

Fermi's Sibyl: Mining the gamma-ray sky for dark matter subhaloes

N. Mirabal^{1,2*}, V. Frías-Martínez³, T. Hassan², and E. Frías-Martínez³

¹ *Ramón y Cajal Fellow*

² *Dpto. de Física Atómica, Molecular y Nuclear, Universidad Complutense de Madrid, Spain*

³ *Telefonica Research, Madrid, Spain*

ABSTRACT

Dark matter annihilation signals coming from Galactic subhaloes may account for a small fraction of unassociated point sources detected in the Second *Fermi*-LAT catalogue (2FGL). To investigate this possibility, we present *Sibyl*, a Random Forest classifier that offers predictions on class memberships for unassociated *Fermi*-LAT sources at high Galactic latitudes using gamma-ray features extracted from the 2FGL. *Sibyl* generates a large ensemble of classification trees that are trained to vote on whether a particular object is an active galactic nucleus (AGN) or a pulsar. After training on a list of 908 identified/associated 2FGL sources, *Sibyl* reaches individual accuracy rates of up to 97.7% for AGNs and 96.5% for pulsars. Predictions for the 269 unassociated 2FGL sources at $|b| \geq 10^\circ$ suggest that 216 are potential AGNs and 16 are potential pulsars (with majority votes greater than 70%). The remaining 37 objects are inconclusive, but none is an extreme outlier. These results could guide future quests for dark matter Galactic subhaloes.

Key words: (cosmology:) dark matter – gamma-rays: observations – galaxies: active – (stars:) pulsars: general

1 INTRODUCTION

The extraordinary success of the *Fermi* mission marks the beginning of the golden age for gamma-ray astrophysics. With 24 months of data, the Second *Fermi*-LAT catalogue (2FGL) lists 1873 sources in the 100 MeV to 100 GeV energy range, of which 886 are AGNs and 108 are pulsars. While *Fermi* has greatly mitigated issues inherent to source localisation in the gamma-ray regime, 269 sources in the 2FGL (15% of the total) remain without obvious counterparts at Galactic latitude $|b| \geq 10^\circ$. Such failure to associate the entire *Fermi* catalogue continues to fuel speculation about the existence of new types of gamma-ray source classes.

Probably the most intriguing potential sources of gamma ray emission are dark matter subhaloes (Diemand et al. 2008; Springel et al. 2008). Numerical cold dark matter (CDM) simulations suggest that galaxies like our own are surrounded by a wealth of small dark matter subhaloes that survived structure formation (Klypin et al. 1999; Moore et al. 1999). Massive subhaloes ($M \geq 10^7 M_\odot$) would correspond to “classical” dwarf galaxies. Less massive ones would be optically elusive and might only be revealed as gamma-ray point sources when weakly interact-

ing massive particles (WIMPs) annihilate to gamma rays (Kuhlen, Madau & Silk 2009). As a result, nearby dark matter subhaloes might be lurking among the unassociated *Fermi* sources at high Galactic latitudes. If found, an annihilation signal from Galactic subhaloes would clinch the first non-gravitational signature of dark matter.

The hunt for dark matter subhaloes in the *Fermi* catalogue is currently underway (Nieto et al. 2011; Belikov, Hooper & Buckley 2011; Zechlin et al. 2011; Ackermann et al. 2012b). Most approaches involve the hypothesised sharp spectral cut-off or step expected at the WIMP mass (Bergström et al. 2005). Assuming that the WIMP mass falls between 100 MeV and 50 GeV, a dark matter subhalo could be detectable in the *Fermi* MeV-GeV band, but would disappear in the GeV-TeV band, effectively creating a TeV dropout.

Here we investigate the possibility of identifying dark matter subhalo candidates using supervised machine learning algorithms. Rather than starting with an *ad hoc* theoretical dark matter spectrum we would like to exploit pattern recognition of known gamma-ray features in associated sources and use this information to locate outliers that might constitute novel emitters. Machine learning algorithms have already been used to study the First *Fermi* LAT Catalogue (1FGL). For example, Ackermann et al. (2012a) inves-

* E-mail: mirabal@gae.ucm.es




Figure 1. Gamma-ray features ranked in order of importance. MeanDecreaseAccuracy measures the difference between accuracy rates before and after permutation of individual features averaged over all trees. Higher percentages indicate more importance.

tigated classification trees and logistic regression to predict classes of unassociated sources in the 1FGL based on a set of gamma-ray features. K-means clustering was also applied to help distinguish individual counterparts within *Fermi* error contours (Mirabal, Nieto, & Pardo 2010).

With an additional year of collected *Fermi* data, the gamma-ray features reported in the 2FGL have improved substantially. In this paper we train the Random Forest classifier (Breiman 2001) on identified/associated *Fermi* objects and build a set of decision trees that provide predictions for high-latitude unassociated *Fermi* objects in the 2FGL. The paper is organised as follows. In section 2 we describe the datasets and the Random Forest algorithm. Section 3 describes the performance of the classifier on unassociated *Fermi* sources. Section 4 details the search for potential outliers. Finally, we provide our conclusions and discuss future work in section 5.

2 DATASETS AND RANDOM FORESTS

Random Forest is an ensemble classifier that grows a large forest of classification trees (Breiman 2001). Decision trees are classification tools that have a tree structure, where each split is based on the information gained considering the elements of the feature space (Quinlan et al. 1994). To classify a new object, each tree in the forest votes on the class. The proportion of votes P that agree on a decision provides a measure of the accuracy of the classification. Random Forest then makes a prediction based on the majority of votes ($P > 0.5$). Random Forest also computes proximities between pairs of objects and produces scaling coordinates (1st and 2nd) that can be used to visualise datasets easily (Cox & Cox 2001). In addition, a number of comparisons have shown that Random Forest is unexcelled in accuracy among current classifiers (Svetnik et al. 2003; Qi et al. 2006; Soto et al. 2011; Richards et al. 2012). The




Figure 2. Properties of *Fermi* features plotted against each other. Top features include index, curve, variability, and 1st and 2nd scaling coordinates (1 and 2 respectively) generated by *Sibyl*. Two distinct classes are clear: AGNs (red) and pulsars (green).

analysis presented here uses the R randomForest package (Liaw & Wiener 2002).

For our dataset, we collected the complete *Fermi* LAT 2FGL catalogue that consists of 1873 sources (100 MeV–100 GeV) of which 1300 are firmly identified/associated and 573 are unassociated sources (Abdo et al. 2012; Ackermann et al. 2011). In total, we consider a list that includes 800 labelled AGNs (BL Lacs and flat-spectrum radio quasars only) and 108 pulsars. There are additional gamma-ray classes in the 2FGL, but AGNs and pulsars are the largest and most common at $|b| \geq 10^\circ$. Thereby we simply consider a bimodality of classes. Novelty detection will be discussed later on. For each of the 908 sources a total of 68 features are reported in the 2FGL. Features include Galactic latitude, Galactic longitude, spectral index (Index), variability, curvature index (Curve), and fluxes in five bands. In addition, we generate four derived features defined by flux ratios $FR_{ij} = Flux_i / Flux_j$ between consecutive bands for 0.1–0.3 GeV (Band 1), 0.3–1 GeV (Band 2), 1–3 GeV (Band 3), 3–10 GeV (Band 4), and 10–100 GeV (Band 5) comparable to the features first introduced by Ackermann et al. (2012a).

To avoid working with too many features that could generate noise in the classifier, we first identify the subset of features that best discriminates what constitutes an AGN or a pulsar. For that purpose, we compute the relevance of each feature towards the target class, rank them by importance, and apply the classifier to a subset of the most relevant ones. Specifically, we use the measure of importance – MeanDecreasedAccuracy – implemented within randomForest (Breiman 2001; Svetnik et al. 2004). Initially, the accuracy rate is computed for each tree as the Random Forest is constructed. The value of a particular feature is then permuted across all the objects while other features are left unchanged, and the accuracy rate is recorded again. The MeanDecreaseAccuracy is the overall percentage decrease in accuracy rate averaged over all trees. If the feature is impor-

Table 1. Predictions and voting percentages for unassociated *Fermi* sources in the 2FGL, ordered by RA

Source	P_{AGN}	P_{Pulsar}	Prediction
2FGL J0004.2+2208	0.974	0.026	AGN
2FGL J0014.3–0509	0.992	0.008	AGN
2FGL J0031.0+0724	0.946	0.054	AGN
2FGL J0032.7–5521	0.998	0.002	AGN
2FGL J0039.1+4331	0.776	0.224	AGN
2FGL J0048.8–6347	0.922	0.078	AGN
2FGL J0102.2+0943	0.998	0.002	AGN
2FGL J0103.8+1324	0.998	0.002	AGN
2FGL J0106.5+4854	0.406	0.594	–
2FGL J0116.6–6153	0.992	0.008	AGN
2FGL J0124.6–2322	0.998	0.002	AGN
2FGL J0129.4+2618	0.820	0.180	AGN
2FGL J0133.4–4408	0.968	0.032	AGN
2FGL J0143.6–5844	1.000	0.000	AGN
2FGL J0158.4+0107	0.990	0.010	AGN

Note: The complete list of predictions is available at <http://www.gae.ucm.es/~mirabal/sibyl.html>

tant, there should be a greater decrease in the accuracy rate compared to the initial one. Figure 1 shows the top most important features ranked by importance. We found that the features that most clearly differentiate AGN and pulsar classes include: Index, Curve, Variability, and Flux Ratios (FR_{12} , FR_{23} , FR_{34} , and FR_{45}). This selection is in general agreement with Ackermann et al. (2012a) who chose similar features for supervised classification of the 1FGL. Additional features showed considerably smaller values in their importance (MeanDecreaseAccuracy) and are thus not considered in the analysis.

In order to construct and train *Sibyl*¹, we start with the 800 AGNs and 108 pulsars. However, given the highly imbalanced nature of the sets, we replicate the pulsar sample to attain a closer size as the AGN class (Ling & Li 1998; Chawla 2005). Practically, the content of the datasets have not changed but the replication mechanism adds weight to the minority sample and achieves improved performance in the classifier.

After matching the AGNs and pulsar datasets, we generate 100 alternate training and testing sets built from randomly selected objects (2/3 and 1/3 of the sample respectively). We next produce Random Forest models with 500 trees for each training set. For validation, individual performance is evaluated at each of the 100 testing sets. Accuracy rates are computed directly by comparing the class predicted by *Sibyl* with the true class for each object in the testing sets. On average, *Sibyl* achieves an accuracy rate of 97.1% based on majority voting (97.7% for AGNs and 96.5% for pulsars). Inclusion of absolute Galactic latitude $|b|$ in the classifier lowered AGN and pulsar accuracy rates slightly to 97.4% to 95.5% respectively. Since pulsars tend to be situated along the Galactic plane and AGN are more numerous at high Galactic latitude, it is possible that using Galactic latitude as a feature could introduce a tiny bias against AGN near the Galactic plane and pulsars away from it (Ackermann et al. 2012a). Generally, most of the misclassifications occur when less than 70% of the individual trees

($P < 0.7$) agree on a classification. Figure 2 displays the outstanding separation between AGNs and pulsars, which explains the high accuracy rates obtained by *Sibyl*.

3 APPLICATION TO UNASSOCIATED SOURCES

The designation of 2FGL sources usually falls into three categories: identified, associated, and unassociated. A firm identification of a gamma-ray source can only be established through contemporaneous temporal variability, similar spatial morphology, or equivalent pulsation at other wavelengths. An association only requires positional correspondence of a possible counterpart with a 2FGL source. Unassociated sources lack a formal counterpart at other wavelengths.

Here, we consider a fourth category to designate 2FGL sources: “prediction”. Our objective is twofold: to predict the class of high-latitude unassociated *Fermi* objects in the 2FGL; and to produce a list of outliers that could be explored as potential dark matter subhaloes. For each of the 269 unassociated *Fermi* sources at $|b| \geq 10^\circ$, *Sibyl* provides a prediction that the object is an AGN (P_{AGN}) or a pulsar (P_{Pulsar}) based on individual votes polled from all trees in the decision forest.

Since we want to isolate outliers that might constitute dark matter subhalo candidates, we only accept *Sibyl* predictions whenever $P > 0.7$ *i.e.*, at least 70% of the trees agree on the final decision. Otherwise, the object remains without a prediction. Such threshold value is set based on the results explained in Section 2. In total, *Sibyl* predicts 216 objects to be AGN and 16 to be pulsars. The resulting predictions and percentages of voting agreements are listed in Table 1. Finally, the remaining 37 objects left without a firm prediction are the focus of our outlier study in the next section. It is important to note that under some specific circumstances, dark matter subhaloes could mimic the spectral properties of certain pulsars (Baltz, Taylor & Wai 2007; Zechlin et al. 2011), we discuss this possibility further in 5.

¹ In ancient Greece, a sibyl was a person or agency considered to be a source of predictions or oracles.




Figure 3. Distribution of outlyingness for the 37 objects without firm predictions (shaded circles) and the 232 predicted by *Sibyl* (open circles). The top outliers are summarised in Table 2.

Table 2. Top outliers among high-latitude unassociated sources in the 2FGL

Source	P_{AGN}	Outlyingness
2FGL J0953.6–1504	0.658	9.0
2FGL J0418.9+6636	0.574	7.2
2FGL J1710.0–0323	0.500	7.1
2FGL J0533.9+6759	0.336	6.6
2FGL J0336.0+7504	0.476	6.2

4 SEARCH FOR DARK MATTER SUBHALOES IN THE 2FGL

In order to better understand the nature of the remaining 37 objects we want compute their outlyingness, which is a measure of how far away an object is from its closest class. Apart for predicting an object’s class, Random Forest computes the proximity of each predicted *Fermi* object n to every element k within each class $\sum_{class} prox(n, k)$. Formally, each individual proximity $prox(n, k)$ is computed as the fraction of trees in which elements n and k fall in the same terminal node (Breiman 2001; Liaw & Wiener 2002). The outlyingness of an element n is calculated as the reciprocal sum of the squared proximities to all objects within its class. This outlying measure is normalised by subtracting the median and dividing by the absolute deviation from the median (Liaw & Wiener 2002). Larger outlyingness values are common in objects that are extremely different from the average, which could correspond to dark matter subhaloes. Figure 3 shows the distribution of outlyingness for the 37 objects without firm predictions. For comparison, we also plot the outlyingness for the remaining 232 objects that were predicted by *Sibyl* in the previous section. Additionally, Table 2 lists the five objects with the largest outlyingness.

Given that outlyingness values much greater than 10 usually indicate novel cases (Breiman 2001), there is no strong indication of novelties (significant outliers) among

the 37 objects without firm predictions. We find that the top five outliers have an average flux of 1.1×10^{-9} ph cm $^{-2}$ s $^{-1}$ (1–100 GeV). Unassociated source fluxes at high latitudes range from 7.7×10^{-9} to 1.1×10^{-10} ph cm $^{-2}$ s $^{-1}$. Thus, they are not necessarily the faintest sources in the dataset. On the other hand, the mean photon index of sources in Table 2 is 2.2 ± 0.3 , while photon indices in the unassociated sample range from 1.1 to about 3.0. Inspection of individual features in this manner yields limited insight into what makes these outliers stand out from the rest of the sample. As mentioned before, the exploration of the entire feature space is precisely where the supervised learning algorithm excels. Unfortunately, *Sibyl* cannot assess by itself whether the outlyingness is due to an anomaly in the data taking process, a simple variation within known *Fermi* classes, or a true novel source class such as dark matter annihilation in Galactic subhaloes.

5 DISCUSSION AND CONCLUSIONS

We have presented the outcome of the Random Forest predictor *Sibyl*. The results show that machine learning algorithms provide a reasonable route not only to predict unassociated AGNs/pulsars in the 2FGL, but also to produce a list of sources with unusual features that could be explored as potential dark matter subhalo candidates. After training on 908 identified/associated *Fermi* objects, *Sibyl* has been applied to predict the class of unassociated *Fermi* sources in the 2FGL. Out of 269 unassociated sources at high latitudes, we have found that 216 are AGN candidates and 16 are considered potential pulsars with prediction accuracy rates greater than 96.5%. *Sibyl* has also produced a list of 37 outlier objects; however, none of these exhibits significant outlyingness that can be directly connected to new gamma-ray classes (including dark matter subhaloes) at this point. We emphasise again that our results are strict predictions based on pattern recognition and thus a rigorous source identification process will have to localise actual counterparts at other wavelengths.

The results leave some room, albeit very small, to accommodate dark matter subhaloes or alternative source classes in the 2FGL. These pockets could be targeted to exhaust all possibilities. Looking forward, zooming in on a reduced group of sources might be a wise observational strategy. For obvious reasons, the set of objects with the largest outlyingness could be a reasonable place to conduct a dedicated survey. If dark matter consists of particles with a mass below 60 GeV (Hooper & Linden 2011), dark matter subhaloes might also be camouflaging among the ranks of predicted pulsars as their spectral signature could be similar to the pronounced spectral cut-off predicted by certain dark matter models. However, a number of these sources could be old radio-quiet pulsars which will complicate the search for a counterpart (Kerr et al. 2012).

There are a number of issues that need further exploration. For instance, the predictions are heavily dependent on the robustness of the spectral parameters listed in the 2FGL. Most machine learning algorithms lack a proper treatment of uncertainties in each of the features considered (Carroll et al. 2006; Morgan et al. 2012). Inclusion of uncertainties as individual features in *Sibyl* did not yield

improved performances in our predictions. With additional years of flight, *Fermi* will likely keep improving the accuracy of the gamma-ray features. However, attempts should be made to account for feature errors properly. In a forthcoming paper, we plan to explore a more refined breakdown into further *Fermi* subclasses (Hassan et al. 2012). There are at least four AGN subclasses in the 2FGL comprising BL Lacs, flat-spectrum radio quasars, misaligned AGNs, and Seyfert galaxies (Ackermann et al. 2011). Pulsars can be further partitioned into radio-loud, radio-quiet, and millisecond pulsars (Abdo et al. 2010).

Ultimately, the main reason that a large *Fermi* fraction remains unassociated to begin with has to do with the quality of localisations in the gamma-ray band. At faint flux levels, it becomes ever more difficult to associate a *Fermi* source with a particular counterpart. The best association procedures rely on positional coincidences and correlations with flat-spectrum radio sources (Ackermann et al. 2011). None the less, considering the results presented here and the scatter in gamma-ray flux it seems likely that many of the unassociated sources at high latitude are AGNs or mid-latitude pulsars with somewhat fainter radio fluxes than their brighter cousins.

Without a major breakthrough in localisations, the actual counterparts of most unassociated *Fermi* objects will be difficult to pinpoint in the short term. Machine learning algorithms can help narrow the options. Eventually, we will see significant improvement in localisations, particularly for Galactic sources, courtesy of the future Cherenkov Telescope Array (CTA) that will achieve enhanced angular resolution above 25 GeV (CTA Consortium 2011).

ACKNOWLEDGMENTS

N.M. acknowledges support from the Spanish government through a Ramón y Cajal fellowship and the Consolider-Ingenio 2010 Programme under grant MultiDark CSD2009-00064. We thank Pablo Saz Parkinson for helpful conversations. We also thank the referee for useful suggestions and comments on the manuscript.

REFERENCES

- Abdo A. A. et al., 2010, *ApJS*, 187, 460
 Abdo A. A. et al., 2011, *ApJS*, 199, 31
 Ackermann M. et al., 2011, *ApJ*, 743, 171
 Ackermann M. et al., 2012a, *ApJ*, submitted (arXiv:1108.1202)
 Ackermann M. et al., 2012b, *ApJ*, submitted (arXiv:1201.2691)
 Baltz E. A., Taylor J. E., Wai L. L., 2007, *ApJ*, 659, L125
 Belikov A. V., Hooper D., Buckley M. R., 2011, preprint (arXiv:1111.2613)
 Bergström L., Bringmann T., Eriksson M., Gustafsson M., 2005, *PRL*, 94, 131301
 Breiman L., 2001, *Machine Learning*, 45, 5
 Carroll R. J., Ruppert D., Stefanski L. A., Crainiceanu C., 2006, *Measurement Error in Nonlinear Models: A Modern Perspective*. Chapman and Hall. London
 Chawla N., 2005, in Maimon O, Rokach L., eds, *Data Mining and Knowledge Discovery Handbook*, New York, p. 853
 Cox T. F., Cox, M.A.A., 2001, *Multidimensional Scaling*. Chapman and Hall. London
 CTA Consortium, 2011, *Exp. Astron.*, 32, 193
 Diemand J. et al., 2008, *Nature*, 454, 735
 Hassan T., et al., 2012, in preparation
 Hooper D., Linden T., 2011, *Phys. Rev. D*, 83, 083517
 Kerr M. et al., 2012, *ApJL*, 748, L2
 Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, *ApJ*, 522, 82
 Kuhlen M., Madau P., Silk J., 2009, *Science*, 325, 970
 Liaw A., Wiener M., 2002, *R News*, 2/3, 18
 Ling C. X., Li C., 1998, in Agrawal R., Stolorz P. E., Piatetsky-Shapiro G., eds, *Proc. of the Fourth International Conference on Knowledge Discovery and Data Mining (KDD-98)*, New York, p. 73
 Mirabal N., Nieto D., Pardo S., 2010, preprint (arXiv:1007.2644)
 Moore B. et al., 1999, *ApJ*, 524, L19
 Morgan A. N., Long J., Richards J. W., Broderick T., Butler N., Bloom J. S., 2012, *ApJ*, 746, 170
 Nieto D., Martínez V., Mirabal N., Barrio J. A., Satalecka K., Pardo S., Lozano I., 2011, preprint (arXiv:1110.4744)
 Qi Y. et al., 2006, *Proteins*, 63, 490
 Quinlan J. R., 1993, *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers, San Francisco, CA
 Richards J. W., Homrighausen D., Freeman P. E., Schafer C. M., Pznanski D., 2012, 419, 1121
 Soto V., Frías-Martínez V., Konstan J., Conejo R., Marzo J. L., Oliver N., 2011, eds, *Proc. User Modeling, Adaptation and Personalization (UMAP 2011)*, Berlin, p. 377
 Springel V. et al., 2008, *MNRAS*, 391, 1685
 Svetnik V. et al., 2003, *J. Chem. Inf. Comput. Sci.*, 43, 1947
 Svetnik V., Liaw A., Tong C., Wang T., 2004, in Roli F., Windeatt T., eds, *Proc. Multiple Classifier Systems 2004*, Berlin, p. 334
 Zechlin H.-S., Fernandes M. V., Elsässer D., Horns D., 2012, *A&A*, 538, 93