



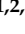




Article

Enriching Artificial Intelligence Explanations with Knowledge Fragments

Jože M. Rožanec^{1,2,3,*} , Elena Trajkova⁴ , Inna Novalija² , Patrik Zajec^{1,2} , Klemen Kenda^{1,2,3} , Blaž Fortuna³  and Dunja Mladenec² 

¹ Jožef Stefan International Postgraduate School, Jamova 39, 1000 Ljubljana, Slovenia

² Jožef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia

³ Qlector d.o.o., Rovšnikova 7, 1000 Ljubljana, Slovenia

⁴ University of Ljubljana, Faculty of Electrical Engineering, Tržaška c. 25, 1000, Ljubljana, Slovenia

* Correspondence: joze.rozanec@ijs.si; (J.M.R.)

Abstract: Artificial Intelligence models are increasingly used in manufacturing to inform decision-making. Responsible decision-making requires accurate forecasts and an understanding of the models' behavior. Furthermore, the insights into models' rationale can be enriched with domain knowledge. This research builds explanations considering feature rankings for a particular forecast, enriching them with media news entries, datasets' metadata, and entries from the Google Knowledge Graph. We compare two approaches (embeddings-based and semantic-based) on a real-world use case regarding demand forecasting.

Keywords: Explainable Artificial Intelligence; Human-Centric Artificial Intelligence; Smart Manufacturing; Demand Forecasting; Industry 4.0; Industry 5.0

1. Introduction

The industry is the part of the economy that is concerned with a highly mechanized and automated production of material goods [1]. Since the beginning of industrialization, technological breakthroughs allowed to achieve increasing manufacturing efficiencies and led to paradigm shifts. The latest such shifts are known as Industry 4.0 and Industry 5.0 [2,3], and aim to foster and increase the digitalization, networking, and automation of the manufacturing processes. Furthermore, these, in turn, enable tighter integration between physical and cyber domains (e.g., through cyber-physical systems [4], and digital twins [5]), increased deployment of intelligence (e.g., through Artificial Intelligence solutions [6]), to achieve autonomy and automation. These changes are expected to shorten the development and manufacturing periods, increase the manufacturing efficiency and sustainability [7], and achieve greater flexibility. Furthermore, emphasis on making such technologies and applications safe, trustworthy, and human-centric is a crucial characteristic of the nascent Industry 5.0 paradigm [8,9].

The increasing digitalization and the democratization of Artificial Intelligence have enabled such models to automate and provide assistance on a wide range of tasks. Machine Learning models can learn from historic data to predict future outcomes (e.g., estimate future demand [10]) or perform certain tasks (e.g., identify manufacturing defects [11] or automate manual tasks with robotic assistance [12]). Nevertheless, it is crucial to establish a synergic and collaborative environment between humans and machines, where humans can flourish in developing their creativity while delegating monotonous and repetitive tasks to machines [9]. Furthermore, such collaboration requires transparency to develop trust in Artificial Intelligence applications. Research exploring the models' rationale behind the predictions and how the insights must be provided to the user as an explanation is known as Explainable Artificial Intelligence.

Insights obtained regarding the models' rationale behind a prediction can be enriched with domain knowledge to provide a better-contextualized explanation to the user [13,14].

To encode domain knowledge, researchers have resorted to using semantic technologies, such as ontologies and knowledge graphs. In order to integrate semantic technologies into the explanation crafting process, the structure of the explanation must be defined, and a procedure established to (a) extract semantic meaning from the explanation, (b) use such information to query external sources (e.g., open knowledge graphs), and (c) enrich the explanation with new information and insights obtained from the external sources.

In this research, we extend the work performed at [15] and [16]. Given a set of demand forecasts, a mapping between features and concepts that define an ontology and a hierarchy of concepts, and a set of keywords associated with each concept, we recourse to a wikification process to extract wiki concepts based on the keywords mentioned above. The wiki concepts are then used to query external sources (e.g., open knowledge graphs or open datasets) and rank query results based on their semantic similarity computing the Jaccard distance. Our approach was tested in the domain of demand forecasting and validated on a real-world case study, using models we developed as part of the European Horizon 2020 projects FACTLOG¹ and STAR². We evaluate our results through two metrics (precision@K and Ratio of Diverse Entries (RDE@K)) to assess how precise and diverse the knowledge fragments are.

The rest of this paper is structured as follows: Section 2 presents related work, Section 3 describes the use case we used and the implementation we followed to test our concept, Section 4 provides the results we obtained and their evaluation. Finally, in Section 5, we provide our conclusions and outline future work.

2. Related Work

2.1. Industry 4.0 and Industry 5.0

Industry 4.0 was coined in 2011 by the German government initiative aimed at developing advanced production systems to increase the productivity and efficiency of the manufacturing industry. The initiative was soon followed by national initiatives across multiple countries, such as the USA, UK, France, Canada, Japan, and China, among others [17–19]. Industry 4.0 was conceived as a technological revolution adding value to the whole manufacturing and product lifecycle. Part of such a revolution is the concept of a smart and integrated supply chain, which aims to reduce delivery times and reduce information distortion across suppliers and manufacturers [20]. The benefits mentioned above are achieved by enhancing the demand forecasting and optimizing the organization and management of materials, suppliers, and inventory [21,22]. To that end, digital twins of existing processes can be created to simulate, test *what-if* scenarios, and enhance them without disruption of the physical operations [23,24].

On top of the aforementioned integrated supply chain, the Industry 4.0 paradigm emphasizes the redesign of the human role in manufacturing, leveraging new technological advancements and capabilities [20]. This aspect is further evolved in Industry 5.0. Industry 5.0 is a value-driven manufacturing paradigm that underscores the relevance of research and innovation to develop human-centric approaches when supporting industry operation [25]. Human-centricity in manufacturing must take into account the skills that are unique to the human workers, such as critical thinking, creativity, and domain knowledge while leveraging machine strengths (e.g., high efficiency in performing repetitive tasks [3,26,27]). Human-centricity can be realized through (a) a systemic approach focused on forging synergic, two-way relationships between humans and machines, (b) the use of digital twins at a systemic level, and (c) the adoption of artificial intelligence at all levels [9,28].

Artificial Intelligence can be essential to achieving human-machine collaboration. In particular, Active Learning and Explainable Artificial Intelligence can be used to complement each other. Active Learning is the sub-field of Artificial Intelligence, concerned with retrieving specific data and leveraging human knowledge to satisfy a particular learning

¹ <https://www.factlog.eu/>

² <http://www.star-ai.eu/>

goal. On the other hand, Explainable Artificial Intelligence is the sub-field of Artificial Intelligence concerned with providing insights into the inner workings of a model regarding its outcome so that the user can learn about its underlying behavior. This way, the human can act as a teacher to Artificial Intelligence models and learn from them through Explainable Artificial Intelligence. This two-way relationship can lead to a trusted collaboration [29,30].

2.2. Demand Forecasting

Demand forecasting is a key component of manufacturing companies. Precise demand forecasts allow to set correct inventory levels, price the products and plan future operations. Any improvements in such forecasts directly translate to the supply chain performance [31]. Demand depends on characteristics that are intrinsic to the product [32] (e.g., elasticity or configuration), and external factors [33] (e.g., particular sales conditions).

Researchers developed multiple schemas to classify demand according to its characteristics. Among them we find the ABC inventory classification system [34], the XYZ analysis [35], and the quadrant proposed by Syntetos et al. [36]. The ABC inventory classification system classifies the items based on their decreasing order of annual dollar volume. The XYZ analysis classifies items according to their consumption patterns: (X) constant consumption, (Y) fluctuating (usually seasonal) consumption, and (Z) irregular consumption. Finally, Syntetos et al. divide demand into four quadrants based on demand size and demand occurrence variability. It has been recognized that Artificial Intelligence models can provide accurate demand forecasts based on past demand and complementary data. Different methods are appropriate for demands with different characteristics [37]. Demand forecasting models are usually framed as time series forecasting problems using supervised regression models or specialized models to learn patterns and forecast future values. Research related to the automotive industry has reported using complementary sources such as unemployment rate [38], inflation rate [39] or Gross Domestic Product [40], and a variety of algorithms, such as the Multiple Linear Regression [41], Support Vector Machine [42], and Neural Networks [43].

The ability to leverage a wide range of complementary data sources is a specific advantage of such models against humans in light of their visual working memory, short-term memory, and capacity to process variables [44]. Furthermore, Artificial Intelligence models avoid multiple cognitive biases to which humans are prone to [45]. Nevertheless, planners must approach such forecasts with critical thinking since they hold responsibility for decisions based on such forecasts. They must understand the models' rationale behind the forecast [46–48], take into account information that could signal adjustments to the forecast are needed and make such adjustments when needed [49–51]. Furthermore, while transparency regarding the models' underlying rationale can be in some cases required by law [52] (e.g., the General Data Protection Regulation(GDPR) [53] or the Artificial Intelligence Act [54]), it also provides a learning opportunity, which is key to employees' engagement [55].

2.3. Explainable Artificial Intelligence

Explainable Artificial Intelligence in a sub-field of Artificial Intelligence research, concerned with how the models' behavioral aspects can be translated into a human interpretable form to understand causality, enhance trustworthiness, and develop confidence [56,57]. Techniques are usually classified based on their complexity (degree of interpretability), their scope (global or local), and whether they are model-agnostic [46]. Regarding their degree of interpretability, models are usually considered black-box (opaque) or white-box (transparent). The source of models' opacity can be due to the inherent properties of the algorithms, their complexity, or due to the explicit requirement to avoid exposing its inner workings (e.g., trade secret), [58,59].

Demand forecasting can be framed as a regression problem, and thus explainability techniques developed for this kind of supervised learning can be used to unveil the models' inner workings. One such technique is LIME [60] (Local Interpretable Model-agnostic

Explanations), which provides a model-agnostic approach to estimate features' relevance for each forecast. LIME creates a linear model to approximate the model's behavior at a particular forecasting point. Then, it estimates the features' relevance by measuring how much predictions change upon features' perturbation. Similar approaches were later developed to ensure explanations were deterministic [61], or take into account non-linear relationships between the features [62,63]. Other frequently cited methods are SHAP [64,65] (SHapley Additive exPlanations, which leverages coalitional game theory to estimate the contributions of individual feature values), LACE [66] (Local Agnostic attribute Contribution Explanation, which leverages SHAP and LIME to provide an explanation through local rules), LoRE [67] (LOcal Rule-based Explanation, which crafts an explanation extracting a decision rule and a set of counterfactual rules), Anchors [68] (which establishes a set of precise rules that explain the forecast at a local level), and Local Foil Trees [69] (which identifies a disjoint set of rules that result in counterfactual outcomes).

While understanding why the model issued a particular forecast is of utmost importance, explanations crafting can be guided using domain knowledge, enriched with complementary data or insights. An approach to enhance explanations given specific domain knowledge was developed by Confalonieri et al. [70], who demonstrated that decision trees built for explainability were more understandable if built considering domain knowledge. Enrichment with complementary data was researched by Panigutti et al. [71], where high fidelity to the forecasting model was achieved by enriching the explainability rules with semantically encoded domain knowledge. Semantic enrichment or recourse to graph representations to bind multiple insights was explored by several authors [14,15,72]. Finally, Rabold et al. [73] devised means to create enhanced, multimodal explanations by relating visual explanations to logic rules obtained through an inductive logic programming system.

While many metrics have been devised to assess the quality of the outcomes produced by explainability techniques [74–77], explanations must also be evaluated and measured from a human-centric perspective. Attention must be devoted to ensuring the explanations convey meaningful information to different user profiles according to their purpose [78–80], that they promote curiosity to increase learning and engagement [81], and provide means to develop trust through exploration [82]. Furthermore, it is desired that the explanations are actionable [83,84], and inform conditions that could change the forecast outcome [85]. The aspects mentioned above are frequently evaluated through qualitative interviews and questionnaires, think-aloud, or self-explanations [81,86]. Methods such as tracking participants' eye movement patterns, measuring users' response time to an explanation, the quantification of answers provided, or the number of explanations required by the user to learn have also been proposed [81,86–88].

3. Use case

This research was developed for a demand forecasting use case. It extends work performed at [15], and [16], taking into account four data sources: (i) data provided by a European original equipment manufacturer targeting the global automotive industry market; (ii) news media entries provided by Event Registry [89], a well-established Media Events Retrieval System that provides real-time mainstream media monitoring; (iii) the EU Open Data Portal [90]; and (iv) the Google Knowledge Graph (Google KG) [91]. We detail the procedure in Fig. 1.

$$JaccardIndex(A, B) = \frac{|A \cap B|}{|A \cup B|} \quad (1)$$

$$JaccardDistance = 1 - JaccardIndex \quad (2)$$

In [15] we studied how explanations given to the user could be enriched with external sources, searching for news media entries and metadata regarding datasets. Searches were performed based on keywords that described features' semantic abstractions and new

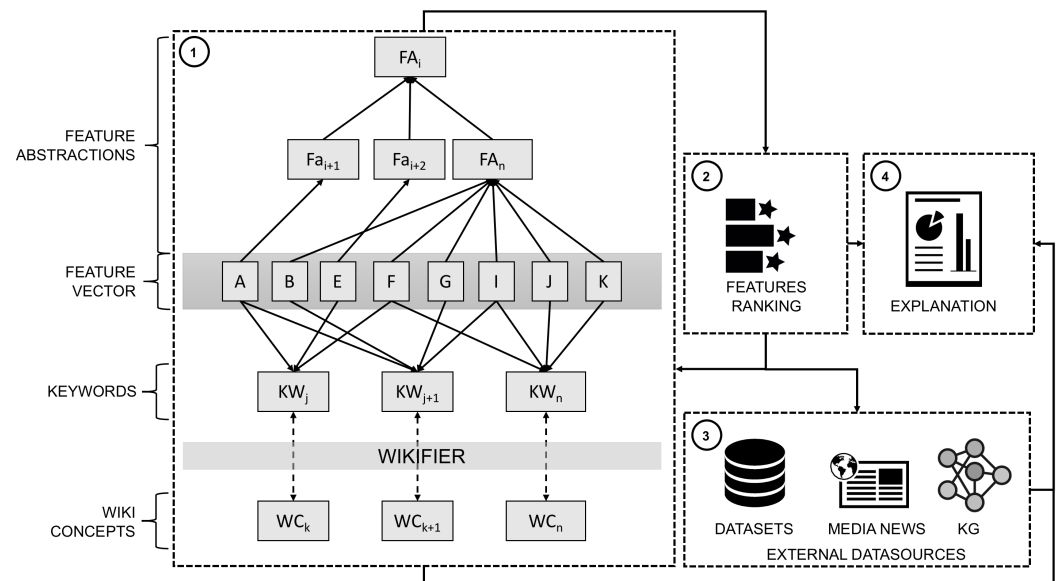


Figure 1. High-level diagram of the components taken into account and the procedure followed to craft the explanations.

keywords identified from retrieved news media entries. The entries were then ranked based on the word movers distance ([92]) between embeddings. In this research, we adopted a different strategy and (i) computed wiki concepts for the features' abstraction keywords (see Table 1), (ii) ranked news media entries based on the Jaccard distance (see Eq. 2) between relevant reference wiki concepts and the ones obtained by wikifying the news media events [93], (iii) queried and ranked external data sources based on reference wiki concepts and the most important concepts that emerged from news media events, and (iv) further enriched the explanations by adding most relevant entries from the Google KG queried with the most relevant wiki concepts obtained from news media events. However, we found that most frequent wiki concepts in media news referred to persons and places, which were not informative to the task at hand. Therefore, we decided to filter them out to enhance the quality of the outcomes.

Feature keywords	Wiki concepts
Car Sales Demand	Car Demand
New Car Sales	Car Sales
Vehicle Sales	Vehicle
Car Demand	Car Demand
Automotive Industry	Automotive Industry
Global GDP Projection	Gross Domestic Product Gross World Product
Global Economic Outlook	Economy World economy
Economic Forecast	Forecasting Economy
Unemployment Rate	Unemployment
Unemployment Numbers	Unemployment
Unemployment Report	Unemployment
Employment Growth	Employment
Long-term Unemployment	Unemployment
Purchasing Managers' Index	Manager (Gaelic games)

Table 1. Mapping between feature keywords and wiki concepts. Wikification was performed invoking the service available at <https://wikifier.org>. Note that while most concepts were accurate, the wikification of the *Purchasing Managers' Index* produced a wrong concept.

4. Evaluation and Results

199

	Metric	Embeddings-based approach	Semantics-based approach
Media Events	Average Precision@1	0,97	0,95
	Average Precision@3	0,97	0,95
	RDE@1	0,30	0,38
	RDE@3	0,11	0,14
Media Events' K&WC	Average Precision@1	0,77	0,71
	Average Precision@3	0,78	0,72
	RDE@1	0,14	0,01
	RDE@3	0,09	0,01
External Datasets	Average Precision@1	0,56	0,68
	RDE@1	0,41	0,43
Google KG	Average Precision@1	NA	0,76
	Average Precision@3	NA	0,46
	RDE@1	NA	0,15
	RDE@3	NA	0,09

Table 2. Results we obtained by analyzing the forecast explanations created for 56 products over three months. The best results are bolded. *Media Events*, *Media Events' Keywords*, *External Datasets* and *Google KG* correspond to contextual information displayed for each forecast explanation. *Media Events' K&WC* stands for Media Events' Keywords and Wiki Concepts.

Our primary interest in this research was to use semantic technologies and external data sources to enrich the explanations with (a) contextual information (e.g., events informed in media news that could explain a given forecast), (b) datasets' metadata that could serve to future model enrichments, and (c) new relevant concepts obtained from integrations with semantic tools and knowledge graphs. We realized (c) for this research by retrieving wiki concepts associated with media news pieces of interest and the results obtained from the Google KG.

$$RDE = \frac{\text{Unique Entries}}{\text{Total Listed Entries}} \quad (3)$$

To evaluate the outcomes of such enrichment, we used two metrics to assess whether the entries were precise and diverse: (a) Average Precision@K, and the Ratio of Diverse Entries (RDE@K) (see Eq. 3). The first metric allowed us to measure how much of the information we displayed was related to the underlying model's features. Given that the users rely on trusted automation [94], we consider that exact results are required to increase the users' trust in the underlying application. Furthermore, considering that curiosity is related to the workers' engagement, we consider a diverse set of entries is preferred to foster it. We quantify such diversity through the RDE@K metric. While the entries do not repeat themselves in a single forecast explanation, nothing prevents having the duplicate entries listed through different forecast explanations. We consider that the best case could avoid repeated entries to maximize users' learning. Nevertheless, different strategies could be adopted, framing the entries display as a recommender system problem and considering users' implicit and explicit feedback to rank and decide whether and where to display them [95–97].

We present our results in Table 2. In the Table we compare results obtained from our previous work (*Embeddings-based approach*) [15] and the *Semantics-based approach* described in Section 3. We found that the former approach had a better performance when considering the media events' diversity, trading a slight decay in precision compared with the embeddings-based approach. Regarding the Media Events' K&WC, the embeddings-based approach achieved better for both diversity and precision. While differences in precision were contested, they were pronounced regarding diversity. We consider it is natural to have a less numerous set of wiki concepts compared to keywords. Nevertheless, we consider that the metric values could be improved. On the other hand, the semantics-based approach was much more precise when recommending datasets and displayed a slightly higher diversity. Finally, when evaluating the entries related to the Google KG, we observed that the diversity was similar to the one obtained for media news with the embeddings-based

approach. The first results were precise, but the precision dropped by 0,30 points when considering $k=3$. We found that those considered erroneous were mostly related to the economy or automotive industry when analyzing the entries. Nevertheless, they were not useful for the explanations at hand since they referenced banks (e.g., Bank of America or Deutsche Bank) or prominent figures (e.g., Edward Fulton Denison, who pioneered the measurement of the United States' Gross Domestic Product, or Kathleen Wilson-Thompson, independent director at Tesla Motors Inc. at the time of this writing). On the other hand, we considered the entries were accurate when they referenced companies from the automotive sector (e.g., Faurecia, Rivian Automotive Inc., Polestar or Vinfast) or related to it (e.g., *Plug Power* which develops hydrogen fuel cell systems to replace conventional batteries or the *Flinkster* carsharing company).

Finally, considering the results regarding Media Events' K&WC and external datasets, we consider that the embeddings-based approach should be used to retrieve new keywords and concepts displayed to the users. In contrast, the semantics-based approach provides the best concepts that lead to better results when searching for external datasets' metadata.

5. Conclusions and Future Work

Along with the increasing adoption of Artificial Intelligence in manufacturing, explainability techniques must be developed to ensure users' can learn the models' behaviors. Furthermore, explanations provided to the user can be enriched with additional insights that foster the users' curiosity, resulting in an exploratory dynamic towards the artificial intelligence application and domain-specific problems and enabling the development of trust towards such applications. One way to achieve this goal is to enrich the explanations with information obtained from external sources to augment users' knowledge and help them make responsible decisions.

This research explored augmenting explanations by incorporating media news, datasets' metadata, and information queried from open knowledge graphs. Furthermore, we compared two approaches (based on embeddings and wiki concepts) to rank the data sources' entries and retrieve new concepts and keywords from them. We found that results were similar when considering media events but differed for new keywords and concepts (embeddings-based approach was best) and external datasets (semantic-based approach was best). Finally, while the integration with the Google KG proved informative, new strategies must be explored to increase the precision of results when showing multiple entries.

In future work, we would like to explore how critical components of an explanation encoded in a graph structure can be used to enhance the explanations and whether users' feedback can lead to better explanations if we frame them as a recommender system problem.

Author Contributions: Conceptualization, J.M.R.; methodology, J.M.R.; software, E.T. and J.M.R.; validation, J.M.R.; formal analysis, J.M.R.; investigation, J.M.R.; resources, J.M.R., P.Z. and E.T.; data curation, J.M.R.; writing—original draft preparation, J.M.R.; writing—review and editing, J.M.R. and K.K.; visualization, J.M.R.; supervision, D.M. and B.F.; project administration, K.K and I.N. ; funding acquisition, K.K, B.F. and D.M. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Slovenian Research Agency and the European Union's Horizon 2020 program projects FACTLOG under grant agreement H2020-869951, and STAR under grant agreement number H2020-956573.

Data Availability Statement: No data was released.

Acknowledgments: This document is the property of the STAR consortium and shall not be distributed or reproduced without the formal approval of the STAR Management Committee. The content of this report reflects only the authors' view. The European Commission is not responsible for any use that may be made of the information it contains.

Abbreviations

The following abbreviations are used in this manuscript:

Google KG	Google Knowledge Graph
LACE	Local Agnostic attribute Contribution Explanation
LIME	Local Interpretable Model-agnostic Explanations
LoRE	LOcal Rule-based Explanation
Media Events' K&WC	Media Events' Keywords and Concepts
RDE	Ratio of Diverse Entries
SHAP	SHapley Additive exPlanations
XAI	Explainable Artificial Intelligence

References

- Lasi, H.; Fettke, P.; Kemper, H.G.; Feld, T.; Hoffmann, M. Industry 4.0. *Business & information systems engineering* **2014**, *6*, 239–242.
- Erro-Garcés, A. Industry 4.0: defining the research agenda. *Benchmarking: an international journal* **2019**.
- Maddikunta, P.K.R.; Pham, Q.V.; Prabadevi, B.; Deepa, N.; Dev, K.; Gadekallu, T.R.; Ruby, R.; Liyanage, M. Industry 5.0: A survey on enabling technologies and potential applications. *Journal of Industrial Information Integration* **2021**, p. 100257.
- Lu, Y. Cyber physical system (CPS)-based industry 4.0: A survey. *Journal of Industrial Integration and Management* **2017**, *2*, 1750014.
- Shafto, M.; Conroy, M.; Doyle, R.; Glaessgen, E.; Kemp, C.; LeMoigne, J.; Wang, L. Draft modeling, simulation, information technology & processing roadmap. *Technology Area* **2010**, *11*.
- Arinez, J.F.; Chang, Q.; Gao, R.X.; Xu, C.; Zhang, J. Artificial intelligence in advanced manufacturing: Current status and future outlook. *Journal of Manufacturing Science and Engineering* **2020**, *142*.
- Ghobakhloo, M. Industry 4.0, digitization, and opportunities for sustainability. *Journal of cleaner production* **2020**, *252*, 119869.
- Martynov, V.V.; Shavaleeva, D.N.; Zaytseva, A.A. Information technology as the basis for transformation into a digital society and industry 5.0. In Proceedings of the 2019 International Conference "Quality Management, Transport and Information Security, Information Technologies"(IT&QM&IS). IEEE, 2019, pp. 539–543.
- Rožanec, J.M.; Novalija, I.; Zajec, P.; Kenda, K.; Tavakoli, H.; Suh, S.; Veliou, E.; Papamartzivanos, D.; Giannetsos, T.; Menesidou, S.A.; et al. Human-Centric Artificial Intelligence Architecture for Industry 5.0 Applications. *arXiv preprint arXiv:2203.10794* **2022**.
- Rožanec, J.M.; Kažič, B.; Škrjanc, M.; Fortuna, B.; Mladenčić, D. Automotive OEM demand forecasting: A comparative study of forecasting algorithms and strategies. *Applied Sciences* **2021**, *11*, 6787.
- Trajkova, E.; Rožanec, J.M.; Dam, P.; Fortuna, B.; Mladenčić, D. Active Learning for Automated Visual Inspection of Manufactured Products. *arXiv preprint arXiv:2109.02469* **2021**.
- Bhatt, P.M.; Malhan, R.K.; Shembekar, A.V.; Yoon, Y.J.; Gupta, S.K. Expanding capabilities of additive manufacturing through use of robotics technologies: A survey. *Additive manufacturing* **2020**, *31*, 100933.
- Dhanorkar, S.; Wolf, C.T.; Qian, K.; Xu, A.; Popa, L.; Li, Y. Who needs to know what, when?: Broadening the Explainable AI (XAI) Design Space by Looking at Explanations Across the AI Lifecycle. In Proceedings of the Designing Interactive Systems Conference 2021, 2021, pp. 1591–1602.
- Dragoni, M.; Donadello, I. A Knowledge-based Strategy For XAI: The Explanation Graph.
- Rožanec, J.M.; Fortuna, B.; Mladenčić, D. Knowledge graph-based rich and confidentiality preserving Explainable Artificial Intelligence (XAI). *Information Fusion* **2022**, *81*, 91–102.
- Rožanec, J.M.; Zajec, P.; Kenda, K.; Novalija, I.; Fortuna, B.; Mladenčić, D. XAI-KG: knowledge graph to support XAI and decision-making in manufacturing. In Proceedings of the International Conference on Advanced Information Systems Engineering, Springer, 2021, pp. 167–172.
- Majstorovic, V.D.; Mitrovic, R. Industry 4.0 programs worldwide. In Proceedings of the International Conference on the Industry 4.0 model for Advanced Manufacturing. Springer, 2019, pp. 78–99.
- Bogoviz, A.V.; Osipov, V.S.; Chistyakova, M.K.; Borisov, M.Y. Comparative analysis of formation of industry 4.0 in developed and developing countries. In *Industry 4.0: Industrial Revolution of the 21st Century*; Springer, 2019; pp. 155–164.
- Raj, A.; Dwivedi, G.; Sharma, A.; de Sousa Jabbour, A.B.L.; Rajak, S. Barriers to the adoption of industry 4.0 technologies in the manufacturing sector: An inter-country comparative perspective. *International Journal of Production Economics* **2020**, *224*, 107546.
- Frank, A.G.; Dalenogare, L.S.; Ayala, N.F. Industry 4.0 technologies: Implementation patterns in manufacturing companies. *International Journal of Production Economics* **2019**, *210*, 15–26.
- Ghobakhloo, M. The future of manufacturing industry: a strategic roadmap toward Industry 4.0. *Journal of Manufacturing Technology Management* **2018**.
- Zheng, T.; Ardolino, M.; Bacchetti, A.; Perona, M. The applications of Industry 4.0 technologies in manufacturing context: a systematic literature review. *International Journal of Production Research* **2021**, *59*, 1922–1954.
- Qi, Q.; Tao, F.; Zuo, Y.; Zhao, D. Digital twin service towards smart manufacturing. *Procedia Cirp* **2018**, *72*, 237–242.
- Rožanec, J.M.; Lu, J.; Rupnik, J.; Škrjanc, M.; Mladenčić, D.; Fortuna, B.; Zheng, X.; Kiritsis, D. Actionable cognitive twins for decision making in manufacturing. *International Journal of Production Research* **2022**, *60*, 452–478.

25. Xu, X.; Lu, Y.; Vogel-Heuser, B.; Wang, L. Industry 4.0 and Industry 5.0—Inception, conception and perception. *Journal of Manufacturing Systems* **2021**, *61*, 530–535. 334
26. Nahavandi, S. Industry 5.0—A human-centric solution. *Sustainability* **2019**, *11*, 4371. 335
27. Demir, K.A.; Döven, G.; Sezen, B. Industry 5.0 and human-robot co-working. *Procedia computer science* **2019**, *158*, 688–695. 337
28. Industry 5.0: Towards more sustainable, resilient and human-centric industry. <https://op.europa.eu/en/publication-detail/-/publication/468a892a-5097-11eb-b59f-01aa75ed71a1/>. Accessed: 2022-03-15. 338
29. Weitz, K.; Schiller, D.; Schlagowski, R.; Huber, T.; André, E. "Do you trust me?" Increasing user-trust by integrating virtual agents in explainable AI interaction design. In Proceedings of the Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents, 2019, pp. 7–9. 340
30. Honeycutt, D.; Nourani, M.; Ragan, E. Soliciting human-in-the-loop user feedback for interactive machine learning reduces user trust and impressions of model accuracy. In Proceedings of the Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, 2020, Vol. 8, pp. 63–72. 343
31. Moroff, N.U.; Kurt, E.; Kamphues, J. Machine Learning and statistics: A Study for assessing innovative demand forecasting models. *Procedia Computer Science* **2021**, *180*, 40–49. 346
32. Purohit, D.; Srivastava, J. Effect of manufacturer reputation, retailer reputation, and product warranty on consumer judgments of product quality: A cue diagnosticity framework. *Journal of consumer psychology* **2001**, *10*, 123–134. 348
33. Callon, M.; Méadel, C.; Rabeharisoa, V. The economy of qualities. *Economy and society* **2002**, *31*, 194–217. 350
34. Teunter, R.H.; Babai, M.Z.; Syntetos, A.A. ABC classification: service levels and inventory costs. *Production and Operations Management* **2010**, *19*, 343–352. 351
35. Scholz-Reiter, B.; Heger, J.; Meinecke, C.; Bergmann, J. Integration of demand forecasts in ABC-XYZ analysis: practical investigation at an industrial company. *International Journal of Productivity and Performance Management* **2012**. 353
36. Syntetos, A.A.; Boylan, J.E.; Croston, J. On the categorization of demand patterns. *Journal of the operational research society* **2005**, *56*, 495–503. 355
37. Rožanec, J.M.; Mladenčić, D. Reframing demand forecasting: a two-fold approach for lumpy and intermittent demand. *arXiv preprint arXiv:2103.13812* **2021**. 357
38. Brühl, B.; Hülsmann, M.; Borscheid, D.; Friedrich, C.M.; Reith, D. A sales forecast model for the german automobile market based on time series analysis and data mining methods. In Proceedings of the Industrial Conference on Data Mining. Springer, 2009, pp. 146–160. 359
39. Vahabi, A.; Hosseininia, S.S.; Alborzi, M. A Sales Forecasting Model in Automotive Industry using Adaptive Neuro-Fuzzy Inference System (Anfis) and Genetic Algorithm (GA). *management* **2016**, *1*, 2. 362
40. Ubaidillah, N.Z. A STUDY OF CAR DEMAND AND ITS INTERDEPENDENCY IN SARAWAK. *International Journal of Business & Society* **2020**, *21*. 364
41. Dwivedi, A.; Niranjana, M.; Sahu, K. A business intelligence technique for forecasting the automobile sales using Adaptive Intelligent Systems (ANFIS and ANN). *International Journal of Computer Applications* **2013**, *74*. 366
42. Wang, X.; Zeng, D.; Dai, H.; Zhu, Y. Making the right business decision: Forecasting the binary NPD strategy in Chinese automotive industry with machine learning methods. *Technological Forecasting and Social Change* **2020**, *155*, 120032. 368
43. Chandriah, K.K.; Naraganahalli, R.V. RNN/LSTM with modified Adam optimizer in deep learning approach for automobile spare parts demand forecasting. *Multimedia Tools and Applications* **2021**, pp. 1–15. 370
44. Halford, G.S.; Baker, R.; McCredden, J.E.; Bain, J.D. How many variables can humans process? *Psychological science* **2005**, *16*, 70–76. 372
45. Barnes Jr, J.H. Cognitive biases and their impact on strategic planning. *Strategic Management Journal* **1984**, *5*, 129–137. 373
46. Adadi, A.; Berrada, M. Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE access* **2018**, *6*, 52138–52160. 374
47. Arrieta, A.B.; Díaz-Rodríguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; García, S.; Gil-López, S.; Molina, D.; Benjamins, R.; et al. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion* **2020**, *58*, 82–115. 376
48. Confalonieri, R.; Coda, L.; Wagner, B.; Besold, T.R. A historical perspective of explainable Artificial Intelligence. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **2021**, *11*, e1391. 378
49. Davydenko, A.; Fildes, R.A.; Trapero Arenas, J. Judgmental adjustments to demand forecasts: Accuracy evaluation and bias correction **2010**. 381
50. Davydenko, A.; Fildes, R. Measuring forecasting accuracy: The case of judgmental adjustments to SKU-level demand forecasts. *International Journal of Forecasting* **2013**, *29*, 510–522. 383
51. Alvarado-Valencia, J.; Barrero, L.H.; Önköl, D.; Dennerlein, J.T. Expertise, credibility of system forecasts and integration methods in judgmental demand forecasting. *International Journal of Forecasting* **2017**, *33*, 298–313. 385
52. Goodman, B.; Flaxman, S. European Union regulations on algorithmic decision-making and a "right to explanation". *AI magazine* **2017**, *38*, 50–57. 387
53. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). <https://eur-lex.europa.eu/eli/reg/2016/679/oj>. Accessed: 2022-03-15. 389

-
54. Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>. Accessed: 2022-03-15.
55. Anitha, J. Determinants of employee engagement and their impact on employee performance. *International journal of productivity and performance management* **2014**.
56. Emmert-Streib, F.; Yli-Harja, O.; Dehmer, M. Explainable artificial intelligence and machine learning: A reality rooted perspective. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **2020**, *10*, e1368.
57. Schwalbe, G.; Finzel, B. XAI Method Properties: A (Meta-) study. *arXiv preprint arXiv:2105.07190* **2021**.
58. Chan, L. Explainable AI as Epistemic Representation. *Overcoming Opacity in Machine Learning* **2021**, p. 7.
59. Müller, V.C. Deep Opacity Undermines Data Protection and Explainable Artificial Intelligence. *Overcoming Opacity in Machine Learning* **2021**, p. 18.
60. Ribeiro, M.T.; Singh, S.; Guestrin, C. "Why should i trust you?" Explaining the predictions of any classifier. In Proceedings of the Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, 2016, pp. 1135–1144.
61. Zafar, M.R.; Khan, N.M. DLIME: A deterministic local interpretable model-agnostic explanations approach for computer-aided diagnosis systems. *arXiv preprint arXiv:1906.10263* **2019**.
62. Hall, P.; Gill, N.; Kurka, M.; Phan, W. Machine learning interpretability with h2o driverless ai. *H2O. ai*. URL: <http://docs.h2o.ai/driverless-ai/latest-stable/docs/booklets/MLIBooklet.pdf> **2017**.
63. Sokol, K.; Flach, P. LIMETree: Interactively Customisable Explanations Based on Local Surrogate Multi-output Regression Trees. *arXiv preprint arXiv:2005.01427* **2020**.
64. Lundberg, S.M.; Lee, S.I. A unified approach to interpreting model predictions. *Advances in neural information processing systems* **2017**, *30*.
65. Strumbelj, E.; Kononenko, I. An efficient explanation of individual classifications using game theory. *The Journal of Machine Learning Research* **2010**, *11*, 1–18.
66. Pastor, E.; Baralis, E. Explaining black box models by means of local rules. In Proceedings of the Proceedings of the 34th ACM/SIGAPP symposium on applied computing, 2019, pp. 510–517.
67. Guidotti, R.; Monreale, A.; Ruggieri, S.; Pedreschi, D.; Turini, F.; Giannotti, F. Local rule-based explanations of black box decision systems. *arXiv preprint arXiv:1805.10820* **2018**.
68. Ribeiro, M.T.; Singh, S.; Guestrin, C. Anchors: High-Precision Model-Agnostic Explanations. In Proceedings of the AAAI, 2018, Vol. 18, pp. 1527–1535.
69. van der Waa, J.; Robeer, M.; van Diggelen, J.; Brinkhuis, M.; Neerinx, M. Contrastive explanations with local foil trees. *arXiv preprint arXiv:1806.07470* **2018**.
70. Confalonieria, R.; Galliana, P.; Kutza, O.; Porellob, D.; Righetia, G.; Troquarda, N. Towards Knowledge-driven Distillation and Explanation of Black-box Models. In Proceedings of the International Workshop on Data meets Applied Ontologies in Explainable AI (DAO-XAI 2021), 2021.
71. Panigutti, C.; Perotti, A.; Pedreschi, D. Doctor XAI: an ontology-based approach to black-box sequential data classification explanations. In Proceedings of the Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 2020, pp. 629–639.
72. Lécué, F.; Abeloos, B.; Anctil, J.; Bergeron, M.; Dalla-Rosa, D.; Corbeil-Letourneau, S.; Martet, F.; Pommellet, T.; Salvan, L.; Veilleux, S.; et al. Thales XAI Platform: Adaptable Explanation of Machine Learning Systems-A Knowledge Graphs Perspective. In Proceedings of the ISWC Satellites, 2019, pp. 315–316.
73. Rabold, J.; Deininger, H.; Siebers, M.; Schmid, U. Enriching visual with verbal explanations for relational concepts—combining LIME with Aleph. In Proceedings of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases. Springer, 2019, pp. 180–192.
74. Lakkaraju, H.; Kamar, E.; Caruana, R.; Leskovec, J. Interpretable & explorable approximations of black box models. *arXiv preprint arXiv:1707.01154* **2017**.
75. Nguyen, A.p.; Martínez, M.R. On quantitative aspects of model interpretability. *arXiv preprint arXiv:2007.07584* **2020**.
76. Rosenfeld, A. Better metrics for evaluating explainable artificial intelligence. In Proceedings of the Proceedings of the 20th international conference on autonomous agents and multiagent systems, 2021, pp. 45–50.
77. Amparore, E.; Perotti, A.; Bajardi, P. To trust or not to trust an explanation: using LEAF to evaluate local linear XAI methods. *PeerJ Computer Science* **2021**, *7*, e479.
78. Samek, W.; Müller, K.R. Towards explainable artificial intelligence. In *Explainable AI: interpreting, explaining and visualizing deep learning*; Springer, 2019; pp. 5–22.
79. Pedreschi, D.; Giannotti, F.; Guidotti, R.; Monreale, A.; Pappalardo, L.; Ruggieri, S.; Turini, F. Open the black box data-driven explanation of black box decision systems. *arXiv preprint arXiv:1806.09936* **2018**.
80. El-Assady, M.; Jentner, W.; Kehlbeck, R.; Schlegel, U.; Sevastianova, R.; Sperrle, F.; Spinner, T.; Keim, D. Towards XAI: Structuring the Processes of Explanations. In Proceedings of the ACM Workshop on Human-Centered Machine Learning, 2019.
81. Hsiao, J.H.w.; Ngai, H.H.T.; Qiu, L.; Yang, Y.; Cao, C.C. Roadmap of designing cognitive metrics for explainable artificial intelligence (XAI). *arXiv preprint arXiv:2108.01737* **2021**.

-
82. Hoffman, R.R.; Mueller, S.T.; Klein, G.; Litman, J. Metrics for explainable AI: Challenges and prospects. *arXiv preprint arXiv:1812.04608* **2018**. 450 451
 83. Keane, M.T.; Smyth, B. Good counterfactuals and where to find them: A case-based technique for generating counterfactuals for explainable AI (XAI). In Proceedings of the International Conference on Case-Based Reasoning, Springer, 2020, pp. 163–178. 452 453
 84. Keane, M.T.; Kenny, E.M.; Delaney, E.; Smyth, B. If only we had better counterfactual explanations: five key deficits to rectify in the evaluation of counterfactual XAI techniques. *arXiv preprint arXiv:2103.01035* **2021**. 454 455
 85. Verma, S.; Dickerson, J.; Hines, K. Counterfactual Explanations for Machine Learning: A Review. *arXiv preprint arXiv:2010.10596* **2020**. 456 457
 86. Mohseni, S.; Zarei, N.; Ragan, E.D. A multidisciplinary survey and framework for design and evaluation of explainable AI systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* **2021**, *11*, 1–45. 458 459
 87. Lage, I.; Ross, A.S.; Kim, B.; Gershman, S.J.; Doshi-Velez, F. Human-in-the-loop interpretability prior. *arXiv preprint arXiv:1805.11571* **2018**. 460 461
 88. Sovrano, F.; Vitali, F. An Objective Metric for Explainable AI: How and Why to Estimate the Degree of Explainability. *arXiv preprint arXiv:2109.05327* **2021**. 462 463
 89. Leban, G.; Fortuna, B.; Brank, J.; Grobelnik, M. Event registry: learning about world events from news. In Proceedings of the Proceedings of the 23rd International Conference on World Wide Web, 2014, pp. 107–110. 464 465
 90. of the European Union, P.O. EU Open Data Portal: The official portal for European data. <https://data.europa.eu>, 2020. Accessed: 2020-12-15. 466 467
 91. Noy, N.; Gao, Y.; Jain, A.; Narayanan, A.; Patterson, A.; Taylor, J. Industry-scale knowledge graphs: Lessons and challenges. *Queue* **2019**, *17*, 48–75. 468 469
 92. Kusner, M.; Sun, Y.; Kolkin, N.; Weinberger, K. From word embeddings to document distances. In Proceedings of the International conference on machine learning, 2015, pp. 957–966. 470 471
 93. Brank, J.; Leban, G.; Grobelnik, M. Annotating documents with relevant wikipedia concepts. *Proceedings of SiKDD* **2017**. 472
 94. Lee, J.D.; See, K.A. Trust in automation: Designing for appropriate reliance. *Human factors* **2004**, *46*, 50–80. 473
 95. Kilani, Y.; Alhijawi, B.; Alsarhan, A. Using artificial intelligence techniques in collaborative filtering recommender systems: Survey. *International Journal of Advanced Intelligence Paradigms* **2018**, *11*, 378–396. 474 475
 96. Karimi, M.; Jannach, D.; Jugovac, M. News recommender systems—Survey and roads ahead. *Information Processing & Management* **2018**, *54*, 1203–1227. 476 477
 97. Sidana, S.; Trofimov, M.; Horodnytskyi, O.; Laclau, C.; Maximov, Y.; Amini, M.R. User preference and embedding learning with implicit feedback for recommender systems. *Data Mining and Knowledge Discovery* **2021**, *35*, 568–592. 478 479