

McGILL UNIVERSITY

**On the
Automated Regional Classification
of Early Breast Cancers**

Joseph Szymborski

Faculty of Medicine
Division of Experimental Medicine
McGill University, Montréal
December, 2017

Supervised by:
Prof. Luke M. McCAFFREY

A thesis submitted to McGill University in partial fulfilment
of the requirements of the Master's degree.

© Joseph Szymborski, 2017

Abstract

The harbinger of a far more lethal and difficult-to-treat disease, early breast cancer is a key, but often overlooked, point of study. Early breast cancers are often difficult to detect and classify, and the subjective nature of the histopathological methods used to detect them often result in dissenting diagnoses between clinicians. These issues can be mitigated through the use of quantitative computational methods, although existing solutions do not account for the heterogeneity of tumours, and often do not operate in manner that is transparent to the clinician; precluding their utility in a clinical setting.

This work presents two models which address the concerns of tumour intraheterogeneity by detecting, classifying and annotating regions within fields of breast biopsy sections that appear to belong to early lesions. The PPreCOGG model, which is a GPU-accelerated texture-based classifier, is able to generate pixel-resolution annotations of cell-patterning that is characteristic of early lesions with robust accuracy ($\approx 94.3\%$ average accuracy in synthetic benchmarks). DeepDuct is a deep learning model that provides accurate and transparent localisation and classification of lesions using gradient-based class activation maps (Grad-CAM). These two models illustrate that it is possible to develop clinically relevant classifiers that can achieve robust accuracy and account for tumour heterogeneity and model transparency.

Résumé

Signe annonciateur d'une maladie létale et difficile à traiter, le cancer du sein au stade précoce est un sujet d'étude qui malgré son importance, est souvent négligé. Le cancer du sein au stade précoce est difficile à détecter et à classifier. Les méthodes histopathologiques utilisées pour détecter ces cancers sont subjectives, et résultent souvent en des diagnostics variables dépendant du professionnel responsable du. L'utilisation de méthodes informatiques quantitatives peut atténuer ces problèmes, mais les solutions existantes ne prennent pas en compte l'hétérogénéité intratumorale. Le fonctionnement souvent obscur de ces méthodes constitue également un obstacle à leur utilisation par des professionnels dans un contexte clinique.

Nous avons développé deux méthodes informatiques afin de résoudre le problème lié à l'hétérogénéité intra-tumorale. Ces dernières sont basées sur la détection, la classification et l'annotation de régions situées dans des sections de biopsies mammaires qui correspondent à des lésions précoces. Le modèle « PPreCOGG » est un classificateur, accéléré par la GPU, qui utilise les textures pour permettre l'annotation de régions contenant des cellules présentant des caractéristiques de cancers du sein précoces. Cette annotation est basée sur la résolution de pixels et permet une détection très précise des régions concernées (environ 94,3% d'exactitude moyenne sur les références artificielles). « DeepDuct » correspond à un modèle d'apprentissage profond permettant une localisation et une classification précise des lésions, via l'utilisation de l'algorithme « Grad-CAM ».

Ces deux méthodes démontrent qu'il est possible de développer des classificateurs capables de détecter l'hétérogénéité intra-tumorale de façon précise dans le cancer du sein, tout en restant faciles d'utilisation pour les cliniciens.

Acknowledgements

Habemus Thesim!

Eminentissimum ac reverendissimum Documentum, et cetra...

(Scribimus latinum sine intellectum)

First and foremost, I'd like to thank the nameless stranger that is responsible for enforcing the deadline for the submission of this thesis; but for their unwavering absolutism this thesis would exist in a perpetual state of being nearly done.

The rest follows in no particular order.

I'd like to express my deep gratitude to my supervisor, Prof. Luke McCaffrey, who graciously offered his unwaivering support, undivided attention, and inspired insight to myself and all others in our lab. Thank you for this opportunity, on which I'll always think warmly of and be better for.

Fellow students and fellows of McCaffrey lab have, without hesitation or question, taken the time to answer dumb questions, show me their magic ways, and commiserate over the seemingly innumerable and sometimes unimaginable ways an experiment could go wrong. They're also all people I'm very proud to call my dear friends. Thank you Andrew, Daya, Carlis, Christina, Li-Ting, Maia, Ruba, Sara, Sudipa, Virginie and our newest additions Aislynn and Rebecca. Un grand merci à Sophie et Virginie pour avoir corrigé mon résumé français!

Some of my oldest friends and most esteemed peers, thanks to the DRMTM for their invaluable support: AD, AL, AV, GK, & PH. #4MOREYEARS.

My family has always supported me in all things, so I'd like to thank my parents, grandparents, and brother for never failing to do all they can and more to see me succeed.

I'd like to acknowledge Arthur Samuel and Marvin Minsky, whose seminal works in machine learning, and their contributions to \TeX and the invention of the confocal microscope, respectively, make them cognitive fathers of this thesis. This thesis was only made possible thanks to the open-source communities that built and maintain the software for everything from the typesetting of this document to the machine learning libraries that power the models in this thesis.

Finally, this thesis is dedicated to the memory of the two laptops that gave up their magic smoke in the name of science and this thesis.

To my Dell Studio 1555 and Asus Zenbook UXA1:

Do not go gentle into that goodnight

Rage, rage against the dying of the backlight

Preface

This thesis was written in accordance with the guidelines outlined by the Faculty of Graduate and Postdoctoral Studies of McGill University.

Contributions

The author wrote the entire thesis, accompanied by editorial comments by Dr. Luke M. McCaffrey. Experimental design, and data analysis contributing to the final results shown in this thesis were performed by the author under the supervision and guidance of Dr. Luke M. McCaffrey.

Licensing

All figures contained in this work belong to their respective copyright owners, and are used here with their permission. The text of this thesis and many of the figures are licensed under the Creative Commons Attribution 4.0 International License. Details about figure copyright, licensing and permission for their use can be found in Appendix B.

Contents

Abstract	i
Résumé	ii
Acknowledgements	iii
Preface	v
Contributions	v
Licensing	v
1 Introduction	2
1.1 An Anatomy of Early Breast Cancers	2
1.1.1 Stages of Early Breast Cancer Progression	2
1.1.2 Intrinsic Molecular Subtypes of Cancer	5
1.2 Current Practices for the Diagnosis of Early Breast Cancers	6
1.3 Machine Learning and the Medical Context	8
1.3.1 Instance-Based Algorithms & their Medical Applications	9
1.3.2 Perceptron-Based Algorithms & their Medical Applications	10
1.3.3 Pattern-based Features for Classifying Medical Imaging	15
2 Design & Methods	17
2.1 PPreCOGG Model	17
2.1.1 PPreCOGG Model Design	17
2.1.2 Software Dependencies	17
2.1.3 Datasets	18
2.1.4 Software and Data Availability	18
2.2 DeepDuct Model	21
2.2.1 Model Design	21
2.2.2 Software Dependencies	21
3 Results	24
3.1 PPreCOGG: A Model for the Per-Pixel Classification of Early Breast Lesions via Gabor Features	24
3.1.1 PPreCOGG Classifies Sub-Regions of Perceptibly Distinct Textures in Synthetic Benchmarks	24

3.1.2	PPReCOGG Classifies Sub-Regions of Different Neoplastic Phenotypes in Synthetic Benchmarks with High Accuracy	27
3.1.3	PPReCOGG Effectively Classifies Sub-Regions of Different Neoplastic Phenotypes in Human Biopsy Samples	27
3.2	DeepDuct: A Deep-Learning Approach to Regional Breast Cancer Classification using Grad-CAM	30
3.2.1	Accurate Classification of Breast Neoplasms in the BreakHis Dataset	30
3.2.2	Activation Mapping of BreakHis ConvNet using Grad-CAM	30
4	Discussion	35
4.1	Robust Regional Diagnosis with PPRECOGG	35
4.1.1	Model Performance	35
4.1.2	Future Directions	36
4.2	DeepDuct as a Model Transparent Regional Classifier	37
4.2.1	Explanations for Neural Network Predictions are Essential for Clinical Use of CADe/x Models	38
4.2.2	Dataset Considerations	39
4.2.3	Future Directions & Improvements	39
	Appendix A: BreakHis Class Codes and Abbreviations	50
	Appendix B: Licensing Information	51

1. Introduction

1.1 An Anatomy of Early Breast Cancers

Breast cancer is both a common and lethal disease, having earned the dubious distinction of being both the most common and second most fatal cancer amongst females in Canada and around the world (Canadian Cancer Society's Advisory Committee on Cancer Statistics, 2015). Breast cancer most commonly arises in the epithelium of the mammary gland's many lactiferous ducts, which form a network that delivers to the nipple the milk that is secreted by the lobules of the mammary gland; which is another common origin of breast carcinomas (*Figure 1.1*). The epithelium of the lactiferous duct is highly organised, with well-defined tissue and cell polarity that is integral to the structure and function of the duct. The tube-like lactiferous duct is comprised of an epithelial inner-layer, which forms a hollow lumen and is surrounded by an outer layer of myoepithelial cells which express smooth-muscle actin (SMA) whose muscle-like contractile properties biomechanically deliver milk along the duct in response to hormonal signalling (Hamperl, 1970).

1.1.1 Stages of Early Breast Cancer Progression

When diagnosing a suspected early breast cancer, pathologists analyse needle-core biopsies with the aim of identifying and classifying any lesions that may be present. Classification of lesions allow medical professionals to better understand the nature of the particular disease, what treatment is most appropriate, and what statistical outcomes are associated with the lesion.

The early stages of breast cancer manifest as pre-invasive, hyper-proliferative lesions that exhibit progressive and gradual deterioration of this epithelial organisation. These lesions belong to four histologically distinct classes: Usual Hyperplasia (UH), Flat Epithe-

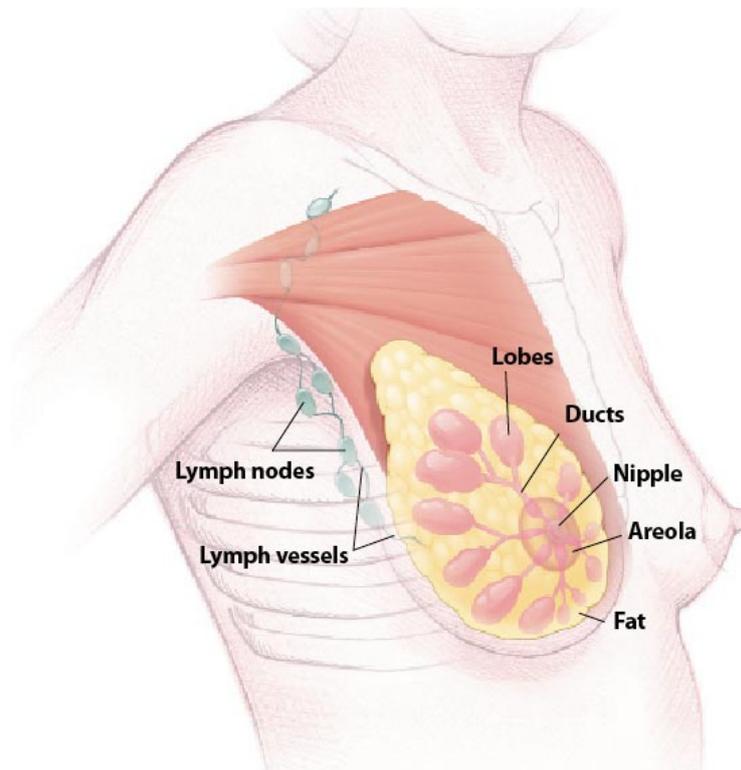


Figure 1.1: Diagram of the human mammary gland and its ducts and lobules. (National Institutes of Health, 2010)

lial Atypia (FEA), Atypical Ductal Hyperplasia (ADH) (or Atypical Lobular Hyperplasia [ALH] when referring to the less common lobular lesion), and Ductal Carcinoma *In Situ* (DCIS).

Ductal or lobular hyperplasias that do not present with abnormal tissue architecture or dysplasia are classified as Usual Hyperplasia (UH), or alternatively Proliferative Disease without Atypia (PDWA). These lesions confer a relative risk of later developing breast cancer as high as 1.9, although this increase in risk is not considered sufficient to warrant any prophylactic measures, including increased follow-up (Mommers *et al.*, 2001). While UH is traditionally believed to progress serially through ADH, DCIS and ultimately IDC due to early Loss of Homozygosity (LOH), more recent cytokeratin immunophenotype and genetic hybridisation analysis has contested the evolutionary relationship between UH and other proliferative breast lesions (O’Connell *et al.*, 1994; Boecker *et al.*, 2002).

ADH lesions are neoplasias of the lactiferous duct that exhibit subtle dysplasia (as evidenced by nuclear hyperchromaticity), and can form micropapillary or cribriform patterns (Page *et al.*, 1959; Dion *et al.*, 2016). Of the estimated one million instances of

benign breast cancer detected in the USA each year, 10% are classified as ADH (Simpson, 2009). While these lesions have been long-known and extensively proven to impart a low relative risk of malignant disease, recent long-term follow-up studies have shown that one in eight individuals will develop more advanced (local or invasive) breast cancers ten years after their diagnosis. This proportion increases to 46% in individuals with more than one atypical foci twenty-five years after diagnosis (Hartmann *et al.*, 2015).

Arising in the terminal duct-lobule unit of the breast, FEA lesions are a purported precursor to early low-grade ductal carcinomas, and in this regard are similar to ADH. Unlike ADH however, FEA lesions are far more uncommon, never present with complex architectural patterns (thus the indication “flat”), and are characterised by multi-layered dilated ascini often made-up of columnar cells (Pinder, 2017). While ADH is suspected to arise from FEA lesions due their frequent coincidence, FEA is not independently associated with a long-term increased risk of breast cancer, leaving the matter unclear (Bombonati & Sgroi, 2011; Lerwill, 2008; Acott & Mancino, 2016).

Benign early lesions go on to progress into localised malignant disease, which in the lactiferous duct is termed ductal carcinoma *in situ* (DCIS). DCIS is classified as a Stage 0 cancer and accounts for 20% of all diagnosed breast cancers in the USA in 2003; representing a 500% increase in occurrence over 20 years (Bleicher, 2013; Kerlikowske, 2010).

While DCIS has a relatively low average standardised mortality ratio (SMR) of 1.8, an estimated 30-50% of cases reoccur as invasive breast cancers (Narod *et al.*, 2015; Page *et al.*, 1982; Betsill *et al.*, 1978). When further stratified by how well the lesion is differentiated, individuals with lesions classified as poorly differentiated (using the European Pathologists Working Group guidelines) have recurrence rates above 60% (Badve *et al.*, 1998). At this early stage of cancer progression, the apical domain of the luminal epithelium has begun to shrink, resulting in abnormally small lumen (a phenotype referred to as “luminal collapse” herein). Our understanding of the processes by which transformed mammary duct epithelium undergoes luminal collapse is still developing, but recent studies have described a mechanism by which luminal tension is lost as myosin II and RhoA activity is greatly decreased at the luminal membrane of DCIS lesions (Halaoui *et al.*,

2017).

The lesion becomes an invasive ductal carcinoma (IDC, or ILC in the lobular instances) when epithelial cells breach the surrounding myoepithelial layer of the duct and infiltrate into extra-cellular matrix (ECM). By this stage, cellular polarity is entirely disrupted and the apical membrane domain has completely disappeared.

1.1.2 Intrinsic Molecular Subtypes of Cancer

While the stage of a breast lesion may indicate its progression, even lesions of the same stage are heterogenous in their gene expression and their response to therapy. To address this heterogeneity, microarray studies of recent years have identified five “intrinsic” molecular subtypes through unsupervised classification of the gene expression data (Perou *et al.*, 2000; Prat *et al.*, 2015).

Understanding the molecular subtype is an important tool for medical practitioners, as breast cancer lesions of different subtypes are associated with different patient outcomes and respond differently to treatment (Engstrøm *et al.*, 2013; Rouzier *et al.*, 2005). An alternative to microarray studies, molecular subtypes can be inferred by testing biopsy tissue for the co-occurrence of a number of molecular markers, as determined by routine immunohistochemical(IHC) studies. These markers range from hormone receptors (estrogen/progesteron receptors, and human epidermal growth factor receptor 2) to proliferation markers (Ki67) (See *Figure 1.2*).

Molecular subtypes of pre-invasive DCIS lesions have been shown to be detectable in a similar fashion to their invasive counter-parts, although relative frequencies between the two being significantly different (Tamimi *et al.*, 2008; Clark *et al.*, 2011). Notably, triple-negative and basal-like phenotypes are occur very rarely in DCIS when compared to invasive disease.

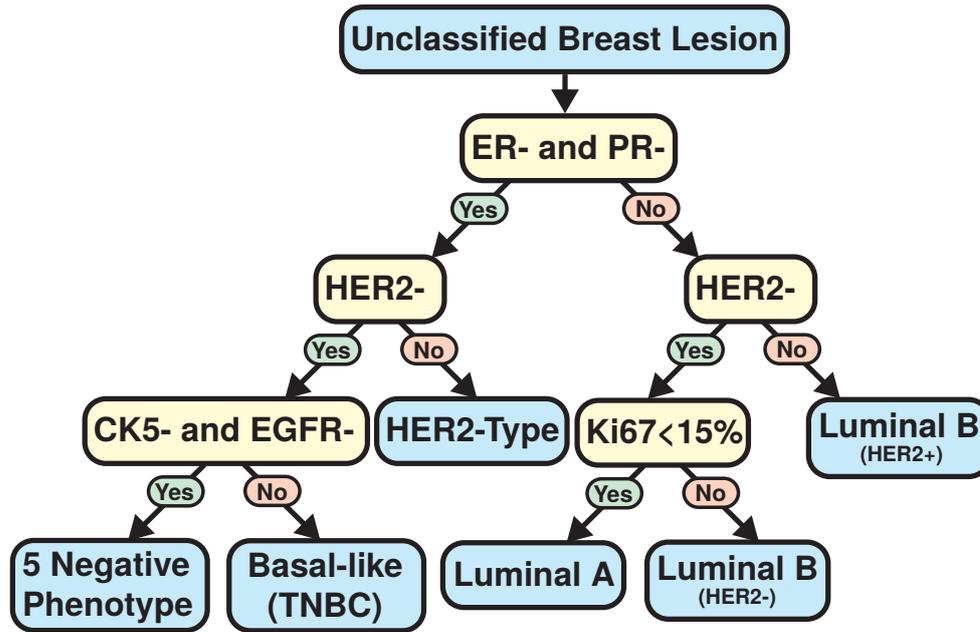


Figure 1.2: Flowchart for the determination of the molecular subtype of breast lesions. Molecular subtype can be determined by staining for the estrogen, progesteron, epithelial growth factor receptors, and human epidermal growth factor receptor 2 (ER, PR, EGFR, and HER2, respectively), as well as the basal/myoepithelial marker cytokerin 5 (CK5) and proliferation marker Ki67. Adapted from Engstrøm *et al.* (2013).

1.2 Current Practices for the Diagnosis of Early Breast Cancers

Patients which exhibit symptoms of mammary neoplasmas, or are otherwise suspected to be at elevated risk due to factors such as age or family history, are subjected to routine screening by mammography. The sensitivity with which mammography is able to detect cancerous lesions, however, is often reduced in young patients or patients whose breast tissue is dense. In the former group, the distinction between invasive and pre-invasive lesions is often unclear, while the dense tissue of the latter group can obscure and mask possible lesions (Ayvaci *et al.*, 2014). In such cases, screening via ultrasound can be a viable alternative, as it offers greater sensitivity (Nothacker *et al.*, 2009).

Suspected lesions detected during the process of screening are diagnosed by histopathological analysis of tissue biopsies (National Comprehensive Cancer Network, 2017). Tissue obtained through biopsy (core-needle, surgical, or otherwise) are sectioned onto glass slides, fixed and stained with relevant histological stains such as hematoxilin and eosin

(H&E).

Despite the standardised and continually refined methods and guidelines clinical pathologists rely on to identify breast cancer lesions from histological sections, there is a great deal of inconsistency and uncertainty that is becoming increasingly apparent. While some types of breast cancer (such as high-grade DCIS and LCIS) are more consistently and reliably identified than others, inter-observer agreement between clinical pathologists is mixed. In a retrospective study, agreement between pathologists for ADH, FEA, and low-grade DCIS regions was only moderate (0.44, 0.47, and 0.47 Cohen’s κ statistic, respectively) (Gomes *et al.*, 2014).

Early breast lesions are associated with increased risk of invasive recurrence, and present important challenges for diagnosis by histopathology. Notably, in a consultation with clinical pathologists, a majority had cited distinguishing atypical ductal hyperplasias (ADH) from usual epithelial hyperplasias (UEH) and ductal carcinoma *in situ* (DCIS) as the most common challenge among their breast biopsy consultations (Putti *et al.*, 2005). These distinctions are significant, as outcomes between these lesions are very different, and are the primary consideration when determining what treatment, if any, to pursue.

Challenges like these can be mitigated in part by computationally assisted detection and diagnosis (CAdE/CADx) software, which analyse medical images in a reproducible and quantitative manner with the aim of making the interpretation of these data by clinicians a less complex and subjective practice. While CAdE software is sometimes used to aid in the screening of breast mammographies, challenges such as dimensional complexity has historically prevented the use of CAdE/x to help interpret histological data (Rangayyan *et al.*, 2007; Madabhushi, 2009).

Despite there being many breast cancer classification models described in the literature, a major limitation of these models is that they fail to accommodate intra-tumour heterogeneity as they commonly adopt whole-field classification modalities (Pareja *et al.*, 2017; Weigelt *et al.*, 2010). Whole-field classification models similarly do not offer insight into so-called “borderline”, which are lesions containing regions exhibiting features of multiple early lesions (Masood & Rosa, 2011).

1.3 Machine Learning and the Medical Context

Machine learning is the study of algorithms that can generalise solutions to problem-spaces without having been explicitly programmed, but instead by “learning” from experiences, employing techniques from the fields of computational statistics, artificial intelligence, and mathematical optimisation.

One of its very first applications, Arthur Samuels coined the term “machine learning” to describe an automated Checkers-playing computer programme he had devised (Samuel, 1959). The Samuels Checkers programme demonstrated that it was possible to have machines solve problems by implementing learning algorithms as opposed to programming solutions in “minute detail” in situations where doing so may be unreasonably onerous or even entirely infeasible.

Medical applications of machine learning can be found early in its history, with Earl Hunt’s application of his Concept Learning System (CLS) for the purpose of medical diagnosis and prognosis as early as 1966 (Hunt *et al.*, 1966). Hunt recognised and stated that machine-learning techniques such as his CLS approach are particularly well-suited for analysing the often large amounts of data collected by medical tests, obviating time-consuming and expensive specialised investigations. The data generated by medical imaging is, in particular, both very rich and difficult to analyse in an efficient, reproducible manner. Machine learning solutions continues to be employed and further innovated for the purpose of better understanding and extracting information from medical imaging (Wernick *et al.*, 2010).

Machine-learning systems for medical applications are preferably of high accuracy and transparent to physicians in its methods, such that unexpected decisions are offered with an explanation that a physician can choose to agree or disagree with (Kononenko, 2001).

The field of machine learning has given rise to a multitude of algorithms, and many of them have been applied to various clinically relevant models of detection and prognosis. Two families of machine learning that have been particularly important to the clinical context will here be surveyed, namely instance-based and perceptron-based machine-learning algorithms.

1.3.1 Instance-Based Algorithms & their Medical Applications

Instance-based learning (IBL) algorithms are machine-learning algorithms that compare features from previous examples to unknown inputs to determine solutions. This is in contrast to other machine-learning algorithm families that generate internal generalised models of a problem-space.

Among IBL algorithms, the k -nearest neighbors (k -NN) and support vector machines (SVMs) are notable for having been widely used for a wide gamut of medical applications.

The k -Nearest Neighbours Algorithm

First described by Fix & Hodges while at the US Air Force as a technical report in 1951, and later formalised by Cover & Hart, the k -NN algorithm is one of the early and fundamental machine learning algorithms, and is used in countless applications today.

The k -NN classification algorithm begins with its training step, whereby an n_F -dimensional feature space is created, where n_F is the number of features per trained data point, all while keeping note of what class each data point in the training feature space belongs to. Subsequent classification steps involves obtaining the features for the unknown data and searching the feature-space generated in the training step for its k nearest features in terms of Euclidean distance of the features, where k is an odd integer. The unknown data is then classified as belonging to the same category as the majority of the k nearest features.

Considered a “lazy” machine learning algorithm, k -NN defers heavy computation from the training stage, where no additional feature processing is required, to the classification stage, which make use of memory-complex search algorithms. Data structures such as k -dimensional (k -d) trees, however, can reduce the memory and computational complexity of these searches (Otaïr, 2013). This, in turn, results in a reduction of time-complexity; in the case of k -d trees, search is performed in $O(\log n)$ time on average.

The k -NN algorithm has long been used for the purpose of clinical diagnostics, with early studies being applied to microcalcification detection systems for mammography, classification of aggressivity of brain tumours, and diagnosis of pigmented skin lesions

(Dhawan *et al.*, 1996; Decaestecker *et al.*, 1997; Dreiseitl *et al.*, 2001). Clinical diagnostic tools based on k -NN classifiers continue to be studied, with innovation focused primarily on algorithm performance and feature engineering (Dhahbi *et al.*, 2015).

1.3.2 Perceptron-Based Algorithms & their Medical Applications

The perceptron has a rich and storied history in the field of artificial intelligence, beginning with Rosenblatt's first descriptions of the algorithm in 1957. At its root, the perceptron is an algorithm to learn some linear binary classifier $f(x)$. The linear classifier trained by a perceptron takes a vector of inputs $x = [x_0, x_1, \dots, x_N]$ and then computes a weighted sum from a weight vector $w = [w_0, w_1, \dots, w_N]$ and some arbitrary bias value b . Sums greater than some threshold θ (typically, $\theta = 0$), are classified as belonging to the first category ($f(x) = 0$), or as belonging to the second ($f(x) = 1$). Alternatively expressed,

$$f(x) = \begin{cases} 0 & \text{if } \left(\sum_{i=0}^N x_i \cdot w_i \right) + b < \theta \\ 1 & \text{if } \left(\sum_{i=0}^N x_i \cdot w_i \right) + b \geq \theta \end{cases}$$

Training of the perceptron occurs by optimising the weights, which is in turn achieved by adding to the weight vector a correction value defined as the product of the difference of the desired and observed values and the input vector. A learning rate η can be multiplied by the correction value to scale the corrections made to weights at each learning step, preventing over-corrections. For an initial weight vector $w^{t=0}$, whose elements are initialised as zero or with random values, to be trained against a training input x_{train} with desired output d , we define a newly trained weight vector $w^{t=1}$ as follows:

$$w^{t=1} = w^{t=0} + \eta (d - f(x_{\text{train}})) x_{\text{train}}$$

This weight optimisation is repeated until the weights converge. Perceptrons were met with much enthusiasm, but Minsky & Papert's landmark 1969 publication *Perceptrons* outlined (among other limitations) the inability of perceptrons to solve problems that are

not linearly-separable, such as the exclusive-or (XOR) functions and the then infeasible computational requirements demanded by complex perceptron models. The limitations of perceptrons illustrated by Minsky & Papert and a period of depressed AI research due to reduced funding and optimism regarding the promise of the field known as the “AI Winter” resulted in perceptrons being largely ignored until the later description of the backpropagation algorithm and the subsequent rise of multi-layer networks (Hendler, 2008).

Artificial Neural Networks

Artificial Neural Networks (ANNs) are systems that overcome the linearity limitations of single-perceptron networks by connecting multiple “hidden” layers of perceptrons (*i.e.*: multi-layer perceptron networks). The perceptrons that make up ANNs are renamed “neurons” in this context and their connections, “synapses”. This new nomenclature is in reference to the biological analogy of the nervous system from which ANNs were first conceived (Kleene, 1956).

ANNs gained in popularity when Paul Werbos described a method for accelerating ANN training through an algorithm for the backwards propagation of errors (backpropagation) from the output layer through to the connected neurons of the network (Werbos, 1982). Hardware limitations, which Minsky & Papert had forewarned the field about in *Perceptrons*, would continue to hamper their utility and adoption until advances in parallel computing, as well as in hardware in the form of graphics processing units (GPUs) would significantly speed up ANN calculations (Simard *et al.*, 2005; Luo *et al.*, 2005).

ANN models have been described throughout the years to help predict breast cancer outcomes. A model described by Floyd *et al.* offered prognostic predictions given a number of features radiologists observed after mammogram such as mass size and calcification morphology (Figure 1.3) which slightly outperformed trained physicians (Floyd *et al.*, 1994). Another ANN model had been described which out-performed TNM staging (a standard staging system which considers primary tumour size, lymph node status, and metastasis) in predicting 5-year outcomes of breast cancer patients, given patient record

information (age, race, payment method, etc...) as well as clinical diagnostic information (hormone status, necrosis, histological grade, etc...) (Burke *et al.*, 1997).

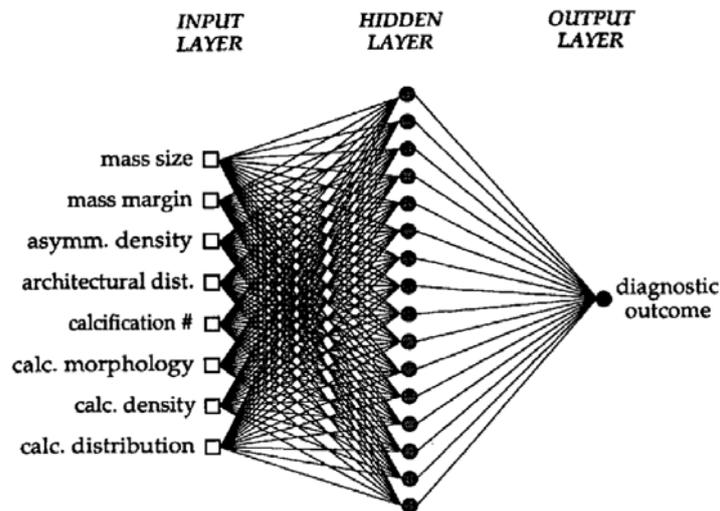


Figure 1.3: Architecture of the ANN described by Floyd *et al.* (1994). Given data interpreted by radiologists from mamograms, this multi-layer network classified benign and malignant tumours with accuracy rates that marginally outperformed trained physicians.

Convolutional Neural Networks

At the heart of an increasing amount of modern CADe/x solutions are the use of convolutional neural networks (CNNs or ConvNets) (Cheng *et al.*, 2016). ConvNets are artificial neural networks that use many hidden layers that typically either apply convolutional or pooling operations at each neuron.

This mix of convolutional and pooling layers was inspired by Hubel & Wiesel who described what he called “simple” and “complex” cells in the visual cortexes of cats and monkeys. Simple cells are stimulated by patterns observed within a specific receptive field by separating the field into inhibitory and excitatory parts, comparable to the role which convolutional layers play. Complex cells, however, have no such fields and responded to stimulus on any part of the field. The behaviour of complex cells is here comparable to pooling layers, which consolidates the output of a hidden layer so that it may be passed on to a single neuron (Hubel & Wiesel, 1968; LeCun *et al.*, 2010).

Models, such as most ConvNets, which make use of multiple hidden layers are de-

scribed as being “deep”, and have historically presented a significant computational challenge. An early landmark applications of ConvNets for image recognition, LeCun’s character recognition model saw much success but required then-impractical computing resources to process images of resolutions much greater than 32×32 pixels. Recent developments in general-purpose computing on graphics processing units (GPGPU), however, have made ConvNets, whose use had previously been regarded as “unrealistic”, practical for a wide-range of applications many years after their first conception (Simard *et al.*, 2005; Crick, 1989). An example of being able to expand on LeCun’s ConvNet architecture due to advances in hardware, VGG16 is a sixteen weight-layer ConvNet architecture (*Figure 2.4*) that has been engineered by (and named for) the Visual Geometry Group at the University of Oxford and obtains robust accuracy rates in general image recognition tasks (Simonyan & Zisserman, 2014). The VGG16 model has been shown to generalise very well to a number of different datasets, and is particularly well suited for localisation tasks, having been awarded first and second place in the classification & localisation task of the ImageNet ILSVRC2014 contest (Russakovsky *et al.*, 2015). This makes the use of the VGG16 architecture well-suited for the function of localisation and classification of medical imagery for the purpose of computational detection and diagnosis.

ConvNets have been successfully used to create very accurate models for the classification of breast cancer lesions. Binary models for classifying benign and malignant lesions, as well as multi-class models for distinguishing between multiple subtypes of breast lesions from H&E stained biopsy slides have established with very high ($> 90\%$) accuracy (Wei *et al.*, 2017; Han *et al.*, 2017). These models, however, are severely limited in that they classify whole imaging fields as belonging to a single class. These approaches entirely ignore the heterogeneous nature of breast lesions and are entirely “black-boxes” for clinicians, offering no added dimensions of information and little understanding as to why the model has interpreted a lesion the way it has. To mitigate this limitation, a classifier would be required to identify, classify and annotate sub-regions that exhibit characteristics of early lesions in medical images of breast biopsies; transparently offering insights into the classifications being made. One such method to so is to generate localisation

annotations with the Gradient-weighted Class Activation Mapping algorithm.

Localisation-Augmented Visualisation of Convolutional Neural Network Using Grad-CAM

Recent work by Selvaraju *et al.* has made the interpretation of ConvNets much more clear by visually annotating inputs with general localisations of objects identified by the model. This technique, named Gradient-weighted Class Activation Mapping (Grad-CAM), can produce heatmaps of areas within an input image that, according to a given ConvNet model, are likely to belong to a given class. Grad-CAM is generalisable to most ConvNet architectures, and does not require to be trained on example localisation annotations.

Grad-CAM has been since used for the dual purpose of localising classified regions and better understanding differences between classes in practical applications ranging from plant stress phenotyping to classifying colorectal polyps (Ghosal *et al.*, 2017; Korbar *et al.*, 2017).

Transfer-Learning for Resource Efficient Training of Neural Networks

Two common limitations of adapting convolutional networks to domain-specific tasks such as classifying medical imagery for computer-aided detection are the large dataset and computational power requirements. These two limitations can be largely addressed by the process of transfer learning, which uses an existing convolutional architecture that has been previously trained (“pre-trained”) on a sufficiently generalised dataset appropriate for the target task (Pan & Yang, 2010). The ImageNet ILSVRC2014 dataset is an example of a widely-adopted, readily-available, and comprehensive general-purpose dataset that is commonly the basis of pre-trained models used for transfer-learning image classification tasks (*Figure 2.3a*) (Russakovsky *et al.*, 2015).

Transfer learning has in-fact been used in a number of computer-aided detection, ranging from thoraco-abdominal lymph node detection and interstitial lung disease classification from chest X-ray and CT scan imaging to classification of skin cancers from dermatoscope imagery (Shin *et al.*, 2016; Esteva *et al.*, 2017).

Transfer learning uses the weights of the many hidden layers of the pre-trained network. The last fully-connected layer of the network is removed from the architecture and a new linear classifier for the network is trained on the new dataset using the pre-trained hidden layers as features.

A limitation of transfer learning is that the later, more specialised, hidden layers of the pre-trained network can lead to reduced accuracy of the model if the original dataset the pre-trained layer was trained against is extremely different from the new dataset. This challenge is usually met by an additional process known as fine-tuning, which continues to train the hidden layers of the pre-trained network against the new dataset using backpropagation (Yosinski *et al.*, 2014).

1.3.3 Pattern-based Features for Classifying Medical Imaging

Features that attempt to quantify properties of patterns have long been studied for the stated purpose of developing classifiers for medical imaging (Mangasarian *et al.*, 1990; Meyer-Baese, 2004).

Pattern detection is particularly applicable to classifying lesions from histology sections of breast tumours, where cell patterning is very significant. The observation that early lesions exhibit distinct cell patterning unique to themselves are reported in the first descriptions of hyperplasia and carcinoma *in situ* of both the mammary duct and lobules by Page *et al.* (1982). The descriptions provided of the early lesion are strongly based on cellular architecture and patterning; distinguishing, for example, ductal and lobular carcinomas *in situ* (DCIS and LCIS) from atypical ductal and lobular hyperplasias (ADH and ALH) by “...round, regular spacing” in the former and their absence in the latter, sometimes exhibiting “...swirls or streaming” Page *et al.* (1982).

Gabor Filters as Texture Features

The Gabor filter is a powerful tool for calculating texture-based features from images (Laine & Fan, 1993). The Gabor filter comes from a family of so-called “wavelet-transforms”, which have been shown to model the manner in which simple cells of the

mammalian visual cortex are stimulated by edges in observed visual fields (Marčelja, 1980). Gabor filters have a number of parameters, such as orientation and bandwidth, that allow them to discriminate between textures differently. Applying the Gabor filter to an image at various resolutions and with several filter orientation parameters has been shown to produce the best results when discriminating between textures (Unser, 1995).

The following thesis describes two models, the PPreCOGG model and the DeepDuct model, which are machine learning models that address the issue of intratumour heterogeneity by classifying subregions within histopathological sections of biopsies of breast lesions.

A model for the **per-pixel recognition of cancers using oriented Gabors on the GPU**, PPreCOGG uses the k -nearest neighbours algorithm to produce high-resolution annotations of diagnostically-relevant cell patterning using texture-based features. The PPreCOGG model achieves a robust rate of accuracy with an average of $\approx 94.3\%$ of pixels being correctly classified on synthetic validation classification tasks, and also is demonstrated to effectively identify sub-regions exhibiting characteristic neoplastic cell patterning in images of human early breast lesions.

The DeepDuct model uses a pre-trained deep convolutional neural network model (namely, VGG16) fine-tuned on a dataset comprised of histological images of breast biopsies classified across eight different lesion types (the BreakHis dataset, *Figure 2.3b*), and combines it with the Grad-CAM algorithm to provide general localisation of the various lesions identified, while providing an opportunity to better understand the model.

2. Design & Methods

2.1 PPreCOGG Model

2.1.1 PPreCOGG Model Design

The PPreCOGG model is a k -nearest-neighbours (k -NN) trained on Gabor features extracted from human breast tumour tissue stained immunofluorescently for E-cadherin.

Gabor features were extracted in a similar fashion as Melendez *et al.* (2008) and is described by *Figure 2.1*. Namely, for each pixel in an image, six windows of increasing size (3×3 , 5×5 , 9×9 , 17×17 , 33×33 , 65×65) centred on the pixel are defined. Each window is then filtered through four Gabor kernels with quarter-turn orientations (*i.e.*: $\theta = \{\frac{1}{2}\pi, \pi, \frac{3}{2}\pi, 2\pi\}$). Each Gabor kernel also has a sinusoidal wavelength of 0.25 pixels ($\lambda = 0.25$), which has been previously described as providing good discrimination in general-purpose texture classification (Manjunath & Ma, 1996).

The mean and the standard deviation of the resulting Gabor energies are then added to a vector for the relevant pixel. This results in a total of 48 features per pixel (6 windows \times 4 orientations \times [1 mean + 1 standard deviation] = 48 feature per pixel).

Gabor feature vectors extracted from the unknown image (i_u) are classified according to a known set of classes $C = \{\text{ADH}, \text{DCIS}\}$ extracted from images that are representative of these known classes ($I_C = \{I_{\text{ADH}}, I_{\text{DCIS}}\}$), as illustrated in *Figure 2.2*. Classification was performed using the k -NN algorithm.

2.1.2 Software Dependencies

The implementation of the k -NN algorithm by the Scikit-Learn (sklearn) python library was used for the purpose of classifying pixels (Pedregosa *et al.*, 2011). Multidimensional Scaling was also implemented by sklearn. Gabor kernels were generated by the OpenCV

library via python language-bindings (Bradski, 2000) and convolved against the extracted windows on the GPU using the Theano library (Al-Rfou *et al.*, 2016). Training and classification tasks were performed on consumer hardware with an nVIDIA GTX 1050M GPU and Intel i7-7700HQ CPU with stock 2.8 GHz clockspeed.

Three-dimensional plots depicted in *Figure 3.1* were generated using the Plot.ly python bindings from Plotly Technologies Inc. (2015). Accuracy and classification plots were created with the matplotlib python library (Hunter, 2007).

2.1.3 Datasets

Brodatz textures were obtained from “*Textures a photographic album for artists and designers*” (Brodatz, 1966).

Human sample tissue was provided by Dr. Atilla Omeroglu in accordance with the guidelines of the McGill Institutional Review Board (IRB; A03-M24-15A) and immunofluorescently labelled by Ruba Halaoui as described in (Halaoui *et al.*, 2017).

2.1.4 Software and Data Availability

The software and documentation for the PPreCOGG model is available at:

<https://github.com/jszym/pprecogg/>

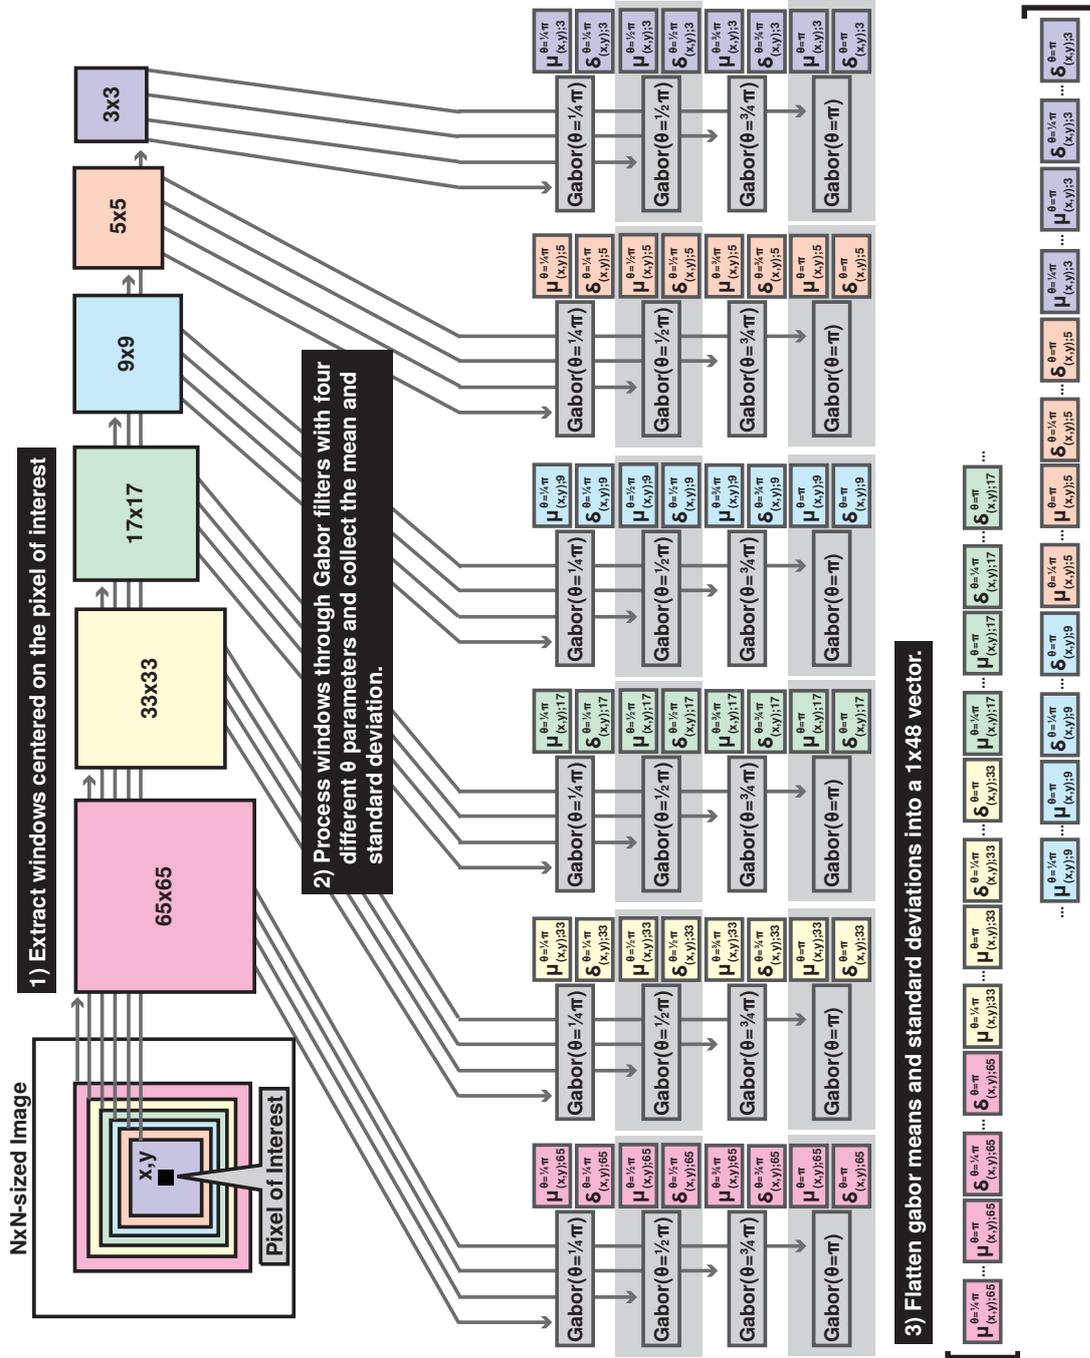


Figure 2.1: Diagram illustrating the method by which Gabor features are extracted from images in the PPRECOGG model.

To reduce computational complexity, images are resized to a resolution of 256×256 pixels. Computation time of training is further reduced by extracting the features of a random sampling of one-quarter of the population of pixels (*i.e.*: For an $N \times N$ image where $N = 256$, $\lfloor N^2/4 \rfloor = 16,384$ pixels) from the total population .

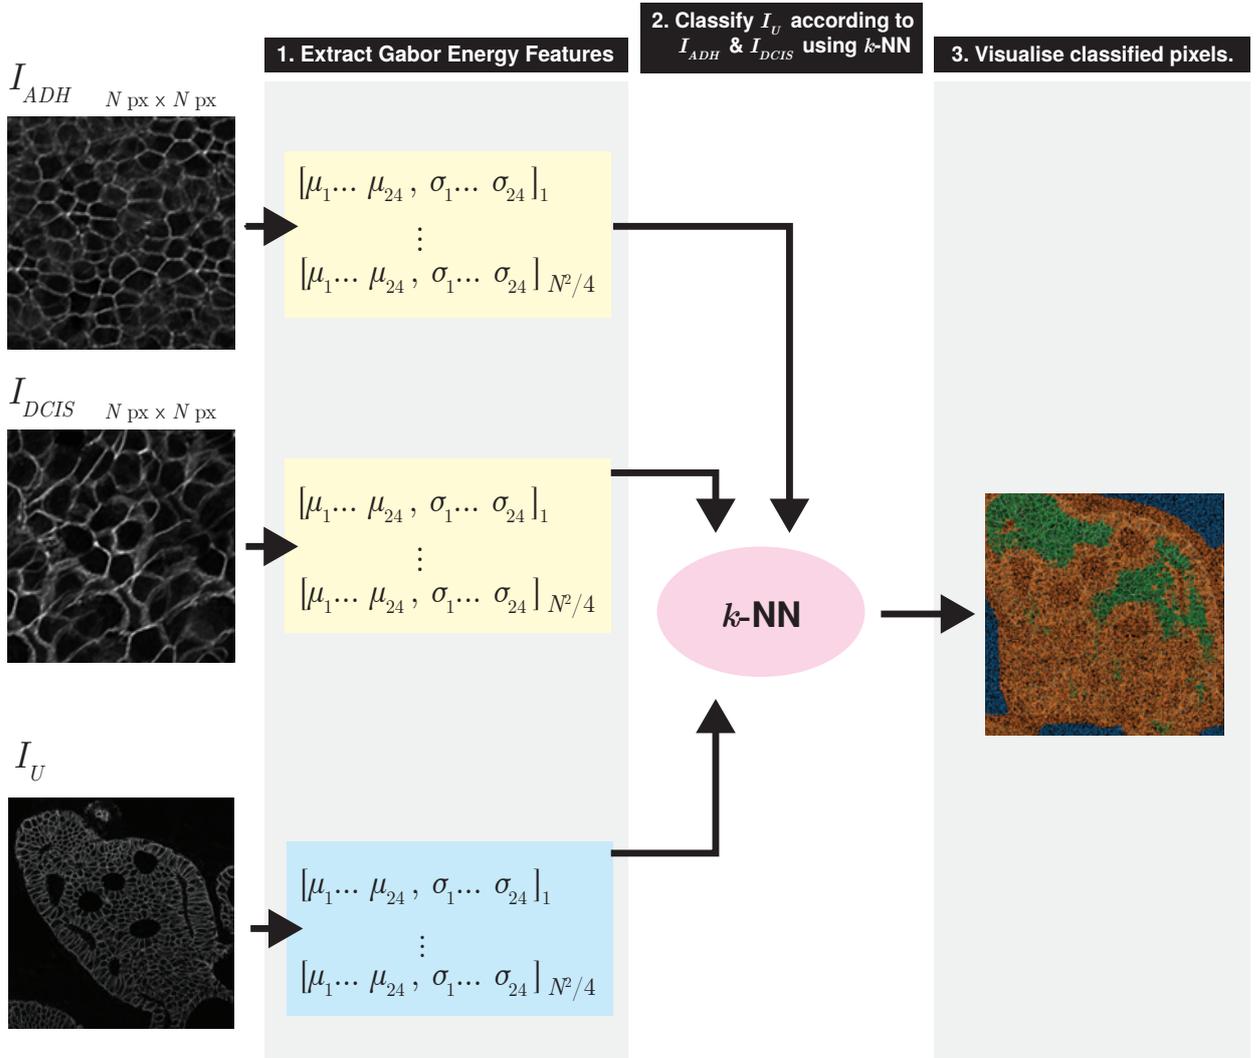


Figure 2.2: Schematic overview of the PPRECOGG model, wherein Gabor features are extracted from images of known classes (I_{ADH} , I_{DCIS}), as well as an image with regions belonging to different classes (I_U). Pixels are classified according to the class of their Gabor features, which are classified using the k -NN algorithm.

2.2 DeepDuct Model

2.2.1 Model Design

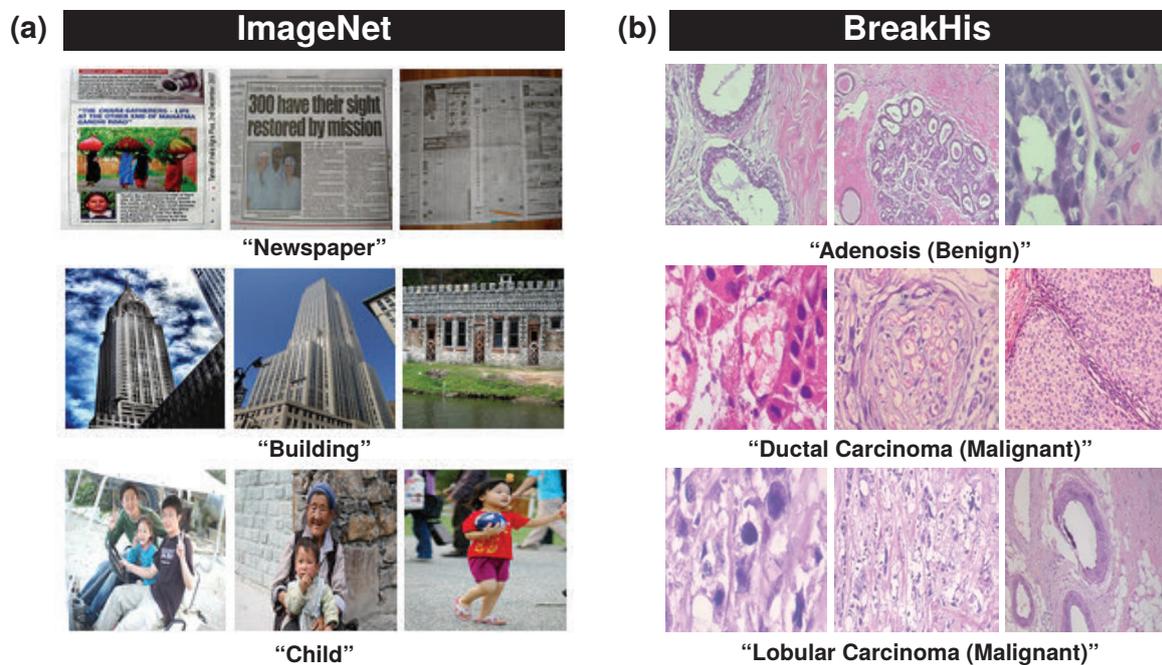
The DeepDuct model begins with a so-called “off-the-shelf” ConvNet architecture; namely the VGG16 model pre-trained on the general-purpose ImageNet dataset (Simonyan & Zisserman, 2014; Deng *et al.*, 2009; Russakovsky *et al.*, 2015). The pre-trained VGG16 model was fine-tuned on the BreakHis dataset, a dataset comprised of approximately 8,000 images of H&E stained human mammary biopsy sections classified according to World Health Organisation (WHO) guidelines (Spanhol *et al.*, 2016; Sunil R. Lakhani, 2012). Abbreviations for the class names present in the BreakHis dataset are used throughout this manuscript, and within the model itself. Refer to Appendix A for a legend of these class codes.

2.2.2 Software Dependencies

The VGG16 model was implemented with Keras using TensorFlow as the back-end (Chollet *et al.*, 2015; Abadi *et al.*, 2016). A Keras/TensorFlow implementation of Grad-CAM implemented in the Keras-Vis library was used to generate attention maps from the BreakHis-trained VGG16 model (Kotikalapudi & contributors, 2017).

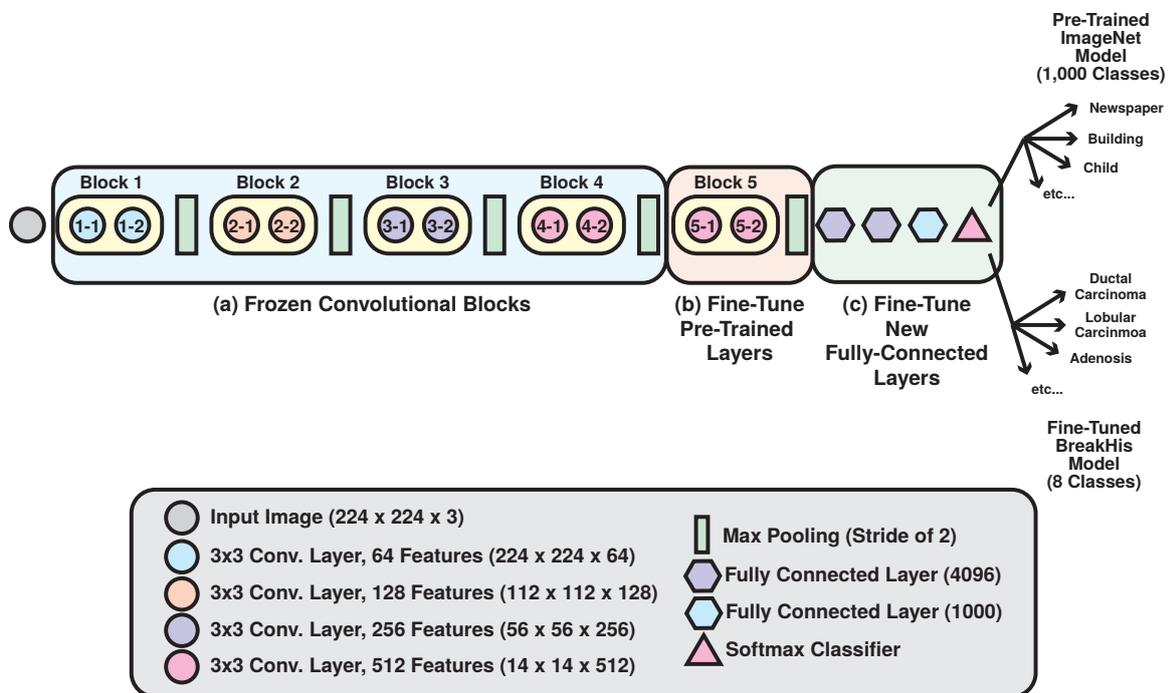
Plots created with the matplotlib and seaborn libraries (Hunter, 2007; Waskom *et al.*, 2014).

Figure 2.3: Examples from the ImageNet and BreakHis Datasets



Legend — Thumbnails from some selected classes of the (a) ImageNet and (b) BreakHis datasets. The DeepDuct model makes use of a ConvNet model pre-trained on the ImageNet dataset and repurposed for classifying images according to the BreakHis dataset via transfer learning. The ImageNet ILSVRC2014 dataset is comprised of $\approx 150,000$ images belonging to 1,000 classes, and the BreakHis dataset is comprised of 8,000 images belonging to 8 classes.

Figure 2.4: Schematic of the VGG16 ConvNet Architecture and its Fine-Tuning



Legend — The VGG16 model is comprised of 16 weight layers, making up five convolutional blocks and a fully-connected classifier. Fine-tuning a pre-trained VGG16 model involves freezing the first four blocks (a), continuing to train the fifth block (b) against the new dataset (in this case, BreakHis) through backpropagation, and finally training a new fully-connected classifier against the new dataset (c).

3. Results

3.1 PPreCOGG: A Model for the Per-Pixel Classification of Early Breast Lesions via Gabor Features

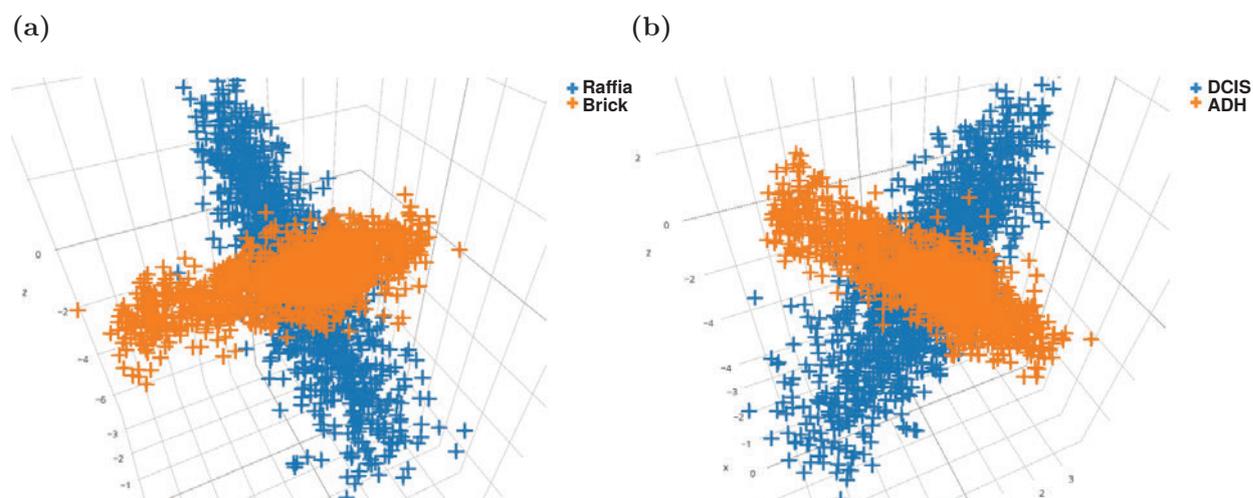
3.1.1 PPreCOGG Classifies Sub-Regions of Perceptibly Distinct Textures in Synthetic Benchmarks

To assess the baseline ability of the PPreCOGG model to distinguish natural patterns from one-another, the PPreCOGG model was trained and benchmarked against the Brodatz textures (namely the “Raffia” and “Brick” textures) compiled in the USC-SIPI image dataset, which is a common image dataset used in the evaluation of image processing and texture recognition community (Weber, 1997). The brick and raffia textures were selected as they are visually distinct, and test images comprised of sub-regions of the two are perceptibly distinct upon visual inspection (*Table 3.1a*).

Multidimensional scaling (MDS) embeddings allow us to project the forty-eight-dimensional Gabor energy features calculated from the Brick and Raffia Brodatz textures onto three dimensions (*Figure 3.1a*). These feature embeddings reveal that while they intersect, the feature space is divided into two distinct planes, each belonging to one of the two classes.

The PPreCOGG model achieved high average accuracy rates of 90% in benchmarks consisting of the Raffia and Brick Wall Brodatz textures (*Table 3.1*).

Figure 3.1: MDS Embedding of Gabor Energy Features from Brodatz and Early Human Breast Lesion Datasets



Legend — Three-dimensional multidimensional scaling (MDS) embedding of Gabor energy features from known classes of the (a) Brodatz dataset and (b) the early human breast lesion dataset.

Table 3.1: Accuracy of the PPreCOGG Model Trained on Brodatz Textures

		Test Image	Test Image	Test Image
		One	Two	Three
(a)	Original Test Image			
(b)	Ground Truth			
(c)	Classified Image			
(d)	Raffia Accuracy Map			
(e)	Brick Accuracy Map			
(f)	Accuracy	90.94%	85.29%	94.00%

Legend — (a) Test images comprised of composites of the Brodatz textures entitled “Raffia” (pg. D84) and “Brick Wall” (pg. D94). (b) The ideal classification (or “ground truth”) of the original test image, where black pixels code for the raffia texture and white codes for the brick texture. (c) A random sampling of pixels classified by the PPreCOGG model. Pixels classified as Raffia are coded by blue points, and and pixels classified as Brick are coded by orange pixels. (d) and (e) Classified pixels are here compared to and overlaid on their ground truth. Pixels which are correctly classified are coded in green, while false-positives are coded in red. (f) Accuracy of the model, as calculated by the quotient of the number of correctly classified pixels and the total number of classified pixels.

3.1.2 PPreCOGG Classifies Sub-Regions of Different Neoplastic Phenotypes in Synthetic Benchmarks with High Accuracy

Similar benchmarks were performed on test patterns composed of images of early lesions (ADH and DCIS) from human patient samples which had been immunofluorescently labelled for E-cadherin. These synthetic benchmarks are meant to simulate and quantitatively measure the efficiency of the PPreCOGG model in the task of classifying whole fields into sub-regions which exhibit cell patterning characteristic to certain early lesions.

MDS embeddings of the 48-dimensional Gabor energy features for the human samples reveal results very similar to the embeddings of the features of the Brodatz textures; two distinct but intersecting planes (*Figure 3.1b*).

The PPreCOGG model achieved high accuracy rates on the human lesion benchmarks, achieving an accuracies ranging from 93.17% to 96.00% across all test images (*Table 3.2*).

3.1.3 PPreCOGG Effectively Classifies Sub-Regions of Different Neoplastic Phenotypes in Human Biopsy Samples

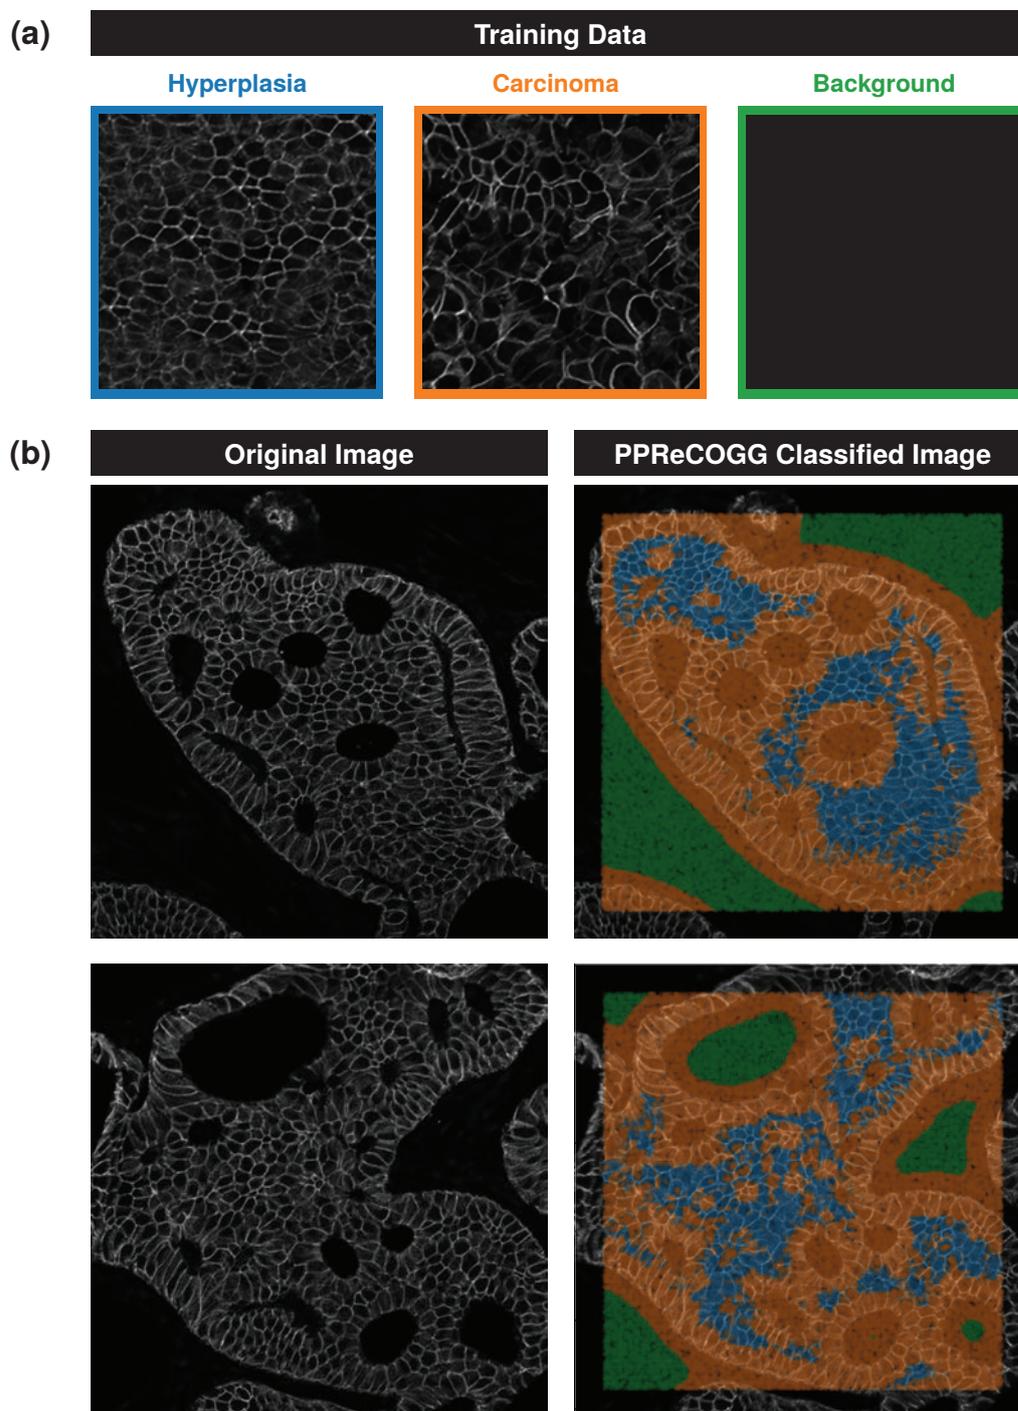
PPreCOGG model trained on E-cadhering patterning found in early breast lesion phenotypes (same features as those used in §3.1.2, see *Figure 3.2a*). This model was subsequently used to classify the pixels in images of early human breast lesions in with E-cadherin had been immunofluorescently labelled (*Figure 3.2b*). PPreCOGG effectively classifies sub-regions within the image that exhibit different cell-patterning characteristics found in early lesions.

Table 3.2: Accuracy of the PPreCOGG Model Trained on Textures Derived from E-cadherin Staining of Human Lesions

		Test Image	Test Image	Test Image
		One	Two	Three
(a)	Original Test Image			
(b)	Ground Truth			
(c)	Classified Image			
(d)	Hyperplasia Accuracy Map			
(e)	Carcinoma Accuracy Map			
(f)	Accuracy	93.17%	93.60%	96.00%

Legend — (a) Test images comprised of composites of textures derived from E-cadherin staining of human lesions exhibiting characteristic hyperplastic or carcinomic cell patterning. (b) The ideal classification (or “ground truth”) of the original test image, where black pixels code for the hyperplastic texture and white codes for the carcinomic texture. (c) A random sampling of pixels classified by the PPreCOGG model. Pixels classified as hyperplasia are coded by blue points, and and pixels classified as carcinoma are coded by orange pixels. (d) and (e) Classified pixels are here compared to and overlaid on their ground truth. Pixels which are correctly classified are coded in green, while false-positives are coded in red. (f) Accuracy of the model, as calculated by the quotient of the number of correctly classified pixels and the total number of classified pixels.

Figure 3.2: Human Lesions as Classified by the PPreCOGG Model



Legend — Sections of human breast biopsies were immunofluorescently stained for E-cadherin and classified using the PPreCOGG model trained on two different early transformed phenotypes, and one background control. (a) Three representative fields of the 512 pixel by 512 pixel images used to train the PPreCOGG model. (b) Fields of breast E-cadherin labelled human breast lesion (left column) were classified by the trained PPreCOGG model (right column). Blue and orange pixels are classified as belonging to a hyperplastic carcinomic regions, respectively. Green pixels are classified as background signal.

3.2 DeepDuct: A Deep-Learning Approach to Regional Breast Cancer Classification using Grad-CAM

3.2.1 Accurate Classification of Breast Neoplasms in the BreakHis Dataset

Transfer-learning the VGG16 on the unmodified BreakHis dataset results in acceptable overall classification accuracy (70%), however closer inspection reveals bias towards one of the classes (ductal carcinoma) due to imbalances in the number of examples between classes (*Figure 3.3a*). Oversampling the dataset such that all classes have an equal number of training examples resulted in a small improvement in overall classification accuracy (72%), and shows a demonstrable reduction in bias (*Figure 3.3b*).

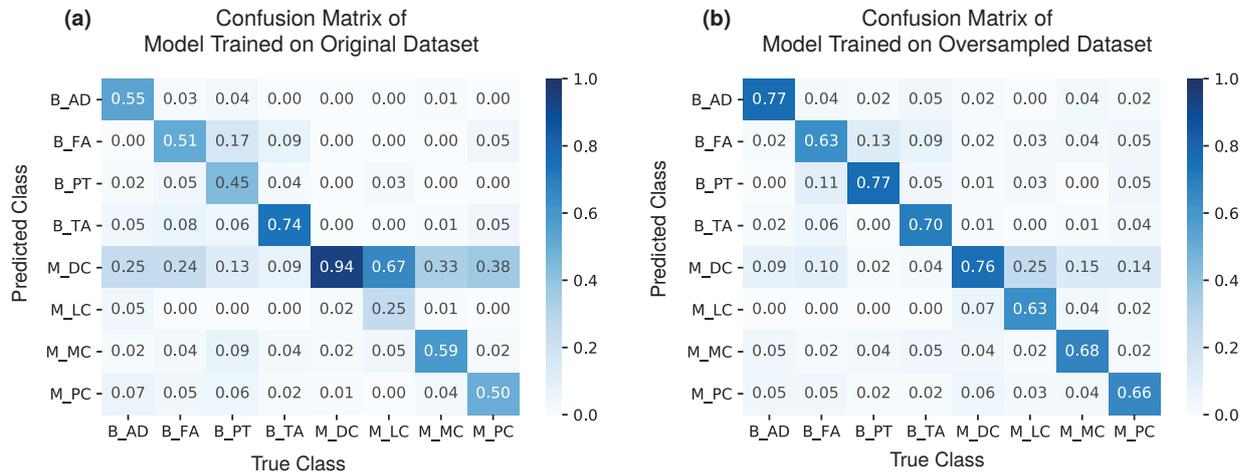
Notably, lobular carcinomas were mistaken for ductal carcinomas in just over two-thirds of the validation set, and correctly identified a quarter of lobular carcinoma examples. Outside the sampling bias, possible explanations for this high false-positive rate include the morphological similarity between the two lesions (lobular carcinomas typically exhibit rounder cells), as well as the low resolution of the dataset, which can degrade visual information of subtle differences. After oversampling the BreakHis dataset, the accuracy and ductal carcinoma false-positive rates have nearly replaced one another, with lobular carcinomas being correctly classified in 63% of cases and false ductal carcinoma classifications in a quarter of cases. Ductal carcinoma false-positives were in fact halved, on average, across nearly all classes after oversampling.

3.2.2 Activation Mapping of BreakHis ConvNet using Grad-CAM

The Grad-CAM algorithm was used to compute class activation maps from the BreakHis fine-tuned VGG16 ConvNet model (*Figure 3.4*). In some cases, regions containing artefacts are highly activated and coincide with high-certainty (*i.e.*: high predicted probability) false-positives (*Figure 3.4e*). It is also not uncommon for low-certainty false-positives to be activated by similar regions across classes (*Figure 3.4f*). Curiously, in many cases

correct predictions are made that are largely activated by the stroma, rather than the epithelium (*Figure 3.4c*).

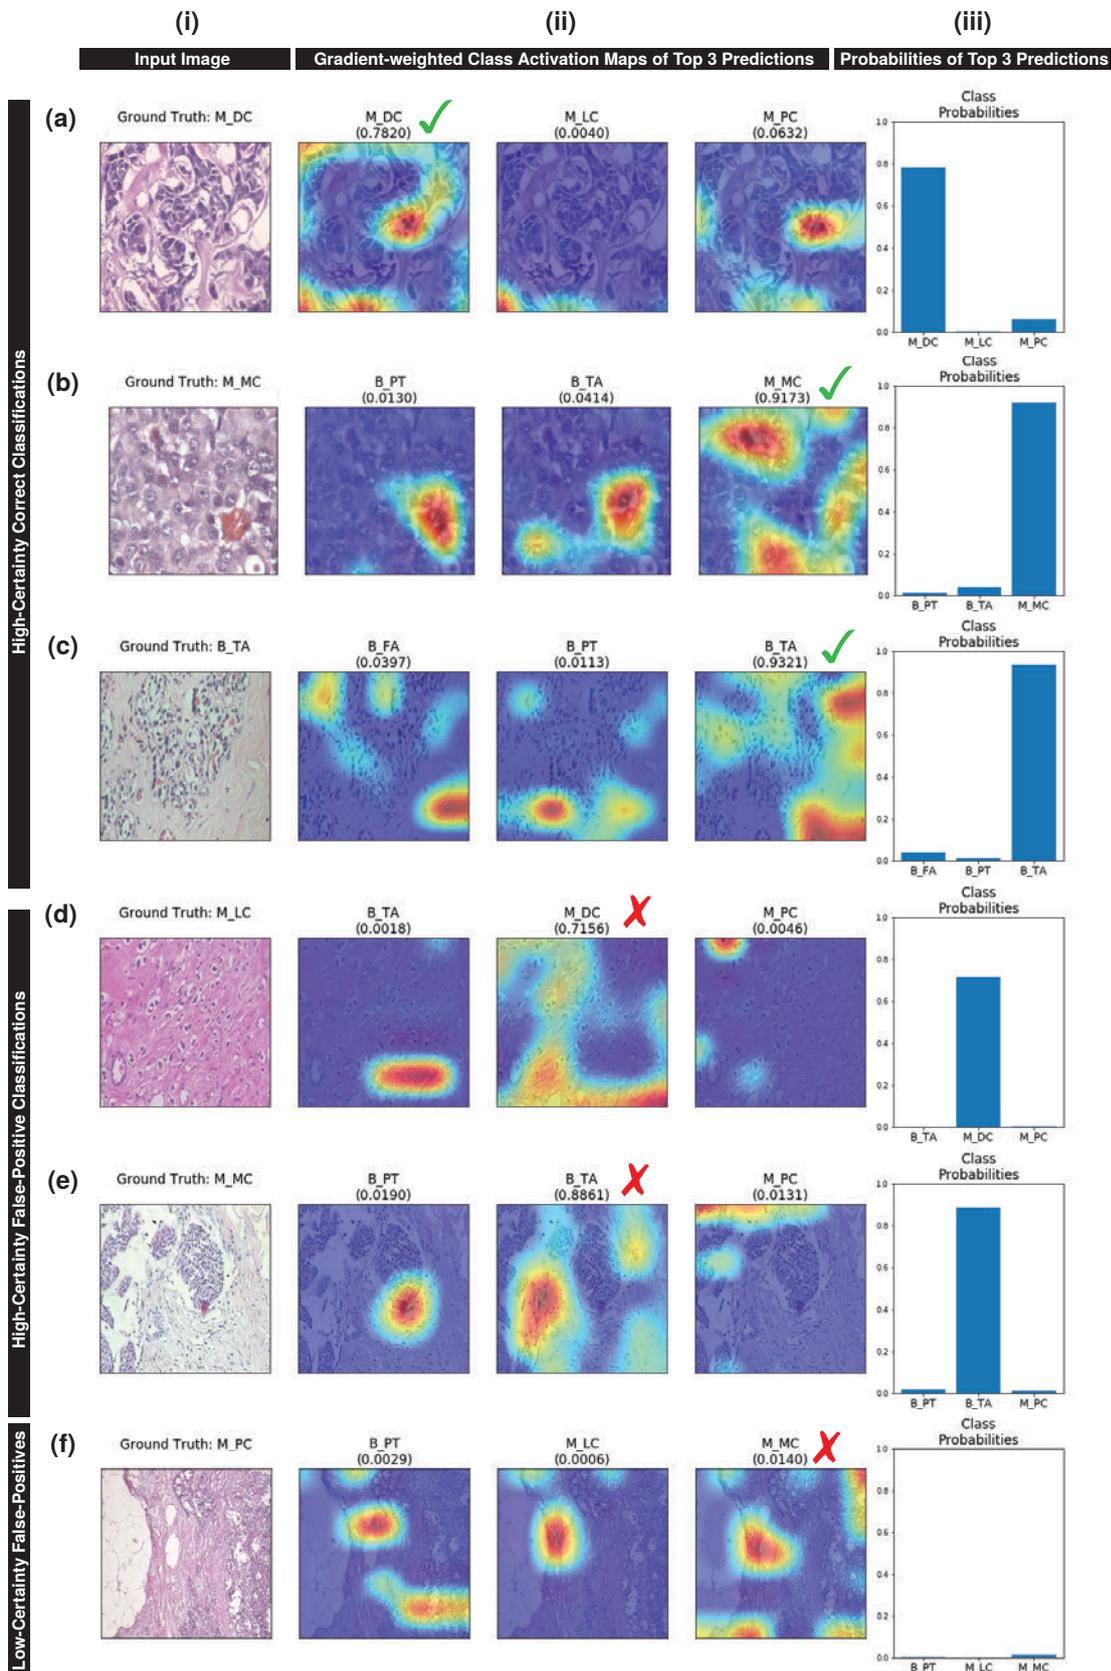
Figure 3.3: Confusion Matrix of VGG16 Model Trained on the Original and Oversampled BreakHis Datasets



Legend — Confusion matrices of the VGG16 model either trained on (a) the original BreakHis dataset as provided by its creators, or (b) a modified version of the BreakHis dataset oversampled such that each class has the same number of training examples. The model trained on the original dataset achieved an average accuracy of 70% on the validation set across all classes, while the oversampled model saw a modest increase for a final average accuracy of 72%. See Appendix A for class code legend used in this figure.

3.2. DEEPDUCT: A DEEP-LEARNING APPROACH TO REGIONAL BREAST CANCER CLASSIFICATION USING GRAD-CAM

Figure 3.4: Activation Maps of the BreakHis Fine-Tuned VGG16 ConvNet



Legend — Selected examples of activation maps of the BreakHis fine-tuned VGG16 model described herein, as calculated by the Grad-CAM algorithm. See Appendix A for a legend of class abbreviations. Column (i) displays the original input image and its ground truth classification. Column (ii) shows the top three predictions of the ConvNet and their activation maps overlaid on the input image; as well as the probability of the classification as reported by the model’s softmax layer (shown in brackets). Column (iii) reports the probabilities of the top three predictions in a bar chart. Top probabilities are indicated with a checkmark (✓) when they match the ground truth (correct classification), and an “x” mark (✗) when they do not (false-positives). Rows (a–c) are examples of correct classifications made by the ConvNet model where the top prediction is of high probability. Rows (d–e) are examples of incorrect classification made by the ConvNet model where the top prediction is of high probability (false-positives). Row (f) is an example of an incorrect classification where none of the classes are reported to be of high-probability.

4. Discussion

The subjective nature of interpreting histopathology, inter-observer disagreement, and reportedly difficult to distinguish early lesions have necessitated reliably replicatable quantitative methods. Previously proposed automated solutions typically ignore tumour heterogeneity and/or are opaque, so-called “black-box” solutions which make them unsuitable for making clinical decisions. The PPreCOGG model, which implements per-pixel texture-based tissue classification, and the DeepDuct model, which leverage a pre-trained ConvNet and class-activation mapping to generate transparent regional classifications, address the aforementioned hurdles faced by previously described systems for automated diagnosis with the aim of demonstrating how clinically relevant models for automated diagnosis can be realised.

4.1 Robust Regional Diagnosis with PPreCOGG

The PPreCOGG model can effectively recognise sub-regions of cell patterning in immunofluorescent confocal imagery, provided a training set that exemplifies said patterns. The PPreCOGG model has been employed the task of identifying, within images of human breast lesions, cell patterns that are characteristic of early breast lesions. PPreCOGG classification in this task is effective and, thanks to its GPU accelerated implementation, performed in a practical time-frame.

4.1.1 Model Performance

The performance of the PPreCOGG model on the Brodatz texture synthetic benchmark attest to the models proficiency at texture recognition tasks between textures that are easily distinguished by human observers upon casual visual inspection (*Table 3.1*). The accuracy of the PPreCOGG model on the Brodatz dataset is comparable to a similar

general-purpose k -NN texture segmentation model based on Gabor features described by Melendez *et al.* (2008), which achieved an average accuracy of $\approx 90\%$ across their synthetic Brodatz benchmarks. Distinguishing the two models is PPRECOGGs GPU acceleration, which allows the model to classify far larger resolution images on inexpensive hardware when compared to the Melendez *et al.* model (512×512 pixels versus 32×32 pixels). PPRECOGG can be easily scaled to far higher resolutions with high-end GPUs. PPRECOGGs true utility, however, is most clearly evinced in the human breast lesion synthetic benchmarks, whereby texture recognition was performed on the test images (Table 3.2a) after being trained on a dataset of human lesions. In these visually challenging tasks, PPRECOGG’s accuracy is equal to or greater than those observed in the visually distinct Brodatz texture benchmarks.

A partial explanation for PPRECOGGs efficiency in both visually distinct and visually challenging texture recognition tasks is provided by the MDS embeddings of the underlying Gabor features of the training set for both the Brodatz and Human Breast Lesion datasets. The MDS embedding of both datasets are remarkably similar, with both classes in each case forming distinct but intersecting planes when scaled to three-dimensional space (Figure 3.1). The distance between the feature-spaces of each class defines the degree to which it is possible for PPRECOGG to distinguish between them. This is largely due to PPRECOGGs reliance on the k -nearest neighbour algorithm for the classification of features.

4.1.2 Future Directions

While the PPRECOGG model readily recognises textural sub-regions within clinical samples, any clinical utility of the PPRECOGG model is dependent on and currently precluded by an immature and incomplete training set. In order for the PPRECOGG model to offer meaningful interpretation and classification of early lesions, a rich dataset is required to capture the many textural manifestations of early lesions. ADH and DCIS lesions are not homogeneously or universally comprised of single textures, and so a sufficiently large and comprehensive dataset is required before the PPRECOGG model can be used to identify

the many faces of early breast lesions.

In addition to a complete training dataset, it is possible to extend the current model to recognise sub-regions according to the identity of neighbouring sub-regions; such that some sub-region identified as belonging to some texture class A would only be reported as belonging to sub-type X if neighbouring regions belong to some texture class B but not C . Rule-sets for these contextual classifications can be learned through random-forest models trained on annotated images, manually according to existing pathology guidelines, or some combination of the two.

The underlying conventional machine-learning algorithm that is the basis of the PPreCOGG model does shape the nature of the conclusions that can be drawn from its output. Namely, the PPreCOGG model is a manifestation of our current understanding of the histopathology of early breast lesions. While this approach results in highly desirable and much needed quantitative and reproducible interpretation of the pathology of biopsy tissue, PPreCOGG as a consequence does not implement feature learning. This is in contrast to models based on neural network algorithms, which forego feature engineering for hidden layers which discover them independently through optimisation. Careful inspection of the hidden layers of the neural network can potentially lead to understanding of early lesion pathology interpretation previously overlooked or otherwise unknown, however such interpretation is nuanced and often provide incomplete “snapshots” of the internal state of the network (Erhan *et al.*, 2010; Zeiler & Fergus, 2013). The DeepDuct model was designed with these concerns in mind.

4.2 DeepDuct as a Model Transparent Regional Classifier

The DeepDuct model described herein provides a proof-of-concept framework for the localisation of breast lesions from H&E staining that does not rely on manual feature selection, transparently reports explanations for its predictions via class activation maps and allows for the potential discovery of new features that could inform future manual diagnosis.

Neural networks have the advantage of dynamically “learning” and optimising features, as opposed to relying on a manual process of feature engineering that is often driven by limited powers of intuition and conventional knowledge. The automated process of feature engineering allows for the potential discovery of new underlying concepts previously not described in the literature that fundamentally define a class from others, and also prevents assumptions and misconceptions from biasing features and the resulting inaccuracies.

4.2.1 Explanations for Neural Network Predictions are Essential for Clinical Use of CADe/x Models

Whole-image ConvNet classification models fine-tuned on the BreakHis dataset have been previously described reporting high-accuracy, as well as patch-based whole-slide classifiers which apply whole-image classification to small patches of an imaged slide resulting, in a form of tumour localisation (Han *et al.*, 2017; Wang *et al.*, 2016). None of these models, however, offer the same extent of transparency and resolution offered by the activation maps provided by the Grad-CAM algorithm used in the DeepDuct model. Existing models of breast lesion classification and localisation remain “black-box” solutions to end-users, particularly those without in-depth knowledge of deep learning algorithms.

Unique among deep-learning based breast lesion classifiers, DeepDuct reports which regions of the input image have lead the model to classify the image as it had. This simultaneously allows for general localisation of classified objects and a glimpse into the internal state of the model, informing and not prescribing a diagnosis. This transparency is essential for a model to see use in contexts such as clinical settings where acting on predictions in blind faith is not an option due to the high-risk associated with the decisions being made. Models that implement “explanations” for their predictions have indeed been shown to increase end-user trust in model predictions, as well as help identify false-positive predictions made by a given model (Ribeiro *et al.*, 2016). To the author’s understanding, DeepDuct is the first application of such explanatory algorithms to the classification and localisation of breast lesions from medical imaging.

4.2.2 Dataset Considerations

The BreakHis dataset, while covering a number of relevant lesion types with a significantly large number of examples for each type, presents some important challenges. Firstly, despite the multiple magnifications provided, images in the BreakHis dataset are not of high-resolution, taken with a digital camera with pixel size of $6.5\mu\text{m}$ and resolution of 480 TV lines. Secondly, the number of examples is extremely imbalanced between classes, with as much of a 7.5-fold difference between the least represented class and the most represented class.

While low-resolution images can be useful for learning so-called “global features”, they’ve proven to be problematic when distinguishing differences between objects with similar high-level features, as is the case between two H&E images exhibiting different lesion subtypes. This problem is illustrated well in the description of Baidu’s Deep Image model, whereby similar objects (such as insects of the same species) can only be distinguished from one another when higher-resolution images are considered in the model (Wu *et al.*, 2015). Training the DeepDuct model on higher resolution datasets would address this concern. High-resolution datasets of breast lesions do exist, but many offer too few examples (INESCTC) or do not offer histological type information outside grade (CAMELYON16).

As described early, imbalances in the number of examples provided per class in the BreakHis dataset had lead to a strong bias towards over-represented classes (*Figure 3.3a*). This bias was addressed by oversampling all under-represented classes by duplicating examples until all classes in the training set contained the same number of examples (*Figure 3.3b*).

4.2.3 Future Directions & Improvements

Implementing the DeepDuct model on smartphones and tablets would afford clinicians low-cost, mobile tools for the annotated classification of breast histology slides through use of commercial or 3D-printable smartphone-microscope adapters (Roy *et al.*, 2014). A less computationally-complex mobile DeepDuct implementation would be required to

account for the limited resources available on the platform. This is typically achieved either by a networked server-client model supported by computation in the cloud, or by replacing the deep, resource-heavy VGG16 model with a more shallow mobile-oriented model such as SqueezeNet (Iandola *et al.*, 2016). While the former is limited by patient-privacy compliance and network connectivity, the latter requires retraining the network on a shallower ConvNet architecture with potential losses in accuracy.

Regional convolutional neural networks (R-CNNs), such as Facebook’s Mask R-CNN, have been developed to provide pixel-resolution object detection in complex scenes (He *et al.*, 2017). Using a Mask R-CNN model trained on a breast lesion dataset can provide higher resolution lesion detection than existing patch-based breast lesion models (Wang *et al.*, 2016).

Bibliography

- Abadi, Martín, Agarwal, Ashish, Barham, Paul, Brevdo, Eugene, Chen, Zhifeng, Citro, Craig, Corrado, Gregory S., Davis, Andy, Dean, Jeffrey, Devin, Matthieu, Ghemawat, Sanjay, Goodfellow, Ian J., Harp, Andrew, Irving, Geoffrey, Isard, Michael, Jia, Yangqing, Józefowicz, Rafal, Kaiser, Lukasz, Kudlur, Manjunath, Levenberg, Josh, Mané, Dan, Monga, Rajat, Moore, Sherry, Murray, Derek Gordon, Olah, Chris, Schuster, Mike, Shlens, Jonathon, Steiner, Benoit, Sutskever, Ilya, Talwar, Kunal, Tucker, Paul A., Vanhoucke, Vincent, Vasudevan, Vijay, Viégas, Fernanda B., Vinyals, Oriol, Warden, Pete, Wattenberg, Martin, Wicke, Martin, Yu, Yuan, & Zheng, Xiaoqiang. 2016. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. *CoRR*, **abs/1603.04467**.
- Acott, A. A., & Mancino, A. T. 2016. Flat epithelial atypia on core needle biopsy, must we surgically excise? *The American Journal of Surgery*, **212**(6), 1211–1213.
- Al-Rfou, Rami, Alain, Guillaume, Almahairi, Amjad, Angermueller, Christof, Bahdanau, Dzmitry, Ballas, Nicolas, Bastien, Frédéric, Bayer, Justin, Belikov, Anatoly, Belopolsky, Alexander, & et al. 2016. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, **abs/1605.02688**(May).
- Ayvaci, Mehmet US, Alagoz, Oguzhan, Chhatwal, Jagpreet, Munoz del Rio, Alejandro, Sickles, Edward A., Nassif, Houssam, Kerlikowske, Karla, & Burnside, Elizabeth S. 2014. Predicting invasive breast cancer versus DCIS in different age groups. *BMC Cancer*, **14**(Aug), 584.
- Badve, Sunil, A’hern, Roger P, Ward, Ann M, Millis, Rosemary R, Pinder, Sarah E, Ellis, Ian O, Gusterson, Barry A, & Sloane, John P. 1998. Prediction of local recurrence of ductal carcinoma in situ of the breast using five histological classifications: A comparative study with long follow-up. *Human Pathology*, **29**(9), 915–923.
- Betsill, William L., Rosen, Paul Peter, Lieberman, Philip H., & Robbins, Guy F. 1978. Intraductal Carcinoma: Long-Term Follow-up After Treatment by Biopsy Alone. *JAMA*, **239**(18), 1863–1867.

- Bleicher, Richard J. 2013. Ductal carcinoma in situ. *Surgical Clinics of North America*, **93**(2), 393–410.
- Boecker, Werner, Moll, Roland, Dervan, Peter, Buerger, Horst, Poremba, Christopher, Diallo, Raihanatou Ina, Herbst, Hermann, Schmidt, Ansgar, Lerch, Markus M, & Buchwalow, Igor B. 2002. Usual ductal hyperplasia of the breast is a committed stem (progenitor) cell lesion distinct from atypical ductal hyperplasia and ductal carcinoma in situ. *The Journal of Pathology*, **198**(4), 458–467.
- Bombonati, Alessandro, & Sgroi, Dennis C. 2011. The molecular pathology of breast cancer progression. *The Journal of Pathology*, **223**(2), 308–318.
- Bradski, G. 2000. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*.
- Brodatz, Phil. 1966. *Textures a photographic album for artists and designers*. Dover Publ.
- Burke, Harry B., Goodman, Philip H., Rosen, David B., Henson, Donald E., Weinstein, John N., Harrell, Frank E., Marks, Jeffrey R., Winchester, David P., & Bostwick, David G. 1997. Artificial neural networks improve the accuracy of cancer survival prediction. *Cancer*, **79**(4), 857–862.
- Canadian Cancer Society's Advisory Committee on Cancer Statistics. 2015. *Canadian Cancer Statistics 2015*.
- Cheng, Jie-Zhi, Ni, Dong, Chou, Yi-Hong, Qin, Jing, Tiu, Chui-Mei, Chang, Yeun-Chung, Huang, Chiun-Sheng, Shen, Dinggang, & Chen, Chung-Ming. 2016. Computer-Aided Diagnosis with Deep Learning Architecture: Applications to Breast Lesions in US Images and Pulmonary Nodules in CT Scans. *Scientific Reports*, **6**(Apr), srep24454.
- Chollet, François, *et al.* 2015. *Keras*. <https://github.com/fchollet/keras>.
- Clark, S E, Warwick, J, Carpenter, R, Bowen, R L, Duffy, S W, & Jones, J L. 2011. Molecular subtyping of DCIS: heterogeneity of breast cancer reflected in pre-invasive disease. *British Journal of Cancer*, **104**(1), 120–127.
- Cover, T., & Hart, P. 1967. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, **13**(1), 21–27.
- Crick, Francis. 1989. The recent excitement about neural networks. *Nature*, **337**(6203), 129–132.
- Decaestecker, Christine, Salmon, Isabelle, Dewitte, Olivier, Camby, Isabelle, Van Ham, Philippe, Pasteels, Jean-Lambert, Brotchi, Jacques, & Kiss, Robert. 1997. Nearest-neighbor classification for identification of aggressive versus nonaggressive low-grade

- astrocytic tumors by means of image cytometry-generated variables. *Journal of Neurosurgery*, **86**(3), 532–537.
- Deng, J., Dong, W., Socher, R., Li, L. J., Li, Kai, & Fei-Fei, Li. 2009 (June). ImageNet: A large-scale hierarchical image database. *Pages 248–255 of: 2009 IEEE Conference on Computer Vision and Pattern Recognition*.
- Dhahbi, Sami, Barhoumi, Walid, & Zagrouba, Ezzeddine. 2015. Breast cancer diagnosis in digitized mammograms using curvelet moments. *Computers in Biology and Medicine*, **64**(Supplement C), 79–90.
- Dhawan, A. P., Chitre, Y., & Kaiser-Bonasso, C. 1996. Analysis of mammographic microcalcifications using gray-level image structure features. *IEEE Transactions on Medical Imaging*, **15**(3), 246–259.
- Dion, L., Racin, A., Brousse, S., Beltjens, F., Cauchois, A., Levêque, J., Coutant, C., & Lavoue, V. 2016. Atypical epithelial hyperplasia of the breast: state of the art. *Expert Review of Anticancer Therapy*, **16**(9), 943–953.
- Dreiseitl, Stephan, Ohno-Machado, Lucila, Kittler, Harald, Vinterbo, Staal, Billhardt, Holger, & Binder, Michael. 2001. A Comparison of Machine Learning Methods for the Diagnosis of Pigmented Skin Lesions. *Journal of Biomedical Informatics*, **34**(1), 28–36.
- Engstrøm, M. J., Opdahl, S., Hagen, A. I., Romundstad, P. R., Akslen, L. A., Haugen, O. A., Vatten, L. J., & Bofin, A. M. 2013. Molecular subtypes, histopathological grade and survival in a historic cohort of breast cancer patients. *Breast Cancer Research and Treatment*, **140**(3), 463–473.
- Erhan, Dumitru, Courville, Aaron, & Bengio, Yoshua. 2010. Understanding representations learned in deep architectures.
- Esteva, Andre, Kuprel, Brett, Novoa, Roberto A., Ko, Justin, Swetter, Susan M., Blau, Helen M., & Thrun, Sebastian. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, **542**(7639), nature21056.
- Fix, Evelyn, & Hodges, Jr. 1951. *Discriminatory Analysis - Nonparametric Discrimination: Consistency Properties*.
- Floyd, Carey E., Lo, Joseph Y., Yun, A. Joon, Sullivan, Daniel C., & Kornguth, Phyllis J. 1994. Prediction of breast cancer malignancy using an artificial neural network. *Cancer*, **74**(11), 2944–2948.
- Ghosal, Sambuddha, Blystone, David, Singh, Asheesh K., Ganapathysubramanian, Baskar, Singh, Arti, & Sarkar, Soumik. 2017. Interpretable Deep Learning applied to Plant Stress Phenotyping. *arXiv:1710.08619 [cs, stat]*, Oct. arXiv: 1710.08619.

- Gomes, Douglas S., Porto, Simone S., Balabram, Débora, & Gobbi, Helenice. 2014. Inter-observer variability between general pathologists and a specialist in breast pathology in the diagnosis of lobular neoplasia, columnar cell lesions, atypical ductal hyperplasia and ductal carcinoma in situ of the breast. *Diagnostic Pathology*, **9**(1), 121.
- Halaoui, Ruba, Rejon, Carlis, Chatterjee, Sudipa June, Szymborski, Joseph, Meterissian, Sarkis, Muller, William J., Omeroglu, Atilla, & McCaffrey, Luke. 2017. Progressive polarity loss and luminal collapse disrupt tissue organization in carcinoma. *Genes & Development*, **31**(15), 1573–1587.
- Hamperl, H. 1970. *The Myothenia (Myoepithelial Cells)*. Springer, Berlin, Heidelberg. Page 161–220.
- Han, Zhongyi, Wei, Benzhen, Zheng, Yuanjie, Yin, Yilong, Li, Kejian, & Li, Shuo. 2017. Breast Cancer Multi-classification from Histopathological Images with Structured Deep Learning Model. *Scientific Reports*, **7**(1), 4172.
- Hartmann, Lynn C., Degnim, Amy C., Santen, Richard J., Dupont, William D., & Ghosh, Karthik. 2015. Atypical Hyperplasia of the Breast — Risk Assessment and Management Options. *The New England Journal of Medicine*, **372**(1), 78–89.
- He, Kaiming, Gkioxari, Georgia, Dollár, Piotr, & Girshick, Ross. 2017. Mask R-CNN. *arXiv:1703.06870 [cs]*, Mar. arXiv: 1703.06870.
- Hendler, James. 2008. Avoiding another AI winter.
- Hubel, D. H., & Wiesel, T. N. 1968. Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, **195**(1), 215–243.
- Hunt, Earl B., Marin, Janet, & Stone, Philip J. 1966. *Experiments in Induction*. Academic Press.
- Hunter, J. D. 2007. Matplotlib: A 2D graphics environment. *Computing In Science & Engineering*, **9**(3), 90–95.
- Iandola, Forrest N., Han, Song, Moskewicz, Matthew W., Ashraf, Khalid, Dally, William J., & Keutzer, Kurt. 2016. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5MB model size. *arXiv:1602.07360 [cs]*, Feb. arXiv: 1602.07360.
- Kerlikowske, Karla. 2010. Epidemiology of Ductal Carcinoma In Situ. *Journal of the National Cancer Institute. Monographs*, **2010**(41), 139–141.
- Kleene, S. C. 1956. *Representation of Events in Nerve Nets and Finite Automata*. Vol. 34. Princeton University Press.

- Kononenko, Igor. 2001. Machine learning for medical diagnosis: history, state of the art and perspective. *Artificial Intelligence in Medicine*, **23**(1), 89–109.
- Korbar, Bruno, Olofson, Andrea M, Mirafior, Allen P, Nicka, Catherine M, Suriawinata, Matthew A, Torresani, Lorenzo, Suriawinata, Arief A, & Hassanpour, Saeed. 2017. Looking Under the Hood: Deep Neural Network Visualization to Interpret Whole-Slide Image Analysis Outcomes for Colorectal Polyps. *Pages 821–827 of: Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. IEEE.
- Kotikalapudi, Raghavendra, & contributors. 2017. *keras-vis*. <https://github.com/raghakot/keras-vis>.
- Laine, A., & Fan, J. 1993. Texture classification by wavelet packet signatures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **15**(11), 1186–1191.
- LeCun, Y., Kavukcuoglu, K., & Farabet, C. 2010. *Convolutional networks and applications in vision*. Page 253–256.
- Lerwill, MF. 2008. Flat epithelial atypia of the breast. *Archives of Pathology & Laboratory Medicine*, **132**(4), 615–621.
- Luo, Zhongwen, Liu, Hongzhi, & Wu, Xincan. 2005. *Artificial neural network computation on graphic process unit*. Vol. 1. Page 622–626 vol. 1.
- Madabhushi, Anant. 2009. Digital pathology image analysis: opportunities and challenges.
- Mangasarian, Olvi L, Setiono, R, & Wolberg, WH. 1990. Pattern recognition via linear programming: Theory and application to medical diagnosis. *Large-scale numerical optimization*, 22–31.
- Manjunath, B. S., & Ma, W. Y. 1996. Texture features for browsing and retrieval of image data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **18**(8), 837–842.
- Marçelja, S. 1980. Mathematical description of the responses of simple cortical cells*. *JOSA*, **70**(11), 1297–1300.
- Masood, Shahla, & Rosa, Marilyn. 2011. Borderline Breast Lesions: Diagnostic Challenges and Clinical Implications. *Advances in Anatomic Pathology*, **18**(3), 190–198.
- Melendez, Jaime, Garcia, Miguel Angel, & Puig, Domenec. 2008. Efficient distance-based per-pixel texture classification with Gabor wavelet filters. *Pattern Analysis and Applications*, **11**(3–4), 365–372.

- Meyer-Baese, Anke. 2004. *Pattern recognition for medical imaging*. Academic Press.
- Minsky, Marvin, & Papert, Seymour. 1969. *Perceptrons*. Perceptrons. M.I.T. Press.
- Mommers, Ellen C.M., Page, David L., Dupont, William D., Schuyler, Peggy, Leonhart, Angelique M., Baak, Jan P.A., Meijer, Chris J.L.M., & Van Diest, Paul J. 2001. Prognostic value of morphometry in patients with normal breast tissue or usual ductal hyperplasia of the breast. *International Journal of Cancer*, **95**(5), 282–285.
- Narod, Steven A., Iqbal, Javaid, Giannakeas, Vasily, Sopik, Victoria, & Sun, Ping. 2015. Breast Cancer Mortality After a Diagnosis of Ductal Carcinoma In Situ. *JAMA Oncology*, **1**(7), 888–896.
- National Comprehensive Cancer Network. 2017 (June). *Breast Cancer Screening and Diagnosis*. Tech. rept. 1.2017.
- National Institutes of Health. 2010. Breast Cancer Basics and You: Introduction. *NIH MedlinePlus*, **5**(2), 17.
- Nothacker, Monika, Duda, Volker, Hahn, Markus, Warm, Mathias, Degenhardt, Friedrich, Madjar, Helmut, Weinbrenner, Susanne, & Albert, Ute-Susann. 2009. Early detection of breast cancer: benefits and risks of supplemental breast ultrasound in asymptomatic women with mammographically dense breast tissue. A systematic review. *BMC Cancer*, **9**(Sep), 335.
- O’Connell, Peter, Pekkel, Vladimir, Fuqua, Suzanne, Osborne, C. Kent, & Allred, D. Craig. 1994. Molecular genetic studies of early breast cancer evolution. *Breast Cancer Research and Treatment*, **32**(1), 5–12.
- Otaïr, Dr Mohammed. 2013. Approximate k-nearest neighbour based spatial clustering using k-d tree. *arXiv:1303.1951 [cs]*, Mar. arXiv: 1303.1951.
- Page, David L, Dupont, William D, Rogers, LW, & Rados, MS. 1959. Atypical hyperplastic lesions of the female breast. *A long-term follow-up study*. *Cancer*.
- Page, David L., Dupont, William D., Rogers, Lowell W., & Landenberger, Marie. 1982. Intraductal carcinoma of the breast: Follow-up after biopsy only. *Cancer*, **49**(4), 751–758.
- Pan, S. J., & Yang, Q. 2010. A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, **22**(10), 1345–1359.
- Pareja, Fresia, Marchiò, Caterina, Geyer, Felipe C., Weigelt, Britta, & Reis-Filho, Jorge S. 2017. Breast Cancer Heterogeneity: Roles in Tumorigenesis and Therapeutic Implications. *Current Breast Cancer Reports*, **9**(1), 34–44.

- Pedregosa, Fabian, Varoquaux, Gaël, Gramfort, Alexandre, Michel, Vincent, Thirion, Bertrand, Grisel, Olivier, Blondel, Mathieu, Prettenhofer, Peter, Weiss, Ron, Dubourg, Vincent, & et al. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, **12**(Oct), 2825–2830.
- Perou, Charles M., Sørlie, Therese, Eisen, Michael B., van de Rijn, Matt, Jeffrey, Stefanie S., Rees, Christian A., Pollack, Jonathan R., Ross, Douglas T., Johnsen, Hilde, Akslen, Lars A., & et al. 2000. Molecular portraits of human breast tumours. *Nature*, **406**(6797), 747–752.
- Pinder, Sarah E. 2017. *Proliferative Breast Lesions*. Springer International Publishing. Page 33–41.
- Plotly Technologies Inc. 2015. *Collaborative data science*.
- Prat, Aleix, Pineda, Estela, Adamo, Barbara, Galván, Patricia, Fernández, Aranzazu, Gaba, Lydia, Díez, Marc, Viladot, Margarita, Arance, Ana, & Muñoz, Montserrat. 2015. Clinical implications of the intrinsic molecular subtypes of breast cancer. *The Breast*, **24**(Nov), S26–S35.
- Putti, T C, Pinder, S E, Elston, C W, Lee, A H S, & Ellis, I O. 2005. Breast pathology practice: most common problems in a consultation service. *Histopathology*, **47**(5), 445–457.
- Rangayyan, Rangaraj M., Ayres, Fábio J., & Leo Desautels, J. E. 2007. A review of computer-aided diagnosis of breast cancer: Toward the detection of subtle signs. *Journal of the Franklin Institute*, **344**(3), 312–348.
- Ribeiro, Marco Tulio, Singh, Sameer, & Guestrin, Carlos. 2016. “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. *arXiv:1602.04938 [cs, stat]*, Feb. arXiv: 1602.04938.
- Rosenblatt, Frank. 1957. The perceptron a perceiving and recognizing automaton. *tech. rep., Technical Report 85-460-1*.
- Rouzier, Roman, Perou, Charles M., Symmans, W. Fraser, Ibrahim, Nuhad, Cristofanilli, Massimo, Anderson, Keith, Hess, Kenneth R., Stec, James, Ayers, Mark, Wagner, Peter, & et al. 2005. Breast Cancer Molecular Subtypes Respond Differently to Preoperative Chemotherapy. *Clinical Cancer Research*, **11**(16), 5678–5685.
- Roy, Somak, Pantanowitz, Liron, Amin, Milon, Seethala, Raja R., Ishtiaque, Ahmed, Yousem, Samuel A., Parwani, Anil V., Cucoranu, Ioan, & Hartman, Douglas J. 2014. Smartphone adapters for digital photomicrography. *Journal of Pathology Informatics*, **5**(Jul).

- Russakovsky, Olga, Deng, Jia, Su, Hao, Krause, Jonathan, Satheesh, Sanjeev, Ma, Sean, Huang, Zhiheng, Karpathy, Andrej, Khosla, Aditya, Bernstein, Michael, & et al. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, **115**(3), 211–252.
- Samuel, A. L. 1959. Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development*, **3**(3), 210–229.
- Selvaraju, Ramprasaath R., Das, Abhishek, Vedantam, Ramakrishna, Cogswell, Michael, Parikh, Devi, & Batra, Dhruv. 2016. Grad-CAM: Why did you say that? *arXiv:1611.07450 [cs, stat]*, Nov. arXiv: 1611.07450.
- Shin, H. C., Roth, H. R., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D., & Summers, R. M. 2016. Deep Convolutional Neural Networks for Computer-Aided Detection: CNN Architectures, Dataset Characteristics and Transfer Learning. *IEEE Transactions on Medical Imaging*, **35**(5), 1285–1298.
- Simard, Patrice Y., Steinkrau, Dave, & Buck, Ian. 2005. *Using GPUs for Machine Learning Algorithms*. IEEE Computer Society. Page 1115–1119.
- Simonyan, K., & Zisserman, A. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR*, **abs/1409.1556**.
- Simpson, Jean F. 2009. Update on atypical epithelial hyperplasia and ductal carcinoma in situ. *Pathology*, **41**(1), 36–39.
- Spanhol, F. A., Oliveira, L. S., Petitjean, C., & Heutte, L. 2016. A Dataset for Breast Cancer Histopathological Image Classification. *IEEE Transactions on Biomedical Engineering*, **63**(7), 1455–1462.
- Sunil R. Lakhani, International Agency for Research on Cancer, Weltgesundheitsorganisation (ed). 2012. *WHO classification of tumours of the breast: views of a working group that convened for a consensus and editorial meeting at the International Agency for Research on Cancer (IARC), Lyon, September 1 - 3, 2011*. 4. ed edn. World Health Organization Classification of tumours. Internat. Agency for Research on Cancer.
- Tamimi, Rulla M, Baer, Heather J, Marotti, Jonathan, Galan, Mark, Galaburda, Laurie, Fu, Yineng, Deitz, Anne C, Connolly, James L, Schnitt, Stuart J, Colditz, Graham A, & et al. 2008. Comparison of molecular phenotypes of ductal carcinoma in situ and invasive breast cancer. *Breast Cancer Research*, **10**(4), R67.
- Unser, M. 1995. Texture classification and segmentation using wavelet frames. *IEEE Transactions on Image Processing*, **4**(11), 1549–1560.

- Wang, Dayong, Khosla, Aditya, Gargeya, Rishab, Irshad, Humayun, & Beck, Andrew H. 2016. Deep Learning for Identifying Metastatic Breast Cancer. *arXiv:1606.05718 [cs, q-bio]*, Jun. arXiv: 1606.05718.
- Waskom, Michael, Botvinnik, Olga, Hobson, Paul, Cole, John B., Halchenko, Yaroslav, Hoyer, Stephan, Miles, Alistair, Augspurger, Tom, Yarkoni, Tal, Megies, Tobias, & et al. 2014. seaborn: v0.5.0 (November 2014). Nov.
- Weber, Allan G. 1997. The USC-SIPI image database version 5. *USC-SIPI Report*, **315**, 1–24.
- Wei, Benzhen, Han, Zhongyi, He, Xueying, & Yin, Yilong. 2017. *Deep learning model based breast cancer histopathological image classification*. IEEE. Page 348–353.
- Weigelt, Britta, Geyer, Felipe C., & Reis-Filho, Jorge S. 2010. Histological types of breast cancer: How special are they? *Molecular Oncology*, **4**(3), 192–208.
- Werbos, Paul J. 1982. *Applications of advances in nonlinear sensitivity analysis*. Springer. Page 762–770.
- Wernick, M. N., Yang, Y., Brankov, J. G., Yourganov, G., & Strother, S. C. 2010. Machine Learning in Medical Imaging. *IEEE Signal Processing Magazine*, **27**(4), 25–38.
- Wu, Ren, Yan, Shengen, Shan, Yi, Dang, Qingqing, & Sun, Gang. 2015. Deep Image: Scaling up Image Recognition. *arXiv:1501.02876 [cs]*, Jan. arXiv: 1501.02876.
- Yosinski, Jason, Clune, Jeff, Bengio, Yoshua, & Lipson, Hod. 2014. *How transferable are features in deep neural networks?* Curran Associates, Inc. Page 3320–3328.
- Zeiler, Matthew D., & Fergus, Rob. 2013. Visualizing and Understanding Convolutional Networks. *CoRR*, **abs/1311.2901**.

Appendix A: BreakHis Class Codes and Abbreviations

Abbreviations for the class names present in the BreakHis dataset are used throughout this manuscript, and within the model itself. These class names refer to the breast lesion types outlined in the WHO guidelines for the classification of breast tumours (Sunil R. Lakhani, 2012).

Benign Tumours

B_AD Benign Adenosis

B_FA Benign Fibroadenoma

B_TA Benign Tubular Adenoma

B_PT Benign Phylodes Tumour

Malignant Tumours

M_DC Malignant Ductal Carcinoma

M_LC Malignant Lobular Carcinoma

M_MC Malignant Mucinous Carcinoma

M_PC Malignant Papillary Carcinoma

Appendix B: Licensing Information

The text of this document, as well as *Figures 1.2, 2.1, 2.2, 2.4, 3.1, 3.2, 3.3* and *Table 3.2* are released under the Creative Commons Attribution 4.0 International License. Attribution is to be made to Joseph Szymborski and/or this thesis.

Figure 1.1 is released by the National Institutes of Health into, and here used under the auspices of, the Public Domain.

Figure 1.3 is copyrighted material that is owned by or exclusively licensed to John Wiley & Sons Inc. or one of its group companies (each a “Wiley Company”) or handled on behalf of a society with which a Wiley Company has exclusive publishing rights in relation to a particular work (collectively “WILEY”). It is here used with permission under license number 4241600451088, received on 3rd December, 2017.

Images used for the creation of some test images whose thumbnails appear in *Table 3.1* belong to Dover Publishing, Inc.. Permission to use these images are given in the original publication (Brodatz, 1966) and is reproduced below:

This book belongs to the Dover Pictorial Archive Series. You may use the designs and illustrations for graphics and crafts applications, free and without special permission, provided that you include no more than three in the same publication or project. (For permission for additional use, please write to Permissions Department, Dover Publications, Inc. 31 East 2nd Street, Mineola, N.Y. 11501)

However, republication or reproduction of any illustrations by any other graphic service, whether it be in a book or in any other design resource, is strictly prohibited.

Thumbnails in *Figure 3.4* and *Figure 2.3* contain images which belong to the BreakHis dataset. Images from the BreakHis dataset were obtained from the website of the Laboratory of Vision, Robotics and Imaging at the Federal University of Parana (<https://web.inf.ufpr.br/vri/databases/breast-cancer-histopathological-database-breakhis/>) and are licensed under the Creative Commons Attribution 4.0 International License.



Creative Commons Legal Code

Attribution 4.0 International

Official translations of this license are available [in other languages](#).



Creative Commons Corporation (“Creative Commons”) is not a law firm and does not provide legal services or legal advice. Distribution of Creative Commons public licenses does not create a lawyer-client or other relationship. Creative Commons makes its licenses and related information available on an “as-is” basis. Creative Commons gives no warranties regarding its licenses, any material licensed under their terms and conditions, or any related information. Creative Commons disclaims all liability for damages resulting from their use to the fullest extent possible.

Using Creative Commons Public Licenses

Creative Commons public licenses provide a standard set of terms and conditions that creators and other rights holders may use to share original works of authorship and other material subject to copyright and certain other rights specified in the public license below. The following considerations are for informational purposes only, are not exhaustive, and do not form part of our licenses.

Considerations for licensors: *Our public licenses are intended for use by those authorized to give the public permission to use material in ways otherwise restricted by copyright and certain other rights. Our licenses are irrevocable. Licensors should read and understand the terms and conditions of the license they choose before applying it. Licensors should also secure all rights necessary before applying our licenses so that the public can reuse the material as expected. Licensors should clearly mark any material not subject to the license. This includes other CC-licensed material, or material used under an exception or limitation to copyright. [More considerations for licensors](#).*

Considerations for the public: *By using one of our public licenses, a licensor grants the public permission to use the licensed material under specified terms and conditions. If the licensor’s permission is not necessary for any reason—for example, because of any applicable exception or limitation to copyright—then that use is not regulated by the license. Our licenses grant only permissions under copyright and certain other rights that a licensor has authority to grant. Use of the licensed material may still be restricted for other reasons, including because others have copyright or other rights in the material. A licensor may make special requests, such as asking that all changes be marked or described. Although not required by our licenses, you are encouraged to respect those requests where reasonable. [More considerations for the public](#).*

Creative Commons Attribution 4.0 International Public License

By exercising the Licensed Rights (defined below), You accept and agree to be bound by the terms and conditions of this Creative Commons Attribution 4.0 International Public License (“Public License”). To the extent this Public License may be interpreted as a contract, You are granted the Licensed Rights in consideration of Your acceptance of these terms and conditions, and the Licensor grants You such rights in consideration of benefits the Licensor receives from making the Licensed Material available under these terms and conditions.

Section 1 – Definitions.

- a. **Adapted Material** means material subject to Copyright and Similar Rights that is derived from or based upon the Licensed Material and in which the Licensed Material is translated, altered, arranged, transformed, or otherwise modified in a manner requiring permission under the Copyright and Similar Rights held by the Licensor. For purposes of this Public License, where the Licensed Material is a musical work, performance, or sound recording, Adapted Material is always produced where the Licensed Material is synched in timed relation with a moving image.
- b. **Adapter's License** means the license You apply to Your Copyright and Similar Rights in Your contributions to Adapted Material in accordance with the terms and conditions of this Public License.
- c. **Copyright and Similar Rights** means copyright and/or similar rights closely related to copyright including, without limitation, performance, broadcast, sound recording, and Sui Generis Database Rights, without regard to how the rights are labeled or categorized. For purposes of this Public License, the rights specified in Section [2\(b\)\(1\)-\(2\)](#) are not Copyright and Similar Rights.
- d. **Effective Technological Measures** means those measures that, in the absence of proper authority, may not be circumvented under laws fulfilling obligations under Article 11 of the WIPO Copyright Treaty adopted on December 20, 1996, and/or similar international agreements.
- e. **Exceptions and Limitations** means fair use, fair dealing, and/or any other exception or limitation to Copyright and Similar Rights that applies to Your use of the Licensed Material.
- f. **Licensed Material** means the artistic or literary work, database, or other material to which the Licensor applied this Public License.
- g. **Licensed Rights** means the rights granted to You subject to the terms and conditions of this Public License, which are limited to all Copyright and Similar Rights that apply to Your use of the Licensed Material and that the Licensor has authority to license.
- h. **Licensor** means the individual(s) or entity(ies) granting rights under this Public License.
- i. **Share** means to provide material to the public by any means or process that requires permission under the Licensed Rights, such as reproduction, public display, public performance, distribution, dissemination, communication, or importation, and to make material available to the public including in ways that members of the public may access the material from a place and at a time individually chosen by them.
- j. **Sui Generis Database Rights** means rights other than copyright resulting from Directive 96/9/EC of the European Parliament and of the Council of 11 March 1996 on the legal protection of databases, as amended and/or succeeded, as well as other essentially equivalent rights anywhere in the world.
- k. **You** means the individual or entity exercising the Licensed Rights under this Public License. **Your** has a corresponding meaning.

Section 2 – Scope.

- a. **License grant.**
 1. Subject to the terms and conditions of this Public License, the Licensor hereby grants You a worldwide, royalty-free, non-sublicensable, non-exclusive, irrevocable license to exercise the Licensed Rights in the Licensed Material to:
 - A. reproduce and Share the Licensed Material, in whole or in part; and
 - B. produce, reproduce, and Share Adapted Material.
 2. Exceptions and Limitations. For the avoidance of doubt, where Exceptions and Limitations apply to Your use, this Public License does not apply, and You do not need to comply with its terms and conditions.
 3. Term. The term of this Public License is specified in Section [6\(a\)](#).
 4. Media and formats; technical modifications allowed. The Licensor authorizes You to exercise the Licensed Rights in all media and formats whether now known or hereafter created, and to make technical modifications necessary to do so. The Licensor waives and/or agrees not to assert any right or authority to forbid You from making technical modifications necessary to exercise the Licensed Rights, including technical modifications necessary to circumvent Effective Technological Measures. For purposes of this Public License, simply making modifications authorized by this Section [2\(a\)\(4\)](#) never produces Adapted Material.
 5. Downstream recipients.
 - A. Offer from the Licensor – Licensed Material. Every recipient of the Licensed Material

automatically receives an offer from the Licensor to exercise the Licensed Rights under the terms and conditions of this Public License.

- B. **No downstream restrictions.** You may not offer or impose any additional or different terms or conditions on, or apply any Effective Technological Measures to, the Licensed Material if doing so restricts exercise of the Licensed Rights by any recipient of the Licensed Material.
6. **No endorsement.** Nothing in this Public License constitutes or may be construed as permission to assert or imply that You are, or that Your use of the Licensed Material is, connected with, or sponsored, endorsed, or granted official status by, the Licensor or others designated to receive attribution as provided in Section [3\(a\)\(1\)\(A\)\(i\)](#).

b. Other rights.

1. Moral rights, such as the right of integrity, are not licensed under this Public License, nor are publicity, privacy, and/or other similar personality rights; however, to the extent possible, the Licensor waives and/or agrees not to assert any such rights held by the Licensor to the limited extent necessary to allow You to exercise the Licensed Rights, but not otherwise.
2. Patent and trademark rights are not licensed under this Public License.
3. To the extent possible, the Licensor waives any right to collect royalties from You for the exercise of the Licensed Rights, whether directly or through a collecting society under any voluntary or waivable statutory or compulsory licensing scheme. In all other cases the Licensor expressly reserves any right to collect such royalties.

Section 3 – License Conditions.

Your exercise of the Licensed Rights is expressly made subject to the following conditions.

a. Attribution.

1. If You Share the Licensed Material (including in modified form), You must:
 - A. retain the following if it is supplied by the Licensor with the Licensed Material:
 - i. identification of the creator(s) of the Licensed Material and any others designated to receive attribution, in any reasonable manner requested by the Licensor (including by pseudonym if designated);
 - ii. a copyright notice;
 - iii. a notice that refers to this Public License;
 - iv. a notice that refers to the disclaimer of warranties;
 - v. a URI or hyperlink to the Licensed Material to the extent reasonably practicable;
 - B. indicate if You modified the Licensed Material and retain an indication of any previous modifications; and
 - C. indicate the Licensed Material is licensed under this Public License, and include the text of, or the URI or hyperlink to, this Public License.
2. You may satisfy the conditions in Section [3\(a\)\(1\)](#) in any reasonable manner based on the medium, means, and context in which You Share the Licensed Material. For example, it may be reasonable to satisfy the conditions by providing a URI or hyperlink to a resource that includes the required information.
3. If requested by the Licensor, You must remove any of the information required by Section [3\(a\)\(1\)\(A\)](#) to the extent reasonably practicable.
4. If You Share Adapted Material You produce, the Adapter's License You apply must not prevent recipients of the Adapted Material from complying with this Public License.

Section 4 – Sui Generis Database Rights.

Where the Licensed Rights include Sui Generis Database Rights that apply to Your use of the Licensed Material:

- a. for the avoidance of doubt, Section [2\(a\)\(1\)](#) grants You the right to extract, reuse, reproduce, and Share all or a substantial portion of the contents of the database;
- b. if You include all or a substantial portion of the database contents in a database in which You have

Sui Generis Database Rights, then the database in which You have Sui Generis Database Rights (but not its individual contents) is Adapted Material; and

- c. You must comply with the conditions in Section [3\(a\)](#) if You Share all or a substantial portion of the contents of the database.

For the avoidance of doubt, this Section [4](#) supplements and does not replace Your obligations under this Public License where the Licensed Rights include other Copyright and Similar Rights.

Section 5 – Disclaimer of Warranties and Limitation of Liability.

- a. **Unless otherwise separately undertaken by the Licensor, to the extent possible, the Licensor offers the Licensed Material as-is and as-available, and makes no representations or warranties of any kind concerning the Licensed Material, whether express, implied, statutory, or other. This includes, without limitation, warranties of title, merchantability, fitness for a particular purpose, non-infringement, absence of latent or other defects, accuracy, or the presence or absence of errors, whether or not known or discoverable. Where disclaimers of warranties are not allowed in full or in part, this disclaimer may not apply to You.**
- b. **To the extent possible, in no event will the Licensor be liable to You on any legal theory (including, without limitation, negligence) or otherwise for any direct, special, indirect, incidental, consequential, punitive, exemplary, or other losses, costs, expenses, or damages arising out of this Public License or use of the Licensed Material, even if the Licensor has been advised of the possibility of such losses, costs, expenses, or damages. Where a limitation of liability is not allowed in full or in part, this limitation may not apply to You.**
- c. The disclaimer of warranties and limitation of liability provided above shall be interpreted in a manner that, to the extent possible, most closely approximates an absolute disclaimer and waiver of all liability.

Section 6 – Term and Termination.

- a. This Public License applies for the term of the Copyright and Similar Rights licensed here. However, if You fail to comply with this Public License, then Your rights under this Public License terminate automatically.
- b. Where Your right to use the Licensed Material has terminated under Section [6\(a\)](#), it reinstates:
 1. automatically as of the date the violation is cured, provided it is cured within 30 days of Your discovery of the violation; or
 2. upon express reinstatement by the Licensor.For the avoidance of doubt, this Section [6\(b\)](#) does not affect any right the Licensor may have to seek remedies for Your violations of this Public License.
- c. For the avoidance of doubt, the Licensor may also offer the Licensed Material under separate terms or conditions or stop distributing the Licensed Material at any time; however, doing so will not terminate this Public License.
- d. Sections [1](#), [5](#), [6](#), [7](#), and [8](#) survive termination of this Public License.

Section 7 – Other Terms and Conditions.

- a. The Licensor shall not be bound by any additional or different terms or conditions communicated by You unless expressly agreed.
- b. Any arrangements, understandings, or agreements regarding the Licensed Material not stated herein are separate from and independent of the terms and conditions of this Public License.

Section 8 – Interpretation.

- a. For the avoidance of doubt, this Public License does not, and shall not be interpreted to, reduce, limit, restrict, or impose conditions on any use of the Licensed Material that could lawfully be made without permission under this Public License.
- b. To the extent possible, if any provision of this Public License is deemed unenforceable, it shall be

automatically reformed to the minimum extent necessary to make it enforceable. If the provision cannot be reformed, it shall be severed from this Public License without affecting the enforceability of the remaining terms and conditions.

- c. No term or condition of this Public License will be waived and no failure to comply consented to unless expressly agreed to by the Licensor.
- d. Nothing in this Public License constitutes or may be interpreted as a limitation upon, or waiver of, any privileges and immunities that apply to the Licensor or You, including from the legal processes of any jurisdiction or authority.

Creative Commons is not a party to its public licenses. Notwithstanding, Creative Commons may elect to apply one of its public licenses to material it publishes and in those instances will be considered the “Licensor.” The text of the Creative Commons public licenses is dedicated to the public domain under the [CC0 Public Domain Dedication](#). Except for the limited purpose of indicating that material is shared under a Creative Commons public license or as otherwise permitted by the Creative Commons policies published at creativecommons.org/policies, Creative Commons does not authorize the use of the trademark “Creative Commons” or any other trademark or logo of Creative Commons without its prior written consent including, without limitation, in connection with any unauthorized modifications to any of its public licenses or any other arrangements, understandings, or agreements concerning use of licensed material. For the avoidance of doubt, this paragraph does not form part of the public licenses.

Creative Commons may be contacted at creativecommons.org.

Additional languages available: [Bahasa Indonesia](#), [Deutsch](#), [français](#), [hrvatski](#), [italiano](#), [Nederlands](#), [norsk](#), [polski](#), [suomeksi](#), [svenska](#), [te reo Māori](#), [Türkçe](#), [українська](#), [العربية](#), [日本語](#). Please read the [FAQ](#) for more information about official translations.

**JOHN WILEY AND SONS LICENSE
TERMS AND CONDITIONS**

Dec 03, 2017

This Agreement between Mr. Joseph Szymborski ("You") and John Wiley and Sons ("John Wiley and Sons") consists of your license details and the terms and conditions provided by John Wiley and Sons and Copyright Clearance Center.

License Number	4241600451088
License date	Dec 03, 2017
Licensed Content Publisher	John Wiley and Sons
Licensed Content Publication	Cancer
Licensed Content Title	Prediction of breast cancer malignancy using an artificial neural network
Licensed Content Author	Carey E. Floyd,Joseph Y. Lo,A. Joon Yun,Daniel C. Sullivan,Phyllis J. Kornguth
Licensed Content Date	Dec 1, 1994
Licensed Content Pages	5
Type of use	Dissertation/Thesis
Requestor type	University/Academic
Format	Print and electronic
Portion	Figure/table
Number of figures/tables	1
Original Wiley figure/table number(s)	Figure 1
Will you be translating?	No
Title of your thesis / dissertation	On the Automated Regional Classification of Early Breast Cancers
Expected completion date	Dec 2017
Expected size (number of pages)	60
Requestor Location	Mr. Joseph Szymborski [REDACTED] Montreal, QC [REDACTED] Canada Attn: Mr. Joseph Szymborski
Publisher Tax ID	EU826007151
Billing Type	Invoice
Billing Address	Mr. Joseph Szymborski [REDACTED]

Montreal, QC [REDACTED]
Canada
Attn: Mr. Joseph Szymborski

Total

0.00 CAD

[Terms and Conditions](#)

TERMS AND CONDITIONS

This copyrighted material is owned by or exclusively licensed to John Wiley & Sons, Inc. or one of its group companies (each a "Wiley Company") or handled on behalf of a society with which a Wiley Company has exclusive publishing rights in relation to a particular work (collectively "WILEY"). By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the billing and payment terms and conditions established by the Copyright Clearance Center Inc., ("CCC's Billing and Payment terms and conditions"), at the time that you opened your RightsLink account (these are available at any time at <http://myaccount.copyright.com>).

Terms and Conditions

- The materials you have requested permission to reproduce or reuse (the "Wiley Materials") are protected by copyright.
- You are hereby granted a personal, non-exclusive, non-sub licensable (on a stand-alone basis), non-transferable, worldwide, limited license to reproduce the Wiley Materials for the purpose specified in the licensing process. This license, **and any CONTENT (PDF or image file) purchased as part of your order**, is for a one-time use only and limited to any maximum distribution number specified in the license. The first instance of republication or reuse granted by this license must be completed within two years of the date of the grant of this license (although copies prepared before the end date may be distributed thereafter). The Wiley Materials shall not be used in any other manner or for any other purpose, beyond what is granted in the license. Permission is granted subject to an appropriate acknowledgement given to the author, title of the material/book/journal and the publisher. You shall also duplicate the copyright notice that appears in the Wiley publication in your use of the Wiley Material. Permission is also granted on the understanding that nowhere in the text is a previously published source acknowledged for all or part of this Wiley Material. Any third party content is expressly excluded from this permission.
- With respect to the Wiley Materials, all rights are reserved. Except as expressly granted by the terms of the license, no part of the Wiley Materials

may be copied, modified, adapted (except for minor reformatting required by the new Publication), translated, reproduced, transferred or distributed, in any form or by any means, and no derivative works may be made based on the Wiley Materials without the prior permission of the respective copyright owner. **For STM Signatory Publishers clearing permission under the terms of the [STM Permissions Guidelines](#) only, the terms of the license are extended to include subsequent editions and for editions in other languages, provided such editions are for the work as a whole in situ and does not involve the separate exploitation of the permitted figures or extracts,** You may not alter, remove or suppress in any manner any copyright, trademark or other notices displayed by the Wiley Materials. You may not license, rent, sell, loan, lease, pledge, offer as security, transfer or assign the Wiley Materials on a stand-alone basis, or any of the rights granted to you hereunder to any other person.

- The Wiley Materials and all of the intellectual property rights therein shall at all times remain the exclusive property of John Wiley & Sons Inc, the Wiley Companies, or their respective licensors, and your interest therein is only that of having possession of and the right to reproduce the Wiley Materials pursuant to Section 2 herein during the continuance of this Agreement. You agree that you own no right, title or interest in or to the Wiley Materials or any of the intellectual property rights therein. You shall have no rights hereunder other than the license as provided for above in Section 2. No right, license or interest to any trademark, trade name, service mark or other branding ("Marks") of WILEY or its licensors is granted hereunder, and you agree that you shall not assert any such right, license or interest with respect thereto
- NEITHER WILEY NOR ITS LICENSORS MAKES ANY WARRANTY OR REPRESENTATION OF ANY KIND TO YOU OR ANY THIRD PARTY, EXPRESS, IMPLIED OR STATUTORY, WITH RESPECT TO THE MATERIALS OR THE ACCURACY OF ANY INFORMATION CONTAINED IN THE MATERIALS, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OF MERCHANTABILITY, ACCURACY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, USABILITY, INTEGRATION OR NON-INFRINGEMENT AND ALL SUCH WARRANTIES ARE HEREBY EXCLUDED BY WILEY AND ITS LICENSORS AND WAIVED BY YOU.
- WILEY shall have the right to terminate this Agreement immediately upon breach of this Agreement by you.

- You shall indemnify, defend and hold harmless WILEY, its Licensors and their respective directors, officers, agents and employees, from and against any actual or threatened claims, demands, causes of action or proceedings arising from any breach of this Agreement by you.
- IN NO EVENT SHALL WILEY OR ITS LICENSORS BE LIABLE TO YOU OR ANY OTHER PARTY OR ANY OTHER PERSON OR ENTITY FOR ANY SPECIAL, CONSEQUENTIAL, INCIDENTAL, INDIRECT, EXEMPLARY OR PUNITIVE DAMAGES, HOWEVER CAUSED, ARISING OUT OF OR IN CONNECTION WITH THE DOWNLOADING, PROVISIONING, VIEWING OR USE OF THE MATERIALS REGARDLESS OF THE FORM OF ACTION, WHETHER FOR BREACH OF CONTRACT, BREACH OF WARRANTY, TORT, NEGLIGENCE, INFRINGEMENT OR OTHERWISE (INCLUDING, WITHOUT LIMITATION, DAMAGES BASED ON LOSS OF PROFITS, DATA, FILES, USE, BUSINESS OPPORTUNITY OR CLAIMS OF THIRD PARTIES), AND WHETHER OR NOT THE PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY PROVIDED HEREIN.
- Should any provision of this Agreement be held by a court of competent jurisdiction to be illegal, invalid, or unenforceable, that provision shall be deemed amended to achieve as nearly as possible the same economic effect as the original provision, and the legality, validity and enforceability of the remaining provisions of this Agreement shall not be affected or impaired thereby.
- The failure of either party to enforce any term or condition of this Agreement shall not constitute a waiver of either party's right to enforce each and every term and condition of this Agreement. No breach under this agreement shall be deemed waived or excused by either party unless such waiver or consent is in writing signed by the party granting such waiver or consent. The waiver by or consent of a party to a breach of any provision of this Agreement shall not operate or be construed as a waiver of or consent to any other or subsequent breach by such other party.
- This Agreement may not be assigned (including by operation of law or otherwise) by you without WILEY's prior written consent.

- Any fee required for this permission shall be non-refundable after thirty (30) days from receipt by the CCC.
- These terms and conditions together with CCC's Billing and Payment terms and conditions (which are incorporated herein) form the entire agreement between you and WILEY concerning this licensing transaction and (in the absence of fraud) supersedes all prior agreements and representations of the parties, oral or written. This Agreement may not be amended except in writing signed by both parties. This Agreement shall be binding upon and inure to the benefit of the parties' successors, legal representatives, and authorized assigns.
- In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall prevail.
- WILEY expressly reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.
- This Agreement will be void if the Type of Use, Format, Circulation, or Requestor Type was misrepresented during the licensing process.
- This Agreement shall be governed by and construed in accordance with the laws of the State of New York, USA, without regards to such state's conflict of law rules. Any legal action, suit or proceeding arising out of or relating to these Terms and Conditions or the breach thereof shall be instituted in a court of competent jurisdiction in New York County in the State of New York in the United States of America and each party hereby consents and submits to the personal jurisdiction of such court, waives any objection to venue in such court and consents to service of process by registered or certified mail, return receipt requested, at the last known address of such party.

WILEY OPEN ACCESS TERMS AND CONDITIONS

Wiley Publishes Open Access Articles in fully Open Access Journals and in Subscription journals offering Online Open. Although most of the fully Open Access journals publish open access articles under the terms of the Creative Commons Attribution (CC BY) License only, the subscription journals and a few of the Open Access Journals offer a choice of Creative Commons Licenses. The license type is

clearly identified on the article.

The Creative Commons Attribution License

The [Creative Commons Attribution License \(CC-BY\)](#) allows users to copy, distribute and transmit an article, adapt the article and make commercial use of the article.

The CC-BY license permits commercial and non-

Creative Commons Attribution Non-Commercial License

The [Creative Commons Attribution Non-Commercial \(CC-BY-NC\)License](#) permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.(see below)

Creative Commons Attribution-Non-Commercial-NoDerivs License

The [Creative Commons Attribution Non-Commercial-NoDerivs License](#) (CC-BY-NC-ND) permits use, distribution and reproduction in any medium, provided the original work is properly cited, is not used for commercial purposes and no modifications or adaptations are made. (see below)

Use by commercial "for-profit" organizations

Use of Wiley Open Access articles for commercial, promotional, or marketing purposes requires further explicit permission from Wiley and will be subject to a fee. Further details can be found on Wiley Online Library <http://olabout.wiley.com/WileyCDA/Section/id-410895.html>

Other Terms and Conditions:

v1.10 Last updated September 2015

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.
