Exploring Word Embedding in Modeling Risk Perception

Claudio Proietti Mercuri[1], Jonas Benjamin Krieger[2], and Rui Jorge Almeida[3,4]

[1] School of Business and Economics
Maastricht University, Maastricht, The Netherlands
c.proiettimercuri@alumni.maastrichtuniversity.nl
[2] The United Nations University - Maastricht Economic and Social Research Institute on Innovation and Technology (UNU-MERIT), Maastricht, The Netherlands
krieger@merit.unu.edu
[3] Dept. of Quantitative Economics, School of Business and Economics
Maastricht University, Maastricht, The Netherlands
[4] Dept. of Data Analytics and Digitization, School of Business and Economics
Maastricht University, Maastricht, The Netherlands.
rj.almeida@maastrichtuniversity.nl

Abstract. Risk perception models attempt to capture how individuals or groups perceive and evaluate risks, focusing on subjective assessments rather than quantitative probabilities of occurrence. The aim is to identify influencing factors that may affect risk perception with the greater aim of better understanding (cognitive) decision-making processes in risk scenarios. To operationalise these insights, we apply word embedding techniques to quantify and assess risk perception information embedded in textual data. This approach, based on large corpora and machine learning models, allows predicting risk perception from the semantic content of language(s). This paper complements initial approaches of using word embedding vector space to forecast risk perceptions scores. We deepened that understanding by examining how the use of similar words to given risk terms in a vector space, or context words can provide additional information regarding their semantic use. Finally, we applied arithmetic operations to incorporate cultural and geographical contexts into risk perception. The results show that adding more words to create context reduces performance of the models, while using arithmetic operations can provide better forecasts of risk perceptions, and that they may also be used to further explore cultural or geographic variations in risk perceptions.

Keywords: Risk Perception · Word2Vec · Word Embedding · Analogies

1 Introduction

The study of risk perception(s) has proven to be important and useful across various domains such as psychology, sociology and economics, as individual risk perceptions may significantly influence human behaviour [23]. Risk perception does not necessarily refer to the probability of an adverse event occurring, but rather to the judgement and evaluation of hazards to which they, their institutions and/or the environment are exposed, taking into account individual experiences and beliefs [22, 20]. Understanding risk perceptions is essential to be able to address any existing knowledge gaps between experts' risk assessments and possibly conflicting misconceptions and false beliefs (often those of laypeople) in a targeted manner to facilitate better-informed (risk) decision-making [17, 7].

Conceptually, explanatory models should reflect the multidimensionality of risk perception and consider the fact that a single 'objective' measure of risk can hardly be defined [24]. This also translates to a strong psychological component that may influence how we perceive and evaluate risks based on our character, cultural background, benefit perception, etc. [25, 28, 1, 11]. Here, various psychological dimensions and variables have been identified that influence individuals' risk assessment [11, 13, 26]. Here, psychological approaches highlight the role of cognitive biases, such as the reliance on readily available information, a broader understanding of risk, or potential benefits, which may influence risk perception [27, 7]. By analysing this rather complex construct of 'risk perception', on the one hand risk literature has been able to identify overarching themes; for example, that natural hazards - despite their potentially drastic effects - are generally rated relatively low compared to technical hazards [28]. On the other hand, for certain domains specific influencing parameters could be extracted and identified, e.g. in the area of public health, factors such as trust in government, health organisations, individual feelings of concern and media exposure [8].

A common tool for 'measuring' and 'analysing' risk perception is the psychometric model of risk, which identifies nine risk dimensions which are commonly used in risk literature [11, 15, 1]. The preferred method of analysis for measuring risk perception have been surveys in which individuals rate different risk sources based on the aforementioned risk dimensions. However, classical surveys present the major drawback of being time-consuming and costly, leading to scalability challenges. An alternative approach has been to use semantic representations to describe knowledge and associations that individuals have with various objects [21]. This approach is based on a renowned linguistic insight where the distribution of words in language reflects an individual's knowledge and associations across various objects and concepts [4, 5, 2]. When using this method, a concept such as "football" is represented by a high-dimensional vector. This allows for measuring the distance of the football vector to the vectors of other concepts, events, and entities. These word vector models have shown to be successful among several areas of application; such as understanding perceptions of food healthiness and category learning [12, 31] or in predicting human judgement regarding multiple domains [21]. They have also been applied to historical data on gender stereotypes and leadership perceptions [2, 6], social media analysis for real-time monitoring of COVID-19 perception [9], and the mapping of risk semantics [29].

After the first 'proof-of-concept' has been provided that word embedding models can be used to predict risk perception [5]. Thereby, word embedding model(s) aim(s) to incorporate the semantic meanings and individuals' associations with risk(s). Hereby, we further study this novel method, where we try to incorporate specific influencing parameters on risk perception. For this, we combine words similar to the risk sources with them to provide context and additional information related to the semantic use of a risk source. Additionally, we define arithmetic operations to include geographical contextual information on the risk terms. For that, we specifically examine the information encapsulated in the 300 dimensions of each vector and assess the impact of several manipulations of the vector space on the accuracy of risk perception predictions. Furthermore, we present how word embedding model(s) capture(s) variation in risk perception and in the words people associate with risk sources across different geographical regions.

2 Embedding Based Risk Perception Model

We propose a word embedding vector for predicting risk perception, extending the risk perception model in [5]. We consider a word2vec model that learns a vector representation of each word using a shallow neural network architecture [16]. The word2vec model is based on the skip-gram-model, that consists of a input layer, a projection layer and an output layer to predict the likelihood of a word appearing in a given context, i.e. sequence of words in a corpus $C(w)$. The context can take the form of a single word or a collection of words [30]. The objective function is to maximize the probability of predicting the target word w given the context words $C(w)$ and parameter set Θ:

$$\arg\max_{\Theta} \prod_{w \in \text{Text}} p(w|C(w); \Theta). \tag{1}$$

This word2vec model is trained in an unsupervised manner on very large data sets, such as news, that allows the model to learn complex word relationships. These relationships led to the idea of testing the soundness of embedding spaces via four-word proportional analogies, typically written as $w_1 : w_2 :: w_3 : w_4$. Each word is represented by a vector, where it is assumed that the vector differences between each pair are roughly equal: $w_2 - w_1 \approx w_4 - w_3$. Using vector arithmetic, we obtain $w_4 \approx w_3 + w_2 - w_1$. A system selects the word with the maximum cosine similarity to all the possible words in $C(w)$ - typically excluding w_1, w_2 and w_3 - assuming all word vectors are unit length [14]. These analogies have been shown to work more reliably for certain types of analogical word relationships [10, 18].

In this paper we propose a series of words to provide context to the risk term. We then use word relationships, similarity of words through the word2vec model, and a set of arithmetic operations to include contextual information to risk terms. Specifically, we use a pre-trained model on Google news from [19]. This model represents information that each individual is exposed to from news and is shown to affect risk assessment [7, 26, 8]. Hence, information obtained from the pre-trained model represents the latent influence of news on risk perceptions. For the word2vec model, following [5], we define a set of risk sources, where risk source i has n ratings provided by each participant in a survey study. We then map each risk source i in a 300-dimensional vector.

The participants in the survey study assess the riskiness of each risk source on a scale from -100 (safe) to +100 (risky). Ratings for a specific risk source i can be denoted as $y_{i1}, y_{i2}, ..., y_{in}$. For each risk source, the average rating $\bar{y}_i$ is computed to obtain a single continuous measure of risk:

$$\bar{y}_i = \frac{1}{n} \sum_{k=1}^{n} y_{ik}, \tag{2}$$

where y_{ik} is the rating of the kth participant for risk source i. The purpose is to model the ratings $\bar{y}_i$, assigned to each risk source i, with the 300-dimension vector x_{in} with $n \in \{1, 2, \ldots, n\}$, which represents the value of the source i on the dimension j of the corresponding 300-dimensional Word2Vec vector [5].

The proposed estimation of risk is a high dimensional regression problem since the number of independent variables (300 values in each vector) is larger than the number of observations (200 ratings representing the 200 risk sources). In this paper we focus our attention on support vector regressors (SVRs) and ridge regression to address this high dimensionality. We apply support vector regressions with radial basis function kernels (SVR-RBF):

$$K(x_{im}, x_{jn}) = e^{-\gamma |x_{im} - x_{jn}|^2}, \tag{3}$$

where γ is the spread of the kernel and x_{im} and x_{jn} represent the mth and nth dimensions of the 300-dimensional vectors associated with risk sources i and j, respectively. For each risk source i, the Ridge objective function is:

$$\underset{\beta}{\text{minimize}} \ \left(|\bar{y}_i - x_i \beta|^2 + \lambda |\beta|^2 \right), \tag{4}$$

where x_i is the 300-dimensional vector of risk source i, y_i is the corresponding risk rating, β is the vector of model parameters, and λ is the regularization parameter. Hyperparameters SVR-RBF and Ridge are optimized using cross-validation and grid search, respectively. The model's performance is assessed using R^2 and RMSE or MSE. For each metric, we perform a paired t-test using the Shapiro-Wilk test, using the psychometric results as baseline.

We first use the conceptual model to replicate [5], to show that word embedding performs comparably with the psychometric approach with human interviews. We then study whether the predictive accuracy of Word2Vec in terms of risk perception can be enhanced. In order to provide more context and semantic meaning, we adjust the word embedding to capture the individuals' psychological associations to specific words. The predictive performance of the model is examined by combining the 300-dimensional vector representation of the corresponding risk source together with the vector representations of the n-closest words (with $n = \{1, 5, 10\}$ to the given risk source, using cosine similarity.

In the second part of this study, we consider another aspect of word embedding, namely the ability to use arithmetic operations (or analogies) to modify the semantic meaning the vector representation of a risk. The survey participants reside in the United States, indicating that their ratings represent the lay judgment within this culture. We propose to use arithmetic operations to transform the semantic representation towards a distinct geographical region, namely an European perspective. This is performed by adding word-vectors *'Europe', 'EU', 'European Union'*, to reflect usage patterns or cultural contexts associated with a risk source in the European context. If there are differences in risk perceptions, distinct words would be associated with the risk sources.

3 Exploring Risk Perception on Different Risk Sources

3.1 Data Description

We apply the proposed method to the data consisting of 200 risk sources generated by participants residing in the US without any category restrictions. After creating these risks sources, they were rated twice in two datasets [5]. The first rating employed the semantic vector, where participants assessed the riskiness of each risk source on a scale from -100 (safe) to +100 (risky). The second rating utilized the psychometric approach, in which participants evaluated each risk source based on nine dimensions in a seven-point scale: voluntariness, immediacy of death, knowledge to the person exposed to the risk, knowledge to science, controllability, novelty, the catastrophic potential of the risk, the potential for fatal consequences, and the amount of dread associated with the risk source [11].

Table 1 presents the summary statistics of the dataset and the two approaches distinguishing between the set of independent variables (i.e. the semantic 300-dimensional vectors and the psychometric 9-dimensional vector) and y representing the dependent variable participant ratings. Rating y, has a high mean and standard deviation, suggesting a high level of disagreement among participants in rating risk sources.

Table 1. Summary Statistics

Study	Count	Mean	Std. Dev.	Min.	Median	Max.
Embedding (X)	200.00	-0.00655	0.18123	-0.52171	-0.00649	0.50319
Psychometric (X)	200.00	3.73364	1.17692	1.47003	3.67197	6.18617
Rating (y)	200.00	14.75664	46.61838	-84.30986	21.19405	93.03356

Figure 1 presents a t-distributed stochastic neighbor embedding (t-SNE) for the 200 risk sources. We observe that semantically similar words are close to each other in the vector space. For instance, region 1 clearly encapsulates risk sources related to natural disasters. Furthermore, slightly above region 1, words such as *cold, sun, ice* are related to natural effects although they are not perceived as natural disasters. In contrast, region 2 encapsulates risk sources strongly associated with criminal activities. Similarly, regions 3 and 4 represent diseases and related concepts, with words associated with physical dysfunctions (region 4). In region 5, which highlights outdoor sports, words scattered around the region refer to other sports or to similar vehicles used in region 5.

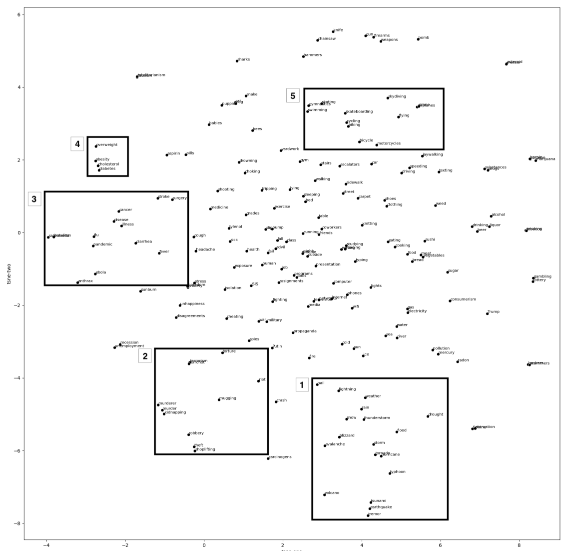

Fig. 1. Scatter-plot: Risk Sources

3.2 Inclusion of Context

In the first study we enhance the ability of word embedding to capture the psychological associations individuals attribute to specific words. The predictive performance of the model is examined by employing the 300-dimensional vector representation of the corresponding risk source together with the vector representations of the n-closest words to the given risk source. The aim is to provide the model with more context – semantic meaning – to improve its accuracy.

Table 2 presents the out-of-sample comparison between the 1-5-10 most similar words – without the original risk source – and the benchmarks, psychometric approach and Bhatia [3]. Even though we increase the predictors from 300 to 3,000 variables, adding the closest words results in worse performance than using the risk word. Adding the original risk source as a term leads to an improvement in the overall RMSE score for both models, with ridge demonstrating superior predictive capabilities. In the latter it is obtained an overall performance comparable to the two benchmarks utilized.

Table 2. Out-of-sample Accuracy: Psychometric and Bhatia (2019)'s study approaches vs the risk source + 1, 5, 10 closest words

Model - Approach	R^2	RMSE	**Approach**	R^2	RMSE
Ridge - Psyc	**0.75**	**22.22**		**0.75**	**22.22**
Ridge - Bhatia	**0.75**	**22.14**		**0.75**	**22.14**
Ridge - 1W	0.69*	24.68*	R& 1W	0.75***	22.53***
Ridge - 5W	0.69**	24.95**	R& 5W	0.72**	23.44**
Ridge - 10W	0.69**	24.73**	R& 10W	0.71***	23.98***
SVR-RBF - Psyc	**0.85**	**17.13**		**0.85**	**17.13**
SVR-RBF - Bhatia	**0.76**	**21.97**		**0.76**	**21.97**
SVR-RBF - 1W	0.66*	25.84*	R& 1W	0.72*	23.49**
SVR-RBF - 5W	0.65*	26.53*	R& 5W	0.67*	25.44*
SVR-RBF - 10W	0.62*	27.62	R& 10W	0.64*	26.84*

We next compare the model's performance for each risk source word. These results are not included in tables or figures for space considerations. Both models tend to undervalue (overvalue) the risk ratings when they are positively rated with ratings between zero and 100 (negatively rated with ratings between -1 to -100). The SVR-RBF model underpredicts more often compared to the Ridge model, especially for risk sources with higher values. The Ridge model, on the other hand, leads to more extreme residuals, indicating higher sensitivity to certain data features. There are some cases where model predictions are unintuitive, as an example words 'Trump' and 'heroin'. This is potentially due to the differences in participants' risk perceptions, and the model's objective to assess an average risk perception in 2.

3.3 Arithmetic Operations

We next investigate the impact of algebraic operations in the embedding space on obtained results. The aim is to acknowledge that individuals from different continents or countries may perceive the same risk source differently. For this purpose, we shift the semantic space toward different geographical areas, and investigate the model's ability to capture variations in risk perception and in the word associated to each risk source.

The proposed arithmetic operations extend the word embedding space by adding different words associated with Europe and subtracting an equal number of words related to the U.S. We further incorporate risk-related words, with and without the U.S. and E.U. words. Finally, a comparative test was conducted to assess the impact of using seemingly unrelated words such as "fruit, piano, kitten" versus risk-related words like "risk, danger, perception". The results are shown Table 3. The trend observed is consistent across all tests: the model's performance began to decline as more words are included, as in the previous section. This is due to the alteration each dimension undergoes as different words are added or subtracted. Such manipulation could mislead the model by changing the statistical occurrence of the examined word.

Table 3. Influence of Algebra Operations on Model Performance, model evaluation

Model	R^2 Test	RMSE Test	Words Used	# Words
Ridge	0.75	22.14	Risk	1
SVR-RBF	0.79	20.22	Risk	1
Ridge	0.75	22.14	USA	1
SVR-RBF	0.79	20.28	USA	1
Ridge	0.75	22.14	Fruit	1
SVR-RBF	0.79	20.26	Fruit	1
Ridge	0.75	22.14	USA, Risk, American, Danger, Perception	5
SVR-RBF	0.73	22.98	USA, Risk, American, Danger, Perception	5
Ridge	0.75	22.14	butterfly, bicycle, umbrella, piano, kitten	5
SVR-RBF	0.71	23.99	butterfly, bicycle, umbrella, piano, kitten	5
Ridge	0.75	22.14	USA, Risk, American, Danger, Perception, Threat, Exposure, Consequence, Insecurity, Unpredictability	10
SVR-RBF	0.55	29.95	USA, Risk, American, Danger, Perception, Threat, Exposure, Consequence, Insecurity, Unpredictability	10

Table 3 shows that adding risk-related words such as "risk, danger, perception" shifted all risk source vectors towards these new vector representations. However, these added words did not provide additional information to the model regarding the precise meaning of each risk source. Instead, they only contributed to a general representation of risk or danger to each source. This poses a problem since the risk sources analyzed, being participant-generated, encompass a wide variety of types, ranging from "Donald Trump" to "cholesterol". Therefore a tailored list of words is manually created, each one associated with a specific source. For instance, for the risk source "anxiety" the related word is "mental_health", and for "anthrax" the related word is "biological_warfare". Each risk source vector was then augmented by adding its corresponding "tailored" word-vector, with the goal of providing more context and representation for the original word.

The aforementioned analysis leads to slight improvements in the results compared to forecasting each rating using the original vector representation of the risk source. Table 4 depicts the difference – in metric results – between the original dataset and the "tailored" dataset. A negative R^2 or a positive RMSE indicates an enhancement of model performance due to the application of the proposed methodology.

We next study the impact of shifting vector spaces towards "American" or "European" perspectives on risk source predictions and perception, using the

Table 4. "Tailored" word embedding analysis, model difference - evaluation results

Model	R^2 Test	RMSE Test	R^2 Training	RMSE Training
Ridge	-0.00776	0.35201	-0.01774	1.87505
SVR-RBF	0.00771	-0.34655	0.00999	-4.57773

addition of “E.U.”, “U.S.” or “USA” word-vectors to risk source vectors. Table 5 presents the performance of SVR-RBF and Ridge models when each risk source is shifted toward the “E.U.”, “U.S.” or “USA” semantic space. Specifically, it shows the difference in results between using only the original dataset and from training the model with the original dataset and then testing with the "American" or "European" dataset. If R^2 (RMSE) is negative (positive), it indicates an improvement of the model performance when using the new method. According to these results, R^2 is mostly unaltered, while the RMSE metric is highly sensitive to semantic changes. For the “U.S.” dataset, RMSE is negative, indicating a deterioration of model performance when the new dataset is employed for predicting the target variable. Conversely, the shifts towards the "E.U." and "USA" contexts resulted in an enhancement of the RMSE metric. The differences can be explained by [10, 18].

Table 5. “American” & “European” word embedding, model difference – evaluation results

W-vector	Model	R^2 Test	RMSE Test	R^2 Training	RMSE Training
U.S.	Ridge Regression	0.12203**	-6.07677***	-0.00043	-16.91446
	SVR-RBF	0.04450*	-2.93738**	0.00000	-24.80804
E.U.	Ridge Regression	-0.02710**	0.37595***	-0.00043	-10.46174
	SVR-RBF	-0.13678*	6.85712*	0.0000	-15.01354
USA	Ridge Regression	-0.11477**	5.28553***	-0.00043	-5.55216
	SVR-RBF	0.07505*	-4.23009**	0.0000	-26.10074

The second part of the analysis investigates visually the geographical shift in the semantic space through a word-cloud. This word-cloud represents the words in the Word2Vec vocabulary that are closest to the newly defined “European” and “American” vectors. The output of this analysis is shown in Figure 2, representing the original risk sources, the “European” risk sources, and the “American” risk sources, respectively. In all of the plots the font size of the words is linearly proportional to the calculated risk association. In this way, bigger words, in the word clouds have an higher association with risk perception.

(a) Original risk

(b) U.S. risk

(c) E.U. risk

Fig. 2. Word cloud with high risk association

In Figure 2 some words appear consistently across all three representations. This is not unexpected as the shift in word embedding space is not dominant, a fact that has been corroborated by the previous computational findings. However, even though certain words appear in all three word-clouds, their risk associations differ since the size of the word varies. Thus a small change in the risk perception is already capture from this aspect. For instance, the term "gangland_wars" is larger in the original dataset than in the European one, and it is even larger in the American cloud. This potentially manifests a distinct perception and frequency of this risk in Europe compared to the American viewpoint.

In relation to the "European" word-cloud, compared to the original dataset, it is noticeable that the term "bombing" appears with greater frequency, and there are unique words such as "suicide". This could be interpreted as the European viewpoint placing more emphasis on these aspects of risk, as opposed to other domains. In comparison with the "American" dataset, there are fewer shared words, and those that are shared vary in size, further substantiating the tangible shift that has occurred. The "American" word-cloud displays words with the largest font size compared to the other two. It incorporates concepts more related to gangs, massacres, and migrants, which are effectively part of the American narrative. The word "republican_dissident" also emphasises the geographic shift that has occurred.

4 Conclusions and Future Work

Risk perception models analyse how individuals or groups perceive and evaluate risks, focusing on subjective assessments of the severity and likelihood of risks. By understanding how different risks are perceived by the public, we can aid designing better communication strategies and policies that are tailored to incorporate public perceptions and concerns.

In this paper we explored the use of word embedding models to infer perceptions of risk from textual data. We studied two different procedures, which included generating similar word lists for each risk source and performing algebraic operations on word vectors. Particularly the use of algebraic operations on word embedding offers potential to better represent risk perceptions, albeit with certain constraints. While adding or subtracting words can augment semantic meaning, our study demonstrated that performance could deteriorate beyond a specific limit. However, altering vector representation with context-specific words, or shifting the semantic space towards a particular geographical perspective, led to discernible changes in word associations. This suggests that algebraic operations can affect model accuracy and potentially improve the representation of risk perceptions in more targeted manipulations in certain cases. While previous methods highly relied on the use of surveys – which limited their scalability – this novel approach could offer great promise.

As future work we are looking into risk perception scores for text passages or social media posts, this is an area of application that could have a lasting impact on future risk communication and allow for much more personalised messaging.

References

1. Alhakami, A.S., Slovic, P.: A psychological study of the inverse relationship between perceived risk and perceived benefit. Risk analysis **14**(6), 1085–1096 (1994)
2. Bhatia, N., Bhatia, S.: Changes in gender stereotypes over time: A computational analysis. Psychology of Women Quarterly **45**(1), 106–125 (2021)
3. Bhatia, S.: The semantic representation of prejudice and stereotypes. Cognition **164**, 46–60 (2017)
4. Bhatia, S.: Associative judgment and vector space semantics. Psychological review **124**(1), 1 (2017b)
5. Bhatia, S.: Predicting risk perception: New insights from data science. Management Science **65**(8), 3800–3823 (2019)
6. Bhatia, S., Olivola, C.Y., Bhatia, N., Ameen, A.: Predicting leadership perception with large-scale natural language data. The Leadership Quarterly **33**, 101535 (2022)
7. Bouder, F.: Towards a concept of genetic risk tolerance: a risk analysis perspective. Journal of Risk Research **21**(2), 190–205 (2018)
8. Dryhurst, S., Schneider, C.R., Kerr, J., Freeman, A.L., Recchia, G., van der Bles, A.M., van der Linden, S.: Risk perceptions of covid-19 around the world. Journal of Risk Research **23**(7-8), 994–1006 (2020)
9. Dyer, J., Kolic, B.: Public risk perception and emotion on twitter during the covid-19 pandemic. Applied Network Science **5**, 1–32 (2020)
10. Finley, G., Farmer, S., Pakhomov, S.: What analogies reveal about word vectors and their compositionality. In: Proceedings of the 6th joint conference on lexical and computational semantics (* SEM 2017). pp. 1–11 (2017)

11. Fischhoff, B., Slovic, P., Lichtenstein, S., Read, S., Combs, B.: How safe is safe enough? a psychometric study of attitudes toward technological risks and benefits. Policy Sciences **9**(2), 127–152 (1978)
12. Gandhi, N., Zou, W., Meyer, C., Bhatia, S., Walasek, L.: Knowledge representations in health judgments. In: Proceedings of the Annual Meeting of the Cognitive Science Society (2020)
13. Johnson EJ, T.A.: Affect, generalization, and the perception of risk. J. Personality Soc. Psych. **45**, 20 (1983)
14. Levy, O., Goldberg, Y.: Neural word embedding as implicit matrix factorization. Advances in neural information processing systems **27** (2014)
15. Loewenstein, G., Weber, E., Hsee, C., Welch, N.: Risk as feelings. Psychological Bulletin **127**, 267–286 (2001)
16. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)
17. Morgan M, Fischhoff B, B.A., C, A.: Risk Communication: A Mental Models Approach, chap. Chapter 6, Development and Evaluation of Communications. Cambridge University Press (2001)
18. Nissim, M., van Noord, R., Van Der Goot, R.: Fair is better than sensational: Man is to doctor as woman is to doctor. Computational Linguistics **46**(2), 487–497 (2020)
19. Rehurek, R., Sojka, P.: Gensim–python framework for vector space modelling. NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic **3**(2) (2011)
20. Renn, O., Benighaus, C.: Perception of technological risk: insights from research and lessons for risk communication and management. Journal of Risk Research **16**, 293–313 (2013)
21. Richie, R., Zou, W., Bhatia, S.: Predicting high-level human judgment across diverse behavioral domains. Collabra: Psychology **5**(1) (2019)
22. Rohrmann, B., Renn, O.: Cross Cultural Risk Perception - A Survey of Empirical Studies, chap. Chapter 1, Risk Perception Research - An Introduction. Kluwer Academic Publishers (2000)
23. Slovic, P.: Perception of risk. Science **236**(4799), 280 (1987)
24. Slovic, P.: The Perception of Risk, chap. Chapter 25, Trust, Emotion, Sex, Politics and Science: Surveying the Risk-assessment Battlefield. Taylor & Francis (2000)
25. Slovic, P.: The Perception of Risk, chap. Chapter 24, Insensitivity to the Value of Human Life: A Study of Psychophysical Numbing. Taylor & Francis (2000)
26. Slovic, P., Weber, E.: Perception of risk posed by extreme events. Tech. rep., University of Oregon (2002)
27. Tversky, A., Kahneman, D.: Judgment under uncertainty: Heuristics and biases. Science **185**(4157), 1124–1131 (1974)
28. Wachinger, G., Renn, O., Bianchizza, C., Coates, T., De Marchi, B., Domènech, L., Jakobson, I., Kuhlicke, C., Lemkow, L., Pellizzoni, L.: Risk perception and natural hazards. CapHaz-Net WP3 Report (2010)
29. Wulff, D., Mata, R.: On the semantic representation of risk. Science advances **8**, eabm1883 (2022)
30. Zhang, Y., Jin, R., Zhou, Z.: Understanding bag-of-words model: a statistical framework. International journal of machine learning and cybernetics **1**, 43–52 (2010)
31. Zou, W., Bhatia, S.: Categories for natural objects. In: Proceedings of the Annual Meeting of the Cognitive Science Society. vol. 43 (2021)