

**This is an electronic reprint of the original article.
This reprint *may differ* from the original in pagination and typographic detail.**

Author(s): Ylitalo, Anna-Kaisa; Särkkä, Aila; Guttorp, Peter

Title: What we look at in Paintings : A Comparison Between Experienced and Inexperienced Art Viewers

Year: 2016

Version:

Please cite the original version:

Ylitalo, A.-K., Särkkä, A., & Guttorp, P. (2016). What we look at in Paintings : A Comparison Between Experienced and Inexperienced Art Viewers. *Annals of Applied Statistics*, 10(2), 549-574. <https://doi.org/10.1214/16-AOAS921>

All material supplied via JYX is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

WHAT WE LOOK AT IN PAINTINGS: A COMPARISON BETWEEN EXPERIENCED AND INEXPERIENCED ART VIEWERS

BY ANNA-KAISA YLITALO^{*,1}, AILA SÄRKKÄ^{†,2} AND PETER GUTTORP^{‡,§}

*University of Jyväskylä**,
Chalmers University of Technology and University of Gothenburg[†],
University of Washington[‡] and Norwegian Computing Center[§]

How do people look at art? Are there any differences between how experienced and inexperienced art viewers look at a painting? We approach these questions by analyzing and modeling eye movement data from a cognitive art research experiment, where the eye movements of twenty test subjects, ten experienced and ten inexperienced art viewers, were recorded while they were looking at paintings.

Eye movements consist of stops of the gaze as well as jumps between the stops. Hence, the observed gaze stop locations can be thought of as a spatial point pattern, which can be modeled by a spatio-temporal point process. We introduce some statistical tools to analyze the spatio-temporal eye movement data, and compare the eye movements of experienced and inexperienced art viewers. In addition, we develop a stochastic model, which is rather simple but fits quite well to the eye movement data, to further investigate the differences between the two groups through functional summary statistics.

1. Introduction. Eye movements are outcomes of cognitive processes in the human brain, and can be recorded with high spatial and temporal resolution by computerized eye trackers. Eye movements provide valuable information about cognitive processes [Duchowski (2002), Rayner (1998, 2009)], and tracking of them is analyzed in a range of different areas, such as language and music reading [Kinsler and Carpenter (1995), Rayner (1998)], psychology [Findlay (2009)], and marketing research [Nagasawa, Yim and Hongo (2005)].

The first measurements of eye movements were made in 1879 independently by M. Lamare in France and Ewald Hering in Germany [Wade (2010)]. Both researchers used an acoustic approach, and noticed, much to everyone's surprise, that in reading text the gaze moves in jerks between points of rest. These extremely rapid movements are called *saccades*. They are essentially involuntary, so once a saccade starts it cannot be interrupted [Findlay (2009)]. Thus, researchers can predict saccade lengths from very few observations early in the saccade [Komogartsev, Ryu and Koh (2009)]. The points of rest, periods in which the gaze

Received October 2014; revised December 2015.

¹Supported in part by the Finnish Doctoral Programme in Stochastics and Statistics and by the Academy of Finland (Project number 275929).

²Supported by the Knut and Alice Wallenberg Foundation.

Key words and phrases. Coverage, intensity, point process, shift function, transition probability.

is staying relatively still around a location of the target space, are called *fixations* [Barlow (1952)]. Eye movements can thus be represented as an alternating sequence of fixations and saccades. In this paper we will mainly focus on sequences of fixation locations and times in the target space, and call such a sequence of observations a fixation process. Our targets are pictures of paintings and the data consist of recorded eye movements of subjects on the paintings. We will introduce some new statistical tools and a model to analyze the fixation process by means of group comparisons.

The fixation process is regarded as a spatio-temporal point process. We are aware of only few studies that use point process methodology to analyze fixation processes. Barthelmé et al. (2013) use inhomogeneous Poisson processes to model fixation locations, and Engbert et al. (2015) perform some preliminary analysis of clustering of the fixations by using inhomogeneous pair correlation function. Furthermore, they construct a dynamic model for saccade selection by discretizing the space.

We are interested in studying how people look at art. The first such study was carried out using two movie cameras by Buswell (1935). His findings indicated that much of the theory of how people look at art needed reformulation. Another interesting eye movement study related to arts is brought out by Locher (2006). He developed a two-stage model for art viewing, where the first stage of viewing of a painting was to obtain a general view of its structure and semantics. In Locher's study, subjects describing a painting orally while viewing it started to describe the painting only after a couple of seconds after they had begun to look at it. At the second stage of the model the viewer focuses on some interesting features, which are analyzed from an aesthetic point of view.

In this paper, we concentrate on comparing the fixation process of inexperienced and experienced art viewers, *novices* and *non-novices* for short. Already in his pioneering work Buswell (1935) compared the fixation durations of 61 non-novices and 117 novices (originally "art experts" and "lay persons") by comparing the group averages based on the 25 first fixations. He repeated the experiment for seven paintings, and in each case he concluded that the non-novices made shorter fixations than the novices. Nevertheless, when comparing the fixations in the areas of the subdivided picture, he did not report any major differences between the two groups. However, there are some studies where some differences have been found. In Kristjanson and Antes (1989) ten paintings were shown to a group of non-novices and a group of novices ("artists" and "non-artists" in the original paper). All the subjects were familiar with three of the paintings (same for all subjects), and the non-novices group was familiar with another three paintings, while the remaining paintings were unknown to all subjects. The only statistically significant difference between the non-novice and novice groups found was that non-novices made more fixations on the parts of the paintings that were not centers of interest when viewing unfamiliar paintings compared to novices. The authors concluded that the non-novices made a more thorough investigation of all areas

of the unfamiliar paintings than the novices. The final conclusion was that it is conceivable that the non-novices extract a different amount or different kind of information during each fixation, even though the pattern of fixations may vary little from the pattern of novices.

Vogt and Magnussen (2007) used eye movement patterns to investigate how non-novices and novices (“artists” and “artistically untrained people” in the original paper) view art. Participants were shown realistic and abstract works of art under two conditions: one asking them to free scan the paintings, and the other asking them to memorize them. Participants’ eye movements were tracked as they either looked at the images or tried to memorize them, and their recall for the memorized images was recorded. The researchers found no differences in the fixation frequency or duration between picture types for non-novices and novices. However, across the two conditions, the novices had more short fixations while free scanning the works, and fewer long fixations while trying to memorize. Non-novices followed the opposite pattern. In addition, non-novices spend more time than novices by scanning the areas which were not defined as regions of interest. There was no statistically significant difference in the recall of the images across groups, except that non-novices recalled abstract images better than novices and remembered more pictorial details.

Here, we investigate whether we can find any differences between how non-novices and novices view a particular painting by using a new set of tools to study eye movements. As mentioned earlier, the fixation process is described by a spatio-temporal point process, and intensities of the point processes and distributions of the fixation and saccade durations are used to compare the eye movements of novices and non-novices. We are mainly concentrating on the eye movements of 20 subjects, 10 novices and 10 non-novices, on the painting *Koli landscape* by Eero Järnefelt [Sundell (1986)] but give some results based on the other five paintings as well. For example, we investigate whether the fixation duration distributions of the novice and non-novice groups remain the same from painting to painting. Finally, based on the data analysis, we construct a rather simple stochastic model that further helps to describe the fixation process and to investigate differences between the non-novices and novices in terms of functional summary statistics.

The paper is organized as follows. The experimental setup and the data from the cognitive art research experiment as well as the results from the preliminary data analysis are described in Section 2. In Section 3, we show the results of a comparison of the two groups, novices and non-novices. A stochastic model is constructed in Section 4, and it is used to make further comparisons of the groups for the Järnefelt painting in Section 5. Finally, in Section 6 we discuss the results.

2. Data and data analysis.

2.1. *Art experiment.* Twenty test subjects participated in an experiment where the task was to observe six pictures of paintings on a computer screen, each during three minutes, and describe orally the mood in each painting while their eye

movements and voice were recorded. The paintings can be seen in Figure 1. The eye movements were recorded by the SMI iView X™ Hi-Speed eye tracker with a temporal resolution of 500 Hz. The resolution of the target screen was 1024×768 pixels and the distance between a participant's head and the screen was around 85 cm. A forehead rest was used in order to avoid redundant movement of the head. The data were collected by Pertti Saariluoma and Sari Kuuva (University of Jyväskylä) and María Álvarez Gil (University of Salamanca) with technical help from Jarkko Hautala and Tuomo Kujala (University of Jyväskylä). All subjects were students at the University of Jyväskylä at the time of the study. Ten of the subjects were either art students (8) or students who had studied art history and frequently visited art exhibitions (2). The remaining ten subjects were students who did not have art as their major or their hobby. We call the participants in the first group *non-novices* and participants in the latter group *novices*. Five of the participants were men (three novices and two non-novices) and 15 women (seven novices and eight non-novices).

In this article our main purpose is to introduce some new tools for eye movement analysis. Therefore, instead of analyzing all the six paintings in detail, we concentrate mainly on the eye movements on one of the paintings, namely, the Järnefelt painting Koli landscape, shown in Figure 1(a). The resolution of the image of the painting is 770×768 pixels, hence, it does not fill the whole computer screen and there are white areas on both sides of the painting. For some subjects, some of the fixations were located in these white areas outside the painting (Figure 2, right) and excluded from the analysis. We will treat saccades going outside the image as missing values.

2.2. Data description. We first look at the eye movement data as a whole and later (Section 3) compare the eye movements of novices and non-novices. An important way of describing the viewing of the painting is to look at a smoothed plot of all fixation locations of all subjects, which represents the viewing foci of the painting. Technically, when describing the locations of fixations by a spatial point process this is an estimate of the intensity of the fixation process, marginalized over time. We use the R package `spatstat` for the intensity estimation [Baddeley, Rubak and Turner (2015)]. A chi-square test for quadrat counts clearly rejects the hypothesis of constant intensity ($p < 0.001$). The areas of the painting that have been visited most by the subjects are shown in Figure 3 (left). We can see that the trunk of the main tree going down almost the entire length of the painting and the top of the small tree next to it seem to be the areas of particular interest in the painting. In Figure 3 (right) one can see the least visited areas, the edges of the painting in this case.

The fixation process varies a lot between subjects, and hence it is interesting to look at some summary statistics for the data. The supplemental article by Ylitalo, Särkkä and Guttorp (2016a) includes a table which shows the total number of fixations and the number of short fixations (less than 40 ms) inside the picture with

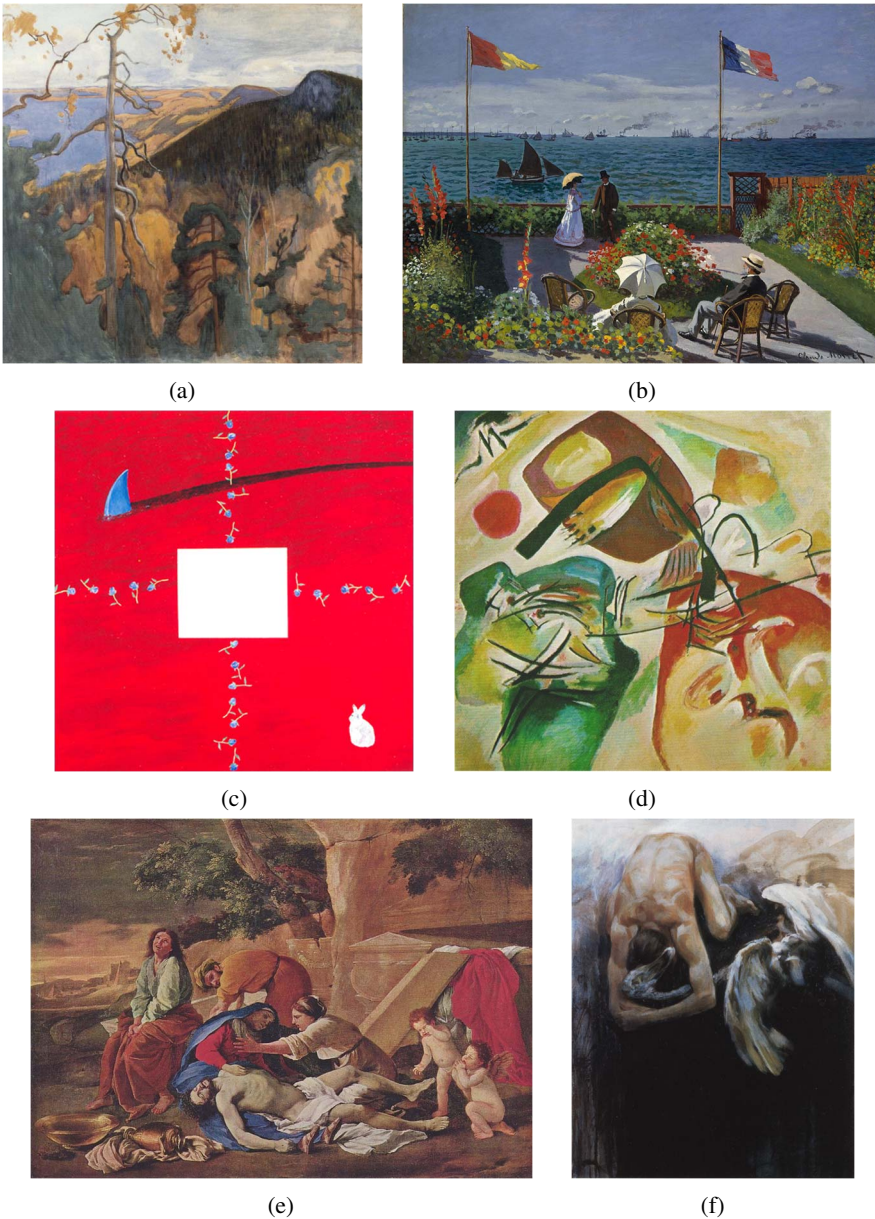


FIG. 1. The stimulus paintings used in the experiment. (a) Eero Järnefelt—Koli-landscape (the turn of 19th and 20th century), source: Sundell (1986)/public domain. (b) Claude Monet—Terrasse à Sainte-Adresse (1867), source: The Metropolitan Museum of Art/public domain. (c) Risto Suomi—La croisée des destins (1988), source: Mikkola (1997). Reprinted with kind permission of Risto Suomi. (d) Wassily Kandinsky—Picture with a Black Arch (1912), source: WikiArt.org/public domain. (e) Nicolas Poussin—La Lamentation sur le Christ (17th century), source: The Yorck Project (2002)/public domain. (f) Pasi Tammi—Poem Forces to Kneel Down (1999), source: Anonymous (2000). Reprinted with kind permission of Pasi Tammi.

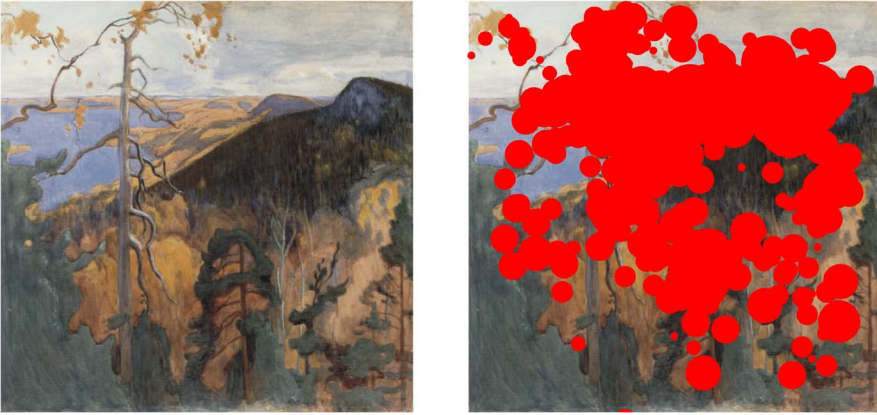


FIG. 2. (Left image) Koli landscape by Eero Järnefelt from the turn of 19th and 20th century. (Right image) Fixations (red spots) for Subject 15, a male non-novice. The duration of each fixation is proportional to the size of the spot.

means, medians, and standard deviations of fixation durations for each subject. The number of fixations during the three minutes time period varies between 326 and 770, and the median fixation duration varies between 150 ms and 438 ms. The mean fixation duration is always greater than the median, indicating a right-skewed distribution of the fixation duration.

Short fixations are not included in our analysis because they are believed to be microsaccades, or eye movements within a fixation [Manor and Gordon (2003)]. The histogram on the left side of Figure 4 indicates that the duration distribution may be a mixture of very short durations and regular ones. When excluding the short fixations, the remaining durations seem to follow a gamma distribution. The

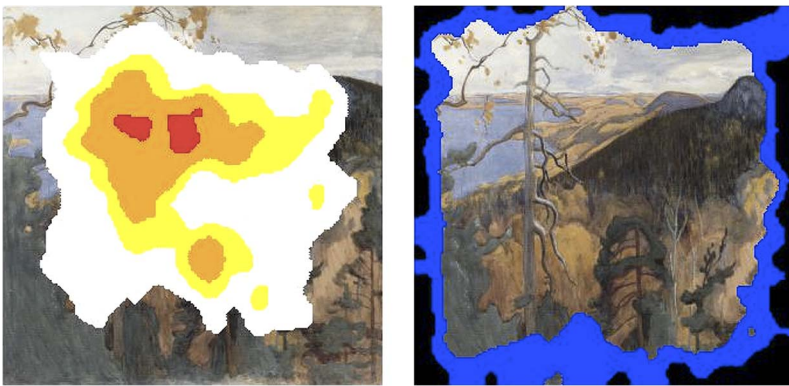


FIG. 3. Areas of the Järnefelt painting that were visited most (left) and least (right) by the gaze of the subjects. (Left image) Top 50% white, top 20% yellow, top 10% orange, and top 1% red. (Right image) Bottom 30% blue and bottom 10% black.

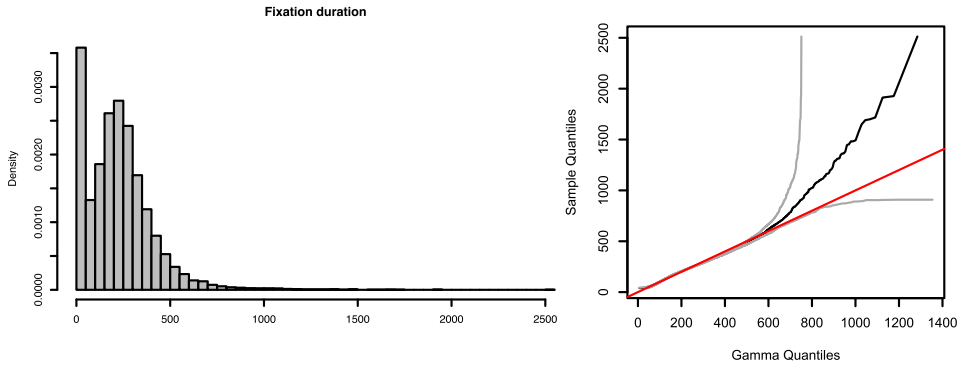


FIG. 4. (Left panel) Distribution of fixation durations for all subjects when short fixations are included. (Right panel) Gamma quantile–quantile plot for fixation durations exceeding 40 ms with asymptotic 95% confidence region (grey) and the line (red) of unit slope through the origin corresponding to perfect agreement of empirical and theoretical quantiles.

right panel shows a gamma quantile–quantile plot [Wilk and Gnanadesikan (1968)] of the regular (longer than 40 ms) fixation durations. The red line of slope 1 falls inside the simultaneous confidence band [Doksum and Sievers (1976)], indicating that the durations can be described well by a gamma distribution.

As mentioned earlier, the complete eye movement process consists of fixations and saccades, the latter being the rapid movements between the fixations. The supplemental article by Ylitalo, Särkkä and Guttorp (2016b) shows summary statistics for saccade durations and saccade lengths, that is, the distances between consecutive fixation locations. The distribution of saccade lengths as well as the distribution of saccade durations are skewed to the right, and are also well described by gamma distributions. The temporal dependence between fixation durations is very weak, as judged by the autocorrelation functions and ARIMA time series fitting.

2.3. Temporal analysis. Art is considered to be a subjective field. How one views artwork is individually unique, and reflects one’s experience, knowledge, preference, and emotions. According to Locher’s two-stage model [Locher (2006)], exploration of a picture starts with a global survey of the pictorial field in order to get an initial overall impression of the structural arrangement and semantic meaning of the composition. The second phase of an aesthetic episode consists of visual scrutiny or focal analysis of interesting pictorial features detected initially to satisfy cognitive curiosity and to develop aesthetic appreciation of the display. Some pioneering investigations into visual exploratory behavior of paintings by Buswell (1935) and Yarbus (1967), and subsequent studies on the informative details of an image by Antes (1974) and Mackworth and Morandi (1967), revealed that observers focus their gaze on specific areas of the image, rather than in a random fashion. The areas receiving high intensities of fixations were interpreted as

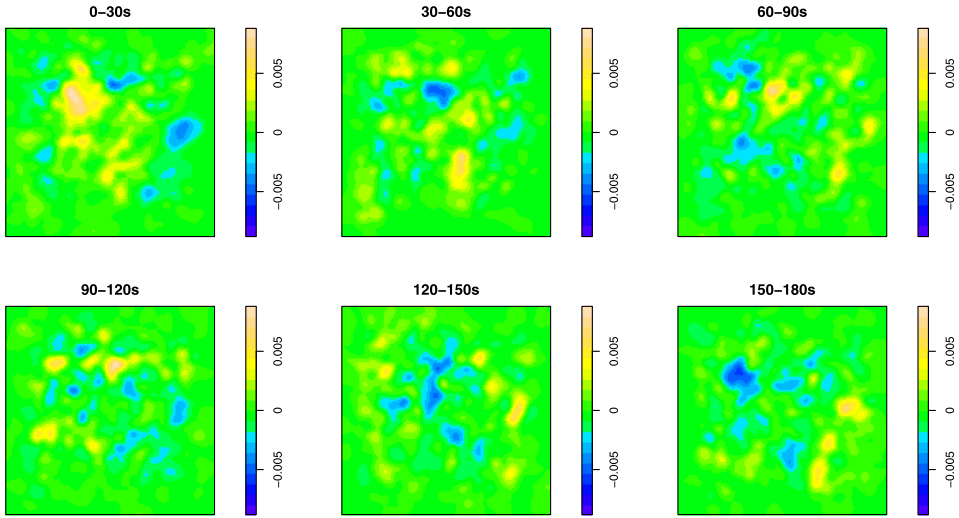


FIG. 5. Residual intensity surfaces in 30 s intervals. The more yellow the color, the higher the intensity compared to the mean intensity surface; the more blue the color, the lower the intensity.

guiding the observer's interest in informative elements of the image [Henderson and Hollingworth (1999)].

Since people do not look at all parts of the painting, but rather focus on some features of interest, the fixation process cannot be assumed to be spatially stationary. Furthermore, the features people look at may vary in time. To investigate the latter, we estimated the intensity surface of fixation locations based on the fixations of all subjects, and plotted residual intensities in 30 s intervals (Figure 5). The residual intensity is the difference between the estimated intensity surface and the mean intensity surface of the six intervals. The intensity in location $x \in I$, where I is the area of the painting, is estimated using the edge-corrected kernel estimator [see, e.g., Gatrell et al. (1996)]

$$(1) \quad \hat{\lambda}_h(x) = \frac{\sum_{i=1}^n h^{-2} K(h^{-1}(x - x_i))}{\int_I h^{-2} K(h^{-1}(x - u)) du},$$

where x_i , $i = 1, \dots, n$, are the fixation locations, h a bandwidth (a smoothing parameter), and K a kernel function, here the standard two-dimensional Gaussian density. The bandwidth for the mean intensity estimation is selected by applying the cross-validation approach described in Diggle (1985) and Berman and Diggle (1989), and that bandwidth (here 17) is the same for all intervals.

One can notice in Figure 5 that during the first 30 s, the gaze is more concentrated on the tall pine tree on the left-hand side of the painting and after that more on the small tree in the middle than on average during the whole three minutes time period. During the second and the third minutes fixations are spread out a

little more than during the first minute. In addition, during the last minute the areas which were of interest at the beginning (the two trees) are now being avoided. Note that here we have only made some visual observations. For a more formal comparison of intensity surfaces we could use the test we introduce in Section 3 below.

We also investigate the difference in fixation duration distributions between the different 30 s time intervals. Figure 6 shows a graphical way to compare distributions, a shift function [Doksum and Sievers (1976)] defined as follows. If X has the cumulative distribution function F , written $X \sim F$, and Y the cumulative distribution function G , $Y \sim G$, then the shift function is defined as

$$\Delta(x) = G^{-1}(F(x)) - x,$$

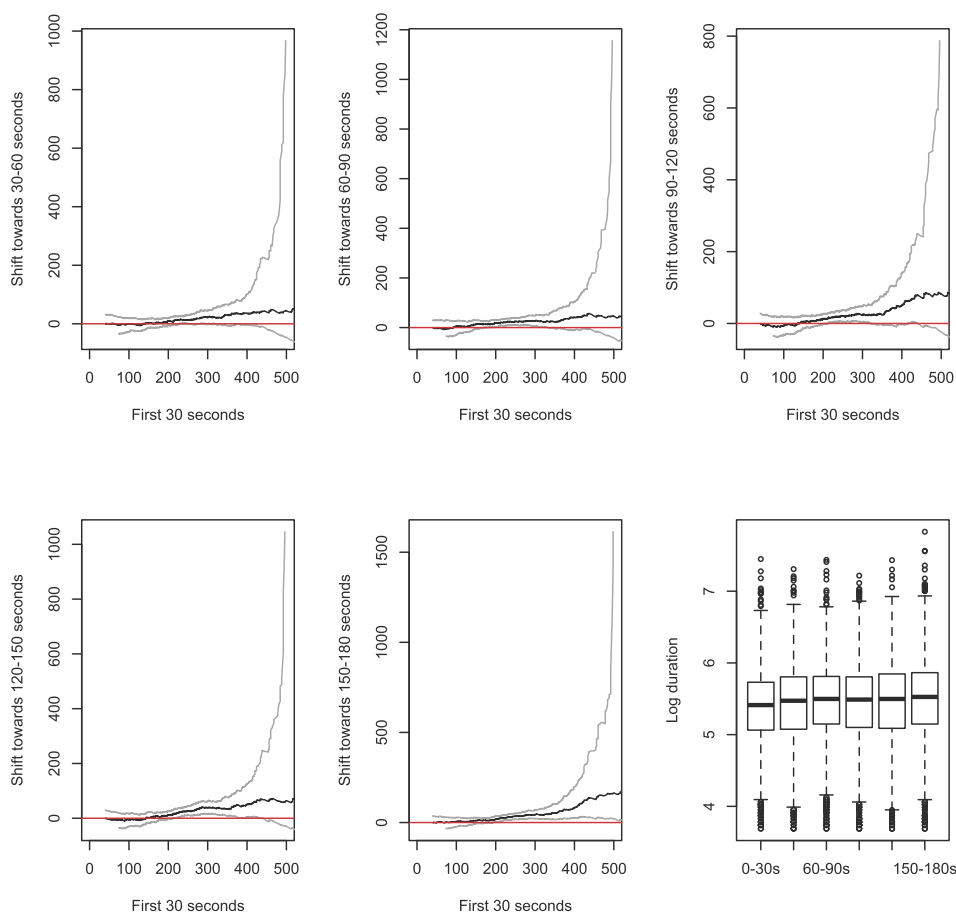


FIG. 6. Fixation duration shift function estimates in 30 s intervals, relative to the first 30 s. The red line corresponds to equal distributions, and the grey lines form 95% asymptotic simultaneous confidence bands. The lower right panel is a box plot comparison of the distributions on a log scale.

and has the property that $X + \Delta(X) \sim G$. The idea of the shift function is to compare two distributions based on how much one distribution needs to be shifted in order to coincide with the other. In general, this may depend on the abscissa at which the comparison is made. We have the simplest case when the second distribution is simply the first shifted by a constant amount. The shift function is then just a horizontal line at the level of the shift (in particular, if the distributions are the same, then the shift function is a horizontal line at the level zero). If a location-scale model is appropriate, then the shift function is a straight line with slope related to the scale change. The shift function is easily estimated using empirical distribution functions, and simultaneous confidence intervals based on the distribution of the Kolmogorov–Smirnov statistic are given in [Doksum and Sievers \(1976\)](#). If a horizontal line at level 0 falls inside the band, then the two distributions are statistically indistinguishable. The shift function estimates in [Figure 6](#) represent the comparison of the second through the sixth 30 s fixation distributions to the first 30 s distribution. It seems that the later time intervals tend to have slightly longer fixations between about 200 ms and 450 ms than the earlier intervals. This is also illustrated by the side-by-side box plots in the lower right-hand panel of the figure.

3. Comparison of novices and non-novices. In the previous section we introduced some tools to describe the fixation process. Now these tools are applied to compare the fixation processes of novices and non-novices. First, we compare the fixation intensity surfaces of the two groups visually, and then construct a test for more formal comparison. Second, we compare the fixation duration distributions in the two groups using shift plots. Finally, we investigate whether the fixation duration distributions change within the two groups if the painting is replaced by another.

3.1. Visual comparison of intensity surfaces. To compare the spatial patterns of fixations, we first estimate the overall intensity for each group by using equation (1); see the top row in [Figure 7](#). We see that both groups concentrate on the middle part of the painting and do not look closely at the edges, and that the fixations of non-novices seem to be slightly more concentrated than the fixations of novices. Non-novices look a lot at the large pine and one of its branches (vertical alignment), while novices look at the same branch of the tree but continue to the right of the branch (horizontal alignment). We also plot the areas of the painting that are visited most (top 5% in yellow and top 1% in red) for each group ([Figure 7](#), bottom row). This verifies the finding of non-novices' interest in the vertical trunk while novices have been concentrating more on the horizontal branch.

A more detailed investigation, where the time interval is divided into 30 s intervals [see supplemental article [Ylitalo, Särkkä and Guttorp \(2016c\)](#)], reveals that especially early on the non-novices are more concentrated on the large tree and its branch than the novices. The gaze of novices starts to spread out much earlier than

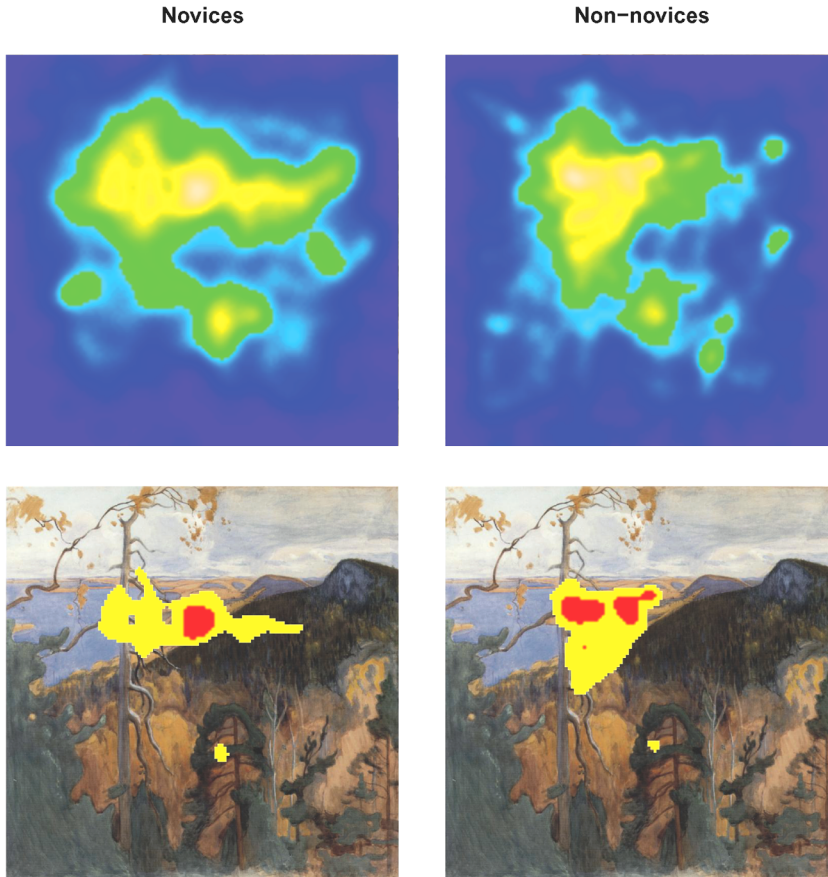


FIG. 7. Top row: The estimated overall intensity surface with bandwidth 20 for novices (on the left) and with bandwidth 16 for non-novices (on the right). Bottom row: Top 5% (yellow) and top 1% (red) intensities, novices vs. non-novices.

the gaze of non-novices and novices have more hot spots than non-novices. The gaze of non-novices mainly stays on the large pine tree and in the end also on the small tree on the right side of the large one, while novices investigate other areas of the painting as well.

3.2. *Test for comparing two intensity surfaces.* In order to test whether two intensity surfaces differ significantly from each other, that is, whether the subjects in the two groups have been looking at different areas of the painting, we apply the method developed by Kelsall and Diggle (1995a, 1995b). They considered a ratio of two kernel density estimates and used a Monte Carlo test for testing if the ratio is constant, that is, whether the kernel estimated intensity surfaces can be considered similar. This ratio is used in epidemiology to explore whether some diseased cases are randomly located among healthy controls and is therefore often

called a relative risk function. The logarithmic relative risk function is defined as $\rho(x) = \log(\lambda_1(x)/\lambda_2(x))$, where λ_i 's are computed using equation (1), and can reveal areas where the intensity of fixations differs between novices and non-novices.

By conditioning on the observed number of fixations, we may view the fixation locations of novices and non-novices as independent samples from the densities $f_1(x)$ and $f_2(x)$, where

$$\hat{f}_i(x) = \frac{\hat{\lambda}_i(x)}{\int_I \hat{\lambda}_i(x) dx},$$

for $i = 1, 2$. We can now define a logarithmic density ratio

$$r(x) = \log \frac{f_1(x)}{f_2(x)} = \rho(x) - c,$$

where $c = \log\left(\frac{\int_I \lambda_1(x) dx}{\int_I \lambda_2(x) dx}\right)$ [see, e.g., Wakefield, Kelsall and Morris (2000), Kelsall and Diggle (1995b)]. We have

$$\hat{r}(x) = \log \frac{\hat{f}_1(x)}{\hat{f}_2(x)},$$

which contains all the information about the spatial variation in $\rho(x)$.

According to Kelsall and Diggle (1995b), the choice of the kernel function is not critical, but choosing the smoothing parameter h is. Furthermore, it is not obvious whether one should use the same bandwidth for both densities when computing the ratio. Kelsall and Diggle (1995b) suggest that if the densities are expected to be nearly equal, then the bandwidths should be equal as well to reduce bias. They also warn that if the sample sizes are very different, then one can get poor results when using the same bandwidth for both samples. Furthermore, Bailey and Gatrell (1995) suggest that the kernel estimate in the denominator should be deliberately oversmoothed (a larger bandwidth should be used). In addition, Hazelton (2007) shows that edge correction terms are not needed if a common bandwidth is used.

We test the null hypothesis that the intensity surfaces are equal, that is, that the logarithmic density ratio $r(x) = 0$, by the following Monte Carlo test [Barnard (1963)]. Ten subjects were randomly selected from the twenty subjects to form the novice group and the remaining ten subjects formed the non-novice group. (Hence, the possible dependence between the fixation locations within a subject was taken into account.) By doing this random grouping m times, we obtain m log-density surfaces $\hat{r}_1, \dots, \hat{r}_m$ under the null hypothesis. As introduced by Kelsall and Diggle (1995b), we use the test statistic

$$T_j = \int_I \hat{r}_j(x)^2 dx, \quad j = 0, \dots, m.$$

The p -value is then $p = \frac{k+1}{m+1}$, where k is the number of test statistics T_j which are larger than or equal to the observed value of the test statistic, T_0 .

Since the number of fixations vary quite a lot between the subjects, we estimate the bandwidths separately for the two groups by using the cross-validation method by Diggle (1985) and Berman and Diggle (1989) with Ripley's isotropic edge correction [Ripley (1988)]. The estimated bandwidth for the novices was approximately 20 and for the non-novices 16. We obtained $T_0 = 59,604$ and $p = 0.108$ after random grouping of subjects 10,000 times. This means that the difference between the intensity surfaces of novices and non-novices is not statistically significant even though there are some dissimilarities in the intensities; see Figure 7. The largest differences between the intensities of the two groups are near the edges of the painting, where there are very few fixations of either group.

We made the same comparisons for the other five paintings and obtained similar results; see supplemental article Ylitalo, Särkkä and Guttorp (2016d). None of the tests gave any significant result, and Fisher's combined probability test resulted in $p = 0.392$ ($\chi^2 = 12.685$, $df = 12$). Therefore, we did not find any overall differences when comparing the intensities of the two groups. However, it seems that the more abstract paintings could reveal some differences between novices and non-novices; see results related to Kandinsky's painting in Ylitalo, Särkkä and Guttorp (2016d).

3.3. Comparing duration distributions. To compare the fixation duration distributions of the novice and non-novice groups, we again look at the shift plot; for the Järnefelt painting see Figure 8. Here, the novices have fewer short fixations (duration less than 125 ms), and more fixations between about 175 and 600 ms than non-novices. Similar results were found for the remaining five paintings; see Ylitalo, Särkkä and Guttorp (2016d). This is in line with the early findings of Buswell (1935), and may mean that novices more rarely only glance at a point on the painting but rather spend more time on each point compared to non-novices.

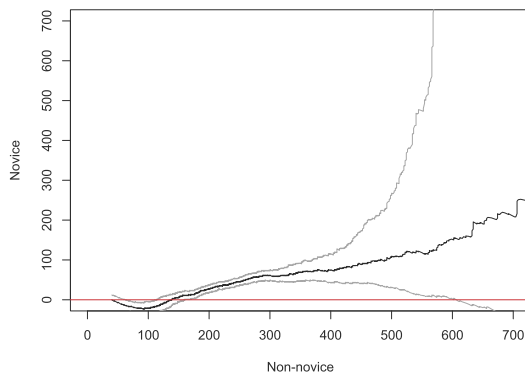


FIG. 8. Shift plot of non-novice distribution vs. novice distribution for fixation duration. The red line indicates equal distributions, and the grey lines form simultaneous asymptotic 95% confidence bands.

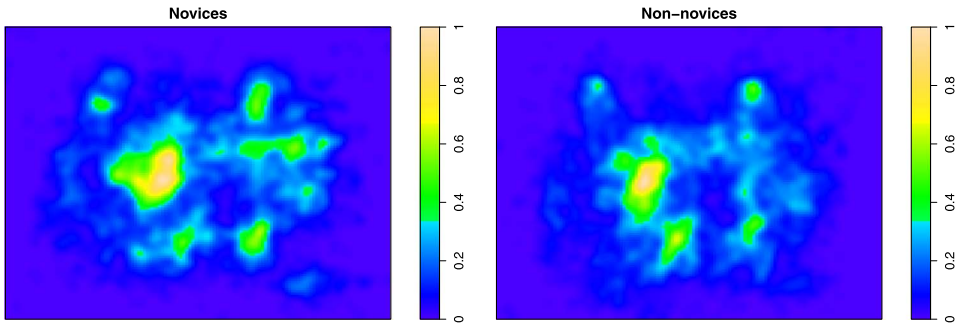


FIG. 9. The estimated overall intensity surface with fixed bandwidth 13 for novices (on the left) and for non-novices (on the right) for the Monet painting.

3.4. *Pairwise comparison of paintings.* Since several paintings are available, we are also able to investigate whether the fixation process is similar for different paintings, that is, whether the fixation process stays stable when looking at paintings. We are particularly interested in whether we detect any differences between the behavior in the novice and non-novice groups. We started by comparing the Järnefelt painting and the painting Terrace at Sainte-Adresse by Claude Monet (1867) shown in Figure 1(b) since both are landscapes, and hence similar to each other. The estimated intensity surface of the fixation locations for Monet's painting is also shown separately for novices and non-novices in Figure 9. The intensity surfaces of the two groups look fairly similar; see also the formal test in Ylitalo, Särkkä and Guttorp (2016d).

Since the pictures are completely different, the intensity surfaces of fixations for the paintings are also different and comparing them would be meaningless. However, it is of interest to compare the fixation duration distributions of novices and non-novices between two paintings in order to see whether the way of looking at paintings remains stable between different paintings. The shift function comparisons of the duration distributions for novices and non-novices are plotted in Figure 10 for the Järnefelt–Monet pair. We can see that for non-novices there is no significant difference between the paintings, while for novices there are more fixations that last 200–350 ms on the Järnefelt painting than on the Monet painting. We made similar comparisons for the other pairs of paintings as well; see supplemental article Ylitalo, Särkkä and Guttorp (2016e). Here, only the results based on the comparisons of the fixation duration distributions for the most similar pairs of paintings, namely, Suomi and Kandinsky which are abstract paintings [see Figures 1(c) and (d)], and Poussin and Tammi which have people in focus [see Figures 1(e) and (f)], are included. For novices, we see a difference for each three pairs of paintings. For non-novices, however, we can only see a clear difference for the Suomi–Kandinsky pair, but there is no difference for the other two pairs of paintings. This indicates that the non-novices, being used to looking at paintings,

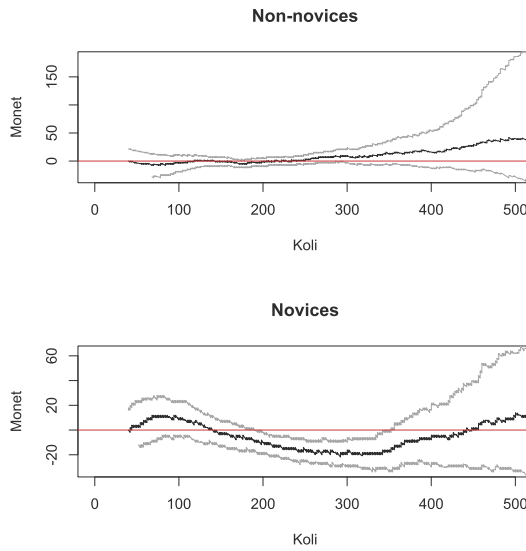


FIG. 10. *Shift function comparison of fixation duration distributions between the Järnefelt and Monet paintings for non-novices (upper panel) and novices (lower panel). The red lines correspond to equal distributions and the grey lines form 95% asymptotic simultaneous confidence bands.*

are more consistent in their fixation durations between paintings, but for novices, the duration of fixations depends on the painting.

4. Stochastic model. Above, we used intensity surfaces and shift plots to compare the eye movements of novices and non-novices. Next, we will construct a simple spatio-temporal model that describes the dynamics of the fixation process. We will fit the model separately to the novice and non-novice groups for the Järnefelt painting and see if we can find some further differences between the groups by comparing the fitted models in terms of functional summary statistics and their model-based simultaneous envelopes.

To build a spatio-temporal model for the fixation process, we need to model fixation locations, fixation durations, saccade durations, and saccade lengths. (Saccade durations are needed to fill the whole three minutes time period.) As building blocks we use the tools discussed in Section 2. Fixation locations are modeled as a realization of a spatial point process, which is characterized by its intensity function. Furthermore, appropriate distributions can be fitted to the fixation and saccade durations and for saccade lengths. After having estimated the intensity surface and the distributions mentioned above, we can simulate from the resulting model and compare the behavior of the model to the behavior in the data, and finally compare the models fitted to the two groups.

A natural reference or null model for spatial point pattern data, locations of fixations in our case, is the homogeneous Poisson process, where the points are

Novices

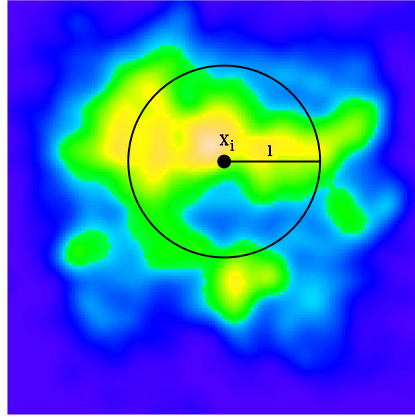


FIG. 11. The next fixation location is chosen from the black circle at distance l from the current fixation x_i according to the estimated intensity surface.

located uniformly and independently of each other on the study area. In our situation, however, only the more interesting parts of the painting are visited, and we cannot assume that the final set of locations of fixations is a realization of a homogeneous process. A homogeneous Poisson process is therefore not a realistic reference process for the fixation locations. Next, we define our own reference model for the fixation process.

Given an intensity surface, the first fixation location is drawn from the probability distribution on the observation window by normalizing the intensity to a bivariate distribution. The length l of the next saccade is drawn from a saccade length distribution for the group in question. Then, the exact location of the new fixation is decided by taking a point at distance l away from the current fixation location according to the intensity surface. In other words, the new location is chosen from the conditional intensity of points at the given distance from the current location; see Figure 11.

Concerning the temporal evolution, it seems natural to draw each fixation duration independently from the distribution of durations.

5. Model fit and comparison between novices and non-novices for the Järnefelt painting. In what follows, we will first give the particular choices of the distributions mentioned in the previous section, and introduce some summary statistics that could be used to assess the goodness of fit of the fitted models. Finally, we fit the model to the data for the Järnefelt painting, separately for novices and non-novices, simulate 200 realizations of each process, and check the goodness of fit of the fitted models by comparing the summary statistics estimated from the data to those estimated from the simulations.

5.1. *Choice of distributions.* The intensity surfaces of the fixation locations are estimated separately for each group and the bandwidth is selected by applying the cross-validation approach implemented in the R package `spatstat` [Baddeley, Rubak and Turner (2015)]. The same bandwidth is used for the data and for the simulations. The first location is drawn from the intensity surface estimated from the locations of the first fixations of the subjects in the group. Alternatively, the intensity surface estimated based on all the fixation locations in the group could be used. We believe that the choice of the first fixation location is not crucial.

Given the summary statistics computed from the data (Section 2.2), an appropriate distribution for the duration of fixations and the duration of saccades is the gamma distribution. Therefore, we model the fixation durations by sampling from the estimated gamma distribution of fixation durations in the group. The duration of each saccade is modeled by drawing from the common distribution of saccade durations estimated from all the subjects, as saccades are involuntary, once started, and therefore not expected to vary depending on viewing experience. The intensity surface, as well as the fixation duration and saccade duration distributions, stay the same throughout the observation period.

For saccade lengths we first fit a Gamma distribution separately for novices and non-novices and then use truncated versions of these distributions in order to avoid jumps outside the painting. The truncation point depends on the current location: if the process is at location x_i , then the longest jump it can take is to the furthest corner of the painting. This distance is called $l_{i,\max}$ and serves as the truncation point for the truncated gamma distribution from which the next jump is sampled. We noticed, however, that the truncated gamma distribution alone does not catch the long jumps that the gaze makes when the person moves her/his attention to another location. Therefore, we include such long jumps into the model and define the distribution for saccade lengths to be a mixture of two distributions: When the process is at location x_i , the length of the next jump is sampled from the uniform distribution $U(l_{i,\max}/2, l_{i,\max})$ with probability p and from the truncated Gamma distribution with truncation point $l_{i,\max}$ with probability $1 - p$. After some experimentation, we fixed the probability p for sampling the saccade length from the uniform distribution above to 0.2 for both groups.

5.2. *Summary statistics.* To get some idea of how the gaze jumps between different areas of the painting in the data and in the simulations, we divide the painting into quarters and compute the transition frequencies between them. Each quarter is called a state and identified with numbers 1–4 (from upper left to lower right quarter).

In order to see how much of the painting is viewed, we use two functional summary statistics, namely, convex hull coverage and ball union coverage defined as follows. Let, as before, I represent the observation window and $x_i \in I$ be the

i th fixation of the fixation process $\{X(t)\}$. An ordered set of n_t fixations up to a fixed time t is denoted by $x = \{x_1, \dots, x_{n_t}\}$. The convex hull of this point set is

$$C_x(t) = \left\{ \sum_{i=1}^{n_t} \alpha_i x_i : \alpha_i \geq 0 \text{ for all } i \text{ and } \sum_{i=1}^{n_t} \alpha_i = 1 \right\}$$

and the relative area of the convex hull is denoted by

$$AC_x(t) = \begin{cases} 0, & \text{if } n_t < 3, \\ \frac{\text{area}(C_x(t))}{\text{area}(I)}, & \text{if } n_t \geq 3. \end{cases}$$

This is called the convex hull coverage of the fixation process. Hence, the convex hull coverage is computed every time a new fixation appears.

Let $b(x_i, R)$ be a ball of fixed radius R centered in a fixation $x_i \in I$. Define

$$U_x(t) = \left\{ \bigcup_{i=1}^{n_t} b(x_i, R) \cap W \right\},$$

and call

$$AU_x(t) = \frac{\text{area}(U_x(t))}{\text{area}(I)}$$

the ball union coverage of the fixation process.

In addition, to measure the total length of the distance the gaze has moved, we use the scanpath length function [Noton and Stark (1971)], which is the sum of the saccade lengths up to time t . Hence, the scanpath length can be defined as

$$L(t) = \sum_{i=1}^{n_t-1} l_i \mathbf{1}(t_{i+1} \leq t),$$

where t_{i+1} is the time when the fixation $i + 1$ at location x_{i+1} takes place and l_i is the length of the jump just before that fixation. That length is the Euclidean distance between two successive fixations, that is, the length of the saccade.

5.3. Model fit. The estimated model-based intensity surfaces for the groups are different; see Figure 12. As mentioned earlier, both groups look a lot at the large pine tree and the top of the small tree next to it. The novice group tends to follow the branch of the large tree, while the non-novice group follows the trunk.

The convex hull coverage, ball union coverage with disc radius 35 pixels, and scanpath length for non-novices and novices are shown and compared to the models in Figure 13. The 95% simultaneous (rank) envelopes were created by using the R package `sptest` [Myllymäki et al. (2015, 2016)]. Figure 13 shows that the coverage statistics based on the model describe the data quite well. The model-based scanpath length is, however, less variable and somewhat shorter in the simulations than in the data. It seems that the model is not able to catch all variation

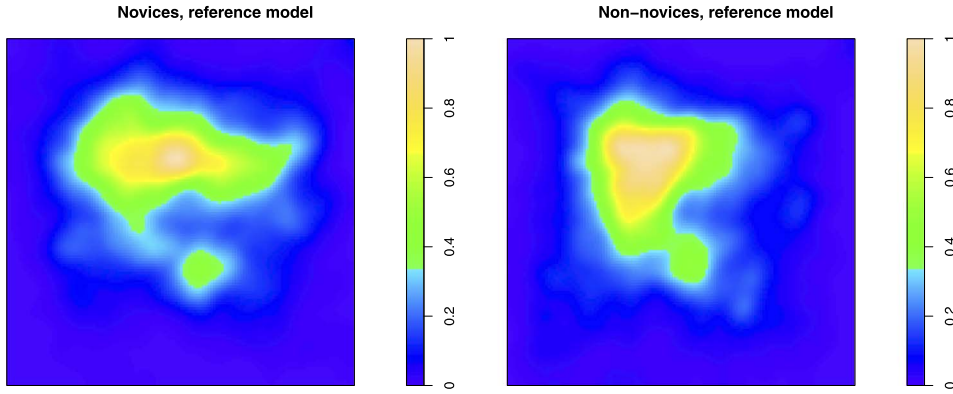


FIG. 12. Scaled overall intensity surfaces of the reference model for novices (left) and for non-novices (right).

related to the scanpath length. In the beginning (during the first 30 seconds for novices and during the first minute for non-novices), the model describes the data well, but later on variation in the data becomes too large to be caught by the model.

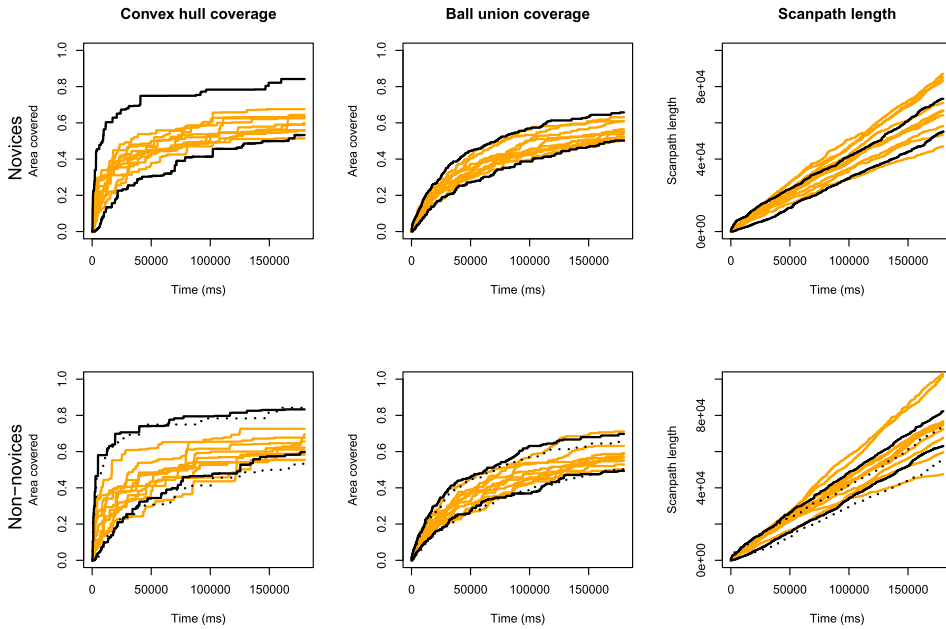


FIG. 13. Convex hull coverage (left), ball union coverage (middle), and scanpath length (right) curves for novices (top) and non-novices (bottom). Orange lines represent the subjects and black solid lines represent 95% simultaneous envelopes estimated from 200 simulated realizations of the reference model of the group in question. Dotted lines (bottom) represent 95% simultaneous envelopes for novices.

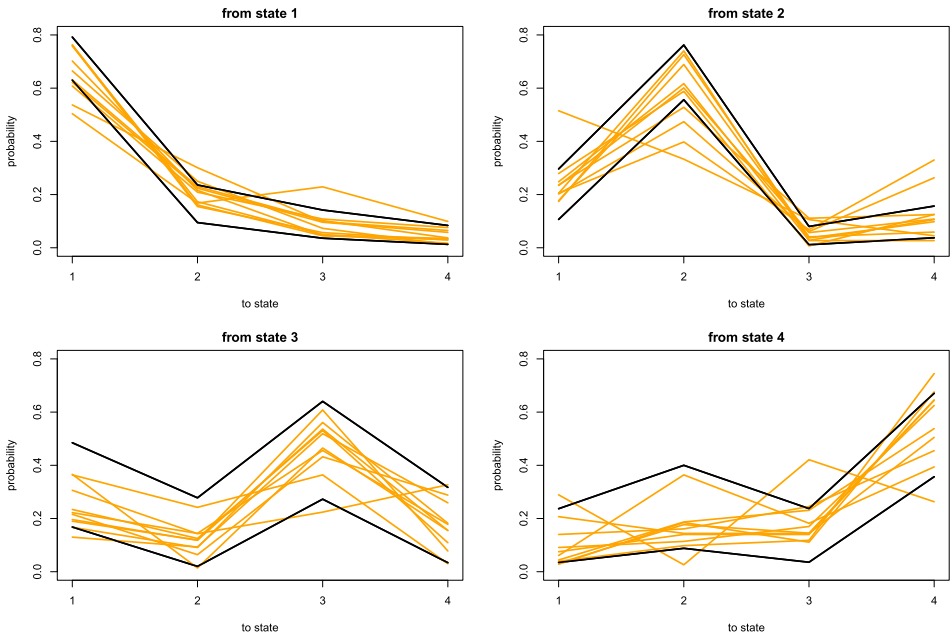


FIG. 14. Transition probabilities for novices with 95% simultaneous envelopes estimated from 200 simulated realizations of the reference model.

One should also note that the scanpath length is strongly affected by the number of fixations and saccade lengths, which both vary quite a lot in our data; see supplemental articles [Ylitalo, Särkkä and Guttorp \(2016a\)](#) and [Ylitalo, Särkkä and Guttorp \(2016b\)](#).

For both the novice data and the non-novice data, the estimated transition probabilities between different quarters of the painting indicate that the model does an adequate job (Figures 14 and 15). Recall that the quarters of the painting are called states and numbered from 1 to 4 (from upper left to lower right quarter).

The model seems to fit quite well to the fixation process for both groups. The fitted models are quite similar, but also some differences can be found. As can be seen in Figure 13, the model-based coverage summaries are quite similar for the two groups. However, the modeled gaze of non-novices seems to move slightly more than the gaze of novices according to the scanpath length (see Figure 13, bottom right plot). This is somewhat surprising, since the intensity surface of non-novices seems to be more compact than the one of novices, which would indicate that the gaze of non-novices takes shorter jumps in smaller areas than the gaze of novices. However, fixation durations of novices are longer than the ones of non-novices in general, which means that non-novices tend to make more jumps during the three minute period compared to novices. Thus, it is possible that the gaze of non-novices makes a larger number of short jumps than the gaze of novices, which results in the gaze traveling a longer path than the gaze of novices.

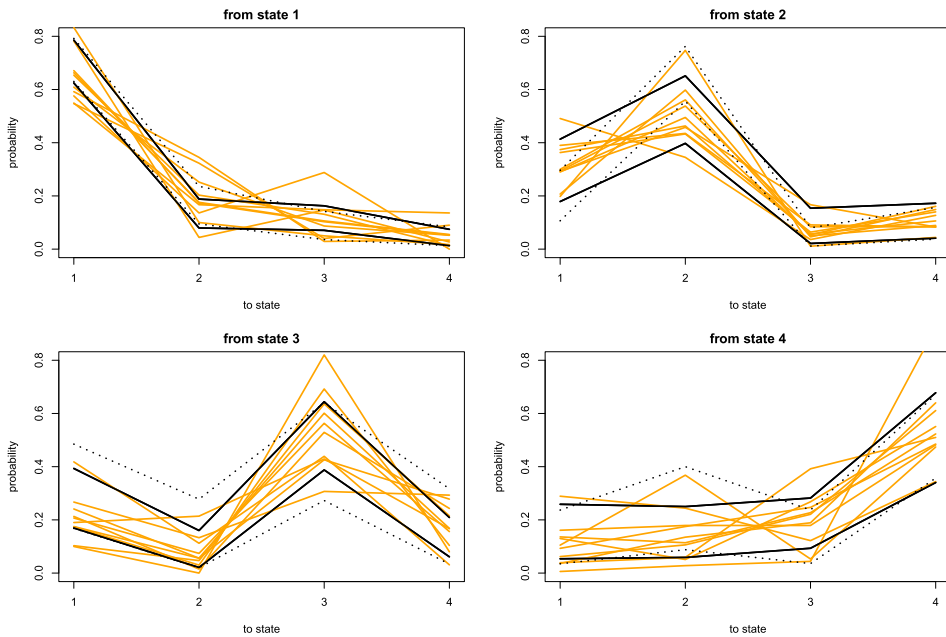


FIG. 15. Transition probabilities for non-novices with 95% simultaneous envelopes estimated from 200 simulated realizations of the reference model. Dotted lines represent 95% simultaneous envelopes for novices.

6. Conclusions and discussion. In this study we have analyzed eye movement data from twenty subjects who had looked at six different paintings. We have been particularly interested in whether non-novices and novices look at paintings in the same way, both when it comes to which parts of the painting they look at and how long the gaze typically stays in one spot. By regarding the fixation process as a spatio-temporal point process and by using a new set of tools we can study the differences from a slightly different perspective than what has been done earlier. Our main purpose has been to see how useful our new approach is and, therefore, we have mainly concentrated on one of the six paintings that we have had available, namely, the Järnefelt painting *Koli landscape*. However, some results based on the other paintings are included.

By looking at the intensity surfaces of the point patterns formed by the locations of fixations we have been able to see some differences between the eye movements of novices and non-novices on the Koli painting. The area that is visited most is more concentrated and more stable in time for the non-novice group than for the novice group. We also constructed a test to compare the intensity surfaces of novices and non-novices. Even though the observed difference was not statistically significant, there was some indication that the intensities cannot be considered equal. By comparing fixation duration distributions instead of bare means we were able to see a clear difference between the two groups under study. Non-novices

tend to have more short fixation durations than novices, meaning that they may only glance at some areas of the painting which novices do not tend to do. This was confirmed by the analysis based on the other paintings.

To have some idea whether the painting itself affects how novices and non-novices look at it, we compared the fixation processes of each group on three pairs of paintings, Järnefelt–Monet, Suomi–Kandinsky, and Poussin–Tammi. The intensity surface of fixations vary of course from painting to painting and it is not meaningful to compare them on different paintings, but it does make sense to compare the duration distributions. An interesting observation was that in the non-novice group the distribution of duration of fixations is more consistent from painting to painting than in the novice group.

Ideally, to be able to have a more detailed comparison of the eye movements of novices and non-novices, we would like to be able to describe the complete fixation process on a painting, that is, where, when, and how long the gaze of a person stays in different parts of the painting. Therefore, we have introduced a simple reference model for the fixation process. The model is a spatio-temporal point process model, where we used the fixation intensity and duration distributions estimated from the data. In order to restrict how far the gaze usually jumps, we also estimated the saccade length distribution from the data. To mimic the tendency that after having stayed on one area of the target painting for some time the gaze jumps to another area on the painting, the model suggests a long jump with some (small) probability p .

We fitted the model to the non-novice and novice groups separately for the eye movements on the Järnefelt painting and saw that the model fits quite well for both groups, except when measured by the scanpath length. The structure of the fixation process (described by our model) is similar in nature for non-novices and novices. The model-based simultaneous envelopes for the functional summary statistics can be used like confidence intervals to compare the two groups. However, since the envelopes overlap, we were not able to reveal any significant differences between the groups based on the model-based summaries.

Our idea here was to find a rather simple model that captures the most important features of the fixation process, and the reference model is quite good for this purpose. However, if the goal is to understand the complete dynamics of the fixation process, then the model constructed here is not good enough since it does not capture that the fixation process is changing in time. We noticed, for example, in Section 2 that both the intensity surface and the fixation duration distribution vary in time. Despite this, we used the same intensity surface and the same gamma distributions for the fixation durations and saccade lengths during the whole time period. We made some experiments by using six intensity surfaces (for each 30 second interval) and two distributions for saccade lengths (one for the first 30 seconds and one for the rest of the three minute time interval) in the reference model, but these modifications did not improve the goodness of fit of the model.

One possible way to improve the model could be to take the order of fixations into account. The fixation process could be regarded as a realization of a sequential point process which can be used to access the nonstationarity of the fixation process. The first author of this paper is currently working on that issue [see Penttinen and Ylitalo (in press)].

To conclude, our new set of tools allowed us to make more detailed comparisons between how novices and non-novices look at paintings than reported in the literature. As far as we know, such comparisons of fixation intensity surfaces, duration distributions, and coverages have not been presented earlier in eye movement studies.

Acknowledgments. The authors are grateful to Pertti Saariluoma, Sari Kuuva, María Álvarez Gil, Jarkko Hautala, and Tuomo Kujala for providing the data and to Antti Penttinen for helpful comments and discussions. The authors thank the two anonymous reviewers, Associate Editor, and Editor for their valuable comments that helped to significantly improve the paper.

The mobility funding provided by the University of Jyväskylä for the first author (AKY) is highly appreciated.

SUPPLEMENTARY MATERIAL

Supplement A: Statistics for fixations inside the picture (DOI: [10.1214/16-AOAS921SUPPA](https://doi.org/10.1214/16-AOAS921SUPPA); .pdf). Table includes information about the fixations, such as total number of fixations and mean fixation duration, for each subject.

Supplement B: Statistics for saccades inside the picture (DOI: [10.1214/16-AOAS921SUPPB](https://doi.org/10.1214/16-AOAS921SUPPB); .pdf). Table includes information about the saccades, such as mean saccade duration and mean saccade length, for each subject.

Supplement C: The most visited areas during the six 30 second intervals for novices and non-novices (DOI: [10.1214/16-AOAS921SUPPC](https://doi.org/10.1214/16-AOAS921SUPPC); .pdf). Figures show the top 5% and top 1% intensities in 30 second intervals (0–30 s), (30–60 s), (60–90 s), (90–120 s), (120–150 s), and (150–180 s) for novices and non-novices.

Supplement D: Results for the comparison of novices and non-novices for all six paintings (DOI: [10.1214/16-AOAS921SUPPD](https://doi.org/10.1214/16-AOAS921SUPPD); .pdf). Results for the intensity surface and fixation duration distribution comparisons between novices and non-novices for all six paintings used in the experiment.

Supplement E: Results for the groupwise fixation duration distribution comparisons for three pairs of paintings (DOI: [10.1214/16-AOAS921SUPPE](https://doi.org/10.1214/16-AOAS921SUPPE); .pdf). Results for the fixation duration distribution comparisons within novice and non-novice groups for the three pairs of paintings: Järnefelt–Monet, Kandinsky–Suomi, and Poussin–Tammi.

REFERENCES

- ANONYMOUS (2000). Brochure of Art Centre Salmela. Mäntyharju, Finland.
- ANTES, J. R. (1974). The time course of picture viewing. *J. Exp. Psychol.* **103** 62–70.
- BADDELEY, A., RUBAK, E. and TURNER, R. (2015). *Spatial Point Patterns: Methodology and Applications with R*. Chapman & Hall/CRC Press, London. Available at <http://www.taylorandfrancis.com/books/details/9781482210200/>.
- BAILEY, T. C. and GATRELL, A. C. (1995). *Interactive Spatial Data Analysis*. Longman, Harlow.
- BARLOW, H. B. (1952). Eye movements during fixation. *J. Physiol.* **116** 290–306.
- BARNARD, G. A. (1963). Contribution to the discussion of professor Bartlett's paper. *J. Roy. Statist. Soc. Ser. A* **25** 294.
- BARTHELMÉ, S., TRUKENBROD, H., ENGBERT, R. and WICHMANN, F. (2013). Modelling fixation locations using spatial point processes. *J. Vis.* **13** 1–34.
- BERMAN, M. and DIGGLE, P. (1989). Estimating weighted integrals of the second-order intensity of a spatial point process. *J. Roy. Statist. Soc. Ser. B* **51** 81–92. [MR0984995](#)
- BUSWELL, G. T. (1935). *How People Look at Pictures—A Study of the Psychology of Perception in Art*. The Univ. Chicago Press, Chicago, IL.
- DIGGLE, P. J. (1985). A kernel method for smoothing point process data. *J. R. Stat. Soc. Ser. C*. **34** 138–147.
- DOKSUM, K. A. and SIEVERS, G. L. (1976). Plotting with confidence: Graphical comparisons of two populations. *Biometrika* **63** 421–434. [MR0443210](#)
- DUCHOWSKI, A. T. (2002). A breadth-first survey of eye tracking applications. *Behav. Res. Methods Instrum. Comput.* **34** 455–470.
- ENGBERT, R., TRUKENBROD, H. A., BARTHELMÉ, S. and WICHMANN, F. A. (2015). Spatial statistics and attentional dynamics in scene viewing. *J. Vis.* **15** 1–17.
- FINDLAY, J. M. (2009). Saccadic eye movement programming: Sensory and attentional factors. *Psychol. Res.* **73** 127–135.
- GATRELL, A. C., BAILEY, T. C., DIGGLE, P. J. and ROWLINGSON, B. S. (1996). Spatial point pattern analysis and its application in geographical epidemiology. *Transactions of the Institute of British Geographers, New Series* **21** 256–274.
- HAZELTON, M. L. (2007). Kernel estimation of risk surfaces without the need for edge correction. *Stat. Med.* **27** 2269–2272.
- HENDERSON, J. M. and HOLLINGWORTH, A. (1999). High-level scene perception. *Annu. Rev. Psychol.* **50** 243–271.
- KELSALL, J. E. and DIGGLE, P. J. (1995a). Kernel estimation of relative risk. *Bernoulli* **1** 3–16. [MR1354453](#)
- KELSALL, J. E. and DIGGLE, P. J. (1995b). Non-parametric estimation of spatial variation in relative risk. *Stat. Med.* **14** 2335–2342.
- KINSLER, V. and CARPENTER, R. H. S. (1995). Saccadic eye movements while reading music. *Vis. Res.* **35** 1447–1458.
- KOMOGARTSEV, O. V., RYU, Y. S. and KOH, D. H. (2009). Quick models for saccade amplitude prediction. *Journal of Eye Movement Research* **1** 1–13.
- KRISTJANSON, A. F. and ANTES, J. R. (1989). Eye movement analysis of artists and nonartists viewing paintings. *Vis. Arts Res.* **15** 21–30.
- LOCHER, P. J. (2006). The usefulness of eye movement recordings to subject an aesthetic episode with visual art to empirical scrutiny. *Psychol. Sci.* **48** 106–114.
- MACKWORTH, N. H. and MORANDI, A. J. (1967). The gaze selects informative details within pictures. *Atten. Percept. Psychophys.* **2** 547–552.

- MANOR, B. R. and GORDON, E. (2003). Defining the temporal threshold for ocular fixation in free-viewing visuocognitive tasks. *J. Neurosci. Methods* **128** 85–93.
- MIKKOLA, K., ED. (1997). Risto Suomi. Amos Anderson Art Museum, publications, new series, no 25.
- MYLLYMÄKI, M., MRKVICKA, T., GRABARNIK, P., SEIJO, H. and HAHN, U. (2016). Global envelope tests for spatial processes. *J. Roy. Statist. Soc. Ser. B*. DOI:10.1111/rssb.12172.
- MYLLYMÄKI, M., GRABARNIK, P., SEIJO, H. and STOYAN, D. (2015). Deviation test construction and power comparison for marked spatial point patterns. *Spat. Stat.* **11** 19–34. MR3311854
- NAGASAWA, S., YIM, S. and HONGO, H. (2005). Feasibility study on marketing research using eye movement: An investigation of image presentation using an “eye camera” and data processing. *J. Adv. Comput. Intell. Intell. Informa.* **9** 440–452.
- NOTON, D. and STARK, L. (1971). Scanpaths in saccadic eye movements while viewing and recognizing patterns. *Vis. Res.* **11** 929–942.
- PENTTINEN, A. and YLITALO, A.-K. (in press). Deducing self-interaction in eye movement data using sequential spatial point processes. *Spatial Statistics*. DOI:10.1016/j.spasta.2016.03.005.
- THE YORCK PROJECT (2002). 10.000 Meisterwerke der Malerei, DVD-ROM, 2002. ISBN 3936122202. Distributed by DIRECTMEDIA Publishing GmbH.
- RAYNER, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychol. Bull.* **124** 372–422.
- RAYNER, K. (2009). Eye movements and attention in reading, scene perception, and visual search. *Q. J. Exp. Psychol., A Hum. Exp. Psychol.* **62** 1457–1506.
- RIPLEY, B. D. (1988). *Statistical Inference for Spatial Processes*. Cambridge Univ. Press, Cambridge. MR0971986
- SUNDELL, D. (1986). *Eero Järnefelt (1863–1937), Retretti 25.5.–21.9.1986*. Retretti.
- VOGT, S. and MAGNUSSEN, S. (2007). Expertise in pictorial perception: Eye-movement patterns and visual memory in artists and laymen. *Perception* **36** 91–100.
- WADE, N. J. (2010). Pioneers of eye movement research. *I-Perception* **1** 33–68.
- WAKEFIELD, J. C., KELSALL, J. E. and MORRIS, S. E. (2000). Clustering, cluster detection, and spatial variation in risk. In *Spatial Epidemiology: Methods and Applications* (P. Elliot, J. C. Wakefield, N. G. Best and D. J. Briggs, eds.) 128–152. Oxford Univ. Press, London.
- WILK, M. B. and GNANADESIKAN, R. (1968). Probability plotting methods for the analysis of data. *Biometrika* **55** 1–17.
- YARBUS, A. L. (1967). *Eye Movements and Vision*. Plenum Press, New York.
- YLITALO, A.-K., SÄRKKÄ, A. and GUTTORP, P. (2016a). Supplement to “What we look at in paintings: A comparison between experienced and inexperienced art viewers.” DOI:10.1214/16-AOAS921SUPPA.
- YLITALO, A.-K., SÄRKKÄ, A. and GUTTORP, P. (2016b). Supplement to “What we look at in paintings: A comparison between experienced and inexperienced art viewers.” DOI:10.1214/16-AOAS921SUPPB.
- YLITALO, A.-K., SÄRKKÄ, A. and GUTTORP, P. (2016c). Supplement to “What we look at in paintings: A comparison between experienced and inexperienced art viewers.” DOI:10.1214/16-AOAS921SUPPC.
- YLITALO, A.-K., SÄRKKÄ, A. and GUTTORP, P. (2016d). Supplement to “What we look at in paintings: A comparison between experienced and inexperienced art viewers.” DOI:10.1214/16-AOAS921SUPPD.
- YLITALO, A.-K., SÄRKKÄ, A. and GUTTORP, P. (2016e). Supplement to “What we look at in paintings: A comparison between experienced and inexperienced art viewers.” DOI:10.1214/16-AOAS921SUPPE.

A.-K. YLITALO
DEPARTMENT OF MUSIC
AND
DEPARTMENT OF MATHEMATICS AND STATISTICS
UNIVERSITY OF JYVASKYLA
P.O. BOX 35
FI-40014
FINLAND
E-MAIL: anna-kaisa.ylitalo@jyu.fi

A. SÄRKKÄ
MATHEMATICAL SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
AND UNIVERSITY OF GOTHENBURG
SE-412 96 GOTHENBURG
SWEDEN
E-MAIL: aila@chalmers.se

P. GUTTORP
DEPARTMENT OF STATISTICS
UNIVERSITY OF WASHINGTON
BOX 354322
SEATTLE, WASHINGTON 98195-4322
USA
AND
NORWEGIAN COMPUTING CENTER
P.O. BOX 114 BLINDERN
NO-0314 OSLO
NORWAY
E-MAIL: peter@stat.washington.edu