

Fully Automated GrowCut-based Segmentation of Melanoma in Dermoscopic Images



Leyton Ho^{1,2*}

Early diagnosis of malignant skin lesions is vital to maximize survival rates of patients with melanoma. Currently, diagnosis of melanoma, a common and deadly form of skin cancer, is done primarily by biopsy, an invasive and costly procedure. A new technique called dermoscopy has been proposed as a pre-biopsy melanoma risk evaluation tool. The accuracy and efficiency of clinical diagnosis can be improved with fully automated dermoscopic diagnosis. Developing such a diagnosis tool is particularly important for melanoma and can be achieved through analysis of skin lesion images using a three-step process of segmentation, feature extraction and classification. In this paper, a new method for automated segmentation of skin lesions is proposed. The method first preprocesses the input image using white balance correction and noise and artifact reduction. Subsequently, it identifies skin abnormalities using a learned healthy-appearing skin intensity model. It then detects the skin lesions using a trained lesion shape classifier, and finally, refines the lesion boundary using GrowCut-based delineation. The method was trained and validated on a publicly available dataset of skin lesions and compared to well-known baseline methods including Otsu Thresholding, Active Contours, and K-Means Clustering. The algorithm achieved a median accuracy of 0.94 and a median Dice Coefficient of 0.87, representing a significant improvement ($p < 0.001$) over the baseline methods. The proposed algorithm is more accurate than established segmentation methods, fully automated, and can be used in an automatic skin cancer diagnosis system with high accuracy and efficiency.

INTRODUCTION

Skin cancer is the most common type of cancer in the United States, affecting one in five Americans over the course of a lifetime (Robinson, 2005). Among skin cancers, melanoma is one of the most dangerous forms, accounting for the vast majority of deaths associated with skin cancer. It will affect an estimated 91,270 people in the United States in 2018, leading to an estimated 9,320 deaths (American Cancer Society, 2018). Although melanoma is a prevalent and potent form of cancer, patients often make full recoveries if it is discovered and treated early. The malignant tumor is simply removed through a relatively straightforward surgery. If melanoma is discovered when it is still localized, patients have a survival rate of 98.5%. This drops to 62% when the disease spreads to the lymph nodes and 18% when it metastasizes to other organs (National Cancer Institute, 2017). Thus, it is extremely important to diagnose malignant melanoma early.

Currently, diagnosis of melanoma is done primarily by biopsy, an invasive and costly procedure (American Cancer Society, 2017). A new technique called dermoscopy has been proposed as a pre-biopsy melanoma risk evaluation tool that can be applied by a wide range of physicians, including family practice doctors (Herschorn, 2012). Dermoscopy is a non-invasive method that uses microscopes to amplify details of skin lesion photographs, such as the colors and microstructures of the skin, the dermoepidermal junction (the area of tissue joining the epidermal and dermal layers of skin), and the papillary dermis (uppermost layer of the dermis). Compared to inspection of cutaneous lesions by the naked eye, this method can increase physicians' confidence in their referral accuracy to dermatologists thereby reducing unnecessary biopsies. The early phases of malignant melanoma, however, share many clinical features with atypical or unusual looking non-malignant moles, also known as dysplastic nevi. As a result, diagnostic accuracy has been shown to range between 50-75% (Stanganelli, 2017). A commonly used standard for pre-biopsy melanoma risk evaluation is described by the ABCDE rule (Abbasi et al., 2004). This rule defines five common characteristics of malignant lesions: asymmetry (A), border irregularity (B), multiple colors (C), diameter greater than six millimeters (D), and enlargement (E). By looking for these five characteristics in dermoscopic images, physicians can evaluate the risks of malignancy of the lesions. Such evaluation, however, is inherently imperfect due to differences in human interpretation of the lesions. To enable faster, more accessible, and more effective evaluation of melanoma risks, there is a need for automatic computerized processing of skin lesions. The

¹ Palo Alto High School, 50 Embarcadero Rd, Palo Alto, CA 94301

² Stanford University Laboratory of Quantitative Imaging and Artificial Intelligence, 1265 Welch Rd, Stanford, CA 94305

*To whom correspondence should be addressed:
leyton.ho@gmail.com



Except where otherwise noted, this work is licensed under <https://creativecommons.org/licenses/by/4.0>

doi:10.22186/jyi.36.1.11-17




Figure 1. Eight images showcasing the variety of ISIC16 dataset.




Figure 2. An example of an original color image of a lesion from ISIC 16 dataset (A) and its grayscale version used by the algorithm after pre-processing with ground truth segmentation outline in blue (B).

goal is a workflow where medical technicians take photographs of suspect lesions, upload them to a remote or local computer and receive a diagnosis and confidence measure. This would lead to widespread and low-cost availability of high-quality melanoma diagnosis.

The automated processing and diagnosis of melanoma consists of detection and delineation of skin lesions followed by extraction of established measurements such as in the ABCDE rule. Fully automated segmentation provides both detection of the lesion in the image and identification of its boundaries. Automated segmentation can be combined with such rules to provide automated malignancy risk estimation for a lesion. Black box approaches to lesion classification, such as end-to-end machine learning algorithms, perform diagnosis without providing insight or explanation for the conclusion (Esteva et al., 2017). In some instances, such approaches can lead to the right conclusions based on wrong assumptions. For example, a deep learning approach to skin lesion segmentation algorithm was more likely to call a tumor malignant in the presence of a ruler in the image. This correlated with an increased likelihood that a lesion was cancerous, because when dermatologists are looking at a lesion that they think might be a tumor, they use the ruler to take an accurate measurement of its size (Patel, 2017). On the other hand, diagnostic pipelines with explicit segmentation and feature identification (such as the ABCDE rule) can show a physician the steps leading to, and provide reasoning for, the generated prediction. This work focuses on improving automated image segmentation usable in a multitude of diagnostic

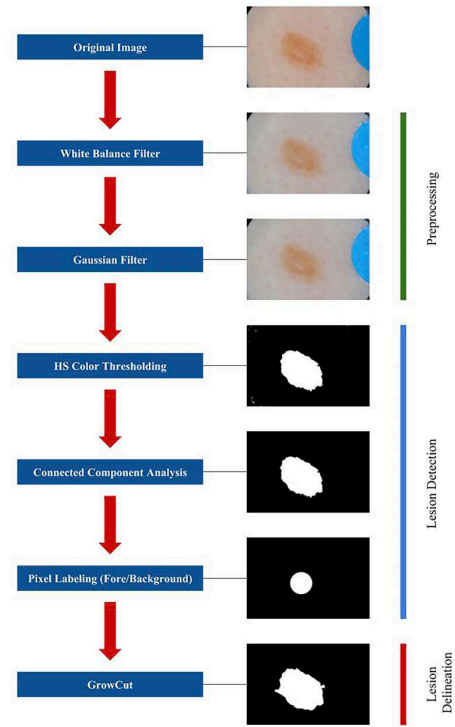


Figure 3. Automated Segmentation Algorithm flowchart.

algorithms.

This paper proposes a new fully automated skin lesion segmentation algorithm, using a pipeline of (a) image pre-processing, (b) lesion detection using a normal-appearing skin intensity model and lesion shape classifier, and (c) refined boundary delineation using the GrowCut segmentation algorithm (Vezhnevets & Konouchine, 2005). The new pipeline with GrowCut, a cellular automaton technique to find homogeneous structures, was investigated to determine the efficacy of this approach. This research demonstrates that a cellular automaton approach supplemented with image pre-processing and good initial object identification leads to better segmentation. The proposed algorithm is implemented using MATLAB and evaluated on a publicly available dataset of 379 images. The results showed that it is significantly better than three well-known baseline segmentation algorithms: Otsu Thresholding, Active Contours and K-Means Clustering, which use simpler approaches to achieve delineation. In the following sections, the proposed algorithm, its evaluation, and results of comparison with baseline techniques using standard error measures are described in detail.

MATERIALS AND METHODS

Dataset

The accuracy of the proposed segmentation algorithm was evaluated in MATLAB using a skin lesion image dataset from the International Skin Imaging Collaboration (“ISIC Archive,” 2016). Specifically, the ISIC 2016 (Figure 1) challenge dataset (ISIC16)

is used, which consists of 900 images with ground truth segmentations (Figure 2), to train the analytic models, and 379 images to test the effectiveness of the algorithm using a variety of different performance metrics (Gutman et al., 2016). The lesion appearances in the dataset vary greatly in attributes like color, texture, uniformity of lesions, and location of lesions.

Automated Segmentation Algorithm

The proposed segmentation algorithm (Figure 3) consists of the following steps: pre-processing, lesion detection, and outline delineation. Pre-processing reduces noise and artifacts in the image, lesion detection finds the center of the lesion and initial lesion radius, and outline delineation refines the lesion boundaries and produces the final segmentation.

Pre-processing

The input image is pre-processed to remove artifacts, such as small skin tone variation, hair and bandages. This allows later stages of the algorithm to focus on the lesion, thus optimizing its accuracy. Two filters are applied in this first stage of the algorithm. First, a white balance correction filter is employed to standardize the image intensities and eliminate the negative effect of illumination variation across images. Second, a Gaussian filter is applied to smooth out the image even more, making it easier to process. The equation of a Gaussian function in two dimensions is:

$$G(x,y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

where x is the distance from the origin in the horizontal axis, y is the distance from the origin in the vertical axis, and σ is the standard deviation (sigma) of the Gaussian distribution.

In the proposed algorithm, the Gaussian function is applied to kernels of 15 pixels by 15 pixels with sigma set to three pixels. Initial experiments were performed to select the kernel size and sigma value. The accuracy of segmentations under the Dice Coefficient was tested using kernel sizes ranging from 8x8 to 20x20 and sigma values between one to four. It was determined that while changing the kernel size made a negligible difference in smoothing out the image, a kernel size of 15x15 pixels was the most accurate of the sizes tested. A sigma of three pixels ensured that the outermost pixels of each block did not receive too much weight, and it was noted to reduce noise in the image well with the chosen kernel size.

The Gaussian filter sets each pixel to the weighted average of its neighbors, with the central pixels receiving more weight than the outermost pixels. Giving more weight to the central pixels results in gentler smoothing than a normal weighted average. The resulting image blur removes small objects but preserves boundaries and edges (Figure 4).

Automated Lesion Detection

Hue and Saturation Learning-based Healthy-Appearing Skin Segmentation

The pre-processed image is then color thresholded (Figure 5) to remove all skin-like pixels. Thresholding is the process where pix-




Figure 4. Original image (A) and pre-processed image (B) from ISIC16 dataset (index 65).




Figure 5. Original image (A) and color thresholded image (B) from ISIC16 dataset (index 24).



Figure 6. An example of original (A), color thresholded (B), and selected connected-component (C) images, ISIC16 (index 440).

els within a specific range of hue and (separately) saturation values are labeled as “not-skin,” and the intersection of these two not-skin pixel masks is found. The hue and saturation (HS) color channels are used instead of the RGB color channels due to inherently large variance of pixel intensity of the lesions in the RGB channels. The HS channels exhibit greater clustering and uniformity in the dataset. The algorithm determines optimal hue and saturation threshold ranges to segment out skin through three-fold validation, that is, training on a third of the data set and validating on the remaining two-thirds in three independent experiments. For each fold of 300 images, the ground truth image is used to randomly sample 1000 pixels of each image that are outside the ground truth mask and away from the border of the image. These pixels are highly likely to be skin pixels. The median and standard deviation of the entire set are then computed. The median is used instead of the mean because using the median ensures that any outlier lesion pixels sampled do not affect the range of intensities calculated, so non-skin pixels mistakenly sampled do not affect the computation. The trained threshold values are then validated on the remaining 600 images in the training data. The accuracy measure is taken after connected component analysis, and based on whether the correct object has been identified as the lesion. This process is performed

on the other two folds, and the three median and standard deviation values are analyzed and converted into a range of intensities to be used as the threshold.

Connected-Component Analysis

After HS thresholding, connected-component analysis is performed on the remaining objects in the image to isolate the lesion. First, the distance from the center of each object to the center of the image is calculated. Images taken for dermoscopy focus on the lesion under consideration. Hence, the lesion is always centrally located and a significant portion of the image. This observation is used to remove objects near the edge of images, which tend to be bandages or labels.

Next, small objects that could not be the lesion are filtered out. Often, after HS thresholding, insignificant circular objects are scattered throughout the image. These objects can be removed from analysis by eliminating all objects smaller than a third of the area of the largest object in the image.

Lastly, specific properties of each remaining object, namely diameter, circumference, and area are extracted. These three properties have known relationships in perfect circles, thus each object is analyzed and the most circular object is identified. After identification, all other objects and noise in the image are removed. The result is an object that is identified as part of the lesion with high confidence (Figure 6). The full lesion will have varied textures and color throughout, and the full shape may be different from the object identified, but this object provides a good starting point for the GrowCut algorithm to act on.

GrowCut-based Delineation

In their work, Vezhnevets and Konouchine (2005) propose a new segmentation technique called GrowCut (Vezhnevets & Konouchine, 2005). The GrowCut algorithm treats an input image as a cellular automaton, where each pixel in the image is considered a cell of the automaton and given a label (Berto & Tagliabue, 2017). For segmentation, the cells are labeled as “foreground” (value = 1) for pixels that are part of the lesion, “background” (value = -1) for skin and other non-lesion objects, or “undetermined” (value = 0) for the remaining pixels. The algorithm then applies automata evolution where the cells compete to capture neighboring cells, thereby changing the labels of undetermined cells and increasing the number of like cells.

Following the lesion detection, the GrowCut method is used to complete the segmentation. All pixels that are identified as part of the lesion are labeled as foreground (value = 1). The background pixels are determined through a series of steps:

1. The input image is converted to grayscale and the location where the change in pixel intensity is the highest is identified.
2. The distance from the center of the identified lesion to this pixel is calculated and set as the lesion radius.
3. All pixels further away from the center of the identified lesion than the lesion radius are labeled as background (value = -1).
4. The remaining pixels are labeled as undetermined (value = 0).
5. Finally, these labels are fed into the GrowCut method and the segmentation is produced.

Skin Lesion Segmentation Using Baseline Algorithms

Other segmentation algorithms are briefly described in this section. These methods are used as baseline methods to compare against our new algorithm in the results section.

Baseline GrowCut Method

The baseline GrowCut method is a cellular automata, which can be described generally as an algorithm working on a lattice of sites (pixels). The cellular automaton is a triplet:

$$A = (S, N, \delta)$$

where S is a non-empty state set, N is the neighborhood system, and $\delta : S^N \rightarrow S$ is the local transition rule, which calculates the state of the system at the next time step. In the baseline implementation of GrowCut, the nearest neighbor system is used for N.

Baseline GrowCut assumes that the lesions are centrally located in the images. Each image is pre-processed using Median and Gaussian filtering to remove noise. Then, a circle generated in the center of the image is assumed to be the lesion and labeled as foreground. Through empirical testing, a radius of 90 pixels was found to give optimal results on the training data set. Finally, steps 3 through 5 as outlined in the “GrowCut-based Delineation” section are performed and the segmentation is obtained.

Active Contours

Active contours use a deformable curve (called a “snake”) to delineate an object outline from a 2D image through energy minimization (Kass, 1988). The curve, or contour, actively moves around the image under the influence of internal spline forces that push it towards image features like edges, while external constraint forces keep the snake near the desired local minimum. Together, by mini-

Table 1. Segmentation of performance measures and how they are calculated. TP = True positive, the number of pixels correctly labeled as foreground (value = 1). TN = True negative, the number of pixels correctly labeled as background (value = 0). FP = False positive, the number of pixels wrongly labeled as foreground. FN = False negative, the number of pixels wrongly labeled as background.

Metric	Formula	Description
Accuracy	$(TP+TN)/(TP + TN +FP + FN)$	Number of pixels labeled correctly divided by the total number of pixels.
Dice Coefficient	$2TP/(TP + FP + TP + FN)$	Focuses on the accuracy of the foreground pixels, and penalizes for wrong labels.
Sensitivity (True Positive Rate)	$TP/(TP + FN)$	Measures the proportion of foreground pixels that are labeled correctly, not considering background pixels.
Positive Predictive Value	$TP/(TP + FP)$	Probability that a pixel labeled as foreground is actually foreground.
Jaccard Index	$TP/(TP + FN + FP)$	Similar to the Dice Coefficient, but penalizes more for wrong predictions.

Table 2. Median performance metrics for proposed algorithm and baseline methods of segmentation on testing data set.

Algorithm	Accuracy	Dice Coefficient	Sensitivity (True Positive Rate)	Positive Predictive Value	Jaccard Index
Proposed Algorithm	0.94	0.87	0.9	0.96	0.77
Baseline GrowCut	0.92	0.76	0.66	1.0	0.61
Active Contours	0.85	0.75	0.96	0.85	0.59
K-Means Clustering	0.87	0.54	0.47	0.97	0.37
Otsu Thresholding	0.87	0.70	0.71	0.89	0.54

mizing the energy function, the forces act to discern the boundary of objects in the image. All pixels inside the boundary are labeled as the lesion. The energy function for the contour can be expressed by the following formula:

$$E_{snake}^* = \int_0^1 E_{snake}(v(s))ds = \int_0^1 (E_{int}(v(s)) + E_{image}(v(s)) + E_{con}(v(s)))ds$$

where $v(s)$ represents the position of the contour, E_{int} represents the internal energy of the bending spline, E_{image} represents the image forces, and E_{con} represents the external constraint forces.

K-Means Clustering

K-means is one of the simplest unsupervised learning algorithms (Hartigan & Wong, 1979). The main idea is to initially separate the data set into k clusters and find the center point of each cluster, known as the centroid of the cluster. The algorithm then assigns each data point to the nearest centroid to create new clusters. This is repeated until the clusters are stable. In the case of skin lesion segmentation, k-means clustering deals with two clusters: foreground and background. Although it has been proven that the procedure terminates, it is not guaranteed to find the optimal configuration and is sensitive to initial choice of centroids. In our use of k-means, initial centroids were chosen that are most obviously foreground and background from the previous steps.

Formally, given a set of pixels (x_1, \dots, x_n) , k-means clustering aims to partition the n pixels into k ($k = 2$ in this instance) sets $S = \{S_1, \dots, S_k\}$ so as to minimize the within-cluster sum of squares (variance). The goal is to minimize the objective function:

$$J = \sum_{j=1}^k \sum_{i=1}^n \|x_i^{(j)} - c_j\|^2$$

where $\|x_i^{(j)} - c_j\|^2$ is the distance measure between a pixel $x_i^{(j)}$ and the cluster center c_j (“Clustering - K-means,” n.d.).

The algorithm is as follows:

1. Select two points, one representing the centroid of the foreground cluster and one representing the centroid of the background cluster.
2. Assign each pixel in the image to the nearest cluster.
3. When all pixels are assigned, recalculate the position of the centroids of the foreground and background.
4. Repeats steps 2 and 3 until the centroids do not move.

Otsu’s Method

Otsu’s Method uses a simple threshold on the grayscale image pixel intensities (Otsu, 1975). The image is first converted to grayscale and an optimal threshold value is found to separate the most frequent clusters of pixel intensities. All pixels with intensities below the threshold are labeled as foreground and all pixels with intensities above the threshold are labeled as background. The formula for computing the Otsu threshold is:

$$\sigma_\omega^2(t) = \omega_0(t)\sigma_0^2 + \omega_1(t)\sigma_1^2(t)$$

where $\sigma_\omega^2(t)$ represents the weighted sum of variances of the two classes (foreground and background). The weights (ω_0 and ω_1) represent the probabilities that the two classes will be separated by the threshold (t), while σ_0^2 and σ_1^2 represent the variance of the classes. Otsu’s method finds the threshold value that minimizes $\sigma_\omega^2(t)$.

Statistical Analysis

The Wilcoxon signed rank test was used to determine whether there is a statistically significant improvement of the median accuracy of the proposed algorithm compared to each baseline algorithm (Wilcoxon, 1945).

RESULTS

The accuracy of the overall segmentation was evaluated in terms of whether each pixel is correctly labeled according to the ground truth image, which defines the correct result. Several commonly used metrics are summarized in Table 1. For the proposed segmentation algorithm and each baseline method, the five metrics described in Table 1 were measured and the median across the testing image set was calculated and presented in Table 2. For these metrics, higher values represent a segmentation result more similar to that of the ground truth image.

The performance metrics shown in Table 2 and the Dice Coefficient box plot (Figure 7) demonstrate that the proposed algorithm is superior to all of the baseline methods (described above), outperforming all baseline methods in the three most important accuracy measures: accuracy, Jaccard index, and Dice Coefficient. The median accuracy of 0.94 demonstrates the strong performance and consistency of the proposed algorithm. The median accuracy of the proposed segmentation algorithm was found to be a statistically significant improvement ($p < 0.001$, Wilcoxon signed rank

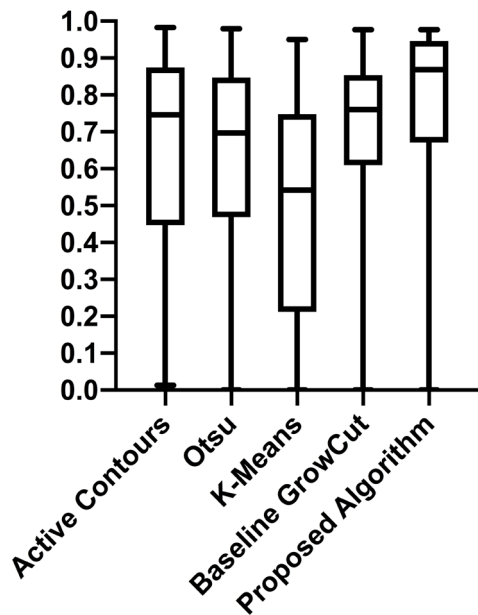


Figure 7. Dice Coefficient boxplot demonstrating how the proposed algorithm outperforms baseline methods.



Figure 8. Example of a case where the algorithm fails due to presence of two lesions in an image. Original image (A) with an overlaid segmentation contour (B), and ground truth segmentation contour (C) in blue, ISIC16 training dataset (index 285).

test) over the four other implemented segmentation methods. A p-value of less than 0.01 is sufficient evidence to conclude a statistically significant difference (Fenton & Neil, 2012). Thus, it can be concluded that the proposed GrowCut-based algorithm is superior to all of the baseline methods. The ISIC16 dataset contains images with a variety of different lighting and defects, proving the merit of the proposed algorithm. The algorithm, however, struggles with some irregular images, as shown by the outlier segmentations in Figure 7. The Sensitivity and Positive Predictive Value measures, while being useful to determine the accuracy of foreground pixels, only measure one aspect of the segmentation and thus, are not representative of the entire segmentation.

DISCUSSION

Accurate automated segmentation of skin lesions in dermoscopic images could assist in the development of reliable computerized diagnosis systems and help decrease the burden of manual evaluation of this pre-biopsy exam. This work proposes a new automated segmentation pipeline based on cellular automata, which on a common dataset, performs favorably compared to the baseline

methods. Specifically, the results demonstrate that the algorithm provides a statistically significant improvement to the accuracy of segmentation, as measured by the Wilcoxon signed rank test, while maintaining high sensitivity and positive predictive value. Currently, further research is required to provide sufficient support that automated diagnostic systems will be able to replace human observation. To garner that support, the next steps will require coupling this segmentation method with a diagnostic algorithm to show that an automated diagnostic pipeline for melanoma can be more effective than manual diagnosis and is capable of implementation in clinics.

Examination of the images that had poor segmentations revealed that they shared some common characteristics. Most prominently, the lesion borders were hard to discern and the edges needed to be inferred using incomplete edge fragments. Secondly, there were also a few images that looked like two separate lesions but which needed to be combined to form the ground truth lesion (Figure 8). These outlier images reduce the accuracy of the algorithm and could be addressed in future work. Solutions may include implementing a deeper feature analysis for all identified objects, aiming to identify characteristics that could differentiate the lesion-like objects or enhancing the algorithm for cases with multiple lesion-like objects.

This paper proposes a novel algorithm for addressing the problem of automatic lesion segmentation for skin lesion dermoscopy. This research can lead to better diagnosis capability for skin cancer. More broadly, any image-based automated diagnostic technique will benefit from improved segmentation.

ACKNOWLEDGMENTS

The author would like to thank Dr. Daniel Rubin for the opportunity to investigate this algorithm within his lab. This work was only possible with his help and encouragement. The author would also like to thank Dr. Alfiya Galimzianova, the mentor who guided this work and provided valuable help and feedback.

REFERENCES

- Abbasi NR, Shaw HM, Rigel DS, et al. Early diagnosis of cutaneous melanoma: revisiting the ABCD criteria. *The Journal of the American Medical Association* (2004), 292:22, 2771–2776.
- American Cancer Society. (2018). Key Statistics for Melanoma Skin Cancer. Retrieved September 15, 2018, from <https://www.cancer.org/cancer/melanoma-skin-cancer/about/key-statistics.html>
- American Cancer Society. (2017). Tests for Melanoma Skin Cancer. Retrieved December 30, 2017, from <http://www.cancer.org/cancer/melanoma-skin-cancer/detection-diagnosis-staging/how-diagnosed.html>.
- Berto F and Tagliabue J. Cellular Automata. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2017). Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/fall2017/entries/cellular-automata/>
- Covalic. (2016). Retrieved December 31, 2017, from <https://challenge.kitware.com/#challenge/560d7856cad3a57cfd481ba>
- Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* (2017), 542:7639, 115–118.
- Fenton NE and Neil M. *Risk Assessment and Decision Analysis with Bayesian Networks*. CRC Press, 2012.
- Gutman D, Codella NCF, Celebi E, et al. (2016, May 4). *Skin Lesion Analysis*

toward Melanoma Detection: A Challenge at the International Symposium on Biomedical Imaging (ISBI) 2016, hosted by the International Skin Imaging Collaboration (ISIC). *arXiv [cs.CV]*. Retrieved from <http://arxiv.org/abs/1605.01397>

Hartigan JA and Wong MA. *Algorithm AS 136: A K-Means Clustering Algorithm. Journal of the Royal Statistical Society. Series C, Applied Statistics* (1979), 28:1, 100–108.

Herschorn A. Dermoscopy for melanoma detection in family practice. *Canadian Family Physician Medecin de Famille Canadien* (2012), 58:7, 740–745, e372–e378.

ISIC Archive. (n.d.). Retrieved December 31, 2017, from <https://isic-archive.com/>

Kass M. Active Contours.pdf. *International Journal of Computer Vision* (1988).

Matteucci M. (n.d.). Clustering - K-means. Retrieved January 28, 2018, from https://home.deib.polimi.it/matteucc/Clustering/tutorial_html/kmeans.html

National Cancer Institute. (2017). Melanoma of the Skin - Cancer Stat Facts. Retrieved December 30, 2017, from <https://seer.cancer.gov/statfacts/html/melan.html>

Otsu N. (1975). from Gray-Level Histograms. *Planning Perspectives: PP*, 285, 296.

Patel NV. (2017, December 11). Why Doctors Aren't Afraid of Better, More Efficient AI Diagnosing Cancer. Retrieved August 2, 2018, from <https://amp.thedailybeast.com/why-doctors-arent-afraid-of-better-more-efficient-ai-diagnosing-cancer>

Robinson JK. Sun exposure, sun protection, and vitamin D. *The Journal of the American Medical Association* (2005), 294:12, 1541–1543.

Stanganelli I. (2017, November 17). Dermoscopy: Overview, Technical Procedures and Equipment, Color. Retrieved December 31, 2017, from <https://emedicine.medscape.com/article/1130783-overview>

Vezhnevets V and Konushin V. (2005). "GrowCut" - Interactive Multi-Label ND Image Segmentation by Cellular Automata. Graphicon. 1.

Wilcoxon F. Individual Comparisons by Ranking Methods. *Biometrics Bulletin* (1945), 1:6, 80–83.