



Published in final edited form as:

Proc IEEE Int Symp Biomed Imaging. 2018 April ; 2018: 34–38. doi:10.1109/ISBI.2018.8363517.

COMPUTER-AIDED DETECTION OF PATTERN CHANGES IN LONGITUDINAL ADAPTIVE OPTICS IMAGES OF THE RETINAL PIGMENT EPITHELIUM

Jianfei Liu, HaeWon Jung, and Johnny Tam

Ophthalmic Genetics and Visual Function Branch, National Eye Institute, National Institutes of Health, Bethesda, MD, USA

Abstract

Retinal pigment epithelium (RPE) defects are indicated in many blinding diseases, but have been difficult to image. Recently, adaptive optics enhanced indocyanine green (AO-ICG) imaging has enabled direct visualization of the RPE mosaic in the living human eye. However, tracking the RPE across longitudinal images on the time scale of months presents with unique challenges, such as visit-to-visit distortion and changes in image quality. We introduce a coarse-to-fine search strategy that identifies paired patterns and measures their changes. First, longitudinal AO-ICG image displacements are estimated through graph matching of affine invariant maximal stable extremal regions in affine Gaussian scale-space. This initial step provides an automatic means to designate the search ranges for finding corresponding patterns. Next, AO-ICG images are decomposed into superpixels, simplified to a pictorial structure, and then matched across visits using tree-based belief propagation. Results from human subjects in comparison with a validation dataset revealed acceptable accuracy levels for the level of changes that are expected in clinical data. Application of the proposed framework to images from a diseased eye demonstrates the potential clinical utility of this method for longitudinal tracking of the heterogeneous RPE pattern.

Index Terms

Adaptive Optics Retinal Imaging; Pictorial structure; Belief Propagation; Indocyanine Green; Superpixel

1. INTRODUCTION

The retinal pigment epithelium (RPE) is an important tissue layer that nourishes and maintains retinal visual cells necessary for vision. RPE defects can result in loss of vision in disease [1]. Recently, we have shown that adaptive optics enhanced indocyanine green (AO-ICG) imaging can be used to visualize the RPE cell mosaic in living human eye [2]. The RPE mosaic has a heterogeneous appearance consisting of dark, gray, and bright patches (Fig. 1). Detection of changes to this underlying pattern can potentially reveal changes to the RPE indicating disease progression. This paper aims to develop a computer-aided detection framework to identify such changes.

Detection of pattern changes on longitudinal AO-ICG images is challenging for a number of reasons. First, variation in image quality or the locations of the imaged areas themselves can

occur across a time window of several months. These are exacerbated by local distortions due to scanning or eye motion-induced artifacts, and also errors in the registration of overlapping images within the same imaging session. Finally, the RPE pattern is often degraded by noise. To overcome these challenges, we developed a coarse-to-fine search strategy to detect pattern changes. Image displacement is first estimated through graph matching of affine invariant maximal stable extremal regions (MSER) [3] in affine Gaussian scale-space. Next, patterns are represented as superpixels and then simplified to pictorial structures. Constrained by image displacement, local superpixel displacements are estimated through a tree-based belief propagation. Finally, these displacements guide superpixel pairing, which can be used to measure their pattern changes in longitudinal AO-ICG images.

2. METHODS

2.1 Image Displacement Estimation

The purpose of this step is to estimate image displacements between longitudinal AO-ICG images. These displacements are used to initialize the estimation of local pattern displacements described in the next section, and to automate the determination of search ranges.

AO-ICG images are first smoothed to reduce noise using total variation flow [4]. The MSER detector is then used to search for dark and bright regions in the smoothed images (green ellipses, Fig. 1). SIFT feature descriptor [5] is employed to describe each MSER, and image displacements are estimated by matching feature descriptors. However, it is inaccurate to directly establish feature descriptors on MSERs because image distortion often causes the same MSER to deform in the AO-ICG image pair, (e.g. Fig. 1, white arrows). To reduce image distortion errors, we adapt affine Gaussian scale-space [6,7] to estimate local affinity and compute feature descriptors on the affine invariant MSERs.

Let $I(\mathbf{p})$, $\mathbf{p} = (x, y)^T$, be an AO-ICG image,

$$L(\mathbf{p}; \Phi_t) = g(\mathbf{p}; \Phi_t) * I(\mathbf{p}). \quad (1)$$

Here, Φ_t is a symmetric positive semi-definite matrix,

$$g(\mathbf{p}; \Phi_t) = \frac{1}{2\pi\sqrt{\det\Phi_t}} e^{-\mathbf{p}^T \Phi_t \mathbf{p}/2}. \quad (2)$$

Assuming $S(\mathbf{p}_1)$ and $S(\mathbf{p}_2)$ are two corresponding MSERs with \mathbf{A} as their affine transformation matrix, we can get $\mathbf{p}_2 = \mathbf{A}\mathbf{p}_1$. Supposing $\Phi_{t,1}$ and $\Phi_{t,2}$ are their affine Gaussian scale matrices, then according to affine Gaussian scale-space [6]:

$$\Phi_{t,1} = \mathbf{A}^T \Phi_{t,2} \mathbf{A}. \quad (3)$$

To estimate Φ_p , we apply the second moment matrix \mathbf{M} to compute local affinity.

$$\mathbf{M}(\mathbf{p}; \Phi_t) = g(\mathbf{p}; \Phi_s) * (\nabla L(\mathbf{p}; \Phi_t)) \cdot (\nabla L(\mathbf{p}; \Phi_t))^T. \quad (4)$$

Here, the integration scale matrix is defined as $\Phi_s = \alpha \Phi_t$ where α is a constant. Again, we can derive

$$\mathbf{M}_2 = \mathbf{A}^{-T} \mathbf{M}_1 \mathbf{A}^{-1}. \quad (5)$$

According to Eq. (3) and (5), we can get $\Phi_t = \beta \mathbf{M}^{-1}$, where β is a constant. The affine transform matrix \mathbf{A} can be computed in terms of \mathbf{M} ,

$$\mathbf{A} = \mathbf{M}_2^{-1/2} \mathbf{R} \mathbf{M}_1^{1/2}. \quad (6)$$

\mathbf{R} is an orthogonal matrix which includes rotation and mirror transformations. Corresponding MSERs can be rewritten as

$$\mathbf{p}_2 = \mathbf{M}_1^{-1/2} \mathbf{R} \mathbf{M}_1^{1/2} \mathbf{p}_1 \Rightarrow \mathbf{M}_2^{1/2} \mathbf{p}_2 = \mathbf{R} \mathbf{M}_1^{1/2} \mathbf{p}_1. \quad (7)$$

It is important that they should contain the same image contents under affine image distortion. We can adjust corresponding MSERs as $\bar{\mathbf{p}}_1 = \mathbf{M}_1^{1/2} \mathbf{p}_1$ and $\bar{\mathbf{p}}_2 = \mathbf{M}_2^{1/2} \mathbf{p}_2$ because corresponding adjusted MSERs only have orthogonal transforms which are affine invariant. This allows us to build accurate SIFT feature descriptors on the adjusted MSERs.

Graph matching [8] is used to find affine invariant MSER correspondences. Besides similarity values between SIFT feature descriptors, graph matching also combines geometric compatibility between MSER correspondences and uniqueness of correspondence into a unified framework. Together, these improve the matching in the presence of image distortions. Fig. 1 illustrates the final MSER corresponding results (red lines). Image displacements can be represented as $\mathcal{L}_g = \{\mathbf{l}_{g,1}, \dots, \mathbf{l}_{g,m}\}$, where $\mathbf{l}_{g,i}$ is the displacement vector between i -th corresponding MSERs.

2.2 Superpixel Displacement Computation

Image displacements provide an initial estimation of pattern correspondences in longitudinal AO-ICG images. Unfortunately, they often produce inaccurate correspondences due to local image distortions, which thereby necessitates a more refined computation of local pattern displacements. However, instead of computing all possible point correspondences, we introduce a tree-based belief propagation method that utilizes image displacements to initialize the search range of pattern correspondences, significantly reducing computational cost.

The heterogeneous pattern can be represented as a collection of smaller regions with homogeneous intensity. Homogeneous regions can be extracted using superpixels [9]. Estimation of local image displacements can thus be solved via determination of superpixel pairings between two images (e.g. Figs 2A and 2B). This problem can again be solved using graph matching [8]. However, graph matching used in estimating image displacement enforces unique correspondence, and the problem here involves many-to-one superpixel correspondence. Instead, we treat this problem as the minimization of pictorial structure through belief propagation [10], and use computed superpixel displacements to guide the subsequent pairing of superpixels.

Let $G = (V, E)$ be the graph structure in Fig. 2A, and $V = \{v_1, \dots, v_n\}$ their node representations, where v_i corresponds to i -th superpixel. E includes all adjacent relationships between two superpixels in Fig. 2A. A configuration $\mathcal{L}_s = \{l_{s,1}, \dots, l_{s,n}\}$ specifies a displacement vector for each superpixel $v_i \in V$. Superpixel displacement computation can be described as:

$$\overline{\mathcal{L}}_s = \arg \min_{\mathcal{L}_s} \left(\sum_{v_i \in V} D(l_{s,i}, I) + \sum_{(v_i, v_j) \in E} F(l_{s,i}, l_{s,j}) \right), \quad (8)$$

where $D(l_{s,i}, I)$ measures how well the superpixel in Fig. 2A matches Fig. 2B when moved by $l_{s,i}$. To reduce the computational cost, the census transform [11] is used as the similarity measurement for $D(l_{s,i}, I)$. $F(l_{s,i}, l_{s,j}) = |l_{s,i} - l_{s,j}|^2$ measures how continuous two displacement vectors from adjacent superpixels are with each other.

However, minimizing an arbitrary graph of Eq. (8) is an NP-hard problem. Additionally, the continuity assumption is often not fulfilled between adjacent superpixels. We reduce the graph structure in Eq. (8) into a minimum-spanning tree [12] (Fig. 2C). The root is arbitrarily selected (red circle, Fig. 2C). Weights of the graph structure are census transform values because we assume that adjacent superpixels with similar visual appearance are more likely to have similar displacement vectors in Fig. 2A.

Dynamic programming is used to efficiently minimize Eq. (8) on the tree structure. Starting with a leaf node v_j , the matching value of the best displacement vector of $l_{s,j}$ given the displacement vector $l_{s,i}$ of its parent node v_i is

$$Y_j(l_{s,i}) = \min_{l_{s,j}} \left(D(l_{s,j}, I) + F(l_{s,i}, l_{s,j}) \right). \quad (9)$$

Here, $l_{s,i}$ is initialized by an image displacement vector $l_g \in \mathcal{L}_g$ in the previous section where the MSER is close to the current superpixel. Meanwhile, the search range of $l_{s,j}$ is also limited by $\max(\mathcal{L}_g)$. Therefore, we can apply image displacements to initialize the computation and automate the determination of search ranges. Iteratively, we can compute the remaining tree nodes, other than the root.

$$Y_j(l_{s,i}) = \min_{l_{s,r}} \left(D(l_{s,j}, l) + F(l_{s,i}, l_{s,j}) + \sum_{v_c \in C_j} Y_c(l_{s,j}) \right). \quad (10)$$

Here, the last term is known as we always compute child nodes first. Eventually, we reach the root. If all its child nodes are known, we can compute its displacement vector as

$$\bar{l}_{s,r} = \arg \min_{l_{s,r}} \left(D(l_{s,j}, l) + \sum_{v_c \in C_j} Y_c(l_{s,j}) \right). \quad (11)$$

Starting from the root, we can work reversely to determine the displacement vectors of all tree nodes until all leaf nodes are reached by changing the computation from ‘min’ in Eq. (9) and (10) to ‘argmin’. Superpixel displacements are illustrated as white arrows in Fig. 2D.

2.3 Determination of Pattern Changes

Each superpixel in the first visit (Fig. 2A) is translated to the second visit through its displacement vector. This can cover a set of superpixels in the second visit. We choose the one with the largest number of overlapping pixels as its counterpart. The mean intensities m_1 and m_2 are computed for the paired superpixels. If $m_2 > m_1 + \gamma$, pattern changes are marked as having changed from dark to bright; $m_2 < m_1 - \gamma$, bright to dark; otherwise, stable. Here, γ is the sensitivity threshold. Since we are interested in detecting only significant changes (e.g. from bright to dark as opposed to from dark to darker, and vice versa), we empirically set $\gamma = 80$ in this paper.

2.4 Dataset Collection

Research procedures adhered to the tenets of the Declaration of Helsinki and were approved by the Institutional Review Board of the National Institutes of Health. Since AO-ICG was first introduced in 2016 [2], there is limited AO-ICG data available. To the best of our knowledge, we present for the first time longitudinal AO-ICG data in four subjects (3 healthy, 1 patient with late-onset retinal degeneration). The time between imaging sessions was 3–4 months for the healthy subjects and 12 months for the patient. A total of 669 overlapping small field-of-view videos were recorded in these subjects, corrected for eye motion artifacts, and combined to create larger RPE images. We created validation datasets from the healthy subjects to evaluate the robustness and accuracy of the detection framework. For each subject, we selected two RPE regions (6 regions total). First, for each RPE region, we generated an image pair, each with a unique distortion due to eye motion (400×400 pixels). Second, for each RPE region, we artificially introduced two types of modifications to the pattern: (1) changing intensity values by 150 (0–255 intensity range; bright regions were made darker, and dark regions brighter), and (2) inverting the intensities (255– I). For the test dataset, we collected image pairs (1000×1000 pixels) in two RPE regions for each of the four subjects.

3. EXPERIMENTAL RESULTS

3.1 Validation Datasets

Detection accuracy was defined as the ratio between total areas of superpixels in the first image that successfully paired counterparts in the second image (whether or not a change was truly made), divided by the total area of the unchanged image plane. No changes were detected across all six image pairs in the case of the first validation dataset, in which the only differences between images were different sets of image distortion due to eye motion and scanning. This establishes that our approach is robust to local distortion. For the second validation dataset, detection accuracies with increasing levels of changes are summarized in Table 1 (mean \pm standard deviation). The detection accuracy is higher in the case of changing intensities vs. inverting the values. This is potentially due to the census transform being more robust to a constant intensity change than to the inverse change. The overall drop in accuracy with increasing area change ratio is likely due to the creation of substantially different trees due to splitting or merging of superpixel areas resulting from the two operations (change and invert). These results indicate high accuracy for area change ratios within 15%. This drop-off in accuracy at larger change ratios is a tradeoff that is needed for the gain in computational speeds realized by the simplification of the superpixels to the tree structure.

3.2 Test Dataset

For the test dataset, the average computation time was 14.4 ± 2.6 seconds for 1000×1000 pixel AO-ICG images. As expected, detection results on six images pairs from three healthy subjects showed minimal changes across visits (area change ratio of $0.2\pm 0.3\%$, mean \pm standard deviation). Detected changes were typically found near the boundaries of images (Fig 3A–C). There were no regions observed with intensity increase in three healthy subjects, while intensity decrease was observed in two image pairs. In contrast, several changes in the RPE pattern were found in the diseased eye with both intensity increases and decreases (Figs. 3D–3F). The overall area change ratio for this patient was 5.2%. Based on our validation results, we estimate that our algorithm can detect with high accuracy changes to the RPE mosaic when the visit-to-visit area change ratio is within 15%. Therefore, these changes are likely to be indicative of disease progression.

4. CONCLUSION AND FUTURE WORK

In this paper, we developed a computer-aided detection framework to identify pattern changes in a novel dataset of longitudinal AO-ICG images. This framework was based on a coarse-to-fine strategy to estimate image displacement and uses a simple, computationally-efficient method to determine correspondences between superpixel regions across different visits. Our proposed algorithm is robust against local distortions that arise from the combination of scanning and eye motion and is accurate for total area changes across visits of up to 15%. As expected, evaluation in a test dataset of healthy subjects revealed little to no change. Application to two image pairs from a patient revealed a 5.2% change to the RPE mosaic across a one-year time window, suggestive of the rate of disease progression.

The proposed framework sometimes yielded false matches when there was significant local image displacement due to errors in registration between overlapping images or degradation of image quality in portions of the image, particularly for the patient data. Such detection errors could be potentially due to the insufficient discriminative capability of the census transform. Further testing on additional patient datasets will help to further improve the robustness of this algorithm. In the future, development of a more discriminative feature descriptor that is computationally inexpensive will lead to improvements in performance. In addition, evaluation of this framework in a larger AO-ICG database with more healthy subjects and patients will also be important. All in all, longitudinal evaluation of changes to the RPE together with tracking of other retinal neurons [8] may lead to new insights about the onset and progression of blinding retinal diseases.

Acknowledgments

This research was supported by the Intramural Research Program of the NIH, National Eye Institute. We would also like to thank Catherine Cukras, Angel Garced, Gloria Babilonia-Ayukawa, and Denise Cunningham for assistance with clinical procedures.

References

1. Sparrow JR, Hicks D, Hamel CP. The Retinal Pigment Epithelium in Health and Disease. *Current Molecular Medicine*. Jul; 2005 10(9):802–823.
2. Tam J, Liu J, Dubra A, Fariss R. In Vivo Imaging of the Human Retinal Pigment Epithelial Mosaic Using Adaptive Optics Enhanced Indocyanine Green Ophthalmoscopy. *Investigative Ophthalmology & Visual Science*. Aug; 2016 57(10):4376–4384. [PubMed: 27564519]
3. Nister D, Stewenius H. Linear Time Maximally Stable Extremal Regions. *Proc 10th European Conference on Computer Vision*. Oct.2008 5303:183–196.
4. Steidl G, Weickert J, Brox T, Mrazek P, Welk M. On the Equivalence of Soft Wavelet Shrinkage, Total Variation Diffusion, Total Variation Regularization, and SIDs. *SIAM Journal on Numerical Analysis*. May; 2004 42(2):686–713.
5. Lowe DG. Distinctive Image Features from Scale-invariant Keypoints. *International Journal of Computer Vision*. Nov; 2004 60(2):91–110.
6. Lindeberg T. Generalized Gaussian Scale-space Axiomatics Comprising Linear Scale-space, Affine Scale-space and Spatiotemporal Scale-space. *Journal of Mathematical Imaging and Vision*. May; 2011 40(1):36–81.
7. Mikolajczyk K, Schmid C. Scale & Affine Invariant Interest Point Detectors. *International Journal of Computer Vision*. Oct; 2004 60(1):63–86.
8. Liu J, Jung H, Tam J. Accurate Correspondence of Cone Photoreceptor Neurons in the Human Eye Using Graph Matching Applied to Longitudinal Adaptive Optics Images. *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Sep.2017 10434:153–161.
9. Achanta R, Shaji A, Smith K, Lucchi A, Fua P, Susstrunk S. SLIC Superpixels Compared to State-of-art Superpixel Methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. May; 2012 34(11):2274–2282. [PubMed: 22641706]
10. Felzenszwalb P, Huttenlocher D. Pictorial Structures for Object Recognition. *International Journal of Computer Vision*. Jan; 2005 61(1):55–79.
11. Zabih R, Woodfill J. Non-parametric Local Transforms for Computing Visual Correspondence. *Proc 3rd European Conference on Computer Vision*. May.1994 801:151–158.
12. Prim RC. Shortest Connection Networks And Some Generalizations. *Bell System Technical Journal*. Nov; 1957 36(6):1389–1401.

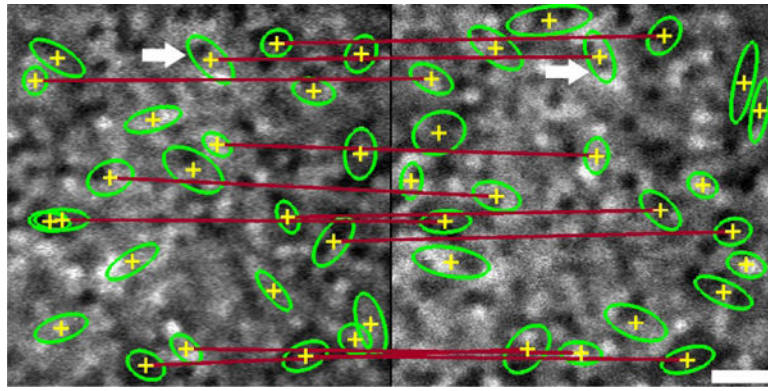


Figure 1.

Image displacement computation in a longitudinal AO-ICG image pair by graph matching of affine invariant MSER correspondences. Green ellipses: MSERs; yellow +', MSER centroids; red lines: MSER feature correspondences. Local image distortion is illustrated by different ellipse shapes and angles corresponding to the same region (e.g. white arrow). Scale bar: 50 μ m.

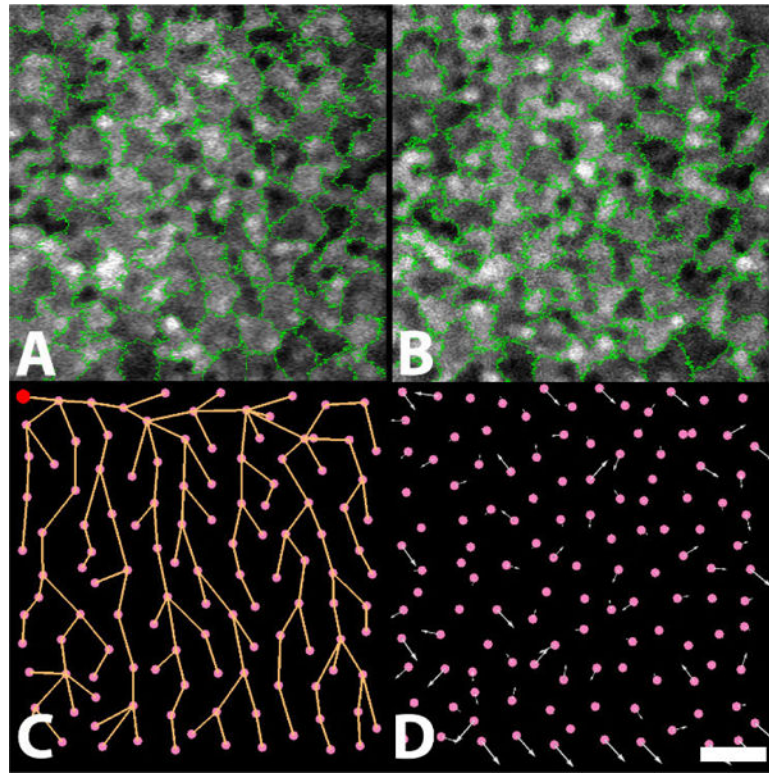


Figure 2. Superpixel displacement computation on a longitudinal image pair. Superpixel decompositions on the AO-ICG image from the (A) first and (B) second visits; (C) tree structure of superpixels from the first visit (red circle, tree root); (D) superpixel displacements of the first visit with respect to the second visit (arrows; lengths indicate amounts of displacement). Scale bar: 50μm.

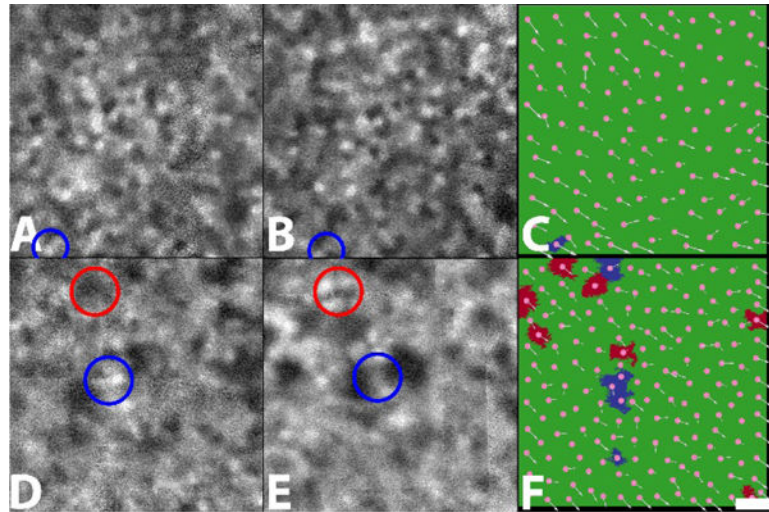


Figure 3.

Detection results of pattern changes on a healthy subject (top row) and a patient with late-onset retinal degeneration (bottom row). Left column: AO-ICG images from the first visit; center column: the second visit after 3 months (top row) and 12 months (bottom row); right column: pattern change maps overlaid with superpixel displacements, where intensity increase (red), decrease (blue), and stable (green). Superpixel displacements indicate the local region movements between AO-ICG images from the first visit to the second. Three regions of pattern changes are highlighted by circles. Scale bar: 50 μ m.

Table 1
Detection accuracy of relative area ratio versus area change ratio of manual intensity operations.

Operation	Area change ratio				
	5%	10%	15%	20%	25%
Change (%)	98.6 ±0.9	97.2 ±1.2	94.9 ±1.9	91.5 ±3.1	85.7 ±2.7
Invert (%)	98.4 ±0.4	96.1 ±0.6	92.2 ±1.1	88.3 ±3.1	83.9 ±3.4