





Article

A Maximum Entropy Procedure to Solve Likelihood Equations

Antonio Calcagni * , Livio Finos , Gianmarco Altoé  and Massimiliano Pastore 

Department of Developmental and Social Psychology, University of Padova, 35131 Padova, Italy; livio.finos@unipd.it (L.F.); gianmarco.altoe@unipd.it (G.A.); massimiliano.pastore@unipd.it (M.P.)

* Correspondence: antonio.calcagni@unipd.it; Tel.: +39-049-827-6524

Received: 20 May 2019; Accepted: 14 June 2019; Published: 15 June 2019



Abstract: In this article, we provide initial findings regarding the problem of solving likelihood equations by means of a maximum entropy (ME) approach. Unlike standard procedures that require equating the score function of the maximum likelihood problem at zero, we propose an alternative strategy where the score is instead used as an external informative constraint to the maximization of the convex Shannon's entropy function. The problem involves the reparameterization of the score parameters as expected values of discrete probability distributions where probabilities need to be estimated. This leads to a simpler situation where parameters are searched in smaller (hyper) simplex space. We assessed our proposal by means of empirical case studies and a simulation study, the latter involving the most critical case of logistic regression under data separation. The results suggested that the maximum entropy reformulation of the score problem solves the likelihood equation problem. Similarly, when maximum likelihood estimation is difficult, as is the case of logistic regression under separation, the maximum entropy proposal achieved results (numerically) comparable to those obtained by the Firth's bias-corrected approach. Overall, these first findings reveal that a maximum entropy solution can be considered as an alternative technique to solve the likelihood equation.

Keywords: maximum entropy; score function; maximum likelihood; binary regression; data separation

MSC: 62F30; 62J12; 62P25

1. Introduction

Maximum likelihood is one of the most used tools of modern statistics. As a result of its attractive properties, it is useful and suited for a wide class of statistical problems, including modeling, testing, and parameters estimation [1,2]. In the case of regular and correctly-specified models, maximum likelihood provides a simple and elegant means of choosing the best asymptotically normal estimators. Generally, the maximum likelihood workflow proceeds by first defining the statistical model which is thought to generate the sample data and the associated likelihood function. Then, the likelihood is differentiated around the parameters of interest by getting the likelihood equations (score), which are solved at zero to find the final estimates. In most simple cases, the maximum likelihood solutions are expressed in closed-form. However, analytic expressions are not always available for most complex problems and researchers need to solve likelihood equations numerically. A broad class of these procedures include Newton-like algorithms, such as the Newton–Raphson, Fisher-scoring, and quasi Newton–Raphson algorithms [3]. However, when the sample size is small, or when the optimization is no longer convex as in the case of more sophisticated statistical models, the standard version of Newton–Raphson may not be optimal. In this case, robust versions should instead be used [4]. A typical example of such a situation is the logistic regression for binary data, where maximum

likelihood estimates may no longer be available, for instance, when the binary outcome variable can be perfectly or partially separated by a linear combination of the covariates [5]. As a result, the Newton–Raphson is unstable with inconsistent or infinite estimates. Other examples include small sample sizes, large numbers of covariates, and multicollinearity among the regressor variables [6]. Different proposals have been made to solve these drawbacks, many of which are based on iterative adjustments of the Newton–Raphson algorithm (e.g., see [7,8]), penalized maximum likelihood (e.g., see [9]), or the homotopy-based method (e.g., see [10]). Among them, bias-corrected methods guarantee the existence of finite maximum likelihood estimates by removing first-order bias [11], whereas homotopy Newton–Raphson algorithms, which are mostly based on Adomian’s decomposition, ensure more robust numerical convergences in finding roots of the score function (e.g., see [12]).

Maximum entropy (ME)-based methods have a long history in statistical modeling and inference (e.g., for a recent review see [13]). Since the seminal work by Golan et al. [14], there have been many applications of maximum entropy to the problem of parameter estimation in statistics, including autoregressive models [15], multinomial models [16], spatial autoregressive models [17], structural equation models [18], the co-clustering problem [19], and fuzzy linear regressions [20]. What all these works share in common is an elegant estimation method that avoids strong parametric assumptions on the model being used (e.g., error distribution). Differently, maximum entropy has also been widely adopted in many optimization problems, including queueing systems, transportation, portfolio optimization, image reconstruction, and spectral analysis (for a comprehensive review see [21,22]). In all these cases, maximum entropy is instead used as a pure mathematical solver engine for complex or ill-posed problems, such as those encountered when dealing with differential equations [23], oversampled data [24], and data decomposition [25].

The aim of this article is to introduce a maximum entropy-based technique to solve likelihood equations as they appear in many standard statistical models. The idea relies upon the use of Jaynes’ classical ME principle as a mathematical optimization tool [22,23,26]. In particular, instead of maximizing the likelihood function and solving the corresponding score, we propose a solution where the score is used as the data constraint to the estimation problem. The solution involves two steps: (i) reparametrizing the parameters as discrete probability distributions and (ii) maximizing the Shannon’s entropy function w.r.t. to the unknown probability mass points constrained by the score equation. Thus, parameter estimation is reformulated as recovering probabilities in a (hyper) simplex space, with the searching surface being always regular and convex. In this context, the score equation represents all the available information about the statistical problem and is used to identify a feasible region for estimating the model parameters. In this sense, our proposal differs from other ME-based procedures for statistical estimation (e.g., see [27]). Instead, our intent is to offer an alternative technique to solve score functions of parametric, regular, and correctly specified statistical models, where inference is still based on maximum likelihood theory.

The remainder of this article is organized as follows. Section 2 presents our proposal and describes its main characteristics by means of simple numerical examples. Section 3 describes the results of a simulation study where the ME method is assessed in the typical case of logistic regression under separation. Finally, Section 4 provides a general discussion of findings, comments, and suggestions for further investigations. Complementary materials like datasets and scripts used throughout the article are available to download at <https://github.com/antcalcagni/ME-score>, whereas the list of symbols and abbreviations adopted hereafter is available in Table 1.

Table 1. List of symbols and abbreviations used throughout the manuscript.

ME	maximum entropy
NR	Newton–Raphson algorithm
NFR	bias corrected Newton–Raphson algorithm
y	sample of observations
\mathcal{Y}	sample space
θ	$J \times 1$ vector of parameters
$\hat{\theta}$	estimated vector of parameters
$\tilde{\theta}$	reparameterized vector of parameters under ME
$f(y; \theta)$	density function
$l(\theta)$	likelihood function
$\mathcal{U}(\theta), \mathcal{U}(\tilde{\theta})$	score function
\mathbf{z}	$K \times 1$ vector of finite elements for $\tilde{\theta}$
\mathbf{p}	$K \times 1$ vector of unknown probabilities for $\tilde{\theta}$
$\hat{\mathbf{p}}$	vector of estimated probabilities for $\tilde{\theta}$

2. A Maximum Entropy Solution to Score Equations

Let $\mathbf{y} = \{y_1, \dots, y_n\}$ be a random sample of independent observations from the parametric model $\mathcal{M} = \{f(y; \theta) : \theta \in \Theta, y \in \mathcal{Y}\}$, with $f(y; \theta)$ being a density function parameterized over θ , $\Theta \subseteq \mathbb{R}^J$ the parameter space with J being the number of parameters, and \mathcal{Y} the sample space. Let

$$l(\theta) = \sum_{i=1}^n \ln f(y_i; \theta)$$

be the log-likelihood of the model and

$$\mathcal{U}(\theta) = \nabla_{\theta} l(\theta) = (\partial l / \partial \theta_1, \dots, \partial l / \partial \theta_j, \dots, \partial l / \partial \theta_J)$$

the score equation. In the regular case, the maximum likelihood estimate (MLE) $\hat{\theta}$ of the unknown vector of parameters θ is the solution of the score $\mathcal{U}(\theta) = \mathbf{0}_J$. In simple cases, $\hat{\theta}$ has closed-form expression but, more often, a numerical solution is required for $\hat{\theta}$, for instance by using iterative algorithms like Newton–Raphson and expectation-maximization.

In the maximum likelihood setting, our proposal is instead to solve $\mathcal{U}(\theta) = \mathbf{0}_J$ by means of a maximum entropy approach (for a brief introduction, see [28]). This involves a two step formulation of the problem, where θ is first reparameterized as a convex combination of a numerical support with some predefined points and probabilities. Next, a non-linear programming (NLP) problem is set with the objective of maximizing the entropy of the unknown probabilities subject to some feasible constraints. More formally, let

$$\tilde{\theta} = (\mathbf{z}_1^T \mathbf{p}_1, \dots, \mathbf{z}_j^T \mathbf{p}_j, \dots, \mathbf{z}_J^T \mathbf{p}_J)^T \quad (1)$$

be the reparameterized $J \times 1$ vector of parameters of the model \mathcal{M} , where \mathbf{z}_j is a user-defined vector of $K \times 1$ (finite) points, whereas \mathbf{p}_j is a $K \times 1$ vector unknown probabilities obeying to $\mathbf{p}_j^T \mathbf{1}_K = 1$. Note that the arrays $\mathbf{z}_1, \dots, \mathbf{z}_J$ must be chosen to cover the natural range of the model parameters. Thus, for instance, in the case of estimating the population mean $\mu \in \mathbb{R}$ for a normal model $N(\mu, \sigma^2)$ with σ^2 known, $\mathbf{z}_\mu = (-d, \dots, 0, \dots, d)^T$ with d as large as possible. In practice, as observations \mathbf{y} are available, the support vector can be defined using sample information, i.e., $\mathbf{z}_\mu = (\min(\mathbf{y}), \dots, \max(\mathbf{y}))^T$. Similarly, in the case of estimating the parameter $\pi \in [0, 1]$ of the Binomial model $\text{Bin}(\pi, n)$, the support vector is $\mathbf{z}_\pi = (0, \dots, 1)^T$. The choice of the number of points K of \mathbf{z} can be made via sensitivity analysis although it has been shown that $K \in \{5, 7, 11\}$ is usually enough for many regular problems (e.g., see [27,29]). Readers may refer to [27,30] for further details.

Under the reparameterization in Equation (1), $\mathcal{U}(\theta) = \mathbf{0}_J$ is solved via the following NLP problem:

$$\begin{aligned} & \underset{(\mathbf{p}_1, \dots, \mathbf{p}_J)}{\text{maximize}} && \mathcal{H}(\mathbf{p}_1, \dots, \mathbf{p}_J) \\ & \text{subject to:} && \mathcal{U}(\tilde{\theta}) = \mathbf{0}_J \\ & && \mathbf{p}_1^T \mathbf{1}_K = 1 \\ & && \vdots \\ & && \mathbf{p}_J^T \mathbf{1}_K = 1, \end{aligned} \tag{2}$$

where $\mathcal{H}(\mathbf{p}) = -\sum_{j=1}^J \mathbf{p}_j^T \log \mathbf{p}_j$ is the Shannon’s entropy function, whereas the score equation $\mathcal{U}(\tilde{\theta})$ has been rewritten using the reparameterized parameters $\tilde{\theta}$. The problem needs to recover $K \times J$ quantities which are defined in a (convex) hyper-simplex region with J (non-) linear equality constraints $\mathcal{U}(\tilde{\theta}_1), \dots, \mathcal{U}(\tilde{\theta}_J)$ (consistency constraints) and linear equality constraints $\mathbf{p}_1^T \mathbf{1}_K, \dots, \mathbf{p}_J^T \mathbf{1}_K$ (normalization constraints). The latter ensure that the recovered quantities $\hat{\mathbf{p}}_1, \dots, \hat{\mathbf{p}}_J$ are still probabilities. Note that closed-form solutions for the ME-score problem do not exist and solutions need to be attained numerically.

In the following examples, we will show how the ME-score problem can be formulated in the most simple cases of estimating a mean from normal, Poisson, and gamma models (Examples 1–3) as well as in more complex cases of estimating parameters for logistic regression (Example 4).

2.1. Example 1: The Normal Case

Consider the case of estimating the location parameter $\mu \in \mathbb{R}$ of a Normal density function with σ^2 known. In particular, let

$$\mathbf{y} = (2.61, 4.18, 3.40, 3.73, 3.63, 2.41, 3.76, 3.93, 4.66, 1.59, 4.51, 2.77)^T$$

be a sample of $n = 12$ drawn from a population with Normal density $N(\mu, \sigma_0^2)$ with $\sigma_0^2 = 1$ known. Our objective is to estimate μ using the information of \mathbf{y} . Let

$$l(\mu) = (\sigma_0^2)^{-1} \|\mathbf{y} - \mu \mathbf{1}_n\|^2$$

be the log-likelihood of the model where constant terms have been dropped and

$$\mathcal{U}(\mu) = (\sigma_0^2)^{-1} (\mathbf{y}^T \mathbf{1}_n - n\mu)$$

be the corresponding score w.r.t. μ . To define the associated ME-score problem to solve $\mathcal{U}(\mu) = 0$, first let $\mu_{\text{ME}} = \mathbf{z}^T \mathbf{p}$ with \mathbf{z} and \mathbf{p} being $K \times 1$ vector of supports and unknown probabilities. In this example,

$$\mathbf{z} = (1.59, 2.10, 2.61, 3.13, 3.64, 4.15, 4.66)^T$$

with $K = 7$, $z_1 = \min(\mathbf{y})$, and $z_K = \max(\mathbf{y})$. Given the optimization problem in (2), in this case \mathbf{p} can be recovered via the Lagrangean method, as follows. Let

$$\mathcal{L}(\lambda_0, \lambda_1, \mathbf{p}) = -\mathbf{p}^T \log \mathbf{p} - \lambda_0 (1 - \mathbf{p}^T \mathbf{1}_K) - \lambda_1 \left((\sigma_0^2)^{-1} (\mathbf{y}^T \mathbf{1}_n - n(\mathbf{z}^T \mathbf{p})) \right) \tag{3}$$

be the Lagrangean function, with λ_0 and λ_1 being the usual Lagrangean multipliers. The Lagrangean system of the problem is

$$\frac{\partial \mathcal{L}(\lambda_0, \lambda_1, \mathbf{p})}{\partial \mathbf{p}} = -\log(\mathbf{p}) - 1 - \lambda_0 - \lambda_1 n \mathbf{z} = \mathbf{0}_K \tag{4}$$

$$\frac{\partial \mathcal{L}(\lambda_0, \lambda_1, \mathbf{p})}{\partial \lambda_0} = 1 - \mathbf{p}^T \mathbf{1}_K = 0 \tag{5}$$

$$\frac{\partial \mathcal{L}(\lambda_0, \lambda_1, \mathbf{p})}{\partial \lambda_1} = (\sigma_0^2)^{-1} (\mathbf{y}^T \mathbf{1}_n - n(\mathbf{z}^T \mathbf{p})) = 0. \tag{6}$$

Solving \mathbf{p} in Equation (4), by using Equation (6), we get the general solutions for the ME-score problem:

$$\hat{\mathbf{p}} = \frac{\exp(-\mathbf{z} \hat{\lambda}_1 n (\sigma_0^2)^{-1})}{\exp(-\mathbf{z} \hat{\lambda}_1 n (\sigma_0^2)^{-1})^T \mathbf{1}_K} \tag{7}$$

where the quantity in the denominator is the normalization constant. Note that solutions in Equation (7) depend on the Lagrangean multiplier $\hat{\lambda}_1$, which needs to be determined numerically [31]. In this particular example, we estimate the unknown Lagrangean multiplier using a grid-search approach, yielding to $\hat{\lambda}_1 = -0.024$. The final solutions are

$$\hat{\mathbf{p}} = (0.087, 0.101, 0.117, 0.136, 0.159, 0.185, 0.215)^T$$

with $\hat{\mu}_{ME} = \mathbf{z}^T \hat{\mathbf{p}} = 3.432$, which corresponds to the maximum likelihood estimate of $\hat{\mu}_{ML} = \frac{1}{n} \mathbf{y}^T \mathbf{1}_n = 3.432$, as expected.

2.2. Example 2: The Poisson Case

Consider the simple case of estimating $\lambda \in \mathbb{R}^+$ of a Poisson density function. Let

$$\mathbf{y} = (5, 7, 7, 4, 4, 8, 15, 7, 7, 4, 7, 3, 8, 5, 4, 7)^T$$

be a sample of $n = 16$ drawn from a Poisson density $\text{Pois}(\lambda)$ and $\mathcal{U}(\lambda) = -n + (\mathbf{y}^T \mathbf{1}_n) / \lambda$ be the score of the model. The reparameterized Poisson parameter is $\lambda_{ME} = \mathbf{z}^T \mathbf{p}$, with support being defined as follows:

$$\mathbf{z} = (0.00, 3.75, 7.50, 11.25, 15.00)^T,$$

where $K = 5$ and $z_K = \max(\mathbf{y})$. Note that, since the Poisson parameter λ is bounded below by zero, we can set $z_1 = 0$. Unlike the previous case, we cannot determine $\hat{\mathbf{p}}$ analytically. For this reason, we need to solve the ME-score problem:

$$\begin{aligned} & \underset{\mathbf{p}}{\text{maximize}} && -\mathbf{p}^T \log(\mathbf{p}) \\ & \text{subject to:} && \mathbf{p}^T \mathbf{1}_K \\ & && -n + (\mathbf{y}^T \mathbf{1}_n) / (\mathbf{z}^T \mathbf{p}), \end{aligned} \tag{8}$$

via the augmented Lagrangean adaptive barrier algorithm as implemented in the function `constrOptim.nl` of the R package `alabama` [32]. The algorithm converged successfully in few iterations. The recovered probabilities are as follows:

$$\hat{\mathbf{p}} = (0.184, 0.256, 0.283, 0.247, 0.034)^T$$

with $\hat{\lambda}_{ME} = 6.375$, which is equal to the maximum likelihood solution $\hat{\lambda}_{ML} = \frac{1}{n} \mathbf{y}^T \mathbf{1}_n = 6.375$, as expected.

2.3. Example 3: The Gamma Case

Consider the following random sample

$$\mathbf{y} = (0.09, 0.35, 0.98, 0.20, 0.44, 0.13, 0.25, 0.48, 0.09, 0.45, 0.03, 0.06, 0.18, 0.26, 0.79, 0.36, 0.26)^T$$

drawn from a Gamma density $\text{Ga}(\alpha, \rho)$ with $\alpha \in \mathbb{R}^+$ being the scale parameter and $\rho \in \mathbb{R}^+$ the rate parameter. The log-likelihood of the model is as follows:

$$l(\alpha, \rho) = -((\alpha - 1) \log(\mathbf{y})^T \mathbf{1}_n - (\mathbf{y}^T \mathbf{1}_n \rho) + n\alpha \log(\rho) - n \log(\Gamma(\alpha)))$$

where $\Gamma(\cdot)$ is the well-known gamma function. The corresponding score function equals to

$$\begin{aligned} \mathcal{U}(\alpha) &= -\mathbf{y}^T \mathbf{1}_n + n\alpha\rho^{-1} \\ \mathcal{U}(\rho) &= \log(\mathbf{y})^T \mathbf{1}_n + n \log(\rho) - n\psi(\alpha), \end{aligned}$$

with $\psi(\alpha) = \frac{\partial}{\partial \alpha} \log(\Gamma(\alpha))$ being the digamma function, i.e., the derivative of the logarithm of the gamma function evaluated in α . The re-parameterized gamma parameters are defined as usual $\tilde{\alpha}_{\text{ME}} = \mathbf{z}_\alpha^T \mathbf{p}_\alpha$ and $\tilde{\rho}_{\text{ME}} = \mathbf{z}_\rho^T \mathbf{p}_\rho$ whereas the supports can be determined as $\mathbf{z}_\alpha = (0, \dots, \bar{\alpha} + \delta)$ and $\mathbf{z}_\rho = (0, \dots, \bar{\rho} + \delta)$, with δ being a positive constant. Note that the upper limits of the support can be chosen according to the following approximations: $\bar{\alpha} = 1/2M$ and $\bar{\rho} = \bar{\alpha}/\bar{y}$, with $M = \log(\bar{y}) - \sum_i \log(y_i)/n$ and $\bar{y} = \sum_i y_i/n$ [33]. In the current example, the supports for the parameters are:

$$\mathbf{z}_\alpha = (0.00, 1.12, 2.24, 3.35, 4.47)^T \quad \text{and} \quad \mathbf{z}_\rho = (0.00, 1.91, 3.82, 5.73, 7.64)^T,$$

where $K = 5$, $\bar{\alpha} = 1.47$, $\bar{\rho} = 4.64$, and $\delta = 3$. The ME-score problem for the gamma case is

$$\begin{aligned} &\underset{\mathbf{p}}{\text{maximize}} && -\mathbf{p}_\alpha^T \log(\mathbf{p}_\alpha) - \mathbf{p}_\rho^T \log(\mathbf{p}_\rho) \\ &\text{subject to:} && \mathbf{p}_\alpha^T \mathbf{1}_K \\ &&& \mathbf{p}_\rho^T \mathbf{1}_K \\ &&& -\mathbf{y}^T \mathbf{1}_n + (n\mathbf{z}_\alpha^T \mathbf{p}_\alpha)(\mathbf{z}_\rho^T \mathbf{p}_\rho)^{-1} \\ &&& \log(\mathbf{y})^T \mathbf{1}_n + n \log(\mathbf{z}_\rho^T \mathbf{p}_\rho) - n\psi(\mathbf{z}_\alpha^T \mathbf{p}_\alpha), \end{aligned} \tag{9}$$

which is solved via an augmented Lagrangean adaptive barrier algorithm. The algorithm required few iterations to converge and the recovered probabilities are as follows:

$$\hat{\mathbf{p}}_\alpha = (0.290, 0.261, 0.222, 0.164, 0.063)^T \quad \text{and} \quad \hat{\mathbf{p}}_\rho = (0.058, 0.138, 0.208, 0.270, 0.327)^T.$$

The estimated parameters under the ME-score formulation are $\hat{\alpha}_{\text{ME}} = 1.621$ and $\hat{\rho}_{\text{ME}} = 5.103$ which equal to the maximum likelihood solutions $\hat{\alpha}_{\text{ML}} = 1.621$ and $\hat{\rho}_{\text{ML}} = 5.103$.

2.4. Example 4: Logistic Regression

In what follows, we show the ME-score formulation for logistic regression. We will consider both the cases of simple situations involving no separation—where maximum likelihood estimates can be easily computed—and those unfortunate situations in which separation occur. Note that in the latter case, maximum likelihood estimates are no longer available without resorting to the use of a bias

reduction iterative procedure [7]. Formally, the logistic regression model with p continuous predictors is as follows:

$$\begin{aligned} \pi_i &= (1 + \exp(-\mathbf{X}_i\boldsymbol{\beta}))^{-1} \\ y_i &\sim \text{Bin}(\pi_i), \end{aligned} \tag{10}$$

where \mathbf{X} is an $n \times p$ matrix containing predictors, $\boldsymbol{\beta}$ is a $p \times 1$ vector of model parameters, and \mathbf{y} is an $n \times 1$ vector of observed responses. Here, the standard maximum likelihood solutions $\hat{\boldsymbol{\beta}}$ are usually attained numerically, e.g., using Newton–Raphson like algorithms [5].

No separation case. As an illustration of the ME-score problem in the optimal situation where no separation occurs, we consider the traditional Finney’s data on vasoconstriction in the skin of the digits (see Table 2) [34].

Table 2. Finney’s data on vasoconstriction in the skin of the digits. The response Y indicates the occurrence ($Y = 1$) or non-occurrence ($Y = 0$) of the vasoconstriction.

Volume	Rate	Y
3.70	0.825	1
3.50	1.090	1
1.25	2.500	1
0.75	1.500	1
0.80	3.200	1
0.70	3.500	1
0.60	0.750	0
1.10	1.700	0
0.90	0.750	0
0.90	0.450	0
0.80	0.570	0
0.55	2.750	0
0.60	3.000	0
1.40	2.330	1
0.75	3.750	1
2.30	1.640	1
3.20	1.600	1
0.85	1.415	1
1.70	1.060	0

In the Finney’s case, the goal is to predict the vasoconstriction responses as a function of volume and rate, according to the following linear term [34]:

$$\text{logit}(\pi_i) = \beta_0 + \beta_1 \log(\text{Volume}_i) + \beta_2 \log(\text{Rate}_i), \tag{11}$$

with $\text{logit} : [0, 1] \rightarrow \mathbb{R}$ being the inverse of the logistic function. In the maximum entropy framework, the model parameters can be reformulated as follows:

$$\boldsymbol{\beta}_{\text{ME}} = (\mathbf{z}^T \otimes \mathbf{I}_{p+1}) \text{vec}(\mathbf{P}^T), \tag{12}$$

where \mathbf{z} is a $K \times 1$ vector of support points, \mathbf{I}_{p+1} is an identity matrix of order $p + 1$ (including the intercept term), \mathbf{P} is a $(p + 1) \times K$ matrix of probabilities associated to the p parameters plus the intercept, \otimes is the Kronecker product, whereas $\text{vec}(\cdot)$ is a linear operator that transforms a matrix into a column vector. Note that in this example $p = 2$ and $K = 7$, whereas the support $\mathbf{z} = (-10, \dots, 0, \dots, 10)^T$ is defined to be the same for both predictors and the intercept (the bounds

of the support have been chosen to reflect the maximal variation allowed by the logistic function). Finally, the ME-score problem for the Finney's logistic regression is:

$$\begin{aligned} & \underset{\text{vec}(\mathbf{P})}{\text{maximize}} && -\text{vec}(\mathbf{P})^T \log(\text{vec}(\mathbf{P})) \\ & \text{subject to:} && \text{vec}(\mathbf{P})^T \mathbf{1}_{p(K+1)} \\ & && \mathbf{X}^T (\mathbf{y} - \boldsymbol{\pi}), \end{aligned} \quad (13)$$

where \mathbf{X} is the $n \times (p + 1)$ matrix containing the variables rate, volume, and a column of all ones for the intercept term, and $\boldsymbol{\pi} = (1 + \exp(-\mathbf{X}\boldsymbol{\beta}_{\text{ME}}))^{-1}$, with $\boldsymbol{\beta}_{\text{ME}}$ being defined as in Equation (12). Solutions for $\hat{\mathbf{P}}$ were obtained via the augmented Lagrangean adaptive barrier algorithm, which yielded the following estimates:

$$\hat{\mathbf{P}} = \begin{bmatrix} 0.000 & 0.004 & 0.062 & 0.159 & 0.220 & 0.263 & 0.293 \\ 0.000 & 0.001 & 0.099 & 0.178 & 0.224 & 0.247 & 0.251 \\ 0.205 & 0.201 & 0.190 & 0.170 & 0.137 & 0.085 & 0.013 \end{bmatrix},$$

where the third line of $\hat{\mathbf{P}}$ refers to the intercept term. The final estimated coefficients are

$$\begin{aligned} \hat{\beta}_{0_{\text{ME}}} &= -2.875 \\ \hat{\beta}_{1_{\text{ME}}} &= 5.179 \\ \hat{\beta}_{2_{\text{ME}}} &= 4.562, \end{aligned}$$

which are the same as those obtained in the original paper of Pregibon et al. [34].

Separation case. As a typical example of data under separation, we consider the classical Fisher iris dataset [35]. As generally known, the dataset contains fifty measurements of length and width (in centimeters) of sepal and petal variables for three species of iris, namely setosa, versicolor, and virginica [36]. For the sake of simplicity, we keep a subset of the whole dataset containing two species of iris (i.e., setosa and virginica) with sepal length and width variables only. Inspired by the work of Lesaffre and Albert [35], we study a model where the response variable is a binary classification of iris, with $Y = 0$ indicating the class virginica and $Y = 1$ the class setosa, whereas petal length and width are predictors of Y . The logistic regression for the iris data assumes the following linear term:

$$\text{logit}(\pi_i) = \beta_0 + \beta_1 \text{length}_i + \beta_1 \text{width}_i, \quad (14)$$

where model parameters can be reformulated as in Equation (12), with $K = 7$, $p = 2$, and \mathbf{z} being centered around zero with bounds $z_1 = -25$ and $z_K = 25$. The ME-score problem for the iris dataset is the same as in (13) and it is solved using the augmented Lagrangean adaptive barrier algorithm. The recovered $\hat{\mathbf{P}}$ is

$$\hat{\mathbf{P}} = \begin{bmatrix} 0.228 & 0.226 & 0.215 & 0.190 & 0.137 & 0.001 & 0.001 \\ 0.000 & 0.039 & 0.040 & 0.158 & 0.218 & 0.257 & 0.285 \\ 0.000 & 0.000 & 0.000 & 0.037 & 0.210 & 0.329 & 0.426 \end{bmatrix},$$

where the intercept term is reported in the third line of the matrix. The estimates for the model coefficients are reported in Table 3 (ME, first column). For the sake of comparison, Table 3 also reports the estimates obtained by solving the score of the model via bias-corrected Newton–Raphson (NRF, second column) and Newton–Raphson (NR, third column). The NRF algorithm uses the Firth's correction for the score function [7] as implemented in the R package `logistf` [37]. As expected, the NR algorithm fails to converge reporting divergent estimates. By contrast, the NRF procedure

converges to non-divergent solutions. Interestingly, the maximum entropy solutions are more close to NRF estimates although they differ in magnitude.

Table 3. Estimates for the iris logistic regression: ME (maximum entropy), NRF (biased-corrected Newton–Raphson), NR (Newton–Raphson). Note that the NRF algorithm implements the Firth’s bias correction [7].

	ME	NRF	NR
β_0	17.892	12.539	445.917
β_1	−10.091	−6.151	−166.637
β_2	12.229	6.890	140.570

3. Simulation Study

Having examined the ME-score problem with numerical examples for both simple and more complex cases, in this section, we will numerically investigate the behavior of the maximum entropy solutions for the most critical case of logistic regression under separation.

Design. Two factors were systematically varied in a complete two-factorial design:

- (i) the sample size n at three levels: 15, 20, 200;
- (ii) the number of predictors p (excluding the intercept) at three levels: 1, 5, 10.

The levels of n and p were chosen to represent the most common cases of simple, medium, and complex models, as those usually encountered in many social research studies.

Procedure. Consider the logistic regression model as represented in Equation (10) and let n_k and p_k be distinct elements of sets n and p . The following procedure was repeated $Q = 10,000$ times for each of the $n \times p = 9$ combinations of the simulation design:

1. Generate the matrix of predictors $\mathbf{X}_{n_k \times (1+p_k)} = [\mathbf{1}_{n_k} | \tilde{\mathbf{X}}_{n_k \times p_k}]$, where $\tilde{\mathbf{X}}_{n_k \times p_k}$ is drawn from the multivariate standard normal distribution $N(\mathbf{0}_{p_k}, \mathbf{I}_{p_k})$, whereas the column vector of all ones $\mathbf{1}$ stands for the intercept term;
2. Generate the vector of predictors β_{1+p_k} from the multivariate centered normal distribution $N(\mathbf{0}_{1+p_k}, \sigma \mathbf{I}_{1+p_k})$, where $\sigma = 2.5$ was chosen to cover the natural range of variability allowed by the logistic equation;
3. Compute the vector π_{n_k} via Equation (10) using $\mathbf{X}_{n_k \times (1+p_k)}$ and β_{p_k} ;
4. For $q = 1, \dots, Q$, generate the vectors of response variables $\mathbf{y}_{n_k}^{(q)}$ from the binomial distribution $\text{Bin}(\pi_{n_k})$, with π_{n_k} being fixed;
5. For $q = 1, \dots, Q$, estimate the vectors of parameters $\hat{\beta}_{1+p_k}^{(q)}$ by means of Newton–Raphson (NR), bias-corrected Newton–Raphson (NRF), and ME-score (ME) algorithms.

The entire procedure involves a total of $10,000 \times 3 \times 3 = 90,000$ new datasets as well as an equivalent number of model parameters. For the NR and NRF algorithms, we used the `glm` and `logistf` routines of the R packages `stats` [38] and `logistf` [37]. By contrast, the ME-score problem was solved via the augmented Lagrangean adaptive barrier algorithm implemented in `constrOptim.nl` routine of the R package `alabama` [32]. Convergences of the algorithms were checked using the built-in criteria of `glm`, `logistf`, and `constrOptim.nl`. For each of the generated data $\{\mathbf{y}, \mathbf{X}\}_{q=1, \dots, Q}$, the occurrence of separation was checked using a linear programming-based routine to find infinite estimates in the maximum likelihood solution [39,40]. The whole simulation procedure was performed on a (remote) HPC machine based on 16 cpu Intel Xeon CPU E5-2630L v3 1.80 GHz, 16×4 GB Ram.

Measures. The simulation results were evaluated considering the averaged bias of the parameters $\hat{B} = \frac{1}{Q}(\beta^{(k)} - \hat{\beta}^{(k)})^T \mathbf{1}$, its squared version \hat{B}^2 (the square here is element-wise), and the averaged variance of the estimates $\hat{V} = \frac{1}{Q} \text{Var}(\hat{\beta}^{(k)})$. They were then combined together to form the mean square error (MSE) of the estimates $\text{MSE} = \hat{V} + \hat{B}^2$. The relative bias $RB = (\hat{\beta}_j^{(k)} - \beta_j^0) / |\beta_j^0|$ was also computed for each predictor $j = 1, \dots, J$, (β^0 indicates the population parameter). The measures were computed for each of the three algorithms and for all the combinations of the simulation design.

Results. Table 4 reports the proportions of separation present in the data for each level of the simulation design along with the proportions of non-convergence for the three algorithms. As expected, NR failed to converge when severe separation occurred, for instance, in the case of small samples and large number of predictors. By contrast, for NRF and ME algorithms, the convergence criteria were always met. The results of the simulation study with regards to bias, variance, and mean square error (MSE) are reported in Table 5 and Figure 1. In general, MSE for the three algorithms decreased almost linearly with increasing sample sizes and number of predictors. As expected, the NR algorithm showed higher MSE than NRF and ME, except in the simplest case of $n = 200$ and $p = 1$. Unlike for the NR algorithm, with increasing model complexity ($p > 1$), ME showed a similar performances of NRF both for medium ($n = 50$) and large ($n = 200$) sample sizes. Interestingly, for the most complex scenario, involving a large sample ($n = 200$) and higher model complexity ($p = 10$), the ME algorithm outperformed NRF in terms of MSE. To further investigate the relationship between NRF and ME, we focused on the latter conditions and analyzed the behavior of ME and NRF in terms of relative bias (RB, see Figure 2). Both the ME and NRF algorithms showed RB distributions centered about 0. Except for the condition $N = 200 \wedge P = 10$, where ME showed smaller variance than NRF, both the algorithms showed similar variance in the estimates of the parameters. Finally, we also computed the ratio of over- and under-estimation r as the ratio between the number of positive RB and negative RB, getting the following results: $r_{\text{ME}} = 1.18$ (over-estimation: 54%), $r_{\text{NRF}} = 0.96$ (over-estimation: 49%) for the case $N = 200 \wedge P = 5$ and $r_{\text{ME}} = 1.12$ (over-estimation: 53%), $r_{\text{NRF}} = 0.91$ (over-estimation: 47%) for the case $N = 200 \wedge P = 10$.

Overall, the results suggest the following points:

- In the simplest cases with no separation (i.e., $N = 50 \wedge P = 1$, $N = 200 \wedge P = 1$, $N = 200 \wedge P = 5$), the ME solutions to the maximum likelihood equations were the same as those provided by standard Newton–Raphson (NR) and the bias-corrected version (NRF). In all these cases, the bias of the estimates approximated zero (see Table 5);
- In the cases of separation, ME showed comparable performances to NRF, which is known to provide the most efficient estimates in the case of logistic model under separation: Bias and MSE decreased as a function of sample size and predictors, with MSE being lower for ME than NRF in the case of $N = 200 \wedge P = 5$ and $N = 200 \wedge P = 10$;
- In the most complex scenario with a large sample and higher model complexity ($N = 200 \wedge P = 5$, $N = 200 \wedge P = 10$), ME and NRF algorithms showed similar relative bias, with ME estimates being less variable than NRF in $N = 200 \wedge P = 10$ condition. The ME algorithm tended to over-estimate the population parameters, by contrast NRF tended to under-estimate the true model parameters.

Table 4. Simulation study: proportions of separation occurred in the data and non-convergence (nc) rates for NR, NRF, ME algorithms.

n	p	Separation	nc _{NR}	nc _{NRF}	nc _{ME}
15	1	0.333	0.085	0.000	0.000
50	1	0.002	0.002	0.000	0.000
200	1	0.000	0.000	0.000	0.000
15	5	0.976	0.237	0.000	0.000
50	5	0.771	0.771	0.000	0.000
200	5	0.000	0.000	0.000	0.000
15	10	1.000	0.002	0.000	0.000
50	10	0.949	0.950	0.000	0.000
200	10	0.013	0.013	0.000	0.000

Table 5. Simulation study: averaged bias, squared averaged bias, and MSE for NR, NRF, ME algorithms.

n	p	NR				NRF				ME			
		\hat{B}	\hat{V}	\hat{B}^2	MSE	\hat{B}	\hat{V}	B^2	MSE	\hat{B}	\hat{V}	B^2	MSE
15	1	-5.54	236.70	30.67	267.36	0.22	0.35	0.05	0.40	-1.17	6.28	1.37	7.64
50	1	-0.13	3.42	0.02	3.44	-0.00	1.41	0.00	1.41	-0.12	1.99	0.01	2.00
200	1	0.03	0.11	0.00	0.11	0.00	0.10	0.00	0.10	0.03	0.11	0.00	0.11
15	5	10.68	1553.37	113.98	1667.33	-1.22	3.00	1.50	4.49	0.20	5.32	0.04	5.36
50	5	7.46	1918.18	55.65	1973.78	-0.44	2.20	0.20	2.39	-0.11	1.45	0.01	1.46
200	5	0.24	1.58	0.06	1.64	0.01	0.50	0.00	0.50	0.12	0.42	0.02	0.44
15	10	-0.97	177.40	0.95	178.35	-0.13	4.82	0.02	4.84	-0.38	8.10	0.14	8.24
50	10	2.80	1490.39	7.83	1498.20	-0.07	1.23	0.00	1.23	-0.02	1.53	0.00	1.53
200	10	0.66	15.29	0.43	15.72	0.02	0.86	0.00	0.86	0.10	0.48	0.01	0.50




Figure 1. Simulation study: averaged bias, squared averaged bias, and mean squared error (MSE) for Newton–Raphson (NR), bias-corrected Newton–Raphson (NRF), maximum entropy (ME) algorithms. Note that the number of predictors p is represented column-wise (outside) whereas the sample sizes n is reported in the x-axis (inside). The measures are plotted on logarithmic scale.

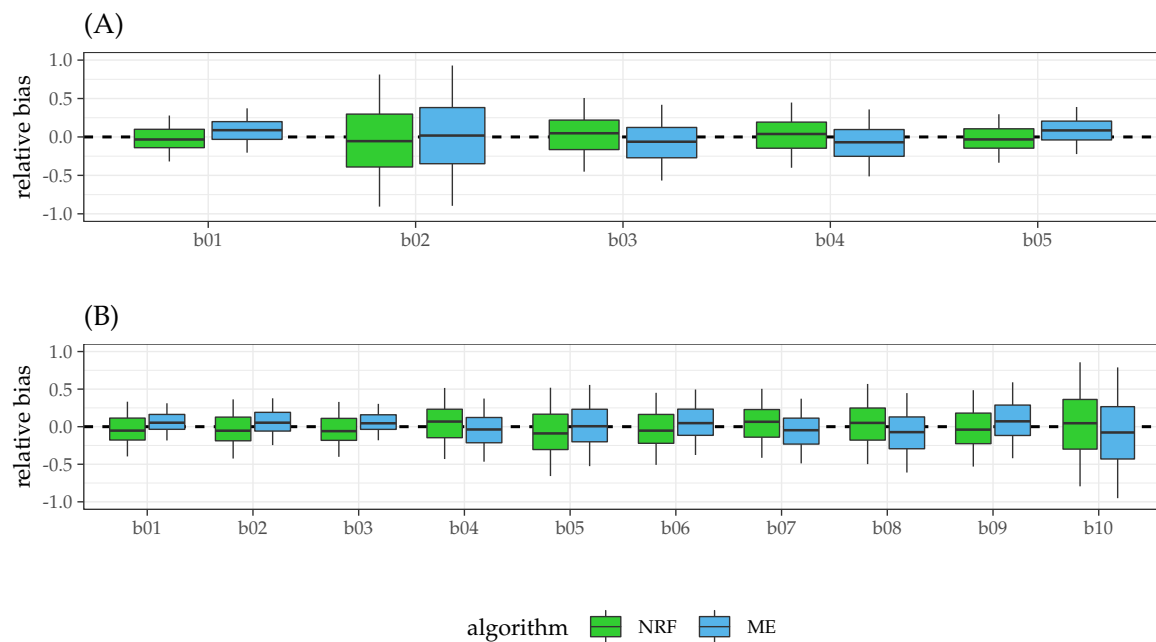


Figure 2. Simulation study: relative bias for NRF and ME algorithms in the conditions $N = 200 \wedge P = 5$ (A) and $N = 200 \wedge P = 10$ (B). Note that plots are paired vertically by predictor. Rate of over-estimation (under-estimation): (A) ME = 0.54 (0.46), NRF = 0.49 (0.51); (B) ME = 0.53 (0.47), NRF = 0.47 (0.53).

4. Discussion and Conclusion

We have described a new approach to solve the problem $\mathcal{U}(\theta) = \mathbf{0}$ in order to get $\hat{\theta}$ in the context of maximum likelihood theory. Our proposal took the advantages of using the maximum entropy principle to set a non-linear programming problem where $\mathcal{U}(\theta)$ was not solved directly, but it was used as informative constraint to maximize the Shannon's entropy. Thus, the parameter θ was not searched over the parameter space $\Theta \subset \mathbb{R}^J$, rather it was reparameterized as a convex combination of a known vector \mathbf{z} , which indicated the finite set of possible values for θ , and a vector of unknown probabilities \mathbf{p} , which instead needed to be estimated. In so doing, we converted the problem $\mathcal{U}(\theta) = \mathbf{0}$ from one of numerical mathematics to one of inference, where $\mathcal{U}(\theta)$ was treated as one of the many pieces of (external) information we may have had. As a result, the maximum entropy solution did not require either the computation of the Hessian of second-order derivatives of $l(\theta)$ (or the expectation of the Fisher information matrix) or the definition of initial values, as is required by Newton-like algorithms θ^0 . In contrast, the maximum entropy solution revolved around the reduction of the initial uncertainty: as one adds pieces of external information (constraints), a departure from the initial uniform distribution \mathbf{p} results, implying a reduction of the uncertainty about θ ; a solution is found when no further reduction can be enforced given the set of constraints. We used a set of empirical cases and a simulation study to assess the maximum entropy solution to the score problem. In cases where the Newton–Raphson is no longer correct for θ (e.g., logistic regression under separation), the ME-score formulation showed results (numerically) comparable with those obtained using the Bias-corrected Newton–Raphson, in the sense of having the same or even smaller mean square errors (MSE). Broadly speaking, these first findings suggest that the ME-score formulation can be considered as a valid alternative to solve $\mathcal{U}(\theta) = \mathbf{0}$, although further in-depth investigations need to be conducted to formally evaluate the statistical properties of the ME-score solution.

Nevertheless, we would like to say that the maximum entropy approach has been used to build a solver for maximum likelihood equations [22,23,26]. In this sense, standard errors, confidence levels, and other likelihood based quantities can be computed using the usual asymptotic properties of maximum likelihood theory. However, attention should be directed at the definition of the support points \mathbf{z} since they need to be sufficiently large to include the true (hypothesized) parameters we are

looking for. Relatedly, our proposal differs from other methods, such as generalized maximum entropy (GME) or generalized cross entropy (GCE) [20,27], in two important respects. First, the ME-score formulation does not provide a class of estimators for the parameters of statistical models. By contrast, GME and GCE are estimators belonging to the exponential family, which can be used in many cases as alternatives to maximum likelihood estimators [28]. Secondly, the ME-score formulation does not provide an inferential framework for θ . While GME and GCE use information theory to provide the basis for inference and model evaluation (e.g., using Lagrangean multipliers and normalized entropy indices), the ME-score formulation focuses on the problem of finding roots for $\mathcal{U}(\theta) = \mathbf{0}$. Finally, an open issue which deserves greater consideration in future investigations is the examination of how the ME-score solution can be considered in light of the well-known maximum entropy likelihood duality [41].

Some advantages of the ME-score solution over Newton-like algorithms may include the following: (i) model parameters are searched in a smaller and simpler space because of the convex reparameterization required for θ ; (ii) the function to be maximized does not require either the computation of second-order derivatives of $l(\theta)$, or searching for good initial values θ^0 ; (iii) additional information on the parameters, such as dominance relations among the parameters, can be added to the ME-score formulation in terms of inequality constraints (e.g., $\theta_j < \theta_t, j \neq t$). Furthermore, the ME-score formulation may be extended to include a priori probability distributions on θ . While in the current proposal, the elements of \mathbf{z}_j have the same probability to occur, the Kullback–Leibler entropy might be used to form a Kullback–Leibler-score problem, where $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_J)^T$ are adequately weighted by known vectors of probability $\mathbf{w} = (\mathbf{w}_1, \dots, \mathbf{w}_J)^T$. This would offer, for instance, another opportunity to deal with cases involving penalized likelihood estimations.

In conclusion, we think that this work yielded initial findings in the solution of likelihood equations from a maximum entropy perspective. To our knowledge, this is the first time that maximum entropy is used to define a solver to the score function. We believe this contribution will be of interest to all researchers working at the intersection of information theory, data mining, and applied statistics.

Author Contributions: conceptualization, A.C. and L.F.; methodology, A.C.; software, A.C., M.P.; formal analysis, A.C. and M.P.; visualization, M.P.; writing—original draft preparation, A.C.; writing—review and editing, A.C., L.F., G.A., M.P.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Cox, D.R. *Principles of Statistical Inference*; Cambridge University Press: Cambridge, UK, 2006.
2. Stigler, S.M. The epic story of maximum likelihood. *Stat. Sci.* **2007**, *22*, 598–620.
3. Tanner, M.A. *Tools for Statistical Inference*; Springer: Berlin, Germany, 2012.
4. Commenges, D.; Jacqmin-Gadda, H.; Proust, C.; Guedj, J. A newton-like algorithm for likelihood maximization: The robust-variance scoring algorithm. *arXiv* **2006**, arXiv:math/0610402.
5. Albert, A.; Anderson, J.A. On the existence of maximum likelihood estimates in logistic regression models. *Biometrika* **1984**, *71*, 1–10.
6. Shen, J.; Gao, S. A solution to separation and multicollinearity in multiple logistic regression. *J. Data Sci.* **2008**, *6*, 515.
7. Firth, D. Bias reduction of maximum likelihood estimates. *Biometrika* **1993**, *80*, 27–38.
8. Kenne Pagui, E.; Salvan, A.; Sartori, N. Median bias reduction of maximum likelihood estimates. *Biometrika* **2017**, *104*, 923–938.
9. Gao, S.; Shen, J. Asymptotic properties of a double penalized maximum likelihood estimator in logistic regression. *Stat. Probabil. Lett.* **2007**, *77*, 925–930.
10. Abbasbandy, S.; Tan, Y.; Liao, S. Newton-homotopy analysis method for nonlinear equations. *Appl. Math. Comput.* **2007**, *188*, 1794–1800.

11. Cordeiro, G.M.; McCullagh, P. Bias correction in generalized linear models. *J. R. Stat. Soci. Ser. B Methodol.* **1991**, *53*, 629–643.
12. Wu, T.M. A study of convergence on the Newton-homotopy continuation method. *Appl. Math. Comput.* **2005**, *168*, 1169–1174.
13. Golan, A. *Foundations of Info-Metrics: Modeling and Inference with Imperfect Information*; Oxford University Press: Oxford, UK, 2017.
14. Golan, A.; Judge, G.; Robinson, S. Recovering information from incomplete or partial multisectoral economic data. *Rev. Econ. Stat.* **1994**, *76*, 541–549.
15. Golan, A.; Judge, G.; Karp, L. A maximum entropy approach to estimation and inference in dynamic models or counting fish in the sea using maximum entropy. *J. Econ. Dyn. Control* **1996**, *20*, 559–582.
16. Golan, A.; Judge, G.; Perloff, J.M. A maximum entropy approach to recovering information from multinomial response data. *J. Am. Stat. Assoc.* **1996**, *91*, 841–853.
17. Marsh, T.L.; Mittelhammer, R.C. Generalized maximum entropy estimation of a first order spatial autoregressive model. In *Spatial and Spatiotemporal Econometrics*; Emerald Group Publishing Limited: Bingley, UK, 2004; pp. 199–234.
18. Ciavolino, E.; Al-Nasser, A.D. Comparing generalised maximum entropy and partial least squares methods for structural equation models. *J. Nonparametric Stat.* **2009**, *21*, 1017–1036.
19. Banerjee, A.; Dhillon, I.; Ghosh, J.; Merugu, S.; Modha, D.S. A generalized maximum entropy approach to bregman co-clustering and matrix approximation. *J. Mach. Learn. Res.* **2007**, *8*, 1919–1986.
20. Ciavolino, E.; Calcagni, A. A Generalized Maximum Entropy (GME) estimation approach to fuzzy regression model. *Appl. Soft Comput.* **2016**, *38*, 51–63.
21. Kapur, J.N. *Maximum-Entropy Models in Science and Engineering*; John Wiley & Sons: Hoboken, NJ, USA, 1989.
22. Fang, S.C.; Rajasekera, J.R.; Tsao, H.S.J. *Entropy Optimization and Mathematical Programming*; Springer Science & Business Media: Berlin, Germany, 2012; Volume 8.
23. El-Wakil, S.; Elhanbaly, A.; Abdou, M. Maximum entropy method for solving the collisional Vlasov equation. *Phys. A Stat. Mech. Appl.* **2003**, *323*, 213–228.
24. Bryan, R. Maximum entropy analysis of oversampled data problems. *Eur. Biophys. J.* **1990**, *18*, 165–174.
25. Calcagni, A.; Lombardi, L.; Sulpizio, S. Analyzing spatial data from mouse tracker methodology: An entropic approach. *Behav. Res. Methods* **2017**, *49*, 2012–2030.
26. Sukumar, N. Construction of polygonal interpolants: a maximum entropy approach. *Int. J. Numer. Methods Eng.* **2004**, *61*, 2159–2181.
27. Golan, A.; Judge, G.G.; Miller, D. *Maximum Entropy Econometrics: Robust Estimation with Limited Data*; Wiley: New York, NY, USA, 1996.
28. Golan, A.; Judge, G.; Miller, D. The maximum entropy approach to estimation and inference. In *Applying Maximum Entropy to Econometric Problems*; Emerald Group Publishing Limited: Bingley, UK, 1997; pp. 3–24.
29. Papalia, R.B. A composite generalized cross-entropy formulation in small samples estimation. *Econom. Rev.* **2008**, *27*, 596–609.
30. Ciavolino, E.; Calcagni, A. A generalized maximum entropy (GME) approach for crisp-input/fuzzy-output regression model. *Qual. Quant.* **2014**, *48*, 3401–3414.
31. Golan, A. Maximum entropy, likelihood and uncertainty: A comparison. In *Maximum Entropy and Bayesian Methods*; Springer: Berlin, Germany, 1998; pp. 35–56.
32. Varadhan, R. *Alabama: Constrained Nonlinear Optimization*; R package version 2015.3-1; R Core Team: Vienna, Austria, 2015.
33. Choi, S.C.; Wette, R. Maximum likelihood estimation of the parameters of the gamma distribution and their bias. *Technometrics* **1969**, *11*, 683–690.
34. Pregibon, D. Logistic regression diagnostics. *Ann. Stat. J.* **1981**, *9*, 705–724.
35. Lesaffre, E.; Albert, A. Partial separation in logistic discrimination. *J. R. Stat. Soc. Ser. B* **1989**, *51*, 109–116.
36. Fisher, R.A. The use of multiple measurements in taxonomic problems. *Ann. Eugen.* **1936**, *7*, 179–188.
37. Heinze, G.; Ploner, M. *Logistf: Firth's Bias-Reduced Logistic Regression*; R package version 1.23; R Core Team: Vienna, Austria, 2018. .
38. R Core Team. *R: A Language and Environment for Statistical Computing*; R Foundation for Statistical Computing: Vienna, Austria, 2019.

39. Konis, K. Linear Programming Algorithms for Detecting Separated Data in Binary Logistic Regression Models. Ph.D. Thesis, Department of Statistics, University of Oxford, Oxford, UK, 2007. Available online: <https://ora.ox.ac.uk/objects/uuid:8f9ee0d0-d78e-4101-9ab4-f9cbceed2a2a> (accessed on 14 June 2019).
40. Konis, K. *SafeBinaryRegression: Safe Binary Regression*; R package version 0.1-3; R Core Team: Vienna, Austria, 2013.
41. Brown, L.D. *Fundamentals of Statistical Exponential Families: With Applications in Statistical Decision Theory*; Lecture Notes-Monograph Series; Institute of Mathematical Statistics: Hayworth, CA, USA, 1986; Volume 9.



© 2019 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).