IMPLICATIONS OF HIERARCHIES FOR RSO RECOGNITION, IDENTIFICATION, AND CHARACTERIZATION

Matthew P. Wilkins[(1)], Avi Pfeffer[(2)], Brian Ruttenberg[(3)], Paul W. Schumacher[(4)], Moriba K. Jah[(5)]

[(1)] *Applied Defense Solutions, P.O. Box 1102, Columbia, MD 21044, 410-715-0005, mwilkins@applieddefense.com*
[(2) (3)] *Charles River Analytics, Inc. 625 Mt. Auburn St., Cambridge, MA 02138 USA, 617-491-3474, {apfeffer,bruttenberg}@cra.com*
[(4)] *Air Force Research Laboratory, 550 Lipoa Parkway, Kihei, HI 96753 paul.schumacher@us.af.mil*
[(5)] *Air Force Research Laboratory, 3550 Aberdeen Ave. SE, Kirtland AFB, NM 87117 moriba.jah.1@us.af.mil*

Abstract: *In our previous work, we demonstrated that hierarchical (taxonomical) trees can be used to depict hypotheses in a Bayesian object recognition and identification process using Figaro, an open source probabilistic programming language. We assume in this work that we have appropriately defined a satellite taxonomy that allows us to place a given space object (RSO) into a particular class of object without any ambiguity. Such a taxonomy allows one to assess the probability of assignment to a particular class by determining how well the object satisfies the unique criteria of belonging to that class. Furthermore, tree-based taxonomies delineate unique signatures by defining the minimum amount of information required to positively identify a RSO. Because of these properties of taxonomic trees, we can now explore the implications of RSO taxonomic trees for model distance metrics and sensor tasking. In particular, we seek to exploit the fact that taxonomic trees provide a model "neighborhood" that can be used to initiate a Monte Carlo or Multiple Hypothesis algorithm. We contend this feature of taxonomies will provide a quantifiable metric for model distances and the explicit number of models that should be considered, both of which currently do not exist. Additionally, the discriminating characteristics of taxonomic classes can be used to determine the kind of data and the associated sensor that needs to be tasked to acquire that data. We also discuss the concept of multiple interacting hierarchies that provide deeper insight into how object interact with one another.*

Keywords: *resident space object, taxonomy, identification, characterization, cataloging*

1. Introduction

Our work in this paper is motivated by a series of rhetorical questions arising out of common problems faced by the Space Situational Awareness (SSA) community. In general, we face situations where we have sparse data of varying quality from a variety of sensor types that all provide small pieces of the SSA puzzle. The SSA community has often attempted to assemble those puzzle pieces using filters that rely upon a choice of dynamic, environment, and sensor models to perform orbit determination (OD) and space object identification (SOI), sometimes known as positive object identification (POI). For each and every batch of data that needs to be processed, it is up to the user of the filter to choose the models and other associated input

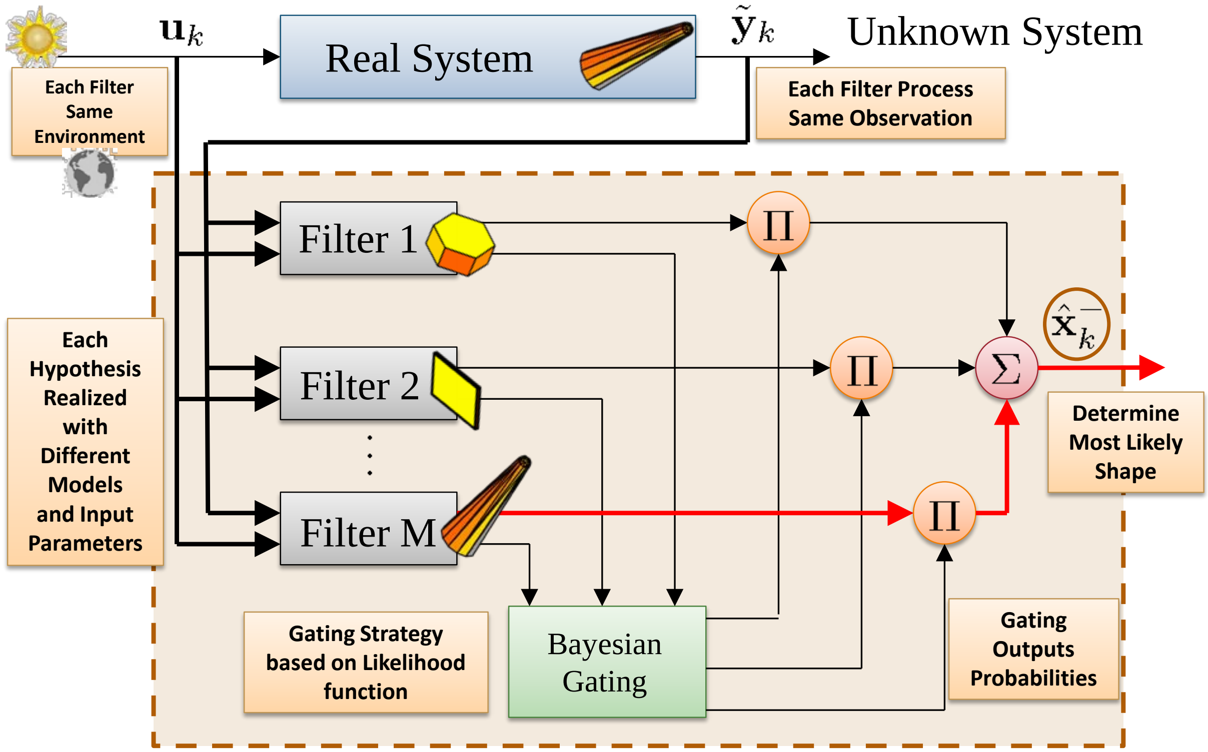

Figure 1: Multiple hypothesis approach to determine the most likely set of models and input parameters that best match the observed data.

bject (RSO) the data represents. But how do you know what hypothesis to choose in order to initialize a particular OD algorithm? Should you choose multiple hypotheses to cover all of the possibilities as is commonly done in the multiple hypothesis approach depicted in Figure 1? How do you know if your choices are an appropriately complete set of all possible hypotheses or just a random sampling? Can these initial hypotheses be chosen in a statistically rigorous algorithmic manner rather than relying on any one individual user's opinion (expert or otherwise)?

In our previous work, we demonstrated that hierarchical (taxonomical) trees can be used to depict hypotheses in a Bayesian object recognition and identification process using Figaro, an open source probabilistic programming language. [1] Provided that we have appropriately defined a satellite taxonomy that allows us to place a given RSO into a particular class of object without any ambiguity, one can assess the probability of assignment to a particular class by determining how well the object satisfies the unique criteria of belonging to that class. Ultimately, tree-based taxonomies delineate unique signatures by defining the minimum amount of information required to positively identify a RSO.

Let us take the rhetorical questions one step further. Before you go collect observational data in the first place, how do you know what sensor to task in order to increase the likelihood of positive object identification? How do you know what observational data are the missing pieces of the POI puzzle? Can sensor tasking be automated in a statistically rigorous algorithmic manner for POI and track custody of RSOs? This work will discuss how hierarchies, otherwise

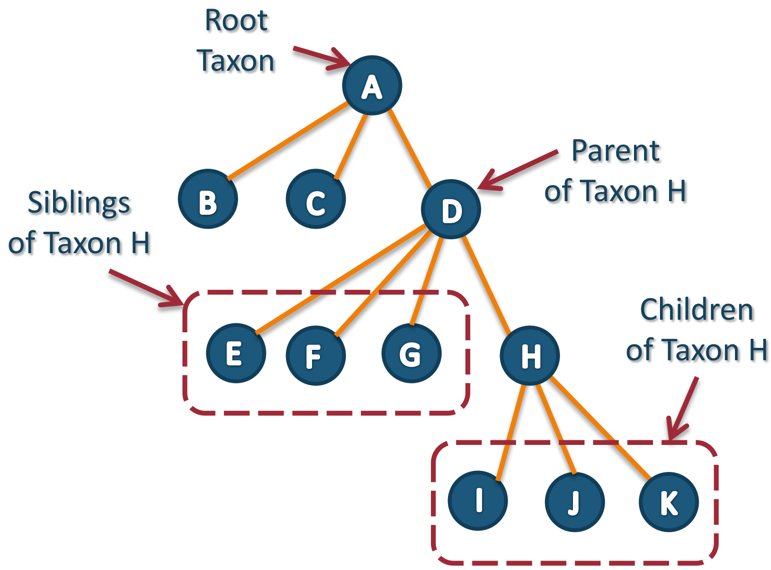

Figure 2: Taxonomy fundamentals

known as taxonomies, can help address the issues that we have posited. Our hope is that the reader will gain an appreciated for the benefits of organizing information in a hierarchical manner to tackle some of the most pressing issues facing the SSA community.

2. Hierarchy (Taxonomy) Fundamentals

Hierarchical structures can be found in many domains, such as social networks. [2] In vision tasks, hierarchical models such as latent Dirichlet allocation (LDA) [3] are frequently used to model the hierarchical generation of features, regions, or objects in an image [4] [5]. Once learned, the hierarchy can be used to classify an object into labels with specific semantic meaning, or general labels if such specificity is not possible [6] [7]. This type of hierarchical reasoning and classification can be found in other domains as well, such as text [8] or audio [9] classification. The best example of a taxonomy that enjoys almost universal scientific acceptance is the classical Linnaean biological taxonomy. [10] A strength of Linnaean taxonomy is that it can be used to organize the different kinds of living organisms, simply and practically. Every species can be given a unique binomial name based upon a set of defining characteristics and features. This uniqueness and stability are a result of the acceptance by biologists specializing in taxonomy, not merely of the binomial names themselves. We seek to provide a similar formal nomenclature system through a defined tree-based taxonomy structure for RSOs.

As depicted in Figure 2, each categorization, beginning with the most general or inclusive root of the hierarchical tree, at any level is called a *taxon*. Each taxon will have a set of uniquely distinguishing features that will allow one to place a given object into a specific group without any ambiguity. When a new object does not fall into a specific taxon that is already defined, the entire tree structure will need to be evaluated to determine if a new taxon should be created. Each taxon can have one or more children that are called siblings and have characteristics that are more refined than its parent. That taxon at the end of a particular branch of the hierarchical tree (sometimes called the leaf taxon) represents an individual space object. If one can assign an object to a leaf taxon, then one has accomplished the process of positive object identification. The collection of characteristics from the root taxon down to the bottom of a branch of the taxonomical tree creates a unique signature. One can assess the probability of assignment to a particular taxon by determining how well the object satisfies the unique criteria of belonging to that taxon and its associated parent taxa. Therefore, we can use taxonomic trees in a Bayesian process to assign prior probabilities to each of our object recognition and identification hypotheses.

3. Hierarchical Reasoning Tool

Probabilistic programming has recently developed as a potential solution to the problems associated with the creation of large and complex hierarchical models. [11] Models created with these languages tend to be highly modular, reusable, and can be reasoned with using a built in suite of algorithms that work on any model created in the language. Using the probabilistic programming library (PPL) Figaro™ by Charles River Analytics (CRA) [12], we have embarked upon a new approach to data correlation and aggregation using a space object taxonomy for automatically recognizing and classifying a space object called the Hierarchical Reasoning Tool (HRT). The Figaro™ PPL provides a uniform holistic language for building probabilistic models and asking queries about space objects rather than a stove pipe set of dynamic and sensor models for specific classes of objects. If you want to represent a new aspect of the domain, you don't have to create a model from scratch. Simply use existing constructs of the Figaro™ language and automatically apply the reasoning algorithms. This approach provides the ability to perform orbit determination using a variety of methods including particle filtering, which is a general method that applies to a wide range of situations that would guarantee the preservation of the orbit error characteristics in highly non-linear, stochastic, dynamical environments.

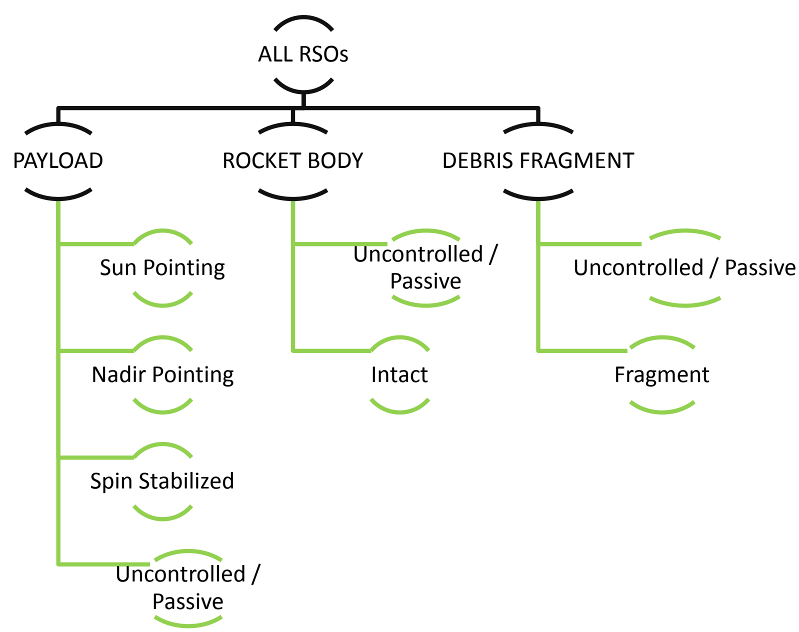

Figure 3: HRT rudimentary hierarchy for proof of concept

In the current proof-of-concept, we utilize a very rudimentary hierarchy that is broken up into three major taxa: payloads, rocket bodies, and debris fragments. Each of these taxa have characteristics defining object shape (intact or fragment) along with a set of possible attitude states (sun pointing, nadir pointing, spin stabilized, and uncontrolled). Objects that are uncontrolled are also deemed passive while objects with a controlled attitude were deemed active. This hierarchy is graphically depicted in Figure 3. Using a multiple model adaptive estimator (MMAE) similar to the one depicted in Figure 1, simulated right ascension and declination angles along with light curve observations were processed by the MMAE and the instantaneous likelihoods generated by the MMAE were fed into our Hierarchical Reasoning Tool (HRT). Initial hypotheses were chosen to match the various combinations of characteristics depicted by the hierarchy in Figure 3. We provide without proof the results of the MMAE in Figure 4 as a starting point to discuss the HRT.

The HRT results presented in Figure 4 represent the instantaneous best guess for the object type based upon the priors and the evidence provided by the MMAE. In these examples, there is no feedback from the HRT to the MMAE. We start from a condition where you assume no knowledge of what type of object you are tracking. As a consequence, the HRT assumes an initial probabilistic distribution of the object types equivalent to the real population distribution as shown in Table 1. We have found that the MMAE must assume a uniform initial distribution. This may be due to underlying stiff dynamics preventing the MMAE from adapting to an initial distribution that places too much emphasis on an incorrect object type. If the initial MMAE probabilities are too far off from uniform or the true hypothesis, then the MMAE will converge to an incorrect hypothesis. This results in incorrect evidence being asserted to the HRT causing it to believe an incorrect hypothesis. This highlights the importance of future work to incorporate confidence metrics and trustworthiness of evidence.

The initial MMAE result from the very first observation is that there is a very low probability for non-payload objects. Thus, such objects are ruled out nearly immediately by the MMAE. Sun, spin, and uncontrolled payloads are not discriminated yet, and still have significant probabilities. The HRT conclusion from first MMAE result is that non-payload objects are discriminated. With

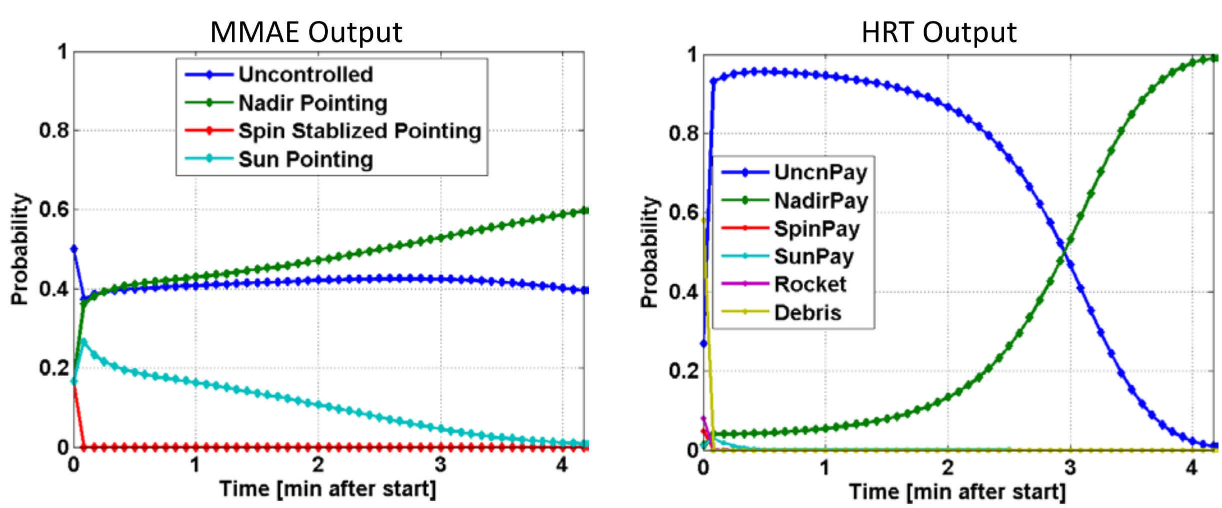

Figure 4: Hierarchical Reasoning Tool demonstration scenario results

Table 1. Initial RSO Population Statistics

		Probability of Occurrence			
Characteristic	**Value**	**RSO**	**Payload**	**Rocket Body**	**Debris Fragment**
Shape	Regular / Intact	0.34	1	0	0
	Rocket Body / Intact	0.08	0	1	0
	Debris / Fragment	0.58	0	0	1
Attitude Control	Active	0.07	0.21	0	0
	Passive	0.93	0.79	1	1
Spin State	Uncontrolled	0.93	0.79	1	1
	Spin Stabilized	0.024	0.14	0	0
	Nadir Pointing	0.023	0.035	0	0
	Sun Pointing	0.023	0.035	0	0

the remaining payload objects being non-discriminated, the most likely object is an uncontrolled payload with probability = ~0.9 strictly based upon population data. As the MMAE continues to process observational data, sun pointing objects are eliminated. We note that sun pointed payloads are a minority in space, which combined with lower probability MMAE results, the HRT discriminates sun payloads early.

As the MMAE slowly begins to discriminate the remaining two objects, uncontrolled and nadir objects, the HRT continues to collect probability evidence. The MMAE has difficulty deciding between the active and passive state. This is reflected in the HRT taking time to decide between those two states. Note that the active/passive decision is “the final straw” for the HRT. Once the building evidence shifts the HRT’s prior likelihoods from uncontrolled/passive to active, there is already enough evidence to choose the nadir pointing hypothesis (the correct solution). Note that the uncontrolled payload initially has a high probability since the HRT realizes that there are many more uncontrolled-payloads than nadir-pointing payloads. The MMAE must present discriminating evidence to show that the correct result is actually nadir-pointing. This evidence is presented slowly, as even the simulation completes the MMAE is still showing a 60/40 split between nadir and uncontrolled payloads. This shows how the HRT can tolerate and react to changing MMAE evidence. We have found that the HRT can consistently come to a conclusion regarding which hypothesis is the correct hypothesis even when the MMAE cannot make a firm decision based upon the dynamical models, environment models, and sensor models alone.

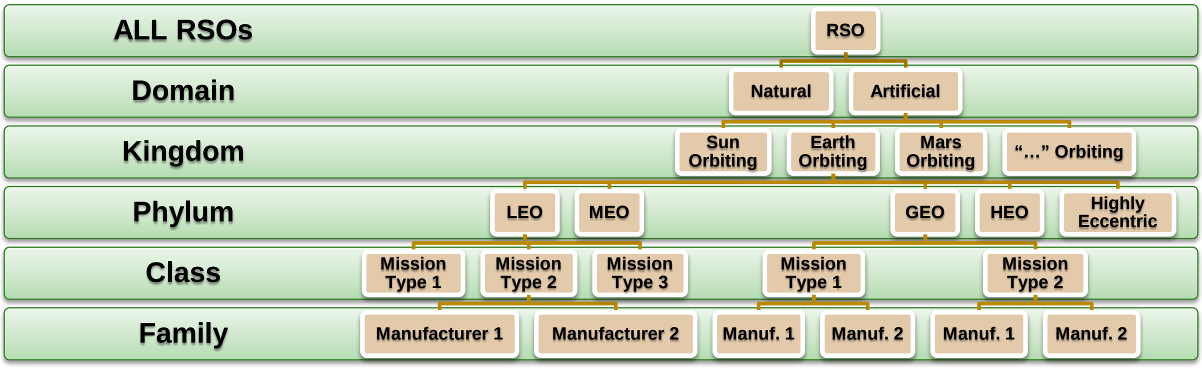

Figure 5: Proposed RSO taxonomy

3. A More Complete Hierarchy of RSOs

In our previous work, we formally described the process of how we decided to construct a resident space object taxonomy. [1] Alternate concepts for constructing an RSO hierarchy exist but do not directly lend themselves to the type of analysis. [13] [14] While we did not originally explore in great detail the distinctiveness of particular satellite busses [15] [16] [17] [18] [19], we now seek to explore the utility of such details for the notion of model distance metrics and sensor tasking requirements. As we previously mentioned, each taxon provides a set of distinct characteristics that allows one to assign an object to that class without ambiguity. The ancestral nodes of a given taxon provide more general criteria. For example, a generic GEO satellite class might be represented by a simple box-wing shape model. However, a sub-class for communication satellites might include one or more antenna shape features.

In order to construct a more robust taxonomy, we utilized the active payload database from the Union of Concerned Scientists (UCS). [20] The UCS Satellite Database is a listing of the more than 1000 operational satellites currently in orbit around Earth. This database conveniently categorizes each payload according to country of origin, owner operator, orbit class, mission class, and manufacturer using open-source information. Figure 7 represents the resulting hierarchy in terms of the broad categorizations down to the manufacturer level. If one populates this hierarchy with actual values, one quickly generates a large and unwieldy hierarchy as depicted in Figure 6. These graphical depictions do not extend to the leaf taxa of individual objects because there would need to be 20,000 to 30,000 taxa at the bottom of the tree.

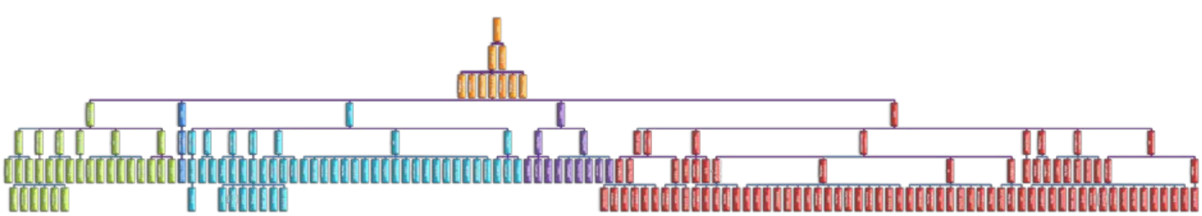

Figure 6: Expanded RSO taxonomy showing how quickly hierarchies can become large and unwieldy

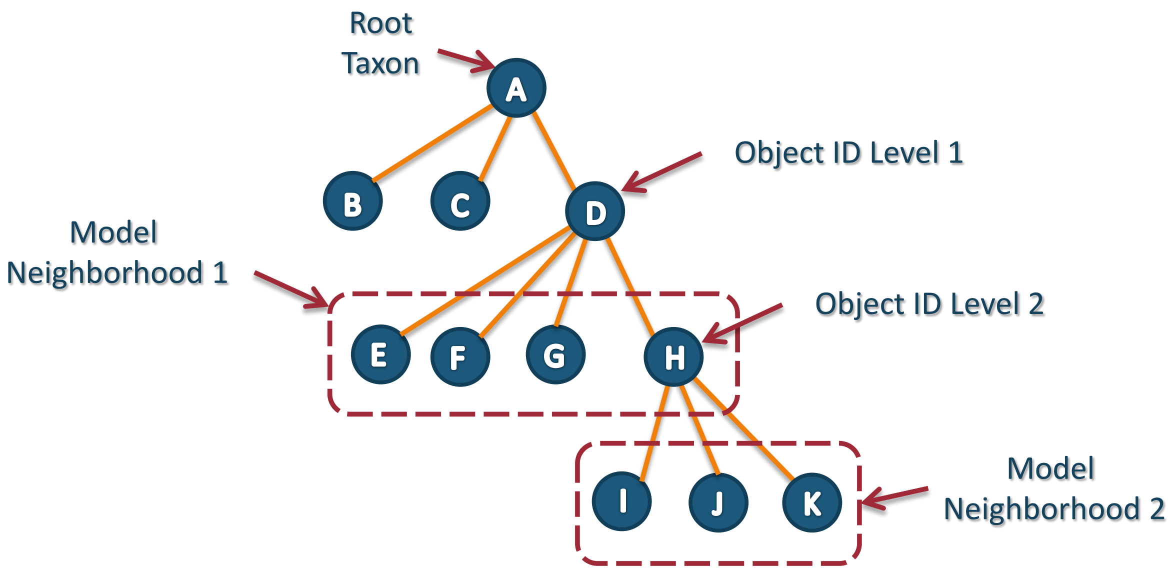

Figure 7: Hypothesis and model neighborhoods

4. Implications of Hierarchies for Model Neighborhoods

Referring back to our rhetorical questions, how do you know what hypothesis to choose in order to initialize a particular OD algorithm? Should you choose multiple hypotheses to cover all of the possibilities as is commonly done in the multiple hypothesis approach depicted in Figure 1? How do you know if your choices are an appropriately complete set of all possible hypotheses or just a random sampling? Consider Taxon D depicted in Figure 5, which has child taxa E, F, G, and H. If one has enough observational evidence to assign an object to Taxon D, or what we term object ID level 1, then the child nodes automatically becomes your model neighborhood. Each taxon in that model neighborhood represents a hypothesis. Each characteristic or feature of that taxon may have a stochastic representation which can also be parametrically varied. If one gathers additional observational evidence and is now able to assign an RSO to Taxon H, then taxa I, J, and K become the new model neighborhood that should be considered.

Instead of relying on experience or intuition to initiate a Monte Carlo or Multiple Hypothesis method, we now have a rigorous method to select neighboring models for consideration. We contend that this will significantly reduce the computational burden of multiple hypothesis methods because one will only consider a realistic set of variations. Furthermore, using the taxonomic tree branches, we can express a model distance metric that is rigorous and quantifiable. A model distance metric would be useful for optimal search strategies to aid in decision making. This is an open question of which various techniques can be proposed. Examples of model distance metrics could be the difference in probabilistic assignment to a hierarchical class, the statistics of individual model characteristics can describe the spread of classes, and using the count of branches and number of different levels between classes could all be considered. Future work will quantitatively explore these model distance metrics.

5. Implications of Hierarchies for Sensor Tasking

Hierarchies also present a method for sensor tasking. Consider the individual characteristic and feature variations between taxonomic classes. Recall again that these classes have been rigorously defined in such a way that one can always place an RSO into a class without ambiguity. If one has assigned an RSO to Taxon D in Figure 5 as before, instead of generating and enumerating multiple hypotheses, one could seek to obtain additional observational evidence to make the determination of which more specific child taxon (i.e. E, F, G, H) to which the object belongs. In order to accomplish this, one simply must examine the characteristic variations between taxa as they tell you exactly the kind of data required to make the determination. Instead of falling back on traditional methodologies where we try and make do with the data at hand, we now have a method to state unequivocally that specific additional data is required to appropriately classify an object. We contend that this sensor tasking metric does not exist in a rigorous form today. As such, we contend that hierarchies have obvious benefits for Space Object Identification (SOI) tasking. However, it is still an open question whether these types of hierarchies will be necessary and sufficient for both SOI and traditional metric tasking. Future work will explore this notion of sensor tasking using hierarchies in a quantifiable way.

6. Implications of Multiple Interacting Hierarchies

While we have described some significant benefits of utilizing hierarchical taxonomies to organize resident space objects, we must also consider some potential consequences of this approach. From our initial work, we embedded the RSO central body and orbit regime as a discriminating categorization at the most general level of the taxonomy. This allowed us to use the detectability of an RSO by a sensor as a discriminator. That is to say, sensors are typically designed to detect RSOs in a particular orbit regime thus the mere fact you are examining data from a particular sensor says something about the allowable orbit regimes. However, this type of construct can lead to a taxonomical hierarchy that is very large and deep, which poses computational challenges while searching complex tree structures for characteristic matches.

Furthermore, by the nature of the satellite business where new space vehicles are launched on almost a daily basis, whatever hierarchy one chooses can and should be evolving to account for new characteristics and features of these objects. Additional complications arise when one considers closely spaced objects that are each represented by large and deep hierarchical depictions. In order to disambiguate these objects rigorously, there may be interactions between these hierarchies that need to be properly accounted for.

Another approach to the problem of characterizing RSOs using a taxonomy is to acknowledge that there may be multiple instances of an RSO hierarchy (i.e. multiple detections) that can be said to interact with a terrain hierarchy (i.e. the sensor field of view). Instead of one large hierarchy that incorporates both orbit regimes and RSO characteristics, we will employ a separate taxonomy for each. Consider a simple thought example where we have a vehicle hierarchy that includes a taxon representing a car and a truck along with a terrain hierarchy with taxa representing highways, freeways, and residential roads. If one detects multiple trucks on a highway, one could assess the probability that those trucks belong to a convoy instead of acting independently of each other. If one happens to be looking at an airport freeway and detect multiple yellow cars, one might come to the conclusion that they are simply a line of taxi cabs rather than assume they are part of a convoy. This is the advantage of multiple interacting hierarchies. Utilizing this concept we can consider some larger concepts of how RSOs might be interacting with one another.

6. Summary and Conclusions

In this work, we showed that taxonomic hierarchies provide a unique and minimal RSO signature. Using these hierarchies, we can model the appearance of objects as well as reason with observations in a Bayesian construct. As a proof-of-concept, we demonstrated how even a rudimentary hierarchy can aid in space object identification using our Hierarchical Reasoning Tool (HRT). We went on to describe a more complete and robust RSO hierarchy that organizes sensor data, dynamic model choices, as well as the physical characteristics of the RSOs. By virtue of organizing information in a hierarchical manner, we discussed how one can generate hypotheses over neighborhoods of the model space as described by the parent-child relationships in the taxonomy. These hypotheses can be ranked according to their probabilistic confidence level based upon the priors and asserted evidence. Furthermore, hierarchies present a capability to efficiently task sensors for optimal information gain because we now know what the missing pieces of the SOI puzzle are before we task a sensor. We are no longer bound to a sensor tasking modality that says high priority objects must be observed by the first sensor that can detect the object. Instead, we can intelligently decide to collect observational data from only those sensors that will increase the probability of assignment to a more detailed taxon. We also discussed how multiple interacting hierarchies can potentially reduce the computational burden associated with large hierarchies and provide a deeper insight into how objects interact with one another.

7. References

[1] M. P. Wilkins, A. Pfeffer, P. W. Schumacher and M. K. Jah, "Towards and Artificial Space Object Taxonomy," in *2013 AAS/AIAA Astrodynamics Specialist Conferece*, Hilton Head, SC, 2013.

[2] M. Gupte, P. Shankar, J. Li, S. Muthukrishnan and L. Iftode., "Finding Hierarchy in Directed Online Social Networks," in *Proceedings of the 20th International Conference on World Wide Web*, 2011.

[3] D. M. Blei, A. Y. Ng and M. I. Jordan, "Latent Dirichlet Allocation," *Journal of Machine Learning*, vol. 3, pp. 993-1022, 2003.

[4] A. S. Bakhtiari and N. Bouguila, "A Hierarchical Statistical Model for Object Classification," in *IEEE International Workshop on Multimedia Signal Processing*, 2010.

[5] L. Fei-Fei and P. 2. Perona, "A Bayesian Hierarchical Model for Learning Natural Scene Categories," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2005.

[6] A. Sadovnik and T. Chen, "Hierarchical Object Groups for Scene Classification," in *IEEE International Conference on Image Processing (ICIP)*, 2012.

[7] M. Marszalek and C. Schmid, "Semantic Hierarchies for Visual Object Recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2007.

[8] A. Sun and E.-P. Lim, "Hierarchical Text Classification and Evaluation," in *IEEE International Conference on Data Mining*, 2001.

[9] J. J. Burred and A. Lerch, "A Hierarchical Approach to Automatic Musical Genre Classification," in *Proceedings of the 6th International Conference on Digital Audio Effects*, 2003.

[10] "http://en.wikipedia.org/wiki/Linnaean_taxonomy," [Online]. [Accessed 1 March 2012].

[11] B. Ruttenberg, M. P. Wilkins and A. Pfeffer, "Hierarchical Reasoning with Proababilistic Programming," in *Submitted to the 28th Conference on Artificial Intelligence (AAAI-14)*, Quebec City, Quebec, Canada, 2014.

[12] "Figaro Software and Tutorial," [Online]. Available: Available for download directly from Charles River Analytics: https://www.cra.com/commercial-solutions/probabilistic-modeling-services.asp. [Accessed 26 June 2013].

[13] C. Fruh, T. Schildknecht, M. Jah, T. Kelecy and P. Kervin, "Development of an Initiail Taxonomy and Classification Scheme for Artificial Space Objects," in *Sixth European Conference on Space Debris*, Darmstadt, Germany, 2013.

[14] C. Frueh, M. Jah, E. Valdez, P. Kervin and T. Kelecy, "Phylogenetic Taxonomy For Artificial Space Objects," in *Proceedings of the 24th Space Flight Mechanics Meeting, Paper #14-334*, Sante Fe, NM, 2014.

[15] T. Payne, S. Gregory and N. Houtkooper, "Long Term Analysis of GEO Photometric Signatures," in *2003 AMOS Technical Conference Proceedings*, Wailea, HI, 2003.

[16] T. Payne and S. Gregory, "Three-Dimensional Analysis of GEO Photometric Signatures," in *2004 AMOS Technical Conference*, Wailea, HI, 2004.

[17] T. Payne, S. Gregory and K. Luu, "Electro-optical Signatures Comparisons of Geosynchronous Satellites," in *Aerospace Conference*, Big Sky, MT, 2006.

[18] T. Payne, S. Gregory and K. Luu, "SSA Analysis of GEOS Photometric Signature Classifications and Solar Panel Offsets," in *2006 AMOS Tech Conference*, Wailea, Maui, HI, 2006.

[19] T. Payne, S. Gregory, J. Tombasco, K. Luu and L. Durr, "Satellite Monitoring, Change Detection, and Characterization Using Non-resolved Electro-Optical Data from a Small Aperture Telescope," in *2007 AMOS Tech Conference*, Wailea, Maui, HI, 2007.

[20] U. o. C. Scientists. [Online]. Available: http://www.ucsusa.org/satellite_database. [Accessed 1 February 2014].

[21] A. Pfeffer, "IBAL: A Probabilistic Rational Programming Language," in *Proceedings of International Joint Conference on Artificial Intelligence*, 2001.

[22] N. D. Goodman, V. K. Mansinghka, D. Roy, K. Bonawitz and J. B. Tenenbaum, "Church: a Language for Generative Models," in *Proceedings of Uncertainty in Artificial Intelligence*, 2008.