A Human-Centric Environment (HCE) Framework for Sustainable Production in a Bakery

Luca Laboccetta[1], Giorgio Terracina[1], Francesco Calimeri[1], Simona Perri[1], Massimiliano Ruffolo[2], Davide Iacopino[2], Marta Maria[2] and Salvatore Iiritano[2]

[1]*University of Calabria, 87036 Rende, Italy*
[2]*Revelis S.r.l., 87036 Rende, Italy*

Abstract

Employee well-being is an increasingly crucial factor for sustainability and efficiency in modern industry, especially within the context of Industry 5.0, which places humans at the center of production processes. This paper presents some work in progress towards the development of an innovative framework that aims to monitor operators' physical stress and well-being and then optimize task assignment in an industrial bakery environment. In particular, in this paper, we concentrate on the description of the general architecture and on the modules for stress prediction and management. This is part of a more general project developed under the NRRP MUR initiative FAIR (Future AI Research): Green-aware AI. Smartwatches and environmental sensors are the main sources of time-series biometric data. This data is used to predict future stress levels through an AI-based approach. These predictions are then used to dynamically reassign or pause operators. The goal is to minimize stress and prevent overload while maintaining adherence to the production plan. Preliminary results of some experiments conducted at a real Industrial Bakery Factory over a three-month period in the production, oven, and packaging departments, demonstrated how the integration of Internet of Everything (IoE) systems and AI can improve employee health and operational efficiency.

Keywords

Industry 5.0, Human-Centric AI, Stream Reasoning, Answer Set Programming (ASP), Employee Well-being, Smart Factory, IoT (Internet of Things)

1. Introduction

The landscape of modern manufacturing is rapidly evolving, driven by the principles of Industry 4.0 and increasingly, Industry 5.0. While Industry 4.0 focused on automation and data exchange, Industry 5.0 extends this vision by placing human well-being at the core of industrial processes, emphasizing sustainability and human-centric approaches. In this paradigm, the health and psychophysical state of employees are no longer secondary considerations but integral components of operational efficiency and long-term business success. This shift necessitates advanced technological solutions capable of real-time monitoring and adaptive management of human factors in dynamic industrial environments. Specifically, the ability to accurately assess and predict critical human states, such as employee stress levels, becomes paramount for proactive interventions and optimized resource allocation. This work is devoted to present some ongoing work on the design and development of a system, named InCoP (Industry 5.0 Collaborative Platform), which is part of the activities funded by the NRRP MUR initiative FAIR (Future AI Research): Green-aware AI. The general aim is to create a versatile and scalable technology platform capable of supporting the implementation of Industry 5.0 principles in various industrial sectors. The platform is designed to be modular and adaptable, allowing the integration of

2nd Workshop on Green-Aware Artificial Intelligence, 28th European Conference on Artificial Intelligence (ECAI 2025), October 25–30, 2025, Bologna, Italy

*Corresponding author.
†These authors contributed equally.
luca.laboccetta@unical.it (L. Laboccetta); giorgio.terracina@unical.it (G. Terracina); francesco.calimeri@unical.it (F. Calimeri); simona.perri@unical.it (S. Perri); massimiliano.ruffolo@revelis.eu (M. Ruffolo); davide.iacopino@revelis.eu (D. Iacopino); marta.maria@revelis.eu (M. Maria); salvatore.iiritano@revelis.eu (S. Iiritano)
0009-0005-7437-9505 (L. Laboccetta); 0000-0002-3090-7223 (G. Terracina); 0000-0002-0866-0834 (F. Calimeri); 0000-0002-8036-5709 (S. Perri)

new features and technologies according to the specific needs of each application context. Main goals of the platform consist in: *(i)* improve the efficiency and flexibility of the production process; *(ii)* create a smarter production system that can learn from data and continuously improve its performance; *(iii)* enable mass customization of products; *(iv)* demonstrate the potential impact of Industry 5.0 through the definition of metrics and the evaluation of results obtained in the pilot company and other application contexts.

To address these complex requirements, two interconnected fields of Artificial Intelligence appear to be particularly relevant. On the one hand, Machine/Deep Learning techniques [1] are crucial for discerning subtle patterns in physiological data and for forecasting future human states like stress. On the other hand, Stream Reasoning (SR) [2] has emerged as a promising field to address the complexities of real-time data analysis and deductive inference over continuous, high-volume data streams.

In the context of machine/deep learning approaches, Recurrent Neural Networks (RNNs) stand out as a highly suitable class of neural networks for processing sequential data like time series, given their inherent ability to maintain memory of past observations and leverage this history for predictions. Among RNN architectures, Long Short-Term Memory (LSTM) [3] and Gated Recurrent Unit (GRU) networks are particularly effective due to their mechanisms for handling vanishing gradients and capturing long-range dependencies, making them state-of-the-art for tasks such as stress prediction, activity recognition, and general physiological monitoring.

While machine/deep learning models excel at pattern recognition and prediction, their seamless integration with symbolic reasoning systems poses unique challenges and opportunities for hybrid AI approaches.

As far as symbolic reasoning systems is concerned, Stream Reasoning (SR) [2] has emerged as a fundamental paradigm in the real-time data analysis landscape, where the incessant flow of information demands not only the processing of data streams, but also the application of logical reasoning to extract meaningful insights. This dual requirement is particularly relevant in domains with the presence of heterogeneous and dynamic data streams, such as Smart Cities, the Internet of Things (IoT), or Healthcare, where the need for timely and accurate decision making can be critical. In recent years, the field has witnessed the development of various approaches to address the challenges posed by the dynamic nature of (possibly big) data streams; indeed, different SR approaches have been proposed [4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18] in contexts like Data Stream Management Systems (DSMS), Complex Event Processing (CEP), Semantic Web, and Knowledge Representation and Reasoning (KRR).

In particular, the present paper focuses on the Human-Centric Environment (HCE) framework for sustainable production in a bakery, which is a part of the overall InCoP platform. Interestingly, HCE integrates symbolic and sub-symbolic reasoning. Specifically, it uses Recurrent Neural Networks (RNNs) to anticipate stress levels and dynamic task reassignment through Stream Reasoning to react in real time to potentially harmful situations. The Stream Reasoning tool used in this project is DP-sr [19, 18].

Through preliminary experiments, we aim to demonstrate how AI-driven stress monitoring, augmented by machine learning (ML)-based stress prediction, can effectively optimize task assignments and foster a safer, more efficient, and sustainable production process.

2. Description of the HCE Framework

In this Section, we first describe the general Human-Centric Environment (HCE) framework (shown in Figure 1), and then concentrate on some of its main modules, namely data generation and collection modules, stress prediction, and dynamic task re-assignment.

Input data come from smartwatches worn by employees, suitably anonymized to avoid privacy issues, and from environmental sensors strategically placed in the production departments. Data are ingested and stored for subsequent elaboration. In parallel, acquired measurements are immediately queued and fed to the stress prediction module. The module's output is stored for next-step processing. The dynamic task reassignment module continuously monitors ingested data and updates operator assignments if needed. Alerts are sent back to individual smartwatches to privately notify workers when they need to

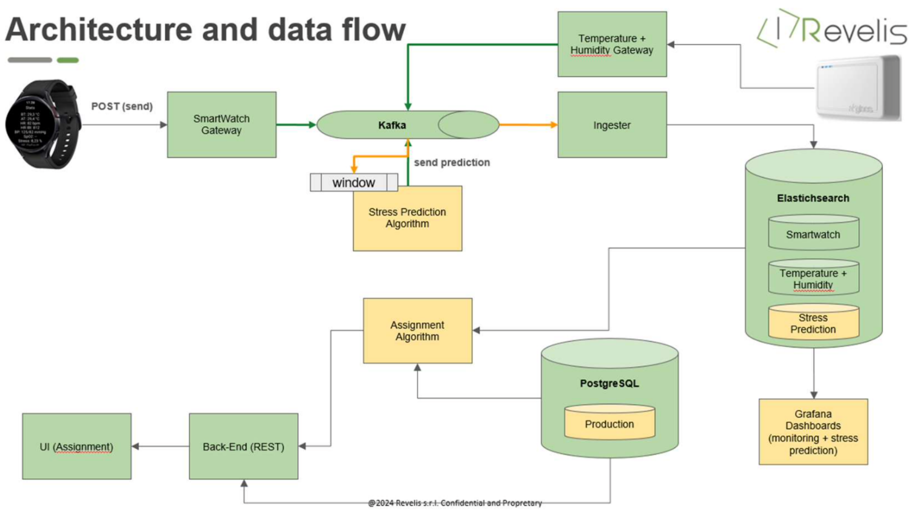

Figure 1: Architecture of the HCE Framework

switch tasks or rest.

It is worth pointing out that the framework shown in Figure 1 is only a part of the more general architecture composing the InCoP platform. The design and description of the complete framework is part of our future work.

2.1. Data Generation and Collection

The implementation of the system in the pilot bakery involved the installation of Internet of Everything (IoE) devices in three key departments: production, ovens, and packaging. Involved devices include:

- **Smartwatches:** Worn by employees, these devices continuously monitor various biometric parameters, including heart rate, blood pressure, blood oxygen level, body temperature, and estimated stress level.
- **Environmental Sensors (Airgloss):** Strategically placed in the production departments, these sensors detect real-time environmental conditions such as temperature and humidity.

As far as data flow is concerned, data from smartwatches and environmental sensors are sent via gateways to a Kafka-based messaging system. An Ingester component reads this data and stores it in an Elasticsearch database.

2.2. Stress prediction

The Stress Prediction Algorithm is a dedicated component responsible for forecasting operators' future stress levels through a specialized Machine Learning (ML) model. This model leverages historical stress trends, utilizing the 12 observations immediately preceding the prediction point. Forecasts are generated autoregressively for 24 subsequent measurements, operating iteratively by incorporating predicted values as new input data for subsequent predictions. The system dynamically acquires new measurements directly from the Kafka queue for prediction, with the output subsequently fed back into the same Kafka queue. The Ingester component then processes this output, transferring it to a dedicated Elasticsearch index containing future stress level predictions. This index serves as the source for both Grafana dashboards (for monitoring) and the dynamic assignment module (for operational adjustments).

Recurrent Neural Networks (RNNs) were identified as the most suitable class for time series analysis. Their inherent structure allows them to retain a memory of past observations, thereby effectively identifying relationships between input sequences and corresponding outputs.

Testing involved multiple neural network architectures, typically concatenating a recurrent section with a dense network segment. The recurrent part specifically explored different configurations of well-known recurrent layers, such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU). The dense portion consisted of sequential layers, with each output vector passed through a LeakyReLU activation function to enhance training efficiency by improving gradient propagation and mitigating issues like the vanishing gradient problem. To combat overfitting and improve generalization, Dropout and Batch Normalization layers were also integrated. Training management benefited from an EarlyStopping mechanism, which halted training if no predictive improvements were observed over a specified number of epochs on a validation set. The optimization target for training was the mean square error loss function.

The final architecture comprised 4 LSTM layers (with decreasing neuron counts), interspersed with Batch Normalization and Dropout (excluding the initial layer), followed by 3 fully connected dense layers (decreasing to 8 neurons). This final layer produces 8 outputs, as the network simultaneously predicts the next stress level for each of the 8 smartwatches based on their respective historical stress levels. The entire stress prediction component was developed in Python, utilizing the Tensorflow and Keras libraries for neural network implementation. The model is summarized in Figure 2.

Stress levels are classified into categories (Low: 0-50, Medium: 50-70, High: 70-100) based on literature-defined thresholds, emphasizing a well-being-oriented approach rather than medical diagnosis.

```
model = Sequential()
model.add(LSTM(64, return_sequences=True, activation="relu", input_shape=(look_back, 8)))
model.add(LSTM(48, return_sequences=True, activation="relu"))
model.add(BatchNormalization())
model.add(Dropout(0.2))
model.add(LSTM(32, return_sequences=True, activation="relu"))
model.add(BatchNormalization())
model.add(Dropout(0.2))
model.add(LSTM(24, activation="relu"))
model.add(Dense(24))
model.add(LeakyReLU(alpha=0.1))
model.add(Dense(16))
model.add(LeakyReLU(alpha=0.1))
model.add(Dense(8, activation='sigmoid'))
model.compile(loss='mse', optimizer='adam')
```

Figure 2: Forecasting model implementation

2.3. Dynamic task re-assignment

The Assignment Algorithm dynamically manages operator tasks to ensure a balance between physical and mental well-being and production efficiency. The workflow of this module is depicted in Figure 3. The module leverages the following data inputs:

- The real-time and predicted stress levels of operators (received from MongoDB every 5 minutes)
- The weekly shift rotation schedules

The logic program within the DP-sr system dynamically adjusts operator assignments based on these integrated data streams. Specifically, key rules in the program include:

- If an operator records four consecutive measurements with a **medium stress level**, they are reassigned to a department where the average stress level of operators is low.
- If an operator records four consecutive measurements with a **high stress level**, they are put into a "Rest" state until their stress level returns to Medium or Low.
- In order to keep compliance of the production plan, if necessary some other, not stressed, operator is moved to the department just left by the stressed operator.

The dual objectives of this module are to ensure compliance with the production plan, thereby avoiding delays and inefficiencies, and to safeguard worker well-being by preventing physical and mental overload. The DP-sr system processes the relevant data and then stores the updated operator assignments back into MongoDB for system-wide access and visualization.

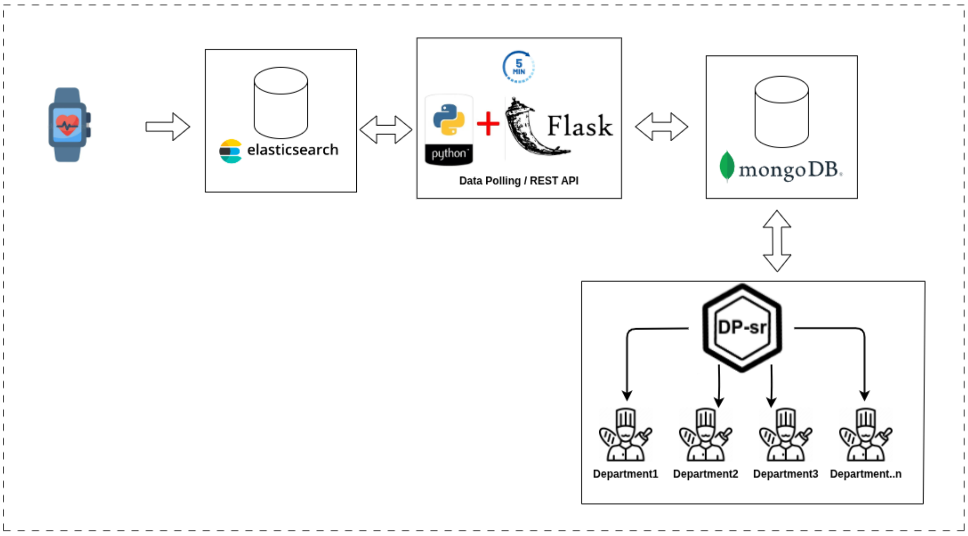

Figure 3: Workflow of the dynamic assignment algorithm

3. Results and Discussion

With a preliminary implementation of the HCE framework, we conducted a three-month experimentation at a pilot Industrial Bakery. These preliminary experiments yielded significant results, demonstrating the tangible benefits of a human-centric AI approach in a real-world industrial setting. The evaluation focused on both business and technical Key Performance Indicators (KPIs).

3.1. Business KPIs Achieved

The project aimed to maximize productivity and ensure worker well-being, with the following outcomes:

- **Maximize Cost Reduction (Equipment Overall Effectiveness):** Improved from an AS-IS value of 60% to an achieved value of 70% (with a TO-BE target of 80%).
- **Maximize Personnel Productivity (Prevent accidents and safety risks):** Maintained an AS-IS value of 60%, with a TO-BE target of 70%. While not increased, the system contributed to maintaining safety levels.
- **Digital Transformation (Architecture integration):** Achieved the TO-BE target of 1 from an AS-IS value of 0.4, indicating successful integration of data sources, people, and smart objects into a collaborative ecosystem.
- **Performance Management (Worker satisfaction, idle time, production delays):** Improved from an AS-IS value of 60% to an achieved value of 70% (with a TO-BE target of 80%), reflecting increased worker satisfaction and optimized time utilization.

Overall, the system's interventions led to increased safety at work through proactive management of fatigue and stress, resulting in a reduction in accident risk and improved worker well-being. Dynamic task management also contributed to increased productivity by adapting worker roles to their physical and psychological conditions, reducing fatigue-related errors.

3.2. Technical KPIs Achieved

From a technical point of view, the system demonstrated robust performance and scalability:

- **Detection of workers' vital parameters:** The system successfully achieved its TO-BE value, providing 12 detections per hour with 5-minute timestamps.
- **Temperature and humidity monitoring:** Similarly, environmental parameters were detected at the target frequency of 12 detections per hour.
- **Scalability and Interoperability (Number of sources, people, and built-in smart objects):** The system achieved its TO-BE target of 8, from an AS-IS value of >5, indicating successful integration of a diverse set of data sources and smart objects.

These technical achievements underscore the robustness and reliability of the data collection and processing pipeline, crucial for real-time SR.

3.3. Barriers Faced and Lessons Learned

The implementation phase encountered several challenges:

- **Sensor malfunction and connectivity issues:** Environmental sensors provided inaccurate readings, and connectivity between devices and the central system was sometimes unstable. Solutions involved regular maintenance, periodic calibrations, and corporate Wi-Fi enhancements with backup solutions.
- **Compliance and privacy concerns:** Staff raised concerns about data privacy, especially sensitive health information. This was addressed through awareness sessions, GDPR compliance (anonymization, restricted access), and transparent communication.
- **Device adoption resistance and data quality:** Some workers were reluctant to use smartwatches, and collected data could be incomplete or noisy. This required ongoing support, targeted training, and data validation techniques.

Key lessons learned emphasize the importance of continuous staff education, robust data management and regulatory compliance, technological robustness (preventative maintenance for hardware), adaptability of AI algorithms, and effective internal communication among stakeholders. These lessons are vital for the scalability and long-term success of such complex Industry 5.0 projects.

4. Conclusions and Future Work

This paper presented some preliminary results on the ongoing design, development and testing of the InCoP platform. In particular, we focused on the Human-Centric Environment Framework developed for a bakery towards sustainable production. Preliminary experimental results demonstrated the practical feasibility and tangible benefits of using a mixture of ML and ASP-based SR approaches to address operator stress and dynamic task assignment. The proposed approach resulted promising to enhance employee well-being and operational efficiency. The achieved business and technical KPIs validate the system's effectiveness in a real-world industrial context, highlighting improvements in productivity, worker satisfaction, and safety. Future work include not only the design and implementation of other components of the InCoP platform, but also further refinements in data acquisition and elaboration in order to further improve KPIs and staff engagement in application development and testing.

Declaration on Generative AI

The author(s) confirm that no Generative AI tools were used in the creation of this article.

Acknowledgments

We acknowledge financial support from projects: PNRR MUR PE0000013-FAIR, Fa.Per.M.E. (CUP H53C22000640006), NutriDieMMe (CUP H53C22000940001). Francesco Calimeri is member of the Gruppo Nazionale Calcolo Scientifico-Istituto Nazionale di Alta Matematica (GNCS-INdAM).

References

[1] I. H. Sarker, Deep learning: A comprehensive overview on techniques, taxonomy, applications and research directions, SN Comput. Sci. 2 (2021) 420. URL: https://doi.org/10.1007/s42979-021-00815-1. doi:10.1007/S42979-021-00815-1.

[2] D. Dell'Aglio, E. D. Valle, F. van Harmelen, A. Bernstein, Stream reasoning: A survey and outlook, Data Sci. 1 (2017) 59–83. URL: https://doi.org/10.3233/DS-170006. doi:10.3233/DS-170006.

[3] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Comput. 9 (1997) 1735–1780. URL: https://doi.org/10.1162/neco.1997.9.8.1735. doi:10.1162/neco.1997.9.8.1735.

[4] D. F. Barbieri, D. Braga, S. Ceri, E. D. Valle, M. Grossniklaus, C-SPARQL: a continuous query language for RDF data streams, Int. J. Semantic Comput. 4 (2010) 3–25. URL: https://doi.org/10.1142/S1793351X10000936. doi:10.1142/S1793351X10000936.

[5] D. L. Phuoc, M. Dao-Tran, J. X. Parreira, M. Hauswirth, A native and adaptive approach for unified processing of linked streams and linked data, in: L. Aroyo, C. Welty, H. Alani, J. Taylor, A. Bernstein, L. Kagal, N. F. Noy, E. Blomqvist (Eds.), The Semantic Web - ISWC 2011 - 10th International Semantic Web Conference, Bonn, Germany, October 23-27, 2011, Proceedings, Part I, volume 7031 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 370–388. URL: https://doi.org/10.1007/978-3-642-25073-6_24. doi:10.1007/978-3-642-25073-6_24.

[6] J. Hoeksema, S. Kotoulas, High-performance distributed stream reasoning using s4, in: Ordring Workshop at ISWC 2011, 2011, pp. 1–8.

[7] T. Pham, M. I. Ali, A. Mileo, C-ASP: continuous asp-based reasoning over RDF streams, in: M. Balduccini, Y. Lierler, S. Woltran (Eds.), Logic Programming and Nonmonotonic Reasoning - 15th International Conference, LPNMR 2019, Philadelphia, PA, USA, June 3-7, 2019, Proceedings, volume 11481 of *Lecture Notes in Computer Science*, Springer, 2019, pp. 45–50. URL: https://doi.org/10.1007/978-3-030-20528-7_4. doi:10.1007/978-3-030-20528-7_4.

[8] A. Mileo, A. Abdelrahman, S. Policarpio, M. Hauswirth, Streamrule: A nonmonotonic stream reasoning system for the semantic web, in: W. Faber, D. Lembo (Eds.), Web Reasoning and Rule Systems - 7th International Conference, RR 2013, Mannheim, Germany, July 27-29, 2013. Proceedings, volume 7994 of *Lecture Notes in Computer Science*, Springer, 2013, pp. 247–252. URL: https://doi.org/10.1007/978-3-642-39666-3_23. doi:10.1007/978-3-642-39666-3_23.

[9] T. M. Do, S. W. Loke, F. Liu, Answer set programming for stream reasoning, in: C. J. Butz, P. Lingras (Eds.), Advances in Artificial Intelligence - 24th Canadian Conference on Artificial Intelligence, Canadian AI 2011, St. John's, Canada, May 25-27, 2011. Proceedings, volume 6657 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 104–109. URL: https://doi.org/10.1007/978-3-642-21043-3_13. doi:10.1007/978-3-642-21043-3_13.

[10] M. Gebser, T. Grote, R. Kaminski, T. Schaub, Reactive answer set programming, in: J. P. Delgrande, W. Faber (Eds.), Logic Programming and Nonmonotonic Reasoning - 11th International Conference, LPNMR 2011, Vancouver, Canada, May 16-19, 2011. Proceedings, volume 6645 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 54–66. URL: https://doi.org/10.1007/978-3-642-20895-9_7. doi:10.1007/978-3-642-20895-9_7.

[11] J. Calbimonte, Ó. Corcho, A. J. G. Gray, Enabling ontology-based access to streaming data sources, in: P. F. Patel-Schneider, Y. Pan, P. Hitzler, P. Mika, L. Zhang, J. Z. Pan, I. Horrocks, B. Glimm (Eds.), The Semantic Web - ISWC 2010 - 9th International Semantic Web Conference, ISWC 2010, Shanghai, China, November 7-11, 2010, Revised Selected Papers, Part I, volume 6496 of *Lecture Notes in*

Computer Science, Springer, 2010, pp. 96–111. URL: https://doi.org/10.1007/978-3-642-17746-0_7. doi:10.1007/978-3-642-17746-0_7.
[12] D. Anicic, P. Fodor, S. Rudolph, N. Stojanovic, EP-SPARQL: a unified language for event processing and stream reasoning, in: S. Srinivasan, K. Ramamritham, A. Kumar, M. P. Ravindra, E. Bertino, R. Kumar (Eds.), Proceedings of the 20th International Conference on World Wide Web, WWW 2011, Hyderabad, India, March 28 - April 1, 2011, ACM, 2011, pp. 635–644. URL: https://doi.org/10.1145/1963405.1963495. doi:10.1145/1963405.1963495.
[13] H. Beck, M. Dao-Tran, T. Eiter, LARS: A logic-based framework for analytic reasoning over streams, Artif. Intell. 261 (2018) 16–70. URL: https://doi.org/10.1016/j.artint.2018.04.003. doi:10.1016/J.ARTINT.2018.04.003.
[14] H. R. Bazoobandi, H. Beck, J. Urbani, Expressive stream reasoning with laser, in: C. d'Amato, M. Fernández, V. A. M. Tamma, F. Lécué, P. Cudré-Mauroux, J. F. Sequeda, C. Lange, J. Heflin (Eds.), The Semantic Web - ISWC 2017 - 16th International Semantic Web Conference, Vienna, Austria, October 21-25, 2017, Proceedings, Part I, volume 10587 of *Lecture Notes in Computer Science*, Springer, 2017, pp. 87–103. URL: https://doi.org/10.1007/978-3-319-68288-4_6. doi:10.1007/978-3-319-68288-4_6.
[15] H. Beck, T. Eiter, C. Folie, Ticker: A system for incremental asp-based stream reasoning, Theory Pract. Log. Program. 17 (2017) 744–763. URL: https://doi.org/10.1017/S1471068417000370. doi:10.1017/S1471068417000370.
[16] X. Ren, O. Curé, H. Naacke, G. Xiao, Bigsr: real-time expressive RDF stream reasoning on modern big data platforms, in: N. Abe, H. Liu, C. Pu, X. Hu, N. K. Ahmed, M. Qiao, Y. Song, D. Kossmann, B. Liu, K. Lee, J. Tang, J. He, J. S. Saltz (Eds.), IEEE International Conference on Big Data (IEEE BigData 2018), Seattle, WA, USA, December 10-13, 2018, IEEE, 2018, pp. 811–820. URL: https://doi.org/10.1109/BigData.2018.8621947. doi:10.1109/BIGDATA.2018.8621947.
[17] F. Calimeri, M. Manna, E. Mastria, M. C. Morelli, S. Perri, J. Zangari, I-dlv-sr: A stream reasoning system based on I-DLV, Theory Pract. Log. Program. 21 (2021) 610–628. URL: https://doi.org/10.1017/S147106842100034X. doi:10.1017/S147106842100034X.
[18] L. Laboccetta, E. Mastria, F. Calimeri, N. Leone, S. Perri, G. Terracina, Towards effective asp-based stream reasoning: Facilitate the reasoning over patterns of events, in: PPDP, ACM, 2024, pp. 18:1–18:12.
[19] F. Calimeri, E. Mastria, S. Perri, Dp-sr: a purely declarative programming framework for stream reasoning, 2024. URL: https://sites.google.com/unical.it/dp-sr.