learn parameters β and γ such that:

$$\hat{r}_{ij}^{BALSE} = \sigma(\beta(\mathcal{R}_j - \gamma)) \hat{r}_{ij}^{ALS} + (1 - \sigma(\beta(\mathcal{R}_j - \gamma))) \hat{r}_{ij}^{LASSO}$$

where:

- \mathcal{R}_j is the number of ratings of the item j ;
- \hat{r}_{ij}^{ALS} is the rating prediction of ALS model for user i on item j ;
- \hat{r}_{ij}^{LASSO} is the rating prediction of LASSO model for user i on item j ;
- $\sigma : x \mapsto 1/(1 + e^{-x})$ is the sigmoid function.

The intuition behind this formula is the following: we want to find a threshold γ such that when the number of ratings of item j verifies $\mathcal{R}_j \gg \gamma$, BALSE mimics ALS, e.g., $\hat{r}_{ij}^{BALSE} \approx \hat{r}_{ij}^{ALS}$, while when $\mathcal{R}_j \ll \gamma$, i.e. in a cold-start setting, BALSE mimics LASSO, e.g., $\hat{r}_{ij}^{BALSE} \approx \hat{r}_{ij}^{LASSO}$. β is just a scaling parameter that indicates how sharp the passage from LASSO to ALS will be. Formally, we want to estimate the parameters β and γ that minimize:

$$\sum_{i,j|r_{ij} \neq 0} (\hat{r}_{ij}^{BALSE} - r_{ij})^2.$$

This formula is differentiable with respect to γ , thus it makes its optimization easier. It can be seen as a soft switch between the two possible predictions (ALS and LASSO), according to the number of ratings of the item. The parameters β and γ are learned using gradient descent.

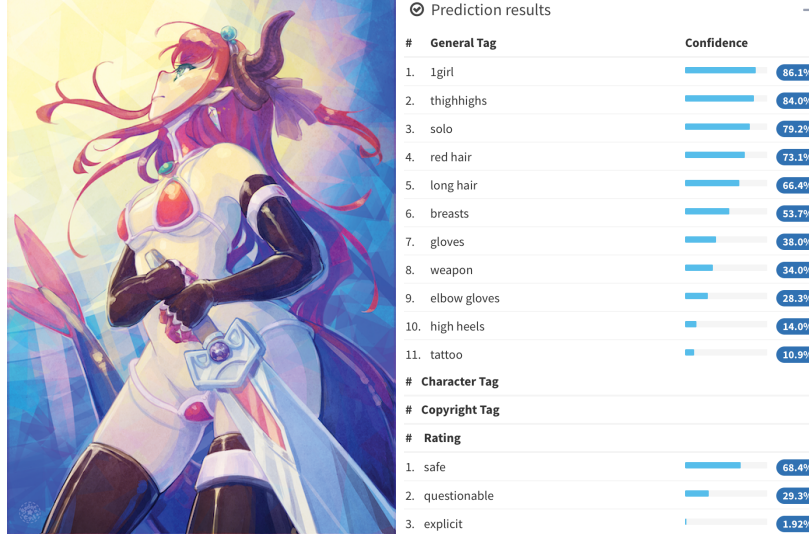


Figure 2. Example of tag prediction on a manga illustration. ©Solène Pichereau, <http://sedeto.fr>.

V. EXPERIMENTS AND RESULTS

A. Mangaki dataset

Mangaki⁴ [25] is a website where people can rate items that represent either manga or anime, and receive recommendations based on their ratings. Mangaki can be seen as an open source version of Movielens [26] for manga and anime. The Mangaki dataset is a 2079×9979 anonymized matrix of 334390 ratings from 2079 users on 9979 items. 80% of the items have a poster.

Users can either rate an item with {favorite, like, neutral, dislike} if they watched it, or {willsee, wontsee} if they did not watch it, i.e. testify whether they want to watch it or not, based on the content presented: poster, possibly synopsis, or some statistics.

B. Models

The models considered in this benchmark are:

- ALS: alternate least squares with weighted λ -regularization from [24], that ignores posters;
- LASSO: regularized linear regression using ratings and the tag predictions from Illustration2Vec, that is content-based;
- BALSE: the proposed method.

In practice, we use $\lambda = 0.1$ and rank $r = 20$ for every ALS model trained and $\alpha = 0.01$ for every LASSO model trained. Ratings are mapped into custom values: (favorite, like, neutral, dislike) = (4, 2, 0.1, -2) and (willsee, wontsee) = (0.5, -0.5). The Steins gate is optimized using gradient descent with exponential decay implemented in TensorFlow. The learning rate starts at 0.9 and decays every 20 steps with

a base of 0.997. All the code is available on our GitHub repository⁵.

C. 5-fold cross validation

We perform a 5-fold cross validation over the triplets (i, j, r_{ij}) of the database, keeping 30% of the train set as a validation set. Therefore, our data is split into a train set (56%), a validation set (24%) and a test set (20%).

The vanilla models ALS and LASSO are trained on both the train set and the validation set. For BALSE, the ALS and LASSO blocks are first trained using the train set only, and the Steins gate parameters β and γ are trained using the validation set⁶, in order to prevent overfitting. For the final predictions of BALSE, blending is performed using the learned β and γ parameters, and the vanilla ALS and LASSO models. Finally, the root mean squared error (RMSE) is computed over the test set.

We distinguish the performance of all three models on three sets: the whole test set, a set of little-known items that received less than 3 ratings in the train and validation set (that represents 1000 ratings, therefore 3% of the test set), and cold-start items, i.e. items that were never seen in the train and validation sets.

D. Results

BALSE achieves a comparable performance than ALS overall, but substantially improves the recommendations on little-known items, see Table I.

The learned parameter γ of the Steins gate was less than 1, see Figure 3, which means that items having at least 1 rating can start to rely on ALS (their ratings) more than LASSO

⁵<https://github.com/mangaki/balse>

⁶Please also note that in Steins gate, the number of ratings R_j of item j is computed over the train set.

⁴<https://mangaki.fr>

Table I
RESULTS OF RMSE ON VARIOUS SUBSETS OF THE TEST SET.

RMSE	Whole test set	Little-known items (3% least rated)	Cold-start items (0 rating)
ALS	1.15681 \pm 0.004	1.29269 \pm 0.029	1.50047 \pm 0.03500
LASSO	1.44444 \pm 0.002	1.31331 \pm 0.036	1.37909 \pm 0.05600
BALSE	1.14954 \pm 0.004	1.22714 \pm 0.036	1.34533 \pm 0.04500



Figure 3. This is Steins gate’s choice: $\gamma = 0.79040$. For items having at least one rating, it is better to rely more on the ratings predicted by ALS than by LASSO.

(their poster) for the predictions. However, BALSE provides better predictions than ALS for cold-start items, because ALS was not trained on them in the train set, therefore it outputs constant predictions.

ALS converges after 10 iterations. Steins gate takes 15k iterations to converge. LASSO is the bottleneck of the proposed approach because one LASSO model should be trained per user that appears in the train set.

E. Explanation of user taste

Using the tags, it is possible to provide an explanation of the taste of any user i using the preference matrix P learned by LASSO, because the columns of P_i are labeled with tags.

LASSO has been appreciated for its explainability [23]: the row preferences of each user are sparse, allowing to capture the tags that explain best the ratings of every user.

As an example, for a certain user among the authors, LASSO or BALSE report that his six most preferred tags are: *kneehighs*, *cat*, *serafuku*⁷, *twin braids*, *japanese clothes* and *angry* whereas his six most disliked tags are: *pleated skirt*, *standing*, *silver hair*, *window*, *torn clothes* and *skirt*. Using this information, LASSO or BALSE can explain a recommendation: “We recommend to you the anime *Chivalry of a Failed Knight*, because there is a girl with *twin braids*, *serafuku* and *japanese clothes*” or a warning: “You might not like the anime *The Asterisk War: The Academy City on the Water* because there is a girl with a *pleated skirt*, even though there are *kneehighs* and *serafuku*.”

VI. CONCLUSION AND FUTURE WORK

We proposed BALSE, a model for recommending anime and manga that makes use of information that is automatically extracted from posters. We showed that our model

performs better than the baseline models, especially in the item cold-start scenario.

This paper is a proof a concept and the obtained results are very encouraging. Indeed, the blending Steins gate is such that any improvement made on any block would improve the overall performance of the approach. As future work, we plan to replace blocks in our architecture with more complex models: Illustration2Vec could be replaced with residual networks [27], ALS could be replaced with factorization machines [28] or co-factorization [4], LASSO could be replaced with Localized Lasso [29], a variant that works well for few samples, many features. We also to integrate more side information, for instance the drawing style of the image, or tags coming from open databases such as AniDB⁸ or AniList⁹, in order to improve the explanation of the users’ preferences.

For the sake of simplicity, we mapped the categorical ratings like, dislike, etc. to ad-hoc values, but we could instead use ordinal regression methods [30]. However, they require more computation to be trained properly. Ensemble methods that blend more than two models could be considered [18], or that rely not also on the number of ratings provided for a certain item, but on the number of ratings provided by a certain user, or the number of works that contain a certain tag.

Here, we mitigated the problem of item cold-start recommendation through the use of extra information on the item side. Obviously, similar results could be obtained for the user cold-start problem, provided enough data is available to describe the users.

Using BALSE, recommender systems can automatically replenish their database, where new items go through the tag prediction track and the explainable model in order to justify the recommendations for their first users, and automatically go to the main track when sufficient ratings have been collected.

ACKNOWLEDGMENTS

This work was carried out while Florian Yger was a visiting researcher at RIKEN Center for AIP, and Kévin Cocchi and Thomas Chalumeau were interns at Mangaki. We would like to thank Nicolas Hug and Étienne Simon for their helpful comments and Solène Pichereau for kindly providing the example illustration of Figure 2.

⁸<http://anidb.net>

⁹<https://anilist.co>

⁷Serafuku means “Japanese school uniform”.

REFERENCES

- [1] C. C. Aggarwal, *Recommender systems*. Springer, 2016.
- [2] M. Kula, “Metadata embeddings for user and item cold-start recommendations,” *arXiv preprint arXiv:1507.08439*, 2015.
- [3] G. Alexandridis, G. Siolas, and A. Stafylopatis, “Parvecmf: A paragraph vector-based matrix factorization recommender system,” *arXiv preprint arXiv:1706.07513*, 2017.
- [4] Y. Fang and L. Si, “Matrix co-factorization for recommendation with rich side information and implicit feedback,” in *Proceedings of the 2nd International Workshop on Information Heterogeneity and Fusion in Recommender Systems*. ACM, 2011, pp. 65–69.
- [5] J. Delparte, A. Karatzoglou, T. Matuszczyk, and S. Canu, “Socially enabled preference learning from implicit feedback data,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer Berlin Heidelberg, 2013, pp. 145–160.
- [6] T. Nedelec, E. Smirnova, and F. Vasile, “Specializing joint representations for the task of product recommendation,” *arXiv preprint arXiv:1706.07625*, 2017.
- [7] Y.-D. Kim and S. Choi, “Scalable variational bayesian matrix factorization with side information,” in *Artificial Intelligence and Statistics*, 2014, pp. 493–502.
- [8] M. Xu, R. Jin, and Z.-H. Zhou, “Speedup matrix completion with side information: Application to multi-label learning,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2301–2309.
- [9] P. Covington, J. Adams, and E. Sargin, “Deep neural networks for youtube recommendations,” in *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 2016, pp. 191–198.
- [10] A. Van den Oord, S. Dieleman, and B. Schrauwen, “Deep content-based music recommendation,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2643–2651.
- [11] J. Wei, J. He, K. Chen, Y. Zhou, and Z. Tang, “Collaborative filtering and deep learning based recommendation system for cold start items,” *Expert Systems with Applications*, vol. 69, pp. 29–39, 2017.
- [12] S. Biswas, L. V. Lakshmanan, and S. B. Ray, “Combating the cold start user problem in model based collaborative filtering,” *arXiv preprint arXiv:1703.00397*, 2017.
- [13] J. Bobadilla, F. Ortega, A. Hernando, and J. Bernal, “A collaborative filtering approach to mitigate the new user cold start problem,” *Knowledge-Based Systems*, vol. 26, pp. 225–238, 2012.
- [14] L. Zhao, Z. Lu, S. J. Plan, and Q. Yang, “Matrix factorization+ for movie recommendation,” in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. AAAI Press, 2016, pp. 3945–3951.
- [15] R. Burke, “Hybrid recommender systems: Survey and experiments,” *User modeling and user-adapted interaction*, vol. 12, no. 4, pp. 331–370, 2002.
- [16] F. Roda, A. Costa, and L. Liberti, “Optimal recommender systems blending,” in *Proceedings of the International Conference on Web Intelligence, Mining and Semantics*. ACM, 2011, p. 60.
- [17] M. Jahrer, A. Töschner, and R. Legenstein, “Combining predictions for accurate recommender systems,” in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2010, pp. 693–702.
- [18] J. Sill, G. Takács, L. Mackey, and D. Lin, “Feature-weighted linear stacking,” *arXiv preprint arXiv:0911.0460*, 2009.
- [19] Y. Koren, “The bellkor solution to the netflix grand prize,” *Netflix prize documentation*, vol. 81, pp. 1–10, 2009.
- [20] H. Wang, Q. Wu, and H. Wang, “Learning hidden features for contextual bandits,” in *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*. ACM, 2016, pp. 1633–1642.
- [21] M. Saito and Y. Matsui, “Illustration2vec: A semantic vector representation of illustrations,” in *SIGGRAPH Asia Technical Briefs*, 2015.
- [22] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [23] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.
- [24] Y. Zhou, D. Wilkinson, R. Schreiber, and R. Pan, “Large-scale parallel collaborative filtering for the netflix prize,” *Lecture Notes in Computer Science*, vol. 5034, pp. 337–348, 2008.
- [25] J.-J. Vie, C. Lailly, and S. Pichereau, “Mangaki: an anime/manga recommender system with fast preference elicitation,” Tech. Rep., 2015.
- [26] F. M. Harper and J. A. Konstan, “The movielens datasets: History and context,” *ACM Trans. Interact. Intell. Syst.*, vol. 5, no. 4, pp. 19:1–19:19, Dec. 2015. [Online]. Available: <http://doi.acm.org/10.1145/2827872>
- [27] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [28] S. Rendle, “Factorization machines,” in *Data Mining (ICDM), 2010 IEEE 10th International Conference on*. IEEE, 2010, pp. 995–1000.
- [29] M. Yamada, K. Takeuchi, T. Iwata, J. Shawe-Taylor, and S. Kaski, “Localized lasso for high-dimensional regression,” *arXiv preprint arXiv:1603.06743*, 2016.
- [30] F. Pedregosa, F. Bach, and A. Gramfort, “On the consistency of ordinal regression methods,” *Journal of Machine Learning Research*, vol. 18, no. 55, pp. 1–35, 2017.