

SHAPE-CONSTRAINED UNCERTAINTY QUANTIFICATION IN UNFOLDING STEEPLY FALLING ELEMENTARY PARTICLE SPECTRA

BY MIKAEL KUUSELA^{*,†} AND PHILIP B. STARK

*École Polytechnique Fédérale de Lausanne and
University of California, Berkeley*

The high energy physics unfolding problem is an important statistical inverse problem in data analysis at the Large Hadron Collider (LHC) at CERN. The goal of unfolding is to make nonparametric inferences about a particle spectrum from measurements smeared by the finite resolution of the particle detectors. Previous unfolding methods use ad hoc discretization and regularization, resulting in confidence intervals that can have significantly lower coverage than their nominal level. Instead of regularizing using a roughness penalty or stopping iterative methods early, we impose physically motivated shape constraints: positivity, monotonicity, and convexity. We quantify the uncertainty by constructing a nonparametric confidence set for the true spectrum, consisting of all those spectra that satisfy the shape constraints and that predict the observations within an appropriately calibrated level of fit. Projecting that set produces simultaneous confidence intervals for all functionals of the spectrum, including averages within bins. The confidence intervals have guaranteed conservative frequentist finite-sample coverage in the important and challenging class of unfolding problems for steeply falling particle spectra. We demonstrate the method using simulations that mimic unfolding the inclusive jet transverse momentum spectrum at the LHC. The shape-constrained intervals provide usefully tight conservative inferences, while the conventional methods suffer from severe undercoverage.

1. Introduction. This paper studies shape-constrained statistical inference in the high energy physics (HEP) *unfolding problem* (Prosper and Lyons, 2011; Cowan, 1998; Blobel, 2013; Kuusela and Panaretos, 2015). This is an important statistical inverse problem arising at the Large Hadron Collider (LHC) at CERN, the European Organization for Nuclear Research, near Geneva, Switzerland. LHC measurements are affected by the finite reso-

^{*}Supported by the Swiss National Science Foundation grant no. 200021-153595.

[†]Now at the University of Chicago.

MSC 2010 subject classifications: Primary 62G15, 62G07; secondary 45Q05, 62P35

Keywords and phrases: Poisson inverse problem, finite-sample coverage, high energy physics, Large Hadron Collider, Fenchel duality, semi-infinite programming

lution of the particle detectors. This causes the observations to be “smeared” or “blurred” versions of the true physical spectra. The unfolding problem is the inverse problem of making inferences about the true spectrum from the noisy, smeared observations.

The unfolding problem can be formalized using indirectly observed Poisson point processes (Kuusela and Panaretos, 2015). It is an ill-posed Poisson inverse problem (Antoniadis and Bigot, 2006; Reiss, 1993). Let M and N be Poisson point processes with intensity functions f_0 and g_0 and state spaces T and S , respectively. We assume that the state spaces are compact intervals of real numbers. Depending on the analysis, they could represent the energies, momenta, or production angles of particles, for example. Let M represent the true, particle-level collision events, and N the smeared, detector-level events. The intensity functions f_0 and g_0 are related by the Fredholm integral equation

$$(1) \quad g_0(s) = (Kf_0)(s) = \int_T k(s, t) f_0(t) dt, \quad s \in S.$$

The kernel $k(s, t)$, which represents the response of the particle detector, is

$$(2) \quad k(s, t) \equiv p(Y = s | X = t, X \text{ observed})P(X \text{ observed} | X = t),$$

where X is a physical particle-level event and Y the corresponding smeared, detector-level event. In this paper, we take $k(s, t)$ to be known. The goal of unfolding is to make inferences about the true intensity f_0 from a single realization of the smeared process N .

Unfolding is used in dozens of LHC analyses annually. Existing unfolding methods, implemented in the ROOUNFOLD (Adye, 2011) software framework, have serious practical limitations. These methods first discretize the continuous problem in Equation (1) using histograms (Cowan, 1998, Chapter 11) and then regularize the estimation problem either by stopping the expectation-maximization (EM) iteration early (D’Agostini, 1995) or by using certain variants of Tikhonov regularization (Höcker and Kartvelishvili, 1996; Schmitt, 2012). Both approaches are ad hoc and produce biased estimates with formal uncertainties that can grossly underestimate the true uncertainty.

More specifically, let $\{T_i\}_{i=1}^p$ and $\{S_i\}_{i=1}^n$ be partitions of the true space T and the smeared space S using histogram bins, and let

$$\boldsymbol{\lambda}_0 \equiv \left[\int_{T_1} f_0(t) dt, \dots, \int_{T_p} f_0(t) dt \right]^T \quad \text{and} \quad \boldsymbol{\mu}_0 \equiv \left[\int_{S_1} g_0(s) ds, \dots, \int_{S_n} g_0(s) ds \right]^T$$

be the expected number of true events and of smeared events in each bin, respectively. The two are related by $\boldsymbol{\mu}_0 = \mathbf{K}\boldsymbol{\lambda}_0$, where the elements of the response matrix \mathbf{K} are

$$(3) \quad K_{i,j} \equiv \frac{\int_{S_i} \int_{T_j} k(s,t) f_0(t) dt ds}{\int_{T_j} f_0(t) dt}, \quad i = 1, \dots, n, \quad j = 1, \dots, p.$$

This matrix depends on the shape of f_0 within each true bin. Since f_0 is unknown, the matrix \mathbf{K} is in practice constructed by replacing f_0 with a simulated ansatz obtained using a Monte Carlo (MC) event generator. After discretization, the unfolding problem amounts to estimating $\boldsymbol{\lambda}_0$ in the Poisson regression problem $\mathbf{y} \sim \text{Poisson}(\mathbf{K}\boldsymbol{\lambda}_0)$, where $\mathbf{y} \equiv [N(S_1), \dots, N(S_n)]^T$ is the vector of bin counts corresponding to the histogram of smeared observations.

The Tikhonov-regularized unfolding techniques then make a Gaussian approximation to the Poisson likelihood and estimate $\boldsymbol{\lambda}_0$ by solving the optimization problem

$$(4) \quad \min_{\boldsymbol{\lambda} \in \mathbb{R}^p} (\mathbf{y} - \mathbf{K}\boldsymbol{\lambda})^T \hat{\mathbf{C}}^{-1} (\mathbf{y} - \mathbf{K}\boldsymbol{\lambda}) + \delta P(\boldsymbol{\lambda}),$$

where $\hat{\mathbf{C}} \equiv \text{diag}(\mathbf{y})$ is an estimate of the covariance of \mathbf{y} , $\delta > 0$ is a regularization parameter, and $P(\boldsymbol{\lambda})$ is a penalty term to regularize the otherwise ill-posed problem. Two Tikhonov-regularized unfolding techniques are common in LHC data analysis: The TUNFOLD variant (Schmitt, 2012) uses the penalty term $P(\boldsymbol{\lambda}) = \|\mathbf{L}(\boldsymbol{\lambda} - \boldsymbol{\lambda}^{\text{MC}})\|_2^2$, where $\boldsymbol{\lambda}^{\text{MC}}$ is a Monte Carlo prediction of $\boldsymbol{\lambda}_0$ and \mathbf{L} is typically a discretized second derivative operator. The singular value decomposition (SVD) variant (Höcker and Kartvelishvili, 1996) replaces the difference $\boldsymbol{\lambda} - \boldsymbol{\lambda}^{\text{MC}}$ with the binwise ratios of $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda}^{\text{MC}}$. EM with early stopping (D'Agostini, 1995) starts the iteration from $\boldsymbol{\lambda}^{\text{MC}}$ and then stops before convergence, which regularizes the solution (iterating to convergence produces undesired oscillations). All these methods regularize the problem by biasing the solution towards a Monte Carlo prediction of $\boldsymbol{\lambda}_0$. The uncertainty of the resulting point estimate $\hat{\boldsymbol{\lambda}}$ is then quantified by estimating its binwise standard errors using error propagation, ignoring the discretization and regularization biases.

The biggest problem with the current unfolding methods is that their uncertainty estimates are unrealistically small. In HEP applications, it is crucial to have confidence intervals with good frequentist coverage properties. But it is well understood that constructing confidence intervals using only the variability of a point estimate generally leads to undercoverage if the point estimate is biased. These issues have been studied extensively,

for instance in the spline smoothing literature (Ruppert, Wand and Carroll, 2003, Chapter 6). Kuusela and Panaretos (2015) show that standard uncertainty quantification techniques can lead to dramatic undercoverage in the unfolding problem. Moreover, using a Monte Carlo prediction first to discretize and then to regularize the problem introduces an unnecessary model-dependence and an additional uncertainty that is extremely difficult to quantify rigorously.

Spectra that fall steeply over many orders of magnitude are especially difficult to estimate well. For such spectra, the response matrix \mathbf{K} depends strongly on the assumed Monte Carlo model, since the spectrum varies substantially within each true bin. Current techniques regularize estimates of such spectra by requiring the unfolded solution to be close to a steeply falling Monte Carlo prediction. Trying to reduce the Monte Carlo dependence by using a global second-derivative penalty that does not depend on λ^{MC} does not work well because the true solution has a highly heterogeneous second derivative, large on the left and small on the right.

Energy, invariant mass and transverse momentum spectra of particle interactions are typically steeply falling (assuming that the analysis is carried out above relevant low-energy thresholds). Recent LHC analyses that involve unfolding steeply falling spectra include the measurement of the differential cross section of jets (CMS Collaboration, 2013a), top quark pairs (CMS Collaboration, 2013b), the W boson (ATLAS Collaboration, 2012), and the Higgs boson (CMS Collaboration, 2016). Expected key outcomes of LHC Run II, currently underway, include more precise measurements of these and related steeply falling differential cross sections. A better method for unfolding such spectra is needed.

This paper develops a new method for forming rigorous confidence intervals in the unfolding problem using an approach that is particularly well suited to steeply falling spectra. Instead of regularizing using a roughness penalty or by stopping EM early, we use shape constraints on the spectrum: positivity, monotonicity, and convexity. Such constraints are an intuitive and physically justified way to impose regularity in a wide range of unfolding analyses, including those mentioned above, without the need to bias the solution towards a Monte Carlo prediction nor to choose a roughness measure or regularization parameter. The use of shape constraints to estimate a steeply falling spectrum is demonstrated in Figure 1.

Shape constraints are easy to incorporate into uncertainty quantification using *strict bounds* (Stark, 1992), which construct simultaneous confidence intervals by solving optimization problems over the set of all functions that satisfy the shape constraints and agree adequately with the data (the misfit

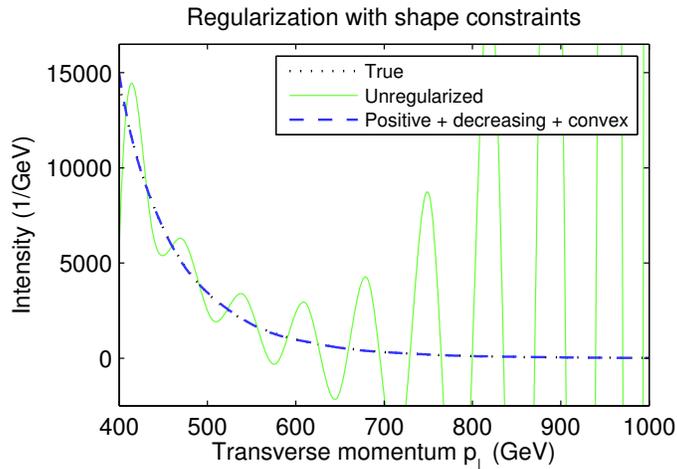


FIG 1. Regularization of a steeply falling elementary particle spectrum using shape constraints. Here the spectrum is the inclusive jet differential cross section described in Section 5.1. Unfolding these data is an ill-posed problem as demonstrated by the solid curve, which shows a typical unregularized point estimate (here a cubic spline that maximizes the likelihood). Shape constraints provide a natural way to rule out such unphysical solutions. If the spline coefficients are constrained so that the estimate is positive, decreasing and convex (dashed line), it becomes hard to distinguish from the true spectrum (dotted line). (These point estimates are presented to demonstrate how well shape constraints regularize the ill-posed problem. The uncertainty quantification method proposed in this paper does not involve point estimates or splines.)

criterion and misfit level are chosen to yield the desired coverage probability). For the unfolding problem, the intervals are computed using a semi-discrete forward mapping, allowing the spectrum to vary arbitrarily within the true bins, subject only to the shape constraints. This approach enables us to form confidence intervals for λ_0 with *guaranteed finite-sample simultaneous frequentist coverage*, avoiding the model-dependence and discretization and regularization biases of existing methods. The proposed intervals are conservative, but still usefully sharp. Figure 2 shows an example of the resulting intervals for the inclusive jet transverse momentum spectrum studied in Section 5. Deriving such intervals and methods to compute them are the main contributions of this paper. To do so, we extend the strict bounds methodology of Stark (1992) to Poisson noise and develop a new way of imposing monotonicity and convexity constraints conservatively.

The rest of this paper is structured as follows: Section 2 sketches the physical and statistical background of this work. Section 3 provides an outline the proposed inference method. Section 4 explains in detail how to construct

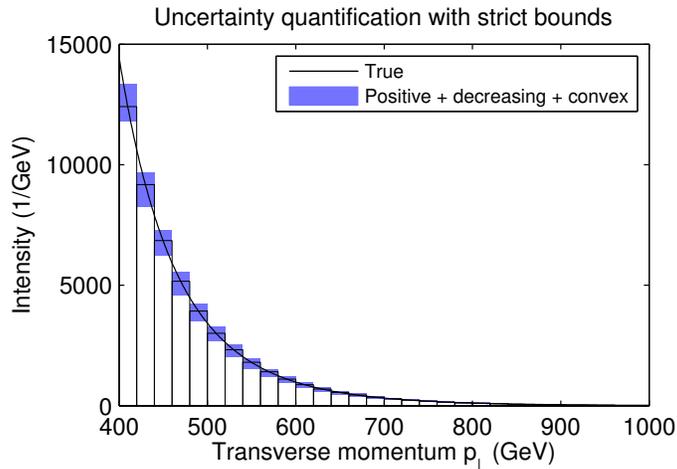


FIG 2. Shape-constrained strict bounds for unfolding the inclusive jet transverse momentum spectrum described in Section 5.1. The confidence intervals are formed by considering all functions that satisfy the shape constraints (here positivity, monotonicity, and convexity) and that fit the observations within a tolerance calibrated to give the specified simultaneous confidence level (here 95%). By construction, the finite-sample simultaneous coverage of these intervals is guaranteed to be at least 95%.

and compute the shape-constrained strict bounds. Section 5 illustrates the bounds in a simulation study and demonstrates that the existing methods can fail to achieve their nominal coverage, even in a typical unfolding scenario. Section 6 concludes. Appendices give derivations and implementation details.

2. Context and background.

2.1. *LHC experiments and unfolding.* The Large Hadron Collider (LHC) is the world's largest and most powerful particle accelerator. It collides protons at velocities close to the speed of light to study the properties and interactions of elementary particles created in these collisions. The collision events are recorded using four massive underground particle detectors, called ATLAS, ALICE, CMS, and LHCb. ATLAS and CMS are general-purpose experiments, while ALICE focuses on heavy-ion and LHCb on bottom-quark physics. Together, these experiments produce approximately 30 petabytes of data annually.

Figure 3 shows the structure of a typical HEP experiment (here the CMS detector, CMS Collaboration (2008)). Particles collide at the center of the experiment. The collision point is surrounded by layers of detectors, which

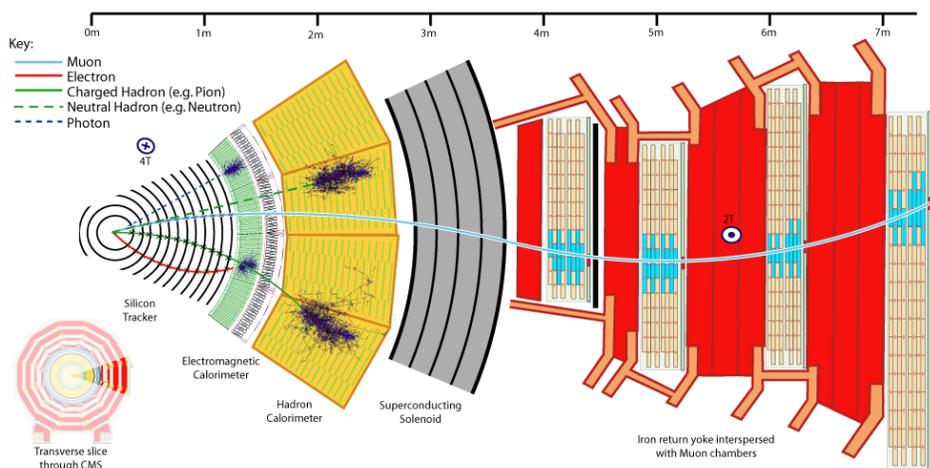


FIG 3. Structure of the CMS experiment at the LHC (Barney, 2004). The experiment uses layers of subdetectors to identify the particles created in the proton-proton collisions and to estimate their properties. Copyright: CERN, for the benefit of the CMS Collaboration.

record information that can be used to reconstruct the collision event. The innermost layer is the *tracker*, which measures the trajectories of charged particles. The next layer is the *electromagnetic calorimeter*, which measures the energies of electrons, positrons, and photons. This is followed by the *hadron calorimeter*, which measures the energies of hadrons—particles composed of quarks. The outermost layer consists of *muon detectors*, which identify and trace muons. The whole detector system is in a strong magnetic field that bends the trajectories of charged particles, enabling their momenta to be measured. Each type of particle leaves a characteristic signature in the various subdetectors; see Figure 3. Pattern recognition algorithms are then used to combine information from the subdetectors to reconstruct the collision event and identify physics objects of interest.

The unfolding problem arises because of the finite resolution of the LHC particle detectors. Let X_i be the true value of some quantity measured in the detector, for instance, the energy of a particle hitting the electromagnetic calorimeter. Due to detector effects, the observed value of this quantity, say, Y_i , includes stochastic noise: the conditional density of Y_i given X_i , $p(Y_i|X_i)$, has support on a continuum of values around X_i . If the noise were additive, i.e., $Y_i = X_i + \varepsilon_i$, with ε_i independent of X_i , then unfolding becomes a deconvolution problem. Usually the appropriate noise models are more complicated though, and methods that can handle general smearing kernels are needed.

It follows from the law of rare events that the number of collisions in HEP experiments is, to a good approximation, Poisson distributed. The particle-level quantities X_1, \dots, X_τ can be treated as a Poisson point process, from which it follows that the smeared observations Y_1, \dots, Y_ξ also comprise a Poisson point process (we may have $\xi < \tau$ due to the limited efficiency of the detector). This leads to the indirectly observed Poisson point process model in Section 1; see also [Kuusela and Panaretos \(2015, Section 3\)](#).

Unfolding is used extensively in LHC data analysis, but not all analyses need to be performed in the unfolded space. In particular, searches for new particles and phenomena usually can be carried out as hypothesis tests in the smeared space. This is in contrast with *measurement analyses* where the goal is to estimate some physically relevant particle spectrum, which is usually most naturally done in the unfolded space. Typically, the main goal in unfolding analyses is to estimate the spectrum in a form that allows other scientists to use it as input to further analyses. This could include hypothesis tests in the unfolded space, combinations of unfolded measurements, or inferring some physical parameter from one or more unfolded results. For example, parton distribution functions ([Forte and Watt, 2013](#)), which quantify the internal structure of a proton, are estimated using simultaneous fits to several unfolded spectra ([NNPDF Collaboration, 2015](#)).

A more detailed introduction to LHC experiments and the role of unfolding in their data analysis is given in [Kuusela and Panaretos \(2015, Section 2\)](#).

2.2. Previous work. Shape-constrained statistical inference goes back to the seminal work of [Grenander \(1956\)](#), who derived the nonparametric maximum likelihood estimator of a density subject to a monotonicity constraint. There is now a large, diverse literature on the topic; see [Robertson, Wright and Dykstra \(1988\)](#) and [Groeneboom and Jongbloed \(2014\)](#) for overviews. Interest has focused on shape-constrained point estimation of density or regression functions using direct observations, and on the nonstandard asymptotic properties of such estimators; see, e.g., [Groeneboom, Jongbloed and Wellner \(2001\)](#) and the references therein. Similarly, there is an extensive literature on nonparametric deconvolution problems; see [Meister \(2009\)](#) for a comprehensive introduction. But using shape constraints to regularize deconvolution-type problems has received only limited attention; examples include [Wahba \(1982\)](#); [Stark \(1992\)](#); [Carroll, Delaigle and Hall \(2011\)](#); [Pflug and Wets \(2013\)](#).

Constructing nonparametric confidence intervals for functions is challenging. When the intervals are based on smoothed estimators, it is difficult to take into account the regularization bias ([Hall and Horowitz, 2013](#); [Ruppert,](#)

Wand and Carroll, 2003, Chapter 6). In fact, under the usual assumptions it is not possible to obtain nonparametric correct-coverage confidence bands that adapt to the smoothness of the underlying function (Low, 1997; Genovese and Wasserman, 2008). This difficulty can be circumvented using shape constraints (Low, 1997; Cai, Low and Xia, 2013), but so far most approaches have been asymptotic (Banerjee and Wellner, 2001; Groeneboom and Jongbloed, 2015). For direct observations, finite-sample shape-constrained confidence intervals with various degrees of adaptivity have been obtained by Hengartner and Stark (1995); Dümbgen (1998, 2003); Davies, Kovac and Meise (2009); Cai, Low and Xia (2013). Rust and Burrus (1972) sketch a construction similar to ours for indirect observations, but do not work out the details or provide a concrete application. In follow-up work, Pierce and Rust (1985) provide an algorithm for various shape constraints, but only for a discretized version of the nonparametric problem. Stark (1992) provides a general recipe for constructing constrained finite-sample confidence intervals in indirect nonparametric problems and serves as the methodological basis for our work.

The two main approaches to unfolding in experimental HEP are Tikhonov regularization (Höcker and Kartvelishvili, 1996; Schmitt, 2012), which goes back to the work of Tikhonov (1963) and Phillips (1962), and D’Agostini iteration (D’Agostini, 1995), which applies EM (Dempster, Laird and Rubin, 1977) in the discrete Poisson regression model but stops before convergence, to provide regularization. D’Agostini iteration has been discovered before in various fields, including optics (Richardson, 1972), astronomy (Lucy, 1974), and positron emission tomography (Shepp and Vardi, 1982; Lange and Carson, 1984; Vardi, Shepp and Kaufman, 1985). It was popularized in HEP by D’Agostini (1995), but had already been studied earlier by Kondor (1983) and Mülthei and Schorr (1987a,b, 1989). Recent proposals for solving the unfolding problem include Bayesian estimation (Choudalakis, 2012), empirical Bayes estimation (Kuusela and Panaretos, 2015), and EM iteration with smoothing (Volobouev, 2015). The last two papers also study the coverage properties of confidence intervals for the unfolded spectrum. (D’Agostini (1995) also describes a smoothed version of the EM iteration, but the method is typically applied without smoothing in current LHC analyses.)

Strict bounds confidence intervals have been used primarily in geophysics, solar physics, and density estimation, using constraints of positivity, monotonicity, m -modality, and general conical constraints; see Stark (1992); Hengartner and Stark (1992, 1995) and references therein. A similar approach was developed by Burrus (1965) for applications in nuclear spectroscopy. Burrus (1965); Burrus and Verbinski (1969); O’Leary and Rust (1986);

Rust and O’Leary (1994) use an approach related to ours for constructing positivity-constrained confidence intervals in unfolding. Their applications use only a positivity constraint because the underlying spectra contain one or more peaks. In contrast, for differential cross section measurements in HEP, the spectra typically are known to be steeply falling and hence monotone and usually also convex. Earlier work on strict bounds also does not explicitly treat Poisson noise.

3. Outline of the methodology. We propose quantifying the uncertainty of the unfolded particle spectrum using shape-constrained *strict bounds confidence intervals* (Stark, 1992). A rough outline of the approach is as follows:

1. For fixed $\alpha \in (0, 1)$, form a $1 - \alpha$ simultaneous confidence set in the binned smeared space. This is the set of all those smeared histograms that are compatible with the smeared observations.
2. Consider the preimage of this set in the unbinned true space. The preimage is a $1 - \alpha$ simultaneous confidence set for the true spectrum. Intuitively, it is the set of all those unfolded spectra that are consistent with the smeared observations.
3. Take the intersection of this unfolded confidence set with the set of spectra that satisfy the shape constraints. If the true spectrum satisfies the imposed constraints, the intersection remains a $1 - \alpha$ simultaneous confidence set for that spectrum, because spectra that violate the constraints cannot contribute to the coverage. The constraints we consider are positivity; positivity and monotonicity; and positivity, monotonicity, and convexity.
4. Characterize the unbinned shape-constrained confidence set in terms of its projections to given functionals. For any functional, the interval from the smallest value of the functional over the confidence set to the largest value of the functional over the confidence set is a $1 - \alpha$ confidence interval for that functional. Moreover, for any set of functionals whatsoever (even an uncountably infinite set), those intervals are simultaneous $1 - \alpha$ confidence intervals for the functionals, because all the intervals cover whenever the unbinned confidence set covers the true unfolded spectrum—which occurs with probability at least $1 - \alpha$. In particular, one may take the functionals to be a collection of binwise integrals of the spectrum, which yields confidence intervals for binwise averages, with $1 - \alpha$ simultaneous coverage in the binned true space.

The last step of this procedure is conservative and may lead to overcoverage. However, undercoverage, which could lead to seriously faulty inferences, is

not possible. The rest of this section develops this approach more rigorously, specifically for Poisson unfolding.

Let F denote the space of (sufficiently regular) functions on the compact interval $T \subset \mathbb{R}$. The random point measure M on state space T is a *Poisson point process* (Reiss, 1993) with a nonnegative intensity function $f \in F$ if and only if

- (a) $M(B) \sim \text{Poisson}(\lambda(B))$ with $\lambda(B) = \int_B f(t) dt$ for every Borel set $B \subset T$;
- (b) $M(B_1), \dots, M(B_n)$ are independent for pairwise disjoint Borel sets $B_i \subset T$, $i = 1, \dots, n$.

In the unfolding problem, the true, particle-level events are a realization of such a Poisson point process M on state space T with intensity function f_0 . Let G denote the space of (sufficiently regular) functions on the compact interval $S \subset \mathbb{R}$. The smeared, detector-level events are a realization of the Poisson point process N on state space S with a nonnegative intensity function $g_0 \in G$. The two intensity functions are related by $g_0 = K f_0$, where the forward operator $K : F \rightarrow G$ is defined by Equation (1).

HEP data are typically binned for convenience and computational tractability. For steeply falling spectra, the leftmost bins may contain billions of observations: treating them individually is not currently computationally feasible. We assume therefore that the smeared observations are counts in bins $\mathbf{y} \equiv [N(S_1), \dots, N(S_n)]^T$, where $\{S_j\}_{j=1}^n$ is a binning of the smeared space S :

$$(5) \quad S_j \equiv \begin{cases} [S_{j,\min}, S_{j,\max}), & j = 1, \dots, n-1, \\ [S_{j,\min}, S_{j,\max}], & j = n, \end{cases}$$

with $S_{j,\max} = S_{j+1,\min}$, $j = 1, \dots, n-1$. The expected value of these counts is $\boldsymbol{\mu}_0 \equiv \left[\int_{S_1} g_0(s) ds, \dots, \int_{S_n} g_0(s) ds \right]^T$, which has components

$$(6) \quad \mu_{0,j} \equiv \int_{S_j} \int_T k(s,t) f_0(t) dt ds = \int_T k_j(t) f_0(t) dt \equiv K_j f_0, \quad j = 1, \dots, n,$$

where $k_j(t) \equiv \int_{S_j} k(s,t) ds$. We assume here and below that f_0 and k are sufficiently regular that we can change the order of integration.

Let \mathcal{K} denote the semi-discrete linear operator that maps a particle-level intensity function f to the detector-level mean vector $\boldsymbol{\mu}$:

$$(7) \quad \mathcal{K} : F \rightarrow \mathbb{R}^n, \quad f \mapsto [K_1 f, \dots, K_n f]^T.$$

With this notation, our statistical model is

$$(8) \quad \mathbf{y} \sim \text{Poisson}(\boldsymbol{\mu}_0), \text{ with } \boldsymbol{\mu}_0 = \mathcal{K}f_0,$$

with the components of \mathbf{y} independent.

We seek simultaneous confidence intervals for the binned expected value vector of the true process M :

$$(9) \quad \boldsymbol{\lambda}_0 \equiv \left[\int_{T_1} f_0(t) dt, \dots, \int_{T_p} f_0(t) dt \right]^T = [H_1 f_0, \dots, H_p f_0]^T,$$

where $\{T_k\}_{k=1}^p$ is a partition of T into intervals and, for each $k = 1, \dots, p$, $H_k f \equiv \int_{T_k} f(t) dt$. As before, the partition is of the form

$$(10) \quad T_k \equiv \begin{cases} [T_{k,\min}, T_{k,\max}), & k = 1, \dots, p-1, \\ [T_{k,\min}, T_{k,\max}], & k = p, \end{cases}$$

with $T_{k,\max} = T_{k+1,\min}$, $k = 1, \dots, p-1$.

Instead of further discretizing the forward mapping, we construct strict bounds confidence intervals (Stark, 1992) under the semi-discrete mapping \mathcal{K} as follows (cf. the outline in the beginning of this section):

1. For fixed $\alpha \in (0, 1)$, let $\Xi \subset \mathbb{R}^n$ be a $1 - \alpha$ simultaneous confidence set for $\boldsymbol{\mu}_0$ in the smeared space, so that $\mathbb{P}_{f_0}\{\Xi \ni \boldsymbol{\mu}_0\} \geq 1 - \alpha$.
2. The preimage $D \equiv \mathcal{K}^{-1}(\Xi)$ of Ξ under the mapping \mathcal{K} is a $1 - \alpha$ confidence set for the unbinned particle-level intensity function f_0 : $\mathbb{P}_{f_0}\{D \ni f_0\} \geq 1 - \alpha$. (The symbol D is mnemonic for *data*.)
3. Let $C \subset F$ denote the set of functions that satisfy the shape constraints. (The symbol C is mnemonic for *constraint*, and also for *convex cone*, which all our constraint sets are.) If $f_0 \in C$, then $C \cap D$ is also a $1 - \alpha$ confidence set for f_0 : $\mathbb{P}_{f_0}\{C \cap D \ni f_0\} \geq 1 - \alpha$.
4. Whenever $f_0 \in C \cap D$ (which occurs with probability at least $1 - \alpha$), any condition that holds for every $f \in C \cap D$ holds for f_0 . In particular, $H_k f_0 \geq \inf_{f \in C \cap D} H_k f \equiv \underline{\lambda}_k$ and $H_k f_0 \leq \sup_{f \in C \cap D} H_k f \equiv \bar{\lambda}_k$, for every k . Each interval $[\underline{\lambda}_k, \bar{\lambda}_k]$ is a $1 - \alpha$ confidence interval for $\lambda_{0,k}$. Moreover, the set $[\underline{\lambda}_1, \bar{\lambda}_1] \times \dots \times [\underline{\lambda}_p, \bar{\lambda}_p] \subset \mathbb{R}^p$ is a $1 - \alpha$ simultaneous confidence set for $\boldsymbol{\lambda}_0$: Whenever $C \cap D$ covers f_0 , the p intervals $\{[\underline{\lambda}_j, \bar{\lambda}_j]\}_{j=1}^p$ all cover the corresponding components of $\boldsymbol{\lambda}_0$ and $\mathbb{P}_{f_0}([\underline{\lambda}_1, \bar{\lambda}_1] \times \dots \times [\underline{\lambda}_p, \bar{\lambda}_p] \ni \boldsymbol{\lambda}_0) \geq 1 - \alpha$.

This construction is illustrated in Figure 4.

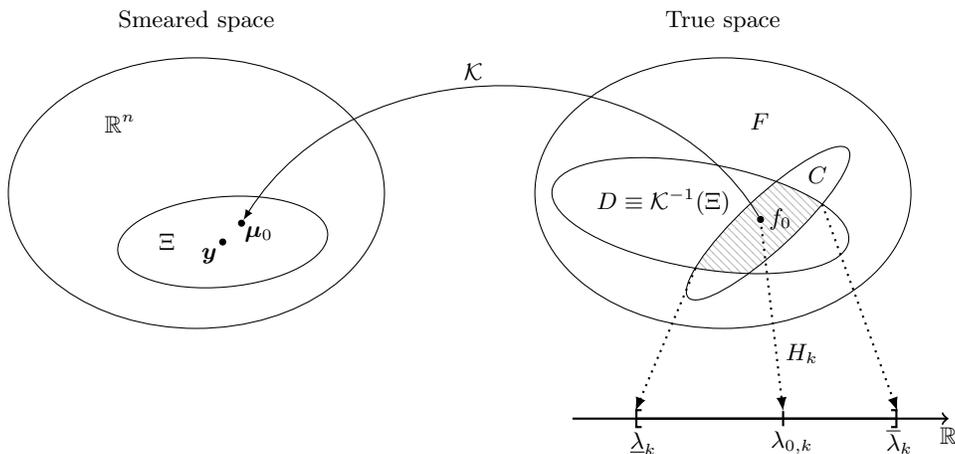


FIG 4. Sketch of strict bounds confidence intervals. The set Ξ in the smeared space is a confidence set for μ_0 based on \mathbf{y} . Its preimage $D \equiv \mathcal{K}^{-1}(\Xi)$ is a confidence set for f_0 . If f_0 satisfies the shape constraints C , then $C \cap D$ is also a confidence set for f_0 . The endpoints of the confidence intervals are given by the extremal values of the functionals of interest H_k over $C \cap D$.

The simultaneous coverage probability of $[\underline{\lambda}_1, \bar{\lambda}_1] \times \cdots \times [\underline{\lambda}_p, \bar{\lambda}_p]$ can be greater than that of $C \cap D$: it has elements that do not correspond to any $f \in C \cap D$. Thus, the intervals are guaranteed to have simultaneous confidence level at least $1 - \alpha$, but the actual coverage probability might be much larger.

Constructing confidence intervals involves solving the optimization problems $\inf_{f \in C \cap D} H_k f$ and $\sup_{f \in C \cap D} H_k f$. Since

$$(11) \quad \sup_{f \in C \cap D} H_k f = - \inf_{f \in C \cap D} -H_k f,$$

we focus on the minimization problem, without loss of generality. We cannot directly solve the infinite-dimensional optimization problem $\inf_{f \in C \cap D} H_k f$. Instead, we follow the approach of Stark (1992) and use Fenchel duality (Luenberger, 1969, Section 7.12) to bound $\inf_{f \in C \cap D} H_k f$ from below by a semi-infinite program with an n -dimensional unknown and an infinite set of constraints. We then discretize the constraints in such a way that every feasible point of the discretized finite-dimensional dual program gives a lower bound for the value of the original program $\inf_{f \in C \cap D} H_k f$, ensuring that the simultaneous confidence level of the resulting intervals remains at least $1 - \alpha$.

4. Shape-constrained unfolding. This section shows how to compute shape-constrained strict bounds confidence intervals in the unfolding

problem. Section 4.1 develops the confidence set Ξ for the smeared mean vector $\boldsymbol{\mu}_0$. Section 4.2 applies Fenchel duality to turn the infinite-dimensional program $\inf_{f \in C \cap D} H_k f$ into a semi-infinite program, whose explicit form for various shape constraints C is derived in Section 4.3. Finally, Section 4.4 explains how to discretize the semi-infinite program in such a way that the confidence level is preserved.

4.1. *Confidence set in the smeared space.* Constructing the confidence set Ξ for the mean $\boldsymbol{\mu}_0$ in the model $\mathbf{y} \sim \text{Poisson}(\boldsymbol{\mu}_0)$ is straightforward. For fixed $\alpha \in (0, 1)$, set $\alpha' = 1 - (1 - \alpha)^{1/n}$. For each $j = 1, \dots, n$, let $[\underline{\mu}_{j,\alpha'}, \bar{\mu}_{j,\alpha'}]$ be a $1 - \alpha'$ confidence interval for $\mu_{0,j}$ constructed using y_j only. Then

$$(12) \quad \begin{aligned} & \mathbb{P}_{f_0}(\underline{\mu}_{1,\alpha'} \leq \mu_{0,1} \leq \bar{\mu}_{1,\alpha'}, \dots, \underline{\mu}_{n,\alpha'} \leq \mu_{0,n} \leq \bar{\mu}_{n,\alpha'}) \\ &= \prod_{j=1}^n \mathbb{P}_{f_0}(\underline{\mu}_{j,\alpha'} \leq \mu_{0,j} \leq \bar{\mu}_{j,\alpha'}) \geq (1 - \alpha')^n = 1 - \alpha, \end{aligned}$$

because the counts y_j in disjoint bins are independent. Thus

$$(13) \quad \Xi \equiv [\underline{\mu}_{1,\alpha'}, \bar{\mu}_{1,\alpha'}] \times \dots \times [\underline{\mu}_{n,\alpha'}, \bar{\mu}_{n,\alpha'}]$$

is a $1 - \alpha$ simultaneous confidence set for $\boldsymbol{\mu}_0$.

The binwise confidence intervals $[\underline{\mu}_{j,\alpha'}, \bar{\mu}_{j,\alpha'}]$ can be formed using the Garwood construction (Garwood, 1936). For bin j , the interval endpoints are

$$(14) \quad \begin{aligned} \underline{\mu}_{j,\alpha'} &= \begin{cases} \frac{1}{2} F_{\chi^2}^{-1}\left(\frac{\alpha'}{2}; 2y_j\right), & y_j > 0, \\ 0, & y_j = 0, \end{cases} \\ \text{and} \\ \bar{\mu}_{j,\alpha'} &= \frac{1}{2} F_{\chi^2}^{-1}\left(1 - \frac{\alpha'}{2}; 2(y_j + 1)\right), \end{aligned}$$

where $F_{\chi^2}^{-1}(\cdot; k)$ is the inverse cumulative distribution function of the χ^2 distribution with k degrees of freedom. These intervals are guaranteed to have confidence level $1 - \alpha'$, although the actual coverage probability is strictly greater than $1 - \alpha'$ for any finite $\mu_{0,j}$, a consequence of the discreteness of the Poisson distribution.

It will be convenient to write the hyperrectangle Ξ as a centered set: $\Xi = \{\tilde{\mathbf{y}} + \boldsymbol{\xi} \in \mathbb{R}^n : \|\text{diag}(\boldsymbol{\ell})^{-1} \boldsymbol{\xi}\|_\infty \leq 1\}$, where, for each $j = 1, \dots, n$, $\tilde{y}_j \equiv (\underline{\mu}_{j,\alpha'} + \bar{\mu}_{j,\alpha'})/2$ and $\ell_j \equiv (\bar{\mu}_{j,\alpha'} - \underline{\mu}_{j,\alpha'})/2$.

4.2. *Strict bounds via duality.* To compute the lower bound $v(\mathcal{P}) \equiv \inf_{f \in C \cap D} H_k f$, we follow the approach of Stark (1992) and apply Fenchel duality. The Fenchel dual (Luenberger, 1969, Section 7.12) of the problem is

$$(15) \quad v(\mathcal{D}) \equiv \sup_{f^* \in C^* \cap D^*} \left\{ \inf_{f \in D} f^*[f] + \inf_{f \in C} (H_k - f^*)[f] \right\},$$

where

$$(16) \quad C^* \equiv \left\{ f^* \in F^* : \inf_{f \in C} (H_k - f^*)[f] > -\infty \right\},$$

$$(17) \quad D^* \equiv \left\{ f^* \in F^* : \inf_{f \in D} f^*[f] > -\infty \right\},$$

and F^* is the algebraic dual space of F , that is, the set of all linear functionals on F . The two problems satisfy weak duality, $v(\mathcal{P}) \geq v(\mathcal{D})$, and, when strict inequality holds, solving the dual problem gives a conservative lower bound. Under regularity conditions (see Luenberger (1969, Section 7.12)), one can establish strong duality, $v(\mathcal{P}) = v(\mathcal{D})$, in which case the bound from the dual is tight. (It is enough that F is normed, the forward functionals $\{K_j\}$ are continuous with respect to the induced dual norm, and the primal problem is feasible; see Stark (1992, Section 10.1).)

The Fenchel dual can be written using a finite-dimensional unknown, which simplifies numerical solution: the set D^* consists of functionals that are linear combinations of the forward functionals K_j ,

$$(18) \quad D^* = \{f^* \in F^* : f^* = \boldsymbol{\nu} \cdot \mathcal{K}, \boldsymbol{\nu} \in \mathbb{R}^n\},$$

where $\boldsymbol{\nu} \cdot \mathcal{K} \equiv \sum_{j=1}^n \nu_j K_j$ (Stark, 1992; Backus, 1970). The dual problem is therefore

$$(19) \quad v(\mathcal{D}) = \sup_{\boldsymbol{\nu} \in \mathbb{R}^n : \boldsymbol{\nu} \cdot \mathcal{K} \in C^*} \left\{ \inf_{f \in D} (\boldsymbol{\nu} \cdot \mathcal{K})[f] + \inf_{f \in C} (H_k - \boldsymbol{\nu} \cdot \mathcal{K})[f] \right\}.$$

The first term can be expressed in closed form: when $f \in D$,

$$(20) \quad \begin{aligned} (\boldsymbol{\nu} \cdot \mathcal{K})[f] &= \boldsymbol{\nu}^T (\tilde{\boldsymbol{y}} + \boldsymbol{\xi}) \geq \boldsymbol{\nu}^T \tilde{\boldsymbol{y}} - |\boldsymbol{\nu}^T \boldsymbol{\xi}| \\ &= \boldsymbol{\nu}^T \tilde{\boldsymbol{y}} - |(\text{diag}(\boldsymbol{\ell})\boldsymbol{\nu})^T (\text{diag}(\boldsymbol{\ell})^{-1}\boldsymbol{\xi})| \\ &\geq \boldsymbol{\nu}^T \tilde{\boldsymbol{y}} - \|\text{diag}(\boldsymbol{\ell})\boldsymbol{\nu}\|_1 \|\text{diag}(\boldsymbol{\ell})^{-1}\boldsymbol{\xi}\|_\infty \\ &\geq \boldsymbol{\nu}^T \tilde{\boldsymbol{y}} - \|\boldsymbol{\nu}\|_1^\ell, \end{aligned}$$

where $\|\boldsymbol{\nu}\|_1^\ell \equiv \|\text{diag}(\boldsymbol{\ell})\boldsymbol{\nu}\|_1$ denotes the weighted 1-norm, and we are considering the n -tuple $\boldsymbol{\nu}$ to be a column vector. Hence,

$$(21) \quad \inf_{f \in D} (\boldsymbol{\nu} \cdot \mathcal{K})[f] \geq \boldsymbol{\nu}^T \tilde{\boldsymbol{y}} - \|\boldsymbol{\nu}\|_1^\ell.$$

When the forward functionals $\{K_j\}$ are linearly independent, an argument similar to that of Stark (1992, Section 5) shows that the lower bound in Equation (21) is sharp.

To summarize, we have established the inequality

$$(22) \quad \inf_{f \in C \cap D} H_k f \geq \sup_{\nu \in \mathbb{R}^n, \nu \cdot \mathcal{K} \in C^*} \left\{ \nu^T \tilde{\mathbf{y}} - \|\nu\|_1^\ell + \inf_{f \in C} (H_k - \nu \cdot \mathcal{K})[f] \right\},$$

which holds with equality under regularity conditions. If the regularity conditions are not satisfied, the solution of the dual problem still gives a valid, conservative bound.

We now turn to the last term in Equation (22). For the shape constraints we consider, C is a convex cone, i.e., C is convex and for all $\gamma \geq 0$, if $f \in C$ then $\gamma f \in C$. It then follows from Stark (1992, Section 6.2) that

$$(23) \quad \begin{aligned} C^* &= \{f^* \in F^* : \inf_{f \in C} (H_k - f^*)[f] = 0\} \\ &= \{f^* \in F^* : (H_k - f^*)[f] \geq 0, \forall f \in C\}, \end{aligned}$$

in which case the dual problem simplifies to

$$(24) \quad \sup_{\nu \in \mathbb{R}^n} \left\{ \nu^T \tilde{\mathbf{y}} - \|\nu\|_1^\ell \right\} \quad \text{subject to} \quad (H_k - \nu \cdot \mathcal{K})[f] \geq 0, \forall f \in C.$$

The lower bound is hence given by a semi-infinite program with an n -dimensional free variable and an infinite set of inequality constraints.

4.3. *Constraints of the dual program.* We consider the following constraints C on the shape of the true intensity f :

- (P) f is positive;
- (D) f is positive and decreasing;
- (C) f is positive, decreasing and convex.

The positivity constraint holds for any intensity function f , while monotonicity and convexity are expected to hold for steeply falling particle spectra.

We now derive the explicit form of the constraint

$$(25) \quad (H_k - \nu \cdot \mathcal{K})[f] \geq 0, \quad \forall f \in C,$$

for each of these sets. First, re-write the left-hand side:

$$(H_k - \nu \cdot \mathcal{K})[f] = H_k f - \sum_{j=1}^n \nu_j K_j f = \int_{T_k} f(t) dt - \sum_{j=1}^n \nu_j \int_T k_j(t) f(t) dt$$

$$(26) \quad = \int_T \underbrace{\left(\mathbf{1}_{T_k}(t) - \sum_{j=1}^n \nu_j k_j(t) \right)}_{\equiv h_k(t)} f(t) dt = \int_T h_k(t) f(t) dt.$$

Hence the dual constraint becomes $\int_T h_k(t) f(t) dt \geq 0$ for all $f \in C$, where $h_k(t) \equiv \mathbf{1}_{T_k}(t) - \sum_{j=1}^n \nu_j k_j(t)$ and $\mathbf{1}_{T_k}$ is the indicator function of T_k .

For simplicity, assume that the space F consists of twice continuously differentiable functions on T (this assumption can be relaxed, at least for constraints (P) and (D)). We furthermore assume that the kernel k defined in Equation (2) is continuous on $S \times T$, so $\{k_j\}$ are continuous on T and $\{h_k\}$ are right-continuous on T . Denote the endpoints of the interval T by T_{\min} and T_{\max} . For each shape constraint, the dual constraint can then be equivalently written as a constraint on h_k :

$$(P) \quad h_k(t) \geq 0, \quad \forall t \in T;$$

$$(D) \quad \int_{T_{\min}}^t h_k(t') dt' \geq 0, \quad \forall t \in T;$$

$$(C) \quad \int_{T_{\min}}^t \int_{T_{\min}}^{t'} h_k(t'') dt'' dt' \geq 0, \quad \forall t \in T \text{ and } \int_T h_k(t) dt \geq 0.$$

The result for the positivity constraint (P) follows directly, while the results for (D) and (C) follow from integration by parts. The derivations are given in Appendix A.

Define

$$(27) \quad k_j^*(t) \equiv \int_{T_{\min}}^t k_j(t') dt' \quad \text{and} \quad k_j^{**}(t) \equiv \int_{T_{\min}}^t \int_{T_{\min}}^{t'} k_j(t'') dt'' dt'.$$

Let $\boldsymbol{\nu} \cdot \mathbf{k}(t) \equiv \sum_{j=1}^n \nu_j k_j(t)$ and define $\boldsymbol{\nu} \cdot \mathbf{k}^*(t)$ and $\boldsymbol{\nu} \cdot \mathbf{k}^{**}(t)$ analogously. Substituting the definition of h_k from Equation (26) gives the following constraints on $\boldsymbol{\nu}$ in the optimization problem $\sup_{\boldsymbol{\nu} \in \mathbb{R}^n} \{\boldsymbol{\nu}^T \tilde{\mathbf{y}} - \|\boldsymbol{\nu}\|_1^\ell\}$:

$$(P) \quad \boldsymbol{\nu} \cdot \mathbf{k}(t) \leq L_k^P(t), \quad \forall t \in T;$$

$$(D) \quad \boldsymbol{\nu} \cdot \mathbf{k}^*(t) \leq L_k^D(t), \quad \forall t \in T;$$

$$(C) \quad \boldsymbol{\nu} \cdot \mathbf{k}^{**}(t) \leq L_k^C(t), \quad \forall t \in T, \text{ and } \boldsymbol{\nu} \cdot \mathbf{k}^*(T_{\max}) \leq T_{k,\max} - T_{k,\min},$$

where the bounding functions on the right-hand side are

$$L_k^P(t) = \mathbf{1}_{T_k}(t),$$

$$L_k^D(t) = \begin{cases} 0, & t < T_{k,\min}, \\ t - T_{k,\min}, & T_{k,\min} \leq t < T_{k,\max}, \\ T_{k,\max} - T_{k,\min}, & t \geq T_{k,\max}, \end{cases}$$

$$L_k^C(t) = \begin{cases} 0, & t < T_{k,\min}, \\ \frac{1}{2}(t - T_{k,\min})^2, & T_{k,\min} \leq t < T_{k,\max}, \\ \frac{1}{2}(T_{k,\max} - T_{k,\min})^2 + (T_{k,\max} - T_{k,\min})(t - T_{k,\max}), & t \geq T_{k,\max}. \end{cases}$$

For positivity, the constraint is piecewise constant; for monotonicity, it consists of two constant parts connected by a linear part; and for convexity, it has a constant and a linear part connected by a quadratic part.

These are explicit expressions for the constraints of the semi-infinite dual program corresponding to the lower bound $\underline{\lambda}_k = \inf_{f \in C \cap D} H_k f$. Similar reasoning shows that the upper bound $\bar{\lambda}_k = \sup_{f \in C \cap D} H_k f = -\inf_{f \in C \cap D} -H_k f$ is bounded from above by $\inf_{\boldsymbol{\nu} \in \mathbb{R}^n} -\{\boldsymbol{\nu}^T \bar{\boldsymbol{y}} - \|\boldsymbol{\nu}\|_1^\ell\}$ subject to the constraints:

$$\begin{aligned} \text{(P)} \quad & \boldsymbol{\nu} \cdot \mathbf{k}(t) \leq -L_k^P(t), \quad \forall t \in T; \\ \text{(D)} \quad & \boldsymbol{\nu} \cdot \mathbf{k}^*(t) \leq -L_k^D(t), \quad \forall t \in T; \\ \text{(C)} \quad & \boldsymbol{\nu} \cdot \mathbf{k}^{**}(t) \leq -L_k^C(t), \quad \forall t \in T, \text{ and } \boldsymbol{\nu} \cdot \mathbf{k}^*(T_{\max}) \leq T_{k,\min} - T_{k,\max}. \end{aligned}$$

4.4. *Discretizing the dual constraints.* To compute the dual bounds requires discretizing the infinite set of constraints. To do so, introduce a grid $\{t_i\}_{i=1}^{m+1}$ on T consisting of $m+1$ grid points, with $m \gg p$. Assume that $t_1 = T_{\min}$, $t_{m+1} = T_{\max}$, and that there is a grid point at each boundary between the true bins $\{T_k\}_{k=1}^p$.

Imagine imposing the constraints just at the grid points $\{t_i\}_{i=1}^{m+1}$. For example, the discretized version of the constraint (P) would be

$$(28) \quad \boldsymbol{\nu} \cdot \mathbf{k}(t_i) \leq \pm L_k^P(t_i), \quad i = 1, \dots, m+1.$$

This is not conservative: the feasible set for Equation (28) is *larger* than that of the original constraint, so the resulting confidence interval could be too short.

To guarantee that the actual confidence level is at least $1 - \alpha$, we need to discretize the constraints in such a way that the discretized feasible set is a subset of the original feasible set for the infinite set of constraints. This requires taking into account the behavior of the constraints between the grid points. We use the grid $\{t_i\}_{i=1}^{m+1}$ to find a convenient upper bound for the left-hand side of the constraint relations and then ensure that this upper bound is everywhere below the right-hand side functions $\pm L_k^P$, $\pm L_k^D$ or $\pm L_k^C$.

Consider the constraint (P) for the lower bound. For each $j = 1, \dots, n$, write $\nu_j = \nu_j^+ - \nu_j^-$ with $\nu_j^+, \nu_j^- \geq 0$, and define the column vectors $\boldsymbol{\nu}^+$ and

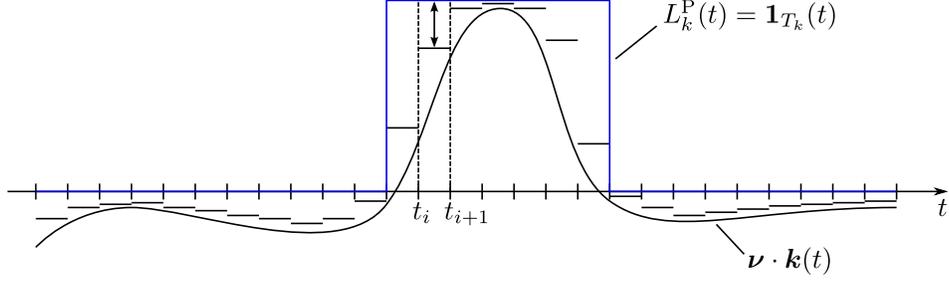


FIG 5. Conservative discretization of the dual constraint $\boldsymbol{\nu} \cdot \mathbf{k}(t) \leq L_k^P(t)$, $\forall t \in T$, using a constant upper bound on each interval $[t_i, t_{i+1})$.

$\boldsymbol{\nu}^-$ in the obvious way, so that $\boldsymbol{\nu} = \boldsymbol{\nu}^+ - \boldsymbol{\nu}^-$. For any $t \in [t_i, t_{i+1})$,

$$\begin{aligned}
 (29) \quad \boldsymbol{\nu} \cdot \mathbf{k}(t) &= \boldsymbol{\nu}^+ \cdot \mathbf{k}(t) - \boldsymbol{\nu}^- \cdot \mathbf{k}(t) \\
 &\leq \sum_{j=1}^n \nu_j^+ \underbrace{\sup_{\xi \in [t_i, t_{i+1})} k_j(\xi)}_{\equiv \bar{\rho}_{i,j}} - \sum_{j=1}^n \nu_j^- \underbrace{\inf_{\xi \in [t_i, t_{i+1})} k_j(\xi)}_{\equiv \underline{\rho}_{i,j}} \\
 &= \sum_{j=1}^n \nu_j^+ \bar{\rho}_{i,j} - \sum_{j=1}^n \nu_j^- \underline{\rho}_{i,j}.
 \end{aligned}$$

This bounds $\boldsymbol{\nu} \cdot \mathbf{k}(t)$ on $[t_i, t_{i+1})$ by a quantity that is constant with respect to t . Since $L_k^P(t)$ is also constant on $[t_i, t_{i+1})$, if we impose the constraints

$$(30) \quad \sum_{j=1}^n \nu_j^+ \bar{\rho}_{i,j} - \sum_{j=1}^n \nu_j^- \underline{\rho}_{i,j} \leq L_k^P(t_i), \quad i = 1, \dots, m,$$

then the original dual constraint will hold: $\boldsymbol{\nu} \cdot \mathbf{k}(t) \leq L_k^P(t)$, $\forall t \in T$. Figure 5 illustrates this construction.

Define

$$(31) \quad \mathbf{A} \equiv \begin{bmatrix} \bar{\rho}_{1,1} & \cdots & \bar{\rho}_{1,n} & -\underline{\rho}_{1,1} & \cdots & -\underline{\rho}_{1,n} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \bar{\rho}_{m,1} & \cdots & \bar{\rho}_{m,n} & -\underline{\rho}_{m,1} & \cdots & -\underline{\rho}_{m,n} \end{bmatrix} \quad \text{and} \quad \tilde{\boldsymbol{\nu}} \equiv \begin{bmatrix} \boldsymbol{\nu}^+ \\ \boldsymbol{\nu}^- \end{bmatrix},$$

and let $\mathbf{b}_k^P \in \mathbb{R}^m$ denote the column vector with components $b_{k,i}^P = L_k^P(t_i) = \mathbf{1}_{T_k}(t_i)$, $i = 1, \dots, m$. Then the discretized dual constraint can be written $\mathbf{A}\tilde{\boldsymbol{\nu}} \leq \mathbf{b}_k^P$.

Since $\boldsymbol{\nu} = \boldsymbol{\nu}^+ - \boldsymbol{\nu}^- = \mathbf{D}\tilde{\boldsymbol{\nu}}$ with $\mathbf{D} \equiv [\mathbf{I}_{n \times n} \quad -\mathbf{I}_{n \times n}]$ and

$$(32) \quad \|\boldsymbol{\nu}\|_1^\ell = \sum_{j=1}^n \ell_j |\nu_j| \leq \sum_{j=1}^n \ell_j (\nu_j^+ + \nu_j^-) = \tilde{\boldsymbol{\ell}}^T \tilde{\boldsymbol{\nu}} \quad \text{with} \quad \tilde{\boldsymbol{\ell}} \equiv \begin{bmatrix} \boldsymbol{\ell} \\ \boldsymbol{\ell} \end{bmatrix},$$

any feasible point of the linear program

$$\begin{aligned} & \sup_{\tilde{\boldsymbol{\nu}} \in \mathbb{R}^{2n}} && (\mathbf{D}^T \tilde{\boldsymbol{y}} - \tilde{\boldsymbol{\ell}})^T \tilde{\boldsymbol{\nu}} \\ & \text{subject to} && \mathbf{A}\tilde{\boldsymbol{\nu}} \leq \mathbf{b}_k^P, \\ & && \tilde{\boldsymbol{\nu}} \geq \mathbf{0}, \end{aligned}$$

yields a conservative lower bound for the k th element of $\boldsymbol{\lambda}_0$ subject to the positivity constraint (P). Similarly, any feasible point of the linear program

$$\begin{aligned} & \inf_{\tilde{\boldsymbol{\nu}} \in \mathbb{R}^{2n}} && -(\mathbf{D}^T \tilde{\boldsymbol{y}} - \tilde{\boldsymbol{\ell}})^T \tilde{\boldsymbol{\nu}} \\ & \text{subject to} && \mathbf{A}\tilde{\boldsymbol{\nu}} \leq -\mathbf{b}_k^P, \\ & && \tilde{\boldsymbol{\nu}} \geq \mathbf{0}, \end{aligned}$$

is a conservative upper bound.

A similar approach allows us to discretize the monotonicity constraint (D) and the convexity constraint (C) conservatively. Deriving these discretizations is somewhat more complicated: we use a first-order Taylor expansion of the kernels $\{k_j\}$ for (D), and a second-order Taylor expansion for (C). For (D), the discretized dual is again a linear program, while for (C) the discretized dual involves optimizing a linear objective function subject to a finite set of nonlinear constraints. Details of the discretizations of (D) and (C) are in Appendix B.

5. Simulation study.

5.1. *Experiment setup.* We demonstrate the shape-constrained strict bounds confidence intervals using a simulation study that mimics unfolding the inclusive jet transverse momentum spectrum (CMS Collaboration, 2011, 2013a) in the Compact Muon Solenoid (CMS) experiment (CMS Collaboration, 2008) at the LHC. A *jet* is a collimated stream of energetic particles, the experimental signature of a quark or a gluon created in proton-proton collisions at the LHC. The jet transverse momentum spectrum describes the average number of jets as a function of their transverse momentum p_\perp , their momentum in the direction perpendicular to the proton beam. The transverse momentum is measured in units of electron volts (eV). Measuring this spectrum is an important test of the Standard Model of particle physics and can be used to constrain free parameters of the theory.

We simulate the data using the particle-level intensity function

$$(33) \quad f_0(p_\perp) = LN_0 \left(\frac{p_\perp}{\text{GeV}} \right)^{-\alpha} \left(1 - \frac{2}{\sqrt{s}} p_\perp \right)^\beta e^{-\gamma/p_\perp}, \quad 0 < p_\perp \leq \frac{\sqrt{s}}{2}.$$

Here $L > 0$ is the integrated luminosity (a measure of the amount of collisions produced in the accelerator, measured in inverse barns, b^{-1}), \sqrt{s} is the center-of-mass energy of the proton-proton collisions, and N_0 , α , β , and γ are positive parameters. This parameterization is motivated by physical considerations and was used in early inclusive jet analyses at the LHC (CMS Collaboration, 2011).

When the jets are reconstructed using calorimeter information, the smearing can be modeled as additive Gaussian noise with zero mean and variance $\sigma(p_\perp)^2$ satisfying

$$(34) \quad \left(\frac{\sigma(p_\perp)}{p_\perp} \right)^2 = \left(\frac{C_1}{\sqrt{p_\perp}} \right)^2 + \left(\frac{C_2}{p_\perp} \right)^2 + C_3^2,$$

where C_i , $i = 1, 2, 3$, are fixed positive constants (CMS Collaboration, 2010). Let p'_\perp denote the smeared transverse momentum. The smeared intensity function is the convolution

$$(35) \quad g_0(p'_\perp) = \int_T N(p'_\perp - p_\perp | 0, \sigma(p_\perp)^2) f_0(p_\perp) dp_\perp, \quad p'_\perp \in S,$$

and the unfolding problem becomes a heteroscedastic deconvolution problem for Poisson point process observations, with the forward kernel given by $k(p'_\perp, p_\perp) = N(p'_\perp - p_\perp | 0, \sigma(p_\perp)^2)$.

At the center-of-mass energy $\sqrt{s} = 7$ TeV and in the central part of the CMS detector, realistic values for the parameters of $f_0(p_\perp)$ are given by $N_0 = 10^{17}$ fb/GeV, $\alpha = 5$, $\beta = 10$, and $\gamma = 10$ GeV and for the parameters of $\sigma(p_\perp)$ by $C_1 = 1$ GeV^{1/2}, $C_2 = 1$ GeV, and $C_3 = 0.05$ (M. Voutilainen, personal communication, 2012). We furthermore set $L = 5.1$ fb⁻¹, which corresponds to the size of the CMS 7 TeV dataset.

The true intensity f_0 obviously satisfies the positivity constraint (P). For the parameter values in our simulations, f_0 is decreasing for $p_\perp \gtrsim 2.0$ and convex for $p_\perp \gtrsim 2.8$. In other words, for intermediate and large p_\perp values—the main focus in inclusive jet analyses (CMS Collaboration, 2013a)—the true intensity satisfies the monotonicity constraint (D) and the convexity constraint (C). In general, physical considerations lead one to expect the p_\perp spectrum to satisfy these constraints, at least for intermediate values of p_\perp .

In the simulations reported here, the true and smeared spaces are $T = S = [400 \text{ GeV}, 1000 \text{ GeV}]$, and we partition both spaces into $n = p = 30$

equal-width bins. The dual constraints are discretized using $m + 1 = 10p + 1$ uniformly spaced grid points. This corresponds to subdividing each true bin into 10 sub-bins. All the experiments were implemented in MATLAB (version R2014a) using its Optimization Toolbox (Mathworks, 2014). Appendix C gives more implementation details. The MATLAB scripts used to produce the results are available at <https://github.com/mkuusela/ShapeConstrainedUnfolding>.

5.2. Results.

5.2.1. *Shape-constrained strict bounds confidence intervals.* Figure 6 shows the true intensity f_0 and the 95 % confidence intervals for λ_0 for the different shape constraints. The true value of λ_0 is shown by the horizontal lines. Results are plotted on both linear and logarithmic scales; the binned quantities were converted to the intensity scale by dividing them by the bin width. The strict bounds confidence intervals cover λ_0 in every bin. The shape constraints have a marked influence on the interval lengths. The positivity-constrained intervals are fairly wide, with zero lower bound in every bin (they are still orders of magnitude shorter than unconstrained intervals). But monotonicity and convexity constraints yield much shorter intervals and correspondingly sharper physical inferences.

Figure 7 shows the dual constraints $\pm L_{10}^P$, $\pm L_{10}^D$ and $\pm L_{10}^C$ (see Section 4.3) and the corresponding optimal solutions for bounding the 10th true bin. Despite the conservative discretization, the optimal solutions can be very close to the constraints. (For the positive lower bound, the optimal solution is $\tilde{\nu} = \mathbf{0}$, which is consistent with the lower bound $\lambda_{10} = 0$.) We also compared the lengths of the conservatively discretized intervals to those of nonconservative intervals, where the dual constraint is only imposed at the grid points $\{t_i\}_{i=1}^{m+1}$ (cf. Equation (28)). We found that the conservative intervals are only slightly wider than the nonconservative ones: For the monotonicity and convexity constraints, the difference was less than 1 % in most bins. For the bin where the difference was the largest, the conservative intervals were 13.2 %, 2.4 %, and 2.0 % longer for the positivity, monotonicity, and convexity constraints, respectively. Evidently, the conservative discretization is not excessively pessimistic.

By construction, the simultaneous coverage probability of the strict bounds confidence intervals in Figure 6 is at least 95 %, but in practice it can be much greater than this. To study the actual coverage probability, we repeated the study for 1,000 independent realizations of the data-generating process. For every realization, the intervals covered λ_0 : in this example, the empirical

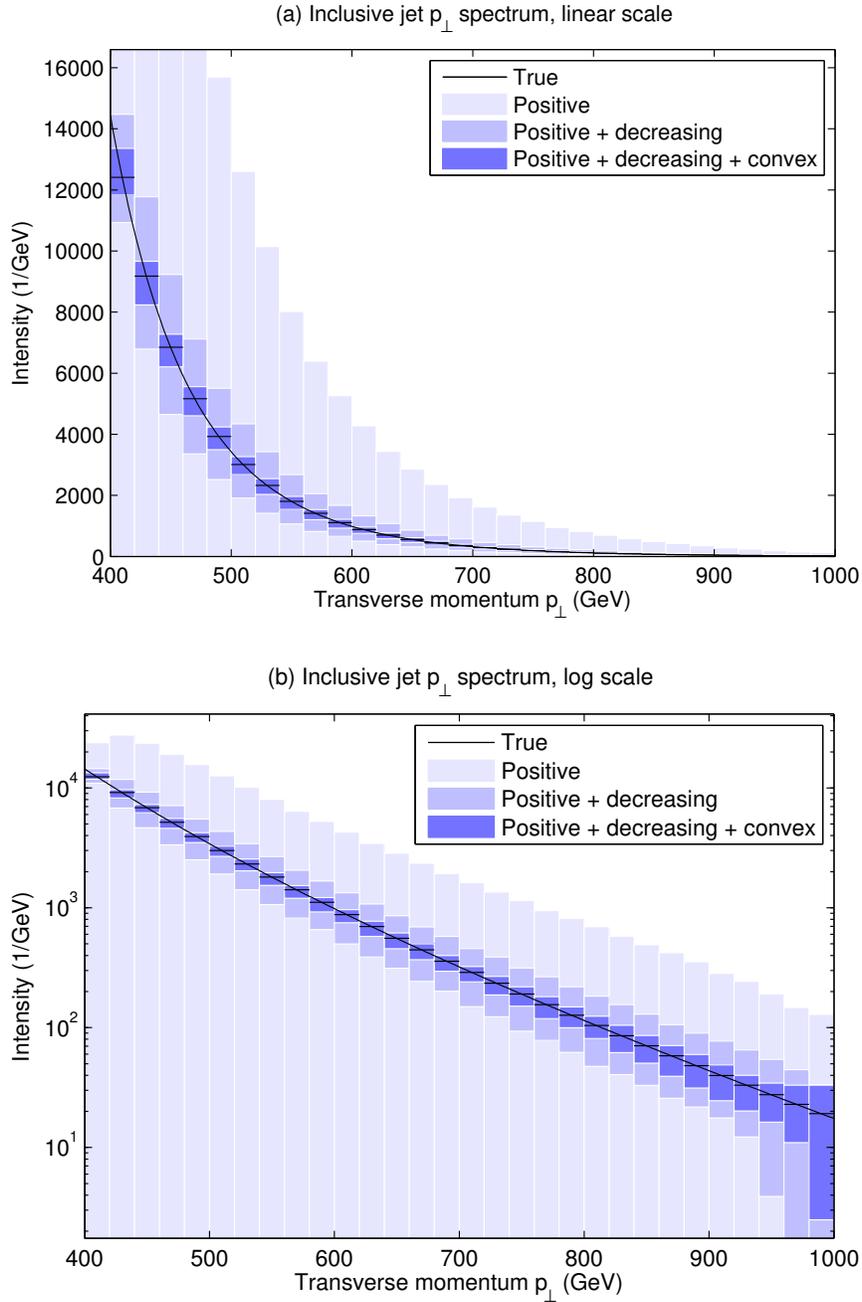


FIG 6. *Shape-constrained strict bounds confidence intervals for the binned inclusive jet transverse momentum spectrum. Figure (a) shows the results on a linear scale and Figure (b) on a logarithmic scale. These intervals are guaranteed to have 95 % simultaneous finite-sample coverage, and indeed do cover the truth across the whole spectrum.*

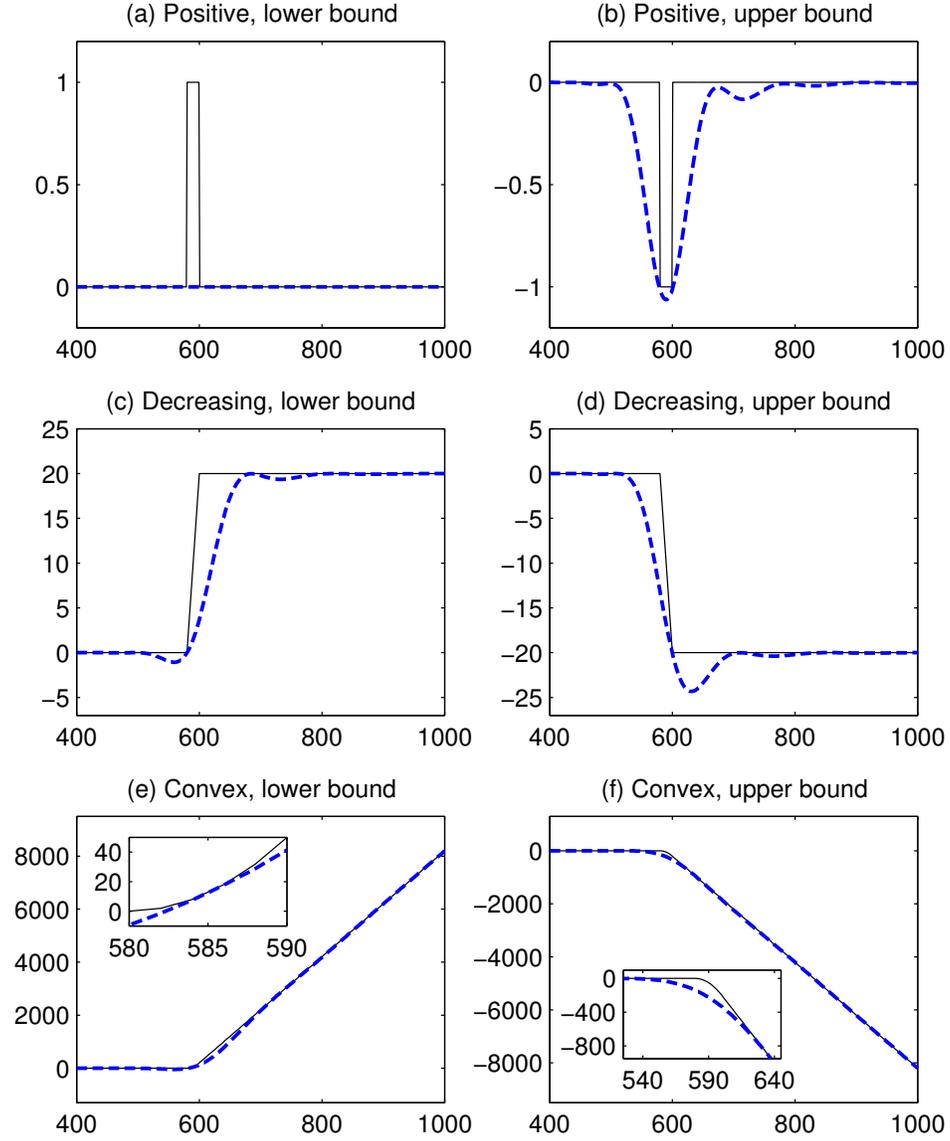


FIG 7. The dual constraints $\pm L_{10}^P$, $\pm L_{10}^D$, and $\pm L_{10}^C$ (solid lines) and the optimal solutions (dashed lines) for bounding the 10th true bin, for the three combinations of shape constraints. Insets in Figures (e) and (f) show the quadratic part of the constraint in greater detail.

TABLE 1

Empirical simultaneous coverage of the 95 % shape-constrained strict bounds for the inclusive jet transverse momentum spectrum of Equation (33), the linearly decreasing spectrum of Equation (36) and the constant spectrum of Equation (37). The columns correspond to the different shape constraints and the rows to the different true spectra. The uncertainties in the parentheses are 95 % Clopper–Pearson intervals for the true coverage probability, to account for the uncertainty in estimating the coverage using 1,000 replications. The shape-constrained intervals are less conservative when the true spectrum is on the boundary of the constraint set.

	Positive	Positive and decreasing	Positive, decreasing and convex
Inclusive jets	1.000 (0.996, 1.000)	1.000 (0.996, 1.000)	1.000 (0.996, 1.000)
Linearly decr.	1.000 (0.996, 1.000)	1.000 (0.996, 1.000)	0.969 (0.956, 0.979)
Constant	1.000 (0.996, 1.000)	0.947 (0.931, 0.960)	0.945 (0.929, 0.958)

simultaneous coverage probability is 100 % (the 95 % Clopper–Pearson confidence interval for the true coverage probability is [0.996, 1.000]).

To demonstrate that there are elements in the constraint set C for which the intervals are less conservative, we also performed the same coverage study for a linearly decreasing intensity function

$$(36) \quad f_0(p_\perp) = \text{const} \cdot (T_{\max} - p_\perp), \quad p_\perp \in [T_{\min}, T_{\max}],$$

and a constant intensity function

$$(37) \quad f_0(p_\perp) = \text{const}, \quad p_\perp \in [T_{\min}, T_{\max}].$$

In both cases, $[T_{\min}, T_{\max}] = [400 \text{ GeV}, 1000 \text{ GeV}]$ and the intensity function was scaled so that the expected total number of particle-level events was the same as for the inclusive jet spectrum. Apart from the particle-level intensity functions, the simulation setup was the same as in the inclusive jet experiments.

To make the coverage study computationally feasible, the convexity-constrained intervals were formed by imposing the dual constraint only on a grid; see Equation (61) in Appendix C.1. Since the intervals with the full conservative discretization are only slightly wider than these nonconservatively discretized intervals (see above), the coverage of the full intervals should not be much higher than the values reported here. The positivity- and monotonicity-constrained intervals were computed using the full conservative discretization.

The empirical simultaneous coverage for the different intensity functions and shape constraints is given in Table 1. The coverage of the monotonicity-constrained intervals is close to the nominal value when the data are gener-

ated from the constant intensity (37). Similarly, the coverage of the convexity-constrained intervals is close to nominal for the linearly decreasing intensity (36) and the constant intensity (37). This suggests that the strict bounds are less conservative when the true intensity f_0 lies on the boundary of the constraint set C , although further studies are needed to better understand how the coverage depends on properties of f_0 .

5.2.2. Undercoverage of existing unfolding methods. Currently, the two most common unfolding methods in LHC data analyses are the SVD variant of Tikhonov regularization (Höcker and Kartvelishvili, 1996) and the D’Agostini iteration (D’Agostini, 1995), which is an EM iteration with early stopping. As explained in Section 1, these estimators are regularized by shrinking the solution towards a Monte Carlo prediction λ^{MC} of the quantity of interest λ_0 . The methods also rely on discretizing the forward mapping using the Monte Carlo event generator. When the Monte Carlo prediction differs from the truth—i.e., essentially always—current procedures for constructing confidence intervals, which only account for the variance of $\hat{\lambda}$ and not for the bias, can have actual coverage probabilities far lower than their nominal confidence levels. This section shows the severity of the problem by unfolding the inclusive jet p_\perp spectrum with the SVD and D’Agostini methods.

In this simulation, the Monte Carlo event generator predicts a spectrum that is a slight perturbation of the true f_0 : the Monte Carlo spectrum f^{MC} is given by Equation (33), with parameters $N_0 = 5.5 \cdot 10^{19}$ fb/GeV, $\alpha = 6$, and $\beta = 12$. The rest of the parameters are set to the same values as before. This spectrum falls off slightly faster than f_0 in both the power-law and the energy cutoff terms. The value of N_0 was chosen so that the overall scale of f^{MC} is similar to f_0 ; this value has little effect on the results (it cancels in SVD unfolding). The discretized smearing matrix \mathbf{K} and the MC prediction λ^{MC} are then obtained by substituting f^{MC} into Equations (3) and (9).

We quantify the binwise uncertainty of $\hat{\lambda}$ using nominal $1 - \alpha$ Gaussian confidence intervals:

$$(38) \quad \left[\hat{\lambda}_k - z_{1-\alpha/2} \sqrt{\widehat{\text{Var}}(\hat{\lambda}_k)}, \hat{\lambda}_k + z_{1-\alpha/2} \sqrt{\widehat{\text{Var}}(\hat{\lambda}_k)} \right], \quad k = 1, \dots, p,$$

where $z_{1-\alpha/2}$ is the $1 - \alpha/2$ standard normal quantile and $\widehat{\text{Var}}(\hat{\lambda}_k)$ is the estimated variance of $\hat{\lambda}_k$. This is essentially what current unfolding software implements (Adye, 2011; Schmitt, 2012). To obtain simultaneous $1 - \alpha$ confidence sets for the whole histogram λ_0 , we adjust the levels of the binwise intervals using Bonferroni’s inequality, setting the nominal level of each interval to $1 - \alpha/p$. In both the SVD and D’Agostini methods, the variances

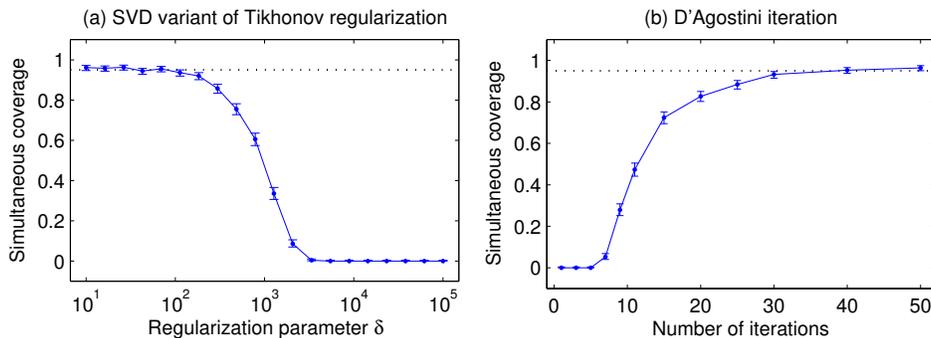


FIG 8. *Simulation estimate of the simultaneous coverage of Bonferroni-adjusted nominal 95 % joint confidence intervals with (a) the SVD variant of Tikhonov regularization and (b) the D’Agostini iteration. The error bars are given by the 95 % Clopper–Pearson intervals and the nominal confidence level is shown by the dotted line. When the regularization is strong, both methods undercover substantially.*

$\widehat{\text{Var}}(\hat{\lambda}_k)$ are estimated using error propagation; for details, see Appendix C, which describes the two methods in more detail.

The coverage properties of the intervals of Equation (38) depend strongly on the regularization strength, which, in the case of Tikhonov regularization, is controlled by the regularization parameter δ (see Equation (4)) and, in the case of the D’Agostini method, by the number of iterations. Figure 8 shows the simultaneous coverage of (nominal) 95 % Bonferroni-corrected normal intervals as a function of the regularization strength estimated from 1,000 independent replications. For weak regularization, the intervals attain their nominal coverage, but for strong regularization, the simultaneous coverage quickly drops to zero.

An obvious question to ask is where along these coverage curves do typical unfolding results lie? Unfortunately, it is impossible to tell: it depends on how good the Monte Carlo predictions f^{MC} are and how the regularization strength is chosen. Currently, most analyses use nonstandard heuristics for choosing the regularization strength, without properly documenting and justifying the criterion used. For instance, ROOUNFOLD documentation (Ade, 2011) recommends simply using four iterations for the D’Agostini method, and many LHC analyses seem to follow this convention. There is no principled reason to use four iterations, which in our simulations would result in *zero* simultaneous coverage.

Even with more principled, data-driven criteria for choosing the regularization strength, reasonable coverage performance is not guaranteed. Fig-

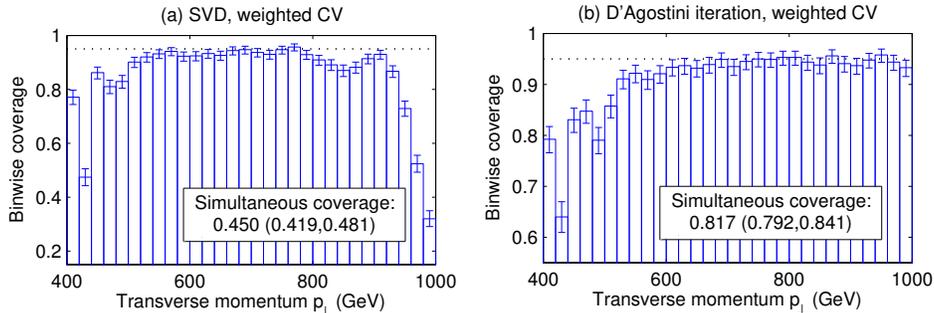


FIG 9. Empirical coverage of the 95 % normal intervals with (a) the SVD variant of Tikhonov regularization and (b) the D’Agostini iteration, when the regularization strength is chosen using weighted cross-validation. The simultaneous coverage is given for Bonferroni-corrected intervals. The uncertainties are the 95 % Clopper–Pearson intervals for the coverage probability based on 1,000 replications. The nominal confidence level is shown by the dotted line. Both methods undercover at small values of the transverse momentum and SVD also at large transverse momenta. In both cases, the simultaneous coverage is below the nominal 95 %, dramatically so for the SVD.

Figure 9 shows binwise and simultaneous empirical coverage of the two methods for 1,000 independent replications when the regularization strength is selected using weighted cross-validation (see Appendix C for details; for D’Agostini, 3 runs where the cross-validation score was still decreasing after 20,000 iterations were omitted from the analysis). All the results are at 95 % nominal confidence. Both methods undercover at small p_{\perp} values; the SVD approach also undercovers at large p_{\perp} values. For the SVD, the smallest binwise coverage is 0.320 (0.291, 0.350) (95 % Clopper–Pearson interval) and for D’Agostini it is 0.640 (0.609, 0.670), much less than the nominal level 0.95. When f^{MC} is further from f_0 , coverage performance becomes even worse. We also tried adding Monte Carlo noise to λ^{MC} , which led to large additional reductions in the coverage.

Caution should be exercised in extrapolating these findings to current LHC unfolding results. Current LHC practice treats the Monte Carlo model dependence as a systematic uncertainty. A typical way to take this uncertainty into account is to compute the unfolded histograms using two or more Monte Carlo event generators and to take the observed differences as an estimate of the systematic uncertainty. The success of this approach depends on how well the Monte Carlo models represent the range of plausible truths and on whether f_0 is in some sense bracketed by these models.

Clearly, the coverage of existing unfolding methods is a delicate function

of the Monte Carlo prediction f^{MC} , the number of true bins p , the regularization method, the regularization strength, and the way these factors are taken into account as systematic uncertainties. It is conceivable that LHC analyses give coverage close to the nominal level, but that would require many fortuitous accidents. It seems impossible to give a rigorous—or even approximate—coverage guarantee for the existing methods. In contrast, shape-constrained strict bounds do not rely on a Monte Carlo prediction of the unknown and do not require the analyst to choose a regularization strength. Moreover, the bounds are guaranteed to have nominal simultaneous coverage for any number of true bins p , provided that f_0 satisfies the shape constraints, a safe assumption for typical steeply falling spectra.

6. Discussion. We have presented a novel approach for unfolding elementary particle spectra that imposes physically justified shape constraints and quantifies the uncertainty of the solution using strict bounds confidence intervals. The resulting intervals have guaranteed simultaneous frequentist coverage whenever the shape constraints hold. This is the case for the important class of unfolding problems with steeply falling spectra. For other classes of problems, other types of shape constraints, such as unimodality, might still provide an attractive way to reduce uncertainty. A natural direction for future work would hence be to extend the approach presented here to m -modal intensities (Hengartner and Stark, 1995). Another useful extension would be to shapes that have a concave part and a convex part with an unknown changepoint between the two.

The type of coverage considered here (simultaneous coverage for the whole histogram λ_0) differs from the binwise coverage traditionally considered in HEP. Binwise confidence intervals can be used to make inferential statements at a single bin, but interpreting the collection of such intervals as a whole is difficult. This issue has also been noted in the HEP literature (Blaubel, 2013, Section 6.6.4). In contrast, simultaneous confidence intervals, such as those in Figure 6, by definition have the property that, under repeated sampling, at least 95 % of the envelopes formed by the intervals will contain the entire true histogram λ_0 . Hence, the confidence envelope can be directly interpreted as a whole.

This is immediately useful for scientists who want to use unfolding results in other analyses. For example, a 5 % goodness-of-fit test of a theory prediction of λ_0 can be performed by simply overlaying the prediction on the 95 % confidence envelope. (Not every spectrum within the envelope fits the data adequately, but any model whose predictions are not wholly contained in the envelope can be rejected at 5 % significance level.) Unfolded

confidence envelopes from two independent experiments can be combined by first making an appropriate multiplicity correction and then following the strict bounds construction of Figure 4 for an identity smearing operator. The strict bounds construction can also be used to extract further information from the unfolded confidence envelopes. For example, one should be able to construct confidence envelopes for parton distribution functions by following the construction for appropriate forward operators and multiplicity-corrected unfolded input spectra. (Think of Figure 4 with the smeared space replaced by the space of unfolded inputs and the true space by the space of parton distribution functions.) To enable this kind of re-use, it may be helpful to compute and publish unfolded confidence envelopes at a variety of confidence levels.

The present work assumes that the smearing kernel $k(s, t)$ is known perfectly, while in fact $k(s, t)$ is usually determined using some auxiliary measurements or simulations, and hence is uncertain. If rigorous uncertainty quantification for $k(s, t)$ is available, it is possible to incorporate that uncertainty into the strict bounds construction (Stark, 1992, Section 9.2). We expect this to be feasible at least when there is a physics-driven parametric model for $k(s, t)$, such as the calorimeter response in Equation (35). Rigorous treatment of the nonparametric case, where estimating $k(s, t)$ essentially becomes a nonparametric quantile regression problem, appears more challenging. Of course, these considerations also affect existing unfolding methods, which treat uncertainty of the smearing kernel using various heuristics.

The user is free to choose the number of smeared bins n and the number of unfolded bins p . As is common in nonparametrics, it is challenging to give general guidelines about how to choose them. However, the proposed approach guarantees conservative coverage for *any* choice of n and p , while the coverage of alternative nonparametric confidence intervals generally depends heavily on the choice of tuning parameters, such as a smoothing parameter or a bandwidth, as illustrated above for the SVD and D’Agostini methods. Our simulations set $n = p$ following the standard practice in HEP, hence enabling a realistic comparison with existing unfolding methods. One might instead choose n to be as large as available computational resources allow (particle physics data are typically not intrinsically binned). One could even eliminate n altogether by constructing an unbinned confidence envelope around the smeared empirical cumulative distribution function using the Dvoretzky–Kiefer–Wolfowitz inequality as in Hengartner and Stark (1992), but this would not be computationally feasible for the large data sets typical in HEP unfolding analyses. In contrast, the choice of p should be guided by physics judgment about the resolution at which one wishes to probe the unfolded

spectrum. In general, the coarser the resolution (i.e. small p), the shorter the confidence intervals will be. The proposed approach is not limited to uniform bin sizes; analogous considerations apply if the bins have unequal widths.

The proposed method potentially has large overcoverage: the intervals may be wider than necessary to attain their nominal simultaneous confidence level. The most important source of slack is presumably the way the set $C \cap D$ is mapped into the finite-dimensional set $[\underline{\lambda}_1, \bar{\lambda}_1] \times \dots \times [\underline{\lambda}_p, \bar{\lambda}_p]$ (see Section 3). In effect, $C \cap D$ is first mapped through $\mathcal{H} : F \rightarrow \mathbb{R}^p, f \mapsto [H_1 f, \dots, H_p f]^T$ and the resulting set $\mathcal{H}(C \cap D)$ is then bounded by a hyperrectangle. Shorter confidence intervals could be obtained by tuning Ξ so that the geometry of $\mathcal{H}(C \cap D)$ matches better to the geometry of hyperrectangles (Stark, 1992, Section 10.2). However, selecting Ξ optimally is an open problem. Alternatively, one could consider bounding $\mathcal{H}(C \cap D)$ using some other geometric form, such as a hyperellipsoid. More generally, the optimality properties of the construction also depend on the discretization of the smeared space (Hengartner and Stark, 1995), and the choice to use histogram binning might not yield optimal asymptotic convergence rates.

Acknowledgments. We warmly thank Victor Panaretos for his support of this work. We are also grateful to Olaf Behnke, Bob Cousins, Tommaso Dorigo, Louis Lyons, Igor Volobouev and Mikko Voutilainen for helpful discussions on unfolding, and to Yoav Zemel, the Associate Editor and the two anonymous reviewers for insightful and detailed feedback on the manuscript. Parts of this work were carried out while MK was visiting the Department of Statistics at the University of California, Berkeley.

APPENDIX A: DERIVATION OF THE DUAL CONSTRAINTS FOR DECREASING AND CONVEX INTENSITIES

This section writes the constraint $\int_T h_k(t) f(t) dt \geq 0, \forall f \in C$, in an equivalent form that does not involve f . Clearly, in the case of the positivity constraint (P), an equivalent constraint is $h_k(t) \geq 0, \forall t \in T$. The derivations for the monotonicity constraint (D) and the convexity constraint (C) employ integration by parts.

A.1. Decreasing intensities. We show that when C is the monotonicity constraint (D),

$$(39) \quad \int_T h_k(t) f(t) dt \geq 0, \forall f \in C \quad \Leftrightarrow \quad \int_{T_{\min}}^t h_k(t') dt' \geq 0, \forall t \in T.$$

Integration by parts yields

$$\begin{aligned}
\int_T h_k(t) f(t) dt &= \int_{T_{\min}}^t h_k(t') dt' f(t) \Big|_{T_{\min}}^{T_{\max}} - \int_T \int_{T_{\min}}^t h_k(t') dt' f'(t) dt \\
(40) \qquad \qquad &= \int_T h_k(t) dt f(T_{\max}) - \int_T \int_{T_{\min}}^t h_k(t') dt' f'(t) dt.
\end{aligned}$$

From this expression it is clear that the right-hand side of Equation (39) implies the left-hand side. To show the opposite implication, assume that $\int_{T_{\min}}^{t^*} h_k(t') dt' < 0$ for some t^* in the interior of T . Then by the continuity of the integral, $\int_{T_{\min}}^t h_k(t') dt' < 0$ for all $t \in (t^* - \delta, t^* + \delta)$ for some $\delta > 0$. Consider a function $d \in C$ that is a strictly positive constant on the interval $[T_{\min}, t^* - \delta]$ and zero on $[t^* + \delta, T_{\max}]$. Substituting d into Equation (40) gives

$$(41) \qquad \int_T h_k(t) d(t) dt = - \int_{t^* - \delta}^{t^* + \delta} \int_{T_{\min}}^t h_k(t') dt' d'(t) dt < 0,$$

a contradiction. Hence $\int_{T_{\min}}^t h_k(t') dt' \geq 0$ for all t in the interior of T and, by the continuity of the integral, for all $t \in T$.

A.2. Convex intensities. We next treat the convexity constraint (C) by showing that

$$\begin{aligned}
(42) \qquad \int_T h_k(t) f(t) dt &\geq 0, \quad \forall f \in C \\
\Leftrightarrow \int_{T_{\min}}^t \int_{T_{\min}}^{t'} h_k(t'') dt'' dt' &\geq 0, \quad \forall t \in T \quad \wedge \quad \int_T h_k(t) dt \geq 0.
\end{aligned}$$

A second application of integration by parts in Equation (40) gives

$$\begin{aligned}
\int_T h_k(t) f(t) dt &= \int_T h_k(t) dt f(T_{\max}) - \int_{T_{\min}}^t \int_{T_{\min}}^{t'} h_k(t'') dt'' dt' f'(t) \Big|_{T_{\min}}^{T_{\max}} \\
&\quad + \int_T \int_{T_{\min}}^t \int_{T_{\min}}^{t'} h_k(t'') dt'' dt' f''(t) dt \\
&= \int_T h_k(t) dt f(T_{\max}) - \int_{T_{\min}}^{T_{\max}} \int_{T_{\min}}^t h_k(t') dt' dt f'(T_{\max}) \\
(43) \qquad \qquad &\quad + \int_T \int_{T_{\min}}^t \int_{T_{\min}}^{t'} h_k(t'') dt'' dt' f''(t) dt.
\end{aligned}$$

From this form one can see that the right-hand side of Equation (42) implies the left-hand side. To show the reverse implication, pick $d \in C$ such that $d(t) = a > 0$ for all $t \in T$. Substituting this into the left-hand side implies that $\int_T h_k(t) dt \geq 0$. Suppose that $\int_{T_{\min}}^{t^*} \int_{T_{\min}}^{t'} h_k(t'') dt'' dt' < 0$ for some t^* in the interior of T . By the continuity of the integral, it follows that $\int_{T_{\min}}^t \int_{T_{\min}}^{t'} h_k(t'') dt'' dt' < 0$ for all $t \in (t^* - \delta, t^* + \delta)$ for some $\delta > 0$. Consider a function $d \in C$ that is linear and strictly decreasing on the interval $[T_{\min}, t^* - \delta]$ and zero on $[t^* + \delta, T_{\max}]$. Substituting d into Equation (43) yields

$$(44) \quad \int_T h_k(t) d(t) dt = \int_{t^* - \delta}^{t^* + \delta} \int_{T_{\min}}^t \int_{T_{\min}}^{t'} h_k(t'') dt'' dt' d''(t) dt < 0,$$

a contradiction. Hence $\int_{T_{\min}}^t \int_{T_{\min}}^{t'} h_k(t'') dt'' dt' \geq 0$ for all t in the interior of T and, by the continuity of the integral, for all $t \in T$.

APPENDIX B: DISCRETIZED CONSTRAINTS FOR DECREASING AND CONVEX INTENSITIES

This section derives conservative discretizations of the dual constraints for the monotonicity constraint (D) and the convexity constraint (C). The strategy follows the procedure of Section 4.4, with appropriate modifications. Discretizing the constraints $\boldsymbol{\nu} \cdot \mathbf{k}^*(t) \leq \pm L_k^D(t)$ or $\boldsymbol{\nu} \cdot \mathbf{k}^{**}(t) \leq \pm L_k^C(t)$ by applying the bound of Equation (29) to k_j^* or k_j^{**} gives up too much because Equation (29) bounds the left-hand side by a constant on each interval $[t_i, t_{i+1})$, while the functions $\pm L_k^D$ or $\pm L_k^C$ on the right-hand side can vary within these intervals. A better approach for the monotonicity constraint (D) is to use a first-order Taylor expansion of k_j^* , which gives a linear upper bound for the left-hand side. For the convexity constraint (C), we use a second-order Taylor expansion of k_j^{**} , which yields a quadratic upper bound.

B.1. Decreasing intensities. We first treat the monotonicity constraint (D). For any $t \in [t_i, t_{i+1})$,

$$(45) \quad \begin{aligned} k_j^*(t) &= k_j^*(t_i) + (k_j^*)'(\xi_j)(t - t_i) \\ &= k_j^*(t_i) + k_j(\xi_j)(t - t_i), \quad \xi_j \in [t_i, t). \end{aligned}$$

This yields

$$\sum_{j=1}^n \nu_j k_j^*(t) = \sum_{j=1}^n \nu_j k_j^*(t_i) + \sum_{j=1}^n \nu_j^+ k_j(\xi_j)(t - t_i) - \sum_{j=1}^n \nu_j^- k_j(\xi_j)(t - t_i)$$

$$\begin{aligned}
&\leq \sum_{j=1}^n \nu_j k_j^*(t_i) + \sum_{j=1}^n \nu_j^+ \sup_{\xi \in [t_i, t_{i+1})} k_j(\xi)(t - t_i) \\
&\quad - \sum_{j=1}^n \nu_j^- \inf_{\xi \in [t_i, t_{i+1})} k_j(\xi)(t - t_i) \\
(46) \quad &= \sum_{j=1}^n \nu_j k_j^*(t_i) + \sum_{j=1}^n \nu_j^+ \bar{\rho}_{i,j}(t - t_i) - \sum_{j=1}^n \nu_j^- \underline{\rho}_{i,j}(t - t_i).
\end{aligned}$$

On $[t_i, t_{i+1})$, this gives a linear upper bound for $\boldsymbol{\nu} \cdot \mathbf{k}^*(t) = \sum_{j=1}^n \nu_j k_j^*(t)$. Since L_k^D is linear on each interval $[t_i, t_{i+1})$, it is enough to enforce the constraint at the endpoints of the interval. By the continuity of L_k^D , we require for each $i = 1, \dots, m$ that

$$(47) \quad \begin{cases} \sum_{j=1}^n \nu_j k_j^*(t_i) \leq \pm L_k^D(t_i), \\ \sum_{j=1}^n \nu_j k_j^*(t_i) + \sum_{j=1}^n \nu_j^+ \bar{\rho}_{i,j} \delta_i - \sum_{j=1}^n \nu_j^- \underline{\rho}_{i,j} \delta_i \leq \pm L_k^D(t_{i+1}), \end{cases}$$

where $\delta_i \equiv t_{i+1} - t_i$. In fact, since $\sum_{j=1}^n \nu_j k_j^*(t)$ is continuous, the first inequality in Equation (47) is redundant: it suffices to enforce the second.

Let $\boldsymbol{\Delta} \equiv \text{diag}(\{\delta_i\}_{i=1}^m)$; let \mathbf{K}^* denote the matrix with elements $K_{i,j}^* = k_j^*(t_i)$, $i = 1, \dots, m$, $j = 1, \dots, n$; and let \mathbf{b}_k^D denote the vector with elements $b_{k,i}^D = L_k^D(t_{i+1})$, $i = 1, \dots, m$. Then $(\mathbf{K}^* \mathbf{D} + \boldsymbol{\Delta} \mathbf{A}) \tilde{\boldsymbol{\nu}} \leq \pm \mathbf{b}_k^D$ is a conservative discretization of the constraint $\boldsymbol{\nu} \cdot \mathbf{k}^*(t) \leq \pm L_k^D(t)$, $\forall t \in T$, where the matrices \mathbf{A} and \mathbf{D} are defined in Section 4.4.

Thus, any feasible point of the linear program

$$\begin{aligned}
&\sup_{\tilde{\boldsymbol{\nu}} \in \mathbb{R}^{2n}} (\mathbf{D}^T \tilde{\boldsymbol{y}} - \tilde{\boldsymbol{\ell}})^T \tilde{\boldsymbol{\nu}} \\
&\text{subject to } (\mathbf{K}^* \mathbf{D} + \boldsymbol{\Delta} \mathbf{A}) \tilde{\boldsymbol{\nu}} \leq \mathbf{b}_k^D, \\
&\quad \tilde{\boldsymbol{\nu}} \geq \mathbf{0},
\end{aligned}$$

yields a conservative lower bound at the k th true bin. A conservative upper bound is given by any feasible point of

$$\begin{aligned}
&\inf_{\tilde{\boldsymbol{\nu}} \in \mathbb{R}^{2n}} -(\mathbf{D}^T \tilde{\boldsymbol{y}} - \tilde{\boldsymbol{\ell}})^T \tilde{\boldsymbol{\nu}} \\
&\text{subject to } (\mathbf{K}^* \mathbf{D} + \boldsymbol{\Delta} \mathbf{A}) \tilde{\boldsymbol{\nu}} \leq -\mathbf{b}_k^D, \\
&\quad \tilde{\boldsymbol{\nu}} \geq \mathbf{0}.
\end{aligned}$$

B.2. Convex intensities. We now treat the convexity constraint (C). For any $t \in [t_i, t_{i+1})$,

$$(48) \quad \begin{aligned} k_j^{**}(t) &= k_j^{**}(t_i) + (k_j^{**})'(t_i)(t - t_i) + \frac{1}{2}(k_j^{**})''(\xi_j)(t - t_i)^2 \\ &= k_j^{**}(t_i) + k_j^*(t_i)(t - t_i) + \frac{1}{2}k_j(\xi_j)(t - t_i)^2, \quad \xi_j \in [t_i, t). \end{aligned}$$

This yields

$$(49) \quad \begin{aligned} \sum_{j=1}^n \nu_j k_j^{**}(t) &= \sum_{j=1}^n \nu_j k_j^{**}(t_i) + \sum_{j=1}^n \nu_j k_j^*(t_i)(t - t_i) + \frac{1}{2} \sum_{j=1}^n \nu_j k_j(\xi_j)(t - t_i)^2 \\ &\leq \sum_{j=1}^n \nu_j k_j^{**}(t_i) + \sum_{j=1}^n \nu_j k_j^*(t_i)(t - t_i) \\ &\quad + \frac{1}{2} \sum_{j=1}^n \left(\nu_j^+ \bar{\rho}_{i,j} - \nu_j^- \underline{\rho}_{i,j} \right) (t - t_i)^2, \end{aligned}$$

where the bound is established as in Equation (46). This bounds $\boldsymbol{\nu} \cdot \mathbf{k}^{**}(t) = \sum_{j=1}^n \nu_j k_j^{**}(t)$ on $[t_i, t_{i+1})$ from above by a parabola. If we ensure that this parabola lies below $\pm L_k^C$ for every $t \in [t_i, t_{i+1})$, the resulting feasible set is a subset of the original set described by infinitely many constraints. We need to require

$$(50) \quad \begin{aligned} \pm L_k^C(t) - \sum_{j=1}^n \nu_j k_j^{**}(t_i) - \sum_{j=1}^n \nu_j k_j^*(t_i)(t - t_i) \\ - \frac{1}{2} \sum_{j=1}^n \left(\nu_j^+ \bar{\rho}_{i,j} - \nu_j^- \underline{\rho}_{i,j} \right) (t - t_i)^2 \geq 0, \quad \forall t \in [t_i, t_{i+1}). \end{aligned}$$

Since L_k^C is either linear or quadratic on $[t_i, t_{i+1})$, the left-hand side is a parabola and we need to ensure that this parabola is positive on the interval $[t_i, t_{i+1})$.

Let $a_{i,k}t^2 + b_{i,k}t + c_{i,k}$ be the parabola corresponding to the left-hand side and let $t_{i,k}^* \equiv -b_{i,k}/(2a_{i,k})$ be the t -coordinate of its vertex. Then $a_{i,k}t^2 + b_{i,k}t + c_{i,k} \geq 0, \forall t \in [t_i, t_{i+1})$, is equivalent to requiring that

$$(51) \quad \begin{cases} a_{i,k}t_i^2 + b_{i,k}t_i + c_{i,k} \geq 0, \\ a_{i,k}t_{i+1}^2 + b_{i,k}t_{i+1} + c_{i,k} \geq 0, \\ \mathbf{1}_{(t_i, t_{i+1})}(t_{i,k}^*)(a_{i,k}(t_{i,k}^*)^2 + b_{i,k}t_{i,k}^* + c_{i,k}) \geq 0. \end{cases}$$

The first two conditions guarantee that the end points of the parabola lie above the t -axis, while the last condition ensures that the vertex is above the t -axis if it is in the interval (t_i, t_{i+1}) . As before, by the continuity of $\sum_{j=1}^n \nu_j k_j^{**}(t)$ and $L_k^C(t)$, the first condition is redundant and can be dropped.

Here $t_{i,k}^*$ depends nonlinearly on $\tilde{\nu}$, so the conservatively discretized optimization problem is not a linear program. Nevertheless, a conservative lower bound at the k th true bin is

$$\begin{aligned} & \sup_{\tilde{\nu} \in \mathbb{R}^{2n}} (\mathbf{D}^T \tilde{\mathbf{y}} - \tilde{\ell})^T \tilde{\nu} \\ \text{subject to } & a_{i,k} t_{i+1}^2 + b_{i,k} t_{i+1} + c_{i,k} \geq 0, \quad i = 1, \dots, m, \\ & \mathbf{1}_{(t_i, t_{i+1})}(t_{i,k}^*) (a_{i,k} (t_{i,k}^*)^2 + b_{i,k} t_{i,k}^* + c_{i,k}) \geq 0, \quad i = 1, \dots, m, \\ & - \sum_{j=1}^n (\mathbf{D}\tilde{\nu})_j k_j^*(T_{\max}) \geq T_{k,\min} - T_{k,\max}, \\ & \tilde{\nu} \geq \mathbf{0}, \end{aligned}$$

where $\tilde{\nu} = \begin{bmatrix} \nu^+ \\ \nu^- \end{bmatrix}$, $\mathbf{D} = [\mathbf{I}_{n \times n} \quad -\mathbf{I}_{n \times n}]$, $t_{i,k}^* = -\frac{b_{i,k}}{2a_{i,k}}$ and the coefficients $a_{i,k}$, $b_{i,k}$ and $c_{i,k}$, which depend on $\tilde{\nu}$, are given by

$$(52) \quad a_{i,k} = \begin{cases} -A_i, & t_i < T_{k,\min}, \\ -A_i + \frac{1}{2}, & T_{k,\min} \leq t_i < T_{k,\max}, \\ -A_i, & t_i \geq T_{k,\max}, \end{cases}$$

$$(53) \quad b_{i,k} = \begin{cases} 2A_i t_i - B_i, & t_i < T_{k,\min}, \\ 2A_i t_i - B_i - T_{k,\min}, & T_{k,\min} \leq t_i < T_{k,\max}, \\ 2A_i t_i - B_i + T_{k,\max} - T_{k,\min}, & t_i \geq T_{k,\max}, \end{cases}$$

$$(54) \quad c_{i,k} = \begin{cases} -A_i t_i^2 + B_i t_i - C_i, & t_i < T_{k,\min}, \\ -A_i t_i^2 + B_i t_i - C_i + \frac{1}{2} T_{k,\min}^2, & T_{k,\min} \leq t_i < T_{k,\max}, \\ -A_i t_i^2 + B_i t_i - C_i - \frac{1}{2} T_{k,\max}^2 + \frac{1}{2} T_{k,\min}^2, & t_i \geq T_{k,\max}, \end{cases}$$

where

$$(55) \quad A_i = \frac{1}{2} \sum_{j=1}^n (\nu_j^+ \bar{\rho}_{i,j} - \nu_j^- \underline{\rho}_{i,j}),$$

$$(56) \quad B_i = \sum_{j=1}^n (\nu_j^+ - \nu_j^-) k_j^*(t_i),$$

$$(57) \quad C_i = \sum_{j=1}^n (\nu_j^+ - \nu_j^-) k_j^{**}(t_i).$$

Similarly, the corresponding upper bound is

$$\begin{aligned}
& \inf_{\tilde{\mathbf{v}} \in \mathbb{R}^{2n}} && -(\mathbf{D}^T \tilde{\mathbf{y}} - \tilde{\boldsymbol{\ell}})^T \tilde{\mathbf{v}} \\
\text{subject to} &&& a_{i,k} t_{i+1}^2 + b_{i,k} t_{i+1} + c_{i,k} \geq 0, \quad i = 1, \dots, m, \\
&&& \mathbf{1}_{(t_i, t_{i+1})}(t_{i,k}^*) (a_{i,k} (t_{i,k}^*)^2 + b_{i,k} t_{i,k}^* + c_{i,k}) \geq 0, \quad i = 1, \dots, m, \\
&&& - \sum_{j=1}^n (\mathbf{D}\tilde{\mathbf{v}})_j k_j^*(T_{\max}) \geq T_{k,\max} - T_{k,\min}, \\
&&& \tilde{\mathbf{v}} \geq \mathbf{0},
\end{aligned}$$

where the coefficients are

$$(58) \quad a_{i,k} = \begin{cases} -A_i, & t_i < T_{k,\min}, \\ -A_i - \frac{1}{2}, & T_{k,\min} \leq t_i < T_{k,\max}, \\ -A_i, & t_i \geq T_{k,\max}, \end{cases}$$

$$(59) \quad b_{i,k} = \begin{cases} 2A_i t_i - B_i, & t_i < T_{k,\min}, \\ 2A_i t_i - B_i + T_{k,\min}, & T_{k,\min} \leq t_i < T_{k,\max}, \\ 2A_i t_i - B_i - T_{k,\max} + T_{k,\min}, & t_i \geq T_{k,\max}, \end{cases}$$

$$(60) \quad c_{i,k} = \begin{cases} -A_i t_i^2 + B_i t_i - C_i, & t_i < T_{k,\min}, \\ -A_i t_i^2 + B_i t_i - C_i - \frac{1}{2} T_{k,\min}^2, & T_{k,\min} \leq t_i < T_{k,\max}, \\ -A_i t_i^2 + B_i t_i - C_i + \frac{1}{2} T_{k,\max}^2 - \frac{1}{2} T_{k,\min}^2, & t_i \geq T_{k,\max}, \end{cases}$$

and A_i , B_i and C_i are given by Equations (55)–(57).

These optimization problems can be solved using standard nonlinear programming algorithms, but some care is needed when choosing the algorithm and its starting point; see Section C.1. Since any feasible point of these programs yields a conservative bound, we need only a good feasible point and not necessarily a global optimum.

APPENDIX C: IMPLEMENTATION DETAILS

This appendix provides implementation details for the unfolding methods used in this paper. Section C.1 focuses on the shape-constrained strict bounds, while Sections C.2 and C.3 describe our implementation of the SVD and D’Agostini methods, respectively.

C.1. Shape-constrained strict bounds. The optimization problems to find the shape-constrained strict bounds involve a relatively high-dimensional solution space, numerical values on very different scales, and fairly complicated constraints. As a result, some care is needed in their numerical

solution, including verifying the validity of the output of the optimization algorithms.

For positivity and monotonicity constraints, where the bounds can be found by linear programming, we used the interior-point linear program solver implemented in the `linprog` function of the MATLAB Optimization Toolbox R2014a (Mathworks, 2014). To find bounds under the convexity constraint, we used the sequential quadratic programming (SQP) algorithm as implemented in the `fmincon` function of the MATLAB Optimization Toolbox R2014a (Mathworks, 2014). (Due to a change in our computing infrastructure, the simulations for Table 1 were run using MATLAB version R2014b, but we do not expect that to affect the results.)

The optimization problems described in Sections 4.4, B.1, and B.2 tend to be numerically unstable when the solver explores large values of $\tilde{\nu}$. We address this issue by imposing an upper bound on $\tilde{\nu}$. For each j , we replace the constraint $\tilde{\nu}_j \geq 0$ with the constraint $0 \leq \tilde{\nu}_j \leq U$, where U is chosen to be large enough that the upper bound is not active at the optimal solution. (Notice that even if the upper bound were active, the solution of the modified problem would still be a valid confidence bound, because the restricted feasible set is a subset of the original feasible set.) Imposing the upper bound substantially improved the stability of the numerical solvers. For the numerical experiments of this paper, U was set to 30 for the positivity constraint, 15 for the monotonicity constraint, and 10 for the convexity constraint.

The solution found by the optimization algorithms can violate the constraints within a preset numerical tolerance. This could make the confidence bound optimistic rather than conservative. To ensure that this does not happen, we verify the feasibility of the solution returned by the optimization algorithm. If it is infeasible, we iteratively scale ν^+ down and ν^- up until the solution becomes feasible. Typically very little fine-tuning of this kind was required to obtain a feasible point.

The SQP algorithm requires a good feasible starting point. To find one, we first solve the linear program corresponding to the nonconservative discretization

$$(61) \quad \sum_{j=1}^n \nu_j k_j^{**}(t_i) \leq \pm L_k^C(t_i), \quad i = 1, \dots, m + 1.$$

We then scale the solution as described above to make it feasible for the conservative discretization; the result is then used as the initial feasible point for SQP. Since the SQP iteration is prohibitively CPU intensive for a large-scale coverage study, the nonconservative discretization (61) was also

each bin by the reciprocal of its estimated variance. That is, we choose δ to minimize

$$(65) \quad \text{CV} = \sum_{i=1}^n \frac{(y_i - \hat{\mu}_i^{-i})^2}{y_i} = \sum_{i=1}^n \frac{1}{y_i} \left(\frac{y_i - \hat{\mu}_i}{1 - H_{i,i}} \right)^2,$$

where $\hat{\mu}_i^{-i}$ is the estimate of the i th smeared bin obtained using $\mathbf{y}^{-i} = [y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_n]^T$. This choice of δ aims to minimize the squared prediction error in the smeared space, as described for example in O'Sullivan (1986). Ruppert, Wand and Carroll (2003, Chapter 5), Hansen (1998, Chapter 7), and Vogel (2002, Chapter 7) review alternative techniques for choosing the regularization parameter.

An estimate of the covariance of $\hat{\boldsymbol{\lambda}}$ (ignoring the data-dependence of δ and $\hat{\mathbf{C}}$) is $\widehat{\text{Cov}}(\hat{\boldsymbol{\lambda}}) = \mathbf{K}^+ \text{diag}(\mathbf{y})(\mathbf{K}^+)^T$. If the distribution of $\hat{\boldsymbol{\lambda}}$ is approximately Gaussian, its binwise uncertainties can be quantified using the intervals of Equation (38) with $\widehat{\text{Var}}(\hat{\lambda}_j) = \widehat{\text{Cov}}(\hat{\boldsymbol{\lambda}})_{j,j}$.

C.3. D'Agostini iteration. The D'Agostini iteration (D'Agostini, 1995) uses the EM algorithm (Dempster, Laird and Rubin, 1977; Shepp and Vardi, 1982; Vardi, Shepp and Kaufman, 1985) to find a maximum likelihood solution in the discrete forward model $\mathbf{y} \sim \text{Poisson}(\mathbf{K}\boldsymbol{\lambda}_0)$. Given a starting point $\boldsymbol{\lambda}^{(0)} > \mathbf{0}$, the $(t+1)$ th step of the iteration is

$$(66) \quad \lambda_j^{(t+1)} = \frac{\lambda_j^{(t)}}{\sum_{i=1}^n K_{i,j}} \sum_{i=1}^n \frac{K_{i,j} y_i}{\sum_{k=1}^p K_{i,k} \lambda_k^{(t)}}.$$

The solution is regularized by stopping the algorithm before it converges to a maximum likelihood estimate. In the ROOUNFOLD implementation (Abye, 2011), the iteration starts at a MC prediction of the unfolded histogram, $\boldsymbol{\lambda}^{(0)} = \boldsymbol{\lambda}^{\text{MC}}$.

We choose the number of iterations using weighted cross-validation, minimizing Equation (65). Due to the nonlinearity of the D'Agostini iteration, the second equality in Equation (65) no longer holds. As a result, evaluating the cross-validation function requires computing n point estimates, one for each left-out smeared bin, which is computationally demanding.

Using linearization, we estimate the covariance of $\boldsymbol{\lambda}^{(t+1)}$ by

$$(67) \quad \widehat{\text{Cov}}(\boldsymbol{\lambda}^{(t+1)}) = \mathbf{J}^{(t+1)} \text{diag}(\mathbf{y})(\mathbf{J}^{(t+1)})^T,$$

where $\mathbf{J}^{(t+1)}$ is the Jacobian of $\boldsymbol{\lambda}^{(t+1)}$ evaluated at \mathbf{y} . Let $\varepsilon_j \equiv \sum_{i=1}^n K_{i,j}$ and

$$(68) \quad M_{i,j}^{(t)} \equiv \frac{\lambda_j^{(t)}}{\varepsilon_j} \frac{K_{i,j}}{\sum_{k=1}^p K_{i,k} \lambda_k^{(t)}}.$$

Then the elements of the Jacobian are (Adye, 2011)

$$(69) \quad J_{j,i}^{(t+1)} = \frac{\partial \lambda_j^{(t+1)}}{\partial y_i} = M_{i,j}^{(t)} + \frac{\lambda_j^{(t+1)}}{\lambda_j^{(t)}} J_{j,i}^{(t)} - \sum_{k=1}^p \sum_{l=1}^n y_l \frac{\varepsilon_k}{\lambda_k^{(t)}} M_{l,j}^{(t)} M_{l,k}^{(t)} J_{k,i}^{(t)},$$

with $J_{j,i}^{(0)} = 0$ for all j, i . The estimated variances for Equation (38) are $\widehat{\text{Var}}(\hat{\lambda}_j) = \widehat{\text{Cov}}(\hat{\boldsymbol{\lambda}})_{j,j} = \widehat{\text{Cov}}(\boldsymbol{\lambda}^{(N_{\text{iter}})})_{j,j}$, where N_{iter} is the number of iterations. The resulting intervals ignore the data-dependence of N_{iter} and treat $\hat{\boldsymbol{\lambda}}$ as Gaussian.

REFERENCES

- ADYE, T. (2011). Unfolding algorithms and tests using RooUnfold. In *Proceedings of the PHYSTAT 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding* (H. B. PROSPER and L. LYONS, eds.). CERN-2011-006 313–318.
- ANTONIADIS, A. and BIGOT, J. (2006). Poisson inverse problems. *The Annals of Statistics* **34** 2132–2158.
- BACKUS, G. (1970). Inference from Inadequate and Inaccurate Data, I. *Proceedings of the National Academy of Sciences* **65** 1–7.
- BANERJEE, M. and WELLNER, J. A. (2001). Likelihood Ratio Tests for Monotone Functions. *The Annals of Statistics* **29** 1699–1731.
- BARNEY, D. (2004). CMS-doc-4172. <https://cms-docdb.cern.ch/cgi-bin/PublicDocDB/ShowDocument?docid=4172>. Retrieved 21.1.2014.
- BLOBEL, V. (2013). Unfolding. In *Data Analysis in High Energy Physics: A Practical Guide to Statistical Methods* (O. Behnke, K. Kröninger, G. Schott and T. Schörner-Sadenius, eds.) 187–225. Wiley.
- BURRUS, W. R. (1965). Utilization of a priori information by means of mathematical programming in the statistical interpretation of measured distributions. Oak Ridge National Laboratory, ORNL-3743.
- BURRUS, W. R. and VERBINSKI, V. V. (1969). Fast-neutron spectroscopy with thick organic scintillators. *Nuclear Instruments and Methods* **67** 181–196.
- CAI, T. T., LOW, M. G. and XIA, Y. (2013). Adaptive confidence intervals for regression functions under shape constraints. *The Annals of Statistics* **41** 722–750.
- CARROLL, R. J., DELAIGLE, A. and HALL, P. (2011). Testing and Estimating Shape-Constrained Nonparametric Density and Regression in the Presence of Measurement Error. *Journal of the American Statistical Association* **106** 191–202.
- CHOU DALAKIS, G. (2012). Fully Bayesian unfolding. arXiv:1201.4612v4 [physics.data-an].
- CMS COLLABORATION (2008). The CMS experiment at the CERN LHC. *Journal of Instrumentation* **3** S08004.
- CMS COLLABORATION (2010). Measurement of the Inclusive Jet Cross Section in pp Collisions at 7 TeV. CMS-PAS-QCD-10-011, available at <http://cds.cern.ch/record/1280682>.
- CMS COLLABORATION (2011). Measurement of the Inclusive Jet Cross Section in pp Collisions at $\sqrt{s} = 7$ TeV. *Physical Review Letters* **107** 132001.
- ATLAS COLLABORATION (2012). Measurement of the transverse momentum distribution of W bosons in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *Physical Review D* **85** 012005.

- CMS COLLABORATION (2013a). Measurements of differential jet cross sections in proton-proton collisions at $\sqrt{s} = 7$ TeV with the CMS detector. *Physical Review D* **87** 112002.
- CMS COLLABORATION (2013b). Measurement of differential top-quark-pair production cross sections in pp collisions at $\sqrt{s} = 7$ TeV. *The European Physical Journal C* **73** 2339.
- NNPDF COLLABORATION (2015). Parton distributions for the LHC run II. *Journal of High Energy Physics* **1504** 040.
- CMS COLLABORATION (2016). Measurement of differential cross sections for Higgs boson production in the diphoton decay channel in pp collisions at $\sqrt{s} = 8$ TeV. *The European Physical Journal C* **76** 13.
- COWAN, G. (1998). *Statistical Data Analysis*. Oxford University Press.
- D'AGOSTINI, G. (1995). A multidimensional unfolding method based on Bayes' theorem. *Nuclear Instruments and Methods A* **362** 487–498.
- DAVIES, P. L., KOVAC, A. and MEISE, M. (2009). Nonparametric Regression, Confidence Regions and Regularization. *The Annals of Statistics* **37** 2597–2625.
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* **39** 1–38.
- DÜMBGEN, L. (1998). New Goodness-Of-Fit Tests and Their Application to Nonparametric Confidence Sets. *The Annals of Statistics* **26** 288–314.
- DÜMBGEN, L. (2003). Optimal confidence bands for shape-restricted curves. *Bernoulli* **9** 423–449.
- FORTE, S. and WATT, G. (2013). Progress in the determination of the partonic structure of the proton. *Annual Review of Nuclear and Particle Science* **63** 291–328.
- GARWOOD, F. (1936). Fiducial limits for the Poisson distribution. *Biometrika* **28** 437–442.
- GENOVESE, C. and WASSERMAN, L. (2008). Adaptive confidence bands. *The Annals of Statistics* **36** 875–905.
- GREEN, P. J. and SILVERMAN, B. W. (1994). *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach*. Chapman & Hall.
- GRENANDER, U. (1956). On the theory of mortality measurement: Part II. *Scandinavian Actuarial Journal* **1956** 125–153.
- GROENEBOOM, P., JONGBLOED, G. and WELLNER, J. A. (2001). Estimation of a Convex Function: Characterizations and Asymptotic Theory. *The Annals of Statistics* **29** 1653–1698.
- GROENEBOOM, P. and JONGBLOED, G. (2014). *Nonparametric Estimation Under Shape Constraints: Estimators, Algorithms and Asymptotics*. Cambridge University Press.
- GROENEBOOM, P. and JONGBLOED, G. (2015). Nonparametric confidence intervals for monotone functions. *The Annals of Statistics* **43** 2019–2054.
- HALL, P. and HOROWITZ, J. (2013). A simple bootstrap method for constructing nonparametric confidence bands for functions. *The Annals of Statistics* **41** 1892–1921.
- HANSEN, P. C. (1998). *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*. SIAM.
- HANSEN, P. C. (2010). *Discrete Inverse Problems: Insight and Algorithms*. SIAM.
- HENGARTNER, N. W. and STARK, P. B. (1992). Conservative finite-sample confidence envelopes for monotone and unimodal densities. Technical Report No. 341, Department of Statistics, University of California, Berkeley.
- HENGARTNER, N. W. and STARK, P. B. (1995). Finite-sample confidence envelopes for shape-restricted densities. *The Annals of Statistics* **23** 525–550.
- HÖCKER, A. and KARTVELISHVILI, V. (1996). SVD approach to data unfolding. *Nuclear Instruments and Methods in Physics Research A* **372** 469–481.

- KONDOR, A. (1983). Method of convergent weights – An iterative procedure for solving Fredholm’s integral equations of the first kind. *Nuclear Instruments and Methods* **216** 177–181.
- KUUSELA, M. and PANARETOS, V. M. (2015). Statistical unfolding of elementary particle spectra: Empirical Bayes estimation and bias-corrected uncertainty quantification. *The Annals of Applied Statistics* **9** 1671–1705.
- LANGE, K. and CARSON, R. (1984). EM reconstruction algorithms for emission and transmission tomography. *Journal of Computer Assisted Tomography* **8** 306–316.
- LOW, M. G. (1997). On nonparametric confidence intervals. *The Annals of Statistics* **25** 2547–2554.
- LUCY, L. B. (1974). An iterative technique for the rectification of observed distributions. *Astronomical Journal* **79** 745–754.
- LUENBERGER, D. G. (1969). *Optimization by Vector Space Methods*. John Wiley & Sons.
- MATHWORKS (2014). Optimization Toolbox User’s Guide Release 2014a.
- MEISTER, A. (2009). *Deconvolution Problems in Nonparametric Statistics*. Springer.
- MÜLTHEI, H. N. and SCHORR, B. (1987a). On an iterative method for a class of integral equations of the first kind. *Mathematical Methods in the Applied Sciences* **9** 137–168.
- MÜLTHEI, H. N. and SCHORR, B. (1987b). On an iterative method for the unfolding of spectra. *Nuclear Instruments and Methods in Physics Research A* **257** 371–377.
- MÜLTHEI, H. N. and SCHORR, B. (1989). On properties of the iterative maximum likelihood reconstruction method. *Mathematical Methods in the Applied Sciences* **11** 331–342.
- O’LEARY, D. P. and RUST, B. W. (1986). Confidence Intervals for Inequality-Constrained Least Squares Problems, with Applications to Ill-Posed Problems. *SIAM Journal on Scientific and Statistical Computing* **7** 473–489.
- O’SULLIVAN, F. (1986). A statistical perspective on ill-posed inverse problems. *Statistical Science* **1** 505–527.
- PFLUG, G. C. and WETS, R. J. B. (2013). Shape-restricted nonparametric regression with overall noisy measurements. *Journal of Nonparametric Statistics* **25** 323–338.
- PHILLIPS, D. L. (1962). A Technique for the Numerical Solution of Certain Integral Equations of the First Kind. *Journal of the ACM* **9** 84–97.
- PIERCE, J. E. and RUST, B. W. (1985). Constrained Least Squares Interval Estimation. *SIAM Journal on Scientific and Statistical Computing* **6** 670–683.
- PROSPER, H. B. and LYONS, L., eds. (2011). *Proceedings of the PHYSTAT 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding*. CERN-2011-006.
- REISS, R. D. (1993). *A Course on Point Processes*. Springer-Verlag.
- RICHARDSON, W. H. (1972). Bayesian-based iterative method of image restoration. *Journal of the Optical Society of America* **62** 55–59.
- ROBERTSON, T., WRIGHT, F. T. and DYKSTRA, R. L. (1988). *Order Restricted Statistical Inference*. John Wiley & Sons.
- RUPPERT, D., WAND, M. P. and CARROLL, R. J. (2003). *Semiparametric Regression*. Cambridge University Press.
- RUST, B. W. and BURRUS, W. R. (1972). *Mathematical Programming and the Numerical Solution of Linear Equations*. American Elsevier.
- RUST, B. W. and O’LEARY, D. P. (1994). Confidence Intervals for Discrete Approximations to Ill-Posed Problems. *Journal of Computational and Graphical Statistics* **3** 67–96.
- SCHMITT, S. (2012). TUnfold, an algorithm for correcting migration effects in high energy physics. *Journal of Instrumentation* **7** T10003.

- SHEPP, L. A. and VARDI, Y. (1982). Maximum likelihood reconstruction for emission tomography. *IEEE Transactions on Medical Imaging* **1** 113–122.
- STARK, P. B. (1992). Inference in Infinite-Dimensional Inverse Problems: Discretization and Duality. *Journal of Geophysical Research* **97** 14055–14082.
- TIKHONOV, A. N. (1963). Solution of incorrectly formulated problems and the regularization method. *Soviet Mathematics: Doklady* **4** 1035–1038.
- VARDI, Y., SHEPP, L. A. and KAUFMAN, L. (1985). A statistical model for positron emission tomography. *Journal of the American Statistical Association* **80** 8–20.
- VOGEL, C. R. (2002). *Computational Methods for Inverse Problems*. SIAM.
- VOLOBOUEV, I. (2015). On the Expectation-Maximization Unfolding with Smoothing. arXiv:1408.6500v2 [physics.data-an].
- WAHBA, G. (1982). Constrained regularization for ill posed linear operator equations, with applications in meteorology and medicine. In *Statistical Decision Theory and Related Topics III*, (S. S. Gupta and J. O. Berger, eds.) **2** 383–418. Academic Press.

SECTION DE MATHÉMATIQUES
ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE
EPFL STATION 8, 1015 LAUSANNE
SWITZERLAND
E-MAIL: mikael.kuusela@gmail.com

DEPARTMENT OF STATISTICS
UNIVERSITY OF CALIFORNIA
BERKELEY, CA 94720-3860
UNITED STATES
E-MAIL: stark@stat.berkeley.edu