Automatic Face Recognition of Newborns, Infants, and Toddlers: A Longitudinal Evaluation

Lacey Best-Rowden, Yovahn Hoole, Anil K. Jain
Department of Computer Science and Engineering
Michigan State University
East Lansing, MI USA 48824
Email: {bestrow1, hooleyov, jain}@msu.edu

Abstract—A number of emerging applications requiring reliable identification of children have called attention to whether biometric traits can be utilized as a solution. While biometric traits based on friction ridge patterns (*e.g.* fingerprints, footprints) have been evaluated to some extent, to our knowledge, no effort has been made to evaluate the efficacy of automatic face recognition of young children over useful durations of time. Additionally, there are some applications where only the face images of a child are available, such as identification of missing or abducted children and children shown in sexually exploitive media sequestered by law enforcement. In this paper, we introduce the Newborns, Infants, and Toddlers Longitudinal (NITL) face image database, which was collected by the authors during four different sessions over a period of one year (March 2015 to March 2016) at the Saran Ashram Hospital, Dayalbagh, India. The NITL database contains 314 subjects in total in the age range of 0 to 4 years old. The aim of this paper is to provide a comprehensive evaluation of a state-of-the-art commercial-off-the-shelf (COTS) face matcher on the NITL face image database to investigate the feasibility of face recognition of children as they age. Experimental results show that while available face recognition technology is not yet ready to reliably recognize very young children, face recognition enrolled at 3 years of age or older may be feasible.

I. INTRODUCTION

Reliable identification of young children is a problem of growing interest in a variety of applications such as vaccination and healthcare tracking, civil ID, and child abduction. Current practice for most applications which require a child's identity is to link the child to their parents' identities. In hospitals this is done at birth with an identification bracelet[1] or with information (e.g. child's name, date of birth, and parents' names) provided by the child's guardian or caregiver. However, there are a few key issues (*e.g.* a guardian or neighbor bringing the child for vaccination instead of the parents) that render this approach inadequate; instead, giving a child their own identity would be preferred.

Tracking the vaccinations of young children is an important application requiring child recognition. According to UNICEF, about 6.6 million children die each year from vaccine-preventable diseases, and about 25 million do not receive the proper vaccines [1]. In order to reduce these numbers and improve the health of children worldwide, it is imperative that healthcare organizations have a reliable means of identifying children to know which vaccines they need or have previously been administered. This is particularly a concern for non-governmental organizations (NGOs) operating in remote villages of developing countries where even parents do not have reliable forms of identification.

Swapping of newborns in hospitals is another application that requires recognition of child identities. While security measures currently in place in many hospitals, such as RFID technology,[2] help to make this a rare occurrence in countries such as the United States, being "switched at birth" is still a possibility due to overcrowded hospitals, negligent hospital workers, or even criminal behavior.[3] Because identification bracelets can be lost or maliciously removed, alternative solutions based on biometric recognition are of interest.

Child identification would also be beneficial for civil ID programs. India's Aadhaar program, in particular, is providing unique identification (UID) numbers to *all* of the approximately 1.2 billion residents of India. At a current enrollment of about 100 "crore" (1 billion) individuals and 93% of the adult population,[4] the next focus is on enrolling children aged 0-5 years old.[5] While Aadhaar relies on "core" biometrics, defined as both irises and all ten fingerprints, biometric data is currently not captured until a child is 5 years of age. In future years, a majority of work in maintaining the program will involve enrolling new births. The Indian government has started capturing a face image of each newborn at hospitals and linking it to one of the parents' biometric.[6] At what age to first enroll a child with his own biometrics is still a question left to be answered.

With the growing success of biometric technologies, there is interest in knowing whether *biometrics* can offer better solutions for recognition of newborns, infants, and toddlers.

---

[1]Footprints of newborns are also taken but, to our knowledge, have not been used in practice for child recognition; it is more of a souvenir for the parents.

[2]http://www.iosrjournals.org/iosr-jce/papers/Vol16-issue3/Version-1/F016312832.pdf

[3]http://www.dailymail.co.uk/news/article-3225887/Parents-discovered-newborn-swapped-birth-doctor-sold-trafficking-ring-finally-reunited-lost-son.html

[4]https://uidai.gov.in/

[5]http://economictimes.indiatimes.com/news/economy/policy/now-modi-government-to-give-aadhaar-to-newborns/articleshow/52441170.cms

[6]http://www.hindustantimes.com/india/govt-hospitals-to-soon-start-aadhaar-linked-birth-registration-of-newborns/story-uGAqHm0LDJcj3Yfu7KVfQM.html

Session 1 Mar. 2015	Session 2 Sep. 2015	Session 3 Jan. 2016	Session 4 Mar. 2016		Session 1 Mar. 2015	Session 2 Sep. 2015	Session 3 Jan. 2016	Session 4 Mar. 2016
3 months	0.598	0.543	0.510		**1 year**	0.632	0.300	0.487
4 months	0.545	0.597	0.650		**2.5 years**	0.415	0.473	0.469
5 months	0.539	0.410	0.145		**3 years**	0.723	0.613	0.828

Fig. 1: *Longitudinal face images of six example subjects from the Newborns, Infants, and Toddlers (NITL) face image database which was collected during four different sessions (March 2015, September 2015, January 2016, and March 2016) at the Saran Ashram Hospital, Dayalbagh, India. The age of each subject at the first acquisition is given in red text, and comparison scores from a state-of-the-art COTS matcher between the first image and all subsequent images are given in black. Verification thresholds at 0.1% and 1% false accept rates (FAR) are 0.615 and 0.524, respectively.*

For targeted applications where a child interacts with an operator (*e.g.* a healthcare worker), biometric traits based on friction ridge patterns (*e.g.* fingerprints) are already being attempted. However, there are a number of applications where such controlled capture of biometric traits may not be possible, and only face images are available for use.

Significant advancements in automatic face recognition technology have called attention to whether state-of-the-art algorithms are capable of recognizing the faces of children, and if so, over what duration of time (*i.e.* time gap between enrolled face and query face)? While aging is an issue for any biometric recognition system, it is particularly a concern with children, as early developmental stages involve rapid changes in the face. The aim of this paper is to investigate the performance of state-of-the-art face recognition systems on face images of newborns, infants, and toddlers.[7] To address this question, it is important to first collect a longitudinal face database. The contributions of this paper are as follows.

- A Newborns, Infants, and Toddlers Longitudinal (NITL) face image database which consists of 314 children ages 0 to 4 years old. The NITL database is *longitudinal* in that multiple face images were captured of the same subjects during four different acquisition sessions over the course of one year (March 2015 to March 2016). There are a total of 161 subjects with face images from all four sessions (see Fig. 1).

[7]The generally accepted definitions of newborns, infants, and toddlers are the following: newborns are 0 to 28 days, infants are 1 to 12 months, and toddlers are 1 to 3 years old.

- Evaluation of a state-of-the-art commercial-off-the-shelf (COTS) face matcher on the NITL database to investigate the challenges in recognizing children as their faces undergo changes over the course of one year. The COTS matcher used in this paper was among the top-3 performers in the NIST FRVT 2013 [2]. We also investigate the challenges associated with the uncooperative nature of children which results in significant variations in pose, illumination and expression.

II. Related Work

A number of different biometric modalities have been considered in the literature for automatic recognition of children, including fingerprints [3], [4], [5], face [6], [7], iris [8], footprint [9], palmprint [9], [10], and ballprint [11].[8] As with all biometric systems, the application and user population should drive the choice of biometric trait. Our objective in this paper is to investigate the feasibility of face as a biometric for children. As such, the remainder of this section highlights prior studies on *face* recognition of children.

Prior work on automatic face recognition of children is limited due to: (i) difficulty in collecting longitudinal face data, and (ii) the general perception that face recognition of children is infeasible because of rapid growth rates, especially during the ages of 0-4 years of interest here. Regarding the former, to the best of our knowledge, the only *publicly available* face image databases that include images of children are FG-NET

[8]The ballprint is the region of the foot located under the big toe.

Fig. 2: Data collection effort in Dayalbagh, Agra, India. (a) Parents and children waiting in line outside the data collection room (pediatrician's clinic), (b) parents signing the consent form agreeing to provide their child's biometric data, (c) data capture by the authors at the two data collection stations, and (d) incentive (a bag consisting of staples) handed out to the parents after data collection.

[12] and FaceTracer [13]. However, the FaceTracer database only has one image per child. Most prior work on age-invariant face recognition utilizes the FG-NET database (see [12] for a summary). However, results are typically reported on the *entire* database (*e.g.* [14], [15]), so performance on the children in the database is not known. Additionally, there are only 82 subjects in total, and only 50 of which have *two* face images at ages younger than five years old. The Cross-Age Celebrity Database (CACD) [16] was collected to study facial aging, but it does not contain any subjects younger than 10 years old. Ricanek *et al.* provide a review of multiple face recognition algorithms on longitudinal face images from the In-the-Wild Child Celebrity (ITWCC) database [7], but the average age of the first/youngest image of each subject is 10.2 years with a standard deviation of 3.9 years. Hence, the ITWCC database is not useful for our study.

The NIST FRVT 2013 [2] conducted an evaluation of accuracy dependence on subject age for different age groups. The *baby* and *kid* age groups consisted of images of 57 and 340 subjects ages $[0, 3)$ and $[3, 8)$ years old, respectively. For six of the top-performing face recognition algorithms, error rates (false negative identification) were greater than 60% for the baby age group, and only one algorithm had an error rate less than 50% for the kid age group. For comparison, error rates for the parents age group (ages 30 to 55 years) were around 10% for the top-performers. These results strongly indicate the difficulty in recognizing young children; however, the database is not publicly available.

Motivated by the non-intrusive capture of face images, recent studies have proposed algorithms for face recognition of newborns, specifically geared towards the application of enrolling newborn identities at birth and verifying them prior to leaving the hospital. Face images of newborns (0-3 weeks old) were captured by Tiwari *et al.* [17] and Bharadwaj *et al.* [6]. Both studies concluded that capturing good quality face images of newborns is difficult because of the gross head reflexes, and pose and expression variations. With four images per newborn in a gallery of 86 subjects, Bharadwaj *et al.* [6] reports rank-1 identification accuracy of 78.5% and verification accuracy of 63.4% at 0.1% false accept rate (FAR) (all face images were cropped using *manually* annotated eye and mouth locations).[9] Even though the face images of the same subject are all taken within 24 hours, these error rates are too high to be of practical value. Furthermore, the database has no *longitudinal* aspect since all images of a subject are taken within 24 hours.

Automatic face recognition of newborns, infants, and toddlers deserves more attention from the research community, but it is important to frame the research goals in terms of feasible applications. While face recognition may not be able to reach accuracy requirements for controlled applications like healthcare, it is still necessary to evaluate the performance and identify the challenges, as it may be a useful investigative tool, in a semi-automatic fashion ("human in the loop"), for certain law enforcement applications (*e.g.* identifying missing children).

---

[9]Bharadwaj *et al.* [6] state that the newborn face image database will be made publicly available, but at the time of this submission, the images have not yet been released.
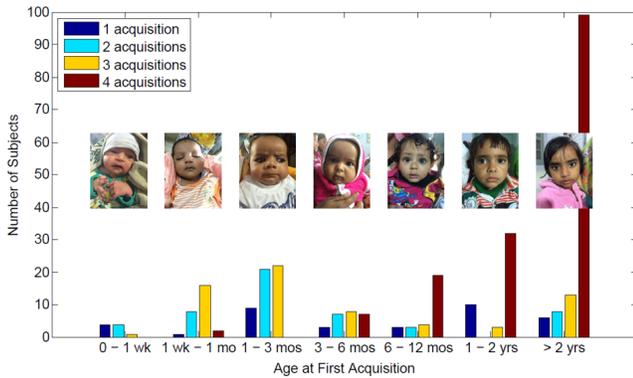
Fig. 3: *Summary of the number of subjects in different age groups, where the age group is determined by the subject's age at the first acquisition (March 2015). An example face image is shown for each of the age groups. The bar colors indicate the number of acquisitions for each age group; for example, among the 125 total subjects older than 2 years at the first acquisition, 99 of them subsequently attended all four sessions of data collection.*

III. Longitudinal Face Image Data Collection

Our longitudinal face data collection was conducted in Dayalbagh, India, as part of an effort to investigate the feasibility of using fingerprints to reliably recognize children ages 0–4 years old. Fingerprints and face images were collected in a pediatrician's free clinic at the Saran Ashram Hospital while the pediatrician, Dr. Bhatnagar, was present. Two different data collection stations were each manned by the authors. Face images were captured using the 8MP rear camera of iPhone 5/5s. The child's name, age, gender, and address and contact number of the child's parents were collected to contact the parents for follow up visits during subsequent phases of data collection. This paper analyzes the longitudinal face database that was collected.

Fig. 2 shows images of children, their accompanying parents, and the data capture stations inside the pediatrician's clinic. Parents were required to sign a consent form (approved by the Dayalbagh's Educational Institute's Ethics Committee, Saran Ashram hospital administration, as well as the authors' institutional review board) giving their consent to provide their child's fingerprint and face images. Face images were captured at one of the two data collection stations, and an incentive (a bag of staples, voucher for the local grocery store, or blanket) was handed out to the parents after data collection was complete.

Face images were collected during four different sessions over the duration of one year: March 2015, September 2015, January 2016, and March 2016. A total of 206 subjects participated in the March 2015 session, of which 178 subjects also returned in March 2016 (one year later), and 161 subjects attended all four sessions. Fig. 3 summarizes the resulting data by age groups and number of acquisitions per subject. During Sessions 2 and 3 of data collection, we further recruited an additional 107 newborns and infants (66 new subjects for Session 2 and 41 new subjects in Session 3). The total number of subjects encountered in our study was 314.

During data collection in the pediatrician's clinic, 3 to 5 face images of each subject were captured sequentially over approximately one minute. Many of the acquired face images are quite unconstrained, particularly with respect to poor illumination conditions in the clinic (see Fig. 2(c)) and uncooperative subjects leading to variations in pose and expression. To meet the throughput requirements of the long line queued up (see Fig. 2(a)), we could not afford to spend more than five minutes or so with each child for the entire biometric data capture (three impressions of each thumb and at least three face images). Hence, face images were captured quickly with little effort to ensure that a high quality face image was captured (especially if the child was screaming or crying). This is the situation in typical operational scenarios we are targeting.

IV. Experiments

To facilitate analysis of both the absolute *age* and the longitudinal *aging* factors, we consider subsets of the NITL database based on the session at which each subject was first encountered (the first face acquisition), which we refer to as the *enrollment* session. Table I details the three subsets of the NITL database used for experiments, namely, (i) newborns, infants, and toddlers enrolled in Session 1 (NIT-S1), (ii) newborns and infants enrolled in Session 2 (NI-S2), and (iii) newborns and infants enrolled in Session 3 (NI-S3).

Face comparison scores were obtained from a state-of-the-art COTS face matcher, denoted COTS-A, which was among the top-3 performers in the NIST FRVT 2013 [2]. All face images were scaled down to 244×326 (one tenth of the original image size) to reduce computational time (there was no loss in accuracy). Fig. 8 shows that poor illumination, motion blur, partial face images, and extreme facial pose caused failure to enroll (FTE) for some face images (*i.e.* the face and/or eyes could not be detected by COTS-A); COTS-A failed to enroll 24 of the 3,144 total face images in the database.[10] Because we captured 3 to 5 face images per subject in each session, we manually selected only three images per subject for the experiments, and found that max fusion across the multiple face images in each session performed better than other fusion schemes (mean and min). Therefore, max fusion is used for all experiments reported here.

We first evaluated the performance of COTS-A on matching face images from the same session to obtain a baseline for intra-session variability. Genuine scores are computed as all pairwise comparisons between face images of the same subject from the same session. Impostor scores are all possible impostor scores from all four sessions (a total of 1,363,885 scores). As expected, the true accept rates are quite high (TARs greater than 93% at 0.1% FAR for all four sessions). This is because intra-session variability is small; face images from the same session of the same subject are often near-duplicates, captured consecutively within a few seconds.

---

[10]If an image failed to enroll, all comparisons with that image are set to the minimum similarity score (resulting in a false reject error for a genuine pair, but a true reject for an impostor pair).
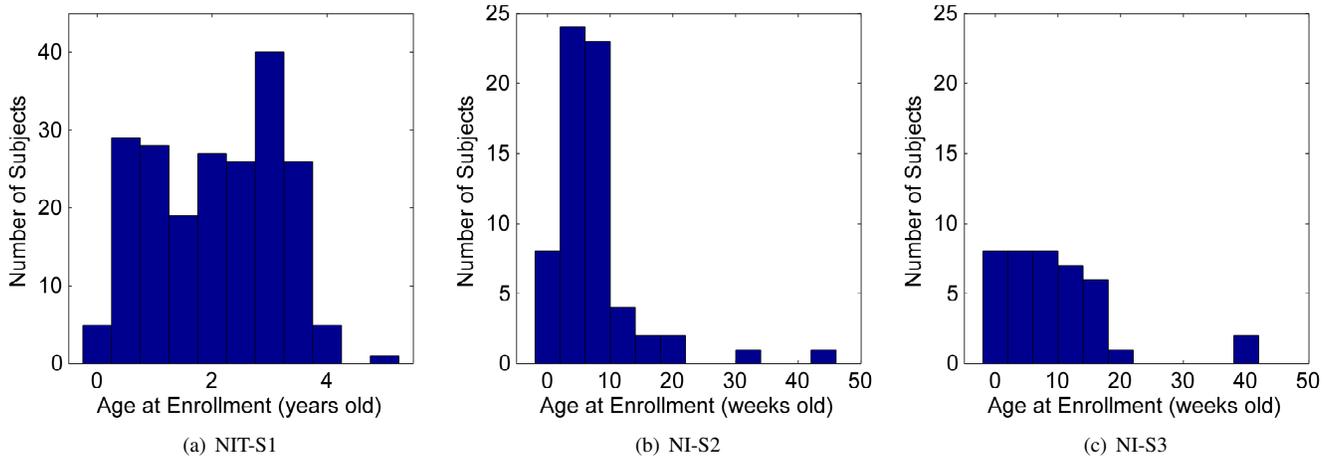
| (a) NIT-S1 | (b) NI-S2 | (c) NI-S3 |

Fig. 4: Distributions of age at enrollment for three subsets of the NITL face image database. (a) NIT-S1: 206 newborns, infants, and toddlers whose first face acquisition was in Session 1, March 2015; (b) NI-S2: 66 newborns and infants whose first face acquisition was in Session 2, September 2015; and (c) NI-S3: 41 newborns and infants whose first face acquisition was in Session 3, January 2016. Note that age in (a) is displayed in years old, while the age in (b) and (c) are in weeks.

Subset	Enrollment Session	# Subjects (males)	Age Range (median age)	# Ret. Sess. 2	# Ret. Sess. 3	# Ret. Sess. 4	Time Gap ($\triangle T$)
NIT-S1	1 (Mar. 2015)	206 (96)	$0 - 5$ (2.0) yrs	169	182	178	12 mos
NI-S2	2 (Sep. 2015)	66 (33)	$0 - 23$ (1.5) mos	*n.a.*	53	51	6 mos
NI-S3	3 (Jan. 2016)	41 (18)	$0 - 32$ (1.8) mos	*n.a.*	*n.a.*	32	2 mos

TABLE I: Summary of the three subsets of the NITL face image database used in our experiments.

To evaluate the longitudinal performance, we compare Session 1 face images of the NIT-S1 subset (206 subjects) to the face images of the same subjects from the subsequent sessions. This allows us to observe performance after 6, 10, and 12 months elapsed time since enrollment. For verification experiments, the impostor distribution is kept consistent and includes all possible impostor scores from comparing Session 1 face images with face images from Sessions 1 through 4, totaling 150,881 scores (after max fusion). For identification experiments, the Session 1 face images make up the gallery, so there are 206 subjects in the gallery.

Fig. 5 shows that the performance for subjects younger than 1 year old at enrollment in Session 1 is much worse than for subjects at least 1 year old.[11] Although the performance is better for older subjects, it is still quite poor; at 0.1% FAR, TAR is only 53% for elapsed time of only 6 months (Session 1 to Session 2). Note that the same operating point results in almost 30% lower TAR for subjects younger than 1 year old. Fig. 5 also shows that Session 2 and Session 4 performance are similar for both verification and identification, suggesting that longitudinal time of 12 months is not much worse than longitudinal time of 6 months. What is peculiar, however, is that Session 3 performance (longitudinal time of 10 months) is much worse than both Session 2 and Session 4. We attribute this to the fact that Session 3 was collected in January 2016 when it was quite cold (winter) in Dayalbagh, India; many

---

[11] We also evaluated the difference between children younger and older than 2 years old, but the performance gap was similar to the difference between children younger and older than 1 year old.

children were wearing warm hats (which we did not ask them to remove), and the sunlight streaming through the open door of the room in the clinic was different during January, as it was cloudy most days.

Because the impostor distribution was a single distribution containing all impostor scores across all sessions for all curves in Fig. 5, we construct a heatmap of mean impostor similarity scores by age group, shown in Fig. 6(a), to further explore the distribution. The higher mean similarity scores along the diagonal elements of the figure indicate that subjects with similar ages tend to falsely match each other more than subjects that are a couple years apart in age.

A. Face Recognition of Newborns and Infants

Same-week Recognition: In Session 2 of data collection, we recruited a new set of 66 newborns and infants (NI-S2 in Table I). As shown in Fig. 4(b), almost all of these subjects were younger than 24 weeks (6 months) old at the first acquisition. Because data collection occurred Monday through Friday, we asked these 66 newborns to return a few days later. This allows us to evaluate the performance of COTS-A on infants and newborns when genuine face images are only separated by two to four days; COTS-A performance in this scenario is extremely poor with TARs of 15.07% and 36.99% at 0.1% and 1% FAR, respectively, and rank-1 identification accuracy of 36.93% against the gallery of only 66 newborns and infants. Note that this COTS-A performance is quite a lot lower than performance reported by Bharadwaj *et al.* (rank-1 accuracy of 78.5% and 63.4% TAR at 0.1% FAR) on their
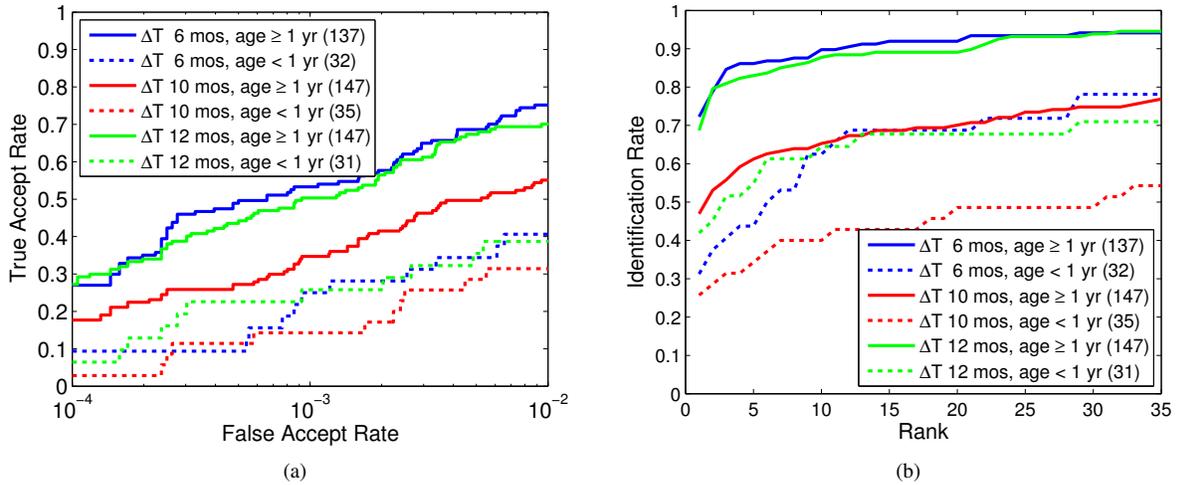
Fig. 5: Longitudinal performance of COTS-A for (a) verification and (b) identification experiments of the 206 total subjects in the NIT-S1 subset of the NITL face image database. The gallery for these experiments consists of face images from Session 1, and the probe sets are face images from Sessions 2, 3, and 4 (with different elapsed times, $\triangle T$). Performance is further broken down by age at Session 1 enrollment (younger than 1 year old vs. at least 1 year old). The numbers in parentheses indicate the number of subjects (out of 206) that returned for each of the subsequent sessions. The number of impostor scores in (a) is 150,881 scores, and the gallery size in (b) is 206 subjects.
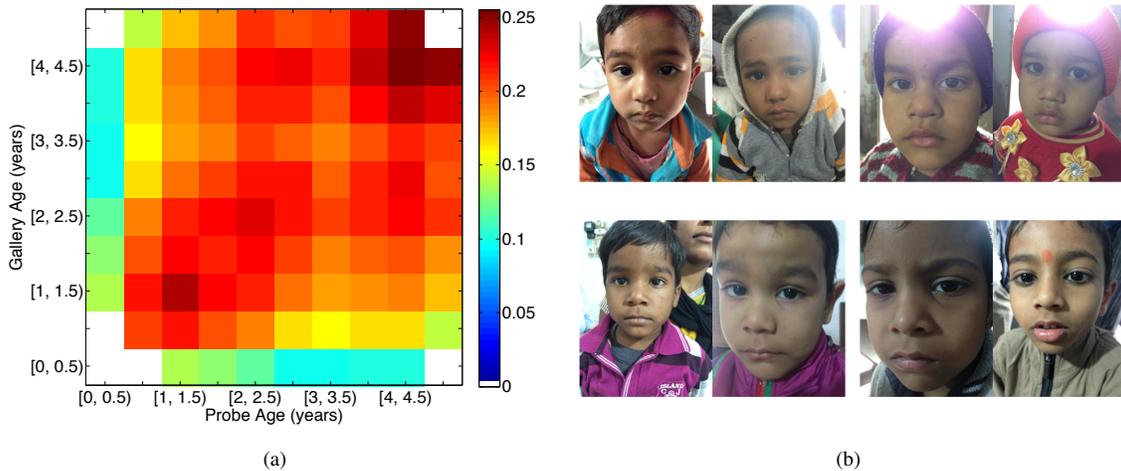

Fig. 6: (a) Heatmap of mean impostor similarity scores binned by different age groups, and (b) example image pairs of false accept errors made by COTS-A in face verification. Higher mean similarity scores along the diagonal elements of the figure in (a) indicate that subjects with similar ages are more likely to falsely match each other than subjects that are a couple years apart in age.

Newborn Face Database [6] which is captured in much more constrained environment and required manual face cropping.

Longitudinal Recognition: Next, we are interested in the performance on newborn and infant face images after time lapse of a few months. For the subset NI-S2, we can look at performance after the subjects have aged 4 months (Session 2 vs. Session 3) and 6 months (Session 2 vs. Session 4). For the subset NI-S3, the performance is only for two months time separation (Session 3 vs. Session 4). For both verification and identification experiments, the gallery consists of the Session 2 images of the NI-S2 subjects and the Session 3 images of the NI-S3 subjects, for a total gallery size of 107 subjects. For verification experiments, the impostor distribution is kept consistent and includes all impostor scores from comparing

the gallery to the remaining sessions of NI-S2 and NI-S3, a total of 19,778 scores (after max fusion).

Fig. 7 shows the resulting performance. Because the impostor distribution is kept consistent for all three curves in Fig. 7(a), higher performance for NI-S3 (pink dotted curve) can be attributed to generally higher *genuine* similarity scores. Conversely, one interesting observation is that NI-S3 performance is the *worst* for identification scenario in Fig. 7(b). One explanation for this peculiarity is that the impostor similarity scores from only the NI-S3 comparisons are higher because the probe images have only aged 2 months so they are more similar in age to the gallery than the NI-S2 probe sets. Hence, the ranks of the true mates (genuine scores) for NI-S3 tend to be worse.
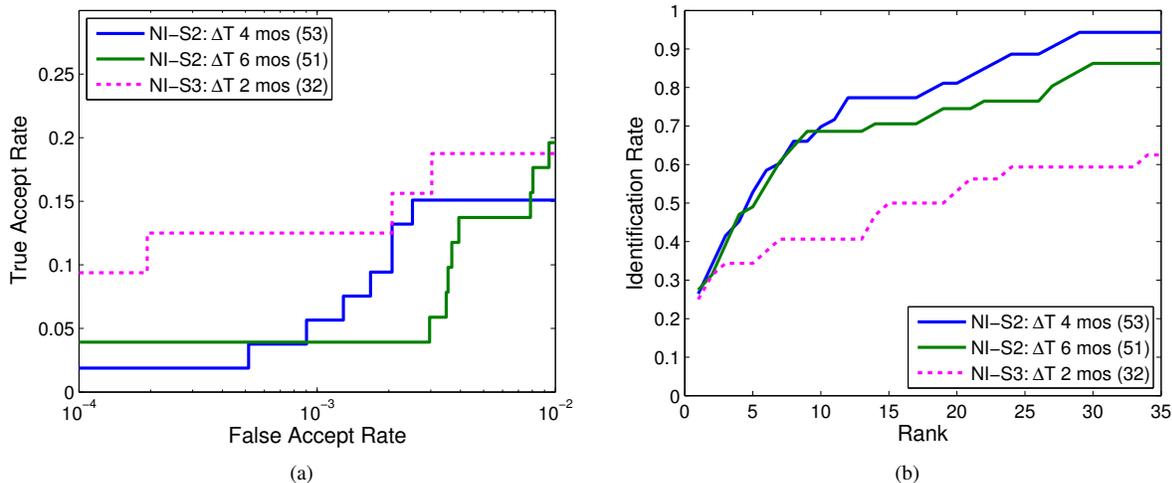
| | (a) | (b) |

Fig. 7: *Longitudinal performance of COTS-A for (a) verification and (b) identification of the 66 and 41 newborns and infants in the NI-S2 and NI-S3 subsets, respectively, of the NITL face image database. The numbers in parentheses indicate the number of subjects (out of 66 or 41) that returned for each of the subsequent sessions. The number of impostor scores used in (a) is 19,778 scores (the same distribution for each of the three ROC curves), and the gallery size in (b) is 107 subjects.*

Age Group at Enrollment (years)	[0, 1) **82**	[1, 2) **34**	[2, 3) **40**	[3, 5) **64**
TAR (%) @ 0.1% FAR for 6 mos. ΔT	15.85	44.12	45.00	60.94

TABLE II: *Verification rates for different age groups at enrollment. Num. of subjects are shown in bold.*

Lastly, we compare the performance for all subjects from the NIT-S1 and NI-S2 subsets with two images at 6 months elapsed time. The verification rates for different age groups (determined by age at enrollment) are shown in Table II for FAR of 0.1% computed from the same impostor distribution of all subjects for all age groups. These results suggest that face recognition beginning at 3 years or older may be feasible.

V. SUMMARY

We have investigated the feasibility of automatic face recognition for newborns, infants, and toddlers to meet growing needs for child recognition in a number of applications ranging from newborn identification in hospitals to identification of missing children. We introduced the Newborns, Infants, and Toddlers Longitudinal (NITL) face image database, which was collected by the authors at the Saran Ashram Hospital, Dayalbagh, India. The database was partitioned to study performance of a state-of-the-art COTS-A with respect to the enrollment age of the child and the time gap between enrolled and query face images. Our results can be summarized as follows: (i) while same-session face recognition has very high accuracy (TAR > 93% at 0.1% FAR), the cross-session performance degrades significantly to 47.93% TAR at 0.1% FAR after the children have aged just 6 months, (ii) the age at enrollment (*e.g.* younger than 1 year vs. older than 1 year old) has more influence on recognition performance than time lapse of 6 or 12 months (Fig. 5), and (iii) as expected, as the age at enrollment increases, the recognition performance improves

(see Table II). While automatic face recognition of young children may be a useful investigative tool (with human in the loop) to help identify, for example, sexually exploited children, the overall accuracies currently do not meet the requirements for most other applications that could benefit from automatic child recognition, such as vaccination tracking and civil ID programs. Future work will involve formal statistical analysis of the longitudinal scores for child face images, following the methodology in recent longitudinal studies of biometric recognition [18], [19], [20].

REFERENCES

[1] "Committing to child survival: A promise renewed," http://www.unicef.org/lac/Committing_to_Child_Survival_APR_9_Sept_2013.pdf, 2013.

[2] P. Grother and M. Ngan, "Frvt: Performance of face identification algorithms," in *NIST Interagency Report 8009*, May 2014.

[3] F. Galton, "Finger prints of young children," *British Association for the Advancement of Science*, 1899.

[4] A. K. Jain, K. Cao, and S. S. Arora, "Recognizing infants and toddlers using fingerprints: Increasing the vaccination coverage," in *IEEE IJCB*, 2014, pp. 1–8.

[5] A. K. Jain, S. S. Arora, L. Best-Rowden, K. Cao, P. S. Sudhish, A. Bhatnagar, and Y. Koda, "Giving infants an identity: Fingerprint sensing and recognition," in *ICTD*, 2016.

[6] S. Bharadwaj, H. S. Bhatt, M. Vatsa, and R. Singh, "Domain specific learning for newborn face recognition," *IEEE Trans. on Information Forensics and Security*, vol. 11, no. 7, pp. 1630–1641, July 2016.

[7] K. Ricanek, S. Bhardwaj, and M. Sodomsky, "A review of face recognition against longitudinal child faces," in *BIOSIG*, 2015.

[8] P. Corby, T. Schleyer, H. Spallek, T. Hart, R. Weyant, A. Corby, and W. Bretz, "Using biometrics for participant identification in a research study: a case report," *J. Am. Medical Informatics Assoc.*, vol. 13, no. 2, pp. 233–235, 2006.

[9] D. Weingaertner, O. R. P. Bellon, L. Silva, and M. N. L. Cat, "Newborn's biometric identification: Can it be done?" in *VISAPP (1)*, 2008, pp. 200–205.

Fig. 8: *Example face images that COTS-A failed to enroll (i.e., incorrect face and/or eye detections) due to large changes in facial pose, poor illumination, and motion blur.*

[10] R. P. Lemes, O. R. P. Bellon, L. Silva, and A. K. Jain, "Biometric recognition of newborns: Identification using palmprints," in *IEEE IJCB*, 2011, pp. 1–6.

[11] J. Kotzerke, A. Arakala, S. Davis, K. Horadam, and J. McVernon, "Ballprints as an infant biometric: A first approach," in *IEEE Workshop on BIOMS*, Oct 2014.

[12] G. Panis, A. Lanitis, N. Tsapatsoulis, and T. F. Cootes, "An overview of research on facial aging using the FG-NET aging database," *IET Biometrics*, May 2015.

[13] N. Kumar, P. N. Belhumeur, and S. K. Nayar, "Facetracer: A search engine for large collections of images with faces," in *Proc. Eur. Conf. Comput. Vis.*, 2008.

[14] Z. Li, U. Park, and A. K. Jain, "A discriminative model for age invariant face recognition," *IEEE Trans. on Information Forensics and Security*, vol. 6, no. 3, pp. 1028–1037, Sep. 2011.

[15] H. Ling, S. Soatto, N. Ramanathan, and D. W. Jacobs, "Face verification across age progression using discriminative methods," *IEEE Trans. on Information Forensics and Security*, vol. 5, no. 1, pp. 82–91, Mar. 2010.

[16] B.-C. Chen, C.-S. Chen, and W. Hsu, "Face recognition and retrieval using cross-age reference coding with cross-age celebrity dataset," *IEEE Trans. on Multimedia*, vol. 17, no. 6, pp. 804–815, Apr. 2015.

[17] S. Tiwari, A. Singh, and S. K. Singh, "Intelligent method for face recognition of infant," *International Journal of Computer Applications*, vol. 52, no. 4, pp. 36–50, 2012.

[18] L. Best-Rowden and A. K. Jain, "A longitudinal study of automatic face recognition," in *International Conference on Biometrics*, 2015.

[19] S. Yoon and A. K. Jain, "Longitudinal study of fingerprint recognition," *Proc. National Academy of Sciences*, vol. 112, no. 28, pp. 8555–8560, Jul. 2015.

[20] P. Grother, J. Matey, E. Tabassi, G. Quinn, and M. Chumakov, "IREX VI. Temporal Stability of Iris Recognition Accuracy," *NIST Interagency Report*, vol. 7948, 2013.