

Habitat-Lite: A GSC Case Study Based on Free Text Terms for Environmental Metadata

Lynette Hirschman,¹ Cheryl Clark,¹ K. Bretonnel Cohen,¹ Scott Mardis,¹ Joanne Luciano,¹ Renzo Kottmann,² James Cole,³ Victor Markowitz,⁴ Nikos Kyrpides,⁵ Norman Morrison,⁶ Lynn M. Schriml,⁷ Dawn Field,⁸ and the Novo Project⁹

Abstract

There is an urgent need to capture metadata on the rapidly growing number of genomic, metagenomic and related sequences, such as 16S ribosomal genes. This need is a major focus within the Genomic Standards Consortium (GSC), and *Habitat* is a key metadata descriptor in the proposed “Minimum Information about a Genome Sequence” (MIGS) specification. The goal of the work described here is to provide a light-weight, easy-to-use (small) set of terms (“Habitat-Lite”) that captures high-level information about habitat while preserving a mapping to the recently launched Environment Ontology (EnvO). Our motivation for building Habitat-Lite is to meet the needs of multiple users, such as annotators curating these data, database providers hosting the data, and biologists and bioinformaticians alike who need to search and employ such data in comparative analyses. Here, we report a case study based on semiautomated identification of terms from GenBank and GOLD. We estimate that the terms in the initial version of Habitat-Lite would provide useful labels for over 60% of the kinds of information found in the GenBank isolation_source field, and around 85% of the terms in the GOLD habitat field. We present a revised version of Habitat-Lite defined within the EnvO Environmental Ontology through a new category, EnvO-Lite-GSC. We invite the community’s feedback on its further development to provide a minimum list of terms to capture high-level habitat information and to provide classification bins needed for future studies.

Introduction

THIS PAPER DISCUSSES the current status of an ongoing effort to create a minimum hierarchical controlled vocabulary for the capture of habitat and environmental metadata on genomics, metagenomics, and 16S ribosomal sequences. This work has two goals. The short-term goal is to develop a light-weight controlled vocabulary (Habitat-Lite) within the EnvO framework to capture high-level habitat and environmental metadata in support of the Genomic Standards Consortium (GSC) Minimal Information about Genome/Metagenome Sequence (MIGS/MIMS) specification (Field et al., 2008a, 2008b). The longer-term goal is to develop a repeatable process for other types of metadata by identifying key terms based on us-

age in databases and the open literature. We will evaluate the coverage, utility, and usability of the key terms and refine the set of terms based on these measures. Additionally, we will develop tools to facilitate the capture of the metadata from free text fields.

This effort originated in the context of the development of the MIGS/MIMS checklist (http://gensc.org/gc_wiki/index.php/MIGS/MIMS), and has also been discussed in the context of the newly established Environment Ontology (EnvO) project (<http://environmentontology.org>—see the GSC EnvO wiki page for ongoing discussion: http://gensc.org/gc_wiki/index.php/EnvO_Project; also see the EnvO SourceForge site: <http://obo.cvs.sourceforge.net/obo/obo/ontology/environmental/>), as part of advocating the use of ontologies

¹Information Technology Center, The MITRE Corporation, Bedford, Massachusetts.

²Microbial Genomics Group, Max Planck Institute for Marine Microbiology and Jacobs University Bremen, 28359 Bremen, Germany.

³Center For Microbial Ecology, Michigan State University, East Lansing, Michigan.

⁴Biological Data Management and Technology Center, Lawrence Berkeley National Laboratory, Berkeley, California.

⁵Department of Energy, Joint Genome Institute, 2800 Mitchell Drive, Walnut Creek, California.

⁶School of Computer Science, University of Manchester, Oxford Road, Manchester, United Kingdom.

⁷Institute for Genome Sciences and Department of Epidemiology and Preventive Medicine, University of Maryland School of Medicine, HSF1, 685 West Baltimore Street, Baltimore, Maryland.

⁸NERC Centre for Ecology and Hydrology, Mansfield Road, Oxford, Oxfordshire, United Kingdom.

⁹<http://environmentontology.org>

in capturing MIGS/MIMS reports. See also GazO, http://gensc.org/gc_wiki/index.php/GAZ_Project, a first step towards an open source gazetteer, constructed on ontological principles, that describes places and place names and the relations between them. This work is informing GSC consensus-building activities, and has led to agreement to adopt the Habitat-Lite terminology for use in the Genomic Contextual Data Markup Language (GCDML) (Kottmann et al., 2008).

There is a strong need for developing methods to facilitate the capture of metadata describing the growing number of genomic and metagenomic projects, including information about isolation source and habitat (Field et al., 2008a; Morrison et al., 2006). The increase in the associated literature is also accelerating, particularly in light of projects such as the Global Ocean Survey (Venter et al., 2004) and the Human MicroBiome (<http://nihroadmap.nih.gov/hmp/>) (Gill et al., 2006), with parallel growth in the relevant databases [see, e.g., Fig. 1 of Morrison et al., 2006, for an illustration of exponential growth in the number of sequences in the International Nucleotide Sequences Database Collaboration (INSDC)]. However, the capture of the metadata associated with these projects remains a major challenge, largely due to the fact that the literature is scattered and the metadata is difficult to find, even by expert manual extraction. Many databases have fields to support the capture of metadata, but such entries are often sparse and are entered as free text, thus lacking standardization in vocabulary and definitions, impeding our ability to perform meaningful comparisons or utilize information from multiple resources. The case studies discussed below illustrate the resulting difficulty in using computational techniques to study the relation between habitat and genotypic or phenotypic properties of organisms (Hunter, 2002, von Mering, 2007)—a key goal of genomic and metagenomic studies.

Our initial work has focused on a specific metadata type, namely habitat. For our purposes here, we define habitat as

“the place or environment where an organism naturally or normally lives and grows.” It is distinguished from “sample source,” which is the environmental context in which a sample is collected, as defined in Morrison et al. (2006). Multiple habitat terms can be associated with a species; by contrast, a sample is associated with a description of its (unique) source. Table 1 shows excerpts from the GOLD database (Liolios et al., 2008); we can see that the “Habitat” field often has multiple entries, in contrast to the “Isolation” field, which describes the specific sample source and is much more detailed. The initial version of Habitat-Lite is aimed at capturing high-level habitat descriptions; ongoing work on the environmental ontology EnvO will provide a much finer-grained set of terms to describe specific environments and sample source information.

The development of Habitat-Lite began with the selection of a small list of widely used high-level terms for describing habitat. We used these terms to “bin” information contained in free text fields for habitat or source information in several key databases. This process enables us to develop measures of coverage, utility, and usability for the term set—for example, how well the controlled vocabulary covers the free text entries, how evenly the entries are distributed across the bins defined by the controlled vocabulary, how well the bins capture useful categories for search, how cost-effectively the controlled vocabulary terms can be used to annotate new data, and how consistent the mappings are across multiple annotators (human or automated). There are trade-offs in this complex space between the detailed information that can be captured with a large well-structured set of terms (e.g., an ontology), versus the time it takes to create a stable set of structures and the cost of acquiring consistent annotation using this much richer terminology, including supporting tools.

The two major data sources chosen for this study contain large numbers of records and descriptors of habitat in free text form. Ideally, we would have looked in the literature to

TABLE 1. HABITAT AND ISOLATION FIELDS FROM THE GOLD DATABASE

<i>Organism</i>	<i>Strain</i>	<i>Phenotype</i>	<i>Habitat</i>	<i>Isolation</i>
<i>Hemophilus influenzae</i> NTHi	PittEE	Pathogen, facultative, nonmotile, rod-shaped	Host	Middle-ear effusion of a child in Pittsburgh
<i>Mycobacterium tuberculosis</i>	H37Ra	Pathogen, aerobe, chemoorganotroph, rod-shaped, nonmotile	Host, TB epidemic	Original human-lung H37 isolate in 1934
<i>Psychrobacter</i> sp.	PRwf-1	Psychrophile, radiation resistant, rod-shaped, nonmotile	Aquatic, soil, Permafrost	
<i>Roseiflexus</i> sp.	RS-1	Filament-shaped, photosynthetic, thermophile, facultative, nonsporulating, motile, rod-shaped	Aquatic, hot spring	Hot spring microbial mat
<i>Lactobacillus reuteri</i>	F275 (JCM 1112)	Probiotic, non-pathogen, rod-shaped, facultative, nonmotile	Intestinal flora	Human isolate that is unable to colonize the intestinal tract of mice
<i>Pseudomonas putida</i>	F1	Aerobe, motile, rod-shaped, non-pathogen	Soil	Polluted creek in Urbana, Illinois, by enrichment culture with ethylbenzene as a sole source of carbon and energy

determine how habitat and isolation source were described. However, for the initial experiments, it was much more efficient to look at fields in existing databases. The two sources were:

1. GenBank (<http://www.ncbi.nlm.nih.gov/Genbank/>): the *isolation_source* field, which captures free text descriptions, in the form entered by submitters to GenBank, related to sample source;
2. Genomes On-Line Database (GOLD) (<http://www.genomesonline.org/>) (Liolios et al., 2008): the *Habitat* field, which captures terms collected from the literature.

Development of Habitat-Lite

As a starting point, one author (D.F.) did a survey for terms used in a number of relevant sources. From this list, she selected a set of high-level terms as a strawman for the first iteration of the Habitat-Lite term list (shown in Table 2). The number of terms was kept small (less than 20), based on discussions with annotators at NCBI (we met with Tatiana Tatusova, Scott Federhen, Karen Clark, and Anji Johnston at NCBI Entrez Genomes, to explore ways to improve the capture of environmental/habitat metadata in GenBank), but could grow in future iterations. Our approach was to identify a set of seed terms, run experiments to determine how well these could “bin” existing entries, determine how useable such a set of terms would be for human and semiautomated annotation, and then iterate, with the goal of producing a consensus-driven “minimal set” of habitat terms that provided good coverage of entries in key resources. Both the original version and the new version of Habitat-Lite using the EnvO-Lite-GSC category are available in .obo format (http://gensc.org/gc_wiki/index.php/Habitat-Lite) and, for example, could be used with OBO-Edit (<http://oboedit.org/>), CoBrA (<http://cobra.umbc.edu/eclipse/>), or the Phenote a notation tool (available at <http://www.phenote.org/>).

The initial list of terms drew on previously published lists of habitat terms used to annotate databases (NCBI Microbial genomes; <http://www.ncbi.nlm.nih.gov/genomes/lproks>.

<http://www.jgi.doe.gov/16s/saiform.php>), on the habitat terms published in the Global Ocean Survey (Nealson and Venter, 2007), on habitat terms used to describe the biases in culture collection strains (Floyd et al., 2005), and on patterns and biases in the complete genome collection (Martiny and Field, 2005); see Supplementary Table 1 for a full listing. These terms were mapped to an early version of the Environment Ontology (<http://environmentontology.org>), as shown in column 3 of Table 2.

Use Cases: Analyses Based on Habitat Data

Habitat-Lite terms were assembled from existing terminologies with the explicit goal of supporting as many use cases as possible—in particular, the ability to “bin” data into interesting categories for purposes of comparison. The use of bins is particularly attractive to biologists, who, for example, wish to extract sequences only associated with “soil bacteria” or “freshwater metagenomes.” In this respect, biologists’ descriptions of “habitat” contrast strongly with those of environmental scientists, who tend to describe habitat in terms of continuous variables.

We are now in the process of assembling use cases to test the coverage of Habitat-Lite. At the fifth GSC meeting, one author (J.C.) presented a small study done on the Ribosomal Database Project (RDP; <http://rdp.cme.msu.edu/>; Cole et al., 2007). The RDP consumes GenBank documents for 16S sequences and maintains them in a highly value-added format. These data are used extensively for contrastive analysis based on environmental factors. To determine both the coverage of environments and the utility of the habitat or environmental information in RDP, a small experiment was carried out in late 2006.

Using information from the INSDC records, one author (J.C.) attempted to manually classify into habitats the 168,911 rRNA sequences marked as *environmental* in RDP release 9.44 (November 2006). The habitat categories that were suggested by Phil Hugenholz (DOE Joint Genome Institute) were modified by splitting *host-associated* into separate categories for *plant* and *animal* (including *human*) *associated*. We first assigned 24.5% of the sequences using their *isolation_source* qualifier. For those sequences without an isolation source tag, or where we were unable to classify based on that tag, we examined the reference titles from the INSDC records and were able to classify another 37.5% of the records. References used by fewer than 150 sequences were not examined because of the effort involved. The remaining 38% of sequences could not be classified because, for the most part, they did not have any habitat information in the INSDC record. Most assignments were made after examination by a single researcher, but spot-checking by a second researcher gave disagreement in assignment for only a small percentage of sequences. By far the biggest category was *animal associated*, and a large fraction of these were *human associated*. The *soil*, *sediment*, and *water* categories also represented large numbers of sequences (see Fig. 1).

A second interesting use case was reported in (von Merling et al., 2007). In this paper, the authors studied the association of preferred habitats for microbial clades and looked for correlations between evolutionary distance and similarity of habitat. The habitat information was taken from free

TABLE 2. INITIAL HABITAT-LITE TERMS AND MAPPINGS TO ENVO (OCT. 2007)

Habitat-Lite Terms for genomes and metagenomes	
1 freshwater habitat	ENVO:00000873
2 marine habitat	ENVO:00000447
3 terrestrial habitat	ENVO:00000446
4 soil	ENVO:00001998
5 water	ENVO:00002006
6 air	ENVO:00002005
7 sediment	ENVO:00002001
8 sludge	ENVO:00002044
9 waste water	ENVO:00002006
10 hot spring	ENVO:00000051
11 hydrothermal vent organism-associated	ENVO:00000215
12 habitat	ENVO:00002032
13 extreme habitat	ENVO:00002020
14 food	ENVO:00002002
15 biofilm	ENVO:00002034
16 microbial mat	ENVO:01000008
17 fossil	ENVO:00002164

**Isolation Source of Environmental Sequences
from RDP Release 9.44
(October 30, 2006)**

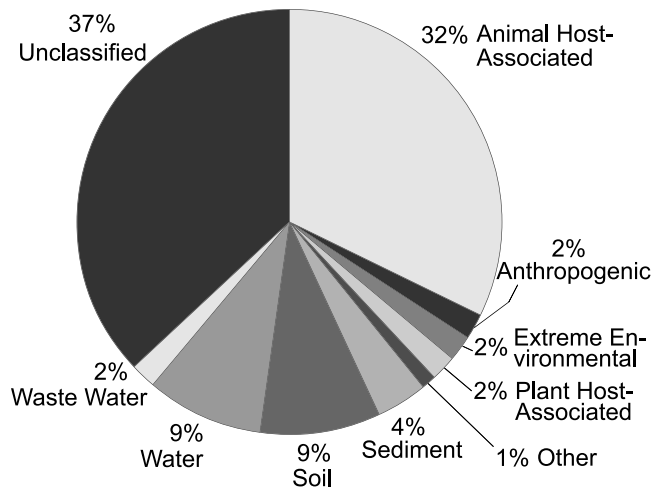


FIG. 1. Categorization of isolation source for environmental sequences from RDP.

text fields of the Greengenes database (Desantis et al., 2006) and the microbial culture collections (Dawyndt et al., 2005). To assess similarity of habitat, the authors manually selected “informative” words found in the annotation of five or more

experiments (von Mering et al., 2007, Tables S2 and S3, supplemental materials) and computed a pairwise similarity score between habitats, based on number of shared keywords. Graphs in Table S2 Figure 2B and 2C show that more habitat “features” are shared among the more closely related organisms, both in terms of taxonomy and molecular similarity.

These use cases illustrate the kinds of information that would be useful to researchers, and also the difficulties of obtaining the information in the absence of a common underlying controlled vocabulary.

GenBank “isolation_source” Entries

To validate and refine the selection of Habitat-Lite terms, examples of habitat or isolation_source information were needed to determine what information was present, and how this information was expressed. An ideal approach would be to extract metadata from the published literature; however, this is quite difficult, because the metadata occurs in many diverse forms, including PDF tables, densely written materials and methods sections, supplementary material, and even in referenced work. Therefore, we took advantage of the large quantity of free text metadata already available—as fields in database records. As a first step, we analyzed the “isolation_source” field from GenBank gene records, which captures, as short free text entries, information about the isolation source of the specific sequence being deposited. John Wilbur (NCBI) provided us with a list of 35,000 distinct isolation_source entries from GenBank gene records as of September 2007—see Table 3 for examples of some entries from

TABLE 3. DISTRIBUTION OF UNIQUE HABITAT-LITE TERMS IN GENBANK ISOLATION SOURCE FIELDS

Class	Table frequency	Data set frequency	Percent total	Example
Organism-associated	14,781	341,003	42.4%	1-year-old male spleen
Water/aquatic	2,008	40,794	5.8%	0 m water at a station in the North Atlantic
Soil	1,115	229,032	3.2%	0–20-cm bulk soil from a mixed forest
Marine	944	3,115,879	2.7%	0.2–0.8 μm fraction from surface sea water
Sediment	723	34,435	2.1%	aquaculture coastal sediments
Terrestrial	594	3,100,550	1.7%	a declining forest
Food	398	4,003	1.1%	(onion)
Sludge	294	9,868	0.8%	1st maturation stage of sludge
Microbial mat	195	9,164	0.6%	a deep sea microbial mat
Waste water	195	5,969	0.6%	activated tannery effluent from treatment plant
Hydrothermal vent	133	3,036	0.4%	14 N Mid Atlantic Ridge Logatchev vent field
Hot spring	121	3,249	0.3%	6–48 celsius region of a hot spring
Extreme	117	4,967	0.3%	a solar saltern
Biofilm	114	3,499	0.3%	aquatic phototrophic biofilm
Freshwater	75	2,609	0.2%	Arctic freshwater lake
Fossil	67	507	0.2%	100,000-year-old fossil
Air	21	768	0.1%	African air sample
Total Habitat-Lite terms	21,896		62.9%	
Total unique terms	34,836			

this field. We were primarily interested in whole genome or metagenomics sequences, but the initial data set consisted of entries for all genes. As a result, frequency counts were heavily skewed towards large metagenomics projects, so we did not use the frequency counts in our analysis. For example, the phrase “locations in the Sargasso Sea, Panama Canal, and the Galapagos Islands” occurred over 3 million times in this data set.

Because of the size of the data set, it was not possible to explore it manually. One of the authors (C.C.) developed a small set of scripts to identify probable classes based on the presence of specific key words in each entry. The key words used for this analysis were based on the original Habitat-Lite terms plus synonyms and, in some cases, specializations. For example, for “waste water” the terms used for matching were “waste water,” “waste-water,” “wastewater,” “sewage,” “sewerage,” etc. For “food,” the terms used for matching included specific kinds of foods, for example, “milk,” “cheese,” “beer,” etc. Similarly, for “organism-associated,” the terms used for matching had to capture the many ways of expressing specific organisms, particularly humans, for example, “patient,” “female,” “subject,” “child,” etc.

Of the almost 35,000 distinct entries in the isolation_source field, some 22,000 (63%) contained specific words or phrases that could be mapped to the 17 Habitat-Lite categories. The bulk of these fell into the Organism-associated category (42%). In addition, we were able to identify over 20% of the entries that were geographic names or temporal expressions or other numerical quantities or identifiers. This enabled us to account for approximately 85% of the entries from isolation_source. The remaining 15% contain low frequency terms—many of them with species information (“wild mulberry”), location information (“Wilson and the Australian Museum),” or information about culture techniques (“top band of HTA gel”).

This pattern-matching approach allowed us to obtain a quick overview of the types of information found in the GenBank isolation_source field. This approach would require significant refinement and/or human intervention if we wished to use it for semiautomated assignment of Habitat-Lite terms to isolation_source entries, for improved search and indexing. In particular, this strategy mapped each entry to a single field, so that, for example, *130 m below sea surface* was mapped only to “marine,” losing the depth information. Similarly, the entry “Marine Biology Laboratory” caused the entry to be associated with the category “Marine”—a plausible inference but certainly not explicit information about habitat.

Habitat Field Entries from the GOLD Database

We next investigated a second data set, which consisted of the Habitat entries from the GOLD database on October 2007. The initial data set consisted of 1455 entries with 2210 terms. Table 1 shows some example GOLD entries, including not only the Habitat field, but also the much more detailed Isolation field. The entries in the Habitat field frequently contained multiple entries that specified the range of known habitats for a specific organism, for example, “Host, TB epidemic” or “Aquatic, Soil, Permafrost.”

Coverage of GOLD terms using Habitat-Lite

First, we looked for exact matches between GOLD Habitat terms and Habitat-Lite terms plus the additional term “aquatic.” This resulted in exact matches for 84% of GOLD Habitat terms. The three most frequent terms (“host,” “aquatic,” and “soil”) covered 75% of GOLD habitat data, while six Habitat-Lite terms were not seen at all in this smaller data set (“air,” “freshwater,” “extreme,” “microbial mat,” “fossil,” “terrestrial”).

Comparison of automated mapping and expert mapping

In the next experiment, we applied the automated mapping used in the GenBank experiment to the unique entries in the GOLD Habitat data, and compared these results to an expert mapping done by one of the authors (D.F.). There were a total of 132 unique entries in the GOLD Habitat field for metagenomes. There was 64% agreement (84 of 132) and 48 cases of differences in the automated mapping versus expert mapping. Most differences were due to a failure in the automated mapping procedure (30 cases, which were not classified or not mapped to the limited controlled vocabulary). Another nine were due to mismatches related to the new category “aquatic” introduced by the expert (five) and four were due to difficulty in classifying between freshwater and water.

The remaining nine discrepancies (shown in Table 4) brought to light interesting problems. Several of the discrepancies pointed out an ambiguity in the classification scheme with respect to “extreme habitat:” terms such as “hot springs,” “permafrost,” and “hypersaline mats” could be classified as “extreme habitat” or into a geographic or environmental feature (“hot springs,” “soil,” “microbial mat”). In another case (“rice paddies”), it is unclear without further context whether the focus was on the *rice* in rice paddies (“organism-associated”) or on the *paddies* (“terrestrial”).

TABLE 4. EXAMPLES OF DISAGREEMENT BETWEEN EXPERT MAPPING AND AUTOMATED MAPPING FOR THE GOLD HABITAT DATA

GOLD Habitat term	Expert mapping	Automated mapping
Mud	terrestrial	Soil
Rice paddies	terrestrial	Organism-associated
Soda lakes	organism-associated	Water
Hot spring	extreme habitat	Hot spring
Permafrost	extreme habitat	Soil
Snow	extreme habitat	Freshwater
Sulfur spring	extreme habitat	Water
Hypersaline mats	microbial mat	Extreme

These examples illustrate well the need for annotation guidelines, to handle situations where a term might be placed in several categories. There are several possible solutions: either there need to be “orthogonal dimensions” that would allow a category like “extreme habitat” to be “checked off” separately from some more specific information about geographic or environmental features or alternatively, there could be a facility to allow a given term to belong to multiple “bins.”

Manual annotation of the GOLD data to two orthogonal bins

The final set of experiments was designed to test the difficulty of the annotation task and to determine whether better annotation could be done by assigning multiple orthogonal terms. As noted above, there are advantages to capturing orthogonal annotations: to preserve richer information for searching, and also to reduce interannotator disagreement. To experiment with this approach, a single author (K.B.C.) annotated the 132 GOLD unique terms using Habitat-Lite in conjunction with an explicit set of guidelines that were meant to ensure that every *Habitat* entry was assigned both a general (*biome*) term and an *environment* term. The guidelines made use of the mappings of the Habitat-Lite terms to the EnvO taxonomy as follows:

1. Assign a child term of *biome* (*freshwater, marine, or terrestrial*).
2. Can the input be assigned a child class of *habitat* (*organism-associated or extreme*)? If so, assign it, and then stop. (This had an undesired effect, which we describe below.)
3. Is the input a food? If so, assign *food*. If not, go to (4).
4. Can the input be assigned a child of *biotic/abiotic* (*biofilms, microbial mat, or fossil*)? If so, assign it, and then stop. If not, go to (5).
5. Can the input be assigned a child class of *hydrographical/physiographic/anthropogenic* (*hot spring, hydrothermal*

vent, or wastewater)? If so, assign it, and then stop. If not, go to (6).

6. Can the input be assigned a child of *environmental substance* (*soil, water, sediment, sludge, or air*)? If so, then assign it.
7. Stop.

The undesired effect of Step (2) was that some inputs that could have been assigned specific terms related to extreme habitats were instead only assigned the more general *extreme* (*habitat*). A simple reordering of the rule might fix this.

The results demonstrate that the annotation task is well within the range of someone with reasonable background in biology. Only 2 out of 132 entries were left unannotated due to lack of domain knowledge: *solfataric fields*, and *self-heated organic materials*. It took approximately 1.5 h to do about 2 * 132 annotations, or around 1.5 terms annotated/minute. Based on this estimate, it would take less than a day’s work to map all of the GOLD Habitat entries to Habitat-Lite.

Discussion

The goals of this work were to create a useful set of high-level terms to capture habitat data, and to develop a methodology that can be applied to similar problems—specifically, to:

1. Determine what descriptors of habitat are recorded and how they are expressed in free text;
2. Determine how well a small set of terms, such as Habitat-Lite, could cover terms found in key resources;
3. Examine the feasibility of (semi-)automated capture of these fields of information for future projects.

Our initial experiments have resulted in a new version of Habitat-Lite (shown in Table 5), based on analysis of the GenBank *isolation_source* field and the *habitat* field in the GOLD

TABLE 5. PROPOSED HABITAT-LITE VERSION 0.2

Top level	Second level	Third level	Example also could be coded for
Choose one or more: Aquatic: freshwater Aquatic: marine	Choose one or more: soil sediment	Free text description, e.g. pinyon-juniper forest soil oxygen-depleted intertidal marine sediment	Terrestrial Aquatic: marine
Aquatic	sludge	Thermophilic methanogenic sludge	Terrestrial?
Terrestrial Air Fossil	waste water hot spring hydrothermal vent	waste water of paper machine hot spring at 70°C the shallow hot vent in Iwo Jima	Aquatic Aquatic, Extreme Aquatic: marine, Extreme Aquatic
	biofilm	biofilm of drinking water distribution system	
Food	microbial mat [Food Ontology or CV]	hot spring microbial mat surface of smear ripened cheese	Aquatic, hot spring Food
Organism Associated	[Species CV] [Anatomy CV, e.g., MIAA]	gut of nitidulid beetle	Organism- Associated
Extreme habitat Other	Select if appropriate	extremely alkaline (pH 12 to 13) groundwater (45.32739 N, 80.40874 W)	Aquatic; Extreme

database. Based on this analysis, we put forward the following recommendations for Habitat-Lite:

- A shift from a “flat” list to one with some structure is necessary.
- The set of terms should support certain inferences useful for search; for example, that a sample labeled *soil* is also *terrestrial*, or that a sample from a *hydrothermal vent* is also *extreme*.
- Consistent annotation requires guidelines for general terms such as *terrestrial* and *aquatic*, to instruct annotators to annotate to the most specific term possible.
- The notion of *extreme environment* is problematic in that it should be annotated **in addition** to a more specific term, such as *hot spring*—thus requiring that certain entries be associated with two Habitat-Lite terms.
- The category *Organism-associated* needs to be subdivided by linking out to other ontologies or controlled vocabularies (specifically, a taxon hierarchy and perhaps a high-level anatomy ontology).
- *Fossil* is an example of a currently infrequently used term, but a candidate for inclusion as a term of “exceptional importance” that could be useful in the future for searching.

The new version of Habitat-Lite has been implemented through the use of the Category mechanism in OBO-Edit with the category “EnvO-Lite-GSC.” We use the name EnvO-Lite-GSC to indicate that the ontology is a “light weight” version based on the full EnvO but tailored to the needs of the Genome Standards Consortium. Other groups that need a specific view of EnvO can use this mechanism, which makes it possible to save a specific view of the full EnvO ontology (available from <http://obo.cvs.sourceforge.net/obo/obo/ontology/environmental/>), drawing on the EnvO identifier space and using same structure of the ontology and the same terms but including only the terms specified by the selected category. Such views can be requested from the EnvO curatorial staff via the EnvO-tracker or the desired view of EnvO can be created locally using OBO-Edit.

The new set of Habitat-Lite terms is structured into two levels: a set of high level terms (first column in the table: *aquatic*, *terrestrial*, *air*, plus *organism-associated*, *food*, *extreme environment*, *fossil*), and a second level of more specific terms (column 2 in Table 5). To maximize capture of information, this version encourages selection of one or more of the high-level terms, one or more of the second-level terms, and recording of the specific information in free text (column 3, Table 5). The free text is shown associated with its level 2 term and in column 4, one or more appropriate top level terms.

To maintain simplicity, there is no obligatory connection/restriction between choice of top-level terms and second-level terms, except for the “food” and “organism-associated” classes. This allows flexibility (e.g., there are both freshwater and salt marshes) with the downside of increased possibility for error or for incomplete annotation. It should be possible to do automated association of high level terms, based on the second level terms, for example, associating “terrestrial” automatically with any annotation of “soil,” “sediment,” or possible new terms such as “sand,” “wood,” “rock,” or “mud.”

The “organism-associated” class should be elaborated by a term describing the organism and an anatomy term for the part of the organism; we will investigate use of a minimal

anatomy ontology, such as Jonathan Bard’s MIAA (Minimal Information about Anatomy; personal communication). The food class for now is just left as free text; it may be possible to use a small specialized food controlled vocabulary or ontology in the future (see http://gensc.org/gc_wiki/index.php/Food_Ontology_Project for discussions about the creation of a food ontology or controlled vocabulary).

Next Steps for Habitat-Lite: Adoption by GOLD, RDP, GCDML

The new version of Habitat-Lite will be tested against the GOLD data and revised to support GOLD (Liolios et al., 2008), IMG (Markowitz et al., 2008a), and IMG/M (Markowitz et al., 2008b). (For IMG, see <http://img.jgi.doe.gov>; for IMG/M, see <http://imgweb.jgi.psf.org/cgi-bin/m/main.cgi>.) GOLD has embraced the adoption of this controlled vocabulary/ontology for its habitat data. Capture of GOLD and IMG habitat data is currently implemented via the Expert Review Web submission form on the Integrated Microbial Genomes (IMG) Web site. All genomes submitted directly into IMG and IMG/M are now required to provide metadata that conforms to the GOLD vocabulary. The RDP (Cole et al., 2007) has also agreed to adopt the revised version of Habitat-Lite. The new version of Habitat-Lite will be supported in GCDML (Kottmann et al., 2008).

Conclusions

These results indicate that it should be possible to produce a list of terms with good high-level coverage for Habitat-Lite. We accept that candidate Habitat-Lite terms provide only very high-level information, and that these terms may be an amalgamation of terms found in different branches of future ontologies, or even among different orthogonal ontologies (e.g., for “organism-associated”). We also recognize that while these terms may provide a useful tool for biologists and databases, they have severe limitations. We emphasize the importance of maximum reporting of information about habitat—in particular, the necessity of preserving free text fields associated with legacy data so that more fine-grained information is never lost, and reanalysis is always possible.

In the long term, our goal is the creation of an interactive metadata checking system (a kind of metadata “spell checker”) that could “read” free text and suggest the correct mapping into a controlled vocabulary/ontology, for user validation or correction, thus ensuring that metadata is comprehensively captured and “binned” at the point of entry.

The use of a combination of Habitat-Lite terms in the short-term, cultural shifts in the way this community annotates to capture more complete descriptions of habitat and isolation source, and future use of ontologies and ontology-aware software will have a measurable benefit on the ability of researchers to effectively reuse ever-growing sources of data for large-scale, downstream analyses.

Toward a minimum information list of habitat terms for use in the GSC

We have posted the initial and revised versions of Habitat-Lite (Table 5) to the GSC Wiki. This list is annotated with recommendations and issues that will be addressed in revising it. We are making an open call for evaluation of this

list of habitat terms in order to develop a consensus-driven version of it that best suits community needs. This terms list will then be implemented in GCDML (Kottmann et al., 2008) and used in the first instance to fill the "Habitat" field of the MGS-compliant Genome Catalogue database (<http://gensc.org>).

Acknowledgments

The work at MITRE (L.H., C.C., K.B.C., S.M., J.L.) has been supported in part by National Science Foundation Grant 0746650, Small Grant for Exploratory Research: Mining Metadata for Metagenomics. We thank Tatiana Tatusova for a number of discussions on an approach to developing a light-weight set of classes for annotation, as well as Scott Federhen, Karen Clark, and Anji Johnston of NCBI. We thank John Wilbur, NCBI, for providing us with the data from the GenBank isolation_source fields. We thank Norman Morrison and Lynn Schriml of the EnvO/Gaz project for critical reads of the manuscript.

Author Disclosure Statement

The authors declare that no competing financial interests exist.

References

- Cole, J.R., Chai, B., Farris, R.J., Wang, Q., Kulam-Syed-Mohideen, S., McGarrell, A.M., et al. (2007). The ribosomal database project (RDP-II): introducing myRDP space and quality controlled public data. *Nucleic Acids Res* **35**, D169–D172.
- Dawyndt, P., Vancanneyt, M., Demeyer, H., and Swings, J. (2005). Knowledge accumulation and resolution of data inconsistencies during the integration of microbial information sources. *IEEE Trans Knowledge Data Eng* **17**, 1111–1126.
- Desantis, T.Z., Hugenholtz, P., and Larsen, N. (2006). GreenGenes, a Chimera-checked 16S rRNA gene database and workbench compatible with ARB. *Appl Environ Microbiol* **72**, 5069–5072.
- Field, D., Garrity, G.M., Gray, T., Morrison, N., Selengut, J.D., Sterk, P., et al. (2008a). The Minimum Information about a Genome Sequence (MIGS) specification. *Nat Biotechnol* **26**, 541–547.
- Field, D., Garrity, G.M., Sansone, S.-A., Sterk, P., Gray, T., Glockner, F.O. (2008b). Meeting Report: The 5th Genomic Standards Consortium Workshop. *OMICS* (this issue).
- Floyd, M.M., Tang, J., Kane, M., and Emerson, D. (2005). Captured diversity in a culture collection: case study of the geographic and habitat distributions of environmental isolates held at the American Type Culture Collection. *Appl Environ Microbiol* **71**, 2813–2823.
- Gill, S.R., Pop, M., Deboy, R.T., Eckburg, P.B., Tumbaugh, P.J., Samuel, B.S., et al. (2006). Metagenomic analysis of the human distal gut microbiome. *Science* **312**, 1355–1359.
- Hunter, L. (2002). Ontologies for programs, not people. *Genome Biol* **3**.
- Kottmann, R., Gray, T., Murphy, S., Kagan, L., Kravitz, S., Lombardot, T., et al. (2008). A standard MIGS/MIMS compliant XML Schema: towards the development of the Genomic Contextual Data Markup Language (GCDML). *OMICS* (this issue).
- Liolios, K., Mavrommatis, K., Tavernarakis, N., and Kyrpides, N.C. (2008). The Genomes OnLine Database (GOLD) in 2007: status of genomic and metagenomic projects and their associated metadata. *Nucleic Acids Res* **36**, D475–D479.
- Markowitz, V.M., Szeto, E., Palaniappan, K., Grechkin, Y., Chu, K., Chen, I.-M.A., et al. (2008a). The Integrated Microbial Genomes (IMG) system in 2007: data content and analysis tool extensions. *Nucleic Acids Res* **36**, D528–D533.
- Markowitz, V.M., Ivanova, N.N., Szeto, E., Palaniappan, K., Chu, K., Dalevi, D., et al. (2008b). IMG/M: a data management and analysis system for metagenomes. *Nucleic Acids Res* **36**, D534–D538.
- Martiny, J.B., and Field, D. (2005). Ecological perspectives on the sequenced genome collection. *Ecol Lett* **8**, 1334–1345.
- Morrison, N., Wood, J.A., Hancock, D., Shah, S., Hakes, L., Gray, T., et al. (2006). Standard annotation of environmental OMICS data: application to the transcriptomics domain. *OMICS* **10**, 172–178.
- Nealson, K.H., and Venter, J.C. (2007). Metagenomics and the Global Ocean Survey: what's in it for us, and why should we care? *ISME* **1**, 185–190.
- Venter, J.C., Remington, K., Heidelberg, J.F., Halpern, A.L., Rusch, D., Eisen, J.A., et al. (2004). Environmental genome shotgun sequencing of the Sargasso Sea. *Science* **304**, 66–74.
- Von Mering, C., Hugenholtz, P., Raes, J., Tringe, S.G., Doerks, T., Jensen, L.J., et al. (2007). Quantitative phylogenetic assessment of microbial communities in diverse environments. *Science* **315**, 1126.

Address reprint requests to:
 Lynette Hirschman
 Information Technology Center
 The MITRE Corporation
 202 Burlington Rd.
 Bedford, MA 01730

E-mail: lynette@mitre.org

This article has been cited by:

1. Peter C Marks, Marc Bigler, Eric B Alsop, Adrien Vigneron, Bart P Lomans, Renato De Paula, Brett Geissler, Nicolas Tsesmetzis. 2018. MetaHCR: a web-enabled metagenome data management system for hydrocarbon resources. *Database* **2018**. . [[Crossref](#)]
2. Evangelos Pafilis, Sune P. Frankild, Julia Schnetzer, Lucia Fanini, Sarah Faulwetter, Christina Pavloudi, Katerina Vasileiadou, Patrick Leary, Jennifer Hammock, Katja Schulz, Cynthia Sims Parr, Christos Arvanitidis, Lars Juhl Jensen. 2015. ENVIRONMENTS and EOL: identification of Environment Ontology terms in text and the annotation of the Encyclopedia of Life. *Bioinformatics* **31**:11, 1872-1874. [[Crossref](#)]
3. Elizabeth R. Ellwood, Betty A. Dunckel, Paul Flemons, Robert Guralnick, Gil Nelson, Greg Newman, Sarah Newman, Deborah Paul, Greg Riccardi, Nelson Rios, Katja C. Seltmann, Austin R. Mast. 2015. Accelerating the Digitization of Biodiversity Research Specimens through Online Public Participation. *BioScience* **65**:4, 383-396. [[Crossref](#)]
4. Curtis Huttenhower, Rob Knight, C. Titus Brown, J. Gregory Caporaso, Jose C. Clemente, Dirk Gevers, Eric A. Franzosa, Scott T. Kelley, Dan Knights, Ruth E. Ley, Anup Mahurkar, Jacques Ravel, Owen White. 2014. Advancing the Microbiome Research Community. *Cell* **159**:2, 227-230. [[Crossref](#)]
5. Kazuhiro Takemoto, Saori Kanamaru, Wenfeng Feng. 2014. Climatic seasonality may affect ecological network structure: Food webs and mutualistic networks. *Biosystems* **121**, 29-37. [[Crossref](#)]
6. Chao Niu, Dong Yu, Yuelan Wang, Hongguang Ren, Yuan Jin, Wei Zhou, Beiping Li, Yiyong Cheng, Junjie Yue, Zhixian Gao, Long Liang. 2013. Common and pathogen-specific virulence factors are different in function and structure. *Virulence* **4**:6, 473-482. [[Crossref](#)]
7. Kaustubh R. Patil, Alice C. McHardy. 2013. Alignment-Free Genome Tree Inference by Learning Group-Specific Distance Metrics. *Genome Biology and Evolution* **5**:8, 1470-1484. [[Crossref](#)]
8. Kazuhiro Takemoto. 2013. Does Habitat Variability Really Promote Metabolic Network Modularity?. *PLoS ONE* **8**:4, e61348. [[Crossref](#)]
9. Pier Buttigieg, Norman Morrison, Barry Smith, Christopher J Mungall, Suzanna E Lewis. 2013. The environment ontology: contextualising biological and biomedical entities. *Journal of Biomedical Semantics* **4**:1, 43. [[Crossref](#)]
10. Renzo Kottmann. Integrated Database Resource for Marine Ecological Genomics 1-6. [[Crossref](#)]
11. H. Teeling, F. O. Glockner. 2012. Current opportunities and challenges in microbial metagenome analysis--a bioinformatic perspective. *Briefings in Bioinformatics* **13**:6, 728-742. [[Crossref](#)]
12. H. Teeling, B. M. Fuchs, D. Becher, C. Klockow, A. Gardebrecht, C. M. Bennke, M. Kassabgy, S. Huang, A. J. Mann, J. Waldmann, M. Weber, A. Klindworth, A. Otto, J. Lange, J. Bernhardt, C. Reinsch, M. Hecker, J. Peplies, F. D. Bockelmann, U. Callies, G. Gerdt, A. Wichels, K. H. Wiltshire, F. O. Glockner, T. Schweder, R. Amann. 2012. Substrate-Controlled Succession of Marine Bacterioplankton Populations Induced by a Phytoplankton Bloom. *Science* **336**:6081, 608-611. [[Crossref](#)]
13. Jack A. Gilbert, Yiming Bao, Hui Wang, Susanna-Assunta Sansone, Scott C. Edmunds, Norman Morrison, Folker Meyer, Lynn M. Schriml, Neil Davies, Peter Sterk, Jared Wilkening, George M. Garrity, Dawn Field, Robert Robbins, Daniel P. Smith, Ilene Mizrahi, Corrie Moreau. 2012. Report of the 13th Genomic Standards Consortium Meeting, Shenzhen, China, March 4-7, 2012. *Standards in Genomic Sciences* **6**:2, 276-286. [[Crossref](#)]
14. Shiri Freilich, Raphy Zarecki, Omer Eilam, Ella Shtifman Segal, Christopher S. Henry, Martin Kupiec, Uri Gophna, Roded Sharan, Eytan Rupp. 2011. Competitive and cooperative metabolic interactions in bacterial communities. *Nature Communications* **2**:1. . [[Crossref](#)]
15. W. J. Sul, J. R. Cole, E. d. C. Jesus, Q. Wang, R. J. Farris, J. A. Fish, J. M. Tiedje. 2011. Bacterial community comparisons by taxonomy-supervised analysis independent of sequence alignment and clustering. *Proceedings of the National Academy of Sciences* **108**:35, 14637-14642. [[Crossref](#)]
16. Susannah G. Tringe, Philip Hugenholtz. The Enduring Legacy of Small Subunit rRNA in Microbiology 123-128. [[Crossref](#)]
17. Melissa Beth Duhaime, Renzo Kottmann, Dawn Field, Frank Oliver Glöckner. 2011. Enriching public descriptions of marine phages using the Genomic Standards Consortium MIGS standard. *Standards in Genomic Sciences* **4**:2, 271-285. [[Crossref](#)]
18. BalaKrishna Kolluru, Sirintra Nakjang, Robert P. Hirt, Anil Wipat, Sophia Ananiadou. 2011. Automatic extraction of microorganisms and their habitats from free text using text mining workflows. *Journal of Integrative Bioinformatics* **8**:2. . [[Crossref](#)]
19. Javier Tamames, Victor de Lorenzo. 2010. EnvMine: A text-mining system for the automatic extraction of contextual information. *BMC Bioinformatics* **11**:1. . [[Crossref](#)]
20. Hongwei Wu, Elliot Moore. 2010. Association analysis of the general environmental conditions and prokaryotes' gene distributions in various functional groups. *Genomics* **96**:1, 27-38. [[Crossref](#)]

21. Konstantinos Liolios, I-Min A. Chen, Konstantinos Mavromatis, Nektarios Tavernarakis, Philip Hugenholtz, Victor M. Markowitz, Nikos C. Kyrpides. 2010. The Genomes On Line Database (GOLD) in 2009: status of genomic and metagenomic projects and their associated metadata. *Nucleic Acids Research* **38**:suppl_1, D346-D354. [[Crossref](#)]
22. R. Kottmann, I. Kostadinov, M. B. Duhaime, P. L. Buttigieg, P. Yilmaz, W. Hankeln, J. Waldmann, F. O. Glockner. 2010. Megx.net: integrated database resource for marine ecological genomics. *Nucleic Acids Research* **38**:Database, D391-D395. [[Crossref](#)]
23. Jack A. Gilbert, Folker Meyer, Rob Knight, Dawn Field, Nikos Kyrpides, Pelin Yilmaz, John Wooley. 2010. Meeting report: GSC M5 roundtable at the 13th International Society for Microbial Ecology meeting in Seattle, WA, USA August 22-27, 2010. *Standards in Genomic Sciences* **3**:3, 235-239. [[Crossref](#)]
24. Elizabeth Glass, Folker Meyer, Jack A Gilbert, Dawn Field, Sarah Hunter, Renzo Kottmann, Nikos Kyrpides, Susanna Sansone, Lynn Schriml, Peter Sterk, Owen White, John Wooley. 2010. Meeting Report from the Genomic Standards Consortium (GSC) Workshop 10. *Standards in Genomic Sciences* **3**:3, 225-231. [[Crossref](#)]
25. Jack A. Gilbert, Folker Meyer, Janet Jansson, Jeff Gordon, Norman Pace, James Tiedje, Ruth Ley, Noah Fierer, Dawn Field, Nikos Kyrpides, Frank-Oliver Glöckner, Hans-Peter Klenk, K. Eric Wommack, Elizabeth Glass, Kathryn Docherty, Rachel Gallery, Rick Stevens, Rob Knight. 2010. The Earth Microbiome Project: Meeting report of the "1st EMP meeting on sample selection and acquisition" at Argonne National Laboratory October 6th 2010. *Standards in Genomic Sciences* **3**:3, 249-253. [[Crossref](#)]
26. Oranmiyan W. Nelson, Scott H. Harrison, George M. Garrity. 2009. Meeting report for SIGS1: First Conference of the Standards in Genomic Sciences eJournal. *Standards in Genomic Sciences* **1**:1, 72-76. [[Crossref](#)]
27. Dawn Field, Peter Sterk, Nikos Kyrpides, Renzo Kottmann, Frank Oliver Glöckner, Lynette Hirschman, George M. Garrity, John Wooley, Paul Gilna. 2009. Meeting Report from the Genomic Standards Consortium (GSC) Workshops 6 and 7. *Standards in Genomic Sciences* **1**:1, 68-71. [[Crossref](#)]
28. John C. Wooley, Dawn Field, Frank-Oliver Glöckner. 2009. Extending Standards for Genomics and Metagenomics Data: A Research Coordination Network for the Genomic Standards Consortium (RCN4GSC). *Standards in Genomic Sciences* **1**:1, 85-90. [[Crossref](#)]
29. Christopher Brewster, Simon Jupp, Joanne Luciano, David Shotton, Robert D Stevens, Ziqi Zhang. 2009. Issues in learning an ontology from text. *BMC Bioinformatics* **10**:S5. . [[Crossref](#)]
30. Amoolya H. Singh, Tobias Doerks, Ivica Letunic, Jeroen Raes, Peer Bork. 2009. Discovering Functional Novelty in Metagenomes: Examples from Light-Mediated Processes. *Journal of Bacteriology* **191**:1, 32-41. [[Crossref](#)]
31. Susannah G Tringe, Philip Hugenholtz. 2008. A renaissance for the pioneering 16S rRNA gene. *Current Opinion in Microbiology* **11**:5, 442-446. [[Crossref](#)]