

*Statistical Challenges for Searches for New Physics
at the LHC*

Kyle Cranmer

*Presented at the PhyStat2005 Conference
Oxford College, United Kingdom
September 12-15, 2005*

September 2005

Physics Department

Brookhaven National Laboratory

P.O. Box 5000
Upton, NY 11973-5000
www.bnl.gov

Notice: This manuscript has been authored by employees of Brookhaven Science Associates, LLC under Contract No. DE-AC02-98CH10886 with the U.S. Department of Energy. The publisher by accepting the manuscript for publication acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

This preprint is intended for publication in a journal or proceedings. Since changes may be made before publication, it may not be cited or reproduced without the author's permission.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.



STATISTICAL CHALLENGES FOR SEARCHES FOR NEW PHYSICS AT THE LHC

KYLE CRANMER

Brookhaven National Laboratory, Upton, NY 11973, USA

e-mail: Kyle.Cranmer@cern.ch

Because the emphasis of the LHC is on 5σ discoveries and the LHC environment induces high systematic errors, many of the common statistical procedures used in High Energy Physics are not adequate. I review the basic ingredients of LHC searches, the sources of systematics, and the performance of several methods. Finally, I indicate the methods that seem most promising for the LHC and areas that are in need of further study.

1 Introduction

The Large Hadron Collider (LHC) at CERN and the two multipurpose detectors, ATLAS and CMS, have been built in order to discover the Higgs boson, if it exists, and explore the theoretical landscape beyond the Standard Model.^{1,2} The LHC will collide protons with unprecedented center-of-mass energy ($\sqrt{s} = 14$ TeV) and luminosity (10^{34} cm⁻²s⁻¹); the ATLAS and CMS detectors will record these interactions with $\sim 10^8$ individual electronic readouts per event. Because the emphasis of the physics program is on discovery and the experimental environment is so complex, the LHC poses new challenges to our statistical methods – challenges we must meet with the same vigor that led to the theoretical and experimental advancements of the last decade.

In the remainder of this Section, I introduce the physics goals of the LHC and most pertinent factors that complicate data analysis. I also review the formal link and the practical differences between confidence intervals and hypothesis testing.

In Sec. 2, the primary ingredients to new particle searches are discussed. Practical and toy examples are presented in Sec. 3, which will be used to assess the most common methods in Sec. 4. The remainder of this paper is devoted to discussion on the most promising methods for the LHC.

1.1 Physics Goals of the LHC

Currently, our best experimentally justified model for fundamental particles and their interactions is the *standard model*. In short, the physics goals of the LHC come in two types: those that improve our understanding of the standard model, and those that go beyond it.

The only particle of the standard model that has not been observed is the Higgs boson, which is key for the standard model's description of the electroweak interactions. The mass of the Higgs boson, m_H , is a free parameter in the standard model, but there exist direct experimental lower bounds and more indirect upper bounds. Once m_H is fixed, the standard model is a completely predictive theory. There are numerous particle-level Monte Carlo generators that can be interfaced with simulations of the detectors to predict the rate and distribution of all experimental observables. Because of this predictive power, searches for the Higgs boson are highly tuned and often employ multivariate discrimination methods like neural networks, boosted decision trees, support vector machines, and genetic programming.^{3,4,5}

While the Higgs boson is key for understanding the electroweak interactions, it introduces a new problem: *i.e.* the *hierarchy problem*. There are several proposed solutions to the problem, one of which is to introduce a new fundamental symmetry, called supersymmetry (SUSY), between bosons and fermions. In practice, the minimal supersymmetric extension to the standard model (MSSM), with its 105 parameters, is not so much a theory as a theoretical framework.

The key difference between SUSY and Higgs searches is that, in most cases, discovering SUSY will not be the difficult part. Searches for SUSY often rely on robust signatures that will show a deviation from the standard model for most regions of the SUSY parameter space. It will be much more challenging to demonstrate that the deviation from the standard model is SUSY and to measure the fundamental parameters of the theory.⁶ In order to restrict the scope of these proceedings, I shall focus LHC Higgs searches, where the issues of hypothesis testing are more relevant.

1.2 The Challenges of LHC Environment

The challenges of the LHC environment are manifold. The first and most obvious challenge is due to the enormous rate of uninteresting background events from QCD processes. The total interaction rate for the LHC is of order 10^9 interactions per second; the rate of Higgs production is about ten orders of magnitude smaller. Thus, to understand the background of a Higgs search, one must understand the extreme tails of the QCD processes.

Compounding the difficulties due to the extreme rate is the complexity of the detectors. The full-fledged simulation of the detectors is extremely computationally intensive, with samples of 10^7 events taking about a month to produce with computing resources distributed around the globe. This computational limitation constrains the problems that can be addressed with Monte Carlo techniques.

Theoretical uncertainties also contribute to the challenge. The background to many searches requires calculations at, or just beyond, the state-of-the-art in particle physics. The most common situation requires a final state with several well-separated high transverse momentum objects (e.g. $t\bar{t}jj \rightarrow b\nu\bar{b}jjjj$), in which the regions of physical interest are not reliably described by leading-order perturbative calculations (due to infra-red and collinear divergences), are too complex for the requisite next-to-next-to-leading order calculations, and are not properly described by the parton-shower models alone. Enormous effort has gone into improving the situation with next-to-leading order calculations and matrix-element-parton-shower matching.^{7,8} While these new tools are a vast improvement, the residual uncertainties are still often dominant.

Uncertainties from non-perturbative effects are also important. For some processes, the relevant regions of the parton distribution functions are not well-measured (and probably will not be in the first few years of LHC running), which lead to uncertainties in rate as well as the shape of distributions. Furthermore, the various underlying-event and multiple-interaction models used to describe data from previous colliders show large deviations when extrapolated to the LHC.⁹ This soft physics has a large impact on the performance of observables such as missing transverse energy.

In order to augment the simulated data chain,

most searches introduce auxiliary measurements to estimate their backgrounds from the data itself. In some cases, the background estimation is a simple sideband, but in others the link between the auxiliary measurement to the quantity of interest is based on simulation. This hybrid approach is of particular importance at the LHC.

While many of the issues discussed above are not unique to the LHC, they are often more severe. At LEP, it was possible to generate Monte Carlo samples of larger size than the collected data, QCD backgrounds were more tame, and most searches were not systematics-limited. The Tevatron has much more in common with the LHC; however, at this point discovery is less likely, and most of the emphasis is on measurements and limit setting.

1.3 Confidence Intervals & Hypothesis Testing

The last several conferences in the vein of *PhyStat 2005* have concentrated heavily on confidence intervals. In particular, 95% confidence intervals for some physics parameter in an experiment that typically has few events. More recently, there has been a large effort in understanding how to include systematic errors and nuisance parameters into these calculations.

LHC searches, in contrast, are primarily interested in 5σ discovery. The 5σ discovery criterion is somewhat vague, but usually interpreted in a frequentist sense as a hypothesis test with a rate of Type I error $\alpha = 2.85 \cdot 10^{-7}$.

There is a formal link between confidence intervals and hypothesis testing: frequentist confidence intervals from the Neyman construction are formally inverted hypothesis tests. It is this equivalence that links the Neyman-Pearson lemma^a to the ordering rule used in the *unified method* of Feldman and Cousins.¹⁰ Furthermore, this equivalence will be very useful in translating our understanding of confidence intervals to the searches at the LHC.

In some cases, this formal link can be misleading. In particular, there is not always a continuous parameter that links the fully specified null hypothesis H_0 to the fully specified alternate H_1 in any physically interesting or justified way. Furthermore,

^aThe lemma states that, for a simple hypothesis test of size α between a null H_0 and an alternate H_1 , the most powerful critical region in the observable x is given by a contour of the likelihood ratio $L(x|H_0)/L(x|H_1)$.

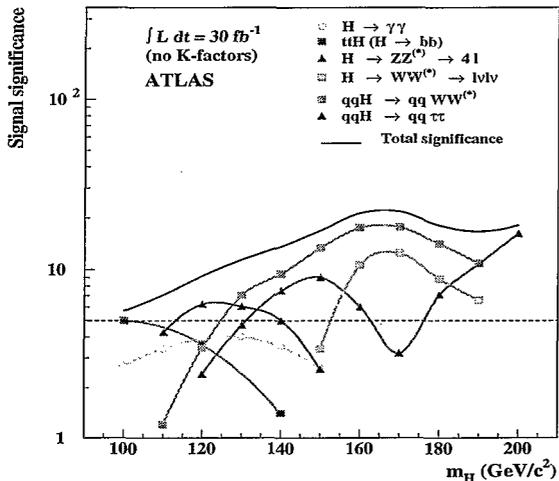


Figure 1. Expected significance as a function of Higgs mass for the ATLAS detector with 30 fb^{-1} of data.

the performance of a method for a 95% confidence interval and a 5σ discovery can be quite different.

2 The Ingredients of an LHC Search

In order to assess the statistical methods that are available and develop new ones suited for the LHC, it is necessary to be familiar with the basic ingredients of the search. In this section, the basic ingredients, terminology, and nomenclature are established.

2.1 Multiple Channels & Processes

Almost all new particle searches do not observe the particle directly, but through the signatures left by the decay products of the particle. For instance, the Higgs boson will decay long before it interacts with the detector, but its decay products will be detected. In many cases, the particle can be produced and decay in many different configurations, each of which is called a *search channel* (see Tab. 1). There are may be multiple signal and background processes which contribute to each channel. For example, in $H \rightarrow \gamma\gamma$, the signal could come from any Higgs production mechanism and the background from either continuum $\gamma\gamma$ production or QCD backgrounds where jets fake photons. Each of these processes have their own rates, distributions for observables, and uncertainties. Furthermore, the uncertainties between processes may be correlated.

In general the theoretical model for a new parti-

cle has some free parameters. In the case of the standard model Higgs, only the mass m_H is unknown. For SUSY scenarios, the Higgs model is parametrized by two parameters: m_A and $\tan\beta$. Typically, the unknown variables are scanned and a hypothesis test is performed for each value of these parameters. The results from each of the search channels can be combined to enhance the power of the search, but one must take care of correlations among channels and ensure consistency.

The fact that one scans over the parameters and performs many hypothesis tests increases the chance that one finds at least one large fluctuation from the null-hypothesis. Some approaches incorporate the number of trials explicitly,¹¹ some approaches only focus on the most interesting fluctuation,¹² and some see this heightened rate of Type I error as the motivation for the stringent 5σ requirement.¹³

2.2 Discriminating Variables & Test Statistics

Typically, new particles are known to decay with certain characteristics that distinguish the *signal* events from those produced by *background* processes. Much of the work of a search is to identify those observables and to construct new discriminating variables (generically denoted as m). Examples include angles between particles, invariant masses, and particle identification criterion. Discriminating variables are used in two different ways: to define a signal-like region and to weight events.

The usage of discriminating variables is related to the test statistic: the real-valued quantity used to summarize the experiment. The test statistic is thought of as being ordered such that either large or small values indicate growing disagreement with the null hypothesis.

A simple “cut analysis” consists of defining a signal-like region bounded by upper- and lower-values of these discriminating variables and counting events in that region. In that case, the test statistic is simply the number of events observed in the signal like region. One expects b background events and s signal events, so the experimental sensitivity is optimized by adjusting the cut values. More sophisticated techniques use multivariate algorithms, such as neural networks, to define more complicated signal-like regions, but the test statistic remains unchanged. In these number counting analyses, the likelihood of

observing n events is simply given by the Poisson model.

There are extensions to this number-counting technique. In particular, if one knows the distribution of the discriminating variable x for the background-only (null) hypothesis, $f_b(m)$, and the signal-plus-background (alternate) hypothesis, $f_{s+b}(m) = [sf_s(m) + bf_b(m)]/(s + b)$, then there is a more powerful test statistic than simply counting events. This is intuitive, a well measured 'golden event' is often more convincing than a few messy ones. Following the Neyman-Pearson lemma, the most powerful test statistic is

$$Q = \frac{L(\mathbf{m}|H_1)}{L(\mathbf{m}|H_0)} \quad (1)$$

$$= \frac{\prod_i^{N_{chan}} Pois(n_i|s_i + b_i) \prod_j^{n_i} \frac{s_i f_s(m_{ij}) + b_i f_b(m_{ij})}{s_i + b_i}}{\prod_i^{N_{chan}} Pois(n_i|b_i) \prod_j^{n_i} f_b(m_{ij})}$$

or equivalently

$$q = \ln Q = -s_{tot} + \sum_i^{N_{chan}} \sum_j^{n_i} \ln \left(1 + \frac{s_i f_s(m_{ij})}{b_i f_b(m_{ij})} \right). \quad (2)$$

The test statistic in Eq. 2 was used by the LEP Higgs Working Group (LHWG) in their final results on the search for the Standard Model Higgs.¹⁴

At this point, there are two loose ends: how does one determine the distribution of the discriminating variables $f(m)$, and how does one go from Eq. 2 to the distribution of q for H_0 and H_1 . These are the topics of the next subsections.

2.3 Parametric and Non-Parametric Methods

In some cases, the distribution of a discriminating variable $f(m)$ can be parametrized and this parametrization can be justified either by physics arguments or by goodness-of-fit. However, there are many cases in which $f(m)$ has a complicated shape not easily parametrized. For instance, Fig. 2 shows the distribution of a neural network output for signal events. In that case kernel estimation techniques can be used to estimate $f(m)$ in a non-parametric way from a sample of events $\{m_i\}$.¹⁵ The technique that was used by the LHWG¹⁴ was based on an adaptive kernel estimation given by:

$$\hat{f}_1(m) = \sum_i^n \frac{1}{nh(m_i)} K \left(\frac{m - m_i}{h(m_i)} \right), \quad (3)$$

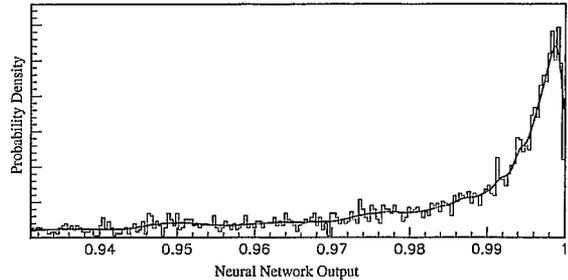


Figure 2. The distribution of a neural network output for signal events. The histogram is shown together with $\hat{f}_1(m)$.

where

$$h(m_i) = \left(\frac{4}{3} \right)^{1/5} \sqrt{\frac{\sigma}{\hat{f}_0(m_i)}} n^{-1/5}, \quad (4)$$

σ is the standard deviation of $\{x_i\}$, $K(x)$ is some kernel function (usually the normal distribution), and $\hat{f}_0(x)$ is the fixed kernel estimate given by the same equation but with a fixed $h(m_i)$

$$h^* = \left(\frac{4}{3} \right)^{1/5} \sigma n^{-1/5}. \quad (5)$$

The solid line in Fig. 2 shows that the method (with modified-boundary kernels) works very well for shapes with complicated structure at many scales.

2.4 The Fourier Transform Technique

Given, $f_s(m)$ and $f_b(m)$ the distribution of $q(x)$ can be constructed. For the background-only hypothesis, $f_b(m)$ provides the probability of corresponding values of q needed to define the single-event pdf ρ_1 .^b

$$\rho_{1,b}(q_0) = \int f_b(m) \delta(q(m) - q_0) dm \quad (6)$$

For multiple events, the distribution of the log-likelihood ratio must be obtained from repeated convolutions of the single event distribution. This convolution can either be performed implicitly with approximate Monte Carlo techniques,¹⁶ or analytically with a Fourier transform technique.¹⁷ In the Fourier domain, denoted with a bar, the distribution of the log-likelihood for n events is

$$\overline{\rho}_n = \overline{\rho}_1^n \quad (7)$$

Thus the expected log-likelihood distribution for background with Poisson fluctuations in the number

^bThe integral is necessary because the map $q(m) : m \rightarrow q$ may be many-to-one.

of events takes the form

$$\rho_b(q) = \sum_{n=0}^{\infty} \frac{e^{-b} b^n}{n!} \rho_{n,b}(q) \quad (8)$$

which in the Fourier domain is simply

$$\overline{\rho_b(q)} = e^{b[\overline{\rho_{1,b}(q)}-1]}. \quad (9)$$

For the signal-plus-background hypothesis we expect s events from the $\rho_{1,s}$ distribution and b events from the $\rho_{1,b}$ distribution, which leads to the expression for ρ_{s+b} in the Fourier domain^c

$$\overline{\rho_{s+b}(q)} = e^{b[\overline{\rho_{1,b}(q)}-1]+s[\overline{\rho_{1,s}(q)}-1]}. \quad (10)$$

This equation generalizes, in a somewhat obvious way, to include many processes and channels.

Numerically these computations are carried out with the Fast Fourier Transform (FFT). The FFT is performed on a finite and discrete array, beyond which the function is considered to be periodic. Thus the range of the ρ_1 distributions must be sufficiently large to hold the resulting ρ_b and ρ_{s+b} distributions. If they are not, the “spill over” beyond the maximum log-likelihood ratio q_{max} will “wrap around” leading to unphysical ρ distributions. Because the range of ρ_b is much larger than $\rho_{1,b}$ it requires a very large number of samples to describe both distributions simultaneously. The implementation of this method requires some approximate asymptotic techniques that describe the scaling from $\rho_{1,b}$ to ρ_b .¹⁸

The nature of the FFT results in a number of round-off errors and limit the numerical precision to about 10^{-16} – which limit the method to significance levels below about 8σ . Extrapolation techniques and arbitrary precision calculations can overcome these difficulties,¹⁸ but such small p -values are of practical little interest.

From the log-likelihood distribution of the two hypotheses we can calculate a number of useful quantities. Given some experiment with an observed log-likelihood ratio, q^* , we can calculate the background-only confidence level, CL_b :

$$CL_b(q^*) = \int_{q^*}^{\infty} \rho_b(q') dq' \quad (11)$$

^cPerhaps it is worth noting that $\overline{\rho(q)}$ is a complex valued function of the Fourier conjugate variable of q . Thus numerically the exponentiation in Eq. 9 requires Euler’s formula $e^{i\theta} = \cos \theta + i \sin \theta$.

In the absence of an observation we can calculate the expected CL_b given the signal-plus-background hypothesis is true. To do this we first must find the median of the signal-plus-background distribution \overline{q}_{s+b} . From these we can calculate the expected CL_b by using Eq. 11 evaluated at $q^* = \overline{q}_{s+b}$.

Finally, we can convert the expected background confidence level into an expected Gaussian significance, $Z\sigma$, by finding the value of Z which satisfies

$$CL_b(\overline{q}_{s+b}) = \frac{1 - \text{erf}(Z/\sqrt{2})}{2}. \quad (12)$$

where $\text{erf}(Z) = (2/\pi) \int_0^Z \exp(-y^2) dy$ is a function readily available in most numerical libraries. For $Z > 1.5$, the relationship can be approximated¹⁹ as

$$Z \approx \sqrt{u - \ln u} \quad \text{with} \quad u = -2 \ln(CL_b \sqrt{2\pi}) \quad (13)$$

2.5 Systematic Errors, Nuisance Parameters & Auxiliary Measurements

Sections 2.3 and 2.4 represent the state of the art for HEP in frequentist hypothesis testing in the absence of uncertainties on rates and shapes of distributions. In practice, the true rate of background is not known exactly, and the shapes of distributions are sensitive to experimental quantities, such as calibration coefficients and particle identification efficiencies (which are also not known exactly). What one would call a *systematic error* in HEP, usually corresponds to what a statistician would refer to as a *nuisance parameter*.

Dealing with nuisance parameters in searches is not a new problem, but perhaps it has never been as essential as it is for the LHC. In these proceedings, Cousins reviews the different approaches to nuisance parameters in HEP and the professional statistical literature.²⁰ Also of interest is the classification of systematic errors provided by Sinervo.²¹ In Sec. 4, the a few techniques for incorporating nuisance parameters are reviewed.

From an experimental point of view, the missing ingredient is some set of auxiliary measurements that will constrain the value of the nuisance parameters. The most common example would be a side-band measurement to fix the background rate, or some control sample used to assess particle identification efficiency. Previously, I used the variable M to denote this auxiliary measurement²²; while Linnemann,¹⁹ Cousins,²⁰ and Rolke, Lopez, and Conrad^{23,24} used y . Additionally, one needs to know

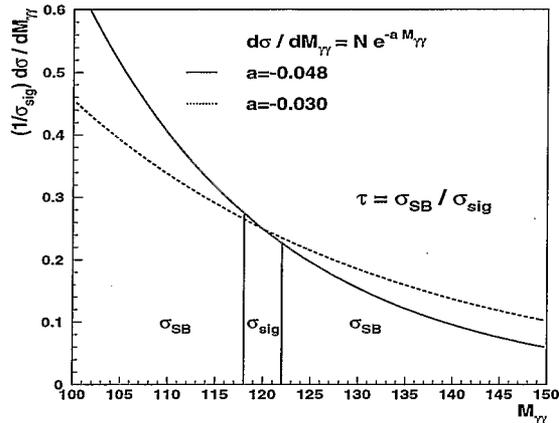


Figure 3. The signal-like region and sideband for $H \rightarrow \gamma\gamma$ in which τ is correlated to b via the model parameter a .

the likelihood function that provides the connection between the nuisance parameter(s) and the auxiliary measurements.

The most common choices for the likelihood of the auxiliary measurement are $L(y|b) = Pois(y|\tau b)$ and $L(y|b) = G(y|\tau b, \sigma_y)$, where τ is a constant that specifies the ratio of the number of events one expects in the sideband region to the number expected in the signal-like region.^d

A constant τ is appropriate when one simply counts the number of events y in an “off-source” measurement. In a more typical case, one uses the distribution of some other variable, call it m , to estimate the number of background events inside a range of m (see Fig. 3). In special cases the ratio τ is independent of the model parameters. However, in many cases (e.g. $f(m) \propto e^{-am}$), the ratio τ depends on the model parameters. Moreover, sometimes the sideband is contaminated with signal events, thus the background and signal estimates can be correlated. These complications are not a problem as long as they are incorporated into the likelihood.

The number of nuisance parameters and auxiliary measurements can grow quite large. For instance, the standard practice at BaBar is to form very large likelihood functions that incorporate everything from the parameters of the unitarity triangle to branching fractions and detector response. These likelihoods are typically factorized into multi-

ple pieces, which are studied independently at first and later combined to assess correlations. The factorization of the likelihood and the number of nuisance parameters included impact the difficulty of implementing the various scenarios considered below.

3 Practical and Toy Examples

In this Section, a few practical and toy examples are introduced. The toy examples are meant to provide simple scenarios where results for different methods can be easily obtained in order to expedite their comparison. The practical examples are meant to exclude methods that provide nice solutions to the toy examples, but do not generalize to the realistic situation.

3.1 The Canonical Example

Consider a number-counting experiment that measures x events in the signal-like region and y events in some sideband. For a given background rate b in the signal-like region, say one can expect τb events in the sideband. Additionally, let the rate of signal events in the signal-like regions – the parameter of interest – be denoted μ . The corresponding likelihood function is

$$L_P(x, y|\mu, b) = Pois(x|\mu + b) \cdot Pois(y|\tau b). \quad (14)$$

This is the same case that was considered in Refs. 20,22,23,24 for $x, y = \mathcal{O}(10)$ and $\alpha = 5\%$. For LHC searches, we will be more interested in $x, y = \mathcal{O}(100)$ and $\alpha = 2.85 \cdot 10^{-7}$. Furthermore, the auxiliary measurement will rarely be a pure number counting sideband measurement, but instead the result of some fit. So let us also consider the likelihood function

$$L_G(x, y|\mu, b) = Pois(x|\mu + b) \cdot G(y|\tau b, \sqrt{\tau b}). \quad (15)$$

As a concrete example in the remaining sections, let us consider the case $b = 100$ and $\tau = 1$. Operationally, one would measure y and then find the value $x_{crit}(y)$ necessary for discovery. In the language of confidence intervals, $x_{crit}(y)$ is the value of x necessary for the $100(1 - \alpha)\%$ confidence interval in μ to exclude $\mu_0 = 0$. In Sec. 4 we check the coverage (Type I error or false-discovery rate) for both L_P and L_G .

Linnemann reviewed thirteen methods and eleven published examples of this scenario.¹⁹ Of the

^dNote that Linnemann¹⁹ used $\alpha = 1/\tau$ instead, but in this paper α is reserved for the rate of Type I error.

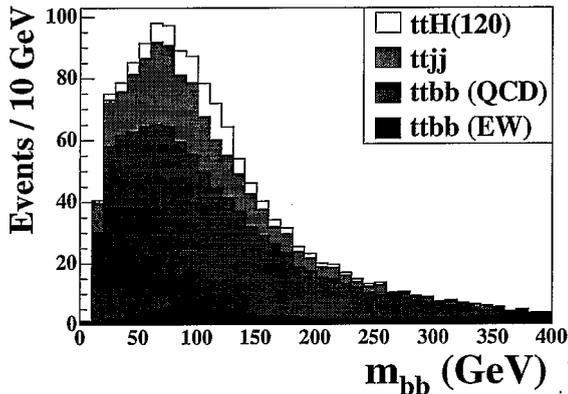


Figure 4. The bb invariant mass spectrum for $t\bar{t}H$ signal and background processes at ATLAS.

published examples, only three (the one from his reference 18 and the two from 19) are near the range of x, y , and α relevant for LHC searches. Linnemann’s review asks an equivalent question posed in this paper, but in a different way: what is the significance (Z in Eq. 12) of a given observation x, y .

3.2 The LHC Standard Model Higgs Search

The search for the standard model Higgs boson is by no means the only interesting search to be performed at the LHC, but it is one of the most studied and offers a particularly challenging set of channels to combine with a single method. Figure 1 shows the expected significance versus the Higgs mass, m_H , for several channels individually and in combination for the ATLAS experiment.²⁵ Two mass points are considered in more detail in Tab. 1, including results from Refs. 1, 25, 26. Some of these channels will most likely use a discriminating variable distribution, $f(m)$, to improve the sensitivity as described in Sec. 2.3. I have indicated the channels that I suspect will use this technique. Rough estimates on the uncertainty in the background rate have also been tabulated, without regard to the classification proposed by Sinervo.

The background uncertainties for the $t\bar{t}H$ channel have been studied in some detail and separated into various sources.²⁶ Figure 4 shows the m_{bb} mass spectrum for this channel.^e Clearly, the shape of the background-only distribution is quite similar to

^eIt is not clear if this result is in agreement with the equivalent CMS result.²⁷

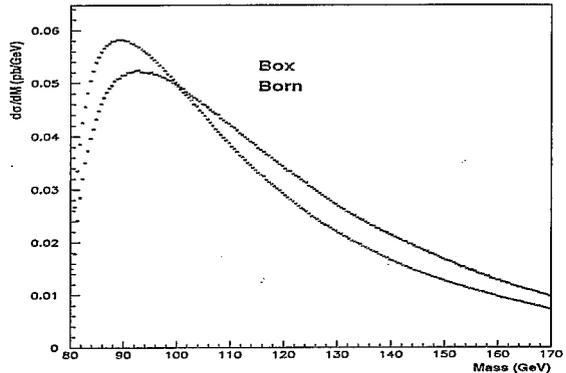


Figure 5. Two plausible shapes for the continuum $\gamma\gamma$ mass spectrum at the LHC.

the shape of the signal-plus-background distribution. Furthermore, theoretical uncertainties and b -tagging uncertainties affect the shape of the background-only spectrum. In this case the incorporation of systematic error on the background rate most likely precludes the expected significance of this channel from ever reaching 5σ .

Similarly, the $H \rightarrow \gamma\gamma$ channel has uncertainty in the shape of the $m_{\gamma\gamma}$ spectrum from background processes. One contribution to this uncertainty comes from the electromagnetic energy scale of the calorimeter (an experimental nuisance parameter), while another contribution comes from the theoretical uncertainty in the continuum $\gamma\gamma$ production. Figure 5 shows two plausible shapes for the $m_{\gamma\gamma}$ spectrum from “Born” and “Box” predictions.

4 Review of Methods

Based on the practical example of the standard model Higgs search at the LHC and the discussion in Sec. 2, the list of admissible methods is quite short. Of the thirteen methods reviewed by Linnemann, only five are considered as reasonable or recommended. These can be divided into three classes: Bayesian methods based on the Cousins-Highland technique, frequentist methods based on the Likelihood principle, and frequentist methods based on the Neyman construction.

4.1 The Cousins-Highland Method

The class of methods frequently used in HEP and commonly referred to as the Cousins-Highland tech-

Table 1. Number of signal and background events for representative Higgs search channels with 30 fb⁻¹ of data. A rough uncertainty on the background rate is denoted as $\delta b/b$, without reference to the type of systematic uncertainty. The table also indicates if the channels are expected to use a weight $f(m)$ as in Eq. 2.

channel	s	b	$\delta b/b$	dominant backgrounds	use $f(m)$	m_H (GeV)
$t\bar{t}H \rightarrow t\bar{t}b\bar{b}$	42	219	$\sim 10\%$	$t\bar{t}j\bar{j}, t\bar{t}b\bar{b}$	Yes	120
$H \rightarrow \gamma\gamma$	357	11820	$\sim 0.1\%$	$\gamma\gamma, j\bar{j}, j\bar{j}$	No	120
$qqH \rightarrow qq\tau\tau \rightarrow qqll\cancel{E}_T$	17	14	$\sim 10\%$	$Z \rightarrow \tau\tau, t\bar{t}$	Yes	120
$qqH \rightarrow qq\tau\tau \rightarrow qqllh\cancel{E}_T$	16	8	$\sim 10\%$	$Z \rightarrow \tau\tau, t\bar{t}$	Yes	120
$qqH \rightarrow qqWW^* \rightarrow qqll\cancel{E}_T$	28.5	47.4	$\sim 10\%$	$t\bar{t}, WW$	Yes	120
$qqH \rightarrow qqWW^* \rightarrow qqll\cancel{E}_T$	262.5	89.1	$\sim 10\%$	$t\bar{t}, WW$	Yes	170
$H \rightarrow ZZ \rightarrow 4l$	7.6	3.1	$\sim 1\%$	$ZZ \rightarrow 4l$	No	170
$H \rightarrow WW \rightarrow ll\cancel{E}_T$	337	484	$\sim 5\%$	$Z \rightarrow \tau\tau, t\bar{t}$	Yes	170

nique (or secondarily Bayes in statistical literature) are based on a Bayesian average of frequentist p -values as found in the first equation of Ref.²⁸. The Bayesian average is over the nuisance parameters and weighted by the posterior $P(b|y)$. Thus the p -value of the observation (x_0, y_0) evaluated at μ is given by

$$p(x_0, y_0|\mu) = \int_0^\infty db p(x_0|\mu, b)P(b|y_0) \quad (16)$$

$$= \int_0^\infty db \int_{x_0}^\infty dx P(x|\mu, b)P(b|y_0) \quad (17)$$

$$= \int_{x_0}^\infty dx P(x|\mu, y_0) \quad (18)$$

where

$$P(x|\mu, y_0) = \int_0^\infty db P(x|\mu, b) \frac{P(y_0|b) P(b)}{P(y_0)} \quad (19)$$

The form in Eq. 16, an average over p -values, is similar to the form written in Cousins & Highland's article. The double-integral is expanded in Eq. 17 and re-written in Eq. 18 to the form that is more familiar to those from LEP Higgs searches.^{16,17} Actually, the dependence on y_0 and the Bayesian prior $P(b)$ shown explicitly in Eq. 19 is often not appreciated by those that use this method.

The specific methods that Linnemann considers correspond to different choices of Bayesian priors. The most common in HEP is to ignore the prior and use a truncated Gaussian for the posterior $P(b|y_0)$, which Linnemann calls Z_N . For the case in which the likelihood $L(y|b)$ is known to be Poisson, Linnemann prefers to use a flat prior, which gives rise to a Gamma-distributed posterior and Linnemann's second preferred method Z_Γ . The method Linnemann calls $Z_{S'}$ can be seen as an approximation of Z_N for

large signals and is what ATLAS used to assess its physics potential.¹ The method not recommended by Linnemann and was critically reviewed in Ref.²⁹.

$$x_{crit}^{S'}(y) = y/\tau + Z\sqrt{y/\tau(1+1/\tau)} \quad (20)$$

4.2 Likelihood Intervals

As Cousins points out, the professional statistics literature seems less concerned with providing correct coverage by construction, in favor of likelihood-based and Bayesian methods. The likelihood principle states that given a measurement x all inference about μ should be based on the likelihood function $L(x|\mu)$. When nuisance parameters are included, things get considerably more complicated.

The profile likelihood function is an attempt to eliminate the nuisance parameters from the likelihood function by replacing them with their conditional maximum likelihood estimates (denoted with a $\hat{\cdot}$). For example, the profile likelihood for L_P in Eq. 14 is given by $L(x, y|\mu_0, \hat{b}(\mu_0))$, with

$$\hat{b}(\mu_0) = \frac{x + y - (1 + \tau)\mu_0}{2(1 + \tau)} + \frac{\sqrt{(x + y - (1 + \tau)\mu_0)^2 + 4(1 + \tau)y\mu_0}}{2(1 + \tau)}. \quad (21)$$

The relevant likelihood ratio is then

$$\lambda_P(\mu|x, y) = \frac{L(x, y|\mu_0, \hat{b}(\mu_0))}{L(x, y|\hat{\mu}, \hat{b})}, \quad (22)$$

where $\hat{\mu}$ and \hat{b} are the unconditional maximum likelihood estimates.

One of the standard results from statistics is that the distribution of $-2 \ln \lambda$ converges to the χ^2 distribution with k degrees of freedom, where k is the

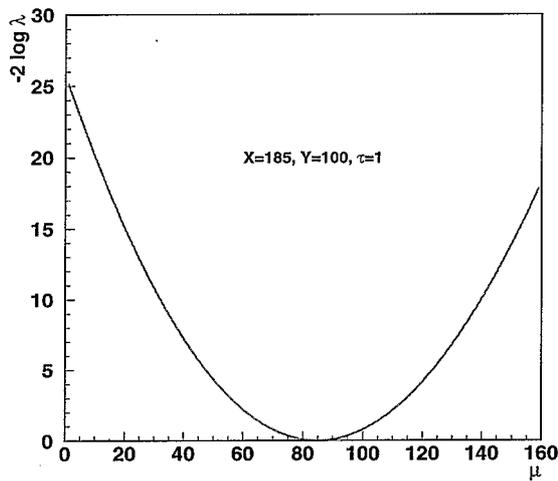


Figure 6. The profile likelihood ratio $-2 \ln \lambda$ versus the signal strength μ for $y = 100$, $\tau = 1$, and $x = x_{crit}(y) = 185$.

number of parameters of interest. In our example $k = 1$, so a 5σ confidence interval is defined by the set of μ with $-2 \ln \lambda(\mu|x, y) < 25$. Figure 6 shows the graph of $-2 \ln \lambda(\mu|x, y)$ for $y = 100$ at the critical value of x for a 5σ discovery.

At PhyStat2003, Nancy Reid presented various adjustments and improvements to the profile likelihood which speed asymptotic convergence properties.³⁰ Cousins considers these methods in more detail from a physicist perspective.²⁰

Only recently was it generally appreciated that the method of MINUIT³¹ commonly used in HEP corresponds to the profile likelihood intervals. The coverage of these methods is not guaranteed, but has been studied in simple cases.^{23,24} These likelihood-based techniques are quite promising for searches at the LHC, but their coverage properties must be assessed in the more complicated context of the LHC with weighted events and several channels. In particular, the distribution of q in Eq. 10 is often highly non-Gaussian.

4.3 The Neyman Construction with Nuisance Parameters

Linnemann's preferred method, Z_{Bi} , is related to the familiar result on the ratio of Poisson means.³² Unfortunately, the form of Z_{Bi} is tightly coupled to the form of Eq. 14, and can not be directly applied to the more complicated cases described above. How-

ever, the standard result on the ratio of Poisson means³² and Cousins' improvement³³ are actually special cases of the Neyman construction with nuisance parameters (without and with conditioning, respectively).

Of course, the Neyman construction does generalize to the more complicated cases discussed above. Two particular types of constructions have been presented, both of which are related to the profile likelihood ratio discussed in Kendall's chapter on likelihood ratio tests & test efficiency.³⁴ This relationship often leads to confusion with the profile likelihood intervals discussed in Sec. 4.2.

The first method is a full Neyman construction over both the parameters of interest and the nuisance parameters, using the profile likelihood ratio as an ordering rule. Using this method, the nuisance parameter is "projected out", leaving only an interval in the parameters of interest. I presented this method at PhyStat2003 in the context of hypothesis testing,^f and similar work was presented by Punzi at this conference.^{22,35} This method provides coverage by construction, independent of the ordering rule used.

The motivation for using the profile likelihood ratio as a test statistic is twofold. First, it is inspired by the Neyman-Pearson lemma in the same way as the Feldman-Cousins ordering rule. Secondly, it is independent of the nuisance parameters; providing some hope of obtaining similar tests.^g Both Punzi and myself found a need to perform some "clipping" to the acceptance regions to protect from irrelevant values of the nuisance parameters spoiling the projection. For this technique to be broadly applicable, some generalization of this clipping procedure is needed and the scalability with the number of parameters must be addressed.^h

The second method, presented by Feldman at the Fermilab conference in 2000, involves a Neyman construction over the parameters of interest, but the nuisance parameters are fixed to the conditional maximum likelihood estimate: a method I will call the *profile construction*. The profile construction is

^fIn simple hypothesis testing μ is not a continuous parameter, but only takes on the values $\mu_0 = 0$ or $\mu_1 = s$.

^gSimilar tests are those in which the critical regions of size α are independent of the nuisance parameters. Similar tests do not exist in general.

^hA Monte Carlo sampling of the nuisance parameter space could be used to curb the curse of dimensionality.²²

an approximation of the full construction, that does not necessarily cover. To the extent that the use of the profile likelihood ratio as a test statistic provides similar tests, the profile construction has good coverage properties. The main motivation for the profile construction is that it scales well with the number of nuisance parameters and that the “clipping” is built in (only one value of the nuisance parameters is considered).

It appears that the CHOOZ experiment actually performed both the full construction (called “FC correct syst.”) and the profile construction (called “FC profile”) in order to compare with the strong confidence technique.³⁶

Another perceived problem with the full construction is that bad over-coverage can result from the projection onto the parameters of interest. It should be made very clear that the coverage probability is a function of both the parameters of interest and the nuisance parameters. If the data are consistent with the null hypothesis for *any* value of the nuisance parameters, then one should probably not reject it. This argument is stronger for nuisance parameters directly related to the background hypothesis, and less strong for those that account for instrumentation effects. In fact, there is a family of methods that lie between the full construction and the profile construction. Perhaps we should pursue a hybrid approach in which the construction is formed for those parameters directly linked to the background hypothesis, the additional nuisance parameters take on their profile values, and the final interval is projected onto the parameters of interest.

5 Results with the Canonical Example

Consider the case $b_{true} = 100$, $\tau = 1$ (*i.e.* 10% systematic uncertainty). For each of the methods we find the critical boundary, $x_{crit}(y)$, which is necessary to reject the null hypothesis $\mu_0 = 0$ at 5σ when y is measured in the auxiliary measurement. Figure 7 shows the contours of L_G , from Eq. 15, and the critical boundary for several methods. The far left curve shows the simple s/\sqrt{b} curve neglecting systematics. The far right curve shows a critical region with the correct coverage. With the exception of the profile likelihood, λ_P , all of the other methods lie between these two curves (*i.e.* all of them under-cover). The rate of Type I error for these methods was evaluated

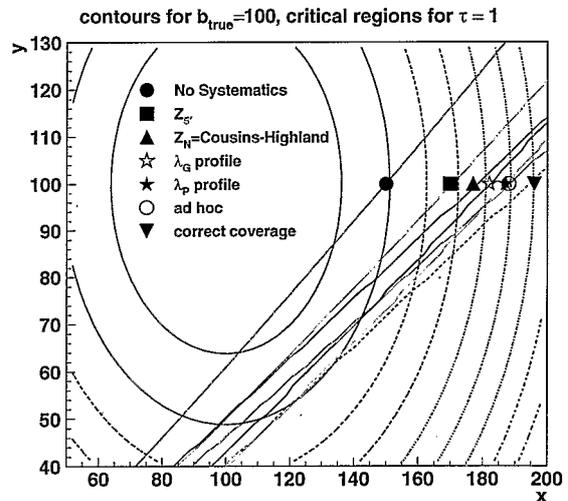


Figure 7. A comparison of the various methods critical boundary $x_{crit}(y)$ (see text). The concentric ovals represent contours of L_G from Eq. 15.

for L_G and L_P and presented in Table 2.

The result of the full Neyman construction and the profile construction are not presented. The full Neyman construction covers by construction, and it was previously demonstrated for a similar case ($b = 100$, $\tau = 4$) that the profile construction gives similar results.²² Furthermore, if the λ_P were used as an ordering rule in the full construction, the critical region for $b = 100$ would be identical to the curve labeled “ λ_P profile” (since λ_P actually covers).

It should be noted that if one knows the likelihood is given by $L_G(x, y|\mu, b)$, then one should use the corresponding profile likelihood ratio, $\lambda_G(x, y|\mu)$, for the hypothesis test. However, knowledge of the correct likelihood is not always available (Sinervo’s Class II systematic), so it is informative to check the coverage of tests based on both $\lambda_G(x, y|\mu)$ and $\lambda_P(x, y|\mu)$ by generating Monte Carlo according to $L_G(x, y|\mu, b)$ and $L_P(x, y|\mu, b)$. In a similar way, this decoupling of true likelihood and the assumed likelihood (used to find the critical region) can break the “guaranteed” coverage of the Neyman construction.

It is quite significant that the Cousins-Highland method, Z_N under-covers, since it is so commonly used in HEP. The degree to which the method under-covers depends on the truncation of the Gaussian posterior $P(b|y)$. Linnemann’s table also shows significant under-coverage (over estimate of the signif-

icance Z). In order to obtain a critical region with the correct coverage, the author modified the region $x_{crit}(y) = x_{crit}^{CH}(y) + C$ and found $C = 16$ provided the correct coverage. A discrepancy of 16 events is not trivial!

Table 2. Rate of Type I error interpreted equivalent $Z\sigma$ for various methods designed for a 5σ test. Monte Carlo events are generated via either L_G or L_P . The critical x for $y = 100$ is also listed for easy comparison.

Method	L_G ($Z\sigma$)	L_P ($Z\sigma$)	$x_{crit}(y = 100)$
No Syst	3.0	3.1	150
$Z_{5'}$	4.1	4.1	171
Z_N C.-H.	4.2	4.2	178
<i>ad hoc</i>	4.6	4.7	188
profile λ_P	5.0	5.0	185
profile λ_G	4.7	4.7	~ 182

Notice that for large x, y the Cousins-Highland result Z_N approaches $Z_{5'}$, where the the critical region is of the form $x_{crit}(y) = y/\tau + n\sqrt{y/\tau}$. Because the boundary is very nearly linear around y_0 , one can find the value of n that gives the proper coverage with a little geometry. In particular, the number n needed to get a $Z\sigma$ test gives

$$x_{crit}(y) = y/\tau + Z\sqrt{1 + 1/\tau m^2}\sqrt{y/\tau} \quad (23)$$

where

$$m = \left(1 + \frac{Z}{2\sqrt{y/\tau}}\right)^{-1} \quad (24)$$

The m^2 factor can be seen as a correction to the $Z_{5'}$ and Z_N (Cousins-Highland) results. Notice that the correction is larger for higher significance tests. As an *ad hoc* method, I experimented with the Cousins-Highland method replacing τ with τm^2 in the posterior $P(b|y)$. The coverage of this *ad hoc* method is better than Z_N , but not exact because x, y are not sufficiently large.

6 Conclusions

I have presented the statistical challenges of searches at the LHC and the current state of the statistical methods commonly used in HEP. I have attempted to accurately portray the complexity of the searches, explain their key ingredients, and provide a practical example for future studies. Three classes of methods, which are able to incorporate all the ingredients, have been identified: Bayesian methods, frequentist methods based on the likelihood principle, and frequentist methods based on the Neyman construction.

The Cousins-Highland method, based on a Bayesian average over the nuisance parameters, shows significant under-coverage in the toy example considered when pushed to the 5σ regime. While Bayesian might not care about coverage, significant under-coverage is undesirable in HEP. Further study is needed to determine if a more careful choice of prior distributions can remedy this situation – especially in more complex situations.

The methods based on the likelihood principle have gained a great deal of attention from HEP in recent years. While the methods appear to do well in the toy example, it requires further study to determine their properties in the more realistic situation with weighted events.

Slowly, the HEP community is coming to grips with how to incorporate nuisance parameters into the Neyman construction. Several ideas for reducing the over-coverage induced by projecting out the nuisance parameters and reducing the computational burden have been presented. A hybrid approach between the full construction and the profile construction should be investigated in more detail.

Finally, it seems that the HEP community is approaching a point where we appreciate the fundamental statistical issues, the limitations of some methods, and the benefits of others. Clearly, the philosophical debate has not ended, but there seems to be more emphasis on practical solutions to our very challenging problems.

Acknowledgments

I would like to thank the many people that helped in preparing this review. In particular, Bob Cousins, Gary Feldman, Jan Conrad, Fredrik Tegenfeldt, Wolfgang Rolke, Nancy Reid, Gary Hill, and Stathes Paganis. I would also like to thank Louis Lyons for his past advice and the invitation to speak at such an enjoyable and productive conference.

This manuscript has been authored by Brookhaven Science Associates, LLC under Contract No. DE-AC02-98CH1-886 with the U.S. Department of Energy. The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, a world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for the U.S. Government purposes.

References

1. ATLAS Collaboration, Detector and physics performance technical design report (volume ii) CERN-LHCC/99-15 (1999).
2. CMS Collaboration, Technical proposal CERN-LHCC/94-38 (1994).
3. H. B. Prosper Advanced Statistical Techniques in Particle Physics, Durham, England, 18-22 Mar 2002.
4. J. H. Friedman, Recent advances in (machine) learning PhyStat2003 (2003).
5. K. Cranmer and R. S. Bowman, *Comput. Phys. Commun.* **167**, 165 (2005).
6. I. Hinchliffe, F. E. Paige, M. D. Shapiro, J. Soderqvist and W. Yao, *Phys. Rev.* **D55**, 5520 (1997).
7. S. Frixione and B. R. Webber, The mc@nlo event generator hep-ph/0207182, (2002).
8. A. Schaliche and F. Krauss, *JHEP* **07**, p. 018 (2005).
9. C. Buttar, D. Clements, I. Dawson and A. Moraes, *Acta Phys. Polon.* **B35**, p. 433 (2004).
10. G. J. Feldman and R. D. Cousins, *Phys. Rev.* **D57**, 3873 (1998).
11. B. Abbott *et al.*, *Phys. Rev.* **D62**, p. 092004 (2000).
12. Y. Gao, L. Lu and X. Wang, Significance calculation and a new analysis method in searching for new physics at the LHC physics/0509174, (2005).
13. G. Feldman, Concluding talk PhyStat2005 (2005).
14. LEP Higgs Working Group, *Phys. Lett.* **B565**, 61 (2003).
15. K. Cranmer, *Comput. Phys. Commun.* **136**, 198 (2001).
16. T. Junk, *Nucl. Instrum. Meth.* **A434**, 435 (1999).
17. H. Hu and J. Nielsen, Analytic Confidence Level Calculations Using the Likelihood Ratio and Fourier Transform CERN 2000-005 physics/9906010, (2000).
18. K. Cranmer, B. Mellado, W. Quayle and Sau Lan Wu, Challenges of Moving the LEP Higgs Statistics to the LHC. PhyStat2003 physics/0312050 (2003).
19. J. Linnemann, Comparison of measures of significance in *PhyStat2003* physics/0312059, (2003).
20. R. Cousins, Treatment of nuisance parameters in high energy physics, and possible justifications and improvements in the statistical literature. PhyStat2005, (2005).
21. P. Sinervo, Definition and treatment of systematic uncertainties in high energy physics and astrophysics PhyStat2003, (2003).
22. K. Cranmer, Frequentist hypothesis testing with background uncertainty PhyStat2003 physics/0310108 (2003).
23. W. A. Rolke and A. M. Lopez, *Nucl. Instrum. Meth.* **A458**, 745 (2001).
24. W. A. Rolke, A. M. Lopez and J. Conrad, *Nucl. Instrum. Meth.* **A551**, 493 (2005).
25. S. Asai *et al.*, *Eur. Phys. J.* **C3252**, 19 (2004).
26. J. Cammin and M. Schumacher, The ATLAS discovery potential for the channel ttH , ($H \rightarrow bb$) ATLAS Note ATL-PHYS-2003-024 (2003).
27. V. Drollinger, Th. Müller, and D. Denegri, Searching for Higgs Bosons in Association with Top Quark Pairs in the $H \rightarrow bb$ Decay Mode CMS NOTE-2001/054 (2001).
28. R. Cousins and V. Highland, *Nucl. Instrum. Meth.* **A320**, 331 (1992).
29. K. Cranmer, P. McNamara, B. Mellado, W. Quayle, and Sau Lan Wu, Confidence level calculations for $H \rightarrow W^+W^- \rightarrow l^+l^- \cancel{p}_T$ for $115 < M_H < 130$ GeV using vector boson fusion ATL-PHYS-2003-008 (2002).
30. N. Reid, Likelihood inference in the presence of nuisance parameters PhyStat2003, (2003).
31. F. James and M. Roos, *Comput. Phys. Commun.* **10**, 343 (1975).
32. F. James and M. Roos, *Nucl. Phys.* **B 172**, 475 (1980).
33. R. Cousins, *Nucl. Instrum. and Meth. in Phys. Res.* **A 417**, 391 (1998).
34. J. Stuart, A. Ord and S. Arnold, *Kendall's Advanced Theory of Statistics, Vol 2A (6th Ed.)* (Oxford University Press, New York, 1994).
35. G. Punzi, Ordering algorithms and confidence intervals in the presence of nuisance parameters PhyStat2005 (2005).
36. D. Nicolo and G. Signorelli Prepared for Conference on Advanced Statistical Techniques in Particle Physics, Durham, England, 18-22 Mar 2002.