Comparing Apples and Orange Cottages

Classifications and Properties

Julia M. Taylor
Computer and Information Technology & CERIAS
Purdue University
West Lafayette, Indiana, USA
jtaylor1@purdue.edu

Victor Raskin
Linguistics & CERIAS
Purdue University
West Lafayette, Indiana, USA
vraskin@purdue.edu

***Abstract*—This paper deals with the rules of good classification and comparison, as well as matching the representation of the results with what has actually been accomplished. The emphasis in machine learning classifications, as well as, sometimes, outside of that paradigm, is almost exclusively on the precision of separating classes from each other, and hardly any effort is made to assess the nature of the classes with regard to their grain size. This results in a considerable disparity between the claimed results and what is really demonstrated, leading in turn to crude solutions to issues and poorly functioning applications. We propose an ontological solution, following the explicit tracing of a conceptual hierarchy underlying the classes. This approach may lead to a variety of solutions that can be compared after classification and similarity studies mature enough to face the issue.**

Keywords-comparison; classification; hierarchy; propert;, similarity; ontology; concept.

I. INTRODUCTION: SLOPPY CLASSIFICATIONS

Research, cognition, reasoning all involve some comparison, classification, similarity. Decisions on the bi-, tri- or multifurcation of a concept are common and inevitable. Statistical methods discover and refine unknown classifiers to divide a large bunch of samples into in- and outliers. Rule-based systems use rules to compare and classify. Are we doing it right? Do we know how to do it right? Is it useful to do it right?

The problem addressed in this paper is the status and (mis)interpretation of classifications and classifiers. Do the researchers have a clear picture of what they actually compare as opposed to what they want or purport to compare?. It addresses primarily the lack of attention and direct research effort in clarifying and codifying this problem—actually, an amazing lack of awareness that exists. Yet, grain size misclassification can have devastating effects on understanding the phenomena and question and the issues with them, as well as recommending precise solutions, pretty universally across research, from political and military solutions to treating bad cells in a patient.

The research question, then, that we are posing here, quite possibly for the first time so explicitly, is how to clarify and raise the precision of a proposed classification in just about any area of research. We will propose that the solution requires an ontological framework and a clear notion of grain size. The paper is not meant as a critique of the status quo with regard to the treatment of classification but rather to inform the diverse communities of scholars of a primising framework for improving that treatment. But first, a couple of examples of classification inexactitude, with consequences. Both come from areas of research where extensive scholarship has been done, including our own, but other than that, all they have in common is something they share with virtually any other area of research, namely, that they do classification and interpret their results.

Recently, we were asked to review a paper (not yet published, it may become officially citable soon) that used a machine-learning approach to separate serious text from satirical one. Not surprisingly, the results were statistically significant. Moreover, one could look at the features that were used to make the classification. A question to be asked is whether one should look at such features to shed light on the properties of satire. A simplified schema of humor classification in Figure 1 helps to see why it is not desirable (see Raskin et al. [1] for a discussion of the state of the art in humor theory and computational humor and for multiple references; cf. Raskin [2])..

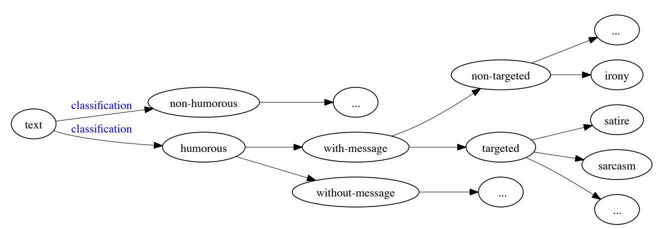

Figure 1. Simplified humor classification.

The question to be asked when one looks at the features is: was the distinction that was caught really the one between non-humorous text and satire? From the figure above, which most humor researchers will consider simplistic perhaps but plausible, satire is not just humorous but also containing a message that is targeted. It is also distinct somehow from irony and sarcasm. None of these features is mirrored on the non-humorous side, and there is a very serious risk of misinterpreting the results of the experiments. In all likelihood, what the classification captures is the distinction, at a much coarser grain size, between humorous and non-humorous text, the latter being a remote ancestor (hyperonym) of satire and impervious to the targeted-message nature of satire.

Similarly, in our own recent work on phishing (see Park et al. [3], Stuart et al. [4], and Park and Taylor [5], we compared *bona-fide*, legitimate email in the Enron corpus with a bunch of known phishing emails. We invested a considerable theoretical and methodological effort in the work and got reportable results. But, we dealt with a situation that is similar to satire, with phishing being the counterpart of satire (see Figure 2).

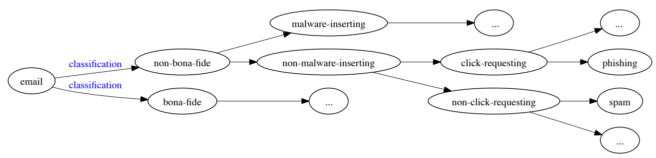

Figure 2. Simplified email classification.

The notions in Figure 2 are better definable than those in Figure 1; so in order to identify phishing, we needed first to have a corpus containing both bona-fide email from non-bona-fide emails. After separating those out, we needed to focus on the non-bona-fide corpus and separate malware-inserting emails from non-inserting; and then on, to a still narrower corpus and separating click requesting from non-click requesting. If the only click-requesting kind is phishing we are home, right? No, actually, there still exists *bona-fide* email that is click-requesting because our graph is actually not a tree but rather a lattice. In any case, we did not provide for any such complications, so we did just the first separation, and we failed to separate phishing from any other kind and sub-kind of non-bona-fide email. Our excuse, if any, is that we did not have enough corpora for the lower divisions. It is the same excuse as in the case of satire above. One does depend heavily on the availability of sufficient corpora but this is not a sufficient excuse for misidentifying the results.

The two examples above are sloppy classifications, and those are ubiquitous. Section II seeks help from adjacent disciplines, namely, the philosophy of science for theory building and research hygiene: as well as from psychology for similarity studies. Section III introduces the Ontological Semantic Technology (OST) whose property-rich ontology is a suitable base for rendering classification more rigorous and precise. We believe that more approaches will be developed to handle various meaning-based data- and text-processing applications, and that will be the time to compare OST to competition. We are not sure, however, that without a similar proper ontological base, a solution is possible. Section IV formalizes the OST approach, with a focus on ontological concepts and properties. The conclusion of Section V puts forth the down-to-earth application of the principle of rigorous comparison and application: where sole-property comparison is impossible or impractical, just explicating the property-set comparison may be a path to success. Given the paucity of effort in ensuring the grain size rigor of classifications and comparisons, the main contribution of this paper is drawing the wide community's attention to the issue of sloppy classifications, especially when the features are used to understand the nature of the crucial role of ontologies as remedy.

II. State of the Art

All research requires definitions, distinctions, comparisons, and classifications. The need to introduce categories and sub-categories is universal. Surprisingly, the state of the art on the precision of classification is minimal: there is no precision metric nor evaluation procedure for doing it right, and there is a definite, if not desperate need in both for virtually any area of research. In this section, we overview research on the philosophy of science that is supposed to contribute to theory building and psychology, mostly, cognitive psychology, on similarity (see references below). We briefly look for help in heuristics as well

A. *Philosophy of Science*

A brief look at the index pages of a couple of new readers in the field (Curd and Cover [6], Balashov and Rosenberg [7]) discovers a shared lack of any mentions of classification, comparison, distinction, separation, or hierarchy as worthy items of discussion. Hardly anything comes up on web and library searches. The last item does emerge in the context of biological classifications, and this should be expected: Linnaean classifications of the animal world, yet another unrelated domain of classifying, along with humor research and phishing mentioned above, have traditionally provided poster-child examples of straightforward sub-classifications, such as shown in Figure 3:

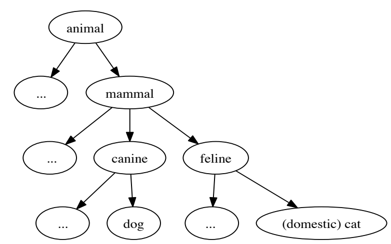

Figure 3. Simplified animal classification.

What Ereshefsky [8] states on the very inside cover, however, is as follows:

> "The question of whether biologists should continue to use the Linnaean hierarchy is a hotly debated issue.Invented before the introduction of evolutionary theory, Linnaeus's system of classifying organisms is based on outdated theoretical assumptions and is thought to be unable to provide accurate biological classifications.
>
> Marc Ereshefsky argues that biologists should abandon the Linnaean system and adopt an alternative that is more in line with evolutionary theory."

The customary advantage of the ancient classification is what was introduced and studied in the 20th-century mathematics as inheritance (Touretzky [9]): mammals inherit all properties of animals and add a few extra properties of their own; canines and felines inherit all of those, and each adds an extra set of additional properties; cats and dogs add another set of properties as well. As a result, a dog collects all the properties from animal (and its superclass, if any) to canine, as well as adding its own properties that other canines do not have.

Perhaps, one explanation of classifications and hierarchies not being actively discussed and researched is that, as per Potochnik and McGill [10],

> "The concept of hierarchical organization is commonplace in science and philosophical treatments of science. Though there are different applications of the concept of hierarchy, our primary focus here is the idea that material composition is hierarchical. Subatomic particles compose atoms, which compose molecules; cells compose tissues, which compose organs, which compose organisms; interbreeding organisms compose populations, which compose communities, which compose ecosystems; and so on. The basic idea is that higher-level entities are composed of (and only of) lower-level entities, but the prevalent concept of hierarchical organization involves stronger claims as well. The compositional hierarchy is often taken to involve stratification into discrete and universal levels of organization. It is also often assumed that levels are nested, that is, that an entity at any level is composed of aggregated entities at the next lower level."

The few references that are there to classifications, hierarchies, and levels in the contemporary philosophy of science seem to be all derived from an almost forgotten classic [Feibleman {11], p. 59], where the very first of the many rules establishing the hierarchy of "integrated levels" states that

> "[e]ach level organises the level or levels below it plus one emergent quality. Thus the integrative levels are cumulative upward. This proposition implies that everything has at least the physical properties and has led to the position of supreme importance of the physical world in science and philosophy."

Very characteristically to this strand, the whole philosophy of levels and hierarchies is limited to the physical world: the last sentence of the quote above limits it to physical objects, typically starting from bottom up with atoms and molecules. Potochnik and McGill [10] follows the same route, even though the paper applies this philosophy to ecology. Ereshefsky [8] is all about biology. So, Attardo and Raskin [12], an additional useful source on humor theory, had to do its own philosophy of science when it needed to establish a hierarchy of abstract levels of representation for a verbal joke in the General Theory of Verbal Humor on the principle of each higher level adding a restriction on the lower level, thus narrowing that latter's scope of included phenomena, as per Figure 4.

The integrative levels theory was, apparently, running high and ambitious in the mid-third of last century (see Bertalanffy and Woodger [13], Novikoff [14], and Bertalanffy [15]0, prompting Feibleman to hope, after Bertalanffy [15], for "a sort of super-science which shall have as its subject-matter the relations between the sciences. The philosophy of science may yet be the source for the development of an empirical field itself consisting of the integrative levels, a sort of meta-empirical field, with its own entities and processes and laws" (Feibleman [11], p. 59]. This has never happened, and this paper is suffering from the lack of helpful pertinent wisdom.

Figure 4. Levels of hierarchy of GTHV.

B. *Formal Ontology*

Some of what the philosophy of science could have delivered was contributed in formal ontology (see, for instance, Guarino and Poli [16]): a rigorous notion of hierarchy and inheritance, with a detailed study of subsumption. Never directly allied with the philosophy of science, it has been a blend of philosophy and logic, not focusing on the building nor application of actual ontologies and thus not involving itself in comparisons and classifications—just contributiung to a solid theoretical foundation for doing it right.

C. *Cognitive Psychology*

The main contribution that cognitive psychology has made for classification and hierarchy is a bit indirect: distinction and classification is closely related to the notion of similarity: note that, in any hierarchy, the subclasses of the same class share all the inherited properties and differ only in those extras that they add, and it was a high-powered strand of research on similarity and properties that flourished in cognitive psychology in several previous decades.

It was, apparently, Gregson [17] that put the measurement of (perceived) similarity on the map of psychological research. His thorough survey of similarity models, spatial and otherwise, did not focus, however, on the foundational notion of property that similarity must be based on, and it was the seminal Tversky [18: p. 330] that did. It proposed the general format for a property-based measurement function of similarity as "$s(a,b) = F(A \cap B, A - B, B - A)$," where "[t]he similarity of a to b is expressed as a function F of three arguments: $A \cap B$, the features that are common to both a and b; A — B, the features that belong to a but not to b; B — A, the features that belong to b but not to a."

In subsequent usage, the formula above has been often traded for a cruder but simpler measure as $A \cap B / A \cup B$, i.e., the intersection of the feature sets of a and b, divided by the union of these sets, standardly normalized to the [0,1] interval.

It is, however, Osherson et al. [19][20] that built a series of similarity models on a couple of subsets of a small animal dataset underlaid by a somewhat greater set of their properties, 48 and 85, respectfully for the latter paper. First having calculated the similarity measurements among the animals from the data set, using the simplified metric above,

Copyright (c) IARIA, 2015. ISBN: 978-1-61208-390-2

they conducted an experiment with 10 human subjects and compared their similarity judgments with those predicted by their model.

There are two aspects of this research that are of a particular interest to us here. First, the origin of the properties used as the foundation of the similarity model: they were compiled by the researchers "empirically" and offered to the 10 subjects with the instruction to detect and suggest the addition of any new property other than the sounds the animals made. An additional property would only be taken into consideration if proposed by more than one subject. We proposed an ontological foundation for our (Taylor and Raskin[19]).

Second, none of Osherson's and his associates' papers over almost two decades, directly based on the animal dataset or following from earlier research on it, had the similarity model as the goal. In fact, the models were obtained to be used as tools in research on human reasoning, such as inductive judgments (Osherson et al. [19]), default probabilities (Osherson et al. [20]), the conjunction fallacy (Tentori et al. [22]). Later related work (Perfors et al. [23], Kemp et al. [24]), using their own variations of animal datasets and properties, applied their models to research the way children learn "domains" and "theories of the world," respectively. We are also interested not so much in similarity models as in the nature of properties that are out there in the world and that people reason with, thus necessitating the need to computerize those properties for the purpose of constructing a meaning-based structure from text and other data.

D. Heuristics

Our best help should have probably come from this step-daughter (hopefully, Cinderella) of mathematics, pretty much completely ignored by other disciplines. The insights in the old classic Polya [25] and the newer classic Pearl [26] should inform theory-building significantly. In fact, heuristics should be the basis of any graduate course or seminar on research methods on top, if not even instead, pure statistics that most universities offer exclusively. Unfortunately, heuristic ideas are hard to pack in off-the-shelf software, and abduction, on which much heuristics rests, has not been able to compete with deduction and induction, instead of its ubiquity and scope, for the minds of scientists and other scholars.

III. Ontology, Hierarchy, and Good Classification

A. Ontology

The ontology comes from our particular approach to computational semantics the Ontological Semantic Technology (OST). The theory-cum-technology is a radical revision and improvement of Nirenburg and Raskin [27];see Raskin et al. [28], Taylor et al. [29], Taylor and Raskin [30], Hempelmann et al. [31], Taylor et al. [32][33]. The centerpiece is indeed the language-independent semi-automatically constructed engineering ontology, as per Gruber [34], consisting of concepts (OBJECTs and EVENTs) linked with a rich system of PROPERTYs. Each supported natural language, e.g., English and Russian shown on Figure 5, has a lexicon with supporting morphological, and syntactic rules, supplemented with phonological rules (not shown), when required by an optional speech recognition functionality.

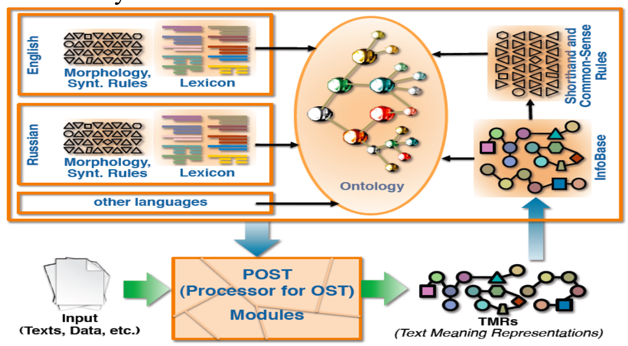

Figure 5. OST Architecture.

A lexicon contains all the lexical entries for the language, each entry with all of its senses. The central components in an entry are the partial syntactic information (SYN-STRUC) and full semantic information (SEM-STRUC). The latter typically "anchors" the sense of a lexical item in the appropriate ontological concept, restricts some of its property fillers if necessary, and binds the variables introduced in the sense's SYN-STRUC. The ontology captures much information about how things are in the single or multi-domain world it serves. The ontology is supplemented by the previously processed information from the InfoBase and by dynamically collected shortcut and commonsense rules on which the machine and a human ontology engineer collaborate.

The OST various processors operate on the anchoring results. When a sentence arrives at its input, the analyzer looks up every word in the lexicon, checks that its usage conforms to the SYN-STRUC, selects the corresponding SEM-STRUC, identifies the event in each clause, and then attempts to match all the other concepts evoked by the words and phrasals as the fillers of the event's properties. The result, the Text Meaning Representation (TMR) of the input sentence is stored in the InfoBase for further usage, including possible correction or challenge by the later arriving text.

B. Hierarchy

The OST ontology, like most real and pseudo-ontologies, is based on subsumption, which means that its IS-A property is privileged to pass on properties from higher-level (parent) nodes to low-level (child) nodes, as shown in Figure 3. The ontology is not a tree, however: rather, it is a lattice. This means occasional complications to inheritance. Also, very rarely, a property to be inherited have to be blocked as, for instance, continuing with the convenient animal world, ostriches and chickens should inherit all properties of birds except the ability to fly.

C. Good Classification

A good classification is a minimal, carefully controlled deviation from the ideal classification, a deviation which occurs only when necessary. The ideal classification is

 ISBN: 978-1-61208-390-2

achieved by a hierarchy, in which each child adds one simple ontological property to the ones it inherits from its parent and, thus, from its entire ancestry. In reality, what is added is most often a set of properties. An ontological approach allows us to be fully aware of what the set consists of and, if necessary, to entice us to separate its property elements out in additional computer experiments. Our ongoing work on composable properties (Taylor and Raskin [35]) promises further progress in this direction.

IV. A BIT OF FORMALISM

For the purposes of this discussion, let us assume that each concept C can be defined as a combination of properties P_1, P_2, …, P_n. Each of these properties can be further restricted by specifying a particular argument to a property. For example, a concept HUMAN (see Figure 6) can have a property GENDER, and for the sake of simplicity, let us assume that the range of this property can be either MALE or FEMALE. In order to define the concept WOMAN, one would need to restrict the property GENDER from its most generic case to that of only FEMALE.

Now, let us consider a more general case. Suppose a concept C1 can have a property P1. Each of its children, C2.i will also have a property P1, restricted by a particular filler -- let us say filler a_i -- as well as inheriting the rest of the properties that C1 inherited, each with the restrictions done for the parent (see Figure 7). Children of C2.i will also have some property, lets call it P2.i, that will be restricted by some fillers, thus introducing a new layer of descendent concepts, all of which will still inherit P1(a_i). Eventually, we will run into a situation where a concept Cj,k is composed by j-1 properties, each of which is restricted by at least once when passed from a ancestor to a descendant. It is possible that concept Cn.i (see Figure 7), which inherited a property P2.3 with filler a8 needs to have a more specific filler than that of a8. Then, the property Pn.i will be the same as property P2.3, but the filler of the property for the children of Cn.i will be have to be children of a8.

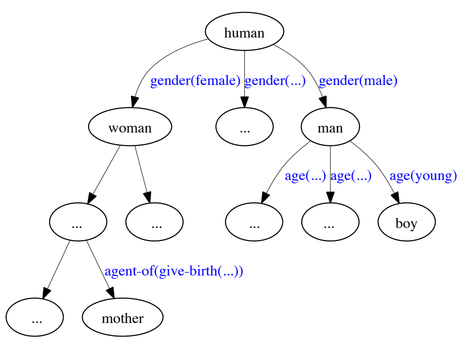

Figure 6. Hierarchy example.

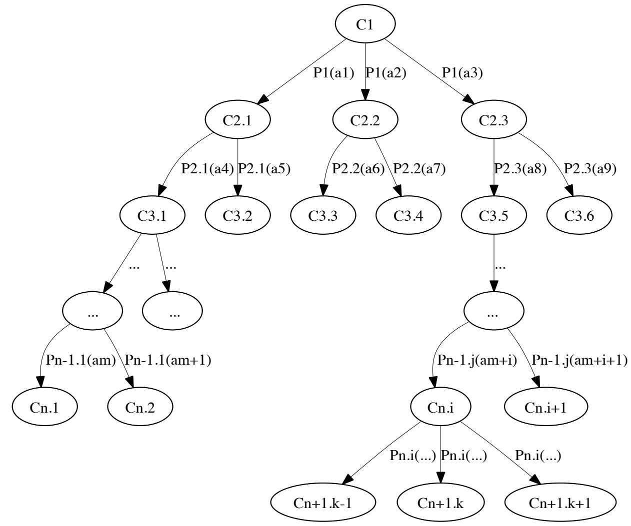

Figure 7. Hierarchy based on properties.

Within such a structure, the concept C3.1 differs from the concept C3.5 by a filler of property P1, as well as by properties uniquely introduced for these concepts, namely, P2.1 and P2.3. Perhaps, it will be clearer in a more concrete example, outlined in Figure 6. The concept MOTHER differs from the concept BOY by a filler of the property GENDER, as well as by the properties that lead to the concept MOTHER, and those that lead to concept BOY. In Figure 6, the only ones that are visible are AGE and AGENT-OF. It is possible that some of these properties are the same – we can use our common sense here—both MOTHER and BOY do have some AGE, and both are likely to be AGENT-OF something. In that case, we will say that property Px and Py are the same.

It is also possible that P1 is entirely inherited by the child concept, without any restriction, but then there will be another property Pi that the child will restrict, otherwise the child and parent concept run the danger of being identical.

If we are not dealing with a tree, but with a lattice, there is more flexibility in building a graph that would absorb the common properties within their parents, whether it would result in multiple parenting scenarios or not. There are many formalisms that allow such lattice to be build and construct it automatically.

Two concepts Ci and Cj, then, can be looked at in terms of 3 sets of properties: one set that is shared between them, one set that Ci has, but not Cj and one set that Cj has but not Ci. To complicate matters, the shared set is likely to have different fillers of the properties, and thus we have to subdivide this set into properties that have identical fillers between Cj and Cj and properties that have different fillers.

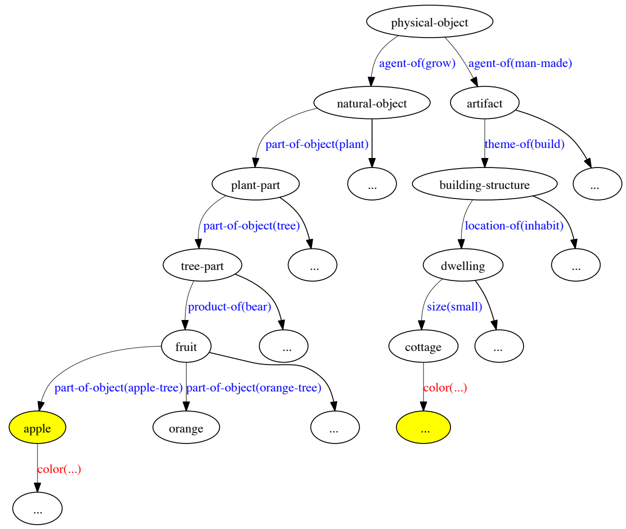

Figure 8. A sample structure to lead to apple and cottage concepts.

As an example, lets us compare apples and orange cottages, as highlighted in Figure 8. The overlapping property set between orange cottage and apple (also orange?) is {AGENT-OF, COLOR}. The other two sets are {PART-OF-OBJECT, PRODUCT-OF} and {THEME-OF, LOCATION-OF, SIZE}. If we are not paying attention, it could be said that a differentiating feature between apples and orange cottages is the fact that one of them is part of a tree, and the other one is not. While that is definitely true, any dwelling is not part of a tree, and neither is any artifact. This happens because the PART-OF-OBJECT property is in the set of properties that only one of the compared concepts has. On the other hand, we could compare orange cottages and apples on the property COLOR (e.g., the color of this apple is exactly like the color of my cottage).

Notice also that the closer the concepts are in their hierarchy, the easier it is to compare them – because, again, the overlapping property set is large. Thus, we can compare apples and oranges on many more parameters than apples and cottages (orange or not). Of course, comparing apples and apples is even better – we can just count the properties in the overlapping set.

If we can rely on such sets of properties and hierarchies, it is easy to see why a whale can be both compared to mammals and fish, even if we (as young children?) may not know where exactly it belongs in the hierarchy.

V. CONCLUSION: CUI BONO?

In this paper, we discussed the hygiene of good classification and comparison and suggested that an ontological foundation would considerably clarify the issue. Even if comparison on one sole property is not attainable or even desirable, and subclasses have to differ from their ontological parents by a set of features, it is useful to be fully aware of it and be prepared to un-bunch them if necessary. It is also crucially important not to misstate nor to misrepresent the result by a hierarchical confusion. Let us not compare apples with orange cottages without figuring out explicitly what properties separate them and their grain size correspondence.

REFERENCES

[1] V. Raskin and J. M. Taylor, “On the tansdisciplinary field of humor research,” Journal of Integrated Design and Process Science 16(3), pp. 133-148, 2013

[2] V. Raskin (ed.), The Primer of Humor Research, Berlin: Mouton de Gruyter, 2008.

[3] G. Park and J. M. Taylor, “Towards text-based phishing detection,” Proc. SDPS-13. Sao Paolo, Brazil, 2013.

[4] L. M. Stuart, G. Park, J. M. Taylor, and V. Raskin, “On identifying phishing emails: Uncertainty in machine and human judgment,” Proc. NAFIPS-14, Boston, MA, 2014.

[5] G. Park, L. M. Stuart, J. M. Taylor, and V. Raskin, “Comparing machine and human ability to detect phishing emails,” Proc. IEEE-SMC-14. San Diego, CA, 2014.

[6] M. Curd, J. A, Cover, and C, Pincock, eds., Philosophy of Science: The Central Issies, 2nd ed. New York: Norton, 2013.

[7] Y. Balashov and A. Rosenberg, eds., Philosophy of Science: Contemporary Readings. London-New York, Routledge, 2002.

[8] M. Ereshefsky, The Poverty of the Linnaean Hierarchy: A Philosophical Study of Biological Taxonomy. Cambridge, UK: Cambridge University Press, 2001.

[9] D. S. Touretzky, The Mathematics of Inheritance Systems. Los Altos, CA: Morgan Kauffman, 1996.

[10] A. Potochnik and Brian McGill, “The limitations of hierarchical organization,” Philosophy of Science 79:1, 2012.

[11] J. K. Feibleman, “Theory of integrative levels,” British Journal for the Philosophy of Science 5 (17), 1954, pp. 59-66.

[12] S. Attardo and V. Raskin, “Script theory revis(it)ed: Joke similarity and joke representation model,” HUMOR: International Journal of Humor Research 4:3-4, 1991.

[13] L. von Bertalanffy and J. H. Woodger, Modern Theories of Development. Oxford: Oxford University Press, 1933.

[14] A. B. Novikoff, “The concept of integrative levels and biology, “ Science 101, 1945, pp. 209-215.

[15] L. von Bertalanffy, “An outline of general system theory,” The British Journal for the Philosophy of Science I-134, 1950.

[16] N. Guarino and R. Poli, Eds., Special Issue: The Role of Formal Ontology in the Information Technology, International Journal of Human and Computer Studies 43:5-6, 1995.

[17] R. Gregson, Psychometrics of Similarity. New York: Academic Press, 1975.

[18] A. Tversky, “Features of similarity,” Psychological Review 84: 4, 1977, pp. 327-352.

[19] D. N. Osherson, O. Wilkie, E. E. Smith, A. Lopez, and E. Shafir, “Category-based induction,” Psychological Review 97:2, 1990, pp. 185-200.

[20] D. N. Osherson, J. Stern, O. Wilkie, N. Stob, and E. E. Smith, “Default probability,” Cognitive Science 15, 1991, pp. 251-269.

[21] J. M. Taylor and V. Raskin, “Understanding and structuring language descriptions: The case of 101 animals,” Proc. IEEE SMC 2012. Seoul, S. Korea, 2012.

[22] K. Tentori, N. Bonini, and D. Osherson, “The conjunction fallacy: A misunderstanding about conjunction?” Cognitive Science 28, 1990, pp. 467-477.

[23] A. Perfors, C. Kemp, and J. B. Tenenbaum, “Modeling the acquisition of domain structure and feature understanding, “ Proc. CogSci-05, 2005.

[24] C. Kemp, J. B. Tenenbaum, T. L. Griffiths, T. Yamada, N Ueda, “Learning systems of concepts with an infinite relational model,” Proc. AAAI-06, 2006.

[25] J. Polya, How to Solve it: A New Aspect of Mathematical Method. Princeton, NJ: Princeton University Press, 1945.

[26] J. Pearl, Heuristics: Intelligent Search Strategies for Computer Problem Solving. New York: Addison-Wesley, 1984.

[27] S. Nirenburg and V. Raskin, Ontological Semantics. Dordrecht: D. Reidel, 2004.

[28] V. Raskin, C. F. Hempelman, J. M. Taylor, “Guessing vs. knowing: The two approaches to semantics in natural language processing,: Proc. Dialogue 2010. Bekasovo/Moscow, Russia, 2010, pp. 642-650.

[29] J. M. Taylor, C. F. Hempelmann, V. Raskin, “On an automatic acquisition toolbox for ontologies and lexicons in ontological semantics,” Proc. ICAI-10, Las Vegas, NE, 2010, pp. 863-869.

[30] J. M. Taylor and V. Raskin, Understanding the unknown: Unattested inpiut in natural language,” Proc. FUZZ-IEEE-11. Taipei, Taiwan, 2011.

[31] C. F. Hempelmann, J. M. Taylor, V. Raskin, “Application-guided ontological engineering, “ Proc. ICAI-10. Las Vegas, NE, 2010.

[32] J. M. Taylor, V. Raskin, C. F. Hempelmann, “From disambiguation failures to common-sense knowledge acquisition: A day in the life of an ontological semantic system,” Proc. WI-IAT-11. Lyon, France, 2011.

[33] J. M. Taylor, V. Raskin, C. F. Hempelmann, “Towards computational guessing of unknown word-meanings: The ontological semantic approach,” Proc. CogSci-11. Boston, MA, 2011.

[34] T. R. Gruber, “Toward principles for the design of ontologies used for knowledge sharing,” in: [14], 1995, pp. 907-928.

[35] J. M. Taylor and V. Raskin, “On the nature of composable properties,” Proc. CCAHI Workshop, AAAI-14, Quebec City, Canada 2014.