

Making Sense of Sensor Data Using Ontology: A Discussion for Road Vehicle Classification

Markus Stocker, Mauno Rönkkö, Mikko Kolehmainen^a

^a*University of Eastern Finland, P.O. Box 1627, 70211 Kuopio, Finland
(markus.stocker@uef.fi, mauno.ronkko@uef.fi, mikko.kolehmainen@uef.fi)*

Abstract: In environmental engineering, sensor measurement is a process undertaken with respect to an environmental domain, i.e. an area of interest, to measure a domain property. Knowledge of a domain, meaning its concepts and relations that hold among them, can be formally represented by means of ontology. Therefore, given an ontology for an environmental domain, it seems reasonable to suggest that sensor data acquisition can be translated into ontological knowledge acquisition. We demonstrate this translation for the domain of road vehicle classification by measurement of vibration. We show how supervised machine learning is applied to learn a function that maps sensor data to ontological concepts. Hence, we abstract from both the physical sensor layer and the sensor data layer by discarding raw measurement data and retaining the knowledge conveyed by these data. We show how rules can be used to infer new domain knowledge, such as vehicle velocity.

Keywords: ontology; sensor data; machine learning; knowledge discovery; environmental ontology learning

1 INTRODUCTION

According to Finkelstein [1982] “measurement is the process of empirical, objective, assignment of numbers to properties of objects or events of the real world in such a way as to describe them”. In environmental engineering, a sensor is a device that transforms the signal of a physical property in the environment of the sensor into sensor data (Bonnet et al. [2001]) and sensor measurement is the recurring application of such transformation for a certain duration and location. The dimensions of measured physical property, duration, location, and purpose of measurement delineate an area of interest, i.e. a domain. In computer science, knowledge of a domain, meaning the concepts of some area of interest and relations that hold among them, can be formally represented by means of ontology, defined as an explicit specification of a conceptualization (Gruber [1993]).

Managing, processing, and making sense of sensor data is an ongoing challenge (Tollefson [2011]; Balazinska et al. [2007]) despite recent advancement in low-cost, low-power, small-size, wireless technology as well as communication protocols, algorithms, and programming models (Akyildiz et al. [2002]; Bharathidasan et al. [2002]); in sensor data management and processing (Carney et al. [2002]; Bonnet et al. [2001]); and in semantic description of sensors, sensor networks, and sensor data (Sheth et al. [2008]; Compton et al. [2009]).

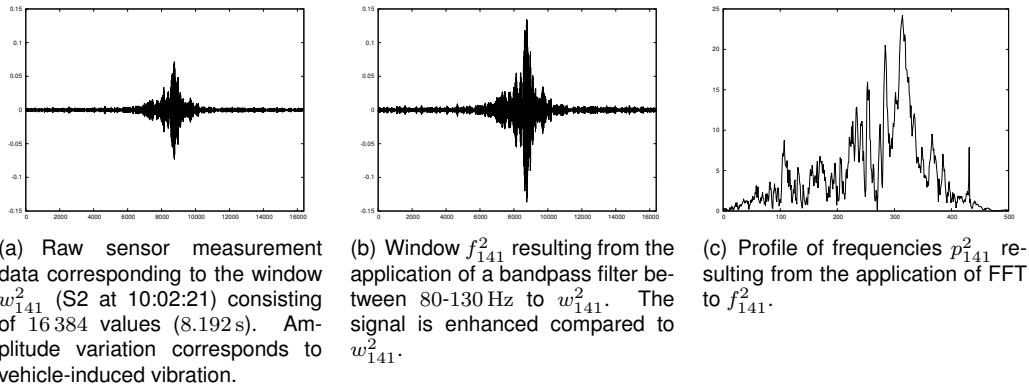


Figure 1: Data processing from raw sensor measurement data to the profile of frequencies for 8.192 s of measurement by S2 starting at 10:02:21.

Sheth et al. [2008] present the Semantic Sensor Web in which sensor data is annotated with spatial, temporal, and thematic semantic metadata deemed to be “essential for discovering and analyzing sensor data”. Compton et al. [2009] review eleven sensor ontologies for the range and expressive power of their concepts, which are used to describe, for instance, the deployment, configuration, components, data, observations, location, accuracy, or sampling frequency of sensors. Terminologies to describe the characteristics of sensors and sensor networks aimed at improving the integration and communication among sensors and networks (Sheth et al. [2008]) have, hence, received considerable attention.

Classifying entities observed by sensors has a long tradition and extensive literature. To name a few relevant studies, using vibration sensors to monitor pavement acceleration and magnetometer sensors to detect vehicles, Bajwa et al. [2011] propose a method to estimate axle count and spacing for trucks. Jackowski and Wantoch-Rekowski [2005] discuss the problem of using neural networks for military vehicle classification using ground vibration. Nooralahiyan et al. [1997] use a directional microphone and a time-delay neural network to classify road vehicles based on their acoustic signature.

Given an ontology for an environmental domain, the aim of this paper is to demonstrate and discuss the acquisition of knowledge about entities observed by sensors when machine learning is used to classify measurement data, and the formal representation of such knowledge in ontology. As such our work relates to the architecture discussed by Liu and Zhao [2005] of a system that can be queried for high-level events without requiring handling of raw signals. It also relates to Stocker et al. [2011] where the authors discuss a method to learn, from measurement data, an atom of a rule in the schema of an environmental ontology.

2 MATERIALS AND METHODS

Road-pavement vibration was measured using three CEF C3M01 accelerometer vibration sensors developed for condition monitoring and machinery maintenance by Control Express Finland (CEF) Oy.¹ (CEF C3M01 sensors are now manufactured by Webrosensor Oy as WBS CM301.²) The sensors – thereafter referred to as S1, S2, and S3 – were

¹<http://www.cef.fi/>

²<http://www.wbs.fi/>

Listing 1: Axioms used to express domain knowledge

Vehicle \sqsubseteq FeatureOfInterest
 LightVehicle \sqsubseteq Vehicle
 HeavyVehicle \sqsubseteq Vehicle
 LightVehicle \sqcap HeavyVehicle $\sqsubseteq \perp$
 DrivingSide \sqsubseteq Property
 DrivingSpeed \sqsubseteq Property

 SensingDevice (S1)
 SensingDevice (S2)
 SensingDevice (S3)

installed at the right side of a road.³ We visually monitored the road by means of a AXIS 211W Wireless Network Camera with an Outdoor Antenna Kit AXIS 211W (Axis Communications [2011]). We acquired vibration data from the three CEF C3M01 sensors and image data from the AXIS camera for a total of six hours on August 30, 2011 between 10 AM and 4 PM. We acquired 42 962 432 measurement values from S1, 42 937 345 from S2, and 42 988 810 from S3. We acquired 25 076 image files from the AXIS camera.

Sensor data was semi-automatically processed to detect and appropriately label the signal corresponding to the vibration induced by vehicles in measurement data. We first processed the AXIS camera files to visually identify vehicle occurrences. Each identified vehicle occurrence was described with metadata for the vehicle type (e.g. personal car), time of the vehicle crossing the approximate location of S2, and driving side (left or right). We found a total of 185 vehicle occurrences (i.e. approximately a vehicle every two minutes on average). Such metadata was used to link occurrences identified in camera data with those (automatically) detected in vibration data. For each vibration sensor s^i whereby $i = 1, 2, 3$ for S1, S2, and S3, respectively, we processed the time series for 6 hours of measurement by means of a window w_j^i of length 16 384, i.e. 8.192 s of measurement values (or 0 in case of missing values) whereby $j = 0, \dots, 21\,599$ is the total number of seconds. Figure 1(a) is a plot of w_{141}^2 , i.e. the values corresponding to 8.192 s of measurement by S2 starting at 10:02:21. To each w_j^i we applied a bandpass filter to suppress frequencies outside 80 Hz and 130 Hz,⁴ leading to a filtered window f_j^i . Figure 1(b) is a plot of f_{141}^2 resulting from the application of the bandpass filter to w_{141}^2 . For each f_j^i we, thereafter, extracted the profile, p_j^i , of frequencies between 80 Hz and 130 Hz computed using Fast Fourier Transform (FFT). Figure 1(c) is a plot of p_{141}^2 , the profile of frequencies resulting from the application of FFT to f_{141}^2 . Finally, we computed the sum, Σp_j^i , for the values of p_j^i . The Σp_j^i , whereby $j = 0, \dots, 21\,599$, form an evenly spaced time series with interval length 1 s and (positive only) peaks of amplitude and width reflecting the magnitude and duration, respectively, of road-pavement vibration as measured by sensor s^i between 10 AM and 4 PM on August 30, 2011. The first derivative of such a time series was calculated to detect the starting (positive rate of change) and ending (negative rate of change) time of a *possible* vehicle occurrence. We empirically determined that 50 and -50 for the positive and negative rate of change, respectively, were appropriate threshold values to detect the starting and ending of vibration possibly induced by vehicles. Note that at this point we could not exclude detecting vibration-like signal that is explained by something other than a vehicle. Hence, we had to link j determined as the time index at which there may be a vehicle within the following 8.192 s with the 185 vehicle occurrences

³Right side with respect to camera perspective

⁴Experimentally, we found that the spectral energy related to road-pavement vibration induced by vehicles strongly lies between 80 Hz and 130 Hz.

identified in camera data. This manual step was executed conservatively, meaning that we linked occurrences only when we were fairly certain to link correctly and we discarded links to vibration data f_j^i that contained amplitude change due to a vehicle occurrence as well as due to one or more other (typically unexplained) reasons.

By linking we, thus, constructed *labelled* p_j^i data which were used to generate training datasets for supervised learning. Such datasets only included training samples occurring between 10:00 AM and 2:50 PM. We evaluated the machine learning classification performance for two tasks: (1) vehicle detection and (2) vehicle classification.⁵ The aim in vehicle *detection* was to classify p_j^i according to whether or not a vehicle occurred in the corresponding window w_j^i (training classes `vehicle` and `no-vehicle`). For this purpose, we additionally extracted 134 p_j^i to serve as training samples without vehicle occurrence (j the index at which w_j^i did not contain a vehicle occurrence). The aim in vehicle *classification* was to classify p_j^i according to light and heavy vehicles (training classes `light` and `heavy`). We used WEKA (Hall et al. [2009], version 3.6.5) to train a Multilayer Perceptron (MLP) neural network classifier.

Given our domain of road vehicle classification by measurement of vibration we extended the Semantic Sensor Network (SSN) Ontology⁶ (W3C Semantic Sensor Network Incubator Group [2009]) to accommodate the class `Vehicle` as domain-specific SSN feature-of-interest; the classes `LightVehicle` and `HeavyVehicle`, both subclasses of `Vehicle`; the disjointness of `LightVehicle` and `HeavyVehicle`, meaning that a vehicle can be either light or heavy, but not both; the two domain-specific SSN properties (ontology classes) `DrivingSide` and `DrivingSpeed`; and the three sensors individuals of `SensingDevice`. Listing 1 provides an overview of the axioms used to express domain knowledge. We (programmatically) populated the ontology with individuals, ω , instances of SSN `Observation`, i.e. `Observation`(ω), with SSN relation `observedBy` to an individual, i , of `SensingDevice`; SSN relation `observationResultTime` for the observation time, j ; and SSN relation `featureOfInterest` with an individual, ψ , instance of `Vehicle`, i.e. `Vehicle`(ψ). Here we considered observations occurring between 2:50 PM and 3:00 PM on August 30, 2011 made by any of the sensors s^i . We used Protégé⁷ (version 4.1) and Jena⁸ (Carroll et al. [2003], version 2.7.0) to (programmatically) manage the ontology, and the Web Ontology Language (OWL) (W3C OWL Working Group [2009]) and Resource Description Framework (RDF) (Manola and Miller [2004]) technologies.

To demonstrate rule-based inference we defined two rules $p \rightarrow q$ of interest to our domain. The first rule stated that the vehicles related to two observations with result time difference below 8s are *same*, i.e. the same physical entity. This rule was motivated by the distance of approximately 45m between consecutive sensors and the average low-volume traffic of our domain. For each observation pair with relations to *same* vehicles, the second rule inferred the velocity of the vehicle. Velocity determined the vehicle's driving side and speed. We used the SSN `Property` classes `DrivingSide` and `DrivingSpeed` to express this knowledge in our ontology. Machine learning inference is demonstrated by feature-of-interest classification of individuals, ψ , instances of `Vehicle`. Here we aimed at specializing the class a vehicle ψ is an instance of to either `LightVehicle` or `HeavyVehicle`, by using data p_j^i for the individual `Observation`(ω) related to ψ . This classification was performed using the machine learning methods and training datasets described above. Hence, for each observation ω occurring between 2:50 PM and 3:00 PM made by any of the sensors s^i we classified p_j^i

⁵Note that the vehicle detection task is different from the detection of (possible) vehicle occurrences in vibration data, discussed above.

⁶<http://www.w3.org/2005/Incubator/ssn/ssnx/ssn>

⁷<http://protege.stanford.edu>

⁸<http://incubator.apache.org/jena/>

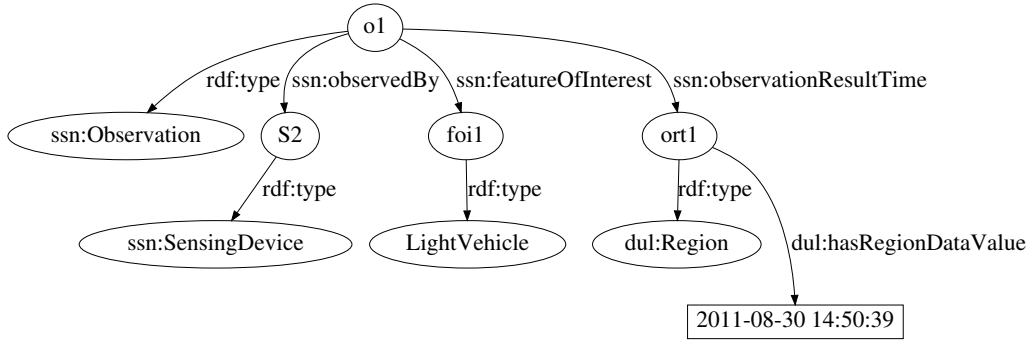


Figure 2: RDF graph describing the observation (o1) made by S2 at 14:50:39 on August 30, 2011 for a personal car classified as light vehicle.

Table 1: Vehicle occurrences with information on right (R) or left (L) driving side, between 2:50 PM and 3:00 PM as identified in camera data and observed by the three sensors. In the ontology, the SSN property `observationResultTime` is used to represent time, and the classes `light vehicle` (LV) or `heavy vehicle` (HV) are used as most specific vehicle type (T) for the SSN `FeatureOfInterest` associated to an observation. Misclassification is highlighted.

Occurrence				S1		S2		S3	
Label	Class	Side	HH:mm	ss	T	ss	T	ss	T
personal-car	LV	R	14:50	36	LV	39	LV	43	HV
van	LV	R	14:53	20	LV	24	LV	27	HV
truck	HV	L	14:57	56	HV	52	HV	49	HV
fire-truck	HV	R	14:58	39	HV	43	LV	46	HV

related to ω using a trained MLP classifier. The result of the classification was, thereafter, formalized in our ontology as either `LightVehicle(ψ)` or `HeavyVehicle(ψ)`, being ψ the vehicle related to ω . We implemented the described rule-based and machine learning inferences as processes performed on the populated ontology, i.e. for observations occurring between 2:50 PM and 3:00 PM.

3 RESULTS

We identified 165 vehicles in camera images between 10:00 AM and 2:50 PM, of which 87 (53%) were detected by S1, 134 (81%) by S2, and 133 (81%) by S3. A total of 10 distinct vehicle types were identified. For the vehicle detection task we grouped the 10 vehicle types into one training class (`vehicle`). For the vehicle classification task we grouped the 10 vehicle types into two training classes, according to vehicle weight (`light` and `heavy`). Classification performance (correctly classified instances) for the vehicle *detection* task resulted to be 92% for S1, 95% for S2, and 96% for S3 (94% on average). Classification performance for the vehicle *classification* task resulted to be 82% for S1, 75% for S2, and 83% for S3 (80% on average). Hence, we argue that it is possible to classify vehicles using vibration measurement data and that the presented methods are suitable to automatically map measurement data to ontological concepts. In so doing we acquire knowledge about vehicles observed by sensors by means of a learned generalization based on training individuals and we formalize such acquired knowledge in ontology. Figure 2 is a visual representation of the RDF graph describing the observation made by the sensing device S2 at result time with data value 2011-08-30

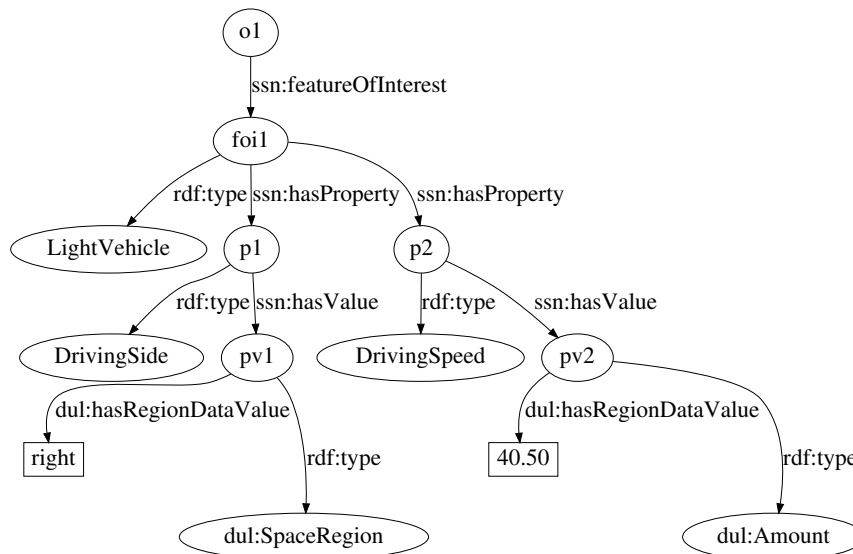


Figure 3: RDF sub-graph for the feature-of-interest (f_{oi1}) related to the observation (o_1) made by S2 at 14:50:39 expanded to include the properties of vehicle driving side and driving speed.

14:50:39 for a light vehicle – driving on the right road-side at speed 40.5 km h^{-1} (Figure 3).

Table 1 summarizes the key elements of the 12 observations corresponding to the 4 vehicles observed by the 3 sensors between 2:50 PM and 3:00 PM. In particular, we show the time (date is August 30, 2011) related to the observation by the SSN property `observationResultTime` and the most specific vehicle type for the SSN `FeatureOfInterest` related to an observation by the SSN property `featureOfInterest`. Time is given in hours and minutes (HH:mm) for an occurrence and in seconds (ss) for an observation. The rule-based reasoning process correctly inferred (not shown) that, for instance, the vehicle related to the observation by S2 at 14:53:24 is `sameAs` the vehicle related to the observation by S3 at 14:53:27 as well as `sameAs` the vehicle related to the observation by S1 at 14:53:20. Further, Figure 3 shows the inferred knowledge for vehicle driving speed and driving side. As we can see in Table 1 machine learning correctly mapped road-pavement vibration measurement data to the specific ontological vehicle class 9 times out of 12 (75%).

We argue that to translate sensor data about entities observed in a sensor network (Figure 1(c)) to symbolic knowledge (Figure 2) has a number of implications. First, we abstract from measurement data. As we have shown, our small sensor network consisting of three sensors acquired approximately 130 million values in six hours of measurement. Such data is of little interest if not for the knowledge conveyed by them. We understand the formalization of such knowledge in ontology as a way to make sense of sensor data. Second, we represent knowledge acquired from sensor networks in ontology. In so doing we allow for an integrated representation of knowledge about what is observed by a sensor network and for rule-based inference of domain knowledge, e.g. vehicle velocity. Hence, ontology facilitates the integration of knowledge acquired in sensor networks and the automatic discovery of new domain knowledge.

4 CONCLUSIONS

For the domain of road vehicle classification by measurement of vibration we have shown, using machine learning to classify measurement data, how knowledge about entities observed in a sensor network can be extracted from sensor data and can be formally represented by means of an ontology language.

Our results showed an acceptable machine learning classification performance for detecting and classifying observed vehicles using sensor data acquired by measurement of vibration. Hence, we used machine learning to translate sensor data about physical entities into symbolic knowledge about the entities, i.e. vehicles, and we formally represented such knowledge in a domain ontology. Rule-based inference was, thereafter, used to infer new knowledge about the observed entities, e.g. vehicle velocity. We have discussed a number of benefits resulting from representing knowledge acquired from sensor networks in ontology, in particular abstraction from measurement data, integrated representation of knowledge about what is observed, and rule-based inference.

In future work it is our aim to extend the presented methodology to work with a heterogeneous sensor network that includes sensors and data of diverse type, for instance the presented vibration sensors and acoustic sensors, and to explore domains other than vehicle classification, in particular also domains in which acquired knowledge in sensor data is not for observed entities, such as here for vehicles, but, e.g., for flux rates, such as for net carbon exchange between ecosystems and the atmosphere, as measured by eddy covariance (Baldocchi et al. [1988]).

ACKNOWLEDGMENTS

We wish to thank Paula Silvonen, research scientist at VTT Technical Research Center of Finland, for her valuable contribution to related work. The infrastructure to access and collect vibration and camera data, as well as the data, are part of research funded by Tekes, the Finnish Funding Agency for Technology and Innovation (funding decision number 40075/09). We thank the Measurement, Monitoring and Environmental Efficiency Assessment (MMEA) research program of the Cluster for Energy and Environment (CLEEN), Finland, for supporting the publication. This work was conducted using the Protégé resource, which is supported by grant LM007885 from the United States National Library of Medicine.

REFERENCES

- Akyildiz, I., W. Su, Y. Sankarasubramaniam, and E. Cayirci. A survey on sensor networks. *Communications Magazine, IEEE*, 40(8):102–114, August 2002.
- Axis Communications. AXIS 211W Network Camera Datasheet, 2011.
- Bajwa, R., R. Rajagopal, P. Varaiya, and R. Kavalier. In-pavement wireless sensor network for vehicle classification. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 85–96, 2011.
- Balazinska, M., A. Deshpande, M. Franklin, P. Gibbons, J. Gray, S. Nath, M. Hansen, M. Liebhold, A. Szalay, and V. Tao. Data Management in the Worldwide Sensor Web. *Pervasive Computing, IEEE*, 6(2):30–40, April-June 2007.
- Baldocchi, D. D., B. B. Hincks, and T. P. Meyers. Measuring Biosphere-Atmosphere Exchanges of Biologically Related Gases with Micrometeorological Methods. *Ecology*, 69(5):1331–1340, October 1988.

- Bharathidasan, A., V. An, and S. Ponduru. Sensor Networks: An Overview. Technical report, Department of Computer Science, University of California, Davis, 2002.
- Bonnet, P., J. Gehrke, and P. Seshadri. Towards Sensor Database Systems. In *Mobile Data Management*, volume 1987 of *Lecture Notes in Computer Science*, pages 3–14. Springer Berlin / Heidelberg, 2001. 10.1007/3-540-44498-X_1.
- Carney, D., U. Çetintemel, M. Cherniack, C. Convey, S. Lee, G. Seidman, M. Stonebraker, N. Tatbul, and S. Zdonik. Monitoring streams: a new class of data management applications. In *Proceedings of the 28th international conference on Very Large Data Bases*, VLDB '02, pages 215–226. VLDB Endowment, 2002.
- Carroll, J. J., I. Dickinson, C. Dollin, D. Reynolds, A. Seaborne, and K. Wilkinson. Jena: Implementing the Semantic Web Recommendations. Technical Report HPL-2003-146, HP Laboratories, Bristol, UK, 2003.
- Compton, M., C. Henson, H. Neuhaus, L. Lefort, and A. Sheth. A Survey of the Semantic Specification of Sensors. In *2nd International Workshop on Semantic Sensor Networks, at 8th International Semantic Web Conference*, 2009.
- Finkelstein, L. Theory and Philosophy of Measurement. In Peter Sydenham and Richard Thorn, editor, *Handbook of Measurement Science, Volume 1, Theoretical Fundamentals*, pages 1–30. John Wiley & Sons, 1982.
- Gruber, T. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.
- Hall, M., E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The WEKA Data Mining Software: An Update. In *SIGKDD Explorations*, volume 11, 2009.
- Jackowski, J. and R. Wantoch-Rekowski. Classification of wheeled military vehicles using neural networks. In *Systems Engineering, 2005. ICSEng 2005. 18th International Conference on*, pages 212–217, 2005.
- Liu, J. and F. Zhao. Towards semantic services for sensor-rich information systems. In *Broadband Networks, 2005. BroadNets 2005. 2nd International Conference on*, volume 2, pages 967–974, October 2005.
- Manola, F. and E. Miller. RDF Primer. Technical Report W3C Recommendation, W3C, 2004.
- Nooralahiyan, A. Y., M. Dougherty, D. McKeown, and H. R. Kirby. A field trial of acoustic signature analysis for vehicle classification. *Transportation Research Part C: Emerging Technologies*, 5(3-4):165–177, 1997.
- Sheth, A., C. Henson, and S. Sahoo. Semantic Sensor Web. *Internet Computing, IEEE*, 12(4):78–83, July-August 2008.
- Stocker, M., M. Rönkkö, F. Villa, and M. Kolehmainen. The relevance of measurement data in environmental ontology learning. In *Environmental Software Systems. Frameworks of eEnvironment*, volume 359 of *IFIP Advances in Information and Communication Technology*, pages 445–453. Springer Boston, 2011.
- Tollefson, J. US launches eco-network. *Nature*, 476(135), August 2011.
- W3C OWL Working Group. OWL 2 Web Ontology Language. Technical Report W3C Recommendation, W3C, 2009.
- W3C Semantic Sensor Network Incubator Group. Semantic Sensor Network Ontology. Technical report, W3C, 2009.