



The Poisson transform for unnormalised statistical models

Simon Barthelme, Nicolas Chopin

► To cite this version:

Simon Barthelme, Nicolas Chopin. The Poisson transform for unnormalised statistical models. *Statistics and Computing*, 2015, 25 (4), pp.767-780. 10.1007/s11222-015-9559-4 . hal-02391270

HAL Id: hal-02391270

<https://hal.science/hal-02391270>

Submitted on 25 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Poisson transform for unnormalised statistical models

Simon Barthelmé, Nicolas Chopin

Abstract

Contrary to standard statistical models, *unnormalised* statistical models only specify the likelihood function up to a constant. While such models are natural and popular, the lack of normalisation makes inference much more difficult. Extending classical results on the multinomial-Poisson transform (Baker, 1994), we show that inferring the parameters of a unnormalised model on a space Ω can be mapped onto an equivalent problem of estimating the intensity of a Poisson point process on Ω . The unnormalised statistical model now specifies an *intensity function* that does not need to be normalised. Effectively, the normalisation constant may now be inferred as just another parameter, at no loss of information. The result can be extended to cover non-IID models, which includes for example unnormalised models for sequences of graphs (dynamical graphs), or for sequences of binary vectors. As a consequence, we prove that unnormalised parametric inference in non-IID models can be turned into a semi-parametric estimation problem. Moreover, we show that the *noise-contrastive estimation* method of Gutmann and Hyvärinen (2012) can be understood as an approximation of the Poisson transform, and extended to non-IID settings. We use our results to fit spatial Markov chain models of eye movements, where the Poisson transform allows us to turn a highly non-standard model into vanilla semi-parametric logistic regression.

Unnormalised statistical models are a core tool in modern machine learning, especially deep learning (Salakhutdinov and Hinton, 2009), computer vision (*Markov random fields*, Wang et al., 2013) and statistics for point processes (Gu and Zhu, 2001), network models (Caimo and Friel, 2011), directional data (Walker, 2011). They appear naturally whenever one can best describe data as having to conform to certain features: we may then define an energy function that measures how well the data conform to these constraints. While this way of formulating statistical models is extremely general and useful, immense technical difficulties may arise whenever the energy function involves some unknown parameters which have to be estimated from data. The reason is that the normalisation constant (which ensures that the distribution integrates to one) is in most cases impossible to compute. This prevents direct application of classical methods of maximum likelihood or Bayesian inference, which all depend on the unknown normalisation constant.

Many techniques have been developed in recent years for such problems, including contrastive divergence (Hinton, 2002; Bengio and Delalleau, 2009),

noise-contrastive estimation (Gutmann and Hyvärinen, 2012) and various forms of MCMC for Bayesian inference (Møller et al., 2006; Murray et al., 2012; Girolami et al., 2013). The difficulty is compounded when unnormalised models are used for non-IID data, either sequential data, or data that include covariates. If the data form a sequence of length n , there are now n normalisation constants to approximate. In our application we look at models of spatial Markov chains, where the transition density of the chain is specified up to a normalisation constant, and again one normalisation constant needs to be estimated per observation.

In the first Section, we show that unnormalised estimation is tightly related to the estimation of point process intensities, and formulate a *Poisson transform* that maps the log-likelihood of a model $\mathcal{L}(\boldsymbol{\theta})$ into an equivalent cost function $\mathcal{M}(\boldsymbol{\theta}, \boldsymbol{\nu})$ defined in an expanded space, where the latent variables $\boldsymbol{\nu}$ effectively estimate the normalisation constants. In the case of non-IID unnormalised models we show further that optimisation of $\mathcal{M}(\boldsymbol{\theta}, \boldsymbol{\nu})$ can be turned into a semi-parametric problem and addressed using standard kernel methods. In the second section, we show that the *noise-contrastive divergence* described in of Gutmann and Hyvärinen (2012) arises naturally as a tractable approximation of the Poisson transform, and that this new interpretation lets us extend its use to non-IID models. (Gutmann and Hyvärinen (2012) call the technique “noise-contrastive estimation”, but we use the term noise-contrastive *divergence* to designate the corresponding cost function.) Finally, we apply these results to a class of unnormalised spatial Markov chains that are natural descriptions of eye movement sequences.

1 Relationship to prior work

Some of the ideas we use here have appeared under different forms in classical statistics, machine learning and spatial statistics. The Poisson transform generalises the multinomial-Poisson transform developed by Baker (1994). It is also a special case of a general family of Bregman divergences introduced by Gutmann and ichiro Hirayama (2011), a special case of another family by Pihlaja et al. (2010), and finally can also be viewed as an empirical version of the generalised Kullback-Leibler divergence for unnormalised measures (Minka, 2005).

Noise-contrastive learning is studied in Gutmann and ichiro Hirayama (2011), although the relationship between logistic regression and estimation has been noted in other places (for example, in the spatial statistics literature, see Baddeley et al., 2010, Baddeley et al., 2014). We go further here in showing that the divergence defined by NCL converges uniformly to the Poisson transform, giving it a new interpretation as an approximate likelihood rather than just a divergence.

Mnih and Kavukcuoglu (2013) and Mnih and Teh (2012a) use the NCL technique in a class of non-IID unnormalised models. However, in the interest of computation time, they ignore normalisation constants. The results given here indicate clearly that neglecting normalisation constants leads in the general case

to non-convergent estimators, as illustrated in Section 4.1. Instead we develop a semi-parametric framework for non-IID estimation, which is both much faster than purely parametric techniques, as well as convergent.

2 The Poisson transform

In this section we show how unnormalised likelihoods can be turned into Poisson process likelihoods at no loss of information. We call the procedure the Poisson transform, as it generalises the Poisson-multinomial transform (Baker, 1994). We give two interpretations, one in terms of upper-bound maximisation, and one in terms of generalised KL divergences. We begin with the IID case, with the generalisation to non-IID data treated further into the text.

2.1 Background on Poisson point processes

Poisson point processes are described at length in Kingman (1993), and we only give here the merest outline. A *Inhomogeneous Poisson point process* (IPP) with intensity function $\lambda(\mathbf{y}) \geq 0$ over space Ω defines a distribution over the set of countable subsets \mathcal{S} of Ω , in such a way that, for any measurable subset $\mathcal{A} \subseteq \Omega$,

$$\# \{\mathcal{S} \cap \mathcal{A}\} \sim \text{Poi}(\lambda_{\mathcal{A}}), \quad \lambda_{\mathcal{A}} = \int_{\mathcal{A}} \lambda(\mathbf{y}) \, d\mathbf{y},$$

assuming $\lambda_{\mathcal{A}} < +\infty$. In words, the number of points to be found in subset \mathcal{A} has a Poisson distribution, with expectation given by the integral of the intensity function within \mathcal{A} ; in discrete spaces the integral may of course be interpreted as a sum. In particular, provided $\int \lambda(\mathbf{y}) \, d\mathbf{y} < +\infty$, the cardinal n of \mathcal{S} is finite, and has a Poisson distribution with expectation equal to the integral of $\lambda(\mathbf{y})$ over the domain (the fact follows from taking $\mathcal{A} = \Omega$). Assuming again $\int_{\Omega} \lambda(\mathbf{y}) \, d\mathbf{y} < +\infty$, the log-likelihood of observing set \mathcal{S} given the intensity function λ is given by:

$$\log p(\mathcal{S}|\lambda) = \sum_{\mathbf{y}_i \in \mathcal{S}} \log \lambda(\mathbf{y}_i) - \int_{\Omega} \lambda(\mathbf{y}) \, d\mathbf{y}. \quad (2.1)$$

2.2 The Poisson transform in the IID case

The Poisson transform is simply stated: when we have n observations from an unnormalised model on Ω , we may treat them as the realisation of a certain point process at no loss of information. This results in a mapping from a likelihood function $\mathcal{L}(\boldsymbol{\theta})$ to another, which we note $\mathcal{M}(\boldsymbol{\theta}, \nu)$, in an expanded space. $\mathcal{M}(\boldsymbol{\theta}, \nu)$ has the same global maximum as $\mathcal{L}(\boldsymbol{\theta})$ and confidence intervals are preserved.

First, the log-likelihood function for n IID observations \mathbf{y}_i from an unnormalised model $p(\mathbf{y}|\boldsymbol{\theta}) \propto \exp\{f_{\boldsymbol{\theta}}(\mathbf{y})\}$ can be written as:

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^n f_{\boldsymbol{\theta}}(\mathbf{y}_i) - n \log \left(\int_{\Omega} \exp\{f_{\boldsymbol{\theta}}(\mathbf{y})\} d\mathbf{y} \right) \quad (2.2)$$

and the ML estimate of $\boldsymbol{\theta}$ is the maximum of $\mathcal{L}(\boldsymbol{\theta})$. We introduce the following alternative likelihood function:

$$\mathcal{M}(\boldsymbol{\theta}, \nu) = \sum_{i=1}^n \{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\} - n \int_{\Omega} \exp\{f_{\boldsymbol{\theta}}(\mathbf{y}) + \nu\} d\mathbf{y} \quad (2.3)$$

which by (2.1) is, up to additive constant $n \log(n)$, the IPP likelihood on Ω for intensity function

$$\lambda(\mathbf{y}) = \exp\{f_{\boldsymbol{\theta}}(\mathbf{y}) + \nu + \log(n)\}.$$

Our first theorem shows that maximum likelihood estimation of $\boldsymbol{\theta}$ via $\mathcal{L}(\boldsymbol{\theta})$ or via $\mathcal{M}(\boldsymbol{\theta}, \nu)$ is equivalent.

Theorem 1. *The set of points $\boldsymbol{\theta}^*$ such that $\boldsymbol{\theta}^* \in \arg \max_{\boldsymbol{\theta} \in \Theta} \mathcal{L}(\boldsymbol{\theta})$ matches the set of points $\tilde{\boldsymbol{\theta}}$ such that $(\tilde{\boldsymbol{\theta}}, \tilde{\nu}) \in \arg \max_{\boldsymbol{\theta} \in \Theta, \nu \in \mathbb{R}} \mathcal{M}(\boldsymbol{\theta}, \nu)$ for some $\tilde{\nu}$. In particular, if $\arg \max_{\boldsymbol{\theta} \in \Theta} \mathcal{L}(\boldsymbol{\theta})$ is a singleton, then so is $\arg \max_{\boldsymbol{\theta} \in \Theta, \nu \in \mathbb{R}} \mathcal{M}(\boldsymbol{\theta}, \nu)$.*

Proof. For a fixed $\boldsymbol{\theta}$, $\mathcal{M}(\boldsymbol{\theta}, \nu)$ admits a unique maximum in ν at $\nu^*(\boldsymbol{\theta}) = -\log \int_{\Omega} \exp\{f_{\boldsymbol{\theta}}(\mathbf{y})\} d\mathbf{y}$, hence $\mathcal{M}(\boldsymbol{\theta}, \nu) \leq \mathcal{M}(\boldsymbol{\theta}, \nu^*(\boldsymbol{\theta})) = \mathcal{L}(\boldsymbol{\theta}) - n$. \square

There are several remarks to make at this stage. First, since $\nu^*(\boldsymbol{\theta}) = -\log \int_{\Omega} \exp\{f_{\boldsymbol{\theta}}(\mathbf{y})\} d\mathbf{y}$, maximising $\mathcal{M}(\boldsymbol{\theta}, \nu)$ can be interpreted as estimating the normalisation constant along with the parameters. There is no estimation cost incurred in treating the normalisation constant as a free parameter, since the global maxima of $\mathcal{L}(\boldsymbol{\theta})$ and $\mathcal{M}(\boldsymbol{\theta}, \nu)$ are the same.

Second, the usual way of computing confidence intervals for $\boldsymbol{\theta}$ is to invert the Hessian of $\mathcal{L}(\boldsymbol{\theta})$ at the mode. We show in the Appendix that the same confidence intervals can be obtained from the Hessian of $\mathcal{M}(\boldsymbol{\theta}, \nu)$ at the mode, so that the Poisson transform does not introduce any over or under-confidence. In addition, the Poisson-transformed likelihood can be used for *penalised* likelihood maximisation (see Application), does not introduce any spurious maxima, and in exponential families it can even be shown to preserve concavity (see Appendix).

Third, at this point we do not yet have a practical way of computing $\mathcal{M}(\boldsymbol{\theta}, \nu)$, since we have assumed that integrals of the form $\int_{\Omega} \exp\{f_{\boldsymbol{\theta}}(\mathbf{y}) + \nu\} d\mathbf{y}$ are intractable. The problem of approximating $\mathcal{M}(\boldsymbol{\theta}, \nu)$ is dealt with in Section 3, where we will see that among other possibilities it can be approximated by logistic regression via noise-contrastive divergence.

Before we deal with practical ways of approximating $\mathcal{M}(\boldsymbol{\theta}, \nu)$, we first generalise the Poisson transform to non-IID data.

2.3 The Poisson transform in the non-IID case

In the non-IID case we still have n datapoints $\mathbf{y}_1 \dots \mathbf{y}_n \in \Omega^n$ but their distribution is allowed to vary. For example the n datapoints might form a Markov chain with (unnormalised) transition density

$$p_{\boldsymbol{\theta}}(\mathbf{y}_t | \mathbf{y}_{t-1}) \propto \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}_t | \mathbf{y}_{t-1})\}$$

which leads to the log-likelihood

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{t=1}^n \left[f_{\boldsymbol{\theta}}(\mathbf{y}_t | \mathbf{y}_{t-1}) - \log \int_{\Omega} \exp \{f_{\boldsymbol{\theta}}(\mathbf{y} | \mathbf{y}_{t-1})\} d\mathbf{y} \right]. \quad (2.4)$$

(The initial point \mathbf{y}_0 is treated as a constant.) Another example is models with covariates \mathbf{x}_i , expressed as $p(\mathbf{y}_i | \mathbf{x}_i, \boldsymbol{\theta}) \propto \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}_i | \mathbf{x}_i)\}$. These two cases are highly similar and for brevity we focus on the sequential case, which we use in our application.

Our first step is to extend the Poisson transform (2.3) to yield a function $\mathcal{M}(\boldsymbol{\theta}, \boldsymbol{\nu})$ where $\boldsymbol{\nu}$ is now a vector of dimension n (one per conditional distribution), $\boldsymbol{\nu} = (\nu_1, \dots, \nu_n)$ and

$$\begin{aligned} \mathcal{M}(\boldsymbol{\theta}, \boldsymbol{\nu}) &= \sum_{t=1}^n \{f_{\boldsymbol{\theta}}(\mathbf{y}_t | \mathbf{y}_{t-1}) + \nu_{t-1}\} \\ &\quad - \int_{\Omega} \left[\sum_{t=1}^n \exp \{f_{\boldsymbol{\theta}}(\mathbf{y} | \mathbf{y}_{t-1}) + \nu_{t-1}\} \right] d\mathbf{y}. \end{aligned} \quad (2.5)$$

Theorem 2. *The set of points $\boldsymbol{\theta}^*$ such that $\boldsymbol{\theta}^* \in \arg \max_{\boldsymbol{\theta} \in \Theta} \mathcal{L}(\boldsymbol{\theta})$ matches the set of points $\tilde{\boldsymbol{\theta}}$ such that $(\tilde{\boldsymbol{\theta}}, \boldsymbol{\nu}^*) = \arg \max_{\boldsymbol{\theta} \in \Theta, \boldsymbol{\nu} \in \mathbb{R}^n} \mathcal{M}(\boldsymbol{\theta}, \boldsymbol{\nu})$.*

Proof. The proof is along the same lines as that of the Theorem 1: maximising $\mathcal{M}(\boldsymbol{\theta}, \boldsymbol{\nu})$ in ν_{t-1} gives $\nu_{t-1}^*(\boldsymbol{\theta}) = -\log \int_{\Omega} \exp \{f_{\boldsymbol{\theta}}(\mathbf{y} | \mathbf{y}_{t-1})\} d\mathbf{y}$, and $\mathcal{M}(\boldsymbol{\theta}, \boldsymbol{\nu}^*(\boldsymbol{\theta})) = \mathcal{L}(\boldsymbol{\theta}) - n$. \square

Note that while $\mathcal{L}(\boldsymbol{\theta})$ involves the sum of n separate integrals, $\mathcal{M}(\boldsymbol{\theta}, \boldsymbol{\nu})$ involves a single integral over a sum. Further, since

$$\nu_{t-1}^*(\boldsymbol{\theta}) = -\log \left(\int_{\Omega} \exp \{f_{\boldsymbol{\theta}}(\mathbf{y} | \mathbf{y}_{t-1})\} d\mathbf{y} \right)$$

the optimal value of ν_{t-1} is a function of \mathbf{y}_{t-1} *only*. This means that we can think of the integration constants as (hopefully smooth) functions of the previous point \mathbf{y}_{t-1} . This leads to the following result: let \mathcal{F} denote an appropriate function space that contains the function $\chi : \Omega \rightarrow \mathbb{R}$ such that

$\chi(\mathbf{u}) = -\log \int_{\Omega} \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}|\mathbf{u})\} d\mathbf{y}$. We introduce the following functional

$$\begin{aligned} \mathcal{M}_{\chi}(\boldsymbol{\theta}, \chi) &= \sum_{t=1}^n \{f_{\boldsymbol{\theta}}(\mathbf{y}_t|\mathbf{y}_{t-1}) + \chi(\mathbf{y}_{t-1})\} \\ &\quad - \int_{\Omega} \sum_t \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}|\mathbf{y}_{t-1}) + \chi(\mathbf{y}_{t-1})\} d\mathbf{y}. \end{aligned} \quad (2.6)$$

Corollary 3. *The set of points $\boldsymbol{\theta}^*$ such that $\boldsymbol{\theta}^* \in \arg \max_{\boldsymbol{\theta} \in \Theta} \mathcal{L}(\boldsymbol{\theta})$ matches the set of points $\tilde{\boldsymbol{\theta}}$ such that $(\tilde{\boldsymbol{\theta}}, \chi) \in \arg \max_{\boldsymbol{\theta} \in \Theta, \chi \in \mathcal{F}} \mathcal{M}_{\chi}(\boldsymbol{\theta}, \chi)$.*

We can use this Corollary to turn inference on unnormalised models into a semiparametric problem, where $\boldsymbol{\theta}$ is estimated parametrically and the normalisation constants are estimated as a non-parametric function $\chi(\mathbf{y}_{t-1})$.

In the formulation used by Corollary 3 there exists possibly (uncountably) many optimal normalisation functions χ , ie. functions that solve $\arg \max_{\boldsymbol{\theta} \in \Theta, \chi \in \mathcal{F}} \mathcal{M}_{\chi}(\boldsymbol{\theta}, \chi)$.

All that is required is that they interpolate the values of the normalisations constants for the various \mathbf{y}_{t-1} in the dataset. To get a unique optimal normalisation function we need to regularise the non-parametric part.

A classical way to solve non-parametric regression problems is to model the non-parametric part as belonging to a Reproducible Kernel Hilbert Space (RKHS), and to add regularisation by including a penalty. The following result shows that penalised non-parametric estimation can be made consistent, and the optimal normalisation function becomes uniquely defined.

Proposition 4. *Let \mathcal{H} denote a RKHS, with kernel function $k(\mathbf{y}, \mathbf{y}')$ and $\|f\|_{\mathcal{H}}$ the corresponding norm. Suppose \mathcal{H} contains one optimal normalisation function, i.e. there exists an $\chi^*(\mathbf{u}) \in \mathcal{H}$, with $\|\chi^*\|_{\mathcal{H}} < \infty$. Then there exists a value $\lambda_0 > 0$ such that the set of maximum likelihood points $\boldsymbol{\theta}^* \in \arg \max_{\boldsymbol{\theta} \in \Theta} \mathcal{L}(\boldsymbol{\theta})$ matches the set of points penalised estimates $\tilde{\boldsymbol{\theta}}$ defined by:*

$$(\tilde{\boldsymbol{\theta}}, \chi) \in \arg \max_{\boldsymbol{\theta} \in \Theta, \chi \in \mathcal{H}} \mathcal{M}_{\chi}(\boldsymbol{\theta}, \chi) - \lambda \|\chi\|_{\mathcal{H}}^2 \quad (2.7)$$

i.e., the penalised non-parametric Poisson estimator is equivalent to the maximum-likelihood estimator.

Proof. The penalised problem is equivalent to the following constrained optimisation problem:

$$\begin{aligned} &\arg \max_{\boldsymbol{\theta} \in \Theta, \chi \in \mathcal{H}} \quad \mathcal{M}_{\chi}(\boldsymbol{\theta}, \chi) \\ &\text{subject to} \quad \|\chi\|_{\mathcal{H}}^2 \leq \rho \end{aligned}$$

for some value ρ dependent on λ (this follows from writing the Lagrangian). By the assumption that there exists an optimal normalisation function in \mathcal{H}

with finite norm, there exists a $\rho_0 < \infty$ such that the constraint is irrelevant and solving the constrained problem above is equivalent to solving the non-penalised problem $\operatorname{argmax}_{\boldsymbol{\theta} \in \Theta, \chi \in \mathcal{F}} \mathcal{M}_\chi(\boldsymbol{\theta}, \chi)$ from Lemma 3. Correspondingly there exists a penalisation parameter $\lambda_0 > 0$ such that the penalised estimate (2.7) matches the non-penalised estimate. \square

Remark 5. For fixed $\boldsymbol{\theta}$, $\operatorname{argmax}_{\chi \in \mathcal{H}} \mathcal{M}_\chi(\boldsymbol{\theta}, \chi) - \lambda |\chi|_{\mathcal{H}}^2$ has a unique solution that can be expressed as $\chi(\mathbf{u}) = \sum \alpha_{t-1} k(\mathbf{u}, \mathbf{y}_{t-1})$

Proof. The result follows from a straightforward application of the Representer Theorem (see Schölkopf and Smola, 2001, page 90). \square

We have only established so far that there exists a value λ_0 so that the penalised non-parametric estimator is equivalent to the ML estimator. We cannot expect to know that value in advance, and so λ_0 needs to be estimated from the data. The following Corollary comes to the rescue:

Corollary 6. *Note $\boldsymbol{\theta}(\lambda), \chi(\lambda)$ the solution for the penalised problem (eq. (2.7)) with regularisation parameter λ . For all $\lambda \leq \lambda_0$, $\mathcal{M}_\chi(\boldsymbol{\theta}(\lambda), \chi(\lambda)) = \mathcal{M}_\chi(\boldsymbol{\theta}(\lambda_0), \chi(\lambda_0))$, i.e. there is no further improvement to the optimal value of the Poisson transform by relaxing the penalty beyond λ_0 .*

Proof. The proof follows again from the constrained formulation. By λ_0 we have already found the optimal solution and there is no point relaxing the constraint further. \square

What the result suggests is that we could start with a high value for λ , perform the optimisation, and reduce the value of λ until the value of $\mathcal{M}_\chi(\boldsymbol{\theta}(\lambda), \chi(\lambda))$ stops improving. We will then have found the most “simple” function that interpolates the normalisation constants. Unfortunately Corollary 6 does not hold for noise-contrastive divergence, and so a different strategy (such as cross-validation) has to be used for selecting λ . We return to the issue in the examples.

3 Practical approximations for the Poisson transform

The Poisson transform gives us an alternative likelihood function for estimation, but one that still involves an intractable integral. In this section we briefly describe some practical approximations. One is based on importance sampling and leads to an unbiased estimate of the gradient (meaning that novel stochastic gradient and approximate Langevin sampling methods are possible). The second is based on logistic regression: we show that the noise-contrastive divergence of Gutmann and Hyvärinen (2012) approximates the Poisson-transformed likelihood. Using that connection, estimation in any non-IID setting can be turned into a semiparametric classification problem.

3.1 Unbiased estimation of the gradient

The first derivatives of $\mathcal{M}(\boldsymbol{\theta}, \nu)$ (eq. 2.3) equal:

$$\begin{aligned} \frac{1}{n} \frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{M}(\boldsymbol{\theta}, \nu) &= \frac{1}{n} \sum_{i=1}^n \frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}}(\mathbf{y}_i) - \int_{\Omega} \frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}}(\mathbf{y}) \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}) + \nu\} d\mathbf{y} \\ \frac{1}{n} \frac{\partial}{\partial \nu} \mathcal{M}(\boldsymbol{\theta}, \nu) &= 1 - \int_{\Omega} \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}) + \nu\} d\mathbf{y} \end{aligned}$$

The integrals on the right hand side can be estimated unbiasedly by Monte Carlo, which is not true in general for the untransformed likelihood. The availability of an unbiased estimator for the gradient means that stochastic gradient algorithms (and their MCMC counterpart, approximate Langevin sampling, Welling and Teh, 2011) can be applied directly. The resulting method has a straightforward interpretation, since we simply adjust ν until $\exp \{f_{\boldsymbol{\theta}}(\mathbf{y}) + \nu\}$ normalises to 1 on average.

3.2 Logistic likelihood as an approximation: IID case

In this section we show how to approximate Poisson-transformed likelihoods, see (2.3) and (2.4), using logistic regression. Reductions to logistic regression appear in many places in the statistical literature. In the context of estimation it is described in the well-known textbook of Hastie et al. (2003) and in detail in Baddeley et al. (2010). The use of logistic regression to estimate *normalisation constants* is described in Geyer (1994). Recently Gutmann and Hyvärinen (2012) introduced a more general theory which they call “noise-contrastive divergence”, and show that logistic regression can be used for joint estimation of parameters and normalisation constants.

The essence of noise-contrastive divergence is to try and teach a logistic classifier to tell true data $\mathcal{S} = \{\mathbf{y}_1, \dots, \mathbf{y}_n\}$, generated from $p_{\boldsymbol{\theta}}(\mathbf{y})$, from random reference data $\mathcal{R} = \{\mathbf{r}_1, \dots, \mathbf{r}_m\}$, generated from some distribution with density $q(\mathbf{r})$. Picking a point \mathbf{u} at random from $\mathcal{S} \cup \mathcal{R}$, and denoting $z = 1$ (resp. $z = 0$) the event that \mathbf{u} comes from \mathcal{S} (resp. \mathcal{R}), one obtains the following log odds ratio:

$$\log \frac{p(z=1|\mathbf{u})}{p(z=0|\mathbf{u})} = \log p_{\boldsymbol{\theta}}(\mathbf{u}) - \log q(\mathbf{u}) + \log(n/m). \quad (3.1)$$

If we assume additionally that $p_{\boldsymbol{\theta}}(\mathbf{y})$ is unnormalised, $p_{\boldsymbol{\theta}}(\mathbf{y}) \propto \exp \{f_{\boldsymbol{\theta}}(\mathbf{y})\}$, one may replace above, in the same spirit as in our Poisson transform, the term $\log p_{\boldsymbol{\theta}}(\mathbf{u})$ by $f_{\boldsymbol{\theta}}(\mathbf{u}) + \nu$, leading to

$$\log \frac{p(z=1|\mathbf{u})}{p(z=0|\mathbf{u})} = f_{\boldsymbol{\theta}}(\mathbf{u}) + \nu - \log q(\mathbf{u}) + \log(n/m). \quad (3.2)$$

This leads to following simple recipe: generate reference data \mathcal{R} , then estimate jointly $(\boldsymbol{\theta}, \nu)$ by fitting the logistic regression (3.2) to the dataset $\mathcal{S} \cup \mathcal{R}$, with points in \mathcal{S} (resp. \mathcal{R}) labelled as $z_i = 1$ (resp. $z_i = 0$).

The obvious connection between our Poisson transform and the noise-contrastive approach is that in both cases the log normalising constant is treated as a free parameter. The following result reveals that this connection is actually deeper.

Theorem 7. *For fixed $\boldsymbol{\theta}$, ν , and $\mathcal{S} = \{\mathbf{y}_1, \dots, \mathbf{y}_n\}$, and under the assumption that $f_{\boldsymbol{\theta}}(\mathbf{y}) - \log q(\mathbf{y}) \leq C(\boldsymbol{\theta})$ for all $\mathbf{y} \in \Omega$, the log-likelihood of the logistic regression defined above:*

$$\begin{aligned} \mathcal{R}^m(\boldsymbol{\theta}, \nu) &= \sum_{i=1}^n \log \left[\frac{n \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\}}{n \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\} + m q(\mathbf{y}_i)} \right] \\ &\quad + \sum_{j=1}^m \log \left[\frac{m q(\mathbf{r}_j)}{n \exp \{f_{\boldsymbol{\theta}}(\mathbf{r}_j) + \nu\} + m q(\mathbf{r}_j)} \right] \end{aligned}$$

is such that

$$\mathcal{R}^m(\boldsymbol{\theta}, \nu) + n \log(m/n) + \sum_{i=1}^n \log q(\mathbf{y}_i) \rightarrow \mathcal{M}(\boldsymbol{\theta}, \nu) \quad (3.3)$$

almost surely as $m \rightarrow +\infty$, relative to the randomness induced by the reference points $\mathcal{R} = \{\mathbf{r}_1, \dots, \mathbf{r}_m\}$.

Proof. See Appendix. \square

The theorem above establishes that $\mathcal{R}^m(\boldsymbol{\theta}, \nu)$ converges to $\mathcal{M}(\boldsymbol{\theta}, \nu)$ pointwise (up to a constant). Uniform convergence (with respect to $\boldsymbol{\theta}$) may be proved under stronger conditions. As a corollary, one obtains that the MLE based on $\mathcal{R}^m(\boldsymbol{\theta}, \nu)$ converges to the intractable MLE of $\mathcal{M}(\boldsymbol{\theta}, \nu)$ as $m \rightarrow +\infty$.

Theorem 8. *Assume that (i) Θ is a bounded set, that (ii) $|f_{\boldsymbol{\theta}}(\mathbf{y}) - \log q(\mathbf{y})| \leq C$ for some $C > 0$ and all $\mathbf{y} \in \Omega$, that (iii) $|f_{\boldsymbol{\theta}}(\mathbf{y}) - f_{\boldsymbol{\theta}'}(\mathbf{y})| \leq \kappa(\mathbf{y}) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|$ for all $\mathbf{y} \in \Omega$ and $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta$, with $\mathbb{E}_q[\kappa] < \infty$, that (iv) there exists $\hat{\boldsymbol{\theta}}$ such that $\mathcal{L}(\hat{\boldsymbol{\theta}}) > \sup_{d(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}) \geq \epsilon} \mathcal{L}(\boldsymbol{\theta})$, for any $\epsilon > 0$. Then for fixed $\mathcal{S} = \{\mathbf{y}_1, \dots, \mathbf{y}_n\}$, and $(\tilde{\boldsymbol{\theta}}_m, \tilde{\nu}_m)$ such that $\mathcal{R}^m(\tilde{\boldsymbol{\theta}}_m, \tilde{\nu}_m) = \sup_{(\boldsymbol{\theta}, \nu) \in \Theta \times \mathbb{R}} \mathcal{R}^m(\boldsymbol{\theta}, \nu)$, one has*

$$\tilde{\boldsymbol{\theta}}_m \rightarrow \hat{\boldsymbol{\theta}} \quad a.s.$$

as $m \rightarrow +\infty$, relative to the randomness induced by the reference points $\mathcal{R} = \{\mathbf{r}_1, \dots, \mathbf{r}_m\}$.

Proof. See Appendix. \square

In particular, the limit of $\tilde{\boldsymbol{\theta}}_m$ as $m \rightarrow +\infty$ has the same properties as the MLE of $\mathcal{L}(\boldsymbol{\theta})$, and thus is consistent, and asymptotically efficient. The theorem above assumes implicitly that the MLE of the logistic regression (with log-likelihood $\mathcal{R}^m(\boldsymbol{\theta}, \nu)$) is well defined, but this is a mild assumption: e.g. if the considered model corresponds to an exponential family, $f_{\boldsymbol{\theta}}(\mathbf{y}) = \boldsymbol{\theta}^T S(\mathbf{y})$, then it is easy to check that $\mathcal{R}^m(\boldsymbol{\theta}, \nu)$ is a concave function of $(\boldsymbol{\theta}, \nu)$.

3.3 Logistic likelihood as an approximation: non IID case

Putting together Theorem 4 and the results in Section 2.3 leads to the following extension of noise-contrastive divergence to non-IID problems. For an unnormalised Markov model $p_{\theta}(\mathbf{y}_t|\mathbf{y}_{t-1}) \propto \exp\{f_{\theta}(\mathbf{y}_t|\mathbf{y}_{t-1})\}$, for data $\mathcal{S} = \{\mathbf{y}_1, \dots, \mathbf{y}_n\}$, generate $m = kn$ reference datapoints \mathbf{r}_{jt} from kernel $q(\mathbf{r}_{jt}|\mathbf{y}_{t-1})$, $j = 1, \dots, k$ (i.e. k points \mathbf{r}_j are generated from ancestor \mathbf{y}_{t-1} , for each t), then fit the semi-parametric logistic regression model that corresponds to the log odds ratio function:

$$\log \frac{p(z=1|\mathbf{u}_{t-1}, \mathbf{u}_t)}{p(z=0|\mathbf{u}_{t-1}, \mathbf{u}_t)} = f_{\theta}(\mathbf{u}_t|\mathbf{u}_{t-1}) + \chi(\mathbf{u}_{t-1}) - \log q(\mathbf{u}_t|\mathbf{u}_{t-1}) + \log(n/m) \quad (3.4)$$

where $(\mathbf{u}_{t-1}, \mathbf{u}_t)$ represents a pair taken at random from $\{(\mathbf{y}_{t-1}, \mathbf{y}_t)\} \cup \{(\mathbf{y}_{t-1}, \mathbf{r}_{jt})\}$. The parameters of this logistic model are vector θ , scalar ν , and function $\chi : \mathcal{Y} \rightarrow \mathbb{R}$, which is why this model is indeed semi-parametric. In practice, fitting such a model is easily achieved using an appropriate regulariser (we use smoothing splines in our application).

The interpretation of the above procedure follows the same lines as in the previous section: for $m \rightarrow +\infty$, the log-likelihood of this logistic model converges to that of the semi-parametric Poisson model defined in Theorem 3; in particular, χ must be seen as an estimator of the (typically smooth) function $\mathbf{y}_{t-1} \rightarrow -\log \int \exp\{f_{\theta}(\mathbf{y}|\mathbf{y}_{t-1})\} d\mathbf{y}$.

More generally, one may extend this approach to other non-IID models. For instance, if $p_{\theta}(\mathbf{y}_t) \propto \exp\{f_{\theta}(\mathbf{y}_t|\mathbf{x}_t)\}$, where \mathbf{x}_t are covariates, then fit the same type of semi-parametric logistic regression as above, but with χ a function of covariates \mathbf{x}_t .

4 Applications to spatial Markov chains

4.1 A toy example

We begin with a toy example that shows how inference based on the Poisson Transform can be implemented in the non-IID case, and show that semi-parametric inference using non-contrastive divergence can be almost as efficient as maximum-likelihood (and much more efficient than completely parametric non-contrastive divergence). In addition, we will see that ignoring normalisation constants as done by Mnih and Teh (2012b) and Mnih and Kavukcuoglu (2013) can lead to severe bias. We have made available a detailed companion document for this section, which includes all the code necessary to replicate our results in R.

Our toy example is a Markov chain in $[-1, 1]$, with transition probability:

$$p_{\theta}(y_t|y_{t-1}) \propto \exp\left\{\theta_1 y_t - \frac{1}{2}\theta_2 (y_t - y_{t-1})^2\right\} \mathbb{I}_{[-1,1]}(y_t). \quad (4.1)$$

We picked this example because it is a simplified version of the spatial




Fig. 4.1: Two realisations from the toy model, a Markov chain constrained to the interval $[-1, 1]$ (eq. 4.1). Red, unfilled dots: $\theta_1 = 0, \theta_2 = 2$. Blue, solid dots: $\theta_1 = -2, \theta_2 = 10$. As expected, the second chain shows a bias towards negative values as well as stronger autocorrelation.

Markov chains we study in the following section. Two realisations from the chain are shown on Fig. 4.1.

In this one-dimensional example it is of course easy to compute the normalisation constant using numerical integration, and thus maximum likelihood inference is possible. To use non-constrastive divergence, we need to pick a reference kernel, and here a uniform, IID distribution does the job quite well: $q(y|y_{t-1}) = \frac{1}{2}\mathbb{I}_{[-1,1]}(y)$.

Positive examples for the logistic regression are formed from actual pairs (y_t, y_{t-1}) , negative examples are formed from pairs (r_{it}, y_{t-1}) , $i = 1, \dots, k$, i.e. one replaces the actual value of y_t with k uniform variates. Thus, there are k reference points per datapoint: $m = kn$.

We note (u_t, u_{t-1}) a generic point (either true data, or reference data). The log-odds for the semi-parametric logistic regression are then, injecting (4.1) into (3.4):

$$\begin{aligned} m_{\theta}(u_t) &= \theta_1 u_t - \frac{1}{2} \theta_2 (u_t - u_{t-1})^2 + \chi(u_{t-1}) + \log(n/m) - \log(1/2), \\ &= \theta_1 u_t + \theta_2 d_t + \chi(u_{t-1}) + \text{cst.} \end{aligned} \quad (4.2)$$

where $d_t = \frac{1}{2}(u_t - u_{t-1})^2$. From a practical perspective, the logistic regression can be performed with u_t and d_t entering as linear effects, and $\chi(u_{t-1})$ as a smooth, nonlinear effect. The constant term may be added as an offset for completeness. (It makes no practical difference since it can be absorbed into $\chi(u_{t-1})$ or the intercept. One needs to include it only if intercepts are penalised.)

The completely parametric variant of (4.2) corresponds to having a different intercept for every value of u_{t-1} . Alternatively, neglecting the normalisation

constants means replacing $\chi(u_{t-1})$ with an intercept term (or, put another way, forcing $\chi(u_{t-1})$ to be constant). Semiparametric inference can be performed using R package `mgcv` (Wood, 2006).

To measure the efficiency of the various estimation methods, we simulated realisations of the chain at a fixed parameter setting of $\theta_1 = -2, \theta_2 = 50$ for increasing n . We also used two different values of k (the ratio of reference points to real data), $k = 10$ and $k = 30$. On each simulation we picked two parameter values at random: $\theta_1 \sim \mathcal{U}(-1, 1), \theta_2 \sim \mathcal{U}(\frac{1}{10}, 10)$, generated n datapoints, and obtained the 4 different estimates. We used 300 repetitions for each value of n and k . Results are shown on Fig. 4.2.

Semiparametric inference performs almost as well as ML. Fully parametric inference is much more variable, although it becomes better for larger values of k . Indeed, theory predicts that it for large enough k it becomes equivalent to ML. The variant of non-contrastive divergence which neglects the normalisation constants performs quite well for θ_2 but shows asymptotic bias in θ_1 . The bias comes from the missing non-linear effect $\chi(u_{t-1})$, which is projected on the linear effect for u_t . This happens because the two are correlated through the dependencies in the chain. Neglecting the normalisation constants then effectively leads to confounding.

Contrary to the ideal Poisson transform (see correlary 6), the non-contrastive divergence approximation is noisy and it is possible to overfit the nonparametric term $\chi(u_{t-1})$. Cross-validation is a valid way of selecting the penalisation level, and here in practice related criteria such as Generalised Cross-Validation and REML work just as well. The results in Fig. 4.2 are obtained using the default criterion (GCV).

4.2 Spatial Markov chains for eye movement data

A perennial problem in spatial statistics is to predict where certain events are likely to take place (for example, cases of malaria in a country) given past occurrences and a set of spatial predictors (for example, availability of mosquito nets). Point process models can be used in such contexts, and one important class of applications is to eye movement data (Barthelmé et al., 2013), where the goal is to predict which locations people will look at in a given visual stimulus (for example a photograph). Eye movements are reliably drawn to certain features in a stimulus, but also exhibit dependencies (Engbert et al., 2014), and the most important of these is that we tend not to move our eyes very much. If we are currently fixating on the bottom-left corner of the screen, it will take a few steps for us to go look in the upper right, even if there is something rather interesting there.

The presence of dependencies motivates the introduction of models of eye movements as *spatial Markov chains*. Here we note \mathbf{y}_t the fixation location at time t , and use a log-linear form for the kernel:

$$p(\mathbf{y}_t | \mathbf{y}_{t-1}) \propto \exp \{s(\mathbf{y}_t) + r(\mathbf{y}_t, \mathbf{y}_{t-1})\} \quad (4.3)$$

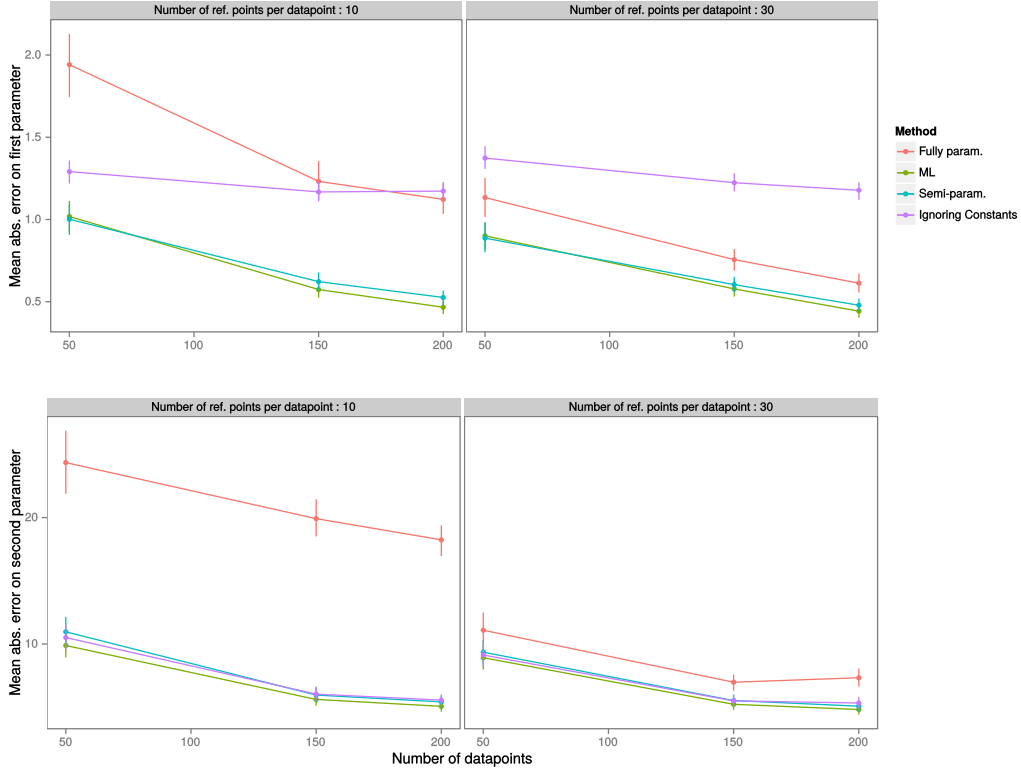


Fig. 4.2: Estimation errors of ML vs three variants of NCD for the one-dimensional Markov chain. The variants are: fully parametric (one ν_i term per datapoint), semi-parametric (normalisation constants are modelled as a smooth function), ignoring constants (logistic regression with a single intercept, as in Mnih and Teh, 2012a). The semi-parametric estimate is almost as good as the ML estimate across the board. The fully parametric estimate performs very poorly when there are few reference points per datapoint (compare the red line across the left and right panels). Finally, neglecting normalisation constants leads to a non-convergent estimator of θ_1 , although performance on θ_2 is very good. See the companion document for a more thorough discussion of this phenomenon.




Fig. 4.3: A sequence of eye movements extracted from the dataset of Kienzle et al. (2009). Fixation locations are in red and successive locations are linked by a straight line.

where $s(\mathbf{y}_t)$ represents purely spatial factors, and $r(\mathbf{y}_t, \mathbf{y}_{t-1})$ is an interaction term that represents spatial dependencies. A well-known factor affecting fixation locations is the centrality bias (Tatler and Vincent, 2009), a preference for looking at central locations, and we take $s(\mathbf{y}_t)$ to be a smooth function of $\|\mathbf{y}_t\|$ (the distance to the center): $s(\mathbf{y}_t) = s(\|\mathbf{y}_t\|)$. Potential interactions between successive locations include a tendency not to stray too far from the current location (Engbert et al., 2014), and a tendency for making movements along the cardinal axes (vertical and horizontal, Foulsham et al., 2008). We therefore further decompose $r(\mathbf{y}_t, \mathbf{y}_{t-1})$ into

$$r(\mathbf{y}_t, \mathbf{y}_{t-1}) = r_{\text{dist}}(\|\mathbf{y}_t - \mathbf{y}_{t-1}\|) + r_{\text{ang}}(\angle(\mathbf{y}_t - \mathbf{y}_{t-1})) \quad (4.4)$$

the sum of a distance and an angular component. We model the unknown functions s , r_{dist} and r_{ang} non-parametrically, using smoothing splines. The corresponding estimators are therefore obtained by penalised likelihood maximisation, and the Poisson transform extends straightforwardly to this case: replace the maximisation of $\mathcal{L}(\boldsymbol{\theta}) - \text{pen}(\boldsymbol{\theta})$ by the maximisation of $\mathcal{M}(\boldsymbol{\theta}, \chi) - \text{pen}(\boldsymbol{\theta})$, where $\boldsymbol{\theta} = (s, r_{\text{dist}}, r_{\text{ang}})$, and χ is a non-parametric function used to estimate the normalising constant, as explained in the previous section.

We use the data of Kienzle et al. (2009), who recorded eye movements while subjects were exploring a set of photographs (Fig. 4.3). There are 14 subjects, each contributing between 600 and 2,000 datapoints. Thanks to the techniques described above, the model described by (4.3) can be turned into a logistic regression, and the R package *mgcv* (Wood, 2006) can be used to estimate the different components using smoothing splines. We used a uniform, IID reference kernel $q(\mathbf{y}_t | \mathbf{y}_{t-1}) = |\Omega|^{-1}$ to produce negative examples, with 20 times as many negative examples as positive. Although the logistic approximation introduces Monte Carlo variance, the estimates are very stable (see Appendix). We fit separate functions for each subject to account for interindividual variability. The results are shown on Fig. 4.4. We replicate known effects from the literature: central locations dominate (although some subjects may display an




Fig. 4.4: Eye movement model. The smooth terms in eq. 4.3 and 4.4 are estimated using smoothing splines by reducing the model to a non-parametric logistic regression. The different panels display the estimated effects of saccade angle (r_{ang}), distance to previous fixation (r_{dist}) and centrality bias (s). Individual subjects are in gray, and the group average is in blue.

off-center bias), and dependencies include both an inhibitory effect of distance and a preference for movements along cardinal orientations.

Once the data have been put into a suitable format, model fitting can be performed in one line of R code (see Appendix) and takes around 5 minutes on a normal desktop. The Poisson transform thus turns an otherwise highly non-standard model into a convenient Generalised Additive Model.

5 Discussion

The Poisson transform suggests a new way of thinking about inference in unnormalised models: if we think of the data as coming from a point process, the integration constant becomes just another parameter to estimate. We have shown that the same idea extends to unnormalised models in the sequential context and in the presence of covariates, in which case parametric estimation may be turned into a semi-parametric problem. Practical approximations of Poisson-transformed likelihoods can be computed using Monte Carlo or using logistic likelihoods that follow from a reinterpretation of noise-contrastive divergence.

Part of the challenge in applying the Poisson transform to models with high-dimensional covariates or dependencies on a high-dimensional vector of past values will be in the design of appropriate kernels for the non-parametric part, which corresponds to conditional normalisation constants. The great advantage of the reduction to logistic regression is that we will be able to leverage the existing literature on nonlinear classification and dimensionality reduction, including recent developments in hashing (Li and König, 2011). Inference in unnormalised models will probably always remain challenging, but we believe the Poisson transform should alleviate some of the difficulties.

A Derivatives of Poisson-transformed likelihoods

The first and second derivatives of $\mathcal{L}(\boldsymbol{\theta})$ and $\mathcal{M}(\boldsymbol{\theta}, \nu)$ are needed in the proofs and we collect them here.

Derivatives of $\mathcal{L}(\boldsymbol{\theta})$:

$$\begin{aligned}
\mathcal{L}(\boldsymbol{\theta}) &= \sum_{i=1}^n f_{\boldsymbol{\theta}}(\mathbf{y}_i) - n \log \left(\int \exp \{f_{\boldsymbol{\theta}}(s)\} ds \right) := \sum_{i=1}^n f_{\boldsymbol{\theta}}(\mathbf{y}_i) - n\phi(\boldsymbol{\theta}) \\
\frac{\partial}{\partial \boldsymbol{\theta}} \phi(\boldsymbol{\theta}) &= \int \frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}}(s) \exp \{f_{\boldsymbol{\theta}}(s) - \phi(\boldsymbol{\theta})\} ds = E_{\boldsymbol{\theta}} \left(\frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}} \right) \\
\frac{\partial^2}{\partial \boldsymbol{\theta}^2} \phi(\boldsymbol{\theta}) &= E_{\boldsymbol{\theta}} \left(\frac{\partial^2}{\partial \boldsymbol{\theta}^2} f_{\boldsymbol{\theta}} \right) + E_{\boldsymbol{\theta}} \left(\left(\frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}} \right) \left(\frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}} \right)^t \right) - E_{\boldsymbol{\theta}} \left(\frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}} \right) E_{\boldsymbol{\theta}} \left(\frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}} \right)^t \\
\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{L} &= \sum_{i=1}^n \frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}}(\mathbf{y}_i) - n \frac{d}{d\boldsymbol{\theta}} \phi(\boldsymbol{\theta}) \\
\frac{\partial^2}{\partial \boldsymbol{\theta}^2} \mathcal{L}(\boldsymbol{\theta}) &= \sum \frac{\partial^2}{\partial \boldsymbol{\theta}^2} f_{\boldsymbol{\theta}}(\mathbf{y}_i) - n \frac{d^2}{d\boldsymbol{\theta}^2} \phi(\boldsymbol{\theta})
\end{aligned}$$

where we have used $E_{\boldsymbol{\theta}}$ as shorthand for the expectation with respect to density $\exp \{f_{\boldsymbol{\theta}}(s) - \phi(\boldsymbol{\theta})\}$.

Derivatives of $\mathcal{M}(\boldsymbol{\theta}, \nu)$:

$$\begin{aligned}
\mathcal{M}(\boldsymbol{\theta}, \nu) &= \sum_{i=1}^n \{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\} - n \int \exp \{f_{\boldsymbol{\theta}}(s) + \nu\} ds \\
\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{M}(\boldsymbol{\theta}, \nu) &= \sum_{i=1}^n \frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}}(\mathbf{y}_i) - n E_{\boldsymbol{\theta}, \nu} \left(\frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}} \right) \\
\frac{\partial}{\partial \nu} \mathcal{M}(\boldsymbol{\theta}, \nu) &= n - n \int \exp \{f_{\boldsymbol{\theta}}(s) + \nu\} ds \\
\frac{\partial^2}{\partial \boldsymbol{\theta}^2} \mathcal{M}(\boldsymbol{\theta}, \nu) &= \sum \frac{\partial^2}{\partial \boldsymbol{\theta}^2} f_{\boldsymbol{\theta}}(\mathbf{y}_i) - n \left(E_{\boldsymbol{\theta}, \nu} \left(\frac{\partial^2}{\partial \boldsymbol{\theta}^2} f_{\boldsymbol{\theta}} \right) + E_{\boldsymbol{\theta}, \nu} \left(\left(\frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}} \right) \left(\frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}} \right)^t \right) \right) \\
\frac{\partial^2}{\partial \nu^2} \mathcal{M}(\boldsymbol{\theta}, \nu) &= -n \int \exp \{f_{\boldsymbol{\theta}}(s) + \nu\} ds \\
\frac{\partial}{\partial \boldsymbol{\theta} \partial \nu} \mathcal{M}(\boldsymbol{\theta}, \nu) &= -n E_{\boldsymbol{\theta}, \nu} \left(\frac{\partial}{\partial \boldsymbol{\theta}} f_{\boldsymbol{\theta}} \right)
\end{aligned}$$

where we have used $E_{\boldsymbol{\theta}, \nu}$ as shorthand for the linear operator $E_{\boldsymbol{\theta}, \nu}(\varphi) = \int \varphi(s) \exp \{f_{\boldsymbol{\theta}}(s) + \nu\} ds$ (which is not an expectation in general).

B Further properties of the Poisson transform

B.1 The Poisson transform preserves confidence intervals

The usual method for obtaining confidence intervals for θ is to invert the Hessian matrix of $\mathcal{L}(\theta)$ at the mode, θ^* :

$$\mathbf{C}_{\mathcal{L}} = \left(-\frac{d^2}{d^2\theta} \mathcal{L} |_{\theta=\theta^*} \right)^{-1}$$

We can show that the same confidence intervals can be obtained from $\mathcal{M}(\theta, \nu)$ at the joint mode, θ^*, ν^* .

At the joint maximum, ν^* normalises the intensity function, and the Hessian of \mathcal{M} equals:

$$\begin{aligned} H &= \begin{bmatrix} \mathbf{H}_{aa} & \mathbf{H}_{ba} \\ \mathbf{H}_{ab} & \mathbf{H}_{bb} \end{bmatrix} = \begin{bmatrix} \frac{\partial^2}{\partial^2\theta} \mathcal{M}(\theta, \nu) & \frac{\partial}{\partial\theta\partial\nu} \mathcal{M}(\theta, \nu) \\ \frac{\partial}{\partial\nu\partial\theta} \mathcal{M}(\theta, \nu) & \frac{\partial^2}{\partial^2\nu} \mathcal{M}(\theta, \nu) \end{bmatrix} \\ &= \begin{bmatrix} \sum \frac{\partial^2}{\partial\theta^2} f_{\theta}(\mathbf{y}_i) - nE_{\theta} \left(\frac{\partial^2}{\partial^2\theta} f \right) - nE_{\theta} \left(\left(\frac{\partial}{\partial\theta} f_{\theta} \right) \left(\frac{\partial}{\partial\theta} f_{\theta} \right)^t \right) & -nE_{\theta} \left(\frac{\partial}{\partial\theta} f_{\theta} \right)^t \\ -nE_{\theta} \left(\frac{\partial}{\partial\theta} f_{\theta} \right) & -n \end{bmatrix} \end{aligned}$$

where again E denotes the expectation with respect to density $\exp \{f_{\theta}(s) - \phi(\theta)\}$.

Inverting $-H$ also yields confidence intervals. By the inversion rule for block matrices, the approximate covariance for θ using $\mathcal{M}(\theta, \nu)$ equals

$$\begin{aligned} \mathbf{C}_{\mathcal{M}}^{-1} &= -(\mathbf{H}_{aa} - \mathbf{H}_{ba}\mathbf{H}_{bb}^{-1}\mathbf{H}_{ab}) \\ &= -\left(\mathbf{H}_{aa} + \frac{1}{n}n^2E_{\theta} \left(\frac{\partial}{\partial\theta} f_{\theta} \right) E_{\theta} \left(\frac{\partial}{\partial\theta} f_{\theta} \right)^t \right) \\ &= -\left[\sum \frac{\partial^2}{\partial\theta^2} f_{\theta}(\mathbf{y}_i) - nE_{\theta} \left(\frac{\partial^2}{\partial^2\theta} f \right) \right. \\ &\quad \left. - nE_{\theta} \left(\left(\frac{\partial}{\partial\theta} f_{\theta} \right) \left(\frac{\partial}{\partial\theta} f_{\theta} \right)^t \right) + nE_{\theta} \left(\frac{\partial}{\partial\theta} f_{\theta} \right) E \left(\frac{\partial}{\partial\theta} f_{\theta} \right)^t \right] \\ &= \mathbf{C}_{\mathcal{L}}^{-1} \end{aligned}$$

B.2 Preservation of log-concavity in exponential families

In exponential families, the log-likelihood is concave, which facilitates inference. The Poisson transform preserves this log-concavity.

In the natural parameterisation, exponential-family models are given by:

$$\mathcal{L}(\theta) = \exp \left\{ \sum_{i=1}^n s(\mathbf{y}_i)^t \theta - \phi(\theta) \right\}$$

with $s(\mathbf{y})$ a vector of sufficient statistics. The second derivative of $\mathcal{L}(\boldsymbol{\theta})$ simplifies to:

$$\begin{aligned} -\frac{1}{n} \frac{\partial}{\partial^2 \boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}) &= \int s(\mathbf{y}) s(\mathbf{y})^t \exp(s(\mathbf{y})^t \boldsymbol{\theta} - \phi(\boldsymbol{\theta})) \\ &= E_{\boldsymbol{\theta}} \{s(\mathbf{y}) s(\mathbf{y})^t\} \end{aligned}$$

a p.s.d. matrix, which establishes concavity.

The second derivatives of $\mathcal{M}(\boldsymbol{\theta}, \nu)$ (Section A) also simplify

$$\begin{aligned} -\frac{1}{n} \frac{\partial^2}{\partial \boldsymbol{\theta}^2} \mathcal{M}(\boldsymbol{\theta}, \nu) &= \exp\{\nu - \nu^*(\boldsymbol{\theta})\} \int s(\mathbf{y}) s(\mathbf{y})^t \exp(s(\mathbf{y})^t \boldsymbol{\theta} - \phi(\boldsymbol{\theta})) \\ -\frac{1}{n} \frac{\partial^2}{\partial \nu \partial \boldsymbol{\theta}} \mathcal{M}(\boldsymbol{\theta}, \nu) &= \exp\{\nu - \nu^*(\boldsymbol{\theta})\} \int s(\mathbf{y}) \exp(s(\mathbf{y})^t \boldsymbol{\theta} - \phi(\boldsymbol{\theta})) \\ -\frac{1}{n} \frac{\partial^2}{\partial \nu^2} \mathcal{M}(\boldsymbol{\theta}, \nu) &= \exp\{\nu - \nu^*(\boldsymbol{\theta})\} \end{aligned}$$

so that the full Hessian \mathbf{H} can be written in block-form as:

$$-\frac{1}{n} \exp\{\nu^*(\boldsymbol{\theta}) - \nu\} \mathbf{H} = \begin{bmatrix} E_{\boldsymbol{\theta}} \{s(\mathbf{y}) s(\mathbf{y})^t\} & E(s(\mathbf{y})) \\ E_{\boldsymbol{\theta}} \{s(\mathbf{y})^t\} & 1 \end{bmatrix} = \mathbf{A}$$

and \mathbf{H} is n.s.d if and only if for all \mathbf{x}, c such that $(\mathbf{x}, c) \neq \mathbf{0}$:

$$\begin{bmatrix} \mathbf{x}^t & c \end{bmatrix} \mathbf{A} \begin{bmatrix} \mathbf{x} \\ c \end{bmatrix} > 0$$

which the following establishes:

$$\begin{aligned} &\begin{bmatrix} \mathbf{x}^t & c \end{bmatrix} \begin{bmatrix} E_{\boldsymbol{\theta}} \{s(\mathbf{y}) s(\mathbf{y})^t\} & E\{s(\mathbf{y})\} \\ E_{\boldsymbol{\theta}} \{s(\mathbf{y})^t\} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ c \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{x}^t & c \end{bmatrix} \begin{bmatrix} E_{\boldsymbol{\theta}} \{s(\mathbf{y}) s(\mathbf{y})^t\} \mathbf{x} + c E_{\boldsymbol{\theta}} \{s(\mathbf{y})\} \\ E_{\boldsymbol{\theta}} \{s(\mathbf{y})^t\} \mathbf{x} + c \end{bmatrix} \\ &= E_{\boldsymbol{\theta}} \{ \mathbf{x}^t s(\mathbf{y}) s(\mathbf{y})^t \mathbf{x} \} + 2 E_{\boldsymbol{\theta}} \{ \mathbf{x}^t s(\mathbf{y}) \} c + c^2 \\ &= E_{\boldsymbol{\theta}} \left[\left(s(\mathbf{y})^t \mathbf{x} + c \right)^2 \right] > 0 \end{aligned}$$

assuming $E_{\boldsymbol{\theta}} \{s(\mathbf{y}) s(\mathbf{y})^t\}$ is p.s.d. for all $\boldsymbol{\theta}$.

B.3 Noise-constrative divergence approximates the Poisson transform (Theorem 7)

We have assumed that

$$f_{\boldsymbol{\theta}}(\mathbf{y}) - \log q(\mathbf{y}) \leq C(\boldsymbol{\theta})$$

for a certain constant $C(\boldsymbol{\theta})$ that may depend on $\boldsymbol{\theta}$, and all $\mathbf{y} \in \Omega$. We rewrite the log-odds ratio as $h(\mathbf{y}) - \log(m)$ where

$$h(\mathbf{y}) := f_{\boldsymbol{\theta}}(\mathbf{y}) + \nu - \log q(\mathbf{y}) + \log(n)$$

does not depend on m ; note $h(\mathbf{y}) \leq \bar{h} := C(\boldsymbol{\theta}) + \nu + \log(n)$. One has:

$$\begin{aligned} \mathcal{R}^m(\boldsymbol{\theta}, \nu) + \log(m/n) &= \sum_{i=1}^n \log \left[\frac{m \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\}}{n \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\} + m q(\mathbf{y}_i)} \right] \\ &\quad + \sum_{j=1}^m \log \left[\frac{m q(\mathbf{r}_j)}{n \exp \{f_{\boldsymbol{\theta}}(\mathbf{r}_j) + \nu\} + m q(\mathbf{r}_j)} \right] \end{aligned}$$

where the first term trivially converges (as $m \rightarrow +\infty$) to

$$\sum_{i=1}^n \{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu - \log q(\mathbf{y}_i)\}.$$

Regarding the second term, one has:

$$\log \left[\frac{m q(\mathbf{r}_j)}{n \exp \{f_{\boldsymbol{\theta}}(\mathbf{r}_j) + \nu\} + m q(\mathbf{r}_j)} \right] = \log \left[1 - \frac{1}{1 + m \exp \{-h(\mathbf{r}_j)\}} \right]$$

where

$$0 \leq \frac{1}{1 + m \exp \{-h(\mathbf{r}_j)\}} \leq \frac{1}{m} \exp(\bar{h}).$$

Since $|\log(1 - x) + x| \leq x^2$ for $x \in [0, 1/2]$, we have, for m large enough, that

$$\left| \log \left[\frac{m q(\mathbf{r}_j)}{n \exp \{f_{\boldsymbol{\theta}}(\mathbf{r}_j) + \nu\} + m q(\mathbf{r}_j)} \right] + \frac{1}{1 + m \exp \{-h(\mathbf{r}_j)\}} \right| \leq \frac{\exp(2\bar{h})}{m^2} \quad (\text{B.1})$$

and

$$\left| \frac{1}{1 + m \exp \{-h(\mathbf{r}_j)\}} - \frac{1}{m} \exp \{h(\mathbf{r}_j)\} \right| \leq \frac{\exp(2\bar{h})}{m^2}$$

and since, by the law of large numbers,

$$\frac{1}{m} \sum_{j=1}^m \exp \{h(\mathbf{r}_j)\} \rightarrow \mathbb{E}_q[\exp \{h(\mathbf{r}_i)\}] = n \int \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}) + \nu\} d\mathbf{y} < +\infty \quad (\text{B.2})$$

almost surely as $m \rightarrow +\infty$, one also has:

$$\sum_{j=1}^m \log \left[\frac{m q(\mathbf{y}_i)}{n \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\} + m q(\mathbf{y}_i)} \right] \rightarrow -n \int \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}) + \nu\} d\mathbf{y}$$

almost surely, since the difference between the two sums is bounded deterministically by $\exp(2\bar{h})/m$.

B.4 Uniform convergence of the noise-constrative divergence (Theorem 8)

We first prove two intermediate results.

Lemma 9. *Assuming that $|f_{\boldsymbol{\theta}}(\mathbf{y}) - \log q(\mathbf{y})| \leq C$ for all $\mathbf{y} \in \Omega$, then there exists a bounded interval I such that, for any $\boldsymbol{\theta}$, the maximum of both functions $\nu \rightarrow \mathcal{M}(\boldsymbol{\theta}, \nu)$ and $\nu \rightarrow \mathcal{R}^m(\boldsymbol{\theta}, \nu)$ is attained in I .*

Proof. Let $\boldsymbol{\theta}$ some fixed value. $\mathcal{M}(\boldsymbol{\theta}, \nu)$ is maximised at $\nu^*(\boldsymbol{\theta}) = -\log \int_{\Omega} \exp\{f_{\boldsymbol{\theta}}(\mathbf{y})\} d\mathbf{y} \in [-C, C]$, since $e^{-C}q \leq f_{\boldsymbol{\theta}} \leq e^C q$. For $\mathcal{R}^m(\boldsymbol{\theta}, \nu)$, using again $e^{-C}q \leq f_{\boldsymbol{\theta}} \leq e^C q$, one sees that $l(\nu) \leq \mathcal{R}^m(\boldsymbol{\theta}, \nu) \leq u(\nu)$, where l and u are functions of ν that diverges at $-\infty$ for both $\nu \rightarrow +\infty$ and $\nu \rightarrow -\infty$; i.e.

$$\begin{aligned} \mathcal{R}^m(\boldsymbol{\theta}, \nu) + \log(m/n) \leq u(\nu) &:= \sum_{i=1}^n \log \left[\frac{m \exp\{C + \nu\}}{n \exp(-C + \nu) + m} \right] \\ &+ \sum_{j=1}^m \log \left[\frac{m}{n \exp\{-C + \nu\} + m} \right] \end{aligned}$$

and the lower bound $l(\nu)$ has a similar expression. Thus one may construct an interval J such that the maximum of function $\nu \rightarrow \mathcal{R}^m(\boldsymbol{\theta}, \nu)$ is attained in J for all $\boldsymbol{\theta}$ (e.g. take J such that for $\nu \in J^c$, $u(\nu) \leq M_l/2$, $l(\nu) \leq M_l/2$, with $M_l = \sup_{\nu} l$). To conclude, take $I = J \cup [-C, C]$. \square

We now establish uniform convergence, but, in light of the previous result, we restrict ν to the interval I defined in Lemma 9.

Lemma 10. *Under the Assumptions that (i) Θ is bounded, that (ii) $|f_{\boldsymbol{\theta}}(\mathbf{y}) - \log q(\mathbf{y})| \leq C$ for all $\mathbf{y} \in \Omega$, that (iii) $|f_{\boldsymbol{\theta}}(\mathbf{y}) - f_{\boldsymbol{\theta}'}(\mathbf{y})| \leq \kappa(\mathbf{y}) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|$ with κ such that $\mathbb{E}_q[\kappa] < \infty$, one has, for fixed $\mathcal{S} = \{\mathbf{y}_1, \dots, \mathbf{y}_n\}$:*

$$\sup_{(\boldsymbol{\theta}, \nu) \in \Theta \times I} \left| \mathcal{R}^m(\boldsymbol{\theta}, \nu) + \log(m/n) + \sum_{i=1}^n \log q(\mathbf{y}_i) - \mathcal{M}(\boldsymbol{\theta}, \nu) \right| \rightarrow 0 \quad (\text{B.3})$$

almost surely, relative to the randomness induced by $\mathcal{R} = \{\mathbf{r}_1, \dots, \mathbf{r}_m\}$.

Proof. Recall that the absolute difference above was bounded by the sum of three terms in the previous Appendix. The first term was

$$\sum_{i=1}^n \left[\log \left[\frac{m \exp\{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\}}{n \exp\{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\} + m q(\mathbf{y}_i)} \right] - \{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu - \log q(\mathbf{y}_i)\} \right]$$

which clearly converges deterministically to 0 as $m \rightarrow +\infty$. In addition, this convergence is uniform with respect to $(\boldsymbol{\theta}, \nu) \in \Theta \times I$, since $|\log x - \log y| \leq c|x - y|$ for $x, y \geq 1/c$, and here, by Assumption (ii),

$$x := \frac{m \exp\{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\}}{n \exp\{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu\} + m q(\mathbf{y}_i)} \geq \frac{m \exp\{-C + \nu\}}{n \exp\{C + \nu\} + m} \geq \exp\{-C + \nu\}$$

and $y = \exp \{f_{\boldsymbol{\theta}}(\mathbf{y}_i) + \nu - \log q(\mathbf{y}_i)\} \geq \exp \{-C + \nu\}$, so both x and y are lower bounded since $\nu \in I$. Similarly $(x - y)$ is bounded by C'/m , where C' is some constant independent of $\boldsymbol{\theta}$.

The second term, see (B.1), was bounded by $\exp(2\bar{h})/m^2$, where \bar{h} , an upper bound of h , may now be replaced by a constant, since $h(\mathbf{y}) := f_{\boldsymbol{\theta}}(\mathbf{y}) + \nu - \log q(\mathbf{y}) + \log(n) \leq C + \nu + \log(n)$ and $\nu \in I$, again by Assumption (ii).

The third term is related to the law of large numbers (B.2) for random variable $H_{(\boldsymbol{\theta}, \eta)}(\mathbf{r}_i) := \exp \{h(\mathbf{r}_i)\}$, which depended implicitly on $(\boldsymbol{\theta}, \eta)$:

$$H_{(\boldsymbol{\theta}, \eta)}(\mathbf{r}_i) = \frac{n \exp \{f_{\boldsymbol{\theta}}(\mathbf{r}_i) + \nu\}}{q(\mathbf{r}_i)}.$$

To obtain (almost surely) uniform convergence, we use the generalised version of the Glivenko-Cantelli theorem; e.g. Theorem 19.4 p.270 in Van der Vaart (2007). From Example 19.7 of the same book, one sees that a sufficient condition in our case is that Θ is bounded (Assumption (i)), and that

$$|H_{(\boldsymbol{\theta}, \eta)}(\mathbf{r}) - H_{(\boldsymbol{\theta}', \eta')}(\mathbf{r})| \leq m(\mathbf{r}) \|\boldsymbol{\xi} - \boldsymbol{\xi}'\|$$

for $\boldsymbol{\xi} = (\boldsymbol{\theta}, \eta)$, $\boldsymbol{\xi}' = (\boldsymbol{\theta}', \eta')$, and m a function such that $\mathbb{E}_q[m] < \infty$. But

$$\begin{aligned} |H_{(\boldsymbol{\theta}, \eta)}(\mathbf{r}) - H_{(\boldsymbol{\theta}', \eta')}(\mathbf{r})| &= \frac{n \exp \{f_{\boldsymbol{\theta}}(\mathbf{r}) + \nu\}}{q(\mathbf{r})} |1 - \exp \{f_{\boldsymbol{\theta}}(\mathbf{r}) + \nu - f_{\boldsymbol{\theta}'}(\mathbf{r}) - \nu'\}| \\ &\leq n e^{C+\nu} |1 - \exp \{f_{\boldsymbol{\theta}}(\mathbf{r}) + \nu - f_{\boldsymbol{\theta}'}(\mathbf{r}) - \nu'\}| \\ &\leq C' \{\kappa(\mathbf{r}) \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| + |\nu - \nu'|\} \\ &\leq C' \{\kappa(\mathbf{r}) + 1\} \|\boldsymbol{\xi} - \boldsymbol{\xi}'\| \end{aligned}$$

by Assumption (ii), and for some constant C' independent of $\boldsymbol{\theta}$, since $|1 - e^x| \leq Kx$ for x, y in a bounded set. One may conclude, since, by Assumption (ii), $\mathbb{E}_q[\kappa] < \infty$. \square

We are now able to prove Theorem 8. Again, let $\boldsymbol{\xi} = (\boldsymbol{\theta}, \nu)$, and rewrite any function of $(\boldsymbol{\theta}, \nu)$ as a function of $\boldsymbol{\xi}$, i.e. $\mathcal{M}(\boldsymbol{\xi})$, $\mathcal{R}^m(\boldsymbol{\xi})$. By e.g. Theorem 5.7 p.45 of Van der Vaart (2007), the uniform convergence B.3 implies that that the maximiser $\hat{\boldsymbol{\xi}}^m$ of $\mathcal{R}^m(\boldsymbol{\theta}, \nu)$ converges to the maximiser $\hat{\boldsymbol{\xi}}$ of $\mathcal{M}(\boldsymbol{\theta}, \nu)$, provided that (a) the maximisation is with respect to $(\boldsymbol{\theta}, \nu) \in \Theta \times I$; and (b) that $\sup_{d(\boldsymbol{\xi}, \hat{\boldsymbol{\xi}}) \geq \epsilon} \mathcal{M}(\boldsymbol{\xi}) < \mathcal{M}(\hat{\boldsymbol{\xi}})$. However, by Lemma 9 one sees that in (a) the same estimators would be obtained by maximising instead with respect to $(\boldsymbol{\theta}, \nu) \in \Theta \times \mathbb{R}$, and (b) is a direct consequence of Assumption (iv) of the theorem, if one takes for $d(\boldsymbol{\xi}, \hat{\boldsymbol{\xi}})$ the supremum norm of $\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}$.

C Additional information on the application

In our application we fit a spatial Markov chain model using logistic regression. Since the procedure involves the generation of a random set of reference points, we incur some Monte Carlo error in the estimates. Estimating the magnitude of

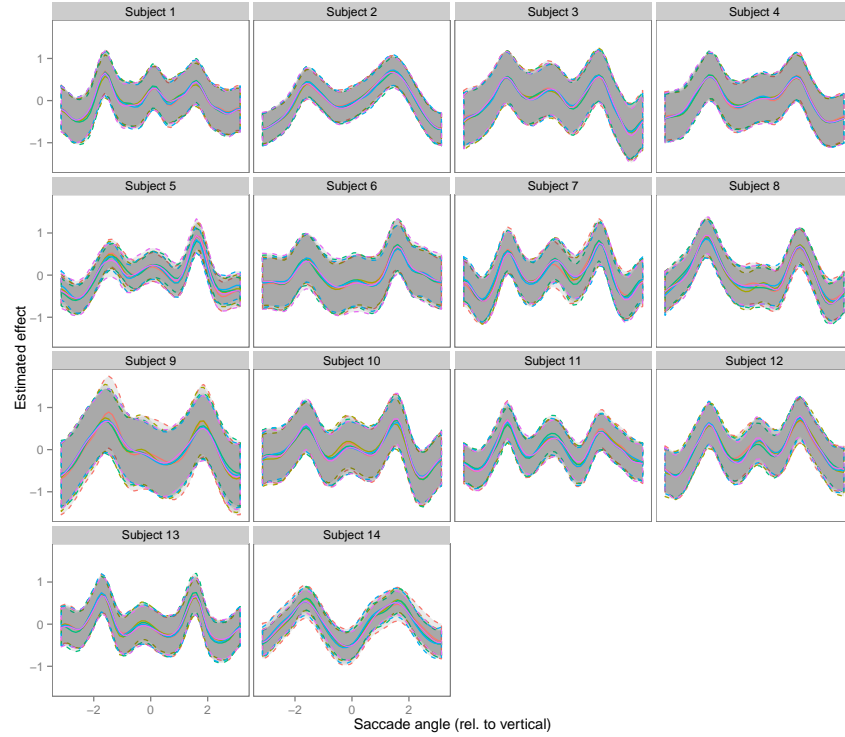


Fig. C.1: Eye movement model: 5 independent replications of the estimates under different sets of random reference points. We show here the estimated effect of saccade angle with an associated 95% pointwise confidence interval. The 5 replicates are in different colours and overlap each other almost completely, showing that 20 reference points per true datapoint are more than enough to produce stable estimates.

the Monte Carlo error is just a matter of running the procedure several times to look at variability in the estimates. We did so over 5 repetitions and report the results in Fig. C.1. For each repetition we plot the estimated smooth effect of saccade angle r_{ang} , along with a 95% confidence band. Since smoothing splines are used, smoothing hyperparameters had to be inferred from the data (using REML, Wood, 2011), and the reported confidence band is *conditional* on the estimated value of the smoothing hyperparameters. The fits and confidence bands are extremely stable over independent repetitions. The *R* command we used was:

```
gam(class ~ s(delta,k=10)+s(dcenter,k=40)+s(fxc.prev,fyc.prev,k=40)
+s(angle,bs="cc",k=20),data=data,family=binomial,method="REML")
```

References

- Baddeley, A., Berman, M., Fisher, N. I., Hardegen, A., Milne, R. K., Schuhmacher, D., Shah, R., and Turner, R. (2010). Spatial logistic regression and change-of-support in poisson point processes. *Electronic Journal of Statistics*, 4(0):1151–1201.
- Baddeley, A., Coeurjolly, J.-F., Rubak, E., and Waagepetersen, R. (2014). Logistic regression for spatial gibbs point processes. *Biometrika*, page ast060.
- Baker, S. G. (1994). The Multinomial-Poisson transformation. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 43(4):495–504.
- Barthelmé, S., Trukenbrod, H., Engbert, R., and Wichmann, F. (2013). Modeling fixation locations using spatial point processes. *Journal of vision*, 13(12).
- Bengio, Y. and Delalleau, O. (2009). Justifying and generalizing contrastive divergence. *Neural computation*, 21(6):1601–1621.
- Caimo, A. and Friel, N. (2011). Bayesian inference for exponential random graph models. *Social Networks*, 33(1):41–55.
- Engbert, R., Trukenbrod, H. A., Barthelmé, S., and Wichmann, F. A. (2014). Spatial statistics and attentional dynamics in scene viewing.
- Foulsham, T., Kingstone, A., and Underwood, G. (2008). Turning the world around: Patterns in saccade direction vary with picture orientation. *Vision Research*, 48(17):1777–1790.
- Geyer, C. J. (1994). Estimating normalizing constants and reweighting mixtures in markov chain monte carlo. Technical Report 568, School of Statistics, University of Minnesota.
- Girolami, M., Lyne, A.-M., Strathmann, H., Simpson, D., and Atchade, Y. (2013). Playing russian roulette with intractable likelihoods. *arxiv 1306.4032*.
- Gu, M. G. and Zhu, H.-T. (2001). Maximum likelihood estimation for spatial models by markov chain monte carlo stochastic approximation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2):339–355.
- Gutmann, M. and ichiro Hirayama, J. (2011). Bregman divergence as general framework to estimate unnormalized statistical models. In Cozman, F. G. and Pfeffer, A., editors, *UAI*, pages 283–290. AUAI Press.
- Gutmann, M. U. and Hyvärinen, A. (2012). Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *J. Mach. Learn. Res.*, 13(1):307–361.
- Hastie, T., Tibshirani, R., and Friedman, J. H. (2003). *The Elements of Statistical Learning*. Springer, corrected edition.

- Hinton, G. E. (2002). Training products of experts by minimizing contrastive divergence. *Neural Comput.*, 14(8):1771–1800.
- Kienzle, W., Franz, M. O., Schölkopf, B., and Wichmann, F. A. (2009). Center-surround patterns emerge as optimal predictors for human saccade targets. *Journal of vision*, 9(5).
- Kingman, J. F. C. (1993). *Poisson Processes (Oxford Studies in Probability)*. Oxford University Press.
- Li, P. and König, A. C. (2011). Theory and applications of b-bit minwise hashing. *Commun. ACM*, 54(8):101–109.
- Minka, T. (2005). Divergence Measures and Message Passing. Technical report, Microsoft Research Technical Report.
- Mnih, A. and Kavukcuoglu, K. (2013). Learning word embeddings efficiently with noise-contrastive estimation. In Burges, C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K., editors, *Advances in Neural Information Processing Systems 26*, pages 2265–2273. Curran Associates, Inc.
- Mnih, A. and Teh, Y. W. (2012a). A fast and simple algorithm for training neural probabilistic language models. In *Proceedings of the 29th International Conference on Machine Learning*, pages 1751–1758.
- Mnih, A. and Teh, Y. W. (2012b). A fast and simple algorithm for training neural probabilistic language models. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*.
- Møller, J., Pettitt, A. N., Reeves, R., and Berthelsen, K. K. (2006). An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika*, 93(2):451–458.
- Murray, I., Ghahramani, Z., and MacKay, D. (2012). MCMC for doubly-intractable distributions.
- Pihlaja, M., Gutmann, M., and Hyvärinen, A. (2010). A family of computationally efficient and simple estimators for unnormalized statistical models. In *UAI 2010, Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence, Catalina Island, CA, USA, July 8-11, 2010*, pages 442–449.
- Salakhutdinov, R. and Hinton, G. E. (2009). Deep boltzmann machines. In *International Conference on Artificial Intelligence and Statistics*, pages 448–455.
- Schölkopf, B. and Smola, A. J. (2001). *Learning with kernels : support vector machines, regularization, optimization, and beyond*. The MIT Press, 1st edition.

- Tatler, B. and Vincent, B. (2009). The prominence of behavioural biases in eye guidance. *Visual Cognition*, 17(6):1029–1054.
- Van der Vaart, A. W. (2007). *Asymptotic Statistics*. Cambridge series in statistical and probabilistic mathematics.
- Walker, S. G. (2011). Posterior sampling when the normalizing constant is unknown. *Communications in Statistics - Simulation and Computation*, 40(5):784–792.
- Wang, C., Komodakis, N., and Paragios, N. (2013). Markov random field modeling, inference & learning in computer vision & image understanding: A survey. *Computer Vision and Image Understanding*, 117(11):1610–1627.
- Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 681–688.
- Wood, S. (2006). *Generalized Additive Models: An Introduction with R (Chapman & Hall/CRC Texts in Statistical Science)*. Chapman and Hall/CRC, 1 edition.
- Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(1):3–36.