Automatic Photo Enhancement Server (HIPIE 2)

Renato Keshet, Carl Staelin, Boris Oicherman, Mani Fischer, Pavel Kisilev, Sagi Schein, Doron Shaked, Marie Vans, Hila Nachlieli, Ruth Bergman, Ron Maurer, Michal Aharon, Alexander Berkovich; HP Labs–Israel; Haifa, Israel
Ran Waidman, Oren Bengigi, Gal Amit, Shlomo Harush, Alon Asher; HP Indigo Division; Rehovot, Israel
Hui Chao, Multimedia Interaction and Understanding Lab, HP Labs; Palo Alto, CA, USA
Darryl Greig, Multimedia Interaction and Understanding Lab, HP Labs; Bristol, UK
Matthew Gaubatz, Print Production Automation Lab, HP Labs; Owego, NY, USA
Steven Simske; Print Production Automation Lab, HP Labs; Fort Collins, CO, USA

Abstract

HIPIE 2 is an internal code name for the HP Smartstream Photo Enhancement Server, *– a robust, scalable, and automatic photo image enhancement application, designed for photo specialty workflows fulfilled using Indigo presses. It is intended for 24/7 operation, without human intervention, and is part of HP Indigo's Smartstream workflow offering. The codename HIPIE stands for the original "HP Indigo Photo Image Enhancement" name, with HIPIE 2 being its second version. This paper describes the various modules of HIPIE 2, giving an overview of the technology inside the product, and its usage.*

(a) (b)

Figure 1. *Image enhancement by the HP Smartstream Photo Enhancement Server. (a) Original image, (b) after processing.*

Introduction

With the advent of the digital press, new types of print jobs are emerging, among them *photo specialty prints*, such as photobooks, customized posters, and calendars. These applications typically handle consumer photographs, which are not as optimized for print as the professional photographs used in traditional jobs. Consumer pictures are often too dark or too bright, not colorful enough, and/or not sharp enough, for instance. As a result, print service providers (PSPs) that fulfill photo specialty workflows need to improve source *image quality* as well as ensuring high print quality.

HIPIE 2 is an automatic image enhancement package, which works as part of a commercial print environment, enhancing each incoming photograph in a fast and reliable way (see Fig. 1). It improves the visual appearance of the photos as they are printed in photo specialty products, without human intervention. It is commercialized by HP Indigo as part of their photo quality offering, and as part of the HP Smartstream workflow offering, under the name "HP Smartstream Photo Enhancement Server".

HIPIE 2 contains the following modules: **System infrastructure**, which manages and coordinate all the different modules; **Image processing engine**, which is responsible for performing the image enhancement tasks; and **Image analysis modules**, which determine the amount and type of enhancements needed for each individual photograph. This paper is organized such that each section covers one of the above modules.

System Infrastructure

Hot-folder system

HIPIE 2 is intended to fit seamlessly into high-throughput photo specialty workflows, which has a number of implications. First, it must be robust, meaning that it must be able to run nonstop for many months, and it must be able to recover automatically from any possible failures. Second, it must be scalable, meaning that it must be able to utilize any number of processors in a distributed computing environment. Third, it must be easy to configure and manage.

HIPIE 2 uses a master-slave architecture and has five components. The *Windows service* is used to (automatically) start and stop the other components of the system, and there is a Windows service component on each host in the cluster. *Hot folder processes* are responsible for "watching" each configured hot folder, and ensuring that each input file is processed appropriately by a worker process. The *queue process* matches work requests from hot folder processes with work offers from worker processes, thus distributing work from the various hot folders throughout the cluster to worker processes. The *worker processes* simply request work from the queue process, enhance the image using the options passed as part of the request, and return the completion status to the requesting hot folder when done. A *logger process* per host in the cluster is used to record log messages from the various components in the system.

In practice we have found that this architecture is very robust and scalable. HIPIE has run in production environments for months on end without restarts, and in a scalability test its performance scaled linearly with the number of processors (48 processors tested).

Control Parameters

The user has access to a number of control parameters, which they can set according to their preferences. The idea is that the user tunes the system, and then HIPIE 2 will work automatically on each input image, according to the "recommendations" set by these controls. The control parameters were not designed for image-by-image tuning operation. The control parameters are: Contrast, Sharpen, Noise Reduction, Shadow Details, Face Color

Correction, Face Make-up, 3D Boosting, Chroma Boost, JPEG Artifact Reduction, Red-Eye Reduction (on/off), and Resolution Enhancement. More on these tasks below.

PDF File Handling and Image Formats

One of the strengths of HIPIE 2 is its ability to deal with PDF input files. The PDF file format was developed by Adobe Systems for application and device independent file processing. A PDF document may contain many different kinds of elements, such as text, fonts, images, tables, forms, etc. Within a document, a tree structure consisting of pages and page elements can be used to pull out enhanceable images. We use the PDF document processing API from Datalogics for most of the PDF handling.

In a typical PDF file, images are encoded as separate objects or elements within the file. This is convenient for enhancement because, even though there may be several images on a page, each image can be handled by HIPIE 2 and enhanced independently from the other images on the page. Additional page analysis, image reconstruction and document optimization are performed for sliced and reused images in the document [1]. An exception is *flattened PDF files*, where all the elements of each page are merged into a single image; in this case, HIPIE 2 performs what we call *unflattening*, which is an image processing operation where the different images within the flat image are detected and separated [2].

HIPIE 2 handles images in JPEG, JPEG 2000, and TIFF formats.

Color Management

Images processed by HIPIE 2 can arrive from a wide variety of sources such as digital cameras, scanners, image editing application (i.e. Photoshop), and others. Consequently, the color domains of images vary. The image color might be encoded in RGB or in CMYK values, and, even within the same domain, the same RGB or CMYK values can be device-dependent (i.e., they may have different color meaning, depending on the device they are representing). In order to cope with this variability, HIPIE transforms the device-dependent color values into sRGB, which is the domain in which the image processing is done. HIPIE 2 uses the industry-standard ICC color management technology for this purpose: The software detects the image color domain from the ICC color profile embedded in it, and can output enhanced images to any user-defined color domain, including CMYK.

Other Features

Nested hot folders. Version 2 enhances an entire directory tree inside the input folder. This feature enables customers to put an entire job in one inner folder, and continue working with their desired workflow mechanism without any change.

Job ticket mode. This optional feature enables tracking and controlling jobs in the enhancement server. This is accomplished by using a unique xml file, which can be read by the server, and which contains information about the job's images - their name and path. When this feature is used, the server will only enhance images that are included in this xml file. After all images are enhanced an output file is created, containing status information ("success", "error", or "not enhanced") on every image . It also contains information about the job – start time, end time, error description.

(a) (b)

Figure 2. *3D boosting. (a) Input image, (b) output.*

SDK. The server's capabilities can be exposed as C++ API. Customers that have special requirements are able to implement them using the SDK.

Image Processing Engine

This section summarizes the various image processing tasks performed by HIPIE 2:

- **Sharpening.** Each input image is sharpened as needed, so that it appears at the right sharpness in print.
- **Noise Reduction.** Noise is estimated, and then reduced, without damaging texture or fine details in the image.
- **3D Boosting.** This is a novel feature [3], whereby shadings are enhanced so to give an impression of increased depth to the objects in the image (see Fig. 2).
- **Contrast Enhancement.** Low-contrast images have their effective dynamic-range stretched; this includes low-exposure images, which are brightened (see Fig. 3).
- **Shadow Detailing.** HIPIE 2 detects the low-exposed regions of the image, and brightens them, without damaging the well-exposed regions of the image (see Fig. 4).
- **Face Make-up.** Another novel feature [4], this makes sure that skin on faces are smoothed, as if the subject applied make-up, reducing the visibility of wrinkles and minor blemishes.
- **Chroma Boosting.** HIPIE 2 enhances overall image chromaticity.
- **JPEG Artifact Reduction.** Blocking and ringing artifacts, typical from strong JPEG compression, are reduced.
- **Face Color Correction.** Skin color is analyzed and corrected for photos with a color cast.
- **Resolution Enhancement.** If the image resolution is too low, compared to the target print size, then the resolution is doubled by a smart interpolation algorithm in order to avoid jagginess or pixelization on edges.
- **Red-Eye Removal.** Optionally, HIPIE contains a module for automatically detecting and removing red-eye effects.

All these tasks are completely automatic, but their strengths can be influenced by the *user controls*, as described above in Section "Control Parameters." In the remainder of this section, we describe how each task is accomplished in HIPIE 2.

(a) (b)

Figure 3. *Brightening by HIPIE 2. (a) Input image, (b) output.*

(a) (b)

Figure 4. *Enhancement of shadow details by HIPIE 2. (a) Input image, (b) output. Notice that, despite the relatively strong brightening of shadowed regions, the bright regions (including sky) remain approximately with their original brightness.*

"Unified Scheme"

The main image processing engine in HIPIE 2 is what we call *the "unified scheme"* [5], a novel algorithm that integrates most image processing tasks (sharpening, noise reduction, local and contrast enhancement, shadow detailing, 3D boosting, make-up, and chroma boosting) into a single module. Instead of just putting together existing modules, the "unified scheme" consists of an integrated algorithm, which performs several tasks simultaneously, without losing the ability to control (or turn off) each task separately. This algorithm is compact (thus fast and simple), yet nevertheless it applies novel image processing techniques and ideas, yielding superior image-quality results.

The heart of the unified scheme is the computation and use of a **photographic mask**. This mask is a grayscale, simplified version of the input image, where all low-contrast details have been removed. Fig. 5 shows an example of a mask. It has sharp delineations of high-contrast edges, but is very smooth within homogeneous regions.

The "unified scheme" consists of two steps: The **multiscale processing** step computes two images: *i)* the photographic mask, and *ii)* the output of *sharpening, denoising, 3D boosting, and make-up* tasks. In short, this is how it is done: The image luminance undergoes a multiscale nonlinear decomposition, which decomposes the image elements into details of different scales. On one hand, by removing all details with contrast below a given high threshold (around 50 graylevels), and reconstructing the image, one obtains the mask. On the other hand, by enhancing the details with contrast above a low threshold and removing the details with contrast below that same threshold, one obtains 3D boosting, sharpening and denoising. This low threshold is computed

Figure 5. *The photographic mask of Fig. 1(a), computed and used by HIPIE 2 in order to produce the output in Fig. 1(b). The mask is the heart of the "unified scheme". Notice how all details and shades are removed (e.g., all face features), but the edges are kept sharp – the removed details/shades are then enhanced by HIPIE 2. The sharp edges of the mask enable brightening dark regions without creating halos, or damaging other regions of the image.*

per-image, as a function of an estimated noise level. Make-up is obtained by decreasing the details on a face-skin region (the computation of which is described in a subsequent section).

The **3D-Preserving contrast enhancement** step performs *global and local contrast enhancement*, including adaptive lighting (shadow lighting and highlight detailing), as well as histogram stretching and brightening/darkening. This step is designed to preserve the "3D perception" of images as one modifies its contrast [as opposed to the traditional straightforward use of a 1D look-up table (LUT), which enhances some regions on the expense of others, thus flattening some regions]. This is accomplished with the assistance of the photographic mask image, by separating the "details" (the difference between the sharpened image and the mask) from the coarse image (the mask itself). The 3D-preserving mechanism then applies a 1D LUT to the mask only, whereas the details are multiplied separately by a factor larger or equal to 1. Then the two channels are added back to produce the output image.

JPEG artifact reduction

The HIPIE JPEG artifact reduction filter [6] is a generalization of Nosratinia's scheme [7], which belongs to the family of over-complete wavelet shrinkage filters (transform-domain filters). While Nosratinia's filter suppresses blocking artifacts very well, it suppresses ringing artifacts much less. Samadani *et. al.* [8] proposed an ad-hoc strategy to improve the ringing suppression by selectively averaging only part of the shifts that are locally less active. Their algorithm considerably improves the ringing suppression over Nosratinia but at a steep computational price. Our algorithm, similarly to Samadani, employs a similar strategy of selective shift averaging, but in different setting which provides similar de-ringing as Samadani but at a considerably lower computational price ($<$ 50% overhead relative to Nosratinia). In HIPIE, JPEG artifact reduction is applied within the decompression phase of JPEG images, by filtering each color channel separately before (potentially) upscaling any downsampled channels.

Face color correction

In the evaluation of image reproduction quality, three color areas are traditionally considered to be the most important: skin, foliage and sky. However, the color reproduction of these areas is often unsatisfactory due to imperfections of color processing of images in digital cameras (such as white balancing). HIPIE 2 includes a solution for automatic correction of some memory colors in digital images, with the emphasis on "preferred" color reproduction.

We follow the method for automatic image analysis and segmentation into Regions of Interest (ROI) according to memory colors that has been developed and described in [9]. The method results in an image-specific model of the memory color regions, which includes a) an 8 bit-per-pixel mask, in which every pixel's value is proportional to the probability that the corresponding image pixel belongs to the memory color area (i.e., skin, face, foliage); and b) the mean and the standard deviation of the distribution of colors in memory color region in the CIELAB color space.

We augment the above method by using a face detection algorithm, which finds occurrences of faces within an input image. We then check whether the color of the pixels within the face correspond to the face skin model. If not, then the cast color of these pixels are moved towards the center of the model.

Resolution Enhancement

When the resolution of the photograph is low, compared to the printing size, then jagginess (or pixelization) may be noticed on contrasted edges within the printed photograph. Blocking artifacts on smoothly varying areas may also appear. In this case, HIPIE 2 enhances the resolution of the photograph (using the Resolution Synthesis algorithm [10]), so edges appear smooth and smooth areas show no blocking artifacts on the print. Resolution enhancement is activated only for low resolution images.

Red-Eye Removal

Red eye is a common defect in images photographed with a flash, whereby the light that is reflected from the eye's retina into the camera lens colorizes the eye pupil in a bright red color. Automatic correction of this defect consists of two stages: The detection, whereby the location and the size of the red eye are determined, and the color correction. In HIPIE 2, we provide an innovative solution, which uses the bounding rectangle as the basis for the creation of the correction mask, while better naturalness is achieved by combining the color information about the red eye with an artificial model of the pupil-iris transition and with the image lightness statistics. The red eye detection is described in the *Image Analysis* section.

Image Analysis

In this section, we list the various image analysis modules in HIPIE 2. The results of this analysis are used in tuning the parameters of the enhancement tasks presented above.

Noise Estimation

HIPIE 2 uses a noise estimation method described in [11]. Its main idea is that local variances in an image are bounded from below by the noise variance. In fact, featureless areas have the same local variance as the noise, whereas the local variances of other areas containing signal features are larger due to the signal contribution. Therefore, we involve only smooth areas in the noise estimation process. Furthermore, we assume spatially varying, intensity-dependent noise typical in digital cameras. In particular, for each input image we estimate the noise profile, that is the noise variance as a function of local lightness. Our algorithm divides each image into non-overlapping blocks, compute

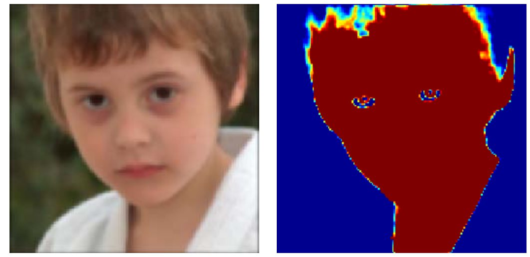

(a) (b)

Figure 6. Skin map. (a) original image, (b) final skin map in face region.

their mean and variance, clusters them into L intervals (typically, L=10) according to the mean, and for each interval, statistically derives a noise variance from a set of up to 30 *lowest* variances in the cluster [11]. We then interpolate the L variance values into the noise profile (LUT) of 256 values for each grey level.

Sharpness Estimation

We define the sharpness of an image to be the sharpness of the sharpest object in the image (since the background may well be out of focus, and still the image appear sharp). Our approach [12] is based on relative frequency content on dominant image features. Assuming that usually the most dominant features in the image are straight lines, relative frequency content is scene-independent (i.e., the relevant subject is always a straight line). The sharpness estimator is then based on the integral of the squared ratio between High Pass and Band Pass IIR (Infinite Impulse Response) filters, computed on the image rows and columns.

Face Detection

The face detection in HIPIE 2 [13] is an adaptation of the HP Labs multi-view face detection technology, which is itself a multi-view extension of the single (portrait) view face detector described by Viola and Jones in [14].

The multi-view detector consists of a parallel set of cascades, each of which corresponds to a different face *view* (such as profile, inverted, 30^o rotated etc), and a patch that passes successfully through every stage of a cascade is labelled as a face in that cascade's view. During normal operation, the detector evaluates all the cascades in parallel and returns either a *not-face* verdict, or the view resulting from a successful cascade evaluation.

Skin Map

HIPIE 2 uses a skin map in order to produce smoother skin results in the output print [15]. Our skip map algorithm learns the skin region in a detected face box, and uses this region to extract a model for the feature (say, color) probability of the captured person skin. Given an image, we start our search with the face detection algorithm mentioned above, and analyze each of the detected faces separately. For each face, a skin probability map is created (Fig. 6).

Sky Detection

Like skin, HIPIE 2 seeks to produce smoother skies [16]. The sky detection in HIPIE 2 is based on the following intuition and common knowledge we have about sky. Sky in an image should cover a large region at the top of the image. Its color should be blue, or gray, and it should have other sky characteristics, e.g., low texture and high luminance. The algorithm looks for large regions which answer those specifications. An important part of this algorithm is the statistical models we have computed for a variety of sky characteristics. In addition to the general sky color model from [17], we have developed a separate color model for blue sky and for gray sky. We also collect statistics for texture and gradient correlation of gray and blue sky. Our experiments have shown that collecting separate statistics for blue and gray sky significantly improves the ability to detect sky.

Red-Eye Detection

The redeye-detection scheme included in HIPIE 2 builds on two previous red-eye algorithms [18, 19], which have different strengths, and leverages advantages of both. The first algorithm [18] uses multi-scale face detection [14], which greatly simplifies the task of finding a red-eye, but can fail if faces with red-eye artifacts cannot be detected first. The second algorithm [19] is attractive because of speed and because it is not tied to a specific facial orientation. One potential problem with this approach that is shared with other classification-based methods [20, 21], however, is that it does not prevent false positives from occurring within faces.

Because the HIPIE 2 red-eye detection algorithm builds on two others, its performance has been evaluated both directly and indirectly. The face-detection based algorithm [18] was recently compared with a number of approaches, both from industry and academia [22]; experiments revealed that this algorithm resulted in the fewest number of false positives and the second highest number of detected artifacts. In the same way, internal testing based on thousands of annotated red-eye images has revealed that the classification-based approach [19] outperforms other industry alternatives. The system in HIPIE 2 combines aspects of these previous algorithms and outperforms both in terms of detection rate and robustness.

Related Work.

Another image enhancement package at HP is the *Real Life Technologies* (RLT) software, developed by the Inkjet division in Vancouver, WA. Evolutionary versions of it have shipped with HP inkjet printers since 2002. HIPIE leverages the code infrastructure, as well as some of the RLT modules, in addition to HPL original algorithms. RLT is primarily designed and optimized for usage in consumer applications, such as Inkjet printer drivers and photo kiosks, while HIPIE is optimized for usage in a commercial print production environment.

Acknowledgment

We would like to thank Ron Gompertz, Morgan Schramm, Julie Hong Ouyang, Paul Brent Fredrickson, Ranjit Bhaskar, and Stephen J. Nichols from HP Inkjet Systems, as well as Dave Berfanger, HP Image Science Team, for the *RLT* code, from which HIPIE developed, as well as the discussions and technical exchange.

We would also like to thank Shimrit Feldman (former HPL-I intern) for her contributions on backlit detection.

References

[1] Hui Chao, Carl Staelin, Sagi Schein, Marie Vans, John Lumley: PDF document restoration and optimization during image enhancement. ACM Symposium on Document Engineering 2008: 150-153

[2] S.J. Simske, D. Li, and J.S. Aronoff, "A Statistical Method for Binary Classification of Images", *DocEng'05*, November 2-4, 2005, Bristol, United Kingdom, pp. 127-129.

[3] Renato Keshet, Doron Shaked and Mani Fischer, "Method And System For Enhancing Image Signals And Other Signals To Increase Perception Of Depth", US patent application 11/888573, filed July 2007.

[4] Mani Fischer and Doron Shaked, "Automatic Digital Face Makeup Tool," filed with the US Patent and Trademark Office, October 31 2008.

[5] Renato Keshet, Doron Shaked, Mani Fischer, Pavel Kisilev and Boris Oicherman, "Unified Spatial Image Processing", US patent application 11/888572, filed July 2007.

[6] R. Maurer and C.-H. Staelin, "Processing an Input Image to Reduce Compression-related Artifacts", *US patent application*, Filed Sept 26, 2007.

[7] A. Nosratinia, "Embedded Post-Processing for Enhancement of Compressed Images", *Proc. Data Compression Conference DCC-99, pp. 62-71, 1999*.

[8] R. Samadani, A. Sundarajan, and A. Said, "Reducing Block DCT Compression Artifacts and Optimizing Quantization Matrices for Improved Artifact Reduction", *HP-Labs Tech Report, HPL-2003-162, August 4, 2003*.

[9] Luo MR, Cui G and Rigg B., "The development of the cie 2000 colour-difference formula: ciede2000." *Colour research and application (2001)* 26: pp. 340-350.

[10] C.B. Atkins, C.A. Bouman, J.P. Allebach, "Optimal image scaling using pixel classification," Hewlett-Packard Labs., Palo Alto, CA; *Proceedings 2001 International Conference on Image Processing*, 2001, Vol. 3, pages 864-867

[11] P. Kisilev, D. Shaked, and S.H. Lim, "Noise and Signal Activity Maps for Better Imaging Algorithms," *ICIP 2007*.

[12] D. Shaked, and I. Tastl, "Sharpness Measure – Towards Automatic Image Enhancement," *ICIP'05*, Sep. 2005.

[13] D. Greig, H. Luo and C. Miller, "Speeding up the Viola-Jones Face Detector", *HP Labs Technical Report HPL-2005-69*, May 2005.

[14] P. Viola, M. Jones, "Robust Real-Time Face Detection", International Journal of Computer Vision 57(2), pp. 137-154, 2004.

[15] H. Nachlieli, G. Ruckenstein, D. Greig, D. Shaked, R. Bergman, C. Staelin, S. Harush, M. Fischer, "Face and skin sensitive automatic image," *US Application: 11/754711*, filed May 29, 2007, published under no 20080298704 on 4 Dec 08.

[16] R. Bergman, H. Nachlieli, G. Ruckenstein, "Perceptual Segmentation of Images," *International patent application: PCT/US2008/009311*, filed July 31, 2008.

[17] O. Martinez Bailac, *Semantic retrieval of memory color content*, PhD Thesis, Universitat Autonoma de Barcelona, 2004.

[18] M. Gaubatz and R. Ulichney, "Automatic Red-Eye Detection and Correction", *Proc. ICIP*, 2002, vol. 1, pp. 804-807.

[19] H. Luo, J. Yen and D. Tretter, "An Efficient Automatic Red-

eye Detection and Correction algorithm", *Proc. ICPR*, 2004, vol. 2, pp. 883-886.
[20] J. Willamowski and G. Csurka, "Probabilistic Automatic Red Eye Detection and Correction"', *Proc. ICPR*, 2006, vol. 3, pp 762-765.
[21] P. Huang, Y. Chien and S. Lai, "Automatic Multi-Layer Red-Eye Detection", *Proc.ICIP*, 2006, pp. 2013 - 2016.
[22] F. Volken, J. Terrier and P. Vandewalle, "Automatic Red-Eye Removal Based on Sclera and Skin Tone Detection", *Proc. IS&T Color in Graphics, Imaging and Vision*, 2006, pp. 359-364.

Author Biography

Renato Keshet *has a BSc degree from IME, Rio de Janeiro, Brazil, and a DSc degree from the Technion–Israel Institute of Technology, both in EE. Renato joined HP Labs–Israel in 1995, where he researches image processing and analysis.*

Carl Staelin *received his PhD in CS from Princeton University (1991) in high performance file system design. Since then he has worked at HP Labs in Palo Alto, CA, and Haifa Israel. His recent work has focused on analysis-based automatic image enhancement. He is a member of IEEE, ACM, and USENIX.*

Boris Oicherman *has received his PhD in colour science from the University of Leeds, UK, and MSc in Digital Imaging from the London College of Communication. Boris is a colour scientist in HP Labs in Haifa, Israel, his research is focused on affective imaging, computational photography and smart colour management.*

Mani Fischer *received his BSc in Computer Science in 1986, MSc and PhD in 1988 and 1992 in Mechanical Engineering, all received from Technion–Israel Institute of Technology. After a post-doc in Cornell University, and seven years at Kulicke & Soffa, Mani joined HP labs in 2002, as a senior researcher in imaging and printing.*

Pavel Kisilev *received his BSc in EE and CS from the Ben Gurion University of the Negev (1994), and his MSc (1998) and PhD (2002) in EE from the Technion. After two years at Tensor Systems Ltd., and two years at the Vision Lab at the Technion, Pavel joined HP Labs–Israel in 2003. His main interests are image and video analysis and processing, and biological signals and systems.*

Sagi Schein *holds a PhD degree in computer graphics and geometric computing from the Technion, Israel Institute of Technology. After working at Qualcomm Israel and KLA-Tencor, Sagi joined HP Labs in March 2006. His main research activity focuses on developing image processing and computer vision on commodity graphics cards (GPUs).*

Doron Shaked *received his BSc in Electrical and Computer Engineering from the Ben Gurion University, Israel (1988), and his DSc in EE from the Technion, Israel (1995). Since 1995 he is with HP Labs–Israel, doing research in Color Printing and Image Processing. Other research interests include: Computer Vision, Pattern Recognition and Shape Analysis.*

Marie Vans *received her PhD in Software Engineering from Colorado State University (1996). Since then she has worked at HP Labs–Israel. Her work has focused on printing and imaging technologies. She is a member of IEEE.*

Hila Nachlieli *received her PhD in Physics from the Technion, Israel, in 2001. She works for HP Labs Israel, and her research interests include image processing, visual inspection tools, and voting theory.*

Ruth Bergman *joined HP Labs in 2001. She received a PhD from MIT in the area of machine learning in 1996. Prior to joining HP she was a member the MIT Lincoln Laboratory (1996-1998) and of NASA's Jet Propulsion Laboratory (1998-2001). Areas of interest include image processing, machine learning, data mining and reasoning.*

Ron Maurer *received his BSc, MSc, and PhD in physics from the Technion (1990, 1994, and 2000). Since 1998, he has worked in HP Labs–Israel. He works on a variety of topics in image processing and analysis.*

Michal Aharon *received her BSc, MSc and PhD from the CS dept., Technion, Israel, in 2001, 2004 and 2006 respectively; during her studies she worked at Intel and IBM. Since her graduation, she has been working as a researcher in HP labs–Israel, focusing mainly on signal processing and color science.*

Alexander Berkovich *received his BSc in EE and BA in physics from the Technion (2007). He is currently a MSc student in the Technion in the EE dept., and working in HP Labs–Israel, in the field of image processing.*

Ran Waidman *received his BSc in Physics and CS at 2000. After working for six years as a software engineer in CTI Squared, he joined HP Indigo in 2005, as a physicist and algorithm developer He is now also in charge of the photo enhancement server software applications.*

Shlomo Harush *received his MSc in physics from the Hebrew university in Jerusalem (1988) and his MBA from Tel Aviv University (1996). Shlomo has working in Research and Development in printing industry since 1989 - First in Indigo Graphic system then in Indigo and now in HP.*

Oren Bengigi *has been developing software over a decade, in a wide variety of software development technologies. He develops and implements color control algorithms and takes part in the development of the HP Photo Enhancement Server at HP Indigo in Israel. He has a B.Sc. in Computer Science from Bar Ilan University, Israel 2007.*

Gal Amit *received his BSc both in physics and in CS from the University of Tel Aviv (1999) and his Msc in theoretical physics from Tel Aviv University (2003). Since then he has been working as a physicist in the Colors and Algorithms group at HP Indigo, Israel. His work has focused on the development of printing process control and color control algorithms.*

Alon Asher *received his BSc in CS (1999) from the Technion, and MSc in engineering sciences (2006) from the Tel-Aviv University, Israel. Since then he has worked at HP, Indigo division. His work has focused on algorithm research and development for press industrial printing and imaging systems.*

Hui Chao *received her BS in Electrical Engineering from Wuhan University, China, and her PhD in Medical Imaging from Massachusetts Institute of Technology. She has worked in HP Lab, Palo Alto, CA, for eight years. Her work has focused on document and image analysis and repurposing.*

Darryl Greig *received his BSc (hons) in Mathematics from the University of Western Australia (1992) and his PhD in probability theory from the University of Melbourne (1995). He has been working with HP Labs since 1997 on machine learning technologies, image understanding and consumer photo applications.*

Matthew Gaubatz *received his BA in applied mathematics from the University of California, Berkeley (2000) and his PhD in electrical engineering from Cornell University (2006). He has since been with HP Labs working on automatic red-eye removal and analysis of print-based authentication. He is a member of the IEEE Signal Processing Society.*

Steven Simske *received his BS in biomedical engineering from Marquette University (1986), his MS in Biomedical Engineering from Rensselaer Polytechnic Institute (1987) and his PhD in EE from the University of Colorado (1990). He has worked on a wide variety of biomedical and imaging projects for the University of Colorado, Colorado School of Mines, HP and HP Labs since.*