

Illumination-Aware Age Progression

Ira Kemelmacher-Shlizerman¹Supasorn Suwajanakorn¹Steven M. Seitz^{1,2}¹University of Washington²Google Inc.

Figure 1. Given a single input photo of a child (far left) our method renders an image at any future age range between 1 and 80. Note the change in shape (e.g., nose gets longer, eyes narrow) and texture, while keeping the identity (and milk mustache!) of the input person.

Abstract

We present an approach that takes a single photograph of a child as input and automatically produces a series of age-progressed outputs between 1 and 80 years of age, accounting for pose, expression, and illumination. Leveraging thousands of photos of children and adults at many ages from the Internet, we first show how to compute average image subspaces that are pixel-to-pixel aligned and model variable lighting. These averages depict a prototype man and woman aging from 0 to 80, under any desired illumination, and capture the differences in shape and texture between ages. Applying these differences to a new photo yields an age progressed result. Contributions include relightable age subspaces, a novel technique for subspace-to-subspace alignment, and the most extensive evaluation of age progression techniques in the literature¹.

1. Introduction

What will a child look like in 20 years? Age progression, which seeks to “age” photographs of faces, is one of the most intriguing of digital image processing operations. It is also one of the most challenging for a variety of reasons. First, the aging process is non-deterministic, depending on environmental as well as genetic factors that may not be evident in the input photos. Second, facial appearance and recognizability is strongly influenced by hair style, glasses, expression, and lighting, which is variable and unpredictable. Finally, there is relatively little data available

from which to build effective models, as existing age analysis databases are relatively small, low resolution, and/or limited in age range.

Nevertheless, age progression techniques have enjoyed significant success in helping to solve missing children cases, where subjects have been recognized many years later based on age progressed images. Described as “part art, part science, and a little intuition” [30], these images are produced by forensic artists who combine a background in art, physical anthropology, and expertise with image editing software to simulate the appearance of a person later in life [13]. Aging photos of very young children from a single photo is considered the most difficult case of all, where age progression beyond a few years is considered impractical [25]. We focus specifically on this very challenging case.

Our approach takes a single photo as input and automatically produces a series of age-progressed outputs between 1 and 80 years of age. Figure 1 shows an example result. Our approach has three primary contributions. First, we present the first fully-automated approach for age progression that operates “in the wild”, i.e., without strong constraints on lighting, expression, or pose. Second, we present some of the first compelling (and most extensive) results for aging babies to adults. And third, we introduce a novel illumination-aware age progression technique, leveraging illumination modeling results [1, 31], that properly account for scene illumination and correct surface shading without reconstructing 3D models or light source directions.

We build on prior work on age progression, notably, the seminal work of Burt and Perrett [6], who created convincing average male faces for several ages (in the range of 20-

¹<http://grail.cs.washington.edu/aging/>

54) by aligning and averaging photos together. A new query photo was then age progressed by adding to it the difference in shape and texture between the average of the desired target age, and the average for the age corresponding to the query. Their approach required manual alignment. Subsequent aging work in the computer vision literature introduced more automation, often using Active Appearance Models [17] or detecting fiducials [32]. Additional improvements included texture modeling for wrinkles [37] and person-specific models [35, 29]. More details can be found in these excellent survey papers [9, 33]. Early face analysis methods proposed to synthesize new faces using image based models, e.g., [7], but did not focus on aging and uncalibrated conditions. There are now several commercial programs that will age photos taken with a webcam or mobile phone. Typically, however, these programs operate effectively only for photos of adults or older children; [23] requires a minimum age of 18, ageme.com lists 7 as the low range, the popular AgingBooth iphone app suggests a minimum age of 15. Furthermore, both commercial offerings and state-of-the-art methods from the research literature still require frontal, simply-lit faces, with neutral expression [9].

There is a body of work on automatic age estimation, e.g., [26, 11, 18]. They, however, did not pursue age progression or other synthesis applications.

Our results set a new bar for age-progression research, demonstrated by a comprehensive evaluation of prior art (the first of its kind in the age progression literature), and an extensive comparison to “ground truth,” via large scale user studies on Amazon Mechanical Turk, as described in Section 4. The key components that make this advance possible are first, a new database consisting of thousands of photos of people spanning age (0 to 100), variable lighting, and variable pose and expression (Section 2.1). Second, *relightable* average images that capture changes in facial appearance and shape across ages, in an illumination invariant manner (Section 2.2). And third, a novel technique for aligning illumination subspaces that enables capturing and synthesizing age transformations (Section 3).

2. Building an Aging Basis

As we age, our faces undergo changes in shape and appearance. The transformation from child to adult is dominated by craniofacial growth, in which the forehead slopes backward, the head expands, and the lower portion of the face extends downward [8]. Changes in later years are dominated by growth of the nose, narrowing of the eyes, the formation of wrinkles and other textural changes.

One of the most compelling ways to model and view these changes across people is by creating a sequence of composite faces, where each composite is the average of several faces of the same gender and age. This idea dates back more than two centuries; Galton [10] generated aver-

age images by taking several exposures of portraits on the same photographic plate. Bensen and Perrett [2] showed that dramatically better composites can be obtained by first aligning facial features (208 fiducials) and warping the images to a reference prior to averaging. Producing composites for aging studies is hampered, however, by the lack of good photographic data for young children, as existing databases are relatively small, low resolution, and limited in age range [9]. In the remainder of this section, we introduce an approach for creating and analyzing a large dataset of human faces across ages, based on thousands of photos from the Internet.

2.1. Data collection

To analyze aging effects we created a large dataset of people at different ages, using Google image search queries like “Age 25”, “1st grade portrait,” and so forth. We additionally drew from science competitions, soccer teams, beauty contests, and other websites that included age/grade information. The resulting databases spans 0 to 100, pooled into 14 age groups (we call them *clusters*), separated by gender. The clusters correspond to ages 0, 1, 2-3, 4-6, 7-9, 10-12, 13-15, 16-24, 25-34, 35-44, 45-56, 57-67, 68-80 and 81-100. The total number of photos in the dataset is 40K and each cluster includes, on average, 1500 photos of different people in the same age range. This database captures people “in the wild” and spans a large range of ages.

2.2. Aligned, re-lightable averages

To obtain dense correspondence between the photos in each cluster, we use the “collection flow” method [15], which enables accurate dense correspondence across images with large illumination variation. The input to the collection flow method are aligned and warped to frontal photos for which we use the pipeline of [16]. Figure 2 shows the average image for each age, and the average of flow-warped photos using collection flow. Note how much sharper the flow-aligned averages look. While these aligned averages can appear remarkably lifelike, the lighting is dull and unrealistic, as it is averaged over all images in the collection. We instead produce *relightable* average images, which may be re-illuminated from any direction with realistic shading effects. We propose to *match* the lighting of any new input image I by first pose-aligning the image [16], and projecting that image onto every age subspace. Specifically, for an age cluster j with flow-aligned average A_j , we compute a rank-4 basis via singular value decomposition on the *flow-aligned* images, i.e., $M_j = U_j D_j V_j^T$ where M_j is $f \times p$ the matrix representation of the cluster’s flow-aligned photos (f is the number of photos and p number of pixels in each photo). As described in [15], this rank-4 approximation retains the lighting and shading of the input photos, but *neutralizes* the changes due to identity and facial expres-

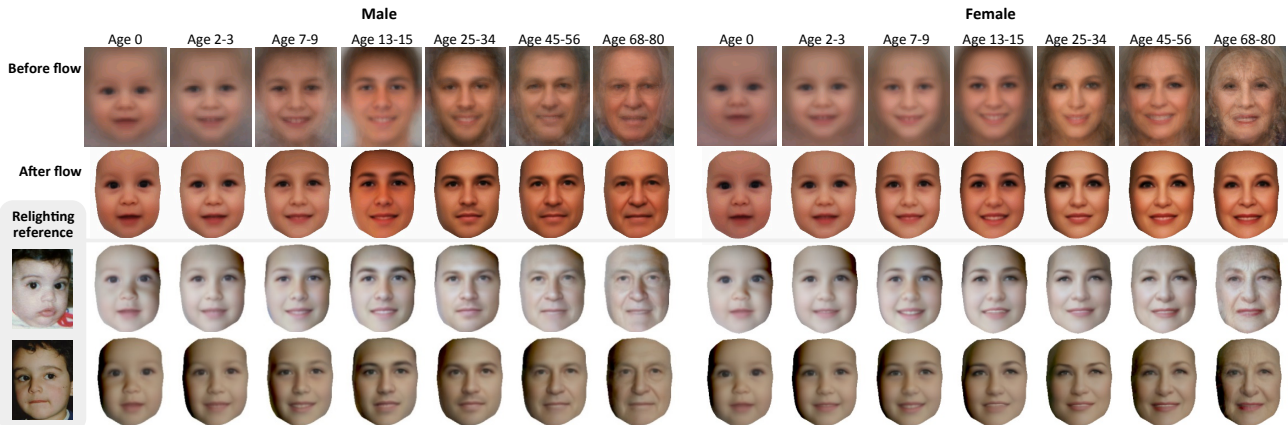


Figure 2. Average images of people at different ages. Each image represents an average of about 1500 individuals. Results in the top row are aligned only to place the eyes, nose, and mouth in rough correspondence. The second row shows averages after pixel-to-pixel alignment. These are much sharper, but the tone is variable, the lighting is unnatural, and subtle shape differences (e.g., wrinkles) are averaged out (to see it zoom-in to the last column). The bottom two rows show *re-lit* averages, matched to two reference frames (far left) with opposite lighting directions. The re-lit results have proper shading, are tone-matched to allow easier comparison across ages, and reveal 3D shape changes (note the nose and forehead).

sion, producing a set of images in nearly perfect alignment with a common, average face pose. Next solving

$$\min_{\alpha} \|I - \alpha V_j^T\|^2 \quad (1)$$

for the coefficients α yields a *re-lit* average that matches the illumination of I :

$$A_j^I = \alpha V_j^T \quad (2)$$

(V_j is truncated to $\text{rank}=4$). Figure 2 (rows 3-4) shows this capability. Two key advantages of relighting are that 1) it generates a more realistic set of average images, bringing out fine details that are only visible with proper shading, and 2) we can *align* the lighting across the set of averages, to enable comparing changes at different ages. We use this relighting capability to estimate flow across clusters as described below.

2.3. Illumination Subspace Flow

We have so far focused on aligning photos within each age cluster. Next, we show how to estimate flow across age clusters, to measure face shape changes over time. Each cluster has many photos under different illumination conditions and thus captures an *illumination subspace*, representing how an average person at a particular age appears under all illuminations [1]. A key contribution of our paper is how to align two such illumination subspaces V_i and V_j .

We seek the (single) optical flow field that aligns V_i and V_j . Our insight is to use relighting for flow estimation. As shown in Fig. 2 (last column), relighting brings out 3-dimensional shape differences that are otherwise invisible when averaging many photos. We therefore propose an optical flow method that optimizes over many different light-

ing conditions. The challenge here is twofold: 1) each illumination subspace represents a continuum of different images, and 2) their coefficient space is not aligned, i.e., any physical lighting direction may map to different lighting coefficients in each illumination subspace.

We introduce a solution that can be easily implemented within the traditional two-view optical flow framework. Let K be the number of database images in the union of clusters i and j . For each image I_k in this union, we project it to each of the two illumination subspaces resulting in an average image A_i^k and A_j^k . The resulting set of images $\{A_i^k\}_{k=1}^K$ can be represented as a single K -channel image \mathbf{A}_i , and similarly for \mathbf{A}_j . Unlike the original illumination subspaces V_i and V_j , these two multi-channel images are *illumination-aligned*; the k^{th} channel of \mathbf{A}_i and \mathbf{A}_j have the same lighting. Hence, our method can work with any optical flow algorithm that supports multiple channel images (including more complex methods like SIFT flow [22, 12]) to compute the **lighting-aware flow** field.

When K is large, a smaller representative set of images can be chosen using either discrete sampling, clustering, or dimensionality reduction techniques. We leveraged the fact that the illumination subspaces are low-dimensional [14] (V_i is 4D) and computed an orthogonal 24D basis (two 4D clusters times 3 color channels) for the K images using PCA. Each basis vector (mean + principle vector) was weighted in proportion to its principle value (we modified [21] to support weighted multi-channel images).

2.4. Age Transformations

To align all age clusters, we compute subspace flow between each pair of successive age clusters i and $i + 1$.

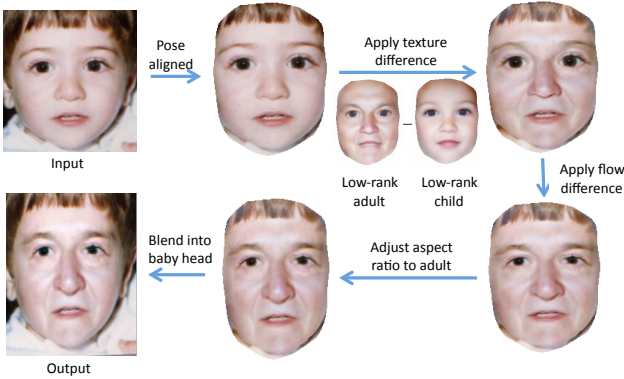


Figure 3. Steps of illumination-aware age progression.

Longer range flows between more disparate ages i and j are obtained by concatenation of the flow fields between i and $i + 1$, $i + 1$ and $i + 2$, ..., $j - 1$ and j . This concatenation approach gives more reliable flow fields than direct, pairwise flow computation between i and j . These flows enable estimating differences in texture and shape between the different age groups, as we describe in the next section.

3. Illumination-Aware Age Progression

Given an input photo of a 2 year old, we can render her at age 60 by computing the difference in flow and texture between the cluster of ages 2-3 (source) and cluster of ages 57-67 (target) and applying it to the input photo. This task is challenging for images “in the wild,” as it requires taking into account variations in lighting, pose, and identity. Illumination and shading are inherently 3D effects that depend upon light source direction and surface shape, e.g., as the nose becomes more angular, its shading should change in a manner that depends on light source direction. We show, however, that it is possible to utilize our rank-4 relightable aging basis to work entirely in the 2D domain, without reconstructing 3D models.

To age progress a face photo we perform the following steps, as illustrated in Figure 3.

Pose correction: the input face is warped to approximately frontal pose using the alignment pipeline of [16] (step 1 in the figure). Denote the aligned photo I .

Texture age progress: Relight the source and target age cluster averages to match the lighting of I as described in Section 2.2, yielding A_s^I and A_t^I . Compute flow $F_{\text{source-input}}$ between A_s^I and I and warp A_s^I to the input image coordinate frame, and similarly for $F_{\text{target-input}}$. This yields a pair of illumination matched projections, J_s and J_t both warped to input. The texture difference $J_t - J_s$ is added to the input image I .

Flow age progress: Apply flow from source cluster to target cluster $F_{\text{target-source}}$ mapped to the input image, i.e., apply $F_{\text{input-target}} \circ F_{\text{target-source}}$ to the texture-modified image

$I + J_t - J_s$. For efficiency, we precompute bidirectional flows from each age cluster to every other age cluster.

Aspect ratio progress: Apply change in aspect ratio, to account for variation in head shape over time. Per-cluster aspect ratios were computed as the ratio of distance between the left and right eye to the distance between the eyes and mouth, averaged over the fiducial point locations of images in each of the clusters.

We also allow for differences in skin tone (albedo) by computing a separate rank-4 subspace and projection for each color channel.

4. Experiments

We now describe implementation details, results, and evaluation based on a large scale user study.

Implementation details For all flow computations, we modified Ce Liu’s [21] implementation (based on Brox et al. [4] and Bruhn et al. [5]) to work with weighted multi-channel photos. We used the following parameters $\alpha = 0.005$, ratio= 0.85, minWidth= 20, nOuterFPIterations= 10, nInnerFPIterations= 1, nSORIterations= 20. We used random SVD [34] for fast low rank computations. Processing the photo database required 30 minutes (on 14 compute nodes) per age cluster of 300 photos, including flow, averages, and subspace computation. Given the precomputed aging basis, age progression of a new input photo takes 0.1 seconds. For blending aged faces into adult heads we estimate fiducials in the adult head photo (computed during pose correction), to match fiducials between the input and target photos, and then run graph cuts to find an optimal seam followed by poisson blending to blend the aged face into the adult head photo [3].

Cropped progression results Figures 1 and 4 show age progressed images generated automatically using our method. The input images were taken from the FGNET database [17] and were not part of the training set used to create the flow and texture age differences. The results shown here focus on extremely challenging photos of children, with examples that cover a wide range of face types and imaging conditions: neutral, smiling, and laughing facial expressions, frontal and non-frontal pose, lower quality scans as well as higher quality photos, female and male children and a variety of lighting conditions. All results are cropped to the face area to show the raw output of the method. Note how the face shape changes with age in these sequences, e.g., the nose stretches, eyes narrow, and wrinkles appear. Textural changes include facial hair, “shadows” in male faces, eye makeup in female faces, and stronger eyebrows. Many more examples can be found in the supplementary material.



Figure 4. Age progression results. For each input image we automatically generate age progressed images for a variety of ages. Note the realistic progression results even with strong directional lighting, non-frontal pose, and non-neutral expressions.

4.1. Evaluation

We performed a large scale user study on Mechanical Turk, the most extensive of its kind in the age progression literature. In particular, we had human subjects compare our results to *every prior age progression result* we could find in the literature, and to ground truth (photos of 82 people at different ages). Each subject was shown a photo of a person at age X (e.g., 4), and two additional photos: A) a photo of the same person at an older age Y (e.g., 25), and B) our age-progressed result. The user was asked which of A or B is more likely to be the same person at age Y. They also had the option of selecting “both are equally likely” or “neither is likely.” Please refer to the supplementary material for a screenshot of the interface and exact wording. The order of our result and the ground truth was randomly chosen to prevent order bias. All photos were cropped to the face area only. If the progressed image at age Y is generated from the reference at age X, it will have the same lighting and expression. To avoid this similarity bias, our age progression result was generated not from the reference shown the user, but instead from a photo of the same person at the closest age to the reference.

Comparison with ground-truth We ran our method on

every photo in the FGNET dataset, and compared to every older photo available for each person. FGNET consists of photos of the same person over time, and several span baby to adult, resulting in a total of 2976 comparisons. Each user was presented three images: a photo of the subject at age X, an older photo at age Y, and an age progressed photo at age Y. They were asked to specify which of the latter two photos was more likely the same person at age Y by choosing: photo A, photo B, both are equally likely, or neither is likely to be the same person at age Y. Each comparison was evaluated by 3 different people, and 12 comparisons were left blank, making the total number of comparisons we received 8916. The number of unique workers was 72. The results are as following: we received 3288 votes (out of 8916, i.e., 37%) that our result is more likely, 3901 (44%) that ground truth is more likely, 1303 (15%) that both are equally likely, and 424 (5%) that neither is likely.

This result is so surprising that it led us to question how proficient humans are at this task, i.e., maybe we are just not good at face recognition across large age differences. To test this hypothesis, we conducted a perceptual study in which each user was shown two real (ground truth) images of the same person, separated by at least 5 years, and asked




Figure 5. Comprehensive comparison to prior work, plotting user study ratings of our method vs. all 120 results from prior work. Blue cells (> 0.55) are where our method scored higher, red cells (< 0.45) have prior method(s) scoring higher, and gray cells are ambiguous. Our method excels for aging children, while prior techniques that target adults perform better for that category.

to specify if it is the same or a different person. We used all pairs (at least 5 years apart) of each person on FGNET, and repeated each test three times on Mechanical Turk (8928 tests in total). The results indicate that people are generally good at recognizing adults across different age ranges, but poor at recognizing children after many years. In particular, across children aged 0-7, participants performed barely better than chance (57%) at recognition for roughly 10 year differences, at chance for 20 years (52%), and worse than chance for 50 years (33%). See supplementary material for the full details of the experiment and results. These studies point to the limits of human evaluation for assessing age progression results.

Ground-truth-blended comparisons While the Mechanical Turk study focuses on cropped faces, we also experimented with blending age progressed faces onto the ground truth head; representative results are shown in Figure 7 (additional results appear in the supplementary material). In each case, we take an input photo in the 0-3 age range and compare the ground truth image at each age (right) with our result (left). We blended our result into the ground truth head, using the process described earlier in this section. (We also include unblended results cropped to only the face area in the supplementary material.) The similarity is impressive, especially given that each sequence (column) is produced from a single baby photo. Note that the facial expression and lighting are fixed from the baby photo and therefore differ from the ground truth. As a strawman, we also blended the input child’s face onto the older ground

truth for comparison (Figure 6 (b)); clearly age progressing the input face prior to blending yields much more realistic composites.

Comparison to prior work We compared our results to *all* prior papers that demonstrate age progression results, with the exception of Lanitis et al. [17] whose results do not specify ages. These papers are: (p1) [37], (p2) [35], (p3) [33], (p4) [27], (p5) [28], (p6) [19], (p7) [20], (p8) [36].

While we’re most interested in long range age progression of very young children, for comparison we ran our method on every result we found in these papers (including adults and older children). The number of age progression results in papers p1-p8 was: 56, 2, 8, 5, 7, 4, 30 and 8 respectively, for a total of 120 comparisons. Each comparison was performed by 10 workers, and there were on average 13 unique workers per paper. Figure 5 plots the results of the user study: the x-axis is the input age group and the y-axis is the output age group. The score is calculated as follows: as in the ground-truth experiment, workers were asked to choose one of the four options. We added 1 point when our result was chosen, 0.5 when “both are likely” was chosen, and 0 when a result from prior work was chosen. The score was then normalized by number of responses per cell (we did not include examples for which the option “neither” was chosen here, as the ground truth evaluation captures similar statistics). As can be seen from Figure 5, our approach almost uniformly outperforms prior work for aging young children, and clearly dominates for aging children to adult. The one “red” box corresponds to an age change of only three years. Note, that there are no prior results in the literature for aging children beyond age 25; we are the first to attempt this task. On the other hand, techniques that focus on modeling older people (modeling wrinkles, hair color, etc.) do better for that category. Note that all previous works typically focus on one of the two age ranges: child to teenager or adult to older person, while our method is general and spans ages 0 to 100 (e.g., Fig. 4). While beyond the scope of this paper, incorporating wrinkles or hair lightening models could yield further improvements in the upper age ranges.

Very few age progression papers address young children [17, 19, 32, 35, 36], and those that do include only a handful of results. See supplementary material for a figure that compares our results to all results in the literature for children under 9 years of age.

Figure 6 (a) compares our results to Perrett et al.’s FaceTransformer tool at <http://morph.cs.st-andrews.ac.uk/Transformer/> and the PsychoMorph tool by Tiddeman et al. [38] at the Face Research Lab website <http://www.faceresearch.org/demos/>. As can be seen, they do not perform well on young children. As a baseline we also compare to applying only the aspect ratio change to the input face (compare columns 2 and 5). Both of these

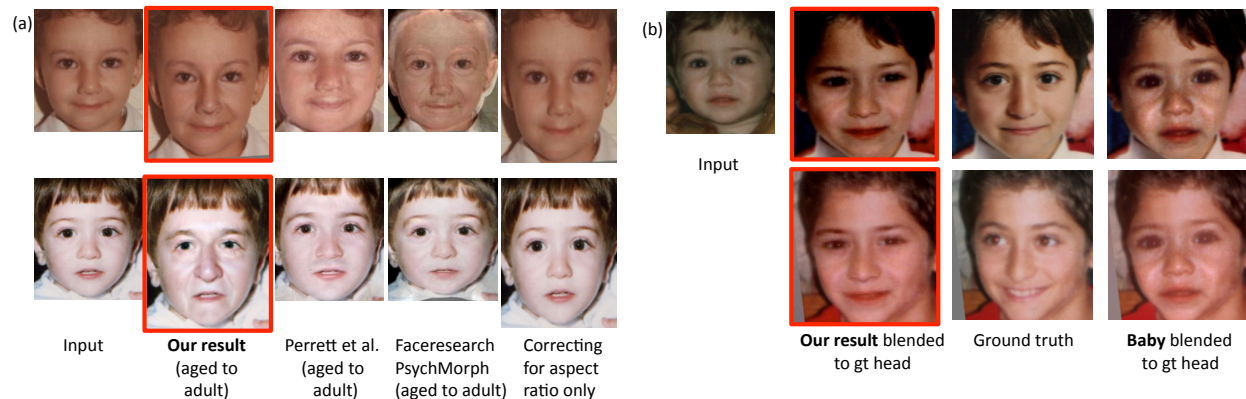


Figure 6. Comparison to other methods: (a) to Perrett et al. and FaceResearch online tool, (b) to mapping the baby’s face (far left) onto the ground truth (column 3) to produce a blended result (far right). The aged results (column 2) look much more similar to the ground truth, indicating that simply blending a face into a head of an older person does not produce a satisfactory age progression, additional shape and texture changes must be added.

tools require manual placement of facial features, whereas our approach is fully automated.

5. Conclusion

We presented a method for automatic age progression of a single photo to any age between 1 and 80, by leveraging thousands Internet photos across age groups. The method works remarkably well, in particular for the challenging case of young children, for which few prior results have been demonstrated. A key contribution is the ability to handle photos “in the wild,” with variable illumination, pose, and expression. Future improvements include: modeling wrinkles and hair whitening [37] to enhance realism for older subjects; output a set of progressed images per single input, building on face editing techniques, e.g., [24]; having a database of heads and upper torsos of different ages to composite our result onto.

Acknowledgements We thank Google and Intel for supporting this research.

References

- [1] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2):218–233, 2003.
- [2] P. Benson and D. Perrett. Extracting prototypical facial images from exemplars. *Perception*, 22:257–262, 1993.
- [3] D. Bitouk, N. Kumar, S. Dhillon, P. N. Belhumeur, and S. K. Nayar. Face swapping: Automatically replacing faces in photographs. In *ACM Trans. on Graph.*, 2008.
- [4] T. Brox, A. Bruhn, N. Papenbergh, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *ECCV*, pages 25–36, 2004.
- [5] A. Bruhn, J. Weickert, and C. Schnrr. Lucas/kanade meets horn/schunck: Combining local and global optic flow methods. *IJCV*, 61:211–231, 2005.
- [6] D. Burt and D. Perrett. Perception of age in adult caucasian male faces—comp graphic manipulation of shape and color information. *Pr. Royal S. London*, 259:137–143, 1995.
- [7] T. Ezzat and T. Poggio. Facial analysis and synthesis using image-based models. In *FG*, pages 116–121, 1996.
- [8] L. G. Farkas. *Anthropometry of the Head and Face*. 1994.
- [9] Y. Fu, G. Guo, and T. Huang. Age synthesis and estimation via faces: A survey. *PAMI*, 32(11):1955–1976, 2010.
- [10] F. Galton. Composite portraits. *Nature*, 18:97–100, 1878.
- [11] G. Guo and G. Mu. Simultaneous dimensionality reduction and human age estimation via kernel partial least squares regression. In *CVPR*, 2011.
- [12] T. Hassner. Viewing real-world faces in 3d. *ICCV*, 2013.
- [13] H. Heafner. Age-progression technology and its application to law enforcement. In *SPIE*, pages 49–55, 1996.
- [14] I. Kemelmacher-Shlizerman and S. M. Seitz. Face reconstruction in the wild. In *ICCV*, pages 1746–1753, 2011.
- [15] I. Kemelmacher-Shlizerman and S. M. Seitz. Collection flow. In *CVPR*, pages 1792–1799, 2012.
- [16] I. Kemelmacher-Shlizerman, E. Shechtman, R. Garg, and S. M. Seitz. Exploring photobios. *SIGGRAPH*, 30(4), 2011.
- [17] A. Lanitis, C. J. Taylor, and T. F. Cootes. Toward automatic simulation of aging effects on face images. *PAMI*, 24, 2002.
- [18] C. Li, Q. Liu, J. Liu, and H. Lu. Learning ordinal discriminative features for age estimation. In *CVPR*, 2012.
- [19] Y. Liang, C. Li, H. Yue, and Y. Luo. Age simulation in young face images. In *Bioinf. and Biomed. Eng.*, 2007.
- [20] Y. Liang, Y. Xu, L. Liu, S. Liao, and B. Zou. Transactions on edutainment vi. chapter A multi-layer model for face aging simulation, pages 182–192. 2011.
- [21] C. Liu. *Beyond Pixels: Exploring New Representations and Applications for Motion Analysis*. PhD thesis, MIT, 2009.
- [22] C. Liu, J. Yuen, A. Torralba, J. Sivic, and W. T. Freeman. Sift flow: Dense correspondence across different scenes. In *ECCV*, pages 28–42. 2008.
- [23] MerrillLynch. <http://faceretirement.merrilledge.com/>, 2013.



Figure 7. Comparison to ground truth images. In each case a single photo of a child (top) is age progressed (left) and compared to photos of the same person (right) at the corresponding age (labeled at left). The age progressed face is composited into the ground truth photo to match the hairstyle and background (see supplementary material for comparisons of just the face regions). Facial expression and lighting are not matched to the ground truth, but retained from the input photo. Note how well the age progressed face matches the ground truth face, given that the full sequence is synthesized from a single baby photo.

- [24] U. Mohammed, S. J. D. Prince, and J. Kautz. Visualization: generating novel facial images. *ACM Trans. Graph.*, 28(3):57:1–57:8, July 2009.
- [25] NCMEC. Age progression. Technical report, National Center for Missing and Exploited Children, 2010.
- [26] B. Ni, Z. Song, and S. Yan. Web image mining towards universal age estimator. *Proc. ACM Multimedia*, 2009.
- [27] U. Park, Y. Tong, and A. K. Jain. Face recognition with temporal invariance: A 3d aging model. In *FG*, 2008.
- [28] E. Patterson, A. Sethuram, M. Albert, and K. Ricanek. Comparison of synthetic face aging to age progression by forensic sketch artist. In *Vis. Imag. Proc.*, pages 247–252, 2007.
- [29] P. Paysan. *Statistical modeling of facial aging based on 3D scans*. PhD thesis, University of Basel, 2010.
- [30] A. Prince. Age progression, forensic and medical artist, <http://aurioleprince.wordpress.com/>, 2013.
- [31] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *SIGGRAPH*, 2001.
- [32] N. Ramanathan and R. Chellappa. Modeling age progression in young faces. In *CVPR*, volume 1, pages 387–394, 2006.
- [33] N. Ramanathan, R. Chellappa, and S. Biswas. Age progression in human faces : A survey. *J. of Vis. Lang. Comp.*, 2009.
- [34] V. Rokhlin, A. Szlam, and M. Tygert. A randomized algorithm for pca. *SIAM J. Mat. Anal.*, 31(3):1100–1124, 2009.
- [35] K. Scherbaum, M. Sunkel, H.-P. Seidel, and V. Blanz. Prediction of individual non-linear aging trajectories of faces. *EUROGRAPHICS*, (3):285–294, 2007.
- [36] C.-T. Shen, W.-H. Lu, S.-W. Shih, and H.-Y. Liao. Exemplar-based age progression prediction in children faces. In *IEEE Int. Symp. on Multimedia*, pages 123–128, 2011.
- [37] J. Suo, S.-C. Zhu, S. Shan, and X. Chen. A compositional and dynamic model for face aging. *PAMI*, 2010.
- [38] B. Tiddeman, M. Stirrat, and D. Perrett. Towards realism in facial transformation: results of a wavelet mrf method. *Computer Graphics Forum, Eurographics*, 24, 2005.