



Article

Enriching Artificial Intelligence Explanations with Knowledge Fragments

Jože Rožanec ^{1,2,3,*}, Elena Trajkova ⁴, Inna Novalija ², Patrik Zajec ^{1,2}, Klemen Kenda ^{1,2,3}, Blaž Fortuna ³ and Dunja Mladenčić ²

¹ Jožef Stefan International Postgraduate School, Jamova 39, 1000 Ljubljana, Slovenia; patrik.zajec@ijs.si (P.Z.); klemen.kenda@ijs.si (K.K.)

² Jožef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia; inna.koval@ijs.si (I.N.); dunja.mladenic@ijs.si (D.M.)

³ Qlector d.o.o., Rovšnikova 7, 1000 Ljubljana, Slovenia; blaz.fortuna@qlector.com

⁴ Faculty of Electrical Engineering, University of Ljubljana, Tržaška c. 25, 1000 Ljubljana, Slovenia; trajkova.elena.00@gmail.com

* Correspondence: joze.rozanec@ijs.si

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 the 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. The embeddings-based approach measures the similarity between relevant concepts and retrieved media news entries and datasets' metadata based on the word movers' distance between embeddings. The semantic-based approach recourses to wikification and measures the Jaccard distance instead. The semantic-based approach leads to more diverse entries when displaying media events and more precise and diverse results regarding recommended datasets. We conclude that the explanations provided can be further improved with information regarding the purpose of potential actions that can be taken to influence demand and to provide "what-if" analysis capabilities.

Keywords: explainable artificial intelligence; human-centric artificial intelligence; smart manufacturing; demand forecasting; Industry 4.0; Industry 5.0



Citation: Rožanec, J.; Trajkova, E.; Novalija, I.; Zajec, P.; Kenda, K.; Fortuna, B.; Mladenčić, D. Enriching Artificial Intelligence Explanations with Knowledge Fragments. *Future Internet* **2022**, *14*, 134. <https://doi.org/10.3390/fi14050134>

Academic Editors: Massimo Stella and Giulio Rossetti

Received: 7 April 2022

Accepted: 26 April 2022

Published: 29 April 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

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 have allowed to achieve increasing manufacturing efficiencies and have 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 by [15,16]. 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 FACTLOG (<https://www.factlog.eu/>, (accessed on 27 April 2022)) and STAR (<http://www.star-ai.eu/>, (accessed on 27 April 2022)). We evaluate our results through three metrics (precision@K and Ratio of Diverse Entries (RDE@K), and coverage regarding five questions proposed in [17]) 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, and 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 [18–20]. 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 [21]. The benefits mentioned above are achieved by enhancing the demand forecasting and optimizing the organization and management of materials, suppliers, and inventory [22,23]. 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 [24,25].

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 [21]. 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 [26]. 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 [3,27,28]). 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 [9,29].

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 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 [30,31].

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 [32]. Demand depends on characteristics that are intrinsic to the product [33] (e.g., elasticity or configuration), and external factors [34] (e.g., particular sales conditions).

Researchers have developed multiple schemas to classify demand according to its characteristics. Among them, we find the ABC inventory classification system [35], the XYZ analysis [36], and the quadrant proposed by Syntetos et al. [37]. 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 was 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 [38]. 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 the unemployment rate [39], inflation rate [40] or gross domestic product [41], and a variety of algorithms, such as multiple linear regression [42], support vector machine [43], and neural networks [44].

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 [45]. Furthermore, artificial intelligence models avoid multiple cognitive biases to which humans are prone to [46]. 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 [47–49], take into account information that could signal that adjustments to the forecast are needed and make such adjustments when needed [50–52]. Furthermore, while transparency regarding the models' underlying rationale can be, in some cases, required by law [53] (e.g., the general data protection regulation (GDPR) [54] or the artificial intelligence act [55]), it also provides a learning opportunity, which is key to the employees' engagement [56].

2.3. Explainable Artificial Intelligence

Explainable artificial intelligence is 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 [57,58]. Techniques are usually classified based on their complexity (degree of interpretability), their scope (global or local), and whether they are model-agnostic [47]. 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), [59,60].

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 the *Local Interpretable Model-agnostic Explanations* (LIME) [61], which provide 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 that explanations were deterministic [62], or take into account non-linear relationships between the features [63,64]. Other frequently cited methods are the *Shapley additive explanations* (SHAP, which leverages coalitional game theory to estimate the contributions of individual feature values) [65,66], *Local Agnostic attribute Contribution Explanation* (LACE, which leverages SHAP and LIME to provide an explanation through local rules) [67], *Local Rule-based Explanation* (LoRE, which crafts an explanation extracting a decision rule and a set of counterfactual rules) [68], *anchors* [69] (which establishes a set of precise rules that explain the forecast at a local level), and *local foil trees* [70] (which identifies a disjoint set of rules that result in counterfactual outcomes).

Understanding why the model issued a particular forecast is of utmost importance [71]. Explanation 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. [72], 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. [73], 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 [14,15,74]. Finally, Rabold et al. [75] 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 [76–79], 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 [80–82] such that they promote curiosity to increase learning and engagement [83], and provide means to develop trust through exploration [84]. Furthermore, it is desired that the explanations are actionable [85,86], and inform conditions that could change the forecast outcome [87]. The aspects mentioned above are frequently evaluated through qualitative interviews and questionnaires, think-aloud, or self-explanations [83,88]. 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 [17,83,88,89].

3. Use Case

This research was developed for a demand forecasting use case, exploring how to provide comprehensive forecast explanations. In particular, we consider the most relevant features regarding a particular forecast and enrich it based on domain knowledge and

information retrieved from external sources. The explanations aim to (a) provide context regarding real-world events that could describe feature values observed for a particular forecast, (b) assist in exploring datasets that can be used to enrich the dataset used to train the model, and (c) discover new relevant concepts by integrating with knowledge graphs. Knowledge regarding which factors drive a particular forecast and real-world events provides the user with valuable context for responsible decision making. Exploring complementary datasets is valuable to data scientists and machine learning engineers, who can iterate upon the existing model and enhance it. Furthermore, such information can also be valuable to other profiles, widening the understanding of forces driving the forecast or the influence of such particular demand. Such understanding can also be widened with insights obtained from knowledge graphs.

This research extends work performed in [15,16,90], 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 [91], a well-established media events retrieval system that provides real-time mainstream media monitoring; (iii) the EU Open Data Portal [92]; and (iv) the Google Knowledge Graph (Google KG) [93]. We detail the procedure in Figure 1.

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

$$JaccardDistance = 1 - JaccardIndex \tag{2}$$

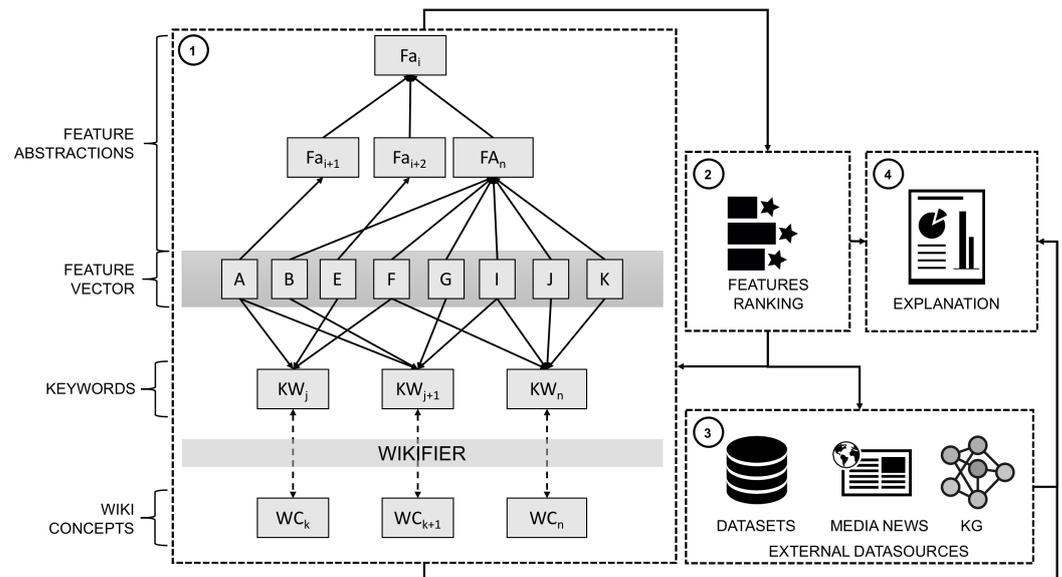


Figure 1. High-level diagram of the components taken into account and the procedure followed to craft the explanations. (1) Given a feature vector, feature abstractions are created, to encode their semantic mining at higher levels. At the same time, keywords are associated to them, whose wiki-concept correspondence is obtained using a wikifier. (2) Given a feature ranking for a particular prediction, (3) complementary data is searched in external datasources to enrich the explanation. (4) Finally, an explanation is assembled based on relevant information gathered concerning the forecast.

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 keywords identified from retrieved news media entries. The entries were then ranked based on the word movers distance ([94]) 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 Equation (2))

between relevant reference wiki concepts and the ones obtained by wikifying the news media events [95], (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 (Media Events' Keywords and Wiki Concepts (Media Events' K&WC)). However, we found that the 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.

Table 1. Mapping between feature keywords and wiki concepts. Wikification was performed invoking the service available at <https://wikifier.org> (accessed on 27 April 2022). Note that while most concepts were accurate, the wikification of the *Purchasing Managers' Index* produced a wrong concept.

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 Gross Domestic Product 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)

4. Evaluation and Results

Our primary interest in this research was to use semantic technologies and external data sources to enrich the explanations. We consider that explanations can be enriched with (a) information contextual to existing domain knowledge, (b) metadata regarding datasets, which can be used to enrich future models, and (c) relevant concepts obtained by integrating with semantic tools and knowledge graphs. We realized (a) by querying events informed in media news, which could potentially explain a given forecast. Furthermore, we realized (b) by querying open dataset portals, and (c) 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 RDE@K (see Equation (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 [96], 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. One such strategy could be to frame the entries displayed as a recommender system problem and consider users' implicit and explicit feedback to rank and decide whether and where to display them [97–99].

We present our results in Table 2. We compare the results obtained in this research (the *semantics-based approach* described in Section 3) with those obtained in our previous work (*embeddings-based approach* [15]). We found that the former approach had 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 performance 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 $k = 3$. Furthermore, 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 that 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).

Considering the results regarding Media Events' K&WC and external datasets, we consider that the embeddings-based approach is best to obtain new keywords and concepts later 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.

Finally, based on the work of [17], we evaluated the goodness of the overall explanations provided to the user based on a score similar to the degree of explainability, computed as the coverage of five questions related to the main archetypes from the abstract meaning representation theory [100]: (i) Why? (ii) How? (iii) What for? (iv) What if? (v) When? Our explanations achieved a score of 0.6, providing information regarding (i) the factors leading to such a forecast (*why is such a forecast issued?*), (ii) how such an outcome can be changed (relevant, actionable aspect), and (iii) when will the forecast demand take place and when did the relevant events influencing the demand occur. The explanations can be improved with information regarding the purpose of the action (increasing or decreasing the demand) and providing capabilities to explore "what-if" scenarios.

Table 2. Results we obtained by analyzing the forecast explanations created for 56 products over three months. The best results are shown in bold. *Media events*, *Media events' keywords*, *external datasets* and *Google KG* correspond to contextual information displayed for each forecast explanation. *NA* is used for entries regarding the results concerning the Google KG for the embeddings-based approach, given that the research we compare against did not provide such results.

	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

5. Conclusions and Future Work

Along with the increasing adoption of artificial intelligence in manufacturing, explainability techniques must be developed to ensure that 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 toward artificial intelligence application and domain-specific problems, enabling the development of trust toward 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). We found that the embeddings-based and semantic-based approaches could complement each other regarding keywords and concepts extracted from media events. While the embeddings-based approach resulted in better scores regarding keywords and concepts shown to the user, the semantics-based approach led to better results when searching for datasets' metadata. 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.

We consider that such an approach toward building explanations can be extended to other use cases. Among its strengths, we consider the capability to explore real-world events impacting the forecasting model and the ability to identify complementary data that can be used to enhance the forecasting model. On the other hand, the approach requires manually associating each feature with higher-level concepts and a set of keywords characterizing them. This requires that the features used to train the model have a meaning intelligible to

the human—a requirement always satisfied by handcrafted features. Another limitation is the need to specify a time frame within which feature-related events are expected to happen, given the features' time of occurrence.

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.R.; methodology, J.R.; software, E.T. and J.R.; validation, J.R.; formal analysis, J.R.; investigation, J.R.; resources, J.R., P.Z. and E.T.; data curation, J.R.; writing—original draft preparation, J.R.; writing—review & editing, J.R. and K.K.; visualization, J.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 were 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.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

GDPR	General Data Protection Regulation
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

References

- Lasi, H.; Fettke, P.; Kemper, H.G.; Feld, T.; Hoffmann, M. Industry 4.0. *Bus. Inf. Syst. Eng.* **2014**, *6*, 239–242. [[CrossRef](#)]
- Erro-Garcés, A. Industry 4.0: defining the research agenda. *Benchmarking Int. J.* **2019**, *28*, 1858–1882. [[CrossRef](#)]
- 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. *J. Ind. Inf. Integr.* **2021**, *26*, 100257. [[CrossRef](#)]
- Lu, Y. Cyber physical system (CPS)-based industry 4.0: A survey. *J. Ind. Integr. Manag.* **2017**, *2*, 1750014. [[CrossRef](#)]
- Shafto, M.; Conroy, M.; Doyle, R.; Glaessgen, E.; Kemp, C.; LeMoigne, J.; Wang, L. Draft modeling, simulation, information technology & processing roadmap. *Technol. Area* **2012**, *32*, 1–38.
- Arinez, J.F.; Chang, Q.; Gao, R.X.; Xu, C.; Zhang, J. Artificial intelligence in advanced manufacturing: Current status and future outlook. *J. Manuf. Sci. Eng.* **2020**, *142*, 110804. [[CrossRef](#)]
- Ghobakhloo, M. Industry 4.0, digitization, and opportunities for sustainability. *J. Clean. Prod.* **2020**, *252*, 119869. [[CrossRef](#)]
- 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), Sochi, Russia, 23–27 September 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* **2022**, arXiv:2203.10794.
- Rožanec, J.M.; Kažič, B.; Škrjanc, M.; Fortuna, B.; Mladenčić, D. Automotive OEM demand forecasting: A comparative study of forecasting algorithms and strategies. *Appl. Sci.* **2021**, *11*, 6787. [[CrossRef](#)]
- Trajkova, E.; Rožanec, J.M.; Dam, P.; Fortuna, B.; Mladenčić, D. Active Learning for Automated Visual Inspection of Manufactured Products. *arXiv* **2021**, arXiv:2109.02469.
- 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. *Addit. Manuf.* **2020**, *31*, 100933. [[CrossRef](#)]

13. 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, Virtual Event, 28 June–2 July 2021; pp. 1591–1602.
14. Dragoni, M.; Donadello, I. *A Knowledge-Based Strategy for XAI: The Explanation Graph*; IOS Press: Amsterdam, The Netherlands, 2022.
15. Rožanec, J.M.; Fortuna, B.; Mladenčić, D. Knowledge graph-based rich and confidentiality preserving Explainable Artificial Intelligence (XAI). *Inf. Fusion* **2022**, *81*, 91–102. [CrossRef]
16. 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: Berlin/Heidelberg, Germany, 2021; pp. 167–172.
17. Sovrano, F.; Vitali, F. An Objective Metric for Explainable AI: How and Why to Estimate the Degree of Explainability. *arXiv* **2021**, arXiv:2109.05327.
18. 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: Berlin/Heidelberg, Germany, 2019; pp. 78–99.
19. 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: Berlin/Heidelberg, Germany, 2019; pp. 155–164.
20. 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. *Int. J. Prod. Econ.* **2020**, *224*, 107546. [CrossRef]
21. Frank, A.G.; Dalenogare, L.S.; Ayala, N.F. Industry 4.0 technologies: Implementation patterns in manufacturing companies. *Int. J. Prod. Econ.* **2019**, *210*, 15–26. [CrossRef]
22. Ghobakhloo, M. The future of manufacturing industry: a strategic roadmap toward Industry 4.0. *J. Manuf. Technol. Manag.* **2018**, *29*, 910–936. [CrossRef]
23. Zheng, T.; Ardolino, M.; Bacchetti, A.; Perona, M. The applications of Industry 4.0 technologies in manufacturing context: A systematic literature review. *Int. J. Prod. Res.* **2021**, *59*, 1922–1954. [CrossRef]
24. Qi, Q.; Tao, F.; Zuo, Y.; Zhao, D. Digital twin service towards smart manufacturing. *Procedia Cirp* **2018**, *72*, 237–242. [CrossRef]
25. 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. *Int. J. Prod. Res.* **2022**, *60*, 452–478. [CrossRef]
26. Xu, X.; Lu, Y.; Vogel-Heuser, B.; Wang, L. Industry 4.0 and Industry 5.0—Inception, conception and perception. *J. Manuf. Syst.* **2021**, *61*, 530–535. [CrossRef]
27. Nahavandi, S. Industry 5.0—A human-centric solution. *Sustainability* **2019**, *11*, 4371. [CrossRef]
28. Demir, K.A.; Döven, G.; Sezen, B. Industry 5.0 and human-robot co-working. *Procedia Comput. Sci.* **2019**, *158*, 688–695. [CrossRef]
29. Industry 5.0: Towards More Sustainable, Resilient and Human-Centric Industry. Available online: <https://op.europa.eu/en/publication-detail/-/publication/468a892a-5097-11eb-b59f-01aa75ed71a1/> (accessed on 15 March 2022).
30. 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 19th ACM International Conference on Intelligent Virtual Agents, Paris, France, 2–5 July 2019; pp. 7–9.
31. 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 AAAI Conference on Human Computation and Crowdsourcing, Virtul, 26–28 October 2020; Volume 8, pp. 63–72.
32. Moroff, N.U.; Kurt, E.; Kamphues, J. Machine Learning and statistics: A Study for assessing innovative demand forecasting models. *Procedia Comput. Sci.* **2021**, *180*, 40–49. [CrossRef]
33. Purohit, D.; Srivastava, J. Effect of manufacturer reputation, retailer reputation, and product warranty on consumer judgments of product quality: A cue diagnosticity framework. *J. Consum. Psychol.* **2001**, *10*, 123–134. [CrossRef]
34. Callon, M.; Méadel, C.; Rabeharisoa, V. The economy of qualities. *Econ. Soc.* **2002**, *31*, 194–217. [CrossRef]
35. Teunter, R.H.; Babai, M.Z.; Syntetos, A.A. ABC classification: service levels and inventory costs. *Prod. Oper. Manag.* **2010**, *19*, 343–352. [CrossRef]
36. Scholz-Reiter, B.; Heger, J.; Meinecke, C.; Bergmann, J. Integration of demand forecasts in ABC-XYZ analysis: Practical investigation at an industrial company. *Int. J. Product. Perform. Manag.* **2012**, *61*, 445–451. [CrossRef]
37. Syntetos, A.A.; Boylan, J.E.; Croston, J. On the categorization of demand patterns. *J. Oper. Res. Soc.* **2005**, *56*, 495–503. [CrossRef]
38. Rožanec, J.M.; Mladenčić, D. Reframing demand forecasting: a two-fold approach for lumpy and intermittent demand. *arXiv* **2021**, arXiv:2103.13812.
39. 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: Berlin/Heidelberg, Germany, 2009; pp. 146–160.
40. 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*, 1–7. [CrossRef]
41. Ubaidillah, N.Z. A study of car demand and its interdependency in sarawak. *Int. J. Bus. Soc.* **2020**, *21*, 997–1011. [CrossRef]
42. Dwivedi, A.; Niranjana, M.; Sahu, K. A business intelligence technique for forecasting the automobile sales using Adaptive Intelligent Systems (ANFIS and ANN). *Int. J. Comput. Appl.* **2013**, *74*, 1–7. [CrossRef]

43. 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. *Technol. Forecast. Soc. Chang.* **2020**, *155*, 120032. [CrossRef]
44. Chandriah, K.K.; Naraganahalli, R.V. RNN/LSTM with modified Adam optimizer in deep learning approach for automobile spare parts demand forecasting. *Multimed. Tools Appl.* **2021**, *80*, 26145–26159. [CrossRef]
45. Halford, G.S.; Baker, R.; McCredde, J.E.; Bain, J.D. How many variables can humans process? *Psychol. Sci.* **2005**, *16*, 70–76. [CrossRef] [PubMed]
46. Barnes, J.H., Jr. Cognitive biases and their impact on strategic planning. *Strateg. Manag. J.* **1984**, *5*, 129–137. [CrossRef]
47. Adadi, A.; Berrada, M. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access* **2018**, *6*, 52138–52160. [CrossRef]
48. 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. *Inf. Fusion* **2020**, *58*, 82–115. [CrossRef]
49. Confalonieri, R.; Coda, L.; Wagner, B.; Besold, T.R. A historical perspective of explainable Artificial Intelligence. *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.* **2021**, *11*, e1391. [CrossRef]
50. Davydenko, A.; Fildes, R.A.; Trapero Arenas, J. *Judgmental Adjustments to Demand Forecasts: Accuracy Evaluation and Bias Correction*; The Department of Management Science, Lancaster University: Lancaster, UK, 2010.
51. Davydenko, A.; Fildes, R. Measuring forecasting accuracy: The case of judgmental adjustments to SKU-level demand forecasts. *Int. J. Forecast.* **2013**, *29*, 510–522. [CrossRef]
52. Alvarado-Valencia, J.; Barrero, L.H.; Önköl, D.; Dennerlein, J.T. Expertise, credibility of system forecasts and integration methods in judgmental demand forecasting. *Int. J. Forecast.* **2017**, *33*, 298–313. [CrossRef]
53. Goodman, B.; Flaxman, S. European Union regulations on algorithmic decision-making and a “right to explanation”. *AI Mag.* **2017**, *38*, 50–57. [CrossRef]
54. 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). Available online: <https://eur-lex.europa.eu/eli/reg/2016/679/oj> (accessed on 15 March 2022).
55. 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. Available online: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206> (accessed on 15 March 2022).
56. Anitha, J. Determinants of employee engagement and their impact on employee performance. *Int. J. Product. Perform. Manag.* **2014**, *63*, 308–323.
57. Emmert-Streib, F.; Yli-Harja, O.; Dehmer, M. Explainable artificial intelligence and machine learning: A reality rooted perspective. *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.* **2020**, *10*, e1368. [CrossRef]
58. Schwalbe, G.; Finzel, B. XAI Method Properties: A (Meta-) study. *arXiv* **2021**, arXiv:2105.07190.
59. Chan, L. Explainable AI as Epistemic Representation. Available online: https://aisb.org.uk/wp-content/uploads/2021/04/AISB21_Opacity_Proceedings.pdf#page=9 (accessed on 25 April 2022).
60. Müller, V.C. Deep Opacity Undermines Data Protection and Explainable Artificial Intelligence. Available online: http://explanations.ai/symposium/AISB21_Opacity_Proceedings.pdf#page=20 (accessed on 25 April 2022).
61. Ribeiro, M.T.; Singh, S.; Guestrin, C. “Why should i trust you?” Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; pp. 1135–1144.
62. Zafar, M.R.; Khan, N.M. DLIME: A deterministic local interpretable model-agnostic explanations approach for computer-aided diagnosis systems. *arXiv* **2019**, arXiv:1906.10263.
63. Hall, P.; Gill, N.; Kurka, M.; Phan, W. Machine Learning Interpretability with H₂O Driverless AI. 2017. Available online: <http://docs.h2o.ai/driverless-ai/latest-stable/docs/booklets/MLIBooklet.pdf> (accessed on 25 April 2022).
64. Sokol, K.; Flach, P. LIMETree: Interactively Customisable Explanations Based on Local Surrogate Multi-output Regression Trees. *arXiv* **2020**, arXiv:2005.01427.
65. Lundberg, S.M.; Lee, S.I. A unified approach to interpreting model predictions. Available online: <https://proceedings.neurips.cc/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf> (accessed on 25 April 2022).
66. Strumbelj, E.; Kononenko, I. An efficient explanation of individual classifications using game theory. *J. Mach. Learn. Res.* **2010**, *11*, 1–18.
67. Pastor, E.; Baralis, E. Explaining black box models by means of local rules. In Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing, Limassol, Cyprus, 8–12 April 2019; pp. 510–517.
68. Guidotti, R.; Monreale, A.; Ruggieri, S.; Pedreschi, D.; Turini, F.; Giannotti, F. Local rule-based explanations of black box decision systems. *arXiv* **2018**, arXiv:1805.10820.
69. Ribeiro, M.T.; Singh, S.; Guestrin, C. Anchors: High-Precision Model-Agnostic Explanations. In Proceedings of the AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018; Volume 18, pp. 1527–1535.
70. Van der Waa, J.; Robeer, M.; van Diggelen, J.; Brinkhuis, M.; Neerinx, M. Contrastive explanations with local foil trees. *arXiv* **2018**, arXiv:1806.07470.

71. Rožanec, J.; Trajkova, E.; Kenda, K.; Fortuna, B.; Mladenčić, D. Explaining Bad Forecasts in Global Time Series Models. *Appl. Sci.* **2021**, *11*, 9243. [[CrossRef](#)]
72. 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), Bratislava, Slovakia, 18–19 September 2021.
73. Panigutti, C.; Perotti, A.; Pedreschi, D. Doctor XAI: An ontology-based approach to black-box sequential data classification explanations. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, 27–30 January 2020; pp. 629–639.
74. 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, Auckland, New Zealand, 26–30 October 2019; pp. 315–316.
75. Rabold, J.; Deiningner, 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: Berlin/Heidelberg, Germany, 2019; pp. 180–192.
76. Lakkaraju, H.; Kamar, E.; Caruana, R.; Leskovec, J. Interpretable & explorable approximations of black box models. *arXiv* **2017**, arXiv:1707.01154.
77. Nguyen, A.p.; Martínez, M.R. On quantitative aspects of model interpretability. *arXiv* **2020**, arXiv:2007.07584.
78. Rosenfeld, A. Better metrics for evaluating explainable artificial intelligence. In Proceedings of the 20th international Conference on Autonomous Agents and Multiagent Systems, Virtual, 3–7 May 2021; pp. 45–50.
79. Amparore, E.; Perotti, A.; Bajardi, P. To trust or not to trust an explanation: using LEAF to evaluate local linear XAI methods. *PeerJ Comput. Sci.* **2021**, *7*, e479. [[CrossRef](#)]
80. Samek, W.; Müller, K.R. Towards explainable artificial intelligence. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*; Springer: Berlin/Heidelberg, Germany, 2019; pp. 5–22.
81. 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* **2018**, arXiv:1806.09936.
82. El-Assady, M.; Jentner, W.; Kehlbeck, R.; Schlegel, U.; Sevastjanova, 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, Glasgow, UK, 4 May 2019.
83. 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* **2021**, arXiv:2108.01737.
84. Hoffman, R.R.; Mueller, S.T.; Klein, G.; Litman, J. Metrics for explainable AI: Challenges and prospects. *arXiv* **2018**, arXiv:1812.04608.
85. 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: Berlin/Heidelberg, Germany, 2020; pp. 163–178.
86. 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* **2021**, arXiv:2103.01035.
87. Verma, S.; Dickerson, J.; Hines, K. Counterfactual Explanations for Machine Learning: A Review. *arXiv* **2020**, arXiv:2010.10596.
88. Mohseni, S.; Zarei, N.; Ragan, E.D. A multidisciplinary survey and framework for design and evaluation of explainable AI systems. *Acm Trans. Interact. Intell. Syst. (Tiis)* **2021**, *11*, 1–45. [[CrossRef](#)]
89. Lage, I.; Ross, A.S.; Kim, B.; Gershman, S.J.; Doshi-Velez, F. Human-in-the-loop interpretability prior. *arXiv* **2018**, arXiv:1805.11571.
90. Rozanec, J.M. Explainable demand forecasting: A data mining goldmine. In Proceedings of the Companion Proceedings of the Web Conference 2021, Ljubljana, Slovenia, 19–23 April 2021; pp. 723–724.
91. Leban, G.; Fortuna, B.; Brank, J.; Grobelnik, M. Event registry: Learning about world events from news. In Proceedings of the 23rd International Conference on World Wide Web, Seoul, Korea, 7–11 April 2014; pp. 107–110.
92. Publications Office of the European Union. EU Open Data Portal: The Official Portal for European Data. Available online: <https://data.europa.eu> (accessed on 15 December 2020).
93. Noy, N.; Gao, Y.; Jain, A.; Narayanan, A.; Patterson, A.; Taylor, J. Industry-scale knowledge graphs: Lessons and challenges. *Queue* **2019**, *17*, 48–75. [[CrossRef](#)]
94. Kusner, M.; Sun, Y.; Kolkin, N.; Weinberger, K. From word embeddings to document distances. In Proceedings of the International Conference on Machine Learning, Lille, France, 6–11 July 2015; pp. 957–966.
95. Brank, J.; Leban, G.; Grobelnik, M. Annotating Documents with Relevant Wikipedia Concepts. Available online: https://ailab.ijs.si/Dunja/SiKDD2017/Papers/Brank_Wikifier.pdf (accessed on 25 April 2022).
96. Lee, J.D.; See, K.A. Trust in automation: Designing for appropriate reliance. *Hum. Factors* **2004**, *46*, 50–80. [[CrossRef](#)]
97. Kilani, Y.; Alhijawi, B.; Alsarhan, A. Using artificial intelligence techniques in collaborative filtering recommender systems: Survey. *Int. J. Adv. Intell. Paradig.* **2018**, *11*, 378–396. [[CrossRef](#)]
98. Karimi, M.; Jannach, D.; Jugovac, M. News recommender systems—Survey and roads ahead. *Inf. Process. Manag.* **2018**, *54*, 1203–1227. [[CrossRef](#)]

-
99. 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 Min. Knowl. Discov.* **2021**, *35*, 568–592. [[CrossRef](#)]
 100. Michael, J.; Stanovsky, G.; He, L.; Dagan, I.; Zettlemoyer, L. Crowdsourcing question-answer meaning representations. *arXiv* **2017**, arXiv:1711.05885.