hep-ph/2512.04160

Try the code yourself!

Resummed Distribution Functions

making perturbation theory positive and normalized

Radha Mastandrea

rmastand@uchicago.edu

In collaboration with Rikab Gambhir

LCTP Spring Symposium: Theoretical Physics and AI

05/19/2026

Why should I care about this talk?

1. We introduce a new framework for resummation in QCD with a single constraint of unitarity

2. We provide a new definition of theory uncertainties

3. We perform a fit to ALEPH collider data to get a reasonable measurement of α using only fixed-order calculations up to O(α^3)

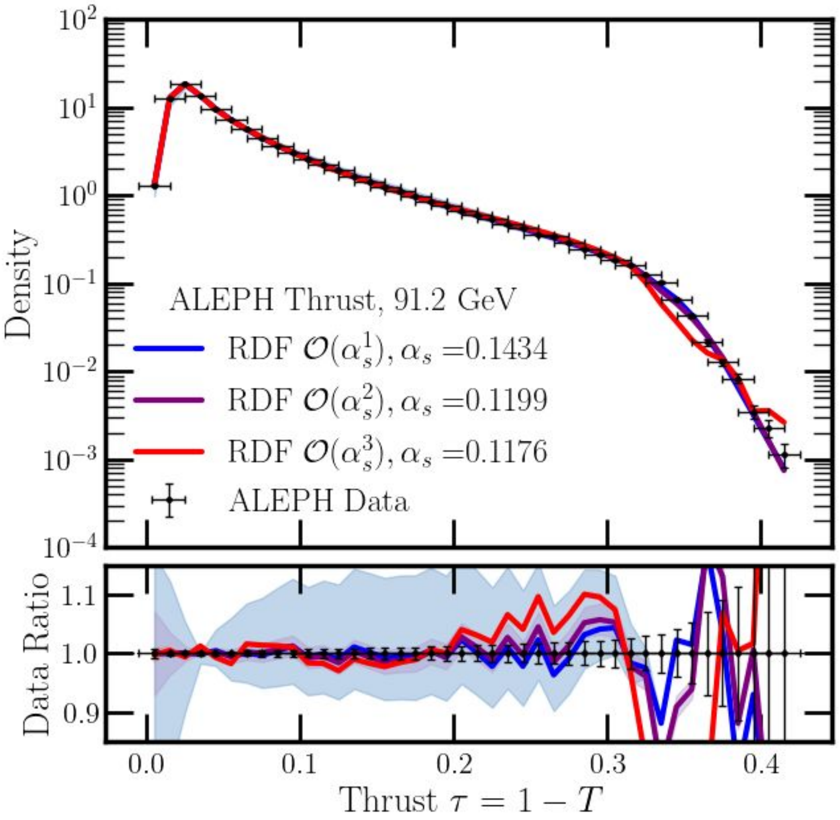

The goal:
make predictions in QCD

With perturbative techniques, we can methodically calculate observables to a fixed order *M* in α.

$$p_{\mathrm{FO}}(x|\alpha) = p_0(x) + \alpha p_1(x) + \alpha^2 p_2(x) + \ldots + \alpha^M p_M(x) + \mathcal{O}\left(\alpha^{M+1}\right)$$

Eg. $ee \longrightarrow jj$

NLO α^1

NNLO α^2

The goal:
make predictions in QCD

With perturbative techniques, we can methodically calculate observables to a fixed order *M* in α.

$$p_{\mathrm{FO}}(x|\alpha) = p_0(x) + \alpha p_1(x) + \alpha^2 p_2(x) + \ldots + \alpha^M p_M(x) + \mathcal{O}\left(\alpha^{M+1}\right)$$

Eg. $ee \to jj$

NLO α^1

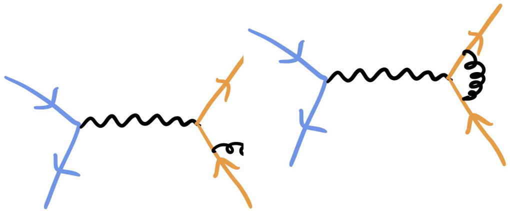

NNLO α^2

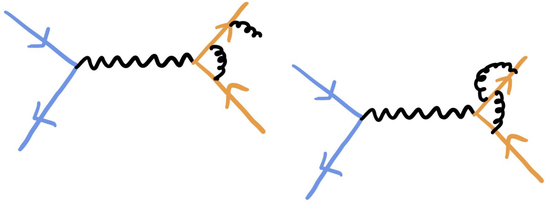

The problem:
pQCD is not well-behaved

A fixed-order calculation can defy unitarity.

These calculations will be ill-behaved, possibly negative, infinite or unnormalized.

The solution:
the Resummed Distribution Function (RDF)

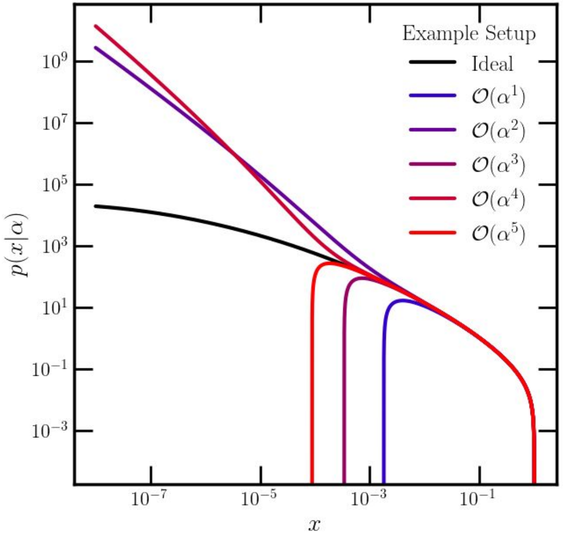

We inject theory knowledge to a fixed order α^M…

RDF

…into the RDF…

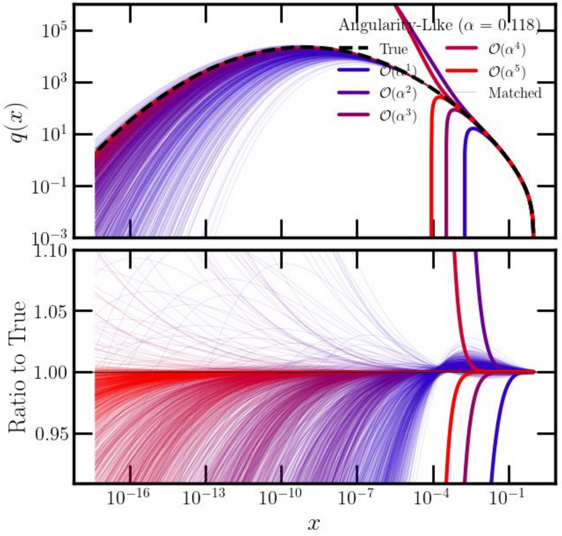

…to get a family of PDFs that contains all the theory knowledge to α^M and has a meaningful notion of uncertainties in α

The solution: the Resummed Distribution Function (RDF)

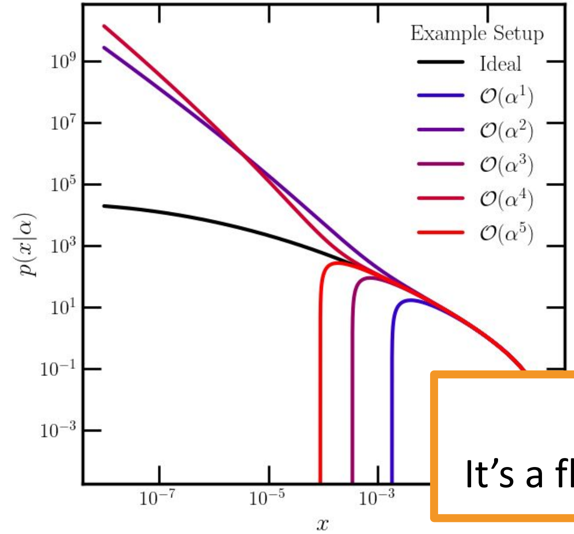

RDF

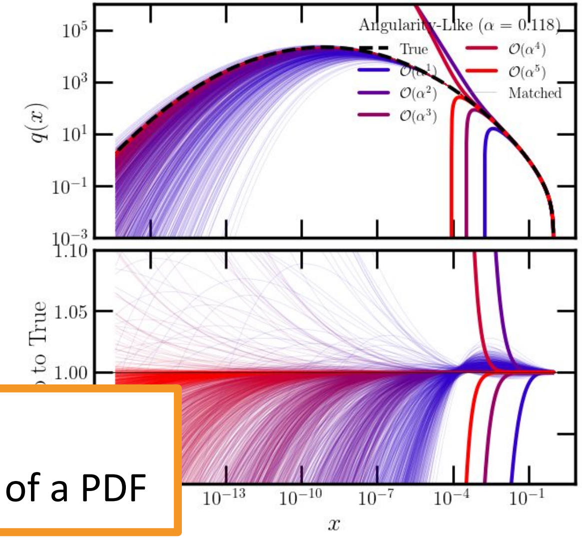

The RDF is not a black box!
It’s a flexible, interpretable parametrization of a PDF

We inject theory knowledge to a fixed order α^M…

…into the RDF…

…to get a family of PDFs that contains all the theory knowledge to α^M and has a meaningful notion of uncertainties in α

What does the RDF look like?

Given a fixed-order PDF, we can construct all valid probability distributions for an observable *t* that contain all the relevant physics up to a certain order of α^M as

$$\boxed{q(t|\alpha) = f(t,\alpha) \cdot \exp\left[-\int_0^t dt'\, f(t',\alpha)\right]}$$

What does the RDF look like?

Given a fixed-order PDF, we can construct all valid probability distributions for an observable t that contain all the relevant physics up to a certain order of α^M as

$$\boxed{q(t|\alpha) = f(t,\alpha) \cdot \exp\left[-\int_0^t dt'\, f(t',\alpha)\right]}$$

The CDF of q is $Q(t|\alpha) = 1 - e^{-\int_0^t dt'\, f(t',\alpha)}$

We want $Q(0\,|\,\alpha) = 0$ and $Q(\infty^*\,|\,\alpha) = 1$. This is satisfied if:

1. f is positive $\longrightarrow$ define $f(t,\alpha) = \exp[-g(t,\alpha)]$
2. f has a diverging integral $\longrightarrow$ satisfied for most g

*Most observables that we care about are bounded. So we can define $t \equiv \log\left(\frac{1}{x}\right)$ for a bounded observable x.

What does the RDF look like?

Given a fixed-order PDF, we can construct all valid probability distributions for an observable *t* that contain all the relevant physics up to a certain order of α^M as

$$\boxed{q(t|\alpha) = f(t,\alpha)\cdot\exp\left[-\int_0^t dt'\, f(t',\alpha)\right]}$$

with $f(t,\alpha) = \exp[-g(t,\alpha)]$

It is helpful to split *g* into two components:

$$g(t,\alpha) = \boxed{-\log(g^*(t,\alpha))} + \boxed{g_{\text{Analytic}}(t,\alpha)}$$

This captures the leading-order behavior in α

This captures the higher order behavior in α

But what is *g*?

It is helpful to split *g* into two components:

$$g(t,\alpha) = \boxed{-\log(g^*(t,\alpha))} + \boxed{g_{\text{Analytic}}(t,\alpha)}$$

This captures the leading-order behavior in α

This captures the higher order behavior in α

g can be any* sufficiently flexible parameterized function of *t* and α!

A simple but effective choice is to choose polynomials, i.e.

$$g^*(t,\alpha),\;\; g_{\text{Analytic}}(t,\alpha) = \sum_{m=0,n=0}^{M,N} g_{mn}\alpha^m t^n \Theta(t-\theta_{mn})$$

*As long as *g* is bounded from above and is analytic up to single logs in α. These conditions are generally easy to satisfy.

A few caveats about the RDF

Is this actually resummation?

It's *a* resummation.

The RDF takes a fixed-order expansion in α and reorganizes it into an expansion in $\alpha^{m-m^*}\frac{p_m(t)}{p_{m^*}(t)} \ll 1$ where m^* is the lowest nontrivial order of α in the fixed-order calculation.

This is not generally equal to standard N^{n}LL resummation. However, if the fixed-order calculation includes all of the relevant logs (i.e running of α, non-global logarithms) to the desired order, then the RDF family will contain the standard resummation.

Does the RDF give us the "true" QCD?

No, the RDF is no free lunch.

The RDF only models the perturbative physics that you put into it!

The RDF doesn't model true asymptotic QCD, as it was constructed to be analytic (and therefore convergent) in α.

The RDF doesn't model any nonperturbative physics.

How do we embed a fixed-order calculation into the RDF?

We want the physics to be equivalent at each order of α in the fixed-order calculation $p_{\text{FO}}(x|\alpha) = p_0(x) + \alpha p_1(x) + \alpha^2 p_2(x) + \ldots + \alpha^M p_M(x)$ and the RDF. We can do this by matching the fixed-order calculation to the Taylor-expanded RDF.

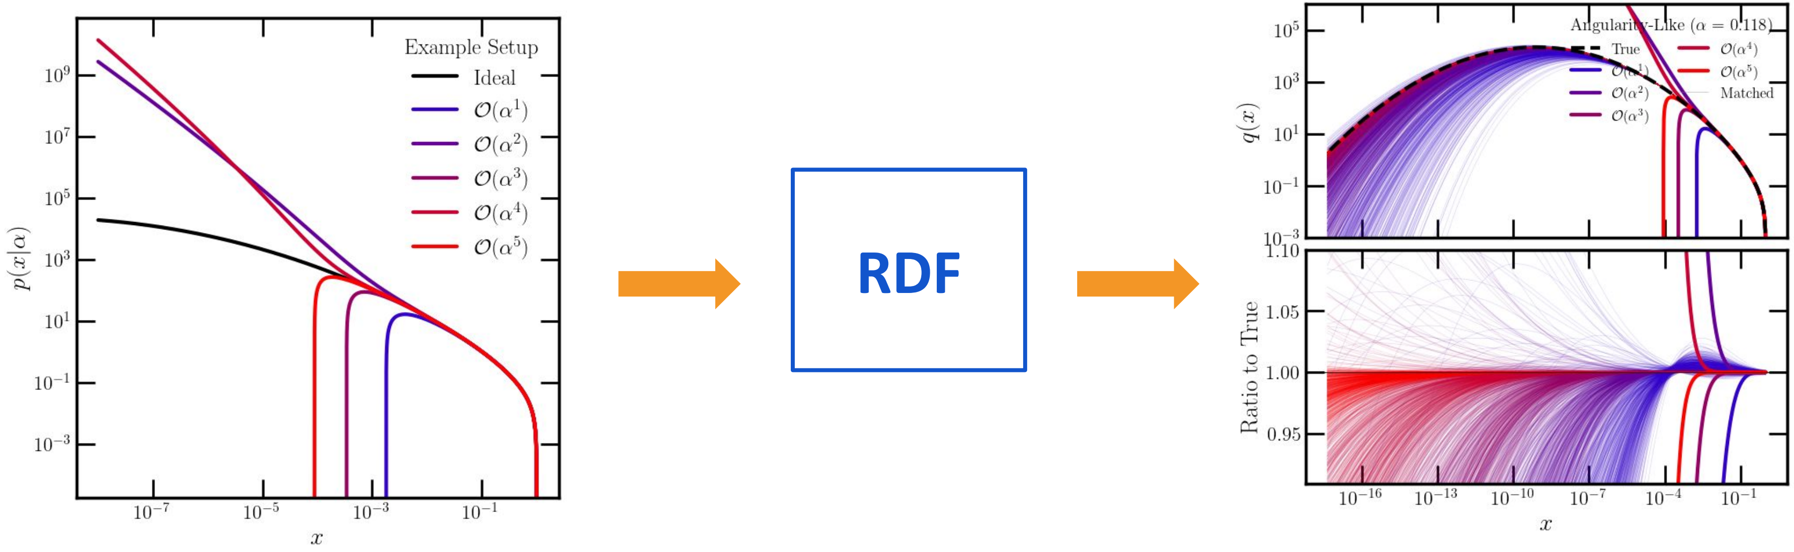

How do we embed a fixed-order calculation into the RDF?

We want the physics to be equivalent at each order of α in the fixed-order calculation $p_{\mathrm{FO}}(x|\alpha) = p_0(x) + \alpha p_1(x) + \alpha^2 p_2(x) + \ldots + \alpha^M p_M(x)$ and the RDF. We can do this by matching the fixed-order calculation to the Taylor-expanded RDF.

For a polynomial *g*, we can just equate coefficients at each order of α to recover

$$g(t,\alpha) = -\log\left(p_{m*}(t)\alpha^{m^*}\right) + \left[\sum_{k\geq 1}\frac{1}{k}\left(\left[-\sum_{m>m^*}^{M}\frac{p_m(t)}{p_{m^*}(t)}\alpha^{m-m^*}\right]^k - \left[\sum_{m=m^*}^{M-m^*}\int_0^t dt'\, p_m(t')\alpha^m\right]^k\right) + \mathcal{O}\left(\alpha^{M+1-m^*}\right)\right]$$

where m^* is the lowest nontrivial order of α.

How do we embed a fixed-order calculation into the RDF?

This leaves the *g* coefficients proportional to α^{M+1} unspecified — this defines the RDF family.

We'll see how to put these degrees of freedom to work soon…

For a polynomial *g*, we can just equate coefficients at each order of α to recover

$$g(t,\alpha) = -\log\Big(p_{m*}(t)\alpha^{m^*}\Big) + \left[\sum_{k\geq 1}\frac{1}{k}\left(\left[-\sum_{m>m^*}^{M}\frac{p_m(t)}{p_{m^*}(t)}\alpha^{m-m^*}\right]^k - \left[\sum_{m=m^*}^{M-m^*}\int_0^t dt'\, p_m(t')\alpha^m\right]^k\right) + \mathcal{O}\left(\alpha^{M+1-m^*}\right)\right]$$

where m^* is the lowest nontrivial order of α.

Matching example: an exponential observable

Let's consider an exponentially distributed observable with a true PDF given by

$$p_{\text{Exponential}}(t|\alpha) = \alpha e^{-\alpha t}$$

But we don't know the true PDF — we only know our fixed order calculation

$$p_{\text{FO}}^{(\text{Exp})}(t|\alpha) = \alpha \sum_{m=0}^{M-1} \frac{1}{m!}(-\alpha t)^m + \mathcal{O}(\alpha^{M+1})$$

Matching example: an exponential observable

Let's consider an exponentially distributed observable with a true PDF given by

$$p_{\text{Exponential}}(t|\alpha) = \alpha e^{-\alpha t}$$

But we don't know the true PDF — we only know our fixed order calculation

$$p_{\text{FO}}^{(\text{Exp})}(t|\alpha) = \alpha \sum_{m=0}^{M-1} \frac{1}{m!}(-\alpha t)^m + \mathcal{O}(\alpha^{M+1})$$

Matching is straightforward in this case. We can read off

$$m^* = 1$$

$$p_{m^*}(t) = 1$$

$$p_m(t) = \frac{(-t)^{m-1}}{(m-1)!}$$

Matching example: an exponential observable

Let's consider an exponentially distributed observable with a true PDF given by

$$p_{\text{Exponential}}(t|\alpha) = \alpha e^{-\alpha t}$$

But we don't know the true PDF — we only know our fixed order calculation

$$p_{\text{FO}}^{(\text{Exp})}(t|\alpha) = \alpha \sum_{m=0}^{M-1} \frac{1}{m!}(-\alpha t)^m + \mathcal{O}(\alpha^{M+1})$$

Matching is straightforward in this case. We can read off

$$m^* = 1$$

$$p_{m^*}(t) = 1$$

$$p_m(t) = \frac{(-t)^{m-1}}{(m-1)!}$$

$$\Rightarrow$$

$$g^*(t,\alpha) = \alpha$$

$$g_{\text{Analytic}}(t,\alpha) = 0 + \mathcal{O}(\alpha^M)$$

Matching example: an exponential observable

Let's consider an exponentially distributed observable with a true PDF given by

$$p_{\text{Exponential}}(t|\alpha) = \alpha e^{-\alpha t}$$

But we don't know the true PDF — we only know our fixed order calculation

$$p_{\text{FO}}^{(\text{Exp})}(t|\alpha) = \alpha \sum_{m=0}^{M-1} \frac{1}{m!}(-\alpha t)^m + \mathcal{O}(\alpha^{M+1})$$

After matching, we recover the RDF

$$q^{(\text{Exp})(M)}(t|\alpha) = \alpha \exp\left(-\alpha t\right) + \mathcal{O}(\alpha^{M+1})$$

In this case, we recover the underlying distribution — and a theory error!

Matching example: an exponential observable

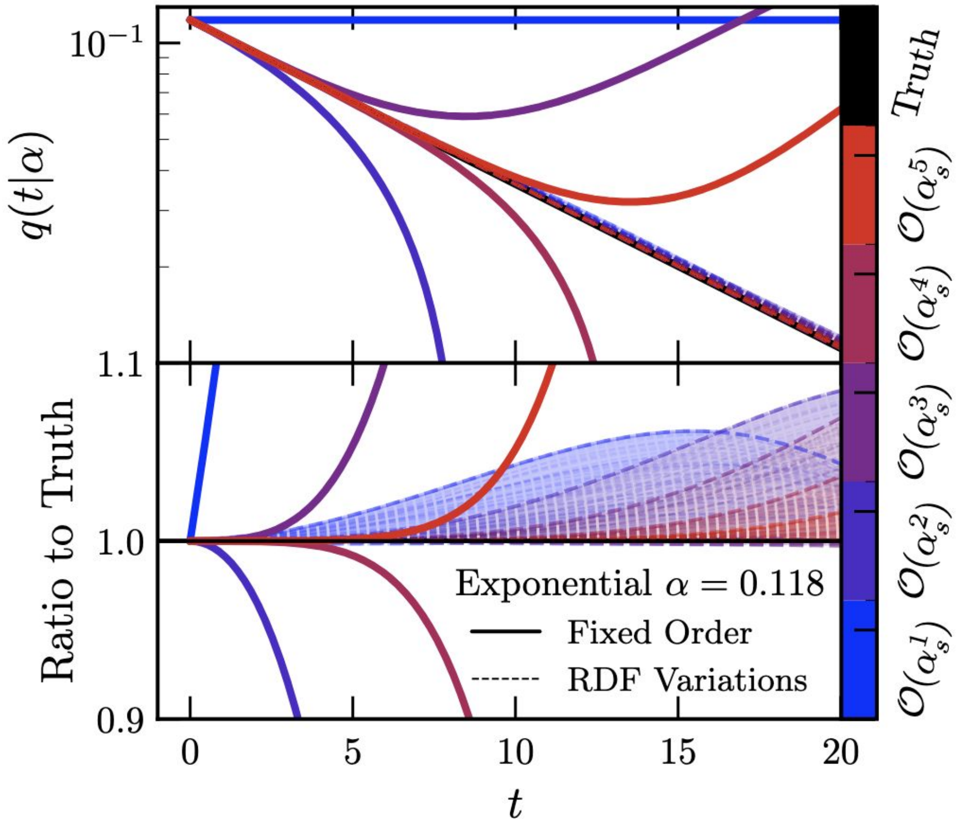

$$q^{(\mathrm{Exp})(M)}(t|\alpha) = \alpha \exp(-\alpha t) + \mathcal{O}(\alpha^{M+1})$$

We can sample the unspecified *g* coefficients from a prior $g_{mn} \sim \mathcal{N}(0, \frac{1}{m!n!})$ to construct uncertainty envelopes.

As *M* increases, the envelopes shrink $\rightarrow$ a fixed order calculation with more knowledge gives a more precise RDF!

Matching example: jet angularities

Let's consider the jet angularity $\lambda^{(\beta)} = \sum_i z_i \left(\frac{\theta_i}{R}\right)^\beta$

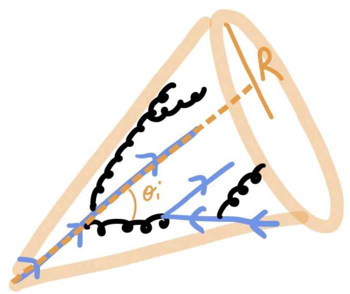

We'll take $t = \log(1/\lambda)$

We have a fixed-order calculation $$p_{\mathrm{FO}}(t) = \left(\frac{\alpha_s C_F}{\pi\beta}\right)\left(2t + 2\log\left(1 - e^{-t}\right) + 3e^{-t} - \frac{3}{2}\right)\Theta\left(t - \log 2\right)$$

$$\downarrow t \to \infty \qquad \text{(Soft and Collinear)}$$

$$= \left(\frac{\alpha_s C_F}{\pi\beta}\right)(2t)$$

Matching is a little harder in this case, but we can do it with Mathematica…

θ_i is measured with respect to the Winner-Take-All axis (roughly the axis of the hardest constituent)

Matching example: jet angularities

When the dust settles, we recover:

$$q(t|\alpha) = \left[\left(\frac{\alpha_s C_F}{\pi\beta}\right)\left(2t + 2\log\left(1-e^{-t}\right) + 3e^{-t} - \frac{3}{2}\right)\Theta\left(t - \log 2\right)\right]$$
$$\times \exp\left[-\left(\frac{\alpha_s C_F}{\pi\beta}\right)\left[t^2 - (\log 2)^2 + 2\left(\mathrm{Li}_2\left(e^{-t}\right) - \mathrm{Li}_2\left(\tfrac{1}{2}\right)\right) - 3\left(e^{-t} - \tfrac{1}{2}\right) - \frac{3}{2}\left(t - \log 2\right)\right]\right]$$

$$\downarrow t \to \infty \qquad \text{(Soft and Collinear)}$$
$$= \left(\frac{\alpha_s C_F}{\pi\beta}\right)(2t)\exp\left(-\frac{\alpha_s C_F}{\pi\beta}t^2\right)$$

RDF contributions

$$p_{\mathrm{FO}}(t) = \left(\frac{\alpha_s C_F}{\pi\beta}\right)\left(2t + 2\log\left(1-e^{-t}\right) + 3e^{-t} - \frac{3}{2}\right)\Theta\left(t - \log 2\right)$$
$$\downarrow t \to \infty \qquad \text{(Soft and Collinear)}$$
$$= \left(\frac{\alpha_s C_F}{\pi\beta}\right)(2t)$$

Matching example: jet angularities

When the dust settles, we recover:

$$q(t|\alpha) = \left[\left(\frac{\alpha_s C_F}{\pi\beta}\right)\left(2t + 2\log\left(1-e^{-t}\right) + 3e^{-t} - \frac{3}{2}\right)\Theta\left(t - \log 2\right)\right]$$
$$\times \exp\left[-\left(\frac{\alpha_s C_F}{\pi\beta}\right)\left[t^2 - (\log 2)^2 + 2\left(\mathrm{Li}_2\left(e^{-t}\right) - \mathrm{Li}_2\left(\tfrac{1}{2}\right)\right) - 3\left(e^{-t} - \tfrac{1}{2}\right) - \frac{3}{2}\left(t - \log 2\right)\right]\right]$$

$$\downarrow t \to \infty \qquad \text{(Soft and Collinear)}$$
$$= \left(\frac{\alpha_s C_F}{\pi\beta}\right)(2t)\exp\left(-\frac{\alpha_s C_F}{\pi\beta}t^2\right)$$

RDF contributio

The RDF normalizes the fixed-order prediction with a meaningful Sudakov suppression

$$p_{\mathrm{FO}}(t) = \left(\frac{\alpha_s C_F}{\pi\beta}\right)\left(2t + 2\log\left(1-e^{-t}\right) + 3e^{-t} - \frac{3}{2}\right)\Theta\left(t - \log 2\right)$$
$$\downarrow t \to \infty \qquad \text{(Soft and Collinear)}$$
$$= \left(\frac{\alpha_s C_F}{\pi\beta}\right)(2t)$$

Matching example: jet angularities

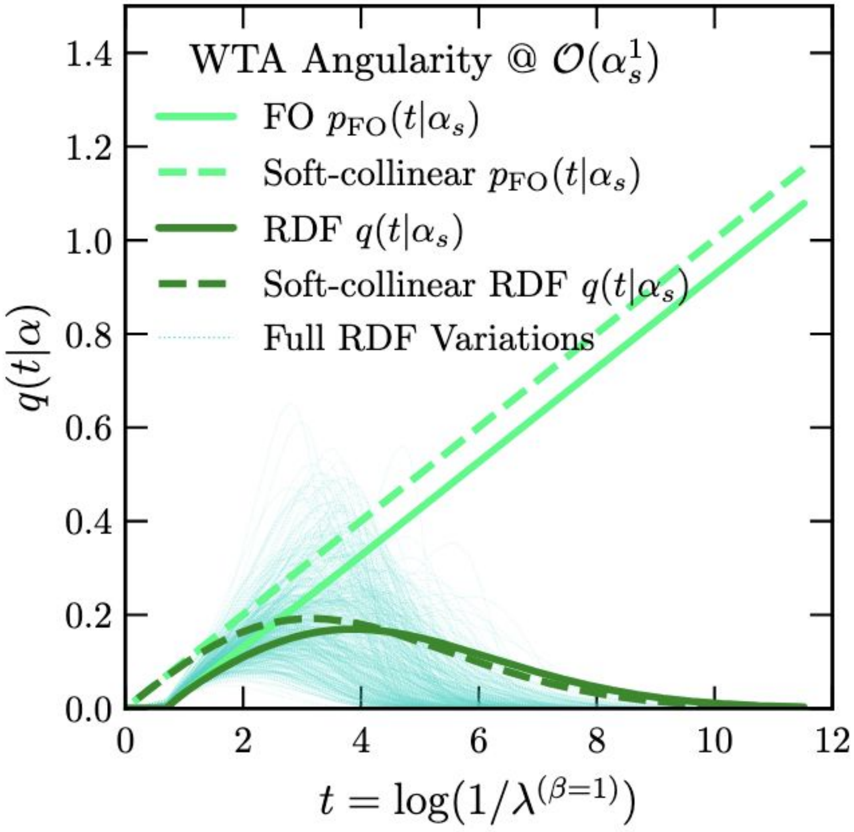

The RDF normalizes the fixed-order prediction with a meaningful Sudakov suppression

Why stop at one angularity?

We start from the O(α^1) differential cross section [Larkoski, Thaler], [Larkoski, Moult, Neill]

$$p_{\mathrm{FO}}(\lambda_\alpha, \lambda_\beta) = \frac{2\alpha_s C_F}{\pi(\alpha-\beta)} \frac{1}{\lambda_\alpha \lambda_\beta} \Theta(\lambda_\alpha^\beta - \lambda_\beta^\alpha)\Theta(\lambda_\beta - \lambda_\alpha)$$

Why stop at one angularity?

We start from the O(α^1) differential cross section [Larkoski, Thaler], [Larkoski, Moult, Neill]

$$p_{\mathrm{FO}}(\lambda_\alpha, \lambda_\beta) = \frac{2\alpha_s C_F}{\pi(\alpha-\beta)} \frac{1}{\lambda_\alpha \lambda_\beta} \Theta(\lambda_\alpha^\beta - \lambda_\beta^\alpha)\Theta(\lambda_\beta - \lambda_\alpha)$$

Working through the RDF framework, we recover

O(α^1) is captured by the unspecified *g* function

$$q(t_\alpha, t_\beta|\alpha_s) = \frac{2\alpha_s C_F}{\pi(\alpha-\beta)} \left(1+\mathcal{O}(\alpha_s)\right) e^{-\alpha_s \frac{C_F}{\pi\beta} t_\beta^2 - \mathcal{O}(\alpha_s)} \Theta_{\alpha\beta}$$

How? We can always decompose a multivariate PDF into a chain of single-variable PDFs

$$p(t_1, t_2, \ldots, t_k|\alpha) = p(t_1|\alpha) \cdot p(t_2|t_1, \alpha) \cdot \ldots \cdot p(t_k|t_1, t_2, \ldots, t_{k-1}, \alpha)$$

Why stop at one angularity?

We start from the O(α^1) differential cross section [Larkoski, Thaler], [Larkoski, Moult, Neill]

$$p_{\mathrm{FO}}(\lambda_\alpha, \lambda_\beta) = \frac{2\alpha_s C_F}{\pi(\alpha-\beta)} \frac{1}{\lambda_\alpha \lambda_\beta} \Theta(\lambda_\alpha^\beta - \lambda_\beta^\alpha) \Theta(\lambda_\beta - \lambda_\alpha)$$

Working through the RDF framework, we recover

O(α^1) is captured by the unspecified *g* function

$$q(t_\alpha, t_\beta | \alpha_s) = \frac{2\alpha_s C_F}{\pi(\alpha-\beta)} \left(1 + \mathcal{O}(\alpha_s)\right) e^{-\alpha_s \frac{C_F}{\pi\beta} t_\beta^2 - \mathcal{O}(\alpha_s)} \Theta_{\alpha\beta}$$

which we can compare with the actual leading log (LL) result

$$q^{\mathrm{LL}}(t_\alpha, t_\beta | \alpha_s) = \frac{2\alpha_s C_F}{\pi(\alpha-\beta)} \left(1 + \frac{2\alpha_s C_F}{\pi(\alpha-\beta)} \frac{(t_\beta - t_\alpha)(\beta t_\alpha - \alpha t_\beta)}{\beta}\right) e^{-\alpha_s \frac{C_F}{\pi}\left(\frac{t_\beta^2}{\beta} + \frac{(t_\alpha - t_\beta)^2}{\alpha-\beta}\right)} \Theta_{\alpha\beta}$$

Why stop at one angularity?

We start from the O(α^1) differential cross section [Larkoski, Thaler], [Larkoski, Moult, Neill]

$p \ldots \lambda_\alpha)$

We performed matching to α^1, and so our results are **equivalent** to α^1 and **consistent** at higher orders — no free lunch!

Working through the RDF framework, we recover

O(α^1) is captured by the unspecified *g* function

$$q(t_\alpha, t_\beta|\alpha_s) = \frac{2\alpha_s C_F}{\pi(\alpha-\beta)}\left(1 + \mathcal{O}(\alpha_s)\right) e^{-\alpha_s \frac{C_F}{\pi\beta} t_\beta^2 - \mathcal{O}(\alpha_s)}\,\Theta_{\alpha\beta}$$

which we can compare with the actual leading log (LL) result

$$q^{\mathrm{LL}}(t_\alpha, t_\beta|\alpha_s) = \frac{2\alpha_s C_F}{\pi(\alpha-\beta)}\left(1 + \frac{2\alpha_s C_F}{\pi(\alpha-\beta)}\frac{(t_\beta - t_\alpha)(\beta t_\alpha - \alpha t_\beta)}{\beta}\right) e^{-\alpha_s \frac{C_F}{\pi}\left(\frac{t_\beta^2}{\beta} + \frac{(t_\alpha - t_\beta)^2}{\alpha-\beta}\right)}\,\Theta_{\alpha\beta}$$

Why stop at one angularity?

We start from th … ki, Moult, Neill]

With a clever choice of g = -log(f), we could match the LL result

$$\boxed{f_\alpha^{\rm LL} = \frac{\beta}{\alpha t_\beta - \beta t_\alpha}\Theta_{\alpha\beta} + \frac{2\alpha_s C_F}{\pi(\alpha-\beta)}(t_\alpha - t_\beta)\Theta_{\alpha\beta}}$$

Working through the RDF framework, we recover

O(α^1) is captured by the unspecified g function

$$q(t_\alpha, t_\beta|\alpha_s) = \frac{2\alpha_s C_F}{\pi(\alpha-\beta)}\left(1 + \mathcal{O}(\alpha_s)\right) e^{-\alpha_s \frac{C_F}{\pi\beta} t_\beta^2 - \mathcal{O}(\alpha_s)}\Theta_{\alpha\beta}$$

which we can compare with the actual leading log (LL) result

$$q^{\rm LL}(t_\alpha, t_\beta|\alpha_s) = \frac{2\alpha_s C_F}{\pi(\alpha-\beta)}\left(1 + \frac{2\alpha_s C_F}{\pi(\alpha-\beta)}\frac{(t_\beta - t_\alpha)(\beta t_\alpha - \alpha t_\beta)}{\beta}\right) e^{-\alpha_s \frac{C_F}{\pi}\left(\frac{t_\beta^2}{\beta} + \frac{(t_\alpha - t_\beta)^2}{\alpha-\beta}\right)}\Theta_{\alpha\beta}$$

Why stop at one angularity?

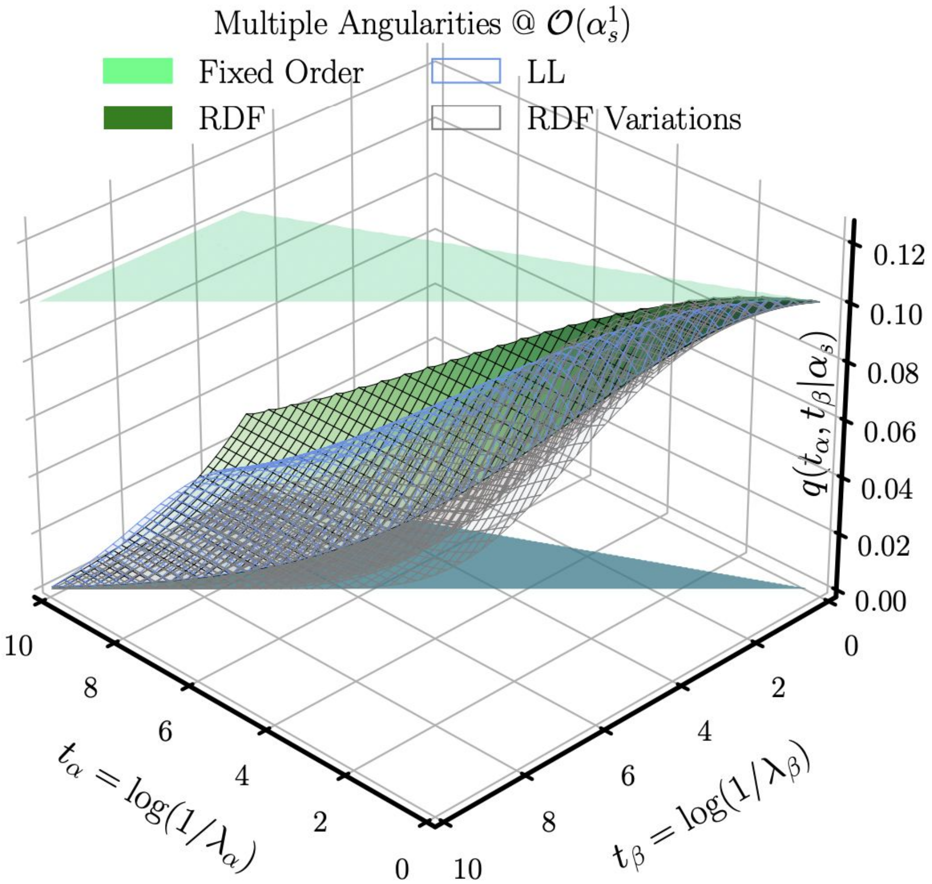

The RDF (green and grey surfaces) normalizes the fixed-order distribution (top green plane).

The RDF family variations capture the LL solution (blue surface).

Example: jet thrust

Let’s consider the jet thrust $T = \max_{\vec{n}} \frac{\sum_i \vec{p_i} \cdot \vec{n}}{\sum_i |\vec{p_i}|}$

We’ll take t = -log(2(1-T)) = -log(2τ)

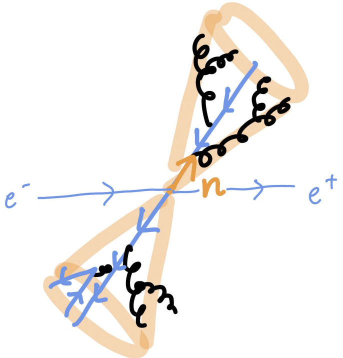

In this case, the fixed-order pdf comes from MC simulation EERAD3v2 [Gehrmann-De Ridder, Gehrmann, Glover, Heinrich], which gives us the binned thrust.

Instead of analytically matching, we can numerically match by minimizing the MSE loss between the fixed-order calculation and the Taylor-expanded RDF:

$$\mathrm{Loss}(\alpha, \phi) = \frac{1}{2} \sum_{\mathrm{Bin}_i} \frac{\left| \sum_{m=m*}^{M} \frac{1}{m!} \frac{\partial^m \mathrm{RDF}}{\partial \alpha^m}(\mathrm{Bin}_i, \alpha, \phi) - \mathrm{Target}(\mathrm{Bin}_i) \right|^2}{\mathrm{Error}(\mathrm{Bin}_i)^2}$$

Example: matching jet thrust to α^1

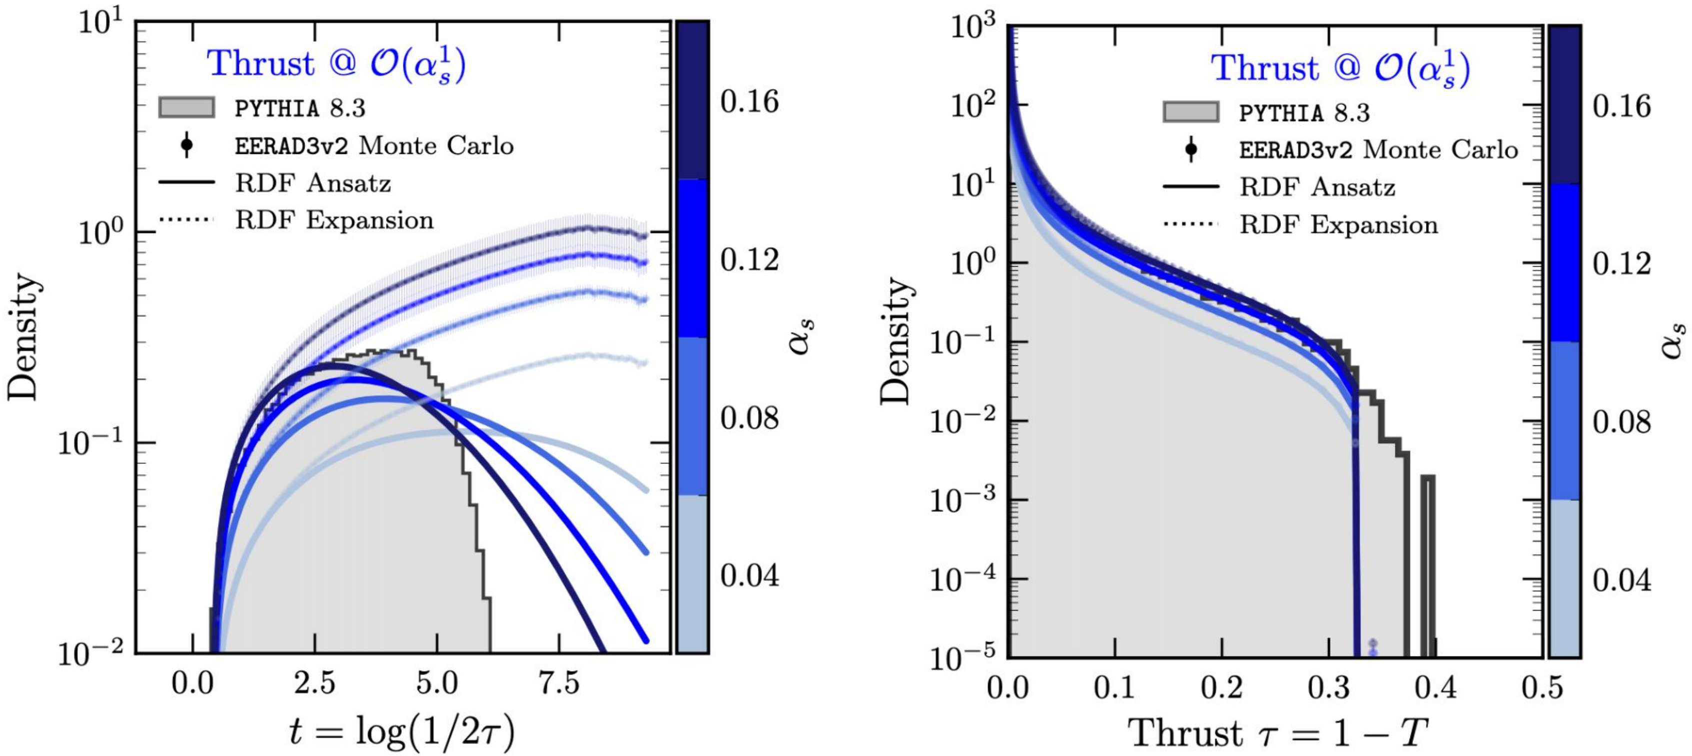

The RDF expansion beautifully matches the fixed-order MC!

We don't match Pythia, which contains nonperturbative physics

Example: matching jet thrust to α^2

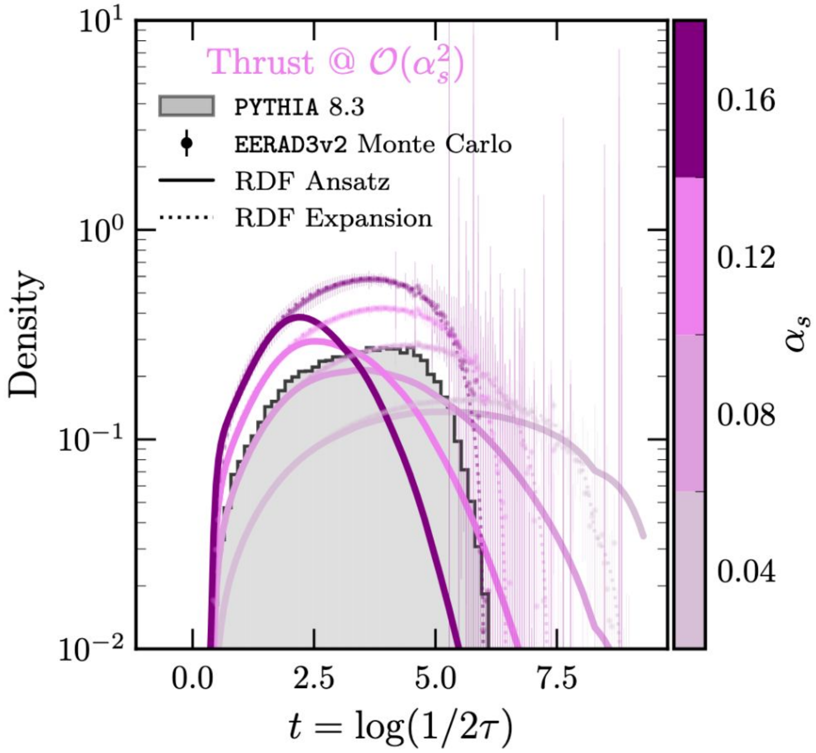

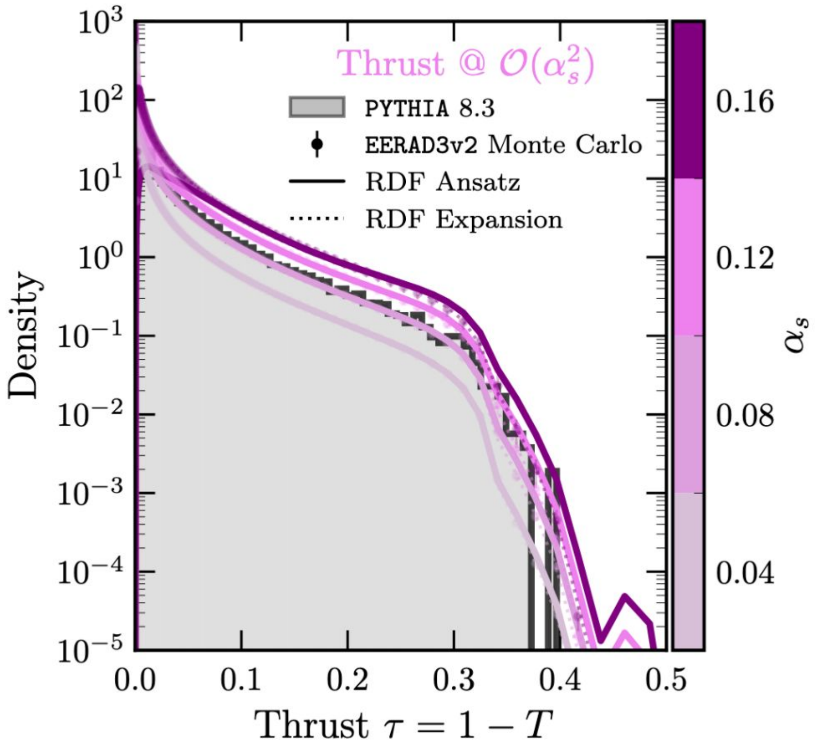

The RDF expansion beautifully matches the fixed-order MC!

We don't match Pythia, which contains nonperturbative physics

Example: matching jet thrust to α^3

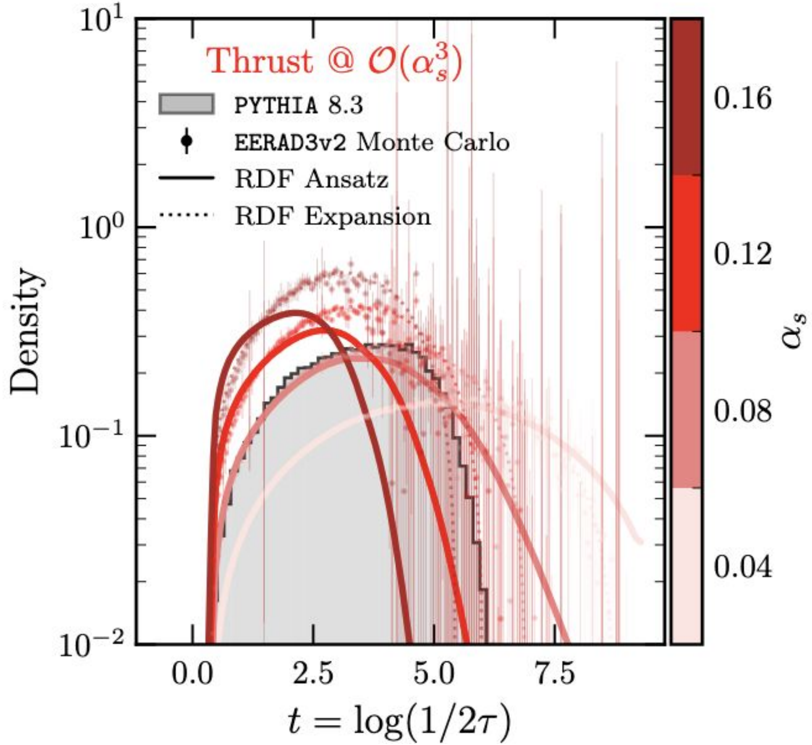

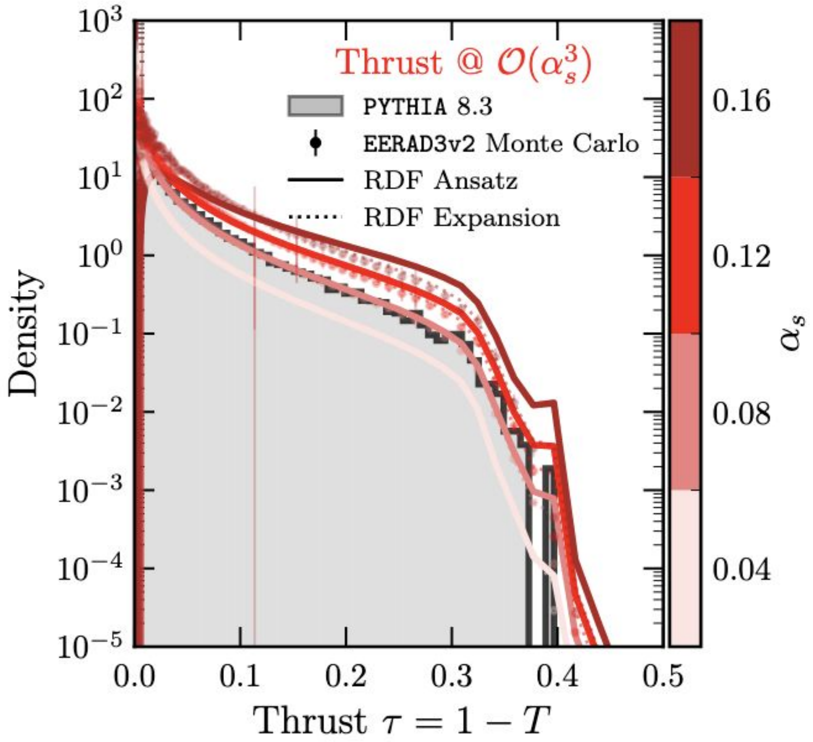

The RDF expansion beautifully matches the fixed-order MC!

We don't match Pythia, which contains nonperturbative physics

What about direct experimental applications of the RDF?

We can do precision measurements*
of the strong coupling constant α!

[hep-ex/0105070]

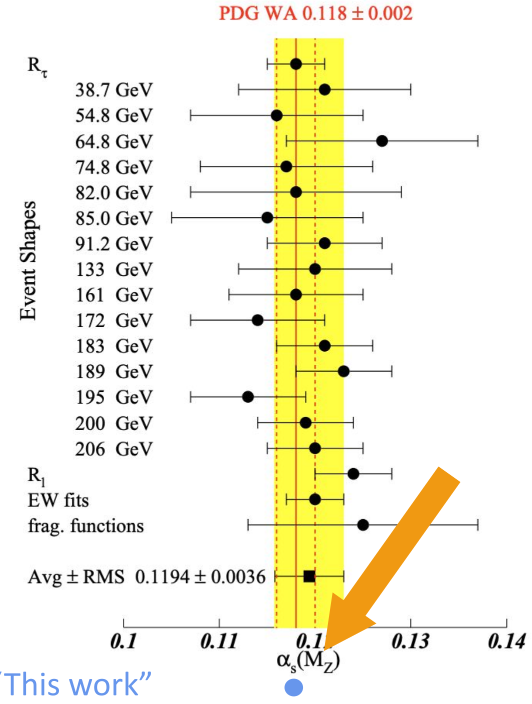

"This work"

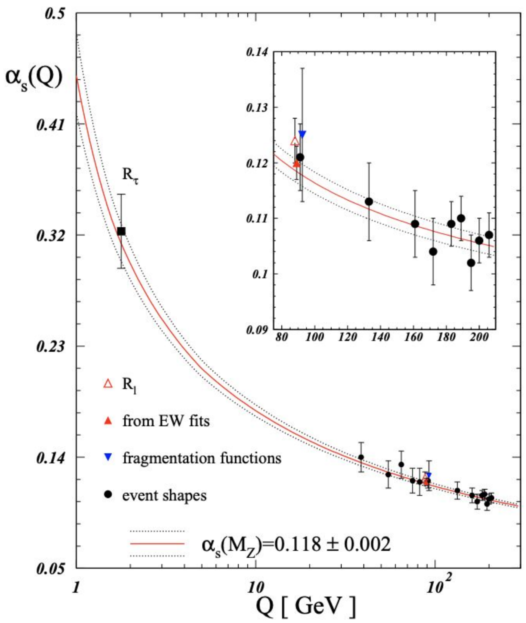

*In the following slides, we do not consider hadronization or experimental systematics, which would be necessary for a true measurement.

α extraction procedure

1. Numerically match g^* and $g_{Analytic}$ to a fixed-order calculation or MC. Freeze!

α extraction procedure

1. Numerically match *g** and $g_{Analytic}$ to a fixed-order calculation or MC. Freeze!
2. Add *nuisance parameters* ν to the g functions governed by priors*

$$g(t,\alpha) = \underbrace{g_{\text{Matched}}(t,\alpha,\phi)}_{\mathcal{O}(\alpha_s^M)} + \underbrace{g_{\text{Nuisance}}(t,\alpha;\nu)}_{\mathcal{O}(\alpha_s^{M+1})}$$

*These priors are *not* dependent on the renormalization scale, although they might in a future version of the RDF.

α extraction procedure

1. Numerically match g^* and $g_{Analytic}$ to a fixed-order calculation or MC. Freeze!
2. Add *nuisance parameters* ν to the g functions governed by priors*

$$g(t,\alpha) = \underbrace{g_{\text{Matched}}(t,\alpha,\phi)}_{\mathcal{O}(\alpha_s^M)} + \underbrace{g_{\text{Nuisance}}(t,\alpha;\nu)}_{\mathcal{O}(\alpha_s^{M+1})}$$

Since the higher-order RDF *g* function is completely unconstrained, we'll want to regulate the *g* coefficients with a reasonable prior.

The choice of this prior directly determines our theory uncertainties.

$$\boxed{\frac{1}{2}\sum_k \frac{|\nu_k - \mu_k|^2}{\sigma_k^2}}$$

$$\boxed{\mu_k = 0} \quad \boxed{\sigma_k = \sigma m! n!}$$

*These priors are *not* dependent on the renormalization scale, although they might in a future version of the RDF.

α extraction procedure

1. Numerically match *g** and $g_{Analytic}$ to a fixed-order calculation or MC. Freeze!
2. Add *nuisance parameters* ν to the g functions governed by priors

$$g(t,\alpha) = \underbrace{g_{\text{Matched}}(t,\alpha,\phi)}_{\mathcal{O}(\alpha_s^M)} + \underbrace{g_{\text{Nuisance}}(t,\alpha;\nu)}_{\mathcal{O}(\alpha_s^{M+1})}$$

3. Define the likelihood function between the RDF and detected data as

$$-\log\mathcal{L}(\alpha,\nu) = \frac{1}{2}\sum_{\text{Bin}_i}\frac{|\text{RDF}(\text{Bin}_i,\alpha,\phi,\nu) - \text{data}(\text{Bin}_i)|^2}{\text{error}(\text{Bin}_i)^2} + \frac{1}{2}\sum_k\frac{|\nu_k-\mu_k|^2}{\sigma_k^2}$$

α extraction procedure

1. Numerically match *g** and *g*$_{Analytic}$ to a fixed-order calculation or MC. Freeze!
2. Add *nuisance parameters* ν to the g functions governed by priors

$$g(t,\alpha) = \underbrace{g_{\text{Matched}}(t,\alpha,\phi)}_{\mathcal{O}(\alpha_s^M)} + \underbrace{g_{\text{Nuisance}}(t,\alpha;\nu)}_{\mathcal{O}(\alpha_s^{M+1})}$$

3. Define the likelihood function between the RDF and detected data as

$$-\log\mathcal{L}(\alpha,\nu) = \frac{1}{2}\sum_{\text{Bin}_i}\frac{|\text{RDF}(\text{Bin}_i,\alpha,\phi,\nu)-\text{data}(\text{Bin}_i)|^2}{\text{error}(\text{Bin}_i)^2} + \frac{1}{2}\sum_k\frac{|\nu_k-\mu_k|^2}{\sigma_k^2}$$

4. Find the α (and ν*)* that minimize the profile likelihood ratio $\log\mathcal{L}(\alpha) = \log\frac{\mathcal{L}(\alpha,\hat{\nu})}{\mathcal{L}(\hat{\hat{\alpha}},\hat{\hat{\nu}})}$

   Define confidence intervals by $-2\Delta\mathcal{L} = 1$

Testing α extraction: an exponential observable

For each order α^M, we

1. Numerically match $g*$ and $g_{Analytic}$ to

$$p_{\mathrm{FO}}^{(\mathrm{Exp})}(t|\alpha) = \alpha \sum_{m=0}^{M-1} \frac{1}{m!}(-\alpha t)^m + \mathcal{O}(\alpha^{M+1})$$

   which fixes up to g_{M7}

2. Augment the RDF with coefficients $g_{\{M+1\}7}$ — these are the nuisance parameters ν

3. Define the likelihood function between the RDF and pseudodata

4. Minimize the profile likelihood ratio

Testing α extraction: an exponential observable

For each order α^M, we

1. Numerically match g^* and $g_{Analytic}$ to

$$p_{\mathrm{FO}}^{(\mathrm{Exp})}(t|\alpha) = \alpha \sum_{m=0}^{M-1} \frac{1}{m!}(-\alpha t)^m + \mathcal{O}(\alpha^{M+1})$$

   which fixes up to g_{M7}

2. Augment the RDF with coefficients $g_{\{M+1\}7}$ — these are the nuisance parameters ν

3. Define the likelihood function between the RDF and pseudodata

4. Minimize the profile likelihood ratio

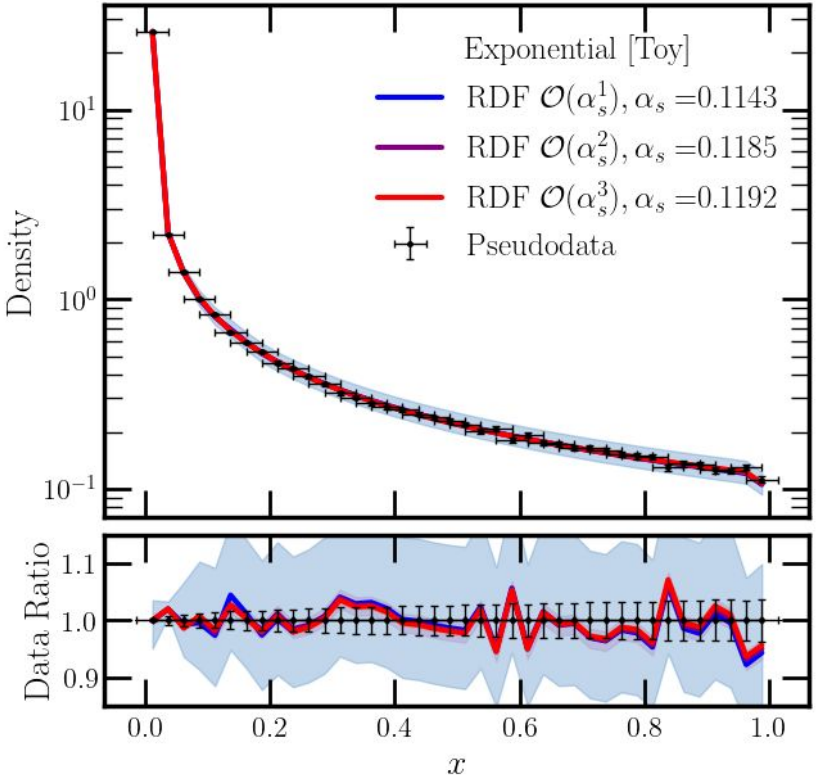

Testing α extraction: an exponential observable

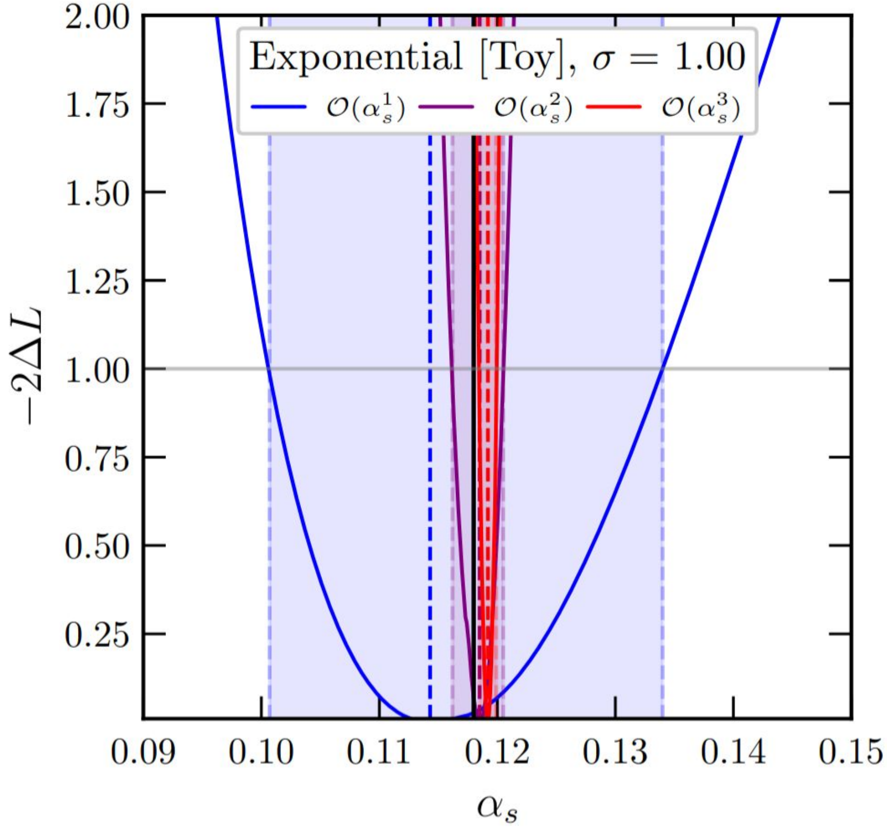

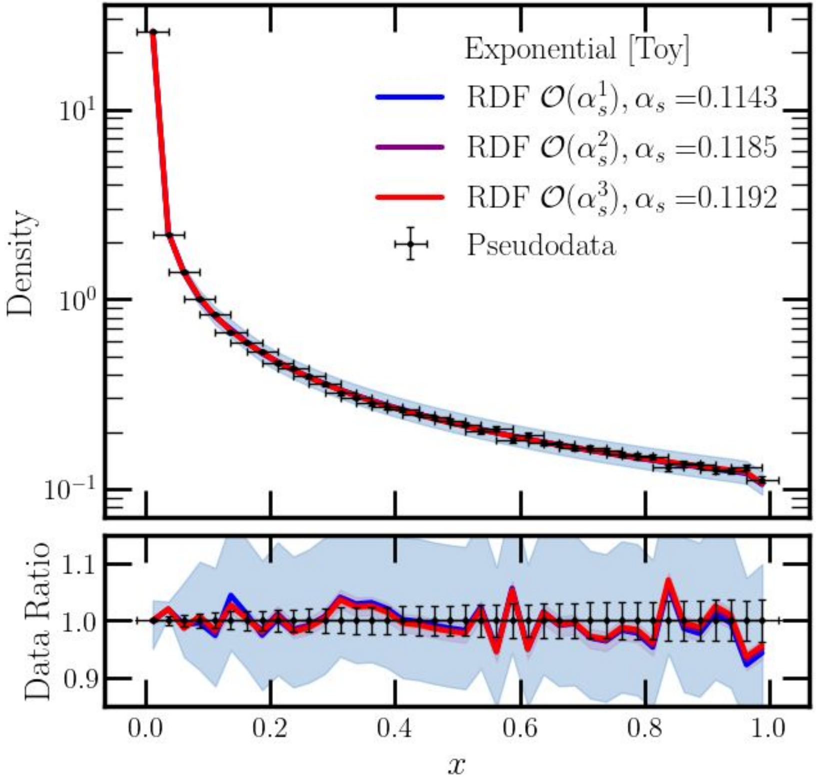

Testing α extraction: an exponential observable

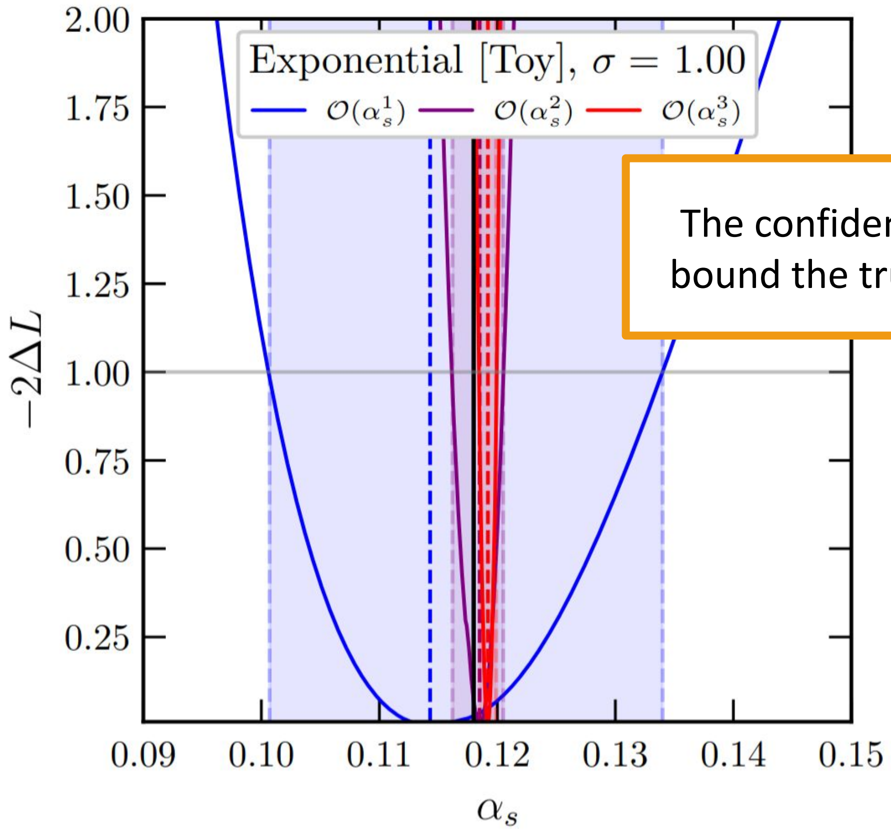

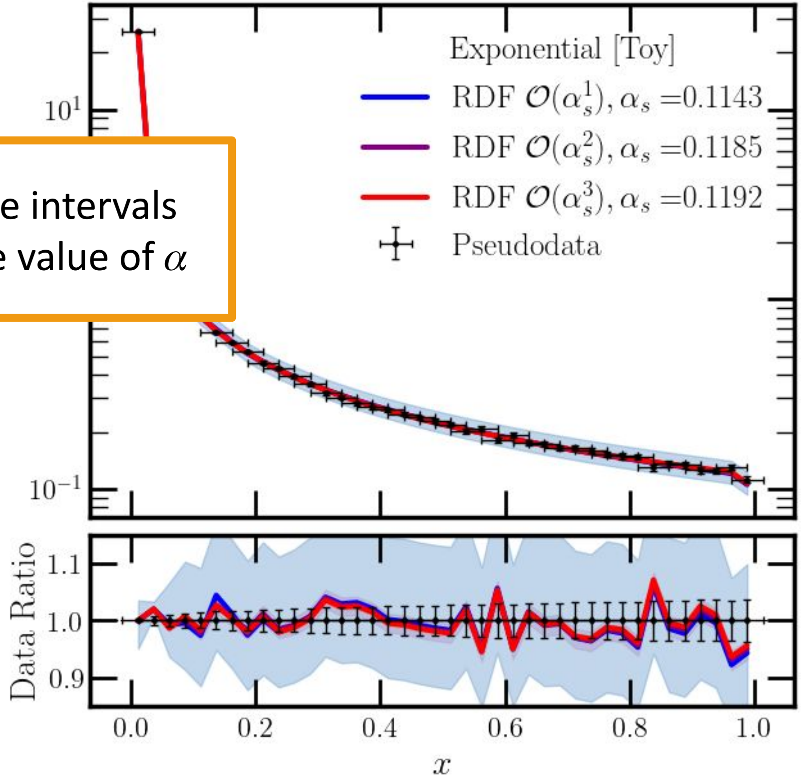

The confidence intervals bound the true value of α

Testing α extraction: an exponential observable

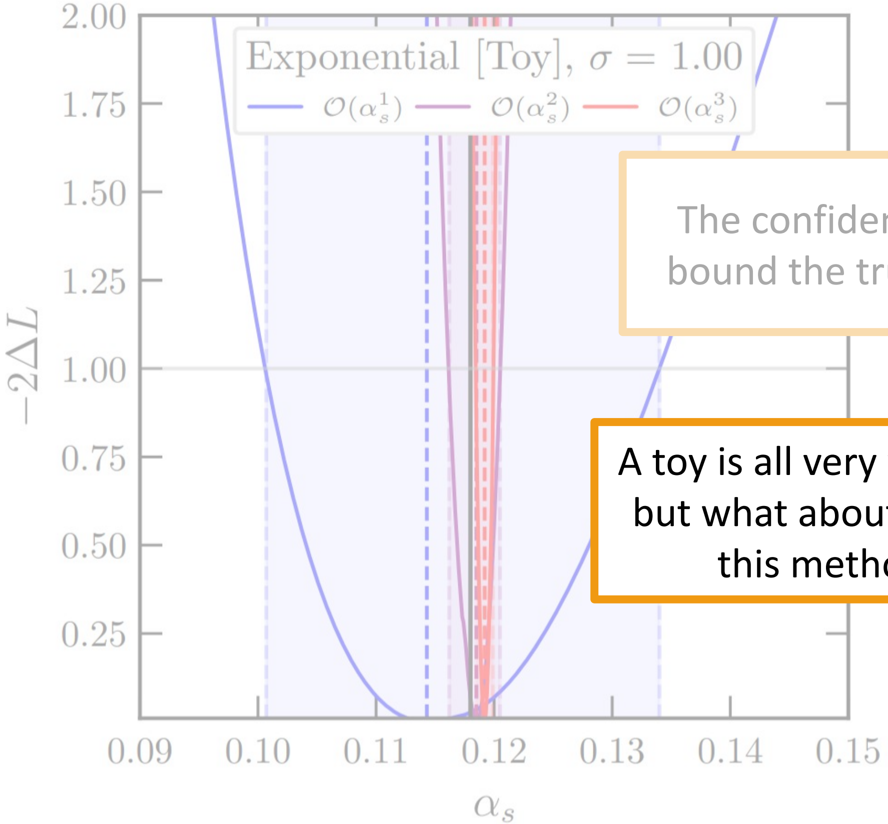

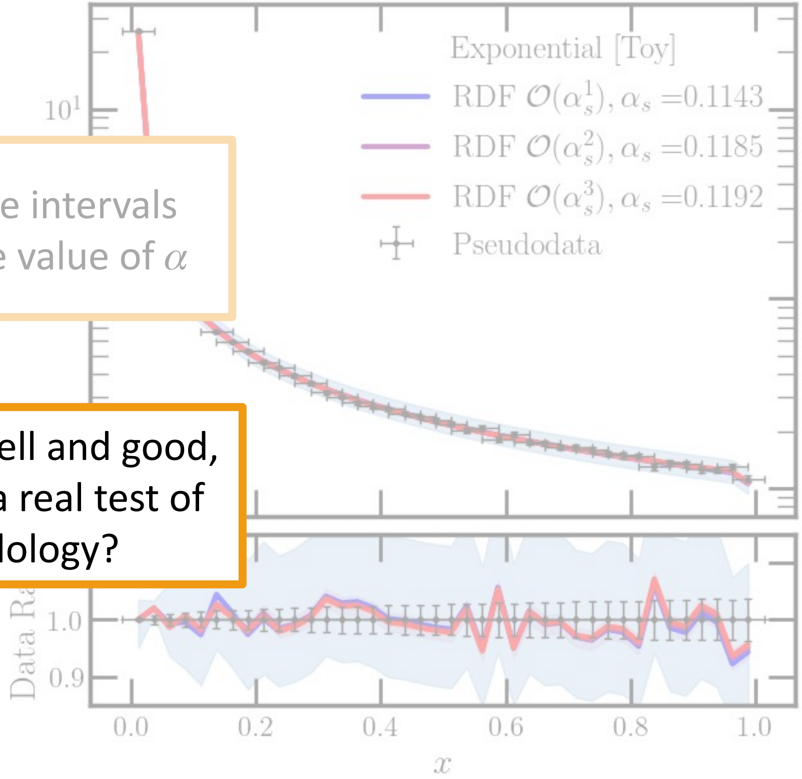

The confidence intervals bound the true value of α

A toy is all very well and good, but what about a real test of this methodology?

Precision measurements at the ALEPH experiment

The ALEPH (active 1980-2002) was a general-purpose Standard Model (and beyond) experiment hosted at LEP.

It was similar to CMS and ATLAS in design and physics scope, but it observed e^+e^- collisions instead of *pp* collisions.

e^+e^- collisions give a clean environment to make a precision measurement of α. A very common way to do this is by fitting to the thrust distribution. Let's try it out!

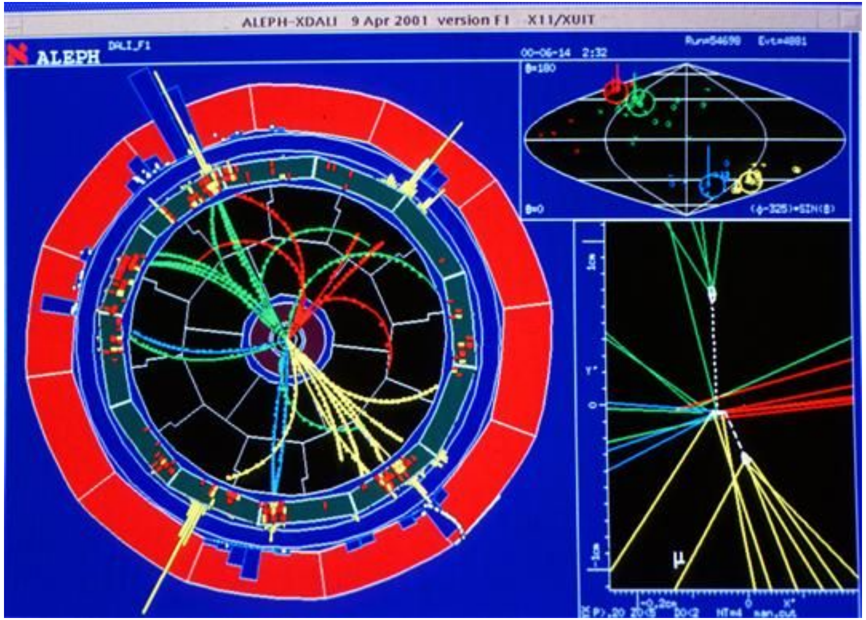

Image

Applying α extraction: fitting to ALEPH data

For each order α^M, we

1. Numerically match g^* and $g_{Analytic}$ to the EERAD3v2 prediction (which we showed earlier in this talk), which fixes up to g_{M7}

2. Augment the RDF with coefficients $g_{\{M+1\}7}$ — these are the nuisance parameters ν

3. Define the likelihood function between the RDF and ALEPH data

4. Minimize the profile likelihood ratio

Applying α extraction: fitting to ALEPH data

For each order α^M, we

1. Numerically match g^* and $g_{Analytic}$ to the EERAD3v2 prediction (which we showed earlier in this talk), which fixes up to g_{M7}

2. Augment the RDF with coefficients $g_{\{M+1\}7}$ — these are the nuisance parameters ν

3. Define the likelihood function between the RDF and ALEPH data

4. Minimize the profile likelihood ratio

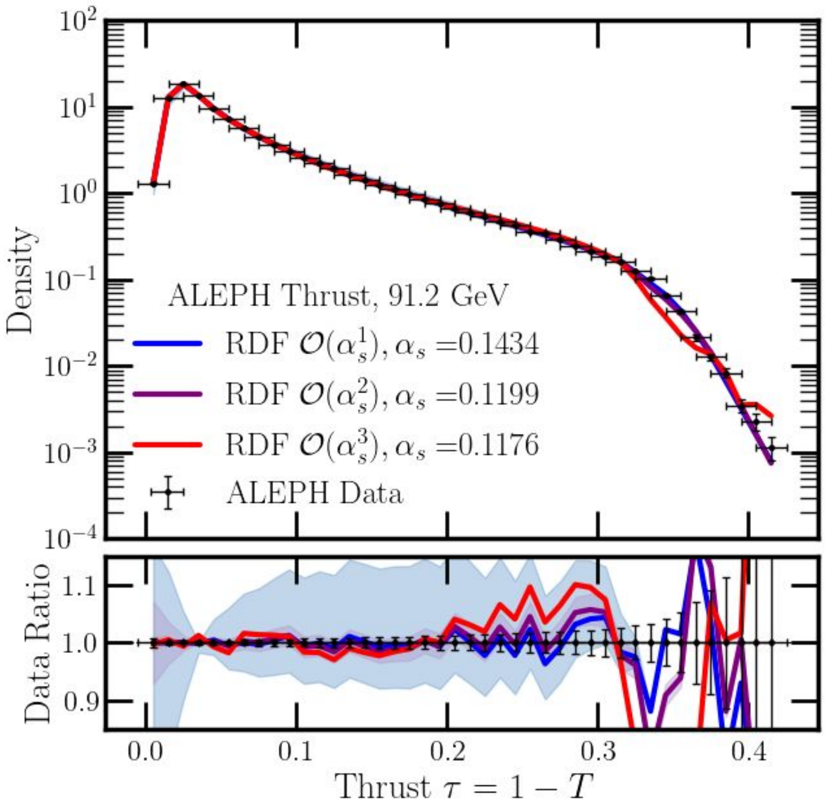

Applying α extraction: fitting to ALEPH data

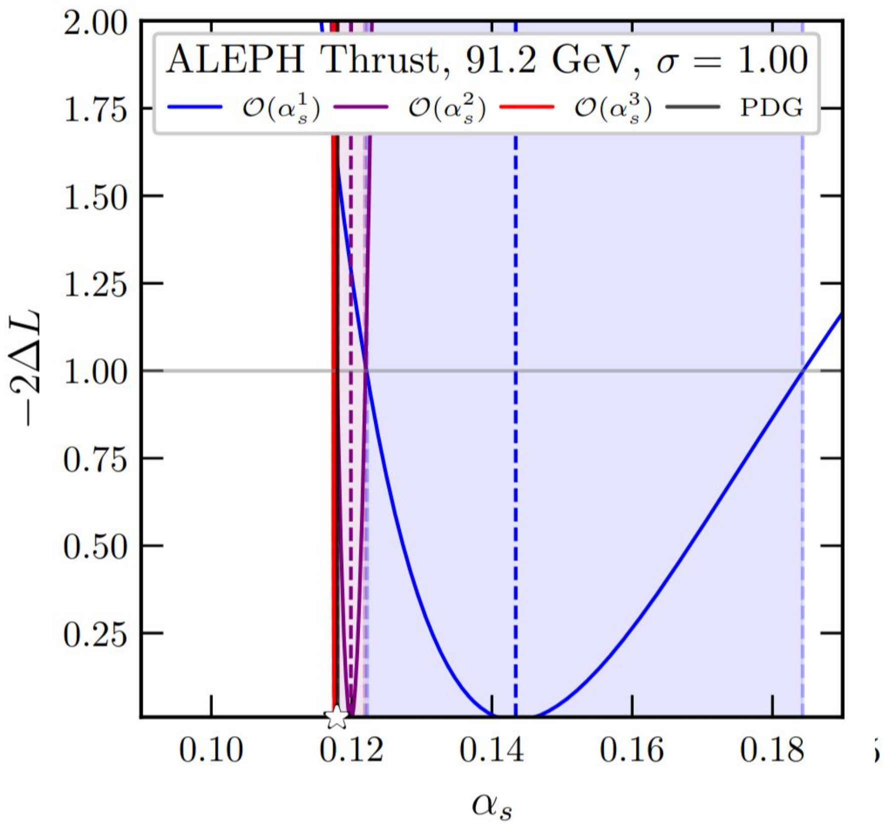

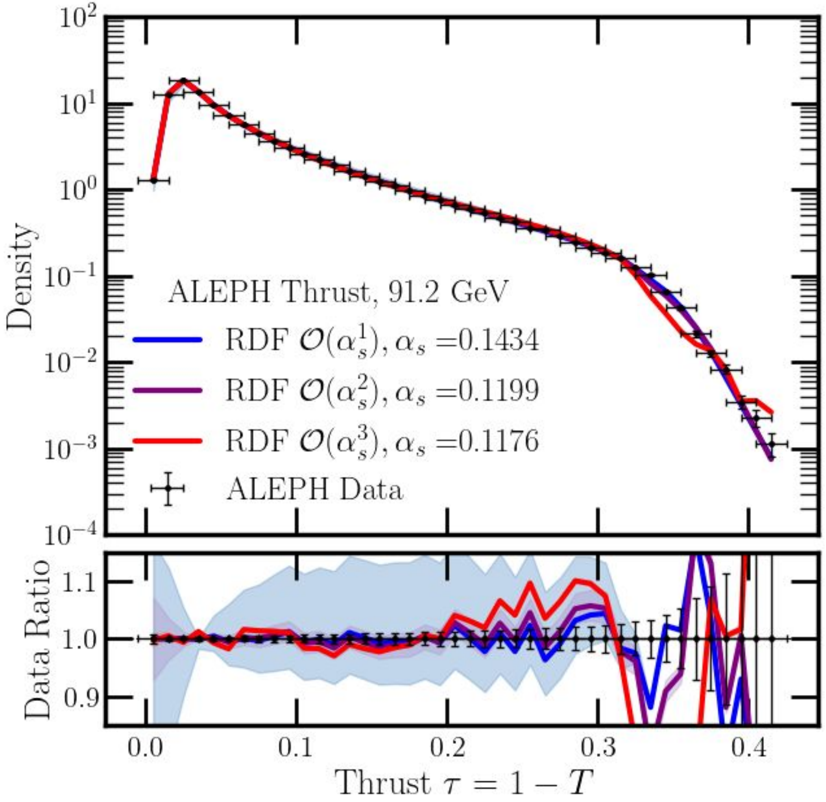

Applying α extraction: fitting to ALEPH data

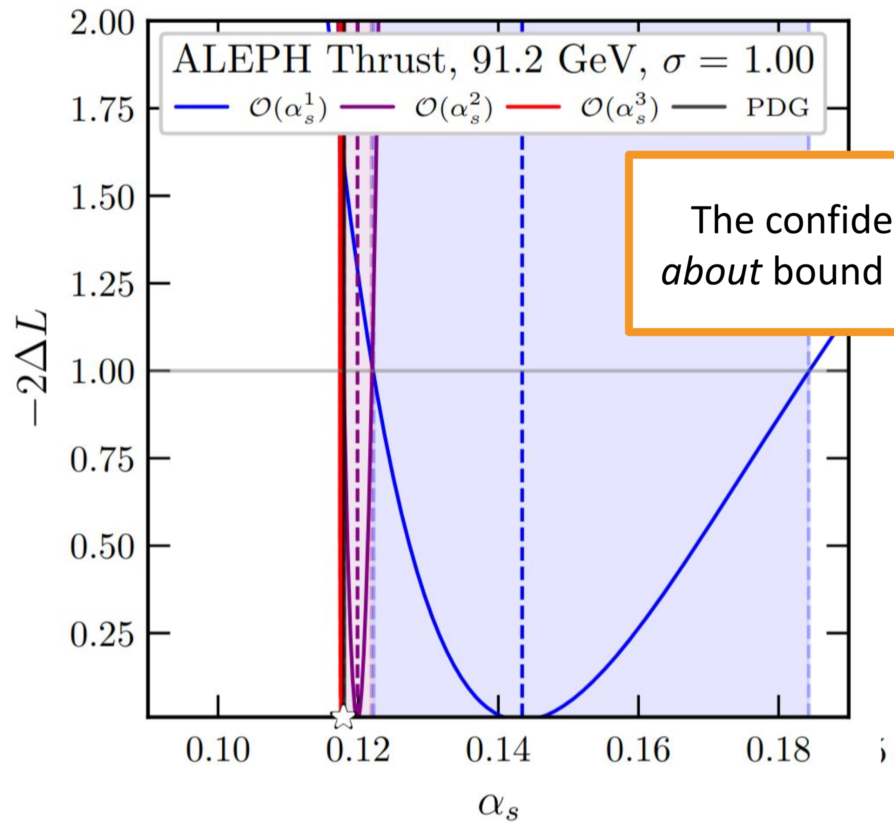

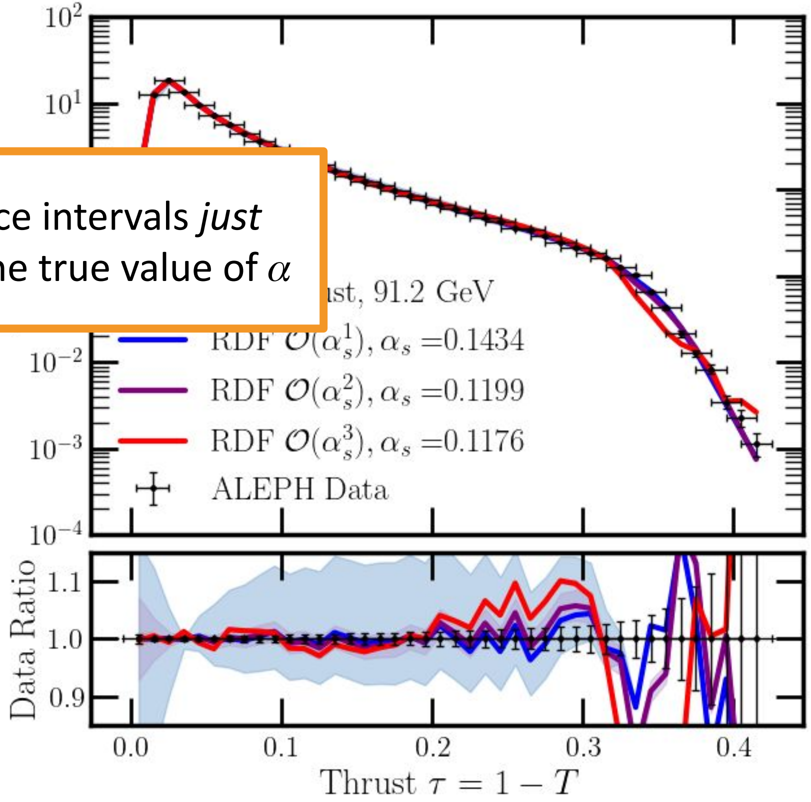

The confidence intervals *just about* bound the true value of α

Applying α extraction: fitting to ALEPH data

	RDF Prior σ			
Order	$\sigma = 0.5$	$\sigma = 1.0$	$\sigma = 2.0$	LEP I+II [Becher, Schwartz]
Thrust $\mathcal{O}(\alpha_s^1)$	$0.1555^{+0.0189}_{-0.0139}$	$0.1434^{+0.0409}_{-0.0212}$	$^{*}\mathit{0.19}^{+0.0000}_{-0.0595}$	0.1142 ± 0.0297
Thrust $\mathcal{O}(\alpha_s^2)$	$0.1202^{+0.0010}_{-0.0009}$	$0.1199^{+0.0019}_{-0.0020}$	$0.1249^{+0.0067}_{-0.0053}$	0.1152 ± 0.0068
Thrust $\mathcal{O}(\alpha_s^3)$	$0.1202^{+0.0002}_{-0.0001}$	$0.1176^{+0.0002}_{-0.0002}$	$0.1164^{+0.0003}_{-0.0002}$	0.1164 ± 0.0033

The O(α^2) *and* O(α^3) RDFs allow us to extract a value of α that is consistent with the world average, using only fixed-order information!

The results are generally robust with respect to the prior for the nuisance parameters.

Comparison to other methods for defining uncertainties

Method	Description	Compared to RDF
Generic scale variation uncertainties	Define uncertainty bands by recomputing results at $[0.5, 2] \times \mu_R$	Arbitrary and without a probabilistic or physical interpretation
Theory Nuisance Parameter [Tackmann]	Parametrize higher order uncertainties in terms of physical theoretical parameters (TNPs)	Model-specific; requires assumptions on what the TNP should be
Cacciari-Houdeau and variations [Bonvini], [Bagnaschi, Cacciari, Guffanti, Jenniches]	Estimate truncation uncertainties through a Bayesian approach that assumes all series coefficients are bounded	Less fine-grained — constraints are enforced for powers of α, while the RDF also uses t as a degree of freedom

Conclusions: why use the RDF?

We have introduced the Resummed Distribution Function, a flexible framework that encodes fixed-order calculations into probability distributions.

The RDF is not a free lunch: it only models the perturbative physics that you put in. Nevertheless, it meaningfully carries out resummation with a single constraint: unitarity.

With the RDF, we can do real QCD — uncertainties and all — with only fixed-order calculations.

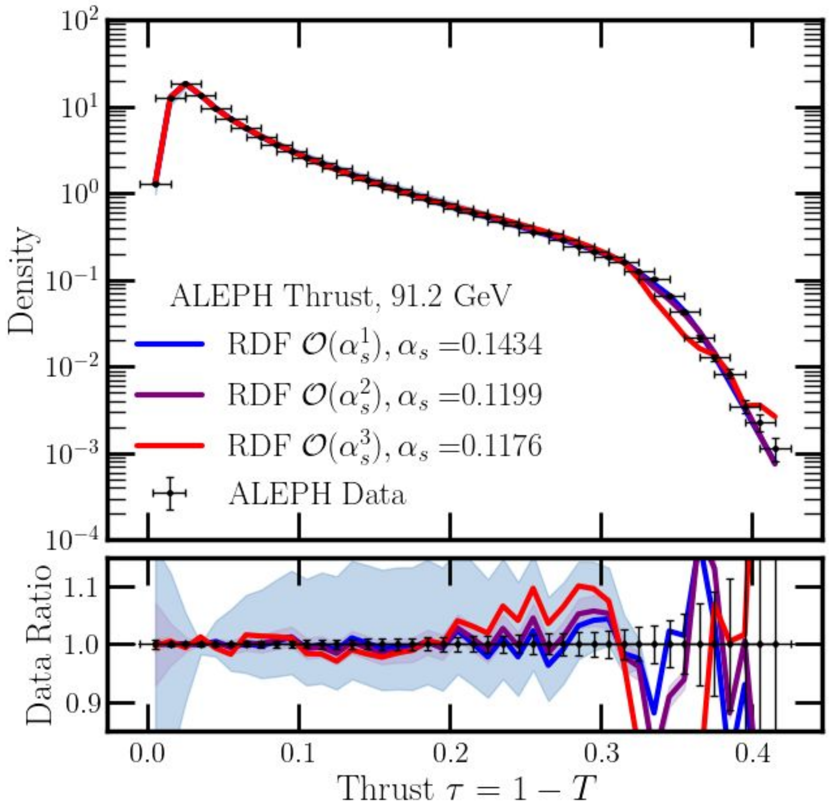

Where could you use the RDF?

- Testing this framework on other event observables (particularly the α extraction)
- Choosing more reasonable priors for the nuisance parameters in α extraction
- Embedding renormalization and factorization scales into the RDF framework to better resemble the typical way of generating theory uncertainties
- In progress: extracting α from jet cross-section ratios in data
- Wild idea: replacing event generators with this framework…

Try the code yourself! [link]

Backup slides

Matching example: exponential toy

$$p_{\text{Exponential}}(t|\alpha) = \alpha e^{-\alpha t}$$

$$p_{\text{FO}}^{(\text{Exp})}(t|\alpha) = \alpha \sum_{m=0}^{M-1} \frac{1}{m!}(-\alpha t)^m + \mathcal{O}(\alpha^{M+1})$$

$$g(t,\alpha) = -\log(\alpha) + \left\{ \sum_k \frac{1}{k} \left(\left[-\sum_{m>1}^{M-1} \frac{(-\alpha t)^{m-1}}{(m-1)!} \right]^k - \left[\sum_{m=1}^{M-2} \int_0^t dt' \frac{(-t')^{m-1}}{(m-1)!} \alpha^m \right]^k \right) + \mathcal{O}\left(\alpha^M\right) \right\} \tag{4.4}$$

$\downarrow$ (Carrying out the integral in the third term)

$$= -\log(\alpha) + \left\{ \sum_k \frac{1}{k} \left(\left[-\sum_{m=2}^{M-1} \frac{(-\alpha t)^{m-1}}{(m-1)!} \right]^k - \left[\sum_{m=1}^{M-2} -\frac{(-\alpha t)^m}{m!} \right]^k \right) + \mathcal{O}\left(\alpha^M\right) \right\} \tag{4.5}$$

$$= -\log(\alpha) + \left\{ \sum_k \frac{(-1)^k}{k} \left(\left[\sum_{m'=1}^{M-2} \frac{(-\alpha t)^{m'}}{(m')!} \right]^k - \left[\sum_{m=1}^{M-2} \frac{(-\alpha t)^m}{m!} \right]^k \right) + \mathcal{O}\left(\alpha^M\right) \right\} \tag{4.6}$$

$$= -\log(\alpha) + \mathcal{O}\left(\alpha^M\right). \tag{4.7}$$

$$g^*(t,\alpha) = \alpha$$

$$g_{\text{Analytic}}(t,\alpha) = 0 + \mathcal{O}(\alpha^M).$$

$$q^{(\text{Exp})(M)}(t|\alpha) = g^*(t,\alpha) \exp\left[-g_{\text{Analytic}}(t,\alpha) - \int_0^t dt'\, g^*(t',\alpha) \exp\left(-g_{\text{Analytic}}(t',\alpha)\right) \right]$$
$$= \alpha \exp\left(-\alpha t\right) + \mathcal{O}(\alpha^{M+1}).$$

Analytic and numeric *g* functions

Analytic

$$g^*(t,\alpha),\;\; g_{\text{Analytic}}(t,\alpha) = \sum_{m=0,n=0}^{M,N} g_{mn}\alpha^m t^n \Theta(t-\theta_{mn})$$

Numeric

$$g(t,\alpha) = -\log\left(\sum_{m=m^*}^{M} \frac{\alpha^m}{m!}\left|\sum_{n=0}^{N} g^*_{mn}\frac{t^n}{n!}\Theta_{T_m}(t-\theta^*_m)\right|_{T_m}\right) + \sum_{m=0,n=0}^{M-m^*,N} g_{\text{Analytic}_{mn}}\frac{t^n}{n!}\frac{\alpha^m}{m!}\Theta_{T_m}(t-\theta_{\text{Analytic}_m})$$

Differences:

- ❖ Explicit factorial scaling of the g_{mn} coefficients for numerical stability
- ❖ Smoothed sigmoid-like absolute values and theta functions with learnable temperatures

α extraction: Rayleigh toy

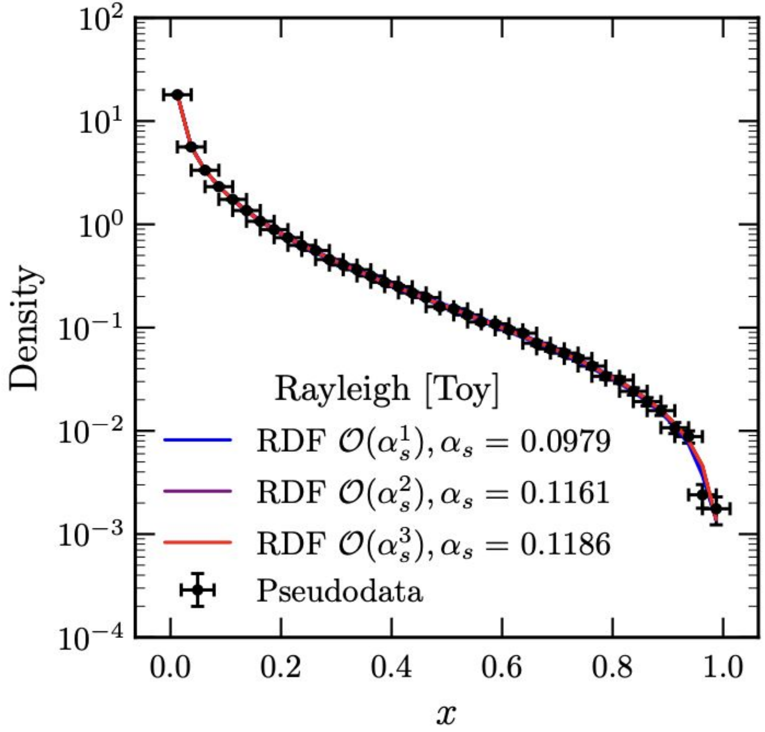

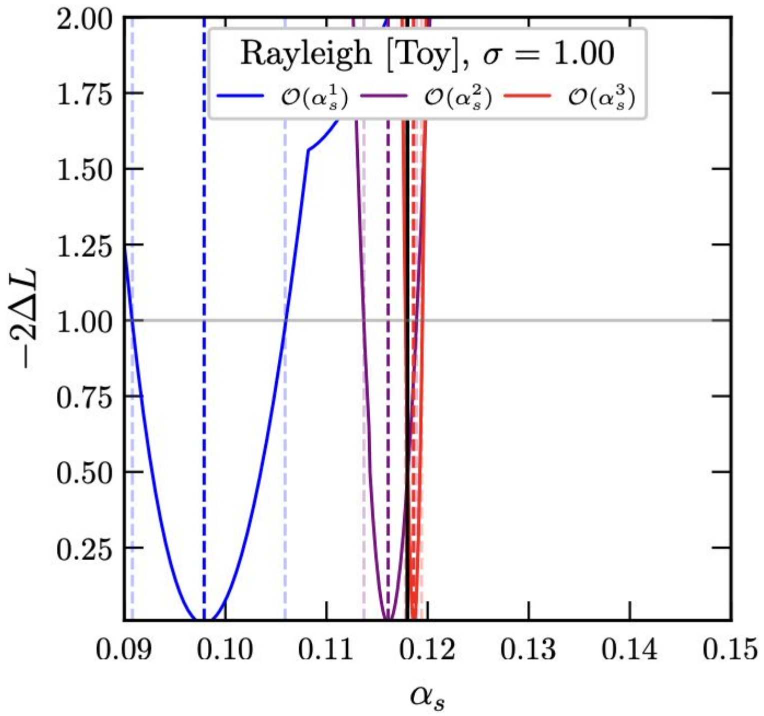

α extraction: ALEPH data with other priors