

Photo Dating by Facial Age Aggregation

Jakub Paplham

Czech Technical University in Prague

paplhjak@fel.cvut.cz

Vojtěch Franc

Czech Technical University in Prague

xfrancv@fel.cvut.cz

Abstract

We introduce a novel method for *Photo Dating* which estimates the year a photograph was taken by leveraging information from the faces of people present in the image. To facilitate this research, we publicly release *CSFD-1.6M*¹, a new dataset containing over 1.6 million annotated faces, primarily from movie stills, with identity and birth year annotations. Uniquely, our dataset provides annotations for multiple individuals within a single image, enabling the study of multi-face information aggregation. We propose a probabilistic framework that formally combines visual evidence from modern face recognition and age estimation models, and career-based temporal priors to infer the photo capture year. Our experiments demonstrate that aggregating evidence from multiple faces consistently improves the performance and the approach significantly outperforms strong, scene-based baselines, particularly for images containing several identifiable individuals.

1. Introduction

Determining when a photograph was taken is a task of considerable practical interest in computer vision. It provides significant value for organizing vast, chronologically unsorted digital archives where embedded metadata is often missing or unreliable. In fields such as digital forensics and journalism, estimating the capture date of an image is a powerful tool for verifying authenticity, helping to construct accurate event timelines and to detect manipulated or synthetically generated media by exposing temporal inconsistencies. The emergence of commercial tools aimed at the general public, such as the MyHeritage PhotoDaterTM [17], highlights the broader industrial interest in this problem. Current automated methods for image dating analyze general visual content. Researchers have developed models that learn temporal signals from era-specific objects such as cars and clothing, trends in photographic color processing, or overall visual style [3, 9, 14, 16, 25, 26]. While effective

in certain scenarios, relying on such cues has limitations and overlooks what is arguably the most informative signal in many photographs: the human faces. The face provides a widely available and uniquely strong temporal signal, especially in contexts like genealogy and photo archiving, where analog scans often have absent or misleading metadata.


In this work, we introduce *facial-age-based image dating*, a novel approach to the task of estimating the capture year of a photograph. Our method leverages the rich temporal information contained within human faces, provided at least one person in the image can be identified and their birth year is known. The core mechanism is to infer the capture year by adding the visually estimated age of a person to their known birth year. We propose a formal probabilistic framework that aggregates this evidence from all identifiable individuals to produce a single, robust date estimate. The framework formally combines visual evidence from state-of-the-art face recognition and age estimation approaches with career-based temporal priors, if available. To evaluate our method's core strength; aggregating evidence from multiple people; existing benchmarks are unsuitable. We therefore construct CSFD-1.6M, a new large-scale dataset designed specifically for multi-person image dating. Derived from movie and TV show stills, it contains over 1.6 million annotated faces, each linked to a known identity, birth year, and an approximate image capture year.

Contributions:


- We formalize a new *face-based* method for image dating.
- We release a new large-scale dataset, and show its value for both our method, and general age estimation.
- We demonstrate that age-based evidence with career priors outperforms strong baselines, and that aggregating evidence from multiple faces improves accuracy.

The remainder of the paper is organized as follows. In Section 2 we review the related work. Section 3 presents our probabilistic model. Section 4 describes the novel dataset constructed to evaluate the approach. Section 5 reports experimental results, and Section 6 concludes with discussion.


¹Data and code are available at the [project github](#).




(a) Image x with detected faces $\mathbf{f} = (f_1, f_2, f_3)$ and the true identities. Still from *Joachim, Put It in the Machine* (1974), © Barrandov Studio. Image included for non-commercial research purposes under fair use.



(b) Recognized identities $\mathbf{i} = (i_1, i_2, i_3)$ in the detected faces $\mathbf{f}$. Knowing the birth year $b(i_j)$ for a recognized identity i_j allows us to compute the year posterior $p(y | f_j, i_j) = p_{\text{age}}(a = y - b(i_j) | f_j)$ using the posteriors $p_{\text{age}}(a | f_j)$ from an age estimation model.



(c) Joint career prior $p(y | \mathbf{i})$ for the recognized identities $\mathbf{i}$.



(d) Final posterior distribution $p(y | \mathbf{f})$ of the image capture year.

Figure 1. **Overview of the proposed dating framework.** Our method is demonstrated on an example image from CSFD-1.6M. (a) The input image with detected faces and their ground-truth identities noted. (b) For each detected face, we show the matched reference identity portrait and the corresponding year posterior, which is derived from the model’s age estimate. (c) The temporal information from the recognized identities is aggregated into a joint movie prior for the capture year. (d) The final capture year posterior is obtained by combining the joint prior from (c) with the individual year posteriors from (b) using the model described in Equation (1).

2. Related Work

Our research builds upon three distinct but related areas in computer vision: the general task of image capture date estimation, facial age estimation, and the use of contextual information in multi-person scenes. We review each in turn.

Image Capture Date Estimation The task of automatically estimating a photograph capture year has been approached from several angles, primarily focusing on non-person-specific visual cues. One line of work learns temporal signals from the entire image, using general cues such

as the physical properties of color film stocks [20]. More recent methods [3, 14, 16] leverage large-scale datasets to train deep networks to find specific temporal cues in general visual content, such as vehicles or buildings. A second line of work uses human appearance as a temporal signal, but does so by modeling period-specific style. Salem *et al.* [25] and Ginosar *et al.* [9] use high school yearbooks to learn trends in clothing and hairstyles. In these approaches, the identity of any given person is irrelevant; their face and attire serve only as an anonymous sample of a collective, era-defining appearance. Our work fundamentally differs. We propose that the specific, unique identity of a person, when

combined with their birth year, provides a much stronger and more direct temporal constraint than style alone.

Facial Age Estimation Estimating a person’s age from a facial image is a core component of our model. This is a well-studied problem, with progress driven by large-scale datasets like IMDB-WIKI [24]. The dominant approach is to train deep neural network for age regression or classification [6–8, 19, 21, 22, 24]. Despite significant progress, the recent NIST Face Analysis Technology Evaluation (FATE) reports that state-of-the-art models still exhibit considerable error, highlighting the difficulty of the task [11]. Our work does not aim to advance age estimation itself; instead, we leverage a state-of-the-art age predictor as a probabilistic component, providing a distribution of ages to reason over.

Reasoning in Multi-Person Images While most face analyses treat faces in isolation, a few studies have explored using context from multiple people to improve performance on a given task. Li *et al.* [12], for instance, use the co-occurrence of other individuals in a photo to improve identity recognition. This is especially valuable when dealing with historical archives, where image quality poses a significant challenge [13]. These works use multi-person context to solve for identity. In contrast, our work leverages the identities and their apparent ages as joint evidence to infer a single, shared latent variable: the image capture year.

3. Methodology

3.1. Probabilistic Model

Notation Let $x \in \mathcal{X}$ be an image taken in year $y \in \mathcal{Y}$ that contains a sequence $\mathbf{f} = (f_1, f_2, \dots, f_n)$ of n detected faces, where each $f_j \in \mathcal{F}$ is a crop extracted from x . Let $\mathcal{I} = \{\text{id}_1, \text{id}_2, \dots, \text{id}_m\}$ be a finite set of $m \geq n$ identities. Let $\mathbf{i} = (i_1, i_2, \dots, i_n)$ be an assignment of identities to the faces $\mathbf{f}$, where $i_j \in \mathcal{I}$ is the identity assigned to face f_j . As a person can appear at most once in an image, we enforce that all assigned identities are unique. Let $\mathcal{A}_n$ be the set of all valid identity assignments for the n faces, $\mathcal{A}_n = \{(i_1, \dots, i_n) \in \mathcal{I}^n \mid i_j \neq i_k \text{ for } j \neq k\}$. Lastly, let $b : \mathcal{I} \rightarrow \mathbb{Z}$ be a mapping of each identity $\text{id} \in \mathcal{I}$ to their birth year.

Model Our goal is to derive the probability distribution $p(y \mid x)$ of the year an image x was taken. We assume that the visual information relevant to the year y is entirely encapsulated by the sequence of n detected faces, $\mathbf{f} = (f_1, \dots, f_n)$. Therefore, we can write: $p(y \mid x) = p(y \mid \mathbf{f})$. The task then becomes to model and compute $p(y \mid \mathbf{f})$. We express the probability by marginalizing over all possible

identity assignments $\mathbf{i} \in \mathcal{A}_n$ and using the chain rule,

$$p(y \mid \mathbf{f}) = \sum_{\mathbf{i} \in \mathcal{A}_n} p(y, \mathbf{i} \mid \mathbf{f}) = \sum_{\mathbf{i} \in \mathcal{A}_n} p(y \mid \mathbf{f}, \mathbf{i}) p(\mathbf{i} \mid \mathbf{f}). \quad (1)$$

The summation over the assignment set $\mathcal{A}_n$ can be computationally intractable if the candidate identity set $\mathcal{I}$ is large. We discuss an approach to managing this in Section 5.

3.2. Simplifying Assumptions

Modelling $p(\mathbf{i} \mid \mathbf{f})$ For simplicity, we assume that given the identities $\mathbf{i} = (i_1, \dots, i_n)$, the appearances of the face crops $\mathbf{f} = (f_1, \dots, f_n)$ are conditionally independent, and the appearance of f_j only depends on its identity i_j :

$$p(\mathbf{f} \mid \mathbf{i}) = \prod_{j=1}^n p(f_j \mid i_j). \quad (2)$$

The assumption in Eq. (2) implies that factors such as shared lighting or camera parameters across faces in the same image are not explicitly modelled when the identities are known. This allows us to write $p(\mathbf{i} \mid \mathbf{f})$ for $\mathbf{i} \in \mathcal{A}_n$ as

$$\begin{aligned} p(\mathbf{i} \mid \mathbf{f}) &= \frac{p(\mathbf{f} \mid \mathbf{i}) p(\mathbf{i})}{p(\mathbf{f})} = \frac{\left(\prod_{j=1}^n p(f_j \mid i_j) \right) p(\mathbf{i})}{p(\mathbf{f})} \\ &= \frac{\left(\prod_{j=1}^n p(f_j \mid i_j) \right) p(\mathbf{i})}{\sum_{\mathbf{i}' \in \mathcal{A}_n} \left(\prod_{k=1}^n p(f_k \mid i'_k) \right) p(\mathbf{i}')}. \end{aligned} \quad (3)$$

Modelling $p(y \mid \mathbf{f}, \mathbf{i})$ We assume conditional independence of face appearances given the year y and identities $\mathbf{i}$, where each face f_j depends only on y and its identity i_j :

$$p(\mathbf{f} \mid y, \mathbf{i}) = \prod_{j=1}^n p(f_j \mid y, i_j). \quad (4)$$

The term $p(y \mid \mathbf{f}, \mathbf{i})$ can then be expressed as:

$$\begin{aligned} p(y \mid \mathbf{f}, \mathbf{i}) &= \frac{p(\mathbf{f} \mid y, \mathbf{i}) p(y \mid \mathbf{i})}{p(\mathbf{f} \mid \mathbf{i})} \\ &= \frac{\left(\prod_{j=1}^n p(f_j \mid y, i_j) \right) p(y \mid \mathbf{i})}{\sum_{y' \in \mathcal{Y}} \left(\prod_{k=1}^n p(f_k \mid y', i_k) \right) p(y' \mid \mathbf{i})} \\ &= \frac{\left(\prod_{j=1}^n \frac{p(y \mid f_j, i_j) p(f_j \mid i_j)}{p(y \mid i_j)} \right) p(y \mid \mathbf{i})}{\sum_{y' \in \mathcal{Y}} \left(\prod_{k=1}^n \frac{p(y' \mid f_k, i_k) p(f_k \mid i_k)}{p(y' \mid i_k)} \right) p(y' \mid \mathbf{i})}. \end{aligned} \quad (5)$$

Note A point of consideration is the term $p(f_j \mid i_j)$, which appears in the formulation of both $p(\mathbf{i} \mid \mathbf{f})$ and $p(y \mid \mathbf{f}, \mathbf{i})$. In a strict generative model where year y influences appearance per $p(f_j \mid y, i_j)$, the term $p(f_j \mid i_j)$

would ideally be derived as $\sum_{y' \in \mathcal{Y}} p(f_j | y', i_j) p(y' | i_j)$. Our current model treats $p(f_j | i_j)$ as a quantity that can be defined or estimated distinctly within each respective component. We acknowledge this pragmatic simplification.

3.3. Implementation Choices

The practical application of this model involves specific choices for its components, which we detail below.

Year Posterior via Age Estimation The year estimation posterior $p(y | f_j, i_j)$ is derived from an age estimation model that outputs a distribution $p_{\text{age}}(a | f_j)$ over age a given the face crop f_j . Knowing the birth year $b(i_j)$ for identity i_j , we set $p(y | f_j, i_j) = p_{\text{age}}(a = y - b(i_j) | f_j)$. To this end, we use a state-of-the-art age estimation model `cvut_002`, which consists of a ViT-B/16 backbone trained with cross-entropy on a proprietary dataset, that ranks competitively in the [NIST FATE](#) challenge [11].

Face Likelihood via Recognition Embeddings The face likelihood $p(f_j | i_j)$ is modeled in the ArcFace [5] embedding space. Let $\mathbf{e}(f_j) \in \mathbb{R}^{512}$ be the unit-norm embedding of face f_j . For an identity $\text{id}_k \in \mathcal{I}$, the likelihood is given by a von Mises-Fisher (vMF) distribution, $p(f_j | i_j = \text{id}_k) = C_D(\kappa) \exp(\kappa \boldsymbol{\mu}_k^T \mathbf{e}(f_j))$. The mean embedding $\boldsymbol{\mu}_k$ for each identity id_k is the normalized average of its N_k portrait embeddings, $\{\mathbf{e}_1, \dots, \mathbf{e}_{N_k}\}$, computed as $\boldsymbol{\mu}_k = \sum_{i=1}^{N_k} \mathbf{e}_i / \|\sum_{i=1}^{N_k} \mathbf{e}_i\|_2$. The shared concentration parameter κ is estimated from the same profile portraits. First, we generate a set of intra-identity cosine similarities using a leave-one-out scheme: for each identity with $N > 1$ portraits, we iteratively form a temporary prototype from $N - 1$ embeddings and compute its similarity to the held-out embedding. The mean of all such similarities across all identities gives us the mean resultant length $\bar{R}$. We then estimate κ using the approximation [2], $\kappa = \bar{R}(D - \bar{R}^2) / (1 - \bar{R}^2)$, where $D = 512$ is the dimensionality of the embedding space.

Identity and Temporal Priors For the identity assignment prior $p(\mathbf{i})$, we use an uninformative uniform distribution over the set of all possible assignments $\mathcal{A}_n$. We model the joint temporal prior as the product of priors for individual identities, $p(y | \mathbf{i}) \propto \prod_{j=1}^n p(y | i_j)$. The strength and source of this prior knowledge can vary significantly in real-world scenarios. To evaluate the performance and the sensitivity to the prior choice of our framework, we define five distinct formulations for the individual priors $p(y | i_j)$:

1. **Uniform Prior** p_U : A *uniform* distribution over the 100 years after birth; the minimal-information scenario.
2. **Decade Prior** p_D : Probability is uniform within a decade; probability mass per decade is proportional to image counts. A prior obtainable from external data.

3. **Movie Prior** p_M : An empirical distribution of *movies* per year in our dataset; a weak oracle prior for ablation.
4. **Image Prior** p_I : An empirical distribution of *images* per year in our dataset; a strong oracle prior for ablation.
5. **Combination Prior** p_C : A convex *combination* of the oracle and uniform priors, $p_C = \lambda p_I + (1 - \lambda) p_U$, used to analyze the model sensitivity to prior strength.

3.4. Handling Unknown Identities (Open-Set Setup)

In any real-world application, an image may contain unknown individuals or faces that cannot be confidently identified. Our framework can be extended to this open-set scenario by introducing a dedicated out-of-distribution (OOD) identity, id_{OOD} , into the identity set $\mathcal{I}$. This OOD identity is modeled with uninformative distributions: its face likelihood $p(f_j | i_j = \text{id}_{\text{OOD}})$ is uniform over the embedding space (a vMF distribution with $\kappa = 0$), and the temporal prior $p(y | i_j = \text{id}_{\text{OOD}})$ and the temporal posterior $p(y | f_j, i_j = \text{id}_{\text{OOD}})$ are uniform over all years $\mathcal{Y}$. This allows the model to assign unknown faces to id_{OOD} , and prevents them from corrupting the final year estimate. It is equivalent to ignoring low-confidence matches. The performance of this formulation is evaluated in Figure 7. For other experiments, we operate in a closed-set setting.

4. The CSFD-1.6M Dataset

To develop and evaluate our model, we collect a novel large-scale dataset for image dating, which we name CSFD-1.6M. The dataset’s primary characteristic is its focus on images containing multiple individuals within a scene, each with an associated identity and temporal label. This is necessary for studying the contextual relationships central to our model, a task for which existing benchmarks, such as IMDB-WIKI [24], are unsuitable as they are curated for single-face analysis, providing annotations for only one face per image.

Data Collection The dataset is constructed by scraping the Czecho-Slovak Movie Database (CSFD). We gathered high-resolution images, each linked to a specific production from which we extract a release year to serve as the temporal label. We acknowledge that the release is an approximation of the true capture date, a source of temporal uncertainty common to datasets in this domain. For each image, we also collect links to the unique profiles of all pictured public figures, providing their birth years and portraits.

Identity Annotation Pipeline We use an automated pipeline to assign detected faces to the known identities linked to an image. For an image with n faces detected by SCRFD-10GF [10] and k linked identities:

1. **Representation:** We compute ArcFace [5] embeddings for all faces. A prototype embedding for each identity

Statistic	Value
Total Scraped Images	1,493,357
Total Detected Faces	3,655,604
Faces Matched to an Identity	1,921,164
Final Faces (Matched & Dated)	1,653,871
Final Images	1,143,213
Total Annotatable Identities	75,952
Unique Identities in Final Set ⁴	46,523

Table 1. Statistics of the CSFD-1.6M creation pipeline.

is created by averaging embeddings from their profile portraits, following the standard practice [5].

2. **Constrained Assignment:** The assignment between the n faces and k identities is solved using the Hungarian algorithm on a cosine similarity matrix. We employ a standard formulation that pads the cost matrix to handle unbalanced cases $n \neq k$ and allows for partial matching². To ensure high-quality matches, costs for invalid pairings are set to infinity. A pair is invalid if its similarity is below 0.2 or if the true age differs from the visually estimated age by more than $20 + 0.25\text{age}_{\text{est}}$ years, allowing for greater error in age estimation for older subjects³.

We do not filter faces by attributes such as size or resolution, leaving such preprocessing to downstream applications.

Dataset Statistics The pipeline results in a dataset containing over 1.6 million faces, each with an assigned identity, birth year, and image year. Its key characteristic are multi-person scenes. Statistics of the dataset are detailed in Table 1, Figure 2 and the supplementary material.

Comparison to Existing Datasets Table 2 contrasts CSFD-1.6M with IMDB-WIKI, its closest large-scale analogue. Our dataset is not only substantially larger but, critically, consists of multi-person scenes required by our model. While no web-scraped dataset is free of label noise, our annotation pipeline leverages the explicit links between images and identity profiles provided by the database, aiming for higher-quality identity assignments. Other datasets designed for age estimation are discussed in Section 5.4.

Public Release We release the CSFD-1.6M annotations to support further research in this area. The release includes all annotations: image URLs, face bounding boxes, identity assignments, and birth years. We also provide the pre-

²We confirmed the robustness of this pipeline by manually annotating a stratified sample of 1000 faces from images containing 1 to 11 faces. The algorithm correctly found 997 matches and produced no incorrect matches.

³Removing the age constraint yielded only 1% additional matches on the full dataset, which we postulate are of lower quality.

⁴Many (29,429) of the 75,952 valid profiles have no gallery images.

Attribute	IMDB-WIKI	CSFD-1.6M
Size (Faces)	~523k	~1.6M
Identities	~20k	~46k
Multi-face	Single-face annotations	Core feature

Table 2. Comparison of CSFD-1.6M with IMDB-WIKI [24].

computed ArcFace embeddings and per-face age posterior distributions. This allows the research community to focus directly on modeling the contextual information shared between faces, as proposed here, without the significant overhead of data collection and feature extraction.

5. Experiments and Results

In this section we present the experimental evaluation of our proposed model. We first describe the baseline and oracle models used for comparison, then detail the evaluation protocol, and finally present and discuss the results.

5.1. Compared Models

We evaluate four distinct models designed to ablate different components of our probabilistic framework. This allows us to quantify the impact of face identification errors, the assignment strategy, and the inclusion of temporal priors.

Oracle Model This model serves as an *oracle* for our task. It assumes perfect knowledge of the identities in the image. For a given image with faces $\mathbf{f}$, we use the single ground-truth identity assignment $\mathbf{i}_{\text{GT}}$. This model isolates the performance of the age estimation and temporal priors, completely removing any error from the face identification.

Full Probabilistic Model (Full) This model is a practical implementation of the framework proposed in Section 3. It approximates the posterior $p(y | \mathbf{f})$ by marginalizing over a tractable subset of all identity assignments. This subset is constructed from candidate pools that cover 99% of the probability mass for each face, capturing the most plausible joint assignments yet remaining computationally feasible.

Top-1 Assignment Model (Top-1) This model is a computationally efficient approximation of the full model. Instead of marginalizing, it determines the single *joint* identity assignment $\mathbf{i}_{\text{Top-1}}$ that has the highest posterior probability, $\mathbf{i}_{\text{Top-1}} = \arg \max_{\mathbf{i} \in \mathcal{A}_n} p(\mathbf{i} | \mathbf{f})$. The final year distribution is then calculated conditioned entirely on this single assignment. This differs from a simple greedy approach as it finds the most likely valid assignment for the group as a whole.

Naive Baseline (Naive) This model serves as a simple *baseline* to demonstrate the value of the probabilistic framework. It performs a greedy, *independent* identification for




Figure 2. Overview of the CSFD-1.6M dataset statistics.

each face, assigning it to the single best-matching identity regardless of the other face assignments. It then calculates the final year distribution by taking the product of the individual year posteriors from the age estimator, $p(y | \mathbf{f}) \propto \prod_{j=1}^n p(y | f_j, i_j)$, notably without the use of any temporal prior $p(y | \mathbf{i})$. This model reflects a straightforward but theoretically less grounded approach to the problem.

Scene-Based Baseline (Scene) To benchmark the face-based approach against the established state-of-the-art, we implement the direct prediction paradigm [3, 14, 16], modeling $p(y | x)$ from the full scene context. This paradigm has been shown to be more effective than alternative approaches like image retrieval [14]. Our Scene baseline consists of a ViT-B/16⁵ backbone that takes an entire image as input and outputs a probability distribution over the years.⁶ Unlike the face-based methods, which leverage pre-trained models, this baseline is trained directly on CSFD-1.6M, as detailed in the protocol below. This represents a best-case scenario for the *Scene* model and a strong baseline.

5.2. Evaluation Protocol

We evaluate the models using the *Mean Absolute Error* (MAE) between the predicted year and the ground-truth year. For all evaluated models the final year prediction is derived as the median of the posterior distribution $p(y | \cdot)$.

Dataset Splits For the *Scene* baseline, we perform 5-fold cross-validation. For each fold, we use a 60/20/20 train/validation/test split. The splits are created to be *movie-disjoint*⁷, ensuring that all images from the same movie belong to only one split. We report the mean MAE across the 5 folds for the *Scene* baseline. In contrast, the face-based models are trained on external datasets. Therefore, they are evaluated on the entirety of the CSFD-1.6M dataset.

⁵Identical architecture to the age estimation model.

⁶We did not benchmark the older method [20] as it is not competitive.

⁷An identity-disjoint split was not feasible, as the high interconnectedness of actors clustered the vast majority of the dataset into a single split.

Scene Baseline Training The *Scene* model is trained for 200 epochs using the Adam optimizer ($\text{lr} = 10^{-4}$), employing early stopping based on the validation set MAE with a patience of 10 epochs. The input images are resized so their shorter side is 384 pixels, from which a 384x384 crop is taken. For training, we apply standard data augmentations.

5.3. Photo Dating Results

This section presents the quantitative results of our experiments. Our central finding is that the face-based models consistently outperform the strong, in-domain *Scene* baseline for images containing more than one person, validating the hypothesis that aggregating facial information (when available) is a superior strategy to general scene analysis.

Ablation of Information Sources To understand the relative importance of the two primary signals in our model; the age-based posteriors $p(y | i_j, f_j)$ and the temporal prior $p(y | \mathbf{i})$; we conduct an ablation study shown in Figure 4. We evaluate the performance using only the temporal prior and only the aggregated age posteriors (equivalent to our *Naive* baseline). The results show that both components are individually informative. Crucially, the full model, which combines both signals, consistently outperforms using either source in isolation. For real-world applications where a specific temporal prior may be unknown or difficult to obtain, our results validate that the age-based posterior alone still serves as a strong, standalone temporal signal. The full framework provides a principled way to further improve this baseline estimate whenever a prior becomes available.

Impact of Temporal Priors Figure 3 details the effect of different temporal priors on the *Full* model. The results demonstrate the value of using an informative prior. The realistic *Decade Prior* (p_D) provides a significant improvement over both the *Naive* model (which uses no temporal prior) and the uninformative *Uniform Prior* (p_U). Notably, while the *Scene* baseline is superior for single-face images, the *facial-age-based* approach surpasses it for all images containing two or more faces. This occurs despite the *Scene*

Benchmark	Pre-training Dataset		
	ImageNet	IMDB-WIKI	CSFD-1.6M
AgeDB	7.05 ± 0.29	6.34 ± 0.25	5.25 ± 0.21
AFAD	3.19 ± 0.04	3.10 ± 0.03	3.04 ± 0.03
MORPH	2.98 ± 0.05	2.88 ± 0.07	2.76 ± 0.05
UTKFace	4.84 ± 0.08	4.64 ± 0.06	4.08 ± 0.03
CLAP2016	5.87	4.89	3.52

Table 3. **Age estimation results.** MAE $\downarrow$ ($\pm$ std) on five benchmarks after pre-training a ResNet-101 model on different datasets.

model being trained directly on this dataset.

Sensitivity to Prior Strength We analyze the model sensitivity to prior information in Figure 5. The plot shows the MAE for the *Full* model using the combination prior p_C , as λ interpolates from a purely uniform prior ($\lambda = 0$) to the oracle image prior ($\lambda = 1$). At $\lambda = 0$, performance matches the *Naive* baseline. As λ increases, performance steadily improves, demonstrating that the framework effectively leverages prior information. Even a weak prior signal (e.g., $\lambda = 0.1$) provides a substantial performance gain.

Ablation of Marginalization To understand the contribution of each part of our framework we ablate the components in Figure 6 using the realistic decade prior (p_D) for all face-based models. The *Oracle* model establishes a clear performance bound. The key finding is that our *Full* model, which marginalizes over possible identity assignments, offers only a marginal improvement over the much simpler *Top-1* approximation. This suggests that the assignment posterior is sharply peaked, likely due to the high quality of the ArcFace embeddings, making the *Top-1* approach a highly effective and computationally tractable alternative.

Analysis of Prediction Bias We analyze the systematic bias of our predictions in Figure 8. The plot shows the distribution of the prediction error (Predicted Year - True Year). The models exhibit a consistent negative bias, predicting capture dates that are on average a few years earlier than the release year. This offset is an expected artifact of our data source and reflects the inherent lag between the film production (photos are taken) and its eventual release.

Robustness to Unknown Identities Real-world images often contain a mix of known and unknown individuals. We evaluate this open-set setup in Figure 7, which shows how the MAE is affected by the number of *known* and *unknown* faces. The error consistently decreases as the number of known identities increases. Conversely, while the error increases with more unknown identities, this degradation is gradual. This confirms that the framework is effective at preventing unknown faces from corrupting the estimate.




Figure 3. **Performance comparison of temporal priors.** MAE for the *Full* model using different priors, and MAE of the *Naive* and *Scene* baselines. The results highlight the benefit of informative priors. The realistic *Decade Prior* (p_D) consistently outperforms the *Naive* baseline and the uninformative *Uniform Prior* (p_U). The oracle priors (p_I , p_M , p_C), which use statistics from the test set, are included to illustrate an upper bound on performance.




Figure 4. **Ablation of information sources.** MAE of predictions made using only a temporal prior (p_D or p_I), only the age posterior (p_{age} , equivalent to the *Naive* model), and the full model combining both (p_D & p_{age} or p_I & p_{age}). The results demonstrate that both the temporal prior and the age posterior are informative signals. Combining them using the proposed aggregation consistently yields lower error than using either source in isolation.

5.4. Age Estimation Performance

To validate CSFD-1.6M as a high-quality resource, we evaluate its effectiveness as a pre-training corpus for age estimation. Following the protocol from [22], we compare a ResNet-101 model initialized with ImageNet weights [4] to models further pre-trained on either IMDB-WIKI [24] or CSFD-1.6M. Each model is then fine-tuned and evaluated on five datasets: AgeDB [15], AFAD [18], MORPH [23], UTKFace [27], and CLAP2016 [1]. The resulting MAE is reported in Table 3. The results demonstrate that pre-training on CSFD-1.6M consistently and substantially out-




Figure 5. **Impact of prior strength (λ) on model performance.** MAE of the *Full* model using the combination prior p_C , as λ interpolates between a purely uniform prior ($\lambda = 0$) and the full oracle image prior ($\lambda = 1$). The results show that even weak prior information yields substantial performance gain over the baseline.




Figure 6. **Ablation of model components using a realistic prior.** MAE of the *Oracle*, *Full*, and *Top-1* models with a decade prior (p_D), compared against the *Naive* and *Scene* baselines. The *Full* and *Top-1* models outperform the strong *Scene* baseline, demonstrating the power of aggregating evidence from multiple faces even when restricted to the weak realistically obtainable prior p_D .

performs pre-training on ImageNet and IMDB-WIKI across all five evaluation benchmarks. This validates CSFD-1.6M not only for the proposed dating task but also as a valuable resource for the broader task of facial age estimation.

6. Conclusion

In this work, we introduced a novel method of *Photo Dating by Facial Age Aggregation*, demonstrating that faces are a powerful source of temporal information. We presented a probabilistic model to formally combine visual evidence from the apparent age and career-based temporal priors. A key contribution is the creation and public release of CSFD-1.6M, a large-scale dataset of over 1.6 million annotated faces. Our experiments validate our approach. We have




Figure 7. **Open-set evaluation.** MAE of the *Full* model using prior p_D stratified on the number of *known* and *unknown* identities.




Figure 8. **Analysis of prediction bias.** Bias (Predicted - True) for the *Full* model using the p_D prior, and the *Naive* model. For both models we show a shaded band representing the central 90% of prediction errors. Both models exhibit a consistent negative bias.

shown that aggregating evidence from multiple faces consistently improves dating accuracy, outperforming a strong model based on the state-of-the-art [3, 14, 16]. Our work provides a foundation for future research into face-based image dating and temporal analysis of media.

Limitations & Ethical Considerations The probabilistic model in Section 3 relies on simplifying assumptions for tractability. The dataset is limited by the inherent temporal noise from using movie release years as capture dates. We release our contributions to foster research in temporal media analysis and encourage their responsible application, however, we recognize the associated ethical duties. The use of public data, the demographic bias and the potential for dual-use are significant concerns. To mitigate these and support reproducibility, we don't distribute the images, but only the annotations, pre-computed features and URLs.

Acknowledgements

This research was supported by the Grant Agency of the Czech Technical University in Prague, grant No. SGS23/176/OHK3/3T/13, by the Grant agency of the Ministry of Interior Czech Republic, project FACIS grant. No. VJ02010041, and by the National Recovery Plan funded project MPO 60273/24/21300/21000 CEDMO 2.0 NPO.

References

- [1] E. Agustsson, R. Timofte, S. Escalera, X. Baro, I. Guyon, and R. Rothe. Apparent and real age estimation in still images with deep residual regressors on appa-real database. In *12th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, 2017. IEEE, 2017. 7
- [2] Arindam Banerjee, Inderjit S. Dhillon, Joydeep Ghosh, and Suvrit Sra. Clustering on the unit hypersphere using von mises-fisher distributions. *Journal of Machine Learning Research*, 6(46):1345–1382, 2005. 4
- [3] Alexandra Barancová, Melvin Wevers, and Nanne van Noord. Blind dates: Examining the expression of temporality in historical photographs. *Proceedings http://ceur-ws.org ISSN*, 1613:0073, 2023. 1, 2, 6, 8
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 7
- [5] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. 4, 5
- [6] Raúl Díaz and Amit Marathe. Soft labels for ordinal regression. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4733–4742, 2019. 3
- [7] Bin-Bin Gao, Chao Xing, Chen-Wei Xie, Jianxin Wu, and Xin Geng. Deep label distribution learning with label ambiguity. *IEEE Transactions on Image Processing*, 26(6):2825–2838, 2017.
- [8] Bin-Bin Gao, Hong-Yu Zhou, Jianxin Wu, and Xin Geng. Age estimation using expectation of label distribution learning. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 712–718. International Joint Conferences on Artificial Intelligence Organization, 2018. 3
- [9] S. Ginosar, K. Rakelly, S. M. Sachs, B. Yin, C. Lee, P. Krähenbühl, and A. A. Efros. A century of portraits: A visual historical record of american high school yearbooks. *IEEE Transactions on Computational Imaging*, 3(3):421, 2017. 1, 2
- [10] Jia Guo, Jiankang Deng, Alexandros Lattas, and Stefanos Zafeiriou. Sample and computation redistribution for efficient face detection. In *International Conference on Learning Representations*, 2022. 4
- [11] Kayee Hanaoka, Mei Ngan, Joyce Yang, George W. Quinn, Austin Hom, and Patrick Grother. Face analysis technology evaluation: Age estimation and verification. Technical Report NISTIR 8525, National Institute of Standards and Technology, Gaithersburg, MD, 2024. 3, 4
- [12] Haoxiang Li, Jonathan Brandt, Zhe Lin, Xiaohui Shen, and Gang Hua. A multi-level contextual model for person recognition in photo albums. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1297–1305, 2016. 3
- [13] Vikram Mohanty, David Thames, Sneha Mehta, and Kurt Luther. Photo sleuth: Identifying historical portraits with face recognition and crowdsourced human expertise. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 10(4): 1–36, 2020. 3
- [14] Adrià Molina, Pau Riba, Lluís Gomez, Oriol Ramos-Terrades, and Josep Lladós. Date estimation in the wild of scanned historical photos: An image retrieval approach. In *Document Analysis and Recognition – ICDAR 2021*, pages 306–320, Cham, 2021. Springer International Publishing. 1, 2, 6, 8
- [15] Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou. Agedb: the first manually collected, in-the-wild age database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshop*, 2017. 7
- [16] Eric Müller, Matthias Springstein, and Ralph Ewerth. “when was this picture taken?” – image date estimation in the wild. In *Advances in Information Retrieval*, pages 619–625, Cham, 2017. Springer International Publishing. 1, 2, 6, 8
- [17] MyHeritage. What is photodater™ and how does it work?, 2023. Accessed: July 2, 2025. 1
- [18] Zhenxing Niu, Mo Zhou, Le Wang, Xinbo Gao, and Gang Hua. Ordinal regression with multiple output cnn for age estimation. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4920–4928, 2016. 7
- [19] Zhenxing Niu, Mo Zhou, Le Wang, Xinbo Gao, and Gang Hua. Ordinal regression with multiple output cnn for age estimation. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4920–4928, 2016. 3
- [20] Frank Palermo, James Hays, and Alexei A. Efros. Dating historical color images. In *Computer Vision – ECCV 2012*, pages 499–512, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. 2, 6
- [21] Hongyu Pan, Hu Han, Shiguang Shan, and Xilin Chen. Mean-variance loss for deep age estimation from a face. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5285–5294, 2018. 3
- [22] Jakub Paplham and Vojtech Franc. A call to reflect on evaluation practices for age estimation: Comparative analysis of the state-of-the-art and a unified benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1196–1205, 2024. 3, 7
- [23] K. Ricanek and T. Tesafaye. Morph: a longitudinal image database of normal adult age-progression. In *7th International Conference on Automatic Face and Gesture Recognition (FGR06)*, pages 341–345, 2006. 7

- [24] Rasmus Rothe, Radu Timofte, and Luc Van Gool. Deep expectation of real and apparent age from a single image without facial landmarks. *International Journal of Computer Vision*, 126(2-4):144–157, 2018. [3](#), [4](#), [5](#), [7](#)
- [25] Tawfiq Salem, Scott Workman, Menghua Zhai, and Nathan Jacobs. Analyzing human appearance as a cue for dating images. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–8, 2016. [1](#), [2](#)
- [26] Lorenzo Stacchio, Alessia Angeli, Giuseppe Lisanti, Gustavo Marfia, et al. Searching for cultural relationships through deep learning models. In *CEUR WORKSHOP PROCEEDINGS*, pages 1–11. CEUR-WS, 2022. [1](#)
- [27] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017. [7](#)