

Photos Are All You Need for Reciprocal Recommendation in Online Dating

James Neve
Univeristy of Bristol
Bristol, United Kingdom
james.neve@bristol.ac.uk

Ryan McConville
Univeristy of Bristol
Bristol, United Kingdom
ryan.mcconville@bristol.ac.uk

ABSTRACT

Recommender Systems are algorithms that predict a user’s preference for an item. Reciprocal Recommenders are a subset of recommender systems, where the items in question are people, and the objective is therefore to predict a bidirectional preference relation. They are used in settings such as online dating services and social networks. In particular, images provided by users are a crucial part of user preference, and one that is not exploited much in the literature. We present a novel method of interpreting user image preference history and using this to make recommendations. We train a recurrent neural network to learn a user’s preferences and make predictions of reciprocal preference relations that can be used to make recommendations that satisfy both users. We show that our proposed system achieves an F1 score of 0.87 when using only photographs to produce reciprocal recommendations on a large real world online dating dataset. Our system significantly outperforms on the state of the art in both content-based and collaborative filtering systems.

CCS CONCEPTS

• **Information systems** → **Social recommendation**; *Decision support systems*; *Personalization*; • **Computer systems organization** → *Neural networks*.

KEYWORDS

Recommender Systems, Reciprocal Recommender Systems, Recurrent Neural Networks, Social Recommendation

1 INTRODUCTION

Recommender Systems (RS) are personalisation tools that are used by online services to recommend items to users [20]. RSs usually make recommendations by computing a preference score between 0 and 1 that represents the extent to which a particular user would like a particular item. This is done by using explicit preference information (for instance, a user’s profile where they have specified their preference) or implicit preference information, such as a user’s purchase history. RSs have become increasingly advanced over the last decade, and most popular online services such as *Amazon* and *Netflix* use them to enhance their users’ experience [3].

Reciprocal Recommender Systems (RRSs) are a subtype of RSs that recommend people to other people. They are commonly used in online dating and social services. While RSs make recommendations based on a unidirectional preference relation involving an inanimate item, RRSs are inherently more complex in that they must make recommendations based on both sides of a bidirectional preference relation. Applying a conventional recommender system

to a reciprocal environment results in recommendations that are only satisfying to one of the two users involved in the interaction. RRS design also involves a number of considerations that are not involved in unidirectional recommendation. For instance, a popular product being repeatedly recommended is not usually a problem, but a popular user appearing in everyone’s recommendations often represents a negative experience for that user [10]. RRSs are often adapted from RSs, where two unidirectional preferences are computed and then combined into a single preference score that represents the preference of two users for each other.

RSs (and therefore RRSs) are often categorised as content-based or collaborative filtering systems. Content-based systems make recommendations based on a user’s preference for specific aspects of an item. These preferences are sometimes explicit, but are more usually inferred implicitly from a user’s preference for previous items [1]. Collaborative filtering algorithms use correlations between multiple users to make recommendations, often by extracting latent factors from a preference matrix of users and items, and inferring preference for those latent factors through historical preferences. Historically, collaborative filtering algorithms have been more effective than content-based algorithms [1]. However, in content-rich environments, the reverse can be true. Content-based filtering algorithms also tend to be more effective at solving the *Cold-Start Problem* [12, 16], where it is difficult for the system to make effective recommendations for new users because of the lack of preference history.

Online dating services and social networks often provide a content-rich environment, with users making decisions about whom to express preference to based on a great many factors, including image data, free text profiles and categorical data such as age and job. In particular, image data is extremely important to modern social services, with many such as *Instagram* using images as the basis for interactions. This has been demonstrated by informal research from industry¹. In this paper, we present a novel recommender system, *Temporal Image-Based Reciprocal Recommender* (TIRR), that uses a Recurrent Neural Network (RNN) to interpret a user’s history of preferences for images, and make predictions about their future preferences in order to make recommendations. This is a significant improvement on the only other previous image-based RRS, *ImRec*[18], in the sense that it outperforms both *ImRec* (previously the state of the art in content-based reciprocal recommendation) and also the current state of the art collaborative filtering solutions.

In addition to the advantages in terms of its improvement in the ROC curve on cross-validation, TIRR is also an advance of the field in the sense that it provides a unified system that predicts matches directly, as opposed to two separate predictions of unidirectional

¹<https://www.guern.net/docs/psychology/okcupid/weexperimentonhumanbeings.html>

preferences followed by an aggregation. There is some doubt as how to combine two unidirectional scores into a single bidirectional score in a way that is fully representative of two users' bidirectional preference for each other; TIRR solves this by predicting the bidirectional relation end to end.

The system was tested using a popular online dating service. We used 200000 users and approximately 800000 expressions of preference combined split across train and test sets.

The original contributions of this paper are therefore threefold:

- (1) We present a content-based RRS, TIRR, that makes recommendations based on historical sequences of images utilizing Siamese networks and LSTMs.
- (2) Previous RRSs predict two unidirectional preferences and then aggregate them; the end-to-end algorithm detailed in this paper predicts the probability of a match directly.
- (3) Based on tests using real-world data, TIRR outperforms not only other content-based RRSs but also the state-of-the-art collaborative filtering RRS.

2 RELATED WORKS

This section contains a review of other academic works that formed the background for this study. This includes works on RRSs, content-based recommendation and on recurrent neural networks for understanding image-based histories.

2.1 Reciprocal Recommender Systems

RRSs are recommender systems used for person-to-person matching, in settings such as online dating, social networks [6] and job recommendation [22]. They are complex in the sense that they need to consider the preferences of both sides.

The earliest RRS in the literature is RECON [21]. This is a content-based recommender system designed by Pizzato et al. based on recommendation using categorical data such as age and hobbies. For two users x and y , the algorithm calculates the preferences of the two users P_x and P_y as vectors based on their historical preferences for individual attributes. Using these historical preferences, RECON estimates unidirectional preferences $Q_{x,y}$ and $Q_{y,x}$ and combines them using the harmonic mean into a single bidirectional preference relation that represents the projected preference of the two users for each other.

RCF [27] extended reciprocal recommendation with an implementation of a collaborative filtering system. RCF uses nearest-neighbour-based recommendation, where for candidate users x and y it calculates the similarity between x and the other users that have liked y and vice-versa. RCF demonstrated improvements in a number of areas over RECON, and was at the time considered the best in class reciprocal recommender system.

Subsequently, a number of systems have made improvements to both collaborative and content-based systems, in addition to designing hybrid systems that exploit the best of both subtypes. For example, Kleinermann et al. improved on RCF by reducing the bias of user popularity on recommendations [10]. ImRec demonstrated that image-based recommendation was more effective than recommendation based on categorical profile data [18].

2.2 Content-Based Recommender Systems

Content-based recommender systems make recommendations based on users' preferences for specific content on a service. This might be structured content such as the category of an item, or it might be unstructured content such as images and free text description.

Recommendations based on unstructured data appear most often in news recommendation [13, 24]. This is a rich area for research because news articles are often written with set structures that make them easier to process, and also because of concerns about serendipity and recommendations reinforcing existing echo chambers. Papers such as [2] demonstrate the capacity for deep learning systems based on freetext information to make effective recommendations.

Examples of content-based recommender systems basing their results on images is much less common. Lei et al. used user preferences to train a model based on ImageNet [5] that predicted user preference for one of two images [14]. This trains a network to map both users and images into the same space by generating embeddings for both, with images that the user preferred being close to the user in the space, and images the user did not like being further away. User preference for subsequent images can then be predicted by relative distance from the user.

Another example is *DeepStyle* [17], which uses a Siamese Network to predict user preference for clothes based on images. DeepStyle uses pairs of positive and negative samples with user preference as the output to differentiate between the two images. This can then be used to make predictions about whether a user might like a new image by comparing it to an existing liked image.

2.3 Recurrent Neural Networks

This paper uses Recurrent Neural Networks (RNNs) to interpret time series data for recommendation. RNNs contain loops, which feed the output of a network back into current neurons. This means that they implement a concept of *memory*: they store computed results, and these results have an impact on subsequent predictions. Each step therefore incorporates information from the previous steps into the prediction.

Standard RNNs are particularly good at processing short sequences, but their memory is *short-term memory*: when training them using longer sequences, the early items in the sequence have very little impact on the final prediction. This is known as the *vanishing gradient problem* [8]. This also exists in deep neural networks, where early layers learn very slowly when trained with backpropagation.

Various architectures have been proposed to overcome this limitation. One that has been particularly successful in allowing RNNs to hold and use information for longer is the *Long Short-Term Memory Network* (LSTM) [9]. A LSTM uses a *forget gate* comprised of a Sigmoid function that determines whether information is kept or not: a value close to 0 results in the information being forgotten by the network, whereas closer to 1 results in the information being stored. This allows for much longer sequences to be processed, which is particularly useful in the field of recommendation, where long sequences of user behaviour are common. RNNs have been used successfully in recommender systems to incorporate time

series data into recommendations [23, 26], but not in reciprocal recommender systems.

3 METHODOLOGY

In this section, we describe a model that produces predictions about user preferences based on the RNN interpretation of user history. The RNN-based model uses a pre-trained siamese network at its core, so we describe that independently first, and then its use in the context of the RNN.

3.1 Problem Formulation

The online dating service we used currently only supports heterosexual relationships. We can therefore assume that for a set of users of one gender $X = \{x_1, x_2, \dots, x_{|X|}\}$ there is a set of candidate users for recommendation $Y = \{y_1, y_2, \dots, y_{|Y|}\}$. A user may have an ordered history of preference expressions for users of length n , for example, $Sx = \{Sx_{t_0}^{y_i}, Sx_{t_1}^{y_j}, Sx_{t_m}^{y_k}, \dots, Sx_{t_n}^{y_l}\}$ where $Sx_{t_m}^{y_i} \in Y$ represents the expression of positive or negative preference of user x for the user y_i at time t_m .

In our reciprocal system, our objective is to estimate $R^{x,y}$, the reciprocal preference score that represents the projected degree of preference of two users for each other. We consider that $R^{x,y}$ is a function of the historical preferences of x and y as well as the two users themselves, and train a model to predict it using all of this information:

$$R^{x,y} = f(S^x, S^y, x, y; \theta) \quad (1)$$

Where θ represents the parameters of the model. Note that contrary to most previous approaches to RRSs, our approach trains a single model to predict reciprocal preference using all of the information, as opposed to combining the results of two models predicting unidirectional preference. Also note that the reciprocal preference is symmetrical i.e. $R^{x,y} = R^{y,x}$.

3.2 Service & Data

The data for our model was provided by a popular online dating service with several million registered users. On this service, the user experience is streamlined so that everyone goes through the same process of interaction.

A user x finds other users by searching, or by viewing recommendations on a list page. From the list page, they can view profiles with images, text and categorical data. If x finds a user y that they want to interact with, they can send a *Like*. In our algorithm, this *Like* is used as a unidirectional indicator of preference.

User y can choose whether or not to reciprocate this *Like*. If they do reciprocate, this is considered a *Match*; if not it is considered a *Dislike*. These represent bidirectional indicators of preference or negative preference respectively. Users who have *Matched* can subsequently message each other, and potentially agree to meet in person.

As we wanted to focus on an algorithm that measured personal attractiveness of users for each other, we made the decision to exclude images from the dataset that did not include user faces using automatic face detection. It is common for deep learning based on faces to also include cropping an affine transformation

of features, but preliminary experiments showed that this did not improve our results.

We also made a number of exclusions in order to increase the reliability of the dataset. We excluded users who had been removed from the service for any reason (often these users are not using the service correctly, which implies that they are not expressing preferences based on their own intuition). Although for privacy reasons we are unable to release the dataset used in our experiments, we do hope that the algorithm will be reproduced on other services.

3.3 Siamese Network Unidirectional User Preference Learning




Figure 1: The architecture of the Siamese network used to learn unidirectional user preferences, which forms a component of TIRR.

In this section, we briefly describe the Siamese network [11] we used to learn unidirectional user preferences, a building block of our proposed method, originally included as a component of *ImRec* [18]. We will utilize this in a novel way to *ImRec* to demonstrate superior performance.

3.3.1 Siamese Network Concept. Siamese networks are commonly used in object recognition [25] and tracking [4], where they have excelled for face verification in scenarios with relatively little training data, known as *one-shot* or *few-shot learning*. As shown in Figure 1, a Siamese network consists of two symmetrical CNNs with shared weights, and a loss function based on the outputs of these two networks and a ground truth label.

The network is trained from tuples of the form (y_a, y_n, y_p) from Y where y_a and y_p are photos of y that have been *Liked* by x and y_n is a photo of y that has been *Disliked* by x . Using y_a as an anchor, two pairs are made from the tuple; (y_a, y_p) is a positive pair where the expected output is 1.0 and (y_a, y_n) as the negative pair where the expected output is 0.0. From these pairs, the network is trained to differentiate between a Liked and a Disliked image, given another Liked image.

The key to this is that the network uses shared weight parameters W for the training and inference process. We map y_1 and y_2 to h_{y1} and h_{y2} using W , which are two points in a 128 dimensional space. We calculate the difference between the two points as follows:

$$D_W(y_1, y_2) = |h_{y1} - h_{y2}| \quad (2)$$

Relating the Siamese network back to our original problem formulation in Section 3.1, this gives us a basis for estimating a unidirectional preference relation $P^{x,y}$ based on two images, one image

from x 's preference history S_k^x , and the current user y , solving the problem:

$$p^{x,y} = g(S_k^x, y; \theta) \quad (3)$$

Layer	Size-in	Size-out	Kernel	Param
input		100x100x3		0
conv1	100x100x3	100x100x3	7x7x3	444
maxpooling1	100x100x3	34x34x3	3x3	
normalization1	34x34x3	34x34x3		12
conv2	34x34x3	34x34x64	3x3x64	1792
maxpooling2	12x12x64	12x12x64	3x3	
normalization2	12x12x64	12x12x64		256
conv3	34x34x3	12x12x192	2x2x192	49344
maxpooling3	12x12x64	4x4x192	3x3	
conv4	4x4x192	4x4x384	2x2x384	295296
maxpooling4	4x4x384	2x2x384	3x3	
conv5	2x2x384	2x2x256	1x1x256	98560
conv6	2x2x256	2x2x256	3x3x256	590080
maxpooling5	2x2x256	1x1x256	3x3	
flatten	1x1x256	256		
dense1	256	256		65792
dense2	256	128		32896

Table 1: The layers of the CNN used as the symmetrical part of the Siamese Network.

3.3.2 Network Layers. Table 1 shows the architecture of the CNN that makes up the two symmetrical branches of the Siamese network. The small convolution kernels used have been shown to effectively identify facial features in deep convolutional networks [19]. The network was trained using an Adam optimiser, with a learning rate of 0.0001.

The output of the network is a value between 0 and 1 expressed by a Sigmoid function, representing whether x is more likely to *Like* or *Dislike* y based on the two images. In the next section, we extend this to use the whole of x 's preference history S_x and show how this becomes an even more effective predictor of preference.

3.3.3 Loss Function. The Siamese network was trained using *binary crossentropy*. This is a standard loss function used in training neural networks, the formula for which is given below. In the following equation, Y is the binary variable representing Like and Dislike, $D_W(y_1, y_2)$ is the embedded distance between two images, g is a neural network and $g(D_W(y_1, y_2))$ is the predicted probability of $D_W(y_1, y_2)$ resulting in a Like.

$$L(y_1, y_2) = -(Y \log(g(D_W(y_1, y_2))) + (1 - Y) \log(1 - g(D_W(y_1, y_2)))) \quad (4)$$

Binary crossentropy was shown experimentally to result in the highest effectiveness metrics for the network.

We note that it is also common to train Siamese networks using a *contrastive loss* function, which uses a *margin* m to increase the network's error when it misclassifies two very similar images. In most situations where a Siamese network is applied, such as face detection, misclassifying two very similar images is as incorrect

as misclassifying two very different images. However, this is not true in our application, where user preferences are not necessarily categorical, and similar images are more likely to be liked by a user than very different images. This explains why binary crossentropy may have been more effective in our tests.

The contrastive loss function is defined below, where the terminology used is the same as in equation 4 and m is the margin:

$$L(y_1, y_2) = (1 - Y) \frac{1}{2} (D_W(y_1, y_2))^2 + Y \frac{1}{2} (\max(0, m - D_W(y_1, y_2)))^2 \quad (5)$$

3.4 Incorporating RNNs for Learning User Preference History

The Siamese network described above, when trained on unidirectional preference, is an effective if elementary model. In this section, we describe the RNN we use to interpret the user history based on the results of the Siamese network.

The output of the Siamese network is a point in 128-dimensional space that represents the preference of a user x for an image y_k based on comparison with the anchor image y_a . Based on initial experimental work, we chose an LSTM-based RNN architecture to interpret the time series of images. The *forget gate* of the LSTM is particularly intuitive in this case. For a state s_t at time t , a forget gate described by f_t , a write gate i_t and a candidate write \tilde{s}_t derived from the input and the previous state, the next state is described by the equation:

$$s_t = f_t \odot s_{t-1} + i_t \odot \tilde{s}_t \quad (6)$$

We might intuitively expect that preferences expressed by users would change over time, and the forget behaviour of the LSTM allows us to model this, with the input for the state s_t of the LSTM modelling the preferences of user x being the user S_t^x , and the final input at $s_{|S^x|+1}$ being the user y whom we wish to estimate x 's preference for.

The LSTM is visualised in Figure 2. Because users have variable length preference histories, we fill the histories of users with shorter histories with dummy images and use a masking layer to filter them. The LSTM and subsequent dense neural network form a representation in 256-dimensional space of the user's preference as a time series.

Layer	Size-in	Size-out	Kernel	Param
input		128x15		0
LSTM	128x15	128	128	128
dense1	128	256	1	128
concat	128x2	256	1	256
dense2	256	128	1	256
output	128	1	1	128

Table 2: The layers of TIRR following the mapping of images into 128-dimensional space by the pre-trained Siamese network

Specifically, the network consists of an input layer, which accepts a maximum of 15 outputs from Siamese networks in 256-dimensional space concatenated together. Experiments determined

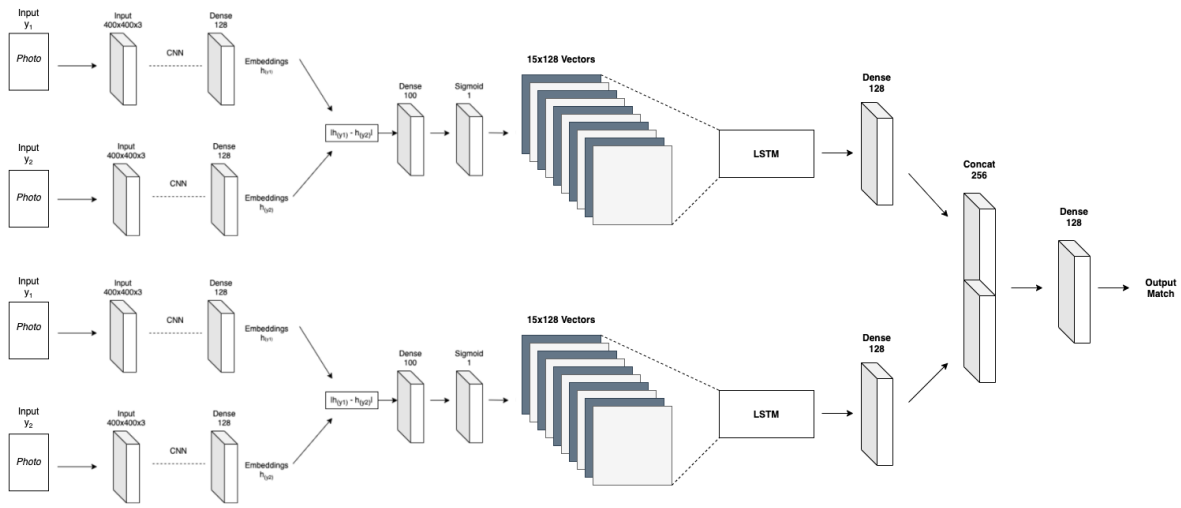


Figure 2: TIRR: the architecture to predict matches using an LSTM to interpret historical preference data on user photos.

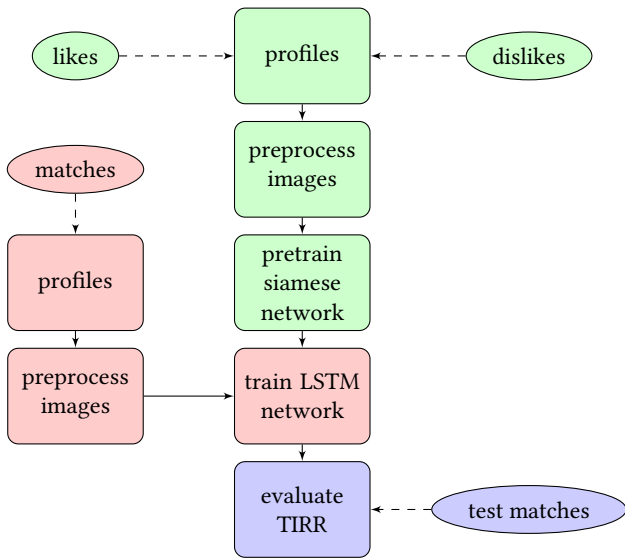


Figure 3: The process by which TIRR is trained. Three independent datasets used represented by different colours.

that more than this did not significantly alter the performance of the network. The layers are described in Table 2. If a user has fewer preferences expressed than this, the earlier images are filled with zeroes, and the network learns to interpret this as dummy data. Following the LSTM, the network consists of a single dense layer of 128 neurons, and then a dropout layer with a dropout rate of 0.4. The network was trained with an Adam optimiser with a learning rate of 0.0001.

3.5 Training and Match Prediction

This section describes training the network to predict matches between two users. As described in Section 3.2, our objective is to differentiate between interactions consisting of bidirectional expressions of preference, *Matches*, and unidirectional expressions of negative preference, *Dislikes*.

The full training process is visualised in Figure 3. Our experiments determined that the network trained extremely slowly when trained in its full form from an initial randomised state, and we therefore pre-trained the Siamese network segment of the network using one dataset, shown in green. The subsequent training of the full system on matches was done using a separate dataset, shown in red. The final evaluation was done using a third dataset, shown in blue. In addition, Neve et al. demonstrated that the Siamese Network training was more effective when two networks were trained separately on male and female data [18]. As the service providing our data currently only supports heterosexual dating, this split does not decrease the usefulness of the application in this case.

Training for the Siamese networks were based on 500000 triplets (y_a, y_p, y_n) sampled from 250000 users split evenly over male and female images. Images were cropped and centered on the faces of users before training. Other methods of preprocessing such as affine transformations, which have been shown to improve the predictive power of other networks [15] did not have any impact on performance. The Siamese networks were trained to predict unidirectional preferences i.e. y_p was an image x had *Liked* (but not necessarily with reciprocity) and y_n was an image x had *Disliked*.

Following convergence of the Siamese network, the LSTM network was trained based on the preference histories of 100000 users to predict *Matches* and *Like-Dislike Tuples*. This dataset was separate from the dataset used to train the Siamese network. Histories were capped at one year, because of concerns that changes to the service’s design and search algorithm over time might have an effect on user preferences. They were also capped to a maximum

of 15 preferences, because initial experiments showed that longer sequences did not improve accuracy, and because some outlier users express thousands of preferences, which results in an unreasonable increase in training and prediction times.

Finally, the LSTM was validated on a separate dataset of 20000 *Matches* and *Like-Dislike Tuples*. There was no overlap in preference expression between the three datasets. There was overlap between the users contained in these datasets, but as in a real-world situation the system would be trained based on users on the service and subsequently used to make predictions for those users in addition to new users, testing in this way is valid and representative.

4 RESULTS & EVALUATION

In this section, we present the results for TIRR compared to the current state of the art in both content-based and collaborative filtering.

4.1 Evaluation Metrics for Reciprocal Environments

RRSs generally use similar metrics for success as standard machine learning models: evaluation via the *ROC Curve* and the related metrics *Precision*, *Recall* and their combined metric *F1 Score*. However, because of the requirement for reciprocal success, their definitions are a little different in this scenario, so we present them here as defined by Pizzato et al. in [20]. In the following equations, R is the set of recommended users, RL is the set of recommended users who matched with each other, and RN is the set of recommended users where at least one expressed negative preference for the other.

$$Precision = \frac{|RL|}{|RL| + |RN|} \quad (7)$$

$$Recall = \frac{|RL|}{|R|} \quad (8)$$

$$F1 = \frac{2 * Precision * Recall}{Precision + Recall} \quad (9)$$

As the models predict a value between 0.0 and 1.0 that represents the strength of the mutual preference relation, the ROC curves in this section are drawn by moving a threshold between these two values and plotting the true and false positive rates.

4.2 Siamese Network Results

We first describe the results generated by the pretrained Siamese network. This network provides a basis for the main user preference prediction model, as the output embeddings from this network provide an input for the RNN.

The ROC curve for the Siamese network is displayed in Figure 4 as the blue line (the green dotted line is the 1-0 reference line). This curve was drawn based on a test set of 20000 samples not in the original training dataset. In general the network is capable of differentiating between a single *Liked* image and a single *Disliked* image based on an anchor image. The curve itself is slightly erratic, but this is not entirely unexpected: a single anchor image is unlikely to have enough information to differentiate between positive and negative preference.

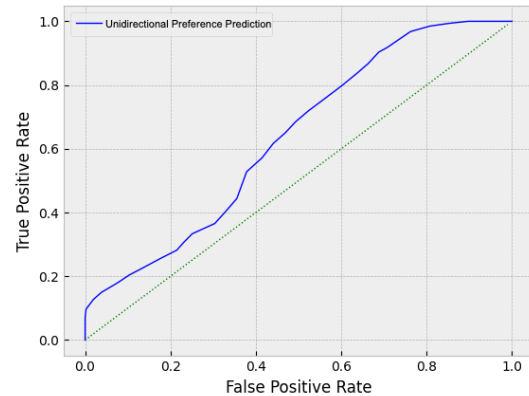


Figure 4: Pretrained Siamese Network ROC curve. This forms a building block of both ImRec and our proposed TIRR model.

As the model by itself is not directly the source of the recommendations, it would not be appropriate to compare it to other recommender systems. For this reason, we present this model without a point of comparison. However, in Section 4.3 we will compare two approaches that use this model as a building block for reciprocal recommendation.

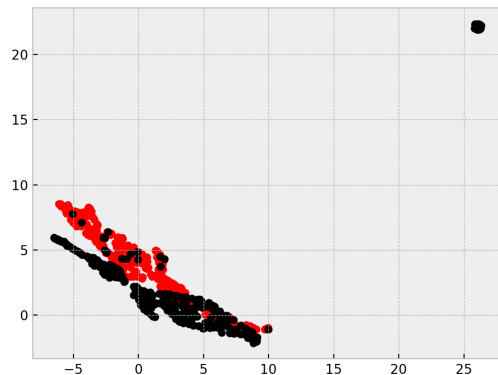


Figure 5: UMAP embeddings of the pretrained Siamese network forming part of TIRR. The red points represent *Liked* images while black points represent *Disliked* images.

As the embeddings in 128-dimensional space from the output of the Siamese network form the input of the RNN. It is therefore useful to visualise these embeddings. In order to do this, we use *Uniform Manifold Approximation and Projection for Dimensionality Reduction* (UMAP) to reduce the 128-dimensional vectors to two-dimensional vectors for visualisation. This visualisation is displayed in Figure 5.

In this visualisation, the black datapoints represent *Disliked* images and the red datapoints represent *Liked* images. It is clear from the visualisation that the embeddings are separable to some extent, even in two dimensions. The anomalous black cluster in the top

right of the image represents heavily distorted or very poor quality images, or images misclassified by the face detection algorithm (i.e. images that do not contain a face). These tend to be almost universally *Disliked*.

4.3 TIRR vs Content-Based Algorithms

As described in Section 1, recommender systems are divided into content-based algorithms and collaborative filtering algorithms.




Figure 6: Content Based Algorithm ROC Curves demonstrating the significant improvement in AUC with TIRR.

Figure 6 displays a comparison of TIRR with other content-based algorithms. As described in Section 2, *RECON* [21] is an algorithm that identifies a user’s implicit preferences for categorical data, and *ImRec* [18] is an algorithm that uses images to make predictions without the RNN-based component of TIRR, instead using a Random Forest and aggregation function.

RECON struggled to generate effective recommendations on our dataset. As *RECON* was also evaluated on a private dataset, without comparing the datasets directly, it is difficult to establish why this is, but one possibility is that modern dating services place a higher emphasis on visual content than services did ten years ago, at the time *RECON* was developed. *ImRec* performs better than *RECON*, but performs significantly worse than our proposed method *TIRR*. The key difference between *TIRR* and *ImRec* is the RNN-based process that allows *TIRR* to interpret historical and time-series data in order to make predictions, whereas *ImRec* treats user preferences in a global way, with no ability to capture individual users preferences.

Algorithm	F1 Score	Precision	Recall	AUC
<i>RECON</i>	0.61	0.56	0.68	0.51
<i>ImRec</i>	0.71	0.60	0.88	0.65
<i>TIRR</i>	0.87	0.86	0.88	0.91

Table 3: Results based on best F1 score for content-based algorithms. Here we can see that our proposed method *TIRR* significantly outperforms the other approaches.

The AUC and maximum F1 score for the three algorithms is described in Table 3. The scores are based on the threshold that

gave the best F1 score in the training set, used in the test set. We consider that this significant improvement of our proposed method *TIRR* derives from the ability of our algorithm to interpret a user’s history of preferences for images over time, and take account of a user’s potentially shifting preferences, whereas *Imrec* provides a global model across all users without distinguishing more than one preference per user at a time, and *RECON* doesn’t make use of images at all.

The table also lists the precision and recall at the points where the best F1 score was recorded. While F1 is an excellent measure of overall performance of an algorithm, the individual precision and recall numbers and their balance are particularly important in RS research because precision tends to influence the trust users have in the RS, which in turn affects their use of it [7]. It is noteworthy that while *ImRec* was relatively successful at predicting which image a user would like, its precision was relatively low in comparison with other algorithms, whereas *TIRR* has very high precision, and is therefore more likely to be trusted and used.

4.4 TIRR vs Collaborative Filtering

In addition to comparing *TIRR* to other content-based RRSs, we also ran tests comparing it to the current best-in-class collaborative filtering algorithms, *RCF* and *LFRR*.




Figure 7: ROC Curves showing the performance of the content-based *TIRR* against the current state of the art collaborative filtering algorithm *LFRR*.

LFRR is a collaborative filtering algorithm based on latent factor models trained by stochastic gradient descent, and *RCF* is a neighbourhood-based collaborative filtering algorithm. *TIRR* outperformed both of these algorithms on our test dataset, although by a slimmer margin than its lead on current content-based filtering algorithms. Nonetheless, this represents a significant advancement in the field of reciprocal recommendation, as in services where images prominently used, our algorithm is likely to be more effective than current collaborative filtering methods.

Table 4 lists the peak performance metrics for the two algorithms. In addition to the higher F1 score, *TIRR* also has a comparable balance of precision and recall to *LFRR*.

Algorithm	F1 Score	Precision	Recall	AUC
LFRR	0.86	0.86	0.85	0.90
TIRR	0.87	0.86	0.88	0.91

Table 4: Results based on best F1 score for the TIRR and LFRR algorithms. Here we can see that the content-based TIRR improves upon the collaborative filtering-based LFRR.

5 CONCLUSIONS

In this paper, we presented a novel algorithm to interpret user preference history using *only* photos and make predictions about future preferences for reciprocal recommendation. We demonstrated that this can effectively be used as a predictor for the probability of mutual preference between two users, and therefore forms the basis for an effective recommender system. We also demonstrated that our algorithm outperforms state of the art reciprocal recommender systems in offline tests using a large dataset from a dating service with real users.

This research demonstrates the value of including historical preference in reciprocal recommendation. Previous research has demonstrated the value of using RNNs to interpret sequences of preferences in user-item recommendation, but this is the first time it has been used in reciprocal recommendation. The improvement over a similar algorithm that does not use sequences of data shows the value of this approach.

Finally, the model itself represents a significant advance in the field of content-based reciprocal recommendation. The model's success also allows us to draw interesting conclusions about the significance of photos in online dating, given their strong predictive power in this dataset. It also provides interesting insight into the potential power of content-based algorithms in online dating: while in many fields, they are outperformed by collaborative filtering, the algorithm presented in this paper performs better on evaluation metrics than the current state-of-the-art collaborative filtering algorithm.

REFERENCES

- [1] Charu Aggarwal. 2016. *Recommender Systems: The Textbook* (1st ed.). Springer, London, England.
- [2] Trapit Bansal, David Belanger, and Andrew McCallum. 2016. Ask the GRU: Multi-task Learning for Deep Text Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys 2016)*. ACM, New York, NY, 107–114. <https://doi.org/10.1145/2959100.2959180>
- [3] Robert Bell and Yehuda Koren. 2007. Lessons from the Netflix prize challenge. *ACM SIGKDD Explorations Newsletter - Special issue on visual analytics* 9, 2 (Dec. 2007), 75–79. <https://dl.acm.org/citation.cfm?id=1345465>
- [4] Luca Bertinetto, Jack Valmadre, João Henriques, Andrea Vedaldi, and Philip Torr. 2016. Fully-Convolutional Siamese Networks for Object Tracking. In *Proceedings of the 2016 European Conference on Computer Vision (ECCV 2016)*. Springer, 850–865. https://link.springer.com/chapter/10.1007/978-3-319-48881-3_56
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2009)*. IEEE, Miami, FL, 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
- [6] Jianming He and Wesley Chu. 2010. A Social Network-Based Recommender System (SNRS). *Data Mining for Social Network Data* 12 (May 2010), 47–74. https://link.springer.com/chapter/10.1007/978-1-4419-6287-4_4
- [7] Jonathan Herlocker, Joseph A. Konstan, L. Terveen, and John Riedl. 2004. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems* 22, 1 (Jan. 2004). <https://doi.org/10.1145/963770.963772>
- [8] Sepp Hochreiter. 1998. The Vanishing Gradient Problem During Learning Recurrent Neural Nets and Problem Solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 6, 2 (May 1998), 107–116. <https://www.worldscientific.com/doi/abs/10.1142/s0218488598000094>
- [9] Sepp Hochreiter and Jurgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (May 1997), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- [10] Akiva Kleinerman, Ariel Rosenfeld, Francesco Ricci, and Sarit Kraus. 2018. Optimally balancing receiver and recommended users' importance in reciprocal recommender systems. In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18)*. ACM, New York, NY, 131–139. <https://doi.org/10.1145/3240323.3240349>
- [11] G Koch, R Zemel, and R Salakhutdinov. 2015. Siamese neural networks for one-shot image recognition. In *Proceedings of the 2015 ICML Deep Learning workshop*. ICML.
- [12] Xuan Nhat Lam, Thuc Vu, Trong Duc Le, and Anh Duc Duong. 2008. Addressing cold-start problem in recommendation systems. In *Proceedings of the 2nd international conference on Ubiquitous information management and communication (ICUIMC '08)*. ACM, New York, NY, 208–211. <https://doi.org/10.1145/1352793.1352837>
- [13] Ken Lang. 1995. NewsWeeder: Learning to Filter Netnews. In *Proceedings 12th International Conference on Machine Learning, (ICML 1995)*. 331–339. <https://pdfs.semanticscholar.org/26fd/e7f657b41be65a0b975615508f4f100e3a0a.pdf>
- [14] Chenyi Lei, Dong Liu, Weiping Li, Zheng-Jun Zha, and Houqiang Li. 2016. Comparative Deep Learning of Hybrid Representations for Image Recommendations. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2016)*. IEEE, 2545–2553. http://openaccess.thecvf.com/content_cvpr_2016/html/Lei_Comparative_Deep_Learning_CVPR_2016_paper.html
- [15] Yoad Lewenberg, Yoram Bachrach, Sukrit Shankar, and Antonio Criminisi. 2016. Predicting Personal Traits from Facial Images Using Convolutional Neural Networks Augmented with Facial Landmark Information. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI 2016)*. AAAI, 4365–4366. <https://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/viewPaper/12384>
- [16] Jovian Lin, Kazunari Sugiyama, Min-Yen Kan, and Tat-Seng Chua. 2013. Addressing cold-start in app recommendation: latent user models constructed from twitter followers. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval (SIGIR '13)*. ACM, New York, NY, 283–292. <https://doi.org/10.1145/2484028.2484035>
- [17] Qiang Liu, Shu Wu, and Liang Wang. 2017. DeepStyle: Learning User Preferences for Visual Recommendation. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2017)*. ACM, New York, NY, 841–844. <https://dl.acm.org/doi/abs/10.1145/3077136.3080658>
- [18] James Neve and Ryan McConville. 2020. ImRec: Learning Reciprocal Preferences Using Images. In *Proceedings of the Fourteenth ACM Conference on Recommender Systems (RecSys '2020)*. ACM, New York, NY, 170–179. <https://doi.org/10.1145/3383313.3411476>
- [19] OM Parkhi, A Vedaldi, and A Zisserman. 2015. Deep face recognition. *BMVA (August 2015)*, 1–12. <https://ora.ox.ac.uk/objects/uuid:a5f2e93f-2768-45bb-8508-74747f85cad1>
- [20] Luiz Pizzato, Tomasz Rej, Joshua Akehurst, Irena Koprinska, Kalina Yacef, and Judy Kay. 2013. Recommending people to people: the nature of reciprocal recommenders with a case study in online dating. *User Model User-Adap Inter* 23, 5 (Nov. 2013), 447–488. <https://link.springer.com/article/10.1007/s11257-012-9125-0>
- [21] Luiz Pizzato, Tomek Rej, Thomas Chung, Irena Koprinska, and Judy Kay. 2010. RECON: a reciprocal recommender for online dating. In *Proceedings of the fourth ACM conference on Recommender systems (RecSys '10)*. ACM, New York, NY, 207–214. <https://doi.org/10.1145/1864708.1864747>
- [22] Zheng Siting, Hong Wenxing, Zhang Ning, and Yang Fan. 2012. Job recommender systems: A survey. In *Proceedings of the 7th International Conference on Computer Science & Education (ICCSE '12)*. IEEE, Melbourne, VIC, Australia, 920–924. <https://doi.org/10.1109/ICCSE.2012.6295216>
- [23] Bartłomiej Twardowski. 2016. Modelling Contextual Information in Session-Aware Recommender Systems with Neural Networks. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. ACM, New York, NY, 273–276. <https://doi.org/10.1145/2959100.2959162>
- [24] Robin van Meteren and Maarten van Someren. 2000. Using Content-Based Filtering for Recommendation. In *Proceedings of the ECML 2000 Workshop: Matching Learning in Information Age, (ECML 2000)*. 47–56. http://users.ics.forth.gr/~potamias/mlnia/paper_6.pdf
- [25] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. Matching Networks for One Shot Learning. *Advances in Neural Information Processing Systems* 29 (2016). <http://papers.nips.cc/paper/6068-learning-feed-forward-one-shot-learners.pdf>
- [26] Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alexander Smola, and How Jing. 2017. Recurrent Recommender Networks. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (WSDM '17)*. ACM, New York, NY, 495–503. <https://doi.org/10.1145/3018661.3018689>
- [27] Peng Xia, Benyuan Liu, Yizhou Sun, and Cindy Chen. 2015. Reciprocal Recommendation System for Online Dating. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM '15)*. ACM, New York, NY, 234–241.