

Management of acoustic metadata for bioacoustics



Marie A. Roch ^{a,*}, Heidi Batchelor ^b, Simone Baumann-Pickering ^b, Catherine L. Berchok ^c,
Danielle Cholewiak ^c, Ei Fujioka ^d, Ellen C. Garland ^c, Sean Herbert ^b, John A. Hildebrand ^b,
Erin M. Oleson ^c, Sofie Van Parijs ^c, Denise Risch ^c, Ana Širović ^b, Melissa S. Soldevilla ^c

^a San Diego State University, Dept. of Computer Science, 5500 Campanile Dr., San Diego, CA 92182-7720, USA

^b Scripps Institution of Oceanography, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0205, USA

^c NOAA National Marine Fisheries Service, Fisheries Science Centers, 1315 East-West Highway, Silver Spring, MD 20910, USA

^d Marine Geospatial Ecology Lab, Nicholas School of the Environment, Duke University, PO Box 90328, Durham, NC 27708, USA

ARTICLE INFO

Article history:

Received 26 July 2015

Received in revised form 1 December 2015

Accepted 9 December 2015

Available online 17 December 2015

Keywords:

Bioacoustics

Metadata

Call spatiotemporal database

Environmental data access

ABSTRACT

Recent expansion in the capabilities of passive acoustic monitoring of sound-producing animals is providing expansive data sets in many locations. These long-term data sets will allow the investigation of questions related to the ecology of sound-producing animals on time scales ranging from diel and seasonal to inter-annual and decadal. Analyses of these data often span multiple analysts from various research groups over several years of effort and, as a consequence, have begun to generate large amounts of scattered acoustic metadata. It has therefore become imperative to standardize the types of metadata being generated. A critical aspect of being able to learn from such large and varied acoustic data sets is providing consistent and transparent access that can enable the integration of various analysis efforts. This is juxtaposed with the need to include new information for specific research questions that evolve over time. Hence, a method is proposed for organizing acoustic metadata that addresses many of the problems associated with the retention of metadata from large passive acoustic data sets. A structure was developed for organizing acoustic metadata in a consistent manner, specifying required and optional terms to describe acoustic information derived from a recording. A client-server database was created to implement this data representation as a networked data service that can be accessed from several programming languages. Support for data import from a wide variety of sources such as spreadsheets and databases is provided. The implementation was extended to access Internet-available data products, permitting access to a variety of environmental information types (e.g. sea surface temperature, sunrise/sunset, etc.) from a wide range of sources as if they were part of the data service. This metadata service is in use at several institutions and has been used to track and analyze millions of acoustic detections from marine mammals, fish, elephants, and anthropogenic sound sources.

© 2015 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

A large variety of marine organisms, including marine mammals, fishes, and invertebrates, produce species-specific acoustic signals, or calls (Anorim, 2006; Hawkins, 1986; Richardson et al., 1995; Versluis et al., 2000). Knowledge of the occurrence of these calls has been valuable in increasing our understanding of the biology and ecology of these often visually elusive organisms (e.g. Aguilar de Soto et al., 2011; Baumann-Pickering et al., 2014; Hernandez et al., 2013; McDonald et al., 2006; Oleson et al., 2007c; Risch et al., 2013; Simpson et al., 2005; Širović et al., 2004). The marine bioacoustics community has invested considerable resources in developing tools to detect, classify, track, localize, and determine the

density of animals based on calls (e.g. Barlow and Taylor, 2005; Blackwell et al., 2013; Deecke and Janik, 2006; Erbe and King, 2008; Gillespie et al., 2013; Kandia and Stylianou, 2006; Marques et al., 2009; Mellinger et al., 2011; Nosal, 2013; Zimmer, 2011). These calls are recorded on a variety of fixed (e.g. moored instruments, bottom seafloor packages) and mobile (e.g. towed arrays, autonomous underwater vehicles, animal tags) instrument configurations.

The tools for analyses of bioacoustic data sets, whether automated, manual, or some combination thereof, can provide a range of information about the calling animals and their environments such as signal characteristics, temporal patterns in vocal behavior, source levels, density estimates, measurements of anthropogenic noise, etc. It is often possible to infer biologically relevant information, such as daily and seasonal activity patterns over potentially large temporal and spatial scales. Information derived from these recordings such as detections of calling animals and the methods used for detection is considered to

* Corresponding author.

E-mail address: marie.roch@sdsu.edu (M.A. Roch).

be metadata (information describing the data) of the acoustic record. However, these metadata have frequently been generated and stored in idiosyncratic formats on the computers of individual researchers. While combining multiple datasets or the results from multiple analyses might often lend more power to the ability to interpret patterns in the data, consistent metadata formats and mechanisms of retrieval are required to remove the formidable obstacles that often hinder merging results from multiple analyses.

The proliferation of such tools for producing large quantities of metadata poses a new set of data management challenges as well as providing exciting opportunities for the bioacoustics community to ask new types of questions in a data rich environment. By pooling metadata from multiple sources, the scope of study that can be undertaken can be significantly expanded, but care must be taken to ensure that the data and methods are compatible.

Of particular importance in these metadata associated with acoustic detections is documentation of the data processing method applied to a given dataset: What portion of the data were analyzed? What was the target and methodology of the analysis? Which detections were gathered in a systematic manner and which were opportunistic or incidental? The methods require enough detail to determine whether studies are compatible. For example, consider combining two studies that used different signal-to-noise ratio thresholds for detecting animals with similar call source levels; for some analysis questions, this difference should be factored into the analysis of the combined data to prevent bias. In most cases, the study with the lower SNR threshold will detect animals from farther away, thus increasing the area over which animals are monitored. Assuming a spherical spreading model (Urlick, 1983 pp. 100–101) and a lower detection threshold of half the acoustic pressure (6 dB) will result in increasing the radius of the monitoring area by a factor of 2, with a corresponding increase in area by a factor of 4. Studies testing hypotheses related to call rates would need to take into account the number of calls detected with respect to the monitored area while those considering characteristics of calls should consider that there would be frequency-dependent differences in the attenuation of the received signal.

Indicating what portion of the data were analyzed is important for constructing valid inferences and is a separate issue from the actual recording regime of the instrument. It is common to subsample data from long-deployment passive acoustic monitoring data sets. The decision of what portion of the acoustic data to analyze can be thought of as a secondary stage of sample design or survey effort, and in this article, we will refer to it as analysis effort.

One must also indicate the species and calls for which systematic analysis effort is conducted. Studies focusing on a single species may not typically record these types of details, especially when all of the data are consistently analyzed due to manageable data size or efficient automated classifiers. However, specification of the details of systematic analysis effort facilitates the retrieval of records appropriate to a researcher's question and is critical for the selection of metadata from data repositories containing diverse analysis effort. It should be noted that in many fields, researchers will record opportunistic or incidental detections that are not part of their systematic effort. Analogs to these type of detections exist in other types of survey studies such as visual point transect (Buckland, 2006) and trap surveys (Buckland et al., 2006). Examples include an ornithologist noting a rare species of bird when moving from one point transect to another or an entomologist electing to perform several opportunistic net sweeps to collect additional samples around a bee trap. In both cases, additional information can be gained from the analysis of these incidental detections or animals and they should be retained. However, during analysis, they must be distinguished from data that were obtained in a systematic manner. Systematic observations are necessary for well-founded inference about spatio-temporal patterns.

Data analysis over large, spatially and temporally varied acoustic data requires *consistency*, which is the first key feature of our approach. This means that standardized names describing data types in the metadata must be selected along with constrained sets of values that can be stored. As an example of this, one may elect to store a species' common name, scientific name, or a coded value representing the species name such as the taxonomic serial numbers provided by the Integrated Taxonomic Information System (ITIS Organization, 2014). Similarly, one might elect to specify that acoustic sampling rate be measured in Hz or kHz; specification of units is necessary to effectively query metadata.

A hierarchy of concepts can be provided by grouping names together under the umbrella of a name that describes the group (sometimes called a frame or structure). An example of this is to use the name “parameters” to describe a collection of settings for a detection algorithm. Names, their values, and hierarchical structure form the basis of an ontology (McGuinness, 2003), a definition of how data are encoded and related to one another.

Consistency must be balanced with the need to be *extensible*. As the body of knowledge about species increases, new questions are posed. An acoustic signal that was considered at one time to be stereotyped may be found to have categorical or graded variations (e.g. Risch et al., 2013 recently showed that minke whale, *Balaenoptera acutorostrata*, thump trains had more pulse structure than previously thought), and researchers may wish to study those variations with respect to individuals, activity state, context, or ecosystem pressures such as habitat loss. In addition, researchers with different goals, analysis techniques, and working in a variety of habitats may have varying needs. Consequently, our goal was to define a system to capture acoustic metadata that is both consistent and extensible.

In this article, we focus on acoustic metadata for marine mammals and anthropogenic sources (e.g. shipping, naval operations, and oil exploration). We have used this type of approach to analyze the calls of numerous species of cetaceans on multiple datasets collected throughout the Pacific, merging results from over 36 years of analysis effort (Baumann-Pickering et al., 2014; Širović et al., 2015).

While the developmental effort focused on sounds from the marine environment, the methods have been extended to include the terrestrial environment with few modifications. Preliminary unpublished work conducted by Peter Wrege and Sara Keen at Cornell's Bioacoustics Research Program on calls from African forest elephants, *Loxodonta cyclotis*, has shown that this can be done without changes to the data representation (personal commun. Sara Keen). The only change to the implementation required was to update the subset of taxonomic serial numbers (ITIS Organization, 2014) stored in our database to include the family *Elephantidae*. The current implementation has expanded this to all species described in ITIS. In cases where altitude is needed (e.g. bird flight calls), our marine-centric name depth would need to be changed or negative depth values could be used.

We describe a set of metadata structuring rules that we call Tethys and provide a brief introduction to the Tethys Metadata Workbench, an implementation of this data framework that includes a server program and client libraries. The Tethys Metadata Workbench can manipulate the metadata as well as access a large variety of Internet-available geophysical, biological, and astronomical data sources. The workbench is designed to be used by individual laboratories. A web-services-based server permits data exchange between research groups, and summary data can be exported into the Ocean Biogeographic Information System – Spatial Ecological Analysis of Megavertebrate Populations (OBIS-SEAMAP; Halpin et al., 2009).

We are developing data representations for instrument deployments and calibration information, acoustic detection, classification, and localization data, and supplemental information. In this paper, we restrict our discussion to metadata related to instrument deployments




Fig. 1. Overview of workflow. Raw acoustic signals are analyzed via a variety of software packages to produce metadata describing animal calls and other aspects of the acoustic environment. The implementation described in Section 4 is capable of processing output from a wide variety of formats. These metadata are stored in a data repository along with details about the analysis protocol and the recording packages. Scientists then request these metadata through interfaces available for several programming languages. This same interface also provides access to other Internet-available data products, which can then be combined with the analysis results to provide broader context for data interpretation. (Sea surface height anomaly image courtesy NOAA Southwest Fisheries Science Center Environmental Research Division.)

and call detection and classification. A general overview of the system use can be seen in Fig. 1.

2. Metadata management

2.1. Background

There are several data standards (both authoritative and community-driven) that have been proposed over the last several years that address portions of the concepts required for acoustic metadata, and when feasible, the proposed system borrows concepts from these existing systems rather than defining new ones. Specifications of equipment have been detailed in *International Standards Organization's ISO-19115 (2003)* and in several of the *Open Geospatial Consortium's (OGC, 1994)* standards. We adopt portions of these standards, but they are not designed for biological applications and do not cover many needed aspects such as describing species, their calls, analysis effort, etc.

Within a biodiversity context, Darwin Core (Wieczorek et al., 2009; Wieczorek et al., 2012) is designed to provide a geospatial inventory of animals. The *Federal Geographic Data Committee (2014)* has a management plan that includes numerous themes for different types of data, including one that encompasses marine biota (*Marine and Coastal Spatial Data Subcommittee, 2012*). Ecologists have designed the *Ecological Metadata Language (EML, Feigraus et al., 2005)* to provide metadata

for ecology measurements. As with the aforementioned standards that focus on equipment, these standards do not specify many attributes that are important to bioacoustics researchers, such as documentation of analysis effort and methods for describing acoustic signals produced by animals.

An information system that attempts to integrate many types of data standards is the *Integrated Ocean Observing System (IOOS <http://www.ioos.noaa.gov>)*. IOOS uses a variety of data formats, ranging from some of the *Open Geospatial Consortium* standards to network common data format (*National Oceanic and Atmospheric Administration, 2013*). Guan et al. (2014) have developed an IOOS standard for passive acoustic recordings that addresses issues related to equipment and the objectives for collecting the acoustic data such as target species or the question being addressed in the collection effort. Currently, there are ongoing efforts by Fornwall et al. (2012) to extend IOOS to include visual line transect survey data and Fujioka et al. (2014) have developed methods for incorporating both visual observation and passive acoustic monitoring into OBIS-SEAMAP (<http://seamap.env.duke.edu>; Halpin et al., 2009). OBIS-SEAMAP is designed to track summary information rather than detailed information such as per call timing, call parameters, detailed localization information, etc.

The Tethys metadata schema is designed to standardize metadata needed for representing bioacoustic analysis, such as analysis effort, spectral and temporal parameters of biological, physical, and anthropogenic signals, localization data derived from synchronized acoustic

arrays, instrument properties such as calibration information, recording protocols such as duty cycle and sampling rate, and spatial effort from mobile platforms such as track lines.

2.2. Data model

Acoustic detection events must specify the attributes of each event and the context in which the event took place. Event attributes might include categorical or descriptive labels, measurements, time of occurrence, and location. This can be linked to both an analytical context describing the process used to produce measurements as well as to a data context describing the data from which the detected events were derived. This type of data lends itself to networked data models, where one piece of data is related to another via a structural linkage. Networked models are especially suitable when data are heterogeneous and can vary significantly from one event to another as they permit the encoding of complicated relationships. In the example of Fig. 2, multiple levels of organization are associated with example calls. The first detection contains a substructure with call parameters that includes the points of a time/frequency contour whereas the second call detected stores a spectrogram image of the call.

This contrasts with the relational database model where each record must have the same form, which gained popularity after the landmark work of Codd (1970). Relational models can handle limited types of diversity through the decomposition of tables into normal form, but this quickly becomes unwieldy in the face of heterogeneous data. While relational databases offer efficiency in

many cases, these models are challenged by loosely structured data and a number of recent projects have looked to alternative ways to organize data that do not fit well into the traditional relational database format (e.g. Chang et al., 2008; Leavitt, 2010).

One method of conceptualizing a network of information is an extensible markup language (XML) document, the approach taken in the OGC, ISO, and IOOS passive acoustic standards. XML documents consist of a hierarchical tree of structuring elements that encapsulate data (Connolly et al., 2007). An element consists of a pair of identical tags (names), enclosed in angled brackets, between which data or other elements are encapsulated. The latter tag has a slash (/) to indicate the end of the element, e.g. `<Latitude> 32.7 </Latitude>`. Elements may contain other elements: `<DeploymentDetails> <Latitude> 32.7 </Latitude> ... </DeploymentDetails>`. Throughout this article, we *italicize* tag names when they appear in the main text and frequently refer to them as elements, indicating that we expect the tag to have associated data. The linkages between items occur either implicitly from the hierarchy or from data that reference other portions of a document tree by means of a network path (sequence of tags leading to a specific element) or an identifier that references a value associated with a specific element. The downside of network models is that the user must know the network organization in order to effectively extract information from it. Consequently, standardizing the hierarchy of data organization is important.

While both the OGC and ISO standards currently do not represent bioacoustic metadata, we could have adopted a larger subset of either standard in the description of instrument deployments. However,




Fig. 2. A network data model provides linkages between different types of data. In this abstraction that loosely mirrors the data organization used by Tethys, a relationship is shown between an instrument deployment and a group of acoustic detection events that are associated with the deployment. Network models allow for heterogeneous structure. In this example, two different detected events are represented. Call parameters were recorded for the first call detected, while the call spectrogram was recorded for the second call detection.

these standards were designed for querying instrumentation in heterogeneous sensor networks. We found that for many bioacousticians, the amount of additional information that would need to be provided both for assembling deployment data and subsequent queries was overwhelming. Instead, we took inspiration from the portions of those standards that were relevant to the task at hand. We concur with Graybeal et al.'s (2012) proposition that it is difficult to provide an ontology that is all things to all users and that mediation can be used to translate between useful ontologies.

2.3. Data consistency and extensibility

XML provides a mechanism for describing data relationships, but additional mechanisms are required to ensure consistency. A common

method to specify this for XML is to use an XML schema (Walmsley, 2002), a specification document (written in XML) that provides an ordered list of data elements that should be in a document as well as simple integrity constraints on the data themselves. The top level element includes a number of sub-elements that permit the denotation of detail. For example, for the metadata describing acoustic events, the child element related to a specific detection analysis effort would describe how detections were made, including details such as the analysis effort, methods used for the analysis, who conducted the analysis, etc. (Fig. 3). Elements in Fig. 3 that contain the circle-plus symbol, $\oplus$, have child elements that further describe their parent. In some cases, a range of numbers will be shown along a link from parent to child, indicating the number of times that an element may be repeated. Extensibility is handled by permitting some elements to have arbitrary children.




Fig. 3. Top-level view of the schema for a detections document. All detections documents have a top-level *Detections* element. The stacked squares connecting *Detections* to its children indicate a sequence of elements: *Description*, *DataSource*, *Algorithm*, etc. Mandatory elements are denoted by bold lines. The majority of elements provide structure for child elements (not shown here), such as a group of elements that describe the systematic analysis effort. Other elements contain data such as the *UserID* which identifies the person or entity that submitted the group of detections to the database. Each element has a data type. With the exception of *UserID*, which has an XML primitive type for alphanumeric data, elements in this figure are types that we have defined elsewhere in the schema.

3. Tethys schemata

Tethys schemata contain a number of functional groups: 1) instrument deployments, 2) detection and classification, 3) groupings of instruments that are referred to as ensembles, 4) localization, and 5) instrument calibrations. We restrict ourselves to describing in detail the first two which are the most mature. For brevity, details of the schemata that are reasonably self-explanatory and not critical to representing the metadata will not be discussed, but the complete schemata are available on the Tethys web site (<http://tethys.sdsu.edu>). These schemata form an implementation agnostic method for representing data that can permit data interchange between different groups using software that implement the schemata. Note that while the implementation developed in Section 4 is built upon an XML database, any database capable of constructing XML from its internal representation of the metadata could be used. As we are interested in developing an ontology that can pool metadata from multiple acoustic metadata servers, the discussion will center on a representation level that enables interoperability.

3.1. Instrument deployments

The Deployment schema (Fig. 4) contains information about specific deployments of an instrument. The goal of this schema is to provide information sufficient to understand where, when, and how acoustic data have been collected.

Each deployment is uniquely identified by a key that consists of three elements: *Project*, *DeploymentID*, and either *Site* or *Cruise*. *Project* is a name that identifies a group of deployments. Examples of such names might include geographic regions, study names, or specific mitigation or monitoring projects. The *DeploymentID* is simply an integer such as the Nth deployment at a specific site. The final element of the deployment key is either the name of a *Site* for a fixed deployment or a *Cruise* (mission) for mobile platforms such as towed arrays or gliders. Collectively, the values of these elements must be unique within the set of deployment documents.

Instrument provides a brief description of the instrument, permitting a specification of an instrument *Type* (e.g. vendor/model) and an *ID* (e.g. serial number) that identifies a specific instrument of the specified type. The instrument's sensors are described in a separate element *Sensors* that has three types of child elements: *Audio*, *Depth*, and a generic *Sensor*. Zero or more children of each type of sensor are permitted although they must be grouped together and in the order specified by the schema. Repetitions allow for instruments with multiple instances of a specific type of sensor such as a multiple hydrophone instrument. All sensor types have a common set of elements used to describe them: a *SensorID* (e.g. serial number), *Geometry* relative to the platform position represented as an x, y, and z offset in meters, a *Name*, and textual *Description*. In addition to the common elements, audio sensor elements contain *hydrophoneID* and *preamplifierID*, and the generic sensor contains a *Type* to describe the sensor type and a *Properties* that can contain arbitrary child elements, thus permitting the specification of any type of sensor (e.g. compass, temperature).

The information about how audio sensors are configured throughout a deployment is stored in *SamplingDetails*. Each hydrophone that is active during a deployment must have one or more associated *Channel* elements in *SamplingDetails* that identifies the hydrophone within the set of audio sensors as well as the channel in the recording that is associated with this deployment. The *Channel* element contains children that specify the details about sampling (sample rate and quantization bits), gain (in dB for calibrated gain and relative numbers for uncalibrated gain such as a number on a dial), and duty cycle. These details are encapsulated in *Regimen* elements that contain a timestamp indicating when the settings were applied. All Tethys timestamps are represented in universal coordinated time (UTC) and encoded using the W3C Consortium's subset of the ISO8601 standard described by Wolf and

Wicksteed (1997). *Regimen* elements may be repeated, thus permitting the representation of complex sampling regimes such as recording with varying duty cycles during the same deployment. Note that calibration data are specified in a different schema. Complete details are beyond the scope of this paper, but each calibration document includes a sensitivity curve, a timestamp of when the calibration was performed, and an indication of whether the calibration was of a specific hydrophone, preamplifier, or complete “end-to-end” system. The schema does not currently support directivity patterns, although adding this would be straightforward.

A *QualityAssurance* element permits the documentation of the methodology and results of quality control checks that may be performed to ensure the integrity of the acoustic data associated with a deployment. It contains three types of elements: *Description*, *ResponsibleParty*, and *Quality*. The *Description* allows text-based descriptions of the quality assurance process through sub-elements *Objectives*, *Abstract*, and *Method*. *ResponsibleParty* provides contact information for the person who is responsible for the quality assurance process. The final element *Quality* permits the annotation of the quality control information itself and may be repeated. Each *Quality* element instance contains mandatory elements for *Start* and *End* times) as well as a *Category* element that allows the denotation of the data as unverified, good, compromised, or unusable. Optional elements provide further details about the data. A *FrequencyRange* element contains elements to denote the bandwidth to which the quality element applies (*Low_Hz*, *High_Hz*). One or more *Channel* elements indicate the channels to which the annotation applies. Finally, a comment allows for a textual description of this instance, such as “Instrument noise on disk writes.”

The *Data* element contains information related to measurements that will have occurred during the deployment. While Tethys is designed for specifying metadata about a recording rather than archiving the acoustic recording data, it is important that the system tracks the storage location of the acoustic data to permit users to locate acoustic data for further analysis, verification, or other uses. One or more uniform resource indicator (URI) elements are used to specify the acoustic data location. URIs are Internet identifiers and are divided into classes for different types of resources (Joint W3C/IEFT URI Planning Interest Group, 2002) such as web pages, digital object identifiers (DOIs), etc. This provides flexibility for the user to specify anything from sophisticated descriptions of where the data are published to something as simple as the description “file cabinet A, disk 12.” While the latter is not a valid URI, the text-based description permits such useful annotations when more sophisticated systems are not in place: these can easily be recognized as not being a valid URI. In some instances, users may have different copies of the data such as one that has come directly off the instrument in a proprietary format and data that have been processed in some way such as decimation. Additional elements permit these data locations to be specified. The *Data* element also lets one specify trackline information (*Track*) for mobile platforms, either via a set of URIs or a list of timestamped samples that can contain longitude and latitude measurements, depth, pitch, and roll.

Information about deployment and recovery is contained in *DeploymentDetails* and *RecoveryDetails*. They contain information about the instrument at the beginning and end of deployment: time and date, latitude and longitude, depth, and contact personnel. Portions of this follow the OpenGIS SensorML specification (Botts and McKee, 2003).

3.2. Detections

The Detections schema (Fig. 3) specifies how acoustic detections are identified and to which data they are attributable. The *Description* element contains several optional child elements (*Objectives*, *Abstract*, and *Method*) that permit a textual abstract describing the analysis effort. For example, when detecting an International Union for Conservation of Nature (IUCN) red-listed species such as North Pacific right whales




Fig. 4. Top level of Deployments schema.

(*Eubalaena japonica*), one might specify that the objective is to detect every potential call, whereas for more common species, the objective might be to minimize false positive detections. Brief summaries of the project can be listed in the *Abstract*, and a high-level description of the detection methods can be specified in *Method*.

DataSource contains elements that associate a detection document with a specific instrument deployment or ensemble set of deployments such as one might use in a wide-aperture array. This permits linkages between the detections and one or more instrument deployment documents so that information about the recording can be accessed. The identity of the person who is responsible for submitting the document is stored in *UserID*.

Algorithm has the purpose of describing the detection process in enough detail to make it reproducible. Like the *Description* element described above, a text-based summary can be specified in *Method*. The *Software* algorithm, its *Version* number, and a list of any *SupportSoftware* that may be required to run the detection algorithm should be listed. As an example, a user may be using a PAMGUARD module (Gillespie et al., 2008) to detect baleen whale calls. The module would be specified in the *Software/Version* elements as this is the algorithm that provides the detections, and PAMGUARD would be specified in the *SupportSoftware*. When analysts are responsible for identifying acoustic events as opposed to an automated algorithm, we recommend using “analyst” as the *Method* value and noting the software they are using for their annotation in *Software*. Most algorithms require the setting of parameters, such as analysis window length, thresholds, kernels, etc. To be reproducible, these parameters must be documented. This is a situation where the need for extensibility is evident; one cannot plan for the types of parameters that future algorithms may require. As a consequence, a *Parameters* element accepts any type of element (Fig. 5).

The *QualityAssurance* element is similar to the one that appears in the deployment schema in that it has child elements *Description* and *ResponsibleParty*. It differs, however, in that *Quality* elements are assigned to individual detections as discussed later.

The final three child elements of *Detections* are related to analysis effort and detections: *Effort*, *OnEffort*, and *OffEffort*. These describe what the detection method was systematically trying to find (*Effort*), the detections found with systematic analysis effort (*OnEffort*), and opportunistic or incidental detections for which there was no systematic analysis effort (*OffEffort*). Specification of effort consists of denoting the portions of the acoustic record that were examined and target signals to be detected. This is critical to analyzing acoustic detections in a meaningful way. As an example, suppose one queried for detections of blue whale (*Balaenoptera musculus*) B calls over 10 years of recordings at a location, but had only conducted analysis during August and September

of each year. Without consideration of systematic acoustic analysis effort, one could draw the erroneous conclusion that blue whales were not present the rest of year.

The children of the *Effort* element (Fig. 6) contain a *Start* and *End* time that must fall within the interval during which the instrument was deployed. This refers specifically to the time span over which the data are analyzed and should not to be confused with the effort related to data collection. This is followed by a *Kind* list which indicates for what the analyst, detector, or classifier is searching. Each *Kind* consists of the mandatory elements *SpeciesID*, *Call*, and *Granularity*.

SpeciesID specifies a taxonomic identifier for the species that is under consideration. The identifier consists of a taxonomic serial number (TSN) from the US government-sponsored Integrated Taxonomic Identification System (ITIS Organization, 2014) which maps to a taxon name. ITIS TSNs are unique positive integers and can be mapped by the Tethys Metadata Workbench to and from scientific and common names as well as local names or abbreviations specified by the user. Upon occasion, one may not know the exact identity of a species when a call is detected. In these cases, the recommended practice is to select a *SpeciesID* that is wide enough to encompass the call. As an example, an echolocation click detector might not report to the species level. In such cases, the suborder Odontoceti (TSN 180404) may be assigned as the *SpeciesID*. XML allows element attributes to provide additional information about an element, and we use the attribute *Group* to denote subpopulations as well as species where the identity is unclear. As an example, McDonald et al. (2009) recorded a beaked whale near Cross Seamount for which the species is not yet known. To represent effort for these calls, we use family Hyperoodontidae (TSN 770799) and specify *Group* “BWC” to denote the unknown beaked whale echolocation signals first observed at Cross Sea Mount. For anthropogenic data that can typically be associated with human designed technology such as pile-driving, shipping, explosions, etc., we use *Homo sapiens* (TSN 180092) as the species name and identify the activity as if it were a type of call. As Tethys permits the storing of physical phenomena, we reserve TSN-10 (to avoid collisions with future ITIS TSNs that are all positive) to specify other abiotic phenomena such as sounds from earthquakes, rain, etc., and use the *Call* element to denote the phenomenon.

Call indicates the type of the species' acoustic signal. Values for *Call* are not restricted as the literature does not always agree on how particular calls should be named. In the supplemental materials, we provide a table of recommended call types for cetacean species with which the authors work. A call type is a categorical description of an acoustic signal produced by an animal and covers both stereotyped calls such as a southern right whale (*Eubalaena australis*) “gunshot” (Clark, 1982), and a general description of highly variable signals such as delphinid

<pre> <Detections> ... <Algorithm> <Method>analyst</Method> <Software>Raven Pro</Software> <Version>1.4 build 48</Version> <Parameters> <SpectrogramWindowSize>512</Spec...> <Brightness>50</Brightness> <Contrast>50</Contrast> <SpectrogramLength_s>100</Spec...> </Parameters> </Algorithm> ... </Detections> </pre>	<pre> <Detections> ... <Algorithm> <Method>tonal detector</Method> <Software>silbido</Software> <Version>1.1</Version> <Parameters> <Framing> <Advance_ms>2</Advance_ms> <Length_ms>8</Length_ms> </Framing> <Threshold_db>10</Threshold_db> </Parameters> </Algorithm> ... </Detections> </pre>
---	--

Fig. 5. Examples of preserving information about how detections were made. Left box shows specification for an analyst annotating data with Raven (Cornell Univ., Ithaca NY) whereas the right box shows the specification for an automated whistle detector.




Fig. 6. Specification of systematic analysis effort within the Detections schema. *Start* and *End* times are followed by a list of one or more elements of *Kind*, specifying the species being detected, specific call type (and optional subtype), and an indication of the level of detail (granularity) for the detection.

(Cetacea: Delphinidae) whistles (Steiner, 1981). Calls are typically assigned to categories based on their temporal and spectral characteristics either using automatic classifiers or by trained analysts. The designation “Undescribed” can be used to annotate calls that have not yet been described in the literature, while “All” indicates that the call type was not noted, simply that the detector was looking for all call types of the given species. It is also possible to specify effort for variants, or subtypes, of a call. Subtypes are discussed in detail later in the context of individual detections.

Granularity describes the level of detail used in the search and can contain one of three values: call, encounter, and binned. Call indicates that the detector will be providing the start time and possibly the end time of each individual call. Encounter is used when the start and end times of a set of calls are to be denoted. An example of this is when a group of calling animals travels through an area. The start and end times of the set of detections denote the period over which the calling animals were within detection range. We do not specify a threshold for distinguishing one encounter from the next as this may vary depending upon the behavior of species. Binned is used when specifying presence/absence or the number of detections within a given time bin. Binned specifications require the presence of XML attribute *BinSize_m* indicating the bin size in minutes. We restrict this to being divisible into a 24 h day and the implied bin edges for detections associated with this granularity are restricted to aligning with day boundaries (e.g. a *BinSize_m* of 15 m is allowable whereas 17 m is not). These restrictions ensure that bins are aligned with one another from one day to the next and permits simple computation of daily call counts or other aligned units. As an example of how relaxing this restriction could be problematic, a 17 m boundary starting at midnight on the first day of recording would span the time from 23:49 of the first day to 00:05 of the second. If the on-duty portion of duty-cycled data does not align with these edges, care must be taken in analysis to not overestimate systematic analysis effort or the resultant detections.

The final two elements of a Detections document are the *OnEffort* and *OffEffort* elements. Both of these can contain sequences of *Detection* elements. Any on effort detection must correspond to the parameters specified in the systematic analysis effort (*Effort*). Detections within the off-effort section are unrestricted.

Each *Detection* (Fig. 7) entry conveys information about a specific detection. At a minimum, it requires a *Start* time and a *SpeciesID*. For binned detections, *Start* should fall on a bin boundary, and if desired, the number of detected calls can be stored with *Count*. *End* times are mandatory for encounter granularity, and otherwise optional. The *Input_File* element is an optional URI providing a link to the audio file from which the detection originated. This is a convenience element for file naming conventions that do not permit filenames to be automatically derived from the *Start* time and the *DataSource*.

At times, it is useful to be able to reference a specific detection (e.g. for referencing components of a song, or a call that was localized) and the optional *Event* element provides a mechanism to provide an alphanumeric identifier that is unique within the set of detections described by a specific detection. For example, Northeast Pacific blue whale AB song (Oleson et al., 2007a,b) might consist of an A call followed by two B calls. Identifiers could consist of sequentially numbered detections or the timestamps indicating when the analyst or algorithm made the detection. To denote the song, one would document the individual vocalizations separately and a song event that references each individual detection: A, B, and B, by their *Event* identifiers.

Call types can be listed. As binned and acoustic encounter effort can encompass multiple call types, more than one call can be listed. While this element is optional to provide for detectors that do not distinguish call types (usually used with an effort of “All”), it is highly recommended to populate this element when it is known. The *Parameters* element allows the recording of optional information about a call or set of calls and contains many of the parameters that bioacousticians are likely to collect (Fig. 8). In addition, *Parameters* can contain information about categorical variants via *Subtype* (e.g. the “A” and “B” type echolocation clicks of Pacific white-sided dolphins in Soldevilla et al., 2010) and measurements associated with the detector/classifier. Detection and classification systems ultimately make categorical decisions that are frequently associated with a measurement score (e.g. a likelihood score, linear discriminant analysis projection value, etc.). Some systems attempt to determine the probability that a decision is correct. Tethys provides optional elements *Score* and *Confidence* for the respective storage of classifier scores and probability of correctness. For species that produce so many calls that they are no longer distinguishable from




Fig. 7. Schema for individual detections.

one another, some researchers have developed indices that indicate the strength of the signal (e.g. fin whale 20 Hz calls; Širović et al., 2004), and in such cases an alternate use for *Score* is to store the index value. Related to these elements is the detection *QualityAssurance* element, which simply shows whether or not this detection has been reviewed and is restricted to the values of unverified, valid, and invalid. We recommend treating the *QualityAssurance* value invalid as a transient value and that these detections be removed once the classification verification is complete. If not, queries for detections should explicitly filter out invalid

detections to prevent these false detections from being included in the analysis.

The majority of the parameters associated with detected events are temporal-frequency measurements with which bioacousticians are familiar, such as minimum and maximum frequency content, received level, duration, etc. For individual calls, these measurements are self-explanatory. When measured for groups of detections (binned or encounter analysis effort), these are assumed to represent typical or mean measurements within the specified time period, and the




Fig. 8. Call parameters. A number of standard measurements of call characteristics are allowed with the possibility of providing project-specific measurements.

Description and *Algorithm* sections should contain information that clarifies the measurement methods used.

The *EventRef* element can be used to reference other events and will contain values from *Event* elements in different detections. An example use would be to reference individually detected calls that comprise the aforementioned blue whale AB song. A final element *UserDefined* can be used to note measurements that are of interest to researchers addressing specific questions. Sample items that could be stored include odontocete burst-pulse interclick intervals or the time-frequency content of whistles.

The final three optional elements of a *Detection* are *Image*, *Audio*, and *Comments*. The first two allow storage of a small amount of data from the acoustic record in image or audio formats and are intended to permit storage of illustrative examples. *Comments* provide analysts or software the capability of noting any interesting or unusual aspects of the detection.

4. Implementation

The schemata described above represent a proposed method for specifying passive acoustic metadata. Such relationships can be implemented using a variety of database technologies. The Tethys Metadata Workbench provides a reference implementation of the Tethys schemata that uses Oracle Corporation's freely available open source Berkeley DBXML (Redwood Shores, CA) data store. While complete details are beyond the purview of this article, we provide a brief overview of the current implementation and capabilities. Details on version 1.0 (current version is 2.3) of the Tethys workbench can be found in Roch et al. (2013).

4.1. Architecture

The Tethys workbench uses a client-server architecture where analysts and scientists connect to a server from their own workstations. Communication is via a representational state transfer (RESTful) paradigm (Fielding, 2000). Messages are sent using hypertext transport protocol (HTTP) by client programs as web requests and the server sends a response. Prior to the introduction of RESTful architectures, web servers needed to maintain detailed information about current data associated with each client conducting sequential operations. RESTful architectures encode this information into the web address that a client program uses, simplifying the web server design.

Implementation of this architecture is built using the open source Python CherryPy framework (CherryPy Team, 2014). A set of client libraries have been developed for Java, Python, and Matlab that permit users to create, retrieve, update, and delete information from the server. Client libraries for the R programming language are in development.

The server currently must be run on 64-bit Windows (Microsoft Corporation, Redmond, WA) platforms to support importing metadata from proprietary Microsoft formats such as Excel workbooks and the Access database. Client software is highly portable and has been tested on Windows and open source Linux (Linux Foundation, San Francisco, CA) platforms. For Windows platforms, an installation program allows the user to select the appropriate components to install on their computer.

4.2. Adding metadata

Utilities exist to translate inputs from sources such as databases, spreadsheets, comma separated value (CSV) files, etc. An XML file must be provided to map names from the data source (e.g. column headers in spreadsheet/CSV files, or table attribute names in a relational database) to those used in the Tethys schemata. Data translation is driven by the contents of the data source translation map file; there is no programming required (Fig. 9).

```
<Mapping>
  <Name>SIO.SWAL.Detections.Analyst.v1</Name>
  ...
  <Directives>
    <Detections>
      ...
      <Effort>
        <Sheet name="MetaData">
          <Entry>
            <Source> [Effort Start] </Source>
            <Kind> DateTime </Kind>
            <Dest> Start </Dest>
          </Entry>
          <Entry>
            <Source> [Effort End] </Source>
            <Kind> DateTime </Kind>
            <Dest> End </Dest>
          </Entry>
        </Sheet>
        <Sheet name="Effort">
          <Kind>
            <Entry>
              <Source> [Species Code] </Source>
              <Kind> speciescode </Kind>
              <Dest> SpeciesID </Dest>
            </Entry>
            <Entry>
              <Source> [Call] </Source>
              <Kind> Calltype </Kind>
              <Dest> Call </Dest>
            </Entry>
            <Entry>
              <Attribute>
                <Source> [BinSize_m] </Source>
                <Dest> BinSize_m </Dest>
              </Attribute>
              <Source> [Granularity] </Source>
              <Dest> Granularity </Dest>
            </Entry>
          </Kind>
        </Sheet>
      </Effort>
      ...
    </Directives>
  </Mapping>
```

Fig. 9. Excerpt from a data source translation map designed to translate metadata stored in a multiple sheet workbook (e.g. Excel, Microsoft Corp., Redmond, WA) to an XML document compliant with the Tethys detection schema. This section handles the specification of analysis effort. For emphasis, context sensitive reserved tags are shown as **bold** text and control the translation. Other tags (e.g. *Effort*) will appear in the generated XML. The **Sheet** element specifies the name of the sheet from which data are to be translated. Workbook sheets are expected to have column headers naming each type of data. Data in these columns are translated to elements through the use of **Entry** elements. The first **Entry** maps the column containing the text "Effort Start" on sheet *MetaData* to the *Start* element and will produce an element with the start date value. This is repeated for each entry in the sheet. Details are available in the Tethys manual (<http://tethys.sdsu.edu>).

If one prefers to generate XML directly from detection algorithms, the provided Nilus library can be called from the user's detection code to generate documents that conform to Tethys schemata. Nilus is written in Java and is accessible from any language that supports a Java interface (e.g. Matlab, R) and does not require installation of other Tethys products. Regardless of the import source, Berkeley DBXML validates the document against the appropriate schema and the server provides additional constraint validations that cannot be specified by the XML schema (e.g. ensuring that detection *End* times are after *Start* times).

4.3. Querying metadata

Queries to the database are written in the XQuery language (Walmsley, 2006), and the Matlab client library provides numerous

functions for querying the server without needing to learn XQuery. The Matlab client also provides visualization capabilities (e.g. Fig. 10). The query facility also supports data retrieval of solar and lunar ephemeris data from the NASA Jet Propulsion Laboratory's Horizons system (Giorgini et al., 1996) as well as data from any data source using NOAA's Environmental Research Division's Data Access Program (ERDDAP; Simons, 2011) protocol which provides access to a large variety of planetary data sources such as the NOAA Tropical Atmospheric Ocean (TAO) buoys, the United Kingdom's Met Office surface temperature anomaly data, or NASA's Ocean Color.

5. Discussion

The Tethys schemata have been used to represent nearly 300 years of combined systematic analysis effort across multiple species and many instrument deployments, producing millions of detections via analyses of the authors. Tethys has been used in work by several of the authors of this manuscript, ranging from production of reports on long-term presence of cetaceans to research on classification and formal spatiotemporal analysis (Baumann-Pickering et al., 2014; Roch et al., 2015; Širović et al., 2015). The ability to query information related to consistent descriptions of instrumentation, detections, and physical phenomena greatly enhances the researcher's ability to analyze large, complex data sets. Ecosystems vary at multiple year time scales (e.g. influences of El Niño and the Pacific Decadal Oscillation; Mantua et al., 1997; Rasmussen and Carpenter, 1982). Therefore, methods that focus on long-term ecosystem trends must have the potential to analyze and visualize data over multiple years or decades. As an example, the visualization of the multi-year diel and detection datasets used to produce Fig. 10 were queried and plotted with approximately 5 minutes of programming effort by an experienced Tethys user and a few minutes of runtime for each.

While Tethys provides flexible schemata for describing a variety of situations, it is expected that the schemata will continue to evolve over time. The flexible nature of XML schema specifications allows Tethys users to define new and complex data representations when they need them, such as additional time-frequency information (Fig. 2) that is not currently part of the Tethys schemata.

Over time, it is possible that the use of some researcher-defined extensions to schemata may become prevalent. Examples of this include parameters that may provide insight to a call that has a graded structure. We expect such extensions to be incorporated into the schemata. Currently, the procedure for this consists of proposing the extension

with the Tethys development team, but should the use of the Tethys schemata become more widespread, more formal procedures will be developed as part of a standards process. Schemata changes that simply involve extensions will not require any changes to existing XML documents nor to any code that queries the Tethys server. Changes to element names would require modifications to the XML, and the mechanisms for doing this are dependent upon the implementation but are a tractable problem.

A difficult area for the current detection schema is the representation of song. We follow the definition of Berwick et al. (2011), where a song consists of sequences of distinct sound types that may be repeated. These sequences can be organized into higher organizational structures. A simple example of this is the aforementioned Northeast Pacific blue whale AB song, whereas the humpback whale (*Megaptera novaeangliae*) has a more complex song with multiple organizational levels (Payne and McVay, 1971). To represent the hierarchical nature of song, elements could be nested (e.g. adding structure to the *Parameters/UserDefined* element of a detection). However, doing so would make it difficult to perform analysis on the individual song units. As a consequence, the currently recommended strategy for representing song is to create individual detections for each song unit and to create detections for song structure that reference the individual units and/or higher levels of organization such as phrases. As song has not been extensively documented using Tethys in any of the authors' organizations, the efficacy of this recommended approach has yet to be demonstrated and further refinements are likely to be needed before representation of song is more formally introduced into the detections schema.

Any system capable of representing complex data relationships has a learning curve, and potential users should not expect immediate productivity. The current set of client libraries has been developed for analysts who will be performing complex statistical analyses across large datasets, and the focus is on providing programmatic interfaces, the ability to organize data for statistical analysis, and visualization capabilities to enable users to address sophisticated hypotheses. The Tethys documentation (<http://tethys.sdsu.edu>) provides detailed instructions on getting started with the system including a "cookbook" guide for getting started with the Matlab client.

6. Conclusions

The Tethys schemata and workbench implementation are a needed extension to the capabilities for processing large bioacoustic data sets. The current implementation has been used to annotate and derive




Fig. 10. Automated detections of blue whale B calls (left) and echolocation clicks from odontocetes (right) over multiple years from an instrument deployed in the Southern California Bight (33.514° N, 119.25° W) showing seasonal and occasional diel patterns. Dark blue shading represents detections. Light blue shaded rectangular regions indicate periods for which there was no analysis effort, and the gray hourglass pattern on the odontocete plot shows nighttime as retrieved from the Horizons ephemeris system, making the diel activity pattern apparent.

information from millions of cetacean, pinniped, fish, and elephant call detections as well as anthropogenic sources from most ocean basins over nearly a decade of deployments. These were generated by a variety of algorithms and numerous users, thus demonstrating their ability to represent long-term bioacoustic data in a productive manner. Tethys is suitable for research in a wide variety of applications such as density and abundance estimates, long-term, seasonal, and diel patterns, social network analysis, etc.

The Tethys schemata, metadata server/client workbench, sample metadata, and associated documentation are freely available from <http://tethys.sdsu.edu>.

Acknowledgments

The authors wish to thank Katie Sieving, Rebecca Trout Fryxell, and the anonymous reviewers whose comments enhanced the clarity of this manuscript. We thank Daniel Hwang for his contributions to earlier versions of the workbench implementation, Lisa Munger for contributions to the schemata discussions and early prototypes of some of the visualization modules, Karlina Merckens for her assistance in developing standards for call names, and members of the Scripps Whale Acoustics Lab for being early adopters of the Tethys Metadata Workbench. Funding for this project was provided by The National Oceanographic Partnership Program, award N00014-11-1-0697 (The Office of Naval Research, Michael Weise, and The Bureau of Ocean Energy Management, Jill Lewandowski and Jim Price) and The United States Navy Living Marine Resources Program, award N00014-15-1-2299 (Anu Kumar).

Appendix A. Suggested call names for cetaceans

Recommended names for calls produced by cetaceans can be found online at <http://dx.doi.org/10.1016/j.ecoinf.2015.12.002>.

References

- Aguilar de Soto, N., Madsen, P.T., Tyack, P., Arranz, P., Marrero, J., Fais, A., Revelli, E., Johnson, M., 2011. No shallow talk: cryptic strategy in the vocal communication of Blainville's beaked whales. *Mar. Mamm. Sci.* E75–E92.
- Anorim, M.C.P., 2006. Diversity of sound production in fish. In: Collin, S.P., Møller, P., Kapoor, B.G. (Eds.), *Communication in Fishes*. Science Publishers, Enfield, NH, pp. 71–105.
- Barlow, J., Taylor, B.L., 2005. Estimates of sperm whale abundance in the Northeastern Temperate Pacific from a combined acoustic and visual survey. *Mar. Mamm. Sci.* 21, 429–445.
- Baumann-Pickering, S., Roch, M.A., Brownell Jr., R.L., Simonis, A.E., McDonald, M.A., Solsona-Berga, A., Oleson, E.M., Wiggins, S.M., Hildebrand, J.A., 2014. Spatio-temporal patterns of beaked whale echolocation signals in the North Pacific. *PLoS ONE* 9, 17.
- Berwick, R.C., Okanoya, K., Beckers, G.J.L., Bolhuis, J.J., 2011. Songs to syntax: the linguistics of birdsong. *Trends Cogn. Sci.* 15, 113–121.
- Blackwell, S.B., Nations, C.S., McDonald, T.L., Greene, C.R., Thode, A.M., Guerra, M., Macrander, A.M., 2013. Effects of airgun sounds on bowhead whale calling rates in the Alaskan Beaufort Sea. *Mar. Mamm. Sci.* 29, E342–E365.
- Botts, M., McKee, L., 2003. A sensor model language: moving sensor data onto the Internet. *Sensors* 20, 30–33.
- Buckland, S.T., 2006. Point-transect surveys for songbirds: robust methodologies. *Auk* 123, 345–357.
- Buckland, S.T., Summers, R.W., Borchers, D.L., Thomas, L.E.N., 2006. Point transect sampling with traps or lures. *J. Appl. Ecol.* 43, 377–384.
- Chang, F., Dean, J., Ghemawat, S., Hsieh, W.C., Wallach, D.A., Burrows, M., Chandra, T., Fikes, A., Gruber, R.E., 2008. Bigtable: a distributed storage system for structured data. *ACM Trans. Comput. Syst.* 26, 1–26.
- CherryPy Team, 2014. CherryPy – A Minimalist Python Web Framework Accessed July 28, 2014 <http://www.cherrypy.org>.
- Clark, C.W., 1982. The acoustic repertoire of the southern right whale, a quantitative analysis. *Anim. Behav.* 30, 1060–1071.
- Codd, E.F., 1970. A relational model of data for large shared data banks. *Commun. ACM* 13, 377–387.
- Connolly, D., Khare, R., Rifkin, A., 2007. The evolution of web documents: the ascent of XML. *World Wide Web J.* 2, 119–128.
- Deecke, V.B., Janik, V.M., 2006. Automated categorization of bioacoustic signals: avoiding perceptual pitfalls. *J. Acoust. Soc. Am.* 199, 645–653.
- Erbe, C., King, A.R., 2008. Automatic detection of marine mammals using information entropy. *J. Acoust. Soc. Am.* 124, 2833–2840.
- Federal Geographic Data Committee, 2014. National Geospatial Data Asset Management Plan: Portfolio Management Implementation Plan for the OMB Circular A-16 Supplemental Guidance (Steering Committee Approved: March 20, 2014) Accessed August 1, 2014 <http://www.fgdc.gov/policyandplanning/a-16/ngda-management-plan>.
- Fegraus, E.H., Andelman, S., Jones, M.B., Schildhauer, M., 2005. Maximizing the value of ecological data with structured metadata: an introduction to ecological metadata language (EML) and principles for metadata creation. *Bull. Ecol. Soc. Am.* 86, 158–168.
- Fielding, R.T., 2000. Architectural Styles and the Design of Network-based Software Architectures (PhD Dissertation) Computer Science, The University of CA, Irvine, CA, p. 162.
- Fornwall, M., Gisinier, R., Simmons, S.E., Moustahfid, H., Canonico, G., Halpin, P., Goldstein, P., Fitch, R., Weise, M., Cyr, N., Palka, D., Price, J., Collins, D., 2012. Expanding biological data standards development processes for US IOOS: visual line transect observing community for mammal, bird, and turtle data. US IOOS Summit Community White Paper. NOAA, Washington DC, p. 6.
- Fujioka, E., Soldevilla, M.S., Read, A.J., Halpin, P.N., 2014. Integration of passive acoustic monitoring data into OBIS-SEAMAP, a global biogeographic database, to advance spatially-explicit ecological assessments. *Ecol. Inform.* 21, 59–73.
- Gillespie, D., Gordon, J., McHugh, R., McLaren, D., Mellinger, D.K., Redmond, P., Thode, A., Trinder, P., Deng, X.-Y., 2008. PAMGUARD: semiautomated, open source software for real-time acoustic detection and localisation of cetaceans. *Proc. Inst. Acoustics. Inst. Acoust.*, Southampton, UK, pp. 54–62.
- Gillespie, D., Caillat, M., Gordon, J., White, P., 2013. Automatic detection and classification of odontocete whistles. *J. Acoust. Soc. Am.* 134, 2427–2437.
- Giorgini, J.D., Yeomans, D.K., Chamberlin, A.B., Chodas, P.W., Jacobson, R.A., Keesey, M.S., Lieske, J.H., Ostro, S.J., Standish, E.M., Wimberly, R.N., 1996. JPL's on-line solar system data service. *Bull. Am. Astron. Soc.* 28, 1158.
- Graybeal, J., Isenor, A.W., Rueda, C., 2012. Semantic mediation of vocabularies for ocean observing systems. *Comput. Geosci.* 40, 120–131.
- Guan, S., Moustahfid, H., Milan, A., Mize, J., 2014. A Metadata Convention for Passive Acoustic Recordings, Version 1.0. US Integrated Ocean Observing System, Silver Spring, MD, p. 47.
- Halpin, P.N., Read, A.J., Fujioka, E., Best, B.D., Donnelly, B., Hazen, L.J., Kot, C., Urian, K., LaBrecque, E., Dimatteo, A., Cleary, J., Good, C., Crowder, L.B., Hyrenbach, K.D., 2009. OBIS-SEAMAP the world data center for marine mammal, sea bird, and sea turtle distributions. *Oceanography* 22, 104–115.
- Hawkins, A.D., 1986. Underwater sound and fish behaviour. In: Pitcher, T.J. (Ed.), *The Behaviour of Teleost Fishes*. Croom and Helm, London, UK, pp. 114–151.
- Hernandez, K.M., Risch, D., Cholewiak, D.M., Dean, M.J., Hatch, L.T., Hoffman, W.S., Rice, A.N., Zemeckis, D., Van Parijs, S.M., 2013. Acoustic monitoring of Atlantic cod (*Gadus morhua*) in Massachusetts Bay: implications for management and conservation. *ICES J. Mar. Sci.* 70, 628–635.
- International Standards Organization, 2003. Geographic Information - Metadata. International Standards Organization, Geneva, p. 140.
- ITIS Organization, 2014. Integrated Taxonomic Information System (ITIS) Accessed July 4, 2014 <http://www.itis.gov>.
- Joint W3C/IETF URI Planning Interest Group, 2002. Uniform Resource Identifiers (URIs), URLs, and Uniform Resource Names (URNs): Clarifications and Recommendations. In: Mealling, M., Dennenberg, R. (Eds.), *Request for Comments series*. Internet Engineering Task Force, p. 11.
- Kandia, V., Stylianou, Y., 2006. Detection of sperm whale clicks based on the Teager-Kaiser energy operator. *Appl. Acoust.* 67, 1144–1163.
- Leavitt, N., 2010. Will NoSQL databases live up to their promise? *Computer* 43, 12–14.
- Mantua, N.J., Hare, S.R., Zhang, Y., Wallace, J.M., Francis, R.C., 1997. A Pacific interdecadal climate oscillation with impacts on salmon production. *Bull. Am. Meteorol. Soc.* 78, 1069–1079.
- Marine and Coastal Spatial Data Subcommittee, 2012. Coastal and Marine Ecological Classification Standard. In: US Government report number: FGDC-STD-018-2012, Committee F.G.D. (Ed.) Federal Geographic Data Committee, Reston, VA, USA, p. 343.
- Marques, T.A., Thomas, L., Ward, J., DiMarzio, N., Tyack, P.L., 2009. Estimating cetacean population density using fixed passive acoustic sensors: an example with Blainville's beaked whales. *J. Acoust. Soc. Am.* 125, 1982–1994.
- McDonald, M.A., Mesnick, S.L., Hildebrand, J.A., 2006. Biogeographic characterisation of blue whale song worldwide: using song to identify populations. *J. Cetac. Res. Manage.* 8, 55–65.
- McDonald, M.A., Hildebrand, J.A., Wiggins, S.M., Johnston, D.W., Polovina, J.J., 2009. An acoustic survey of beaked whales at Cross Seamount near Hawaii. *J. Acoust. Soc. Am.* 125, 624–627.
- McGuinness, D.L., 2003. Ontologies come of age. In: Fensel, D., Hendler, J.A., Lieberman, H., Wahlster, W. (Eds.), *Spinning the Semantic Web: Bringing the World Wide Web to its Full Potential*. MIT Press, Cambridge, MA, pp. 171–196.
- Mellinger, D.K., Martin, S.W., Morrissey, R.P., Thomas, L., Yosco, J.J., 2011. A method for detecting whistles, moans, and other frequency contour sounds. *J. Acoust. Soc. Am.* 129, 4055–4061.
- National Oceanic and Atmospheric Administration, 2013. IOOS, Contribute Your Data Accessed January 10, 2014 http://www.ioos.noaa.gov/data/contribute_data.html.
- Nosal, E.-M., 2013. Methods for tracking multiple marine mammals with wide-baseline passive acoustic arrays. *J. Acoust. Soc. Am.* 134, 2383–2392.
- Oleson, E.M., Calambokidis, J., Barlow, J., Hildebrand, J.A., 2007a. Blue whale visual and acoustic encounter rates in the southern California bight. *Mar. Mamm. Sci.* 23, 574–597.
- Oleson, E.M., Calambokidis, J., Burgess, W.C., McDonald, M.A., LeDuc, C.A., Hildebrand, J.A., 2007b. Behavioral context of call production by eastern North Pacific blue whales. *Mar. Ecol. Prog. Ser.* 330, 269–284.

- Oleson, E.M., Wiggins, S.M., Hildebrand, J.A., 2007c. Temporal separation of blue whale call types on a southern California feeding ground. *Anim. Behav.* 74, 881–894.
- Open Geospatial Consortium, 1994. OGC Standards and Supporting Documents Accessed June 1, 2010 <http://www.opengeospatial.org/standards>.
- Payne, R.S., McVay, S., 1971. Songs of humpback whales. *Science* 173, 585–597.
- Rasmuson, E.M., Carpenter, T.H., 1982. Variations in tropical sea-surface temperature and surface wind fields associated with the southern oscillation El-Nino. *Mon. Weather Rev.* 110, 354–384.
- Richardson, W.J., Greene, J.C.R., Malme, C.I., Thomson, D.H., 1995. *Marine Mammals and Noise*. Academic Press, San Diego, CA.
- Risch, D., Clark, C.W., Dugan, P.J., Popescu, M., Siebert, U., Van Parijs, S.M., 2013. Minke whale acoustic behavior and multi-year seasonal and diel vocalization patterns in Massachusetts Bay, USA. *Mar. Ecol. Prog. Ser.* 489, 279–295.
- Roch, M.A., Baumann-Pickering, S., Batchelor, H., Hwang, D., Sirovic, A., Hildebrand, J.A., Berchok, C.L., Cholewiak, D., Munger, L.M., Oleson, E.M., Van Parijs, S., Risch, D., Soldevilla, M.S., 2013. Tethys: a workbench and database for passive acoustic metadata. *Oceans* 2013 (5 pp.).
- Roch, M.A., Stinner-Sloan, J., Baumann-Pickering, S., Wiggins, S.M., 2015. Compensating for the effects of site and equipment variation on delphinid species identification from their echolocation clicks. *J. Acoust. Soc. Am.* 137, 22–29.
- Simons, R.A., 2011. ERDDAP - The Environmental Research Division's Data Access Program Accessed Feb 1, 2012 <http://coastwatch.pfeg.noaa.gov/erddap>.
- Simpson, S.D., Meekan, M., Montgomery, J., McCauley, R., Jeffs, A., 2005. Homeward sound. *Science* 308, 221–221.
- Širović, A., Hildebrand, J.A., Wiggins, S.M., McDonald, M.A., Moore, S.E., Thiele, D., 2004. Seasonality of blue and fin whale calls and the influence of sea lee in the Western Antarctic Peninsula. *Deep-Sea Res. II Top. Stud. Oceanogr.* 51, 2327–2344.
- Širović, A., Rice, A., Chou, E., Hildebrand, J.A., Roch, M.A., 2015. Seven years of blue and fin whale call abundance in Southern California. *Endanger. Species Res.* 28, 61–75.
- Soldevilla, M.S., Wiggins, S.M., Hildebrand, J.A., 2010. Spatio-temporal comparison of Pacific white-sided dolphin echolocation click types. *Aquat. Biol.* 9, 49–62.
- Steiner, W.W., 1981. Species-specific differences in pure tonal whistle vocalizations of five western North Atlantic dolphin species. *Behav. Ecol. Sociobiol.* 9, 241–246.
- Urick, R.J., 1983. *Principles of Underwater Sound*. McGraw Hill.
- Versluis, M., Schmitz, B., von der Heydt, A., Lohse, D., 2000. How snapping shrimp snap: through cavitating bubbles. *Science* 289, 2114–2117.
- Walmsley, P., 2002. *Definitive XML Schema*. Prentice Hall PTR, Upper Saddle River, NJ.
- Walmsley, P., 2006. *XQuery*. O'Reilly, Farnham, UK.
- Wieczorek, J., Döring, M., De Giovanni, R., Robertson, T., Vieglais, D., 2009. Darwin Core Accessed July 3, 2013 <http://tdwgng.tdwg.org/standards/>.
- Wieczorek, J., Bloom, D., Guralnick, R., Blum, S., Döring, M., Giovanni, R., Robertson, T., Vieglais, D., 2012. Darwin Core: an evolving community-developed biodiversity data standard. *PLoS ONE* 7, 8.
- Wolf, M., Wicksteed, C., 1997. Date and Time Formats Accessed July 1, 2015 <http://www.w3.org/TR/NOTE-datetime>.
- Zimmer, W.M.X., 2011. *Passive Acoustic Monitoring of Cetaceans*. Cambridge University Press, Cambridge.