

RECEIVED: November 17, 2014

REVISED: January 13, 2015

ACCEPTED: January 27, 2015

PUBLISHED: February 18, 2015

Jet-images: computer vision inspired techniques for jet tagging

Josh Cogan, Michael Kagan, Emanuel Strauss and Ariel Schwartzman

*SLAC National Accelerator Laboratory,
Menlo Park, CA 94028, U.S.A.*

E-mail: joshgc@slac.stanford.edu, makagan@slac.stanford.edu,
estrauss@slac.stanford.edu, sch@slac.stanford.edu

ABSTRACT: We introduce a novel approach to jet tagging and classification through the use of techniques inspired by computer vision. Drawing parallels to the problem of facial recognition in images, we define a *jet-image* using calorimeter towers as the elements of the image and establish jet-image preprocessing methods. For the jet-image processing step, we develop a discriminant for classifying the jet-images derived using Fisher discriminant analysis. The effectiveness of the technique is shown within the context of identifying boosted hadronic W boson decays with respect to a background of quark- and gluon-initiated jets. Using Monte Carlo simulation, we demonstrate that the performance of this technique introduces additional discriminating power over other substructure approaches, and gives significant insight into the internal structure of jets.

KEYWORDS: Jets, Hadronic Colliders

ARXIV EPRINT: [1407.5675](https://arxiv.org/abs/1407.5675)

Contents

1	Introduction	1
2	Algorithm	2
2.1	Jets as images	2
2.2	Jet-image preprocessing	3
2.3	Jet-image processing: constructing the discriminant	6
3	Samples	7
4	Case study: W boson jet tagging	8
5	Conclusion	11
A	Fisher jet images	12

1 Introduction

Modern TeV scale colliders, such as the LHC [1], are capable of producing highly boosted jets from the decay of heavy particles. The large momentum and mass of the heavy particle may lead to a boosted system in which the heavy particle’s decay products merge into a single jet. In order to discover and study such systems, it is vital to discriminate boosted heavy particle jets from those of QCD initiated quark/gluon jets (which we will refer to as QCD jets for the remainder of the text). The techniques of jet substructure, which exploit the internal structure of the jet, have been developed for this task. Contemporary experimental and theoretical work in the field of jet substructure includes, but is not limited to, subjet finding and jet tagging / jet flavor identification algorithms [2–12], jet grooming and pileup mitigation [13–15], template and matrix element techniques [16, 17], energy flow and event shapes [18–24], and probabilistic jet clustering [25, 26]; more inclusive summaries can be found in references [27–30].

Typical approaches to jet flavor identification and jet tagging begin with a simplified analytic expression of the system’s behavior, such as the two-pronged nature of the hadronic W boson decay, or the large particle multiplicity of gluon initiated jets relative to quark initiated jets. A variable which capitalizes on this analytic knowledge and is ideally not subject to large theoretical uncertainties is derived. In contrast, our method does not rely on an analytic description of the system, such as colorflow information, but instead uses a catalog of jets produced by Monte Carlo simulation as a training sample to learn the discriminating features between several categories of jets. The method begins by putting jets through a series of preprocessing steps to build a consistent representation of the jets.

After preprocessing, we employ a fast, linear, and powerful method for feature extraction and physical interpretation.¹

Our methods for preprocessing and discrimination are inspired by techniques in the field of computer vision. Most notably, our task is similar to that of facial recognition, and we have developed jet-specific analogs to the algorithms used in facial recognition. As such, the work presented here serves to both build a connection between the fields of jet substructure and computer vision, and as an introduction to the use of computer vision techniques in the analysis of hadronic final states. The basic algorithm presented here is thus serving as a general proof of principle, and can be more deeply developed in future studies.

In section 2 we present the “jet-image” algorithm, which both describes the definition of the jet-image, the jet-image preprocessing steps, and the jet-image processing (that is, the methods of discrimination). Section 3 describes the samples used for our performance studies of the algorithm, and section 4 presents the case study where the algorithm is employed for discriminating hadronically decaying W jets from QCD jets and is compared to other jet substructure techniques. In section 5 we discuss the information that can be gained from our approach, and further possible application and developments.

2 Algorithm

All operations described bellow are performed in Python (v2.7.2) using the NUMPY [31] and SCIPY [32] libraries. In addition, many of the classifiers are built using SCIKIT-LEARN [33], and the figures are made using MATPLOTLIB [34] and ROOT [35].

2.1 Jets as images

For this study, we restrict the inputs to jet reconstruction strictly to calorimeter towers and, with little loss in generality, approximate the calorimeter as a single layered grid of towers² with spacing $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ spanning $[-2.5, 2.5]$ ³ in η and $[0, 2\pi]$ in ϕ . Using common jet-finding algorithms [38–42], we identify the angular location of the jets (as the jet axis) and save all towers within a $2R$ by $2R$ square, centered on that location for further examination. This yields the $(2\frac{R}{0.1} + 1) \times (2\frac{R}{0.1} + 1)$ grid⁴ of towers which we denote the *jet-image*, an example of which can be seen in figure 1a. Henceforth, we will refer to calorimeter towers which are included in a jet-image as pixels. It is important to note that with this jet-image representation the similarity of two jets can be established using the dot product between two jets. That is, an $N \times N$ pixel jet-image A can be recast

¹Sophisticated non-linear multivariate techniques are also an option but often suffer from long training and testing times, and provide less physical intuition to the underlying discriminating features

²As practically all the particles in a jet deposit their energy in the calorimeters, the idealized calorimeter object sums and stores the energy of all the particles entering a given tower, excluding neutrinos and muons.

³This η coverage was chosen as it roughly corresponds to the region which is most commonly used for physics analysis by the ATLAS [36] and CMS [37] experiments as it is covered by both calorimeter and tracking detectors.

⁴The additional tower ensures that if the jet axis is at the center of a tower, all calorimeter towers within a radius of R are contained in the grid

into a N^2 dimensional vector \bar{A} , where the elements $A_{i,j} = \bar{A}_{i+N \times (j-1)}$ with $i, j \in \{1 \dots N\}$, and the dot product between jet-images A and B is

$$A \cdot B = \sum_{k=1}^{N^2} \bar{A}_k \bar{B}_k \quad (2.1)$$

While this study focuses on jet-images built from calorimeter towers, any inputs which can be mapped to a pixelated image could be used to build jet images. For instance, jet-images could be built from tracks, calorimeter clusters, truth particles, etc., all of which could in principle use much finer pixel granularity due to the higher resolution of the inputs. Once the jet-images are built from the inputs, the algorithms described in the proceeding text are directly applicable.

Jet-images have several properties that make them useful inputs to jet flavor tagging techniques. First, each image has the same dimensionality — every jet can be described by a fixed set of numbers. This is an important property for the application of image processing techniques and for training classification algorithms. Second, as a low-level jet description, the jet-image uses all available information for later discriminatory techniques rather than compressing the information into a set of derived variables. Currently, the information used per pixel is the total deposited transverse energy but could be extended to include further information, such as the longitudinal segmentation of energy deposits in the calorimeter. However, we leave these extension for future studies. Third, the similarity between two jets can be easily and quickly computed using standard linear algebra computations. While the concept of analyzing jets as images is not new (see for instance [43–46]), this paper presents a computer vision inspired framework for processing and interpretation of jet-images.

Having re-cast the representation of the calorimeter information, we introduce the jet-image preprocessing steps and the discrimination techniques used for flavor identification.

2.2 Jet-image preprocessing

Approaches to facial recognition (see [47] and references therein for review) attempt to classify a face in an image by learning the structure of the face, i.e. learning the expected distribution of pixel intensities from a set of example images. The ability to accurately recognize a face in an image can be sensitive to lighting conditions, intensity, shadows, the orientation of the faces, facial expressions, noise, etc. These features introduce large variations in the pixel intensities but are not correlated with different categories and thus obfuscate the underlying discriminatory distributions. Image preprocessing steps are used to standardize the images and greatly improve later classification performance.

Analogously, our approach to jet tagging attempts to learn the structure of the transverse energy distribution of the pixels of the jet-image. The use of preprocessing is paramount to accurate jet-image classification, and as such we have developed a set of preprocessing steps specifically for this context. Many of these steps have parallels in facial recognition, which should become evident from the description. The jet-image preprocessing proceeds as follows:

1. **Noise Reduction:** The effects of noise in the jet-image from pileup are reduced using trimming [13]. Trimming is used due to its simplicity, the fact that it uses subjects which are used a later stages in the algorithm, and because it is seen to greatly reduce the impact of pileup on the algorithm performance. Different jet grooming techniques (see [29, 42] and references therein) which mitigate the effects of pileup could also be used.
2. **Point of Interest Finding:** The positions of the leading regions of transverse energy deposition are located using the subjects resulting from the trimming procedure. Subjects are ordered by p_T to decide the relevance of the points of interest. The number of subjects required depends on the alignment procedure (step 3), but at least one subject must be identified. This step, in the case of a 2-prong decay such as a W-jet, is similar to locating the eyes in an image of a face.
3. **Alignment:** The jet-image is aligned such that the relative pixel location of primary features are always the same. This step exploits symmetries of the $\eta - \phi$ space and removes variations in jets which do not improve classification accuracy. This is analogous to aligning the eyes of a face within an image to always lie in the same pixels. There are three components to jet-image alignment:
 - *Rotation:* Rotation is performed to remove the stochastic nature of the decay angle relative to the $\eta - \phi$ coordinate system. This alignment can be done very generally, by determining the principle axis [48] of the original image and rotating the image around the jet-energy centroid such that the principle axis is always vertical. Alternatively, process specific information can be used. For two-body decay processes (such as the hadronic decay of a W boson) the direction connecting the axes of the leading two subjects can be rotated until the leading subject is directly above the subleading subject. However using process specific information can lead to acceptance loss, for example when the two decay products merge into a single subject.
 - *Translation:* Once rotated, the jet-image is translated such that the jet-image energy centroid or leading subject is always centered in the same pixel. This procedure is paramount to adjusting the position of the eyes in a picture, anchored so as to always appear in the same pixels.
 - *Reflection:* Once translated, the image is reflected over the vertical axis such that the side of the image with maximal transverse energy always appears on the right side of the image. This ensures that the hardest radiation always appears in similar locations and can be exploited by the training procedure for classification.
4. **Equalization:** Jet-images are normalized such that the dot product of a jet-image with itself (i.e. the sum of squared pixel contents) is equal to one. This removes the absolute energy scale dependence of jet-images, thus allowing for comparisons of

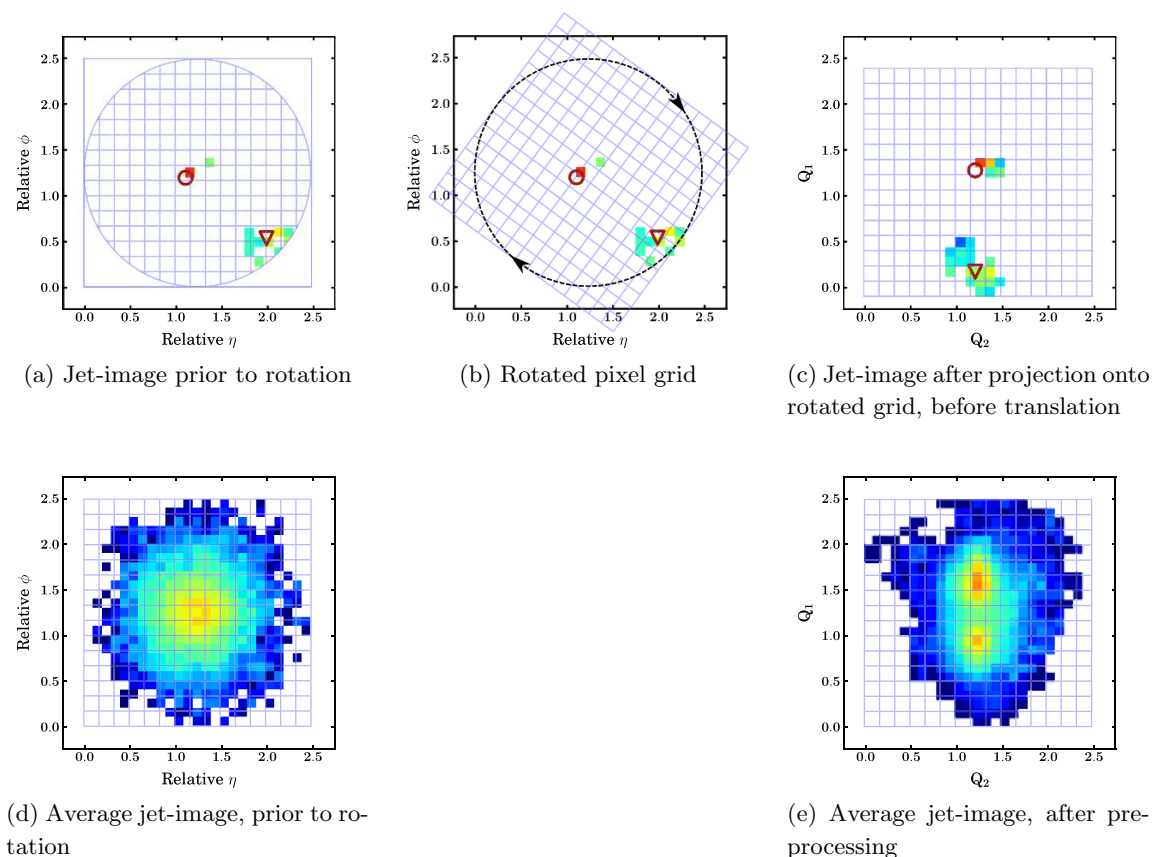


Figure 1: The preprocessing of jet-images and the impact on the average jet-image for W jets in which the leading jet with p_T between 200 and 250 GeV. Note that the grid in figure 1c appears shifted down to represent the jet-image before translation, which is subsequently translated such that the leading subjet lies in the location ($Q_1 \sim 1.5, Q_2 \sim 1.25$) as see in the final average jet-image of figure 1e.

jet-images with different energies. This step is analogous to the standardizing the lighting conditions of images.

5. **Binning:** In many cases, the expected jet-images may vary significantly with a known variable; in this case, the variable can divide a class of jet-images into a set of sub-classes with more uniform jet-images. For instance, if the total transverse energy of the jet-image or the ΔR between the subjets causes significant variations, jet images can be binned into different ranges of the variable. This is analogous to separating images based on the facial expression. A different discriminant can then be trained separately for each sub-class.

An example of the image preprocessing with jets from hadronically decaying W bosons can be seen in figure 1 plotted using the η and ϕ coordinates of the pixels relative to the jet axis (before rotation) and using the rotated coordinate system Q_1 and Q_2 after

rotation. The average of a large sample of W jet-images without preprocessing can be seen in figure 1d, where there is no clear sign of the two-prong decay structure of the W jet. Figures 1a, 1b, and 1c illustrate the different stages of preprocessing on a single W jet-image. Finally, figure 1e shows the average of the same large sample of W jet-images after preprocessing with the two-prong structure of the W decay clearly visible.

This preprocessing prescription is generic enough to be applied to most jet flavor identification problems. However, the impact of the preprocessing steps on the system under consideration must be carefully considered and in some cases tailored. For example, while trimming, or more generally jet grooming, is important for LHC experiments which expect large amounts of pileup, it could be dropped for use in a cleaner experimental environment. Similarly, the rotation steps can be adjusted depending on the system under study (i.e. two-prong decays or three-prong decays).

2.3 Jet-image processing: constructing the discriminant

Having preprocessed the jet images, we have a representation which can be thought of as a vector specifying a coordinate in a high-dimensional space. Linear classification of the jets is performed by projecting the vector onto the discriminate direction which ideally maximizes the separation between different classes of jets.

A number of linear classifiers can be used for discrimination. In facial recognition, Principle Components Analysis (PCA) [49] and Fisher's Linear Discriminant (FLD) [50] are common choices and their specified directions are known as *eigenfaces* and *fisherfaces* [51] respectively. While both methods have been investigated within the context jet-images, FLD is observed to perform significantly better and thus this paper focuses on FLD exclusively. FLD identifies the plane in the high dimensional feature space which maximizes the separation between the jet classes and simultaneously minimizes the scatter within each jet class.⁵ Since FLD uses knowledge of the within-class variations, it is less influenced by variations present in both classes than PCA. In addition, we employ a regularized FLD implementation [52] which mitigates the impact of statistical fluctuations in the training sample and reduces overfitting.

FLD is trained using a set of preprocessed example jet-images from two classes and produces a discriminant F , which we denote the *Fisher-jet*, that has the same dimensionality as the example jet-images and thus can be viewed as a jet-image itself. Discrimination between classes for a jet-image, A , is then achieved by projecting A onto the Fisher-jet. That is,

$$D[A] = \sum_{k=1}^{N^2} \bar{F}_k \cdot \bar{A}_k \quad (2.2)$$

where $D[A]$ is the discriminant output for the jet-image A , and \bar{F} and \bar{A} are the vector representations of F and A , respectively, as discussed in section 2.1.

FLD produces a robust classifier which is fast to apply, fast to derive, and can itself be interpreted as a type of image. Therein lies much of the technique's power. For any system

⁵In general, FLD can be trained on M classes, with $M \geq 2$, resulting in $M - 1$ Fisher discriminants for separating the M classes

under study, the discriminant reduces the dimensionality of the problem, such that the properties of a jet can easily be expressed numerically and a cut on the discriminant value D can be used to separate classes. Since the discriminant output is computed by performing a projection of a jet-image onto the Fisher-jet, positive values in the Fisher-jet identify features which are indicative of a jet of class A whereas class B is highlighted by negative coefficients. The magnitude of the coefficient indicates the classification strength of the feature. To understand the behavior of the discriminant, the Fisher-jet image reveals what differences are being exploited in a simple and intuitive way. This is an important aspect of the technique. Such a visualization can be used to convince ourselves that the algorithm is detecting relevant differences between the two classes. Furthermore, the ability to visualize the classifier means that this approach is also useful as a step in exploratory analysis, and could be used to guide the construction of top-down analytical solutions to a problem. This aspect of our technique contrasts with widely used multivariate techniques in high-energy physics, such as neural networks and boosted decision trees, which predominantly use variables derived from the jet's kinematics and substructure rather than directly using the jet constituents.

3 Samples

Monte Carlo samples are produced with PYTHIA8 (v8.176) [53], MADGRAPH (v1.5.11) [54], and HERWIG++ (v2.6.3) [55], with proton-proton collisions at $\sqrt{s} = 8$ TeV. In all cases, the effect of pileup from multiple collisions in one crossing of the collider beams is simulated by adding the energy of particles from additional PYTHIA8 “minBias” interactions. The number of additional pileup interactions per bunch crossing is a random variable sampled from a Poisson distribution with mean μ . In this paper we consider the cases of no pileup and $\mu = 30$.

The signal sample used for these studies is di-boson WW production, with one W decaying to a muon and neutrino, and one W decaying to a pair of quarks. The matrix element and shower for the boosted W samples are calculated with PYTHIA8. The boost of the system is controlled by applying a loose cut on the matrix-element \hat{p}_T , followed by a harder cut on the leading reconstructed jet p_T such that it falls within a 50 GeV-wide p_T window (e.g. $200 < p_T < 250$ GeV). The minimum and maximum \hat{p}_T was set 50 GeV below and above the jet p_T interval being studied. The corresponding light and gluon jet background is simulated using PYTHIA8 with $W + \text{jets}$, where the W decays to a muon and a neutrino. Again, a loose cut on \hat{p}_T is used to set a generator-level boost for the system, followed by a harder cut on the leading reconstructed jet. Similar samples are generated with HERWIG++ for validation studies.

Jet finding is performed with Fastjet (v3.03) [56] on the massless four-vectors defined by the calorimeter towers. For each generated event, the towers are clustered using the Cambridge/Aachen algorithm [40] with a radius $R = 1.2$ and only jets with $p_T > 25$ GeV are kept. We apply trimming with $k_T R = 0.3$ subjets and $f_{cut} = 5\%$ to reduce the impact of pileup interactions. This serves as the noise reduction step introduced in section 2.2, and no additional pileup subtraction was performed.

4 Case study: W boson jet tagging

In this section we demonstrate the performance of this technique in classifying boosted $W \rightarrow qq'$ jets versus QCD jets. This serves both as an illustration of the power of the technique, and as a pedagogical example.

The inputs are grids of 25×25 pixels (i.e the jet-images of section 2.1), centered around the reconstructed jet axes, as described in section 3. Only the leading jet in each event is used and it is required to have mass $M \in [65, 95]$ GeV. For image preprocessing, we use trimming to reduce noise, the leading 2 subjets to align the image, and translation (based on the jet-image energy centroid), reflection, and equalization to uniformize the jet-images.

After preprocessing, it is observed that the jet-images varied significantly both with the jet transverse momentum (p_T) and the ΔR between the leading two subjets (ΔR_{jj}). Neither of these variations are surprising, since jet p_T differences cause wide variations in the energy distribution within the jet and differences in ΔR_{jj} lead directly to differences in the separation between features in the jet-image. As a result, the jet-images are binned in both p_T and ΔR_{jj} . We use 50 GeV bins for jets with $p_T \in [200, 500]$ GeV, and the ΔR_{jj} binning of each p_T bin is ($\Delta R_{jj} < 0.4$, $\Delta R_{jj} \in [0.4, 0.6]$, $\Delta R_{jj} \in [0.6, 0.8]$, $\Delta R_{jj} \in [0.8, 1.0]$, $\Delta R_{jj} \in [1.0, 1.2]$, $\Delta R_{jj} > 1.2$). The ΔR_{jj} binning is chosen such that resolvable differences in subjet position on the order of the subjet radius would not end up in the same bin and smear the jet images, but different optimizations could be explored.⁶

After preprocessing both the W and QCD jets, we train a FLD using these samples as signal and background respectively. A separate FLD is trained for each bin of $(p_T, \Delta R_{jj})$. It should be noted that equal numbers of W and QCD jets are used in the training of each FLD, as this was seen to produce better results than using different numbers of events in each sample (or equivalently different priors for each sample).

The resulting discriminant Fisher-jet, recast to a 25×25 image, can be seen in figure 2a for a single bin ($p_T \in [250, 300]$ GeV, $\Delta R_{jj} \in [0.6, 0.8]$), while the discriminant output when this Fisher-jet is applied to samples of W and QCD jets can be seen in figure 2b. Jet images of several bins of p_T and ΔR can be found in appendix A. As noted earlier, the sign of the pixel value in the Fisher-jet indicates which class the pixel helps identify, as seen in figure 2a where red pixels indicate importance for W-jets while blue features imply importance for QCD jets. We can see in the figure that the presence of a hard primary subjet located at ($Q_1 \sim 1.5, Q_2 \sim 1.25$) is almost irrelevant for discrimination, as both classes have hard primary subjets. However, the presence of a second hard subjet near ($Q_1 \sim 0.75, Q_2 \sim 1.25$) is a strong indicator of a W-jet. Finally, we see that energy appearing around the two subjets is a strong indication of a QCD jet, which is essentially telling us that QCD subjets tend to be wider and have radiation surrounding the subjets.

The background rejection vs. signal efficiency curves for the FLD, computed using the 1-D likelihood ratios of the output distribution of the FLD for W-jets and QCD jets, can be seen in figure 3a, along with the rejection vs. efficiency curves observed when using N-subjettiness (τ_2/τ_1) [7, 8] computed analogously with the 1-D likelihood ratios.

⁶For instance, a p_T dependent ΔR binning may prove beneficial for the very highly boosted cases in which most W jets have small ΔR separation.




Figure 2: A Fisher’s linear discriminant presented as an image (left) and the distributions of the discriminant output when applied to W-jets and Light-jets (right), when the FLD is trained on jets with $p_T \in [250, 300]$ GeV, mass $M \in [65, 95]$ GeV, and separation between subjects of $\Delta R \in [0.6, 0.8]$.

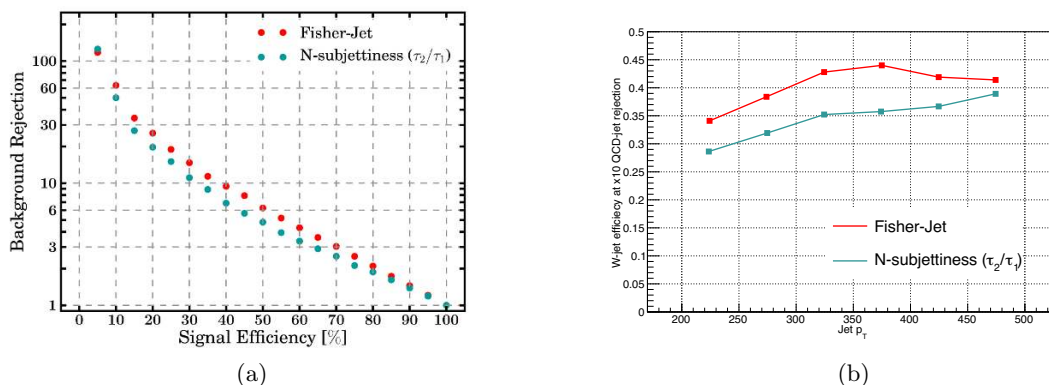


Figure 3: Background rejection vs signal efficiency curves obtained by training discriminants on samples with $p_T \in [250, 300]$ GeV (left), and the W-jet efficiency at fixed $\times 10$ QCD jet rejection versus jet p_T (right).

For the rejection vs. efficiency curve in figure 3a Fisher-jets are trained on jets satisfying $p_T \in [250, 300]$ in 6 bins of ΔR_{jj} , and a combined 1D likelihood ratio distribution is computed by taking the likelihood ratio for each jet computed with respect to appropriate ΔR_{jj} bin and merging these likelihood ratio values into a single distribution. The N-subjettiness distributions are not binned in ΔR_{jj} as this did not show any improvements in performance. Figure 3b shows the efficiency of W jets at a fixed QCD jet rejection of 10 as a function of jet p_T for the FLD (combining the 6 bins of ΔR_{jj} for each jet p_T bin) and for N-subjettiness. It can be seen that FLD outperforms N-subjettiness for the full range of jet p_T examined.

It should be noted that the output of FLD and N-subjettiness are correlated, as shown

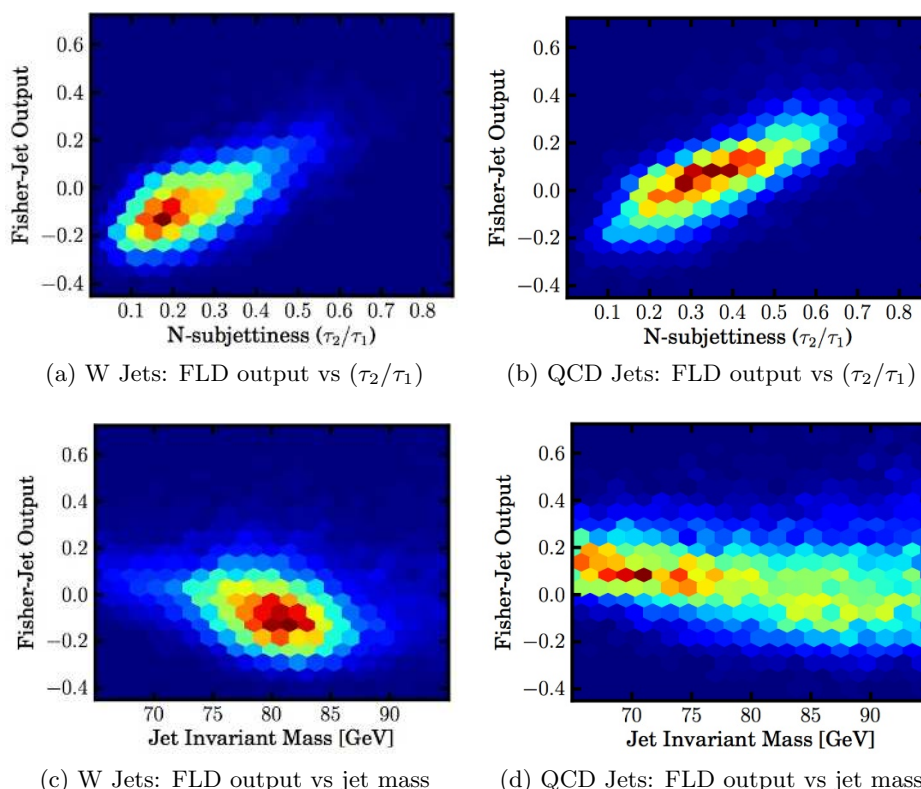


Figure 4: The Fisher-jet discriminant output for the bin $p_T \in [250, 300]$ GeV and $\Delta R \in [0.6, 0.8]$ plotted against: (4a) the output of N-subjettiness for W-jets, (4b) the output of N-subjettiness for QCD jets, (4c) the invariant mass of the jet for W-jets, and (4d) the invariant mass of the jet for QCD jets.

in figures 4a and 4b for W and QCD jets respectively, with a correlation coefficient of approximately 0.7 for both W and QCD jets. Thus, the Fisher-jet approach is able to combine in a linear way the information comprising the jet effectively, and capture much of the information of N-subjettiness and more. On the other hand, mass, which relies on quadratic relationships between the inputs, is a simple quantity which FLD does not reproduce, as shown in figures 4c and 4d for W and QCD jets respectively. Since the Fisher-jet output is only slightly correlated with mass, with a correlation coefficient of approximately -0.25 for both W and QCD jets indicating a small degree of anti-correlation, the performance of the classifier does not change dramatically whether it is applied to a small window around the W mass, or to a sample without jet mass cuts.

To investigate the effect of pileup, which essentially acts as a source of noise within the jet-image, the Fisher-jets are trained on samples without pileup and subsequently applied to statistically independent samples with pileup.⁷ No significant degradation in performance

⁷The reason for such an approach is that the samples without pileup are the best representation of the underlying physics (in the parlance of computer vision, these are the clearest pictures without noise) and are observed to give the best discrimination performance even in samples with noise due to pileup.

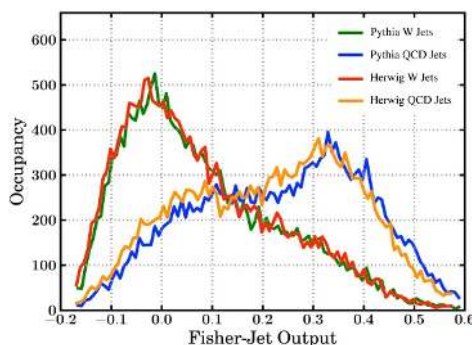


Figure 5: The Fisher-jet discriminant output, trained on a sample of PYTHIA8 W and QCD jets applied to PYTHIA8 and HERWIG++ samples.

is observed, likely due to the application of trimming. More in depth studies of pileup impact on the Fisher-jet approach are left for future studies.

To check the generator dependence of the classifier, a second sample of W and QCD jets are generated using HERWIG++ which implements an independent description of the hard sub-process and subsequent showering. The Fisher-jet trained on the PYTHIA8 samples is then applied to the HERWIG++ samples. The Fisher-jet discriminant output distributions for both W -jets and QCD jets from both PYTHIA8 and HERWIG++ are shown figure 5 for jets with $p_T \in [200, 250]$ GeV and $\Delta R \in [0.6, 0.8]$, and are seen to be extremely similar between both generators.

5 Conclusion

We present a generic algorithm to extract discriminating information between different classes of jets using parallels to techniques used in the field of computer vision.⁸ The algorithm uses a representation of jets as images, applies preprocessing techniques to construct a consistent set of jet-images, and applies a Fisher linear discriminant, which has been trained on a collection of example jets, to process the jet-images. The application of the Fisher discriminant projects different classes of jets to opposite ends of the discriminant output spectrum, thereby allowing for separation of the classes. The resulting algorithm is a linear function of the inputs and is easy to understand, with little prior knowledge needed in the tuning of the discriminant and does not rely on an analytic description of the system studied. This method is seen to be competitive with N -subjettiness when separating hadronically decaying boosted W bosons from QCD jets.

Importantly, the resulting algorithm can be used as a tool to visually discover differences in samples with no complete analytical description. The jet-images and the Fisher-jet can be trivially visualized, allowing one to study the important features of different jet classes and to gain insight into the features which are most discriminating between different classes. In this way, the method provides a starting place to provide simple answers to high dimensional problems.

⁸The software used in this analysis is available at <https://github.com/makagan/JetImages>.

This computer vision inspired approach stands out in its flexibility and its ease of inspection. While this study focuses on W vs QCD jets, initial studies have show that the method can be easily applied to the $H \rightarrow b\bar{b}$ system and provide powerful discrimination. The algorithm could be adapted for use in other problems of jet discrimination, such as identifying hadronically decaying tops. More generally, many problems in experimental physics contain a correspondence to the imaging problem discussed here, such as categorizing the ring images produced in Super-K [57], where our methods are likely directly applicable. In addition, since a consistent set of jet-images is produced by the pre-processing steps, non-linear classifiers could easily be used for jet-image processing to improve the discrimination performance.

Jet-images provide a connection between the fields of particle physics and computer vision, and is a departure from the ways that experimental high energy physics currently uses machine learning algorithms. However, this is merely a first step, while many of the techniques of the vast field of computer vision have yet to be explored in the context of particle physics and may prove powerful within this domain.

Acknowledgments

We would like to thank Jon Butterworth, Matthew Schwartz, Peter Skands, and Jesse Thaler for their helpful feedback. We would also like to thank Su Dong and the SLAC ATLAS group for their insightful discussions on the topic. This work is supported by the US Department of Energy (DOE) Early Career Research Program and grant DE-AC02-76SF00515.

A Fisher jet images




Figure 6: Fisher linear discriminants presented as jet-images, or Fisher-jets, for several p_T and ΔR_{jj} bins for jets with mass $M \in [65, 95]$ GeV.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] L. Evans and P. Bryant, *LHC Machine*, 2008 *JINST* **3** S08001 [[INSPIRE](#)].
- [2] M.H. Seymour, *Searches for new particles using cone and cluster jet algorithms: A comparative study*, *Z. Phys.* **C 62** (1994) 127.
- [3] T. Plehn and M. Spannowsky, *Top Tagging*, *J. Phys.* **G 39** (2012) 083001 [[arXiv:1112.4441](#)] [[INSPIRE](#)].
- [4] J.M. Butterworth, A.R. Davison, M. Rubin and G.P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *AIP Conf. Proc.* **1078** (2009) 189 [[arXiv:0809.2530](#)] [[INSPIRE](#)].
- [5] D.E. Kaplan, K. Rehermann, M.D. Schwartz and B. Tweedie, *Top Tagging: A Method for Identifying Boosted Hadronically Decaying Top Quarks*, *Phys. Rev. Lett.* **101** (2008) 142001 [[arXiv:0806.0848](#)] [[INSPIRE](#)].
- [6] J. Thaler and L.-T. Wang, *Strategies to Identify Boosted Tops*, *JHEP* **07** (2008) 092 [[arXiv:0806.0023](#)] [[INSPIRE](#)].
- [7] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N-subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [[INSPIRE](#)].
- [8] J. Thaler and K. Van Tilburg, *Maximizing Boosted Top Identification by Minimizing N-subjettiness*, *JHEP* **02** (2012) 093 [[arXiv:1108.2701](#)] [[INSPIRE](#)].
- [9] J. Gallicchio et al., *Multivariate discrimination and the Higgs + W/Z search*, *JHEP* **04** (2011) 069 [[arXiv:1010.3698](#)] [[INSPIRE](#)].
- [10] J. Gallicchio and M.D. Schwartz, *Quark and Gluon Tagging at the LHC*, *Phys. Rev. Lett.* **107** (2011) 172001 [[arXiv:1106.3076](#)] [[INSPIRE](#)].
- [11] Y. Cui, Z. Han and M.D. Schwartz, *W-jet Tagging: Optimizing the Identification of Boosted Hadronically-Decaying W Bosons*, *Phys. Rev. D* **83** (2011) 074023 [[arXiv:1012.2077](#)] [[INSPIRE](#)].
- [12] T. Plehn, G.P. Salam and M. Spannowsky, *Fat Jets for a Light Higgs*, *Phys. Rev. Lett.* **104** (2010) 111801 [[arXiv:0910.5472](#)] [[INSPIRE](#)].
- [13] J.-E. Daum, U. Harst and M. Reuter, *Running Gauge Coupling in Asymptotically Safe Quantum Gravity*, *JHEP* **01** (2010) 084 [[arXiv:0910.4938](#)] [[INSPIRE](#)].
- [14] M. Cacciari, G.P. Salam and G. Soyez, *The Catchment Area of Jets*, *JHEP* **04** (2008) 005 [[arXiv:0802.1188](#)] [[INSPIRE](#)].
- [15] S.D. Ellis, C.K. Vermilion and J.R. Walsh, *Recombination Algorithms and Jet Substructure: Pruning as a Tool for Heavy Particle Searches*, *Phys. Rev. D* **81** (2010) 094023 [[arXiv:0912.0033](#)] [[INSPIRE](#)].
- [16] M. Backovic, J. Juknevič and G. Perez, *Boosting the Standard Model Higgs Signal with the Template Overlap Method*, *JHEP* **07** (2013) 114 [[arXiv:1212.2977](#)] [[INSPIRE](#)].
- [17] D.E. Soper and M. Spannowsky, *Finding physics signals with shower deconstruction*, *Phys. Rev. D* **84** (2011) 074002 [[arXiv:1102.3480](#)] [[INSPIRE](#)].

- [18] A. Hook, M. Jankowiak and J.G. Wacker, *Jet Dipolarity: Top Tagging with Color Flow*, *JHEP* **04** (2012) 007 [[arXiv:1102.1012](#)] [[INSPIRE](#)].
- [19] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [20] S.D. Ellis, C.K. Vermilion, J.R. Walsh, A. Hornig and C. Lee, *Jet Shapes and Jet Algorithms in SCET*, *JHEP* **11** (2010) 101 [[arXiv:1001.0014](#)] [[INSPIRE](#)].
- [21] L.G. Almeida et al., *Substructure of high- p_T Jets at the LHC*, *Phys. Rev. D* **79** (2009) 074017 [[arXiv:0807.0234](#)] [[INSPIRE](#)].
- [22] C.F. Berger, T. Kucs and G.F. Sterman, *Event shape/energy flow correlations*, *Phys. Rev. D* **68** (2003) 014012 [[hep-ph/0303051](#)] [[INSPIRE](#)].
- [23] J. Gallicchio and M.D. Schwartz, *Seeing in Color: Jet Superstructure*, *Phys. Rev. Lett.* **105** (2010) 022001 [[arXiv:1001.5027](#)] [[INSPIRE](#)].
- [24] M. Jankowiak and A.J. Larkoski, *Angular Scaling in Jets*, *JHEP* **04** (2012) 039 [[arXiv:1201.2688](#)] [[INSPIRE](#)].
- [25] D. Kahawala, D. Krohn and M.D. Schwartz, *Jet Sampling: Improving Event Reconstruction through Multiple Interpretations*, *JHEP* **06** (2013) 006 [[arXiv:1304.2394](#)] [[INSPIRE](#)].
- [26] S.D. Ellis, A. Hornig, T.S. Roy, D. Krohn and M.D. Schwartz, *Qjets: A Non-Deterministic Approach to Tree-Based Jet Substructure*, *Phys. Rev. Lett.* **108** (2012) 182003 [[arXiv:1201.1914](#)] [[INSPIRE](#)].
- [27] A. Altheimer et al., *Boosted objects and jet substructure at the LHC. Report of BOOST2012, held at IFIC Valencia, 23rd-27th of July 2012*, *Eur. Phys. J. C* **74** (2014) 2792 [[arXiv:1311.2708](#)] [[INSPIRE](#)].
- [28] G.P. Salam, *Towards Jetography*, *Eur. Phys. J. C* **67** (2010) 637 [[arXiv:0906.1833](#)] [[INSPIRE](#)].
- [29] A. Abdesselam et al., *Boosted objects: A Probe of beyond the Standard Model physics*, *Eur. Phys. J. C* **71** (2011) 1661 [[arXiv:1012.5412](#)] [[INSPIRE](#)].
- [30] A. Altheimer et al., *Jet Substructure at the Tevatron and LHC: New results, new tools, new benchmarks*, *J. Phys. G* **39** (2012) 063001 [[arXiv:1201.0008](#)] [[INSPIRE](#)].
- [31] T. Oliphant et al., *NumPy: Open source numerical computation tools for Python*, <http://www.numpy.org/> (1995).
- [32] E. Jones et al., *SciPy: Open source scientific tools for Python*, <http://www.scipy.org/> (2001).
- [33] F. Pedregosa et al., *Scikit-learn: Machine learning in Python*, *J. Machine Learning Res.* **12** (2011) 2825.
- [34] J.D. Hunter, *Matplotlib: A 2d graphics environment*, *Comput. Sci. Eng.* **9** (2007) 90.
- [35] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, *Nucl. Instrum. Meth. A* **389** (1997) 81 [[INSPIRE](#)].
- [36] ATLAS collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *2008 JINST* **3** S08003 [[INSPIRE](#)].
- [37] CMS collaboration, *The CMS experiment at the CERN LHC*, *2008 JINST* **3** S08004 [[INSPIRE](#)].

- [38] S. Catani, Y.L. Dokshitzer, M.H. Seymour and B.R. Webber, *Longitudinally invariant K_t clustering algorithms for hadron hadron collisions*, *Nucl. Phys. B* **406** (1993) 187 [INSPIRE].
- [39] S.D. Ellis and D.E. Soper, *Successive combination jet algorithm for hadron collisions*, *Phys. Rev. D* **48** (1993) 3160 [hep-ph/9305266] [INSPIRE].
- [40] Y.L. Dokshitzer, G.D. Leder, S. Moretti and B.R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [hep-ph/9707323] [INSPIRE].
- [41] M. Cacciari, G.P. Salam and G. Soyez, *The Anti- $k(t)$ jet clustering algorithm*, *JHEP* **04** (2008) 063 [arXiv:0802.1189] [INSPIRE].
- [42] J. Shelton, *TASI Lectures on Jet Substructure*, arXiv:1302.0260 [INSPIRE].
- [43] I. Volobouev, *FFTJet: A Package for Multiresolution Particle Jet Reconstruction in the Fourier Domain*, arXiv:0907.0270 [INSPIRE].
- [44] V. Rentala, W. Shepherd and T.M.P. Tait, *Tagging Boosted W s with Wavelets*, *JHEP* **08** (2014) 042 [arXiv:1404.1929] [INSPIRE].
- [45] I.M. Dremin, G.K. Eyyubova, V.L. Korotkikh and L.I. Sarycheva, *Two-dimensional discrete wavelet analysis of multiparticle event topology in heavy ion collisions*, *Indian J. Phys.* **85** (2011) 39 [arXiv:0711.1657] [INSPIRE].
- [46] J.W. Monk, *Wavelet Analysis: Event De-noising, Shower Evolution and Jet Substructure Without Jets*, arXiv:1405.5008 [INSPIRE].
- [47] W. Zhao, R. Chellappa, P. J. Phillips and A. Rosenfeld, *Face recognition: A literature survey*, *ACM Comput. Surv.* **35** (2003) 399.
- [48] J.J. Sylvester, *A demonstration of the theorem that every homogeneous quadratic polynomial is reducible by real orthogonal substitutions to the form of a sum of positive and negative square*, *Phil. Mag.* **4** (1852) 138.
- [49] K. Pearson, *On lines and planes of closest fit to systems of points in space*, *Phil. Mag.* **2** (1901) 559.
- [50] R.A. Fisher, *The use of multiple measurements in taxonomic problems*, *Annals Eugen.* **7** (1936) 179 [INSPIRE].
- [51] P.N. Belhumeur, J.P. Hespanha and D.J. Kriegman, *Eigenfaces vs. fisherfaces: Recognition using class specific linear projection*, *IEEE Trans. Patt. Anal. Machine Intel.* **19** (1996) 711.
- [52] Z. Zhang, G. Dai, C. Xu and M.I. Jordan, *Regularized discriminant analysis, ridge regression and beyond*, *J. Machine Learning Res.* **11** (2010) 2199.
- [53] T. Sjöstrand, S. Mrenna and P.Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [arXiv:0710.3820] [INSPIRE].
- [54] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer and T. Stelzer, *MadGraph 5 : Going Beyond*, *JHEP* **06** (2011) 128 [arXiv:1106.0522] [INSPIRE].
- [55] M. Bahr et al., *HERWIG++ Physics and Manual*, *Eur. Phys. J. C* **58** (2008) 639 [arXiv:0803.0883] [INSPIRE].
- [56] M. Cacciari, G.P. Salam and G. Soyez, “Fastjet: A software package for jet finding in pp and e+e- collisions.” <http://fastjet.fr/> (2005).
- [57] SUPER-KAMIOKANDE collaboration, Y. Fukuda et al., *The Super-Kamiokande detector*, *Nucl. Instrum. Meth. A* **501** (2003) 418 [INSPIRE].