

JOINT AND INDIVIDUAL ANALYSIS OF BREAST CANCER HISTOLOGIC IMAGES AND GENOMIC COVARIATES

BY IAIN CARMICHAEL¹, BENJAMIN C. CALHOUN², KATHERINE A. HOADLEY³,
MELISSA A. TROESTER⁴, JOSEPH GERADTS⁵, HEATHER D. COUTURE⁶,
LINNEA OLSSON⁷, CHARLES M. PEROU⁸, MARC NIETHAMMER⁹, JAN HANNIG^{10,*} AND
J. S. MARRON^{10,†}

¹*Department of Statistics, University of Washington, ic9@uw.edu*

²*Department of Pathology and Laboratory Medicine, University of North Carolina at Chapel Hill, ben.calhoun@unchealth.unc.edu*

³*Department of Genetics, Lineberger Comprehensive Cancer Center, Computational Medicine Program, University of North Carolina at Chapel Hill, hoadley@med.unc.edu*

⁴*Department of Epidemiology, Department of Pathology and Laboratory Medicine, University of North Carolina at Chapel Hill, troester@unc.edu*

⁵*Department of Pathology and Laboratory Medicine, East Carolina University Brody School of Medicine, geradtsj20@ecu.edu*

⁶*Pixel Scientia Labs, heather@pixelscientia.com*

⁷*Department of Epidemiology, University of North Carolina at Chapel Hill, lolsson@live.unc.edu*

⁸*Department of Genetics, Department of Pathology and Laboratory Medicine, University of North Carolina at Chapel Hill, cperou@med.unc.edu*

⁹*Department of Computer Science, University of North Carolina at Chapel Hill, mn@cs.unc.edu*

¹⁰*Department of Statistics, University of North Carolina at Chapel Hill, *jan.hammig@unc.edu; †marron@unc.edu*

The two main approaches in the study of breast cancer are histopathology (analyzing visual characteristics of tumors) and genomics. While both histopathology and genomics are fundamental to cancer research, the connections between these fields have been relatively superficial. We bridge this gap by investigating the Carolina Breast Cancer Study through the development of an integrative, exploratory analysis framework. Our analysis gives insights—some known, some novel—that are engaging to both pathologists and geneticists. Our analysis framework is based on angle-based joint and individual variation explained (AJIVE) for statistical data integration and exploits convolutional neural networks (CNNs) as a powerful, automatic method for image feature extraction. CNNs raise interpretability issues that we address by developing novel methods to explore visual modes of variation captured by statistical algorithms (e.g., PCA or AJIVE) applied to CNN features.

1. Introduction. Histologic images (Figure 1) of tissue morphology have long been utilized in treatment decisions and prognostics for breast cancer. For example, tumor grade is scored by evaluating tumor morphologic features and has high value in predicting outcomes for cancer cases. Recent years have found that genomic assays can offer a second line of evidence to guide treatment and prognostics. While both histopathology and genomics are known to be valuable, they are typically assessed separately in both clinical and research settings. Pathology data is almost immediately available after a tumor is excised, whereas genomic data may not arrive for many weeks. Efforts to integrate these two types of information are typically confined to assessing whether the pathology and the genomic data are concordant in their estimation of progression risk. Unfortunately, this means that information that may be gained by understanding the interaction between histology and genomics has been largely neglected.

Received April 2020; revised October 2020.

Key words and phrases. Multiview data, dimensionality reduction, image analysis, deep learning, interpretability, breast cancer histopathology, gene expression.




FIG. 1. A lower resolution view of a hematoxylin and eosin (H&E) stained one mm core from a breast cancer tissue microarray. The darker purple color (hematoxylin) stains nuclear material, such as DNA, while the pink (eosin) highlights structures, such as connective tissue.

The primary goal of our study is to leverage both pathology and genomic data in a more concerted fashion. Our approach, based on angle-based joint and individual variation explained (AJIVE) (Feng et al. (2018)) allows domain experts to explore how information is shared across histopathology and genetic data. We expect some information to be jointly shared by both data modalities, for example, some tumors may have a high mitotic index in the pathology data and show high expression of proliferation-related genes. Other information may be contained in one modality but not the other, for example, overexpression of a particular oncogene may not manifest morphologically, or some microenvironment features that are pronounced in images, such as mucin or adipose content, may have limited effects on gene expression. By identifying signals that are *joint* and those that are *individual* (present in one data type but not the other), we provide a powerful exploratory analysis framework.

One historical barrier to histologic image analysis is that morphologic features are often described only by qualitative features (e.g., high, medium or low nuclear grade) which are determined by pathologists. Extracting this information is time consuming, expensive, ignores information and is subject to inter-rater variability (Elmore et al. (2015)). Recent quantitative histopathology approaches have begun to address these issues by developing statistical tools for histological image data (Beck et al. (2011), Ash et al. (2018)). For example, *convolutional neural networks* (CNNs) have proven adept at solving predictive tasks in cancer histology (Couture et al. (2018), Liu et al. (2017), Chen et al. (2019)).

Inspired by the success of CNNs for predictive tasks, we leverage CNN features for automatic image feature extraction in our analysis. While CNN features offer representational advantages, they are challenging to interpret and are not traditionally used for exploratory/inferential analyses. This leads to a secondary goal of our study, developing an approach to interpret signals in the histology image data captured by CNN features (*representative patch views*, see Section 3). These representative patch views are critical for bringing different domains together (pathology, genomics and epidemiology) because they help enable experts from these different domains to engage deeply with each other through our analysis.

A limiting factor for this type of research is availability of well-annotated datasets with both digital pathology and genomic data. When both data types are available, studies tend to be of relatively small size. However, the Carolina Breast Cancer Study (CBCS) motivating our investigation includes histologic images and RNA expression data for a large cohort ($n = 1191$) of breast cancer patients diagnosed with invasive breast cancer between 2008 and 2013 (Troester et al. (2017)).

Our analysis discovers immediately interpretable joint and individual signals. The first AJIVE joint component uncovers tumor grade in the pathology and the basal-like subtype in

the genetics. The association between high grade and basal-like tumors is well known in the breast cancer literature (Carey et al. (2006)). The fact this association shows up prominently in the joint information is highly encouraging to domain experts.

The second joint component identifies previously unknown histologic features of luminal B tumors. Understanding the histology of luminal B tumors is immediately relevant to clinicians because genetic based testing can be expensive and time consuming. Histologic grade is one feature that distinguishes luminal A and B tumors, but it has limited accuracy (Allott et al. (2018)). Our analysis suggests that retraction artifacts (among other things) can be used by pathologists to identify luminal B tumors. These histological features are traditionally ignored by pathologists as artifacts of the tumor processing pipeline. These novel insights will be validated in a followup study for use by pathologists in clinical settings.

The histologic individual AJIVE components capture features of the tumor microenvironment including *mucinous carcinoma*, high adipose content, and degraded tissue. These features further validate that we can accurately separate individual information from joint. The genetic individual AJIVE components contain additional genetic subtype information as well as pick up on well-known technical variation.

Section 2 presents the data provided by CBCS as well as the patch based, CNN image features extraction approach. Section 3 presents our approach for interpreting visual modes of variation. Section 4 discusses the integrative analysis, including an overview of AJIVE. Section 5 discusses the results of the joint, image individual and genetic individual AJIVE components. Finally, Section 6 concludes with more discussion about transfer learning and exploratory analysis with deep learning. The Supplementary Material provides additional results and details, including explanation of important tissue structures discussed in the results section, discussing of hyperparameter choices, all figures shown to the pathologists and more details about statistical procedures. Additional visualizations may be found in Supplement B (Carmichael et al. (2021)) or online at <https://marronwebfiles.sites.oasis.unc.edu/AJIVE-Hist-Gene/>.

All correlations and AUC statistics reported in the text of the paper are statistically significant at a level of 0.05 after controlling for multiple testing with the Benjamini–Hochberg procedure (Benjamini and Hochberg (1995)), unless stated otherwise (see Supplement A, Carmichael et al. (2021)).

1.1. Related literature. There is a large literature on dimensionality reduction for multi-block data including classical algorithms (Hotelling (1936), Kettenring (1971), Wold (1985), Yang and Michailidis (2016), Gaynanova and Li (2019)). JIVE Lock et al. (2013) and AJIVE Feng et al. (2018) are some of the first methods to look at both joint as well as individual modes of variation.

Interpretability in deep learning is a growing field (Chen et al. (2018a), Kim et al. (2018), Molnar et al. (2018), Olah et al. (2018), Vellido, Martín-Guerrero and Lisboa (2012), Holzinger et al. (2019)). We explored adapting *saliency map* approaches (Selvaraju et al. (2017), Springenberg et al. (2014), Sundararajan, Taly and Yan (2017), Zeiler and Fergus (2014), Adebayo et al. (2018)) for interpreting the results of our analysis. Unfortunately, none of the methods provided consistently interpretable outputs (potentially due to our use of transfer learning) and raised issues which will be explored in a follow up paper.

Deep learning based predictive analysis of histological images is a growing area (Komura and Ishikawa (2018), Aeffner et al. (2019)) which includes tasks such as classification/regression (Bejnordi et al. (2018), Ilse, Tomczak and Welling (2018), Liu et al. (2017, 2018), Veta et al. (2019), Wang et al. (2016)), semantic segmentation Jiménez and Racoceanu (2019), Mahmood et al. (2019) and microscope augmentation (Chen et al. (2018b)). CNN architectures that integrate genetic (or other) information are also being explored for these predictive tasks (Couture et al. (2018), Mahmood et al. (2018), Srivastava

et al. (2018), Chen et al. (2019)). Other papers train networks to predict genetic mutations directly from histological images (Coudray et al. (2018)). Other studies used nondeep learning based methods to do exploratory, integrative analysis of histological and genetic data (Beck et al. (2011), Wang et al. (2013), Cooper et al. (2015))

A similar joint, exploratory analysis of breast cancer H&E image and gene expression data was performed by (Ash et al. (2018)). Our methods differ from theirs in a number of ways: they only examined joint signals, while we examine both joint and individual signals; they used a sparse CCA, while we use AJIVE; we develop and use the RPVs for image interpretation; they trained an autoencoder, while we use transfer learning. An important result of our paper is that even simple transfer learning effectively captures the important signals in the data.

1.2. *Software and data release.* The code to reproduce the analysis in this paper can be found in Supplement C (Carmichael et al. (2021)) and on github at https://github.com/idc9/breast_cancer_image_analysis. The raw data, for example, H&E images, gene expression data, clinical variables cannot be released publicly due to patient confidentiality concerns. Researchers may request permission to access the raw data used in this study by visiting <https://unclineberger.org/cbcs/for-researchers/>.

We used many of the standard Python data science libraries (Hunter (2007), Jones, Oliphant and Peterson (2014), McKinney (2011), Pedregosa et al. (2011), Van Der Walt, Colbert and Varoquaux (2011), van der Walt et al. (2014), Waskom et al. (2018)). The PyTorch framework is used for all neural network computations, and the pre-trained VGG16 weights are downloaded with the PyTorch vision library (Paszke et al. (2017)). AJIVE computations are done with the `py_jive` package (Carmichael (2020)) which was developed for this project.

2. Data. This section describes the data used in the analysis as well as our image feature extraction pipeline. Tuning parameter selection is discussed in more detail in Section 6 and in Supplement A.

2.1. *Carolina breast cancer study.* The data are from the Carolina Breast Cancer Study, a population-based study of black and white women with invasive breast cancer diagnosed between 2008–2013 in North Carolina. Tumor blocks were collected and cores were transferred from the donor paraffin blocks to prepare tissue microarrays as well as to isolate RNA for gene expression analysis. The current analysis includes $n = 1191$ patients for whom both image and gene expression data were available. Additional details about these data (e.g., descriptive statistics, tissue preparation, gene expression processing) are described in (Troester et al. (2017), Allott et al. (2018)).

For each patient a pathologist reviewed a paraffin-embedded tumor block and marked the area containing the invasive carcinoma. Then, a lab technician extracted a number of circular “cores” which were then transferred into a recipient TMA paraffin block and eosin (H&E) and imaged. Supplement B shows a graphical depiction of this process. The upshot is that, for each patient, we have a median¹ of four H&E stained core images. The images of these cores are roughly square with an average width of approximately 2500 pixels.² An example core image is shown in Figure 1. It is appealing to work with cores and not the much larger *whole slide images* because the cores provide more concentrated tumor cells and are more computationally tractable.³

¹Minimum of one and maximum of eight.

²Min 600, max 3400.

³The whole slide images can be of order $50,000 \times 50,000$ pixels or larger.

Pathologic evaluation of the tumor (including histologic type and grade) was based on the original whole tissue sections. We also compute a number of variables describing image features, such as the proportion of white background and the median intensity of the background pixels.

For each patient, we have the PAM50 gene expression measurements which are 50 genes chosen to distinguish the five clinically relevant, genetic subtypes (basal-like, luminal A, luminal B, molecular HER2 and normal-like) (Parker et al. (2009)). The intrinsic subtype gene list was developed using genes which were consistently expressed within the tumor while minimizing the contribution of the nontumor microenvironment; therefore, the PAM50 genes do not describe the tumor microenvironment (Perou et al. (2000)). We also have variables derived from the PAM50 gene expression, such as proliferation score and risk of recurrence score (ROR-PT) (Parker et al. (2009)).

CBCS provides clinically relevant immunohistochemical variables (ER status, clinical HER2 status and PR status) which are derived from routine methods used in the clinical laboratory.

The 50 gene expression variables are centering and scaled by their standard deviation, resulting in the gene expression data matrix with 1191 samples and 50 features.

2.2. Image processing and patch representation. In order to achieve uniform visual stain density, the raw H&E core images are stain normalized using the procedure described in Macenko et al. (2009). The set of background pixels of each image (i.e., the whitespace in Figure 1) is then estimated via the following procedure. Each image is converted to grayscale; then, a background pixel intensity threshold is estimated with a weighted combination of (0.1) Otsu’s method Otsu (1979) and (0.9) the triangle method (Zack, Rogers and Latt (1977)). The *background mask* (True/False array saying whether or not a pixel is in the background) is then used for a variety of downstream tasks.

Next, we create a *patch-based* representation of each image. Each core image is broken into a grid of 200×200 pixel patches. To make an even grid of patches, the image is first padded with the estimated typical background pixel so its dimensions are divisible by 200. Using the background mask, patches which are more than 90% background are thrown out (Figures 2a and 2b). The background threshold (90%) was selected via manual inspection to be the smallest value such that patches with large amounts of fat and some tissue are still included (Figure 2b).

There are 1191 subjects in the study resulting in 5920 core images from which we extract a total of 761,767 patches. We estimate the channel (red, green, blue values for each



(a) The patch grid for an example core.

(b) A core with high fat content.

FIG. 2. *Patch grid for two example cores, one with low fat content (a) and one with high fat content (b). Black squares indicate patches with too much background that were excluded. The background threshold is selected such that the thin pieces of tissue surrounding fat cells (and other structures with lots of white space such as mucin) are included.*

pixel) mean and standard deviation from the patch dataset. Before being input into the neural network, each pixel channel is mean centered then scaled by the standard deviation.

The choices of patch size, background intensity threshold and background proportion patch threshold are discussed further in Supplement A.

2.3. CNN feature extraction. After the raw images are processed, CNN features are extracted from each patch. We use the last convolutional layer of the pretrained VGG16 architecture (Simonyan and Zisserman (2014)), with an additional *spatial mean pooling* layer added to the end of this architecture, to average out spatial information resulting in 512 CNN features. The network was pretrained on the famous ImageNet dataset and the pretrained weights of the network are downloaded from the torch vision library (Paszke et al. (2017)). The use of pretrained features allows us to avoid training a CNN on our limited dataset, and no additional fine-tuning is performed (see Section 6.2 for discussion).

Finally, core images are represented as an average of their patch features (again, ignoring patches which are over 90% background). Patients are then represented by an average of their cores. Supplement B provides a graphical representation of this process. Each CNN feature is first mean centered and then scaled by its standard deviation, resulting in the image feature data matrix with 1191 samples (each subject) and 512 neural network image features.

3. Representative patch views. A key challenge for doing exploratory/inferential analysis on populations of images using deep learning is interpretability. In this section we explain the novel, broadly applicable RPV method for interpretation of the visual signals captured by CNN (or other) features.

The RPVs assume images⁴ are represented via the patch based approach described in Sections 2.2 and 2.3. In other words, images are broken into a collection of patches; image features are extracted for each patch, and then images are represented as an average of their patch features. Suppose we compute a loadings⁵ vector of image features (e.g., the first PCA component). The RPVs highlight the visual features in one image associated with one end of a loadings vector (e.g., the positive end of PC 1 for the image with the most positive scores).

Figure 3 shows an example RPV for one subject for the negative end of AJIVE joint component 1. The leftmost column shows the four cores for this subject. The rightmost five columns show the top 20 patches for the negative end of the first joint component from the *patch based localization* approach described below in Section 3.1. The second column shows the location on each core of the top 20 patches. The RPVs are multiscale, in the sense that they give insights at both the core level and the patch level.

3.1. Patch based localization. Here, we describe the general approach used to select the representative patches for the RPVs described in the above section. Patch based localization is accomplished by projecting patches onto the loadings vector. Main ideas are illustrated below in the context of PCA and can be extended to other methods (e.g., AJIVE, linear regression, etc). For PCA the goal of the RPVs is to understand *one end of one PCA component*. The RPVs in Section 5 are all obtained using AJIVE.

Let n be the number of subjects (images) in the dataset, m_i be the number of patches for the i th image and p_{ij} be the j th patch for the i th subject. Let $f : \text{image} \rightarrow \mathbb{R}^k$ be a generic feature extraction function outputting k features. Note that f can be any feature extraction

⁴In the CBCS study each subject has a number of core-images and subjects are represented as an average of their images. For exposition purposes we pretend each subject has one image in this section, however, the extension to the multiimage case is clear.

⁵We use the convention that *loadings* correspond to features while *scores* correspond subjects.




FIG. 3. The full representative patch view from Figure 5. The first column shows the four cores for this subject (high resolution is needed to see detail). The 20 patches (Section 3.1) in the last five columns are representative of the visual features associated with one extreme of a mode of variation. The second column shows the patches' locations on the cores. Lymphocytes show up prominently in these patches. Some images (1, 8, 11, 18) predominantly show tumor cells, and others (10, 14, 15) show a roughly even mixture of lymphocytes and tumor cells.

function, for example, in Section 5, f is the CNN feature extraction function described in Section 2. Also, let $z_{ij} := f(p_{ij})$ be the features for patch p_{ij} , $z_i := \frac{1}{m_i} \sum_{j=1}^{m_i} z_{ij}$ be the average patch features for the i th subject and $Z \in \mathbb{R}^{n \times k}$ be the image feature dataset whose rows are the z_i . Finally, let $\tilde{Z} \in \mathbb{R}^{n \times k}$ be the matrix Z after preprocessing (e.g., centering and scaling). Suppose we compute the first PCA component of the processed data; let $\mathbf{v} \in \mathbb{R}^k$ be a loadings vector and $\mathbf{u} = \tilde{Z}\mathbf{v} \in \mathbb{R}^n$ the scores vector.⁶

Consider the positive end of this component, and let i^* be the index of a particular image. We perform *patch based localization*⁷ by projecting every patch of subject i^* onto \mathbf{v} . In detail, let $\tilde{z}_{i^*,j}$ be the features of the j th patch for subject i^* after the processing. Let $s_{i^*,j} := \tilde{z}_{i^*,j}^T \mathbf{v}$ be the scores of this patch for $j = 1, \dots, m_{i^*}$. Now, let $j_{(1)}, \dots, j_{(20)}$ be the indices of the patches with the 20 most positive patch scores (i.e., $s_{i^*,j_{(1)}} \geq s_{i^*,j_{(2)}} \dots$). We call these the *representative patches* for the positive end of this component.

4. Integrative analysis. This section gives an overview of the integrative image and genetic analysis. We assume image features for each patient have been extracted, as described in Section 2.1. The gene expression data have been processed as in Allott et al. (2018), Troester et al. (2017) with additional mean centering and scaling. The first subsection describes the AJIVE analysis, and the second subsection describes the pathology review process for interpreting the AJIVE image modes of variation.

4.1. *Angle-based joint and individual variation explained.* AJIVE is a statistical feature extraction/dimensionality reduction algorithm for multiblock data (Feng et al. (2018)). The

⁶We assume that the scores are the projection of the data onto the loadings vector.

⁷We use the term localization because this method helps identify which regions in the image are playing an important role in the given mode of variation.

goal of AJIVE is to find joint signals, if any exist, which are common to all data blocks as well as individual signals which are specific to each block, if they exist. Here, we give a brief overview of AJIVE for two data blocks.

Consider two data blocks $X \in \mathbb{R}^{n \times d_x}$, $Y \in \mathbb{R}^{n \times d_y}$ on the same set of n observations. AJIVE estimates what variation is *joint* to both data blocks as well as what variation is *individual* to each block. In particular, each matrix is decomposed into a sum of *joint*, *individual* and *error* terms,

$$X = J^x + I^x + E^x \quad \text{and} \quad Y = J^y + I^y + E^y$$

while imposing the following constraints:

- $\text{col-span}(J^x) = \text{col-span}(J^y) := \mathcal{J} \subseteq \mathbb{R}^n$;
- $\text{col-span}(J^x) \perp \text{col-span}(I^x)$ and $\text{col-span}(J^y) \perp \text{col-span}(I^y)$;
- $\text{col-span}(I^x) \cap \text{col-span}(I^y) = \emptyset$.

All subspaces live in \mathbb{R}^n where n is the number of observations. The two joint matrices span the same *joint subspace*, \mathcal{J} . The two individual matrices span subspaces which are orthogonal to the joint subspace. We refer to the rank of the joint subspace as the *joint rank*, R_J , and the rank of the two individual subspaces as the X and Y individual ranks, R_x , R_y .

The mechanics of AJIVE are outlined below for the case of $B = 2$ data blocks.⁸ The properties of the common normalized scores discussed below follow from the fact that they are the subspace flag mean of the PCA scores subspaces (Draper et al. (2014)). We use a different estimate of the block common loadings, L^x , L^y than in the original paper. One of the key statistical procedures in AJIVE is to estimate the joint rank⁹ which is achieved using the Wedin bound and the random direction bound detailed in (Feng et al. (2018)):

1. *Initial signal extraction*: Estimate low rank PCAs of X , Y with ranks r_{init}^x , r_{init}^y (e.g., selected by inspecting the PCA scree plots). Denote this initial PCA of X by U_{init}^x , D_{init}^x , V_{init}^x where $U_{\text{init}}^x \in \mathbb{R}^{n \times r_{\text{init}}^x}$, $V_{\text{init}}^x \in \mathbb{R}^{d_b \times r_{\text{init}}^x}$, similarly for y .

2. *Signal space extraction*: Perform CCA on the PCA scores, U_{init}^x , U_{init}^y . Using the *random direction bound* and the *Wedin bound* estimate the CCA rank, R_J . Let $S^x, S^y \in \mathbb{R}^{n \times R_J}$ be the matrices whose columns are the x/y CCA scores with unit norm. Let $W^x \in \mathbb{R}^{r_{\text{init}}^x \times R_J}$, $W^y \in \mathbb{R}^{r_{\text{init}}^y \times R_J}$ be the matrices whose columns are the CCA x/y loadings. Let $C \in \mathbb{R}^{n \times R_J}$ be the *common normalized scores* which have the property of being proportional to the average of the x/y CCA scores. In other words, the j th column of C is unit norm and is proportional to the average of the j th columns of S^x and S^y . Additionally, the common normalized scores are orthonormal, that is, $C^T C = I_{R_J}$. Finally, let¹⁰ $L^x := V_{\text{init}}^x D_{\text{init}}^{x^{-1}} W^x \in \mathbb{R}^{d_x \times R_J}$ be the *x -common loadings*, similarly for y .

3. *Signal space extraction*: Let $J^x := CC^T X$ be the estimated joint matrix (i.e. projection onto the joint subspace). Let $\tilde{I}^x := (I - CC^T)X \in \mathbb{R}^{n \times d_x}$ (i.e., projection onto the orthogonal complement of the joint subspace). Let R_x be the number of singular values of \tilde{I}^x above the threshold discussed in Section 2.4 of Feng et al. (2018), and let $I^x \in \mathbb{R}^{n \times d_x}$ be the rank R_x SVD approximation of \tilde{I}^x . We denote the PCA of the individual matrix I^x by $U_{\text{individual}}^x \in \mathbb{R}^{n \times R_x}$, $D_{\text{individual}}^x \in \mathbb{R}^{R_x \times R_x}$, $V_{\text{individual}}^x \in \mathbb{R}^{d_x \times R_x}$ which is also of interest, similarly for y .

The outputs of interest in this paper are the following:

⁸The original paper describes the procedure somewhat differently, but this description is equivalent.

⁹This is accomplished by estimating which principal angles between $\text{col-span}(U_{\text{init}}^x)$ and $\text{col-span}(U_{\text{init}}^y)$ are smaller than random in an appropriate sense.

¹⁰Note the j th column of V^x is equivalent to the rank r_{init}^x principal components regression coefficient of the j th column of the common normalized scores, C , regressed on the X matrix.

- The joint rank, R_J .
- The *common normalized scores*, $C \in \mathbb{R}^{n \times R_J}$.
- The *common loadings*,¹¹ $L^x \in \mathbb{R}^{d_x \times R_J}$, $L^y \in \mathbb{R}^{d_y \times R_J}$.
- The $U_{\text{individual}}^x \in \mathbb{R}^{n \times R_x}$ and $V_{\text{individual}}^x \in \mathbb{R}^{d_x \times R_x}$ which are referred to as the *block specific, individual scores and loadings*, similarly for y .

The common loadings, L^x , L^y are different than those in (Feng et al. (2018)). The loadings computed here are the loadings such that $XL^x + YL^y \propto C$, that is, the average of the resulting scores are proportional to the common normalized scores. Computing the loadings in this way ensures that they incorporate joint information only.

REMARK 4.1. It can be checked that the random direction bound is equivalent to the classical Roy's largest root test CCA rank selection method (Johnstone (2008)).

4.1.1. *AJIVE analysis of CBCS data.* An AJIVE analysis is performed on the 1191 subjects in the analysis. The first data block is the 512 CNN image features described in Section 2; the second data block is the 50 PAM50 genes. All other variables discussed below are used to interpret the AJIVE results. The initial signal ranks are 81 (image features) and 30 (genes) and were selected by inspection of the difference of the log-singular values and airing on the side of picking too high a rank. AJIVE estimates a joint rank of 7, image individual rank of 76 and genetic individual rank of 25. The AJIVE diagnostic plot, detailed in Feng et al. (2018), is provided in Supplement A.

4.2. *Pathology review of images.* In close collaboration with pathologists (B.C. and J.G.), we reviewed the first three joint and image individual components at two levels of granularity.

In the first approach, which we refer to as *global sort*, all core images are reviewed in sequence after sorting by the patient scores. Joint components are sorted by common normalized scores, C , and individual components are sorted by block specific scores, $U_{\text{individual}}^{\text{image}}$ (see Section 4.1). After sorting, the images are reviewed in sequence (e.g., from the negative to the positive end) to explore the visual signals captured by a given component. The benefits of the global sort method are: (i) a large number of images are inspected, and (ii) we get a sense of the high level changes¹² as we move along a component from the extreme negative to the middle then to the extreme positive end and (iii) we can see if the trends found in the RPVs (see next paragraph) hold broadly for the entire component. The downsides of this method are that it is time intensive and does not provide explicit information about what visual signals are important in a given image. The H&E images are quite large and complex and finding patterns across a set of images is challenging.

The RPV approach developed in Section 3 extracts more fine-grained information at the patient level. The RPVs of the 15 most negative and 15 most positive subjects are inspected for each component. The RPVs are created with the common loadings L^{image} for the joint components and the block specific individual loadings $V_{\text{individual}}^{\text{image}}$ (see Section 4.1). The number 15 was selected to balance showing "enough" information without taking too much time. The RPVs have the benefit of highlighting a more focused set of visual patterns.

Tables 1 and 3 display the pathologist's observations based on the RPVs at each end of each component. Each column summarizes the pathologist's impression of a clinically relevant

¹¹These were not given names in the Feng et al. (2018) and were computed slightly differently.

¹²In a preliminary analysis where image patches with a large amount of background were not excluded (see Section 2.2), the global sort method on the first few principal components revealed that the primary modes of variation in the data are driven by the raw amount of background. This motivated the exclusion of patches with too much background.

TABLE 1

A pathologist’s summary of important clinical features based on the RPVs of the 15 most extreme subjects for each end of the first three joint components. The “homogeneous” column indicates whether or not the patterns shown in the RPVs were consistent across the 15 subjects

Component	End	Homo- geneous	Tumor cellularity	Tubule formation	Nuclear grade	Adipocytic stroma	Collagenous stroma	Lym- phocytes	Necrosis
1	positive	no	low	yes	1, 2	yes	yes	no	no
	negative	yes	high	no	3	no	limited	yes	yes
2	positive	no	variable	yes	3	focal	yes	few	no
	negative	yes	moderate/high	yes	2	focal	yes	no	no
3	positive	no	variable	yes	3	yes	limited	yes	no
	negative	yes	moderate/high	no	3	no	yes	no	no

histological feature. The *homogeneous* column indicates whether or not there appeared to be a consistent pattern across the reviewed RPVs. The global sort review shows these trends hold for more than just the 15 most extreme images. These observations are key to understanding the connections between the pathology and the genetics.

5. Results. This section discusses the results for the joint AJIVE components (Section 5.1), the image individual (Section 5.2) and genetic individual (Section 5.3). For the sake of time—both the readers’ and the pathologists’—we focus on the top three components from each of the joint, image individual and genetic individual.

The pathology review of the images from the joint and image individual components is described in Section 4.2. While the pathologist reviewed the full RPVs (Figure 3), only mini-RPVs (e.g., Figure 9a) displaying eight patches are shown below in the text of the paper. The full RPVs shown to the pathologists, all AJIVE genetic loadings and all clinical data comparisons are provided in Supplement B. The methodology for clinical data comparisons (e.g., multiple testing control) is discussed in Supplement A.

5.1. Joint image and genetic information. Table 1 summarizes the pathologist’s observations based on the RPVs of the first three joint components. Table 2 shows the association between immunohistochemical status (ER and clinical Her2) and the first three joint components.

5.1.1. First AJIVE joint component. We initially consider the negative and positive extremes of the first joint component separately.

TABLE 2

AUC scores for two immunohistochemical (IHC) variables, ER status and clinical HER2 status (positive vs. negative classes), for first three joint components based on AJIVE common normalized scores. All six of these comparisons are statistically significant. The sign in parentheses indicates which end of the component the IHC positive status cluster on if there was a clear visual separation in the histogram (see Supplement B). For example, ER status positive tumors cluster on the positive end of component 1

Component	ER status	Clinical HER2 status
1	0.883 (+)	0.558
2	0.752 (−)	0.617 (+)
3	0.551	0.777 (−)




FIG. 4. Several representative patches from the most negative subject of joint component 1. The dark, round, purple structures are lymphocyte nuclei; these patches show dense tumor infiltrating lymphocytes which are characteristic of high central grade tumors. The third images in both rows show a significant number of tumor cells, mixed with some lymphocytes. This tumor is a genetically basal-like tumor like most of the other tumors on the negative end of this component (Figure 7b).

From the pathology perspective, two distinct visual patterns show up in the negative end of the first joint component (Figures 4 and 5). Supplement A has a brief explanation of the various tumor structures which are relevant to this paper. The first pattern is dense *tumor infiltrating lymphocytes* (TILs) and is illustrated in Figure 4 which shows the RPV of the most negative subject of the first joint component. The smaller cells, which have hyperchromatic round nuclei and relatively scant cytoplasm (i.e., dark, round, purple structures), are lymphocytes. In particular types of breast cancer, TILs in the intratumoral stroma are associated with




FIG. 5. Patches from a tumor on the negative end of joint components 1. The patches selected for the RPV of this tumor show high nuclear grade cells which are large and irregularly shaped. These are also characteristic of high grade tumors. This tumor is also genetically basal-like.




FIG. 6. The subject with the most positive scores. These RPV patches are comprised primarily of normal breast tissue and contain few tumor cells. The eosinophilic, fibrillar (light pink, stringy) tissue is collagenous stroma which is found in normal breast tissue while the white circles surrounded by purple dots are neoplastic ducts. This tumor is genetically luminal A.

prognosis and may be associated with response to immunomodulatory therapy (Wein et al. (2017)).


The second visual pattern in the negative end of the first joint component is dense, *high nuclear grade tumor cells* and is illustrated in Figure 5. Nuclear grade describes how abnormal the tumor cells look: “low grade” means the tumor cells look similar to regular cells (“well-differentiated”), and “high grade” means the tumor cells look markedly abnormal (“poorly-differentiated”), for example, are enlarged and irregularly shaped (Rosen (2001)).

On the positive end of the first joint component, the pathology review shows subjects whose cores contain mostly normal breast tissue, that is, little tumor tissue. This pattern is illustrated by Figure 6 which shows the subject with the most positive scores for the first joint component. These patches contain few tumor cells and are mostly normal breast structures such as *collagenous stroma* (the light pink, stringy tissue) and ducts.

The first joint component is related to histopathological features including *tumor grade* and *histological type* (ductal vs. lobular). For example, Figure 7a shows that high *grade* tumors cluster on the negative end of the first joint component while low grade tumors cluster on the positive end (AUC = 0.945). Tumor grade incorporates cellular differentiation and other architectural features as an indicator of aggressiveness (Elston and Ellis (2002)). This first component is also statistically significantly related to histological type with *ductal* on the negative end and *lobular* on the positive end (AUC = 0.785).

From the genetics perspective the first joint component strongly tracks the *proliferation score* as well as the contrast between basal-like vs. luminal A tumors. Figure 8 shows the PAM50 joint loadings vector for the first component. Several of the top negative genes (e.g., CCNB1, CENPF, MYC, MKI67) are associated with high tumor cell proliferation and tend to have low expression levels in normal breast tissue. Several of the top positive genes (e.g., MLPH, MMP11) tend to have high expression levels in normal breast tissue. Note that FOXC1 is highly expressed in both basal-like and normal-like breast myoepithelium.

Figure 7b shows that basal-like tumors cluster on the negative end of the first joint component, molecular HER2 and luminal B cluster in the middle while luminal A and normal tumors cluster on the positive end. Note the AUC score for basal-like vs. luminal A is 0.984



(a) Tumors on the negative end tend to be high grade while those on positive end tend to be low grade. (b) The basal-like class is on the negative end of the scores, the HER2/luminal B classes are in the middle and the normal/luminal A classes are on the positive end. (c) Proliferation score is strongly associated with the common normalized scores with a Pearson correlation of -0.82.

FIG. 7. Joint component 1 scores vs. PAM50 molecular subtype (7b), central grade (7a) and proliferation score (7c). The x-axis in each plot shows the subjects' common normalized scores. These figures tell a complementary story. For example, basal-like tumors tend to be more aggressive, high grade and have a high proliferation score.

which is quite high. Luminal B and molecular HER2 are separated from basal-like (AUCs of 0.886 and 0.876). The separation indicates that this joint component is distinguishing more subtle histopathological and molecular features beyond proliferation and cellularity. Figure 7c shows a strong, negative correlation between the first joint component scores and the proliferation score which is a genetic measure indicative of how fast tumor cells grow (Whitfield et al. (2002)).

Strikingly, the first joint component almost perfectly separates ROR-PT which is a combined genetic and pathology based risk of recurrence score (Parker et al. (2009)). Patients with high ROR-PT are clustered on the negative end while patients with a low ROR-PT are clustered on the positive end with an AUC of 0.999.

In addition to genetic phenotypes measured by RNA expression data, as just discussed, we also have immunohistochemistry (IHC) data, a surrogate measure of RNA subtypes and the most common way of classifying tumors in a clinical setting. From the IHC perspective the first joint component is strongly related to ER status and weakly related to clinical HER2




FIG. 8. PAM50 loadings vector from joint component 1. Several of the top negative genes are associated with high tumor cellularity and, typically, have low expression levels in normal breast tissue (consistent with Figures 5 and 4). Several of the top positive genes tend to have high expression levels in normal breast tissue (consistent with Figure 6).

status (see Table 2). Clinical ER negative tumors cluster on the negative end of this component with an AUC of 0.883.

In this first joint component the pathology and genetics tell complementary stories that are familiar to breast cancer experts. The data raise the possibility that this joint component separates tumors based on one or more histologic features associated with tumor grade. These features could include aspects of nuclear atypia (i.e., increased nuclear size, irregular shape, altered chromatin pattern, multiple nucleoli) which are reflected in the nuclear grade. Tumors with a high combined histologic grade also tend to be more cellular and show less tubule or gland formation, as compared to low-grade tumors.

From the genetics perspective, basal-like tumors are on the negative end, molecular HER2/luminal B tumors are in the middle, and luminal A/normal-like tumors are on the positive end. The joint scores are strongly negatively correlated with the proliferation score. The negative genes in Figure 8 are predominantly proliferation regulated genes; however, we note several of the positive genes are often considered basal-specific genes. These genes are also expressed in normal myoepithelium and are representative of the normal ducts still observed within slides of the low grade tumors (Livasy et al. (2006), Heng et al. (2017)).

Aggressive tumors tend to have high tumor cellularity and little benign tissue. In less aggressive tumors there is typically more normal breast tissue. Basal-like tumors tend to be more aggressive and are generally associated with high tumor grade, ER negativity, ductal histology and high proliferation score (Livasy et al. (2006), Troester et al. (2017), Williams et al. (2019)). Luminal B and molecular HER2 tumors tend to be moderately aggressive. Luminal A and normal-like tumors are less aggressive, and it is known these tend to be low grade.

It is promising that this mode of variation turned up in the first joint component. These connections between the underlying genetic drivers and the pathological impressions have both geneticists and pathologists excited about the potential of AJIVE to quantitatively integrate these different aspects of cancer.

5.1.2. *Joint component 2.* From the pathology perspective, the tumors on the negative end of joint component 2 show mostly collagenous stroma surrounded by moderate nuclear grade tumor cells. Figure 9 shows the mini-RPVs of two subjects from the negative end of the second joint component. The positive end of this component was not homogeneous (Table 1).

From the genetics perspective, the negative end of joint component 2 picks out the luminal B subtype (Figure 10b). Looking at the PAM50 loadings vector, ESR1, SLC39A6 are the two most negative genes in the PAM50 loadings (Figure 10a) and are known to be high in clinically ER+ cancers (Parker et al. (2009)). The luminal B observations cluster on the negative end of this direction and are statistically significantly separated from the other PAM50


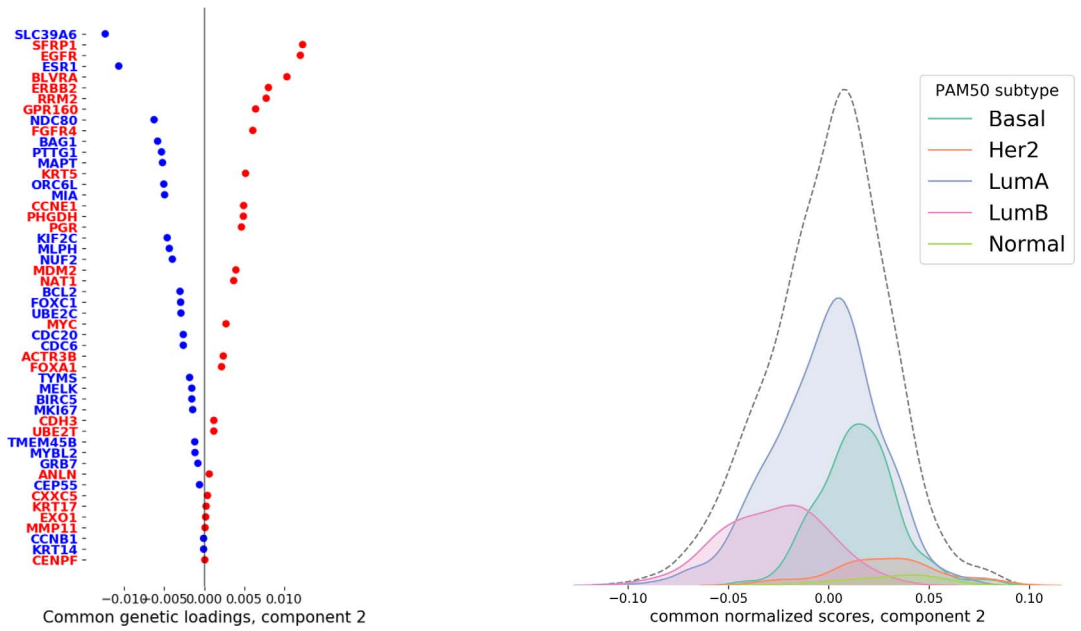


FIG. 9. Two tumors from the negative end of joint component 2. Both RPVs show collagenous stroma (light pink, stringy tissue) surrounded by moderate nuclear grade tumor cells and moderate tumor cellularity. The tumor in (a) is genetically luminal B, and the tumor in (b) is luminal A.



(a) Joint component 2, PAM50 loadings. (b) Joint component 2 scores conditioned on PAM50 subtype.

FIG. 10. The luminal B and luminal A classes are associated with the negative end of joint component 2. Not much is known about this histology of luminal cancers.


subtypes with AUC scores of: basal = 0.905, HER2 = 0.933, luminal A = 0.760, normal = 0.950 (Figure 10b).

From the immunohistochemical perspective the second joint component is moderately related to ER status while weakly related to clinical HER2 status (Table 2). Clinical ER positive tumors cluster on the negative end of this component with an AUC of 0.752.

The pathology perspective of this second joint component appears to pick up on morphological features of luminal B tumors, that is, intratumoral channels of stromal cells which are surrounded by moderate nuclear grade cancer cells. To our knowledge, little is known about the histological features of luminal B tumors.

Pathologists do not currently use stromal features in the diagnosis and classification of tumors. However, tumor stroma and microenvironment (Eiro et al. (2019)) and the stromal features of benign and tumor-adjacent normal tissue (Román-Pérez et al. (2012), Chollet-Hinton et al. (2018)) are areas of active investigation. Interestingly, Beck et al. (2011) used image analysis approaches to demonstrate connections between certain stroma morphological features and patient survival. Recent studies using CNNs have shown that breast biopsies may be accurately classified as malignant solely based on stromal features (Bejnordi et al. (2018)).

5.1.3. Joint component 3. From the genetics perspective, the negative end of joint component 3 picks out molecular HER2. The HER2 observations are separated from the other PAM50 subtypes with AUC scores of: basal = 0.947, luminal A = 0.940, luminal B = 0.833, normal = 0.950. Interestingly, ERBB2 and EGFR have large negative values in the joint loadings vector while GRB7, which is on the same *amplicon* as ERBB2, is almost 0 (Figure 11a). The negative end of this component is also moderately related to clinical HER2 status with an AUC of 0.777 (Table 2). This component is identifying not only clinical HER2 samples (as determined by IHC staining) but more strongly the molecular HER2 subtype of samples (as determined by gene expression). Previous work (Network et al. (2012)) has shown both gene expression and protein and phosphoprotein levels of ERBB2 and EGFR are



(a) Joint component 3, PAM50 loadings. (b) Joint component 3 scores conditioned on PAM50 subtype.

FIG. 11. *The molecular HER2 subtype are associated with the negative end of joint component 3.*

significantly enriched in clinically HER2 samples that are also the molecular HER2 subtype compared to clinical HER2 samples that are Luminal subtypes. This is consistent with the separations we see in joint component 3.

From the pathology perspective, the negative end joint component 3 again shows collagenous stroma, but this time surrounded by high nuclear grade tumor cells (Figure 12). Recall joint component 2 was similar but with moderate grade tumor cells. This third joint component appears to be picking up on morphological features of molecular HER2 tumors. Similar to joint component 2, it is interesting that the stroma appear to play an important role in this component.

5.2. Image individual information. As mentioned in Section 2.1, the PAM50 genes were selected to emphasize genes expressed in tumor epithelium, not genes highly expressed in tumor microenvironment features such as fat cells, collagenous stroma and, in some cases, mucin. Several of these microenvironment features have clear visual signals (e.g., high fat content images have round, clear adipose cells) and show up prominently in the AJIVE individual components.




FIG. 12. *Tumors from the negative end of joint component 3 showing tumor cells surrounded collagenous stroma.*




FIG. 13. Two tumors from the positive end of image individual component 1 showing high fat content.

5.2.1. *Image individual component 1.* All of the images on the positive end of the first image individual component show a very clear theme of tumors with high fat content (Figure 13). High fat content is a strong visual signal, so it makes sense that it shows up as an early individual mode of variation for image data. The negative end of the first image individual component shows tumors with low tumor cellularity and low/moderate grade nuclei (Table 3 and Figure 14).

5.2.2. *Image individual component 2.* The negative end of image individual component 2 clearly picks out *mucinous carcinoma* tumors (Figure 15). Mucinous carcinomas are characterized by tumor cells floating in pools of mucin. These cancers present a very clear visual pattern of dark purple tumor cells surrounded by wispy looking mucin. Mucinous carcinoma is a rare histological subtype which the PAM50 genes Perou et al. (2000) are not designed to identify.

Mucinous carcinomas are typically low-grade, hormone receptor-positive, have a good prognosis and appear to be genetically different from invasive ductal carcinomas of no special type (Di Saverio, Gutierrez and Avisar (2008), Diab et al. (1999), Weigelt et al. (2009), Lacroix-Triki et al. (2010)). Mucinous carcinomas are usually genetically luminal-type (typically luminal A) (Caldarella et al. (2013), Colleoni et al. (2011), Weigelt et al. (2009)). All of the top 15 tumors on the negative end of this component are genetically luminal (12 are luminal A, and three are luminal B). Interestingly, neither the luminal A nor B classes are strongly associated with the individual scores for this component overall; none of the difference in distribution tests (see Supplement A) for luminal A vs. another class were statistically significant (similarly for luminal B). This is consistent with variation appearing in an image individual component.

The positive end of individual component 2 picks out images with moderate cellularity and collagenous stroma surrounded by moderate nuclear grade tumor cells (Figure 16).

TABLE 3

A pathologist's observations of first three image individual components from RPVs of 15 most extreme subjects on either end of the component

Com- ponent	End	Homo- geneous	Tumor cellularity	Tubule formation	Nuclear grade	Adipocytic stroma	Collagenous stroma	Lym- phocytes	Necrosis
1	positive	yes	low	focal	1, 2	yes	limited	few	no
	negative	yes	variable	no	2	no	yes	few	some
2	positive	yes	moderate	no	2, 3	focal	yes	yes	no
	negative	yes	low	no	1	focal	limited	no	no
3	positive	yes	low	yes	2	focal	yes	few	no
	negative	yes	low	no	2	yes	no	no	no




FIG. 14. *Two tumors from the negative end of image individual component 1 showing moderate nuclear grade, variable tumor cellularity, collagenous stroma.*

The positive and negative ends of image individual component 2 show contrasting histologic features. The patches from the tumors on the positive end are entirely filled with a combination of tumor cell aggregates separated by areas of dense collagenous stroma. Adipocytic stroma is absent, and the only optically clear space is in areas of retraction artifacts where tumor cell groups appear to be pulled away from adjacent stroma (a known artifact of histologic preparation in some invasive tumors). The patches from the negative end show extracellular mucin from mucinous carcinomas with low or no tumor cellularity and just a few wispy bands of stromal collagen. The contrasting histology raises the possibility that this component may separate tumors based on one or more of the following features: tumor cellularity, tumor grade, extracellular stromal composition.

5.2.3. Image individual component 3. The negative end of image individual component 3 picks up on tumors whose patches contain a large amount of optically clear space. This includes tumors with high fat content (Figure 17b), where the cells discohesive (Figure 17c) and disrupted tissue sections (Figure 17a). Recall (Section 2.2) that patches with too much background (over 90%) are removed. Therefore, white space surrounding the tumors and large white spaces in the core are unlikely to influence the amount of white space in the patches representing the image. Some of the features seen in the images in Figures 17, 17a and 17c are likely related to technical variation in the tumor fixation/preservation and the quality of the histologic preparation. While the high fat content pattern seen in the positive end of the first image individual component is similar to this component (i.e., it picks up on large amounts of white space), the first component uniformly contains high fat content images in the top 15 images which is unlike this third component.

The positive end of this third component picks on images with a large amount of dense collagenous stroma (Figure 18). This pattern is very clear in all 15 of the most positive subjects'




FIG. 15. *Two tumors from the negative end of image individual component 2 both displaying mucinous carcinomas.*




FIG. 16. Two tumors from the positive end of image individual component 2 both showing collagenous stroma, moderate to high nuclear grade, moderate tumor cellularity.

representative images views (see Supplement A). These tumors have lower tumor cellularity, moderate nuclear grade and have a moderate number of lymphocytes. Similar to the amount of white space, the dense collagenous stroma is a clear visual pattern.

5.3. Genetic individual information. Figure 19a shows the PAM50 loadings vector of the first genetic individual component. This component picks up on overall gene expression levels which is a common source of technical variation. Both the second and third genetic individual components show connections to the PAM50 subtypes based on the clinical data comparisons given in Supplement A albeit with weaker separations than the joint components.

The second genetic individual component identifies additional information, which varies between luminal A and normal, that is not dependent on cell proliferation and seems to be more related to features such as estrogen receptor signaling and keratin expression status. The scores for this second component separate normal-like from luminal A with an AUC of 0.801.

Figure 19b shows a scatter plot of the loadings vector from genetic individual component 2 compared to the normal-luminal A mean difference direction.¹³ Several of the genes on the top left of 19b (*ESR1*, *FOXA1*, *PGR*) are all part of the estrogen signaling pathway (Oh et al. (2006)). Several of the genes in the middle (*CCNB1*, *MYC*, *MKI67*, *TYMS*, *MYBL2*, *CCNE1*) are related to proliferation, suggesting this component is unrelated to proliferation. Several of the genes in the bottom right (*KRT5*, *KRT14*, *KRT17*) are characteristic of normal myoepithelium as well as basal-like breast cancer (Lazard et al. (1993)).

6. Conclusion. This paper develops methods that, using deep learning and AJIVE, give interpretable, simultaneous image and genetic results. For example, the first AJIVE joint component discovers the well-known association between basal-like tumors (genetic) and




FIG. 17. Three tumors from the negative end of image individual component 3 all showing the clear visual pattern of a large amount of optically clear space.

¹³The genes were first scaled by their standard deviation, so this is the naive Bayes classification direction.





FIG. 18. Two tumors from the positive end of image individual component 3. Both show a visually clear pattern of dense collagenous stroma, low tumor cellularity, moderate nuclear grade, some lymphocytes.

high-grade tumors (histology). The second AJIVE joint component finds novel histological features of Luminal B tumors. Inferential and exploratory analysis leveraging deep learning is a promising area which presents many interesting, open questions—some of which are discussed below. These analytical tools enable simultaneous engagement from both the pathology and genetic communities which is critical for the fundamental biomedical interpretations.

Future research should evaluate whether the features learned in this paper can be reproducibly identified by pathologists and/or automated computer vision systems as well as whether these features can be validated in external test sets.

Scaling histological image analysis pipelines to gigapixel *whole slide images* (WSI) is an important future direction. In clinical practice, pathologists use WSIs; the core images used in this paper require additional preparation, are typically only available in some research settings and may ignore important tumor information (e.g., spatial heterogeneity across the



(a) Genetic individual component 1 loadings shows common genetic technical mode of variation not expected to be associated with the images.

(b) Genetic individual component 2 loadings vector compared to normal-luminal A mean difference direction. Luminal A is on the left (negative) while normal-like is on the right (positive).

FIG. 19. Genetic individual components show technical variation as well as additional PAM50 subtype information.

tumor, particularly histological patterns not observed in the sampled region). Analysis of WSIs presents computational challenges as these images are orders of magnitude larger than the core images.

6.1. *Patch representation.* Patch based approaches have shown promise for predictive tasks using deep learning Ilse, Tomczak and Welling (2018). The patch based approach taken in this paper was selected because it: (i) will scale to whole slide images, (ii) can identify localized image information, for example, with the RVPs and (iii) creates a smaller feature set.¹⁴ The approach of averaging patch features ignores some within image heterogeneity. For image-only analysis, methods such as Backenroth et al. (2018), Bishop and Tipping (1998) may be able to capture additional within-image heterogeneity. In the context of multiview data, additional methodology needs to be developed to account for grouped observations (e.g., Pourzanjani et al. (2017)).

6.2. *Transfer learning.* Training a neural network can be time and resource intensive. Furthermore, CNNs often require a large amount of training data to be fit effectively. Transfer learning allows the data analyst to use more powerful neural networks with less data and less time spent tuning CNN parameters (Yosinski et al. (2014), Sharif Razavian et al. (2014)). First, a CNN is trained to solve a different predictive task on a large, external dataset—typically, the famous ImageNet classification task (Deng et al. (2009)). Then, the pretrained network parameters may be *fine-tuned* on the dataset of interest to solve the predictive problem of interest.

The setting of this paper is a bit different. First, we are doing exploratory analysis, not predictive analysis. Second, while we do have labels which could be used to fine-tune the CNN (e.g., the PAM50 subtypes), we do not want to use these labels because then the network would be aware of information which we might want to (re)discover and/or validate in the following analysis. This leaves us with a couple of options to still use transfer learning including: training an unsupervised algorithm (e.g., autoencoders (Kingma and Welling (2013)), *generative adversarial networks* (Goodfellow et al. (2014)) or selfsupervised learning algorithms (Oord, Li and Vinyals (2018), Lu et al. (2019))) or not doing any fine tuning. Ash et al. (2018) trains autoencoder which has some disadvantages: a CNN is required to be trained which is time/resource consuming, a number of new hyperparameters are introduced into the problem and either the data are used twice or external data are needed. We explore the latter option and show, perhaps surprisingly, that pretrained CNN features with no fine tuning are able to capture complex visual signals in a domain vastly different from the one they were originally trained on.

Even in the context of transfer learning, there are many choices to be made about how to extract neural network features from an image, including network architecture, layer (or layers) of the network and feature aggregation (e.g., spatial mean pooling discussed in Section 2.3). For predictive modeling these hyperparameters can be set using an error metric and methods such as cross-validation; however, as discussed in the above paragraph, we do not have such error metrics readily available to guide hyperparameter choices. Preliminary sensitivity analysis showed that the results of our analysis are not particularly sensitive to mild differences in architecture choices. Better methods to select these CNN hyperparameters is an open area of research.

Acknowledgments. We thank the Carolina Breast Cancer Study participants and staff. We also want to acknowledge Robert C. Millikan, founder of the Carolina Breast Cancer Study Phase 3.

¹⁴Passing the full core images through the CNN resulted in 10^6 features.

Funding. Research reported in this publication was supported by a Specialized Program of Research Excellence (SPORE) in breast cancer (P50 CA058223), an award from the Susan G. Komen Foundation (OGUNC1202), the North Carolina University Cancer Research Fund and a Cancer Center Support Grant (P30 CA016086). Iain Carmichael and J. S. Marron were partially supported by NSF Grant IIS-1633074, BIG DATA 2016–2019. Iain Carmichael is currently supported by NSF MSPRF DMS-1902440. Katherine Hoadley was supported by Komen Career Catalyst Grant (CCR16376756).

SUPPLEMENTARY MATERIAL

Supplement A: Additional analysis details and and important biological background (DOI: [10.1214/20-AOAS1433SUPPA](https://doi.org/10.1214/20-AOAS1433SUPPA); .pdf). Discussion of important tissue structures discussed in the results, AJIVE diagnostic plot and clinical data comparison methods (e.g. multiple testing control).

Supplement B: Supplementary visualizations (DOI: [10.1214/20-AOAS1433SUPPB](https://doi.org/10.1214/20-AOAS1433SUPPB); .zip). Additional figures that include all representative patch views shown to pathologists, all AJIVE genetic loadings vectors and all clinical data comparisons. See the file readme.txt for details.

Supplement C: Code to reproduce the analysis (DOI: [10.1214/20-AOAS1433SUPPC](https://doi.org/10.1214/20-AOAS1433SUPPC); .zip). A Python package and scripts that run the analysis.

REFERENCES

- ADEBAYO, J., GILMER, J., MUELLY, M., GOODFELLOW, I., HARDT, M. and KIM, B. (2018). Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems* 9505–9515.
- AEFFNER, F., ZARELLA, M. D., BUCHBINDER, N., BUI, M. M., GOODMAN, M. R., HARTMAN, D. J., LUJAN, G. M., MOLANI, M. A., PARWANI, A. V. et al. (2019). Introduction to digital image analysis in whole-slide imaging: A white paper from the digital pathology association. *J. Pathol. Inform.* **10**.
- ALLOTT, E. H., GERADTS, J., COHEN, S. M., KHOURY, T., ZIRPOLI, G. R., BSHARA, W., DAVIS, W., OMILIAN, A., NAIR, P. et al. (2018). Frequency of breast cancer subtypes among African American women in the AMBER consortium. *Breast Cancer Res.* **20** 12.
- ASH, J., DARNELL, G., MUNRO, D. and ENGELHARDT, B. (2018). Joint analysis of gene expression levels and histological images identifies genes associated with tissue morphology. *BioRxiv* 458711.
- BACKENROTH, D., GOLDSMITH, J., HARRAN, M. D., CORTES, J. C., KRAKAUER, J. W. and KITAGO, T. (2018). Modeling motor learning using heteroscedastic functional principal components analysis. *J. Amer. Statist. Assoc.* **113** 1003–1015. MR3862335 <https://doi.org/10.1080/01621459.2017.1379403>
- BECK, A. H., SANGOI, A. R., LEUNG, S., MARINELLI, R. J., NIELSEN, T. O., VAN DE VIJVER, M. J., WEST, R. B., VAN DE RIJN, M. and KOLLER, D. (2011). Systematic analysis of breast cancer morphology uncovers stromal features associated with survival. *Sci. Transl. Med.* **3** 108ra113. <https://doi.org/10.1126/scitranslmed.3002564>
- BEJNORDI, B. E., MULLOOLY, M., PFEIFFER, R. M., FAN, S., VACEK, P. M., WEAVER, D. L., HERSCORN, S., BRINTON, L. A., VAN GINNEKEN, B. et al. (2018). Using deep convolutional neural networks to identify and classify tumor-associated stroma in diagnostic breast biopsies. *Mod. Pathol.* **31** 1502.
- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. MR1325392
- BISHOP, C. M. and TIPPING, M. E. (1998). A hierarchical latent variable model for data visualization. *IEEE Trans. Pattern Anal. Mach. Intell.* **20** 281–293.
- CALDARELLA, A., BUZZONI, C., CROCETTI, E., BIANCHI, S., VEZZOSI, V., APICELLA, P., BIANCALANI, M., GIANNINI, A., URSO, C. et al. (2013). Invasive breast cancer: A significant correlation between histological types and molecular subgroups. *J. Cancer Res. Clin. Oncol.* **139** 617–623.
- CAREY, L. A., PEROU, C. M., LIVASY, C. A., DRESSLER, L. G., COWAN, D., CONWAY, K., KARACA, G., TROESTER, M. A., TSE, C. K. et al. (2006). Race, breast cancer subtypes, and survival in the Carolina breast cancer study. *JAMA* **295** 2492–2502.
- CARMICHAEL, I. (2020). pyjive: A Python package for AJIVE. Available at https://github.com/idc9/py_jive. <https://doi.org/10.5281/zenodo.4091752>

- CARMICHAEL, I., CALHOUN, B. C., HOADLEY, K. A., TROESTER, M. A., GERADTS, J., COUTURE, H. D., OLSSON, L., PEROU, C. M., NIETHAMMER, M., HANNIG, J. and MARRON, J. S. (2021). Supplement to “Joint and individual analysis of breast cancer histologic images and genomic covariates.” <https://doi.org/10.1214/20-AOAS1433SUPPA>, <https://doi.org/10.1214/20-AOAS1433SUPPB>, <https://doi.org/10.1214/20-AOAS1433SUPPC>.
- CHEN, C., LI, O., TAO, C., BARNETT, A. J., SU, J. and RUDIN, C. (2018a). This looks like that: Deep learning for interpretable image recognition. Preprint. Available at [arXiv:1806.10574](https://arxiv.org/abs/1806.10574).
- CHEN, P.-H. C., GADEPALLI, K., MACDONALD, R., LIU, Y., NAGPAL, K., KOHLBERGER, T., DEAN, J., CORRADO, G. S., HIPPI, J. D. et al. (2018b). Microscope 2.0: An augmented reality microscope with real-time artificial intelligence integration. Preprint. Available at [arXiv:1812.00825](https://arxiv.org/abs/1812.00825).
- CHEN, R. J., LU, M. Y., WANG, J., WILLIAMSON, D. F., RODIG, S. J., LINDEMAN, N. I. and MAHMOOD, F. (2019). Pathomic fusion: An integrated framework for fusing histopathology and genomic features for cancer diagnosis and prognosis. Preprint. Available at [arXiv:1912.08937](https://arxiv.org/abs/1912.08937).
- CHOLLET-HINTON, L., PUVANESARAJAH, S., SANDHU, R., KIRK, E. L., MIDKIFF, B. R., GHOSH, K., BRANDT, K. R., SCOTT, C. G., GIERACH, G. L. et al. (2018). Stroma modifies relationships between risk factor exposure and age-related epithelial involution in benign breast. *Mod. Pathol.* **31** 1085.
- COLLEONI, M., ROTMENSZ, N., MAISONNEUVE, P., MASTROPASQUA, M. G., LUINI, A., VERONESI, P., INTRA, M., MONTAGNA, E., CANCELLO, G. et al. (2011). Outcome of special types of luminal breast cancer. *Ann. Oncol.* **23** 1428–1436.
- COOPER, L. A. D., KONG, J., GUTMAN, D. A., DUNN, W. D., NALISNIK, M. and BRAT, D. J. (2015). Novel genotype-phenotype associations in human cancers enabled by advanced molecular platforms and computational analysis of whole slide images. *Lab. Invest.* **95** 366–376. <https://doi.org/10.1038/labinvest.2014.153>
- COUDRAY, N., OCAMPO, P. S., SAKELLAROPOULOS, T., NARULA, N., SNUDERL, M., FENYÖ, D., MOREIRA, A. L., RAZAVIAN, N. and TSIRIGOS, A. (2018). Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nat. Med.* **24** 1559–1567. <https://doi.org/10.1038/s41591-018-0177-5>
- COUTURE, H. D., WILLIAMS, L. A., GERADTS, J., NYANTE, S. J., BUTLER, E. N., MARRON, J. S., PEROU, C. M., TROESTER, M. A. and NIETHAMMER, M. (2018). Image analysis with deep learning to predict breast cancer grade, ER status, histologic subtype, and intrinsic subtype. *npj Breast Cancer* **4** 30.
- DENG, J., DONG, W., SOCHER, R., LI, L.-J., LI, K. and FEI-FEI, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition* 248–255. IEEE.
- DI SAVERIO, S., GUTIERREZ, J. and AVISAR, E. (2008). A retrospective review with long term follow up of 11,400 cases of pure mucinous breast carcinoma. *Breast Cancer Res. Treat.* **111** 541–547.
- DIAB, S. G., CLARK, G. M., OSBORNE, C. K., LIBBY, A., ALLRED, D. C. and ELLEDGE, R. M. (1999). Tumor characteristics and clinical outcome of tubular and mucinous breast carcinomas. *J. Clin. Oncol.* **17** 1442–1448. <https://doi.org/10.1200/JCO.1999.17.5.1442>
- DRAPER, B., KIRBY, M., MARKS, J., MARRINAN, T. and PETERSON, C. (2014). A flag representation for finite collections of subspaces of mixed dimensions. *Linear Algebra Appl.* **451** 15–32. MR3198905 <https://doi.org/10.1016/j.laa.2014.03.022>
- EIRO, N., GONZALEZ, L. O., FRAILE, M., CID, S., SCHNEIDER, J. and VIZOSO, F. J. (2019). Breast cancer tumor stroma: Cellular components, phenotypic heterogeneity, intercellular communication, prognostic implications and therapeutic opportunities. *Cancers* **11** 664.
- ELMORE, J. G., LONGTON, G. M., CARNEY, P. A., GELLER, B. M., ONEGA, T., TOSTESON, A. N., NELSON, H. D., PEPE, M. S., ALLISON, K. H. et al. (2015). Diagnostic concordance among pathologists interpreting breast biopsy specimens. *JAMA* **313** 1122–1132.
- ELSTON, C. W. and ELLIS, I. O. (2002). Pathological prognostic factors in breast cancer. I. The value of histological grade in breast cancer: Experience from a large study with long-term follow-up. CW Elston & IO Ellis. *Histopathology* 1991; 19; 403–410: AUTHOR COMMENTARY. *Histopathology* **41** 151–151.
- FENG, Q., JIANG, M., HANNIG, J. and MARRON, J. S. (2018). Angle-based joint and individual variation explained. *J. Multivariate Anal.* **166** 241–265. MR3799646 <https://doi.org/10.1016/j.jmva.2018.03.008>
- GAYNANOVA, I. and LI, G. (2019). Structural learning and integrative decomposition of multi-view data. *Biometrics* **75** 1121–1132. MR4041816 <https://doi.org/10.1111/biom.13108>
- GOODFELLOW, I., POUGET-ABADIE, J., MIRZA, M., XU, B., WARDE-FARLEY, D., OZAIR, S., COURVILLE, A. and BENGIO, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems* 2672–2680.
- HENG, Y. J., LESTER, S. C., TSE, G. M., FACTOR, R. E., ALLISON, K. H., COLLINS, L. C., CHEN, Y.-Y., JENSEN, K. C., JOHNSON, N. B. et al. (2017). The molecular basis of breast cancer pathological phenotypes. *J. Pathol.* **241** 375–391.

- HOLZINGER, A., LANGS, G., DENK, H., ZATLOUKAL, K. and MÜLLER, H. (2019). Causability and explainability of artificial intelligence in medicine. *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.* e1312.
- HOTELLING, H. (1936). Relation between two sets of variates. *Biometrika*.
- HUNTER, J. D. (2007). Matplotlib: A 2D graphics environment. *Comput. Sci. Eng.* **9** 90.
- ILSE, M., TOMCZAK, J. M. and WELLING, M. (2018). Attention-based deep multiple instance learning. Preprint. Available at [arXiv:1802.04712](https://arxiv.org/abs/1802.04712).
- JIMÉNEZ, G. and RACOCEANU, D. (2019). Deep learning for semantic segmentation versus classification in computational pathology: Application to mitosis analysis in breast cancer grading. *Front. Bioeng. Biotechnol.* **7** 145.
- JOHNSTONE, I. M. (2008). Multivariate analysis and Jacobi ensembles: Largest eigenvalue, Tracy–Widom limits and rates of convergence. *Ann. Statist.* **36** 2638–2716. [MR2485010 https://doi.org/10.1214/08-AOS605](https://doi.org/10.1214/08-AOS605)
- JONES, E., OLIPHANT, T. and PETERSON, P. (2014). SciPy: Open source scientific tools for Python.
- KETTENRING, J. R. (1971). Canonical analysis of several sets of variables. *Biometrika* **58** 433–451. [MR0341750 https://doi.org/10.1093/biomet/58.3.433](https://doi.org/10.1093/biomet/58.3.433)
- KIM, B., WATTENBERG, M., GILMER, J., CAI, C., WEXLER, J., VIEGAS, F. et al. (2018). Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). In *International Conference on Machine Learning 2673–2682*.
- KINGMA, D. P. and WELLING, M. (2013). Auto-encoding variational Bayes. Preprint. Available at [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- KOMURA, D. and ISHIKAWA, S. (2018). Machine learning methods for histopathological image analysis. *Comput. Struct. Biotechnol. J.* **16** 34–42. <https://doi.org/10.1016/j.csbj.2018.01.001>
- LACROIX-TRIKI, M., SUAREZ, P. H., MACKAY, A., LAMBROS, M. B., NATRAJAN, R., SAVAGE, K., GEYER, F. C., WEIGELT, B., ASHWORTH, A. et al. (2010). Mucinous carcinoma of the breast is genomically distinct from invasive ductal carcinomas of no special type. *J. Pathol.* **222** 282–298.
- LAZARD, D., SASTRE, X., FRID, M. G., GLUKHOVA, M. A., THIERY, J.-P. and KOTELIANSKY, V. E. (1993). Expression of smooth muscle-specific proteins in myoepithelium and stromal myofibroblasts of normal and malignant human breast tissue. *Proc. Natl. Acad. Sci. USA* **90** 999–1003.
- LIU, Y., GADEPALLI, K., NOROUZI, M., DAHL, G. E., KOHLBERGER, T., BOYKO, A., VENUGOPALAN, S., TIMOFEEV, A., NELSON, P. Q. et al. (2017). Detecting cancer metastases on gigapixel pathology images. Preprint. Available at [arXiv:1703.02442](https://arxiv.org/abs/1703.02442).
- LIU, Y., KOHLBERGER, T., NOROUZI, M., DAHL, G. E., SMITH, J. L., MOHTASHAMIAN, A., OLSON, N., PENG, L. H., HIPPI, J. D. et al. (2018). Artificial intelligence-based breast cancer nodal metastasis detection: Insights into the black box for pathologists. *Arch. Pathol. Lab. Med.*
- LIVASY, C. A., KARACA, G., NANDA, R., TRETIKOVA, M. S., OLOPADE, O. I., MOORE, D. T. and PEROU, C. M. (2006). Phenotypic evaluation of the basal-like subtype of invasive breast carcinoma. *Mod. Pathol.* **19** 264.
- LOCK, E. F., HOADLEY, K. A., MARRON, J. S. and NOBEL, A. B. (2013). Joint and individual variation explained (JIVE) for integrated analysis of multiple data types. *Ann. Appl. Stat.* **7** 523–542. [MR3086429 https://doi.org/10.1214/12-AOAS597](https://doi.org/10.1214/12-AOAS597)
- LU, M. Y., CHEN, R. J., WANG, J., DILLON, D. and MAHMOOD, F. (2019). Semi-supervised histology classification using deep multiple instance learning and contrastive predictive coding. Preprint. Available at [arXiv:1910.10825](https://arxiv.org/abs/1910.10825).
- MACENKO, M., NIETHAMMER, M., MARRON, J. S., BORLAND, D., WOOSLEY, J. T., GUAN, X., SCHMITT, C. and THOMAS, N. E. (2009). A method for normalizing histology slides for quantitative analysis. In *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro* 1107–1110. IEEE.
- MAHMOOD, F., YANG, Z., ASHLEY, T. and DURR, N. J. (2018). Multimodal densenet. Preprint. Available at [arXiv:1811.07407](https://arxiv.org/abs/1811.07407).
- MAHMOOD, F., BORDERS, D., CHEN, R., MCKAY, G. N., SALIMIAN, K. J., BARAS, A. and DURR, N. J. (2019). Deep adversarial training for multi-organ nuclei segmentation in histopathology images. *IEEE Trans. Med. Imag.*
- MCKINNEY, W. (2011). Pandas: A foundational Python library for data analysis and statistics. *Python High Perform. Sci. Comput.* **14**.
- MOLNAR, C. et al. (2018). *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*. E-book at <https://christophm.github.io/interpretable-ml-book/>, version dated 10.
- NETWORK, C. G. A. et al. (2012). Comprehensive molecular portraits of human breast tumours. *Nature* **490** 61.
- OH, D. S., TROESTER, M. A., USARY, J., HU, Z., HE, X., FAN, C., WU, J., CAREY, L. A. and PEROU, C. M. (2006). Estrogen-regulated genes predict survival in hormone receptor-positive breast cancers. *J. Clin. Oncol.* **24** 1656–1664.
- OLAH, C., SATYANARAYAN, A., JOHNSON, I., CARTER, S., SCHUBERT, L., YE, K. and MORDVINTSEV, A. (2018). The building blocks of interpretability. *Distill* **3** e10.

- OORD, A. V. D., LI, Y. and VINYALS, O. (2018). Representation learning with contrastive predictive coding. Preprint. Available at [arXiv:1807.03748](https://arxiv.org/abs/1807.03748).
- OTSU, N. (1979). A threshold selection method from gray-level histograms. *IEEE Trans. Syst. Man Cybern.* **9** 62–66.
- PARKER, J. S., MULLINS, M., CHEANG, M. C., LEUNG, S., VODUC, D., VICKERY, T., DAVIES, S., FAURON, C., HE, X. et al. (2009). Supervised risk predictor of breast cancer based on intrinsic subtypes. *J. Clin. Oncol.* **27** 1160.
- PASZKE, A., GROSS, S., CHINTALA, S., CHANAN, G., YANG, E., DEVITO, Z., LIN, Z., DESMAISON, A., ANTIGA, L. et al. (2017). Automatic differentiation in pytorch.
- PEDREGOSA, F., VAROQUAUX, G., GRAMFORT, A. et al. (2011). Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12** 2825–2830. [MR2854348](https://doi.org/10.1162/jmlr.2011.12.1)
- PEROU, C. M., SØRLIE, T., EISEN, M. B., VAN DE RIJN, M., JEFFREY, S. S., REES, C. A., POLLACK, J. R., ROSS, D. T., JOHNSEN, H. et al. (2000). Molecular portraits of human breast tumours. *Nature* **406** 747.
- POURZANJANI, A. A., WU, T. B., JIANG, R. M., COHEN, M. J. and PETZOLD, L. R. (2017). Understanding coagulopathy using multi-view data in the presence of sub-cohorts: A hierarchical subspace approach. In *Machine Learning for Healthcare Conference* 338–351.
- ROMÁN-PÉREZ, E., CASBAS-HERNÁNDEZ, P., PIRONE, J. R., REIN, J., CAREY, L. A., LUBET, R. A., MANI, S. A., AMOS, K. D. and TROESTER, M. A. (2012). Gene expression in extratumoral microenvironment predicts clinical outcome in breast cancer patients. *Breast Cancer Res.* **14** R51. <https://doi.org/10.1186/bcr3152>
- ROSEN, P. P. (2001). *Rosen's Breast Pathology*. Williams & Wilkins, Baltimore.
- SELVARAJU, R. R., COGSWELL, M., DAS, A., VEDANTAM, R., PARIKH, D. and BATRA, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *2017 IEEE International Conference on Computer Vision (ICCV)* 618–626. IEEE.
- SHARIF RAZAVIAN, A., AZIZPOUR, H., SULLIVAN, J. and CARLSSON, S. (2014). CNN features off-the-shelf: An astounding baseline for recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops* 806–813.
- SIMONYAN, K. and ZISSERMAN, A. (2014). Very deep convolutional networks for large-scale image recognition. Preprint. Available at [arXiv:1409.1556](https://arxiv.org/abs/1409.1556).
- SPRINGENBERG, J. T., DOSOVITSKIY, A., BROX, T. and RIEDMILLER, M. (2014). Striving for simplicity: The all convolutional net. Preprint. Available at [arXiv:1412.6806](https://arxiv.org/abs/1412.6806).
- SRIVASTAVA, A., KULKARNI, C., MALLICK, P., HUANG, K. and MACHIRAJU, R. (2018). Building transomics evidence: Using imaging and ‘omics’ to characterize cancer profiles. In *PSB 377–388*. World Scientific, Singapore.
- SUNDARARAJAN, M., TALY, A. and YAN, Q. (2017). Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning, Vol. 70* 3319–3328. [JMLR.org](https://jmlr.org).
- TROESTER, M. A., SUN, X., ALLOTT, E. H., GERADTS, J., COHEN, S. M., TSE, C.-K., KIRK, E. L., THORNE, L. B., MATHEWS, M. et al. (2017). Racial differences in PAM50 subtypes in the Carolina breast cancer study. *J. Natl. Cancer Inst.* **110** 176–182.
- VAN DER WALT, S., COLBERT, S. C. and VAROQUAUX, G. (2011). The NumPy array: A structure for efficient numerical computation. *Comput. Sci. Eng.* **13** 22.
- VAN DER WALT, S., SCHÖNBERGER, J. L., NUNEZ-IGLESIAS, J., BOULOGNE, F., WARNER, J. D., YAGER, N., GOUILLART, E. and YU, T. (2014). scikit-image: Image processing in Python. *PeerJ* **2** e453. <https://doi.org/10.7717/peerj.453>
- VELLIDO, A., MARTÍN-GUERRERO, J. D. and LISBOA, P. J. (2012). Making machine learning models interpretable. In *ESANN 12* 163–172. Citeseer.
- VETA, M., HENG, Y. J., STATHONIKOS, N., BEJNORDI, B. E., BECA, F., WOLLMANN, T., ROHR, K., SHAH, M. A., WANG, D. et al. (2019). Predicting breast tumor proliferation from whole-slide images: The TUPAC16 challenge. *Med. Image Anal.*
- WANG, C., PÉCOT, T., ZYNGER, D. L., MACHIRAJU, R., SHAPIRO, C. L. and HUANG, K. (2013). Identifying survival associated morphological features of triple negative breast cancer using multiple datasets. *J. Am. Med. Inform. Assoc.* **20** 680–687.
- WANG, D., KHOSLA, A., GARGEYA, R., IRSHAD, H. and BECK, A. H. (2016). Deep learning for identifying metastatic breast cancer. Preprint. Available at [arXiv:1606.05718](https://arxiv.org/abs/1606.05718).
- WASKOM, M., BOTVINNIK, O., O’KANE, D., HOBSON, P., OSTBLOM, J., LUKAUSKAS, S., GEMPERLINE, D. C., AUGSPURGER, T., HALCHENKO, Y. et al. (2018). Seaborn (v0.9.0). <https://doi.org/10.5281/zenodo.1313201>
- WEIGELT, B., GEYER, F. C., HORLINGS, H. M., KREIKE, B., HALFWERK, H. and REIS-FILHO, J. S. (2009). Mucinous and neuroendocrine breast carcinomas are transcriptionally distinct from invasive ductal carcinomas of no special type. *Mod. Pathol.* **22** 1401–1414. <https://doi.org/10.1038/modpathol.2009.112>

- WEIN, L., SAVAS, P., LUEN, S. J., VIRASSAMY, B., SALGADO, R. and LOI, S. (2017). Clinical validity and utility of tumor-infiltrating lymphocytes in routine clinical practice for breast cancer patients: Current and future directions. *Front. Oncol.* **7** 156. <https://doi.org/10.3389/fonc.2017.00156>
- WHITFIELD, M. L., SHERLOCK, G., SALDANHA, A. J., MURRAY, J. I., BALL, C. A., ALEXANDER, K. E., MATESE, J. C., PEROU, C. M., HURT, M. M. et al. (2002). Identification of genes periodically expressed in the human cell cycle and their expression in tumors. *Mol. Biol. Cell* **13** 1977–2000.
- WILLIAMS, L. A., HOADLEY, K. A., NICHOLS, H. B., GERADTS, J., PEROU, C. M., LOVE, M. I., OLSHAN, A. F. and TROESTER, M. A. (2019). Differences in race, molecular and tumor characteristics among women diagnosed with invasive ductal and lobular breast carcinomas. *Cancer Causes Control* **30** 31–39. <https://doi.org/10.1007/s10552-018-1121-1>
- WOLD, H. (1985). Partial least squares. In *Encyclopedia of Statistical Sciences*, Vol. 6 (S. Kotz and N. L. Johnson, eds.). Wiley, New York.
- YANG, Z. and MICHAELIDIS, G. (2016). A non-negative matrix factorization method for detecting modules in heterogeneous omics multi-modal data. *Bioinformatics* **32** 1–8. <https://doi.org/10.1093/bioinformatics/btv544>
- YOSINSKI, J., CLUNE, J., BENGIO, Y. and LIPSON, H. (2014). How transferable are features in deep neural networks? In *Advances in Neural Information Processing Systems* 3320–3328.
- ZACK, G. W., ROGERS, W. E. and LATT, S. A. (1977). Automatic measurement of sister chromatid exchange frequency. *J. Histochem. Cytochem.* **25** 741–753.
- ZEILER, M. D. and FERGUS, R. (2014). Visualizing and understanding convolutional networks. In *European Conference on Computer Vision* 818–833. Springer, Berlin.