Consolidation of massive medical emergency events with heterogeneous situational context data sources

Thomas James Tiam-Lee, Rui Henriques, Jose Costa, Vasco Manquinho and Helena Galhardas

INESC-ID and Instituto Superior Técnico, Universidade de Lisboa, Rua Alves Redol 9, Lisbon, Portugal

Abstract

The prevalence, spatiotemporal distribution, and category incidence of medical emergencies are rapidly changing worldwide. The current pandemic context and emerging trends in public health create the need for self-adapting Emergency Medical Services (EMS). Emergency occurrences and responses are intricately dependent on contextual factors, including weather, epidemic context, urban traffic, large-scale events, and demographics. In this context, monitoring emergency occurrences, medical responses, and their situational context is essential to optimize EMS efficiency and efficacy. In this work, we implement best practices in multidimensional database modelling to consolidate emergency event data with public sources of situational context for context-aware data analysis. The resulting design is able to address challenges pertaining to the massive, incomplete, and spatiotemporal nature of emergency event data and the heterogeneity of context sources and their varying spatiotemporal footprints. We present a study case on real-world medical emergency data from Portugal. The results show the efficient retrieval of data structures conducive to spatiotemporal data mining tasks.

Keywords

medical emergency service, situational context, spatiotemporal data, heterogeneous data consolidation, multidimensional data model

1. Introduction

Many societies around the world have nationwide services to respond to emergencies, including individual health emergencies, large-scale disasters, traffic accidents, amongst others. Emergency Medical Services (EMS) generally provide initial triage, the necessary care on-scene, and efficient transportation to health facilities for the subsequent full care delivery [1].

Despite the pivotal worldwide role of EMS, they are increasingly pressured to meet higher service levels due to the rising number of emergencies in major urban centers. They also need to adapt to the ongoing changes in public health caused by shifts in demographics and disease prevalence [2]. This has led to various studies to improve EMS by looking at different situational contexts [3, 4, 5, 6]. Furthermore, EMS face additional challenges in the current pandemic context where ambulance waiting times near hospitals can be considerably high, and emergency requests occur at later complication stages [7]. Finally, a significant number of situational factors that can predispose emergency prevalence are commonly neglected.

In this context, the consolidation of emergency event data with publicly available sources of relevant situational context – including weather, epidemics, large-scale events, urban traffic, demographics or zoning information – is essential to better assess causality factors underlying the spatiotemporal prevalence of emergencies and response inefficiencies. Still, several challenges found at the data consolidation levels limit the applicability of context-aware emergency data analysis. First, there are challenges on the heterogeneity and real-time monitoring of situational context data and the varying geographical-and-temporal footprint of meaningful context occurrences, such as festivities, sports matches, traffic jams, or abnormal weather conditions. Second, the inherent complexity and massive size of spatiotemporal emergency event data, characterized by a multiplicity of stages (call, triage, activation, vehicle dispatch, and so on), arbitrarily-high degree of missing values, the abundance of relevant categories (diagnostics, assistance provided, and type of dispatched vehicles). Third, the need to efficiently handle data analytics with spatial, temporal and emergency-specific drill-down, roll-up, slicing and dicing criteria.

In this work, we apply best practices in multidimensional data modelling to address the above challenges, using real-world EMS data in Portugal[1] as a study case. In particular, we consolidate emergency event data with available sources of situational context for effective and efficient subsequent spatiotemporal data analysis. The result is an integrated data warehouse that supports the

[1]This work is anchored in the pioneer research and innovation project Data2Help, an initiative that aims at developing a set of computational tools to optimize the operations of EMS.

Published in the Workshop Proceedings of the EDBT/ICDT 2022 Joint Conference (March 29-April 1, 2022), Edinburgh, UK

thomas.tiam-lee@tecnico.ulisboa.pt (T. J. Tiam-Lee); rmch@tecnico.ulisboa.pt (R. Henriques); jose.a.costa@tecnico.ulisboa.pt (J. Costa); vasco.manquinho@tecnico.ulisboa.pt (V. Manquinho); helena.galhardas@tecnico.ulisboa.pt (H. Galhardas)
0000-0003-1820-3984 (T. J. Tiam-Lee); 0000-0002-3993-0171 (R. Henriques); 0000-0002-4205-2189 (V. Manquinho); 0000-0002-9330-3910 (H. Galhardas)

efficiently retrieval of relevant multi-source information to assist in various computational tasks. Gathered results from using the data warehouse suggest its relevance in boosting retrieval efficiency, promoting querying efficacy, and deriving data structures conducive to subsequent data mining tasks.

The paper is structured as follows. Section 2 provides background on EMS and introduces the study case along with its major challenges. Section 3 describes the implementation of the multidimensional data modelling solution. Section 4 presents the performance of the data warehouse-based system in terms of efficiency and potential in context-aware data analysis. Section 5 discusses some related work. Finally, concluding remarks and future work are drawn in Section 6.

2. Background

Emergency medical services in mainland Portugal are coordinated by Instituto Nacional de Emergência Médica (INEM)[2]. In most cases, medical emergencies are reported via the 112 number, where specialized medical staff classifies the emergency and dispatches the proper vehicles (ambulance, helicopter, life-support vehicle, among others) along with the medical staff. Each vehicle is equipped to deal with different situations, from light injuries to life-support. In 2019, INEM answered (dispatched) nearly 1.3 million calls (1.2 million vehicles).

INEM stores the data associated with all medical emergencies in a *relational database*. The database maintains procedural-based views of medical emergencies, including their spatiotemporal frame, as well as care-based views comprising emergency diagnoses, provided assistance, and outcomes. Despite the ongoing efforts, as not all dispatched vehicles and staff are from INEM, abundant records associated with the emergency response are missing, particularly those pertaining to the timestamps of site arrival, departure, and hospital redirection.

Until now, INEM has not considered the role of situational context in shaping emergency prevalence and response. Moreover, the assignment of medical emergency vehicles in a preventive way is rarely done, except for special events for which, by law, the event's organization needs to ensure nearby emergency resources (concerts, sports matches, etc.). Figs. 1 and 2 illustrate the effect that festivals and weather factors can bear on the prevalence of specific emergencies. Part of the work reported in this paper is to easily incorporate these situational contexts into medical emergency data analysis. The context-aware predictive modelling of emergency events is essential to support resource allocation and assist in vehicle allocation at large gatherings.

[2]http://www.inem.pt (accessed 2020-05-02)

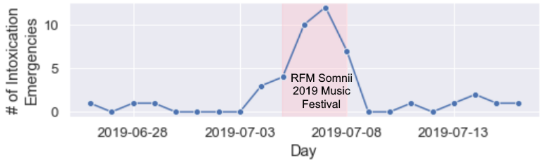

Figure 1: Intoxication cases in the absence and presence of the RFM Somnii festival (1km radius).

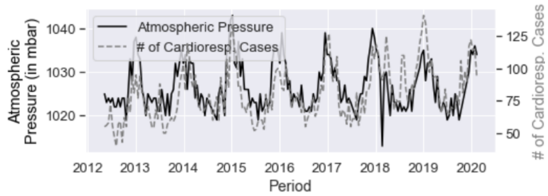

Figure 2: Correlation of cardiorespiratory arrest cases with atmospheric pressure.

3. Proposed Multidimensional Data Model

Principles. To address the challenges introduced in Section 1 and 2, we integrate available sources of emergency and situational context data using a multidimensional schema (Fig. 3) with four major properties of interest.

First, spatial and temporal dimensions with multiple calendric and territorial hierarchies are included to support the specification of spatiotemporal queries. These dimensions capture location and time information and are linked to the different stages of an emergency process (e.g., dispatch, arrival, return) to support process-related queries. The location dimension captures the geographical coordinates as well as the district and municipality that encompasses the location, while the time dimension captures the relevant time components such as the year, month, and minute.

Second, an arbitrarily high number of context-specific dimensions and facts are introduced to capture one-to-many relationships between emergencies and situational context sources in accordance with the spatiotemporal footprint of the monitored large-scale events and sensor measurements. In the emergency response database, we implement dimensions capturing weather, festivity, and sports event information.

Third, multiple fact tables are further instantiated to differentiate between complete and incomplete emergency occurrences, thus supporting the efficiency of stage-specific queries. In this context, we preserve the stance of facts as materializations of the linked dimensions.

Finally, complementary dimensions are further considered to hierarchically describe the typology (diagnostic)

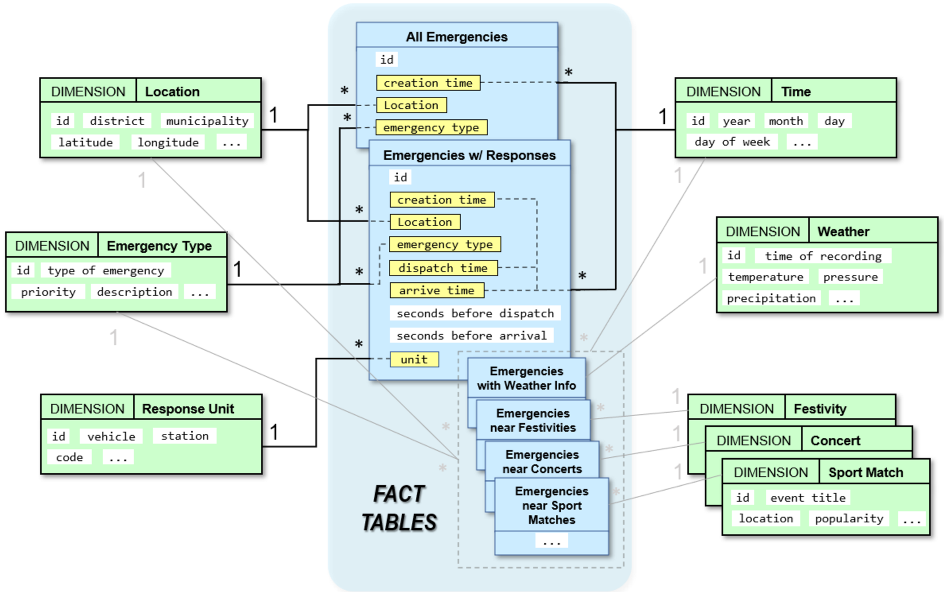

Figure 3: Multidimensional model: emergency records in cloven fact tables linked to spatiotemporal, emergency-specific and context dimensions with multiple hierarchies.

and severity of emergencies and abundant information on the dispatched vehicles and provided assistance. For instance, the unit dimension contains information about the responding vehicles, while the emergency dimension contains information such as the type and severity of the emergency. The introduced spatial, temporal, emergency categorization and contextual dimensions offer hierarchical content to support the incorporation of drill-down, roll-up, slicing and dicing operations within queries. These operations are essential to retrieve data structures, such as multivariate time series, conducive to subsequent descriptive or predictive tasks.

Despite the complexity of context-enriched emergency data, the proposed schema is intrinsically simple. Complementarily to expressive OLAP querying, a service layer is also provided to support parametric queries for advanced context-enriched spatiotemporal analytics.

Online context data sourcing. Ministries, municipalities and weather institutes generally have well-established efforts to gather and publicly provide relevant context data. Hence, periodic routines can be placed to dynamically acquire situational context from structured or/and semi-structured data sources via portals and APIs provided by national-wide initiatives, city Councils, institutes and other entities. Illustrating, the Lisbon city Council stores semi-structured representations of large-scale public events and urban traffic at Lisboa Aberta portal[3]. The spatiotemporal footprint of events and measurements can be autonomously inferred from their category and duration. Thus, the dynamic retrieval of context data and corresponding upload can be done in an automated fashion [8]. In this context, emergency event associations are precomputed once, alleviating the subsequent computational complexity of context-aware analytics.

4. Preliminary Results

This section gathers preliminary results from the proposed data warehousing approach for spatiotemporal emergency data analysis, discussing efficiency gains (section 4.1) and aspects of usefulness and interpretability (section 4.2). Experiments were run using SQL Server on Intel Xeon CPU E3-1230 v6 @ 3.50GHz with 16GB RAM.

4.1. Efficiency

Considering the Portuguese study case introduced in section 2, relevant data retrieval operations were applied on both the INEM relational database and the proposed multidimensional database. The selected operations cover spatiotemporal queries typically involved in data analysis tasks. Each query was executed 10 times to account for variability between runs. A two-tailed paired t-test was performed to assess the statistical significance of the observed differences. Table 1 shows the results. The full list of queries that we used are listed in Appendix A.

[3] http://lisboaaberta.cm-lisboa.pt/index.php/pt/ (accessed 2020-05-02)

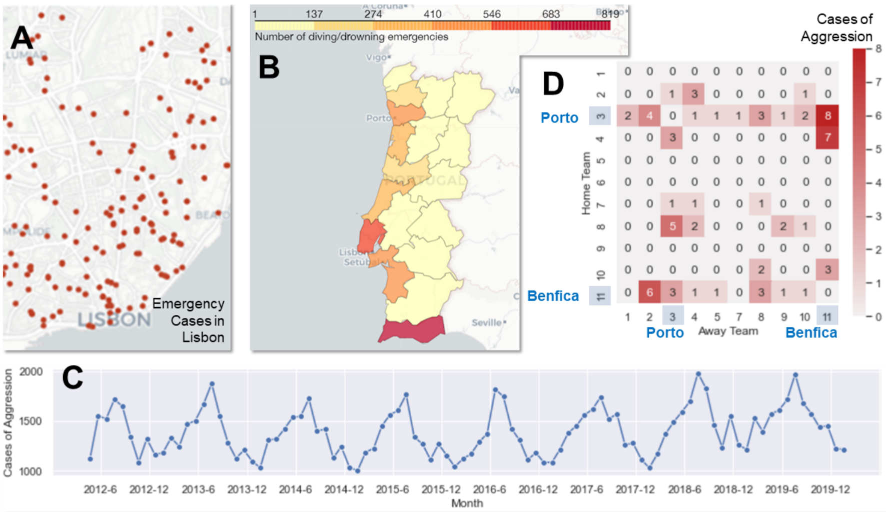

Figure 4: Visualizations generated from context-enriched emergency data. (a) Emergencies in Lisbon on a specified day (with anonymization filter), (b) Number of diving emergencies per district, (c) Number of aggression emergencies per month, (d) Number of aggression emergencies near soccer events grouped by teams playing. The queries used in retrieving the data necessary for generating these visualizations are listed in Appendix B.

Table 1
Comparison of execution times (in seconds) of queries executed over the original and the multidimensional databases.

Query	A	B	C	D	E	F	G	H
Original	0.28	58.90	0.47	71.36	73.77	59.67	72.61	73.01
Multidim.	**0.20**	**0.68**	**0.24**	**0.24**	**1.11**	**0.17**	**0.36**	**0.10**

Execution times on the multidimensional database consistently outperformed those on the relational database with statistical significance ($p < 0.01$). A small improvement could be observed on Query A, involving the retrieval of individual emergencies parameterized by date, due to a proper organization of temporal criteria. Efficiency improvements are observed for queries C, D, and E, which involved aggregating the number of emergencies according to specific criteria of interest, such as municipality and emergency type. Queries D and E involved large numbers of groups, magnifying efficiency differences. Queries B, G, and H involved the aggregation of emergency response times. For the original database, subtraction operations had to be performed between timestamps for these queries. This was not necessary on the multidimensional database as the differences are pre-computed. Overall, the data warehousing process on the context-enriched database improved data retrieval efficiency, supporting demanding analytics and visualization requests.

4.2. Context-aware data analysis support

We show the potential of the data warehousing system in context-aware data analysis support by performing various types of data retrieval operations associated with spatiotemporal data analysis. Fig. 4 provides visualizations of the outputs of context-enriched data queries.

Fig. 4a-b relies on individual emergencies' retrieval and grouping for a given spatial criteria. These simple queries can be structured into four parts: selecting the desirable fact table, joining the dimensions with relevant information to query, specifying the conditions to filter results, and specifying the grouping conditions, if any. This query structure is easy to construct and can easily be modified to accommodate various combinations of parameters. The grouping of related variables in dimensions further makes retrieval of desired information more intuitive. Retrieval of information can easily be done by joining the appropriate dimension rather than going through a list of all possible columns.

Fig. 4c shows the retrieval of time series data by grouping and aggregating desired emergencies into fixed time intervals. To this end, the individual timestamps in the time dimension are mapped into columns containing abundant calendric information. This allows usable queries in place of computationally expensive SQL functions (such as datepart) to extract individual date components, supporting the specification of the desirable time series granularity. The precomputation of values of interest, such as durations along emergency stages, further simplifies the data retrieval process.

The integration of contextual information from external sources for context-aware emergency analysis is further illustrated in Fig. 4d. In this example, aggression events occurring in the spatiotemporal footprint of sports matches are retrieved. This can be easily accom-

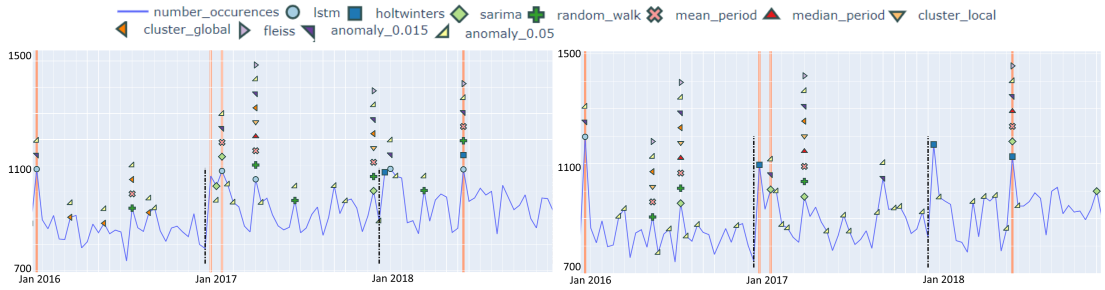

(a) all conditions

(b) intoxication, body pain, trauma, altered consciousness states

Figure 5: Application of various anomaly detection approaches (marked by the symbols above) in identifying unusual spikes in the number of emergencies along with different constraints. The vertical lines correspond to known registered anomalies.

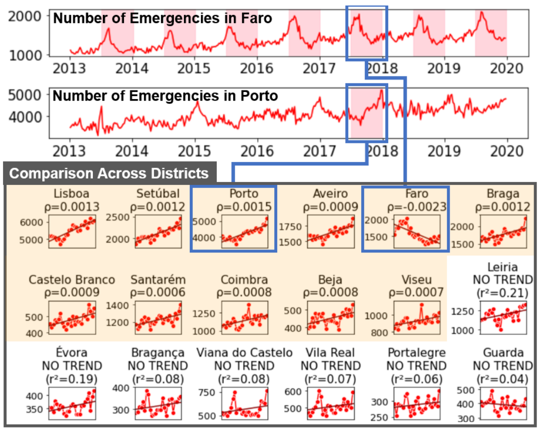

Figure 6: Interesting trend discovery allows us to discover trends that deviate in a particular context. I this case, the number of emergency occurrences in the district of Faro behaves differently (decreases) compared to those of other districts.

plished by identifying the desirable context dimensions and querying the fact entries holding a foreign key to the target sports match. The addition of context-specific fact tables not only ensures integrity by preventing foreign keys to be null; it also improves interpretability by explicitly stating which dimensions can be accessed and guaranteeing that the traced occurrences fall within the time and radius of a context event.

Finally, the multidimensional data warehouse supports the development of more complex data analysis processes. In the Portuguese study case, we present two examples: i) detection of anomalies in which unusual levels of occurrences of certain types of emergencies are extracted from the data (Figure 5) [9], and ii) discovery of interesting trends, where diverging patterns on the rise or drop in the number of cases within localities or emergency types are automatically extracted (Figure 6). These applications show the relevance of the data warehousing system in efficiently supporting convenient data analysis facilities.

5. Related Work

Numerous multidimensional modelling principles have been proposed for heterogeneous data source consolidation with spatiotemporal content [10]. However, in contrast to our work, the association of events based on their spatiotemporal footprint and the handling of missing data is pushed towards the analytical processing stage, hampering query expressivity and efficiency. Classic multidimensional databases have been extended to cater for specific scenarios and use cases. For example, Papadias et al. [11] introduced a framework for indexing spatiotemporal data for the efficient execution of ad-hoc group-bys using a combined spatial and temporal dimension. In the presence of data sources with undefined or partially defined structures, approaches have been proposed to automatically discover dimensions. In the work of Mansmann et al. [12], a data enrichment layer was added to detect structural elements in user-generated Twitter data with data mining techniques. Similarly, Gutiérrez-Batista et al. [13] used hierarchical clustering techniques to extract the multidimensional structure from textual data. Alternatively, spatiotemporal ontologies can be used to this end [14]. Thalhamer et al. [15] introduced the concept of active data warehouses, able to autonomously extract rules on behalf of the data analyst. The use of multidimensional database structures has been documented in various domains such as management [16], education [17], agriculture [18] and epidemiology [19]. In urban data, the handling of large amounts of spatiotemporal content is largely notable [20], and comprises identical challenges that can be addressed by a multidimensional approach, as demonstrated in our work.

6. Concluding Remarks

In this work, we applied best practices in multidimensional database modelling to consolidate emergency events and public sources of situational context to support and boost context-aware emergency data analysis. The solution addressed challenges pertaining to the massive, incomplete, and spatiotemporal nature of emergency event data, as well as challenges related to the sourcing, heterogeneity and varying spatiotemporal footprint of context sources. We further introduced a study case where the data warehousing system yielded statistically significant performance improvements. This work opens up the possibility to support and scale up spatiotemporal data analysis and data mining tasks with heavy context-enriched data retrieval demands, whether of descriptive or predictive nature. Amongst diverse ends, this possibility can be used to assist real-time monitoring and decision tasks in favour of self-adapting EMS able to timely detect and respond to emerging changes. The contributions of this work extend beyond medical emergency domains, as the principles apply to similar spatiotemporal domains with dependencies on situational factors like traffic and utility supply data.

Acknowledgements

This work is supported by Portuguese national funds through FCT, Fundação para a Ciência e Tecnologia (FCT), under projects UIDB/50021/2020, DSAIPA/AI/0044/2018 and DSAIPA/AI/0111/2018.

References

[1] B. S. Roudsari, A. B. Nathens, C. Arreola-Risa, P. Cameron, I. Civil, G. Grigoriou, R. L. Gruen, T. D. Koepsell, F. E. Lecky, R. L. Lefering, et al., Emergency medical service (ems) systems in developed and developing countries, Injury 38 (2007) 1001–1013.

[2] R. Aringhieri, M. E. Bruni, S. Khodaparasti, J. T. van Essen, Emergency medical services and beyond: Addressing new challenges through a wide literature review, Computers & Operations Research 78 (2017) 349–368.

[3] L. Aboueljinane, E. Sahin, Z. Jemai, J. Marty, A simulation study to improve the performance of an emergency medical service: application to the french val-de-marne department, Simulation modelling practice and theory 47 (2014) 46–59.

[4] H. J. Kam, J. O. Sung, R. W. Park, Prediction of daily patient numbers for a regional emergency medical center using time series analysis, Healthcare informatics research 16 (2010) 158.

[5] M. A. Mahmood, J. E. Thornes, F. D. Pope, P. A. Fisher, S. Vardoulakis, Impact of air temperature on london ambulance call-out incidents and response times, Climate 5 (2017) 61.

[6] T. Janchar, C. Samaddar, D. Milzman, The mosh pit experience: Emergency medical care for concert injuries, The American journal of emergency medicine 18 (2000) 62–63.

[7] Y. Katayama, K. Kiyohara, T. Kitamura, S. Hayashida, T. Shimazu, Influence of the covid-19 pandemic on an emergency medical service system: a population-based, descriptive study in osaka, japan, Acute medicine & surgery 7 (2020) e534.

[8] S. Cerqueira, R. Henriques, E. Arsénio, Integrative analysis of traffic and situational context data to support urban mobility planning, in: European Transport Conference 2020, 2020.

[9] L. Silva, H. Galhardas, V. Manquinho, R. Henriques, Uniano: robust and efficient anomaly consensus in time series sensitive to cross-correlated anomaly profiles, in: Proceedings of the 2021 SIAM International Conference on Data Mining (SDM), SIAM, 2021, pp. 324–332.

[10] O. Romero, A. Abelló, A survey of multidimensional modeling methodologies, International Journal of Data Warehousing and Mining (IJDWM) 5 (2009) 1–23.

[11] D. Papadias, Y. Tao, P. Kanis, J. Zhang, Indexing spatio-temporal data warehouses, in: Proceedings 18th International Conference on Data Engineering, IEEE, 2002, pp. 166–175.

[12] S. Mansmann, N. U. Rehman, A. Weiler, M. H. Scholl, Discovering olap dimensions in semi-structured data, Information Systems 44 (2014) 120–133.

[13] K. Gutiérrez-Batista, J. R. Campaña, M.-A. Vila, M. J. Martin-Bautista, Building a contextual dimension for olap using textual data from social networks, Expert Systems with Applications 93 (2018) 118–133.

[14] A. Salguero, F. Araque, C. Delgado, Spatio-temporal ontology based model for data warehousing, in: WSEAS International Conference. Proceedings. Mathematics and Computers in Science and Engineering, 7, World Scientific and Engineering Academy and Society, 2008.

[15] T. Thalhammer, M. Schrefl, M. Mohania, Active data warehouses: complementing olap with analysis rules, Data & Knowledge Engineering 39 (2001) 241–269.

[16] K.-W. Chau, Y. Cao, M. Anson, J. Zhang, Application of data warehouse and decision support system in construction management, Automation in construction 12 (2003) 213–224.

[17] E. M. A. Butt, S. Quadri, E. M. Zaman, Star schema implementation for automation of examination records, in: Proceedings of the International Conference on Frontiers in Education: Computer Science and Computer Engineering (FECS), The Steering Committee of The World Congress in Computer Science, 2012, p. 1.

[18] A. K. Gupta, B. D. Mazumdar, Multidimensional schema for agricultural data warehouse, International Journal of Research in Engineering and Technology 2 (2013) 245–253.

[19] B. Amin, H. Djamila, Towards a spatiotemporal data warehouse for epidemiological surveillance, Instrumentation Mesure Metrologie 18 (2019) 1–7.

[20] E. Zim'nyi, Spatio-temporal data warehouses and mobility data: Current status and research issues, in: 2012 19th International Symposium on Temporal Representation and Reasoning, IEEE, 2012, pp. 6–9.

Appendix A: Selected queries for efficiency evaluation

Original Database	Multidimensional Database
A: Get emergency dates, types and priority level for a given day.	
SELECT time, type, priority FROM emergencies WHERE CAST(time AS date) = <some date>	SELECT time, type, priority FROM all.facttable AS a INNER JOIN emergency.dimension AS b ON a.emergency = b.id WHERE time BETWEEN <some dates>
B: Get the response time and priority of each emergency.	
SELECT arrivaltime-time, priority FROM emergencies WHERE arrivaltime IS NOT NULL and time IS NOT NULL	SELECT arrivalseconds, priority FROM completeresponse.facttable AS a INNER JOIN emergency.dimension AS b ON a.emergency = b.id
C: Get the number of emergencies per priority level.	
SELECT priority, COUNT(*) AS count FROM emergencies GROUP BY priority ORDER BY priority	SELECT priority, COUNT(*) FROM all.facttable AS a INNER JOIN emergency.dimension AS b ON a.emergency = b.id GROUP BY priority ORDER BY priority
D: Get the number of emergencies per emergency type.	
SELECT type, COUNT(*) AS count FROM emergencies GROUP BY type ORDER BY type	SELECT type, COUNT(*) FROM all.facttable AS a INNER JOIN emergency.dimension AS b ON a.emergency = b.id GROUP BY type ORDER BY type
E: Get the number of emergencies per municipality.	
SELECT municipality, COUNT(*) AS count FROM emergencies GROUP BY municipality ORDER BY municipality	SELECT municipality, COUNT(*) FROM all.facttable AS a INNER JOIN emergency.dimension AS b ON a.emergency = b.id GROUP BY municipality ORDER BY municipality
F: Get the number of a specific type of occurrence, per month and year.	
SELECT DATEPART(year, CAST(time AS date)), DATEPART(month, CAST(time AS date)), COUNT(*) FROM emergencies WHERE type = 'DROWNING' GROUP BY DATEPART(year, CAST(time AS date)) DATEPART(month, CAST(time AS date)) ORDER BY DATEPART(year, CAST(time AS date)) DATEPART(month, CAST(time AS date))	SELECT year, month, COUNT(*) FROM all.facttable AS a INNER JOIN time.dimension AS b ON a.time = b.id INNER JOIN emergency.dimension AS c ON a.emergency = b.id WHERE type = 'DROWNING' GROUP BY year, month ORDER BY year, month
G: For occurrences with units, get the average time to activation of the response unit.	
SELECT AVG(dispatchtime-time) FROM emergencies WHERE dispatchtime IS NOT NULL and time IS NOT NULL AND dispatchtime-time > 0	SELECT AVG(dispatchseconds) FROM withunits.facttable WHERE dispatchseconds > 0
H: For occurrences that are complete, get the maximum time to destination.	
SELECT MAX(destinationtime-time) FROM emergencies WHERE destinationtime IS NOT NULL and time IS NOT NULL AND destinationtime-time > 0	SELECT MAX(desintationseconds) FROM completeresponse.facttable WHERE desintationseconds > 0

Appendix B: Selected queries to generate plots

A. Get the location of emergency cases in Lisbon on a specified date (date withheld for privacy)	
SELECT * FROM all.facttable AS a INNER JOIN location.dimension AS b ON a.locationid = b.id WHERE district='Lisboa' AND time BETWEEN <some time>	Select fact table. Add desired dimensions. State filtering conditions.
B. Get the number of diving / drowning emergencies per district	
SELECT district, COUNT(id) AS count FROM all.facttable AS a INNER JOIN location.dimension AS b ON a.locationid = b.id INNER JOIN emergency.dimension AS c ON a.emergencyid = c.id WHERE type = 'DROWNING' GROUP BY district	Select fact table and aggregation. Add desired dimensions. State filtering conditions. Specify aggregation grouping.
C. Get the number of aggression emergencies per month over time	
SELECT year, month, COUNT(id) AS count FROM all.facttable AS a INNER JOIN time.dimension AS b ON a.time = b.id INNER JOIN emergency.dimension AS c ON a.emergencyid = c.id WHERE type = 'AGRESSION' GROUP BY year, month	Select fact table and aggregation. Add desired dimensions. State filtering conditions. Specify aggregation grouping.
D. Get the number of aggression emergencies within 1km radius of football events, grouped by the teams playing	
SELECT team1, team2, COUNT(id) AS count FROM ftbl.facttable AS a INNER JOIN ftbl.dimension AS b ON a.ftbl = b.id INNER JOIN emergency.dimension AS c ON a.emergencyid = c.id WHERE type = 'AGRESSION' GROUP BY team1, team2	Select fact table and aggregation. Add desired dimensions. State filtering conditions. Specify aggregation grouping.

Note: Names of tables and columns have been altered for security purposes.