# Advancing Single-Photon Searches in MicroBooNE with PROfit: Systematic Studies and NC Coherent 1$\gamma$ Sensitivity Optimization

## REU Program at Columbia University – Nevis Laboratories

Scott Anderson[1]

[1]DePauw University

August 5, 2026

## Abstract

The MiniBooNE anomaly is an unsolved problem in neutrino physics regarding the origin of an observed excess of electromagnetic shower events at low energies. MicroBooNE has the capability to identify whether these events stem from electrons or photons. This paper focuses on three studies pertaining to the search for photon events in MicroBooNE. The BAGELS event selection framework and PROfit statistical fitting framework were used for these studies. The first two studies provide a greater understanding of single-photon backgrounds. Both examine whether individual systematic uncertainties could be responsible for previously observed discrepancies between prediction and observation in neutral-current $\pi^0$ events with hadronic activity and two-photon events originating outside of the detector's fiducial volume. These studies were unable to attribute the discrepancies to individual systematic effects. The third study optimizes sensitivity to coherent neutral-current single-photon events. This study details the process of selecting the most sensitive binning variable, bin-edge placement, and the most useful sidebands. A signal-specific binary BDT was trained and tuned for this purpose. The analysis predicts the exclusion of scale factors greater than 2.67 times the Standard Model prediction at 90% confidence with $1 \times 10^{21}$ protons on target. A single-bin configuration using the tuned BDT score was applied to MicroBooNE open data with $9.57 \times 10^{19}$ protons on target and excludes scale factors greater than 3.10 times the Standard Model prediction at 90% confidence. These sensitivities are preliminary because signal cross-section and flux uncertainties, as well as detector systematics, are not yet included.

# Contents

# 1 Introduction

## 1.1 Neutrinos

Neutrinos are notoriously elusive Standard Model (SM) particles belonging to the lepton family. They carry no electric charge and exist in three flavors: electron, muon, and tau. Each neutrino's flavor is identified by the charged lepton produced alongside the neutrino in a charged-current weak interaction. Despite their weakly interacting nature, neutrinos are among the most abundant matter particles in the Universe, produced in processes including nuclear reactions in stars, radioactive decay, supernovae, and cosmic-ray interactions [1]. A unique characteristic of neutrinos is their ability to oscillate, or change flavor, as they travel. The minimal SM originally treated neutrinos as massless. However, the experimental observation of neutrino oscillations demonstrates that neutrinos have nonzero mass [2]. Neutrinos are therefore the lightest known massive elementary particles [3]. Their lack of electric charge, small masses, and inability to interact through the electromagnetic (EM) or strong nuclear forces make them exceptionally difficult to detect. Nonetheless, experiments like MiniBooNE study neutrino behavior and search for new physics that may provide insight into open questions such as the nature of dark matter and the origin of the matter-antimatter asymmetry [1].

## 1.2 The MiniBooNE Anomaly

The MiniBooNE experiment was a mineral-oil Cherenkov detector in the Booster Neutrino Beamline at Fermilab that was built to study neutrino oscillations at a short baseline [4]. Charged particles exceeding the speed of light in that medium produced cones of Cherenkov light, which appeared as rings on the detector's photomultiplier tubes. Although it could effectively reconstruct the energy and direction of the EM activity, MiniBooNE could not reliably distinguish between electron-induced and photon-induced particle cascades. This limitation became problematic when MiniBooNE observed a $4.8\sigma$ excess of EM shower events at low reconstructed energies [4]. MiniBooNE could not determine the origin of this anomaly. The construction of MicroBooNE was motivated by the need to investigate the unresolved excess.

## 1.3 MicroBooNE and LArTPCs

MicroBooNE was constructed in the same beamline as MiniBooNE at Fermilab. It is a liquid-argon time projection chamber (LArTPC) with approximately 85 metric tons of liquid argon in its active volume [5]. Liquid argon provides a cheap, dense, scintillating detector medium that ionizes when charged particles pass through it. An applied voltage across the detector creates an electric field that drifts liberated electrons toward an anode plane on the detector wall. Electrons induce signals and are collected by wires in the anode plane. Figure 1 provides an illustration of this readout process. Combining wire position, charge arrival time, and signal size gives multiple views of the interaction and permits three-dimensional reconstruction and calorimetric energy measurement. Photomultiplier tubes also inform the timing and position of events in the detector [5].

The combination of charge and light information allows MicroBooNE to differentiate between electron and photon events. Electron-induced showers generally originate at the

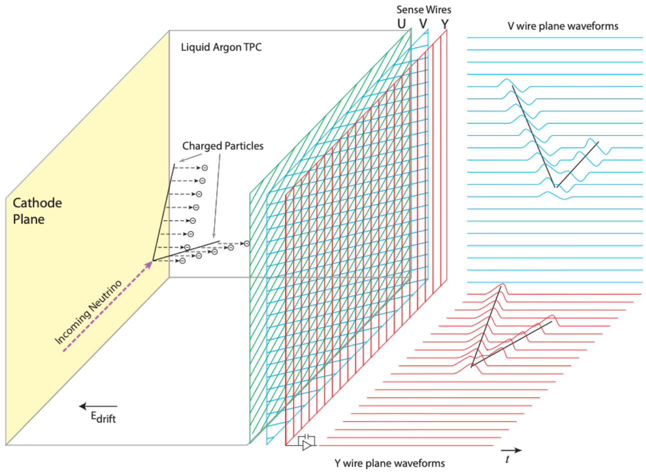


**Figure 1:** Schematic of a liquid-argon time projection chamber. Ionization electrons drift toward the anode wire plane, producing waveforms containing the charge and timing information used to reconstruct particle tracks. Reproduced from Ref. [5].

reconstructed interaction vertex. Having no electric charge, photons can travel through the liquid argon without ionizing it until they interact and create charged particles. Their showers can therefore begin detached from the vertex. Figure 2 displays representative reconstructed electron and photon candidates. This paper is concerned with the photon rather than the electron sector.

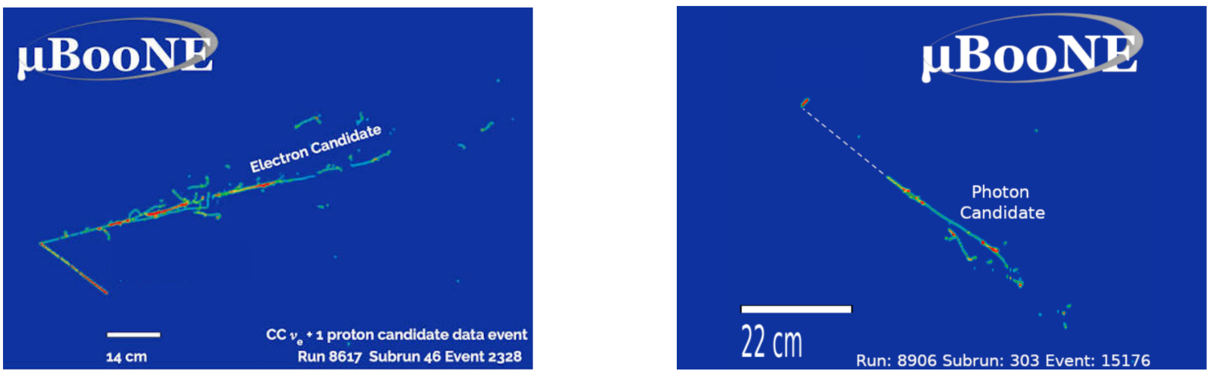


**(a)** Electron-induced EM shower. Adapted from Ref. [6].

**(b)** Photon-induced EM shower. Adapted from Ref. [7].

**Figure 2:** MicroBooNE event-display images illustrating the different origins of EM showers.

## 1.4 Goals of This Work

This project advanced the single-photon search by studying both background and signal events. Through a systematics-centered analysis, the background studies aimed to provide a clearer understanding of two discrepancies in MicroBooNE. The first was a previously reported deficit of neutral-current $\pi^0$ events with at least one reconstructed proton (NC$\pi^0$Np) [8]. The second was an unresolved excess of single-photon events that potentially have a contribution from two-photon events originating outside the fiducial volume (2$\gamma$outFV) [9]. Both studies tested whether cross-section or flux systematic uncertainties could explain the discrepancies through controlled injections into the model. In each case, the goal was to determine whether a systematic variation could alter the target channel without producing correlated distortions in other channels.

The signal study optimized sensitivity to neutral-current coherent single-photon production (NC coherent $1\gamma$). The analysis explored a broad assortment of candidate binning variables, the number and placement of histogram bins, collapsed two-dimensional histograms, and combinations of sidebands to constrain nuisance parameters. A signal-specific boosted decision tree was trained and incorporated into a PROfit template fit. The final configuration was tested with both an Asimov dataset at a projected exposure of $1 \times 10^{21}$ protons on target (POT) and the limited MicroBooNE open dataset at $9.57 \times 10^{19}$ POT. These studies strengthen the MicroBooNE single-photon search by researching backgrounds and increasing sensitivity to a rare single-photon production process.

# 2 Analysis Framework

## 2.1 BAGELS Event Classification

Because a neutrino is not directly detectable, its interaction is reconstructed from the resulting particles. Therefore, these studies rely heavily on an event-selection framework alongside a statistical fitting framework. Broad Analysis of General ELectromagnetic Showers (BAGELS) combines features from other reconstruction frameworks (Wire-Cell, Pandora, gLEE, and LANTERN) to train a multivariate boosted decision tree (BDT) that classifies events into one of twenty topologies. Channel selections are then formed by applying cuts to the BDT probability outputs, with a fixed priority ordering used when an event passes the cut for more than one channel. BDT probability outputs, such as $P_{1\gamma 0p}$, are topological scores, meaning they characterize the reconstructed particles rather than the underlying physics process. Higher scores indicate that the reconstructed particles resemble one photon and no protons. Figure 3 summarizes the topology set.

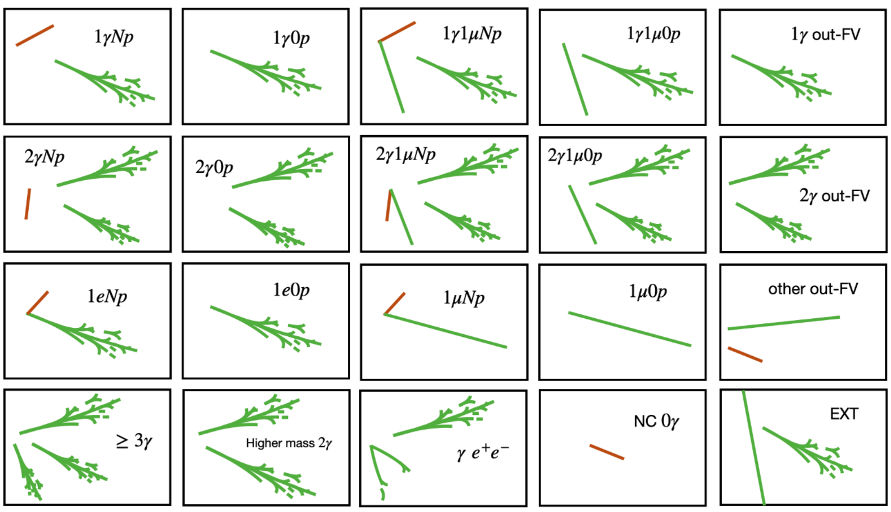


**Figure 3:** The twenty reconstructed event topologies classified by the BAGELS framework.

## 2.2 PROfit Fitting

Once BAGELS assigns events to reconstructed topology channels, PROfit provides the statistical framework used to compare those selected samples with simulation and test physics hypotheses. PROfit is a C++ fitting framework created for neutrino oscillation experiments [10]. Through an XML configuration file, PROfit allows the user to specify the input simulation files, analysis channels, binning edges, model parameters, and systematic uncertainties included in a study. PROfit then produces event spectra, covariance matrices, visualizations of pulled systematic parameters, and best-fit histograms for either real or Asimov data.

During a fit, PROfit varies the relevant physics and systematic parameters to find the prediction that best reproduces the data. A $\chi^2$ statistic quantifies the agreement between the prediction and the data, with a smaller $\chi^2$ indicating better agreement. Systematic uncertainties can change the number of events predicted in each bin and can therefore produce correlated shifts across multiple channels. PROfit uses data from background-dominated sidebands to constrain uncertainties that also affect the signal channel.

Other features of PROfit include controlled systematic injections, which allow the effects of individual systematic parameters to be studied. PROfit also supports template fitting, in which the normalization of a selected prediction component is treated as a free parameter. The framework applies a range of scale factors to this component, calculates $\chi^2$ at each value by minimizing the nuisance parameters, and determines the scale factor that provides the best fit.

# 3 Understanding Single-Photon Backgrounds

## 3.1 Neutral-Current $\pi^0$ Deficit Study

Neutral-current $\pi^0$ production is a significant background in the MicroBooNE single-photon search. These interactions are defined by a neutrino producing a neutral pion without changing flavor. The pion subsequently decays into two photons. This event can resemble a single-photon interaction when one photon leaves the fiducial volume and the other is reconstructed. Moreover, accurately modeling NC$\pi^0$ interactions is necessary for predicting the background in the single-photon channels, $1\gamma 0p$ and $1\gamma$Np.

A previous MicroBooNE analysis observed an approximately 20% deficit of NC$\pi^0$ events containing at least one reconstructed proton (NC$\pi^0$Np), as shown in Figure 4. This study investigated whether the deficit could be reproduced by cross-section or flux systematics. The analysis included four closely related channels: $1\gamma 0p$, $1\gamma$Np, NC$\pi^0 0p$, and NC$\pi^0$Np. In Figure 3, NC$\pi^0 0p$ and NC$\pi^0$Np are labeled $2\gamma 0p$ and $2\gamma$Np, respectively. These labels refer to the same channels. These four channels were selected because they contain one and two photons with and without protons. This combination of closely related channels made it possible to determine whether a systematic uncertainty could reduce the NC$\pi^0$Np prediction without producing similar changes in the other channels. This analysis used the variable `wc_kine_reco_Enu`, Wire-Cell's reconstructed neutrino energy. Each channel's event spectrum was composed of two bins, distinguishing between high and low energies and simplifying interpretation.

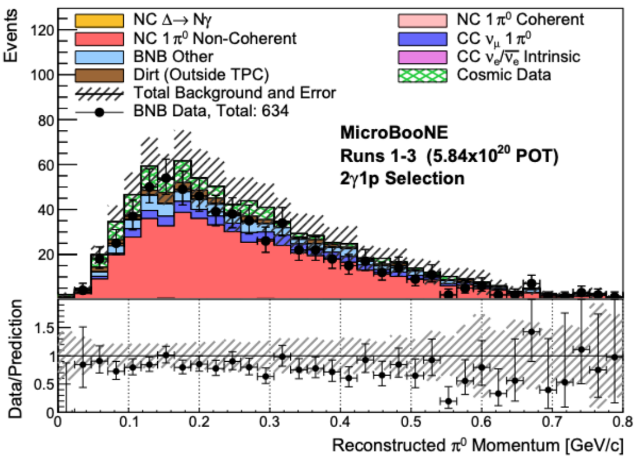


**Figure 4:** Reconstructed $\pi^0$ momentum in the MicroBooNE $2\gamma 1p$ selection. The data lie below the central-value prediction across much of the distribution. Reproduced from Ref. [8].

An isolated 20% deficit was manually injected into an Asimov dataset in the NC$\pi^0$Np channel. PROfit's profile command then varied systematics to find the best fit to this artificial dataset. The most strongly pulled systematics were chosen as candidates for reproducing the deficit. The deficit was then removed, and PROfit's profile command was run again on Asimov data with the leading candidate systematics individually injected at their extremes, $\pm 3\sigma$. For each injection, the fractional change from the central-value prediction was calculated in every bin and channel. Figure 5 visualizes these results in a heatmap.

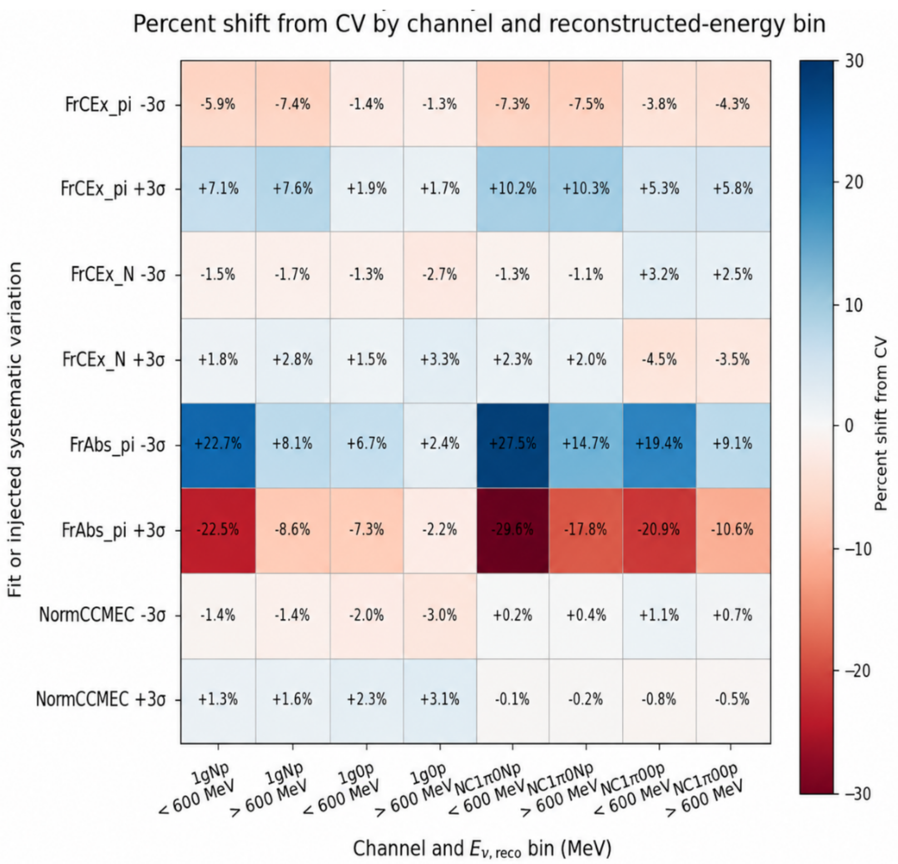


**Figure 5:** Percent shift from the central-value prediction in each reconstructed-energy bin and channel after injecting systematic variations at $\pm 3\sigma$.

The only systematic capable of producing a deficit of the desired magnitude is the cross-section final-state-interaction pion-absorption systematic, $\text{FrAbs}_{\pi}$. However, the effect is not isolated to the $\text{NC}\pi^0\text{Np}$ channel. This study shows that a known systematic uncertainty can imitate the observed $\text{NC}\pi^0\text{Np}$ deficit but cannot naturally reproduce the deficit as a unique change confined to the target channel. Furthermore, this result indicates that the discrepancy is unlikely to arise from a simple isolated shift of one systematic parameter.

## 3.2 Out-of-Fiducial-Volume Excess Study

This section details the second background study. Events that originate outside the detector's fiducial volume (FV) and contain two reconstructed photon showers are classified as $2\gamma$outFV. These events are important to single-photon signal channels because activity entering from outside the fiducial volume can be incompletely reconstructed and contribute to background. As evidenced by Figure 6, a previous MicroBooNE analysis observed an approximately 30% excess of single-photon events near the edge of the detector, which could have a contribution from $2\gamma$outFV events. Similarly to the previous study, this study asked whether known cross-section or flux systematic uncertainties could reproduce the excess without altering predictions in channels for events originating from within the FV.

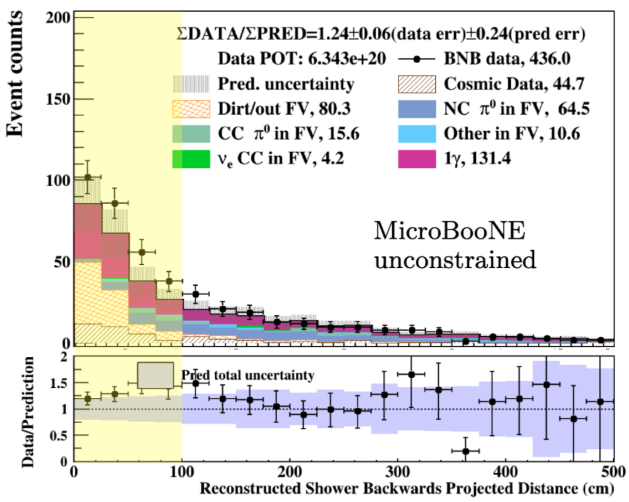


**Figure 6:** Reconstructed shower backwards-projected distance in the MicroBooNE single-photon selection. The region below 100 cm contains an excess relative to the prediction. Adapted from Ref. [9].

Wire-Cell's backwards-projected distance from the shower's starting point to the TPC wall, `wc_reco_backwards_projected_dist`, was used similarly to the analysis shown in Figure 6. Several out-of-fiducial-volume and in-volume photon channels were included so that the response of the target channel could be compared with related event topologies. The horizontal axis of Figure 7 lists all of the channels included in this analysis. Each channel's spectrum was divided into eight bins.

Following the same procedure as the $\text{NC}\pi^0$ deficit study, an Asimov dataset was constructed by manually injecting a 30% excess into the $2\gamma$outFV prediction below 100 cm. The remaining channels were left at their central-value predictions. PROfit's profile command was used to obtain the best fit to this dataset. The most strongly pulled systematics were selected as candidates. The 2$\gamma$outFV excess was removed, and PROfit's profile command was run again on Asimov data with candidate systematics individually injected at $\pm 3\sigma$ to determine how each affected the event rate across all included channels. For each injection, the percent change in the integrated event yield from the central-value prediction to the best fit was calculated for every channel. These results are displayed as a heatmap in Figure 7.

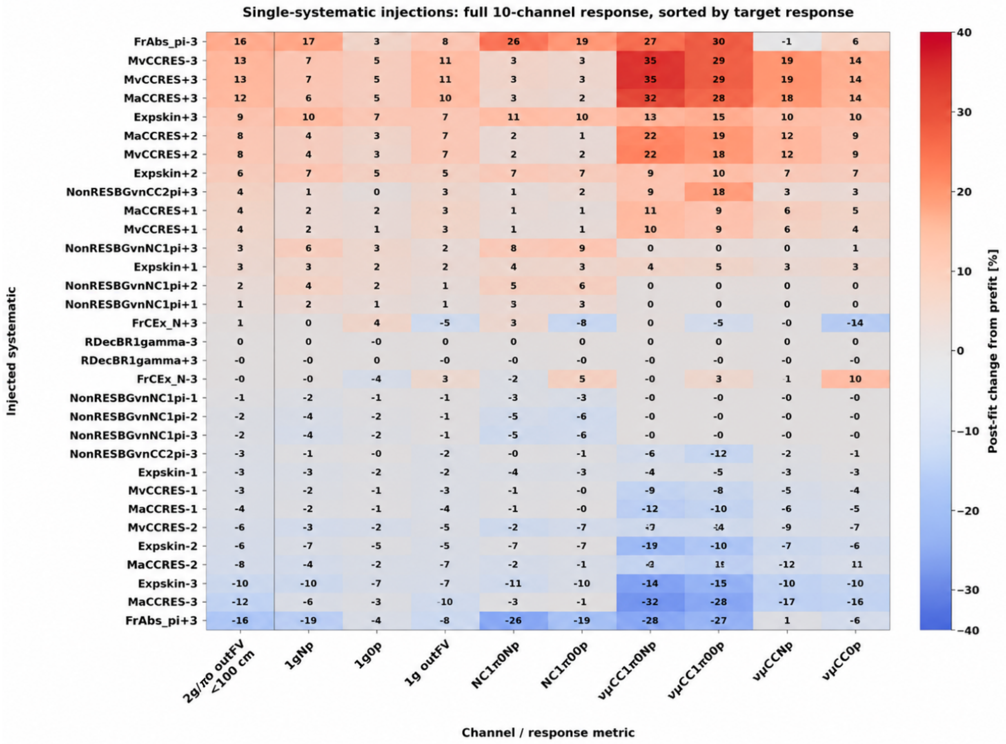


**Figure 7:** Percent change in the post-fit integrated event yield relative to the central-value prediction in each analysis channel after individually injecting the candidate systematics, sorted by the resulting excess in the target channel.

No systematic injection was able to produce the full excess. Variations that increased the 2$\gamma$outFV prediction also altered the photon, neutral-pion, or charged-current channels in the FV. The study found that known systematic uncertainties could contribute to an out-of-FV excess, but no individual systematic effect was isolated to the 2$\gamma$outFV channel. Therefore, none of the tested systematics provide a unique explanation for the observed excess.

# 4 NC Coherent Single-Photon Sensitivity

## 4.1 Signal Definition and Strategy

NC coherent 1$\gamma$ is a rare SM process expected to be a subdominant source of single photons from neutrino-argon interactions. It is characterized by a weak interaction between the neutrino and argon via the $Z$ boson, resulting in a photon and a neutrino with lower energy. Coherent refers to the neutrino interacting with the entire argon nucleus, which remains intact, rather than specific nucleons. Figure 8 illustrates this process.

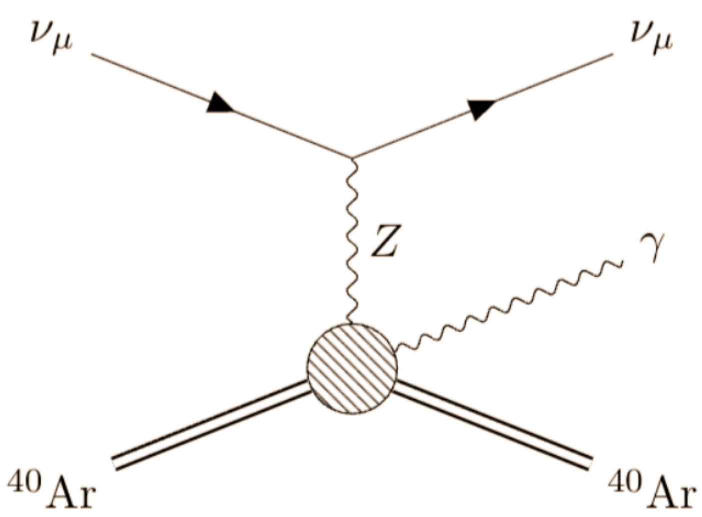


**Figure 8:** Diagram of NC coherent single-photon production on argon.

This theorized process remains to be experimentally observed, although a previous analysis placed constraints on its production rate of photons. Above a factor of 24 times the nominal SM prediction, the rate of NC coherent $1\gamma$ is excluded at a confidence level of 90% [7]. This study aimed to improve MicroBooNE's sensitivity to NC coherent $1\gamma$ production through analysis optimization. The signal channel for this optimization was the BAGELS $1\gamma 0p$ channel, which is defined by events having a probability score for this topology, $P_{1\gamma 0p}$, that exceeds 0.9.

Using PROfit's template fit model, the scale factors on the NC coherent $1\gamma$ event rate were treated as free model parameters, allowing PROfit to compute $\chi^2$ values for a manually determined range of scale factors. Plots of $\Delta\chi^2$ versus the excess scale factor above the nominal prediction were then created for each analysis configuration. The steepness of the curves on these graphs, corresponding to the exclusion power for each configuration, allowed for the comparison between analysis setups. The following sections consider the choice of x-axis binning variable, the number and placement of bins, the use of collapsed 2D histograms, and the selection of the most useful sidebands.

## 4.2 Binning Variable Scan

A total of 957 reconstruction variables from multiple reconstruction frameworks were considered, requiring an efficient method to evaluate their performance without testing each individually. First, a semantic analysis using a regex cascade grouped variables into families. Using the template fit model, PROfit's profile command was then run using spectra binned in a candidate variable from each family. Binning edges were chosen to divide the weighted event distribution equally among eight bins, and the same procedure was applied consistently to each variable. Sensitivity curves were made for each profile run. Each variable's performance was quantified by its $\Delta\chi^2$ value at an excess scale factor above the central-value prediction of 10 ($x = 10$), allowing the variables to be compared with a common metric. Underperforming families were searched for variables with low Spearman rank correlation to the representative variable, indicating that they might contain different information. Weakly correlated variables were tested before eliminating the family from the study. High-performing representative variables prompted further searches within their respective families. In total, 240 profile scans were performed, and glee_trackstub_num_unassociated_hits provided the greatest sensitivity to NC coherent $1\gamma$. glee_trackstub_num_unassociated_hits is a variable counting low-energy isolated depositions that are not included in Pandora track or shower objects. These depositions may be consistent with neutron activity, which would indicate a noncoherent neutrino interaction and help distinguish it from the coherent signal.

BAGELS BDT outputs were then tested as binning variables, and in general, they outperformed individual variables. Being informed by hundreds of features, the BDT outputs are positioned to differentiate the signal better than individual reconstruction variables. The new best-performing variable was the probability score for the same topology as the signal process, $P_{1\gamma 0p}$. Although $P_{1\gamma 0p}$ identifies events that resemble the signal topology, it is not optimized for separating coherent production from other $1\gamma 0p$ processes.

Consequently, a binary classifier dedicated to selecting NC coherent $1\gamma$ events was trained using 1,022 input variables. The classifier is a gradient-boosted decision tree trained with XGBoost [11]. Hyperparameter tuning was performed to improve the original BDT. The eight hyperparameters listed in Table 1 were tuned with Optuna using a tree-structured Parzen estimator [12]. A total of 213 trials were performed, of which 56 ran to completion while the others were pruned. Trial performance was measured by the area under the receiver operating characteristic curve (AUC), the probability that a randomly chosen signal event receives a higher score than a randomly chosen background event. By trial 17, the best AUC was within 0.0002 of the campaign best 0.9961. The lack of significant improvement during later trials indicates that more tuning will likely have modest effects. Figure 9 shows that the majority of unpruned trials have similar performance and plateau within a small range. The figure also demonstrates that more boosting rounds are not likely to substantially improve training, as performance plateaued well before 200 rounds in most completed trials.

**Table 1:** The eight hyperparameters were tuned with Optuna using a tree-structured Parzen estimator.

| Parameter | Range | Scale | Best |
|---|---|---|---|
| max_depth | 2–10 | linear | 7 |
| learning_rate | 0.001–0.3 | log | 0.1439 |
| subsample | 0.5–1.0 | linear | 0.8840 |
| colsample_bytree | 0.5–1.0 | linear | 0.6045 |
| min_child_weight | 1–10 | linear | 4 |
| reg_lambda | 0.01–10 | log | 9.939 |
| reg_alpha | 0.0001–10 | log | 0.0059 |
| scale_pos_weight | 37.0–148.0 | log | 42.07 |

The original BDT and tuned BDT are referred to as BDT v1 and BDT v2, respectively. Figure 10 gives the sensitivity curves of all configurations in this process of determining the best binning variable. Steeper curves correspond to more sensitive configurations, so BDT v2 provides the greatest expected sensitivity among the variables tested.

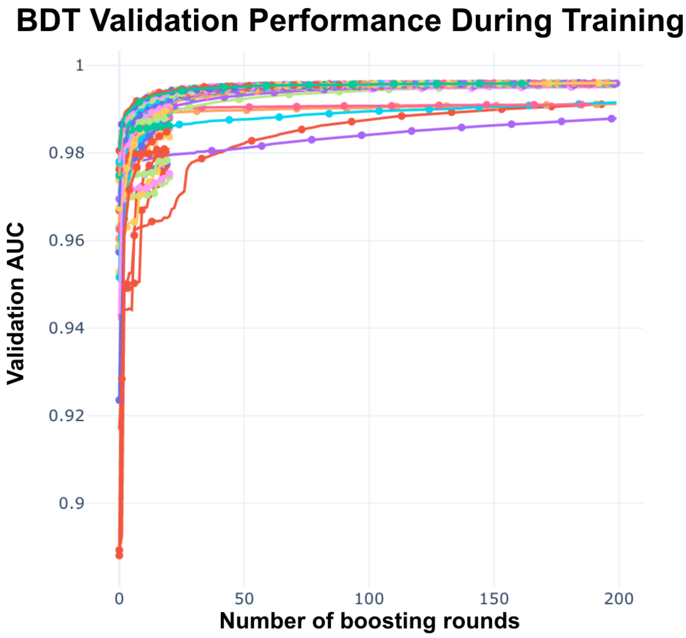


**Figure 9:** BDT validation AUC as a function of the number of boosting rounds.

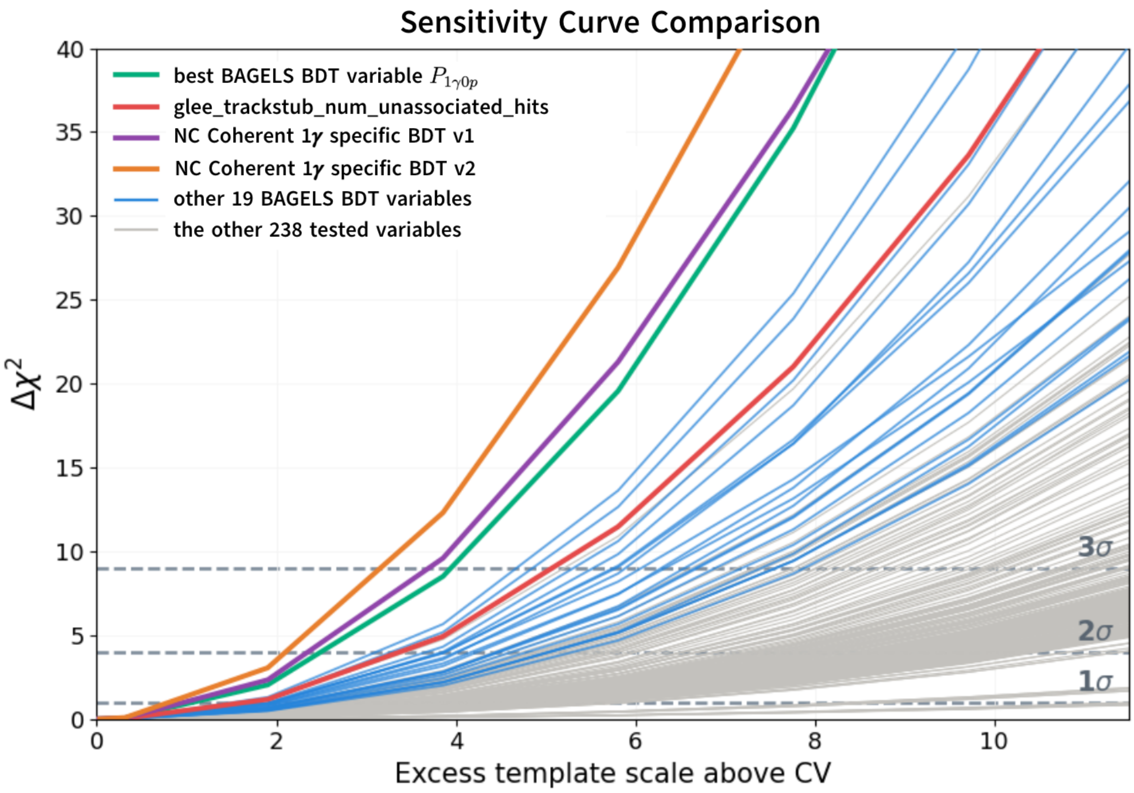


**Figure 10:** Profiles of all the configurations during the process of selecting the optimal binning variable. Interpretation should be limited to relative sensitivity comparisons because these curves were produced using an incorrect $\chi^2$ metric and a since-resolved error in PROfit's treatment of Monte Carlo statistical uncertainties.

### 4.3 Binning Edges and Sideband Optimization

After the winning variable, BDT v2, was chosen, binning edges and sidebands were optimized. Equal-population binning was used for the variable tests, but having different binning edges also affects sensitivity. NumPy was used to divide all events in the $1\gamma 0p$ channel into 2,000 microbins containing an equal number of weighted events, before different density scores were computed for each microbin. These density scores were summed across all microbins, and edges were placed so that each bin contained an equal share of the cumulative density. Using these edges, candidate binning configurations were evaluated by comparing sensitivity curves similar to those in Figure 10. Several density definitions were compared, differing in which sources of uncertainty enter the denominator. The best-performing density function was

$$\frac{s^2}{s+b}, \tag{1}$$

where $s$ and $b$ are the expected signal and background yields in a microbin. Each bin was additionally required to contain at least ten total expected events. This occupancy floor prevents the fit from relying on bins with extremely limited Monte Carlo statistics. Without this floor, the apparent sensitivity increased indefinitely as more bins were implemented. As binning becomes finer, the statistical uncertainty should eventually offset the sensitivity gain from increasing the number of bins, but this threshold was never reached. This effect was an artifact of sparsely populated bins, for which the Gaussian approximation used in the $\chi^2$ calculation becomes unreliable. With the floor in place, two-dimensional binning was also tested but consistently decreased sensitivity to the signal regardless of the second variable paired with the BDT v2 score.

Sideband optimization was performed using ablation studies. Including all 19 BAGELS sidebands in the analysis creates the most sensitive configuration but does little to provide an understanding of the physics constraining the signal. Without any sidebands, the analysis has a $\Delta\chi^2$ value of 53.06 at $x = 10$. With all 19 sidebands, the analysis improves to $\Delta\chi^2 = 65.44$ at $x = 10$. To recover the majority of the total gain in sensitivity from including all sidebands in the analysis, four sidebands were identified as especially important: NC $1\pi^0$Np, NC $1\pi^0 0p$, $\nu_\mu$CCNp, and $\nu_\mu$CC$0p$. An analysis with these sidebands produces a $\Delta\chi^2$ value of 63.46 at $x = 10$. These four sidebands recover 84.0% of the total sensitivity gained from having all sidebands turned on. The NC$\pi^0$ channels constrain the most significant background, $\pi^0$s, while the $\nu_\mu$CC channels constrain flux and cross section systematics on backgrounds because these channels have high statistics.

The optimized configuration contains 11 bins in the BDT v2 score with the four background sidebands described above. At $1 \times 10^{21}$ POT the channel holds 10.8 signal events against 513.8 background events, and all signal fractions below are quoted relative to this selected sample. The bin with the lowest BDT v2 score events contains approximately 75% of the background but only 8.7% of the signal. It is retained because a template fit can use this nearly pure-background region to constrain nuisance parameters in the fit. The highest-score bin contains 3.83 signal events and 7.86 background events. This corresponds to a purity of 32.8% and contains 35.5% of the coherent signal in the channel.

Combined Neyman–Pearson $\chi^2$ was used in the final analysis following the example of the previous MicroBooNE NC coherent $1\gamma$ analysis [7]. For the comparisons that follow, it is important to know that several uncertainties were not incorporated into this analysis because they are not yet available. The treatment of kaon flux uncertainties needs to be updated. Detector systematics were not considered for either the background or signal. Furthermore, the NC coherent $1\gamma$ signal lacks cross-section and flux uncertainties. These omissions reduce uncertainty from the fit, which leads to sensitivities that may be optimistic compared to an analysis in which these uncertainties are included.

## 4.4 Sensitivity Calculation

Table 2 compares the optimized configuration with two baselines at $1 \times 10^{21}$ POT. A single-bin counting analysis in reconstructed neutrino energy reaches $3\sigma$ only at $x = 16.46$. Replacing reconstructed energy with eight quantile bins using the coherent BDT v2 score improves the reach to $x = 4.66$. Optimized binning edges and including the four sidebands further improve it to $x = 3.27$. The full optimization improves the projected $3\sigma$ reach by a factor of approximately five relative to the counting analysis.

**Table 2:** Projected sensitivity at $1 \times 10^{21}$ POT. The excess convention is $x = sf - 1$, where x is scale factor above the central-value, and sf is the rate times the SM prediction. Confidence levels in this table use the two-sided Wilks' approximation.

| Configuration | $\Delta\chi^2(x = 10)$ | $3\sigma$ at $x$ | 90% CL $x$ |
|---|---|---|---|
| Optimized: 11 bins + 4 sidebands | 63.46 | 3.27 | 1.67 |
| Baseline: 8 quantile bins, no sidebands | 35.53 | 4.66 | 2.40 |
| Single bin in $E_\nu^{\text{reco}}$ | 3.03 | 16.46 | 9.49 |

The projected $3\sigma$ result $x = 3.27$ corresponds to a total signal rate of $sf = 4.27$ times the SM prediction. Interpreted as an expected upper limit, the optimized Asimov curve gives $x < 1.67$ at 90% confidence using the two-sided Wilks convention.

## 4.5 Validation and Open-Data Result

The optimized configuration was also applied to the $9.57 \times 10^{19}$ POT MicroBooNE open dataset. However, a single bin was used for this run, and a cut made using the BDT v2 variable was chosen so that there would be 10.0 predicted events in the bin. Of the 10.0 predicted events, 0.90 were coherent signal events. With the open dataset, only five data events were observed. The best fit therefore occurred at the physical boundary $sf = 0$, corresponding to no coherent production and no excess being observed.

Figure 11 compares the previously published result with the observed open-data limit and the predicted sensitivity of the optimized analysis configuration.

The open-data result is a validation limit rather than a measurement at the full MicroBooNE exposure. The open-data run excludes scale factors greater than 3.1 at a 90% confidence level. Both the projected and observed limits remain optimistic because signal cross-section and flux systematics as well as detector systematics for both signal and background are absent. Generating these signal variations is crucial for a direct comparison with the old result. Still, the purity increase from this new analysis is substantial and is expected to preserve an improvement in sensitivity over the previous result once all uncertainties are implemented.

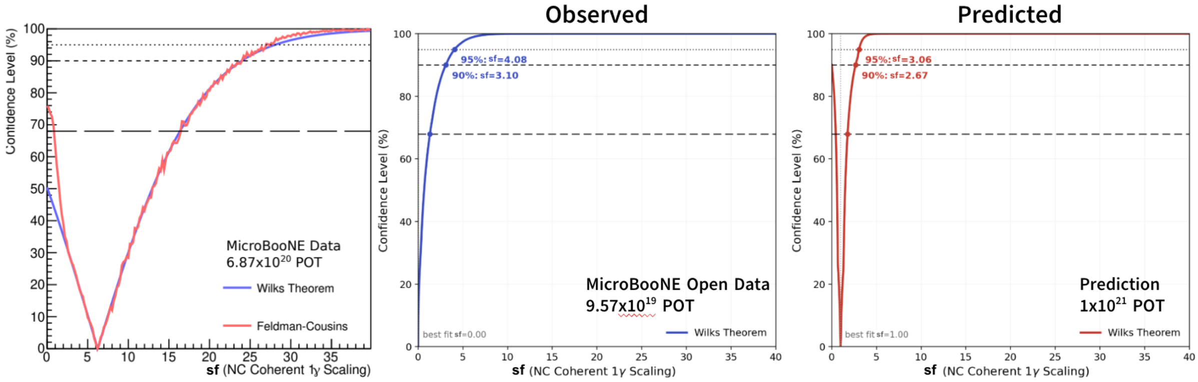


**Figure 11:** Comparison of the previously published MicroBooNE result, the observed result using the MicroBooNE open data, and the predicted sensitivity at $1 \times 10^{21}$ POT. The horizontal axis gives the scale factor, $sf$, applied to the NC coherent $1\gamma$ prediction.

# 5 Summary and Conclusions

This work contributed to the MicroBooNE single-photon search by studying backgrounds and improving signal sensitivity for NC coherent $1\gamma$ events. The NC$\pi^0$Np study found that a cross-section final-state-interaction pion-absorption systematic can reproduce a deficit of comparable magnitude to one previously observed in this channel. This systematic explanation is limited, though, due to the correlated effects it has across different channels. The $2\gamma$outFV study reached a similar conclusion: several nuisance parameter injections partially enhanced the target channel, but none recovered the full 30% excess, and none was unique to the outFV channel.

For the signal study, a dedicated XGBoost classifier was trained to distinguish NC coherent $1\gamma$ production from the backgrounds surviving the $1\gamma0p$ selection. After determining the best binning algorithm, maximizing the number of bins that can maintain 10 events per bin, and including the most constraining sidebands (NC$1\pi^0$Np, NC$1\pi^0 0p$, $\nu_\mu$CCNp, and $\nu_\mu$CC$0p$), the analysis was run with Asimov data for $1 \times 10^{21}$ POT. It excludes a scale factor greater than 2.67 times the SM prediction at a confidence level of 90%. Compared to the $6.87 \times 10^{20}$ MicroBooNE measurement, this analysis is projected to supply approximately a ninefold improvement in sensitivity. A single-bin analysis with MicroBooNE open data was also performed, producing an exclusion comparable to the $1 \times 10^{21}$ POT prediction, albeit with much fewer statistics. While there are limitations regarding the completeness of systematics in these analyses, there is substantial improvement in purity over previous configurations. This improved analysis is a testament to the selection power of BAGELS and the signal-specific BDT v2.

There are a few foreseeable steps that should be taken to further understand and optimize this analysis. To evaluate the validity of this configuration, signal cross-section and flux systematics, as well as detector systematics for signal and background, should be included. These uncertainties will likely reduce the presented configuration's sensitivity to NC coherent $1\gamma$. There is also the possibility that imposing a cut on events with high blip energies could improve this analysis's distinguishing power. Blips are tiny clusters of energy, and their information is not included in the BDTs because they may be mismodeled. Nonetheless, the variable `blip_no_shower_cone_sumE` may have untapped distinguishing power (perhaps other blip variables as well), with many of the NC coherent $1\gamma$ events clustering at lower energies. `blip_no_shower_cone_sumE` is the sum of the blip energies that lie outside of the reconstructed EM shower cones in an event. Events with low total blip energies could indicate a lack of hadronic activity, which could help distinguish coherent signal from noncoherent events. Finally, once the analysis is better understood and optimized, it should be applied to the full MicroBooNE dataset to determine the resulting best-fit scale factor and constraints.

# 6 Acknowledgements

I would like to sincerely thank Dr. Mark Ross-Lonergan and Dr. Lee Hagaman for their supreme guidance, instruction, and kindness throughout this project. Their availability and thoughtfulness greatly improved my experience. I would also like to express extreme gratitude to Dr. Georgia Karagiorgi and Dr. Reshmi Mukherjee for their coordination and leadership of an incredibly rewarding summer program. Amy Garwood and Asia Latt also deserve recognition for their care and helpfulness. Finally, I am grateful to everyone at Nevis, all of whom made it such a welcoming environment for the summer.

This material is based upon work supported by the National Science Foundation under Grant No. PHY-2349438.

# References


[1] Pedro A. N. Machado. "Neutrino Properties and Interactions". In: *Neutrino Physics and Astrophysics*. Ed. by F. W. Stecker. World Scientific, 2023. Chap. 2. DOI: 10.1142/9789811282645_0002. arXiv: 2206.13449 [hep-ph].

[2] Y. Fukuda et al. "Evidence for Oscillation of Atmospheric Neutrinos". In: *Physical Review Letters* 81 (1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003.

[3] S. Navas et al. "Review of Particle Physics". In: *Physical Review D* 110.3 (2024), p. 030001. DOI: 10.1103/PhysRevD.110.030001.

[4] A. A. Aguilar-Arevalo et al. "Updated MiniBooNE Neutrino Oscillation Results with Increased Data and New Background Studies". In: *Physical Review D* 103.5 (2021), p. 052002. DOI: 10.1103/PhysRevD.103.052002. arXiv: 2006.16883 [hep-ex].

[5] R. Acciarri et al. "Design and Construction of the MicroBooNE Detector". In: *Journal of Instrumentation* 12.02 (2017), P02017. DOI: 10.1088/1748-0221/12/02/P02017. arXiv: 1612.05824 [physics.ins-det].

[6] Wanwei Wu. *Latest Results from the MicroBooNE Experiment*. Presentation at the Lake Louise Winter Institute. On behalf of the MicroBooNE Collaboration, FERMILAB-SLIDES-24-0019-V. Feb. 2024. URL: https://lss.fnal.gov/archive/2024/slides/fermilab-slides-24-0019-v.pdf.

[7] P. Abratenko et al. "First Search for Neutral Current Coherent Single-Photon Production in MicroBooNE". In: (2025). arXiv: 2502.06091 [hep-ex].

[8] P. Abratenko et al. "Measurement of Neutral Current Single $\pi^0$ Production on Argon with the MicroBooNE Detector". In: *Physical Review D* 107.1 (2023), p. 012004. DOI: 10.1103/PhysRevD.107.012004. arXiv: 2205.07943 [hep-ex].

[9] P. Abratenko et al. "Inclusive Search for Anomalous Single-Photon Production in MicroBooNE". In: *Physical Review Letters* 136.18 (2026), p. 181806. DOI: 10.1103/89qs-41cp. arXiv: 2502.06064 [hep-ex].

[10] Mark Ross-Lonergan and PROfit Contributors. *PROfit: A Neutrino-Oscillation and Beyond-the-Standard-Model Fitting Framework*. https://github.com/markrosslonergan/PROfit. GitHub repository, accessed August 1, 2026. 2026.

[11] Tianqi Chen and Carlos Guestrin. "XGBoost: A Scalable Tree Boosting System". In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2016, pp. 785–794. DOI: 10.1145/2939672.2939785. arXiv: 1603.02754.

[12] Takuya Akiba et al. "Optuna: A Next-generation Hyperparameter Optimization Framework". In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2019, pp. 2623–2631. DOI: 10.1145/3292500.3330701. arXiv: 1907.10902.