

Table of Contents

Q&A for VBS WW OS leptonic.....	1
Comments on preapproval.....	1
To be added in ANv7.....	1
Slides.....	5
Comments about the paper draft.....	7
v1 25 February 2021.....	7
Comments about the AN.....	8
v6 03 March 2021.....	8
v5 26 January 2021.....	12
v4 13 January 2021.....	13
v3 04 January 2021.....	14
v2 23 November 2020.....	19
v1.....	21
General questions and discussion.....	21
Minutes from 15-09-2020 SMP-VV.....	21

Q&A for VBS WW OS leptonic

Color code for answers

- - Comment is acknowledged and answered
- - Authors are working on answering the comment
- - Comment requires further work to be addressed or need attention from the internal reviewer regarding a specific issue
- - We do not agree with the comment and arguments are given

Comments on preapproval

To be added in ANv7

- Converge and clarify the signal definition choice (for both analyses)
 - ◆ Add one or two bins in the cutbased analysis with events with $m_{jj}:[300-500]$ GeV and $det_{jj}<3.5$

● We have added three bins in the cut-based analysis, defined as follows:

1) $300 \text{ GeV} < m_{jj} < 500 \ \&\& \ 2.5 < det_{jj} < 3.5$

2) $300 \text{ GeV} < m_{jj} < 500 \ \&\& \ det_{jj} > 3.5$

3) $m_{jj} > 500 \ \&\& \ 2.5 < det_{jj} < 3.5$

Such bins have been included in each Zll region. In this way the two analyses share the same phase space definition, hence we can now derive an apple-to-apple comparison. Evaluating the expected significance in the different flavour channel (for each dataset) we get the following results:

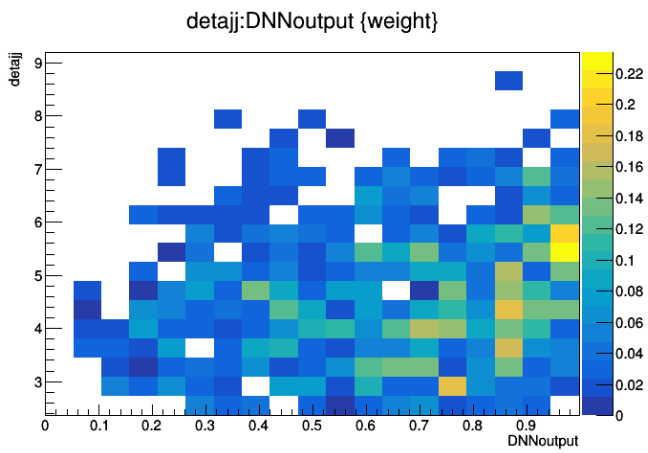


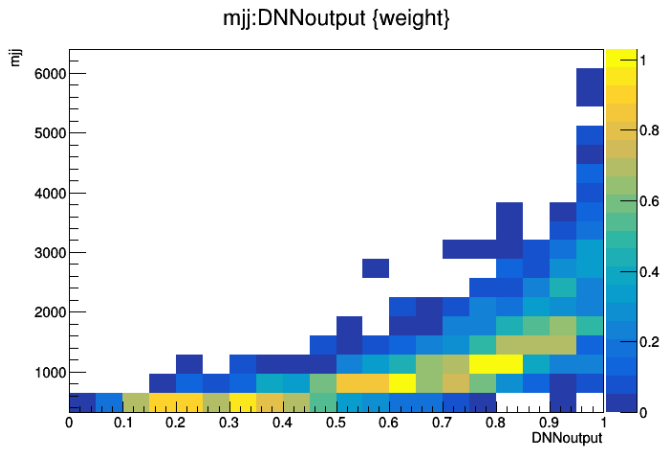
dataset	mjj shape-based analysis	DNN analysis
2016	1.83 sigma	1.89 sigma
2017	1.95 sigma	1.92 sigma
2018	2.82 sigma	2.88 sigma
full Run2	3.79 sigma	3.75 sigma

The mjj shape-based analysis clearly benefits from loosening the VBS-like phase space definition, and so we will use this selection.

- ◆ Check the DNN sensitivity by cutting on $m_{jj}>500\text{GeV}$ and $det_{jj}>3.5$

● We raise the thresholds for mjj from 300 to 500 GeV and for det_{jj} from 2.5 to 3.5. We use 2016 to estimate how this affects the DNN performance. We find that the significance decreases by about 2% with tighter cuts, passing from 2.12 to 2.07 (all leptonic channels included). The plots below represent DNN:mjj (on the left) and DNN:det_{jj} (on the right) for the signal in the Zll < 1 (top row) and Zll > 1 (bottom row).

QandAforVBSOSWW < Main < TWiki




QandAforVBSOSWW < Main < TWiki

As you can see, it is not 100% true that an event with low m_{jj} and low det_{jj} ends in the low score region of the DNN output. This explains why the significance with a tighter cut on m_{jj} and det_{jj} decreases.

- Check the possible anticorrelation between QCD WW and EW WW in the fit

🟢 VBS signal and WW QCD normalisations are 30% anti-correlated, see the correlation matrix below where only scaling parameters are displayed:



We did investigate the reliability of the fit procedure through the use of toys, see Kenneth's question below in "Slides" section.

- Further validation of the datacards
 - ◆ merging ee and mumu categories if it doesn't help

🟢 Merging ee/mumu categories slightly reduces the expected statistical significance (e.g. 2018 data set: 1.78 sigma \rightarrow 1.52 sigma, $m_{jj} > 300$ && $det_{jj} > 2.5$). Although a 15% gain on the expected significance in the SF category does not mean that the combined fit improves by the same amount, since the analysis is mainly driven by the DF category, it could be worth keeping the ee/mumu splitting in the analysis.


- ◆ merging production processes using a scheme like the ones in the figures

🟢 Higgs contribution is not relevant at all in SF categories, due to the very tight m_{ll} cut (>120 GeV), so we decided to neglect such samples. Moreover, merging all different Higgs contributions is not very feasible, since each production mode is affected by different theoretical systematics which are treated separately.

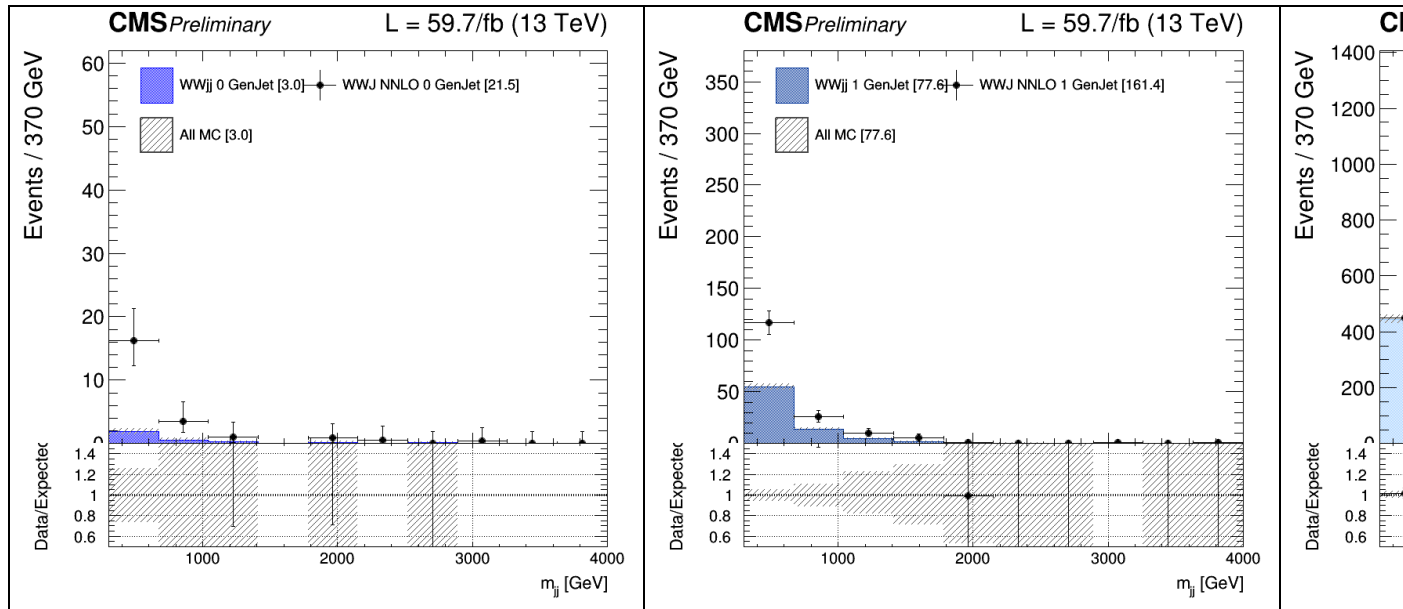
- WW+QCD MC samples: WWJJ vs. WW inclusive.
 - ◆ Show m_{jj} distribution starting at $m_{jj} > 300$ GeV, combine all lepton flavors, Zll regions, and all years 🟡

- ◆ Check the fraction of events of 0, 1, and 2 parton jets at GEN level from existing WW inclusive sample

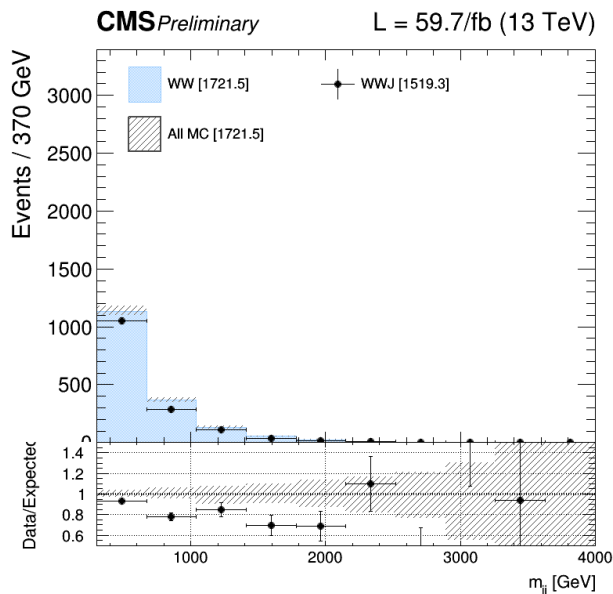
● We compared the MadGraph LO WWJJ sample we are currently employing in the analysis with an inclusive WW NNLO sample generated with powheg (WWJ) [1]. This is a fair comparison, since the QCD precision of the second jet is at LO in both samples. At reco level we observe an overall good agreement in m_{jj} , although the two samples significantly differ in the very first bin. Events have been selected with $m_{jj} > 300$ GeV and $\text{detajj} > 2.5$.



Indeed the fraction of 0/1 gen-jets entering the signal region is much higher for the WWJ sample at low m_{jj} values, whereas the prediction for events with at least 2 gen-jets with $pt > 30$ GeV is basically the same.



We also drew a comparison applying the analysis preselection defined with gen-level variables: no relevant discrepancies are found in the shape of (gen-) m_{jj} .



We could hence replace the WWjj MadGraph sample we are currently employing with the WWJ Powheg one.

[1]

https://cmsweb.cern.ch/das/request?input=dataset%3D%2FWWJto2L2Nu_NNLOPS_TuneCP5_13TeV-powheg-pyth

- Make use of EW LLJJ MC samples instead of EW ZJJ (that overlap with dibosons)
- We are processing the EW LLJJ sample, in the meantime we are cutting on $m_{jj} > 120$ GeV at LHE level, removing the overlap with the semi-leptonic sample.
- Investigate further the surprising agreement for the third jet distribution, using different PS configurations. Share the configuration setup.
- Configuration setup has been shared, see <https://hypernews.cern.ch/HyperNews/CMS/get/SMP-21-001/17.html>

Slides

Q Guillermo s16 is this selection or signal definition ? • signal definition

Q then it is inconsistent between cutbased and DNN ? • not settled yet on the signal definition. Indeed need to compare with same m_{jj} cut

Q Guillermo Wouldn't it make sense to have a more VBS-like definition? How much do you gain by relaxing the cuts? •

Q Aram: When you say you get better performance, do you mean the ROC curve or going all the way to the expected result? • We compare first the ROC curve, to understand qualitatively which models have better performance, and to make a first selection of all the models tested. Then we extract the expected results to have a quantitative measure of the gain of the DNN wrt to m_{jj} .

Q Guillermo Just to make sure you have a gain, you could add a bin with all events with m_{jj} in [300, 500] or detajj in [2.5, 3.5] to have a more fair comparison. Also, you should have a consistent cut between the two channels in order to define a consistent cross-section •

Q Paolo: Personally I don't think going down to a low value is a problem as long as you stay away from the triboson production 🟢 ok

Q Guillelmo s40, for the CRs are you using 1 bin per region? 🟢 The CR you would make would still be dominated by top

Q Guillelmo If it's free floating, it must be very anti-correlated with the signal. Do you really have the ability to separate them? 🟡

Q Kenneth You could do some tests with toys, possibly drawing the data from a biased distribution built from $a \cdot \text{QCD} + b \cdot \text{EW}$. You should see how reliably you recover the values of a and b that you put in vs

🟢 We checked the fit reliability through the use of toys (500 for each configuration), generating data with different a,b values. The fit procedure shows that input parameters are recovered regardless of initial settings.

input parameters	fitted parameters
a = 0.5 ; b = 0.5	a_fit = 0.497 +/- 0.015 ; b_fit = 0.500 +/- 0.012
a = 0.5 ; b = 1	a_fit = 0.510 +/- 0.015 ; b_fit = 0.963 +/- 0.014
a = 0.5 ; b = 2	a_fit = 0.483 +/- 0.016 ; b_fit = 1.994 +/- 0.016
a = 1 ; b = 0.5	a_fit = 0.994 +/- 0.015 ; b_fit = 0.489 +/- 0.012
a = 1 ; b = 1	a_fit = 1.013 +/- 0.016 ; b_fit = 0.970 +/- 0.014
a = 1 ; b = 2	a_fit = 0.992 +/- 0.016 ; b_fit = 1.983 +/- 0.017
a = 2 ; b = 0.5	a_fit = 1.974 +/- 0.015 ; b_fit = 0.475 +/- 0.012
a = 2 ; b = 1	a_fit = 1.985 +/- 0.015 ; b_fit = 0.977 +/- 0.014
a = 2 ; b = 2	a_fit = 1.923 +/- 0.015 ; b_fit = 1.982 +/- 0.017

Q Guillelmo Did you make a check of merging the ee and mm channels? 🟢 We checked that we really don't lose much, we can probably do this, but we need to study it a bit more

Q Guillelmo In the data cards, you should really combine the small processes rather than having them all split. There are quite a few warnings that need to be addressed 🟡

Q Paolo: We should understand the off-shell effects. 🟢 we did try to make a sample of $p p > l v l v j j$ and some tests, in the backup

Q Paolo Is the sample really LO WW+2j only? Did you make some comparison? 🟢 Yes also in backup

Q Kenneth Combine channels and years (at least 2017/8) to have more clear comparison of the gen-level differences 🟡

Q check the fraction of 0,1 partons contributing in the inclusive samples 🟢 ok

Q Paolo s8 Z+2jets EW ==> switch to EW LLJJ samples 🟡

Q Manjit If you compare your sherpa and MadGraph samples, there is a bump in the ratio plot, can you really ignore this? 🟢 It's only a couple of bin, so yes, not so relevant.

Q Manjit s19: How do you choose 80% and 20% for the split? 🟢 roughly yes, maybe not the exact numbers, but the training should be the larger one

Q Manjit You've used mjj for one channel and DNN for another, if you're going to combine them, do you make some compatibility check of the two? 🟡

Q Paolo s13: Surprising you don't see much difference in the PS settings for the third jet. 🟢 We are sure that it was configured correctly. We can share the settings in any case

Comments about the paper draft

v1 25 February 2021

Comments implemented in the v2 of the paper.

Intorduction:

- Fig1: I m not really sure if you need 5 Feynman diagrams to get the point across. 🟢 ok, reduced to one example
- Ln 20: It s a bit odd to cite the 2016 result only. 🟢 added citation of latest result
- Ln 41: Would be better to define QCD-induced production more explicitly earlier 🟢 done

Section 4:

- I think Table 1 is not relevant to have. Nevertheless, we should have a table (or tables) with the (post)fit data, signal, and background yields. 🟡
- This reads more like an AN, and you cite and AN. The likelihood ratio test statistic is used in ~every CMS measurement with a search or low stats measurement. The appropriate citations are the usual profile likelihood ratio ones, not a CMS paper (refer to any published VBS result). 🟢 ok, added correct citation.
- It s also important to stress that your measurement is a search for a process, defined by the signal strength of the process, and that the significance of the signal strength is then quantified by the significance of the likelihood ratio test statistic. 🟢 ok, rephrased.
- Refer to published VBS results and restructure along these lines. 🟢 ok
- Ln 142: You can restore the broken line numbers by wrapping the equation in $\begin{linenomath*}$ and $\end{linenomath*}$ 🟢 fixed
- DeepFlav is this the right way to refer to this? I don t see any working points defined in the paper 🟢 It is referred also as DeepJet, changed.
- Ln 159: You extract the signal strength and then calculate its significance 🟢 ok

Section 5:

- We should see full run2 distributions, the individual years are irrelevant for outsiders. We should have the mjj and DNN distributions for all SRs and CRs. I believe it would also be good in the AN. You can have the split distributions in years in the Appendix, but it's better to show the combined distributions in the main body. Exceptions in cases when studying specific 2017 and 2018 issues. 🟡 ok, adding full run II distributions for SR and CRs.
- Fitting strategy. While it's not completely clear in the AN (as mentioned by Yacine and Kenneth), it's not clear in the paper draft either. I am not sure if I understand 1189 "ttbar and DY normalizations....after being initially left free to float". What does it mean initially? 🟢 For the normalization of the major backgrounds (tt and DY) data driven estimates using control regions are employed. The normalization of top and DY is left to float freely in the fit and con strained by the corresponding control region.

Results section:

- You don't really describe how the two analyses should be considered. Is the DNN one the nominal one? ● Yes, the DNN is the nominal one. In the paper (v2) we are going to quote the results obtained combining DF (DNN) with SF (mjj) categories.
- It's interesting to show the significances and signal strengths per channels and analyses, but not per year. ● Okay, updated table.
- Consider looking at the VVV discovery paper for inspiration on how to treat the two together side by side ● ok

Conclusions:

- You don't have any and you should ● added

Comments about the AN

Here questions regarding the AN are collected and addressed for each available version. Link to the gitlab repository; <https://gitlab.cern.ch/tdr/notes/AN-20-073/-/tree/master>


v6 03 March 2021

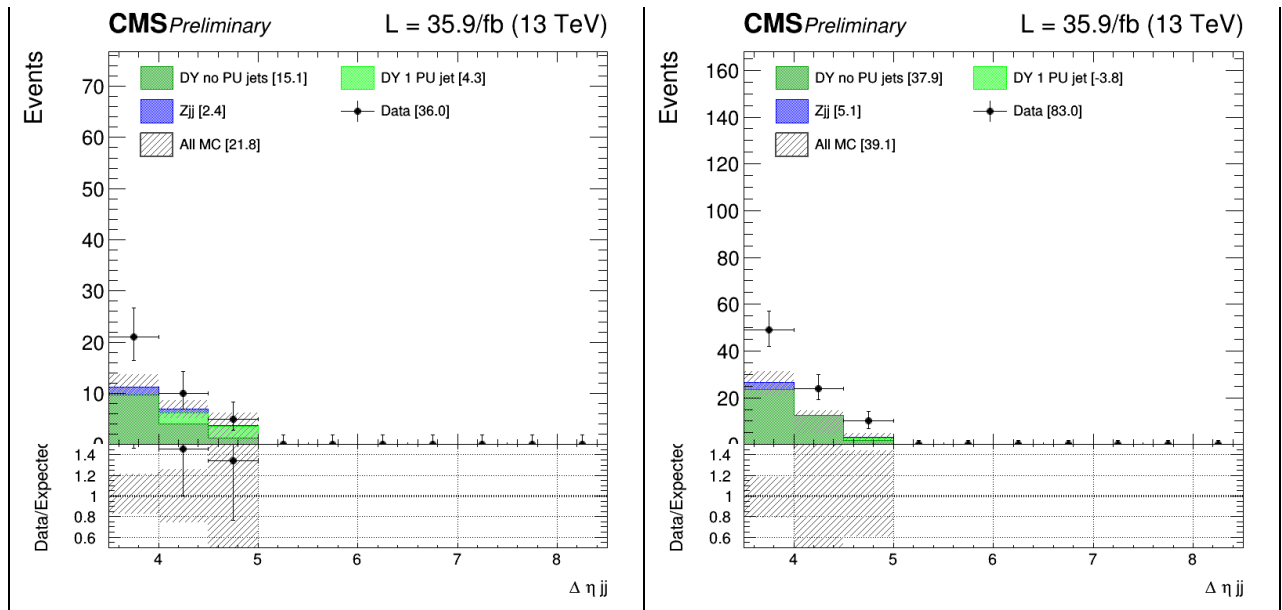
Link to the note: <https://icms.cern.ch/tools/notes/entries/AN/2020/073>

Comments for the pre-approval talk:


- TableSec 7.4, Fig 42 and 43: are the pileup jets defined by an ID or by matching to GEN?
- The DY_PUJets process is defined requesting at least one of the two leading reco jets with $pt > 30$ GeV not being matched to a GEN jet having $pt > 25$ GeV. Therefore, the DY_hardJets sample has both leading jets matched at GEN level. This sentence will be added in the AN as well.
- Are you sure that the issue is pileup, or could it just be mismeasurement outside the tracker? Do you have plots showing bins in jet eta tracker vs. outside? (perhaps both jets $eta < 2.5$, 1 jet $eta < 2.5$, and 2 jets $eta < 2.5$). Would also be kind of interesting to see inside HF or not ($eta > 3$).
- We believe that the issue, mostly visible in the 2016 DY sample, is due to the simulation of the hard radiation and/or to the relative fraction of events w/ and w/o PU jets. Plots of $detajj$ with both jets inside the tracker or at least one outside are shown below:

Two tracker jets ($ee/mumu$):





One tracker jet (ee/mumu):



As you may see, the region with both jets inside the tracker is entirely populated by the DY_hardJets process and shows a large disagreement. As a further cross-check, we did try to use this categorization to determine both DY_hardJets and DY_PUJets normalisations and results are in agreement with the strategy we are using in the AN (slide 7-8 <https://hypernews.cern.ch/HyperNews/CMS/get/SMP-21-001/8/1.html>).

- Fig 47 and 48: the nonprompt background statistics really seem insufficient. I don't think it's a good idea to fit with this background estimation. How many raw data events do you have here? Some possible approaches: - Loosen the ID somehow to have a better sample of events? - Combine all years rather than fitting separately - Derive the shape from a looser region and scale with the ratio of signal region/loose region

• We cannot define a looser selection than the one we are currently using to estimate the fake rate, the definition of lepton's WPs is the loosest possible satisfying the trigger-safe requirement. Moreover, nonprompt leptons are really a marginal background for the SF analysis, we indeed expect 3 events in the full run2 for the muon signal region category, which is basically less than an event per bin in m_{jj} , and 12 in the ee region.

- In section 8, you regularly reference splitting the DY into PU and no PU jet events. How is this defined in the signal region? Purely by splitting events with $\eta_{jj} > 5$ or < 5 ? I assume you also split

the signal region in these bins? Why is this never shown? It would be good to see the signal distributions with the DY colored according to the two contributions.

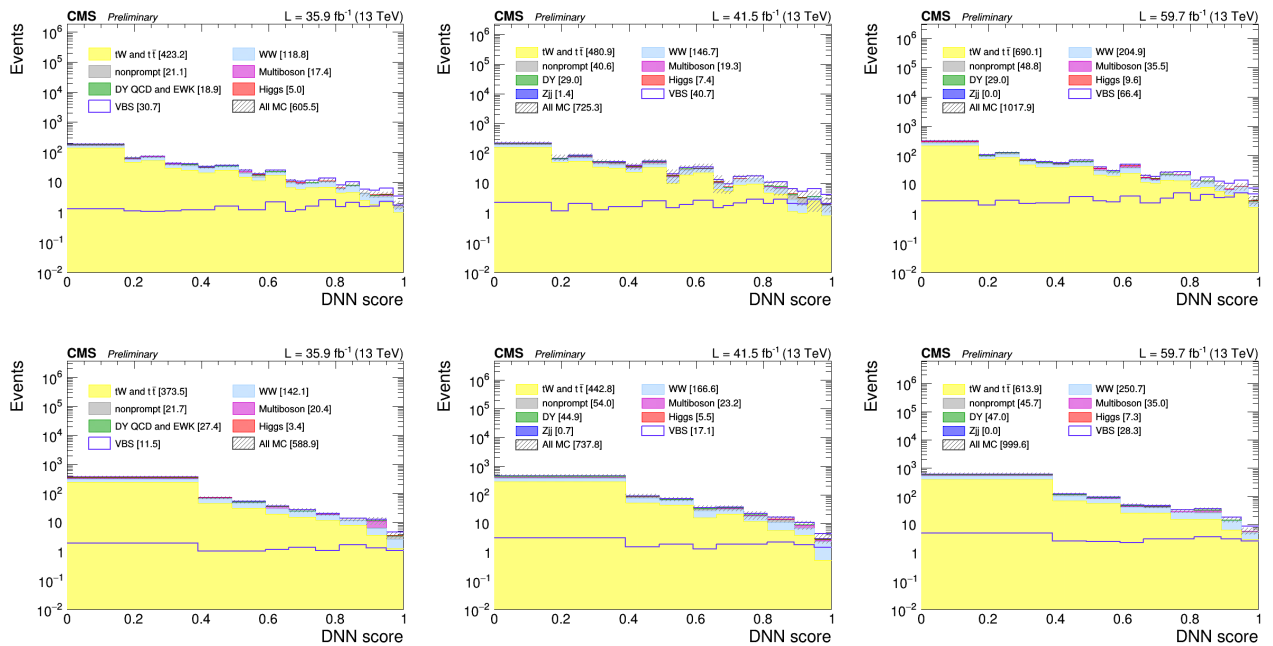
• Our strategy is to treat DY_PUJets and DY_hardJets as two different processes. Each of them has a dedicated control region: $\text{detajj} < 5$ DY CR is enriched with DY_hardJets events, while the other one is mainly populated by the DY_PUJets contribution. There is no detajj splitting in the signal region (see table 8) and the two processes contribute there with the yields determined in their respective CR. Both samples are shown in figures 47-48 (light green = "hard" DY process, dark green = DY with at least 1 PU jet).

• Fig. 49: Why is this the only place that Z EW is referenced? Is it included in other plots but not labelled? Also, DY EW isn't really meaningful since it's not a Drell-Yan process

• We will keep the Zjj sample separated from the pure DY, as it is in the rest of the AN.

• I'm kind of concerned that the stats are so low in the DNN distribution, Fig. 49. This should definitely be rebinned. Ideally the stats of the backgrounds would also be increased.

• We have rebinned the DNN output asking in each bin for at least one signal events, 2 signal + background events and a maximum of 30% of statistical error on background. For minor backgrounds we require a yield > 0 in all bins. The binning has been implemented on 2016 dataset and then applied to the other 2 years. In the figure you can see in top (bottom) row the $Zll < 1$ ($Zll > 1$) signal region for 2016/2017/2018 respectively.



New results are extracted and reported in table below.

year	significance	err. on signal strenght
2016	2.16	-0.48/+0.53
2017	2.33	-0.44/+0.48
2018	3.28	-0.32/+0.34
fullRun2	4.39	-0.24/+0.26

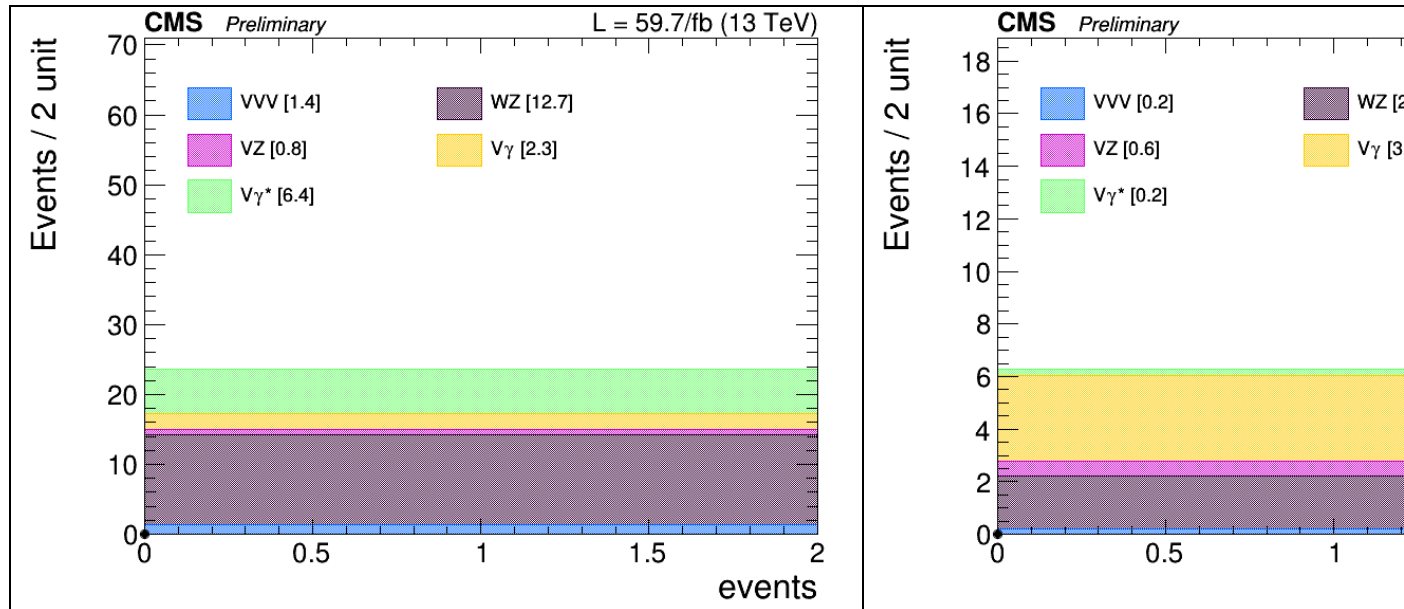
These results are going to be updated in AN v7.

• Can you clarify what DY sample you are using? Quite a lot are listed in the introduction, and the stats don't seem great

• Table 6 shows all DY samples we are employing for the SF analysis, for some of them we are also using available extensions for further increasing the statistics, we will include those as well in the list.

- Is WZ the major source of multiboson background? How many events do you have in the sample, and how many raw events pass the final selection

• Here's the number of events of each process entering in the "multiboson" definition (2018 data set). Plots are drawn in inclusive e/μ , e/e and μ/μ signal regions respectively: while WZ is the major contribution in the different flavour category, it is equally important as $V\gamma$ in the same flavour analysis.



- Did you check the WW+jj samples against the WW inclusive ones? We usually didn't use these VV+2j LO samples in the past, because the matching scale in Pythia gives very hard 3j radiation. It's worth at least checking the impact of using other samples if you have the statistics.

• Here you may find the comparison between inclusive and WWjj sample: the inclusive WW sample is plotted as data while the LO WWjj sample is the solid azure histogram. The dashed grey bands include both MC stat and theory uncertainties and mjj shapes are in agreement within error bars in almost each signal region.

2016: https://mlizzo.web.cern.ch/mlizzo/VBS/WW_Full2016v7/?match=mjj

2017: https://mlizzo.web.cern.ch/mlizzo/VBS/WW_Full2017v7/?match=mjj

2018: https://mlizzo.web.cern.ch/mlizzo/VBS/WW_Full2018v7/?match=mjj

- Are the impact plots up to date (Fig. 50-52)? I don't see all the parameters for the DY as I would expect Could you share the complete impact plot files (a link in the twiki would be enough)? There are various high ranked uncertainties that are purely statistical? This landscape would change with binning changed

• Plots shown in the AN are updated, the r_{vbs} estimation is mainly driven by the DF analysis and that's why SF-related nuisances don't impact much in the VBS measurement. You may find all pages here:

Link to full impact plots mjj analysis:

2016: https://mlizzo.web.cern.ch/mlizzo/VBS/impacts_ANv6/impacts_2016.pdf

2017: https://mlizzo.web.cern.ch/mlizzo/VBS/impacts_ANv6/impacts_2017.pdf

2018: https://mlizzo.web.cern.ch/mlizzo/VBS/impacts_ANv6/impacts_2018.pdf

Full Run2: https://mlizzo.web.cern.ch/mlizzo/VBS/impacts_ANv6/impacts_combination.pdf

Link to full impact plots DNN analysis:

2016:

https://bpinolin.web.cern.ch/bpinolin/VBSOS/plots_ANv6/impact_plots/impacts_combine_SFDF_2016_new_new.pdf

2017:

https://bpinolin.web.cern.ch/bpinolin/VBSOS/plots_ANv6/impact_plots/impacts_combine_SFDF_2017_new_new.pdf

2018:

https://bpinolin.web.cern.ch/bpinolin/VBSOS/plots_ANv6/impact_plots/impacts_combine_SFDF_2018_new_new.pdf

Full Run2:

https://bpinolin.web.cern.ch/bpinolin/VBSOS/plots_ANv6/impact_plots/impacts_combine_SFDF_FullRun2_new_new.pdf

- Did you share your combine cards with Pietro (and us) yet?

🟢 This is the gitlab repository where all datacards have been uploaded:
<https://gitlab.cern.ch/cms-hcg/cadi/smp-21-001>

v5 26 January 2021

Link to the note: <https://icms.cern.ch/tools/notes/entries/AN/2020/073>

In v5 all comments from v4 have been implemented. Main concerns related to v5 are the followings:

- We would expect to gain more sensitivity when using a DNN approach to extract the signal, at this level both m_{jj} and the DNN score show similar results. Is there room for any optimization?
- 🟢 The training procedure now includes both QCD WW and $t\bar{t}$ pair production as backgrounds. Doing so, the expected statistical significance increases by roughly ~15% in the different flavour analysis and when combining all categories together we almost reach 5 expected sigma.
- The same flavour DY control region shows some criticities in data/MC agreement, especially for the 2016 data set. Have you tried to implement a bin-by-bin corrections?
- 🟢 In order to tackle the observed data/MC disagreement we changed paradigm for the same flavour analysis. The new strategy we came up with is based on two main points: 1) Discrepancies strongly depend on detajj and this could be the hint of a PU dependancy; 2) CR and SR need to be as similar as possible. Eventually we split the DY sample into two contributions, one including events in which at least one jet comes from PU and the other one for the remaining "hard" events. Two independent parameters are used to scale their normalizations in the fit procedure. In order to gain sensitivity to these contributions, the DY control region has been divided into 2 detajj bins ($>$ or $<$ than 5). Besides we increased the MET cut up to 60 GeV, as it is for the SR as well. Although "hard"-like events are unlikely to be found in such a high-MET region, the categorization in detajj is suitable for separating the two DY sub-samples and allows a better estimation of their yields.

v4 13 January 2021

Link to the note: <https://icms.cern.ch/tools/notes/entries/AN/2020/073>

Follow up on generators:

We never managed to produce a meaningful sample with POWHEG.

- Could you be more specific on the issues encountered when trying to produce those samples? Perhaps GEN group can be of help? Even if the issues are critical with POWHEG it's worth documenting the studies that you made for reference.

🟢 The issue we encountered with Powheg was related to the sample generation, as it appeared like all events had the same seed. We tried to get in contact with Powheg's authors but we never had a follow-up on that, hence we dropped the study.

- A study of MadGraph
<<https://twiki.cern.ch/twiki/bin/edit/Main/MadGraph?topicparent=Main.QandAforVBSOSWW;nowysiwyg=1>> at Gen level would also be useful. This could be done on NanoGen
<<https://twiki.cern.ch/twiki/bin/view/Main/NanoGen>> pretty easily. We can help you with the configuration, then you just need to generate events and make a few comparison plots of your sensitive variables at Gen level. Since this is the first time this state has been studied, it would make the analysis stronger.

🟢 This has been documented in the AN (see figure 6).

We performed a preliminary study where we compare our signal sample at GEN level (starting from MiniAOD files) with LO VBS W+W- sample generated with Sherpa, along with its built-in parton shower. The Rivet analysis employed for this comparison contains the main cuts which define our signal selection. We considered jets with $pt > 30$ GeV, from which we have further removed leptons with $pt > 10$ GeV contained in their cone ($R = 0.4$). Both samples are affected by an issue affecting the colour reconnection scheme, which results in generating more jets within the pseudorapidity gap of the two tagging jets. In Sherpa, a fix for this problem is available, and the difference in the production rate of the third jet is well visible. Nevertheless, inclusive two-jets distributions agree within a fair 10%, and there are no relevant shape differences affecting m_{jj} , which is our chosen fit variable.

Sherpa vs Madgraph: https://mlizzo.web.cern.ch/mlizzo/Rivet-plots/Sherpa_vs_MadGraph/

Sherpa + PS fix vs Madgraph: https://mlizzo.web.cern.ch/mlizzo/Rivet-plots/Sherpa_Psfix_vs_MadGraph/

- Could you please specify what PS did you use for the Madgraph samples? is it with the default Pythia8 or Herwig? it would be useful to have both. In the case of the Pythia8 it would be useful to check the dipoleRecoil option as well. Would it be possible to update these plots with more statistics?

🟢 The PS used with MadGraph samples is the default Pythia8. Plots with more statistics have been uploaded in the AN (see figures 4 and 5).

- Also the Sherpa PS fix vs Madgraph plots shows large differences mainly in the 3rd jet variables and that's indeed due to the colour reconnection scheme. Even though the checks done previously showed that the cut-based analysis is not affected by the issue, now that you have a DNN approach the conclusion might be different. I would suggest also to check the impact on the DNN with Sherpa-PS-fix to start and with MG5+Herwig when ready.



Follow up on DNN discussion:

We compare the ROC curves obtained applying the models to the analysis samples to estimate the discrimination power of a network wrt to another. As to the overfitting, we check that the loss function evaluated on the validation dataset does not increase with the number of epochs, but decreases or remains stable (as the ones we show in the fig. 8 of the AN). Moreover, we are also considering other two metrics: the recall ($TP/(TP+FN)$) and the precision ($TP/(TP+FP)$). And finally, we also check that the distribution of the DNN score obtained with the training and with the validation samples are overlapped.

- If I understand correctly the optimisation is done by "hand", so you check if the loss function is relatively flat and does not increase with the epoch. Is that correct? have you tried using a more quantitative approach such as Kolmogorov-Smirnov test? This is, I believe, what the SMP-20-013 is using.

🟢 We are implementing, as suggested, the Kolmogorov-Smirnov test in the optimisation procedure of the latest networks to further check the absence of overfitting.

- Figure 10-11-12: I see that the loss function is oscillating with the number of epochs (same pattern with the efficiency and purity). Do you have an explanation for this? On my knowledge, such behaviour is symptomatic of an optimisation oscillating around a saddle point. Maybe you can reduce the learning rate so that the gradient descent doesn't overshoot the minima. Also, I see that (line 380) the LR is automatically optimised as the learning progresses. Could you show a plot of the LR as a function of the epochs? Maybe the oscillation is an artefact of this automation.

🟢 The oscillation pattern you see in the metrics is due to the Cyclical Learning Rate algorithm [1] used in the training. With this method, three parameters are set for the learning rate: a lower and an upper bound and step size. Thus, the learning rate increases from the lower to the upper bound in steps; when reaching the upper bound, the learning rate decreases until the lower bound is touched; the process repeats during all the training. Figure [2] shows an example of the behavior of the learning rate during each iteration of the training. The wave-like behavior of the loss is a consequence of this learning rate oscillation. In particular, the bottom of the wave corresponds to the minimum value of the learning rate, while the top corresponds to the maximum learning rate. The Cyclical Learning Rate helps prevent overfitting and reduces the number of iterations needed to optimize the networks.

[1] <https://arxiv.org/abs/1506.01186>

[2] <https://cernbox.cern.ch/index.php/s/3Efm6UG9XvDViat>

v3 04 January 2021

Link to the note: <https://icms.cern.ch/tools/notes/entries/AN/2020/073>

- The numbers in Tables 9-11 between v2 and v3 have changed quite a lot, the signal is changing by almost 10% in 2016. We really need a more detailed explanation of what changed here. This is still the same selection, without the DNN involved, right? It would really speed up our review to give a breakdown of the impact of individual changes. Just NanoAODv5 --> NanoAODv7 is too vague, we need to know what corrections etc are changing that impact the physics results.

🟢 In addition to the change in NanoAOD version there are two additional modifications: the working point for the muons has been changed, following a similar change in the HWW analysis from which we inherit the object definition. In particular, for muons we have moved from a cut based WP to a WP cutting at 0.8 on the $t\bar{t}H_{mva}$, as described in the AN 2019/125. Also we have moved the bveto from the DeepCSV loose WP to the DeepFlavor loose WP. Both improve sensitivity in almost all categories.

- Table 9-11: How do you treat the negative nonprompt yields? (there is still one negative yield in the new version, there were several in the old).

• At the moment they are go into combine as they are.

- General point: I agree with Yacine's comment that studying the signal with another generator would be useful. I remember studying POWHEG some time ago. Did you conclude that there was an issue with POWHEG?

• We never managed to produce a meaningful sample with POWHEG.

- A study of MadGraph+Herwig at Gen level would also be useful. This could be done on NanoGen pretty easily. We can help you with the configuration, then you just need to generate events and make a few comparison plots of your sensitive variables at Gen level. Since this is the first time this state has been studied, it would make the analysis stronger.

• We performed a preliminary study where we compare our signal sample at GEN level (starting from MiniAOD files) with LO VBS W+W- sample generated with Sherpa, along with its built-in parton shower. The Rivet analysis employed for this comparison contains the main cuts which define our signal selection. We considered jets with $pt > 30$ GeV, from which we have further removed leptons with $pt > 10$ GeV contained in their cone ($R = 0.4$). Both samples are affected by an issue affecting the colour reconnection scheme, which results in generating more jets within the pseudorapidity gap of the two tagging jets. In Sherpa, a fix for this problem is available, and the difference in the production rate of the third jet is well visible. Nevertheless, inclusive two-jets distributions agree within a fair 10%, and there are no relevant shape differences affecting m_{jj} , which is our chosen fit variable.

Sherpa vs Madgraph: https://mlizzo.web.cern.ch/mlizzo/Rivet-plots/Sherpa_vs_MadGraph/

Sherpa + PS fix vs Madgraph: https://mlizzo.web.cern.ch/mlizzo/Rivet-plots/Sherpa_PSfix_vs_MadGraph/

- We think it would be important to make a combined EW+QCD measurement in a fiducial region. Using the shape-based fit for this, with EW WW and QCD WW as signal, should be an easy addition that is appreciated by theorists.

• We are currently working on that and we will soon implement the measurement in the documentation. We have not yet settled on a fiducial volume definition, but we propose to perform the fit in such a way that the fiducial and nonfiducial signal components entering the signal region are scaled together. If we follow this approach the fiducial volume definition does not matter when fitting, and plays a role only when translating the signal strength extracted from the fit into a fiducial cross section. We already were able to fit the EWK+QCD sample as signal, and for that we get an expected result for the signal strength of 1 ± 0.26 . We would like to work on the exact definition of the fiducial region between now and the preapproval.

- Ln 100: There are a lot of definitions of the Zeppenfeld variable. The one you use is sometimes called the centrality (zeta), with the Zeppenfeld variable reserved for $z_{\text{etall}}/z_{\text{etajj}}$. Did you try the zeppenfeld with this definition as well? It would probably be clearer to adopt this language (as in SMP-18-001)

• We have tried to use for the categorization of the signal region $Z_{\text{etall}}/z_{\text{etajj}} = \frac{1}{2} \frac{\text{abs}((\text{lep1} + \text{lep2}) - (\text{jet1} + \text{jet2}))}{|\text{jet1} - \text{jet2}|}$ instead of the usual Z_{ll} (defined at line 100 of the AN). We had a quick test using only different flavor categories and only top control region in the final fit. We tried some different scenarios, splitting the signal region in two categories wrt to $Z_{\text{etall}}/z_{\text{etajj}}$ and changing the cutting value from 0.1 to 0.5 in steps of 0.05. Results are reported in the table below.

cut on $Z_{\text{etall}}/z_{\text{etajj}}$	Significance $Z_{\text{ll}}/z_{\text{etajj}}$
0.1	2.34

0.15	2.43
0.2	2.39
0.25	2.38
0.3	2.36
0.35	2.27
0.4	2.25
0.45	2.21
0.5	2.36

The significance obtained with the usual categorization (i.e. $Z_{\text{ep}} < 1 / Z_{\text{ep}} > 1$) is 2.56. Therefore, the usual categorization has the best performance.

We will adopt, as suggested, the naming convention as in SMP-18-001.

- Sec. 6.1: It's awfully hard to see the improvement in a lot of these plots. Can you show only the region of interest, and plot $\text{abs}(\eta)$ as well so there are more stats to see the performance?

• We have plotted $\text{abs}(\eta)$ of the two leading jets for all the flavor categories (ee, mm, em) in the top [1] and DY [2] control regions. The data/MC agreement in the horns region ($2.5 < |\text{jet}| < 3.2$) is everywhere good. These plots will be included in section 6.1 of ANv4.

[1]

https://fcetorel.web.cern.ch/fcetorel/VBS_OS/test/2017/ControlRegions_jethornscheck_ANv2_v7_100121/top/

[2]

https://fcetorel.web.cern.ch/fcetorel/VBS_OS/test/2017/ControlRegions_jethornscheck_ANv2_v7_100121/DY/

Questions on the impact plots, Fig. 42-44:

- QCDscale_top_2j wasn't shown in the previous version. Is this the shape uncertainty of the top background? Was it just overlooked? Is it not included in the norm param because of the shape effect?
- In the previous version (ANv3), QCDscale_top_2j wasn't accounted for and it's the QCD scale uncertainty related to the top background. Both up and down variations are calculated as the difference between the nominal histogram and the envelope obtained by considering the highest up and down QCD scale variation in each bin. Such uncertainty is treated as a shape effect and the varied distribution is normalised to the nominal integral (that's indeed why it is not included in the rate parameter).
- In the previous version, you had an uncertainty labeled CMS_scale_met , and I was wondering what this is. Is this the JES propagated to the MET or is it the unclustered energy? Did you remove it or did it get pushed further down the ranking?
- In the current version CMS_scale_met is still presented but has been slightly pushed down in the ranking by other uncertainties. It's computed by varying the MET energy scale of PF algorithm candidates which are not clustered into jets and it is properly propagated to other variables which depend on the MET itself. Up and down histograms are then normalised to the nominal one, thus this contribution is treated as a shape effect.
- What is the primary source of the stat uncertainties that are dominant in the impact plots? Is it the stat uncertainty on the nonprompt?
- The primary source of statistical uncertainties in the impact plot is mainly due to the top sample in almost all m_{jj} bins within different flavour categories and to the DY contribution in some flavour categories.
- Can you group these all together and evaluate the total contribution to the uncertainty from the stat group? Similarly, can you do this for the JES? You can do this with a likelihood scan, e.g., following

the instructions here:

<https://indico.cern.ch/event/577649/contributions/2388797/attachments/1380376/2098158/HComb-Tutorial-N>

• Here is shown uncertainties breakdown, performed over a likelihood scan on the Asimov dataset. The total error is split into JES, systematic and statistical contributions; the latter is clearly what limits our analysis: The plots will be included in an appendix of AN version 4.

https://mlizzo.web.cern.ch/mlizzo/VBS/JES_breakdown/

- Where is the nonprompt norm uncertainty? For the combined fit, can you put all the nuisances into the appendix?

• The main uncertainty source on the "Fake" sample is a normalization uncertainty of 30% derived from a closure test in MC. This uncertainty is modeled as a lognormal distribution, separately for events with a subleading electron or muon. They rank 78 and 169 in the combined impacts plot with an effect of 0.5% and 0.2% respectively on the signal strength. We will create an appendix in version 4 of the AN for all nuisances considered in the combined fit.

- Some of your JES and JER uncertainties are pretty one-sided. Can you add a few illustrative examples of the input shapes you use to the AN?

• Overall JES/JER uncertainties seem reasonable, although for some of them Up/Down variations are indeed one-sided in few mjj bins, as it may be observed here for the 2018 dataset:

https://mlizzo.web.cern.ch/mlizzo/VBS/Full2018_v7/JES+JER/

Most impactful JES + JER uncertainties are drawn for main processes, i.e. VBS, top and WW samples, in each signal category. Similar plots are extracted for other years. These plots will be included as an appendix in the new version 4 of the AN.

Questions about DNN approach:

Section 5.0:

- You have mentioned that the datasets should be balanced, so you increased the signal samples weights in training. Does that mean that you include the event weights in a way or another in the DNN training? if so can you be more explicit how this information is incorporated in the DNN?

• Yes, the weights of the events are considered in the DNN training. In particular, the loss computed for each sample is multiplied by the weight associated with it. In this way, the back propagation will behave differently depending on the weight of the events, giving more importance to the events with a higher weight.

At first we consider as weight for each event $XS \cdot \text{lumi} \cdot SF$, and then a balancing is made. This means that the total number of weighted events of the signal dataset should be the same as the one of the backgrounds datasets combined. This is achieved increasing the weight of the signal samples in the training, using as weight: $\text{weight} / \text{mean}(\text{weights})$. While to balance the background we use as weight: $\text{weight} \cdot nS / \text{sum}(\text{weights})$, where nS represents the number of simulated signal events.

- Since you have divided the samples into two datasets one for training and the other for validation, I think it would be good to show the DNN probability distributions for both training and testing to illustrate the absence of overfitting.

• We will include the DNN probability distributions for both training and testing in the updated documentation v4.

- What loss function are you using?

• We are using the binary cross entropy as loss function.

- It seems that you have used only ttbar samples as background. Out of curiosity, have you tried including other backgrounds to see if the discrimination power improves or deteriorates?

• Until now we have considered only ttbar as background, because it is the dominant one in the signal region (its yield is ~10 times the WWqcd one, which is the second most relevant background). We are trying to add in the training also the WW qcd to see if it will improve the network performance.

- It would be nice to see some of the ROC curves you are mentioning in the text.

• We will add the ROCs comparison for mjj and DNN in the version 4 of the AN. Here [1] ([2]) some examples for the low Zll (high Zll) categories for the three years. The DNN performs better than mjj.

[1] <https://bpinolin.web.cern.ch/bpinolin/VBSOS/ROCs/lowZ/2016>

<https://bpinolin.web.cern.ch/bpinolin/VBSOS/ROCs/lowZ/2017>

<https://bpinolin.web.cern.ch/bpinolin/VBSOS/ROCs/lowZ/2018>

[2] <https://bpinolin.web.cern.ch/bpinolin/VBSOS/ROCs/highZ/2016>

<https://bpinolin.web.cern.ch/bpinolin/VBSOS/ROCs/highZ/2017>

<https://bpinolin.web.cern.ch/bpinolin/VBSOS/ROCs/highZ/2018>

Section 5.1:

- Could you substitute the N in the text to reflect the results obtained? As I understand, the DNN optimisation is still ongoing, but it would be good to mention the architecture used to make sense of the results.

• In the v4 of the AN we will fix this. However, we are using neural networks with 2 or 3 hidden layers, and a number of neurons that goes from 50 to 150.

- Maybe this is not important in your case, but have you tried using dropout layers? this has proven to reduce overfitting.

• During the optimisation of a network we try different architectures and tools; we try dropout layers as well. It's true that they help to reduce overtraining, but in some cases they inficiate the performance of the DNN, and therefore in those cases they are discarded.

- You mentioned that a down-weight of mjj/2000 is applied, I am curious to know how this information is used in the DNN.

• We multiply the weights of the events for mjj / 2000 only if mjj \geq 2000 GeV. In this way we give more importance to all the high-mjj events (i.e. the events with mjj > 2000 GeV) during the training process. In the training of the DNN this information is used with a direct rescaling of the loss function. In fact, the loss computed for each sample is multiplied by the weight associated with it. In this way, the back propagation will behave differently depending on the weight of the events, giving more importance to the events with a higher weight.

Section 5.2:

- The strategy consists of choosing the best variables and have a tradeoff between overtraining (line 370) and discrimination power. For the discrimination power, I guess you used the area under ROC, right? Could you provide us with the methodology used to estimate the overfitting?

• We compare the ROC curves obtained applying the models to the analysis samples to estimate the discrimination power of a network wrt to another. As to the overfitting, we check that the loss function evaluated on the validation dataset does not increase with the number of epochs, but decreases or remains stable (as the ones we show in the fig. 8 of the ANv3). Moreover, we are also considering other two metrics: the recall ($TP/(TP+FN)$) and the precision ($TP/(TP+FP)$). And finally, we also check that the distribution of the DNN score obtained with the training and with the validation samples are overlapped.

- In line 367, you mention that an optimal value has to be searched, it would be nice to show more details on that.

• To find the optimal configuration, we started with a small DNN (2 layers with 20 neurons each) and then we trained it with as many variables as possible. If the DNN overtrained, we ranked the variables thanks to the SHAP (see AN-2019/239), considering their importance in terms of impact on the DNN output, and we removed the 2 less important variables. Then we repeated the process (training->ranking->variables removing) until the DNN was not overtrained anymore. If the results in terms of performance were not satisfying, then, we incremented the DNN structure (number of layers and/or neurons) and repeated the process until we found the optimal set of training variables with this new structure. We have repeated all this process until we have found a DNN with a satisfying performance, that means with a ROC curve that shows a better performance with respect to mjj in the whole phase-space

v2 23 November 2020

Link to the note: <https://icms.cern.ch/tools/restplus/relay/piggyback/notes/AN/2020/73/files/2/download>

General comments: * Various references are missing (example: Line 258, Line 289,)

• References are updated.

* Out of curiosity, I see you have mentioned a DNN approach in line 112: are you also considering implanting a DNN analysis besides the cut-based one?

• Yes, we are working in parallel on a DNN approach in the different flavor category to boost the analysis performance.

Section 3:

- Why are you using NanoAODv5 for 16 and 17 datasets? The current version is v7, are you planning to update soon?

• Yes, we are planning to update the analysis, moving it to NanoAODv7 datasets.

* For the signal, you are using MG5 interfaced with Pythia 8, where you require 2 jets in the final state at LO. This could lead to large discrepancies in a case of third jet veto (such the Zll variable), due to a mis-modelling of colour-connection in Pythia 8. You could consider generating WW+3j at the LO with the dipoleRecoil=on in the Pythia 8 settings in order to mitigate this issue. You can find more details on the following links:

- <https://cds.cern.ch/record/2655303/files/ATL-PHYS-PUB-2019-004.pdf>

- <https://arxiv.org/pdf/1812.05118.pdf>
- <https://link.springer.com/article/10.1140/epjc/s10052-020-8326-7>
- <https://indico.cern.ch/event/961185/#24-recent-vbf-recommendation>

I would also recommend using a different parton shower (Herwig++ or 7) as cross-check

The Zll variable should not introduce additional mismodelling in our signal sample, since it's not strictly related to the third jet kinematics. Rather, it describes the polar distribution of the di-lepton system w.r.t. the two tagging jets and, for the signal, we expect to find more activity in the central region. Indeed this is what happens and that's why the $Zll < 1$ category is enriched with signal and has a favourable S/B ratio. As additional evidence to such behaviour, we provide the main jet distributions for the signal sample, evaluated both inclusively in Zll and applying the categorization (example provided for the 2016 dataset -> might be updated with 2017 and 2018):

- em_me inclusive: https://mlizzo.web.cern.ch/mlizzo/VBS/Full2016_ANv1/Zll_inclusive/?match=*em*j
- em_me Zll cut: https://mlizzo.web.cern.ch/mlizzo/VBS/Full2016_ANv1/Zll_categories/?match=*em*j
- ee inclusive: https://mlizzo.web.cern.ch/mlizzo/VBS/Full2016_ANv1/Zll_inclusive/?match=*ee*j
- ee Zll cut: https://mlizzo.web.cern.ch/mlizzo/VBS/Full2016_ANv1/Zll_categories/?match=*ee*j
- mm inclusive: https://mlizzo.web.cern.ch/mlizzo/VBS/Full2016_ANv1/Zll_inclusive/?match=*mm*j
- mm Zll cut: https://mlizzo.web.cern.ch/mlizzo/VBS/Full2016_ANv1/Zll_categories/?match=*mm*j

No differences in the shape of the distributions are visible, meaning that the Zll cut does not affect the third jet kinematics.

* Have you checked if you are effected by the HEM issue in 2018 dataset?

We apply on 2018 datasets the recipe to cure the HEM issue [1]. The effect on our control regions is negligible, as you can see comparing plots where corrections are applied (top [2], DY [3]) to the ones in which they are not (top [4], DY [5]). The checks on HEM issue will be included in section 6.2 of ANv3.

[1] <https://hypernews.cern.ch/HyperNews/CMS/get/JetMET/2000.html>

[2]

https://fcetorel.web.cern.ch/fcetorel/VBS_OS/test/2018/ControlRegions_v6_HEM_v3_141220/top_corrHEM/

[3]

https://fcetorel.web.cern.ch/fcetorel/VBS_OS/test/2018/ControlRegions_v6_HEM_v3_141220/DY_corrHEM/

[4] https://fcetorel.web.cern.ch/fcetorel/VBS_OS/test/2018/ControlRegions_v6_HEM_v3_141220/top/

[5] https://fcetorel.web.cern.ch/fcetorel/VBS_OS/test/2018/ControlRegions_v6_HEM_v3_141220/DY/

Section 5:

* You applied the PUJID only the $2.5 < |jetl| < 3.2$ region, have you tried to apply the pileup id to other eta regions? maybe this would improve the agreement of the very forward jets

● We are already applying a PUJID loose in all the eta range for all jets with $pt < 50$ GeV. In addition to that, in 2017 we require the two leading jets to pass the tight PUJID wp, if their eta is in the range $2.5 < \eta_{jet1} < 3.2$.

* In the note we understand that the jet horns are an issue only in 2017? Have you checked for 2016 and 2018? I do remember that in VBF Higgs we have seen the same issue in 2016 dataset as well.

● We checked both 2016 and 2018 datasets to find if the jet horns issue was affecting them. As to 2018, in both DY[1] and top [2] CR the agreement in the $2.5 < \eta_{jet1} < 3.2$ looks quite good. In 2016 the agreement is a bit worse (DY[3], top[4]), in particular for DY cr in the same flavor categories, but still not comparable to what is observed for 2017 [see fig. 8-12 of the ANv2].

[1] https://fcetorel.web.cern.ch/fcetorel/VBS_OS/test/2018/ControlRegions_v6_jethorns_111220/DY/

[2] https://fcetorel.web.cern.ch/fcetorel/VBS_OS/test/2018/ControlRegions_v6_jethorns_111220/top/

[3] https://fcetorel.web.cern.ch/fcetorel/VBS_OS/test/2016/ControlRegions_jethorns_041220/DY/

[4] https://fcetorel.web.cern.ch/fcetorel/VBS_OS/test/2016/ControlRegions_jethorns_041220/top/

* Also on the same note, have you applied the latest JEC/JES recommendations? If not, you might consider updating to the latest recipe that showed better Data/MC agreement in the horns region.

● We are planning to update soon the analysis from NanoAODv5 to NanoAODv7, which include the latest JEC/JES recommendations (here GT comparison of the two versions [0]).

[0]

https://cms-conddb.cern.ch/cmsDbBrowser/diff/Prod/gts/102X_mc2017_realistic_v7/102X_mc2017_realistic_v8

Section 8:

* Can you be more explicit about the treatment of the theory uncertainty in the VBS signal? from the text it seems as if you varied only the factorisation scale by 1/2 and 2.

● The theory uncertainty on the VBS signal is indeed evaluated by varying the factorisation scale by 1/2 and 2. However, since the normalization of the signal is measured during the fit procedure, we divided the varied histograms by the integral of the nominal one (i.e. the one with $\mu_F = 1$), in order to account for possible modifications affecting only the shape of the distributions.

* Can you also comment on how the experimental uncertainties are correlated across years?

● Experimental uncertainties are kept uncorrelated across the three years, as mentioned in lines 440-442 ANv2.

v1

Empty skeleton, first draft.

General questions and discussion

Minutes from 15-09-2020 SMP-VV

Philip :

- The e/mu regions are still dominated by the top backgrounds, you might consider finding more variables to reduce this. In ATLAS in Run I, this was done with a cut on the m_{T2} variable. Take a look at the corresponding paper and see if this variable would be useful. This was meant to target top quark mass to discriminate against $t\bar{t}$. If i remember correctly the variable was computed with some min, or max of [$MT2(1lv+vb\text{fjet}1)$, $MT2(1lv+vb\text{fjet}2)$].
https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/HIGG-2013-13/fig_20.pdf/

• We are planning to include m_{T2} in the analysis to see if it can help in the suppression of the top backgrounds. We are investigating to understand the definition of the variable.

Paolo :

- You re using the LO MC for Drell-Yan, can you switch to the NLO one?
- The NLO DY sample has not enough statistics to populate the signal region we have defined in the anlysis, thus we use LO HT-binned samples to provide for lack of MC stat in the same flavour categories. We share this approach with the HWW high mass analysis.
- We did try to employ the NLO DY sample instead of the HT binned samples and we observed a general improvement in the high Z_{ll} DY CR. However, this doesn't hold for the low Z_{ll} category, where a significant discrepancy between data and MC is still present.
 - You also process the VBF Z sample, one would expect that this could be significant.
- We included the Z_{jj} sample in the analysis. Still, its contribution seems to be not so significant and it does not cover the data-MC gap.
 - Can you request the signal sample with the Pythia dipole recoil shower (and Herwig)? Perhaps in the UL?
- Working on it.

Yacine :

- On the categorization, you say that the Zeppenfeld variable improves the sensitivity. Did you try it wrt other variables? Have you tried using the Z_{l1} rather than just Z_{ll} ?
- We tried using Z_{l1} (instead of Z_{ll}) to split the signal region in two categories for 2018 ($Z_{l1} < 1$ and $Z_{l1} \geq 1$). The signal purity in region $Z_{ll} < 1$ (expected to have the most favorable S/\sqrt{B}) is not as good as the one in the old $Z_{ll} < 1$ category. Thus we obtain a statistical significance (2.49) worse than the one obtained with the old configuration (3.07).
 - Also, how did you optimize the binning for the m_{jj} ?
- We optimize the binning requiring no empty bins.

MattiaLizzo - 2021-03-01

This topic: Main > QandAforVBSOSWW

Topic revision: r55 - 2021-05-06 - FlaviaCetorelli



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback