

Taming big maritime data to support analytics

George A. Vouros¹, Christos Doulkeridis, Georgios Santipantakis, Akrivi Vlachou

Department of Digital Systems,
University of Piraeus
Greece

Abstract. This article aims to present important challenges and progress towards the management of data regarding the maritime domain for supporting analysis tasks. The article introduces our objectives for big data analysis tasks, motivating our efforts towards advanced data management solutions for mobility data in the maritime domain. The article introduces the data sources to support specific maritime situation awareness scenarios that are addressed in the datAcron² project, presents the overall infrastructure designed for managing and exploiting data for analysis tasks, and presents a representation for integrating data from different sources, revolving around the notion of semantic trajectories: the datAcron ontology.

1 Introduction: Overall objectives to manage mobility data

The datAcron project aims to advance the management and integrated exploitation of voluminous and heterogeneous data-at-rest (archival data) and data-in-motion (streaming data) sources, so as to significantly improve the capacities of surveillance systems to promote safety and effectiveness of critical operations for large numbers of moving entities in large geographical areas.

Challenges throughout the Big Data ecosystem of surveillance systems concern effective detection and forecasting of moving entities' trajectories, as well as recognition and prediction of events due to entities' trajectories and contextual data. Challenges emerge as the number of moving entities and related operations increase at unprecedented scale. This, in conjunction with the demand for more and more frequent data from many different sources and for each of these entities, result in generating vast data volumes of heterogeneous nature, at extremely high rates, whose intertwined exploitation calls for novel big data techniques and algorithms that lead to advanced data analytics: this is a core research issue that we address in the datAcron project.

More concretely, core research challenges in datAcron include:

¹ Corresponding author: georgev@unipi.gr

² datAcron project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 687591 (<http://datacron-project.eu>)

- distributed management and querying of spatiotemporal RDF data-at-rest (archival) and data-in-motion (streaming), following an integrated approach;
- reconstruction and forecasting of moving entities' trajectories in the challenging Maritime (2D space with temporal dimension) and Aviation (3D space with temporal dimension) domains;
- recognition and forecasting of complex events due to the movement of entities (e.g. prediction of potential collision, capacity demand, hot spots / paths);
- interactive visual analytics for supporting human exploration and interpretation of the above.

Technological developments are validated and evaluated in user-defined challenges that aim at increasing the safety, efficiency and economy of operations concerning moving entities in the aviation and maritime domains. The main benefit arising from improved trajectory prediction in the aviation use-case lies in the accurate prediction of complex events, or hotspots, leading to benefits to the overall efficiency of an air-traffic management (ATM) system. On the other hand, discovering and characterizing the activities of vessels at sea are key tasks to Maritime Situational Awareness (MSA) indicators and constitute the basis for detecting/predicting vessel activities, towards enhancing safety, detecting anomalous behaviors, enabling an effective and quick response to maritime threats and risks.

In both domains, semantic trajectories are turned into "first-class citizens": In practice, this forms a paradigm shift towards operations that are built and revolve around the notion of trajectory. For instance, in the MSA world, trajectories are essential for tracking vessels' routes, detecting and analyzing anomalous behaviour and supporting critical decision-making. datAcron considers trajectories as «first class» citizens and aims to build solutions towards managing data that are connected via, and contribute to, enriched views of trajectories: Doing so, datAcron revisits the notion of semantic trajectory and builds on it. Specifically, it is expected that meaningful moving patterns will be computed and be exploited to recognizing and predicting the behaviour and states of moving objects, taking advantage of the wealth of information available in disparate and heterogeneous data sources, integrated in a representation in which trajectories are the main entities.

The objective of this paper is to review challenges and recent progress towards managing big data for supporting analysis tasks regarding moving objects at sea (e.g. for predicting vessels' trajectories, events and/or support visual analytics tasks): Such data may be surveillance data, but also data regarding vessels' characteristics, past events, areas of interest, patterns of movement, etc. These are data from disparate and heterogeneous sources that should be integrated to support maritime' experts awareness of situations, together with the automatic computations of indicators that contribute to such awareness.

The article presents the datAcron maritime use case. It then presents the overall datAcron infrastructure to manage big mobility data, focusing on data management issues. It then presents the datAcron ontology for the representation of maritime data

towards providing integrated views of these data, focusing on the notion of semantic trajectory.

2 Taming Big Data in the Maritime Use Case: Motivation and Challenges

The maritime environment has a huge impact on the global economy and our everyday lives. Specifically, surveillance systems of moving entities at sea have been attracting increasing attention due to their importance for the safety and efficiency of maritime operations. For instance, preventing ship accidents by monitoring vessel activity represents substantial savings in financial cost for shipping companies (e.g., oil spill cleanup) and averts irrevocable damages to maritime ecosystems (e.g., fishery closure). The past few years have seen a rapid increase in research and development of information-oriented infrastructures and systems addressing many aspects of data management and data analytics related to movement at sea (e.g., maritime navigation, marine life). In fact, the correlated exploitation of heterogeneous and large data sources offering voluminous historical and streaming data is considered as an emergent necessity given the (a) wealth of existing data, (b) the opportunity to exploit such data towards building models of entities' movement patterns and understanding the occurrence of important maritime events.

It is indeed true that reaching appropriate Maritime Situation Awareness (MSA) for the decision maker requires processing in real-time a high volume of information of different nature originating from a variety of sources (sensors and humans), with a lack of veracity (uncertain, imprecise, vague, ambiguous, incomplete, conflicting, incorrect, etc), and coming with high velocity. Different types of data are available, which only if properly combined and integrated can provide useful knowledge. However, the correlated exploitation of data from disparate and heterogeneous data sources is a crucial computational issue.

The growing number of sensors (in coastal and satellite networks) makes the sea one of the most challenging environments to be effectively monitored; the need for methods for processing of vessel motion data, which are scalable in time and space, is highly critical for maritime security and safety. For instance, about 12,000 ships/day are tracked in EU waters and about 100,000,000 AIS positions are recorded per month in EU waters (EMSA, 2012). Beyond the volume of data concerning ships' positions obtained from AIS, they might not be always sufficient for the purposes of detection and prediction algorithms. Only if properly combined and integrated with other data acquired from other data/information sources (not only AIS), they can provide useful information and knowledge for achieving the maritime situational awareness in support to the datAcron maritime use case.

The Maritime Use Case for datAcron focuses on the control of fishing activities, as it fulfills many of the requirements for validating the technology to be developed in datAcron: It addresses challenging problems deemed of interest for the maritime operational community in general; it is aligned with the European Union maritime policy and needs in particular; and relies on available datasets (unclassified, shareable) among the teams and others of interest in the research community (e.g. AIS

data, radar datasets, databases of past events, intelligence reports, etc). Moreover, it is of considerable complexity as it encompasses several maritime risks and environmental issues such as environmental destruction and degradation but also maritime accidents, Illegal, Unreported, and Unregulated (IUU) fishing and trafficking problems.

The support for processing, analysis and visualization of fishing vessels at the European scale - although not worldwide, the capability of predicting the movement of maritime objects and the identification of patterns of movement and navigational events, shall improve existing solutions to monitor the compliance to the European common fisheries policy.

Besides the control of fishing activities, another core issue is safety. Fishing, in peace situation, is known as one of most dangerous activity and is regularly ranked in the top five dangerous activities, depending on the years. Safety does not concern only fishing vessels themselves but also the surrounding traffic and more generally all other human activities at sea.

The data to be used in datAcron comprise real and quasi-real data streams as well as archival (or historical) European datasets supporting the fishing scenarios specified. These usually need to be cleaned up from inconsistencies, converted into standard formats, harmonized and summarized.

The following list briefly summarizes typical datasets that are relevant to the datAcron scenarios:

- **Automatic Identification System (AIS)** messages broadcasted by ships for collision avoidance;
- **Marine protected/closed areas**, where fishing and sea traffic may be (temporarily) forbidden;
- **Traffic separation schemes and Nautical charts** useful to define vessel routes;
- **Vessel routes and Fishing areas** estimated from historical traffic data;
- **Registry data** on vessels and ports;
- **Records of past events**, such as incidents and illegal activities reports;
- **Meteorological and oceanographic (METOC) data** on atmospheric and sea state conditions and currents.

Despite the urgent need for the developments of maritime data infrastructures, current information and database systems are not completely appropriate to manage this wealth of data, supporting also the analytics tasks targeted in datAcron. The reasons behind are at least as follows. First, the very large data volumes generated require the development of pre-filtering data integration process that should deliver data synopses in real-time, while maintaining the main spatio-temporal and semantic properties. Next, additional ocean and atmospheric data at global and local scales are often necessary to evaluate events and patterns at sea in the most appropriate way, leading to additional data integration issues [17].

In addition to the above, data measurements have an intrinsic uncertainty, which may be addressed by proper data fusion algorithms and clustering in the preparation/preprocessing phase (by assessing the quality of data themselves) and by combining measurements from complementary sources [17].

3 Big Data Management Challenges in datAcron

As already said, we in datAcron aim at recognizing and forecasting complex events and trajectories from a wealth of input data, both data-at-rest and data-in-motion, by applying appropriate techniques for Big Data analysis. The technical challenges associated with Big Data analysis are manifold, and perhaps better illustrated in [1,2], where the Big Data Analysis Pipeline is presented. As depicted in Figure 1, five major phases (or steps) are identified in the processing pipeline:

- Data Acquisition and Recording
- Information Extraction and Cleaning
- Data Integration, Aggregation, and Representation
- Query Processing, Data Modeling, and Analysis
- Interpretation




Figure 1. Major steps in analysis of Big Data (from [1,2])

Data Acquisition

As already said, large volumes of high-velocity data are created in a streaming fashion, including surveillance data and weather forecasts, that need to be consumed in datAcron. One major challenge is to perform online filtering of this data, in order to keep only the necessary data that contain the useful information. To this end, we apply data summarization techniques on surveillance data, thus keeping only the “critical points” of a moving object’s trajectory, which signify changes in the mobility of the moving object. Such a summarized trajectory is shown in Figure 2, comprising the low-level events detected as critical trajectory points. A research challenge for datAcron is to achieve a data reduction rate above 90%, without compromising the quality of the compressed trajectories and of course the quality of trajectories’ and events’ analysis tasks [14].

Another challenge in the data acquisition phase is to push computation to the edges of the Big Data management system: To achieve this we perform online data summarization of surveillance data on the input stream directly, as soon as it enters the system. Moreover, we employ in-situ processing techniques, near to the streaming

data sources, in order to identify additional low-level events of interest, such as the entrance/leave of moving objects in specific areas of interest (such as protected marine areas), and events requiring cross-streaming processing.

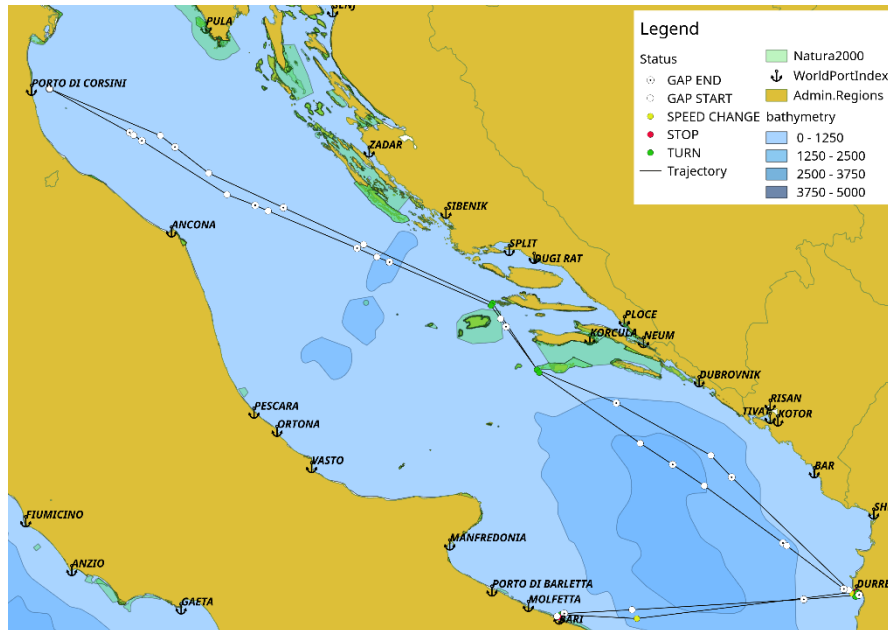


Figure 2. Summarized trajectory: Critical points with specific meaning indicate important low-level events at specific points.

Information Extraction and Cleaning

Given the disparity of data sources exploited in datAcron, with miscellaneous data in various formats for processing and analysis, a basic prerequisite for the subsequent analysis tasks is to extract the useful data and transform it in a form that is suitable for processing. As a concrete example, weather forecasts are provided as large binary files (GRIB format), which cannot be effectively analyzed. Therefore, we extract the useful meteorological variables from these files together with their spatio-temporal information, so that they can be later associated with mobility data. These should be done in operational time (i.e. in milliseconds), enriching the stream(s) of surveillance data.

In addition, surveillance data are typically noisy, contain errors, and are associated with uncertainty. Data cleaning techniques are applied in the streams of surveillance data, in order to reconstruct trajectories with minimum errors that will lead to more accurate analysis results with higher probability. Indicative examples of challenges addressed in this respect include handling delayed surveillance data, dealing with intentional erroneous data (spoofing) or hardware/equipment errors, etc.

Data Integration, Aggregation and Representation

Having addressed data cleaning, the next challenge is how to integrate the heterogeneous data coming from various data sources, in order to provide a unified and combined view. Our approach is to transform and represent all input data in RDF, following a common representation (i.e. the datAcron ontology) shown in the previous section, which has been designed purposefully to accommodate the different data sources. However, data transformation does not suffice by itself. To achieve data integration, we apply online link discovery techniques in order to interlink streaming data from different sources, a task of major significance in datAcron.

In particular, the types of discovered links belong to different categories with most representative ones: (a) moving object with static spatial area, (b) moving object to spatio-temporal variables, (c) moving object to moving object. In the first case, we monitor different relations (enter, exit, nearby) between a moving object and areas of interest, such as protected Natura areas or fishing zones. In the second case, we enrich the points of a trajectory with weather information coming from weather forecasts. Finally, in the last case, we identify relations between moving objects, e.g., two vessels approaching each other, or staying in the same place for unusually long period. By means of link discovery, we derive enriched data representations across different data sources, thereby providing richer information to the higher-level analysis tasks in datAcron.

Query Processing, Data Modeling and Analysis

Another Big Data challenge addressed in datAcron relates to scalable processing of vast-sized RDF graphs that encompass spatio-temporal information. Towards this goal, we design and develop a parallel spatio-temporal RDF processing engine on top of Apache Spark. Individual challenges that need to be solved in this context include RDF graph partitioning, implementing parallel query operators that shall be used by the processing engine, and exploiting the capabilities of Spark in the context of trajectory data.

Complex event detection is also performed in datAcron, where the objective is to detect events related to the movement of objects in real-time.

Last, but not least, particular attention is set towards predictive analytics, namely trajectory prediction and event forecasting. Both short-term and long-term predictions are useful depending on the domain, and in particular for maritime: A hard problem is to perform long-term prediction. For instance, as far as trajectory prediction is concerned, we may distinguish location prediction (where a moving object will be after X hours) and trajectory prediction (what path will a moving object follow in order to reach position P).

Interpretation

To assist the task of human-based interpretation of analysis results, as well as the detection of patterns that may further guide the detection of interesting events – tasks that are fundamental for any Big Data analysis platform- datAcron relies on visual analytics. By means of visual analysis tools, it is possible to perform visual and interactive exploration of moving objects and their trajectories, visualize aggregates or data summaries, and ultimately identify trends or validate analysis results that would be hard to find automatically.

4 Semantic trajectories revisited: An ontology for maritime data to support movement analysis.

Given the significance of trajectories, analysis methods (e.g. for the detection and prediction of trajectories and events), in combination to visual analytics methods, require trajectories to be (a) available at multiple levels of spatio-temporal analysis, (b) easily transformed to spatio-temporal constructs/forms suitable for analysis tasks, and (c) providing anchors for linking contextual information and events related to the movement of any object: Doing so, representation of trajectories at the semantic level aim to provide semantically meaningful integrated views of data regarding the mobility of vessels, at different levels of analysis.

The term contextual information denotes any type of information about entities that affect the behaviour of an object (e.g. weather conditions, or events of special interest), and information about entities that are being affected by the behaviour of an object (e.g. a fishing or a protected area). Moreover, the context of an objects' trajectory may include the trajectories of other objects in its vicinity: As already said, surrounding traffic may entail safety concerns. The association of trajectories to contextual information and events results to enhanced semantic trajectories of moving objects.

Existing approaches for the representation of semantic trajectories suffer from at least one of the following limitations: (a) use of plain textual annotations instead of semantic links to other entities [3,4,5,8,10,11,12], (b) only limited types of events can be represented as resources [3,4,5,6,7], (c) make assumptions of the structure of trajectories, restricting the levels of analysis and representations supported [6],[9], and (d) semantic links between entities are mostly application-specific, rather than generic ones [6],[7].

Motivated by real-life emerging needs in MSA we aim at providing a coherent and generic scheme supporting the representation of semantic trajectories at different levels of spatial and temporal analysis: Trajectories may be seen as single geometries, or arbitrary sequences of moving objects' positions over time, or as sequences of events, or sequences of trajectories segments, each linked with important semantic information: These different levels of representation must be supported by the datAcron ontology.

The datAcron ontology is expressed in RDFS. The main concepts and properties in this ontology are depicted in Figure 3 and presented in the next paragraphs.

Places: The concept `:Place`, is instantiated by the static spatial resources representing places and regions of special interest. Places are related to any type of trajectory (or segment), weather conditions and events (by relations "within" or "nearby"). `:Place` is a generalization of Places of Interest (POIs) related to trajectories and Regions of Interest (ROIs) [13] associated with a stop event of a moving object.

A place is always related to a geometry with the property `:hasGeometry`.

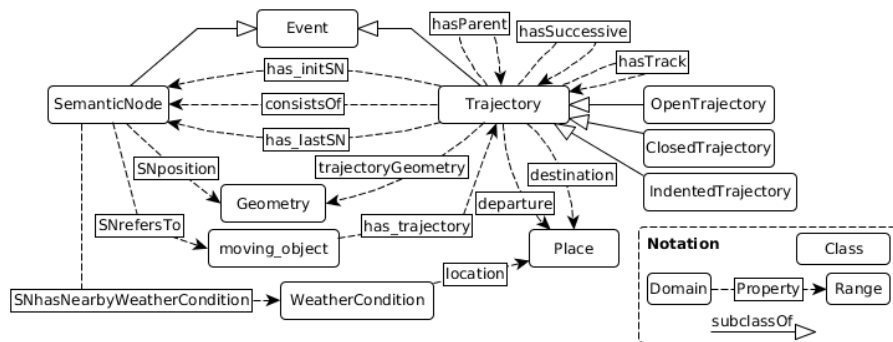


Figure 3. Core concepts and properties

Semantic Nodes: For the representation of moving objects behaviour at varying levels of analysis in space and time, and in order to associate trajectories with contextual information, we use the concept of `:SemanticNode`. A semantic node specifies the position of a moving object in a time period or instant, or a specific set of spatio-temporal positions of a single moving object. In the latter case, it specifies an abstraction/aggregation of movement track positions (e.g. the centroid for a set of positions) and can be associated with a place and a temporal interval when this movement occurred. In both cases, a semantic node may represent the occurrence of an event type.

More importantly, any instance of `:SemanticNode` can be associated with contextual information known for the specific spatio-temporal position or ROI.

Also, the semantic node is associated with weather information regarding node's spatio-temporal extent.

Trajectories: A trajectory is a temporal sequence of semantic nodes or trajectory segments. The main properties relating trajectories to semantic nodes are `:hasInitNode`, `:hasLastNode` representing trajectory initial and last semantic nodes respectively, and `:consistsOf` for relating trajectories to intermediate semantic nodes. The property `:hasNext` relates consecutive semantic nodes in a trajectory to maintain the temporal sequence of nodes.

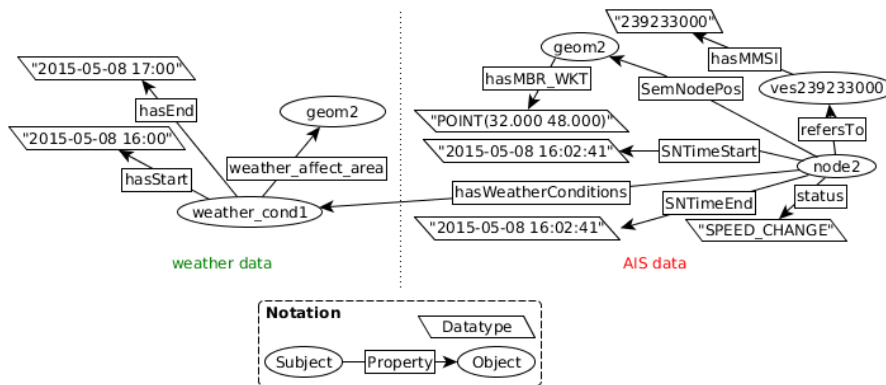


Figure 4. A semantic node instance linked to weather data

Various types of trajectories are supported, such as the `:OpenTrajectory`, where the last semantic (terminal) node is not yet reached, and the `:ClosedTrajectory`, in which the last node is specified. Having said this, it should be noticed that criteria for determining terminal positions are application and domain specific. A trajectory can also be classified as `:intendedTrajectory` specifying a planned or predicted trajectory. Thus, each moving object, at a specific time, may be related to multiple trajectories, actual or intended/predicted ones, and semantic nodes can be reused in different types of trajectories, w.r.t. spatial and temporal granularity.

The properties `:hasParent`, `:hasSuccessive` relate a trajectory with other trajectories, forming a structured trajectory. Specifically, the first property relates a trajectory to its parent (the whole) and the second one relates the successive trajectories (the parts).

For instance, Figure 5 illustrates the trajectory of a vessel through ports “Porto Di Corsini”, “Durrës” and “Bari”, and Figure 6 demonstrates the corresponding structured trajectory, with its trajectory segments.

Trajectory segments have a starting and an ending semantic node, and are associated to a time interval and geometry.

Event: The `:Event` concept is instantiated by spatio-temporal entities representing specific, aggregated or abstracted positions instantiating a specific event pattern. The instantiation of any such event pattern can be part of a preprocessing task on raw data, or can be done by a function applied to the RDF data, resulting to the generation of new triples representing the recognized events. Thus, an event is associated with a set of semantic nodes, which may be in temporal sequence. Each event may be associated with one or more moving objects, and it has spatial, temporal and domain-specific properties, according to the properties of semantic nodes and trajectories. It must be pointed out that a semantic node or a trajectory can be specified to belong to more than one event types (e.g. `:Rendezvous` and `:PackagePicking`).




Figure 5. An example trajectory: The figure shows the activity of a vessel near Durres port. The multiple stop positions with intermediate gap-start, gap-end points could indicate a suspicious behavior.

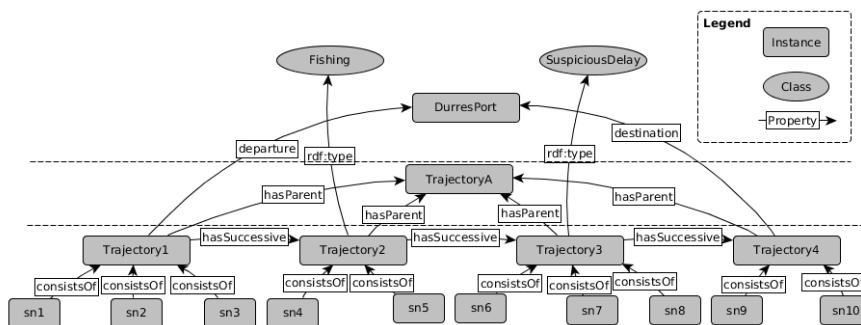


Figure 6. A structured trajectory

Events are distinguished to low-level and high-level events: The former are those detected from raw trajectory data or from time-variant properties of a single moving object, disregarding contextual data. For instance, a `:Turning` or an `:Accelerating` event is a low-level event, since it concerns a specific object and can be detected directly from raw trajectory data. Figures 2 and 5 depicts such low-level events as trajectory “critical” points. High-level (or complex) events are detected or predicted by means of features specifying movement and/or time-variant

properties of moving objects, in addition to contextual ones. For example, the detection of a `:Fishing` event needs the consideration of the type of the vessel and the known fishing regions, in addition to the vessel's raw trajectory.

5 Concluding Remarks

The datAcron project aims to advance the management and integrated exploitation of voluminous and heterogeneous data-at-rest (archival data) and data-in-motion (streaming data) sources, so as to address important challenges towards the management of data regarding the maritime domain for supporting analysis tasks. It is indeed true that vast data volumes of heterogeneous nature, at extremely high rates, whose necessity for intertwined exploitation for supporting analysis tasks in the maritime domain, calls for novel big data techniques and algorithms that lead to advanced data analytics.

Towards achieving its objectives, datAcron considers semantic trajectories to be “first-class citizens”, following the paradigm shift towards operations that are built and revolve around the notion of trajectory. Thus, datAcron revisits the notion of semantic trajectory and builds on it. Specifically, it is expected that meaningful moving patterns will be computed and be exploited to recognizing and predicting the behaviour and states of moving objects, taking advantage of the wealth of information available in disparate and heterogeneous data sources.

Given the significance of trajectories, analysis methods (e.g. for the detection and prediction of trajectories and events), in combination to visual analytics methods, require trajectories to be (a) available at multiple levels of spatio-temporal analysis, (b) easily transformed to spatio-temporal constructs/forms suitable for analysis tasks, and (c) providing anchors for linking contextual information and events related to the movement of any object: Towards these objectives datAcron has devised a representation for trajectories at the semantic level, providing semantically meaningful integrated views of data regarding the mobility of vessels, at different levels of analysis.

Finally, as already mentioned we address issues of all five major phases (or steps) identified in the processing pipeline of a big data architecture:

- Data Acquisition and Recording
- Information Extraction and Cleaning
- Data Integration, Aggregation, and Representation
- Query Processing, Data Modeling, and Analysis
- Interpretation

The design of the datAcron overall architecture reflects these issues, also in close connection to the requirements from the maritime domain.

Our current work focuses on Data Integration, Aggregation and Representation, as well as on Query Processing, Data Modeling, and Analysis. Methods aim towards providing big data solutions to these processing phases, even at operational times.

Acknowledgement: This work is supported by the `datAcron` project, which has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 687591 (<http://datacron-project.eu>)

References

1. Agrawal D. et al. (2012). Challenges and Opportunities with Big Data: A white paper prepared for the Computing Community Consortium committee of the Computing Research Association. <http://cra.org/ccc/resources/ccc-led-whitepapers/>
2. H. V. Jagadish, Johannes Gehrke, Alexandros Labrinidis, Yannis Papakonstantinou, Jignesh M. Patel, Raghu Ramakrishnan, Cyrus Shahabi: Big data and its technical challenges. *Commun. ACM* 57(7): 86-94 (2014)
3. Alvares, L.O., Bogorny, V., Kuijpers, B., de Macêdo, J.A.F., Moelans, B., Vaisman, A.A.: A model for enriching trajectories with semantic geographical information. In: *Proc. of GIS*. p. 22 (2007)
4. Baglioni, M., de Macêdo, J.A.F., Renso, C., Trasarti, R., Wachowicz, M.: Towards semantic interpretation of movement behavior. In: *Advances in GIScience*, pp. 271-288. Springer (2009)
5. Bogorny, V., Kuijpers, B., Alvares, L.O.: ST-DMQL: A semantic trajectory data mining query language. *International Journal of Geographical Information Science* 23(10), 1245-1276 (2009)
6. Bogorny, V., Renso, C., de Aquino, A.R., de Lucca Siqueira, F., Alvares, L.O.: Constant - A conceptual data model for semantic trajectories of moving objects. *Trans. GIS* 18(1), 66-88 (2014)
7. Fileto, R., May, C., Renso, C., Pelekis, N., Klein, D., Theodoridis, Y.: The Baquara2 knowledge-based framework for semantic enrichment and analysis of movement data. *Data Knowl. Eng.* 98, 104-122 (2015)
8. van Hage, W.R., Malaisé, V., de Vries, G., Schreiber, G., van Someren, M.: Combining ship trajectories and semantics with the simple event model (SEM). In: *Proc. of EIMM*. pp. 73-80 (2009)
9. Nogueira, T.P., Martin, H.: Querying semantic trajectory episodes. In: *Proc. of MobiGIS*. pp. 23-30 (2015)
10. Spaccapietra, S., Parent, C., Damiani, M.L., de Macêdo, J.A.F., Porto, F., Vangenot, C.: A conceptual view on trajectories. *Data Knowl. Eng.* 65(1), 126-146 (2008)
11. Yan, Z., Chakraborty, D., Parent, C., Spaccapietra, S., Aberer, K.: Semantic trajectories: Mobility data computation and annotation. *ACM TIST* 4(3), 49 (2013)
12. Yan, Z., Macedo, J., Parent, C., Spaccapietra, S.: Trajectory ontologies and queries. *Transactions in GIS* 12(s1), 75-91 (2008)

13. Baglioni, M., de Macêdo, J.A.F., Renso, C., Trasarti, R., Wachowicz, M.: Towards semantic interpretation of movement behavior. In: *Advances in GIScience*, pp. 271–288. Springer (2009)
14. Patroumpas K., Alevizos E., Artikis A., Vodas M., Pelekis N., Theodoridis Y.: Online event recognition from moving vessel trajectories. *GeoInformatica* 21(2): 389-427 (2017)
15. Ngomo, Axel-Cyrille Ngonga. "Orchid–reduction-ratio-optimal computation of geo-spatial distances for link discovery." In *International Semantic Web Conference*, pp. 395-410. Springer Berlin Heidelberg, 2013.
16. Georgala, Kleanthi, Mohamed Ahmed Sherif, and Axel-Cyrille Ngonga Ngomo. "An Efficient Approach for the Generation of Allen Relations." In *European Conference on Artificial Intelligence (ECAI)*. 2016.
17. Claramunt C., et al: Maritime data integration and analysis: recent progress and research challenges, *EDBT 2017* (to appear)