

Spatial statistics and attentional dynamics in scene viewing

Ralf Engbert

University of Potsdam, Germany
Bernstein Center for Computational Neuroscience Berlin,
Berlin, Germany



Hans A. Trukenbrod

University of Potsdam, Germany
Bernstein Center for Computational Neuroscience Berlin,
Berlin, Germany



Simon Barthelmé

Bernstein Center for Computational Neuroscience Berlin,
Berlin, Germany
University of Geneva,
Geneva, Switzerland



Felix A. Wichmann

Eberhard Karls University of Tübingen, Germany
Bernstein Center for Computational Neuroscience
Tübingen, Germany
Max Planck Institute for Intelligent Systems,
Tübingen, Germany



In humans and in foveated animals visual acuity is highly concentrated at the center of gaze, so that choosing where to look next is an important example of online, rapid decision-making. Computational neuroscientists have developed biologically-inspired models of visual attention, termed saliency maps, which successfully predict where people fixate on average. Using point process theory for spatial statistics, we show that scanpaths contain, however, important statistical structure, such as spatial clustering on top of distributions of gaze positions. Here, we develop a dynamical model of saccadic selection that accurately predicts the distribution of gaze positions as well as spatial clustering along individual scanpaths. Our model relies on activation dynamics via spatially-limited (foveated) access to saliency information, and, second, a leaky memory process controlling the re-inspection of target regions. This theoretical framework models a form of context-dependent decision-making, linking neural dynamics of attention to behavioral gaze data.

Introduction

Research on visual attention models over the past 25 years has resulted in a number of computational models (Borji & Itti, 2013)—using diverse computational

mechanisms—often capable of predicting fixation *locations* for a given input image with reasonable accuracy (Itti, Koch, & Niebur, 1998; Kienzle, Franz, Schölkopf, & Wichmann, 2009; Torralba, Oliva, Castelhano, & Henderson, 2006; Tsotsos et al., 1995). The models compute so-called *saliency maps*, highlighting those parts of an input image that stand out relative to the surrounding areas (Itti & Koch, 2001). However, the human visual system is foveated, i.e., it is only able to acquire high-resolution information from a very limited region surrounding the current gaze position (the fovea). Outside the foveal region, visual acuity falls off rapidly, while the effects of visual crowding increase, so that visual processing in the periphery has very limited resolution (Jones & Higgins, 1947; Levi, 2008; Rosenholtz, Huang, & Ehinger, 2012).

As a consequence, to explore an entire visual scene we must shift our gaze continually to new regions of interest by producing rapid eye movements (saccades) about three to four times per second (Findlay & Gilchrist, 2012). Thus, given the progress on mathematical models of visual attention, there is an increasing need for computational models that bridge the gap between static saliency maps—which a human observer's visual system can only know *after* exploring the entire image with its fovea—and the dynamic principles of saccadic selection underlying the generation of scanpaths by human

Citation: Engbert, R., Trukenbrod, H. A., Barthelmé, S., & Wichmann, F. A. (2015). Spatial statistics and attentional dynamics in scene viewing. *Journal of Vision*, 15(1):14, 1–17. <http://www.journalofvision.org/content/15/1/14>, doi: 10.1167/15.1.14.

observers. Moreover, part of the mismatch between computer-generated saliency maps and actual gaze patterns might be explained by properties of the visuomotor system (Findlay & Walker, 1999). Recently, a number of publications addressed specific aspects of this problem, e.g., different roles for short and long saccades (Tatler, Baddeley, & Vincent, 2006) or return saccades (Ludwig, Farrell, Ellis, & Gilchrist, 2009; Wilming, Harst, Schmidt, & König, 2013). Moreover, behavioral biases might produce an important contribution to eye-movement statistics (Tatler & Vincent, 2009). What is currently missing is an integrative computational model that addresses the key aspects of visuomotor control in a coherent theoretical framework. We set out to develop one possible integrative model.

Spatial patterns of gaze positions carry rich information on the processes of saccadic selection by the human visual system, and this information can be analyzed applying methods from the theory of spatial point processes (Illian, Penttinen, Stoyan, & Stoyan, 2008; Barthelmé, Trukenbrod, Engbert, & Wichmann, 2013). Saliency maps aim at the prediction of two-dimensional (2D) densities of gaze patterns (first-order spatial statistics). However, saliency maps do not contain the rich information about spatial interactions inherent in experimental eye-tracking data: Fixations are interdependent. Second-order statistics provide quantitative tools to investigate interactions in gaze patterns. Such interactions in turn may be used to gain information about the processes (Law et al., 2009) underlying the generation of neighboring gaze positions, which themselves are directly related to models of saccadic selection.

We start with analyzing the spatial statistics of gaze patterns using point process theory (Illian et al., 2008) and show that gaze patterns are characterized by small-scale clustering, in addition to the inhibition-of-return mechanism (Klein, 2000) that is thought to represent the dominant dynamical principle in extant attention models (Itti & Koch, 2001). Next, since these results provide strong constraints for possible neural mechanisms of saccadic selection, we develop a dynamical model for real-time attention allocation and gaze control based on activation-based maps (Engbert, 2012; Engbert, Mergenthaler, Sinn, & Pikovsky, 2011). Finally, the model is compared against a range of statistical null models using methods of spatial statistics.

Methods

Experiment

Stimulus material

A set of 30 randomly selected, natural landscape photographs (color) was presented to human observers

on a 20 in. CRT monitor (Mitsubishi Diamond Pro 2070; frame rate 120 Hz; resolution: 1280 × 1024 pixels; Mitsubishi Electric Corporation, Tokyo, Japan). Images were classified into two categories, natural object-based scenes (image set 1: 15 images) versus images showing abstract natural patterns (image set 2: 15 images). All images were presented centrally with gray borders extending 32 pixels to the top/bottom and 40 pixels to the left/right of the image, since accuracy of eye tracking systems falls off toward the monitor edges.

Task and procedure

Participants were instructed to position their heads on a chin rest in front of a computer screen at a viewing distance of 70 cm. Eye movements were recorded binocularly using an Eyelink 1000 video-based eye-tracker (SR-Research, Osgoode/ON, Canada) with a sampling rate of 1000 Hz. Trials began with a black fixation cross presented on gray background at a random location within the image boundaries. After successful fixation, the fixation cross was replaced by the image for 10 s. Participants were instructed to explore each scene for a subsequent memory test. During the experiment, we presented 30 images twice. Here we limit our analysis to the first presentation.

Participants

We recorded eye movements from 35 participants (20 female, 15 male) aged between 17 and 36 years (mean age: 24 years) with normal or corrected-to-normal vision. Participants were recruited from the University of Potsdam and from a local school (32 students, three pupils). All participants received credit points or 8€ for (about US \$9.50) for participation.

Data preprocessing and saccade detection

We applied a velocity-based algorithm for saccade detection (Engbert & Kliegl, 2003; Engbert & Mergenthaler, 2006). Saccades had a minimum amplitude of 0.5° and exceeded the average velocity during a trial by six standard deviations for at least 6 ms. Eye traces between two successive saccades were tagged as fixations with a mean fixation position averaged across both eyes. Since eye position was determined by the presentation of a fixation cross at the beginning of a trial, we excluded all initial fixations from the data set (image set 1: 525; image set 2: 525). Furthermore, we removed fixations containing a blink or with a blink during an adjacent saccade (image set 1: 580; image set 2: 588). Overall, the number of fixations remaining for further analyses was 13,349 (image set 1) and 12,740 (image set 2).

Spatial statistics

Gaze positions can be interpreted as realizations from a spatial point process (Illian et al., 2008) that can be represented as the random set of points $N = \{x_1, x_2, x_3, \dots\}$ (also called a point pattern). The 2D density (or intensity) λ of the spatial point process is given as the expectation or mean value of the number of points in an observation window B , i.e., $\lambda = E(n(B))$, where $n(\cdot)$ is a counting measure. A process is statistically homogeneous (or stationary), if N and the translated set $N_x = \{x_1 + x, x_2 + x, x_3 + x, \dots\}$ have the same distribution for all x . For a stationary spatial point process, the intensity λ is constant over space. For a nonstationary process, the intensity is a function of location, $\lambda = \lambda(x)$. For the computation of densities from experimental data, we used kernel-density estimates with bandwidth parameters chosen according to Scott's rule (Baddeley & Turner, 2005; Scott, 1992). To compute deviations between 2D densities P_{kl} and Q_{kl} at grid position (k, l) , we used a symmetric version of the Kullback-Leibler divergence derived from information gain (Beck & Schlögl, 1993), i.e.,

$$\Delta_{KLD} = \frac{1}{2} \sum_{k,l} \left(P_{kl} \log \frac{P_{kl}}{Q_{kl}} + Q_{kl} \log \frac{Q_{kl}}{P_{kl}} \right). \quad (1)$$

Second-order statistics (see also the illustrated notes in the Appendix) are based on the pair density $\rho(x_1, x_2)$, which gives the probability $\rho(x_1, x_2) dx_1 dx_2$ of observing points in each of two disks b_1 and b_2 with linear dimensions dx_1 and dx_2 , respectively. Point patterns can be characterized by the pair density, which is typically a function of the pair distance, i.e., $\rho(x_1, x_2) = \rho(r)$ with $r = \|x_1 - x_2\|$, for two arbitrary realizations x_1 and x_2 . Using a kernel-based method, a estimator for the pair density can be written as

$$\hat{\rho}(r) = \sum_{x_1, x_2 \in W}^{\neq} \frac{k(\|x_1 - x_2\| - r)}{2\pi r A_{\|x_1 - x_2\|}}, \quad (2)$$

where $k(\cdot)$ is an appropriate kernel and $A_{\|x_1 - x_2\|}$ denotes an edge correction at distance $\|x_1 - x_2\|$ (Baddeley & Turner, 2005). For numerical computations we used the Epanechnikov kernel (Illian et al., 2008), i.e.,

$$k(x) = \begin{cases} \frac{3}{4h} \left(1 - \frac{x^2}{h^2}\right), & \text{for } -h \leq x \leq h \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

The problem of choosing the bandwidth h appropriately is frequently discussed in the literature (Illian et al., 2008). The bottom line from this discussion is that the behavior of the estimator should be analyzed over a

range of bandwidths. We will run such an analysis below (Figure 2).

The pair correlation function $g(r)$ is a normalization of the pair density with respect to first-order intensity $\hat{\lambda}$, so that the estimator for the pair correlation is given by $\hat{g}(r) = \hat{\rho}(r) / \hat{\lambda}^2$. The interpretation of the pair correlation function for a given point pattern is straightforward. For a random pattern without clustering, the pair correlation function is $\hat{g}(r) \approx 1$ across the full range of distances r . If $\hat{g}(r) > 1$, then pairs of fixations are more abundant than on average at distance r . If $\hat{g}(r) < 1$, then pairs of fixations are less abundant than on average at a distance r . Thus, the pair correlation function $\hat{g}(r)$ measures how selection of a particular point location (i.e., fixation position) is influenced by other fixations at distance r .

Using the inhomogeneous pair correlation function $\hat{g}_{inhom}(r)$, we can remove the first-order inhomogeneity from the second-order spatial statistics, i.e.,

$$\hat{g}_{inhom}(r) = \sum_{x_1, x_2 \in W}^{\neq} \frac{1}{\hat{\lambda}(x_1)\hat{\lambda}(x_2)} \frac{k(\|x_1 - x_2\| - r)}{2\pi r A_{\|x_1 - x_2\|}}. \quad (4)$$

Estimation of $\hat{g}_{inhom}(r)$ involves two steps: First, we estimated the overall intensity $\hat{\lambda}(x)$ for all fixation positions obtained for a given scene. In this procedure we borrow strength from the full set of observations to obtain reliable estimates of the inhomogeneity. Second, we computed the pair correlation function from a single trial with respect to the inhomogeneous density of the full data set.

In case of a given pair correlation function $\hat{g}(r)$, the scalar quantity (Illian et al., 2008)

$$A_g = \int_0^{\infty} (\hat{g}(r) - 1)^2 dr, \quad (5)$$

denoted as PCF deviation in the following, serves as a useful test statistic that quantifies the deviations from randomness for a given point pattern with inhomogeneous density $\hat{\lambda}(x)$. The integral in Equation (5) was evaluated numerically for pair distances r between 0.1° and 5° (image set 1 and 2) and between 0.1° and 3° (Le Meur, Le Callet, Barba, & Thoreau [2006] data; see below).

Results

We conducted an eye-tracking experiment on scene viewing with 35 human observers using 15 object-based natural scenes (image set 1). Resulting gaze data were evaluated using first- and second-order spatial statistics; we found that data exhibit unexpected spatial

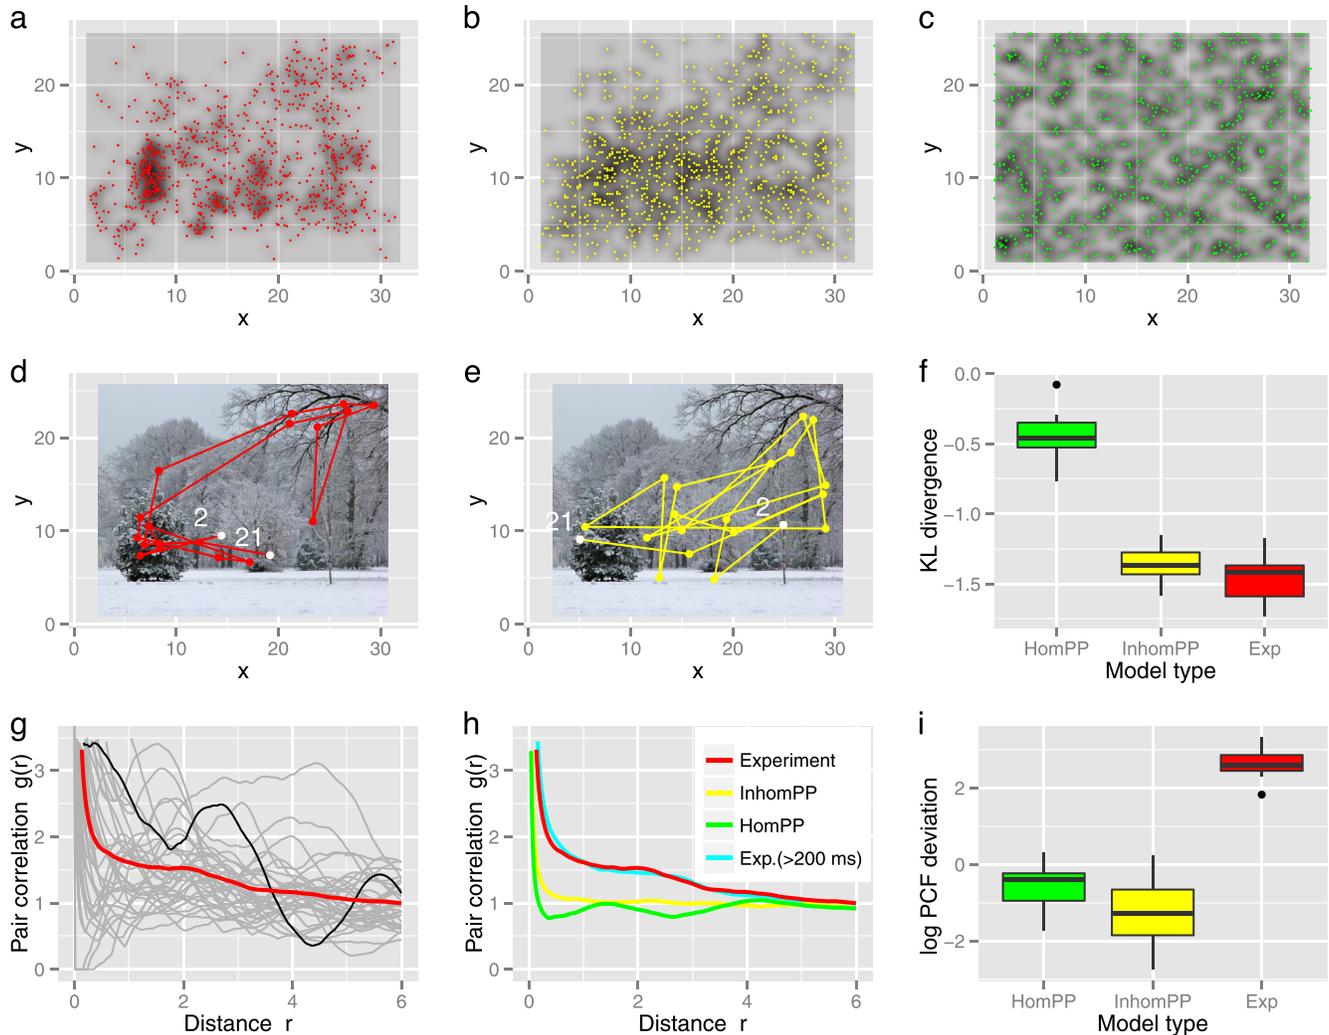


Figure 1. Analysis of pair correlation functions for experimental gaze sequences (image set 1) and for computer-generated surrogate data. (a) Experimental data of gaze positions from human observers (red) and estimated intensity from kernel density estimate (gray levels) for image #2. (b, c) Realizations of gaze positions generated by inhomogeneous and homogeneous point processes, respectively. (d, e) Typical single-trial fixation sequences from experiment (red) and inhomogeneous point process (yellow). (f) Kullback-Leibler divergence (KLD) indicates that the inhomogeneous point process approximates the experimental 2D density of gaze positions. (g) Pair correlation functions (PCFs) for experimental data (single trials: light gray; single trial from (d): black; averaged over trials: red). (h) Mean PCFs for experimental data, inhomogeneous and homogeneous Poisson process. (i) The PCF deviation shows that the experimental data are spatially correlated, while the two surrogate datasets fail to reproduce this statistical pattern.

aggregation (or clustering). We reproduced this finding for a set of 15 abstract natural patterns (image set 2) and for an external dataset (Le Meur et al., 2006) that was made publicly available (Bylinskii, Judd, Durand, Oliva, & Torralba, 2012). Based on these results, we developed a dynamical model for saccadic selection that was evaluated by the spatial-statistics approach introduced in this section.

Spatial statistics and pair correlation function

We began by numerically computing the spatial (2D) density of gaze positions from experimental data

collected for image set 1 (Figure 1a). Fixation positions are indicated by red dots (a total of 930 fixations from 35 observers for image #2). Densities were computed using a 2D kernel density estimator (Baddeley & Turner, 2005; R Core Team, 2013; see Methods) and are visualized by gray shading in the plot. The bandwidth parameter h_{density} for the kernel density estimation was computed according to Scott's rule (Scott, 1992; range from 1.8° to 2.2° for h over all images from image set 1). The obtained 2D density $\hat{\lambda}(x, y)$ is inhomogeneous because of the dependence on position (x, y) . A representative sample trajectory from a single trial is given in Figure 1d, where the second and last fixation of the scanpath is highlighted by white

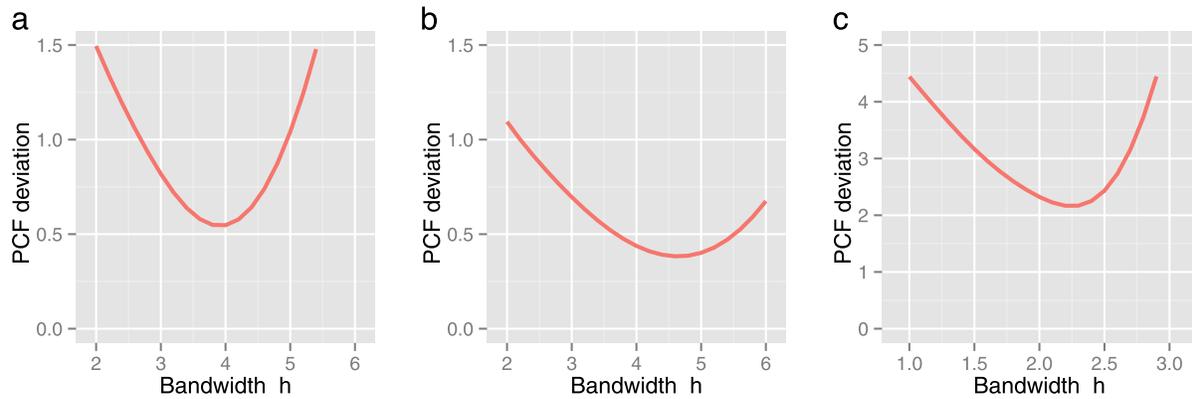


Figure 2. Optimal bandwidth parameter for inhomogeneous pair correlation function (PCF) for three different image sets. For simulated data from the inhomogeneous point process, the PCF deviation Δ_g , Equation (5), was computed as a function of the bandwidth h for the underlying kernel density estimate. The optimal bandwidth corresponds to the position of the minimum of Δ_g . (a) For image set 1, $\hat{h}_1 = 4.0^\circ$. (b) For image set 2, $\hat{h}_2 = 4.6^\circ$. (c) For the Le Meur dataset, $\hat{h}_3 = 2.2^\circ$, due to a much smaller presentation display compared with our experiments.

color and by their serial numbers. The first fixation was omitted, since all trials started at a random position within an image determined by our experimental procedure (see Methods).

The pair correlation function $g(r)$ gives a quantitative summary of interactions in fixation patterns by measuring how distance patterns between fixations differ from what we would expect from independently distributed data (see Appendix). A value of $g(r)$ above 1 for a particular distance r indicates clustering, meaning that there are more pairs of points separated by a distance r than we would expect if fixation locations were statistically independent.

For the estimation of the pair correlation function, short sequences of gaze positions from single experimental trials were considered. However, spatial inhomogeneity of the 2D density was taken into account. To obtain a reliable estimate of the spatial inhomogeneity, the 2D density was estimated from the full data-set (Figure 1a) of all fixations on a given image taken from all participants and trials. It is important to note, however, that in the computation of the kernel density estimate $\hat{\lambda}(x)$ used for the inhomogeneous pair correlation function, Equation (4), an optimal bandwidth parameter h is needed to avoid two possible artifacts: First, if h is very small, then spatial correlations might be underestimated due to overfitting of the inhomogeneity of the density. Second, if h is too large, then spatial correlations might be overestimated, since first-order inhomogeneity is not adequately removed from the second-order spatial statistics. We solved this problem by computing the PCF deviation Δ_g for the inhomogeneous point process for varying values of the bandwidth h (Figure 2). Since the inhomogeneous point process generates uncorrelated

fixations, i.e., $g_{\text{theo}}(r) = 1$, the optimal bandwidth for the dataset corresponds to a minimum of the PCF deviation Δ_g (quantifying the deviation from the ideal value $g(r) = 1$). For image set 1, the optimal value was estimated as $\hat{h}_1 = 4.0^\circ$ (Figure 2a).

Based on this density estimate, we can compute the inhomogeneous pair correlation function $g_{\text{inhom}}(r)$, in which first-order inhomogeneity is removed from the second-order spatial correlations (see Methods). As a result, we obtained pair correlations from individual trials (Figure 1g, gray lines). Deviations from $g_{\text{inhom}}(r) \approx 1$ indicate spatial clustering at a specific distance r . The mean pair correlation function $\bar{g}_{\text{inhom}}(r)$ provides evidence for clustering at small spatial scales with $r < 4^\circ$ (Figure 1g, red line). Such a scale is greater than the foveal zone ($r < 2^\circ$) and might provide an estimate of the size of the *effective* perceptual window in free scene viewing. This result is compatible with earlier findings that the zone of active selection of saccade targets extends beyond the fovea into the parafovea up to eccentricities of 4° (Reinagel & Zador, 1999).

Next, we carried out the same numerical computations for two sets of surrogate data. The surrogate data were generated to test the null hypotheses of complete spatial randomness, both for an inhomogeneous point process with position-dependent intensity $\lambda(x)$ and for a homogeneous point process with constant intensity λ_0 . For the inhomogeneous point process, we sampled from the estimated intensity $\hat{\lambda}(x)$ (Figure 1b), whereas a constant intensity $\hat{\lambda}_0$, obtained from spatial averaging, was used for the homogeneous point process (Figure 1c). Both surrogate datasets are important for checking the reliability of the computation of the pair correlation function for the original data (Figure 1g). First, the inhomogeneous point process gives a flat

mean correlation function with $g(r) \approx 1$ (Figure 1h), which demonstrates the absence of clustering (except for the divergence at very small scales as an effect of numerical computation issues). Thus, the spatial correlations in the experimental data are not a simple consequence of spatial inhomogeneity. Second, the result for the homogeneous point process (Figure 1h) is the same as for the inhomogeneous point process, which indicates that the correction for inhomogeneity needed for computations in Figure 1g does not produce unwanted artefacts due to possible overfitting of the spatial inhomogeneities. We conclude that our experimental data give a clear indication for spatial clustering at length-scales smaller than 4° of visual angle. Additionally, we checked the hypothesis that this effect of spatial clustering might be due to saccadic undershoot and subsequent short correction saccades by excluding all fixations with durations shorter than 200 ms. A related analysis of the PCF indicates no qualitative differences from the original data (Figure 1h).

Results reported so far were obtained for a single, representative image. Over the full set of 15 images, we analyzed the 2D densities using a symmetrized form of the Kullback-Leibler divergence (KLD) based on the concept of information gain (Beck & Schlögl, 1993). For the experimental data, we applied a split-half procedure (first half of participants vs. second half of participants) and computed the KLD between the two experimental densities. The corresponding KLD values demonstrate that the inhomogeneous point process reproduces the 2D density (Figure 1f), while the homogeneous point process clearly fails to approximate the systematic inhomogeneity in the image. Model type had a significant effect on KLD, $\chi^2(2) = 105.4$, $p < 0.01$. Contrasts revealed that (1) spatially inhomogeneous data-sets were different from the homogeneous data, $b = 0.319$, $t(28) = 22.85$, $p < 0.01$, and that (2) experimental data were significantly different from inhomogeneous surrogate data, $b = 0.062$, $t(28) = 2.56$, $p = 0.016$.

Both sets of surrogate data were submitted to an analysis of the pair correlation function, where the deviations from $g(r) \approx 1$ were computed to obtain a PCF measure indicating the amount of spatial correlation averaged over distances (see Methods). Results indicate that the surrogate data produce—as designed—uncorrelated gaze positions (low PCF deviation), while the experimental data by human observers exhibit spatially correlated gaze positions (Figure 1i). Model type had a significant effect on PCF, $\chi^2(2) = 98.4$, $p < 0.01$.

To investigate the reliability of our finding of spatial clustering on short-length scales during natural scene viewing, we investigated two other datasets using the same procedures with an optimal bandwidth adjusted

for each dataset (Figure 2a–c). First, we compared the mean pair correlation functions per image for natural object-based scenes (image set 1; Figure 3a, d) and abstract natural patterns (image set 2; Figure 3b, e). While variability of the pair correlation function might be greater for the abstract scenes than for the object-based scenes (Figure 3d, e; gray lines) the mean pair correlation functions for the images sets (Figure 3d, e; red lines) are very similar. This result indicates that the clustering on short length scales is a robust phenomenon, which does not seem to depend sensitively on scene content. Second, we analyzed a publicly available dataset (Le Meur et al., 2006) from the MIT saliency benchmark (Bylinskii et al., 2012; Judd, Durand, & Torralba, 2012) consisting of 40 participants who viewed 27 color images for 15 s. For these data, the spatial scale of the presentation display for the images (Figure 3c) was considerably smaller than in our experiment. Consequently, we observed pair correlation functions that indicate clustering on smaller scales ($< 1.5^\circ$, Figure 3f). Thus, while scene content does not seem to exert a strong influence on spatial correlations, spatial scale of the image modulates the spatial scale of the clustering of fixations.

We conclude that second-order spatial statistics obtained for the experimental data are significantly different from stochastic processes implementing the assumption of spatial randomness. Furthermore, the mere presence of spatial inhomogeneity in the experimental data cannot explain by itself the observed spatial correlations, which is evident in the results for the inhomogeneous point process. While inhibition-of-return (Klein, 2000) has been discussed frequently as one of the key principles added to saliency maps for saccadic selection (Itti & Koch, 2001), spatial clustering of gaze positions is an additional statistical property that is highly informative on mechanisms of gaze planning, but has been neglected so far. Next, we use these results to develop and test a dynamical model for saccade generation that uses activation field dynamics to reproduce spatial statistics of first- and second-order.

A dynamical model of saccade generation

A key assumption for the model we propose is the combination of two neural activation maps to implement dynamical principles for saccadic selection. First, a fixation map $f(x, y; t)$ is keeping track of the sequence of fixations by inhibitory tagging (Itti & Koch, 2001). Second, an attention map $a(x, y; t)$ that is driven by early visual processing controls the distribution of attention. Physiologically, the assumption of the dynamical maps is supported by the presence of an allocentric motor map of visual space in the primate entorhinal cortex (Killian, Jutras, &

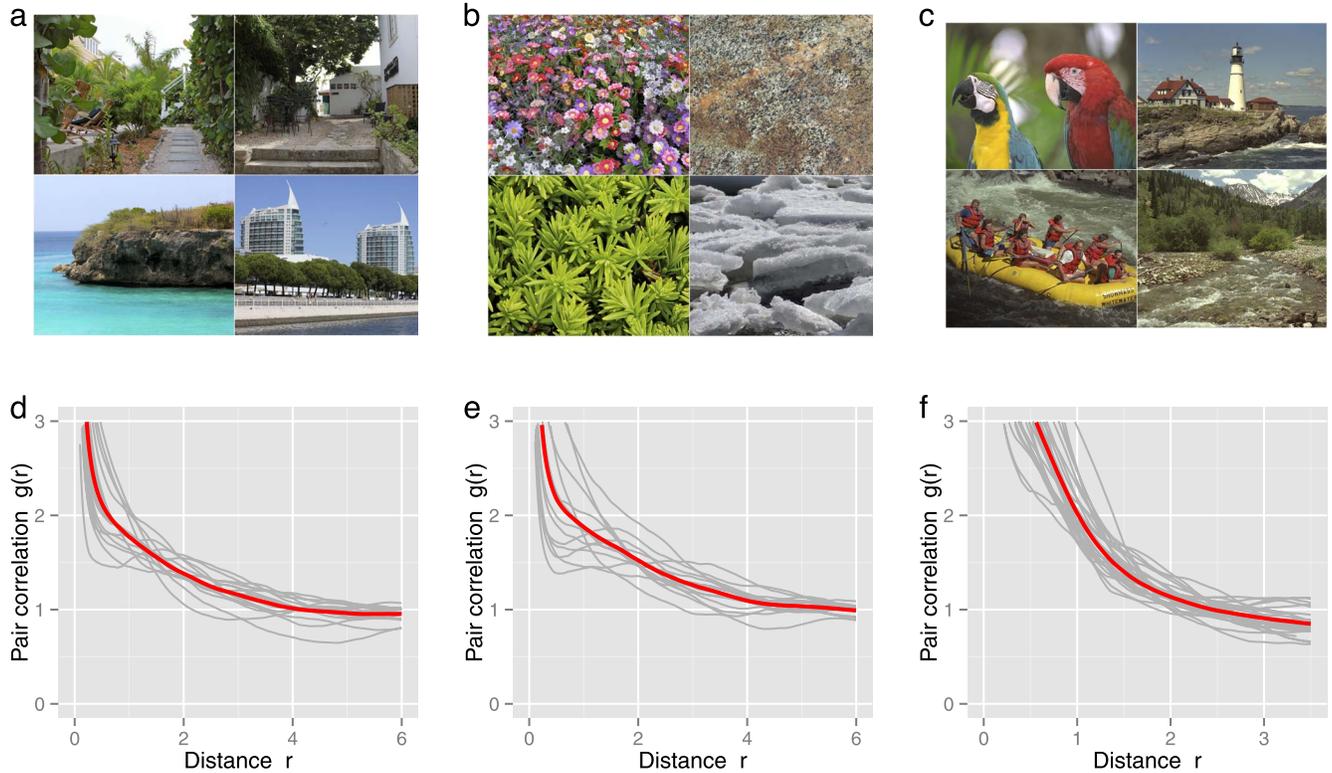


Figure 3. Comparison of the mean pair correlation functions for three different image sets: (a) Natural object-based scenes. (b) Abstract natural patterns. (c) Natural scenes from the Le Meur et al. dataset. Mean pair correlation functions are similar for object-based (d) and abstract (e) natural scenes. The presentation of images on a smaller display in the Le Meur dataset compared to our experiments (f) resulted in a smaller length scale of spatial clustering.

Buffalo, 2012). Moreover, this map is spatially discrete (Stensola et al., 2012) and serves as a biological motivation for the fixation and attention maps in our model.

We implemented activation maps for attention and fixation (inhibitory tagging) on a discrete square lattice of dimension $L \times L$. Lattice points (i, j) have equidistant spatial positions (x_i, y_j) for $i, j = 1, \dots, L$, where $x_i = x_0 + i\Delta x$ and $y_j = y_0 + j\Delta y$. As a consequence, attention and fixation maps are implemented in spatially discrete forms, $\{a_{ij}(t)\}$ and $\{f_{ij}(t)\}$, respectively. For the numerical simulations, time was discretized in steps of $\Delta t = 10$ ms with $t = k \times \Delta t$ and $k = 0, 1, 2, \dots, T$.

If the observer's gaze is at position (x_g, y_g) at time t , then a position-dependent activation change $F_{ij}(x_g, y_g)$ and a global decay proportional to the current activation $-\omega f_{ij}(t)$ are added to all lattice positions to update the activation map at time $t + 1$, i.e.,

$$f_{ij}(t + 1) = F_{ij}(x_g, y_g) + (1 - \omega)f_{ij}(t), \quad (6)$$

where the activation change $F_{ij}(x_g, y_g) \equiv F_{ij}(t)$ is implicitly time-dependent because of the time-dependence of gaze positions, $x_g(t), y_g(t)$. The constant $\omega < 1$ determines the strength of the decay of activation. For the spatial distribution of the activation change $F_{ij}(t)$

we assume a Gaussian profile, i.e.,

$$F_{ij}(t) = \frac{R_0}{\sqrt{2\pi}\sigma_0} \exp\left(-\frac{(x_i - x_g(t))^2 + (y_j - y_g(t))^2}{2\sigma_0^2}\right), \quad (7)$$

with the free parameters σ_0 and R_0 controlling the spatial extent of the activation change and the strength of the activation change, respectively. In our model, the build-up of activation in the fixation map is a mechanism of inhibitory tagging (Itti & Koch, 2001) to reduce the amount of refixations on recently visited image patches.

For the attention map $a_{ij}(t)$ we assume similar dynamics, however, the width of Gaussian activation change $A_{ij}(t)$ is assumed to be proportional to the static saliency map $\{\phi_{ij}\}$. The updating rule for the attention map is given by

$$a_{ij}(t + 1) = \frac{\phi_{ij}A_{ij}(t)}{\sum_{kl} \phi_{kl}A_{kl}(t)} + (1 - \rho)a_{ij}(t), \quad (8)$$

with decay constant $\rho \ll 1$. As a result, the saliency map ϕ_{ij} is accessed locally through a Gaussian aperture with size σ_1 and scale parameter R_1 , similar to Equation

(7). Using the local read-out mechanism, information is provided for the attention map to identify regions of interest for eye guidance.

The fixation map monitors recently visited fixation locations by increasing local activation at the corresponding lattice points (Engbert et al., 2011; Freund & Grassberger, 1992). If the observer's gaze position is at a position corresponding to lattice position (i, j) , then a position-dependent activation change F_{ij} in the form of a Gaussian profile is added locally in each time step, while a global decay proportional to the current activation is applied to all lattice positions. The width of the Gaussian activation σ_0 and the decay ω are the two free parameters controlling activation in the fixation map. For the attention map $a_{ij}(t)$ we assume similar dynamics, including local increase of activation with size σ_1 and global decay ρ . However, the amount of activation change A_{ij} is assumed to be proportional to the time-independent saliency map ϕ_{ij} , so that the local increase of activation is $+ \phi_{ij} A_{ij} / \sum_{kl} A_{kl}$.

Our modeling assumptions are related to specific hypotheses on model parameters. We expect that the size of the Gaussian profile for the attention map is larger than the corresponding size of the fixation map, $\sigma_1 > \sigma_0$, since attention is the process driving eye movements into new regions of visual space, while the inhibitory tagging process should be more localized. A similar expectation can be formulated on the decay constants. Since inhibitory tagging is needed on a longer time scale as a foraging facilitator, we expect a slower decay in the fixation map compared to the attention map, i.e., $\omega < \rho$.

Next, we assume that, given a saccade command at time t , both maps are evaluated to select the next saccade target. First, we apply a normalization of both attention and fixation maps as a general neural principle to obtain relative activations (Carandini & Heeger, 2011). Second, we introduce a potential function as the difference of the normalized maps,

$$u_{ij}(t) = - \frac{[a_{ij}(t)]^\lambda}{\sum_{kl} [a_{kl}(t)]^\lambda} + \frac{[f_{ij}(t)]^\gamma}{\sum_{kl} [f_{kl}(t)]^\gamma}, \quad (9)$$

where the exponents λ and γ are free parameters. However, a value of $\lambda = 1$ is a necessary boundary condition to obtain a model that accurately reproduces the densities of gaze positions. In a qualitative analysis of the model (see Appendix), pilot simulations showed that γ is an important control parameter determining spatial correlations, where $\gamma \approx 0.3$ was used to reproduce spatial correlations observed in our experimental data.

The potential $u_{ij}(t)$, Equation (9), can be positive or negative at position (i, j) . Lattice positions with a positive potential, $u_{ij} > 0$, are excluded from saccadic

selection, since corresponding regions were visited recently with high probability. Among the lattice positions with negative activations, we implemented stochastic selection of saccade targets proportional to relative activations, also known as Luce's choice rule (Luce, 1959). We implemented this form of stochastic selection from the set $S = \{(i, j) | u_{ij} < 0\}$, where the probability $\pi_{ij}(t)$ to select lattice position (i, j) at time t as the next saccade target is given by

$$\pi_{ij}(t) = \max \left(\frac{u_{ij}(t)}{\sum_{(k,l) \in S} u_{kl}(t)}, \eta \right). \quad (10)$$

The noise term η is an additional parameter controlling the amount of noise in target selection.

Numerical simulations of the model

Our computational modeling approach to saccadic selection has been developed to propose a minimal model that captures the types of spatial statistics observed in experimental data. After fixing model parameters, all that is needed to run the model on a particular image is a 2D density estimate of gaze patterns (from experimental data) or a corresponding 2D prediction from one of the available saliency models (Borji & Itti, 2013). To reduce computational complexity in the current study and to exclude potential mismatches between data and model simulations due to the saliency models, we run the model on experimentally realized densities of gaze positions, which is equivalent to assuming an exact saliency model. However, our modeling approach is compatible with future dynamical saliency models that provide time- and position-dependent saliency during a sequence of gaze shifts, thus our model introduces a general dynamical framework and is not tied to using empirical data.

The numerical values of the five model parameters were estimated from experimental data recorded for the first five images of natural object-based scenes (image set 1) using a genetic algorithm approach (Mitchell, 1998; see Appendix, Table 1). The remaining 10 images of image set 1 were used for model evaluations—as the 15 images of image set 2. The objective function for parameter estimation was based on evaluation of first-order statistics (2D density of gaze positions) and the distribution of saccade lengths. In agreement with our first expectation, the estimated optimal values for the spatial extent of the inhibitory tagging process in the fixation map, $\sigma_0 = 2.2^\circ$, is considerably smaller than the corresponding size of the build-up function for the attention map, $\sigma_1 = 4.9^\circ$. Our second expectation was related to the decay constants, which turned out to be

Parameter	Symbol	Mean	Error	Min	Max	Reference
Fixation map						
Activation span [°]	σ_0	2.16	0.11	0.3	10.0	Eq. (7)
Decay	$\log_{10} \omega$	-4.03	0.28	-5.0	-1.0	Eq. (6)
Attention map						
Activation span [°]	σ_1	4.88	0.25	0.3	10.0	Eq. (8)
Decay	$\log_{10} \rho$	-1.18	0.08	-3.0	-1.0	Eq. (8)
Target selection						
Additive noise	$\log_{10} \eta$	-4.04	0.07	-9.0	-3.0	Eq. (10)

Table 1. Model parameters.

larger for the attention map, $\rho = 0.066$, than for the fixation map, $\omega = 9.3 \times 10^{-5}$, so ρ was greater than ω , again as expected. Finally, the noise level in the target map is $\eta = 9.1 \times 10^{-5}$.

An example for the simulation of the model demonstrates the interplay between inhibitory processes from the fixation map and the attention map during gaze planning (Figure 4). The fixation map builds up activation at fixated lattice positions (yellow to red), while the attention map identifies new regions of interest for saccadic selection (blue). These simulations show on a qualitative level how the model implements the interplay of the assumed mechanisms of

inhibitory tagging and saccadic selection of gaze positions (see Supplementary Video).

To investigate model performance qualitatively, we ran simulations for one image (image #6, 930 fixation, Figure 5a) and obtained a number of fixations similar to the experimental data (882 fixations, Figure 5b). Single-trial scanpaths from experiments and simulations are shown additionally in Figure 5a and 5b, respectively). The resulting distributions of saccade lengths indicate that our dynamical model is in good agreement with experimental data, while the two surrogate datasets (homogeneous and inhomogeneous point processes) fail to reproduce the distribution (Figure 5c). An analysis of the pair correlation

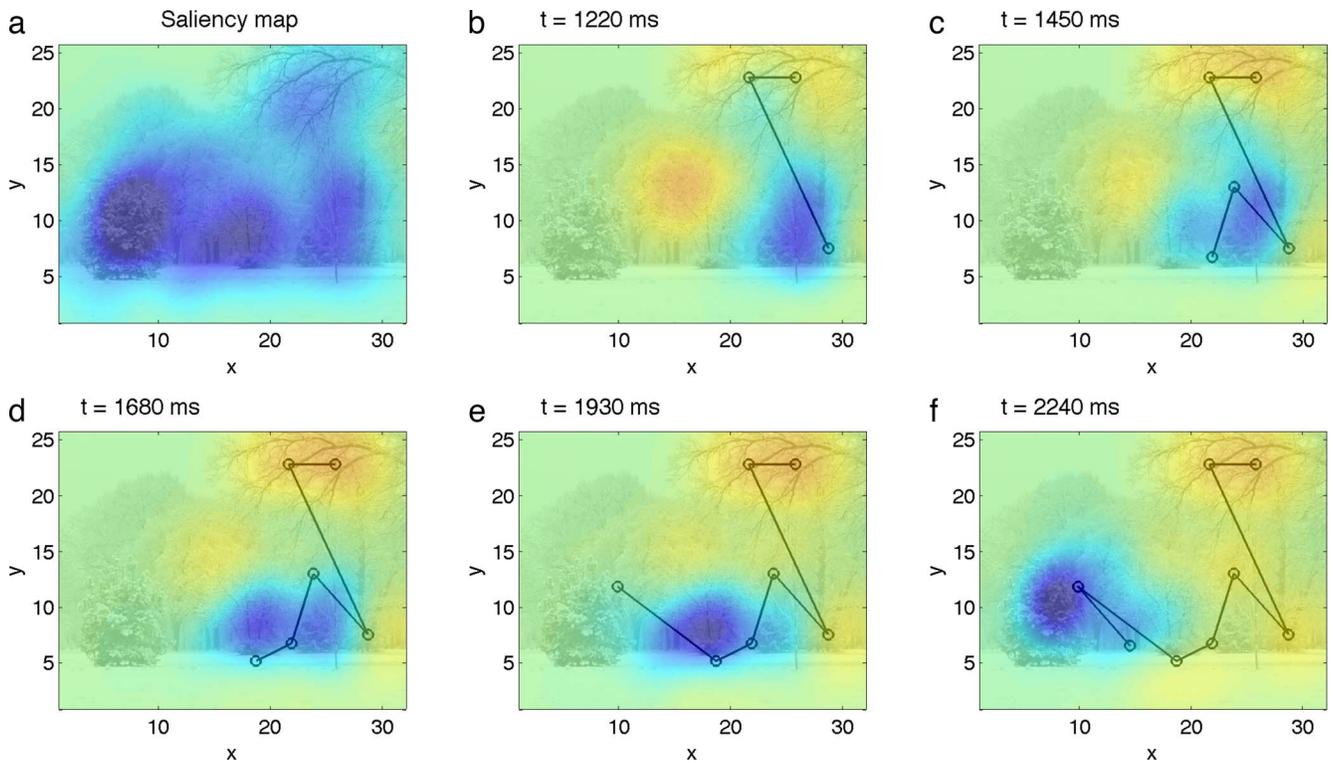


Figure 4. Illustration of a simulated sequence of gaze positions and the activation dynamics of the model. (a) The density of gaze positions (empirical saliency) is used as a proxy for a computed saliency map that drives activation in the attention map. (b–f) Sequence of snapshots of the potential (blue = low, yellow = high). Note that blue color indicates density of fixations in (a), while it refers to low values of the potential in (b)–(f).

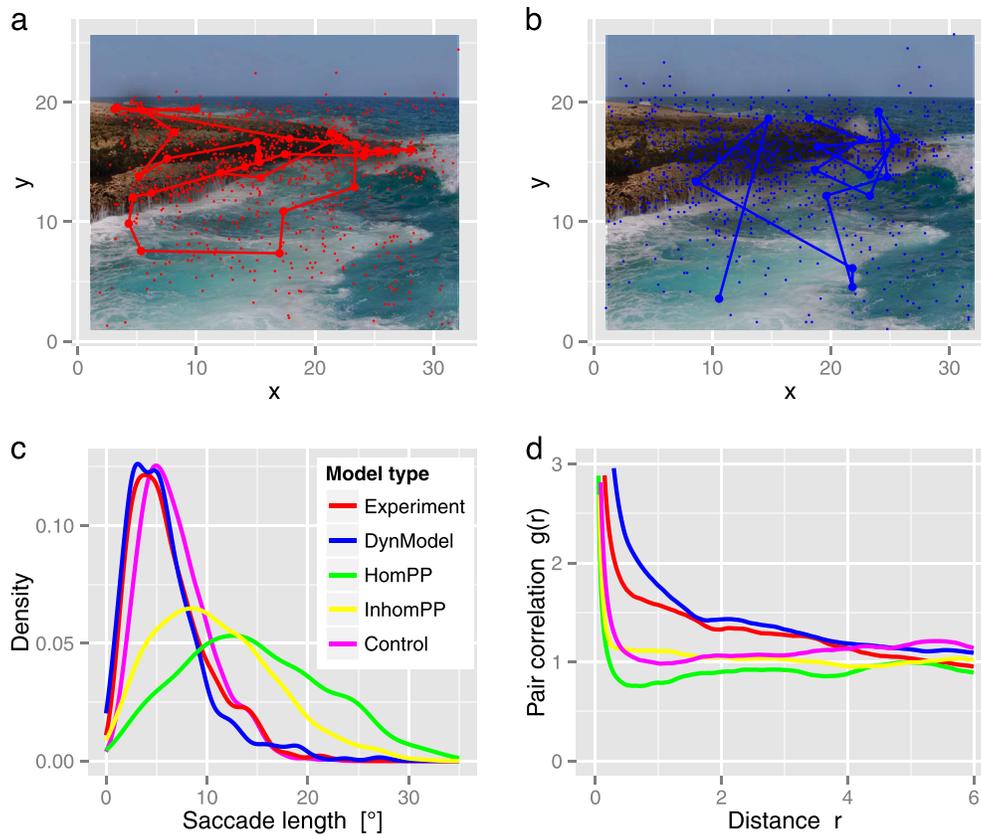


Figure 5. Distribution of saccade lengths and pair correlation functions from model simulations (image #6). (a) Experimental distribution of gaze positions (red dots) and a representative sample trial (red lines). (b) Corresponding plot of simulated data obtained from our dynamical model (blue dots). A single-trial simulation is highlighted (blue line). (c) Distributions of saccade lengths for experimental data (red), dynamical model (blue), homogeneous point process (green), inhomogeneous point process (yellow), and a statistical control model (magenta). (d) Pair correlation functions for the different models. The dynamical model (blue line) produces spatial correlations similar to the experimental data (red line).

functions indicates that the spatial correlations present in the experimental data were approximated by the dynamical model (Figure 5d), however, the two surrogate datasets representing uncorrelated sequences by construction produce qualitatively different spatial correlations.

To investigate the influences of saccade-length distributions on pair correlations, we constructed another statistical control model that had access to the image-specific saccade-length distribution. It is important to note that this model was not introduced as a competitor to the dynamical model, which is able to *predict* saccade-length distributions. The statistical control model approximated the distribution of saccade lengths l and 2D densities of gaze positions x by sampling from the joint probability distribution $p(x, l)$ under the assumption of statistical independence of saccade lengths and gaze positions, i.e., $p(x, l) = p(x)p(l)$. This model, by construction, approximates the distribution of saccade lengths and 2D density of gaze positions (Figure 5c). The simulations indicate, how-

ever, that even the combination of inhomogeneous density of gaze positions and non-normal distribution of saccade lengths used by the statistical control model cannot explain spatial correlations in the experimental data characterized by the pair correlations function (Figure 5d).

For the statistical analysis of model performance on new images, we carried out additional numerical simulations. We fitted model parameters to data obtained for the first five images only (see above) and predicted data for the remaining 10 images by the new simulations to isolate parameter estimation from model evaluation (calculating test errors rather than training errors).

Our simulations show that the dynamical model predicted the 2D density of gaze positions accurately (Figure 6a). The obtained KLD values for the model (blue) were comparable to KLD values calculated by the split-half procedure for the experimental data (red) and to the KLD values obtained for the statistical control models (green = Homogeneous point process,

yellow = Inhomogeneous point process, magenta = Control model). Model type had a significant effect on KLD, $\chi^2(4) = 109.4$, $p < 0.01$. Post hoc comparisons indicated significant effects between all models ($p < 0.01$) except for the comparison between experimental data and the dynamical model ($p = 0.298$) and for the comparison between dynamical model and inhomogeneous point process ($p = 0.679$). In an analysis of the PCF estimated from the same set of simulated data (Figure 6b), the dynamical model (blue) produced deviations from an uncorrelated point process that are in good agreement with the experimental data (red). Model type had a significant effect on PCF deviation, $\chi^2(4) = 91.6$, $p < 0.01$. Post hoc comparisons indicated significant effects for all comparisons ($p < 0.01$) except for the comparison between experimental data and the dynamical model ($p = 0.990$) and between homogeneous and inhomogeneous point processes ($p = 0.996$).

To check the reliability of the results, we performed all corresponding calculations for the images showing abstract natural scenes (image set 2). For these simulations, we used the set of model parameters fitted to the first five images of image set 1, which corresponds to the hypothesis that scene content (object-based scenes vs. abstract natural patterns) does not have a strong impact on spatial correlations in the scanpath data. For the simulated densities (Figure 6c), we reproduced the statistical results from image set 1. Again, model type had a significant effect on KLD, $\chi^2(4) = 162.6$, $p < 0.01$. Post hoc comparisons indicated significant effects between all models ($p < 0.01$) except for the comparison between experimental data and the dynamical model ($p = 0.251$), for the comparison between dynamical model and inhomogeneous point process ($p = 0.784$), and for the comparison between inhomogeneous point process and experimental data ($p = 0.012$). For the pair correlation function (Figure 6d), model type had a significant effect, $\chi^2(4) = 140.6$, $p < 0.01$, and post hoc comparisons indicated significant effects for all comparisons ($p < 0.01$) except for the comparison between experiment and dynamical model ($p = 0.453$) and between homogeneous and inhomogeneous point processes ($p = 0.700$). Thus the main statistical results obtained from image set 1 were reproduced for images of abstract natural patterns (image set 2). These results lend support to the hypothesis that scene content does not have a strong influence on second-order spatial statistics of gaze patterns.

Thus, the dynamical model performed better than any of the statistical models in predicting the average pair correlations. Although one of our statistical control models generated data by using image-specific saccade-length information in addition to the 2D density of gaze position, it could not predict the spatial correlations as accurately as the dynamical model that

was uninformed about the image-specific saccade-length distribution.

Discussion

Current theoretical models of visual attention allocation in natural scenes are limited to the prediction of first-order spatial statistics (2D densities) of gaze patterns. We were interested in attentional dynamics that can be characterized by spatial interactions (as found in the second-order statistics). Using the theory of spatial point processes, we discovered that gaze patterns can be characterized by clustering at small length scales, which cannot be explained by spatial inhomogeneity of the 2D density. We proposed and analyzed a model based on dynamical activation maps for attentional selection and inhibitory control of gaze positions. The model reproduced 2D densities of gaze maps (first-order statistics) and distributions of saccade lengths as well as pair correlations (second-order spatial statistics).

Spatial statistics

While research on the computation of visual saliency has been a highly active field of research (Borji & Itti, 2013), there is currently a lack of computational models for the generation of scanpaths on the basis of known saliency. Inhibition-of-return (Klein, 2000) has been proposed as a key principle to prevent continuing refixation within regions of highest saliency. However, our analysis of the pair correlation function demonstrates that saccadic selection at small length scales is dominated by spatial clustering. Thus, our findings are highly compatible with the view that inhibition-of-return cannot easily be observed in eye-movement behavior in natural scene viewing experiments (Smith & Henderson, 2011). However, the spatial correlations can be exploited to investigate dynamical rules underlying attentional processing in the visual system. Our experimental data show a clear effect of spatial clustering for length scales shorter than about 4° . However, these results are incompatible with the current theory of saliency-based attention allocation combined with an inhibition-of-return mechanism. Future simulations will investigate the predictive power of the model when saliency models are used as input.

Modeling spatial correlations

In a biologically plausible computational model of saccade generation, a limited perceptual span needs to

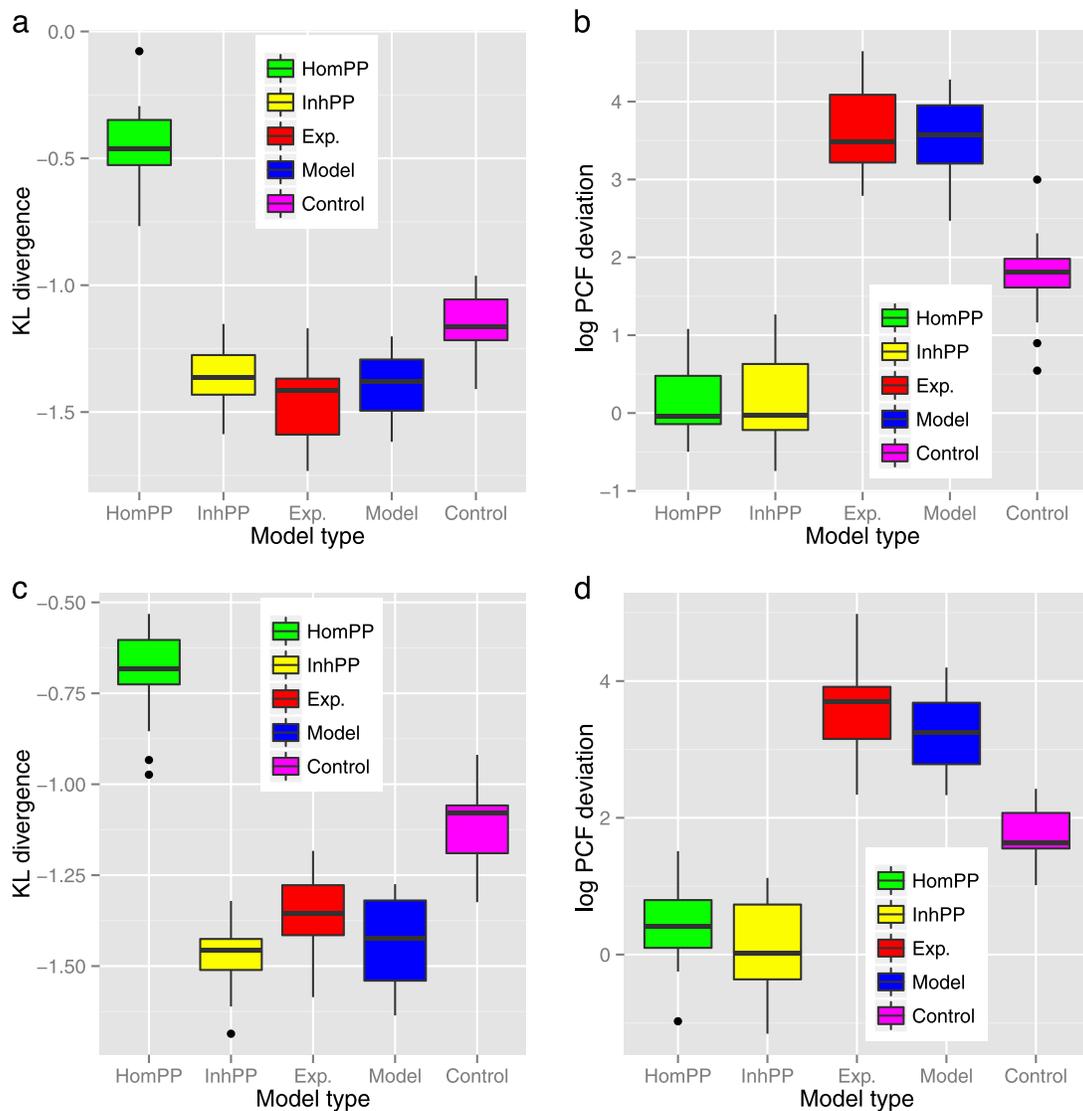


Figure 6. Model predictions on images not used for parameter estimation. Predicted data were generated from the dynamical model for the 10 images of image set 1 not used for parameter estimation (a, b) and for all 15 images from image set 2 (c, d). (a) Model simulations of the KLD for experiments on image set 1 (10 images not used for model parameter estimation) and dynamical model and three different statistical models. For the experimental data, a split-half procedure was applied to compute KLD. (b) Corresponding PCF deviations for the same model-generated and experimental data on image set 1. (c) KLD measures for image set 2. (d) PCF deviation for image set 2.

be implemented for attentional selection (Findlay & Gilchrist, 2012). We addressed this problem by assuming a Gaussian read-out mechanism with local retrieval from the saliency map through a limited aperture, analogous to the limited extent of high-fidelity information uptake through the fovea. We used the experimentally observed density of fixations as a proxy for visual saliency.

First, our results indicate that a very limited attentional span (Gaussian with standard deviation parameter $\sim 4.9^\circ$) of about twice the size of the activation mechanism for tracking the gaze positions ($\sim 2.2^\circ$) is sufficient for saccade planning. This atten-

tional span is efficient, however, since the combination of fixation and attention maps in our model actively drives the model's gaze position to new salient regions computed via normalization of activations (Carandini & Heeger, 2011).

Second, our model correctly predicted spatial clustering of gaze positions at small-length scales. The pair correlation function indicates that there is a pronounced contribution by refixations very close to the current gaze position. This effect is compatible with the distribution of saccade lengths, however, a statistical control model that generated data from statistically independent probabilities of 2D density and

saccade lengths could not reproduce the pair correlations adequately.

Implications for computational models of active vision

During active vision, our visual system relies on frequent gaze shifts to optimize retinal input. Using a second-order statistical analysis we demonstrated that spatial correlations across scanpaths might provide important constraints for computational models of eye guidance. In our dynamical model, spatial clustering at small scales is the result of two principles. First, the fixation map is driven by an activation function with a small spatial extent ($\sim 2.5^\circ$). Second, the time-scale of activation build-up in the fixation map is slow compared to the build-up of activation in the attention map. Both mechanisms permit refixations at positions very close to the current gaze position before the system moves on to new regions of visual space.

Limitations of the current approach

The current work focused on spatial statistics of gaze patterns and we proposed and analyzed dynamical mechanisms of eye guidance in scene viewing. In our model, a Gaussian read-out mechanism for the static empirical saliency map was implemented as a simplification. A more biologically plausible combination of our model of eye guidance with a dynamical saliency model (Borji & Itti, 2013) is a natural extension of the current framework, and the development of such a model is work in progress in our laboratories. Clearly, the current modeling architecture is not limited to input from static saliency maps.

Another simplification is related to the timing of saccades (Nuthmann, Smith, Engbert, & Henderson, 2010). In the current version of our model, we implemented random timing and sampled fixation durations randomly from a predefined distribution. More adequate models of fixation durations, however, will need to include interactions of processing difficulty between fovea and periphery (Laubrock, Cajar, & Engbert, 2013).

Supplemental information

This work includes a supplemental video animation of the model simulations. Experimental data on fixation patterns and computer code for statistical analysis and model simulations will be made available via the Potsdam Mind Research Respository (PMR2, <http://read.psych.uni-potsdam.de/>).

Keywords: scene perception, eye movements, attention, saccades, modeling, spatial statistics

Acknowledgments

This work was supported by Bundesministerium für Bildung und Forschung (BMBF) through the Bernstein Computational Neuroscience Programs Berlin (Project B3, FKZ: 01GQ1001F and FKZ: 01GQ1001B to R. E. and F. A. W., resp.) and Tübingen (FKZ: 01GQ1002 to F.A.W.) and by Deutsche Forschungsgemeinschaft (grants EN 471/13–1 and WI 2103/4–1 to R. E. and F. A. W., resp.).

Commercial relationships: none.

Corresponding author: Ralf Engbert.

Email: ralf.engbert@uni-potsdam.de.

Address: Cognitive Science Program & Department of Psychology, University of Potsdam, Potsdam, Germany.

References

- Baddeley, A., & Turner, R. (2005). spatstat: An R package for analyzing spatial point patterns. *Journal of Statistical Software*, *12*(6), 1–42. Available from <http://www.jstatsoft.org/>
- Barthelmé, S., Trukenbrod, H., Engbert, R., & Wichmann, F. (2013). Modeling fixation locations using spatial point processes. *Journal of Vision*, *13*(12), 1, 1–34, <http://www.journalofvision.org/content/13/12/1>, doi:10.1167/13.12.1. [PubMed] [Article]
- Beck, C., & Schlögl, F. (1993). *Thermodynamics of chaotic systems*. Cambridge, UK: Cambridge University Press.
- Borji, A., & Itti, L. (2013). State-of-the-art in visual attention modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *35*(1), 185–207.
- Bylinskii, Z., Judd, T., Durand, F., Oliva, A., & Torralba, A. (2012). MIT Saliency Benchmark. Available at <http://saliency.mit.edu>
- Carandini, M., & Heeger, D. J. (2011). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, *13*, 51–62.
- Engbert, R. (2012). Computational modeling of collicular integration of perceptual responses and attention in microsaccades. *The Journal of Neuroscience*, *32*(23), 8035–8039.
- Engbert, R., & Kliegl, R. (2003). Microsaccades

- uncover the orientation of covert attention. *Vision Research*, 43, 1035–1045.
- Engbert, R., & Mergenthaler, K. (2006). Microsaccades are triggered by low retinal image slip. *Proceedings of the National Academy of Sciences, USA*, 103, 7192–7197.
- Engbert, R., Mergenthaler, K., Sinn, P., & Pikovsky, A. (2011). An integrated model of fixational eye movements and microsaccades. *Proceedings of the National Academy of Sciences, USA*, 108, E765–E770.
- Findlay, J. M., & Gilchrist, I. D. (2012). Visual attention—a fresh look. *Psychologist*, 25(12), 900–902.
- Findlay, J. M., & Walker, R. (1999). A model of saccade generation based on parallel processing and competitive inhibition. *Behavioral and Brain Sciences*, 22(4), 661–674.
- Freund, H., & Grassberger, P. (1992). The red queen's walk. *Physica A*, 190, 218–237.
- Illian, J., Penttinen, A., Stoyan, H., and Stoyan, D. (2008). *Statistical analysis and modelling of spatial point patterns*. New York: Oxford University Press.
- Itti, L., & Koch, C. (2001). Computational modeling of visual attention. *Nature Reviews Neuroscience*, 2, 1–11.
- Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20, 1254–1259.
- Jones, L. A., & Higgins, G. C. (1947). Photographic granularity and graininess. III. Some characteristics of the visual system of importance in the evaluation of graininess and granularity. *Journal of the Optical Society of America*, 37, 217–263.
- Judd, T., Durand, F., & Torralba, A. (2012). A benchmark of computational models of saliency to predict human fixations. In MIT Tech Report no. MIT-CSAIL-TR-2012-001. Available at <http://dspace.mit.edu/handle/1721.1/68590>
- Kienzle, W., Franz, M. O., Scholköpf, B., & Wichmann, F. A. (2009). Center-surround patterns emerge as optimal predictors for human saccade targets. *Journal of Vision*, 9(5), 7 1–15, <http://www.journalofvision.org/content/9/5/7>, doi:10.1167/9.5.7. [PubMed] [Article]
- Killian, N. J., Jutras, M. J., & Buffalo, E. A. (2012). A map of visual space in the primate entorhinal cortex. *Nature*, 491, 761–764.
- Klein, R. M. (2000). Inhibition of return. *Trends in Cognitive Science*, 4, 138–147.
- Laubrock, J., Cajar, A., & Engbert, R. (2013). Control of fixation duration during scene viewing by interaction of foveal and peripheral processing. *Journal of Vision*, 13(12), 11, 1–20, <http://www.journalofvision.org/content/13/12/11>, doi:10.1167/13.12.11. [PubMed] [Article]
- Law, R., Illian, J., Burslem, D. F. R. P., Gratzner, G., Gunatilleke, C. V. S., & Gunatilleke, I. A. U. N. (2009). Ecological information from spatial patterns of plants: Insights from point process theory. *Journal of Ecology*, 97, 616–626.
- Le Meur, O., Le Callet, P., Barba, D., & Thoreau, D. (2006). State-of-the-art in visual attention modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(5), 802–817.
- Levi, D. M. (2008). Crowding—an essential bottleneck for object recognition: A mini-review. *Vision Research*, 48, 635–654.
- Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.
- Ludwig, C. J., Farrell, S., Ellis, L. A., & Gilchrist, I. D. (2009). The mechanism underlying inhibition of saccadic return. *Cognitive Psychology*, 59(2), 180–202.
- Mitchell, M. (1998). *An introduction to genetic algorithms*. Cambridge, MA: MIT Press.
- Nuthmann, A., Smith, T. J., Engbert, R., & Henderson, J. M. (2010). State-of-the-art in visual attention modeling. *Psychological Review*, 117, 382–405.
- R Core Team. (2013). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <http://www.R-project.org/>
- Reinagel, P., & Zador, A. M. (1999). Natural scene statistics at the centre of gaze. *Network: Computation in Neural Systems*, 10(4), 341–350.
- Rosenholtz, R., Huang, J., & Ehinger, K. A. (2012). Rethinking the role of top-down attention in vision: effects attributable to a lossy representation in peripheral vision. *Frontiers in Psychology*, 3(13), 1–14.
- Scott, D. W. (1992). *Multivariate density estimation. Theory, practice and visualization*. New York: Wiley.
- Smith, T. J., & Henderson, J. M. (2011). Looking back at Waldo: Oculomotor inhibition of return does not prevent return fixations. *Journal of Vision*, 11(1), <http://www.journalofvision.org/content/11/1/3>, doi:10.1167/11.1.3. [PubMed] [Article]
- Stensola, H., Stensola, T., Solstad, T., Friland, K., Moser, M.-B., & Moser, E. I. (2012). The entorhinal grid map is discretized. *Nature*, 492, 72–78.

- Tatler, B. W., Baddeley, R. J., & Vincent, B. T. (2006). The long and the short of it: Spatial statistics at fixation vary with saccade amplitude and task. *Vision Research*, 46(12), 1857–1862.
- Tatler, B. W., & Vincent, B. T. (2009). The prominence of behavioural biases in eye guidance. *Visual Cognition*, 17(6–7), 1029–1054.
- Torralba, A., Oliva, A., Castelhano, M. S., & Henderson, J. M. (2006). Contextual guidance of eye movements and attention in real-world scenes: The role of global features in object search. *Psychological Review*, 113, 766–786.
- Trukenbrod, H. A., & Engbert, R. (2014). ICAT: A computational model for the adaptive control of fixation durations. *Psychonomic Bulletin & Review*, 21(4), 907–934.
- Tsotsos, J. K., Culhane, S. M., Wai, W. Y. K., Lai, Y., Davis, N., & Nuflo, F. (1995). Modeling visual-attention via selective tuning. *Artificial Intelligence*, 78, 507–545.
- Wilmington, N., Harst, S., Schmidt, N., & König, P. (2013). Saccadic momentum and facilitation of return saccades contribute to an optimal foraging strategy. *PLoS Computational Biology*, 9(1), e1002871.

Appendix

Estimation of model parameters

Some of the free parameters of the model were set to fixed values to reduce the number of free parameters and to facilitate parameter estimation. First, saccade timing was outside the primary scope of the current work. Time intervals between two decisions for saccadic eye movements are drawn from a gamma distribution of eighth order (Trukenbrod & Engbert, 2014) with a mean value of $\mu = 275$ ms. Second, we assumed that the build-up of activation is considerably faster in the attention map than in the fixation map by choosing $R_0 = 0.01$ and $R_1 = 1$, i.e., $R_1/R_0 \sim 100$.

Model parameters were estimated by minimization of a loss function combining information on the densities of gaze positions and of saccade lengths,

$$A(\sigma_0, \sigma_1, \omega, \rho, \eta) = \sum_i (p_i^e - p_i^s)^2 + \sum_j (q_j^e - q_j^s)^2, \quad (\text{A11})$$

where p^e and p^s are the experimental and simulated distributions of pair distances between all data points

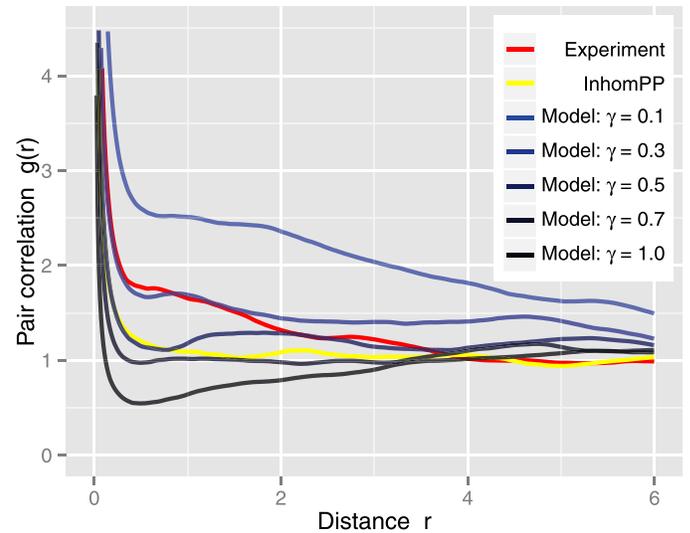


Figure A1. Pair correlation function obtained from simulations for different values of γ (blue lines) in comparison to the experimentally observed PCF (red line) and the result for the inhomogeneous point process (yellow line). All simulations were carried out for image #1 of our data set.

for a given image and q^e and q^s are the distributions of saccade lengths for experimental and simulated data, respectively. The minimum of the objective function A was determined by a genetic algorithm approach (Mitchell, 1998) within a predefined range (Table 1). Mean values and standard errors of the means were computed from five independent runs of the genetic algorithm.

Qualitative analysis of the model

The pair correlation function was the most important statistical concept in model evaluation. In our model, the strength of spatial correlation turned out to be related to the value of the exponent γ in the fixation map of the potential, Equation (9). We performed numerical simulations with the value of parameter γ fixed at different values between 0 and 1 to investigate the dependence of the spatial correlations on this parameter qualitatively (Figure A1). While $\gamma = 1$ produces negatively correlated scanpaths, $g(r) < 1$, at short pair distances r , it is possible to produce even stronger PCF value than in the experimental data for $\gamma < 0.3$. Thus, a single parameter in our model can generate a broad range of second-order statistics.

Some notes on the pair correlation function

The pair correlation function can be used to examine the second-order statistics of a point pattern. We first

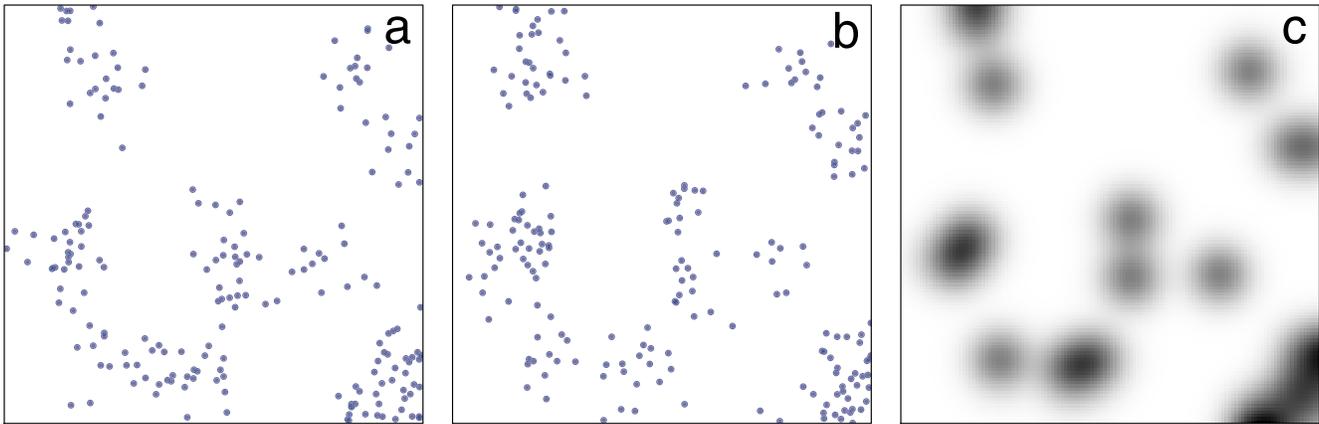


Figure A2. First-order properties of point processes. (a, b) Two samples from a point process (c). The intensity of the point process, $\lambda(x)$, which corresponds to the expected number of points to be found in a small circle around location x . Dark regions indicate high intensity (density).

need to define a few terms. A *point process* is a probability distribution that generates random point patterns: a sample from a point process is a set of observed locations (i.e., fixations, in our case). Therefore, taking two different samples from the same point process will result in two different sets of locations, although the locations may be similar (Figure A2).

First-order statistics: the intensity function. The first-order statistics of a point process are given by its intensity function $\lambda(x)$. The higher the value of $\lambda(x)$, the more likely we are to find points around location x . Figure A2c shows the theoretical intensity function for the point process generating the points in Figures A2a and A2b.

One way to look at the first-order statistics of a point process is via random variables that count how many points fall in a given region. For example, we could define a variable c_A that counts how many points fall within area A , for a given realization of the point process. The expectation of c_A (how many points fall in A on average) is given by the intensity function, i.e.,

$$E(c_A) = \int_A \lambda(x) dx, \quad (\text{A12})$$

where the integral is computed over area A . A slightly different viewpoint is given by the *density function*, which is a normalized version of the intensity function, defined as

$$\bar{\lambda}(x) = \frac{\lambda(x)}{\int_{\Omega} \lambda(x') dx'}, \quad (\text{A13})$$

where the integral in the denominator is over the *observation window* Ω , which in our case corresponds to the monitor (we cannot observe points outside of the observation window). The density function integrates

to 1 over the observation window and represents a probability density: If we now define a random variable z_A that is equal to one, when a (small) area A contains one point and 0 otherwise, we obtain

$$p(z_A = 1) = \int_A \bar{\lambda}(x) dx = \bar{\lambda}(x_A) dA, \quad (\text{A14})$$

where x_A is the center of area A and dA its area. If A is small, then $\bar{\lambda}(x)$ will be approximately constant over A , and the integral simplifies to $\bar{\lambda}(x_A)$ times the volume.

Second-order properties. The first-order properties inform us about how many points can be expected to find in an area, or, in the normalized version, whether we can expect to find a point at all. Second-order properties tell us about *interaction between areas*: whether for example we are more likely to find a point in area A if there is a point in area B .

In the case of the point process (Figure A2), the points are generated independently and do not interact in any way, so that knowing the location of one point tells us nothing about where the other ones will be. As shown in the manuscript, this is not so with fixation locations, which tend to cluster at certain distances.

The second-order statistics of a point process capture such trends, and one way to describe the second-order statistics is to use the pair correlation function. The pair correlation function is derived from the *pair density function* $\rho(x_A, x_B)$, which gives the probability of finding points at *both* location x_A and location x_B . Let us consider two random variables z_A and z_B , which are equal to 1, if there are points in their respective areas A and B , and 0 otherwise (again we assume that the areas are small). The probability that $z_A = 1$ and that $z_B = 1$ individually is given by the density function, Equation (A13). The probability that *both* are equal to one is given by the pair density function,

$$\begin{aligned} p(z_A = z_B = 1) &= \int_A \int_B \rho(x, x') dx dx' \\ &= \rho(x_A, x_B) dA dB. \end{aligned} \quad (\text{A15})$$

The pair density function already answers our question of whether observing a point in A makes it more likely to see one in B , and vice versa. If points are completely independent, then the resulting pair density is given by

$$p(z_A = z_B = 1) = \bar{\lambda}(x_A) \bar{\lambda}(x_B) dA dB. \quad (\text{A16})$$

If the pair density function gives us a different result then an interaction is occurring. Therefore, if we take the ratio of the pair density, Equation (A15) to the product of the densities, we obtain a measurement of deviation from statistical independence, i.e.,

$$c(x, x') = \frac{\rho(x, x')}{\bar{\lambda}(x) \bar{\lambda}(x')} \quad (\text{A17})$$

The resulting object is, however, a complicated, four-dimensional (i.e., two dimensions for x and two dimensions for x') function and in practice it is preferable to use a summary measure, which is the pair correlation function expressing how often pairs of points are found at a distance of ε from each other. The pair correlation function is explained informally in Figure A3.

More formally, the pair correlation function is just an average of $c(x, x')$ for all pairs x, x' that are separated by a distance r , i.e.,

$$\rho(r) = \int_x \int_{x' \in \Omega | d(x, x') = r} c(x, x') dx dx' \quad (\text{A18})$$

In the above equation, the notation $x' \in \Omega | d(x, x') = r$ indicates that we are integrating over the set of all points x' that are on a circle of radius r around x (but still in the observation window Ω).

If we are to estimate $\rho(r)$ from data, we need an estimate of the intensity function, as it appears as a correction in Equation (A17). In addition, since we

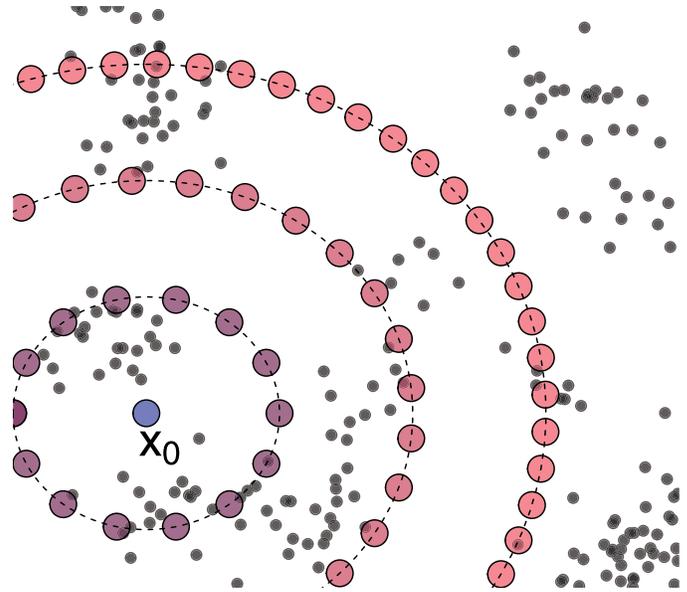


Figure A3. From the pair density function to the pair correlation function. The pair density function $\rho(x_A, x_B)$ describes the probability of finding points at both x_A and x_B in a sample from the point process. As such, it is a four-dimensional function, and hard to estimate and visualize. The pair correlation function (PCF) is a useful summary. To compute the raw pcf, we pick an initial location x_0 (circle) and look at the probability of finding a point both at x_0 and in locations at a distance ε from x_0 (first array of circles around x_0). We do this for various distances (other arrays of circles) to compute the probability of finding pairs as a function of ε . Finally we average over all possible locations x_0 , to obtain the pair correlation function. The pair correlation function therefore expresses how likely we are to find two points at a distance ε from each other.

have only observed a discrete number of points, the estimated pair density function can only be estimated by smoothing, which is why a kernel function needs to be used. We refer readers to (Illian et al., 2008) for details on pair correlation functions.