Dynamic Trust Scoring of Railway Sensor Information

Marcin Lenart, Andrzej Bielecki, Marie-Jeanne Lesot, Teodora Petrisor, Adrien Revault d'Allonnes

Dynamic Trust Scoring of Railway Sensor Information

Marcin Lenart[1,2,3], Andrzej Bielecki[3], Marie-Jeanne Lesot[2], Teodora Petrisor[1] and Adrien Revault d'Allonnes[4,2]

[1] Thales, Campus Polytechnique, Palaiseau, France,
[2] Sorbonne Université, CNRS, Laboratoire d'Informatique de Paris 6, LIP6, F-75005 Paris, France
[3] Chair of Applied Computer Science, Faculty of EAIIB AGH University of Science and Technology, Cracow, Poland
[4] Université Paris 8, LIASD EA 4383, Saint-Denis, France

Abstract. A sensor can encounter many situations where its readings can be untrustworthy and the ability to recognise this is an important and challenging task. It opens the possibility to assess sensors for forensic or maintenance purposes, compare them or fuse their information. We present a proposition to score a piece of information produced by a sensor as an aggregation of three dimensions called reliability, likelihood and credibility into a trust value that take into account a temporal component. The approach is validated on data from the railway domain.

Keywords: information scoring, sensor, trust, reliability, likelihood, credibility

1 Introduction

Information scoring (see e.g. [1]) aims at assessing the quality of available pieces of information and, in particular, the trust that can be put in them. It plays a crucial role in any decision-aid system. For instance, in an information fusion system, equally considering reliable and unreliable sources may severely cripple the results. Sensors are not an exception, the information they produce is often used to get enhanced knowledge about a given situation and the ability to differentiate between them in terms of quality is a much needed feature.

Indeed, sensors do not always produce correct information. There are many situations in which a sensor can fail, e.g. producing out of range values, when encountering unfavourable operating conditions, communication problems or other interferences. Knowing whether the information produced by sensors is trustworthy can be key in many aspects, for instance, to choose the ones with the highest quality level for a given time interval. It can also be used to predict maintenance operations for sensors with decreased quality of information.

As detailed in Section 2, current quality measurements for sensors are mainly based on scoring *reliability* either from meta-data [5, 7] or ground truth evaluation [3, 8]. Other systems include *credibility* to further improve scoring by com-

paring information with other sources [8]. These solutions can suffer from lack of external knowledge (meta-data or ground truth) which can make them unusable.

This paper aims to address these limitations by decreasing the dependence on meta-data or ground truth and incorporating statistical analysis into the computation. It proposes new definitions for three dimensions chosen such that different aspects of the source and the information can be captured.

The paper is organised as follows: Section 2 presents some of the current approaches to score information quality for sensors, in Section 3 the proposed process of information scoring is explained and in Section 4 the approach is illustrated on real-world data. Section 5 concludes the paper and discusses future research directions.

2 Literature Review

This section briefly discusses general Information Quality scoring and describes approaches dedicated to the special case where the considered information is provided by sensors.

General Information Quality Assessment. The task of information scoring is mainly addressed through the decomposition of its quality into components, assessed on different dimensions whose list and definitions vary depending on the author. Some examples are relevance and truthfulness [13], reliability and certainty [10], source-trustworthiness and information-credibility [2], sincerity, competence, intention of the source and plausibility [11] or trust [18, 4, 8, 15], see [17] for a complete list. One recent approach, introduced in [15], considers trust evaluation in a multivalued logic framework based on four dimensions: reliability and competence, which evaluate the source, and plausibility and credibility, which relate to the information content, spanning the range from source to information, from general to contextual and from subjective to objective.

Information Quality Dimensions in the Context of Sensor Measurement Many papers [5, 7, 9, 14] focus on the case of information provided by sensors. They often consider three dimensions, called reliability, contextual reliability and credibility, but vary in the way they are scored, as detailed below.

Reliability is generally understood as the ability of a system to perform its required functions under stated conditions for a specified time. It is an *a priori* assessment of the source.

Different approaches are considered to score reliability. In [7, 5] meta-information on the source are considered, e.g. its specification, protocol or environment. The gathered knowledge is then combined to propose a final reliability score. This approach is limited to the case where valuable meta-data are available.

A second approach to define reliability consists in viewing it as accuracy [8, 3] in the case where ground truth is available, i.e. knowledge about the expected

results. This suffers the same limitation as the previous approach since ground truth is not always available. Blasch [3] views it as a compound notion that aggregates several sub-dimensions. He enriches the previously mentioned approach by considering that reliability requires accurate, confident and timely results. However these three are not always achievable simultaneously, e.g. sometimes having more accurate or confident data leads to longer collecting time, which induces a choice between accuracy and timeliness.

Such an approach also leaves open the question of the aggregation operator to be used to combine the selected components of reliability. Blasch presents a user-driven approach, where these three dimensions are weighted based on a desired utility.

Contextual Reliability aims at changing reliability depending on the task the device is used for and thus the context of each piece of information.

Mercier et al. [12] propose to score reliability in a way where it better reflects the reality of a sensor and its working environment by enriching it with its context. Then, different situations can result in different output qualities for a given sensor. For instance in the case of target recognition [12], the performances of a data acquisition system may depend on weather conditions and on background and target properties, making the reliability of the decision system dependent on the target at hand. A sensor that recognises between three objects (helicopter, aeroplane and rocket) can have different accuracies for each one, effectively creating a vector of three reliabilities with different contexts.

Credibility can be defined as the level of confirmation of a given piece of information by other, independent, sources and constitutes another component of information quality. There are situations where assessing reliability is difficult or even impossible. This is where scoring credibility can provide an alternative or a complement to scoring reliability.

Using a "majority vote" strategy, it is possible to either improve the quality of the acquired piece of information [8] or combine multiple similar and dissimilar sensors to improve the overall quality of calculations by aggregating all outputs into one [16, 9, 8, 14].

Credibility for a piece of information is a relation to other pieces of information provided by independent sources, which ends with two cases, information is either concurring or conflicting [8]. The more pieces of information confirming the given piece of information, the more credible it is. This presents two possibilities of usage: i/ calculating ground-truth-type-of-reference by taking as output the majority of the sensors and then comparing it to evaluated sensor's output [9, 14] or ii/ combining all outputs to determine information quality by grouping sensors according to the feature they measure and evaluating the degree of consensus between them [8]. This approach can suffer limitations if sources are lacking or if their information is not comparable.

Table 1. Example of input data structure and output trust scores for a sensor.

Date	Time	Sensor ID	Message	Trust
11.03.2015	07:24:53	AC1	occupied	0.9
11.03.2015	07:25:40	AC1	**occupied**	**0.3**
11.03.2015	08:23:18	AC1	occupied	0.7
11.03.2015	08:24:08	AC1	clear	0.7
11.03.2015	09:15:23	AC1	occupied	0.8
11.03.2015	09:16:08	AC1	clear	0.8
11.03.2015	09:39:45	AC1	occupied	0.8
11.03.2015	09:40:29	AC1	clear	0.8
11.03.2015	10:22:14	AC1	occupied	0.8
11.03.2015	10:23:03	AC1	clear	0.9

3 Proposed Process for Information Scoring for Sensors

The information scoring model we propose is inspired by the multidimensional approach introduced in [15] which considers source evaluation, using reliability and competence, as well as content evaluation, using plausibility and credibility. We adapt, evaluate and aggregate three among these dimensions for the case of sensors, more precisely railway monitoring sensors. The three dimensions are also aggregated into a single *trust* value, which, in our case, is attributed to a sensor's reading at a given time. The presented approach has, in addition, a dynamic character: to score dimensions for the current log entry, the previous log entries are considered as well as their computed trust values.

This section gives a high level description of the considered data then details the process of dynamically scoring multi-dimensional trust.

3.1 Data Structure

The data structure we use for information scoring has the following characteristics: it is in the form of a log file whose entries contain a date, a time, a sensor id and a value, as shown in Table 1 for a real data example. The entries are event-triggered, i.e. they occur only when an event happens. The possible values represent the different messages given by the sensor that describe the sensor state. In the real data we consider, these messages can be occupied, clear or some type of disturbance. We aim to give a trust evaluation for each log entry, as illustrated in the last column of Table 1. We exhibit a part of data where a deficiency in quality is encountered which corresponds to a decreased trust value (see the second entry in Table 1).

For the computations, some notions need to be specified. We denote $\mathcal{L}$ the complete log set and $\mathcal{L}_s$ the set of log entries produced by sensor s. The notation l corresponds to one log entry defined as a vector containing three values: $l.fullDate$ corresponding to date and time, $l.sensor$ to the sensor id and $l.message$ contains the provided piece of information describing the sensors state. The set of all sensors is denoted $\mathcal{S}$, and the set of all times $\mathcal{T}$.

3.2 Scoring Trust

As explained at the beginning of this section our proposition is similar to [15], adapting and implementing this theoretical proposition to the specific case of sensors. We also consider reliability and credibility, presenting our view of scoring them in Sections 3.3 and 3.5 respectively. Regarding competence, its definition and scoring in the case of sensors appear to require knowledge about the system that is difficult to acquire e.g. external conditions or range of measurements. For instance, if competence is defined as the capacity of the sensor to provide the measurements it was designed for, this value is high when the sensor is working in its optimal conditions. The lack of that knowledge about the sensor and its surroundings makes it marginally useful in our trust calculation. Finally, we propose to replace plausibility with likelihood: whereas plausibility takes into account user background knowledge, likelihood depends only on the log file and takes into account the entry history.

The rest of this section presents our propositions for scoring each dimension, which takes into account meta-information and statistical analysis. *Reliability* is considered as a function of a sensor and time: $r(s,t)$, *likelihood* is related to the log entry: $lkh(l)$ and *credibility* applies to a log entry as well: $cr(l)$. They are finally aggregated into a trust value: $trust(l)$.

3.3 Reliability

Reliability, as a source metric, focuses on the specifics of the sensor, not the measures it provides. It is an *a priori* assessment of the source. This section first discusses the various approaches that can be proposed, organising them as constant vs. dynamic and meta-data-based vs. history-based; it then formalises the proposed definition.

Discussion. A basic approach could consist in making reliability a constant value, e.g. depending on the sensor type or brand: it might be known that specific sensors are of better quality than others and thus *a priori* provide more trustworthy information.

A way to enrich this basic definition is to take into account time and to define a dynamic reliability, for instance considering that this initial reliability value decreases when the sensor becomes older. This approach requires acquiring the knowledge about the obsolescence speed of the sensors, which might be difficult to know.

Note that it is possible to enrich further such a dynamic definition of reliability by taking into account maintenance operations, if their dates and types are known, although their interpretation can be debatable: they can be seen either as increasing reliability by slowing down the sensor ageing, or can be considered as the sign of the sensor needing repairs, casting doubts on the quality of the information it provides.

These approaches rely on the availability of very rich meta-information about the sensors, among which the type, brand, age, obsolescence speed and dates of

maintenance operations. Another source of information that can be exploited to define sensor reliability is offered by the history of its previous outputs, which is available in the log file. Indeed, reliability can be related to the question whether the device is working properly or not, which can be derived from its downtime or from its error messages. However sensor log files usually are event-triggered, which means that the downtime is not reported as such. The next paragraph describes in more detail a reliability definition based on error messages.

Proposed Sensor Reliability Definition. The measure we propose is based on the interpretation according to which the more errors a sensor reports, the less reliable it is: error messages indicate it encounters problems. We propose a dynamical measure that automatically adapts to the current state of the sensor, depending on what happened in its recent history; formally, it is defined as:

$$\begin{aligned} r : \mathcal{S} \times \mathcal{T} &\longrightarrow [0,1] \\ (s,t) &\longmapsto 1 - \frac{|error(recent(\mathcal{L}_s,t))|}{|recent(\mathcal{L}_s,t)|} \end{aligned} \tag{1}$$

where $recent : \mathcal{L} \times \mathcal{T} \to \mathcal{P}(\mathcal{L})$ provides the set of log entries produced by the sensor s in the considered time window t and $error : \mathcal{L} \to \mathcal{P}(\mathcal{L})$ is the function which extracts the set of error entries in this time window.

The definition of the considered time window, which determines the notion of "recent history" and the value of the reference *recent* can take several forms: it can be directly defined as an entry number, indicating the number of previous messages one may want to take into account; it can also be a temporal window, from which the log entries to be considered must be retrieved.

3.4 Likelihood

Likelihood measures how likely a piece of information is, independently of its source, but usually depending on available external information. Its expression varies according to the type of this considered external information.

For instance, in the case where the considered piece of information describes the position of a train on a track, it might be confronted to a train schedule, so as to check the compatibility with this external knowledge.

In the case considered in this paper, as described in Section 3.1, the pieces of information indicate the sensor states. We propose to measure their likelihood according to their compatibility with a model stating the allowed state evolution. Indeed, it can for instance be known that a sensor cannot remain in the 'occupied' state at two consecutive time stamps. A more general state evolution model for our considered sensors is illustrated in Figure 1: the two main states are *occupied* and *clear* and the several error states are distinguished. It can be seen that this sensor type cannot successively report *clear* and *disturbed* but an intermediary message *section disturbed* is used.

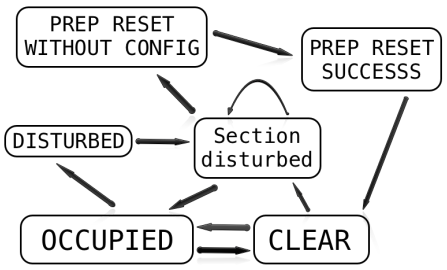

Fig. 1. Example of a state evolution model.

The proposed approach considers two cases: the message flow is compatible with the model or it is not. In the first case, the trust value of the previous log entry is considered. If it is strong, the likelihood will be high; if the log entry was untrustworthy, the likelihood will be lowered accordingly, indicating the fact that it could have been faulty. In the second case, the trust value of the previous log entry is also used to decrease likelihood: when that information was trustworthy, the likelihood will be low, otherwise the information is not considered enough to fully lower the likelihood score. The formal definition of $lkh : \mathcal{L} \to [0, 1]$ is:

$$lkh(l) = \begin{cases} trust(prv(l)), & \text{if } l.message \text{ compatible with } prv(l).message \\ 1 - trust(prv(l)), & \text{otherwise} \end{cases} \tag{2}$$

where $prv : \mathcal{L} \to \mathcal{L}$ returns the single log entry l' which is the entry provided by the same sensor just before the current entry l and $l.message$ is compatible with $l'.message$ when that state evolution is allowed by the model.

3.5 Credibility

Credibility aims to confirm or deny a piece of information, independently of its source, by comparing it with information from other sources. Its expression depends on the type of information provided by other sources.

Discussion. In the case considered in this paper where the piece of information describes the position of a train on tracks, it might be confirmed by its neighbouring source which should have reported the passing train shortly before. To implement this principle, the relative positions of the sensors are required, for instance in the form of sensor network. Figure 2 illustrates such a network: the nodes represent sensors and the lines between them indicate that two sensors are neighbours. For instance, when sensor S2 reports an activity, it means that the train had to pass through sensor S3 and it should have reported that fact with a log entry.

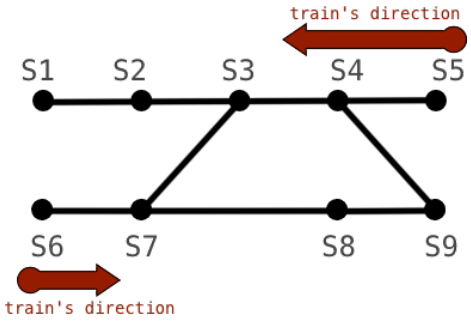

Fig. 2. Representation of the sensor locations on a portion of the railway structure

Formalization. The proposed approach considers scoring credibility of the sensor's state by looking through the recent log entries to find the ones which confirm the event and the ones which contradict it. The previously computed trust values for the considered entries are aggregated, ending with the final fusion of two values representing confirmation and contradiction. Formally, the credibility function cr is thus defined as:

$$\begin{aligned} cr : \mathcal{L} &\longrightarrow [0,1] \\ l &\longmapsto agg_1(agg_2(confirm(l)), agg_3(infirm(l))) \end{aligned} \tag{3}$$

where $confirm : \mathcal{L} \to \mathcal{P}(\mathcal{L})$ returns a set of entries that confirm l; $infirm(l) : \mathcal{L} \to \mathcal{P}(\mathcal{L})$ returns a set of entries that contradict l; agg_1, agg_2 and agg_3 are three aggregation operators applied to the trust scores of their set of logs.

Selection of aggregation. An aggregation operator in general is a function which reduces a set of numbers into a single, meaningful, number. The selection of an operator opens a wide discussion due to the diversity and variety of existing aggregation operators, each with its characteristics and properties, see e.g. [6].

The purpose of $agg_{2,3}$ is to combine the trust values of multiple entries. We propose to discard conjunctive and disjunctive operators, which can be considered as too extreme and to favour compromise operators that allow a compensation effect. As all log entries have the same impact, we consider the average. The agg_1 operator aims at combining the global confirmation trust (c) and the global infirmation trust (i). We require the following behaviour at the boundaries: if $c = 1$ and $i = 0$, the aggregated result must be 1; if $c = 0$ and $i = 1$, it must be 0.

They ensure that a fully confirmed piece of information has the highest credibility score and a fully contradicted information has the lowest credibility score. Therefore agg_1 needs to be asymmetrical. To meet these requirements we propose to define $agg_1 : [0,1] \times [0,1] \to [0,1]$ as:

$$agg_1(c,i) = \frac{1 + c - i}{2} \tag{4}$$

For the sake of simplicity, equation (3) omits its temporal dependence: confirmations and infirmations are looked for in recent entries. The notion of *recent* is equivalent to the one presented in Section 3.3. A too small window can result in "false negative", if the confirmation is earlier and outside the window However a too large window can result in "false positives": the confirmation does not exist but the previous train passage is included.

3.6 Trust

The final step then consists in aggregating the three dimensions: reliability, likelihood and credibility into the trust score.

We propose an approach to divide the overall trust scoring into two phases. First reliability and plausibility are aggregated. Indeed these dimensions both have an abating effect, leading to decrease trust, therefore, we propose to aggregate them using a t-norm, offering a conjunctive behaviour. The implementation described in Section 4 more specifically considers the probabilistic t-norm. Credibility can either increase or decrease trust due to its external factor as an opposite to the first two dimensions. We propose to consider a compromise operator, the weighted average. Trust is thus formally defined as

$$\begin{aligned} trust : \mathcal{L} &\longrightarrow [0,1] \\ l &\longmapsto \alpha \cdot r(l.sensor, l.fullDate) \cdot lkh(l) + (1-\alpha) \cdot cr(l) \end{aligned} \tag{5}$$

where the constant $\alpha \in [0,1]$ is set *a priori* to manipulate the influence of both sides.

4 Illustration on Real Data

This section describes the implementation of the proposed approach for real-world data from the railway domain. Among different sensors, the axle counter (AC) was chosen for its crucial role in maintaining safe and efficient train traffic. The aim is thus to verify the information it produces e.g. "the train is on this part of track", "the train left this part of track", "the sensor is not working properly". The dataset contains 60 axle counters to provide information on the train presence in the different part of tracks. The example of messages produced by AC is presented in Table 1 and all types of messages are included in the graph shown in Figure 1.

This section first describes the experimental protocol and then presents an illustrative example.

4.1 Experimental Protocol

The testing process is challenging due to the lack of a ground truth for the available dataset. We choose to illustrate our scoring by considering the original data as a reference and building a synthetic dataset from it with added random noise.

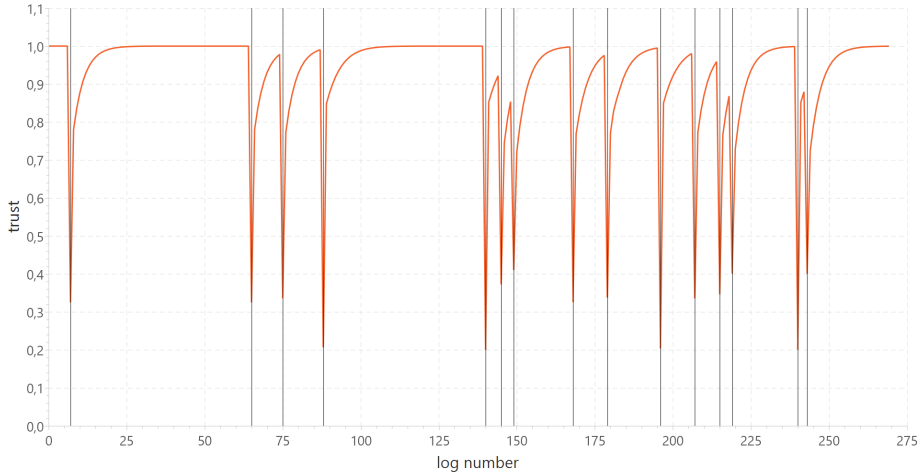

Fig. 3. Trust evolution for a single sensor affected by noise.

We change 5% of the AC states randomly, where the changes mean replacing the message of the log entry with a different one. We constrain the noise injection to preserve the initial distribution of the sensor states, i.e. if the *disturbed* message appears 1% in the dataset and the *clear* message appears around 49%, the same proportions hold in the noisy data.

The initial values for the 3 dimensions are set to 1.0; the window defining recent entries for reliability and credibility is set to consider entries from the previous 10 minutes; for Trust, we set $\alpha = 0.75$ in Equation (5).

4.2 Illustrative Example

Two approaches are considered when testing: applying noise to only one sensor or to multiple sensors.

Single Sensor Subject to Noise. In Figure 3, the trust values (y-axis) for the modified sensor are plotted over the log-entry number (x-axis). The noise is applied only to this device, its positions are highlighted by the vertical lines. It is noticeable that the corrupted entries are recognised, the trust being lowered for these log entries. Also, the trust level does not recover immediately after the decrease but takes time to do so, which is reflected by the introduction of previous logs trust into the computation. The part of the chart around entry number 220 presents one of the cases where the trust value was not able to fully recover, due to encountering another invalid entry which ended with another decrease. This example shows the ability of this tool to properly handle this scenario as well.

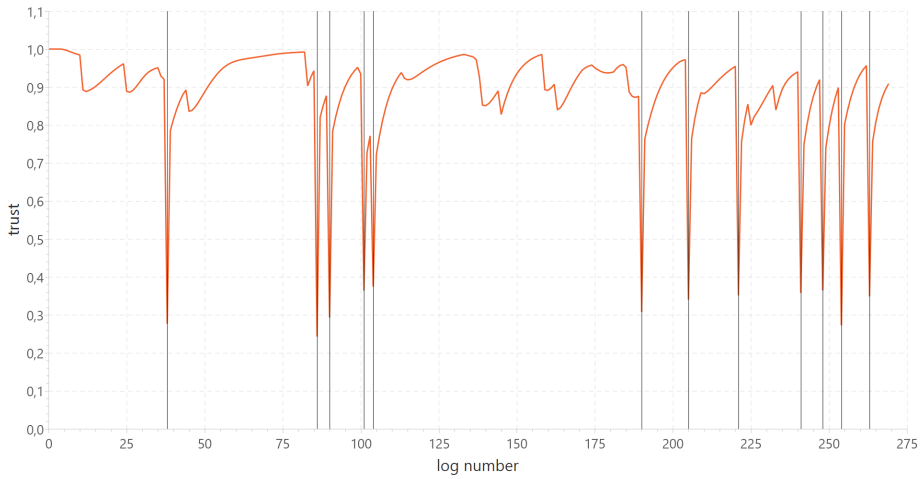

Fig. 4. Trust evolution for one sensor, when noise affects all sensors.

Multiple Sensors Subject to Noise. In this case, noise is applied to all sensors to observe how different sensors affect each other's trust. Figure 4 shows the evolution of trust values for a reference sensor, vertical line show its modified entries, the ones of the other sensors are omitted. The interesting part is the smaller decrease in trust for the entries that were not modified. The explanation for it lies in other sensors and their low trust scores. Due to the correlation between sensors, one of them can influence the other's trust. The level of the decrease depends on the trust value of the correlated sensor. Even though the trust value can decrease for the entries that were not modified, the level of that decrease is noticeably weaker compared to that of modified logs.

5 Conclusion

The variations in information produced by sensors bring out the need for an information quality scoring system taking into account both sensors and their output characteristics. Our approach proposes a modified version of dynamical trust scoring with three dimensions: reliability, likelihood and credibility. The temporal nature of the sensor's signal is considered in the aggregated trust score. We illustrated the proposed approach on a real-world railway dataset. Future works will include performing an experimental validation with statistical study generalising the illustrative example. Another perspective lies in proposing enriched scoring methods for presented dimensions.

Acknowledgements. This work was supported in part by Thales Polska.

References

1. Batini, C., Scannapieco, M.: Data and Information Quality. Springer International Publishing (2016)
2. Besombes, J., Revault d'Allonnes, A.: An extension of STANAG2022 for information scoring. In: Int. Conf. on Information Fusion, FUSION'08. pp. 1–7 (2008)
3. Blasch, E.P.: Derivation of a reliability metric for fused data decision making. In: IEEE National Aerospace and Electronics Conference. pp. 273–280 (2008)
4. Demolombe, R.: Reasoning about trust: A formal logical framework. Trust Management pp. 291–303 (2004)
5. Destercke, S., Buche, P., Charnomordic, B.: Evaluating data reliability: An evidential answer with application to a web-enabled data warehouse. IEEE Trans. Knowl. Data Eng. 25(1), 92–105 (2013)
6. Detyniecki, M.: Fundamentals on aggregation operators Technical report University of California Berkeley. Ph.D. thesis (2001)
7. Florea, M.C., Bossé, É.: Dempster-Shafer Theory: combination of information using contextual knowledge. In: Int. Conf. on Information Fusion, FUSION'09. pp. 522–528. IEEE (2009)
8. Florea, M.C., Jousselme, A.L., Bossé, É.: Dynamic estimation of evidence discounting rates based on information credibility. RAIRO-Operations Research 44(4), 285–306 (2010)
9. Guo, H., Shi, W., Deng, Y.: Evaluating sensor reliability in classification problems based on evidence theory. IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics) 36(5), 970–981 (2006)
10. Lesot, M.J., Delavallade, T., Pichon, F., Akdag, H., Bouchon-Meunier, B., Capet, P.: Proposition of a semi-automatic possibilistic information scoring process. In: Procs. of the 7th Conf. of the European Society for Fuzzy Logic and Technology (EUSFLAT-2011) and LFA-2011. pp. 949–956. Atlantis Press (2011)
11. Lesot, M.J., Revault d'Allonnes, A.: Information quality and uncertainty. In: Kreinovich, V. (ed.) Uncertainty Modeling: Dedicated to Professor Boris Kovalerchuk on his Anniversary, pp. 135–146. Springer International Publishing (2017)
12. Mercier, D., Quost, B., Denœux, T.: Refined modeling of sensor reliability in the belief function framework using contextual discounting. Information Fusion 9(2), 246–258 (2008)
13. Pichon, F., Dubois, D., Denoeux, T.: Relevance and truthfulness in information correction and fusion. Int. Jour. of Approximate Reasoning 53(2), 159 – 175 (2012)
14. Pon, R.K., Cárdenas, A.F.: Data quality inference. In: Procs. of the 2nd Int. workshop on Information quality in information systems. pp. 105–111. ACM (2005)
15. Revault d'Allonnes, A., Lesot, M.J.: Formalising information scoring in a multivalued logic framework. In: Information Processing and Management of Uncertainty in Knowledge-Based Systems: 15th International Conference, IPMU 2014, Proceedings, Part I. pp. 314–324. Springer International Publishing (2014)
16. Rogova, G., Hadzagic, M., St-Hilaire, M.O., Florea, M.C., Valin, P.: Context-based information quality for sequential decision making. In: 2013 IEEE Int. Multi-Disciplinary Conf. on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA). pp. 16–21 (2013)
17. Sidi, F., Panahy, P.H.S., Affendey, L.S., Jabar, M.A., Ibrahim, H., Mustapha, A.: Data quality: A survey of data quality dimensions. In: Procs. of Int. Conf. on Information Retrieval Knowledge Management. pp. 300–304 (2012)
18. Young, S., Palmer, J.: Pedigree and confidence: Issues in data credibility and reliability. In: Int. Conf. on Information Fusion, FUSION'07. pp. 1–8 (2007)