Ecological Informatics 23 (2014) 118–125

Contents lists available at ScienceDirect

Ecological Informatics

journal homepage: www.elsevier.com/locate/ecolinf

Multimedia information retrieval and environmental monitoring: Shared perspectives on data fusion

Alan F. Smeaton [a,*], Edel O'Connor [b], Fiona Regan [c]

[a] Insight Centre for Data Analytics and School of Computing, Dublin City University, Glasnevin, Dublin 9, Ireland
[b] CLARITY: Centre for Sensor Web Technologies, Dublin City University, Glasnevin, Dublin 9, Ireland
[c] MESTECH: Marine Environmental Sensing Technology Hub, Dublin City University, Glasnevin, Dublin 9, Ireland

A R T I C L E I N F O

Article history:
Received 22 January 2013
Received in revised form 16 July 2013
Accepted 21 October 2013
Available online 5 November 2013

Keywords:
Sensor data fusion
Multimedia information retrieval
Trust and reputation framework
Environmental monitoring

A B S T R A C T

Computer-based remote monitoring of our environment is increasingly based on combining data derived from in-situ-sensors with data derived from remote sources, such as satellite images or CCTV. In such deployments it is necessary to continuously monitor the accuracy of each of the sensor data streams so that we can account for sudden failures of sensors, or errors due to calibration drive or biofouling. In multimedia information retrieval (MMIR), we search through archives of multimedia artefacts like video programs, by implementing several independent retrieval systems or agents, and we combine the outputs of each retrieval agent in order to generate an overall ranking. In this paper we draw parallels between these seemingly very different applications and show how they share several similarities. In the case of environmental monitoring we also need some mechanism by which we can establish the trust and reputation of each contributing sensor, though this is something we do not need in MMIR. In this paper we present an outline of a trust and reputation framework we have developed and deployed for monitoring the performance of sensors in a heterogeneous sensor network.

1. Introduction

Our world is full of complex natural phenomena, some of which occur naturally such as our weather, the spread of diseases or various types of ecosystems, and some of which are man-made such as the characteristics of vehicular traffic flow in cities or the spread of online social networks. Whatever the form of complex phenomena, as humans we strive to understand them and in doing so we may be able to take advantage of them, to use them, or simply to further our own knowledge of the world we live in.

Part of understanding a complex phenomenon is being able to assess its state or its "wellness" so we can track it over time to see how it changes or evolves or to see how it reacts to various inputs we may make. All this requires being able to measure it, usually with a set of sensors of various kinds. However, sensors are usually designed so they measure one parameter or viewpoint only and in order to faithfully represent and then interpret a complex phenomenon so we can understand it, we need multiple viewpoints or perspectives whose actual values we combine together to give an overall, holistic overview. For example, the state or wellness of the human body is a function of sensor readings for body temperature, blood pressure, heart rate, respiration, and so on. The polluted or clean state of a river is assessed based on a combination of data values representing water flow and recent rainfall in the river's catchment area, water temperature, nutrient concentration, water turbidity and conductivity, etc. It is clear that only by combining the perspectives offered by different sensors can we get a true picture of complex phenomena that we are interested in, such as the human body or the pollution in a river.

When we use sensors, whether off-the-shelf or newly-developed, to measure some parameter from our natural world we need to be conscious of their shortcomings. Some of the sensors we use can suffer from calibration drift and need regular re-calibration. Some sensors in many natural environments suffer from biofouling which can reduce their accuracy and increase their margin of error. Other kinds of sensors have reliability issues or have a finite lifetime because they use energy or perhaps consume chemical reagents every time they take a reading, or they may just be prone to breaking. It is this scenario of using sensors which can have inherent errors and thus be noisy which we are primarily interested in because these sensors tend to be the newly-developed ones which, because they are new, offer new possibilities for understanding complex natural phenomena. In such scenarios, combining data values in order to get an overview of the underlying system is not always a straightforward process.

This area of combining sensor data readings is sometimes called "*sensor fusion*" and has already been the subject of much research in different fields because combining sensor data has applications in so many areas. At its most basic level, sensor fusion requires normalisation of the data gathered by different sensors because measurement units may be different, nomenclature may be different, timestamps may need to be aligned, or temporal interpolation across the data values from a sensor may be necessary in order to allow direct comparison with another

* Corresponding author.
E-mail address: alan.smeaton@dcu.ie (A.F. Smeaton).

1574-9541/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ecoinf.2013.10.004

sensor's stream of data. This task is sometimes called data cleaning and refers to a basic data processing task which is needed so more advanced processing such as identifying features, can take place as described by Roantree et al. (2009).

Depending on the sensor type, sets of data values from each sensor data stream can usually be grouped together to identify features or higher-level signals, sometimes called events. For example, the data representing an increase in heart rate combined with an increase in respiration rate and an increase in body temperature can be grouped to represent an event corresponding to an increase in metabolism. An increase in water temperature in river water, a decrease in dissolved oxygen and an increase in dissolved nutrients like nitrogen can be grouped together to represent an event corresponding to a decrease in water quality. Such event detection requires its own kind of data processing and is the subject of much investigation, such as that reported at the annual workshop on events held as part of ACM Multimedia Conference each year.

In this article we assume the cleaning and alignment of data from sensors, as well as the detection of events, are already completed and we focus on interpreting and assessing the impact of whatever underlying phenomenon has taken place. For example, the increase in body metabolism mentioned earlier and detected from heart rate, respiration rate and temperature, could have been caused by an infection, a sudden shock, or the person simply taking some exercise. The decrease in river water quality detected from changes in temperature, dissolved oxygen and nutrients could have been caused by an algal bloom. Our particular interest is how to detect and interpret changes such as these when some of the underlying sensors we use are not always trustworthy or reliable and how we can learn lessons and adopt techniques used from another field, to help us.

In Section 2 we provide a literature review of some of the previous work in the use of sensors in environmental monitoring and previous work in sensor data fusion as used in the domain of ecology, and environmental monitoring in particular. In Section 3 we introduce the area of multimedia information retrieval where the task is to locate multimedia artefacts such as video shots from an archive given a vague description of an information need from a user. We show how contemporary approaches to multimedia retrieval are based on combining the results of multiple search agents, with very strong similarities to sensor fusion. In Section 4 we highlight the parallels between sensor fusion and multiple media retrieval, pointing out that these are just two of the many domains in which data from different sensors/agents need to be combined in order to achieve some goal. Our approach to addressing how to manage integrating the data from sometimes inconsistent and unreliable sensors into an overall holistic view of a system is to build a trust and reputation framework which can account for the various reliabilities and accuracies of the different sensors and we present an overview of this contribution in Section 5, which is followed by a conclusion to the article.

2. Sensor fusion in environmental modelling

Environmental monitoring is a topic which for a long time was based on manual observations, laboriously gathered and collated. With the advent of sensor technology, digitisation, and easy communications via the internet, we are now seeing the application of wireless sensor networks (WSNs) to environmental monitoring as well as to many other applications. The principle behind WSNs is quite straightforward. In order to achieve widespread coverage, robustness, and tolerance to errors in sensor readings, the task of sensing an environment is broken down and shared among many independent nodes, each working autonomously but cooperating to collect and share data. Each node in a WSN is self-powered, has local communications to other nodes only, and is coupled to some sensing technology appropriate to whatever is being monitored. Because each individual node is functionally simple, the overall deployment is scalable and cost-effective, though there are challenges in terms of power management, robustness, calibration drift, communications protocols, and security, all of which can be overcome with clever software. Good overviews of WSNs can be found in Akyildiz et al. (2007) and Yick et al. (2008).

Information fusion is not a new topic and has existed in many forms in different application areas, as shown by the fact that there is an International Society of Information Fusion and an annual Information Fusion Conference. Sensor fusion in the context of environmental modelling, involves combining data from sensors, either heterogeneous or homogeneous, in order to provide a more complete, a more accurate overall picture of the underlying ecology which is being sensed. When the sensors whose data are being fused are homogeneous, such as a set of temperature, wind speed or humidity sensors within a local region, then the task is a fairly straightforward mathematical process and techniques based on the Central Limit Theorem (Giraitis and Surgailis, 1990) are among the most commonly used. This states that for a set of M random variables such as sensor readings taken over time, the mean tends to be normally distributed. For fusing the data from two homogeneous sensors, the fused result is a linear combination of the individual measurements but where each is then weighted by its noise variance. So a sensor with a history of erratic readings and thus high noise variance will not be weighted highly and will contribute less to the overall end result.

This fundamental approach to managing the differences in accuracy among sensors has been adapted in many different ways with several variants proposed. A distributed scheme to implement this rather than the more commonly used centralised approach, which requires keeping track of each sensor's performance over time, was developed by Xiao et al. (2005). Other techniques such as Dempster–Shafer theory and using Bayesian networks have also been used in “simple” sensor fusion applications.

Model-data fusion (MDF) is a technique which is particularly appropriate for integrating data in multi-sensor environments and has been applied in areas as diverse as research in ecology and palaeoecology (Peng et al., 2011), to estimating soil moisture content and land surface temperature (Renzullo et al., 2008). MDF provides more than just a quantitative addition of sensor data but a different approach based on maintaining multiple constraints as the observations (data values) from sensors are integrated into the models which are then updated. The paper by Peng et al. (2011) reviews key features of MDF and shows how it can be used for estimating parameter uncertainty and model error identification and it is this ability to identify inconsistencies between model and data, that is one of the strengths of MDF.

In work by Xie et al. (2008), the authors focus on fusing information taken from various sources of imagery in an application for mapping vegetation cover from remote sensed images. While images are more complex than simple point sensors (temperature, nutrient quantity, light level, etc.), different image sources from remote sensing offer a richness of data because the images can have different spectral, spatial, radioactive or temporal characteristics and because cameras are now so cost-effective to deploy, their use is becoming increasingly prevalent in environmental monitoring so this is something we should not ignore.

The paper by Xie et al. presents an overview of how to use remote sensing imagery to classify and map vegetation cover and when remotely sensed images are then combined with point sensor data readings, as we do later in this paper, we have the possibility to compute a really deep understanding of the area being monitored, in a really cost-effective way. However, interpreting such imagery and combining with point sensing data values requires quite extensive domain knowledge as well as knowledge of the area being monitored, and this is expensive to achieve. Several researchers have turned to machine learning techniques from computer science to help. These are software algorithms that take sets of data, such as sensor data or imagery, along with interpretation of that data, known as ground truth, and which then automatically learn the underlying patterns which lead to different interpretations. Software models for the application area are then constructed and these are then used to interpret new, even real time, data streams from sensors. For example, Sanchez-Hernandez et al. (2007)