Date Estimation in the Wild of Scanned Historical Photos: An Image Retrieval Approach

Adrià Molina[0000−0003−0167−8756], Pau Riba[0000−0002−4710−0864], Lluis Gomez[0000−0003−1408−9803], Oriol Ramos-Terrades[0000−0002−3333−8812], and Josep Lladós[0000−0002−4533−4739]

Computer Vision Center and Computer Science Department, Universitat Autònoma de Barcelona, Catalunya {priba,lgomez,oriolrt,josep}@cvc.uab.cat, adria.molinar@e-campus.uab.cat

Abstract. This paper presents a novel method for date estimation of historical photographs from archival sources. The main contribution is to formulate the date estimation as a retrieval task, where given a query, the retrieved images are ranked in terms of the estimated date similarity. The closer are their embedded representations the closer are their dates. Contrary to the traditional models that design a neural network that learns a classifier or a regressor, we propose a learning objective based on the nDCG ranking metric. We have experimentally evaluated the performance of the method in two different tasks: date estimation and date-sensitive image retrieval, using the DEW public database, overcoming the baseline methods.

Keywords: Date Estimation · Historical Photographs · Image Retrieval · Ranking Loss · Smooth-nDCG.

1 Introduction

Historical archives and libraries contain a large variability of document sources that reflect the memory of the past. The recognition of the scanned images of these documents allows to reconstruct the history. A particular type of archival data are historical photographs which are full of evidence that tells us the story of that snapshot in time. One just needs to pay attention to the subtle cues that are found in different objects that appear in the scene: the clothes that people wear, their haircut styles, the overall environment, the tools and machinery, the natural landscape, etc. All of these visual features are important cues for estimating its creation date. Apart from that, texture and color features might also be of great help to accurately estimate of image creation date since photographic techniques have evolved throughout history and have imprinted a specific date fingerprint on them.

Date estimation of cultural heritage photographic assets is a complex task that is usually performed by experts (e.g. archivists or genealogists) that exploit their expert knowledge about all the features mentioned above to provide

precise date estimations for undated photographs. But their manual labor is costly and time consuming, and automatic image date estimation models are of great interest for dealing with large scale archive processing with minimal human intervention.

Most approaches in date estimation for historical images try to directly compute the estimation through classification or regression [5,9,12]. As alternative of these classical approaches, in this paper we present a method for date estimation of historical photographs in a retrieval scenario. Thus, the date estimation of photographs is incorporated in the ranked results for a given query image. This allows to predict the date of an image contextualized regarding the other photographs of the collection. In the worst case, when the exact date is not exactly estimated, the user can obtain a relative ordering (one photograph is older than another one), which is useful in archival tasks of annotating document sources. The proposed model for historical photograph retrieval is based in a novel ranking loss function *smooth-nDCG* based on the Normalized Discounted Cumulative Gain ranking metric, which is able to train our system according to a known relevance feedback; in our case the distance in years between images.

The main idea in our approach relies on optimizing rankings such that the closer is image's date to the query's date for a certain photograph the higher will be ranked. When receiving an unknown image as query the method computes the distances towards a support dataset, consisting of a collection of images with known dates. The date is estimated assuming that the highest ranked images are images from the same date.

In contrast to the literature reviewed in Section 2, our method allows not only to predict but to rank images given a query. This means that considering an image from a certain year our system is capable of retrieving a list of images ordered by time proximity. This may be useful for many applications and systems that rely on retrieving information from a large amount of data.

The rest of the paper is organized as follows: In Section 2 we present the most relevant state of the art related to our work. The key components of the proposed model are described in Section 3, where we describe the learning objectives considered in the training algorithms, and in Section 4, where we outline the architecture and training process of the model. Section 5 provides the experimental evaluation and discussion. Finally, Section 6 draws the conclusions.

2   Related Work

The problem of automatic estimation of the creation date of historical photographs is receiving increased attention by the computer vision and digital humanities research communities. The first works go back to Schindler *et al.* [14,13], where the objective was to automatically construct time-varying 3D models of a city from a large collection of historical images of the same scene over time. The process begins by performing feature detection and matching on a set of input photographs, followed by structure from motion (SFM) to recover 3D points and camera poses. The temporal ordering task was formulated as a constraint satis-

faction problem (CSP) based on the visibility of structural elements (buildings) in each image.

More related with the task addressed in our work, Palermo *et al.* [10] explored automatic date estimation of historical color images based on the evolution of color imaging processes over time. Their work was formulated as a five-way decade classification problem, training a set of one-vs-one linear support vector machines (SVMs) and using hand-crafted color based features (e.g. color co-occurrence histograms, conditional probability of saturation given hue, hue and color histograms, etc.). Their models were validated on a small-scale dataset consisting of 225 and 50 images per decade for training and testing, respectively.

In a similar work, Fernando *et al.* [3] proposed two new hand-crafted color-based features that were designed to leverage the discriminative properties of specific photo acquisition devices: color derivatives and color angles. Using linear SVMs they outperformed the results of [10] as well as a baseline of deep convolutional activation features [2]. On the other hand, martin *et al.* [8] also showed that the results of [10] could be slightly improved by replacing the one-vs-one classification strategy by an ordinal classification framework [4] in which relative errors between classes (decades) are adequately taken into account.

More recently, Müller *et al.* [9] have brought up to date the task of image date estimation in the wild by contributing a large-scale, publicly available dataset and providing baseline results with a state-of-the-art deep learning architecture for visual recognition [15]. Their dataset, Date Estimation in the Wild (DEW), contains more than one million Flickr[1] images captured in the period from 1930 to 1999, and covering a broad range of domains, e.g., city scenes, family photos, nature, and historical events.

Beyond the problem of unconstrained date estimation of photographs, which is the focus of this paper, there are other related works that have explored the automatic prediction of creation dates of certain objects [7,16], or of some specific types of photographs such as yearbook portraits [5,12].

In this work, contrary to all previously published methods, we approach the problem of date estimation from an image retrieval perspective. We follow the work of Brown *et al.* [1] towards differentiable loss functions for information retrieval. More precisely, our model learns to estimate the date of an image by minimizing the the Normalized Discounted Cumulative Gain.

In all our experiments we use Müller *et al.*'s Date Estimation in the Wild [9] dataset. We also share with [9] the use of state of the art convolutional neural networks in contrast to classic machine learning and computer vision approaches, such as [3].

3   Learning Objectives

As mentioned in Section 1 our retrieval system relies on a neural network [6] trained to minimize a differentiable ranking function [1]. Traditionally, information retrieval evaluation has been dominated by the mean Average Precision

---

[1] https://www.flickr.com/

(mAP) metric. In contrast, we will be using the Normalized Discounted Cumulative Gain (nDCG). The problem with mAP is that it fails in measuring ranking success when the ground truth relevance of the ranked items is not binary.

In the task of date estimation the ground truth data is numerical and ordinal (e.g. a set of years in the range 1930 to 1999), thus given a query the metric should not punish equally for ranking as top result a 1-year difference image than a 20-years difference one. This problem is solved in the nDCG metric by using a relevance score that measures how relevant is a certain sample to our query. This allows us to not only to deal with date estimation retrieval but to explicitly declare what we consider a good ranking.

In this section we derive the formulation of the smooth-nDCG loss function that we will use to optimize our date estimation model.

Mean Average Precision: First, let us define *Average Precision* (AP) for a given query q as

$$\mathrm{AP}_q = \frac{1}{|\mathcal{P}_q|} \sum_{n=1}^{|\Omega_q|} P@n \times r(n), \tag{1}$$

where $P@n$ is the precision at n and $r(n)$ is a binary function on the relevance of the n-th item in the returned ranked list, $\mathcal{P}_q$ is the set of all relevant objects with regard the query q and Ω_q is the set of retrieved elements from the dataset. Then, the mAP is defined as:

$$\mathrm{mAP} = \frac{1}{Q} \sum_{q=1}^{Q} \mathrm{AP}_q, \tag{2}$$

where Q is the number of queries.

Normalized Discounted Cumulative Gain: In information retrieval, the *normalized Discounted Cumulative Gain* (nDCG) is used to measure the performance on such scenarios where instead of a binary relevance function, we have a graded relevance scale. The main idea is that highly relevant elements appearing lower in the retrieval list should be penalized. In the opposite way to mAP, elements can be relevant despite not being categorically correct with respect to the query. The *Discounted Cumulative Gain* (DCG) for a query q is defined as

$$\mathrm{DCG}_q = \sum_{n=1}^{|\Omega_q|} \frac{r(n)}{\log_2(n+1)}, \tag{3}$$

where $r(n)$ is a graded function on the relevance of the n-th item in the returned ranked list and Ω_q is the set of retrieved elements as defined above. In order to allow a fair comparison among different queries that may have a different sum of relevance scores, a normalized version was proposed and defined as

$$\mathrm{nDCG}_q = \frac{\mathrm{DCG}_q}{\mathrm{IDCG}_q}, \tag{4}$$

where IDCG_q is the ideal discounted cumulative gain, *i.e.* assuming a perfect ranking according to the relevance function. It is formally defined as

$$\text{IDCG}_q = \sum_{n=1}^{|\Lambda_q|} \frac{r(n)}{\log_2(n+1)} \tag{5}$$

where Λ_q is the ordered set according the relevance function.

As we have exposed, we want to optimize our model at the retrieval list level, but the two classical retrieval evaluation metrics defined above are not differentiable. In the following we define two 'smooth' functions for ranking optimization that are inspired in the mAP and nDCG metrics respectively.

Ranking Function (R). Following the formulation introduced in [11], these two information retrieval metrics can be reformulated by means of the following ranking function,

$$\mathcal{R}(i, \mathcal{C}) = 1 + \sum_{j=1}^{|\mathcal{C}|} \mathbb{1}\{(s_i - s_j) < 0\}, \tag{6}$$

where $\mathcal{C}$ is any set (such as Ω_q or $\mathcal{P}_q$), $\mathbb{1}\{\cdot\}$ is the Indicator function, and s_i is the similarity between the i-th element and the query according to S. In this work we use the cosine similarity as S but other similarities such as inverse euclidean distance should work as well. Let us then define cosine similarity as:

$$S(v_q, v_i) = \frac{v_q \cdot v_i}{\|v_q\|\|v_i\|}. \tag{7}$$

Still, because of the need of a indicator function $\mathbb{1}\{\cdot\}$, with this formulation we are not able to optimize following the gradient based optimization methods and we require of a differentiable indicator function. Even though several approximations exists, in this work we followed the one proposed by Quin *et al.* [11]. Thus, we make use of the sigmoid function

$$\mathcal{G}(x; \tau) = \frac{1}{1 + e^{\frac{-x}{\tau}}}. \tag{8}$$

where τ is the temperature of the sigmoid function. As illustrated in Figure 1 the smaller is τ the less smooth will be the function. For small values the gradient is worse defined but the approximation is better since the sigmoid becomes closer to the binary step function.

Smooth-AP: Following [1] and replacing the term $P@n \times r(n)$ in Eq. 1 by the Ranking Function introduced in Eq. 6, the AP equation becomes

$$\text{AP}_q = \frac{1}{|\mathcal{P}_q|} \sum_{i \in \mathcal{P}_q} \frac{1 + \sum_{j \in \mathcal{P}_q, j \neq i} \mathbb{1}\{D_{ij} < 0\}}{1 + \sum_{j \in \Omega_q, j \neq i} \mathbb{1}\{D_{ij} < 0\}}. \tag{9}$$

where we use $D_{ij} = s_i - s_j$ for a more compact notation. Finally, replacing the indicator function by the sigmoid function (Eq. 8) we obtain a smooth approximation of AP
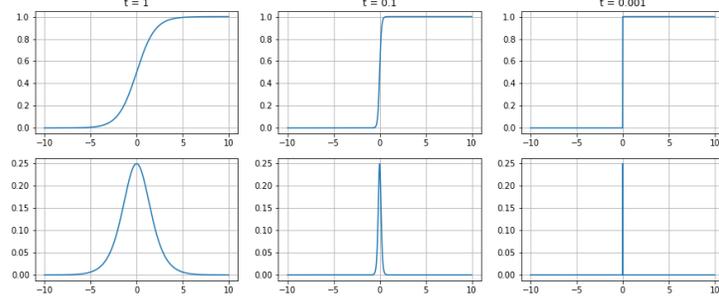
Fig. 1. Shape of the smooth indicator function (top) and its derivative (bottom) for different values of the temperature: 1 (left), 0.1 (middle), 0.001 (right).

$$\mathrm{AP}_q \approx \frac{1}{|\mathcal{P}_q|} \sum_{i \in \mathcal{P}_q} \frac{1 + \sum_{j \in \mathcal{P}_q, j \neq i} \mathcal{G}(D_{ij}; \tau)}{1 + \sum_{j \in \Omega_q, j \neq i} \mathcal{G}(D_{ij}; \tau)}. \tag{10}$$

Averaging this approximation for all the queries in a given batch, we can define our loss as

$$\mathcal{L}_{AP} = 1 - \frac{1}{Q} \sum_{i=1}^{Q} AP_q, \tag{11}$$

where Q is the number of queries.

Smooth-nDCG: Following the same idea as above, we replace the n-th position in Eq. 3 by the ranking function, since it defines the position that the i-th element of the retrieved set, and the DCG metric is expressed as

$$\mathrm{DCG}_q = \sum_{i \in \Omega_q} \frac{r(i)}{\log_2 \left(2 + \sum_{j \in \Omega_q, j \neq i} \mathbb{1}\{D_{ij} < 0\}\right)}, \tag{12}$$

where $r(i)$ is the same graded function used in Eq. 3 but evaluated at element i. Therefore, the corresponding smooth approximation is

$$\mathrm{DCG}_q \approx \sum_{i \in \Omega_q} \frac{r(i)}{\log_2 \left(2 + \sum_{j \in \Omega_q, j \neq i} \mathcal{G}(D_{ij}; \tau)\right)} \tag{13}$$

when replacing the indicator function by the sigmoid one.

The smooth-nDCG is then defined by replacing the original DCG_q by its smooth approximation in Eq. 4 and the loss $\mathcal{L}_{nDCG}$ is defined as

$$\mathcal{L}_{nDCG} = 1 - \frac{1}{Q} \sum_{i=1}^{Q} nDCG_q, \tag{14}$$
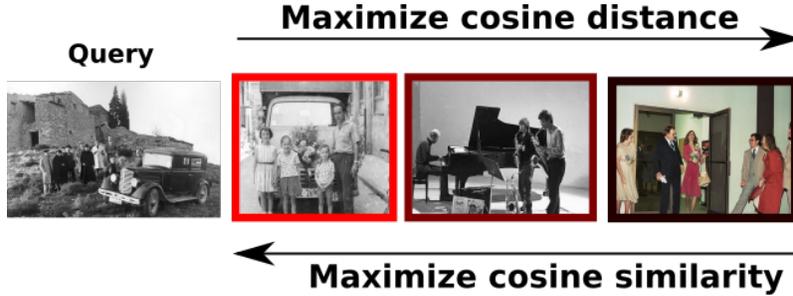
where Q is the number of queries.

Fig. 2. Objective function example, images with closer ground truth should maximize the similarity in the output space.

4   Training process

The main contribution of this paper relies on using the smooth-nDCG loss defined in the previous section for optimizing a model that learns to rank elements with non-binary relevance score. Specifically, our model learns to project images into an embedding space by minimizing the smooth-nDCG loss function. In all our experiments we use a Resnet-101 [6] convolutional neural network pretrained on the ImageNet dataset. As illustrated in Figure 2 the idea is that in the learned embedding space the cosine similarity between image projections is proportional to the actual distance in years. We will discuss in Section 5 how teaching the system to rank images according to the ground truth criteria leads the embedding space to follow a certain organization (illustrated in Figure 7).

In order to obtain such embedded representation of the photographs, we present the training algorithm (see Algorithm 1) that mainly relies on using the equations presented in the previous Section 3 in the output space for optimizing the CNN so it minimizes the smooth-nDCG loss function (see Eq. 14) and, consequently, maximizes the nDCG metric in our rankings, meaning that images will be closer in the output space as they are closer in the ground truth space.

Algorithm 1 Training algorithm for the proposed model.

    Input: Input data $\{\mathcal{X}, \mathcal{Y}\}$; CNN f; distance function D; max training iterations T
    Output: Network parameters w

1: **repeat**
2:     Process images to output embedding $h \leftarrow f_w(\{x_i\}_{i=1}^{N_{batch}})$
3:     Get Distance matrix from embeddings, all vs all rankings $M \leftarrow D(h)$
4:     Calculate relevance from training set Y Eq. 15
5:     Using the relevance score, $\mathcal{L} \leftarrow$ Eq. 14
6:     $w \leftarrow w - \Gamma(\nabla_w \mathcal{L})$
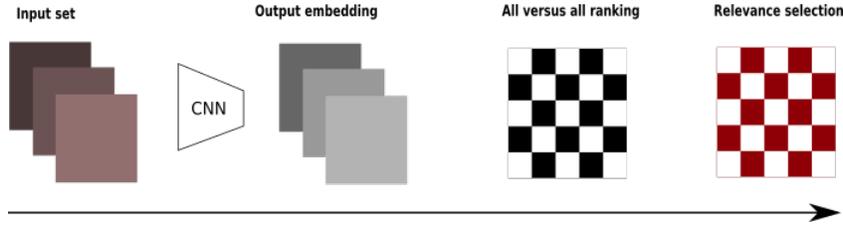7: **until** Max training iterations T

Fig. 3. Proposed baseline. Once ranked the images in the nDCG space we'll be computing smooth-nDCG in order to back-propagate the loss through the CNN.

As shown in Figure 3, once we process the images through the CNN, we compute the rankings for each one of the batch images. We use cosine similarity as distance function, however it could be replaced by any distance function model over the vector space. Since nDCG requires a relevance score for each sample, the design of the relevance function is an important part of the proposed method. A possible relevance function could be just the inverse distance in years as defined in Eq. 16.

Additionally, we decided to consider only the closest decade in the relevance function (illustrated in Figure 4), so every image further than 10 years is considered as relevant as any other. Eq. 15 is the relevance function used in our model (see Eq. 14) but other relevance settings could be, like Eq. 17 in order to exponentially punish more the images that should not be as far as they are.

$$r(n; \gamma) = \max(0, \gamma - |y_q - y_n|) \tag{15}$$

$$r(n) = \frac{1}{1 + |y_q - y_n|} \tag{16}$$

$$r(n) = e^{\frac{1}{1+|y_q - y_n|}} \tag{17}$$

where y_q and $y_n \in \mathcal{Y}$ are the dates of the query and the n-th image in the training set, respectively; and γ, in Eq. 15, is an hyper-parameter that has experimentally set to 10.

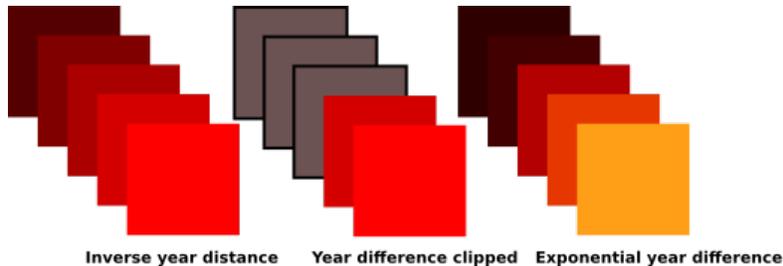Inverse year distance   Year difference clipped   Exponential year difference

Fig. 4. Different examples of relevance functions for ground truth data.

Once trained, minimizing the retrieval loss function, we can predict the year for an unknown image through k-Nearest Neighbors (k-NN) to a support set.

In one hand, we can use a set of N train images as support set; each batch is randomly selected for each prediction. Note that, since prediction relies on the k-th most similar images, the bigger the support set the better should perform the prediction. In Figure 5 we show the date estimation mean absolute error of our method on the DEW training set depending on the k parameter for the k-NN search. Hence, we would be using 10-Nearest Neighbors in all our experiments for the prediction task.
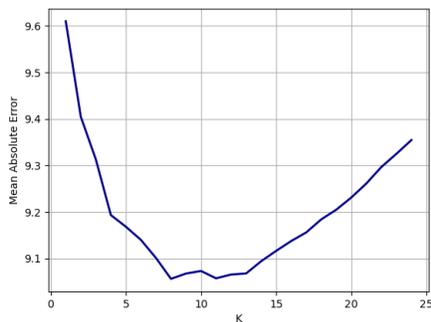

Fig. 5. Mean Absolute Error (MAE) on prediction by k-Nearest Neighbors with respect to each k using a randomly selected support set of 128 images

Additionaly, we can compute the retrievals within the entire training set with Approximate Nearest Neighbours[2], which allow us to compute the retrievals efficiently. As it's shown in Figure 6, the possibility of using the whole train set makes the k parameter more likely to be huge. We don't observe a notable divergence in the results until using a an enormous k parameter, which may indicate quite good result in terms of retrieval since worst responses are at the bottom of the dataset.

5   Experiments

In this section we evaluate the performance of the proposed method in two different tasks: date estimation and date-sensitive image retrieval. In all our experiments we use the DEW [9] dataset, which is composed of more than one million images captured in the period from 1930 to 1999. The test set is balanced across years and contains 1120 images, 16 images per year. The temperature parameter, τ, in Eq. 8, is set to 0.01.

---

[2] Using public Python library ANNOY for Approximate Nearest Neighbours `https://github.com/spotify/annoy`
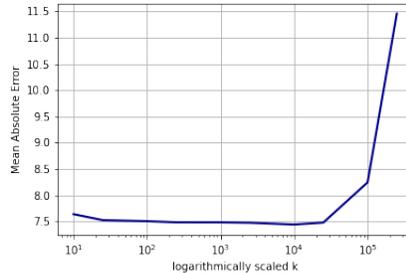
Fig. 6. Mean Absolute Error with different k for k-NN using the entire train set as support set. Note that the horizontal axis grows logarithmically.

For the task of date estimation we use the Mean Absolute Error (MAE) as the standard evaluation metric. Given a set of N images and their corresponding ground truth date annotations (years) $y = \{y_1, y_2, \ldots, y_N\}$, the MAE for a set of predictions $y' = \{y'_1, y'_2, \ldots, y'_N\}$ is calculated as follows:

$$\text{MAE} = \frac{1}{N} \sum_{i}^{N} |y_i - y'_i| \tag{18}$$

For the image retrieval task we use the mAP and nDCG metrics defined in Section 3. In Table 1 we present a comparison of our methods with the state of the art in the task of date estimation and with a visual similarity baseline, that ranks images using the euclidean distance between features extracted from a ResNet CNN [6] pre-trained on ImageNet. As we will discuss in the conclusions section, our approach is not directly comparable to Müller *et al.*'s [9]. On one hand, we are not trying to directly improve the prediction task, but a ranking one; this means that our model has to learn semantics from the images such that it learns a continuous ordination. On the other hand, the way we are estimating image's date is by using a ground truth dataset or support dataset. Nevertheless, we observe that our model isn't too far with respect to the state of the art prediction model. We consider this a good result since estimating the exact year of a image is not the main purpose of the smooth-nDCG loss function. Note that in Table 1 we mention a "cleaned test set". This is because we observed the presence of certain placeholder images indicating they are no longer available in Flickr. We computed the MAE for the weights given by Müller *et al.*[9] [3] for a proper comparison. The presence of this kind of images may have a low impact while evaluating a classification or regression, but they can have a higher impact on predicting with k-NN, since all the retrieved images will be placeholders as well independently of their associated years.

---

[3] This is an extended version of Müller et al. [9] made publibly available by the same authors at `https://github.com/TIB-Visual-Analytics/DEW-Model`

Table 1. Mean absolute error comparison of our baseline model on the test set of the Date Estimation in the Wild (DEW) dataset.

Baseline	MAE	mAP	nDCG
Müller et al. [9] GoogLeNet (regression)	7.5	-	-
Müller et al. [9] GoogLeNet (classification)	7.3	-	-
Müller et al. [9] ResNet50 (classification) (cleaned test set)	7.12	-	-
Visual similarity baseline	12.4	-	0.69
Smooth-nDCG 256 images support set	8.44	0.12	0.72
Smooth-nDCG 1M images support set	7.48	-	0.75
Smooth-nDCG 1M images support set (weighted kNN)	7.52	-	0.75

We present as well a result of predicting the year by weighted k-NN where the predicted year will be the sum of the k neighbours multiplied by their similarities to the query.

Since our model is trained for a retrieval task, we provide further qualitative examples of good and bad rankings. This exemplifies how our loss function is not mean to predict years, but finding good rankings according to query's year. Note we computed nDCG using as relevance the closeness in years between images' ground truth clipped at 10 years (see Eq. 15). Despite it may seem nDCG is so high and mAP so low; there are many subtle bias in this appearance. Since the ground truth can be interpreted as a continuous value (such as time) there are not many matches in any retrieval of N random images for a certain query regarding to query's year compared with how many negative samples are there. Then, mAP hardly punishes results in this situation. In the opposite way, the same thing happens to nDCG. Since nDCG has this 'continuous' or 'scale' sense, any retrieval will have some good samples that satisfies the relevance function (see Eq. 15), being the results way less punished if there aren't many exact responses in the retrieval. Briefly, we could say nDCG brings us a more balanced positive/negative ratio than mAP for randomly selected sets to rank (which is absolutely more realistic than forcing it to be balanced). Additionally, nDCG allows us to approach retrievals that are not based on obtaining a category but a wide number of them, or even retrieving non-categorical information such as date estimation.

Additionally, as we observe in Figure 7, we compute the average cosine similarity between clusters from different groups of years. This shows us how our output space is distributed with respect to the ground truth data. As we expected; clusters for closer years are the most similar between themselves. This is because ranking is nothing but taking the closer points to our query; by forcing our system to take those points that satisfies the relevance function (see Eq. 15), points in the embedding space will have to be organized such that distances are proportional to ground truth's distances.
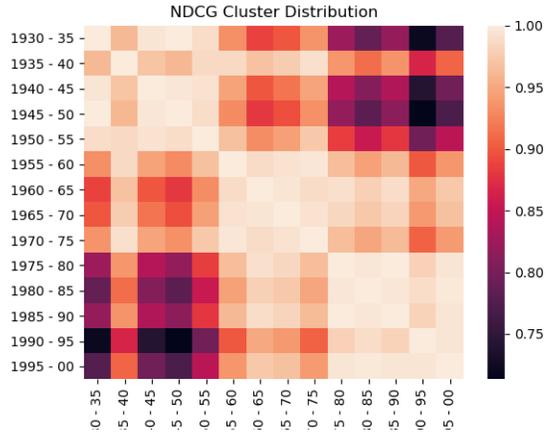
Fig. 7. Clusters average similarity for each 5 year interval. We observe a certain visible pattern in the diagonal; the closer the years the closer the clusters.

Finally, in Figure 9, we compare the top-5 retrieved images for a given query with our model and with the visual similarity baseline model, that ranks images using the distance between features extracted from a ResNet CNN [6] pre-trained on ImageNet. We appreciate that the visual similarity model retrieves images that are visually similar to the query but belong to a totally different date. On the contrary, our model performs better for the same query in terms of retrieving similar dates although for some of the top ranked images the objects layout is quite different to the query. Additionally, we present the results with a mixed model that re-ranks images with the mean ranking between both methods for each image in the test set; in this way we obtain a ranking with visually similar images from the same date as the query.

6   Conclusions

In this paper we have proposed a method for estimating the date of historic photographs. The main novelty of the presented work is that we do not formulate the solution to the problem as a prediction task, but as an information retrieval one. Hence metrics such as mean absolute error do not have to be considered as important as retrieval metrics. Even so, we managed to build an application-level model that can estimate the date from the output embedding space with a decent error according to previous baselines [9]. Since our prediction method is essentially a clustering or k-Nearest Neighbors method we are using a known support set, so predictions relies on already labeled data.

However, considering the output embedding space Figure 7 we conclude that smooth-nDCG function works pretty well for ranking large-scale image data with deeper ground truth structures; unlike smooth-AP [1], smooth-nDCG cares

Fig. 8. Qualitative results for our date estimation model. We show a set of images from the DEW test set and indicate their ground truth year annotation and distance to the query.

Fig. 9. Qualitative results for a retrieval system that combines our date estimation model and a visual similarity model. The top-left image is the query, in each row we show the top-5 retrieved images using a visual similarity model (top), our model (middle), and the combination of the two rankings (bottom).
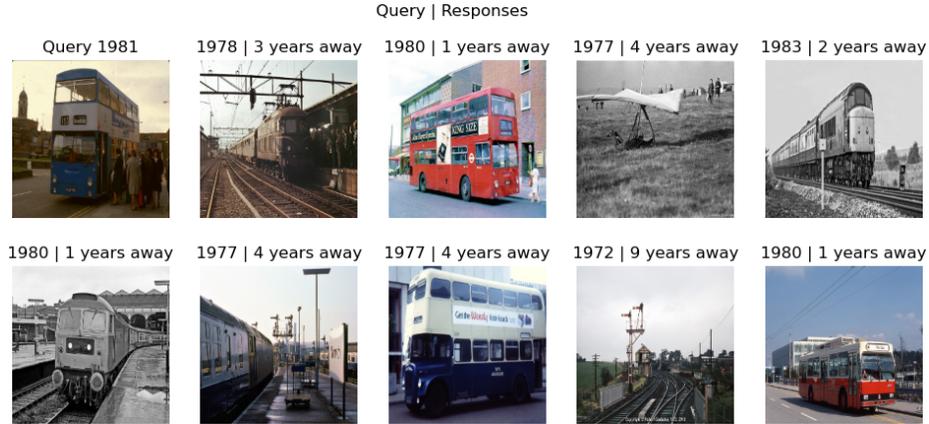
Query | Responses

Query 1981	1978 | 3 years away	1980 | 1 years away	1977 | 4 years away	1983 | 2 years away

1980 | 1 years away	1977 | 4 years away	1977 | 4 years away	1972 | 9 years away	1980 | 1 years away

Fig. 10. Possible biased retrieval due lack of balance in the training set.

about the whole sorted batch retrieved, not only how many good images are, but how are the bad ones distributed along the ranking. This allowed us to retrieval ordered information (such as years), or tree-structured information where different categories can be closer or further depending on how the tree is like.

In the case of date estimation in the wild [9] dataset, we found out some problematic patterns that can lead the retrieval task to certain bias. As there is not a clear balance between categories in the dataset, many classes such as *trains* or *vehicles* may be clustered together regardless the actual dates of the images Figure 10. However, it is not easy to say when a certain category is most likely to be linked to a certain period of time. For example, selfies are a category way more common nowadays than 20 years ago. Something similar may happen with certain categories that could be biasing the retrieval to a pure object classification task; however we found out many good clues, Figures 7, 8, that indicate our embedding for ranking images is working properly.

Despite further research is needed; smooth-nDCG usage for large-scale image retrieval is a promising novel approach with many practical applications. Since Brown *et al.*[1] smoothed us the path to minimize the average precision, we propose smooth-nDCG to model more complex retrieval problems where different labels should not be punished equally according to a certain criteria. As we commented in smooth-nDCG brings a new way of ranking images beyond categorical labels from a neural network approach; then we would like to consider this approach as an application example for larger information retrieval problems with a more complex structure.

Acknowledgment

ture Department of the Generalitat de Catalunya, and the CERCA Program / Generalitat de Catalunya.

References

1. Brown, A., Xie, W., Kalogeiton, V., Zisserman, A.: Smooth-ap: Smoothing the path towards large-scale image retrieval. In: Proceedings of the European Conference on Computer Vision. pp. 677–694 (2020)
2. Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., Darrell, T.: Decaf: A deep convolutional activation feature for generic visual recognition. In: Proceedings of the International Conference on Machine Learning. pp. 647–655 (2014)
3. Fernando, B., Muselet, D., Khan, R., Tuytelaars, T.: Color features for dating historical color images. In: Proceedings of the International Conference on Image Processing. pp. 2589–2593 (2014)
4. Frank, E., Hall, M.: A simple approach to ordinal classification. In: Proceedings of the European Conference on Machine Learning. pp. 145–156 (2001)
5. Ginosar, S., Rakelly, K., Sachs, S., Yin, B., Efros, A.A.: A century of portraits: A visual historical record of american high school yearbooks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 1–7 (2015)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
7. Lee, S., Maisonneuve, N., Crandall, D., Efros, A.A., Sivic, J.: Linking past to present: Discovering style in two centuries of architecture. In: Proceedings of the International Conference on Computational Photography (2015)
8. Martin, P., Doucet, A., Jurie, F.: Dating color images with ordinal classification. In: Proceedings of International Conference on Multimedia Retrieval. p. 447 (2014)
9. Müller, E., Springstein, M., Ewerth, R.: "when was this picture taken?"–image date estimation in the wild. In: Proceedings of the European Conference on Computer Vision. pp. 619–625 (2017)
10. Palermo, F., Hays, J., Efros, A.A.: Dating historical color images. In: Proceedings of the European Conference on Computer Vision. pp. 499–512 (2012)
11. Qin, T., Liu, T.Y., Li, H.: A general approximation framework for direct optimization of information retrieval measures. Information retrieval **13**(4), 375–397 (2010)
12. Salem, T., Workman, S., Zhai, M., Jacobs, N.: Analyzing human appearance as a cue for dating images. In: Proceedings of the IEEE Winter Conference on Applications of Computer Vision. pp. 1–8 (2016)
13. Schindler, G., Dellaert, F.: Probabilistic temporal inference on reconstructed 3d scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1410–1417 (2010)
14. Schindler, G., Dellaert, F., Kang, S.B.: Inferring temporal order of images from 3d structure. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1–7 (2007)
15. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2015)

16. Vittayakorn, S., Berg, A.C., Berg, T.L.: When was that made? In: Proceedings of the IEEE Winter Conference on Applications of Computer Vision. pp. 715–724 (2017)