Case for Automated Detection of Diabetic Retinopathy

Nathan Silberman, Kristy Ahlrich, Rob Fergus and Lakshminarayanan Subramanian

Abstract

Diabetic retinopathy, an eye disorder caused by diabetes, is the primary cause of blindness in America and over 99% of cases in India. India and China currently account for over 90 million diabetic patients and are on the verge of an explosion of diabetic populations. This may result in an unprecedented number of persons becoming blind unless diabetic retinopathy can be detected early.

Aravind Eye Hospitals is the largest eye care facility in the world, handling over 2 million patients per year. The hospital is on a massive drive throughout southern India to detect diabetic retinopathy at an early stage. To that end, a group of $10-15$ physicians are responsible for manually diagnosing over 2 million retinal images per year to detect diabetic retinopathy. While the task is extremely laborious, a large fraction of cases turn out to be normal indicating that much of this time is spent diagnosing completely normal cases.

This paper describes our early experiences working with Aravind Eye Hospitals to develop an automated system to detect diabetic retinopathy from retinal images. The automated diabetic retinopathy problem is a hard computer vision problem whose goal is to detect features of retinopathy, such as hemorrhages and exudates, in retinal color fundus images. We describe our initial efforts towards building such a system using a range of computer vision techniques and discuss the potential impact on early detection of diabetic retinopathy.

1. Introduction

Diabetic retinopathy is the leading cause of blindness in adults around the world today (IDF 2009b). The International Diabetes Foundation reports that India has the largest share of this population with over 50 million people and growing rapidly (IDF 2009a). Health care costs motivated by diabetes are also increasing around the world. In the United States alone, projected costs of 376 billion are expected to rise to 490 billion by 2030 (Unwin et al. 2009). Despite the disease's alarming growth on the global scale, diabetic retinopathy often strikes with few initial symptoms before invoking irreversible damage. Many patients are unaware of the problem before its diagnosed.

Given the gravity of the effects of retinopathy, early detection of the disease is absolutely essential in preventing un-

necessary blindness. Surgical replacement of clouded vitreous with saline, laser photocoagulation to prevent clotting or closing off of damaged blood vessels, and steroid injection are three of the most common methods to treat the disease once discovered. Unfortunately, there is no known cure for diabetic retinopathy and these treatments are management strategies at best.

Aravind Eye Hospitals in South India is the largest eye care facility in the world in terms of surgical volume and number of patients treated. In 2007 alone, they treated over 2.3 million out-patients and performed over 270,000 surgeries. Today, the Aravind Eye Care System encompasses five primary and three managed hospitals. Additionally, they perform a wide-range of telemedicine initiatives that covers thousands of rural villages in India on a yearly basis.

Aravind is currently leading a massive initiative in effort to detect diabetic retinopathy at an early stage by scanning every patient for diabetic retinopathy. To accomplish this goal, Aravind sends groups of physicians on a daily basis to various rural locations in South India to manually inspect patients retinas for diabetic retinopathy. Unfortunately, this task is extremely labour intensive and they have too few ophthalmologists to inspect the roughly 2 million retinal images they receive per year.

We describe our early experiences in working with Aravind Eye Hospitals to build an automated system for detecting diabetic retinopathy. Given that a large fraction of screened patients bear no sign of retinopathy(over 95%), any manual process at the scale that Aravind operates is not only extremely time-consuming for its staff but may also be considered wasteful and counterproductive since time spent inspecting healthy patients might be better used helping those already diagnosed with the disease. To make matters worse, the number of ophthalmologists in a country like India is insufficient to support the growing diabetic population. Most rural regions in India do not have a local ophthalmologists and rural populations do not typically travel to the closest city. Consequently, if an automated detection system could be developed, it would significantly improving the efficiency of ophthalmologists in this initiative and would extend basic retinopathy diagnosis to rural regions without the need for an expert.

In this paper, we discuss our progress building such a system. We have developed a model of the different types of
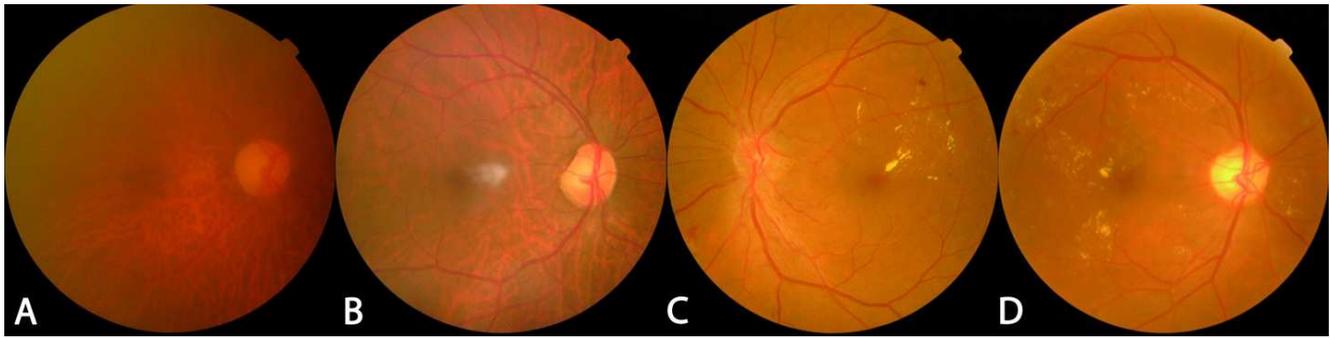
Figure 1: Examples of color fundus images. The first two images show healthy eyes while the last two images contain exudates, manifestations of retinopathy.

retinopathy based on SIFT (Lowe 2004) descriptors. We train a Support Vector Machine classifier to label individual patches of the image and then make an aggregate decision over the entire image based on the patch-level predictions. Our initial evaluation results are promising with regard to the prospect of developing a long term, robust model. The rest of the paper is organized as follows. Section 2 provides a more thorough overview of the Aravind Eye Hospital model and their approach to tackling diabetic retinopathy. Section 3 explains the disease of diabetic retinopathy and how it can be recognized. Section 4 reviews previous attempts to provide automated detection of retinopathy. Section 5 and 6 describe our approach to the problem and our initial evaluation of our efforts.

2. The Aravind Operational Model

To better motivate the need for an automated detection system for diabetic retinopathy, we will briefly describe the existing Aravind operational model and how an automated system can significantly impact their existing operations.

Aravind Eye Hospitals conducts mobile health camps where a group of ophthalmologists and technicians visit new rural location each day. Each mobile team travels with a vehicle typically equipped with a satellite Internet connection for transmitting information to the main hospital. On a typical day, the doctors conduct a village-based rural health camp which attracts between $600 - 1000$ patients from the surrounding area to be screened for eye problems. Due to the size of the country and the number of ophthalmologists involved, each location can only be covered around once every six to twelve months. Given the infrequency of each visit, and the sizable local demand, it is common to see a long line of patients waiting to be screened each day.

Every eye camp consists of a morning and afternoon session. During the morning session, all of the patients are thoroughly screened. In addition to routine eye examinations and tests, Aravind conducts a retinal scan of each patient using an ophthalmoscope to generate a retinal image, which can later be used to diagnose retinopathy. This step must be performed on every patient since many patients are unaware as to whether or not they have diabetes. The retinal images are then transmitted to one of Aravind's main hos-

pitals via satellite link where a group of ophthalmologists manually diagnose the images and provide results within several hours. During the camp's afternoon session, these results are communicated back to the patients who attended the morning session. Every patient diagnosed with a serious case of diabetic retinopathy is encouraged to make an appointment at the main hospital where the retinopathy can be surgically managed.

Over the past few years, we have helped Aravind Eye Hospitals partially in changing the mobile hospital model to setting up rural health centers which connect to the main hospital using Wifi-based Long Distance Networks (WiLD-Net) (Patra et al. 2007). WiLDNet provides very high bandwidth connectivity that enables Aravind to conduct video conferencing sessions between rural patients and doctors in the main hospital. In this model, health workers and nurses at the rural center, collect retinal images of patients and transmit it to the doctor over the wireless link.

Why automated detection?

With regard to a limited medical staff, an automated system can significantly decrease the manual labour involved in diagnosing large quantities of retinal images. While this represents an obvious and significant gain, there is a larger, logistical need for automated and immediate diagnoses in rural settings: patients in rural areas are fundamentally harder to reach than in urban environments. In the mobile hospital setting, if a patient comes for the morning session, there is no guarantee that the same patient will return for the afternoon session unless the need for continued diagnosis is demonstrated. Consequently, if a patient actually has diabetic retinopathy, it is essential to convey the urgency of the situation to the patient as soon as a diagnosis can be reached in order to convince them to travel to a hospital for full treatment. This is a fundamental health care delivery problem in many rural developing regions.

With an automated system, the doctor or local health worker can be made aware of the diabetic retinopathy problem during a single session with a patient. This enables the medical personnel to immediately and visually demonstrate the existing problem to the patient which makes it easier to convince them of the urgency of their situation.

They can also immediately schedule appointments for the patients without further delay for continued diagnosis and follow-up visits at a regular hospital. An automated system also helps local health workers to detect serious diabetic retinopathy cases without the need for local ophthalmology experts. Consequently, in the ideal case, one can perform early detection of diabetic retinopathy cases in rural areas by training health workers and conducting mobile health camps without the need for local experts or even Internet connectivity.

3. Detecting Diabetic Retinopathy

As per the National Institute of Health (NIH 2009), diabetic retinopathy is the most common diabetic eye disease that is caused by changes in the blood vessels of the retina. In some people with diabetic retinopathy, blood vessels may swell and leak fluid. In other people, abnormal new blood vessels grow on the surface of the retina. As illustrated in Figure 1, images of patients with diabetic retinopathy patients can exhibit red and yellow spots which are problematic areas indicative of hemorrhages and exudates. In many retinal image, such as image (C) in Figure 1, a central dark spot represents the *macula* of the eye. The presence of hemorrhages and exudates in this region is indicative of a serious diabetic retinopathy condition that can soon lead to blindness.

In brief, diabetic retinopathy has four stages:

1. *Mild Non-proliferative Retinopathy.* At this early stage, micro-aneurysms may occur. These manifestations of the disease are small areas of balloon-like swelling in the retina's tiny blood vessels.

2. *Moderate Non-proliferative Retinopathy.* As the disease progresses, some blood vessels that nourish the retina are blocked.

3. *Severe Non-proliferative Retinopathy.* Many more blood vessels are blocked, depriving several areas of the retina with their blood supply. These areas of the retina send signals to the body to grow new blood vessels for nourishment.

4. *Proliferative Retinopathy.* At this advanced stage, the signals sent by the retina for nourishment trigger the growth of new blood vessels. These new blood vessels are abnormal and fragile. They grow along the retina and along the surface of the clear, vitreous gel that fills the inside of the eye. By themselves, these blood vessels do not cause symptoms or vision loss. However, they have thin, fragile walls. If they leak blood, severe vision loss and even blindness can result.

From a computer vision perspective, there are four specific indicators within the retinal images. First, *micro-aneurysms* are small red dots in the retinal image which are early indicators of blood clots. These are difficult to detect since they may be captured by at most a handful of pixels in an image and are sometimes indistinguishable from background noise or occluded blood vessels. Second, *exudates* are bright yellow spots which are indicative of concentrated fat deposits in the retina that in turn lead to blood clots and spillage. Third, *hemorrhages* are red spots which refer to
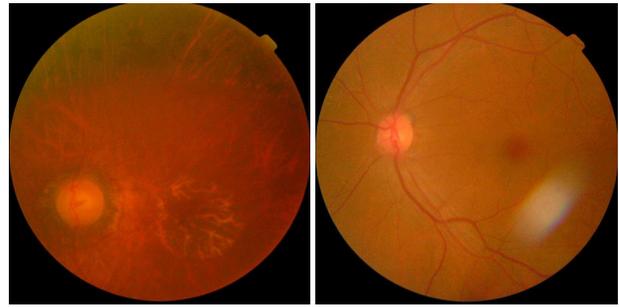

Figure 2: Misleading Retinal Images. The first image contains a patch with the same color intensity as the optic disc and may be mistaken for exudates. However, this component is not related to diabetic retinopathy. The second image contains a large artifact of the ophthalmoscope used to capture the retina and may be confused with retinopathy if examined at too small a scale.

blood stained regions due to bleeding in the retina. Finally, *cotton wool spots* are hazy whitish yellow regions that are caused by damage of nerve fibers.

Detection of these visual elements remains difficult for several reasons. First, hemorrhages and microaneurysms are often difficult to distinguish from blood vessels, even for a trained professional. Larger hemorrhages are of the same color as blood vessels but typically differ geometrically. However, smaller hemorrhages and microaneurysms are often similar in color, geometry and texture to the thin root-like components of the blood vessel network. In a similar vein, exudates and optic discs share similar color spaces. In clear images, the circular geometry of the optic disc can be used to discriminate between it and exudates. For this reason, many previous attempts have first tried to localize and mask the blood vessel network and the optic disc. However, many images are clouded by non-retinopathic features making doing so difficult. Figure 2 demonstrates several examples of non-retinopathic visual elements that can easily be confused with retinopathy. Finally, the color variation among exudates and hemorrhages in a single image is far less than the overall variation among these retinopathy components across all images. To date, no robust color normalization technique has been developed to allow simply color channels and contrast to be sufficient indicators for classifications.

Our work primarily focuses on the detection of exudates in the retinal image and using this mechanism to detect different stages of diabetic retinopathy.

4. Limitations of Existing Works

Several different computational strategies have been used in efforts to solve the problem of automated diabetic retinopathy. These attempts have been limited both by the amount of data available to researchers in this area, and in the variety of methods used to solve the problem. A common theme in the literature of this area has been the need to split the problem up into first identifying the normal features or parts of the

retina, such as the blood vessels, fovea and optic disc, and then attempting to identify and possibly localize exudates and hemorrhages. The need to do so stems from the difficulty in formalizing the difference between similarly colored components such as hemorrhages versus blood vessels and exudates versus optic discs. In certain images, doing so is challenging even for a trained professional.

Much work has been done to attempt to register normal parts of the retina. Pinz et al, (Pinz et al. 1998) use gradient-based techniques and hough transforms to map and localize blood vessels, the optic disc, the scotoma and the fovea. Chaudhuri et al (Chaudhuri et al. 1989) using two dimensional matched filters to map the network of blood vessels in the retina, a technique adopted by many later works. Sinthanayothin et al (Sinthanayothin et al. 1999) find blood vessels by performing PCA on image gradients and inputing the results to a neural network. Additionally, they localize the optic disc through simple intensity variations in image patches and find the fovea through matched filters. Di Wu et al (Di Wu and Bauman 2006) use gabor filters to trace and patch together the blood vessels.

While these methods have had varying success identifying and localizing components of the retina, they typically operate on images with a good amount of contrast, few occlusions of retinal objects and few, if any, manifestations of retinal disease. The last fifteen years has seen, however, a steady increase in the literature that attempts to tackle the problem of detecting not only retinal components in diseased eyes, but also retinopathatic components. Generally speaking, one may split these methods into morphological or rule based methods and machine learning methods which have become increasingly more widespread. Morphological methods include filtering (Saiprasad Ravishankar 2009) and segmentation (Thomas Walter and Erginay 2002).

A number of attempts have been made to use machine learning to automatically locate manifestations of retinopathy. Examples include unsupervised methods such as PCA (Li and Chutatape 2003), k-means clustering (Akara Sopharak and Barman 2009) (Osareh et al. 2001) and gaussian mixture models (Sánchez et al. 2006). Supervised learning techniques have also been attempted including neural networks (Osareh et al. 2001), (García et al. 2009), (G. G. Gardner and Elliott 1996) naive bayes (Sopharak et al. 2008), (Wang et al. 2000) and support vector machines (Chutatape and Zhang. 2004), (Zhang and Chutatape 2005).

Overall, most of these attempts have come from the medical imaging community. There remain, however, many techniques from recent computer vision literature that have not been tried. Conversely, many of the computer-vision based attempts seem insufficient in practical terms because they do not generalize well when presented with the variety of real data observed by practitioners.

5. Automatic Detection of Diabetic retinopathy

The four main constraints for our application were its operating system portability, its ease of use to a non-computer-friendly user, its speed performance and its accuracy in comparison to existing efforts to automatically diagnose diabetic retinopathy.

We chose Java to allow for maximum system portability since anybody who wants to execute the application need only have Java installed on their system. The interface is extremely simple providing the user with the choice of modes: (1) detect retinopathy in a single image, where the results will be displayed onscreen or (2) detect retinopathy for every image in a specified folder, the results of which will be saved as labeled images in a new subfolder. This workflow easily allows a remote medical worker to display clearly to a patient where their eyes are being impacted by the disease or, alternatively, the ability to collect data on many persons at a time and call back only those with retinopathy just minutes later.

Technical Overview: Our detection approach entails training a Support Vector Machine (SVM) classifier to recognize exudates in images. For training data, we were provided with approximately 20,000 unlabeled images from Aravind Eye Hospital. Since these images were unlabeled, we provided pixel-level labels for 41 of these images. These fine-grained labels mark retinopathic manifestations including exudates, hemorrhages, hemes and microaneurysms as well as normal retina components such as blood vessels and optic discs. In addition, we labeled another 3000 images at the image-level where each image-level label specifies whether or not the image contains exudates and/or hemorrhages. For our task, we concentrated on the exudates only. Therefore, a positive label denotes the presence of at least one exudate in an image and a negative label denotes the lack of any exudates.

Preprocessing: Before performing classification, we first preprocess each image in several steps. First, we perform a global color-balancing operation. This results in shifting darker or weakly saturated images to an appearance better amenable to detecting exudates. Next, we increased the luminosity of pixels in the image whose edge energy and "orientedness" (Granlud and Knutsson 1995) in the green channel was high. To extract the edges of each image, we use a customized derivative filter (Simoncelli 1994) that performed far better than standard filters such as Canny and Sobel. Finally, the image is converted to grayscale.

Removal of the Optic Disc: As in several previous works (Sopharak et al. 2008), (Akara Sopharak and Barman 2009), (Saiprasad Ravishankar 2009), we remove the optic disc to aid exudate detection. Unlike previous attempts however, we use an exemplar-based approach. We selected 19 exemplar images of optic discs and computed an average optic disc image. Using this single exemplar, we calculate the normalized correlation against every location in the image at multiple scales. Once we've found the closest match, we mask out a square region at that scale.

Feature Extraction: We extracted SIFT features at 2 scales using overlapping windows of size 13x13 and 25x25. We limited the region of the image from which features were extracted via two mechanisms. The first is a segmentation mask, created during pre-processing that masks out the optic disc and the non-retinal background component of the image. The second is a heuristic region-of-interest in the
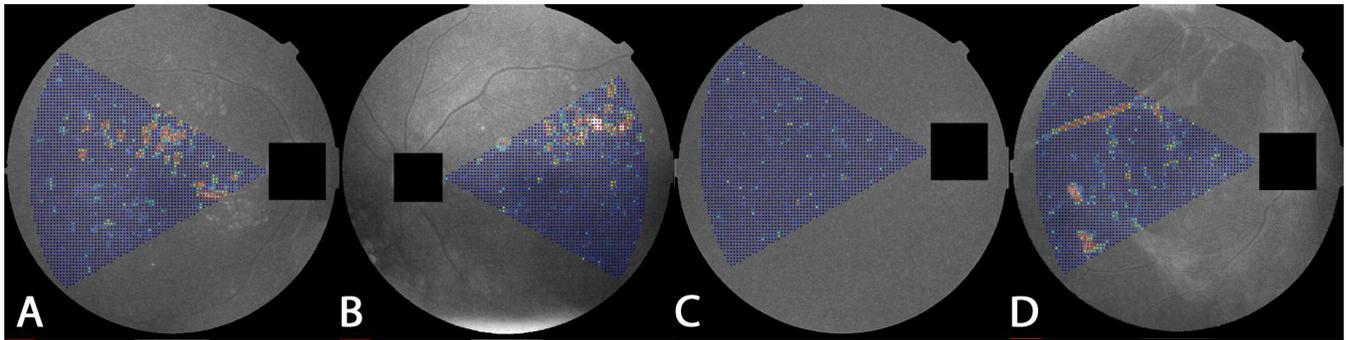
Figure 3: Images of retinas which have undergone patch classification. In each image, the optic disc has been automatically masked by a black square. A conic region-of-interest extends from the optic disc through the rest of the image. Inside each cone, we perform classification on small sub-windows of the image. The confidence of each classification is displayed on a color gradient from blue, which denotes low confidence, to red, which denotes high confidence. (A) and (B) both contain exudates which can be seen by the red (confident) patch classifications. (C) displays a healthy retina and exhibits a low response throughout the region-of-interest. (D) is an example of a misclassification. Notice that the blood-vessel in the center has been incorrectly classified as an exudate.

image which includes a horizontal, conic area of the retina spanning from the location of the optic disc outward towards the center of the retinal. The spatial layout of color fundus images are highly regular in that the major blood vessels branch upwards and downwards from the optic disk. These blood vessel regions were highly confusing to our feature extractor and we found that excluding them did not impact performance as exudates were rarely found in these areas.

A gaussian SVM was trained using 2000 negative and 1309 positive patches. At test time, the SVM classifies each patch of the image and outputs a probability that each patch contains an exudate. In order to provide a single aggregate classification over the entire image, we calculate the fraction of all patches in the image that were classified as exudates and whose confidence score was greater than 93%. If this fraction exceeded a learned threshold, we classified the entire image as containing an exudate.

6. Evaluation

While our ultimate goal was classifying exudates, we also performed an evaluation of our optic disc detection mechanism. We considered a successful masking of the optic disk to be one in which no part of the optic disc escaped our mask. On 1000 images, we correctly located the optic disc 984 times, or 98.4%.

To evaluate our exudate detection pipeline, we extracted SIFT feature descriptors from 1000 images, 87 of which contained exudates. Our evaluation is measured at the image level, rather than the pixel of patch level. The ROC curve for this evaluation is shown in Figure 4. The equal error rate is 87%.

While these results are preliminary, the overall accuracy is promising given our simple model. The false negative rate should ultimately be lowered, however, and we will have to consult with clinicians about where the exact comprise between false negatives and false positives should be. With regard to speed, our system classifies each new image in less
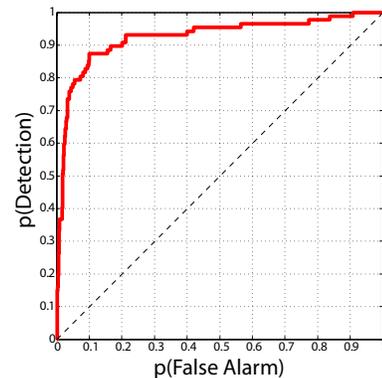

Figure 4: ROC Curve for image-level classification for 1000 images, of which 87 contain exudates. The equal error rate is 87%.

than a second making it faster than the vast majority of efforts to date, some of which may take up to several minutes per image.

7. Initial Deployment Experiences

This project began as an off-shoot of our prior work on deployed WiLDNet (led by the TIER group at UC Berkeley) at Aravind Eye Hospitals to aid their rural telemedicine efforts. At that time, they expressed need for an automated detection system for diabetic retinopathy. At Berkeley, we initially tried using standard image processing techniques to process images but our approaches yielded poor accuracy of detection between normal and abnormal cases.

Later, we began with the modest problem of distinguishing between the retinal images of normal patients with those diagnosed with retinopathy. This too proved to be a hard problem due to significant variations across images along several dimensions not anticipated: variations in retina back-

grounds across patients, images with variable quality and focus (possibly taken by different models of ophthalmoscopes) and spots in retinal images of patients with different types of eye surgeries (cataract, laser). Using a combination of machine learning algorithms, we were able to achieve reasonable accuracy in addressing this problem. We provided a version of this software to Aravind early this year. The immediate feedback they provided was the need for a more advanced system to localize the hemorrhages and exudates.

This paper describes our experiences with building one such system. In July 2009, we provided Aravind with an extended version of our application that was trained using 41 images with pixel-level marking; the marking was performed by ophthalmologists at NYU. Aravind tested our new software across several images and pinpointed specific cases with inaccuracies in detecting hemorrhages and exudates.

Additionally, we found that the training set in our initial system was insufficient in representing the variability of images in the real world. We recently obtained a much larger data set of 20, 000 images. As previously mentioned, we have used 1000 of these images to improve our algorithm. We are currently retraining our algorithms on a much larger scale and have presented some of our preliminary results of the new system in this paper. We hope to deploy the new version of our detection software at the end of this year.

Acknowledgements

References

Akara Sopharak, B. U., and Barman, S. 2009. Automatic exudate detection from non-dilated diabetic retinopathy retinal images using fuzzy c-means clustering. *Sensors* 9(3):2148–2161.

Chaudhuri, S.; Chatterjee, S.; Katz, N.; Nelson, M.; and Goldbaum, M. 1989. Detection of blood vessels in retinal images using two-dimensional matched filters. In *IEEE Transactions on Medical Imaging*, volume 8, 263–269.

Chutatape, O., and Zhang., X. 2004. Detection and classification of bright lesions in color fundus images. In *International Conference on Image Processing*, 139–142.

Di Wu, Ming Zhang, J.-C. L., and Bauman, W. 2006. On the adaptive detection of blood vessels in retinal images. In *IEEE Transactions on Biomedical Engineering*, volume 53.

G. G. Gardner, D. Keating, T. H. W., and Elliott, A. T. 1996. Automatic detection of diabetic retinopathy using an artificial neural network: a screen tool. *British Jouranl of Ophthalmology* 80:940–944.

García, M.; Clara I. Sánchez, J. P.; López, M. I.; and Hornero, R. 2009. Detection of hard exudates in retinal images using a radial basis function classifier. volume 37, 1448–1463.

Granlud, G. H., and Knutsson, H. 1995. *Signal Processing For Computer Vision*. Kluwer Academic Publishers.

IDF. 2009a. Latest diabetes figures paint grim global picture. http://www.idf.org/press_releases.

IDF. 2009b. What is diabetes? http://www.idf.org/about_diabetes.

Li, H., and Chutatape, O. 2003. A model-based approach for automated feature extraction in fundus images. In *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, 394. Washington, DC, USA: IEEE Computer Society.

Lowe, D. G. 2004. Distinctive image features from scale-invariant keypoints. In *International Journal of Computer Vision*, volume 60, 91–110.

NIH. 2009. Facts about diabetic retinopathy disease [nei health information]. http://www.nei.nih.gov/health/diabetic/retinopathy.asp.

Osareh, A.; Mirmehdi, M.; Thomas, B.; and Markham, R. 2001. Automatic recognition of exudative maculopathy using fuzzy c-means clustering and neural networks. *Medical Image Understanding and Analysis* 49–52.

Patra, R.; Nedevschi, S.; Surana, S.; Sheth, A.; Subramanian, L.; and Brewer, E. 2007. Wildnet: Design and implementation of high performance wifi based long distance networks. In *NSDI*.

Pinz, A.; Berngger, S.; Datlinger, P.; and Kruger, A. 1998. Mapping the human retina. *IEEE Transactions on Medical Imaging* 17:606–619.

Saiprasad Ravishankar, Arpit Jain, A. M. 2009. Automatic feature extraction for early detection of diabetic retinopathy in fundus images. CVPR.

Sánchez, C.; Mayo, A.; García, M.; López, M. I. .; and Hornero, R. 2006. Automatic image processing algorithm to detect hard exudates based on mixture models. In *Proceedings of the 28th Annual International Conference of the IEEE EMBS*, volume 1, 4453–6.

Simoncelli, E. 1994. Design of multi-dimensional derivative filters. In *In First International Conference on Image Processing*, 790–793.

Sinthanayothin, C.; Boyce, J. F.; Cook, H.; and Williamson, T. H. 1999. Automated localization of the optic dic, fovea, and retinal blood vessels from digital color fundus images. In *Br J. Opthalmol.*, volume 83, 902–910.

Sopharak, A.; Nwe, K. T.; Moe, Y. A.; Dailey, M. N.; ; and Uyyanonvara., B. 2008. Automatic exudate detection with a naive bayes classifier. In *International Conference on Embedded Systems and Intelligent Technology*, 139–142.

Thomas Walter, Jean-Claude Klein, P. M., and Erginay, A. 2002. A contribution of image processing to the diagnosis of diabetic retinopathy—detection of exudates in color fundus images of the human retina. *IEEE Transactions on Medical Imaging* 21(10):1236–1243.

Unwin, N.; Whiting, D.; Gan, D.; Jacqmain, O.; and Ghyoot, G. 2009. *IDF Diabetes Atlas*. 4 edition.

Wang, H.; Hsu, W.; Goh, K. G.; and Lee, M. L. 2000. An effective approach to detect lesions in color retinal images. In *in IEEE Conference on Computer Vision and Pattern Recognition*, 181–187.

Zhang, X., and Chutatape, O. 2005. Top-down and bottom-up strategies in lesion detection of background diabetic retinopathy. In *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 2*, 422–428. Washington, DC, USA: IEEE Computer Society.