

- 1]Charlotte M. Jones-Todd
- 2]Peter Caie
- 3]Janine B. Illian
- 4]Ben C. Stevenson
- 5]Anne Savage
- 2]David J. Harrison
- 6]James L. Bown

Abstract

[Abstract]Diagnosis and prognosis of cancer is informed by the architecture inherent in cancer patient tissue sections. This architecture is typically identified by pathologists, yet advances in computational image analysis facilitate quantitative assessment of this structure. In this article we develop a spatial point process approach in order to describe patterns in cell distribution within tissue samples taken from colorectal cancer (CRC) patients. In particular, our approach is centered on the Palm intensity function. This leads to taking an approximate-likelihood technique in fitting point processes models. We consider two Neyman-Scott point processes and a void process, fitting these point process models to the CRC patient data. We find that the parameter estimates of these models may be used to quantify the spatial arrangement of cells. Importantly, we observe characteristic differences in the spatial arrangement of cells between patients who died from CRC and those alive at follow-up.

Identifying prognostic structural features in tissue sections of colon cancer patients using point pattern analysis

[

1 Introduction

A fundamental aspect of cancer patient diagnosis and prognosis concerns the assessment by pathologists of histopathological architectural and morphological properties within patient tissue sections. These sections typically comprise both cancerous (tumour), and non-cancerous (e.g., stroma) tissue structures, with regions of each intermixed in space.[1] Pathologists categorise tumours into stages that are associated with the progression of the cancer and patient outcome. Cancer staging is good at predicting population survival statistics but not as accurate at predicting an individual patient's prognosis.[2] This is due, in part, to complex tumour heterogeneity and a lack of histopathological features within the tissue which can be reliably identified and reproducibly quantified by eye. The pattern of invasive growth in colorectal cancer (CRC) has been previously linked to the level of aggressiveness of the disease and patient prognosis.[3, 4, 5] The phenomenon of tumour budding, where small distinct islands of tumour cells are widely dispersed within the stroma at the invasive front of CRC, has been shown to be prognostically significant.[6, 7]

Advances in computational image analysis provide an opportunity for quantitative and objective assessment of tissue morphology. In previous work,[8] the focus of investigation was on the morphological pattern of the tumour and no consideration was given to the spatial arrangement of cells. In particular, image analysis methodology was found to standardise the quantification of histopathological features (e.g., tumour budding, lymphatic vessel density, and lymphatic vessel invasion) within the invasive front of the tissue section.

The importance of the tissue architecture for tumour grading is recognised where in previous work the spatial patterning of cells was characterised and used as an indicator of patient outcome.[1] This was achieved through considering the locations of cells as a point pattern—a realisation of a spatial point process. In particular, this work compared a first-order statistic, the estimated point process intensity, and two second-order statistics, the pair correlation function and Ripley's K-function, between two patient groups: patients whose cancer had and had not metastasised, a major indicator of patient survival. Results, however, showed no differences between patient groups with respect to first- or second-order spatial statistics and the authors recognised that standard statistical spatial point process methods do not adequately capture the spatial architecture of cancerous tissue.

In reality, cancerous tissue is a result of many complex processes resulting in obscure spatial structures capturing the intermixing of cancerous and non-cancerous cells. The morphology of the tissues reflects some of this complexity and as such we propose that these particular

⁰**Abbreviations:** CRC, colon rectal cancer; NSPP, Neyman-Scott point process

methods were perhaps too simplistic and relied on assumptions (e.g., homogeneity) that do not hold and therefore failed to appropriately capture the spatial structure of the tissue sections. Characterising the spatial structure of cells is nontrivial and requires the use of more complex, non-standard, spatial point process methodology. In the following sections we develop spatial point process methodology that characterises the spatial arrangement of tissue using the interpoint distances between cells to inform consideration of the spatial morphology of CRC tissue.

1.1 Characterising the structure of parent-daughter point processes

In order to present the spatial point process methodology and, in particular, the point process models developed herein we consider two types of points, *parent* and *daughter* points. Note that these do **not** refer to the distinction between tumour and stroma cells; rather parent and daughter points are **abstract** constructs. A classic example of a parent-daughter process is the Neyman-Scott point process (NSPP), which is a *cluster process*. Here, the parent points, themselves generated by a homogeneous Poisson process, randomly generate daughters centered at their unobserved locations.[9] We propose an additional process that relies on unobserved parents, which we call a *void process*. In a void process, observed points are a realisation of a homogeneous point process, but those within some fixed distance of an unobserved parent are deleted. To define a void process, we consider two homogeneous Poisson processes X and P . The realisation of P is a pattern that forms the centroid locations, $p \in P$, of circular voids of radius $R > 0$. Through superimposing P on X then any $x \in X$ that fall within the voids are deleted. The resulting pattern is the observed void point pattern, a realisation of the void process (see Appendix A.1 for further details).

Figure 1 plot *i*) shows a realisation of a void process and plots *ii*) and *iii*) show two realisations of NSPPs—a Thomas and a Matérn process, respectively. The difference between a Thomas and a Matérn process is the dispersion of daughter points around their parents. In the case of a Thomas process the spatial locations of daughters around their parents follows a bivariate normal distribution and in the case of a Matérn process the daughters are uniformly distributed in a circle around their parents.

In all the processes discussed above the parent points are a realisation of a homogeneous Poisson point process with intensity D . Each process is characterised by two further parameters that relate to the daughter points. In the case of the void process we define these as λ , the intensity of daughters prior to deletion, and R , the radius of the voids. Hence, the parameter vector of a void process is given by $\theta = (D, R, \lambda)$.

In the case of the Neyman-Scott point process we define the parameter vector as $\theta = (D, \phi, \gamma)$. The number of daughters sired by a parent is assumed to be Poisson with expectation ϕ . In addition, conditional on their parent’s locations, the daughters are scattered in space according to some distribution with parameter γ . In the case of a Thomas process the variance-covariance matrix of the bivariate Gaussian distribution of daughters around their parents is given by $\gamma^2 \mathbf{I}_2$. In the case of a Matérn process, γ is the radius of the circle, centered at the location of a parent point, within which daughters are uniformly scattered.

Figure 1 illustrates how the point density in a circle of radius r centered at an arbitrary observed point, $b(x, r)$ (dashed circles), changes as the distance r increases. When r is small the point density in $b(x, r)$ is high because the existence of one point suggests others are likely to be nearby. Considering a NSPP, when r is large the density falls because the influence of the local cluster is diminished at larger distances. This also holds for void processes, as the existence of one point suggests that there is not a nearby parent and so other points are likely to be nearby as they are safe from deletion. As r increases, the fraction of the circle $b(x, r)$

that intersects with a void is likely to increase, thus decreasing point density within it. For further details see Appendix A.

In the following sections we describe parameter estimation for void, Thomas, and Matérn processes. Thomas processes have previously been fitted via maximisation of the Palm likelihood functions[10, 11] and we extend this approach to void and Matérn processes. In Section 3 we present the results of the application of our approach to a CRC patient data set.[2, 8]




Figure 1: i) A simulated void processes in the unit square. Daughter points are generated by a homogeneous Poisson process and are deleted if they fall within circular voids, of radius $R = 0.075$, centered at parent points (red crosses, unobserved in practice). In areas outwith the voids the daughters (black dots) have intensity given by $\lambda = 300$. The number of parents simulated follow a Poisson distribution with expectation $D = 10$. The dotted circles indicate the simulated voids. Considering an arbitrarily chosen daughter (encircled diamond) then there are more likely to be nearby daughters as that chosen daughter is not within a void. The dashed circles show how the intensity of observed daughters changes for different distances r , radius of the circle, away from an arbitrary point (i.e., the density of daughters is higher nearer the chosen point, but further away a void is more likely to be encountered and therefore the density decays). Plots ii), Thomas, and iii), Matérn, show two simulated Neyman-Scott point processes in the unit square with unobserved parent points (red crosses) siring the observed daughters (black dots). In each case parents are generated by a homogeneous Poisson process with intensity $D = 7$. The number of daughters sired by each parent are IID Poisson with expectation $\phi = 8$. In each case $\gamma = 0.05$. In ii) daughters are dispersed around their parents due to a bivariate normal distribution, $\mathbf{N}_2(\mathbf{0}, \gamma^2 \mathbf{I}_2)$; in iii) the daughters are uniformly distributed around their parents in a circle with radius γ . Here, again, the dashed circles show how the intensity of observed daughters changes for different distances r , radius of the circle, away from an arbitrary point (encircled diamond); the density of daughters is higher nearer the chosen point, but further away points in the sibling cluster are less likely to be encountered and therefore the density decays.

2 A Palm likelihood approach for parameter estimation

In the following sections we detail fitting clustered and void point process models to nuclei locations to model the spatial morphology of CRC tissue sections. To achieve this we consider the Palm intensity of the point patterns formed by the locations of cell nuclei in the CRC tissue sections. The Palm intensity is a function of distance r , characterised by the parameter vector θ , and is denoted $\lambda(r; \theta)$. The Palm intensity gives the expected intensity of a spatial point process at a distance r from an arbitrarily chosen point.[11, 12, 13] Estimation of θ for both the void and the cluster processes is nontrivial. We follow estimation methodology based on the Palm intensity function,[11] to fit a Thomas cluster process and extend this approach to allow parameter estimation for both Matérn cluster processes and void processes. These extensions to standard spatial point process methodology, we believe, are far better equipped than standard techniques to capture the complexities inherent in CRC data.

The following subsections derive the Palm intensity for two-dimensional void processes (Section 2.1), detail the derivation of the Palm intensity for the Thomas process,[11] and generalise this to the Matérn Palm intensity (Section 2.2). The full derivation of these Palm intensities in general d dimensions is given in Appendix A.

2.1 Void point process Palm intensity

The Palm intensity of a void process may be formulated as the product of the global point density of the pattern prior to deletion, λ , and $p_s(r)$, the probability that a potential point at distance r is not within a distance R of an unobserved parent and is therefore safe from deletion.

This probability is related to the geometry of the intersection between two circles of common radius R centered at an observed daughter and a potential daughter point. This is illustrated in more detail in Figure A.1 in the Appendix. The intersection, $I(r)$, between two circles centered at an observed daughter and a potential daughter respectively is the only region we know a parent cannot exist within R of the potential daughter. The remaining area, $A(r)$, of the circle centered at the potential point is the only region which may contain a parent whose void would delete that potential daughter. The potential daughter is safe from deletion if there are no parents in the region $A(r)$. Parents are generated by a homogeneous Poisson point process with intensity D therefore because the number of parents within $A(r)$ is Poisson with expectation $D A(r)$ this occurs with probability

$$p_s(r) = \exp(-D A(r)).$$

In 2 dimensions $A(r) = \pi R^2 - I(r)$. This area can be determined through the use of elementary geometry, detailed in Appendix A.1.[14]

The area of intersection, $I(r)$, between two circles with a common radius R is given by,

$$I(r) = \pi R^2 \text{I} \left(1 - \left(\frac{r}{2R} \right)^2 ; \frac{3}{2}, \frac{1}{2} \right). \quad (1)$$

Here $\text{I}(z; a, b) = \frac{B(z; a, b)}{B(a, b)}$ is the regularised Beta function. The Palm intensity function is then derived as,

$$\begin{aligned} \lambda(r; \theta) &= \lambda p_s(r), \\ &= \lambda \exp \left(-D \pi R^2 \left[1 - \text{I} \left(1 - \left(\frac{r}{2R} \right)^2 ; \frac{3}{2}, \frac{1}{2} \right) \right] \right), \\ &= \lambda \exp \left(-D \pi R^2 \left[1 - \text{F}_{g(r)} \left(\frac{3}{2}, \frac{1}{2} \right) \right] \right), \end{aligned} \quad (2)$$

where $g(r) = 1 - \left(\frac{r}{2R}\right)^2$, and $F_{g(r)}(\cdot, \cdot)$ is the CDF of the Beta distribution. Note when $r = 0 \Rightarrow g(r) = 1 \Rightarrow F_1(\cdot, \cdot) = 1 \Rightarrow \lambda(0; \theta) = \lambda$, the intensity of daughters prior to deletion. In addition, when $r = 2R \Rightarrow g(r) = 0 \Rightarrow F_0(\cdot, \cdot) = 0 \Rightarrow \lambda(0; \theta) = \lambda \exp(-D \pi R^2)$, due to the properties of the CDF. For the full derivation of the Palm intensity (2) see Appendix A.1.

Figure 2, plot i), shows both the empirical Palm intensity function (solid line), [15] and the fitted Palm intensity (dashed line) for the simulated void process shown in Figure 1. The Palm intensity is a piece-wise continuous function with two sub-domains $(0, 2R]$, $[2R, \infty)$. The horizontal asymptote to which the Palm intensity decays is the baseline intensity.

2.2 Neyman-Scott point process Palm intensity

The Palm intensity function of a Neyman-Scott point process (in 2 dimensions) is given by [10]

$$\lambda(r; \theta) = D \phi + \frac{\phi f_y(r; \gamma)}{2 \pi r}, \quad (3)$$

where the parameter D is the intensity of parents. Letting Y denote the distance between two randomly selected sisters (i.e., daughters sired by the same parent) then $f_y(r; \gamma)$ is the PDF of Y and is characterised by the parameter γ pertaining to the form of distribution of daughters around their parents.

The Palm intensity of a modified Thomas process is a special case of Equation (3) where (i) parent locations are realisations of a homogeneous Poisson process, and (ii) daughters are dispersed due to a bivariate normal distribution. [11] The Palm intensity for a modified Thomas process is given by

$$\lambda(r; \theta) = D \phi + \frac{\phi}{4\pi\gamma^2} \exp\left(\frac{-r^2}{4\gamma^2}\right). \quad (4)$$

Thus, $\lambda(r; \theta)$ is the sum of the intensity of non-sister points, $D \phi$, and a Gaussian function describing the intensity due to sister points. For a Thomas process, the PDF of Y is the distance between two normally distributed sisters. For example, in Equation 4,

$$f_y(r; \gamma) = \frac{r \exp(-r^2/(4\gamma^2))}{2 \gamma^2}, \quad (5)$$

where γ , is the parameter describing the Gaussian dispersion of daughters around their parents.

For the Matérn process, $f_y(r; \gamma)$ is now the PDF of the distance between two sisters generated from a uniform distribution within a circle of radius γ . This PDF [16] is given by

$$\begin{aligned} f_y(r; \gamma) &= \frac{4}{B\left(\frac{3}{2}, \frac{1}{2}\right)} \frac{r}{\gamma^3} \times \left[{}_2F_1\left(\frac{1}{2}, -\frac{1}{2}, \frac{3}{2}, 1\right) \gamma - {}_2F_1\left(\frac{1}{2}, -\frac{1}{2}, \frac{3}{2}, \frac{r^2}{4\gamma^2}\right) \frac{r}{2} \right], \\ &= \frac{4r \int_{\frac{r}{2}}^{\gamma} (\gamma^2 - x^2)^{\frac{1}{2}} dx}{B\left(\frac{3}{2}, \frac{1}{2}\right) \gamma^4}. \end{aligned} \quad (6)$$

Here $B(\cdot, \cdot)$ denotes the beta function, and ${}_2F_1(\cdot, \cdot, \cdot, \cdot)$ is the hyper-geometric function. For derivation of this see Appendix A.2. [16]

The empirical and fitted Palm intensities for the simulated Neyman-Scott point processes in Figure 1 are depicted in Figure 2. Both simulated processes in Figure 1, plots ii) and iii), were generated with parameter values $\theta = (D, \phi, \gamma) = (7, 8, 0.05)$. The Palm intensities in each case decay to the same asymptote, the baseline intensity $D\phi$.




Figure 2: Both the empirical[15] (solid lines) and fitted (dashed lines) Palm intensities for the simulated patterns shown in Figure 1. The fitted intensities were estimated using the methods discussed herein. The Palm intensity for the void process, parametrised as $\theta = (D, R, \lambda) = (10, 0.075, 300)$, is shown in plot i). This is a piece-wise continuous function with two sub-domains $(0, 2R]$, $[2R, \infty)$. The Palm intensities for both the simulated Neyman-Scott point processes, parametrised as $\theta = (D, \phi, \gamma) = (10, 8, 0, 05)$, are shown in plot ii). Here, the Matérn cluster process decays at a much steeper rate due again to it being a piece-wise continuous function with two sub-domains $(0, 2\gamma]$, $[2\gamma, \infty)$. The vertical dotted lines in each plot indicate where $r = 2R$ for the void process and $r = 2\gamma$ for the Matérn process.

2.3 The Palm likelihood for point process models

Before defining the Palm likelihood, the complication posed by edge effects should be noted. Edge effects occur because here we do not observe points patterns on R^2 , but rather a subset of R^2 , and this is not accounted for in the above derivation of the Palm intensity. For example, the distance between a point and its nearest neighbour within some region may in fact be greater than the distance between it and its “true” nearest neighbour that just happened to fall outside the study region. In the context of the CRC data discussed here a single image only manages to display a section of tissue, however, the pattern, and hence spatial structure of the cell nuclei, does not simply terminate at the image’s edge. This is an issue that arises in the modelling of any spatial point pattern and there are a number of approaches to address it.[12] Our approach calculates distances between points subject to periodic boundary conditions and as such we are required to set a truncation distance, t .[11, 10] This is set to a distance that is larger than any plausible distance between two sisters, or between parents in the void process case.

Taking an approximate-likelihood approach, the estimator for the parameter vector θ is given by

$$\hat{\theta} = \arg \max_{\theta} L(\theta; \mathbf{r}),$$

where $L(\theta; \mathbf{r})$ is the Palm likelihood.[11] Through the numerical maximisation of $\log(L(\theta; \mathbf{r}))$, with respect to θ , $\hat{\theta}$ is evaluated. The Palm likelihood, $L(\theta; \mathbf{r})$, is given by

$$L(\theta; \mathbf{r}) = \left(\prod_{\{i,j: \|\mathbf{x}_i - \mathbf{x}_j\| < t, i \neq j\}} n \lambda(r; \theta) \right) \exp \left(-n \int_0^t \lambda(r; \theta) 2 \pi r \, dr \right). \quad (7)$$

Here the product is that taken over all n distinct observed pairs of points, x_i, x_j ($i \neq j$), where the distance between them, $\|\mathbf{x}_i - \mathbf{x}_j\|$, is less than the truncation distance t .[10, 11] The integral term may be simplified for each process discussed herein; this results in an objective function that is very computationally efficient to compute. Simplifications of this integral are given for each process below.

- *Void process.* The integral for the void process is derived as,

$$\int_0^t \lambda(r; \theta) 2 \pi r \, dr = \lambda 2 \pi \exp(-D \pi r^2) \int_0^t \exp \left(D \pi r^2 F_{g(r)} \left(1, \frac{1}{2} \right) \right) r \, dr.$$

This is intractable as it contains the CDF of a Beta distribution, $F_{g(r)}(\cdot, \cdot)$, where $g(r) = 1 - \left(\frac{r}{2R}\right)^2$, but algorithms exist for efficient computation.[17]

- *Neyman-Scott point processes.* The integral given in likelihood (7) simplifies to,

$$\int_0^t \lambda(r; \theta) 2 \pi r \, dr = D \phi + \phi F_y(t; \gamma),$$

where $F_y(t; \gamma) = \int_0^t f_y(r; \gamma)$ is the CDF of the distance between two randomly selected sisters. Below we give $F_y(t; \gamma)$ for each case of the Neyman-Scott point process.

- *Matérn cluster process*

$$F_y(t; \gamma) = \frac{t^2}{\gamma^2} \left(1 - \frac{B_{\alpha} \left(\frac{1}{2}, \frac{3}{2} \right)}{B \left(\frac{1}{2}, \frac{3}{2} \right)} \right) + 4 \frac{B_{\alpha} \left(\frac{3}{2}, \frac{3}{2} \right)}{B \left(\frac{1}{2}, \frac{3}{2} \right)},$$

where $B_{\alpha}(\cdot, \cdot)$ and $B(\cdot, \cdot)$ are the incomplete beta function and the beta function respectively. That is, $B_{\alpha}(a, b) = \int_0^{\alpha} u^{(a-1)}(1-u)^{(b-1)} du$ and for $\alpha = 1$ $B(a, b) = B(a, b)$. The derivation of these equations is detailed in Appendix A.2.

– *Thomas cluster process*

Again let Y be the Euclidean distance between two randomly selected daughters (the locations of whom are independent given the the parent locations). The CDF of Y has been previously derived,[10] and for two-dimensional processes is given by

$$F_y(t; \gamma) = P\left(1, \frac{t^2}{4\gamma^2}\right),$$

where $P(\cdot, \cdot)$ is the regularised gamma function.

3 Spatial point process models for CRC data

The data analysed in this article refer to forty-two patients drawn from a wider data set of a pan-Scotland cohort, diagnosed with CRC.[8] At follow up, nineteen of the forty-two patients had died of CRC and the remaining twenty-three were alive. Follow up was a maximum of fourteen and a quarter years post-surgery, see Appendix C Table C2 for further details. In addition, the severity of the cancerous tissue was graded by a pathologist as either Dukes A (least severe), Dukes B, or Dukes C (most severe). All patients graded Dukes A were alive at follow up and those graded Dukes C were all dead at follow up; the patients graded Dukes B consisted of a mixture (approximately 50%) of patients who died from CRC and were alive at follow up.

Each patient had up to fifteen fields of view captured from the invasive front of the tissue section using a x20 objective. Automated immunofluorescence was performed using a DAKO link 48 platform, which labels patient tissue sections with panCytokeratin (tumour cells) and DAPI (all nuclei). Images of the tissue sections were captured using a HistoRx PM-2000 imaging platform (HistoRx Ltd., Branford, CT, USA). Set exposure times and imaging profiles were utilised to capture each fluorophore during image capture. These images were then processed by an automated imaging algorithm implemented in Definiens Developer XD™ that segmented the cancer from stroma and segmented each individual nucleus across an image. A panCytokeratin antibody, whose antigen is expressed in CRC cells, was used to visualise each cancer cell. This fluorescence visualisation was input to train a machine-learning algorithm within the Definiens software to automatically detect all cancer cells and segment them from the stromal cells. Post tumour to stroma segmentation the DAPI fluorescence was utilised to segment each nucleus in the image. Definiens' hierarchical image analysis approach allowed the classification of each nucleus to be only assigned to either a cancer cell or a stromal cell. Finally, the spatial coordinates from the centre of each classified nucleus was exported from Definiens and was used as input for the point process models.

Plot iii) in Figure 3 illustrates, for one patient, the observed point pattern of cells within a tissue section obtained from the digital image shown in plot i). The procedure for going from plot i) to plot iii) in Figure 3 for each slide required the use of an automatic imaging algorithm, detailed above.[2] In summary, distinct regions in the digitised tissue section were first divided into four types through machine learning: (i) tumour, (ii) stroma, (iii) necrosis/lumen and (iv) no tissue. Plot iii) of Figure 3 shows the point pattern formed by the tumour and stroma cell nuclei, black and grey, respectively, of the same tissue section. From this we see that the morphological patterns within tissue sections, at the very least, (i) are non-homogeneous, and (ii) exhibit complex spatial morphology.

Standard point pattern statistics are unable to capture the complex structures inherent in the tissue structure, such as the intermixing of tumour and stroma cells or features that are a result of non-homogeneity due to spatial clustering of cells, spatial intermixing of tumour and stroma, or necrosis. We therefore take a non-standard spatial point process approach using the Palm intensity function to characterise the structure of the tissue sections. We consider



Figure 3: Illustration of one image of a patient’s slide which enables the pinpointing of nuclei. Plot i) is a composite immunofluorescence digital image (red fluorescence highlights tumour cells and blue fluorescence highlights all nuclei in the image). Plot ii) is an image analysis mask overlay from automatic machine learnt segmentation of the digital image: tumour (purple), and stroma (turquoise). Plot iii) is the point pattern formed by the nuclei of the tumour (black) and stroma (grey) cells shown in the previous two images.

maximum Palm likelihood estimation for three point processes: a void process, a Thomas process, and a Matérn process. The parameters of these processes are then considered to represent characteristics of the tissue structure, and are estimated via maximisation of the Palm likelihood, see Section 2.3.

We implemented the methods described in Section 2 in the R package `palm`, [18] which is available on CRAN. Appendix B gives an example of its use in the context of the CRC data. In addition, online supplementary material illustrating the fitting of the model discussed in this article may be found at https://github.com/cmjt/examples/blob/master/CRC_point_process.md.

3.1 Results

Each of the point processes discussed give rise to different structures in point pattern data. The parameters in each process therefore reflect different aspects of the morphology of the tissue sections. Having fitted the models discussed above to the CRC data in this section we investigate whether the estimated point process parameters aid in discriminating between patient mortality outcomes (i.e., dead from disease or alive at follow up). We also determine whether patient staging (i.e., Dukes A, Dukes B and Dukes C) falls in line with the structure inferred by the parameter estimates.

We estimate parameters at the patient level for each tumour and stroma pattern by taking the product of the likelihood in Equation (7) over the multiple patterns (images) for each patient. These parameters are summarised for each i) Dukes staging and ii) patient outcome group by the mean and associated 95% confidence interval, which are calculated using one thousand bootstrap resamples.

We wish to compare parameter estimates between patients who were alive at follow-up and those who were not. This allows us to identify parameters that might enable us to determine patient mortality at follow up. For each parameter of each process, we conducted a permutation test—a nonparametric alternative to a t-test—to determine whether or not there

was a statistically significant difference between the two outcome groups. We computed the difference between the mean values of the parameter in each group and used this as our test statistic. We then enumerated 9999 random allocations of outcome groups to observations, and computed a p-value from the proportion of allocations that resulted in a test statistic with a magnitude at least as large as that we observed. We correct for multiple tests using two correction methods: \star Bonferroni correction and \diamond a false discovery rate correction.[19] The Bonferroni method is notoriously conservative controlling the family-wise error rate (i.e., the overall chance of rejecting the null hypothesis when it is true). The false discovery rate method we use holds for independent p-values with non-negative association controlling the expected proportion of false discoveries amongst the rejected hypotheses.

In addition, we use survival analysis to take into account patient follow up time, fitting Cox proportional hazard models[20] using each estimated parameter as a predictor. Table 4 in Appendix C gives the full results of these models.

Finally, we use our fitted models to explore the suitability of the NSPP and void point processes for capturing the underlying complexities of the cell distributions. We do this by using the empty space function, a functional summary characteristic of point patterns, to compare the fitted models to both the theoretical Poisson process and data simulated with the fitted parameter values.

3.1.1 Estimated parameters of the void process

In the context of spatial point patterns, a void is a region that contains no points where points may be expected, given the general structure of the data. This section assumes that each of the separate point patterns formed by tumour and stroma nuclei are realisations of void processes. Therefore, the void process describing the tumour cell point pattern reflects the patterning of stroma cells and vice versa; note the void processes also reflect the less frequently occurring regions of necrosis/lumen and no tissue.

Table 1 summarises the bootstrap resamples of the fitted void process parameters: parent density, D , child density, λ , and void radii, R . Table 1 show the averages and the associated 95% confidence interval for each Dukes stage and patient mortality group at follow up. The quoted p-values refer to the permutation test introduced above, after correcting for multiple tests.

The permutation tests (using the false discovery rate correction) determined that:

- (i) There is strong evidence (p-value 0.007) to suggest that stroma daughter density, λ , in patients that died was lower than in those that were alive at follow up. In addition, there is weak evidence (p-value 0.086) to suggest that stroma parent density, D , in patients that died was lower than in those that were alive at follow up. This is also indicated by the results of the Cox proportional hazard models where the hazard ratios were estimated respectively as $\hat{H}R = 0.477$, CI (0.277, 0.820), and $HR = 0.797$, CI (0.630, 1.008) for the scaled daughter and parent densities. Under the null hypothesis of no difference between outcome groups, the hazard ratio is equal to 1.
- (ii) There is weak evidence (p-value 0.086) to suggest that tumour daughter density, λ , in patients that died was higher than in those that were alive at follow up. Using Cox proportional hazard the results also indicate that patients with higher (scaled) daughter density are more likely to die, $\hat{H}R = 1.352$, CI (1.086, 1.685).

These results also align with the patients' Dukes grading. This would be as we would expect due to all patients graded C dying from CRC and all patients graded A being alive at follow up. For example, stroma nuclei patterns of patients graded Dukes C had a lower parent and daughter density on average than those graded Dukes A. The difference is less clear when

considering patients graded Dukes B; this we might expect as Grade B patients were a mixture (approximately 50%) of patients who died from CRC and those still alive at follow up.

Table 1: Summaries of the bootstrap resamples for parameters of the void process. These are summarised by the mean value and 95% confidence intervals for both tumour and stroma patterns at each Dukes grade, A, B, and C, and patient mortality at follow up, Alive or Died. The p-values relate to a permutation test as follows: $H_0 : |\mu_{Died} - \mu_{Alive}| = 0$ vs $H_1 : |\mu_{Died} - \mu_{Alive}| \neq 0$. To adjust for multiple comparisons we use both the Bonferroni correction method and a false discovery rate correction. These are denoted by the superscripts \star and \diamond respectively.

	Tumour			Stroma			Tumour			Stroma		
	A	B	C	A	B	C*	Alive	Died	p-value	Alive	Died*	p-value
\hat{D}	12.03	11.17	9.03	6.79	3.88	3.19	10.28	11.85	1.000 \star	5.42	3.41	
2.5%	5.68	8.25	4.78	4.59	3.14	2.35	6.65	8.41	0.686 \diamond	4.07	2.75	0.572 \star
97.5%	20.46	14.28	15.3	9.49	4.54	4.86	14.4	15.67		7.08	4.24	0.086 \diamond
$\hat{\lambda}$	3739.47	3930.64	6856.21	7876.23	3241.88	3421.34	3875.66	5118.35	0.601 \star	5632.16	3089.52	
2.5%	3423.6	3525.96	4989.92	6913.74	2935.34	3010.97	3476.16	4158	0.086 \diamond	4653.86	2832.47	0.007 \star
97.5%	4199.15	4470.04	8575.72	8925.9	3531.47	3884.2	4383.93	6173.21		6687.72	3346.47	0.007 \diamond
\hat{R}	0.20	0.23	0.30	0.22	0.25	0.22	0.23	0.25	1.000 \star	0.25	0.22	
2.5%	0.16	0.21	0.22	0.18	0.21	0.21	0.19	0.22	0.454 \diamond	0.21	0.20	1.000 \star
97.5%	0.25	0.26	0.39	0.28	0.29	0.23	0.26	0.30		0.29	0.25	0.454 \diamond

* The void process did not converge for one patient's stroma pattern; therefore, these results are based on forty-one out of forty-two patients.

3.1.2 Estimated parameters of both Neyman-Scott point processes

In the context of histopathology a Neyman-Scott point process can be thought of as a process giving rise to the clustering of cells within the tissue. Therefore, the assumption here is that the distribution of cells within both the tumour and the stroma are each realisations of a Neyman-Scott point process of two different formulations (i.e., A Thomas process and a Matérn process). Observed cells are daughters and we seek to infer the unobserved parents. Parents thus represent some abstracted developmental process that led to the observed arrangement of daughter cells within the tissue.

Thomas process Table 2 summarises the bootstrap resamples of the fitted Thomas process parameters parent density, \hat{D} , number of daughters per parent, $\hat{\phi}$, and dispersion parameter, $\hat{\gamma}$. These are summarised by the averages and the associated 95% confidence intervals for each Dukes stage and patient mortality group at follow up.

The permutation tests (using the false discovery rate correction) determined that:

- (i) There is strong evidence (p-value 0.007) to suggest that stroma nuclei patterns of patients who died from CRC have on average a lower number of daughters per parent, ϕ . This is also indicated by the results of the Cox proportional hazards model where the hazard ratio was estimated as $\hat{HR} = 0.242$, CI (0.080, 0.732).
- (ii) There is evidence (p-value 0.049) to suggest that tumour parent density, D , is lower in patient's who died from CRC. Using Cox proportional hazard the results also indicate that patients with lower (scaled) parent density are less likely to die, $\hat{HR} = 0.381$, CI (0.155, 0.940).

These results again align with the patients' Dukes grading. For example, stroma nuclei patterns of patients graded Dukes C had on average a lower number of daughters per parent

than those graded Dukes A. The difference is less clear when considering patients graded Dukes B; this we might expect as Grade B patients were a mixture (approximately 50%) of patients who died from CRC and those still alive at follow up.

Table 2: Summaries of the bootstrapped resamples for parameters of the Thomas process: summarised by the mean value and 95% confidence intervals for both tumour and stroma patterns at each Dukes grade, A, B, and C, and patient mortality at follow up, Alive or Died. The p-values relate to a permutation test as follows: $H_0 : |\mu_{Died} - \mu_{Alive}| = 0$ vs $H_1 : |\mu_{Died} - \mu_{Alive}| \neq 0$. To adjust for multiple comparisons we use both the Bonferroni correction method and a false discovery rate correction.[19] These are denoted by the superscripts \star and \diamond respectively.

	Tumour			Stroma			Alive	Tumour			Stroma		
	A	B	C	A	B	C		Died	p-value	Alive	Died	p-value	
\hat{D}	27.95	10.26	8.52	28.97	35.38	30.53	19.23	8.91		27.06	39.70		
2.5%	16.34	8.68	5.23	17.97	21.32	22.37	13.14	6.90	0.198 \star	17.63	25.6	1.000 \star	
97.5%	45.67	11.92	12.05	41.73	52.14	39.27	27.8	10.92	0.049 \diamond	37.47	56.57	0.360 \diamond	
$\hat{\phi}$	78.89	102.12	233.25	179.68	87.35	76.96	98.43	148.36		144.43	67.33		
2.5%	55.95	83.36	149.76	126.49	67.32	61.53	79.59	101.62	1.000 \star	110.29	54.88	0.014 \star	
97.5%	103.11	125.02	312.39	232.7	109.06	95.44	119.03	200.82	0.151 \diamond	177.1	79.79	0.007 \diamond	
$\hat{\gamma}$	0.08	0.09	0.09	0.08	0.09	0.09	0.09	0.08		0.09	0.09		
2.5%	0.06	0.08	0.08	0.07	0.09	0.08	0.08	0.08	1.000 \star	0.08	0.08	1.000 \star	
97.5%	0.09	0.09	0.11	0.10	0.10	0.11	0.10	0.10	0.792 \diamond	0.10	0.10	0.742 \diamond	

Matérn process The bootstrap resamples of the estimated Matérn process parameters for the tumour and stroma patterns are summarised in Table 3. The differences are similar to those noted above when assuming a Thomas cluster process. This is to be expected as both are cluster processes but differ only in their structure of a cluster.

The permutation tests (using the false discovery rate correction) determined that:

- (i) There is strong evidence (p-value 0.01) to suggest that stroma nuclei patterns of patients who died from CRC have on average a lower number of daughters per parent, ϕ . This is also indicated by the results of the Cox proportional hazards model where the hazard ratio was estimated as $\hat{HR} = 0.198$, CI (0.053, 0.740).
- (ii) There is evidence (p-value 0.055) to suggest that tumour parent density, D , is lower in patient's who died from CRC. Using Cox proportional hazard the results also indicate that patients with lower (scaled) parent density are less likely to die, $\hat{HR} = 0.388$, CI (0.156, 0.965).

These results again align with the patients' Dukes grading. For example, stroma nuclei patterns of patients graded Dukes C had on average a lower number of daughters per parent than those graded Dukes A. The difference is less clear when considering patients graded Dukes B; this we might expect as Grade B patients were a mixture (approximately 50%) of patients who died from CRC and those still alive at follow up.

3.2 Assessing model fit using the empty space function

Above we discuss a novel application of spatial point processes in capturing the spatial arrangement of cells. Doing so illustrates the value of analytic methods to aid in the understanding

Table 3: Summaries of the bootstrapped resamples for the Matérn process parameters: summarised by the mean value and 95% confidence intervals for both tumour and stroma patterns at each Dukes grade, A, B, and C, and patient mortality at follow up, Alive or Died. The p-values relate to a permutation test as follows: $H_0 : |\mu_{Died} - \mu_{Alive}| = 0$ vs $H_1 : |\mu_{Died} - \mu_{Alive}| \neq 0$. To adjust for multiple comparisons we use both the Bonferroni correction method and a false discovery rate correction.[19] These are denoted by the superscripts \star and \diamond respectively.

	Tumour			Stroma			Tumour			Stroma		
	A	B	C	A	B	C	Alive	Died	p-value	Alive	Died	p-value
\hat{D}	28.22	10.71	8.86	39.27	34.38	31.34	19.56	9.37	0.277 \star	30.22	40.96	1.000 \star
2.5%	16.17	9.13	5.69	22.58	22.2	24.23	13.1	7.26	0.055 \diamond	20.27	28.39	1.000 \star
97.5%	44.64	12.51	12.36	60.77	48.47	38.01	28.67	11.56		41.62	59.04	0.422 \diamond
$\hat{\phi}$	77.53	97.62	221.04	149.12	81.87	72.89	95.33	140.73	1.000 \star	126.43	63.08	0.029 \star
2.5%	56.18	80.2	136.32	103.63	64.5	60.71	76.7	96.25	0.151 \diamond	99.2	51.32	0.010 \diamond
97.5%	99.38	116.58	295.43	193.5	101.41	88.78	113.57	187.87		154.42	73.91	
$\hat{\gamma}$	0.15	0.16	0.18	0.14	0.18	0.18	0.16	0.16	1.000 \star	0.16	0.18	1.000 \star
2.5%	0.12	0.15	0.15	0.11	0.17	0.17	0.15	0.14	0.686 \diamond	0.14	0.16	0.422 \diamond
97.5%	0.18	0.18	0.2	0.17	0.19	0.2	0.18	0.18		0.18	0.19	

of cancerous tissue structure. We note above that the morphology of cancerous tissue structure is complex; therefore, no single model would be perfectly adequate to describe all the structural features in the data. However, we believe that our proposed methodology is a step towards taking into account the spatial morphology of the nuclei, which has not previously been considered.

In order to informally assess the fit of the models to the CRC data we compare the spatial patterning of the observed data and the data simulated from the fitted model through using the estimated empty space function, $H(r)$. [12] In two-dimensions, this describes the probability that the disc $b(x, r)$ of radius r centered at x is not empty. In the stationary case assumed here (i.e., shifting the pattern by a vector does not affect the distribution of points) the disc may be centered at the origin (i.e., $b(o, r)$).

Figure 4 shows the empty space functions for each considered point process for the image shown in Figure 3. Solid lines represent the estimated empty space function for the fitted model. The grey dotted lines are the estimated empty space functions from patterns simulated with the parameter estimates from the fitted model (i.e., the estimated parameters for the set of images to which Figure 3 belongs). The empty space function is estimated using the Chiu-Stoyan estimator employed in the R package `spatstat`. [21] The dashed lines are the theoretical empty space function for a homogeneous Poisson process, $H(r) = 1 - \exp(-\lambda\pi r^2)$ where λ is point density and r is the interpoint distance. [12]

We assess the model's fit in Figure 4 by looking for overlap between the envelope of the empty space functions for the simulated data, shown by the grey lines, and the empirical empty space function for the fitted model, shown by the solid line. Comparison of the solid line and the dashed line compares the fitted CRC model to a homogeneous Poisson process where the probability of points in a sphere of radius r is proportional to the area of that sphere.

The overlap of the envelope of the empty space functions for the simulated data and the empirical empty space function for the fitted model shown in the left hand plots of Figure 4 indicate that the Thomas, Matérn, and void process all appear to fit the data well. The right hand top and middle rows of Figure 4 indicate that the pattern formed by stroma cells is

slightly less likely to have nearby points at the same distances as a NSPP (either a Thomas or Matérn process).

From Figure 4 we see that the stroma cell pattern more quickly approaches that of a homogeneous Poisson process; this is the case at a distance of 3% of the width and height of the image ($r = 0.03$). This seems reasonable from the pattern of stroma cells shown in Figure 3 and is also reflected by the estimated lower void density for stroma cells (Table 1), which would lead to a pattern with fewer gaps.

A comparison of the empty space function for both the tumour and stroma cell patterns to that of the homogeneous Poisson process (dashed lines) reveals that at short distances the cells exhibit some regularity, whereas only at longer distances does the pattern exhibit clustering. This is expected as in both cases the points represent cell nuclei, and due to the size of the cells at short distances no other cells (points) can exist.

4 Discussion

This article presents a novel application of spatial point process statistics to capture the structure inherent in the spatial arrangement of cells through both void and cluster point processes. We characterise the spatial arrangement of cells within CRC tissue sections by treating their nuclei locus as a spatial point pattern and consider three spatial point processes (void, Thomas, Matérn), taking a Palm likelihood approach for parameter estimation in each case. This estimation is nontrivial and is achieved through extending existing work[10] which uses the interpoint distances between points (cells) to inform consideration of the spatial morphology of the tissue structure.

We illustrate that by using estimated parameters of the Palm intensity function as classifiers of the spatial morphology of cancerous tissue sections that the patterns formed by cells may be informative as to CRC patient survival. Although no one model would perfectly encapsulate the intricacies in the spatial structure of this tissue, we believe our methods open the door to taking a more analytical approach to describe these types of data. Although beyond the scope of this work, other point process models could be used in this setting. For example, the pattern of all nuclei (tumour and stroma) could be considered to be a realisation of a marked point process with cell type as a mark (i.e., characteristic of the point). This approach would consider all cell nuclei to be realisations from the same spatial point process where characteristics of the points (marks i.e., cell type) were realisations of another process, which would encompass the inter- and intra-mark interactions. However, such an approach would presume that the mark process followed some defined formulation. Furthermore, using latent structures in a marked log-Gaussian Cox process, for example, might enable the mark-point dependence to be inferred. Yet, in this case, a dependence between cell types and their location is not likely and not of particular interest. In our work we fit models to each pattern (tumour and stroma) independently to capture the spatial spatial features exhibited in both patterns.

The modelling framework detailed herein provides a standardised methodology that describes and reports the spatial distribution of cells in cancerous tissue sections in a way that avoids observer variability. That is, we do not rely human classifications of tissue morphology but use the pattern formed by the cells themselves to infer patterning in tissue structure. Using this methodology we rely on the ability of the imaging software to, with negligible error, pinpoint locations of the cell nuclei and to correctly identify the type of cell (i.e., tumour or stroma). Our proposed framework indicates that certain parameters of the processes we discuss are useful indicators of patient mortality. These parameters also align with the different Dukes staging the tissue samples were classified as by pathologists. This is not surprising as the grades were given by experts in the field and can be thought to roughly reflect the severity of the tumour; hence, they are a good indicator of patient outcome.

The literature reports morphological features such as tumour budding to be significantly associated with disease survival. However, the lack of consensus on quantification methods and observer variability has led to their exclusion from clinical guidelines.[22] In brief, high tumour budding is found in patients with an infiltrative growth pattern: finger-like protrusions invading widely across the stroma and thus forming large gaps between the cancer protrusions. Both tumour budding and infiltrative growth pattern would be reflective of the tumour cell void patterning described herein. In contrast, low tumour budding, or a pushing border growth pattern, described as a solid tumour mass with little stroma existing between cells, has been correlated with good outcome.[23] However cancerous tissue sections typically lack histopathological features that can be reliably identified and reproducibly quantified by eye. Our work uses an analytical framework to identify potential indicators of patient survival without the need to rely human observation and risk potential bias.

More broadly, this work reflects a growing interest in the use of analytical techniques in recognition of the importance of both spatial structure and spatial variations within cellular data. For example, in some cases it was observed that both the abundance of immune cells and their spatial variation within the tumour are important factors in patient outcome.[24] In addition, a recent review of spatial heterogeneity in cancers outlined the importance and relevance of spatial statistics in describing cellular patterns.[25] Here, we demonstrate that analytical methods perhaps more commonly used in ecological contexts[10, 15] may be used and built upon to aid understanding of the spatial structure of CRC tissue sections.

The methodology we describe above may be used to describe subjectively reported histopathological features such as infiltrative invasion pattern and tumour budding throughout the stroma. This is a known, and common, phenomenon that has been previously highly subjective and not reproducible through manual observation.[3, 22] We believe that the methods detailed herein are better equipped to capture the complexities and spatial heterogeneity inherent within the CRC microenvironment than the currently employed methods. We propose that this analysis enables future work centered on refining the treatment of individual cells, by stratifying individual cells further by biological and/or physical measurements in order to explore more complex questions that point at specific mechanisms of disease progression.

Acknowledgements

We give our sincere thanks to the three anonymous reviewers along with the Editor for their invaluable advice in greatly improving this paper. All suggestions were insightful and constructive resulting in a much clearer article. We greatly appreciate the time and effort all reviewers and the Editor put in.




Figure 4: This figure illustrates for one image of one patient’s slide, that shown in Figure 3, how the empty space function, $H(r)$, can be used to informally assess the model fit in each instance of the Thomas, Matérn, and void process. The dashed black line in each panel shows $H(r)$ for a homogeneous Poisson process (i.e., a constant intensity of points within some given region). We would expect $H(r)$ for our point pattern of cells to deviate from this. This is indeed the case in each panel. The empty space function, $H(r)$, for the tissue pattern is shown by the solid black line, which does deviate from the homogeneous Poisson line. The bottom right plot shows an overlap between the simulated pattern’s empty space function and that of the homogeneous Poisson process for the stroma nuclei. Looking again at Figure 3 this seems somewhat plausible as the stroma cells do seem to “blanket” the image. In each plot the distance r represent the percentage of the image’s width and height scaled to be $\in [0, 1]$. However, to assess model suitability it is not this we are interested in but the overlap between the fitted $H(r)$ for the pattern and the envelope of empty space functions for simulated data shown by the grey dotted lines.

References

- [1] T Mattfeldt and F Fleischer. Characterization of squamous cell carcinomas of the head and neck using methods of spatial statistics. *Journal of Microscopy*, 256(1):46–60, 2014.
- [2] Peter D Caie, Ying Zhou, Arran K Turnbull, Anca Oniscu, and David J Harrison. Novel histopathologic feature identified through image analysis augments stage ii colorectal cancer clinical reporting. *Oncotarget*, 7(28):44381, 2016.
- [3] Rafael S Pinheiro, Paulo Herman, Renato M Lupinacci, Quirino Lai, Evandro S Mello, Fabricio F Coelho, Marcos V Perini, Vincenzo Pugliese, Wellington Andraus, Ivan Cecconello, et al. Tumor growth pattern as predictor of colorectal liver metastasis recurrence. *The American Journal of Surgery*, 207(4):493–498, 2014.
- [4] Kazuaki Tokodai, Hiroto Narimatsu, Akiko Nishida, Kai Takaya, Yasuyuki Hara, Naoki Kawagishi, Eiji Hashizume, and Noriaki Ohuchi. Risk factors for recurrence in stage ii/iii colorectal cancer patients treated with curative surgery: the impact of postoperative tumor markers and an infiltrative growth pattern. *Journal of surgical oncology*, 114(3):368–374, 2016.
- [5] Teppei Morikawa, Aya Kuchiba, Zhi Rong Qian, Mari Mino-Kenudson, Jason L Hornick, Mai Yamauchi, Yu Imamura, Xiaoyun Liao, Reiko Nishihara, Jeffrey A Meyerhardt, et al. Prognostic significance and molecular associations of tumor growth pattern in colorectal cancer. *Annals of surgical oncology*, 19(6):1944–1953, 2012.
- [6] Gregor Rieger, Viktor H Koelzer, Heather E Dawson, Martin D Berger, Marion Hädrich, Daniel Inderbitzin, Alessandro Lugli, and Inti Zlobec. Comprehensive assessment of tumour budding on cytokeratin stains in colorectal cancer. *Histopathology*, 70(7):1044–1051, 2017.
- [7] Linde De Smedt, Sofie Palmans, and Xavier Sagaert. Tumour budding in colorectal cancer: what do we know and what can we do? *Virchows Archiv*, 468(4):397–408, 2016.
- [8] Peter D Caie, Arran K Turnbull, Susan M Farrington, Anca Oniscu, and David J Harrison. Quantification of tumour budding, lymphatic vessel density and invasion through image analysis in colorectal cancer. *Journal of Translational Medicine*, 12(156):1–12, 2014.
- [9] Jerzy Neyman and Elizabeth L Scott. Statistical approach to problems of cosmology. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 20(1):1–43, 1958.
- [10] B. C Stevenson, D. L Borchers, and R. M Fewster. Cluster capture-recapture to account for identification uncertainty on aerial surveys of animal populations. In submission.
- [11] Ushio Tanaka, Yoshiko Ogata, and Dietrich Stoyan. Parameter estimation and model selection for neyman-scott point processes. *Biometrical Journal*, 50(1):43–57, 2008.
- [12] Janine Illian, Antti Penttinen, Helga Stoyan, and Dietrich Stoyan. *Statistical analysis and modelling of spatial point patterns*. John Wiley & Sons, Chichester, 2008.
- [13] Jean-François Coeurjolly, Jesper Møller, and Rasmus Waagepetersen. A tutorial on palm distributions for spatial point processes. *International Statistical Review*, 85(3):404–420, 2017.

- [14] Shengqiao Li. Concise formulas for the area and volume of a hyperspherical cap. *Asian Journal of Mathematics and Statistics*, 4(1):66–70, 2011.
- [15] R. M Fewster, B. C Stevenson, and D. L Borchers. Trace-contrast models for capture–recapture without capture histories. *Statistical Science*, 31(2):245–258, 2016.
- [16] Shu-Ju Tu and Ephraim Fischbach. Random distance distribution for spherical objects: general theory and applications to physics. *Journal of Physics A: Mathematical and General*, 35(31):6557, 2002.
- [17] Armido R Didonato and Alfred H Morris Jr. Algorithm 708: Significant digit computation of the incomplete beta function ratios. *ACM Transactions on Mathematical Software (TOMS)*, 18(3):360–373, 1992.
- [18] B. C Stevenson. *palm: Fitting point process models using the Palm likelihood*, 2017.
- [19] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, 57(1):289–300, 1995.
- [20] Terry M Therneau and Patricia M Grambsch. *Modeling survival data: extending the Cox model*. Springer Science & Business Media, New York, 2013.
- [21] Adrian Baddeley, Ege Rubak, and Rolf Turner. *Spatial Point Patterns: Methodology and Applications with R*. Chapman and Hall/CRC Press, London, 2015.
- [22] Eva Karamitopoulou, Inti Zlobec, Viktor Hendrik Koelzer, Rupert Langer, Heather Dawson, and Alessandro Lugli. Tumour border configuration in colorectal cancer: proposal for an alternative scoring system based on the percentage of infiltrating margin. *Histopathology*, 67(4):464–473, 2015.
- [23] Inti Zlobec, Kristi Baker, Parham Minoo, Shinichi Hayashi, Luigi Terracciano, and Alessandro Lugli. Tumor border configuration added to tnm staging better stratifies stage ii colorectal cancer patients into prognostic subgroups. *Cancer*, 115(17):4021–4029, 2009.
- [24] Sidra Nawaz, Andreas Heindl, Konrad Koelble, and Yinyin Yuan. Beyond immune density: critical role of spatial heterogeneity in estrogen receptor-negative breast cancer. *Modern Pathology*, 28:766–777, 2015.
- [25] Andreas Heindl, Sidra Nawaz, and Yinyin Yuan. Mapping spatial heterogeneity in the tumor microenvironment: a new era for digital pathology. *Laboratory Investigation*, 95:377–384, 2015.
- [26] Terry M Therneau and Thomas Lumley. Package *survival*. *R Top Doc*, 128, 2015.

A The Palm intensity function

This appendix derives the d -dimensional Palm intensity functions for both the void and Matérn processes, Section 2. The application, Section 3, considers only the 2-dimensional case; therefore, above, we only provide the details of the Palm intensities and likelihoods in 2-dimensions. In this appendix we generalise this to consider d -dimensions. Due to this we now consider, for example, the volume of hyperspheres, and not the area of circles.

A.1 A d -dimensional void point process

We define a d -dimensional void process as follows. Let \bar{X} and P be independent homogeneous Poisson processes on R^d . Then the void process is given by $X = \bar{X} \setminus (\bigcup_{p \in P} b(p, R))$, where $b(v, R) = \{u \in R^d : \|u - v\| \leq R\}$ is the ball centered at v with radius $R > 0$.

The probability of a potential point, $x \in \bar{X}$, being outside the area of influence of a parent point, $p \in P$, is related to the geometry of the intersection between hyperspheres of common radius R centered at an observed daughter and a potential daughter. This concept is illustrated in Figure 5 where plot i) shows the intersection, $I(r) = b(d_o, R) \cap p(d_p, R)$, of two circles of radius R centered at d_o and d_p where $r < R$. The filled black circle represents an observed daughter, d_o , and the open circle represents a possible daughter point, d_p . We know that there cannot exist a parent within distance R of the observed daughter; therefore, given the location of the observed daughter the intersection, $I(r)$, cannot contain a parent point. The grey circles represent voids, of radius R , centered at potential parent points. The only area a parent might exist that would delete the potential daughter is in the region $A(r) = b(p_p, R) \setminus I(r)$, see the potential void centered at the grey cross. Any void centered a distance greater than R from the potential daughter will not delete it, see the potential void centered at the grey circled cross. This figure illustrates the geometry in two-dimensions, but this section considers the void process in d -dimensions and in order to calculate the intersection, $I(r)$ between two hyperspheres of common radius R the radius of the hyper-spherical caps, $R \sin(\theta)$, and the height of the hyper-spherical caps $R \cos(\theta)$ is required, where θ is the colatitude angle. Thus, the volume of intersection only depends on the radii of the hyperspheres as well as the distance between their centers r , see Equation 8. In plot ii), Figure 5, $I(r) = 0$ and the location of the observed daughter gives no information about potential parent locations that would affect the potential daughter. Hence, as the distance between daughters increases—from r plot i) to r in plot ii)—then the Palm intensity decays to the baseline intensity.




Figure 5: Plot i) illustrates the geometry of the intersection, $I(r) = b(d_o, R) \cap p(d_p, R)$, of two circles of radius R centered at d_o and d_p where $r < R$. The filled black circle represents an observed daughter, d_o , and the open circle represents a possible daughter point, d_p . Plot ii) illustrates that when $r > R$ the observed point offers no information as to the probability of a parent affecting the potential daughter, because no single void can exist that would delete both points.

To calculate $I(r)$, in Figure 5, the integrand of the volume of a $d - 1$ sphere of radius $R \sin(\theta)$ with height $R \cos(\theta)$ is required. As the hyperspheres are of common radius (R being the radius of a void) $I(r)$ is simply just twice this volume.[14] This is given by,

$$\begin{aligned}
I^d(r) &= 2 \times \int_0^\phi v^{d-1}(R \sin \theta) dR \cos \theta, \\
&= 2 \times \int_0^\phi v^{d-1}(R \sin \theta) R \sin \theta d\theta, \\
&= 2 \times \frac{\pi^{\frac{d-1}{2}}}{\Gamma(\frac{d-1}{2} + 1)} R^d \int_0^\phi \sin^d \theta d\theta, \\
&= 2 \times \frac{\pi^{\frac{d-1}{2}}}{\Gamma(\frac{d-1}{2} + 1)} R^d J_d(\phi), \\
&= \frac{\pi^{\frac{d-1}{2}}}{\Gamma(\frac{d-1}{2} + 1)} R^d B\left(\frac{d+1}{2}, \frac{1}{2}\right) I\left(\sin^2 \phi; \frac{d+1}{2}, \frac{1}{2}\right), \\
&= v^d(R) I\left(\sin^2 \phi; \frac{d+1}{2}, \frac{1}{2}\right), \\
&= v^d(R) I\left(1 - \left(\frac{r}{2R}\right)^2; \frac{d+1}{2}, \frac{1}{2}\right),
\end{aligned} \tag{8}$$

noting that $B(a, b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}$, $\Gamma\left(\frac{1}{2}\right) = \sqrt{\pi}$, and that $(R \cos \theta)^2 + (R \sin \theta)^2 = R^2 \rightarrow (\sin \theta)^2 = 1 - (\cos \theta)^2$ and using the cosine rule leads to $\cos \theta = \frac{r^2 + R^2 - R^2}{2rR} = \frac{r}{2R}$. Here $I(z; a, b) = \frac{B(z; a, b)}{B(a, b)}$ is the regularised Beta function.

Recall that the Palm intensity of the void process can be written as $\lambda_0(r) = \lambda p_s(r)$ where $p_s(r) = \exp(-D A(r))$ is the probability that an arbitrary point has no parent within a some distance R . Here $A(r) = v^d(R) - I(r)$ where $v^d(R)$ is the d -dimensional volume of a hypersphere of radius R . Then the Palm intensity is given by

$$\begin{aligned}
\lambda_0(r) &= \lambda \exp\left(-D v^d(R) \left[1 - I\left(1 - \left(\frac{r}{2R}\right)^2; \frac{d+1}{2}, \frac{1}{2}\right)\right]\right), \\
&= \lambda \exp\left(-D v^d(R) \left[1 - F_{g(r)}\left(\frac{d+1}{2}, \frac{1}{2}\right)\right]\right),
\end{aligned} \tag{9}$$

where $g(r) = 1 - \left(\frac{r}{2R}\right)^2$, and $F_{g(r)}(\cdot, \cdot)$ is the CDF of the Beta distribution. Thus, when $r = 0 \Rightarrow g(r) = 1 \Rightarrow F_1(\cdot, \cdot) = 1 \Rightarrow \lambda_0(0) = \lambda$. In addition, when $r = 2R \Rightarrow g(r) = 0 \Rightarrow F_0(\cdot, \cdot) = 0 \Rightarrow \lambda_0(0) = \lambda \exp(-D v_d(R))$, due to the properties of the CDF. The functional form of this Palm intensity is shown in Figure 6. Letting $d = 2$ in Equation (9) would lead to the Palm intensity given in Equation (2).

A.2 A d -dimensional Matérn point process

Daughters of a Matérn process are uniformly distributed around their parents. The parameter γ in Equation 3 refers to the radius of the sphere centered at a selected parent outwith which we do not observe sired daughters. Figure 1 illustrates the two types of Neyman-Scott point processes simulated with the same value of γ ; for the same value of γ the Palm intensity function for the Matérn process is initially much higher, and decays at a much faster rate to the horizontal asymptote, $D\nu$. As illustrated in Figure 7 $\lambda(r; \theta)$ is a continuous piece-wise monotonic function of two sub-domains, $[0, 2\gamma]$ and $[2\gamma, \infty)$. The common endpoint of the sub-domains, 2γ , relates to the structure of the Matérn process; that is, the distance between two




Figure 6: The functional form of the Palm intensity for the void process in 2 dimensions. The horizontal asymptote is given by $\lambda \exp(-D\pi R^2)$, which is the value that $\lambda_0(r)$ decays to for values of $r \geq 2R$. The Palm intensity at $r = 0$ is simply λ , as for $r = 0$ $g(r) = 1$ in Equation 2, thus the exponential becomes 1. At the value $r = 2R$ the volume of intersection between the spheres encircling an observed daughter and a potential point of radius R is zero, thus the contribution from the CDF of the Beta distribution to $\lambda(r; \theta)$ is zero.

sibling daughters cannot be more than the diameter, 2γ , of a sphere centered at an unobserved parent away from one another. Thus, the probability of observing a sister at a distance r from an arbitrarily chosen daughter pertains to the intersection of the hyperspheres centered at these points, $b(x, \gamma) \cap b(y, \gamma), x \neq y \in N$, where N is the point pattern. When the distance between these points, $r \geq 2\gamma$ then $b(x, \gamma) \cap b(y, \gamma) = 0$.

The d -dimensional version of Equation 6 is given by,

$$\begin{aligned}
 f_y^d(r; \gamma) &= \frac{2d}{B(\frac{d}{2} + \frac{1}{2}, \frac{1}{2})} \frac{r^{d-1}}{\gamma^{d+1}} \left[{}_2F_1\left(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, 1\right) \gamma - {}_2F_1\left(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, \frac{r^2}{4\gamma^2}\right) \frac{r}{2} \right], \\
 &= \frac{2d r^{d-1} \int_{\frac{r}{2}}^{\gamma} (\gamma^2 - x^2)^{\frac{d-1}{2}} dx}{B(\frac{d}{2} + \frac{1}{2}, \frac{1}{2}) \gamma^{2d}}. \tag{10}
 \end{aligned}$$

Here $B(\cdot, \cdot)$ denotes the beta function, and ${}_2F_1(\cdot, \cdot, \cdot, \cdot)$ the hyper-geometric function.

Below we show how this PDF reduces in $d = 2$ and $d = 3$ to forms equivalent to the PDFs of the distances between two randomly selected sisters in the respective dimensions.[12, p.376]

- for $d = 2$ It should be noted that,

$$\begin{aligned}
 \int_{\frac{r}{2}}^{\gamma} (\gamma^2 - x^2)^{\frac{d-1}{2}} dx &= \int_{\frac{r}{2}}^{\gamma} (\gamma^2 - x^2)^{\frac{1}{2}} dx \\
 &= \frac{1}{8} \left(4\gamma^2 \sec^{-1}\left(\frac{2\gamma}{r}\right) - r \sqrt{4\gamma^2 - r^2} \right).
 \end{aligned}$$

therefore,

$$\begin{aligned}
\frac{2 \, d \, r^{d-1} \int_{\frac{r}{2}}^{\gamma} (\gamma^2 - x^2)^{\frac{d-1}{2}} dx}{B(\frac{d}{2} + \frac{1}{2}, \frac{1}{2}) \gamma^{2d}} &= \frac{4 \, r \int_{\frac{r}{2}}^{\gamma} (\gamma^2 - x^2)^{\frac{1}{2}} dx}{B(\frac{3}{2}, \frac{1}{2}) \gamma^4} \\
&= \frac{r \Gamma(2)}{\frac{1}{2} \sqrt{\pi}^2 \gamma^4 2} \left(4 \gamma^2 \sec^{-1} \left(\frac{2\gamma}{r} \right) - r \sqrt{4\gamma^2 - r^2} \right) \\
&= \frac{r}{\pi \gamma^4} \left(4 \gamma^2 \sec^{-1} \left(\frac{2\gamma}{r} \right) - r \sqrt{4\gamma^2 - r^2} \right) \\
&= \frac{r}{\pi \gamma^4} \left(4 \gamma^2 \cos^{-1} \left(\frac{r}{2\gamma} \right) - r \sqrt{4\gamma^2 - r^2} \right),
\end{aligned}$$

noting that $\sec^{-1}(x) = \cos^{-1}(\frac{1}{x})$, and as $B(\frac{3}{2}, \frac{1}{2}) = \frac{\Gamma(\frac{3}{2})\Gamma(\frac{1}{2})}{\Gamma(2)}$, recalling that $\Gamma(\frac{3}{2}) = \frac{1}{2}\sqrt{\pi}$, $\Gamma(\frac{1}{2}) = \sqrt{\pi}$, and $\Gamma(2) = 1$.

•for $d = 3$ It should be noted that,

$$\begin{aligned}
\int_{\frac{r}{2}}^{\gamma} (\gamma^2 - x^2)^{\frac{d-1}{2}} dx &= \int_{\frac{r}{2}}^{\gamma} (\gamma^2 - x^2) dx \\
&= \frac{1}{24} (r - 2\gamma)^2 (r + 4\gamma).
\end{aligned}$$

therefore,

$$\begin{aligned}
\frac{2 \, d \, r^{d-1} \int_{\frac{r}{2}}^{\gamma} (\gamma^2 - x^2)^{\frac{d-1}{2}} dx}{B(\frac{d}{2} + \frac{1}{2}, \frac{1}{2}) \gamma^{2d}} &= \frac{6 \, r^2 \int_{\frac{r}{2}}^{\gamma} (\gamma^2 - x^2) dx}{B(2, \frac{1}{2}) \gamma^6} \\
&= \frac{6 \, r^2}{24 \, B(2, \frac{1}{2}) \gamma^6} (r - 2\gamma)^2 (r + 4\gamma), \\
&= \frac{r^2 \Gamma(\frac{5}{2})}{4 \gamma^6 \Gamma(2) \Gamma(\frac{1}{2})} (r - 2\gamma)^2 (r + 4\gamma) \\
&= \frac{r^2 \frac{3}{4} \sqrt{\pi}}{4 \gamma^6 \sqrt{\pi}} (r - 2\gamma)^2 (r + 4\gamma) \\
&= \frac{3 \, r^2}{16 \, \gamma^6} (r - 2\gamma)^2 (r + 4\gamma) \\
&= \frac{3 \, r^2}{16 \, \gamma^6} \left(\gamma - \frac{r}{2} \right)^2 \left(2\gamma + \frac{r}{2} \right).
\end{aligned}$$

noting that $\Gamma(\frac{5}{2}) = \frac{3}{4}\sqrt{\pi}$, $\Gamma(\frac{1}{2}) = \sqrt{\pi}$, and $\Gamma(2) = 1$.

Upon substitution of the PDF (Equation 10) into the Palm intensity function (Equation 3) simplifications occur—this is also the case for the modified Thomas process, Section 2.3.[10] These simplifications circumvent the numerical instability in $\lambda(r; \boldsymbol{\theta})$ at $r = 0$ as both the numerator and denominator in the second term contain the term r^{d-1} . Thus,

$$\begin{aligned}
\lambda(r; \boldsymbol{\theta}) &= D E_c(\phi) + \frac{[E_s(\phi) - 1] f_y^d(r; \gamma)}{s^d(r)}, \\
&= D E_c(\phi) + \frac{[E_s(\phi) - 1] 2 d r^{d-1}}{B(\frac{d}{2} + \frac{1}{2}, \frac{1}{2}) \gamma^{d+1}} \frac{\Gamma(\frac{d}{2} + 1)}{d \pi^{d/2} r^{d-1}} \\
&\times \left[{}_2F_1\left(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, 1\right) \gamma - {}_2F_1\left(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, \frac{r^2}{4\gamma^2}\right) \frac{r}{2} \right], \\
&= D E_c(\phi) + \frac{2[E_s(\phi) - 1]}{B(\frac{d}{2} + \frac{1}{2}, \frac{1}{2}) \gamma^{d+1}} \frac{\Gamma(\frac{d}{2} + 1)}{\pi^{d/2}} \\
&\times \left[{}_2F_1\left(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, 1\right) \gamma - {}_2F_1\left(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, \frac{r^2}{4\gamma^2}\right) \frac{r}{2} \right], \tag{11} \\
&= D E_c(\phi) + \frac{2[E_s(\phi) - 1]}{\gamma^{d+1} \pi^{\frac{d}{2}}} \frac{\Gamma(\frac{d}{2} + 1)^2}{\Gamma(\frac{d}{2} + \frac{1}{2}) \Gamma(\frac{1}{2})} \\
&\times \left[{}_2F_1\left(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, 1\right) \gamma - {}_2F_1\left(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, \frac{r^2}{4\gamma^2}\right) \frac{r}{2} \right], \\
&= D E_c(\phi) + \frac{2[E_s(\phi) - 1] \Gamma(\frac{d}{2} + 1)^2}{\gamma^{d+1} \pi^{\frac{1}{2}(d+1)} \Gamma(\frac{d}{2} + \frac{1}{2})} \\
&\times \left[{}_2F_1\left(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, 1\right) \gamma - {}_2F_1\left(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, \frac{r^2}{4\gamma^2}\right) \frac{r}{2} \right].
\end{aligned}$$

noting that $B(x, y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}$, and $\Gamma(\frac{1}{2}) = \sqrt{\pi}$.




Figure 7: The functional form of the Palm intensity for both the Matérn (solid curve) and Thomas (dotted curve) variants of the Neyman-Scott point process in 2 dimensions. The horizontal asymptote is given by $D\nu$ for both processes. The difference between the horizontal asymptote and the y -intercept is given by $\nu/(4\pi\sigma^2)$ for the Thomas process, and $2\nu F(\cdot)/(\pi B(\cdot)\gamma^2)$ for the Matérn process. Here $F(\cdot) = {}_2F_1(\frac{1}{2}, \frac{1}{2} - \frac{d}{2}, \frac{3}{2}, 1)$ denotes the hyper-geometric function, and $B(\cdot) = B(\frac{3}{2}, \frac{1}{2})$ denotes the beta function. The point of inflection of the Gaussian term for the Thomas process is given by $\sigma\sqrt{2}$, for the Matérn process 2γ is the point at which the Palm intensity decays to the horizontal asymptote.

B Software implementation

All of the spatial point process models mentioned in this article can be fitted using the R package `palms`, [18] which is on CRAN. The two fitting functions used are `fit.ns()`, to fit either

a Thomas or a Matérn cluster process, and `fit.void()`, to fit a void process model. Below, we outline fitting the models to the CRC data however a more comprehensive tutorial is available online at https://github.com/cmjt/examples/blob/master/CRC_point_process.md.

Let the R object `points` be a list of length N where each element is a $2 \times n$ matrix of the nuclei locations (either stroma or tumour) for one image of one of the patient's N images. Furthermore, let R in each case be the truncation distance as mentioned in Section 3, `lims` be a list of length N where each element is a matrix of the rectangular image (for the matching element in `points`) boundary where columns give the lower and upper bounds respectively, and `c(lower, upper)` be parameter bounds for the parameter D of the void process. Then to fit a Thomas, Matérn, and void process respectively to the point pattern `points` one would run the following R code,

```
library(palm)
## the package palm can be installed from CRAN by running install.packages("palm")
fit.thomas <- fit.ns(points = points, lims = lims, R = R)
fit.matern <- fit.ns(points = points, lims = lims, R = R, disp = "uniform")
fit.void <- fit.void(points = points, lims = lims, R = R,
                    bounds = list(Dc = c(lower, upper)))
```

It is also worth noting that the argument `points` may also be a single matrix of point locations. In this case the `lims` argument must also be a singular matrix giving the bounds of `points`.

C Follow up of CRC patients

Using the R package `survival`[26] we fitted Cox proportional hazards models[20] using each estimated parameter as a predictor. The results of these models are shown in Table 4. In line with the results discussed in Section 3.1 an higher number of daughters per parent in both the Thomas and Matérn process, ϕ , for Stroma patterns were linked with patients not dying from CRC.

Table 5 gives the Dukes staging and site of the tumour along with the number of months post surgery follow up was. Due to the sensitive nature of the data we cannot give more information than provided here.

Table 4: Results of the Cox proportional hazards models fitted using each estimated parameter from the point processes discussed herein as predictors. The estimated coefficients of the predictors are given in the first column, along with their standard errors a p-value associated with the probability that that parameter is zero (second and third column respectively). The predictors (parameters from the point process models discussed in this article) were scaled so that the estimated coefficient of the Cox proportional hazard regression were more readable: void $\hat{\lambda}/1000$, Thomas $\hat{D}/10$, Thomas $\hat{\phi}/100$, Thomas $\hat{\gamma} \times 10$, Matérn $\hat{D}/10$, Matérn $\hat{\phi}/100$, and Matérn $\hat{\gamma} \times 10$. The hazard ratio, HR column, relates to risk of dying given a one unit increase in the associated (scaled) parameter. For example, assuming a Thomas process then for every unit increase in the expected number of daughters per parent/100 in the Stroma pattern then the a patient is less likely to die. The 95% confidence interval (CI) bounds of the hazard ratio are given in the last two columns.

			Estimate	se	p-value	HR	HR CI bounds	
							2.5%	97.5%
void	Tumour	\hat{D}	0.018	0.022	0.424	1.018	0.975	1.063
	Tumour	$\hat{\lambda}$	0.302	0.112	0.007	1.352	1.086	1.685
	Tumour	\hat{R}	-0.127	2.477	0.959	0.881	0.007	113.063
	Stroma	\hat{D}	-0.227	0.120	0.059	0.797	0.630	1.008
	Stroma	$\hat{\lambda}$	-0.741	0.277	0.007	0.477	0.277	0.820
	Stroma	\hat{R}	-2.324	3.314	0.483	0.098	0.000	64.852
Thomas	Tumour	\hat{D}	-0.964	0.460	0.036	0.381	0.155	0.940
	Tumour	$\hat{\phi}$	0.286	0.256	0.263	1.332	0.806	2.199
	Tumour	$\hat{\gamma}$	-0.843	0.981	0.391	0.431	0.063	2.947
	Stroma	\hat{D}	0.077	0.063	0.218	1.080	0.955	1.222
	Stroma	$\hat{\phi}$	-1.418	0.565	0.012	0.242	0.080	0.732
	Stroma	$\hat{\gamma}$	0.800	1.251	0.522	2.226	0.192	25.839
Matérn	Tumour	\hat{D}	-0.947	0.465	0.042	0.388	0.156	0.965
	Tumour	$\hat{\phi}$	0.368	0.285	0.197	1.445	0.826	2.528
	Tumour	$\hat{\gamma}$	-0.495	0.528	0.349	0.610	0.217	1.717
	Stroma	\hat{D}	0.065	0.064	0.306	1.068	0.942	1.210
	Stroma	$\hat{\phi}$	-1.617	0.672	0.016	0.198	0.053	0.740
	Stroma	$\hat{\gamma}$	0.797	0.693	0.250	2.220	0.571	8.632

Table 5: A summary of the forty-two CRC patients discussed above including the site of the tumour and the number of months post surgery follow up occurred. All patients will a follow up of fewer than nine years died during the follow up period. All patients graded as Dukes C died from CRC whereas all patients graded Dukes A were alive at follow up; twelve of the patients graded Dukes B were alive at follow up, eleven had died.

Dukes stage	Tumour site	Sex	Follow up (months)
A	Sigmoid colon	M	144.90
A	Sigmoid colon	M	138.90
A	Ascending colon/hepatic flexure	F	145.40
A	Caecum	F	140.60
A	Ascending colon/hepatic flexure	M	126.20
A	Rectum	M	125.40
A	Sigmoid colon	F	125.80
A	Rectum	M	122.40
A	Sigmoid colon	M	128.30
A	Sigmoid colon	F	120.80
A	Sigmoid colon	M	124.40
B	Sigmoid colon	F	46.70
B	Ascending colon/hepatic flexure	F	37.00
B	Sigmoid colon	F	66.80
B	Rectum	F	32.60
B	Splenic flexure	M	137.80
B	Descending colon	M	26.50
B	Sigmoid colon	F	11.60
B	Ascending colon/hepatic flexure	M	10.20
B	Rectum	M	131.80
B	Sigmoid colon	M	135.40
B	Ascending colon/hepatic flexure	M	108.30
B	Caecum	M	127.70
B	Rectum	M	37.90
B	Rectosigmoid	M	127.60
B	Caecum	F	118.60
B	Descending colon	F	117.20
B	Caecum	F	123.10
B	Rectum	F	118.20
B	Transverse colon	F	112.30
B	Ascending colon/hepatic flexure	F	112.40
B	Rectum	M	111.80
B	Transverse colon	F	109.50
B	Rectum	F	47.30
C	Sigmoid colon	F	20.20
C	Rectosigmoid	F	171.90
C	Rectosigmoid	F	37.60
C	Caecum	F	25.30
C	Rectum	F	10.43
C	Caecum	F	22.73
C	Caecum	M	171.63
C	Caecum	M	23.56