Machine learning qualitatively changes the search for new particles
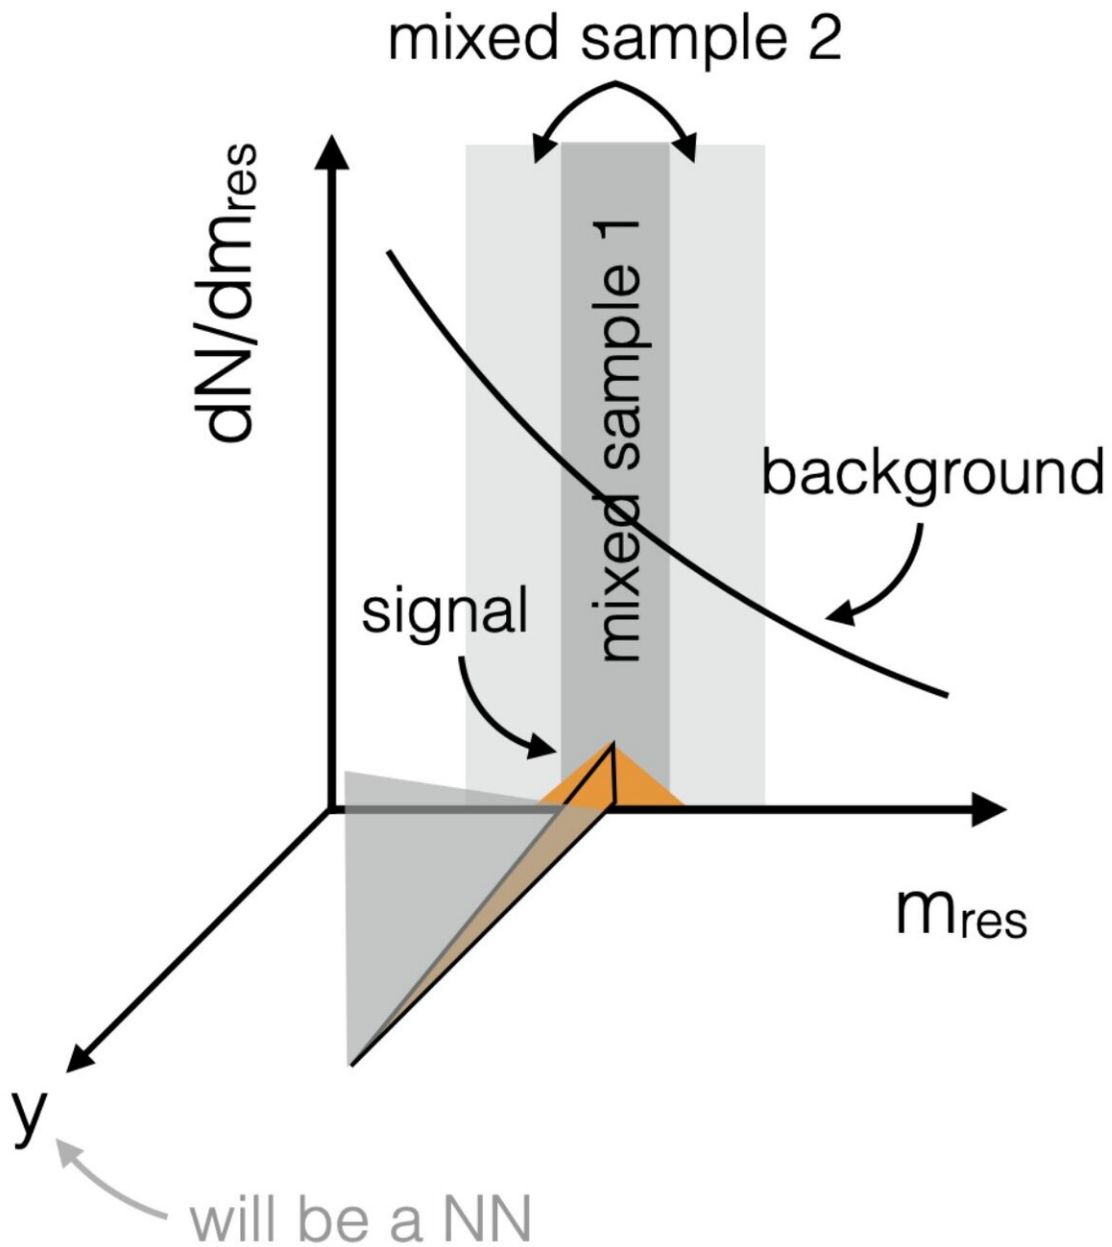
June 17 2020

Figure 1: Diagram illustrating the construction of mixed samples for training a weakly supervised CWoLa classifier in the bump hunt. In the ATLAS search, the resonant feature (mres) is the dijet mass and the other features (y) are the masses of the two jets. Credit: ATLAS Collaboration/CERN

The ATLAS Collaboration at CERN is exploring novel ways to search for new phenomena. Alongside an extensive research program often inspired by specific theoretical models—ranging from quantum black holes to supersymmetry—physicists are applying new model-independent methods to broaden their searches. ATLAS has just released the first model-independent search for new particles using a novel technique called "weak supervision."

Searches for new particles typically start with a specific theoretical model. Given the model's phenomenology and parameters, physicists will simulate how new particles would be produced and decay in the ATLAS detector. They then simulate the Standard Model background processes in order to develop classifiers (with or without machine learning) that separate signals from background. These classifiers determine the best phase-space region of the data to be studied, where a hypothetical signal is expected to be enriched. Finally, physicists will compare the data and background prediction in search of anomalies.

ATLAS' new search uses machine-learning classifiers (neural networks) developed directly on data in order to reduce their dependence on a specific model. This is a significant departure from the standard methods because the data are unlabelled: it is not known if a particular proton–proton collision event is background or signal. This method—known as "weak supervision"—exploits structures in the data without needing per-event labels.

Alongside with this method, the new ATLAS search uses one of the most traditional simulation-independent anomaly detection strategies: the "bump hunt." The goal of a bump hunt is to look for a localized "bump" on top of a smooth background. Such bumps are a generic feature of many models of new particles, where the bump happens at the mass of the new particle. The new search builds on this strong foundation to enhance the sensitivity to a wide variety of hypothetical particles without

specifying their properties ahead of time.

The combination of bump hunt and weak supervision results in an analysis that is mostly free of signal-model and background-model dependence.
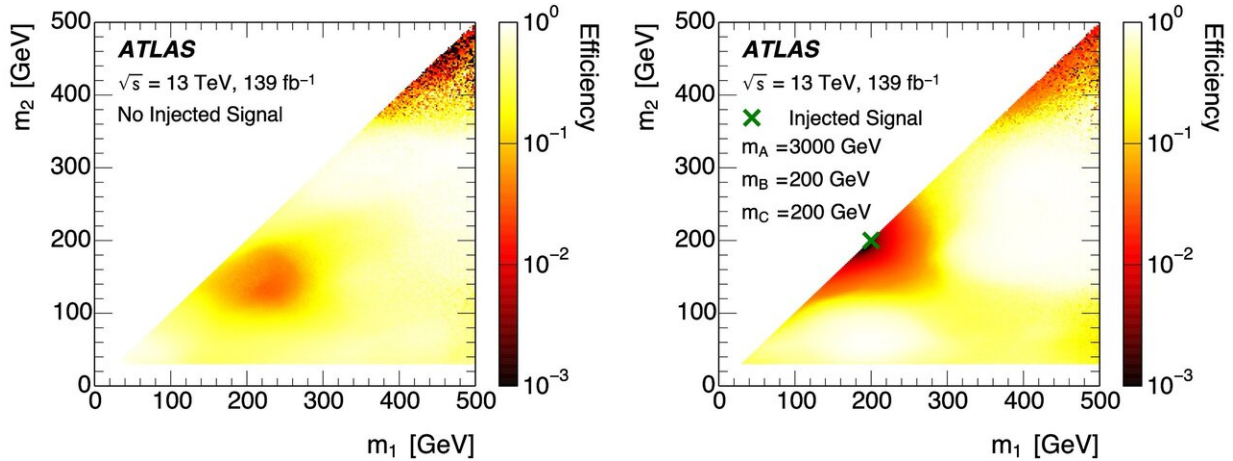

Figure 2: The neural network output in one dijet mass bin. As a two-dimensional function, the output can be readily visualised as an image, where the intensity corresponds to the efficiency of the network output in the dijet mass bin. The left plot has no signal injected and the right plot shows the output when a hypothetical particle at 3 TeV that decays into two other particles at 200 GeV is added to the data. Credit: ATLAS Collaboration/CERN

Detecting anomalies with weak supervision

ATLAS physicists trained neural networks on data using a technique called "Classification without labels" (CWoLa, pronounced "Koala"). In this approach, physicists construct two mixed datasets composed of background and potentially also signal. These are identical except for the relative proportions of the potential signal. While the signal-vs-

background labels are unknown for each event, the neural networks can be trained to differentiate between the two datasets. With sufficient data and a powerful enough classifier, this is actually optimal for distinguishing signal from background.

The CWoLa method is combined with a bump hunt when creating the mixed datasets above, as shown in Figure 1. Signal events would be characterized by a localized resonance region and a sideband region. These regions would have other features (y) that can also be used to train the neural networks. If there is no signal, a neural network would not learn anything and if there is a signal, it may learn to pick it out over the background.

The new ATLAS search is the first application of fully data-driven machine-learning-enhanced anomaly detection. The search examined events with hadronic final states, using the invariant mass of pairs of particle "jets" as the resonant feature and the masses of the individual jets as the features to train the CWoLa classifier. Using this restricted set of features, physicists have successfully established the procedure and have found it is already sensitive to a wide range of new particles.

Physicists were able to train the neural networks while avoiding a statistical trials factor which would reduce the sensitivity of the search from training and testing on the same data. The neural network (Figure 2) is mapped to an efficiency. For example, 10% means that 90% of events have a network output that is lower than this value. In the absence of the signal, the network should not learn anything (as the two mixed datasets should be the same), but there must be a region of low efficiency by design. The right plot of Figure 2 shows that the network is able to identify the injected signal, even though it was not told where to look in advance!
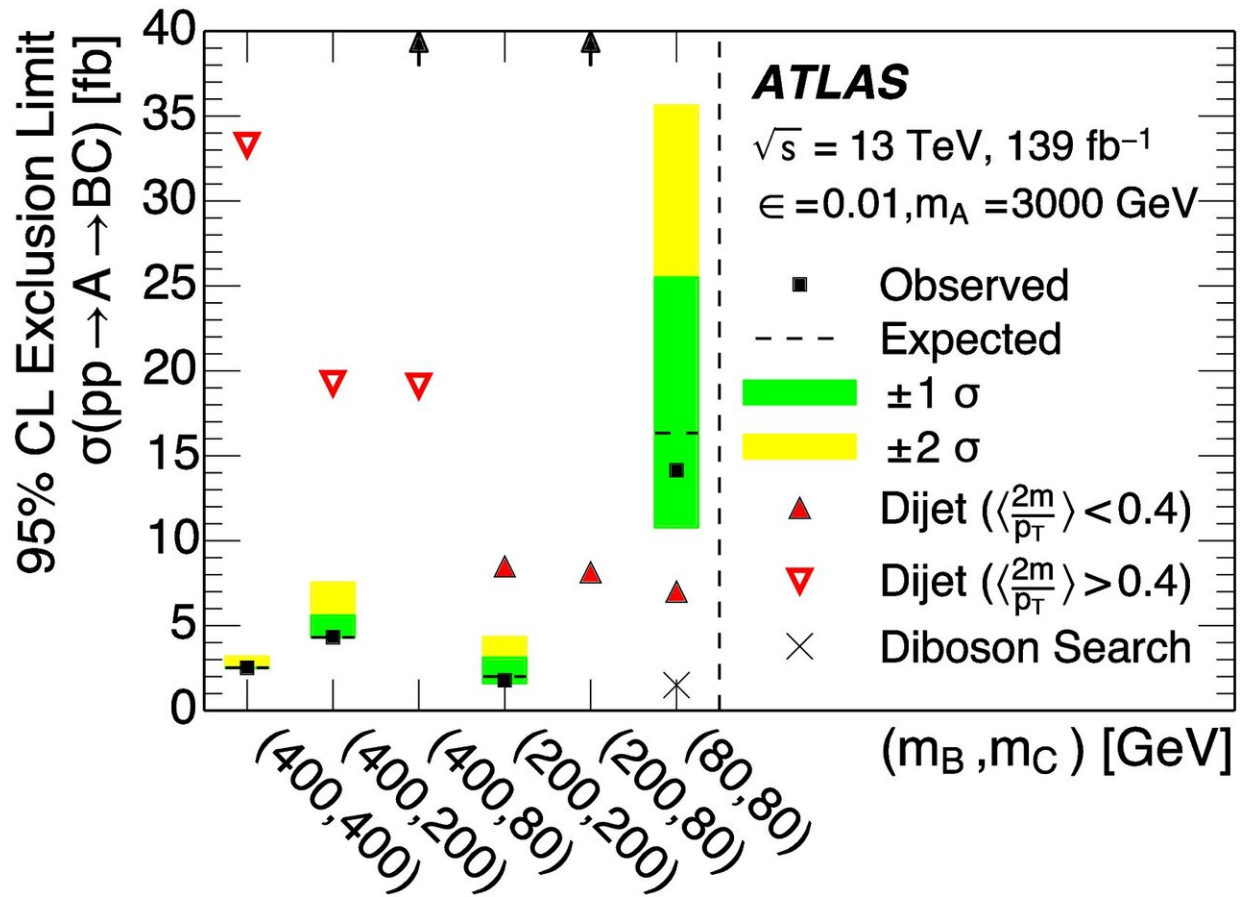
Figure 3: Particular signals are simulated and then added to the data in order to set limits. The models chosen here represent a heavy particle A (with a mass of 3 TeV) decaying to two other new particles B and C with masses written on the horizontal axis. The vertical axis is the limit - lower numbers indicate stronger limits. The new search is compared with two existing results from ATLAS: the inclusive dijet search (red triangles) and a dedicated search for jets produced from W and Z bosons (grey cross). Credit: ATLAS Collaboration/CERN

Providing new precision

The new search did not result in significant evidence for new particles and quantifying what was not found was its own challenge. Usually, physicists can simply ask how much signal would have to be added to

register a significant excess, and then that amount of signal is declared excluded as no excess was observed. Achieving similar exclusions for this analysis required all of the neural networks to be re-trained for each modeled signal type and signal amount.

The resulting limits are presented in Figure 3. Producing this plot required training about 20,000 neural networks! Some signals were harder for the neural networks to find than others, with those in regions with a lot of background proving particularly challenging. For other signals, the new limits are stronger than previous limits and improve upon previous searches in a similar phase space.

Looking to the future

This new approach taken by ATLAS has many possibilities for extensions. The weakly supervised bump hunt could be applied to additional event topologies and more features could be added to broaden the sensitivity to new particles. More complex neural networks may be needed to accommodate higher-dimensional feature spaces and this will require demanding computational resources. ATLAS physicists are also considering a variety of alternative anomaly-detection techniques, which may be able to complement the CWoLa-based search. It is likely that no one method will cover everything—multiple approaches will be needed to ensure broad, robust, and strong sensitivity to new particles.

More information: Dijet resonance search with weak supervision using 13 TeV proton-proton collisions in the ATLAS detector, arXiv:2005.02983 [hep-ex]: arxiv.org/abs/2005.02983

Provided by ATLAS Experiment