GEANT4 Dark Bremsstrahlung from MADGRAPH

Tom Eichlersmith

he/him/his

University of Minnesota

eichl008@umn.edu

Jeremiah Mans, Joe Muse, Michael Revering

January 24, 2024

GEANT4 Dark Bremsstrahlung from MADGRAPH
Tom Eichlersmith
he/him/his
University of Minnesota
eichl008@umn.edu
Jeremiah Mans, Joe Muse, Michael Revering
January 24, 2024

"Into-the-weeds" discussion of paper about this project.

- ▸ paper on inspirehep (read for full context)
- ▸ source code on GitHub

This talk will assume no prior contact with this paper and will spend some time motivating the work and defining the necessary vocabulary.

Background

Background

- Monte Carlo? That sounds fancy
- Weizacker-Williams approximation
- What is dark bremsstrahlung?
- Geant4 and MadGraph description

Monte Carlo Simulation in HEP

Simple Example

LDMX

We are experimenters at LEP and we want to simulate the $e^- e^+ \to \mu^- \mu^+$ process.

1. Calculate an expression for the differential cross section ▸ Physics StackExchange
2. As a quantum process, the outgoing particles will not be fully pre-determined; however, this process is well constrained – only the muon outgoing angles are randomly distributed
   - ϕ uniformly between 0 and 2π
   - θ according to a more complicated function of trigonometric functions
3. Sample values from these angular distributions and then calculate the corresponding muon properties
   - Use the uniformly-distributed stream of random numbers provided by some ▸ RNG and transform it with ▸ Inverse Transform Sampling or ▸ Rejection Sampling
4. Further analysis or simulation (e.g. where do these muons exit DELPHI's muon chambers assuming no other material interaction?)

To provide some context on general simulation a.k.a. "Monte Carlo" in HEP, let's go through a very simple example.

Go through text on slide

- "constrained": we know the initial four-momenta to a (relatively) high degree of certainty, so the outgoing muon four-momenta only have two degrees of freedom.
- More on how nice this form is on the next slide

Key

We have an expression for the distribution we want to sample from that has particular properties.

Step 3 is an important point to highlight: in HEP this sampling step attempting to replicate a physical process is often called "event generation".

Important Point

Since our differential cross section is **separable, integrable, and invertible**, we can use ITS to sample from it.

$$\frac{d\sigma}{d\Omega} = \frac{\alpha^2}{16E^2}\left(1 + \cos^2\theta\right)$$

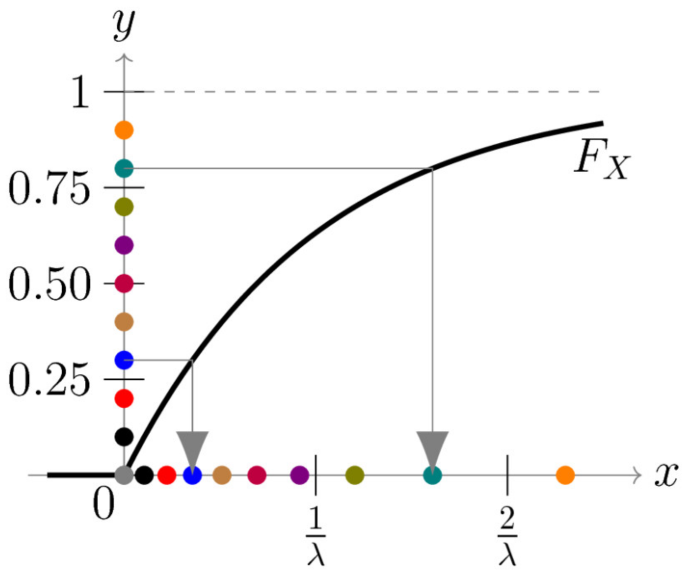

With **Inverse Transform Sampling**

- Quickly and easily transform our uniformly random numbers into our desired distribution

Bad News

A lot of physics distributions are not separable, integrable, or invertible (many of which satisfy *none* of these criteria!)

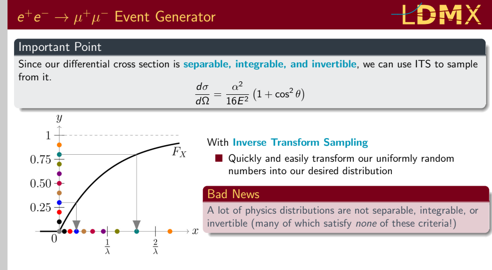

In this simple toy example, our process is very well constrained. We know the initial energies and so the only parts of the muon that are left to be randomly distributed (the two angles) are separated, their distributions are indefinitely integrable, and those indefinite integrals are invertible. This enables us to use ITS.

Unfortunately for us, while this procedure works for many physical processes (and is used widely in simulation libraries used throughout HEP), many processes do not have such simple cross sectional forms.

Foreshadowing...

Without ITS, we will need to fall back to another form of sampling.

Weizäcker-Williams Approximation

LDMX

Physically-Motivated Approximation

- Treat high energy incident particle as photon beam (e.g. massless)
- Assume cross section at minimum momentum transfer is dominant contribution
- Include target particle complexity via form factors

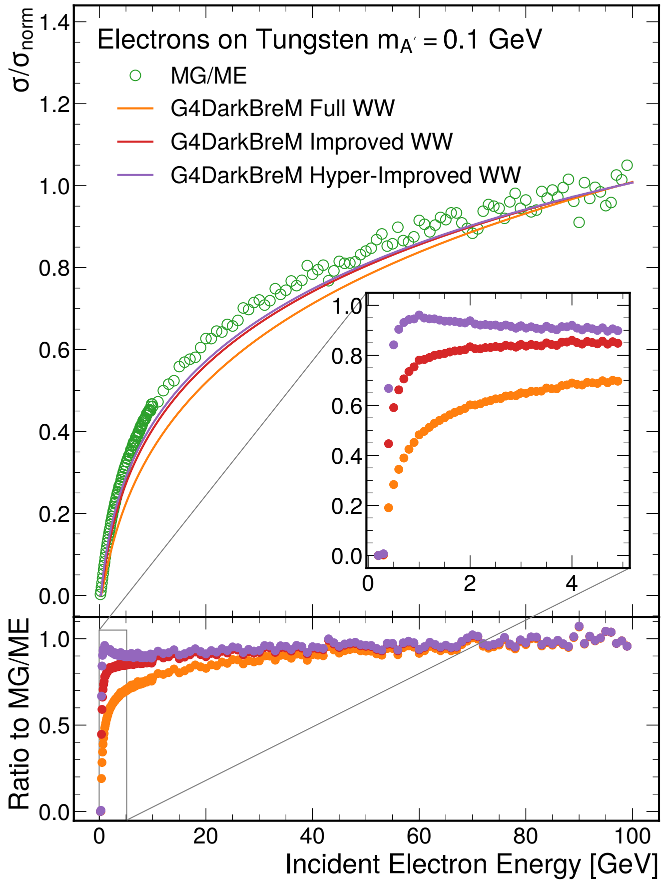

Variety of Implementations

These physically-motivated approximations can still be applied in a variety of mathematical ways.

- doi:10.1103/PhysRevD.8.3109
- ArXiV:9310359
- doi:10.1016/0370-2693(72)90622-3

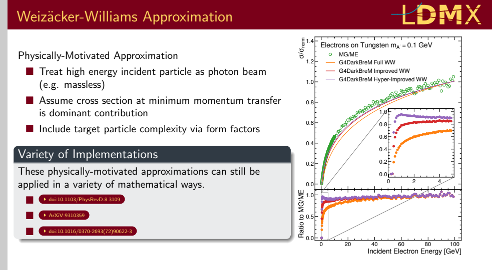

Been around for awhile (since before τ discovery)

Intentionally Vague

1. These two physical approximations can be applied mathematically in a variety of ways
2. My own personal lack of familiarity with the actual derivations

Point of Confusion

Often, we intentionally divide by the high-energy behavior of the cross section because the overall scale of the cross section floats with ϵ. The relative rate between different energy points is what we care about for the process since biasing and re-weighting can be used to handle any overall-scale differences.

Boilerplate Slide on Dark Matter

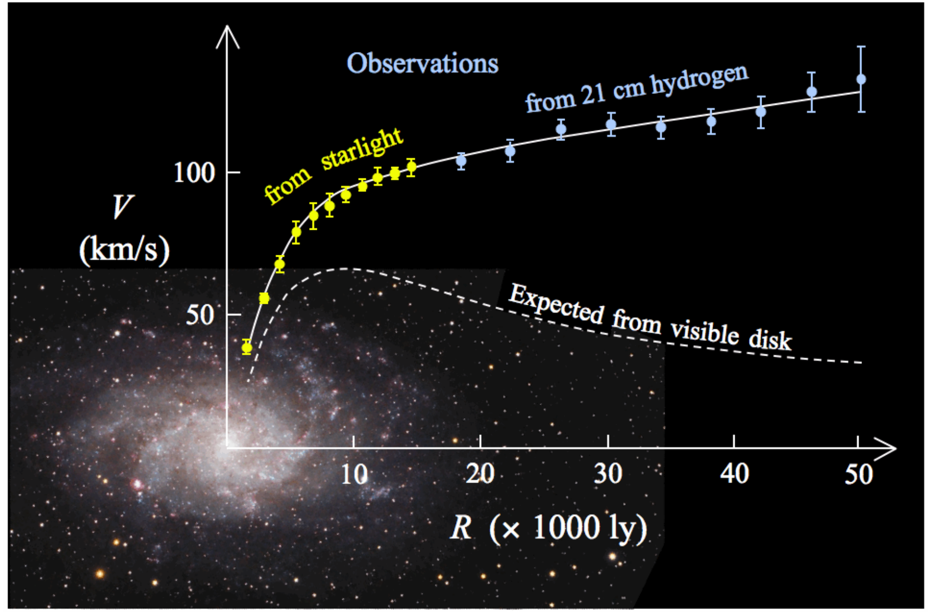

Figure: Galactic rotation curves deviating from the prediction using visible mass.

- Plenty of evidence for "Dark Matter" which is a specific (unidentified) particle
  (one piece – glactic rotation curves – shown to left)
- Thermal relic DM can be represented by a "dark sector" that mixes with standard particles via a heavy spin-1 boson
  (mixing diagram show below)

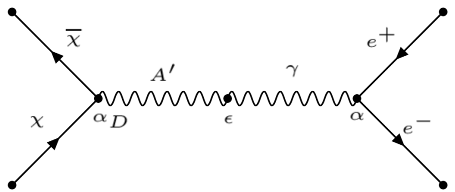

Figure: Feynman diagram for SM-DM mixing.

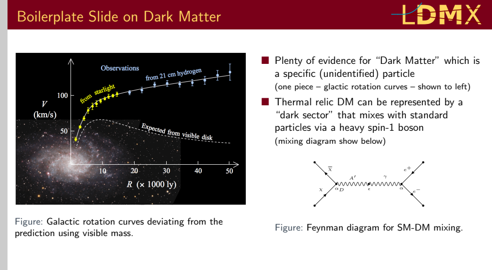

Probably my densest slide since I am explicitly skipping *a lot* of information.

1. We have a lot of evidence for the existence of DM and its existence as a particle. (galactic rotation curves, bullet cluster, CMB and BAO, . . .)
2. Motivates a so-far unseen particle that currently exists in the universe.
3. CMB data implies that this particle existence since that era ⇒ naturally assume existence since big bang
4. "Thermal Relic" that was in thermal equilibrium with standard matter but then "froze out" into its present density
5. Needs some interaction with standard matter allowing mixing
6. For lower mass DM (a.k.a. "Light" DM), we hypothesize existence of heavy spin-1 boson that can then mix with standard photon (higher mass DM can use the standard weak interaction – WIMPs)

Dark Bremsstrahlung

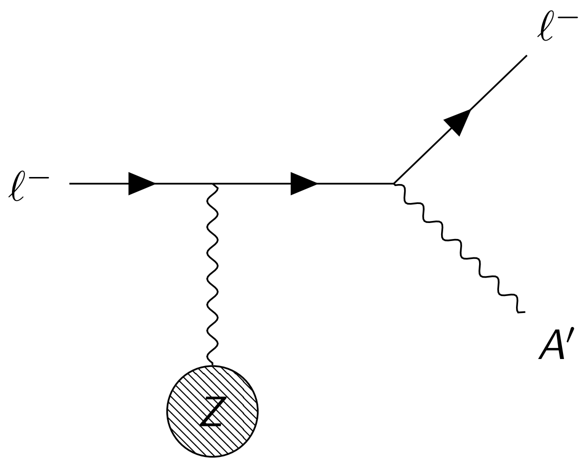

Figure: Feynman diagram for the dark brem process.

Production

With this additional interaction (inducing an effective interaction between leptons and hypothetical dark photon A'), we can produce material in the dark sector in experiments.

Same diagram as bremsstrahlung except emitting a dark photon $\Rightarrow$ **dark brem**

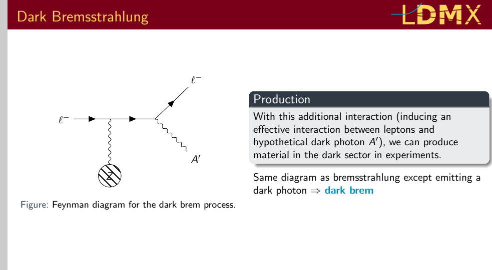

With the additional interaction connecting the dark sector with the standard particles, we can re-diagonalize and find an effective coupling between standard charged particles and the dark photon. This effective coupling scales with the mixing strength ϵ so it could be weak enough to be as-yet unseen in our experiments; nevertheless, it offers a simple mode of producing dark photons within our terrestrial experiments.

While the nucleus appears to simply be standing by in this diagram, it is important to remember that it is *necessary* to conserve four momentum. The pesky nucleus in this $2 \rightarrow 3$ process is a reason for developing a new simulation method.

Tools

Goal: simulate dark brem interactions within our detectors

Geant4

C++ library for simulating particle interactions with bulk materials. Used within many experiments to observe how particles of varying types and energies interact with detector designs.

MadGraph/MadEvent

Warning: experimentalist description. Deduce feynman rules into `Fortran` code (MadGraph) and do a random-sampling integration of the phase space to estimate the cross section and sample particle kinematics (MadEvent).

Tools
Goal: simulate dark brem interactions within our detectors
Geant4
C++ library for simulating particle interactions with bulk materials. Used within many experiments to observe how particles of varying types and energies interact with detector designs.
MadGraph/MadEvent
Warning: experimentalist description. Deduce feynman rules into Fortran code (MadGraph) and do a random-sampling integration of the phase space to estimate the cross section and sample particle kinematics (MadEvent).

Now, our overall goal is to simulate dark brem interactions within our detectors.
Geant4 is widely used (and is used by LDMX and CMS) to do simulations of particle interactions with detectors, so devleoping some method of dark brem simulation within Geant4 is crucial.
MG/ME is helpful for developing some source of "truth" about how the dark brem interaction should behave. For our purposes, it can do estimations of the cross section and sample particle kinematics of events in a way that is "closer" to the theory itself since it actually involves a particle that can represent the nucleus. MG/ME is already capable of simulating the dark brem interaction; however, it is not able to simulate particles interacting with bulk materials.
Note We do not compare to other event generators. You will see that this method is largely unchanged if you would use another generator.

The Problem

The Problem

Those of you more familiar with Geant4 may now be wondering what the issue is. One can define any particles they want and freely define how they interact with any other particles.

The Nucleus

- No general analytic solution
- Standard approximation (WW) not well behaved

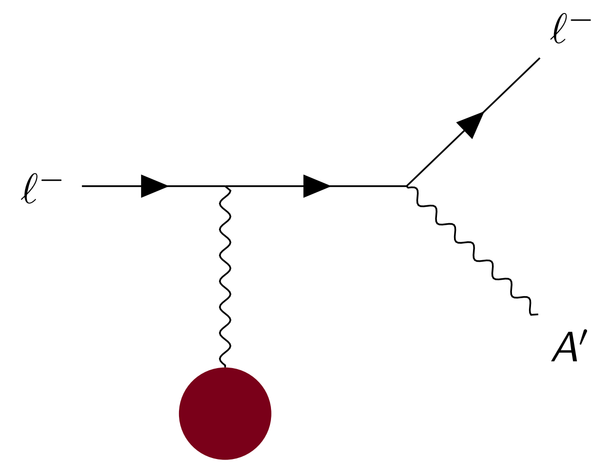

Figure: Feynman diagram for the dark brem process with pesky nucleus highlighted.

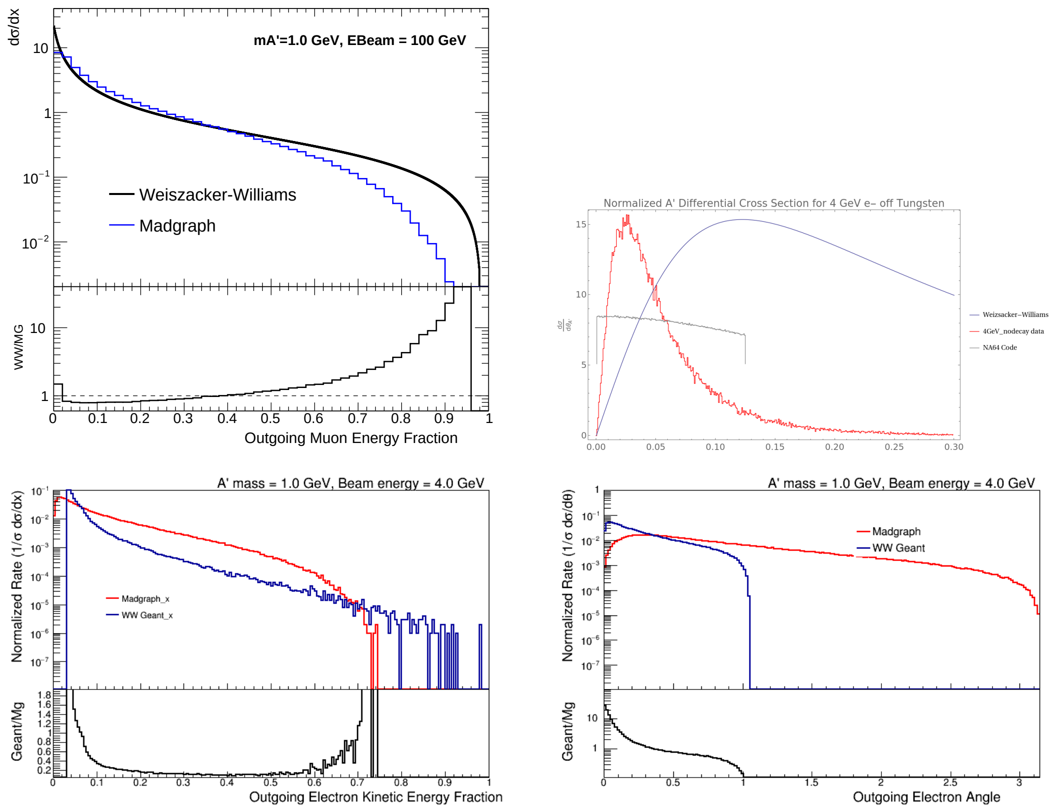

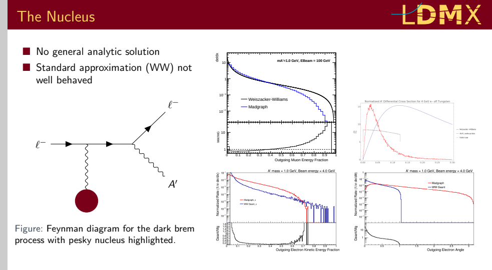

The pesky nucleus is our main source of frustration here.
First and foremost, the presence of the nucleus makes this process a $2 \to 3$ process which does not have a **separable, integrable, and invertible** solution for the outgoing kinematic distributions.
However, it also acts as a spoiler for other potential solutions.

- Standard approximation (WW) ▸ ArXiV 9310350 is not well behaved in the differential case (and is not separable!)
- Another approximation (DMG4) does not attempt to simulate the recoiling electron's angle
- MG/ME would be difficult to run "in-situ" due to different languages and simulation styles
- **Need** variable incident lepton energy so that other processes can happen to this lepton prior to its dark brem (e.g. muon production via Z decay in CMS or loss of energy via standard interaction with detector material)

Caveats on Plots

- From several years ago, specific implementation of WW lost to the sands of time
- "NA64" in upper right is referring to the DMG4 software package which has improved greatly over the years

Method

The nucleus has spoiled the easiest solution – fully implementing an analytic solution or approximation directly into Geant4 – and so we turn our attention to a more complicated solution that can still produce results faithful to MadGraph/MadEvent.

Key Features

In order to implement a custom physics process in Geant4, we need to define two things

1. The cross section as a function of energy for the process so Geant4 can decide when it should occur.
2. What happens when a particle undergoes this process i.e. how does the particle change and what other particles are created? in essence: what are the outgoing kinematics?

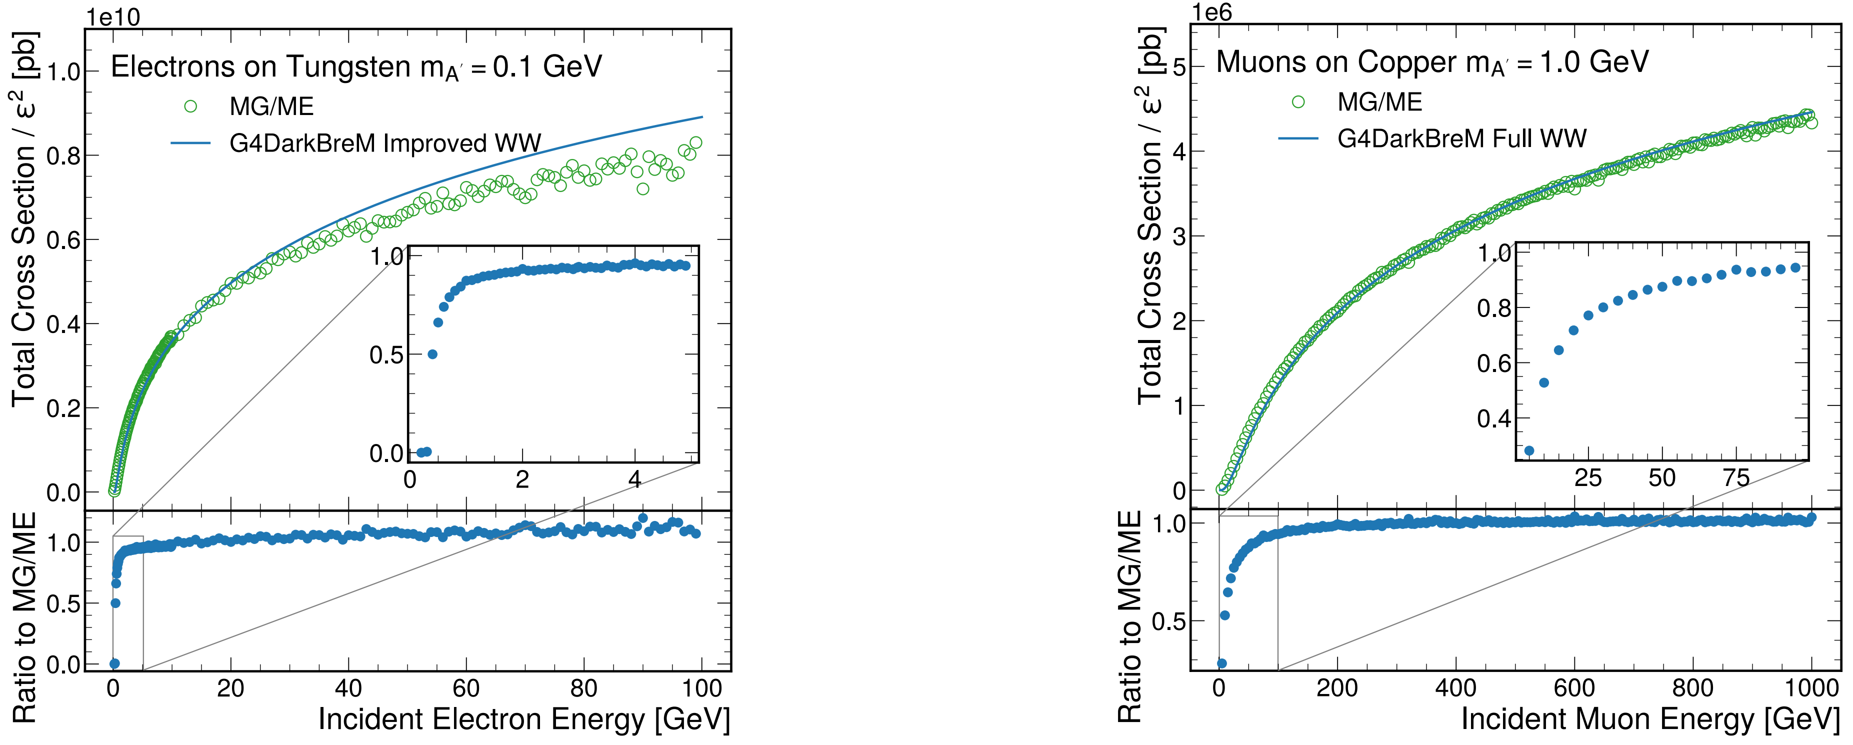

Figure: The WW approximation ends up pretty close to the MG/ME estimation.

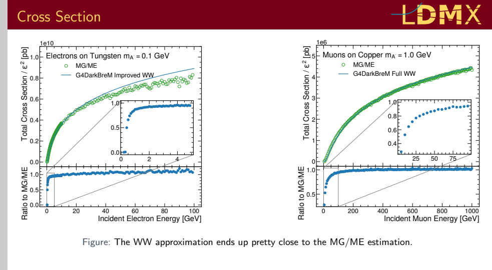

The total cross section as a function of energy can be faithfully modeled (in shape) by the WW.

Ignore

- Dark Photon – ignore accuracy since it is not directly visible to our detectors
- Nucleus – ignore completely assuming it will stay under energy required to initiate more complicated nuclear interactions

Core Idea From an incident energy, generate a sample of the recoiling lepton's kinematics.

Are there distributions to sample that are constants across incident energy?

No – but there are a few that vary slowly

Outgoing Kinematics
Ignore
Dark Photon – ignore accuracy since it is not directly visible to our detectors
Nucleus – ignore completely assuming it will stay under energy required to initiate more complicated nuclear interactions
Core Idea From an incident energy, generate a sample of the recoiling lepton's kinematics.
Are there distributions to sample that are constants across incident energy?
No – but there are a few that vary slowly

With the total cross section being calculated with WW, we can turn our attention to the kinematics of the outgoing particles. This method chooses to ignore some aspects of the full kinematics.

1. Ignore the accuracy of the dark photon kinematics since it is not directly visible to our detectors (although, it ends up being fairly good).
2. Completely ignore the nucleus since we assume that the energy given to it by this process stays below any amount required to initiate more complicated nuclear interactions (e.g. liberating a nucleon).
   - One *could* have the energy given to the nucleus become "deposited energy" in the Geant4 volume that the dark brem occurs within; however, that is not currently implemented since, again, the energy is too small to distinguish itself from regular processes.

How can we get the recoiling lepton's kinematics from a nearly-arbitrary incident energy?
We could sample from a distribution that does not change with incident energy...
Big Credit to Michael Revering for actually discovering these distributions through trial-and-error.

Energy Fraction x

Fraction of incident lepton energy left with recoiling lepton

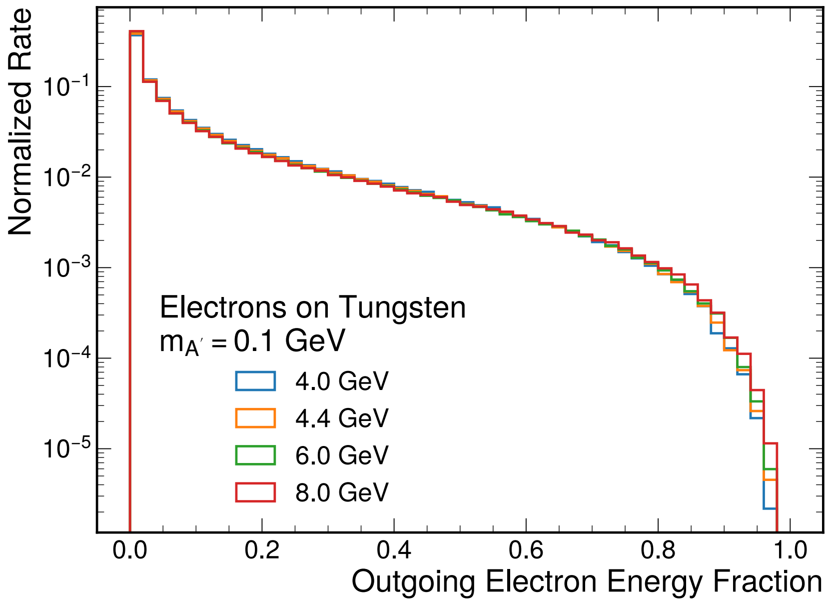

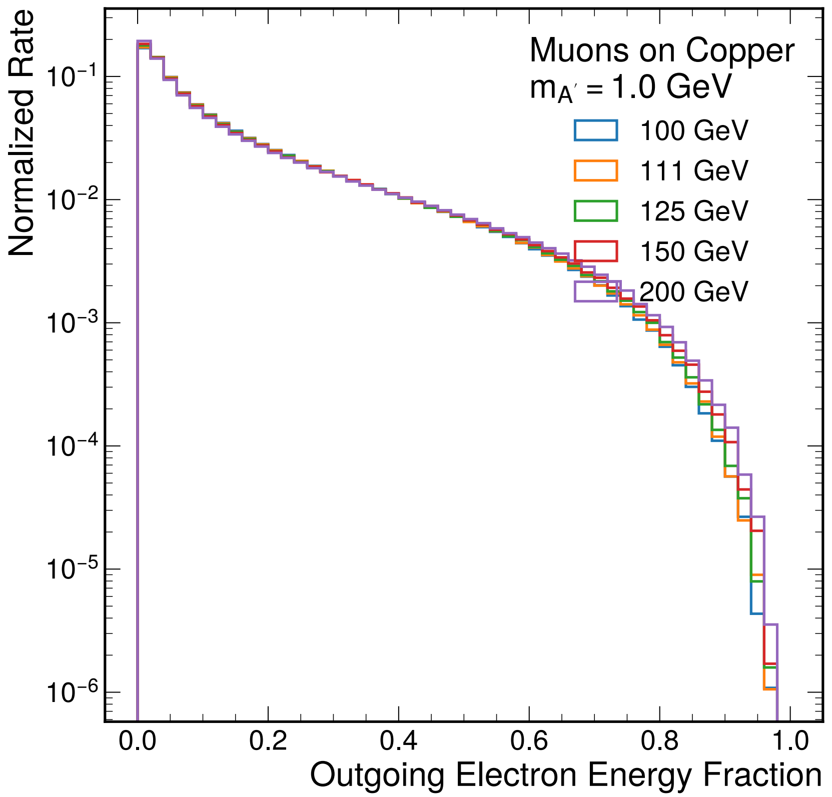

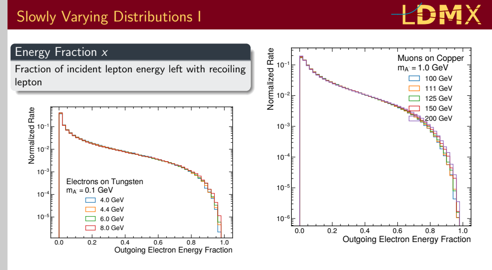

Apologies for the difference in aspect ratio.
The first distribution helpful to us is the energy fraction

$$x = \frac{E_{\text{recoil}}}{E_{\text{incident}}}$$

The shape of this distribution changes extremely slowly even after crossing a factor of two in incident energy.

Example

If we have electrons on tungsten, we could sample x from a distribution created with an 8 GeV beam for any electron energy down to $\sim$ 4 GeV and still maintain good agreement with the true distribution.

Slowly Varying Distributions II

Transverse Momentum p_T

Recoiling lepton's momentum transverse to the incident direction.

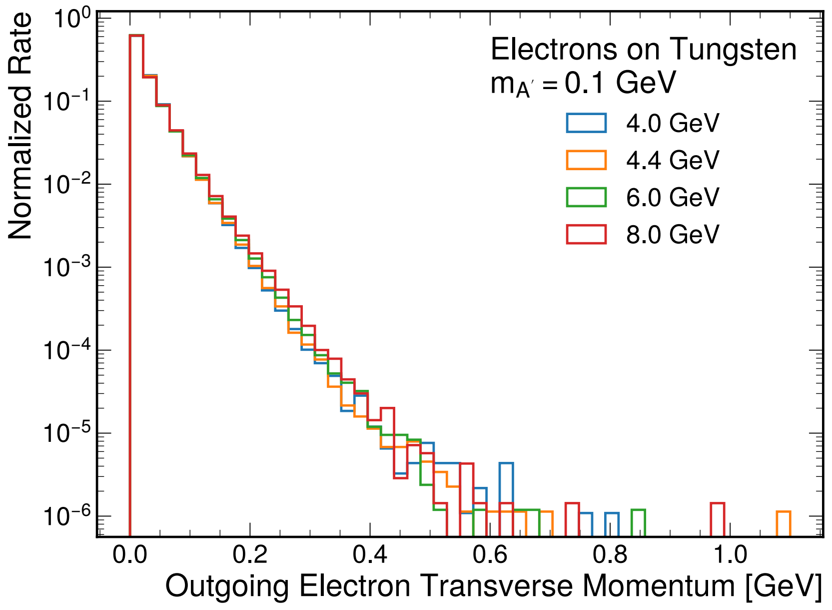

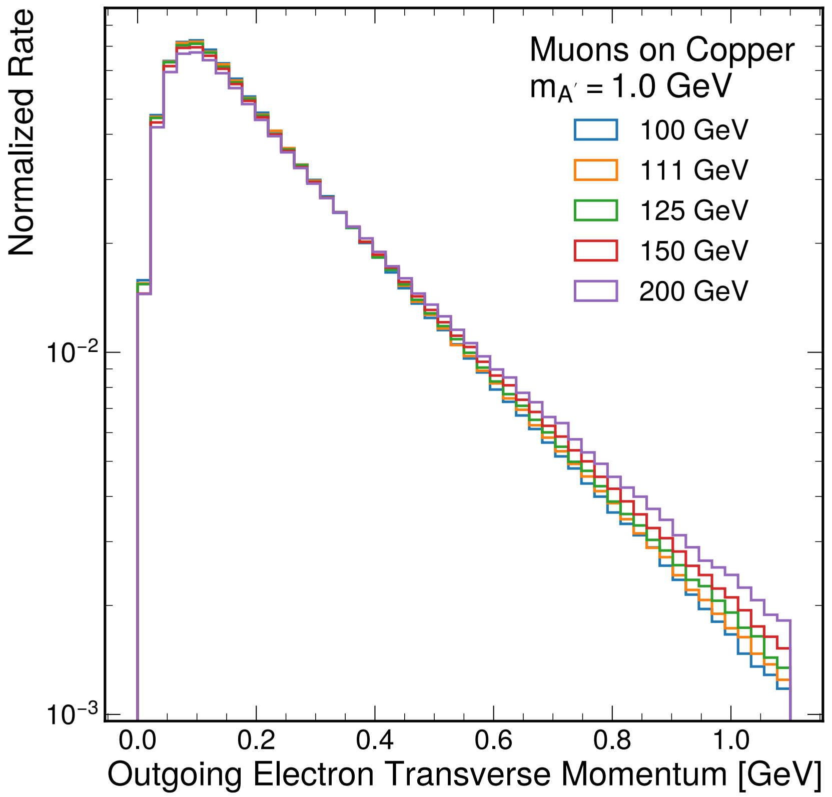

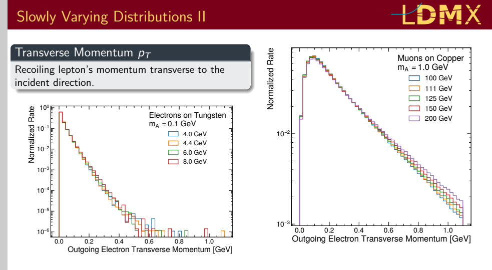

The second distribution helpful to us is the transverse momentum *relative to the incident direction* (important to emphasize in the Geant4 scenario where this incident direction will probably not be aligned with any axes).

While this distribution changes a bit quicker than the x distribution, it can still be usable for incident energies closer to each other (for example within $\sim 10\%$ in the muon case).

Note

The accuracy of the angular distribution (here defined via p_T) is the main limiting factor of this method. As can be seen on the prior slide, the x distribution changes much slower than the p_T distribution (and any other "angular" variables we looked at).

Procedure

MadGraph/MadEvent already do this sampling of these kinematic distributions for us when generating events, so we use those events as a "reference library" from within the Geant4 simulation.

1. (Before Geant4) Generate a library of events at a variety of beam energies
2. When given a lepton by Geant4, select an event from the sample with the closest beam energy above the current energy of the lepton.
3. **Scaling** Use the selected event to define the recoiling lepton's kinematics relative to the incident lepton direction.
4. Rotate the recoiling lepton momentum out of the incident lepton frame.
5. Calculate the dark photon's kinematics by conserving 3-momentum between the incident and recoiling lepton.
   - Can't conserve 4-momentum in all events because we are ignoring the nuclear recoil, choosing to put all this error into the dark photon's kinematics.

Instead of using a the sample of events generated by MG/ME to create a distribution from which to sample, we just use these events directly as a representative sample of these distributions. This "reference library" has only been tested via MG/ME; however, it could be created using a different event generator if the user desires.

Next Up

Step three is a bit vague because that is where there is some ambiguity. I call this the "scaling" and we can do this procedure in a few ways.

Scaling Procedure

1. Keep selecting a new event until $p_T^2 + m_\ell^2 < E^2$ where E is the actual incident lepton's energy.
   - If the reference library has a beam energy "close enough" ($\sim 10\%$) to the incident energy, only one event is needed a vast majority of the time. A warning is printed if the reference library is failing to satisfy this criteria and repeatedly selecting new events.
2. Set the outgoing lepton's energy to xE where x is the energy fraction from the reference library.
3. Set the outgoing lepton's p_T to the p_T taken from the reference library.
4. Calculate the outgoing lepton's p_z using its mass and **choose** $p_z > 0$.

Scaling Procedure

1. Keep selecting a new event until $p_T^2 + m_\ell^2 < E^2$ where E is the actual incident lepton's energy.
   - If the reference library has a beam energy "close enough" ($\sim 10\%$) to the incident energy, only one event is needed a vast majority of the time. A warning is printed if the reference library is failing to satisfy this criteria and repeatedly selecting new events.
2. Set the outgoing lepton's energy to xE where x is the energy fraction from the reference library.
3. Set the outgoing lepton's p_T to the p_T taken from the reference library.
4. Calculate the outgoing lepton's p_z using its mass and choose $p_z > 0$.

This method is called "ForwardOnly" due to its choice of keeping $p_z > 0$ in all events. This choice was made because

1. Difficult to find a longitudinal variable whose distribution changed slowly relative to incident energy
2. The back-scattering ($p_z < 0$) rate is several orders of magnitude below the forward scattering and (in our cases) the back-scattered events are likely to be removed by downstream analysis cuts anyways.

Other Options

The G4DarkBreM package on GitHub actually implements two other scaling methods.

1. "CMScaling": boost into an approximate Center-of-Momentum frame using the incident lepton kinematics from the CoM frame defined by the generated event – seems to distort the longitudinal momentum distribution more in the higher-rate $p_z > 0$ region compared to the Forward Only. Distortion might be worth it if back-scattering events are of importance to you, would require more study.
2. "Undefined": don't do anything to the sampled event – obviously unhelpful but is nice for testing any developments

Validation

Validation

Note

This procedure is very artificial and is not physically motivated outside of the distribution comparisons shown earlier, so it is helpful to look more directly at the effects of this method compared to other methods.

Energy Fraction

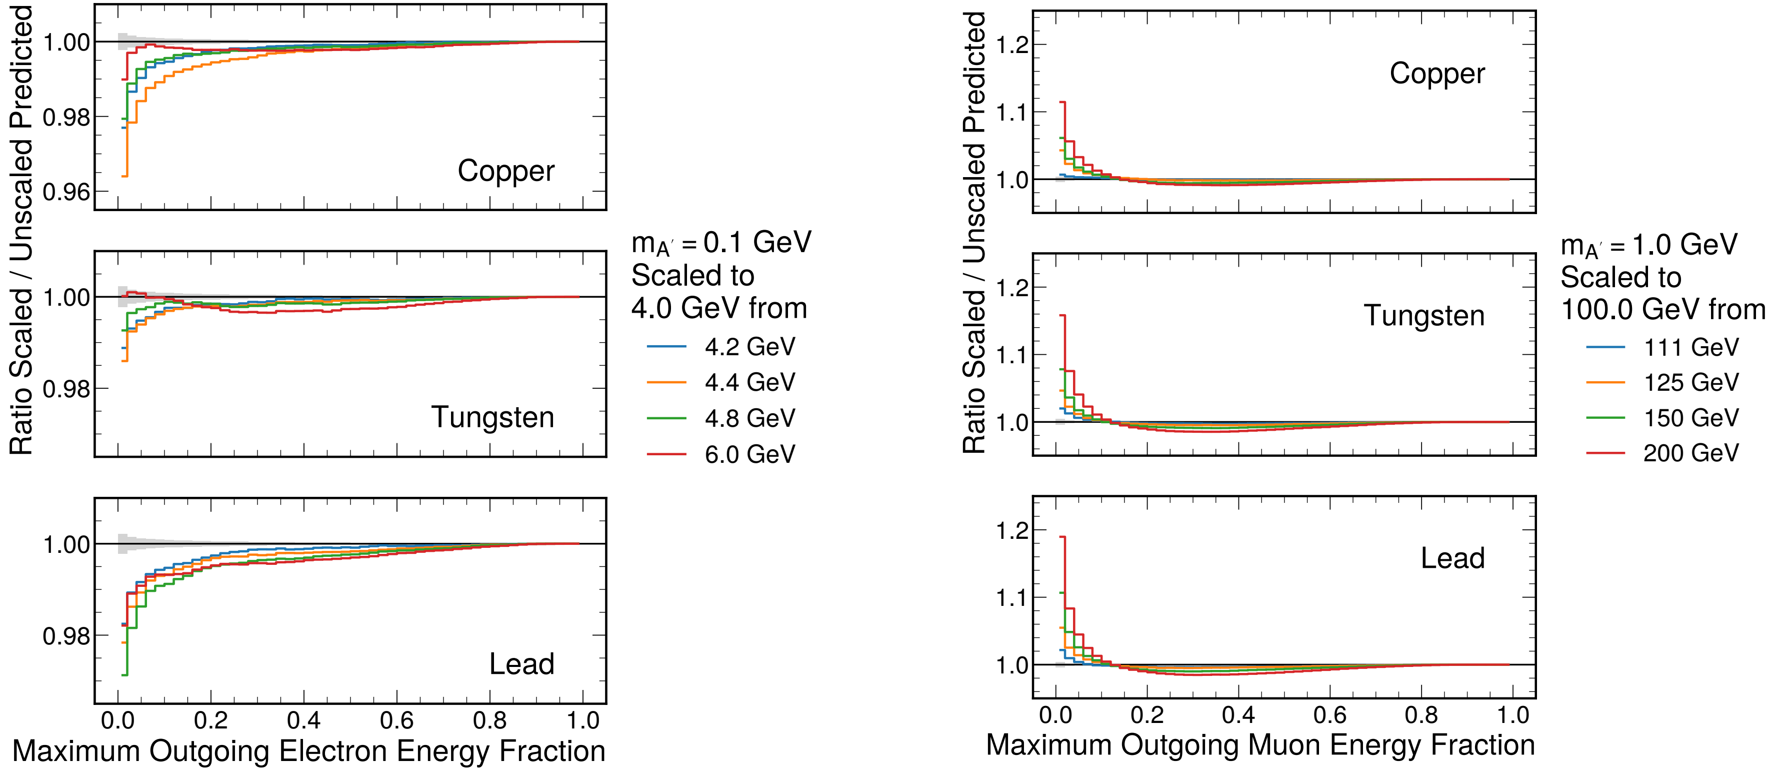

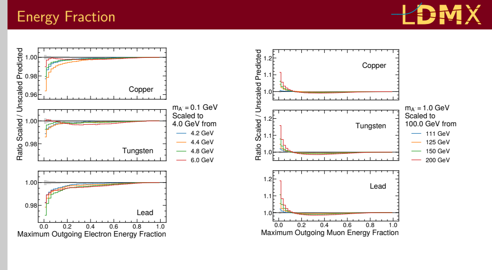

Stay with Me

The plots on this slide are complicated and hold a lot of information, so please stay with me.

Definition

We are showing the *ratio* of the *cumulative* distributions comparing the results of the events after scaling ("Scaled") to the events generated at the actual energy ("Unscaled").

Goal

Taking the ratio after looking at the cumulative distributions allows us to *roughly* read this plot as an estimate on the error of this method.

Staying within $\sim 10\%$ of the incident energy allows the scaled distribution to stay within $\sim 5\%$ of the true distribution even in the bin requiring the least amount of outgoing lepton energy.

Transverse Momentum

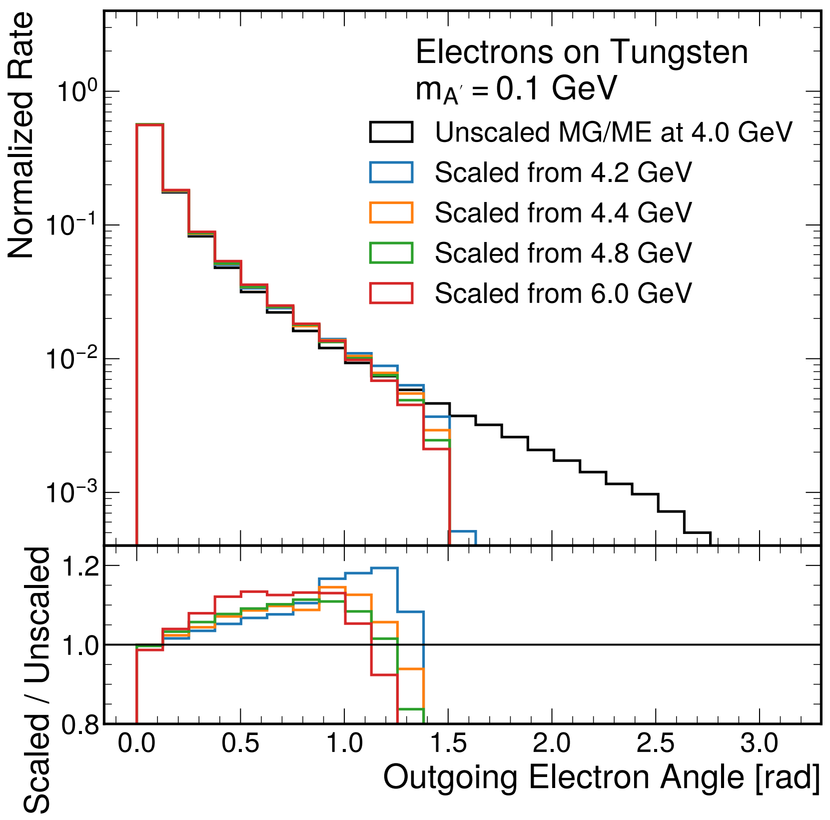

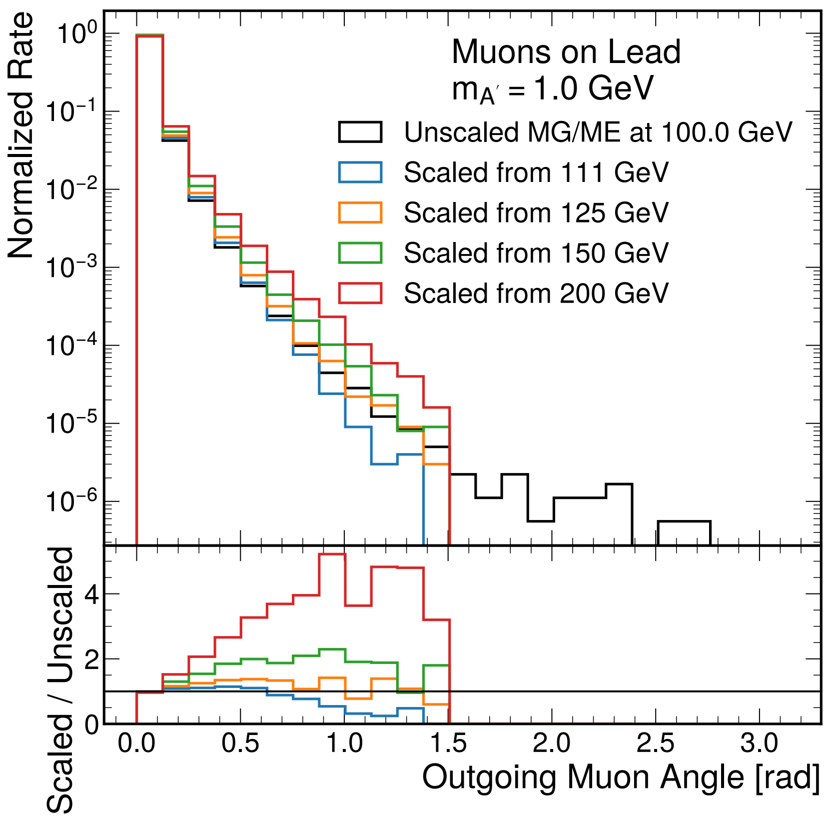

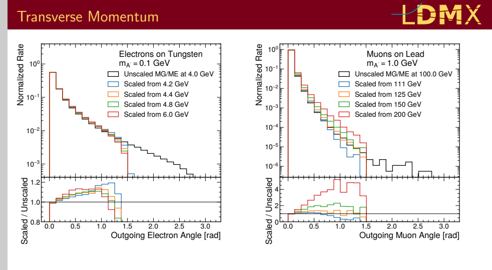

A bit simpler of a comparison – looking at the rates and comparing them with the true/“unscaled” distributions.

Summary

Stays with $\sim$ 20% of the rate in the forward region when scaling from 10% above the actual energy.

Only showing one material per lepton here, but the other materials show the same trend.

Comparison and Capabilities

Now that we've validated this scaling technique, let's look at a more full test of its abilities.

Basic Set Up

- The input reference libraries have energies sampled in 10% increments starting from the beam energy and going down to twice the dark photon mass.
- The Geant4 simulations are just a block of material within a world volume of air.
- Looking at sim-level estimations of detectable quantities; however, these are not the only quantities one could look at.
  1. Recoiling lepton's angle relative to the beam direction
  2. Visible energy (i.e. all energy in particles that are *not* the dark photon)

There is a package that implements a variety of dark matter models and their interactions within the context of Geant4: Dark Matter for Geant4 (DMG4) which we can also try to use.
Comparisons made with DMG4 v1.2 (March 18, 2022), current version is v2.5 (Jan 3, 2023).

Expect to follow MG/ME distribution

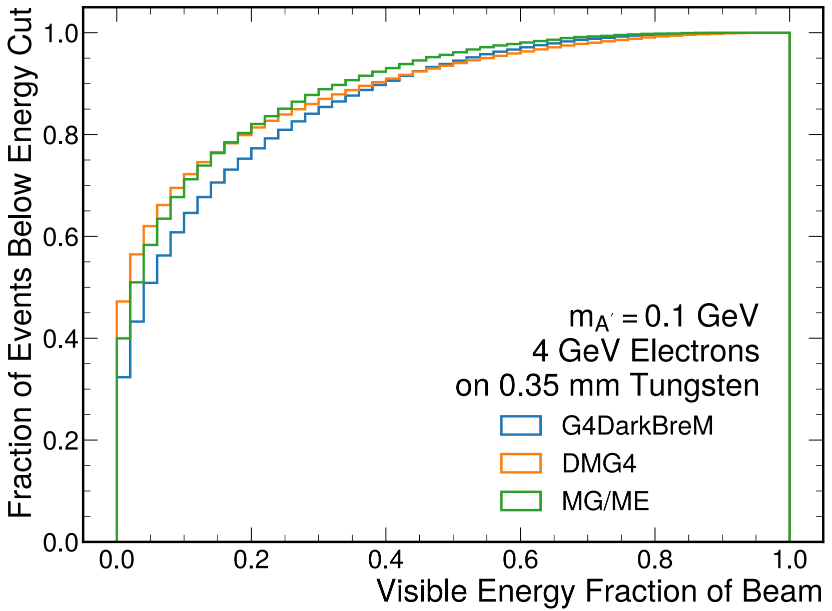

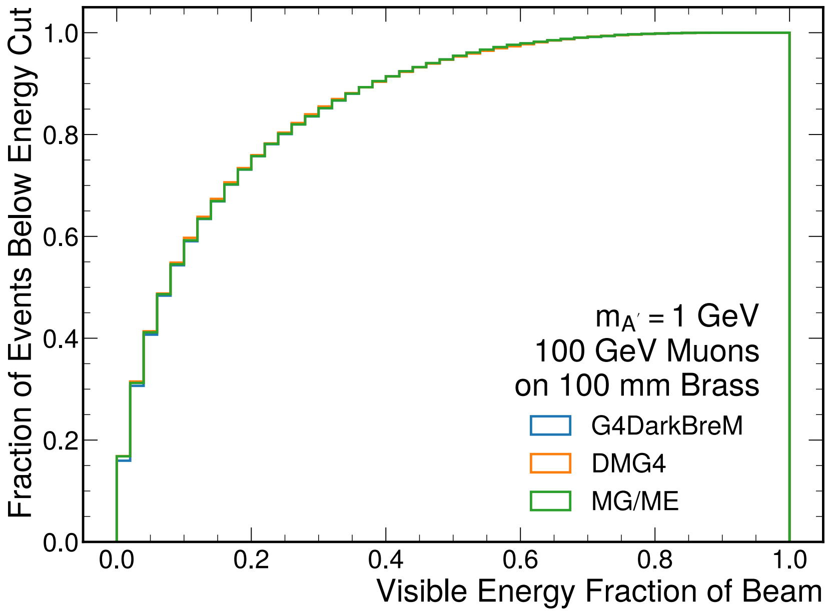

In a thinner target, we expect the incident lepton to not lose much (if any) energy before undergoing the dark brem process. In this case, both G4DarkBreM and DMG4 follow the mono-energetic MG/ME distributions well.

Expect to follow MG/ME distribution

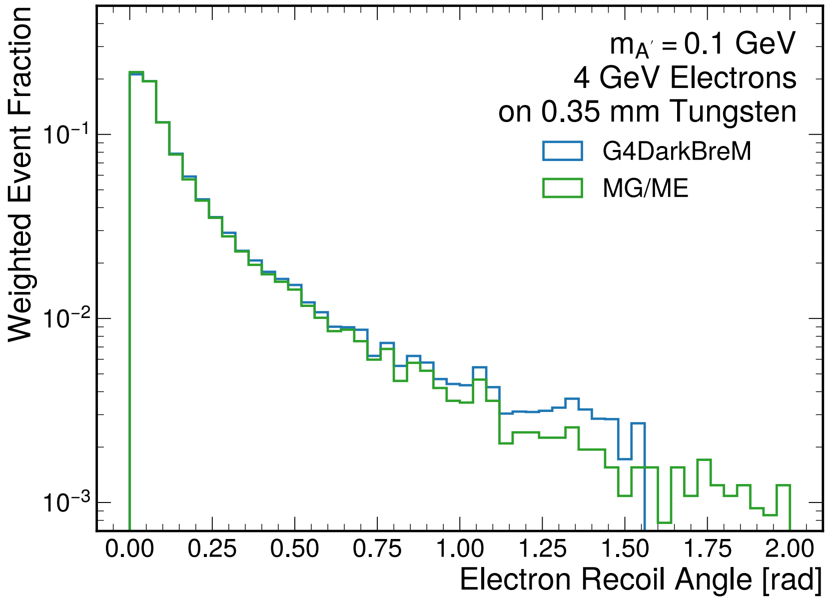

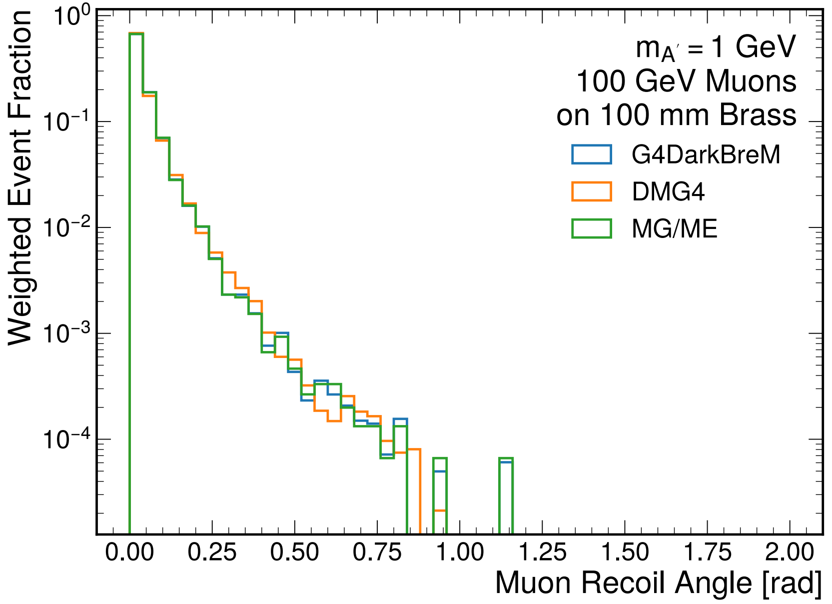

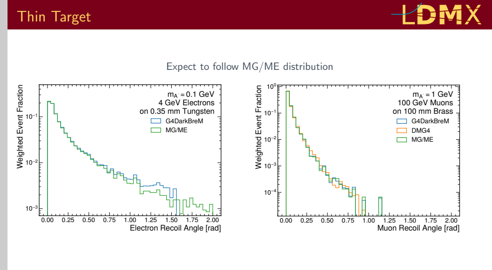

The recoiling angle is also well modeled for both libraries in the muon case. In the electron case, DMG4 *makes no attempt* to model the outgoing electron's angle, instead setting its angle (relative to the incident electron) to zero in all events, so its line is omitted from the plot.

DMG4 Recoiling Electron Angle

This statement was true as of DMG4 v1.2. The code has been drastically refactored and so I wasn't able to easily check if the current version (v2.5) also makes this choice or not. No mention of this is made in the HISTORY file packaged with the code.
Obtain your own copy of DMG4 from
http://mkirsano.web.cern.ch/mkirsano/DMG4.tar.gz

Expect to deviate from monoenergetic MG/ME distribution

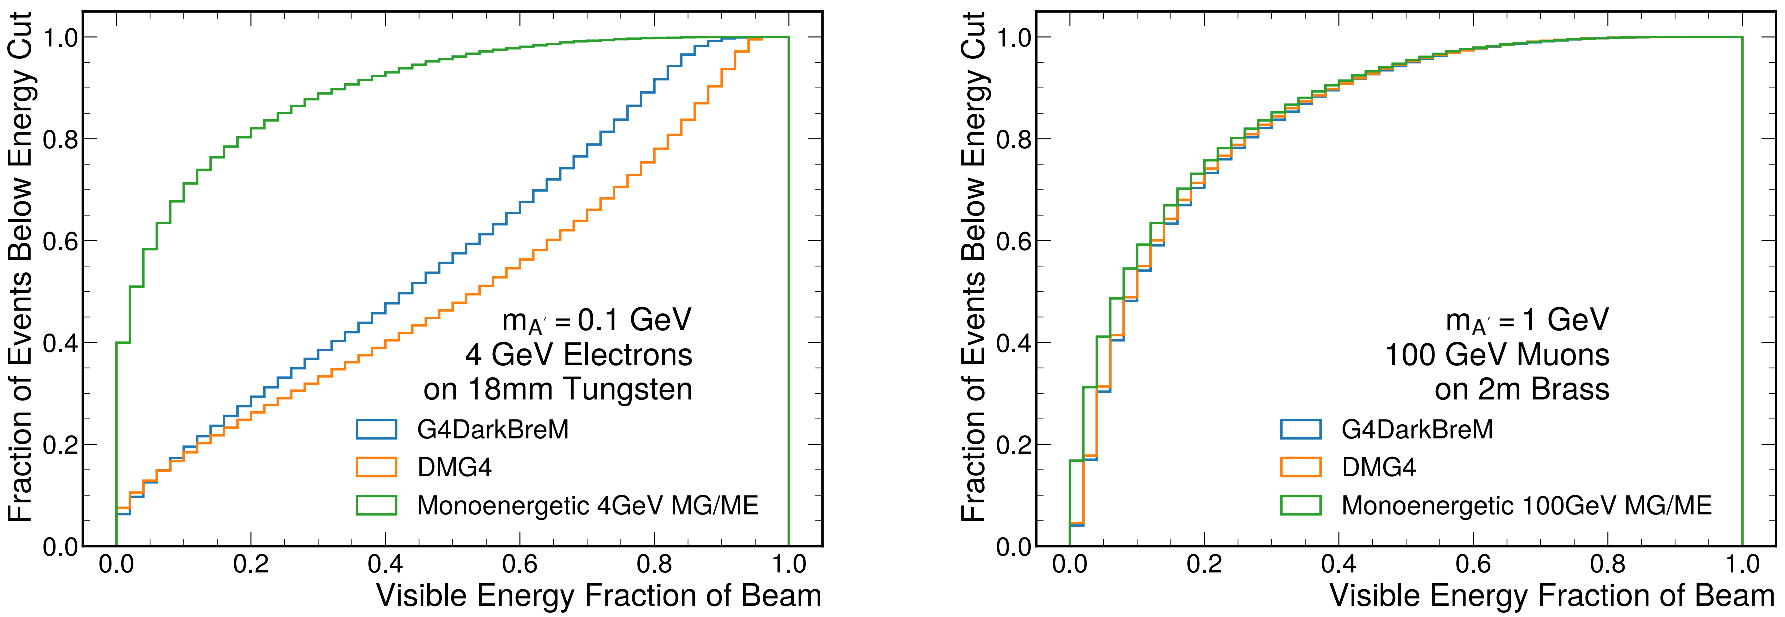

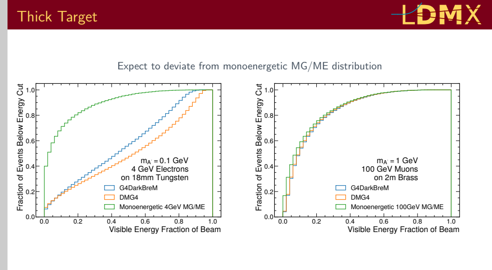

With how infrequent the muon interacts, its distributions do not change much. (They only lose $\sim 0.1\%$ of their kinetic energy in 2m of brass.)
The electrons, however, do start to undergo a shower in the $\sim 5X_0$ target and we can observe a strong separation between the monoenergetic MG/ME and the variable-incident-energy libraries.

Thick Target

Expect to deviate from monoenergetic MG/ME distribution

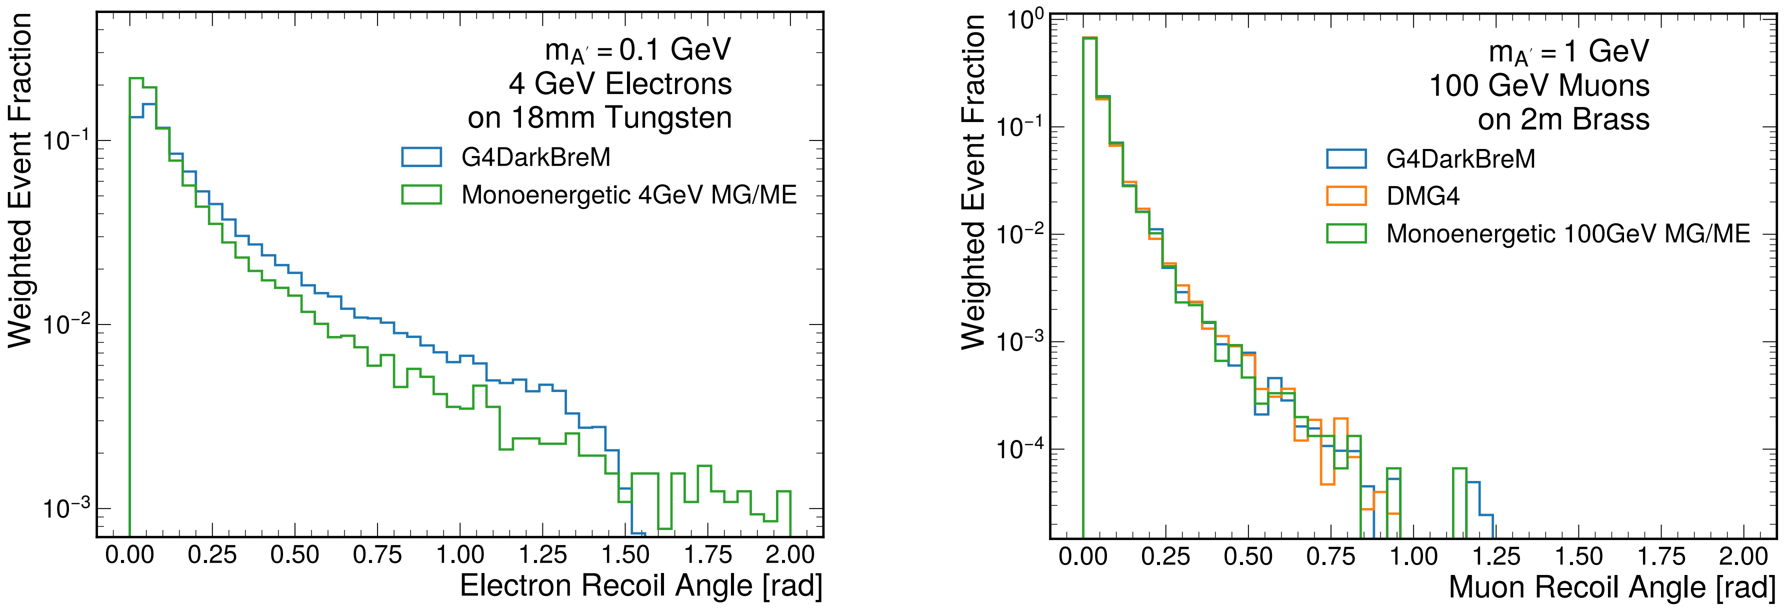

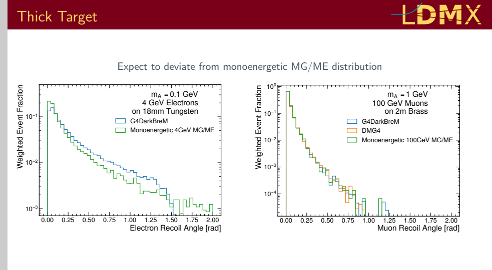

Again, the muon case is slightly more boring, but the electron case shows the separation between the beam direction and the incident direction, broadening the recoiling angle distribution.

Concluding Remarks

Concluding Remarks

Concluding Remarks

Novel technique for modeling the dark brem process within GEANT4
Extends an event generator to near-arbitrary incident energies

Why

- Support a model of the recoiling electron's angle
- Potentially test different dark brem event generators within the same GEANT4 simulation set up
- Potentially useful for other processes similar in topology to dark brem

Concluding Remarks
Novel technique for modeling the dark brem process within GEANT4
Extends an event generator to near-arbitrary incident energies
Why
Support a model of the recoiling electron's angle
Potentially test different dark brem event generators within the same GEANT4 simulation set up
Potentially useful for other processes similar in topology to dark brem

Future Work

G4DarkBreM
Investigate other scaling methods (e.g. using the nucleus reference frame with Lorentz boosts) for improved performance.

DMG4
Support for recoiling electron angle in dark brem process may be validated via this process.

Pythia
Two potential paths both of which could be compared to this method for understanding.

1. Pythia running of MG/ME-generated code.
2. Pythia-native model of dark brem.

Much of this work was motivated by the experimental need to model dark brem from within the GEANT4 simulation instead of simply using the event generated by MG/ME as the primary particles to begin the simulation. Future work would involve making this procedure either more accurate relative to a full matrix-element calculation and/or easier to use (i.e. removing the need for the generation of a reference library).

While I (and many of the other folks who worked on G4DarkBreM) may not take responsibility for this work, I do find it helpful for future researchers to know our thoughts on potential paths towards improvement.

Questions

Questions

Testing and validation of this method was done with a MadEvent4 workspace available on GitHub
▸ LDMX-Software/dark-brem-lib-gen

Properties

- Add a dark photon (S=1, Q=0, mass configurable)
- Add nucleus (S=1/2, Q=0, mass configurable)
- Custom coupling between nucleus and both photons that include inelastic and elastic form factors (A, Z are then parameters)
- Modify source code to give electron mass (prevents divergence of total cross section, not needed if replicating with MG5)
- Lepton beam incident on *stationary* nucleus (energy set equal to its mass in GeV).

Dark Brem in MadGraph/MadEvent

Testing and validation of this method was done with a MadEvent4 workspace available on GitHub

Properties
- Add a dark photon (S=1, Q=0, mass configurable)
- Add nucleus (S=1/2, Q=0, mass configurable)
- Custom coupling between nucleus and both photons that include inelastic and elastic form factors (A, Z are then parameters)
- Modify source code to give electron mass (prevents divergence of total cross section, not needed if replicating with MG5)
- Lepton beam incident on stationary nucleus (energy set equal to its mass in GeV).

WW in Geant4

Michael Revering did the “pure WW” simulation within Geant4 by sampling its 2D differential cross section directly. Since the WW distribution is not separable and it follows orders-of-magnitude changes, he also developed a “capping” function (right) that can help make this simulation computationally feasible.

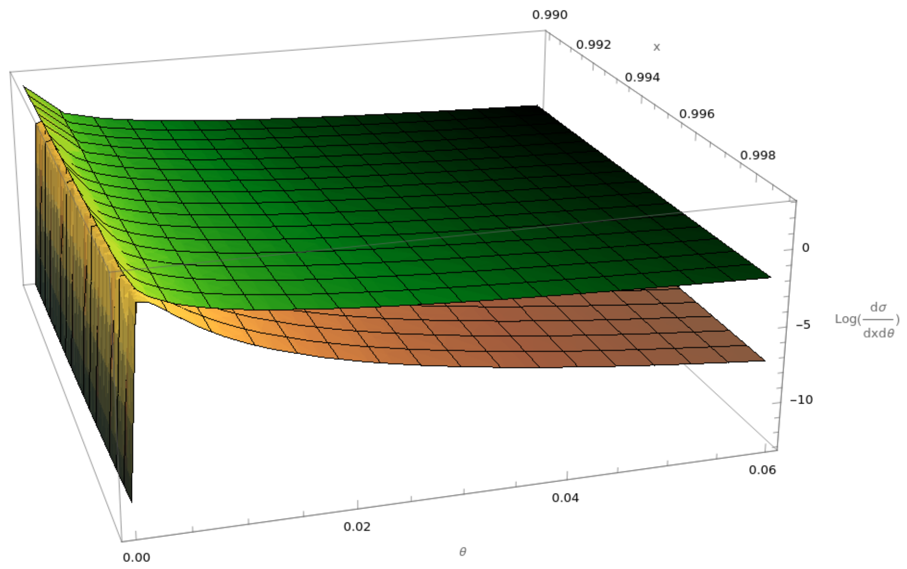


```
TwoDNCap[mA_, x_, θ_, EBeam_] := 1 / ( EBeam^2/16 mA^2 * ( mA/(mA+mmu) (1 - x) + (mmu/(mA+mmu))^2 * x ) * (100 θ + 0.05) )
```

Figure: The “capping” function (green) plotted against the WW distribution (tan) along with its functional form (bottom).

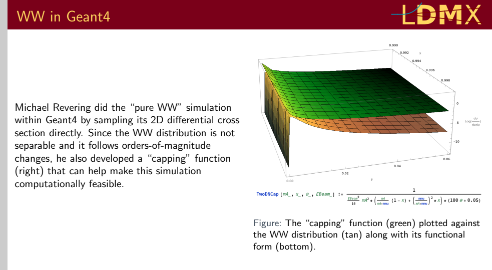

Michael’s Description of this Procedure

As the two variables for the WW distributions aren’t separable, we couldn’t directly produce points following their distribution and instead needed to create a bounding function which was separable, integrable, and invertable, such that we could generate points uniformly in a 3D space (muon energy, muon angle, and probability density) which fully contained the WW distributions. By then accepting points which fall below the probability density of the WW functions, you produce a samples which follow their distributions. This part actually worked fine, although if your bounding function doesn’t match WW well it can be computationally slow, as you have to reject a lot of points which end up in the space above WW. Devising the right function was somewhat tricky, as it peaks very sharply at low energies and small angles, and you need to generate a bounding function that can scale effectively with the muon mass and incoming energy. I think I have some old 2D distribution plots somewhere if you want to include this part as a general difficulty of the function-based method, though I think a clever-enough function selecting process would make this step not a real problem.

Biasing and Filtering

With an embedded process, we can use GEANT4-native biasing and filtering to make sure the dark brem occurs within volumes of interest. A biasing factor used within LDMX to generate dark brem off electrons was found empirically to be

$$B = \frac{1}{\epsilon^2} m_A^{\max(2, \log_{10}(m_A))} \quad \text{where} \quad m_A \text{ is in MeV}$$

This biasing factor allowed for the events to follow shower development and produce well-behaved event weights. ϵ should be kept well below one to avoid the process from happening in unbiased volumes (LDMX uses a value of $\epsilon = 0.01$).
Filtering is also necessary to avoid events without a dark brem or multiple dark brems being put into the output sample. Doing this filtering in a computationally performant way is a complicated issue on its own and is outside the scope of this talk; however, one can achieve the functionality necessary by simply defining a UserTrackingAction whose PostUserTrackingAction looks for dark photons being produced and aborting the event if your requisite criteria are not met.

"follow shower development"

The shape of the distribution of where the dark brem occurs follows the shape of where all shower interactions occur. In LDMX, we impose this by looking at the z-location of the dark brem compared to the z-location of all calorimeter hits.

"well-behaved event weights"

The event weights should follow (roughly) an exponential distribution peaking at $1/B$.
Over-biasing can lead to weights that are all exactly $1/B$ signalling that no other (unbiased) process is being allowed to occur.