



HAL
open science

Unsupervised Change Detection using Joint Autoencoders for Age-Related Macular Degeneration Progression

Guillaume Dupont, Ekaterina Kalinicheva, Jérémie Sublime, Florence Rossant, Michel Pâques

► **To cite this version:**

Guillaume Dupont, Ekaterina Kalinicheva, Jérémie Sublime, Florence Rossant, Michel Pâques. Unsupervised Change Detection using Joint Autoencoders for Age-Related Macular Degeneration Progression. Springer. ICANN 2020, Part II, 1, In press, LNCS 12397. hal-02963848

HAL Id: hal-02963848

<https://hal.science/hal-02963848>

Submitted on 14 Oct 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Unsupervised Change Detection using Joint Autoencoders for Age-Related Macular Degeneration Progression ^{*}

Guillaume Dupont¹, Ekaterina Kalinicheva¹[0000-0001-8332-2491], Jérémie Sublime^{1,2}[0000-0003-0508-8550], Florence Rossant¹[0000-0003-2517-5213], and Michel Pâques³

¹ ISEP, 10 rue de Vanves, 92130 Issy-Les-Moulineaux, France
firstname.lastname@isep.fr

² LIPN - CNRS UMR 7030, 99 av. J-B Clément, 93430 Villetaneuse, France

³ Clinical Imaging Center 1423, Quinze-Vingts hospital, INSERM-DGOS Clinical Investigation Center, Paris, France
mpaques@15-20.fr

Abstract. Age-Related Macular Degeneration (ARMD) is an eye disease that has been an important research field for two decades now. Researchers have been mostly interested in studying the evolution of lesions that slowly causes patients to go blind. Many techniques ranging from manual annotation to mathematical models of the disease evolution bring interesting leads to explore. However, artificial intelligence for ARMD image analysis has become one of the main research focus to study the progression of the disease, as accurate manual annotation of its evolution has proved difficult using traditional methods even for experienced doctors. Within this context, in this paper, we propose a neural network architecture for change detection in eye fundus images to highlight the evolution of the disease. The proposed method is fully unsupervised, and is based on fully convolutional joint autoencoders. Our algorithm has been applied to several pairs of images from eye fundus images time series of ARMD patients, and has shown to be more effective than most state-of-the-art change detection methods, including non-neural network based algorithms that are usually used to follow the evolution of the disease.

Keywords: Change detection · Unsupervised learning · ARMD.

1 Introduction

Dry age-related macular degeneration (ARMD or AMD), a degenerative disease of the retina, is a main cause of irreversible visual loss. It is characterized by a centrifugal progression of atrophy of the retinal pigment epithelium (RPE),

^{*} This study has been approved by a French ethical committee (Comité de Protection des Personnes) and all participants gave informed consent.

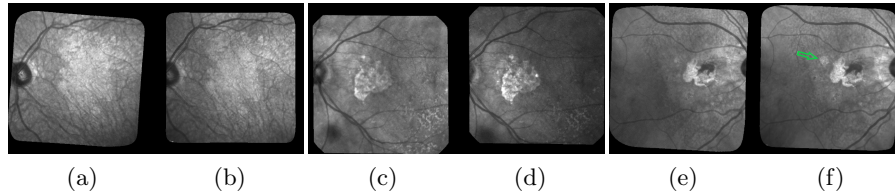


Fig. 1: 3 examples of pairs of images acquired 6 months apart. The GA corresponds to the bright areas. The green arrow in (f) shows a new lesion.

a cellular layer playing a key role in the maintenance of the photoreceptors. Blindness may occur when the fovea (the central part of the eye), is affected. The disease can be diagnosed and monitored using fundus images: ophthalmologists can observe pathologic features such as drusens that occur in the early stages of the ARMD, and evaluate the geographic atrophic (GA) progression in the late stages of degeneration (Fig. 1).

Automatic analysis of dry ARMD fundus images is of high medical interest and this has been an important research field for two decades, for diagnosis [15] or follow up [10],[16] purposes. Imaging modalities are most often color eye fundus images [11, 13, 3], fundus autofluorescence (FAF) [12, 6, 16], and, to a lesser extent, confocal scanning laser ophthalmoscopy (cSLO) in infrared (IR), or optical coherence tomography (OCT) [5]. In this work, we use cSLO images in infrared (IR): this modality is comfortable to acquire for the patients, and has better resolution and contrast than color imaging, an older technology. Our goal is to detect the appearance of new atrophic areas and quantify the growth of GA from pairs of images acquired from follow-up exams and to ultimately propose predictive models of the disease progress.

Figure 1 shows 3 pairs of consecutive images, taken at 6 months interval. The lesions (GA) in the fundus and around the optical disk are the brighter areas. Monitoring the GA progression in these areas is obviously challenging because of the images quality: uneven light, saturation issues, illumination distortion between images, GA poorly contrasted with retinal structures interfering (vessel, optical disk), blur, etc. The difficulty also lies in the high variability of the lesions in terms of shape, size and number. The lesion boundary is quite smooth in some cases (c and d) and very irregular in others (a and b). At any time, new spots can appear (as shown by the green arrow between e and f) and older lesions can merge. All these defaults make the manual delineation task very complex, even for expert ophthalmologists.

In order to assess the disease progression, it is necessary to perform a differential analysis between consecutive images to get the lesion growth, so that the lesion growth can be modelled. In this paper, we propose a fully unsupervised differential analysis method based on a joint autoencoders. Our model does not require labeled images that are difficult to come by in quantity and quality high enough to train a supervised neural network. Our method is applied to pairs

of images of a patient eye fundus time series and aims at efficiently segmenting medically significant changes between the two images: we are interested only in changes in the GA lesions, while meaningless differences and light artefacts should be ignored.

2 Related Works

The following works are most related to our proposed algorithm as they are unsupervised algorithm applied to various eye disease images, including ARMD: In [18], Troglia et al. published an improvement of their previous works realized with Nappo [19] where they use the Kittler and Illingworth (K&I) thresholding method. Their method consists of applying the K&I algorithm on random sub-images of the difference image obtained between two consecutive eye fundus images of a patient with retinopathy. By doing so, they obtain multiple predictions for each pixel and can then make a vote to decide the final class. This approach has the advantage that it compensates for the non-uniform illumination across the image, however it is rather primitive since it does not actually use any Machine Learning and rely on different parameters of the thresholding method to then make a vote. To its credit, even if it achieves a relatively weak precision, it is fully unsupervised like our method. In [13], the authors tackle a similar problematic to ours where they correct eye fundus images by pairs, by multiplying the second image by a polynomial surface whose parameters are estimated in the least-squares sense. In this way, illumination distortion is lessened and the image difference enhances the areas of changes. However, the statistical test applied locally at each pixel is not reliable enough to get an accurate map of structural changes.

Other works related with eye diseases take the different approach of segmenting lesions in individual images instead of looking for changes in pairs of images. In [11], Köse et al. proposed an approach where they first segment all healthy regions to get the lesions as the remaining areas. This approach also requires segmenting separately the blood vessels, which is known to be a difficult task. This method involves many steps and parameters that need to be supervised by the user. In [16], Ramsey et al. proposed a similar but unsupervised method for the identification of ARMD lesions in individual images: They use an unsupervised algorithm based on fuzzy c-means clustering. Their method achieves good performances for FAF images, but it performs less well for color fundus photographs. We can also mention the work of Hussain et al. [7] who proposed another supervised algorithm to track drusen progression for ARMD. It uses a U-Net to segment vessels and detect the optic disc to reduce the region of interest of drusen detection. After that, they detect the drusen using intensity ratio between neighbor pixels.

Other traditional Machine learning algorithms have also been used for GA segmentation such as random forest [3] or k-nearest neighbor classifiers [6]. Feature vectors for these approaches typically include intensity values, local energy, texture descriptors, values derived from multi-scale analysis and distance to the

image center. Nevertheless, these algorithms are supervised: they require training the classifier from annotated data, which brings us back to the difficulty of manually segmenting GA areas.

Apart from medicine, change detection algorithms have been proposed for many different applications such as remote sensing or video analysis. In [1] the authors reveal a method combining principal component analysis (PCA) and K-means algorithm on the difference image. In [8], an architecture relying on joint auto-encoders and convolutional neural networks is proposed to detect non-trivial changes between two images. In [2] the authors propose an autoencoder architecture for anomaly detection in videos.

Finally, as we have seen that quite a few methods rely on segmentation first and change detection after. We can also mention a few noteworthy unsupervised segmentation algorithms used outside the field of medicine: Kanazaki et al. [9] used CNN to group similar pixels together with consideration of spatial continuity as a basis of their segmentation method. Finally, W-Nets using with a soft Normalized-Cut Loss are another option.

3 Data description

In this section, we will provide some details on the data we used.

Our images were all acquired at the Quinze-Vingts National Ophthalmology Hospital in Paris, in cSLO with IR illumination. Patients have been followed-up during a few years, hence we have series of retinal fundus images, often for both eyes, showing the progression of the GA. The average number of images in each series is 13. We used 336 images from 15 patients time series taken between 2007 and 2019. All pictures are in grayscale and vary greatly in size, but the most common size is $650 * 650$ pixels.

As mentioned in the introduction, the images contain many imperfections such as blurs, artifacts and, above all, non-uniform illumination (see Fig. 2). All images were spatially aligned with i2k software⁴. Furthermore, all images are surrounded by black border that contain no useful information. These borders were removed from the segmentation area using a mask.

All images were preprocessed using a new method (not published yet) to reduce the light distortion within images series. This algorithm relies on an illumination / reflectance model and corrects all images of a serie with respect to a common reference image. Uneven illumination generally remains present in every processed image (Fig. 2) but the smooth illumination distortions are compensated. The calculus of the absolute value of the difference between two consecutive images demonstrates the benefit of this algorithm (Fig. 2, last column).

We used 3 different series of images to rate our proposed change detection method: they feature different characteristics in terms of disease progress, lesion

⁴ <https://www.dualalign.com/retinal/image-registration-montage-software-overview.php>




Fig. 2: Illumination correction. The 3 images on the left represent the 2 original consecutive images and their raw difference in absolute value; On the right: the same images after illumination correction and the new difference.




Fig. 3: Autoencoder architecture

shape and size. We developed several user-guided segmentation tools for ophthalmologists to build the validation images with the expected segmentation: These tools allowed them to provide a manual segmentation of some of our dataset images using techniques such as segmentation algorithms, local thresholdings and simple linear interpolation to draw the lesions borders. However, it is worth noting that the task remained long and tedious and couldn't be done for all images as on average, it took 13 minutes to process one single image. Finally, the binary change mask between two consecutive images was obtained by subtraction of the segmentation masks.

4 Proposed algorithm

Our algorithm uses principles from previous remote sensing applications [8], where an unsupervised deep autoencoder was used to automatically detect meaningful changes between two satellite images with the goal of finding new constructions or changes in the landcover, all the while discarding seasonal changes.

In our paper, we use the common issues between satellite imaging and our medical ARMD eye fundus to adapt this deep learning algorithm: both types of images may suffer from lighting issues, noise issues, blurry elements, complex objects present in the images, various intervals between images, and most importantly the common goal of detecting only the changes from specific classes within the images.

While the problematic look similar, remote sensing and medical applications also have specific issues: medical images only have one channel, the textures are

different, and the size and scale of both the images and objects to analyze is quite different. To account for these specificities, we had to modify their AE architecture as described in the next subsection.

4.1 Joint Autoencoder Architecture

As mentioned previously, our algorithm is based on autoencoders. Autoencoders [4] are a type of neural networks whose purpose is to make the output as close as possible to the input. The learning process consists in an encoder part learning some meaningful representations of the original data and a decoder transforming them back into the original data: in a fully convolutional AE, a stack of convolutive layers is applied to the input image in order to extract feature maps (FM) which will then be used to reconstruct the input image.

Usually AEs with dense layers are used to perform a dimensionality reduction followed by a clustering or segmentation. However, in computer vision, fully convolutional AEs are preferred for their ability to extract textures. Examples of such networks include fully convolutional networks (FCNs) or U-Nets [17]. However, in our case we do not use pooling layers, and so we keep the same dimensions as the input and only the depth increases.

Our network (Fig 3) is made of 4 convolutional layers in the encoder of size 16, 16, 32, 32 respectively, and in the same way of 4 convolutional layers of size 32, 32, 16, 16 respectively in the decoder side. We apply a batch normalization and a ReLU activation function at each step of the network except for the last layer of the encoder where we only add the L2 normalization, and also for the last layer of the decoder where we apply a Sigmoid function (see in Figure 3).

4.2 Algorithm steps

Our algorithm is made of four steps. We start by dividing the images into several patches. Then we build the joint autoencoder where it learns how to reconstruct the images, and after we tweak the method by learning the autoencoder to reconstruct not the image itself but the precedent or the future image. The neural networks will learn the changes caused by the non-uniform illumination or noise, but will fail on ARMD progression generating a high reconstruction error (RE), consequently making it possible to detect them. The next subsections will detail some of these steps:

Pre-training Let us consider a series of M images representing the progression of ARMD in a patient’s eye. Each image has the same number of N useful patches, and we sample $\left\lfloor \frac{N}{M} \right\rfloor$ of the patches from every image. This allows us to build a unique autoencoder AE that works for all pairs in the series, and to prevent overfitting. We also apply a gaussian filter to the patches in order to weight the pixels by giving more importance to the center of the patch in the RE calculus.

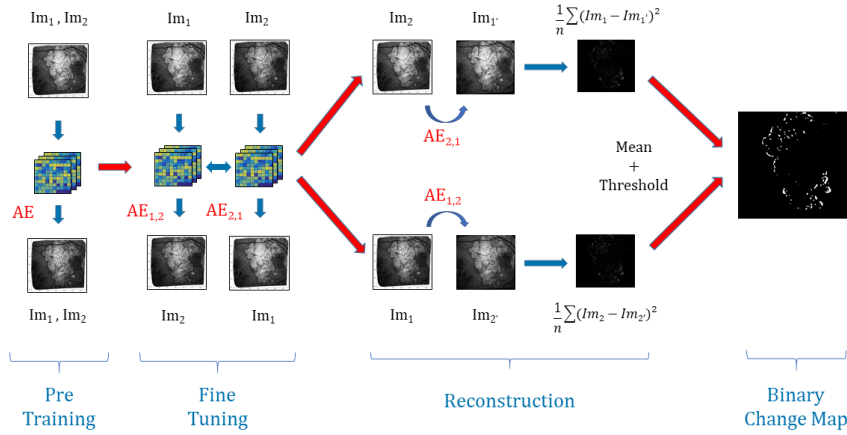


Fig. 4: Structure of the algorithm. *Example for series of two images Im_1 and Im_2 and n the number of patches.*

During the encoding pass of the AE, the model extracts feature maps of N patches of chosen samples with convolutional layers (Figure 4), and then during the decoding pass, it reconstructs them back to the initial ones.

Fine-tuning For every consecutive pair $i, i + 1$ with $i \in \llbracket 1; M - 1 \rrbracket$ of images we are going to build two autoencoders initialized with the weights found in the pre-training part. On one hand $AE_{i,i+1}$, aims to reconstruct patches of Im_{i+1} from patches of Im_i and, on the other hand, $AE_{i+1,i}$ is going to reconstruct patches of Im_i from patches of Im_{i+1} .

The whole model is trained to minimize the difference between: the decoded output of $AE_{i,i+1}$ and Im_{i+1} , the decoded output of $AE_{i+1,i}$ and Im_i , and the encoded outputs of $AE_{i,i+1}$ and $AE_{i+1,i}$, see Figure 4.

This joint configuration where the learning is done in both temporal directions, using joint backpropagation, has empirically proven to be much more robust than using a regular one-way autoencoder [8]. To optimize the parameters of the model, we use the mean squared error (MSE) of the reconstructed patches

Reconstruction and Thresholding Once the models are trained and stabilized, we perform the image reconstruction. For each pair, we note $Im_{i+1'}$ the reconstruction of Im_{i+1} from Im_i with $AE_{i,i+1}$ and likewise we note $Im_{i'}$ the reconstruction of Im_i from Im_{i+1} with $AE_{i+1,i}$. Then, we calculate the reconstruction error RE for every patch between Im_i and $Im_{i'}$ on one side and between Im_{i+1} and $Im_{i+1'}$ on another side. This gives us two images for each pair representing the average REs for $Im_{i'}$ and $Im_{i+1'}$, that we average to get only one. The model will easily learn the transformation of unchanged areas from one image to the other: changes in luminosity and blurring effects. At the same

time, because the changes caused by the disease progression are unique, they will be considered as outliers by the model, and thus will have a high *RE*. Hence, we apply Otsu’s thresholding [14] that requires no parameters and enables us to produce a binary change map (*BCM*) (Figure 4).

5 Experimental results

5.1 Experimental settings

We chose to compare our methods presented in subsection 4.1 with 3 other methods, all on the preprocessed images. We applied all the methods to 3 of our series for which we have a ground truth.

The following parameters were chosen for all convolutional layers of our method: kernel size to 3, stride to 1 and padding to 1. Adam algorithm was used to optimize the models. We set the number of epochs to 8 for the pre-training phase and just 1 for the fine-tuning phase. For both phases, the learning rate was set to 0.0001 and the batch size to 100. Using trials and errors to determine the best parameters, we chose a patch size $P = 13$ and a value $\sigma = 12$ for the Gaussian filters.

The first method that we use for comparison is a simple subtraction of two consecutive images with an application of Otsu’s thresholding on the result. The second comparison is a combination of principal component analysis (PCA) and K-means algorithm on the difference image proposed by Celik et al. in [1], and we apply it to medical images with blocks of size 5. To finish we take a Deep-Learning based approach [9] which uses CNN to group similar pixels together with consideration of spatial continuity. This work by Kanazaki et al. was initially made for unsupervised segmentation, consequently, we apply the algorithm to our images and then do the segmentation subtractions to get binary change maps. The convolutive layers have the same configuration than for our network and we set the *minimum number of labels* to 3.

All algorithms were executed on an Nvidia GPU (RTX TITAN) with 64 GB of RAM and an Intel 9900k. It took about 20 minutes for a series of 8 frames with a patch size P of 13, with the execution time increasing with it.

5.2 Results

The results for patients 1, 3 and 5 are shown in Table 1, as well as Figures 5 and 6 that we added for visual inspection purposes. Note that the scores presented in Table 1 are for the complete series (15 to 20 pairs per patient), while the scores shown in the Figures are for the individual couples of images used in each example.

When looking at Table 1, we can see that the simple difference coupled with Otsu thresholding achieves the best recall results on average, that there is no clear winner for the Precision, and that our proposed method on average has the best F1 Score. Note that we did not use Accuracy because of the strong class

Table 1: Results and comparison of the different approaches. *It contains the means of the recall, the precision and the F1 score for each serie.*

Patient ID	Method	Recall	Precision	F1 score
001 15 images	Diff+Otsu	0.68	0.11	0.16
	Kanezaki	0.32	0.29	0.18
	Celik	0.48	0.28	0.3
	Our method	0.44	0.21	0.26
003 5 images	Diff+Otsu	0.55	0.1	0.17
	Kanezaki	0.2	0.27	0.07
	Celik	0.24	0.33	0.27
	Our method	0.29	0.28	0.28
005 8 images	Diff+Otsu	0.46	0.2	0.26
	Kanezaki	0.2	0.43	0.21
	Celik	0.26	0.37	0.28
	Our method	0.33	0.34	0.32
Total (pa- tients' mean)	Diff+Otsu	0.57	0.14	0.2
	Kanezaki	0.24	0.33	0.15
	Celik	0.33	0.33	0.28
	Our method	0.35	0.28	0.29

imbalance, with a large majority of "no change class" pixels leading to results over 95% of accuracy for all methods that were irrelevant.

Our interpretation of these results is the following: Otsu thresholding applied to the difference between two images has the best recalls because it detects most real change pixels. But the binary change map is also very noisy, corresponding to a high number of false positives (wrongly detected changes) which is confirmed by the very low precision score. This can also be observed on Figure 5-d which is an example of the high number of false positive detected using Otsu thresholding compared with the ground truth in Figure 5-c, or our method result in Figure 5-e.




Fig. 5: Difference + Otsu thresholding VS our approach (AE) on patient 003. **a&b**-Images taken 3 months apart, **c**-groundtruth, **d**-difference with Otsu (F1 score = 0.26), **e**-our method (F1 score = 0.36).

In Figure 6, we compare our approach with the 2 other algorithms relying on more advanced Machine Learning techniques. First we can see that like in Table 1, our approach gets the best F1-score for both patients and pairs of images. Then, we can see that Kanezaki et al. approach achieves over-segmentation in Figure 6-d and under-segmentation in Figure 6-j, which highlights that it is more difficult to parametrize and may require different parameters for each pair of image, which is not the case for both our approach and Celik et al. approach. Finally, regarding the comparison between Celik et al. approach and our proposed method, we can see from Figures 6-e, 6-f, 6-k and 6-l, that also like in Table 1, Celik et al. approach achieves overall good results that are comparable to the ones of our method. However, in the same way that we have better F1-score and accuracy results, the visual results for our methods are also better as the changes we detect in the lesions are cleaner and overall less fragmented into very small elements compared with the ones found by Celik et al. approach. Furthermore, we can see that our method finds changes that are more in the peripheral areas of the lesions, while Celik et al. approach tends to find elements inside existing lesions (6-k) which are of lesser interest from a medical point of view.




Fig. 6: Comparison of 3 methods on patients 001 (1st line) and 005 (2nd line): **a**-Patient 001 April 2017, **b**-Patient 001 October 2017, **c**-ground truth, **d**-Asako Kanezaki approach (F1= 0.17), **e**-Turgay Celik approach (F1= 0.43), **f**-our method (F1= 0.43), **g**-Patient 005 October 2017, **h**-Patient 005 June 2018, **i**-ground truth, **j**-Asako Kanezaki approach (F1= 0.15), **k**-Turgay Celik approach (F1= 0.35), **l**-our method (F1= 0.4).

Overall, we can conclude that both Otsu thresholding and Kanezaki’s approach suffer from risks of over-segmentation detecting a lot of noise, or under-segmentation detecting nothing, both of which are impossible to exploit from a medical point of view. On the other hand, despite somewhat mild recall and precision scores, Celik approach and our method are visually much better at

detecting meaningful changes in ARMD lesions structures. Moreover, we can see that our proposed method has a slightly higher F1-Score and finds structures that are visually better and more interesting from a medical point of view since they tend to be more at the border of existing lesions instead of inside them, and are also less fragmented.

6 Conclusion and future works

In this work, we have introduced a fully unsupervised deep learning new architecture that detects the evolution of ARMD lesions in eye fundus series of images. With a pre-cleaning of the series to remove as much lighting issues as possible, our proposed method is based on an auto-encoder architecture that can detect non-trivial changes between pairs of images, such as the evolution of a lesion, while discarding more trivial changes such as lighting problems or slight texture changes due to different image angles. Our proposed method was applied to 3 real sets of images, and was compared with 3 methods from the state of the art. Despite mild F1-Score results due to various issues, our method has shown to give good enough results for a fully unsupervised algorithm and to perform better than the other methods from the state of the art, and may prove useful to assist doctors in properly detecting the evolution of ARMD lesions by proposing a first raw segmentation of the evolution.

In our future works, we plan on working on an approach that can work on full time series rather than pairs of images. This would require both better lighting correction algorithms as well as more accurate ground-truth but may lead to more interesting models to predict the evolution of ARMD.

References

1. Celik, T.: Unsupervised change detection in satellite images using principal component analysis and k -means clustering. *IEEE Geoscience and Remote Sensing Letters* **6**(4), 772–776 (2009)
2. Chong, Y.S., Tay, Y.H.: Abnormal event detection in videos using spatiotemporal autoencoder. In: Cong, F., Leung, A.C., Wei, Q. (eds.) *Advances in Neural Networks - ISNN 2017 - 14th International Symposium, ISNN 2017, Sapporo, Hakodate, and Muroran, Hokkaido, Japan, June 21-26, 2017, Proceedings, Part II. Lecture Notes in Computer Science*, vol. 10262, pp. 189–196. Springer (2017). https://doi.org/10.1007/978-3-319-59081-3_23
3. Feeny, A.K., Tadarati, M., Freund, D.E., Bressler, N.M., Burlina, P.: Automated segmentation of geographic atrophy of the retinal epithelium via random forests in areds color fundus images. *Computers in biology and medicine* **65**, 124–136 (2015)
4. Hinton, G.E., Salakhutdinov, R.R.: Reducing the dimensionality of data with neural networks. *Science* **313**(5786), 504–507 (2006).
5. Hu, Z., Medioni, G.G., Hernandez, M., Hariri, A., Wu, X., Satta, S.R.: Segmentation of the geographic atrophy in spectral-domain optical coherence tomography and fundus autofluorescence images. *Investigative ophthalmology & visual science* **54**(13), 8375–8383 (2013)

6. Hu, Z., Medioni, G.G., Hernandez, M., Sadda, S.R.: Automated segmentation of geographic atrophy in fundus autofluorescence images using supervised pixel classification. *Journal of Medical Imaging* **2**(1), 014501 (2015)
7. Hussain, M.A., Govindaiah, A., Souied, E., Smith, R., Bhuiyan, A.: Automated tracking and change detection for age-related macular degeneration progression using retinal fundus imaging. pp. 394–398 (06 2018). <https://doi.org/10.1109/ICIEV.2018.8641078>
8. Kalinicheva, E., Sublime, J., Trocan, M.: Change detection in satellite images using reconstruction errors of joint autoencoders. In: *Artificial Neural Networks and Machine Learning - ICANN 2019: Image Processing - 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17-19, 2019, Proceedings, Part III*. pp. 637–648 (2019). https://doi.org/10.1007/978-3-030-30508-6_50
9. Kanezaki, A.: Unsupervised image segmentation by backpropagation. In: *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* (2018)
10. Köse, C., Sevik, U., Gençalioglu, O.: Automatic segmentation of age-related macular degeneration in retinal fundus images. *Computers in biology and medicine* **38** **5**, 611–9 (2008)
11. Köse, C., Sevik, U., Gençalioglu, O., Ikibas, C., Kayikicioğlu, T.: A statistical segmentation method for measuring age-related macular degeneration in retinal fundus images. *Journal of medical systems* **34**, 1–13 (02 2010)
12. Lee, N., Laine, A.F., Smith, R.T.: A hybrid segmentation approach for geographic atrophy in fundus auto-fluorescence images for diagnosis of age-related macular degeneration. In: *2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. pp. 4965–4968. IEEE (2007)
13. Marrugo, A.G., Millan, M.S., Sorel, M., Sroubek, F.: Retinal image restoration by means of blind deconvolution. *Journal of Biomedical Optics* **16**(11), 116016 (2011)
14. Otsu, N.: A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics* **9**(1), 62–66 (Jan 1979). <https://doi.org/10.1109/TSMC.1979.4310076>
15. Priya, R., Aruna, P.: Automated diagnosis of age-related macular degeneration from color retinal fundus images. In: *2011 3rd International Conference on Electronics Computer Technology*. vol. 2, pp. 227–230. IEEE (2011)
16. Ramsey, D.J., Sunness, J.S., Malviya, P., Applegate, C., Hager, G.D., Handa, J.T.: Automated image alignment and segmentation to follow progression of geographic atrophy in age-related macular degeneration. *Retina* **34**(7), 1296–1307 (2014)
17. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)
18. Troglia, G., Alberti, M., Benediktsson, J., Moser, G., Serpico, S., Stefánsson, E.: Unsupervised change-detection in retinal images by a multiple-classifier approach. pp. 94–103 (04 2010)
19. Troglia, G., Nappo, A., Benediktsson, J., Moser, G., Serpico, S., Stefánsson, E.: *Automatic Change Detection of Retinal Images*, vol. 25, pp. 281–284 (01 2010)