

Application of the Iterated Weighted Least-Squares Fit to counting experiments

Hans Dembinski^{a,*}, Michael Schmelling^a, Roland Waldi^b

^a*Max Planck Institute for Nuclear Physics, Heidelberg, Germany*

^b*Rostock University, Rostock, Germany*

Abstract

Least-squares fits are popular in many data analysis applications, and so we review some theoretical results in regard to the optimality of this fit method. It is well-known that common variants of the least-squares fit applied to Poisson-distributed data produce biased estimates, but it is not well-known that the bias can be overcome by iterating an appropriately weighted least-squares fit. We prove that the iterated fit converges to the maximum-likelihood estimate. Using toy experiments, we show that the iterated weighted least-squares method converges faster than the equivalent maximum-likelihood method when the statistical model is a linear function of the parameters and it does not require problem-specific starting values. Both can be a practical advantage. The equivalence of both methods also holds for binomially distributed data. We further show that the unbinned maximum-likelihood method can be derived as a limiting case of the iterated least-squares fit when the bin width goes to zero, which demonstrates the deep connection between the two methods.

1. Introduction

In this paper, we review some theoretical results on least-squares methods, in particular, when they yield optimal estimates. We show how they can be applied to counting experiments without sacrificing optimality. The insights discussed here are known in the statistics community [1, 2], but less so in the

*Corresponding author

high-energy physics community. Standard text books on statistical methods and papers, see *e.g.* [3, 4], correctly warn about biased results when standard variants of the least-squares fit are applied to counting experiments with small numbers of events, but do not show that these can be overcome. The results presented here are of practical relevance for fits of linear models, where the iterated weighted least-squares method discussed in this paper converges faster than the standard maximum-likelihood method and does not require starting values near the optimum.

The least-squares fit is a popular tool of statistical inference. It can be applied in situations with k measurements $\{y_i|i = 1, \dots, k\}$, described by a model with m parameters $\mathbf{p} = (p_j|j = 1, \dots, m)$ that predicts the expectation values $E[y_i] = \mu_i(\mathbf{p})$ for the measurements. The measurements differ from the expectation values by unknown residuals $\epsilon_i = y_i - \mu_i(\mathbf{p})$. The solution $\hat{\mathbf{p}}$ that minimizes the sum $Q(\mathbf{p})$ of squared residuals,

$$Q(\mathbf{p}) = \sum_{i=1}^k (y_i - \mu_i(\mathbf{p}))^2, \quad (1)$$

is taken as the best fit of the model to the data.

More generally, the measurements and the model predictions can be regarded as k -dimensional vectors $\mathbf{y} = (y_i|i = 1, \dots, k)$ and $\boldsymbol{\mu} = (\mu_i|i = 1, \dots, k)$, for which one wants to minimize a distance measure. In Eq. (1), we minimize the squared Euclidean distance. A generalization is the bilinear form

$$Q(\mathbf{p}) = (\mathbf{y} - \boldsymbol{\mu})^T \mathbf{W} (\mathbf{y} - \boldsymbol{\mu}), \quad (2)$$

where \mathbf{W} is a positive-definite symmetric matrix of weights. This variant is called *weighted least squares* (WLS). Eq. (1) is recovered with $\mathbf{W} = \mathbf{1}$. An important special case is when the weight matrix is equal to the inverse of the true covariance matrix \mathbf{C} of the measurements, $\mathbf{W} = \mathbf{C}^{-1}$ with $\mathbf{C} = E[\mathbf{y}\mathbf{y}^T] - E[\mathbf{y}]E[\mathbf{y}]^T$. For uncorrelated measurements, Eq. (2) simplifies to the familiar form

$$Q(\mathbf{p}) = \sum_{i=1}^k (y_i - \mu_i(\mathbf{p}))^2 / \sigma_i^2, \quad (3)$$

with variances $\sigma_i^2 = \text{E}[y_i^2] - \text{E}[y_i]^2$.

Aitken [5] showed in a generalization to the Gauss-Markov theorem [3, p. 152] that minimizing $Q(\mathbf{p})$ with $\mathbf{W} \propto \mathbf{C}^{-1}$ produces an optimal, in the sense as detailed below, solution for linear models $\boldsymbol{\mu}(\mathbf{p}) = \mathbf{X}\mathbf{p}$, where \mathbf{X} is a constant $k \times m$ matrix. The theorem applies when the covariance matrix \mathbf{C} is finite and non-singular. Then, $Q(\mathbf{p})$ has a unique minimum at

$$\hat{\mathbf{p}} = (\mathbf{X}^T \mathbf{C}^{-1} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{C}^{-1} \mathbf{y}. \quad (4)$$

The best fit parameters $\hat{\mathbf{p}}$ in this case are a linear function of the measurements \mathbf{y} with the covariance matrix

$$\mathbf{C}_p = (\mathbf{X}^T \mathbf{C}^{-1} \mathbf{X})^{-1}. \quad (5)$$

If the measurements are unbiased, $\text{E}[\mathbf{y}] = \boldsymbol{\mu}$, this solution is the *best linear unbiased estimator* (BLUE). Like all linear estimators, Eq. (4) is unbiased if the input is unbiased. In addition, it has minimal variance of all linear estimators. This is true for any shape of the data distribution and any sample size. These excellent properties may be compromised in practical applications, since the covariance matrix \mathbf{C} is often only approximately known.

The least-squares approach is often regarded as a special case of the more general *maximum-likelihood* (ML) approach. The ML principle states that the best fit of a model should maximize the likelihood L , which is proportional to the joint probability of all measurements under the model. In practice, it is more convenient to work with $\ln L$ rather than L , so that the product of probabilities turns into a sum of log-probabilities,

$$\ln L(\mathbf{p}) = \ln \prod_{i=1}^k P_i(y_i; \mathbf{p}) = \sum_{i=1}^k \ln P_i(y_i; \mathbf{p}). \quad (6)$$

Here $P_i(y_i; \mathbf{p})$ is the value of the probability density at y_i for continuous outcomes or the actual probability for discrete outcomes. The ML method needs a fully specified probability distribution for each measurement, while the WLS method uses only the first two moments.

The parameter vector $\hat{\mathbf{p}}$ that maximizes Eq. (6) is called the *maximum-likelihood estimate* (MLE). MLEs have optimal asymptotic properties; asymptotic here means in the limit of infinite samples. They are *consistent* (asymptotically unbiased) and *efficient* (asymptotically attaining minimal variance) [3]. In many practical cases of inference, in particular when data are Poisson-distributed, this method is known to produce good estimates also for finite samples. These properties make the ML fit the recommended tool for the practitioner [3, 4].

The WLS fit can be derived as a special case of a ML fit, if one considers normally distributed measurements y_i with expectations μ_i and variances σ_i^2 , where each measurement has the probability density function (PDF)

$$P_i(y_i; \mu_i, \sigma_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(y_i - \mu_i)^2}{2\sigma_i^2}\right). \quad (7)$$

For fixed σ_i^2 , we obtain Eq. (3) from Eq. (6),

$$\ln L(\mathbf{p}) = -\sum_{i=1}^k \frac{(y_i - \mu_i(\mathbf{p}))^2}{2\sigma_i^2} + c \equiv -\frac{1}{2}Q(\mathbf{p}) + c, \quad (8)$$

where the constant term c depends only on the fixed variances σ_i^2 . Constant terms do not affect the location $\hat{\mathbf{p}}$ of the maximum of $\ln L$ and the minimum of Q . We will often drop them from equations.

This derivation shows that for Gaussian PDFs a ML and a WLS fit give identical results when the same fixed variances are used, even if they are not the true variances. This does not hold in general, but is relevant in this context. When data are Poisson-distributed and have small counts, common implementations of the WLS fit are biased as we will show in the following section. The bias does not originate from the skewed shape of the Poisson distribution however, but rather from the fact that the weights are either biased or not fixed.

To these standard methods, we add the *iterated weighted least-squares* (IWLS) fit [1]. It yields maximum-likelihood estimates when data are Poisson or binomially distributed with only the probabilities as free parameters [2]. This extends the strict equivalence between ML and WLS fits to a larger class of problems,

an extension which is highly relevant in practice, since counts in histograms are Poisson distributed, and counted fractions are binomially distributed with the denominator considered fixed. The iterations are used to successively update estimates of the variances σ_i^2 , which are kept constant during minimization.

When IWLS and ML fits are equivalent, which one is recommended? We conducted toy experiments where IWLS and ML fits are carried out numerically, as is common in practice. We found similar convergence rates for both methods when the model is non-linear, and a significantly faster convergence for the IWLS fit if the model is linear. This makes the IWLS fit a useful addition to the toolbox.

We have seen how the WLS fit can be derived from the ML fit under certain conditions. Inversely, we will show that the unbinned ML fit can be derived as a limiting case from the IWLS fit under weak conditions. The derivation shows that the two approaches are deeply connected.

2. Least-Squares Variants In Use

Standard variants of the WLS fit used in practice produce biased estimates when the fit is applied to Poisson-distributed data with small counts. The bias is often attributed to the breakdown of the normal approximation to the Poisson distribution, but it is actually related to how the unknown true variances σ_i^2 in Eq. (3) are replaced by estimates.

We demonstrate this along a simple example. We fit the single parameter μ of the Poisson-distribution

$$P(n; \mu) = e^{-\mu} \mu^n / n!, \quad (9)$$

to k counts $\{n_i | i = 1, \dots, k\}$ sampled from it. The maximum-likelihood estimate for μ can be computed analytically by maximizing Eq. (6). We solve $\partial \ln L / \partial \mu \equiv \partial_\mu \ln L = 0$ for μ and obtain the arithmetic average

$$\hat{\mu} = \frac{1}{k} \sum_{i=1}^k n_i, \quad (10)$$

which is unbiased and has minimal variance. We will now apply variants of the WLS fit to the same problem, which differ in how they substitute the unknown true variance.

Variance computed for each sample. For a single isolated sample, the unbiased estimate of μ is $\hat{\mu}_i = n_i$, with variance $\text{Var}[n_i] = \mu \simeq \hat{\mu}_i = n_i$. This is the origin of the well-known \sqrt{n} -estimate for the standard deviation of a count n . With this variance estimate, we get

$$Q(\mu) = \sum_{i=1}^k (n_i - \mu)^2 / n_i. \quad (11)$$

This form is called Neyman's χ^2 in the statistics literature [4]. Replacing the true variance μ by its sample estimate n_i is an application of the *bootstrap principle* discussed by Efron and Tibshirani [6]. To obtain the minimum, we solve $\partial_\mu Q = 0$ for μ and obtain the harmonic average

$$\frac{1}{\hat{\mu}} = \frac{1}{k} \sum_{i=1}^k \frac{1}{n_i}. \quad (12)$$

The solution is biased and breaks down for samples with $n_i = 0$. The variance estimates here are constant (they do not vary with $\hat{\mu}$), but differ from sample to sample. This treatment ignores the fact that the true variance is the same for all samples in this setup.

Variance computed from model. Another choice is to directly insert $\text{Var}[n_i] = \mu$ in the formula,

$$Q(\mu) = \sum_{i=1}^k (n_i - \mu)^2 / \mu. \quad (13)$$

This form, called Pearson's χ^2 [4], is a conceptual improvement, because μ is the exact but unknown value of the variance. However, the variance μ now varies together with the expectation value μ . Solving $\partial_\mu Q = 0$ for μ yields the quadratic average

$$\hat{\mu} = \sqrt{\frac{1}{k} \sum_{i=1}^k n_i^2}. \quad (14)$$

This estimate is also biased, but can handle samples with $n_i = 0$. The bias may come at a surprise, since we used the exact value for the variance after all. The failure here can be traced back to the fact that the variance estimates $\sigma_i^2 = \mu$ are not fixed during the minimization. A small positive bias on μ in Eq. (13) leads to a second order increase in the numerator, which is overcompensated by a first order increase of the denominator. In other words, the fit tends to increase the variance even at the cost of a small bias in the expectation when given this freedom, because overall it yields a reduction of Q .

Constant variance. Finally, we simply use $\sigma_i^2 = c$, where c is an arbitrary constant,

$$Q(\mu) = \sum_{i=1}^k (n_i - \mu)^2 / c. \quad (15)$$

We solve $\partial_\mu Q = 0$ for μ and obtain the optimal maximum-likelihood estimate Eq. (10) as the solution; the constant c drops out.

This seems counter-intuitive, since we used a constant for all samples instead of a value close or equal to the true variance. However, this case satisfies all conditions of the Gauss-Markov theorem. The expectation values are trivial linear functions of the parameter $\mu_i = \mu$. The variances σ_i^2 are all equal and only need to be known up to a global scaling factor, hence any constant c will do.

We learned that keeping the variance estimates constant during minimization is important, but the estimates should in general be as close to the true variances as possible. An iterated fit can satisfy both requirements.

3. Iterated Weighted Least-Squares

The iterated (re)weighted least-squares methods (IWLS or IRLS) are well known in statistical regression [1], and can be applied to fits with k measurements $\{y_i | i = 1, \dots, k\}$ described by a model with m parameters $\mathbf{p} = (p_j | j = 1, \dots, m)$, which predicts the expectations $E[y_i] = \mu_i(\mathbf{p})$ and variances $\text{Var}[y_i] = \sigma_i^2(\mathbf{p})$ of each measurement. We will discuss the special application where the y_i are entries of a histogram. One then minimizes the sum of squared

residuals

$$Q(\mathbf{p}) = \sum_{i=1}^k (y_i - \mu_i(\mathbf{p}))^2 / \sigma_i^2(\hat{\mathbf{p}}), \quad (16)$$

where the σ_i^2 are constant within one iteration of the fit and computed from the model using the parameter estimate $\hat{\mathbf{p}}$ that minimized $Q(\mathbf{p})$ in the previous iteration. A convenient choice for the first iteration is $\sigma_i^2 = 1$. One iterates until $\hat{\mathbf{p}}$ converges.

In particle physics, we often work with samples drawn from two monoparametric distributions of the exponential family:

- Poisson distribution. Example: fitting a distribution function to a histogram of counts.
- Binomial distribution with fixed number of trials. Example: fitting an efficiency function to two histograms with generated and accepted events.

Charles, Frome, and Yu [2] derived that the IWLS fit gives the exact same result as the ML fit for a family of distributions. We demonstrate this in the appendix for the special distributions discussed here.

The Hessian matrices of second derivatives are also equal up to a constant factor, $\partial_{p_l} \partial_{p_m} Q = -2 \partial_{p_l} \partial_{p_m} \ln L$. The inverse of the Hessian is an estimate of the covariance matrix of the solution, an important uncertainty estimate in practical applications.

We emphasize that the equivalence does not depend on the size of the data sample or on the functional form of the model that predicts the expectation values for the measurements. In particular, when the IWLS fit is applied to histograms, it is not biased by small counts per bin or even empty bins.

3.1. Including systematic uncertainties

A formal discussion of how systematic uncertainties can be handled with the IWLS fit is outside of the scope of this paper, but we note that it can include systematic uncertainties. Barlow [7] discusses how correlated systematic

uncertainties can be handled in a least-squares fit. One minimizes Eq. (2) in each iteration with a matrix

$$\mathbf{C} = \mathbf{C}'(\hat{\mathbf{p}}) + \mathbf{C}_{\text{sys}}(\hat{\mathbf{p}}), \quad (17)$$

where \mathbf{C}' is the current estimate of the stochastic covariance computed from the previous solution, and \mathbf{C}_{sys} is a current covariance matrix that represents the systematic uncertainties of the measurements. The matrix \mathbf{C}_{sys} may be a function of the parameter vector. Like the covariance matrix \mathbf{C}' , it is kept constant during each iteration, and updated between iterations using the current value of $\hat{\mathbf{p}}$. This approach has been successfully applied in a combination of measurements from the CDF and D0 experiments [8].

3.2. IWLS or ML fit?

When the IWLS and the ML fits are equivalent, which one should be used in practice? The two methods produce the same results in analytical problems, but can have different performance in numerical problems. In practice, the extrema of the log-likelihood function $\ln L(\mathbf{p})$ and the weighted least-squares function $Q(\mathbf{p})$ are usually found with a local optimizer, like the MIGRAD algorithm in the MINUIT package [9, 10]. Computing the functions is sometimes expensive, when the fitted data sets are large and the model has many parameters. Numerical methods are therefore judged based on the number of function evaluations required to converge to the optimum within some tolerance. Another criterion is robustness, the ability to converge to the right optimum from a point in the neighborhood of the solution.

To address these points, we conducted toy experiments with Poisson-distributed counts n_i and find that the ML method requires less function evaluations than the IWLS methods in general. However, the rate of convergence of the IWLS method can be greatly accelerated, when the model that computes the count expectation $E[n_i] = \mu_i(\mathbf{p})$ is linear in the parameters, $\mu_i = \mathbf{X}_i \mathbf{p}$, where \mathbf{X}_i is a vector of constants. The maximum of $\ln L(\mathbf{p})$ usually cannot be found analytically in this case, but the minimum of $Q(\mathbf{p})$ is given by Eq. (4) in each iteration of the IWLS fit. When the computing time is dominated by the eval-




Figure 1: Example of a toy data set (points) to test the performance of the ML, IWLS, and L-IWLS fits (see text). *Left:* The curve represent the fit result of the three methods for this data set (which are identical). *Right:* Intermediate parameter states during the optimization for the ML, IWLS, and L-IWLS fits (see text), in each iteration of the respective algorithms until the stopping criterion is reached. The L-IWLS fit often converges quadratically. The IWLS fit is slowed down by the artificial dampening that we introduced to avoid oscillations.

uation of $Q(\mathbf{p})$ or $\ln L(\mathbf{p})$, solving the IWLS fit is faster than the ML fit. The IWLS fit also does not require a problem-specific starting point for the optimization in this case. We call this special variant the L-IWLS fit. All three methods are able to handle fits that have bounded parameters, which are common in particle physics. In our toy experiments, the parameters are bounded to be non-negative. Details are given in the next section.

Whether the ML or the IWLS fits are more robust in the above sense is more difficult to say. No general proofs can be given for either method. Our toy studies suggest the following order of increasing robustness: IWLS, L-IWLS, ML. In some toy experiments, the IWLS methods require many more iterations than average, producing a long tail in the distribution of iteration counts. Such tails are not observed for the ML fit. It is likely, however, than a more sophisticated implementation of the IWLS fit than ours could improve the robustness of this method.

3.3. Performance in toy experiments

We compare the performances of ML, IWLS, and L-IWLS fits in a series of 1000 toy experiments with Poisson-distributed samples. We use a linear model for the expectation with two parameters, $\mu(x, \mathbf{p}) = (p_0 + p_1 x^2)$, with

$x \in [0, 1]$ as an independent variable. For the true parameters $\mathbf{p}_{\text{truth}} = (1, 10)$, we simulate 10 pairs (x_i, n_i) . The x_i are evenly spaced over the interval $[0, 1]$, μ_i is calculated for each x_i based on the true parameters, and finally a random sample n_i is drawn for each μ_i from the Poisson distribution. The model is then fitted to each toy data set using the following three methods. One of the toy experiments is shown in Fig. (1).

- ML fit: Starting from Eq. (A.1) we use the MIGRAD algorithm from the MINUIT package to find the minimum. We pass the exact analytical gradient to MIGRAD for this problem, replacing the numerical approximation that MINUIT uses otherwise. We restrict the parameter range to $p_k \geq 0$ and add an epsilon to μ whenever it appears in a denominator to avoid division by zero.
- IWLS fit: We use Newton’s method to update \mathbf{p} ,

$$\mathbf{p}_{n+1} = \mathbf{p}_n - \mathbf{H}^{-1} \partial_{\mathbf{p}} Q, \quad (18)$$

with the exact analytical gradient $\partial_{\mathbf{p}} Q$ and Hessian \mathbf{H} for this problem. Since the model is linear and the function Q quadratic, Newton’s method yields the exact solution for the given gradient and Hessian matrix, but without taking the boundary condition $p_k \geq 0$ into account. We resolve this in an ad hoc way, by setting negative parameter values are set to zero. Since the covariance matrix is fixed in each Newton step, each step fulfills the requirements of the IWLS method. We update the covariance matrix after each step for the computation of the next step. To check for convergence, we use the MINUIT criterion, which is based on the estimated distance-to-minimum and deviations in the diagonal elements of the inverted Hessian [9].

This approach works very well for most toy experiments, but in some rare cases ($< 1\%$) the solution starts to oscillate indefinitely between two states. We resolve this again in an ad hoc way by averaging the updated




Figure 2: Application of the maximum-likelihood (ML), the iterated weighted least-squares (IWLS) fit and its specialization for linear models (L-IWLS) to 1000 toy experiments with Poisson-distributed samples (see text). *Top row:* Histograms of the two fitted parameters of the model \hat{p}_0 and \hat{p}_1 are shown, overlaid for all three fit methods. The histograms are nearly identical. *Bottom row:* Normalized histograms of the number of evaluations of the model function for the the fit methods in double-logarithmic scale.

parameter vector with the previous one, $\mathbf{p}_{n+1} := (\mathbf{p}_{n+1} + \mathbf{p}_n)/2$ after each Newton step. This slows down the convergence rate drastically, but avoids the oscillations.

- L-IWLS fit: We solve Eq. (4) with the NNLS [11] algorithm as implemented in SciPy [12], and iterate. It solves Eq. (4) under the boundary condition $p_k \geq 0$. To check for convergence, we again use the MINUIT criterion.

We note that our application of the general IWLS fit to a problem with a linear model is artificial. We only do this here to compare all three fitting methods on the same problem. For the IWLS and ML fits, we use the optimistic

starting point $\mathbf{p}_{\text{truth}} = (1, 10)$. The ML and IWLS fits therefore run under ideal conditions compared to the L-IWLS fit, which does not require a specific starting point. In practice, one will usually start with a less ideal starting point, which slows down the convergence of ML and IWLS fits compared to the L-IWLS fit.

In case of the ML and IWLS fits, we increase the call counter for each evaluation of Q or $\ln L$ and each evaluation of their gradients for all values of z by one. In case of the L-IWLS fit, we count one application of the NNLS algorithm as one call, since it requires essentially one computation of the gradient.

The results are shown in Fig. (2). As expected, the results are equal within the numerical accuracies of the numerical algorithms, which stop when MINUIT's standard convergence criterion is reached. This criterion roughly gives a precision of about 10^{-3} in the parameter relative to its uncertainty.

The average number of calls required to converge is different: 19.3 for ML, 23.6 for IWLS, and 4.8 for L-IWLS. The L-IWLS fit is the fastest to converge, requiring only a quarter of the function evaluations of the ML fit. The IWLS fit is the slowest, it requires about 20% more calls on average than the ML fit. This is mainly due to artificial dampening. In cases where the dampening is not needed, the IWLS fit converges as rapidly as the L-IWLS fit. Since we chose a linear model for this performance study in order to compare all three methods, a Newton's step computes the exact solution to the fitting problem for the current covariance matrix estimate.

An investigation shows that the convergence issues of the IWLS fit appear when a parameter of the model is very close to zero. If this is not the case and no dampening is applied, the IWLS and L-IWLS fits produce identical results. MINUIT was designed to handle such cases well and shows a much more stable convergence rate. This suggests that the issues of the IWLS fit can be overcome as well with a more sophisticated implementation, but this comes at the cost of a slower convergence in favorable cases. The overall performance of the IWLS fit will probably not surpass that of the more straight-forward ML fit.

In conclusion, we recommend the L-IWLS fit for linear models and the ML fit for non-linear models.

4. Unbinned maximum-likelihood from IWLS

In the introduction, we reviewed how the WLS fit can be derived as a special case of the ML fit, when measurements are normally distributed with known variance. Alternatively, the WLS fit can also be derived from geometric principles without relying on the ML principle. We will now show that the unbinned ML fit can be derived as a limiting case of the IWLS fit.

For the unbinned ML fit of a known probability density $f(x; \mathbf{p})$ of a continuous stochastic variable x with parameters \mathbf{p} , one maximizes the sum of logarithms of the model density evaluated at the measurements $\{x_i | i = 1 \dots k\}$,

$$\ln L(\mathbf{p}) = \sum_{i=1}^k \ln f(x_i; \mathbf{p}). \quad (19)$$

The maximum is found by solving the system of equations $\partial_{p_j} \ln L(\mathbf{p}) = 0$. The density $f(x; \mathbf{p})$ must be at least once differentiable in \mathbf{p} .

To derive these equations as a limit of the IWLS fit, we assume that $f(x; \mathbf{p})$ is finite everywhere in x , so that the probability density is not concentrated in discrete points.

We start by considering a histogram of k samples x_i . Since the samples are independently drawn from a PDF, the histogram counts n_l are uncorrelated and Poisson-distributed. Following the IWLS approach, we minimize the function

$$Q(\mathbf{p}) = \sum_l \frac{(n_l - kP_l)^2}{k\hat{P}_l} \quad (20)$$

and iterate, where $P_l(\mathbf{p}) = \int_{x_l}^{x_l + \Delta x} f(x; \mathbf{p}) dx$ is the expected fraction of the samples in bin l , and $\hat{P}_l = P_l(\hat{\mathbf{p}})$ is the value based on the fitted parameters $\hat{\mathbf{p}}$ from the previous iteration. Expansion of the squares yields three terms,

$$Q(\mathbf{p}) = \sum_l \frac{n_l^2}{k\hat{P}_l} - 2 \sum_l \frac{n_l P_l}{\hat{P}_l} + k \sum_l \frac{P_l^2}{\hat{P}_l}. \quad (21)$$

The first term is proportional to $1/\Delta x$, but not a function of \mathbf{p} . Therefore it does not contribute to the minimum obtained by solving the equations $\partial_{p_j} Q(\mathbf{p}) = 0$. We drop it in the following and consider only the second and third term, which both are functions of \mathbf{p} .

We investigate the limit $\Delta x \rightarrow 0$. Since $f(x)$ is finite everywhere, we have ultimately either zero or one count in each bin. With $P_l \rightarrow f(x_l; \mathbf{p}) \Delta x$, the second term has a finite limit

$$\sum_l \frac{n_l P_l}{\hat{P}_l} \xrightarrow{\Delta x \rightarrow 0} \sum_{i=1}^k \frac{f(x_i; \mathbf{p}) \Delta x}{f(x_i; \hat{\mathbf{p}}) \Delta x} = \sum_{i=1}^k \frac{f(x_i; \mathbf{p})}{f(x_i; \hat{\mathbf{p}})}, \quad (22)$$

where only bins around the measurements x_i with one entry contribute ($n_l = 1$), and the bin widths cancel. The third term also has a finite limit,

$$\sum_l \frac{P_l^2}{\hat{P}_l} \xrightarrow{\Delta x \rightarrow 0} \sum_l \frac{f^2(x_l; \mathbf{p}) (\Delta x)^2}{f(x_l; \hat{\mathbf{p}}) \Delta x} = \int \frac{f^2(x; \mathbf{p})}{f(x; \hat{\mathbf{p}})} dx. \quad (23)$$

One Δx cancels in the ratio and in the limit $\Delta x \rightarrow 0$ the remaining sum is the very definition of a Riemann integral.

We now consider the derivatives $\partial_{p_j} Q(\mathbf{p})$ in the limit of many iterations. We assume that the iterations converge, so that the previous solution $\hat{\mathbf{p}}$ approaches the next solution \mathbf{p} . We get

$$\begin{aligned} \partial_{p_j} Q(\mathbf{p}) &= -2 \sum_{i=1}^k \frac{\partial_{p_j} f(x_i; \mathbf{p})}{f(x_i; \hat{\mathbf{p}})} + k \int \frac{2f(x; \mathbf{p}) \partial_{p_j} f(x; \mathbf{p})}{f(x; \hat{\mathbf{p}})} dx \\ &\xrightarrow{\hat{\mathbf{p}} \rightarrow \mathbf{p}} -2 \sum_{i=1}^k \frac{\partial_{p_j} f(x_i; \mathbf{p})}{f(x_i; \mathbf{p})} + 2k \partial_{p_j} \int f(x; \mathbf{p}) dx. \end{aligned} \quad (24)$$

The last term vanishes in the limit, because $\int f(x; \mathbf{p}) dx = 1$ is constant.

We finally obtain the equivalence

$$\partial_{p_j} Q(\mathbf{p}) \xrightarrow{\Delta x \rightarrow 0, \hat{\mathbf{p}} \rightarrow \mathbf{p}} -2 \sum_{i=1}^k \frac{\partial_{p_j} f(x_i; \mathbf{p})}{f(x_i; \mathbf{p})} = -2 \sum_{i=1}^k \partial_{p_j} \ln f(x_i; \mathbf{p}) = -2 \partial_{p_j} \ln L(\mathbf{p}). \quad (25)$$

The derivatives are equal up to a constant factor, which means that the solutions of $\partial_{p_j} Q(\mathbf{p}) = 0$ and $\partial_{p_j} \ln L(\mathbf{p}) = 0$ are equal. In other words, the IWLS solution in the limit of infinitesimal bins is found by minimizing the negative log-likelihood of the probability density. The latter is effectively a shortcut to the solution, which does not require iterations.

We showed the equivalence for the case when measurements consist of a single variable x_i per event for simplicity, but it also holds for the general case

of a set of n n -dimensional vectors $\{\mathbf{x}_i | i = 1 \dots k\}$ with $\mathbf{x}_i = (x_{ji} | j = 1 \dots n)$ and a corresponding n -dimensional probability density $f(\mathbf{x}; \mathbf{p})$. In this case, one would repeat the derivation starting from an n -dimensional histogram.

The derivation provides some insights.

- The absolute values of $Q(\hat{\mathbf{p}})$ and $-2 \ln L(\hat{\mathbf{p}})$ at the solution $\hat{\mathbf{p}}$ are not equal. They differ by an (infinite) additive constant.
- The derivatives $\partial_{p_j} Q(\mathbf{p})$ and $-2 \partial_{p_j} \ln L(\mathbf{p})$ differ in general when \mathbf{p} is not the solution $\hat{\mathbf{p}}$, because the second term in Eq. (24) does not vanish for $\mathbf{p} \neq \hat{\mathbf{p}}$.

In practice, the second point means that the MINUIT package produces the same error estimates for the solution $\hat{\mathbf{p}}$ if the HESSE algorithm is used, but not if the MINOS algorithm is used. The HESSE algorithm numerically computes and inverts the Hessian matrix of second derivatives at the minimum, which gives identical results for Q and $-2 \ln L$. The MINOS algorithm scans the neighborhood of the minimum, which for Q and $-2 \ln L$ usually has a different shape.

5. Notes on goodness-of-fit tests

For a goodness-of-fit (GoF) test, one computes a test statistic for a probabilistic model and a set of measurements. The test statistic is designed to have a known probability distribution when the measurements are truly distributed according to the probabilistic model. If the value for a particular model is very improbable, the model may be rejected.

It is well-known that the minimum value $Q(\hat{\mathbf{p}})$ is χ^2 -distributed with expectation $(k - m)$, if the measurements are normally distributed, where k and m are the number of measurements and number of fitted parameters, respectively [3]. This GoF property is so useful and frequently applied, that the function $Q(\mathbf{p})$ is often simply called *chi-square*.

In general, $Q(\hat{\mathbf{p}})$ is not χ^2 -distributed for measurements that are not normally distributed around the model expectations. Approximately, it holds for

Poisson and binomially distributed measurements when counts are not close to zero, and fractions are neither too close to zero or one. Stronger statements can be made about the expectation value of $Q(\hat{\boldsymbol{p}})$. For linear models with m parameters and k unbiased measurements with known covariance matrix \boldsymbol{C} , the expectation of $Q(\hat{\boldsymbol{p}})$ is guaranteed to be

$$\text{E}[Q(\hat{\boldsymbol{p}})] = k - m, \quad (26)$$

regardless of the sample size and the distribution of the measurements, as shown in the appendix. Therefore, the well-known quality criterion that the *reduced* χ^2 should be close to unity, $Q(\hat{\boldsymbol{p}})/(k - m) \simeq 1$, is often useful even if measurements are not normally distributed.

We saw previously that $-2 \ln L(\hat{\boldsymbol{p}})$ differs from $Q(\hat{\boldsymbol{p}})$ by an infinite additive constant, which is a hint that it cannot straight-forwardly replace the latter as a GoF statistic. When used with unbinned data, $\ln L(\hat{\boldsymbol{p}})$ is ill-suited as a GoF test statistic. Heinrich [13] presented striking examples when $\ln L(\hat{\boldsymbol{p}})$ carries no information of how well the model fits the measurements. Cousins [14, 15] gave an intuitive explanation for this fact. The IWLS fit provides a maximum-likelihood estimate for measurements that follow a Poisson or binomial distribution and a GoF test statistic as a side result, which in general is not exactly χ^2 -distributed, but its distribution can often be obtained from a Monte Carlo simulation.

6. Conclusions

An iterated weighted least-squares fit applied to measurements, which are Poisson- or binomially distributed around model expectations, provides the exact same solution as a maximum-likelihood fit. This holds for any model and any sample size. When the two fit methods are equivalent, the maximum-likelihood fit is still recommended, except when the model is linear. In this case, the minimum of the weighted least-squares problem can be found analytically in each iteration, which usually needs less computations overall than numerically maximizing the likelihood and requires no problem-specific starting point. The iterated weighted least-squares fit provides a goodness-of-fit statistic in addition,

while the maximum-likelihood fit usually does not. Of course, a goodness-of-fit statistic can always be separately computed after the optimization, but in case of the maximum-likelihood it requires implementing two functions in a computer program instead of one.

Whether the two fit methods give equivalent results depends only on the probability distribution of the measurements around the model expectations. Here we presented proofs of the equivalence for Poisson and binomial distributions. In the statistics literature [1, 2], more general proofs are given that hold also for some other distributions.

7. Acknowledgments

We thank Bob Cousins for a critical reading of the manuscript and for valuable pointers to the primary statistical literature, and the anonymous reviewer for suggestions to clarify additional points.

Appendix A. Equivalence of ML and ILWS for Poisson-distributed data

A common task is to fit a model to a histogram with k bins, each with a count n_i . Especially in multi-dimensional histograms some bins may have few or even zero entries. This poses a problem for a conventional weighted least-squares fit, but not for a ML fit or an IWLS fit.

A ML fit of a model with m parameters $\mathbf{p} = (p_j | j = 1, \dots, m)$ to a sample of k Poisson-distributed numbers $\{n_i | i = 1, \dots, k\}$ with expectation values $E[n_i] = \mu_i(\mathbf{p})$ is performed by maximizing the log-likelihood

$$\ln L(\mathbf{p}) = \sum_{i=1}^k n_i \ln \mu_i - \sum_{i=1}^k \mu_i, \quad (\text{A.1})$$

which is obtained by taking the logarithm of the product of Poisson probabilities (9) of the data under the model, and dropping terms that do not depend on \mathbf{p} .

To find the maximum, we set the m first derivatives

$$\frac{\partial \ln L}{\partial p_j} = \sum_{i=1}^k \frac{n_i}{\mu_i} \frac{\partial \mu_i}{\partial p_j} - \sum_{i=1}^k \frac{\partial \mu_i}{\partial p_j} \quad (\text{A.2})$$

for $j = 1$ to m to 0. We get a system of equations

$$\sum_{i=1}^k \frac{n_i - \mu_i}{\mu_i} \frac{\partial \mu_i}{\partial p_j} = 0. \quad (\text{A.3})$$

We now approach the same problem as an IWLS fit. The sum of weighted squared residuals is

$$Q = \sum_{i=1}^k \frac{(n_i - \mu_i)^2}{\hat{\mu}_i}, \quad (\text{A.4})$$

where $\hat{\mu}_i$ is the expected variance computed from the model, using the parameter estimate $\hat{\mathbf{p}}$ from the previous iteration. To find the minimum, we again set the m first derivatives to 0 and obtain

$$\frac{\partial Q}{\partial p_j} = -2 \sum_{i=1}^k \frac{n_i - \mu_i}{\hat{\mu}_i} \frac{\partial \mu_i}{\partial p_j} = 0. \quad (\text{A.5})$$

Eq. (A.5) and Eq. (A.3) yield identical solutions in the limit $\hat{\mu}_i \rightarrow \mu_i$, and so do their solutions. The limit is approached by iterating the fit, so that we actually obtain the maximum-likelihood estimate from the IWLS fit. Remarkably, this does not depend on the size of the counts n_i per bin. The equivalence holds even when many bins with zero entries are present. To obtain this result, the $\hat{\mu}_i$ must be constant. If $\hat{\mu}_i$ was replaced by μ_i in Eq. (A.4), extra non-vanishing terms would appear in Eq. (A.5).

As already mentioned, when the expectations are linear functions, the unique analytical solution to Eq. (A.5) is given by Eq. (4), with $\mathbf{C}^{-1} = (\delta_{ij} / \hat{\mu}_i | i, j = 1, \dots, k)$ and $y_i = n_i$. An analytical solution of Eq. (A.3) is not known to the authors. The IWLS fit converges faster than the ML fit in this case.

Appendix B. Equivalence of ML and IWLS for binomially distributed data

Another common task is to obtain an efficiency function of a selection or trigger as a function of an observable. One collects a histogram of generated events with bin contents N_i , and a corresponding histogram of accepted events with bin contents n_i . The N_i are considered as constants here, while the n_i are drawn from the binomial distribution. The goal is to obtain a model function

that best describes the efficiencies ϵ_i that best describe the drawn samples n_i . A single least-squares fit will give biased results when many n_i are close to either 0 or N_i , but not a ML or an IWLS fit.

A ML fit of a model with m parameters $\mathbf{p} = (p_j | j = 1, \dots, m)$ for a sample of k binomially distributed numbers $\{n_i | i = 1, \dots, k\}$ with expectations $E[n_i] = \mu_i(\mathbf{p}) = \epsilon_i(\mathbf{p}) N_i$ is performed by maximizing the log-likelihood

$$\ln L(\mathbf{p}) = \sum_{i=1}^k n_i \ln \mu_i + \sum_{i=1}^k (N_i - n_i) \ln(N_i - \mu_i), \quad (\text{B.1})$$

which is obtained by taking the logarithm of the product of binomial probabilities to observe n_i when $\mu_i = \epsilon_i N_i$ are expected,

$$P(n_i; \mu_i, N_i) = \binom{N_i}{n_i} \epsilon_i^{n_i} (1 - \epsilon_i)^{N_i - n_i} = \binom{N_i}{n_i} \frac{\mu_i^{n_i} (N_i - \mu_i)^{N_i - n_i}}{N_i^{N_i}}, \quad (\text{B.2})$$

and dropping terms that do not depend on \mathbf{p} . A binomial distribution has two parameters (μ_i, N_i) , but it is a monoparametric distribution in this context since the N_i are known and only the μ_i are free parameters.

Again we set the m first derivatives

$$\frac{\partial \ln L}{\partial p_j} = \sum_{i=1}^k \frac{n_i}{\mu_i} \frac{\partial \mu_i}{\partial p_j} - \sum_{i=1}^k \frac{N_i - n_i}{N_i - \mu_i} \frac{\partial \mu_i}{\partial p_j} \quad (\text{B.3})$$

to zero for $j = 1$ to m . The minimum is obtained by solving

$$\sum_{i=1}^k \frac{n_i - \mu_i}{\mu_i(1 - \mu_i/N_i)} \frac{\partial \mu_i}{\partial p_j} = 0. \quad (\text{B.4})$$

For the IWLS fit, we need to minimize the sum

$$Q(\mathbf{p}) = \sum_{i=1}^k \frac{(n_i - \mu_i)^2}{\hat{\mu}_i(1 - \hat{\mu}_i/N_i)}. \quad (\text{B.5})$$

where the variances for the binomial distribution with expectation μ_i are $\sigma_i^2 = \text{Var}[n_i] = N_i \epsilon_i (1 - \epsilon_i) = \mu_i (1 - \mu_i/N_i)$. Again, we replaced μ_i in the variance by the constant estimate $\hat{\mu}_i$ from the previous iteration. Setting the m first derivatives to 0, we obtain

$$\frac{\partial Q}{\partial p_j} = -2 \sum_{i=1}^k \frac{n_i - \mu_i}{\hat{\mu}_i(1 - \hat{\mu}_i/N_i)} \frac{\partial \mu_i}{\partial p_j} = 0 \quad (\text{B.6})$$

Like in the previous case, Eq. (B.4) and Eq. (B.6) yield identical solutions in the limit $\hat{\mu}_i \rightarrow \mu_i$, which is approached by iterating the minimization. Again, we obtain the maximum-likelihood estimate with the IWLS fit.

Like in the previous case, the L-IWLS fit for a linear model converges faster than the ML fit, while the IWLS fit converges more slowly than the ML fit in general.

Appendix C. Expectation of $Q(\hat{\mathbf{p}})$ for linear models

We compute the expectation of Q in Eq. (2), evaluated at the solution $\hat{\mathbf{p}}$ from Eq. (4) for linear models with $E[\mathbf{y}] = \mathbf{X}\mathbf{p}$, where \mathbf{X} is a fixed $k \times m$ matrix, and where the measurements \mathbf{y} have a known finite covariance matrix \mathbf{C} . Similar proofs are found in the literature [16]. The covariance matrix of $\hat{\mathbf{p}}$ is obtained by error propagation with the matrix $\mathbf{M} = (\mathbf{X}^T \mathbf{C}^{-1} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{C}^{-1}$ and $\hat{\mathbf{p}} = \mathbf{M}\mathbf{y}$ as

$$\mathbf{C}_p = \mathbf{M}\mathbf{C}\mathbf{M}^T = (\mathbf{X}^T \mathbf{C}^{-1} \mathbf{X})^{-1}, \quad (\text{C.1})$$

where we used that \mathbf{C}^{-1} and $(\mathbf{X}^T \mathbf{C}^{-1} \mathbf{X})^{-1}$ are symmetric matrices.

The expectation is a linear operator. Since the solution $\hat{\mathbf{p}} = \mathbf{M}\mathbf{y}$ is a linear function of the measurement, we have

$$E[\hat{\mathbf{p}}] = \mathbf{M}E[\mathbf{y}] = \mathbf{M}\mathbf{X}\mathbf{p} = \mathbf{p}, \quad (\text{C.2})$$

in other words, $\hat{\mathbf{p}}$ is an unbiased estimate of \mathbf{p} .

We expand Q evaluated at $\hat{\mathbf{p}}$,

$$Q(\hat{\mathbf{p}}) = \mathbf{y}^T \mathbf{C}^{-1} \mathbf{y} - \mathbf{y}^T \mathbf{C}^{-1} \mathbf{X} \hat{\mathbf{p}} - \hat{\mathbf{p}}^T \mathbf{X}^T \mathbf{C}^{-1} \mathbf{y} + \hat{\mathbf{p}}^T \mathbf{X}^T \mathbf{C}^{-1} \mathbf{X} \hat{\mathbf{p}}, \quad (\text{C.3})$$

which simplifies with $\mathbf{C}_p^{-1} \hat{\mathbf{p}} = \mathbf{X}^T \mathbf{C}^{-1} \mathbf{y}$ and Eq. (C.1) to

$$Q(\hat{\mathbf{p}}) = \mathbf{y}^T \mathbf{C}^{-1} \mathbf{y} - \hat{\mathbf{p}}^T \mathbf{C}_p^{-1} \hat{\mathbf{p}}. \quad (\text{C.4})$$

The scalar result of a bilinear form is trivially equal to the trace of this bilinear form, and a cyclic permutation inside the trace then yields

$$Q(\hat{\mathbf{p}}) = \text{Tr}(\mathbf{C}^{-1} \mathbf{y}\mathbf{y}^T) - \text{Tr}(\mathbf{C}_p^{-1} \hat{\mathbf{p}}\hat{\mathbf{p}}^T). \quad (\text{C.5})$$

We compute the expectation on both sides and get, using linearity of trace and expectation,

$$\mathbb{E}[Q(\hat{\mathbf{p}})] = \text{Tr}(\mathbf{C}^{-1} \mathbb{E}[\mathbf{y}\mathbf{y}^T]) - \text{Tr}(\mathbf{C}_p^{-1} \mathbb{E}[\hat{\mathbf{p}}\hat{\mathbf{p}}^T]). \quad (\text{C.6})$$

The definition of the covariance matrix $\mathbf{C} = \mathbb{E}[\mathbf{y}\mathbf{y}^T] - \mathbb{E}[\mathbf{y}]\mathbb{E}[\mathbf{y}]^T$ is inserted, and vice versa for \mathbf{C}_p . We get

$$\mathbb{E}[Q(\hat{\mathbf{p}})] = \text{Tr}(\mathbf{C}^{-1}\mathbf{C} + \mathbf{C}^{-1}\mathbb{E}[\mathbf{y}]\mathbb{E}[\mathbf{y}]^T) - \text{Tr}(\mathbf{C}_p^{-1}\mathbf{C}_p + \mathbf{C}_p^{-1}\mathbb{E}[\hat{\mathbf{p}}]\mathbb{E}[\hat{\mathbf{p}}]^T). \quad (\text{C.7})$$

The trace of a matrix multiplied with its inverse is equal to the number of diagonal elements, which is k in case of \mathbf{C} and m in case of \mathbf{C}_p . We use this, $\mathbb{E}[\mathbf{y}] = \mathbf{X}\mathbf{p}$, $\mathbb{E}[\hat{\mathbf{p}}] = \mathbf{p}$, and again the linearity of the trace, to get

$$\mathbb{E}[Q(\hat{\mathbf{p}})] = k + \text{Tr}(\mathbf{C}^{-1}\mathbf{X}\mathbf{p}\mathbf{p}^T\mathbf{X}^T) - (m + \text{Tr}(\mathbf{C}_p^{-1}\mathbf{p}\mathbf{p}^T)). \quad (\text{C.8})$$

The remaining traces are identical and cancel,

$$\text{Tr}(\mathbf{C}^{-1}\mathbf{X}\mathbf{p}\mathbf{p}^T\mathbf{X}^T) = \text{Tr}(\mathbf{X}^T\mathbf{C}^{-1}\mathbf{X}\mathbf{p}\mathbf{p}^T) = \text{Tr}(\mathbf{C}_p^{-1}\mathbf{p}\mathbf{p}^T), \quad (\text{C.9})$$

and so we finally obtain the result

$$\mathbb{E}[Q(\hat{\mathbf{p}})] = k - m, \quad (\text{C.10})$$

which is independent of the PDFs that describe the scatter of the measurements \mathbf{y} around the expectation values $\mathbb{E}[\mathbf{y}]$.

References

- [1] J. A. Nelder, R. W. M. Wedderburn, *Generalized Linear Models*, J. R. Statist. Soc. **A135**, 370–384 (1972).
- [2] A. Charles, E.L. Frome, P. L. Yu, *The Equivalence of Generalized Least Squares and Maximum Likelihood Estimates in the Exponential Family*, J. Am. Stat. Assoc. **71**, 169–171 (1976).
- [3] F. James, *Statistical methods in experimental physics*, World Scientific Publishing Company, 2006.

- [4] S. Baker and R. D. Cousins, *Clarification of the use of chi-square and likelihood functions in fits to histograms*, Nuclear Instruments and Methods in Physics Research **221**, 437 (1984).
- [5] A. C. Aitken, *IV. On Least Squares and Linear Combination of Observations*, Proc. R. Soc. Edinburgh **55**, 42 (1935).
- [6] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap*, Monographs on Statistics and Applied Probability No. 57, Chapman & Hall/CRC, Boca Raton, Florida, USA, 1993.
- [7] R. J. Barlow, *Combining experiments with systematic errors*, arXiv.1701.03701 (2017).
- [8] T.A. Aaltonen *et al.* (CDF, D0 collaborations), *Combination of measurements of the top-quark pair production cross section from the Tevatron Collider*, Phys. Rev. **D89**, 072001 (2014).
- [9] F. James and M. Roos, *Minuit - a system for function minimization and analysis of the parameter errors and correlations*, Computer Physics Communications **10**, 343 (1975).
- [10] iminuit team, *MINUIT from Python*, <https://github.com/iminuit/iminuit> (2013), accessed: 2018-03-05.
- [11] C. Lawson and R. Hanson, *Solving Least Squares Problems*, Society for Industrial and Applied Mathematics, 1995, <https://doi.org/10.1137/1.9781611971217>.
- [12] E. Jones *et al.*, *SciPy: Open source scientific tools for Python*, <http://www.scipy.org> (2001), accessed: 2018-03-05.
- [13] J. Heinrich, *Pitfalls of Goodness-of-Fit from Likelihood*, in *Statistical Problems in Particle Physics, Astrophysics, and Cosmology*, edited by L. Lyons, R. Mount, and R. Reitmeyer, p. 52, 2003, arXiv:physics/0310167.

- [14] R.D. Cousins, *Generalization of chisquare goodness-of-fit test for binned data using saturated models, with application to histograms*, http://www.physics.ucla.edu/~cousins/stats/cousins_saturated.pdf (2013), accessed: 2018-07-19.
- [15] R.D. Cousins, *On Goodness-of-Fit Tests*, <http://cousins.web.cern.ch/cousins/ongoodness6march2016.pdf> (2016), accessed: 2018-07-19.
- [16] see, e.g., M. G. Kendall, A. Stuart, *The Advanced Theory of Statistics*, vol 2, Sect. 19.9, 3rd edition, Charles Griffin & Co., 1961.