Quantitative Dermatopathology of 11 Skin Conditions through Transfer Learning of Tissue Photon Interaction Statistical Physics

Sumit Agrawal SUMIT@RESEARCH.SKINCURATE.COM
Sri Phani Krishna Karri SRIKRISHNA@SKINCURATE.COM
Kausik Basak KAUSIK@SKINCURATE.COM
Tamoghna Ojha TAMOGHNA@SKINCURATE.COM
Debdoot Sheet DEBDOOT@SKINCURATE.COM
SkinCurate Research Private Ltd. Kharagpur, 721306

Abstract

Quantitative dermatopathology is being widely employed in clinical practices for precise detection of various skin conditions. Dermatoscopes are placed in contact with the skin for visual observation and this (a) risks cross-contamination between investigated subjects and (b) limits instrument use when investigating infectious lesions and wounds. This paper presents a framework for quantitative dermatopathology using consumer grade camera for contact free imaging of skin conditions. Inductive transfer learning of tissue-photon-interaction (TPI) statistical physics modeled for 11 skin conditions estimates TPI as a spatially localized poisson process of photons sensed by the RGB sensors, and utilizes this knowledge to detect skin condition. Inter-camera and illumination variations are compensated by including a maximum likelihood estimation (MLE) based photon density normalization across all images in the dataset. 8-folded cross-validation experiments over a dataset of 440 images of 11 skin conditions provided the condition identification accuracy of $94.55 \pm 5.05\%$ within the top-3 predictions.

1. Introduction

Skin diseases have been globally reported to have over $1,000$ clinical conditions, ranging across benign like acne to the malignant ones like basal cell carcinoma. Dermatoscope is an assembly of a magnifier, illumination source, a transparent plate and a transparent insulation between the instrument and the skin which is used by clinicians for detailed investigation of a wide variety of skin lesions. Being placed directly on the skin of patients, dermatoscopes may harbour potentially pathogenic bacteria. Alternative current approaches employ consumer grade cameras which provide the opportunity of contact-free imaging ruling out the risks of cross-contamination.

In recent years, there has been an increased demand for an automatic system capable of recognizing the skin disease. An effective computer aided diagnosing system not only helps patients with no or little access to health services but also benefits typical GPs who have received minimal dermatology training.

Quantitative dermatopathology is one such system of using computational techniques for analyzing dermatological images for precision diagnosis. These systems utilize methods contributed to the use of either digital dermatoscopy or consumer-grade camera images under uncontrolled illumination conditions such as: learning based analysis with human-in-loop feature extraction (Razeghi et al., 2013). Automatic classification of skin conditions has not been extensively explored to our knowledge. The conventional approaches generally account to identifying benign vs. malignant (Ganster et al., 2001) but not skin condition (Razeghi et al., 2013), which is an evident clinical challenge and is solved in this paper. The prior art is limited to (i) human-in-loop free skin condition identification, (ii) imaging device and illumination independent functioning, (iii) pixel level lesion identification in the skin condition images. The challenge lies in robustly classifying skin conditions with consumer graded camera imaging, hence leveraging the potential of computer aided-diagnosis. Our method

Appearing in Proceedings of the 2^{nd} Indian Workshop on Machine Learning, IIT Kanpur, India, 2016.

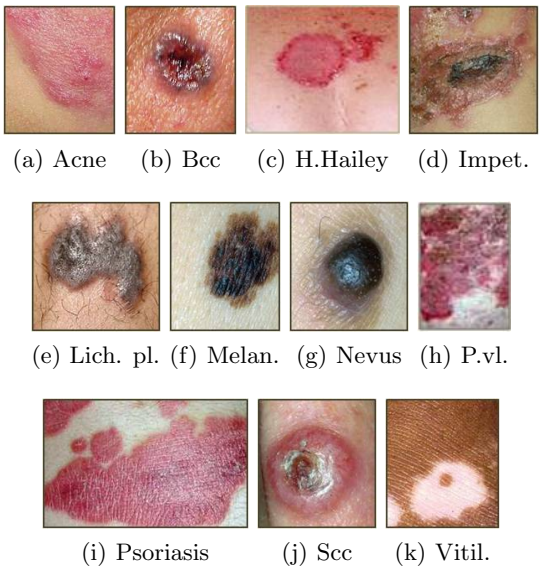

(a) Acne (b) Bcc (c) H.Hailey (d) Impet.

(e) Lich. pl. (f) Melan. (g) Nevus (h) P.vl.

(i) Psoriasis (j) Scc (k) Vitil.

Figure 1. Example images of the 11 skin conditions considered for quantitative dermatopathology.

employs statistical physics model for transfer learning of tissue-photon-interaction to segregate 11 skin disorders. Theortical and experimental studies (Healey & Kondepudy, 1994) well establishes the poisson distribution of photon sensing using standard CCD or CMOS digital sensors. The flow of our solution using transfer learning framework (Pan & Yang, 1994) is exposed as, (i) subjects imaging with consumer grade cameras under uncontrolled illumination (ii) normalizing the photon densities across images (iii) modeling the statistical physics of tissue-photon-interaction (TPI), (iv) transfer learning of TPI for identification of lesions at pixel precision, and (v) classification of the subject image among the 11 skin conditions(Fig. 1).

2. Problem Statement

Let I be an RGB image of the skin acquired by a consumer grade camera under natural illumination. Every pixel at coordinate (x, y) on the image is associated with three intensity values $(I_r(x, y), I_g(x, y), I_b(x, y))$, where I_r, I_g and I_b are the intensities sensed by opto-electronic sensors corresponding to red, green and blue optical spectral bands respectively. These band specific readings are stochastic in nature and follow a statistical physics model, Φ (Healey & Kondepudy, 1994), Φ is defined as a function of the optical response of the camera, spectral incident irradiance pattern and spectral response of TPI model. With a maximum likelihood estimation (MLE) of photon densities across multiple cameras, response of TPI remains the only variable factor affecting Φ. Since the model Φ varies over 11 different skin disorders and a complete range of patients, we propose to learn Φ from the information $(I_r(x, y), I_g(x, y), I_b(x, y))$ from annotated training images using a set of weak learners to have knowledge represented as a set $\Theta = \{\Phi\}$, which then facilitates to create a tissue specific photon interaction statistical physics model capable of segregating 11 skin disorders.

Formal definition: The probability $p(\omega|I, (x, y))$ of detecting a tissue of type ω is the response of a function $H(\omega|\Theta, I, (x, y); \{I\}_{train})$ that uses the knowledge of TPI locally learned on the sample image I and previously learned $\Theta = \{\Phi\}$ using the training set $\{I\}_{train}$.

3. Exposition to the Solution

The given problem can be modeled as an inductive transfer learning problem (Pan & Yang, 1994) where we improve the performance of a machine learning task by including knowledge acquired while solving a related task at an earlier stage. The knowledge gets transferred from *source task* (accomplished on the *source domain*), and is imbibed at the *target task* (accomplished on the *target domain*). Here knowledge acquired at the *source* by modeling of TPI statistical physics, is transferred to the *target* for probabilistically classifying tissues among $\omega \in \Omega = \{$healthy skin, acne rosacea, basal cell carcinoma, hailey hailey, impetigo, lichen planus, melanoma, nevus, pemphigus vulgaris, psoriasis, squamous cell carcinoma, vitiligo$\}$ utilizing the function $H(\omega|\Theta, I; \{I\}_{train})$. Both *source* and *target* domain are the RGB images acquired using consumer grade cameras under uncontrolled illumination, and are MLE normalized. The normalization process, *source* and the *target* tasks are described next:

3.1. MLE Normalization of Images

Our approach is generic to adapt to the variations across the images acquired from multiple consumer grade cameras. In order to incorporate this functionality we correct for (i) *white balance* shifts introduced by uncontrolled illumination conditions using the white-patch approximation rule with green channel as the reference standard; and (ii) *sensor response normalization* across all optical wavelengths across multiple cameras by matching the histogram of each color channel to the MLE of the histogram of the corresponding channel over all images in the dataset.

3.2. Modeling of TPI Statistical Physics

The photon induced voltage *(D)* read out by the camera as a result of the TPI is formulated as: $D =$

$(KR_{ep}T + N_{DC} + N_S + N_R)A + N_Q$, where K is the external quantum efficiency of the sensor (Volts /electron), T is the typical integration time of the sensor. N_{DC}, N_S, N_R and A are the dark current noise, shot noise, readout noise and amplification factor of the readout circuitry respectively. N_Q is the quantization noise of the ADC. The parameter R_{ep} is the rate of photon induced electron generation at a site on the camera photosensor and is defined as (Healey & Kondepudy, 1994)

$$R_{ep} = \int_{\lambda}\int_{y}\int_{x} B(x,y,\lambda)S_r(x,y)q(\lambda)dxdyd\lambda \quad (1)$$

where (x, y) are continuous coordinates on the sensor plane, $q(\lambda)$ is the internal quantum efficiency of the detector (electrons/ Joule) as a function of wavelength of incident radiation λ. $S_r(x, y)$ is the spatial response of the collection site on the sensor. The spectral irradiance pattern $B(x, y, \lambda)$ (Watt / unit area) incident on the sensor is modeled as

$$B(x,y,\lambda) = [R(x,y,\lambda)L(x,y,\lambda) * p(x,y,\lambda)]t(\lambda) \quad (2)$$

where $*$ is the spatial convolution operator, $p(x, y, \lambda)$ is the point-spread-function of the camera optics, and $t(\lambda)$ is the spectral transmission of the optics. $R(x, y, \lambda)$ is the spatially varying spectral reflectance of the surface being imaged and $L(x, y, \lambda)$ is the spatially varying illumination model.

The parameters $N_{DC}, N_S, N_R, N_Q << KR_{ep}T$ for calibrated sensors (Healey & Kondepudy, 1994), thus $D \approx AKR_{ep}T = R_{ep}T'$. The digitally readout voltage (d) from a sensor is a stochastically sensed value of the induced voltage (D) at an instance and is known to be Poisson distributed (Foi et al., 2001) with

$$f(d|R_{ep}, T') \propto \frac{(R_{ep}T')^d e^{-R_{ep}T'}}{d!} \qquad \lambda \in [\lambda_1, \lambda_2] \quad (3)$$

where $R_{ep}T' = E[d] = var(d)$, with $E[\cdot]$ and $var(\cdot)$ respectively representing mathematical expectation and variance operators, and the system is responsive to the range of optical wavelength $\lambda \in [\lambda_1, \lambda_2]$

Each of the three channels of RGB sensor with photon incidence rates $(R_{ep[r]}, R_{ep[g]}, R_{ep[b]})$ follows the Poisson distribution, hence the parameters for these distributions can be estimated as $R_{ep[\lambda]}T' = E[d_{[\lambda]}]$. We estimate $R_{ep[\lambda]}T'$ at a location $(x, y) \in I$ using samples in a neighborhood of $k \times k$ pixel centered at (x, y), at multiple scales (Willett & Nowak, 1994), where each value of k represents a unique scale and $k \in \mathbb{Z}^+$.

Table 1. Accuracy of 11 skin condition identification in 8−folded cross validation experiments. We present mean and variation across folds, max. and min. accuracy.

Top $k-$pred.	mean± std. dev.(%)	max%	min%
1	66.14 ± 9.06	76.36	50.91
2	90.00 ± 5.14	94.54	81.82
3	94.55 ± 5.05	100.0	83.64
4	97.50 ± 3.21	100.0	90.91
5	98.41 ± 1.80	100.0	94.55
6	99.32 ± 0.94	100.0	98.18

In the set of MLE normalized images, only $R(x, y, \lambda)$ in (2) varies over different pathologies and lesions indicating the TPI. Hence, $R_{ep[\lambda]}$ at multiple scales k is an indicator of $R(x, y, \lambda)$ and can thus be appropriately used to model the knowledge of tissue-photon-interaction. Our knowledge set Θ can now be represented as:

$$\Theta = \{\Phi^k_{[\lambda]}\} \qquad \forall[\lambda] \in \{R, G, B\}, k \in \{k_1, ..., k_n\} \quad (4)$$

which is an ordered vector learned at n different scales for the three channels of RGB sensor.

3.3. Inductive Transfer Learning of TPI Model for Quantitative Dermatopathology

The knowledge acquired as Θ in (4) must be tissue specific set, $\{\Theta\}|\omega \ \forall\omega \in \Omega$ given that $R(\cdot)$ in (2) would be tissue specific. The non-parametric distributed clusters formed by $\{\Theta\}|\omega$, can be efficiently learned using supervised non-parametric learners such as ensemble learners. For learning $\{\Theta\}|\omega$ as the model $H(\omega|\Theta, I, (x, y); \{I\}_{train})$, we have used random forest non-parametric supervised learners (Breiman, 2001), which is an ensemble of multiple decision trees jointly forming the model $H(\cdot)$. Hence, the probability $p(\omega|I, (x, y))$ of detecting a tissue of type $\omega \in \Omega$ is the response of the learned forest $H(\omega|\Theta, I, (x, y); \{I\}_{train})$.

4. Experimental Results - Discussions

The experiments were performed for identifying 11 skin conditions using a dataset of 40 images per class, totaling to 440 images. The images were sources from publicly available sources[1] and are archived[2] along with classifications. Each of the images were labeled by a Dermatopathologist with 15 years of experience.

[1] http://www.atlasdermatologico.com.br/browse.jsf

[2] http://www.skincurate.com/dl/ISBI2016_Dataset.zip

Table 2. Lesion area classification accuracy for11 skin conditions computed image wise over 8−folded experiments.

Skin Cond.	mean ± std. dev.(%)	max%	min%
Acne	84.62 ± 13.13	98.70	65.13
B.C.C.	54.98 ± 7.12	66.45	46.55
H.Hailey	66.11 ± 24.11	93.33	27.64
Impetigo	66.00 ± 21.78	85.39	34.20
Lich. pl.	50.24 ± 22.15	78.79	8.62
Melanoma	60.80 ± 24.10	87.21	20.32
N. Nevus	54.22 ± 11.66	70.94	41.55
P.vulgaris	78.33 ± 15.25	97.54	48.47
Psoriasis	61.60 ± 25.48	91.81	19.66
S.C.C.	60.58 ± 13.09	83.09	48.96
Vitiligo	53.41 ± 9.96	71.18	42.38

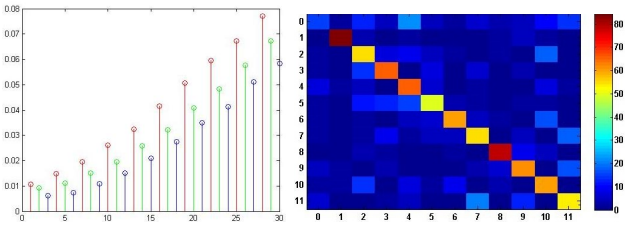

(a) Feature Importance (b) Classification Matrix

Figure 2. Performance of the model in terms of (a) factor importance of each component Φ in the TPI model, and (b) the confusion matrix of classification over 11 skin cond.

An 8-fold cross validation experiment was performed using the dataset with no tested samples shared across folds and preserving the *priori* in each fold to that of the dataset. Table 1 presents the accuracy of skin condition identification in our experiments, where accuracy is reported in terms of true-positive within the top $k-$predictions. We achieve $94.55 \pm 5.05\%$ accuracy in the top 3 predictions, while outperforming the prior art at accuracy of 25.12% (Razeghi et al., 2013) and the average lesion area classification accuracy over all the classes is $62.80 \pm 10.62\%$.

Table 2 and Fig. 2(a) summarize the pixel classification accuracy for each of the 11 skin conditions over 8 folds. Conditions like acne rosacea (Fig. 1(a)) which are visually distinct achieve comparably higher accuracy, while conditions like psoriasis, melanoma, hailey hailey (Fig. 1) exhibit high variance, in conjunction with their large intra-class variance in appearance. Fig. 2(a) shows the (factor) importance map for each of the TPI model constituents. The estimators corresponding to the red wavelengths are observed to be associated with higher importance compared to green and blue, and the importance increases over scales.

5. Conclusion

We have proposed here a method for quantitative dermatopathology using (i) consumer grade camera imaging with uncontrolled illumination and (ii) inductive transfer learning of tissue-photon interaction statistical physics for skin condition identification. We include an MLE based photon density normalization scheme to compensate for inter-camera and illumination variation. We experimentally evaluate its performance over a dataset of 11 commonly encountered skin conditions and report an accuracy of 94.55%. Future scope is to extend this approach over a larger classes of skin conditions imaged with consumer grade cameras.

Dual Submission

This paper is currently under review for *IEEE Journal of Biomedical and Health Informatics.*

References

Breiman, L. *Random forests*, *Mach. Learn.* vol. 45, no. 1, pp. 5-32, 2001.

Foi, A., Trimeche, M., Katkovnik, V. and Egiazarian, K. *Practical poissonian-gaussian noise modeling and fitting for singleimage raw-data*, *IEEE Trans. Image Process.*, vol. 17, no. 10, pp. 1737-1754, 2008

Ganster, H., Pinz, A., Rohrer, R., Wildling, E., Binder, M. and Kittler, H. *Automated melanoma recognition*, *IEEE Trans. Med. Imag.*, vol. 20, no. 3, pp. 233–239, 2001

Healey, G.E. and Kondepudy, R. *Radiometric ccd camera calibration and noise estimation* , *IEEE Trans. Pat. Anal.*, Mach. Intell., vol. 16, no. 3, pp. 267-276, 1994

Pan, S. J. and Yang, Q. *A survey on transfer learning*, *IEEE Trans. Knowledge Data Engg.*, vol. 22, no. 10, pp. 1345-1359, 2010.

Razeghi, O., Fu, H., and Qiu,G. *Building skin condition recogniser using crowd-sourced high level knowledge*, *Med. Imag. Under. and Anal.*, pp. 225230, 2013

Willett, R. M. and Nowak, R. D. *Multiscale poisson intensity and density estimation*, *IEEE Trans. Inf. Th.*, vol. 53, no. 9, pp. 31713187, 2007.