

RECKONition: a NLP-based system for Industrial Accidents at Work Prevention

(Discussion Paper)

Patrizia Agnello¹, Silvia Maria Ansaldi¹, Emilia Lenzi², Alessio Mongelluzzo² and Manuel Roveri²

¹INAIL - Dipartimento Innovazioni Tecnologiche

²Politecnico di Milano - Dipartimento di Elettronica, Informazione e Bioingegneria

Abstract

Extracting patterns and useful information from Natural Language datasets is a challenging task, especially when dealing with data written in a language different from English, like Italian. Machine and Deep Learning, together with Natural Language Processing (NLP) techniques have widely spread and improved lately, providing a plethora of useful methods to address both Supervised and Unsupervised problems on textual information. We propose RECKONition, a NLP-based system for Industrial Accidents at Work Prevention. RECKONition, which is meant to provide Natural Language Understanding, Clustering and Inference, is the result of a joint partnership with the Italian National Institute for Insurance against Accidents at Work (INAIL). The obtained results showed the ability to process textual data written in Italian describing industrial accidents dynamics and consequences.

Keywords

Workplace Safety, Industry 4.0, Natural Language Processing.

1. Introduction

Despite the ever-growing awareness and the advances in the technology and the procedure, the problem of accidents at work represents a relevant issue from both the social and economic point of view. In more detail, accidents at work refer to damages to the health of the worker induced by accidents correlated with the working activities. Unfortunately, such accidents can result into serious damages, such as temporary or permanent reduction or loss of working capacity, or even death. In Italy, this issue is particularly perceived as very relevant since in last few years the number of accidents at work is larger than 500.000 per year, among which more than 1000 resulted in the death of the worker. Addressing this issue in Italy is the primary goal of INAIL, which is National Insurance Institute for Industrial Accidents and Occupational Diseases. More specifically, INAIL's objectives are the reduction of injuries, the protection of workers performing hazardous jobs, and the facilitation of the return to work of people injured at workplace. To achieve these relevant and challenging objectives, INAIL explores

SEBD 2021: The 29th Italian Symposium on Advanced Database Systems, September 5-9, 2021, Pizzo Calabro (VV), Italy

✉ p.agnello@inail.it (P. Agnello); s.ansaldi@inail.it (S. M. Ansaldi); emilia.lenzi@polimi.it (E. Lenzi); alessio.mongelluzzo@polimi.it (A. Mongelluzzo); manuel.roveri@polimi.it (M. Roveri)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

novel technological and societal solutions and paradigms to provide an integrated system of protection, ranging from preventive actions at the workplace to medical services and financial assistance.

Many Natural Language Processing (NLP) applications and solutions can be found in the literature to tackle tasks like Sentiment Analysis or Text Classification, for instance [1]. However, when it comes to the Italian language, and injury prevention in particular, there are no suited solutions available. In this path, the aim of this paper is to introduce RECKONition, a novel Natural Language Processing-based system to extract knowledge from textual descriptions of accidents at work aiming at defining industrial accidents preventive actions at the workplace. The RECKONition system comprises three different modules, i.e., Association Rule Generation, Textual Description Clustering and Textual Description Inference, able to extract relationships between accident events and create groups of homogeneous accident descriptions representing the input for the definition of preventive actions. The RECKONition system has been successfully applied to a real-world database storing the textual descriptions of accidents occurring in Northern Italy from 2013-2018.

The paper is organized as follows. In Section 2 introduce RECKONition system and its three algorithmic modules. In Section 3, 4, and 5 respectively, we detail our implementation of the Association Rules Mining, Natural Language Clustering, and Natural Language Prediction modules. Section 6 summarizes our work and discuss on the possible next steps and improvements

2. The RECKONition system: extracting knowledge from textual descriptions of accidents at work

The architecture of the RECKONition system, which is described in Fig. 1, comprises three different modules: Association Rule Generator, Textual Description Clustering and Textual Description Inference. All these models receive in input a set of textual descriptions of accidents at work and provide in output *association rules* to discover relationships between terms, *clusters of textual descriptions* highlighting groups of accidents at work sharing similarities in their descriptions, and *next sentence predictions* from the textual description of the accidents. All the outputs of the RECKONition system represent valuable tools for the INAIL expert to gain knowledge about occurred accidents at work and define prevention policies. The three modules of RECKONition system are detailed in the next sections.

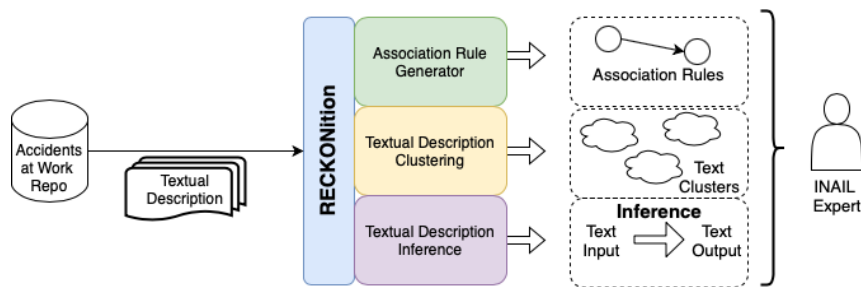


Figure 1: The Architecture of the RECKONition system.

3. Association Rules Mining

Association Rules (ARs) allow to discover relations between variables in large datasets and to graphically show such dependencies in a convenient representation which can be easily interpreted by humans [2] (we show in Figure 2 an example of an AR representation).




Figure 2: An example of a negative association rule graphical representation.

Let us define our *database* D as a collection of *transactions* from the transactions set T : a transaction is defined as an element of our database, i.e., a sentence containing a textual description of the work dynamics leading to an accident. Hence, for each transaction we can define a set of words or a transformation of words, called *items*, which belong to the items set I and constitute our transaction. In this scenario, an association rule $A \Rightarrow B$ indicates that the occurrence of item B is usually observed when also A is found in the transaction, for $A, B \subset I$.

Association Rules mining procedure is based on three measures, namely *Support*, *Confidence*, and *Lift*, whose values define whether a rule is meaningful or not. The support of an item A can be defined as $Supp(A) = P(A) = \frac{|\{t | \forall t \in T \wedge A \in t\}|}{|T|}$, namely, the percentage of transactions where item A is found. Support can also be defined for a rule, e.g., $Supp(A \Rightarrow B)$, to represent the percentage of transactions in the database where both A and B are found. With the confidence of a rule we model the conditional probability of observing the consequent item having observed the antecedent, so it can hence be defined as $Conf(A \Rightarrow B) = \frac{P(AB)}{P(A)}$. The third measure, i.e., the lift, is used to characterize the relationship between the antecedent and the consequent of the association rule. It is defined as $Lift(A \Rightarrow B) = \frac{P(AB)}{P(A)P(B)} = \frac{P(B|A)}{P(B)}$, a lift greater than 1 identifies a positive relationship between the items, i.e., the conditional probability of observing B given that A is found is greater than the probability of observing B in the database.

Apriori algorithm is a well-known association rules mining method in the literature, which leverages the measures we introduced to explore the itemset and build association rules from them [3]. This method only focuses on the subset of those items belonging to the Frequent Itemset (FIS), i.e., those items whose support is greater than a minimum threshold $minsupp$. However, in some settings (e.g., medical descriptions) those items whose occurrence frequency is very low can be the ones carrying more semantical importance (e.g., names of diseases or even specific injuries). Following this intuition, we provided a Python implementation¹ to the algorithm FISinFIS Apriori proposed in [4], where Positive Association Rules (PARs) and Negative Association Rules (NARs) are built from both the Frequent Itemset and the Infrequent Itemset. FISinFIS Apriori algorithm sets a minimum threshold on the Inverse Document Frequency (IDF), defined as $IDF(i) = \log \frac{|T|}{|\{t | t \in T \wedge i \in t\}|}$, to filter those items which are too frequently

¹The code is available at https://github.com/AlessioMongelluzzo/FISinFIS_Apriori_Python

used throughout the database. Moreover, we introduced in our implementation an additional threshold to the IDF as an upper bound for too rarely used items. The rationale behind this decision is to avoid typos and missing values placeholders which constitute a non-negligible part of our dataset. The results coming from the application of the mining algorithm to our work accidents descriptions database allows RECKONition to detect those industrial properties or actions that are more likely related to an injury or, with negative rules, to their prevention, we show in Figure 3 an AR graph example from a mock experiment performed on few sentences we wrote for this purpose.




Figure 3: Example of an association rule graph from FISinFIS Apriori application.

4. Natural Language Clustering

Textual Description Clustering [5] is performed to highlight differences and similarities among accident descriptions, and to group together the similar ones in order to facilitate the identification of appropriate interventions according to the situation described. In the following sections we will show the two different approaches used by our system. We will see both of them use the *K-Medoids* algorithm [6], but differ in the operations performed on the *dataset*.

4.1. Tags Occurrence Clustering

A first approach is based on an *ontology* describing the context of interest (in the case of RECKONition, the metallurgical company) and uses TAGs representing ontological classes to integrate contextual knowledge in the clustering analysis.

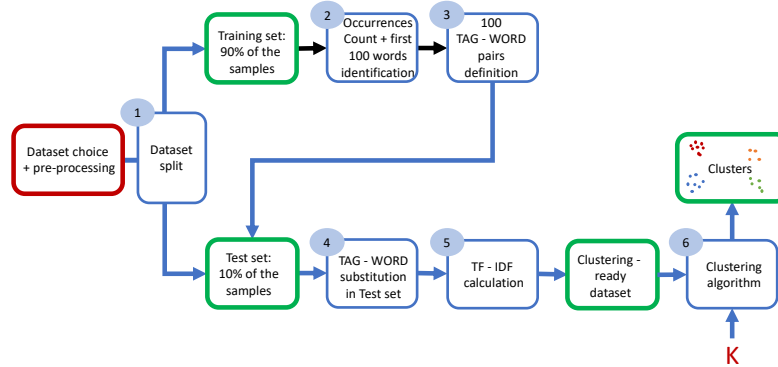


Figure 4: Tags Occurrences Clustering - The process

As you can see in Figure 4, we calculate the occurrences of each word, we identify the hundred most frequent ones, and we define TAG - WORD pairs in the training set; we then perform the hundred substitution in the test set. In this way, words representing similar concepts are read by the algorithm as the same term. As can also be seen from the diagram, before calculating the occurrences, all the necessary pre-processing operations were carried on the entire *dataset*, and in particular the *stop_words* were eliminated. Almost all the remaining words resulted to be descriptive of the context then, and it was possible to associate them with a TAG. In the rare case in which this did not happen, or the substitution introduced ambiguity, the terms remained unchanged. Once performed the substitution, we calculate the Time frequency - Inverse Document Frequency (TF-IDF) [5] of each term in each description. After step 5 we then have a *dataset* ready for clustering and it is composed by 6662 terms as features, whose values correspond to the calculated TF-IDF.

Once we have the *dataset*, we perform *K-Medoids*. Here we report some example of the clusters obtained by setting the $K = n_clusters$ parameter equal to 30, and we show how we are able to separate the different situations described.

```
Cluster 12:
medoid: [USING HOSE]
elements: [USING HOSE]; [USING DRILL]; [USING HAMMER]; [USING CUTTER];...

Cluster 17:
medoid: [GOT OFF FORKLIFT]
elements: [GOT OFF FORKLIFT]; [RIDE FORKLIFT]; [MOVE TRUCK]; [RIDING TRUCK];...
```

What is important to notice is that the algorithm is not only able to group together descriptions containing the exact same words, but it is also able to identify the ones concerning similar incidents in the analysed context. In this perspective, Clusters 12 represents accidents involving working tools, while Clusters 17 represents accidents occurring on board mobile machinery.

4.2. Transformers-Based Clustering

A different clustering approach adopted by RECKONition leverages the novel Encoder-Decoder models with *Attention* and *self-Attention* mechanism to build homogenous groups of accidents descriptions [7], [8]. These techniques comprise several stacked encoder-decoder layers, which can model different syntactic structures with self-Attention modules together with feed-forward neural networks, resulting in millions of trainable parameters: the base version of the Bidirectional Encoder Representations from Transformers (BERT) [9] features 12 encoder-decoder layers, 768 hidden nodes for each layer and 12-attention-heads, resulting in 110M parameters, hence, requiring a very large amount of data to be properly trained, as well as time and computational capacity. RECKONition uses a pre-trained BERT model publicly available², which was trained on 13GB of Italian textual data from Wikipedia and OPUS corpora³, and performs fine-tuning of such model on a subset of the accidents-at-work *dataset* provided by INAIL. The

²The pre-trained model can be found at <https://huggingface.co/dbmdz/bert-base-italian-cased>

³<https://opus.nlpl.eu/>

fine-tuned model is then validated on a hold-out set which is used to perform clustering on unseen textual descriptions: during the forward step of a new sentence, the hidden states of the last decoder layer in the model are extracted and used as numerical features representing the sentence. This procedure is performed for each description in the validation set, so to build a new *dataset* with shape $v \times (768 * n_t)$, being 768 the number of hidden states of a BERT layer, v the number of sentences in the validation set, and n_t the number of length of the tokenized sentences. RECKONition leverages Incremental Principal Component Analysis (IPCA) [10] to reduce the dimensionality of the representative features to a number of principal components so that at least 85% of the variability of data is explained, and performs *K-Medoids* clustering for values of the $n_clusters$ k in the range [2, 100]. Unlike the clustering approach described in Section 4.1, this method can effectively find sentences with highly correlated syntactic structures. We show in the following two examples of clusters obtained with the Transformers-based clustering approach applied to a sample dataset.

```
MEDOID: "RIDING THE FORKLIFT "
["RIDING THE FORKLIFT", "RIDING THE CAR", "RIDING THE FORKLIFT ", "RIDING THE CAR", "RIDING
THE CAR", "RIDING THE BIKE", "RIDING THE FORKLIFT", "REPAIRING THE FORKLIFT"]

MEDOID: "CLIMBING THE LADDER"
["CLIMBING THE LADDER", "CLIMBING THE WALL", "CLIMBING THE LADDER", "CLIMBING ON THE
ROOF", "CLIMBING THE LADDER ", "CLIMBING THE LADDER ", "REMOVING HEAR DEFENDER"]
```

5. Natural Language Prediction

There is plenty of experimental settings where a Language Model (LM) [11], i.e., a model that is able to process textual sequences [12], can be effectively applied. These range from feature extraction (as we did with Trasformers-based clustering in Section 4.2), to text classification, question answering, next sentence prediction, natural language inference, or even text generation. RECKONition comprises a custom LM, which was properly built and trained on the accidents *dataset* provided by INAIL, with the purpose of performing next sentence prediction from the textual description of the accidents dynamics to those of the consequences on the workers involved. The high-level architecture of the LM is shown in Figure 5. We emphasize that the tokenization step [13], which allows to move from symbolical (textual) to numerical representation, requires to define a vocabulary size to fix the number of tokens, which our model sets to 5000 items. The embedding layer is then used to project the sparse representation obtained by the tokenization step to a dense space with size 128. The sequence learning ability of the model resides in the n layer shown in Figure 5, which is composed of two Bidirectional Long Short-Term Memory (LSTM) with 100 units each [14], [15], a flattening layer, two dense layers with 50 neurons each and ReLU activations, and a final Dropout [16] layer with a 50% drop rate before the output dense layer with 5000 neurons (one for each item) and sigmoidal activation. This model can be trained in a supervised fashion [17] minimizing the Binary Cross Entropy loss function with Adaptive Moment Estimation (ADAM) [18]. This is achieved by providing the descriptions of the accidents dynamics as input features, together with the corresponding effects

and injuries on the workers involved as input targets. For a new unseen sentence describing a working scenario, the model will output a probability vector on the items in the vocabulary defining those which better describe the consequence of the input description.




Figure 5: High-level architecture of the language model. s and s' are textual input and output, respectively. Layer t represents the tokenizer, e is an embedding layer, and d a decoder. The main component n is a recurrent neural network model.

We show in the following two examples of textual prediction given some input descriptions, where the output sentence results in a scattered sequence of items due to the limitation on the vocabulary size.

Description: "the worker was climbing the ladder"

Prediction: ["the", "to", "its", "fall", "off", "ankle", "sprained"]

Description: "she was driving her car to go home"

Prediction: ["she", "her", "to", "the", "cervical", "trauma"]

6. Discussion and conclusion

We introduced and described RECKONition system and its application scenario on a real world dataset. The system is currently operating to understand and extract knowledge from the accidents descriptions database and support the industries towards a safer workplace. The results obtained and the continuous dialogue with INAIL's partners have highlighted the unsupervised ability of RECKONition to both cluster work activities and to identify relationship between accidents and consequences. As a next step, RECKONition will be deployed in real factories where proper sensors can produce inputs to the system, allowing real-time predictions to detect hazardous activities and at the same time improving data quality and performances of the system itself.

Acknowledgements

This work has been funded by INAIL within the BRiC/2018, ID09 framework, project RECKON. The authors wish to acknowledge all the other researchers involved in the project: Francesco Braghin (PoliMI, DMEC), Enrico Cagno (PoliMI, DIG), and their research groups. We also thank Cinzia Frascheri (IAL) and Irene Tagliaro (API-TECH).

References

- [1] P. Badjatiya, S. Gupta, M. Gupta, V. Varma, Deep learning for hate speech detection in tweets, in: Proceedings of the 26th international conference on World Wide Web companion, 2017, pp. 759–760.
- [2] R. Agrawal, T. Imieliński, A. Swami, Mining association rules between sets of items in large databases, in: Proceedings of the 1993 ACM SIGMOD international conference on Management of data, 1993, pp. 207–216.
- [3] R. Agrawal, R. Srikant, et al., Fast algorithms for mining association rules, in: Proc. 20th int. conf. very large data bases, VLDB, volume 1215, Citeseer, 1994, pp. 487–499.
- [4] S. Mahmood, M. Shahbaz, A. Guergachi, Negative and positive association rules mining from text using frequent and infrequent itemsets, The Scientific World Journal 2014 (2014).
- [5] J. Leskovec, A. Rajaraman, J. Ullman, Mining of massive datasets cambridge university press, 2014.
- [6] H.-S. Park, C.-H. Jun, A simple and fast algorithm for k-medoids clustering, Expert systems with applications 36 (2009) 3336–3341.
- [7] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using rnn encoder-decoder for statistical machine translation, arXiv preprint arXiv:1406.1078 (2014).
- [8] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, arXiv preprint arXiv:1706.03762 (2017).
- [9] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, arXiv preprint arXiv:1810.04805 (2018).
- [10] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using rnn encoder-decoder for statistical machine translation, arXiv preprint arXiv:1406.1078 (2014).
- [11] Y. Bengio, R. Ducharme, P. Vincent, C. Janvin, A neural probabilistic language model, The journal of machine learning research 3 (2003) 1137–1155.
- [12] I. Sutskever, O. Vinyals, Q. V. Le, Sequence to sequence learning with neural networks, arXiv preprint arXiv:1409.3215 (2014).
- [13] J. J. Webster, C. Kit, Tokenization as the initial phase in nlp, in: COLING 1992 Volume 4: The 15th International Conference on Computational Linguistics, 1992.
- [14] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural computation 9 (1997) 1735–1780.
- [15] A. Graves, N. Jaitly, A.-r. Mohamed, Hybrid speech recognition with deep bidirectional lstm, in: 2013 IEEE workshop on automatic speech recognition and understanding, IEEE, 2013, pp. 273–278.
- [16] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, The journal of machine learning research 15 (2014) 1929–1958.
- [17] C. M. Bishop, Pattern recognition and machine learning, springer, 2006.
- [18] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980 (2014).