



Contents lists available at ScienceDirect

Econometrics and Statistics

journal homepage: www.elsevier.com/locate/ecosta

Analyzing cellwise weighted data

Peter J. Rousseeuw¹

Section of Statistics and Data Science, University of Leuven, Belgium

ARTICLE INFO

Article history:

Received 27 September 2022

Revised 31 January 2023

Accepted 31 January 2023

Available online xxx

Keywords:

Cellwise outliers

Covariance matrix

EM algorithm

Likelihood

Missing values

ABSTRACT

Often the rows (cases, objects) of a dataset have weights. For instance, the weight of a case may reflect the number of times it has been observed, or its reliability. For analyzing such data many rowwise weighted techniques are available, the most well known being the weighted average. But there are also situations where the individual *cells* (entries) of the data matrix have weights assigned to them. An approach to analyze such data is proposed. A cellwise weighted likelihood function is defined, that corresponds to a transformation of the dataset which is called unpacking. Using this weighted likelihood one can carry out multivariate statistical methods such as maximum likelihood estimation and likelihood ratio tests. Particular attention is paid to the estimation of covariance matrices, because these are the building blocks of much of multivariate statistics. An R implementation of the cellwise maximum likelihood estimator is provided, which employs a version of the EM algorithm. Also a faster approximate method is proposed, which is asymptotically equivalent to it.

© 2023 The Author(s). Published by Elsevier B.V. on behalf of EcoSta Econometrics and Statistics.

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

1. Motivation

Often the rows (cases, objects) of a dataset have weights. For analyzing such data many rowwise weighted techniques are available. For instance, the concept of a weighted average is widely known, and has been used extensively in areas such as survey sampling. When the observations are d -variate points $\mathbf{x}_1, \dots, \mathbf{x}_n$ with nonnegative weights w_1, \dots, w_n , their weighted average is simply

$$\bar{\mathbf{x}}_w := \frac{w_1 \mathbf{x}_1 + \dots + w_n \mathbf{x}_n}{w_1 + \dots + w_n} . \quad (1)$$

The weight w_i can arise in different ways. It can be the number of times that \mathbf{x}_i has been observed ('frequency weight'). But a weight does not have to be an integer: the weight w_i can also reflect the reliability or precision of the observation \mathbf{x}_i . Expression (1) is also used outside of statistics, for instance in physics this is the center of gravity of a system with masses w_1, \dots, w_n . Note that the effect of the weights in (1) is relative, in the sense that multiplying all weights by the same constant yields the same result.

E-mail address: peter@rousseeuw.net

¹ ORCID 0000-0002-3807-5353.

Similar expressions are those of the weighted covariance matrix

$$\frac{\sum_{i=1}^n w_i (\mathbf{x}_i - \bar{\mathbf{x}}_w) (\mathbf{x}_i - \bar{\mathbf{x}}_w)^\top}{\sum_{i=1}^n w_i} \quad (2)$$

and of weighted least squares regression, given by

$$\operatorname{argmin}_{\boldsymbol{\theta}} \sum_{i=1}^n w_i r_i^2 \quad (3)$$

in which the r_i are the residuals $y_i - (\theta_0 + \theta_1 x_{i1} + \dots + \theta_p x_{ip})$ with the usual notation.

All of these can be seen as examples of maximizing a weighted likelihood. Let us denote the likelihood of an observation \mathbf{x} by $f(\boldsymbol{\theta}|\mathbf{x})$, where the parameter $\boldsymbol{\theta}$ can be a number, vector, matrix etc. It is often convenient to work with the loglikelihood

$$L(\boldsymbol{\theta}|\mathbf{x}) := \ln f(\boldsymbol{\theta}|\mathbf{x}).$$

When the data are independent and identically distributed (i.i.d.), the loglikelihood of the sample $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ is

$$L(\boldsymbol{\theta}|\mathbf{X}) = \sum_{i=1}^n L(\boldsymbol{\theta}|\mathbf{x}_i).$$

The weighted loglikelihood is given by

$$L(\boldsymbol{\theta}|\mathbf{X}, \mathbf{w}) = \sum_{i=1}^n w_i L(\boldsymbol{\theta}|\mathbf{x}_i) \quad (4)$$

where the weights are combined into the vector $\mathbf{w} = (w_1, \dots, w_n)$, and the corresponding weighted likelihood is

$$f(\boldsymbol{\theta}|\mathbf{X}, \mathbf{w}) = \prod_{i=1}^n f(\boldsymbol{\theta}|\mathbf{x}_i)^{w_i}. \quad (5)$$

There is a fairly large literature on the use of weighted likelihood. [Hu \(1994\)](#) and [Hu and Zidek \(2002\)](#) consider some data points less relevant than others, and wish to diminish their role in order to trade bias for precision. The paper by [O'Hagan et al. \(2019\)](#) focuses in particular on gaussian mixture models, see the function `me.weighted` in the R package `mclust` ([Fraleay et al., 2022](#)). [Magis \(2015\)](#) uses weighted likelihood for item response models. On the other hand, some authors have used weighted likelihood to reduce the effect of potential outliers in estimation, see e.g. [Field and Smith \(1994\)](#) for location and scale, [Agostinelli and Markatou \(1998\)](#) for linear regression, [Croux et al. \(2013\)](#) for ordinal regression, and [Majumder et al. \(2021\)](#) for additional theoretical properties. [Agostinelli and Markatou \(2001\)](#) focused on hypothesis tests in this context.

2. Cellwise weighted likelihood

The weights we have considered so far were all rowwise weights, that is, they were assigned to entire rows of the dataset. However, it is also possible that the individual cells (entries) of the data matrix have weights assigned to them. For instance, the weight of a cell could be indicative of the level of confidence in that particular measurement, or may be related to its reliability or measurement accuracy. It may also be derived from a fuzziness measure or a probability.

How can such cellwise weighted data be analyzed, that is, how can we estimate parameters, carry out tests or other inference on them, and make predictions? As in the previous section we will address this issue by likelihood, first for a single row \mathbf{x} . We assume that each cell of \mathbf{x} has a weight, combined in the weight vector

$$\mathbf{w} = (w_1, \dots, w_d)$$

where $w_j \geq 0$ for $j = 1, \dots, d$. A weight $w_j = 0$ is taken to mean that the corresponding cell x_j is missing. Extending our earlier notation, we will denote the usual observed likelihood ([Little and Rubin, 2020](#)) of a row \mathbf{x} with some missing entries by $f(\boldsymbol{\theta}|\mathbf{x})$ as well.

The question is now whether we can define a sensible likelihood in this setting. First we note that \mathbf{w} may contain ties, that is, $w_j = w_{j'}$ for $j \neq j'$. Let us sort the *unique* nonzero weights as

$$w^{(1)} > w^{(2)} > \dots > w^{(q)} > 0$$

with the number of levels $q \leq d$, and corresponding sets of indices

$$I^{(\ell)} = \{j: w_j = w^{(\ell)}\} \quad \text{for} \quad \ell = 1, \dots, q.$$

We then consider the cumulative index sets

$$J^{(1)} = I^{(1)}$$

$$J^{(2)} = I^{(1)} \cup J^{(2)}$$

$$\dots$$

$$J^{(q)} = I^{(1)} \cup J^{(2)} \cup \dots \cup J^{(q)}$$

so $J^{(1)} \subset J^{(2)} \subset \dots \subset J^{(q)}$. For each $\ell = 1, \dots, q$ we denote by $\mathbf{x}^{(\ell)}$ a new row with components

$$x_j^{(\ell)} = \begin{cases} x_j & \text{for } j \text{ in } J^{(\ell)} \\ \text{NA} & \text{otherwise} \end{cases} \quad (6)$$

for $j = 1, \dots, d$. We now define the weighted loglikelihood as the linear combination

$$L(\boldsymbol{\theta}|\mathbf{x}, \mathbf{w}) := \sum_{\ell=1}^q (w^{(\ell)} - w^{(\ell+1)})L(\boldsymbol{\theta}|\mathbf{x}^{(\ell)}) \quad (7)$$

with the convention $w^{(q+1)} = 0$. For the likelihood itself this becomes

$$f(\boldsymbol{\theta}|\mathbf{x}, \mathbf{w}) = \prod_{\ell=1}^q f(\boldsymbol{\theta}|\mathbf{x}^{(\ell)})(w^{(\ell)} - w^{(\ell+1)}) \quad (8)$$

These formulas look unfamiliar at first, but when one thinks about it they make perfect sense. If all $w_j = 1$ one recovers the usual likelihood, and if all w_j are 0 or 1 it becomes the observed likelihood. When the cell weights are the number of repeated measurements, the sets $J^{(\ell)}$ are intuitive. But as the main benefit we see the ability to work with the accuracy or trustworthiness of individual measurements on a continuous scale.

An i.i.d. sample with n datapoints corresponds to an $n \times d$ matrix \mathbf{X} , and the weights form an $n \times d$ matrix \mathbf{W} . The overall likelihood of the sample then becomes

$$f(\boldsymbol{\theta}|\mathbf{X}, \mathbf{W}) = \prod_{i=1}^n \prod_{\ell=1}^{q_i} f(\boldsymbol{\theta}|\mathbf{x}_i^{(\ell)})(w_i^{(\ell)} - w_i^{(\ell+1)}) \quad (9)$$

with loglikelihood

$$L(\boldsymbol{\theta}|\mathbf{X}, \mathbf{W}) = \sum_{i=1}^n \sum_{\ell=1}^{q_i} (w_i^{(\ell)} - w_i^{(\ell+1)})L(\boldsymbol{\theta}|\mathbf{x}_i^{(\ell)}) \quad (10)$$

These formulas have a practical interpretation. They are equivalent to computing the overall observed likelihood of an artificial dataset $\mathbf{X}^{(\mathbf{W})}$ with row weights, as in (4) and (5). The matrix $\mathbf{X}^{(\mathbf{W})}$ is obtained by ‘unpacking’ the matrix \mathbf{X} according to the weights in \mathbf{W} . This is done by replacing each row \mathbf{x}_i of \mathbf{X} by q_i rows $\mathbf{x}_i^{(\ell)}$ that may contain NA’s and have row weights $v_i^{(\ell)} := (w_i^{(\ell)} - w_i^{(\ell+1)}) > 0$. Rows with a row weight of 0 are left out. This new matrix $\mathbf{X}^{(\mathbf{W})}$ still has d columns but might have up to nd rows. When all $w_{ij} = 1$ we obtain $\mathbf{X}^{(\mathbf{W})} = \mathbf{X}$, and when all w_{ij} are zero or one we recover the incomplete dataset in which the cells x_{ij} with $w_{ij} = 0$ are set to missing. Note that the unpacking transform can also be used outside of the likelihood context.

As an illustration, below are the first 3 rows of a data set \mathbf{X} with four variables, together with the weights of their cells in the matrix \mathbf{W} :

$$\mathbf{X} = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 \end{matrix} \\ \begin{matrix} A \\ B \\ C \\ \dots \end{matrix} & \begin{bmatrix} 2.8 & 5.3 & 4.9 & 7.4 \\ 2.3 & 5.7 & 4.3 & 7.2 \\ 2.5 & 5.1 & 4.4 & 7.6 \\ \dots & \dots & \dots & \dots \end{bmatrix} \end{matrix} \quad \mathbf{W} = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 \end{matrix} \\ \begin{matrix} A \\ B \\ C \\ \dots \end{matrix} & \begin{bmatrix} 0.8 & 1.0 & 0.3 & 0.4 \\ 0.3 & 0.5 & 0.9 & 0.5 \\ 1.0 & 0.6 & 0.0 & 0.7 \\ \dots & \dots & \dots & \dots \end{bmatrix} \end{matrix}$$

Case A has 4 different nonzero weights, so it unpacks into 4 rows of the matrix $\mathbf{X}^{(\mathbf{W})}$ below, all labeled as A. The first of these rows has the real value 5.3 in its second position, corresponding to the only cell in \mathbf{X} with weight $w_{1j} \geq 1.0$, and NA’s elsewhere. Since the next cell will come in at weight 0.8, this first row of $\mathbf{X}^{(\mathbf{W})}$ gets the row weight $v_1^{(1)} = 1.0 - 0.8 = 0.2$ that we see in the column vector on the right hand side. The second row of $\mathbf{X}^{(\mathbf{W})}$ has real values in cells 1 and 2, which are the cells of \mathbf{X} with $w_{1j} \geq 0.4$ so the weight of this row is $v_1^{(2)} = 0.8 - 0.4 = 0.4$ on the right. The third row has three real values, and the fourth row has real values in all of its cells.

Next we unpack row B of \mathbf{X} , which is analogous except that cells 2 and 4 have the same cell weight $w_{22} = 0.5 = w_{24}$ so there are only three different weights, hence row B only yields three rows in $\mathbf{X}^{(\mathbf{W})}$. Indeed, in row 6 of $\mathbf{X}^{(\mathbf{W})}$ the entries 5.7 and 7.2 join at the same time. Finally, row C of \mathbf{X} does have four different cell weights, but the lowest of them is zero. The latter would yield a row of $\mathbf{X}^{(\mathbf{W})}$ consisting exclusively of NA’s, but such uninformative rows are not kept, so C also unpacks into only three rows of $\mathbf{X}^{(\mathbf{W})}$.

$$\mathbf{X}^{(W)} = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & \nu \end{matrix} \\ \begin{matrix} A \\ A \\ A \\ A \\ B \\ B \\ B \\ C \\ C \\ C \\ \dots \end{matrix} & \begin{bmatrix} \text{NA} & 5.3 & \text{NA} & \text{NA} & \\ 2.8 & 5.3 & \text{NA} & \text{NA} & \\ 2.8 & 5.3 & \text{NA} & 7.4 & \\ 2.8 & 5.3 & 4.9 & 7.4 & \\ \text{NA} & \text{NA} & 4.3 & \text{NA} & \\ \text{NA} & 5.7 & 4.3 & 7.2 & \\ 2.3 & 5.7 & 4.3 & 7.2 & \\ 2.5 & \text{NA} & \text{NA} & \text{NA} & \\ 2.5 & \text{NA} & \text{NA} & 7.6 & \\ 2.5 & 5.1 & \text{NA} & 7.6 & \\ \dots & \dots & \dots & \dots & \dots \end{bmatrix} \end{matrix}.$$

One of the important uses of the likelihood function is to compute the maximum likelihood estimator (MLE) of θ . In view of the matrix unpacking interpretation, this is quite feasible. All we have to do is to apply maximum likelihood to the unpacked matrix $\mathbf{X}^{(W)}$ with its row weights. We will call this estimator the *cellwise weighted maximum likelihood estimator* (cwMLE).

For inference it is useful to know the large sample behavior of the estimator. The exact MLE that minimizes the observed likelihood is asymptotically normal under regularity conditions that are similar to those for complete data, as seen in Section 6.1.3 of [Little and Rubin \(2020\)](#) with references to proofs. Therefore the MLE is also consistent. Some algorithms for the MLE, such as the Newton-Raphson algorithm, preserve its asymptotic normality. This is also true for the Fisher scoring algorithm, see e.g. [Jamshidian and Bentler \(1999\)](#), [Jorgenson and Petersen \(2012\)](#), and [Takai \(2020\)](#). The formulas for the asymptotic covariance matrix when using Newton-Raphson or Fisher scoring are given in Section 9.1 of [Little and Rubin \(2020\)](#).

The most popular algorithm for the MLE of incomplete data is the EM algorithm of [Dempster et al. \(1977\)](#). The supplemented EM (SEM) algorithm of [Meng and Rubin \(1991\)](#) also provides, as a byproduct, a numerically stable estimate of the asymptotic covariance matrix of the estimator.

In many situations the observed likelihood is hard to compute because it requires integration, which prevents a closed form. When that happens one can approximate the observed likelihood by Monte Carlo, again yielding asymptotically normal estimates, see e.g. [Sung and Geyer \(2007\)](#) and the references cited therein.

Apart from estimation, the cellwise weighted likelihood can also be used for inference, for instance by applying a likelihood ratio test using Wilks' chi-squared theorem.

3. Covariance from cellwise weighted data

We now apply the technology of the previous section to the ubiquitous multivariate model where the data \mathbf{X} are generated from a gaussian distribution with unknown parameters μ and Σ . We will denote the cellwise weighted MLE (cwMLE) estimates as $\hat{\mu}$ and $\hat{\Sigma}$ with entries $\hat{\mu}_j$ and $\hat{\Sigma}_{jk}$. An R implementation is available which applies the unpacking transform followed by a rowwise weighted implementation of the EM algorithm for location and covariance, which uses iteration.

For rowwise weights we know we can compute the weighted MLE by the explicit formulas (1) and (2) of the rowwise weighted mean and the rowwise weighted covariance. For cellwise weights no explicit formulas for the cwMLE are possible. But can we at least come up with simple explicit expressions that *approximate* the cwMLE? For estimating μ it is natural to consider a cellwise weighted mean (cwMean) $\tilde{\mu}$ given by

$$\tilde{\mu}_j := \frac{\sum_{i=1}^n w_{ij} x_{ij}}{\sum_{i=1}^n w_{ij}} \quad (11)$$

for $j = 1, \dots, d$ in which each coordinate $\tilde{\mu}_j$ uses a different set of weights.

When estimating Σ a natural expression for the entry $\tilde{\Sigma}_{jk}$ would be

$$\tilde{\Sigma}_{jk} := \frac{\sum_{i=1}^n w_{ijk} (x_{ij} - \tilde{\mu}_j)(x_{ik} - \tilde{\mu}_k)}{\sum_{i=1}^n w_{ijk}}. \quad (12)$$

(As we are approximating an MLE, there is no analog of subtracting a degree of freedom in the denominator.) The weight w_{ijk} in (12) depends on both the row number i and the variables j and k . But how should such a weight w_{ijk} be defined? If we think about the construction of the cellwise loglikelihood (10) in section 2 and apply it to the estimation of Σ_{jk} , we note that the components x_{ij} and x_{ik} are only available together in some rows of $\mathbf{X}^{(W)}$, with total weight equal to the lowest of w_{ij} and w_{ik} . Above that level at least one of them becomes NA, so in those terms of (10) row i cannot contribute to the estimation of Σ_{jk} . This reasoning suggests using

$$\tilde{w}_{ijk} := \min(w_{ij}, w_{ik}). \quad (13)$$

Table 1

Average of estimates when there are imprecise data cells.

n	estimates for the μ_j			for diagonal entries of Σ			for the off-diagonal entries of Σ			
	\bar{x}	cwMLE	cwMean	Cov	cwMLE	cwCov	Cov	cwMLE	cwCov	sqrtCov
10	0.004	0.000	0.001	2.782	0.919	0.903	0.007	0.002	0.006	0.007
20	0.005	0.004	0.004	2.802	0.973	0.970	-0.006	0.001	0.002	0.001
50	-0.004	-0.001	-0.001	2.791	0.998	0.998	-0.010	0.003	0.004	0.003
100	0.003	0.000	0.000	2.808	1.014	1.014	-0.002	-0.004	-0.004	-0.004
200	0.000	0.000	0.000	2.803	1.017	1.017	-0.001	0.002	0.002	0.002
500	0.000	0.001	0.001	2.800	1.021	1.021	0.001	0.000	0.000	0.000
1000	0.000	0.000	0.000	2.800	1.021	1.021	-0.001	0.000	0.000	0.000
2000	0.000	0.000	0.000	2.798	1.021	1.021	-0.002	0.000	0.000	0.000
5000	0.000	0.000	0.000	2.801	1.022	1.022	-0.001	0.000	0.000	0.000
10000	0.000	0.000	0.000	2.800	1.022	1.022	0.000	0.000	0.000	0.000

We will call the resulting value of (12) the *cellwise weighted covariance* (cwCov) and denote it as $\tilde{\Sigma}_{jk}$. Note that for the diagonal entries $\tilde{\Sigma}_{jj}$ the weights simply become $\tilde{w}_{ijj} = w_{ij}$. Formulas (12) and (13) are explicit in the original x_{ij} and w_{ij} (no unpacking is required) and allow for fast computation. Due to its entrywise construction the combined matrix $\tilde{\Sigma}$ need not be positive semidefinite (PSD) in general, but we will see that it gives an excellent approximation to $\hat{\Sigma}$ and becomes PSD for increasing sample size.

A different cellwise weighted covariance matrix was proposed by Van Aelst et al. (2011). It also falls in the framework of (12) but uses the weight function

$$\tilde{w}_{ij} = \sqrt{w_{ij}w_{ik}}. \quad (14)$$

We will denote the resulting entries by $\tilde{\Sigma}_{jk}$ forming the matrix $\tilde{\Sigma}$ which we call the *square root covariance matrix* (sqrtCov), which also is not necessarily PSD. Note that the weights used in the diagonal entries $\tilde{\Sigma}_{jj}$ also become $\tilde{w}_{ijj} = w_{ij}$ so the diagonals of $\tilde{\Sigma}$ and $\hat{\Sigma}$ coincide, but there is no obvious relation between their off-diagonal entries.

4. Illustration with imprecise data cells

We now illustrate the behavior of the cellwise weighted estimators of μ and Σ in the previous section. We start by generating n i.i.d. data points according to the standard bivariate normal distribution, so $d = 2$. Next we ‘jitter’ some of the cells in the following way. We randomly draw 20% of the data cells, and add independent noise to them that is normally distributed with mean zero and standard deviation 3. An equivalent way to formulate this jittering scenario is to say that the data cells x_{ij} all have a univariate normal distribution with mean zero, most of them with variance $v_{ij} = 1$ except for a random fraction of 20% of the cells that has variance $v_{ij} = 3^2 + 1 = 10$. The latter cells can be seen as less precise than the remaining 80%.

It is still possible to estimate μ by the classical mean \bar{x} , whose components remain unbiased since all cells have mean zero. But \bar{x} gives every cell the same weight, which does not reflect the differences in precision. Alternatively we could assign weights w_{ij} to the cells that are a decreasing function of the variances, for instance $w_{ij} = 1/v_{ij}^2$.

We ran a small simulation, consisting of 5000 replications for sample sizes n ranging from 10 to 10000. Apart from \bar{x} and the classical covariance matrix Cov we also computed the estimates $\hat{\mu}$ and $\hat{\Sigma}$ obtained by cwMLE, the cwMean vector $\tilde{\mu}$ and cwCov matrix $\tilde{\Sigma}$, and the sqrtCov matrix $\tilde{\Sigma}$. For $n \geq 20$ both the cwCov and sqrtCov matrices were positive definite in all 5000 replications. Table 1 reports average values of the components of all these estimates.

As expected, we see that all three estimators of the $\mu_j = 0$ tend to zero. Also the off-diagonal entries of the covariance estimators tend to zero, which is intuitive due to the symmetries in the data. When estimating the diagonal entries Σ_{jj} the situation is quite different, as the classical Cov goes to $0.80 + 0.20 * 10 = 2.8$. The estimators cwMLE and cwCov stay much closer to 1 because they downweight the imprecise cells.

Table 2 shows the variances of the estimators over the 5000 replications, multiplied by the sample size n . Here we see that the cellwise weighted estimators of μ_j have a much lower variance than the classical mean, which attaches the same weight to the precise and the imprecise cells. This effect is even more pronounced for the estimates of the off-diagonal entries Σ_{jk} where the variance of the classical covariance is inflated more relative to the cellwise weighted estimators. Note that we divided the variances of the estimates of the diagonal entries Σ_{jj} by 2, which would be the lowest achievable variance if all cells were precise. In those columns the entries for the diagonal of the unweighted Cov are much higher than those of cwMLE and cwCov due to the imprecise data cells. Overall, the cellwise weighted estimators performed the best in this mixed precision setting.

Table 2
Variance of estimates when there are imprecise data cells.

n	estimates for the μ_j			for diagonal entries of Σ			for the off-diagonal entries of Σ			
	\bar{x}	cwMLE	cwMean	Cov	cwMLE	cwCov	Cov	cwMLE	cwCov	sqrtCov
10	2.72	1.32	1.26	24.77	1.24	1.12	8.01	1.72	1.33	1.27
20	2.76	1.25	1.23	24.42	1.22	1.20	7.84	1.62	1.45	1.39
50	2.75	1.24	1.23	24.45	1.22	1.22	7.71	1.56	1.51	1.46
100	2.77	1.24	1.24	24.35	1.25	1.25	8.08	1.52	1.50	1.45
200	2.70	1.24	1.23	24.01	1.23	1.23	8.11	1.53	1.56	1.51
500	2.74	1.23	1.23	23.62	1.23	1.23	7.78	1.52	1.56	1.50
1000	2.76	1.25	1.25	24.13	1.24	1.24	8.08	1.52	1.55	1.49
2000	2.79	1.26	1.26	24.03	1.24	1.24	8.03	1.52	1.55	1.50
5000	2.83	1.22	1.22	24.21	1.27	1.27	7.97	1.50	1.54	1.47
10000	2.79	1.26	1.26	24.10	1.29	1.29	7.90	1.51	1.54	1.48

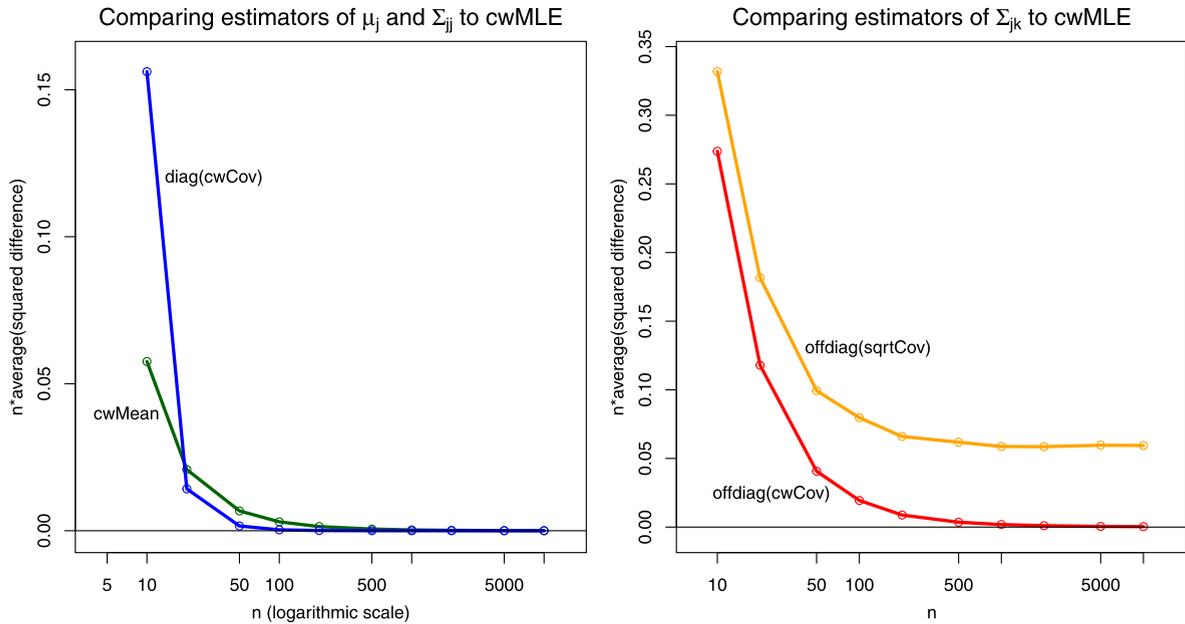


Fig. 1. Data with imprecise cells: plot of (15) for (left) the estimator cwMean and the diagonal entries of cwCov, and (right) the off-diagonal entries of cwCov (lower curve) and of sqrtCov (upper curve).

In Tables 1 and 2 we see that the entries for cwMean and the cwMLE estimator of μ_j are close to each other, especially for large n . In fact, we can see a bit more. From the simulated estimates we also computed

$$n \frac{1}{Md} \sum_{m=1}^M \sum_{j=1}^d (\hat{\mu}_j^{(m)} - \tilde{\mu}_j^{(m)})^2 \tag{15}$$

where $\hat{\mu}_j^{(m)}$ is the estimate in replication m for $m = 1, \dots, M$. The left panel of Figure 1 shows this as a function of n in the curve labeled cwMean. We see that it goes down to zero for increasing n , indicating that cwMean is in fact asymptotically equivalent to the estimate $\hat{\mu}$ of cwMLE. (Note that the Chebyshev inequality implies that $\sqrt{n}(\tilde{\mu}_j - \hat{\mu}_j)$ goes to zero in probability.) We see the same effect for the analogous quantity comparing the diagonal entries of cwCov with $\hat{\Sigma}_{jj}$. The bottom curve in the right panel of Figure 1 compares the off-diagonal entries of cwCov with $\hat{\Sigma}_{jk}$ and goes to zero too. All of this suggests that the combination of cwMean and cwCov is asymptotically equivalent to the cwMLE method, which is intuitively understandable since the construction of the weights (13) mimics the guiding principle of the cellwise weighted likelihood. On the other hand, the upper curve in the right panel does not go to zero, indicating that sqrtCov is not asymptotically equivalent with cwMLE.

The asymptotic equivalence of the pair (cwMean,cwCov) with cwMLE can be exploited in several ways. Since cwMean and cwCov are easy to compute, they could be used as replacements or approximations of cwMLE when cwCov is PSD. And if cwCov is not PSD we can easily regularize it by carrying out the spectral decomposition of cwCov and replacing its nonpositive eigenvalues by a small positive number.

Table 3
MSE multiplication factors of cellwise weighted estimators when the weights are random and uniform on [0,1].

n	estimates for μ		for diagonal of Σ		for off-diagonal of Σ		
	cwMLE	cwMean	cwMLE	cwCov	cwMLE	cwCov	sqrtCov
10	1.31	1.31	1.22	1.22	1.23	1.19	1.03
20	1.33	1.34	1.26	1.25	1.36	1.30	1.11
50	1.30	1.30	1.35	1.35	1.43	1.41	1.21
100	1.34	1.34	1.34	1.34	1.48	1.48	1.26
200	1.35	1.35	1.32	1.32	1.49	1.50	1.26
500	1.31	1.31	1.28	1.28	1.46	1.46	1.24
1000	1.35	1.35	1.31	1.31	1.48	1.48	1.26
2000	1.35	1.35	1.32	1.32	1.52	1.52	1.28
5000	1.31	1.31	1.29	1.29	1.50	1.50	1.27
10000	1.34	1.34	1.37	1.37	1.47	1.47	1.25
∞	1.33	1.33	1.33	1.33	1.50	1.50	1.27

Another possibility is to use cwMean and cwCov (regularized when needed) as initial estimates in the algorithm of cwMLE. This is now an option in the R code. In experiments this reduced the number of iteration steps, while the result was identical. (The simulation yielding [Figure 1](#) didn't use this option, so the effect we saw is not an artifact of the computation.)

5. Results with precise data cells and random weights

We now look at a different situation, where the data cells are actually precise but we use weights that are not constant. We generate i.i.d. data \mathbf{x}_i following a multivariate gaussian random variable X , and let the weights w_{ij} in the matrix \mathbf{W} be i.i.d. according to a random variable W that is independent of X . The latter condition resembles the missing completely at random (MCAR) assumption for missing data. In this setting one can verify that the components $\tilde{\mu}_j$ of cwMean are asymptotically normal, and hence consistent. The asymptotic variance of $\tilde{\mu}_j$ equals that of the unweighted mean (which is also the unweighted MLE) multiplied by the factor

$$V(W) := \frac{E[W^2]}{E[W]^2} \quad (16)$$

which is at least 1 since $E[W^2] - E[W]^2 = \text{Var}[W] \geq 0$, so the asymptotic efficiency $\text{eff} = 1/V(W)$ is at most 1. The same factor $V(W)$ also multiplies the asymptotic variance of the diagonal elements $\tilde{\Sigma}_{jj} = \tilde{\Sigma}_{jj}$. The variances of the off-diagonal entries $\tilde{\Sigma}_{jk}$ are instead multiplied by $V(\tilde{W})$ where $\tilde{W} = \min(W_1, W_2)$ in which W_1 and W_2 are independent copies of W . For the sqrtCov matrix the factor becomes $V(\tilde{\tilde{W}})$ where $\tilde{\tilde{W}} = \sqrt{W_1 W_2}$.

A small simulation was run to illustrate these properties. The data were generated from the bivariate standard gaussian distribution, with 5000 replications for each value of n . The weights were randomly generated according to the uniform random variable W on [0,1]. Using the uniform variable W yields the population factor $V(W) = (1/4)/(1/3) = 4/3 \approx 1.33$ for the asymptotic variance of cwMean and the diagonal entries of cwCov and sqrtCov. For the off-diagonal entries of cwCov we require the distribution of \tilde{W} which has density $f(w) = 2(1-w)I(0 \leq w \leq 1)$ and $E[\tilde{W}] = 1/9$, $E[\tilde{W}^2] = 1/6$ so $V(\tilde{W}) = 3/2 = 1.50$. The computation is a bit harder for the off-diagonal entries of sqrtCov. There $\tilde{\tilde{W}}$ has density $g(\tilde{\tilde{w}}) = 4\tilde{\tilde{w}} \log(1/\tilde{\tilde{w}})I(0 \leq \tilde{\tilde{w}} \leq 1)$ which yields $E[\tilde{\tilde{W}}] = 16/81$ and $E[\tilde{\tilde{W}}^2] = 1/4$ so $V(\tilde{\tilde{W}}) = 81/64 \approx 1.27$.

The entries in [Table 3](#) are the mean squared errors of the estimates for μ_j averaged over $j = 1, 2$, and likewise for the off-diagonal entries Σ_{jk} . Since the unweighted MLE estimators $\bar{\mathbf{x}}$ and Cov are efficient for these data with precise cells, we do not list them here. The MSE of the cellwise weighted estimates for $\tilde{\mu}_j$ should trend to the value $V(F)$. For the diagonal entries Σ_{jj} we divide the MSE by 2 (which is the asymptotic variance of the unweighted estimator) so the result should go to $V(F)$ as well. For the off-diagonal entries, the MSE should trend to $V(\tilde{W})$ for cwMLE and cwCov, and to $V(\tilde{\tilde{W}})$ for sqrtCov.

In [Table 3](#) we see that for $n \geq 100$ the empirical MSE multiplication factors are quite close to their population versions listed in the row $n = \infty$. We also note that the MSE values of the cwMLE location are close to those of cwMean, that those of the diagonal of the cwMLE covariance are close to those of cwCov, and similarly for the off-diagonal entries of these covariances. This confirms our expectation that the asymptotic variances of cwMLE coincide with those of cwMean and cwCov.

The left panel of [Figure 2](#) shows (15) as in [Figure 1](#) and again indicates that cwMean is asymptotically equivalent to the estimator $\hat{\mu}$ of cwMLE. The other curves in [Figure 2](#) reflect that cwCov is asymptotically equivalent to the cwMLE covariance estimator whereas sqrtCov is not.

In the last column of [Table 3](#) we see that the off-diagonal entries of sqrtCov are more efficient than those of cwCov. There are two reasons for this. First, the simulation is for the idealized situation where the \mathbf{X} sample is perfectly gaussian with constant accuracy, and in that situation the unweighted covariance Cov would perform best. And secondly, sqrtCov is more similar to Cov than cwCov is, since its weights are closer to constant. Since weights are only defined up to a factor,

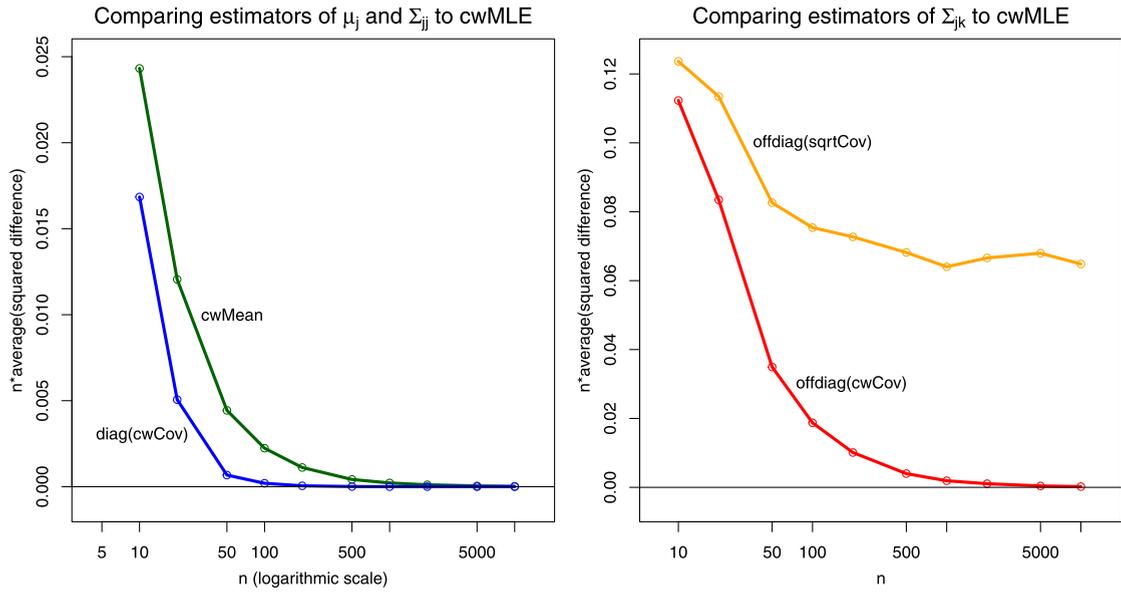


Fig. 2. Precise data cells with random weights: plot of (15) for (left) the estimator cwMean and the diagonal entries of cwCov, and (right) the off-diagonal entries of cwCov (lower curve) and of sqrtCov (upper curve).

how close they are to constant can be measured by their coefficient of variation $cv[W] = \text{Stdev}[W]/E[W]$. It is easily seen that there is a monotone relation between the variance factor and the coefficient of variation:

$$V(W) = cv(W)^2 + 1. \tag{17}$$

In the current setting we have $cv(\tilde{W})^2 = 1/2$ and $cv(\tilde{W})^2 = 17/64 \approx 0.2656$ so the weights of sqrtCov have a lower cv than those of cwCov, and hence yield a lower variance factor. For a constant weight W we would get $cv(W) = 0$ yielding a variance factor of 1, which is the lowest possible.

Repeating the simulation for other dimensions d gave similar results (not shown).

Let us now consider the situation where the dataset \mathbf{X} contains NA's that are missing completely at random. This can be put in our framework by using a matrix \mathbf{W} of cell weights that are 0 or 1, such that the zeroes in \mathbf{W} are placed at the positions of the NA's in \mathbf{X} . In that situation the unpacked matrix $\mathbf{X}^{(W)}$ is just \mathbf{X} and all its row weights are 1. Therefore, the cellwise weighted likelihood coincides with the observed likelihood of the incomplete dataset \mathbf{X} . The cwMLE method then reduces to the MLE of incomplete data, whose computation requires iteration. On the other hand, we can still compute cwMean and cwCov explicitly. Note that sqrtCov coincides with cwCov in this setting, because for zero-one weights w_{ij} and w_{ik} it holds that $\min(w_{ij}, w_{ik}) = \sqrt{w_{ij}w_{ik}}$. Also note that in this situation the entry $\tilde{\Sigma}_{jk}$ of cwCov is just the average of the $(x_{ij} - \tilde{\mu}_j)(x_{ik} - \tilde{\mu}_k)$ over the pairs with both x_{ij} and x_{ik} non-missing.

We have simulated the MCAR setting by generating the weights from a Bernoulli random variable with success probability 0.9, corresponding to 10% of missing values. The \mathbf{X} data were generated as before. We ran 5,000 replications for each sample size n . From the properties of the Bernoulli random variable W we immediately obtain the variance factor $V(W) = E[W^2]/E[W]^2 = 0.9/(0.9)^2 = 1/0.9 \approx 1.11$. Since the distribution of \tilde{W} is Bernoulli with success probability $0.9^2 = 0.81$ we analogously find $V(\tilde{W}) = 1/0.81 \approx 1.23$.

In Table 4 we again see that the limiting behavior takes hold already at low sample sizes. Not surprisingly, the efficiency of the location estimates and the diagonal of the covariance matrices is $1/V(W) = 90\%$ which is the fraction of non-missing cells x_{ij} . Analogously, the efficiency of the off-diagonal of the covariance is $1/V(\tilde{W}) = 81\%$, the percentage of non-missing pairs (x_{ij}, x_{ik}) . Also, Figure 3 illustrates the asymptotic equivalence of the combination of cwMean and cwCov with cwMLE. Since in the MCAR situation cwMLE is just the usual MLE for incomplete data, and cwCov has the simple expression above, this asymptotic equivalence was presumably known before to some.

6. Example

Cellwise weights can be due to varying accuracy or reliability of entries in the data matrix, which differs from the concept of random noise that underlies much of statistics. Often it is assumed that all cells are equally accurate, but this may not be true in reality. A scientific community that cares about the accuracy of data is that of soft computing, and in particular fuzzy numbers. A fuzzy number is a fuzzy set, which is not localized in a single point but has a membership function. The more spread out the fuzzy number, the less accurate the measurement is.

Table 4

MSE multiplication factors of cellwise weighted estimators when the weights are zero-one with the zeroes at MCAR missing values.

n	estimates for μ		for diagonal of Σ		for off-diagonal of Σ	
	cwMLE	cwMean	cwMLE	cwCov	cwMLE	cwCov
10	1.14	1.10	1.14	1.10	1.36	1.15
20	1.13	1.12	1.10	1.09	1.24	1.15
50	1.09	1.09	1.11	1.11	1.23	1.20
100	1.09	1.09	1.13	1.13	1.25	1.24
200	1.09	1.09	1.14	1.14	1.26	1.25
500	1.11	1.11	1.09	1.09	1.24	1.24
1000	1.11	1.11	1.10	1.10	1.26	1.26
2000	1.11	1.11	1.09	1.09	1.28	1.28
5000	1.12	1.12	1.10	1.10	1.22	1.22
10000	1.13	1.13	1.11	1.11	1.24	1.24
∞	1.11	1.11	1.11	1.11	1.23	1.23

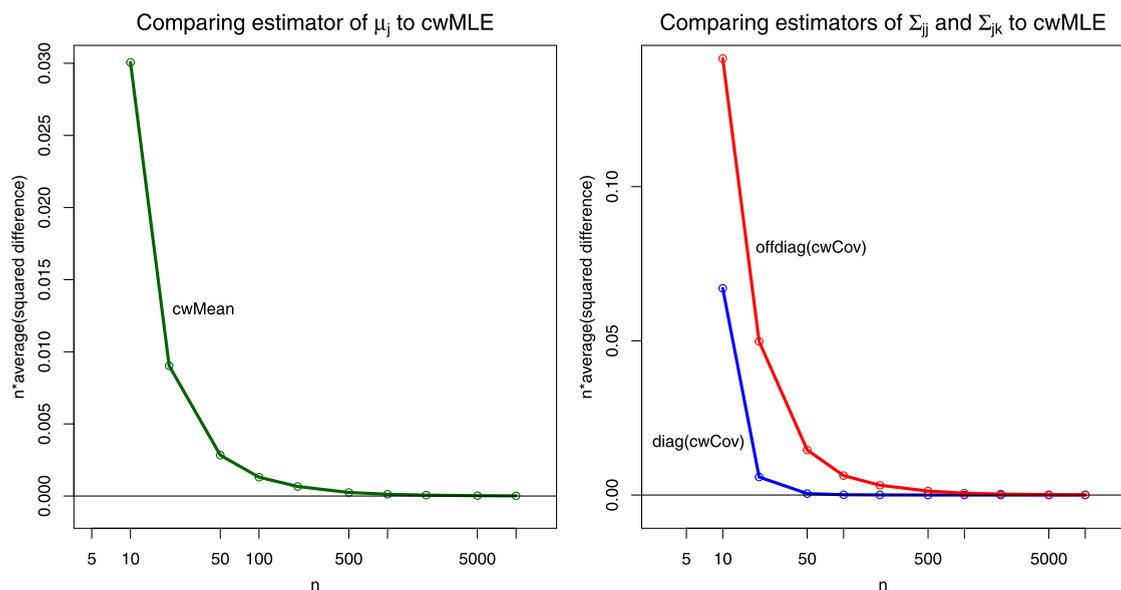


Fig. 3. Data with MCAR missing values: plot of (15) for (left) the estimator cwMean, and (right) the diagonal and off-diagonal entries of cwCov.

As an example we consider a fuzzy dataset of [Hesamian and Akbari \(2019\)](#) about six personality traits of 10 subjects. The data matrix \mathbf{X} is in the left panel of [Table 5](#). The weights in the right panel are the inverse of the length of the support of the membership functions, normalized so the largest weight equals one. Due to its small sample size and lack of detail this dataset is not very interesting in itself, but it offers the opportunity to illustrate some aspects of the methods developed here.

Estimating the covariance matrix of these cellwise weighted data by cwMLE is immediate. We looked at scatterplots of each pair of variables, with the 95% confidence ellipses of cwMLE as well as those of the plain unweighted MLE. In most of these plots the ellipses looked rather similar, but let us consider a pair of variables for which they differ.

[Figure 4](#) plots trait 6 versus trait 3. The data points are shown as crosses, with the horizontal bar proportional to the cell weight of the x-coordinate, and the vertical bar to that of the y-coordinate. These weights vary a lot across the data. We see only 8 crosses rather than 10 because subjects 3 and 4 are tied here, as are subjects 5 and 9. We visualized this by adding up the cell weights of the tied subjects.

The solid blue ellipse represents cwMLE, and we see that its approximation cwCov (dashed line) is quite close to it. Both are quite dissimilar to the red ellipse of the unweighted MLE, which extends further to the left and yields a lower correlation coefficient (0.10 versus 0.32). That the red ellipse extends further to the left is because it gives all coordinates weight one, so the two x-coordinates on the left hand side pull as hard as all the others, unlike in cwMLE which takes their low cell weights into account. The centers of the blue ellipses lie higher than the red one, and the blue ellipses are a bit slanted to the right, mainly due to the large vertical cell weight of the data point in the upper right corner.

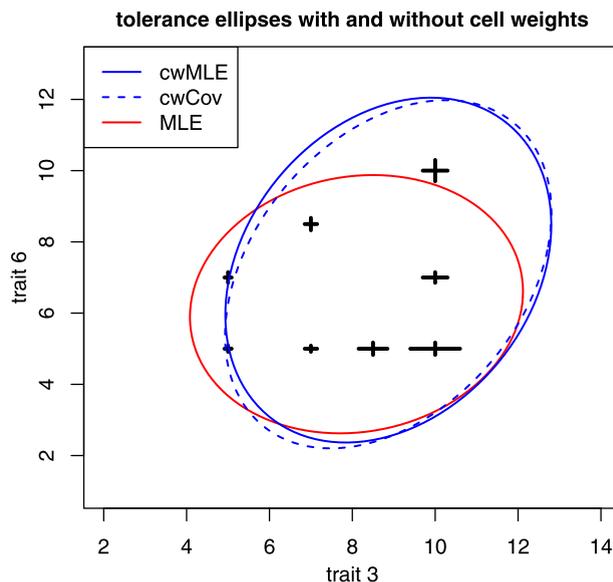


Fig. 4. Plot of variable t6 in Table 5 versus variable t3. The arms of the crosses reflect the cell weight of each coordinate. The 95% tolerance ellipses of cwMLE, cwCov, and the unweighted MLE are shown.

Table 5
Cellwise weighted data on personality traits.

data matrix \mathbf{X}						weight matrix \mathbf{W}					
t1	t2	t3	t4	t5	t6	t1	t2	t3	t4	t5	t6
7	5	7	5	5	5	0.50	0.29	0.50	0.29	0.29	0.29
10	10	10	7	8.5	7	1.00	1.00	1.00	0.50	0.58	0.50
5	5	10	5	5	5	0.29	0.29	1.00	0.29	0.29	0.29
10	10	10	5	5	5	1.00	1.00	1.00	0.29	0.29	0.29
7	7	8.5	5	5	5	0.50	0.50	0.58	0.29	0.29	0.29
10	5	5	8.5	8.5	5	1.00	0.29	0.29	0.58	0.58	0.29
5	7	7	5	5	8.5	0.29	0.50	0.50	0.29	0.29	0.58
10	10	10	10	10	10	1.00	1.00	1.00	1.00	1.00	1.00
8.5	7	8.5	5	5	5	0.58	0.50	0.58	0.29	0.29	0.29
5	10	5	7	5	7	0.29	1.00	0.29	0.50	0.29	0.50

7. Summary and Outlook

When faced with cellwise weighted data one can use the proposed likelihood function, for which it is convenient to apply the unpacking transform to the data. After this transform one can carry out cellwise weighted maximum likelihood estimation (cwMLE) of the parameters, as well as likelihood-based inference.

For the ubiquitous multivariate gaussian model an iterative algorithm for the cwMLE is made available. The faster explicit methods cwMean and cwCov are asymptotically equivalent to the cwMLE and can be seen as approximations, if needed after regularizing cwCov to make it PSD. In simulations the limiting behavior was accurate already at relatively low sample sizes.

A reviewer inquired about non-gaussian data. The likelihood of an alternative model distribution is different but formula (9) of the cellwise weighted likelihood can still be applied, as well as unpacking and the EM algorithm. For instance, the cwMLE can be used for data from a multivariate t -distribution, requiring only a bit more computation time. The approximations cwMean and cwCov are not as general and would obtain a lower statistical efficiency in that situation. Constructing fast approximations specifically tailored to the t -distribution would be harder since there is no explicit formula for its unweighted MLE to begin with.

The main benefit of this note is expected to be in the analysis of data in which the cells are measured with different accuracies, or there are other reasons to assume that the reliability varies across cells. Section 6 gave an example with such data. Other fields where data cells have different accuracies are cDNA arrays (Lawrence et al., 2004) and oligonucleotide arrays (Turro et al., 2007) where credibility intervals for the data values are derived from posterior distributions.

The methodology in this note is also intended to be used in robust statistics, where missing values can be difficult to handle, see e.g. Cantoni and de Luna (2020). A specific application is to the emerging field of cellwise outliers that started with the publication of Alqallaf et al. (2009). There are methods that detect outlying cells, such as the Detect Deviating Cells method (Rousseeuw and Van den Bossche, 2018) or the cellwise MCD method (Raymaekers and Rousseeuw, 2022). Both

of these provide standardized cellwise residuals, which are large for outlying cells. After such a method has run, one can assign weights to the cells based on the size of their standardized cellwise residuals. The approach proposed here can then produce cellwise reweighted estimates. This postprocessing step may benefit the overall stability and accuracy of the final result. It is analogous to the casewise reweighting step that is often carried out after a casewise robust method.

Software availability. An R implementation of the proposed techniques has been incorporated in the `cellwise` package (Raymaekers and Rousseeuw, 2023) on CRAN, with the vignette `cellwise_weights_examples` reproducing the example in Section 6.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

Thanks go to Stefan Van Aelst, Jakob Raymaekers, and the reviewers for helpful comments improving the presentation.

References

- Agostinelli, C., Markatou, M., 1998. A one-step robust estimator for regression based on the weighted likelihood reweighting scheme. *Statistics & Probability Letters* 37, 341–350.
- Agostinelli, C., Markatou, M., 2001. Test of hypotheses based on the weighted likelihood methodology. *Statistica Sinica* 11, 499–514.
- Alqallaf, F., Van Aelst, S., Yohai, V.J., Zamar, R.H., 2009. Propagation of outliers in multivariate data. *The Annals of Statistics* 37, 311–331.
- Cantoni, E., de Luna, X., 2020. Semiparametric inference with missing data: Robustness to outliers and model misspecification. *Econometrics and Statistics* 16, 108–120.
- Croux, C., Haesbroeck, G., Ruwet, C., 2013. Robust estimation for ordinal regression. *Journal of Statistical Planning and Inference* 143, 1486–1499.
- Dempster, A., Laird, N., Rubin, D., 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B* 39 (1), 1–22.
- Field, C., Smith, B., 1994. Robust estimation – a weighted maximum likelihood approach. *International Statistical Review* 62, 405–424.
- Fraley, C., Raftery, A. E., Scrucca, L., Murphy, T. B., Fop, M., 2022. Package `mclust`: Gaussian Mixture Modelling for Model-Based Clustering, Classification, and Density Estimation. CRAN, R package 5.4.10. <https://CRAN.R-project.org/package=mclust>.
- Hesamian, G., Akbari, M.G., 2019. Principal component analysis based on intuitionistic fuzzy random variables. *Computational and Applied Mathematics* 38 (158), 1–14.
- Hu, F., 1994. Relevance Weighted Smoothing and a New Bootstrap Method. The University of British Columbia.
- Hu, F., Zidek, J.V., 2002. The weighted likelihood. *The Canadian Journal of Statistics* 30, 347–371.
- Jamshidian, M., Bentler, P.M., 1999. ML estimation of mean and covariance structures with missing data using complete data routines. *Journal of Educational and Behavioral Statistics* 24, 21–41.
- Jorgenson, B., Petersen, H.C., 2012. Efficient estimation for incomplete multivariate data. *Journal of Statistical Planning and Inference* 142, 1215–1224.
- Lawrence, N., Milo, M., Niranjana, M., Rashbass, P., Soullier, S., 2004. Reducing the variability in cDNA microarray image processing by Bayesian inference. *Bioinformatics* 20 (4), 518–526.
- Little, R., Rubin, D., 2020. *Statistical analysis with missing data* (third edition). John Wiley and Sons, New York.
- Magis, D., 2015. A note on weighted likelihood and Jeffreys modal estimation of proficiency levels in polytomous item response models. *Psychometrika* 80, 200–204.
- Majumder, S., Biswas, A., Roy, T., Kumar Bhandari, S., Basu, A., 2021. Statistical inference based on a new weighted likelihood approach. *Metrika* 84, 97–120.
- Meng, X.-L., Rubin, D.B., 1991. Using EM to obtain asymptotic variance-covariance matrices: The SEM algorithm. *Journal of the American Statistical Association* 86, 899–909.
- O'Hagan, A., Brendan Murphy, T., Scrucca, L., Gormley, I.C., 2019. Investigation of parameter uncertainty in clustering using a Gaussian mixture model via jackknife, bootstrap and weighted likelihood bootstrap. *Computational Statistics* 34, 1779–1813.
- Raymaekers, J., Rousseeuw, P. J., 2022. The cellwise minimum covariance determinant estimator. *arxiv* 2207.13493. <https://arxiv.org/abs/2207.13493>.
- Raymaekers, J., Rousseeuw, P. J., 2023. `cellwise`: Analyzing Data with Cellwise Outliers. R package, CRAN. <https://CRAN.R-project.org/package=cellwise>.
- Rousseeuw, P.J., Van den Bossche, W., 2018. Detecting deviating data cells. *Technometrics* 60, 135–145. <https://doi.org/10.1080/00401706.2017.1340909>
- Sung, Y.J., Geyer, C.J., 2007. Monte Carlo likelihood inference for missing data models. *The Annals of Statistics* 35, 990–1011.
- Takai, K., 2020. Incomplete-data Fisher scoring method with steplength adjustment. *Statistics and Computing* 30, 871–886.
- Turro, E., Bochkina, N., Hein, A.-M., Richardson, S., 2007. BGX: a Bioconductor package for the Bayesian integrated analysis of Affymetrix GeneChips. *BMC Bioinformatics* 8:439, 1–10.
- Van Aelst, S., Vandervieren, E., Willems, G., 2011. Stahel-Donoho estimators with cellwise weights. *Journal of Statistical Computation and Simulation* 81 (1), 1–27.