Artificial Intelligence for Innovation Readiness Assessment

Tobias Eljasik-Swoboda
AI Architect
ONTEC AG
Wien, Austria
Tobias.Eljasik-Swoboda@ontec.at
https://orcid.org/0000-0003-2464-8461

Christian Rathgeber
Product Owner
ONTEC AG
Wien, Austria
Christian.Rathgeber@ontec.at
https://orcid.org/0000-0002-0912-163X

Rainer Hasenauer
Marketing Management Institute
WU Vienna University of Economics and Business
Wien, Austria
rainer.hasenauer@wu.ac.at
https://orcid.org/0000-0002-9894-5349

***Abstract*—There are two main necessities for an innovation to become successful. Firstly, there must be a target market willing to buy the innovation and secondly, the technology must be ready to deliver. The READINESSnavigator is a state of the art software tool that packages an innovation readiness assessment methodology that is proven to increase success chances for innovations. Usage of the READINESSnavigator has shown, that it works best if there is somebody to explain the method. This paper outlines our approach to automate this individual using artificial intelligence. To do so, it outlines the history of innovation readiness assessment methods, artificial intelligence and explains the high-level architecture of an automated innovation coach.**

Keywords—innovation management, readiness assessment, artificial intelligence, market readiness, technology readiness, AI readiness, data readiness.

I. Introduction

The difference between an innovation and an invention is that the innovation can consist out of multiple inventions and is successful [1]. According to the Triple Bottom Line model success includes profitability, ecological acceptability as well as social acceptability [2]. Literature suggests, that there are two fundamental questions for successful innovations [3]: (1) is the envisaged market ready for the innovative technology? This can also be read as *"will they buy it?"* (2) is the technology ready for the market? Simply put *"does it work and can we produce it?"* The obvious sweet spot consists out of a ready market in combination with a ready technology. This creates two follow-up problems. The first problem is the problem of measuring market and technology readiness. The second problem is to synchronize the time and content of technology- as well as market development [3].

To address these two problems, Hasenauer et al. have created a framework to manage innovations [3]. It was evaluated on 57 startups and 26 high-tech products. Adherence to the framework increased the chance to successfully innovate by 30 % and increased the chance to export said innovation by 60 %. *ONTEC A*G packaged this framework into the READINESSnavigator software in collaboration with *INiTS Universitäres Gründerservice Wien GmbH* [4]. In previous works, the READINESSnavigator has been extended with additional readiness models to specifically address issues when innovating in the context of Artificial Intelligence (AI) and data science [5]. The primary feedback on using these models, as well as the READINESSnavigator in general was that it is most useful when there is somebody to explain it to its users. In this paper, we explore possibilities of automating the explaining person by applying AI to the problem.

To do so, this contribution is structured as follows: In section II, the history of readiness models is briefly described. Section III introduces a short description of the state of the art in Artificial Intelligence while Section IV outlines our model to create an automated innovation coach as part of the READINESSnavigator software. Section V draws conclusions and provides an outlook on future works.

II. A Brief History of Readiness Models

NASA started using a technology readiness model internally in 1974. This model was shared with the world after publication in 1989 [6]. It expresses the readiness of any component on a scale from 1 to 9 that model partially overlapping degrees of progress in technology development. Levels 1 and 2 encompass basic technology research. Levels 2 and 3 are reached with research to prove feasibility of the component. Levels 3 to 5 describe the phase of technology development. On levels 5 and 6, the technology is demonstrated while levels 6 to 8 express system and subsystem development. The final levels of 8 and 9 express usage of the developed component in launch and operations.

Dent and Pettit built on this technology readiness model to include market readiness which models how ready an innovation's market is [7]. In their innovation framework, Hasenauer et al. went into much more detail and expressed technology- and market readiness as aggregate of multiple individual readiness levels that reflect different aspects of the innovation [2]. According to Hasenauer et al., market readiness consists out of four sub-readiness levels: Competitive supply readiness expresses the knowledge about and availability of competing products for the innovation project. Demand readiness reflects the knowledge about and amount of existing demand for the innovative product. The customer readiness expresses how likely customers are going to adopt the product while the product readiness expresses how readily the product and its options are prepared for widespread usage. Technology readiness on the other hand consists out of 3 distinct scales: Firstly, the intellectual property rights (IPR) readiness expresses how well the intellectual property of an innovation is protected and if there is freedom to operate (FTO) without violating somebody else's patents in the target market. Integration readiness expresses how well the technology can be integrated in the target environment likely to be encountered at a potential customer. The third pillar of technology readiness expresses how readily available manufacturing capabilities for the product are.

Hasenauer et al.'s research yielded an optimal proportion of market- and technology readiness to successfully innovate. This has been captured in Ontec's READINESSnavigator software that measures an innovations progress and provides tips for innovators on which tasks to focus next within the innovation project [4]. Eljasik-Swoboda et al. have extended the basic market- and technology readiness models with models relevant for data science and AI projects as manufacturing readiness greatly focuses on the production of physical objects [5]. There are existing AI readiness models. Intel proposed a methodology to measure a company's readiness to use AI while Capgemini assessed the readiness of different regions in the world to use AI [8][9]. Both are interesting but not directly relevant for innovation projects. Therefore, Eljasik-Swoboda et al.'s model breaks down AI readiness in two distinct readiness dimensions and data readiness into four dimensions [5]. Specification readiness expresses how well the AI's goals are specified while algorithmic readiness expresses the availability of effective and efficient algorithms to solve the problem at hand. The four data readiness dimensions express challenges with obtaining necessary training data, especially in highly regulated environment such as Europe where the GDPR emphasizes the protection of individual's data. Data existence readiness express if the necessary training data actually exists and if there is a means to capture it. Data format and data quality readiness express how well the data format is understood and if there are inherent quality issues. Such quality issues can be inherent biases that would lead to discriminating AIs. Data legal readiness express if the data can actually be used for the purpose while expert knowledge readiness model how far along one is on capturing the necessary expert data to develop the AI in question.

The dimensions of market- and technology readiness are independent of each other. This is different for AI- and data readiness where logical interdependencies exist. For instance, one cannot have high algorithmic readiness when there is insufficient specification readiness. The same is true for data existence readiness and data format and quality readiness: If the data does not exist, it is difficult to ascertain its format and quality.

After using the READINESSnavigator for a variety of AI innovation projects, almost every innovator gave the feedback, that the model is useful but difficult to understand and that it works best if there is a coach or consultant using the model to explain the process. As skilled personnel is cost intensive, we started work on our approach to create an automated innovation coach as extension of the READINESSnavigator. This approach is the subject of this contribution.

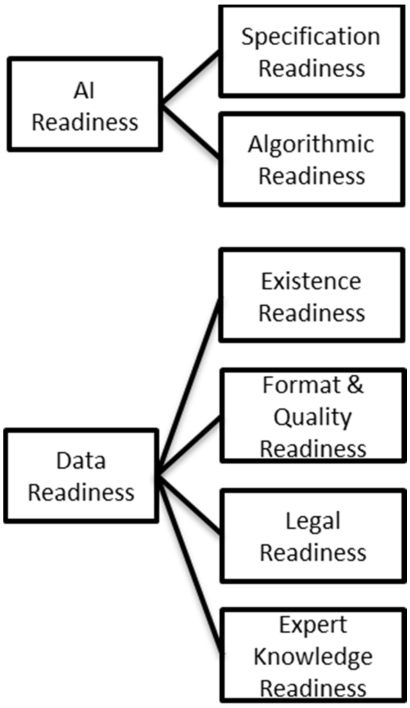

Fig. 1: READINESSnavigator AI and Data readiness levels [5]

III. State of the Art in AI

The term artificial intelligence (AI) is difficult to define because there is no consistent definition for intelligence, no matter whether it is natural, human, artificial or even quantum computing based [10]. AI is mostly defined as creating an artificial agent that performs intelligent decisions in its environment. In general, it is understood as the approach of emulating human decision making and behavior [11]. Literature differentiates between general AI, which essentially simulates the whole range of possible human interactions and weak AI which focuses on specific tasks [12]. Weak AIs have started to outperform human beings on specific tasks, most visibly in beating human world champions in more and more complex games [13]. They however are not able to easily switch from one task to another and reach any reasonable results. Transfer learning is the research field of transferring abilities in one task to another [14]. It largely remains a research field but some promising results in the area for language understanding have been reported [15].

AI has been a field of research since 1956 [16]. Its history is far longer than that with ancient Greeks describing the artificially intelligent automaton Talos [17]. Over the time of research a plethora of problems has been solved using a multitude of approaches. The two most fundamental approaches are to use rules (rule based AI, symbolic AI, expert system) or to apply machine learning to the task [18]. Machine learning can be described as methods to automatically learn patterns from available data to allow for structural forecasting [19]. The most common use cases are to classify objects, continue existing time-series data or to cluster objects. One can of course also create hybrid systems that combine machine learning based methods with rule based methods. This approach has been referred to as *neural-symbolic integration* [20].

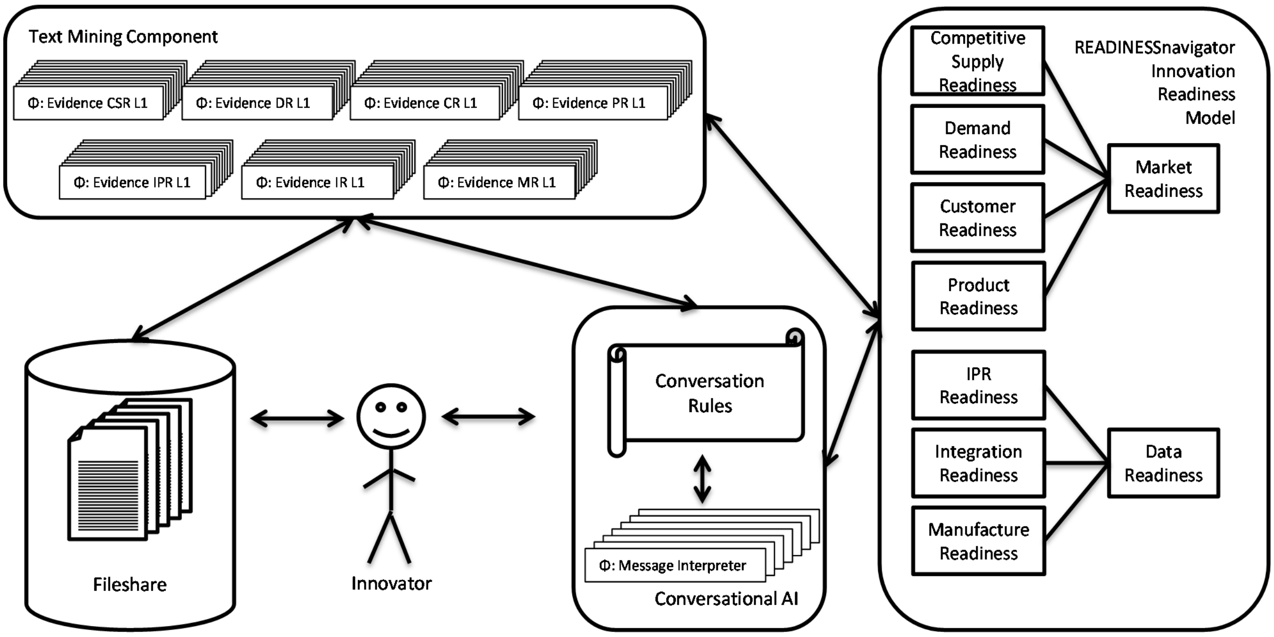

Fig. 2: Illustration of the automated innovation coach: The text mining component interfaces with the file share and provides its results to inform the readiness model as well as the conversational AI. The conversational AI asks its questions based on the results of the text mining component as well as the readiness model.

A popular method to implement machine learning is the use of simulated neural networks. These emulate the functionality of nerve clusters to achieve specific goals [21][22]. The popularity of this approach made them almost synonymous with AI even though other highly effective approaches, such as Support Vector Machines (SVMs) for classification, exist [23]. Deep learning is the idea to stack multiple machine learning approaches but also has become synonymous with having a neural network that uses many layers of neurons, which is somewhat analogous to how the mammalian neocortex is structured [24][25]. There are different settings in which machine learning can occur. When using supervised learning, the machine has labeled examples from which to learn. This is most common for classification problems. In unsupervised learning, the available example data is not labeled. Therefore the machine has to infer the regularities of the dataset on its own. Another interesting approach is referred to as active learning. Here, the learning machine queries a human being for labels of specific data points. Usually, those most different from already known data points are queried to maximize effectiveness. One can of course combine multiple learning methods with each other. For example one can use unsupervised learning to cluster data into sets of high similarity, and then query a human being to label the set as a whole in an active learning process that subsequently enables supervised learning.

A different approach to classify AIs is by the modalities they work with [18]. Such modalities are what these AIs use as input and can generate as outputs. The main modalities are audio, video and text. The latter can be extracted and be used as source to generate the first two modalities. AIs can also be multi-modal encompassing more than one of the modalities.

IV. The Automated Innovation Coach

The major criticism when using the READINESSnavigator is that the underlying model and software are complex and work best if there is an innovation coach or consultant moderating and explaining the process [5]. This section describes our approach to automate this process by using AI techniques to create an automated innovation coach based on the READINESSnavigator model.

Hasenauer et al.'s model, similar to all other readiness models described in section II, measures readiness by checking if certain steps of the innovation process have been performed or not [2]. Usually, there is evidence for performing certain steps, like filing a patent application or creating a ranking of the competition. This evidence is usually stored on a computer, ideally reachable from a network so that multiple collaborators have access to the documents. To take advantage of this evidence, the envisioned AI requires two primary components. The first component is a text analysis system that interfaces with the innovation project's file share and potentially utilized mail boxes. The second component is a conversational AI or chat-bot that converses with users to ascertain an innovation's readiness and to answer questions. In a sense, the automated innovation coach becomes a (virtual) team member and needs access to the project's documents and converses with the team by chat. Figure 2 illustrates the major components of this AI.

After granting access to the documents, identifying the current readiness level becomes a text mining problem that can be expressed as a series of text categorization (TC) problems [26]. The definition of TC is that a classifier $\Phi: (D,C) \rightarrow \{0,1\}$ approximates a target function $\Phi': (D,C) \rightarrow \{0,1\}$ as closely as possible. Here, D is the set of documents while C is the set of categories. To model our problem, the categories of *C* express evidence for the fulfillment of specific readiness levels.

As with most fields of AI, there are two fundamental ways of constructing classifiers: Rule-based and machine-learning based (see section III). As expressed by the READINESSnavigator's data readiness levels, both fields have their drawbacks in either requiring explicitly modeled expert knowledge or examples for certain types of evidence. These are difficult to obtain in sufficient quantity, especially for cutting-edge innovation projects. For instance Nawroth et al. have shown that it can take decades between a medical drug being described in natural language within research papers and it being modeled in a machine-readable fashion within an ontology [27]. There is no reason to believe this happens faster in other domains of research with the exception of computer science.

The C3 suite of TC micro-services aims to minimize required resources and explicitly modeled language when constructing a TC system [28]. They therefore provide a potential basis for constructing the text mining part of the AI as good training data examples for the individual classifiers are likely to be sparse. Especially in the domain of Natural Language Processing (NLP), transfer learning has been demonstrated [15]. Word embeddings capture the meaning of words in a high-dimensional coordinate system using an unsupervised learning process [29]. This makes them a valuable component for any TC application that has to deal with sparse examples as much larger knowledge sources, like the entirety of Wikipedia, can be used to bootstrap language understanding. Another common technique to achieve high effectiveness TC results is to employ classifier committees [26]. In such committees, the results of multiple classifiers are used in a combination function that generates the overall classification decision. The first stage of the AI is therefore going to be constructed using the afore-mentioned state of the art methods at our disposal.

The second part of the automated innovation coach is a conversational agent or chat-bot. It forms a natural language interface to communicate with the READINESSnavigator users. Normally, a contemporary chat-bot tries to understand an order given by the user. For our purpose, the pattern is different, as the automated innovation coach is leading and steering the conversation to ascertain the current readiness levels and give tips on the next tasks. To do so efficiently, the chat-bot part can make use of the results from the text mining part of the AI by asking the user if it correctly understood the evidence within a file, that certain tasks have been performed. Besides ascertaining the innovation's readiness level, this is also a way to further train the TC classifiers used in the text mining part in an active learning fashion by using sentences like: *"I have read in file XYZ.docx that you filed for a patent application. Is that true?"* or *"I read in file ABC.xlsx that you ranked competitors. Have you finished this part of your market research?"* When the user answers *yes* or *no*, the AI can determine which readiness level is currently appropriate. It can also use files *XYZ.docx* and *ABC.xlsx* as new examples to extend the target function for their respective classifiers. This way, the text-mining part of the AI is continuously improved by users interacting with the second part. Improved text mining effectiveness will subsequently increase the quality of questions asked by the conversational AI to the innovator.

When an innovator's project is assessed by the AI, it is only natural for the innovator to want to achieve high scores within the readiness model. To address this, the AI needs to do more than ask if a certain task has been performed but should query the outcome to make sure that the specific level has been reached. Reliably measuring if an explanation is good is actually another TC problem which is very difficult to solve, as good explanations for one innovation might not be good explanations for another innovation. This is similar to the discipline of argument mining and argument stance detection for which has been shown, that topic specificity has a significant impact on performance [28][30]. This means that if the AI was trained on innovations from one technical domain its performance in a completely different field is likely to be negatively impacted. The same is true for human experts that have to work outside their field of expertise. Another difficulty is the fact, that innovation projects carry an inherent fuzziness and uncertainty.

Usually, chat-bots parse messages written by users in two ways: Firstly, they attempt to extract named entities, which are concepts contained in a sentence. For instance *"pick up the glass from the table"* can have the entities *glass* and *table* for which is modeled, that table can be a location. Another task is to detect intent which is, what the user wants the AI to do. This can actually also be modeled as TC problem in which different intents that the AI can fulfill are the categories while the sentences are documents belonging to that category.

Besides asking the user about the innovation's readiness, the AI should also be capable to answer questions about the process and model. This way it can act as an automated glossary and knowledge base.

Creating such a system is a daunting task given the multitude of different classes that the multiple classifiers need to be able to assign documents (or short chat messages) to. If one focuses on market- and technology readiness alone, there are seven distinct readiness dimensions each with nine levels. That are 63 classes for which there can be evidence to have reached them. These classes are also transitively related within their scale. For instance, evidence for being on level three implies that the necessary tasks for level two have also been performed. A similar number of reply-classes within the conversational part of the AI are also likely. Additionally, the AI must be able to understand, if the user has a question and must be able to answer this question. This pushes the number of relevant classes for which sufficient examples and/or rules are needed well into the hundreds. Acquiring this training data is a crucial sub-project of implementing the automated innovation coach. If the automated innovation coach is to make use of other readiness models, like AI- and Data Readiness, the necessary TC classifiers and conversational rules must be created in addition to the existing ones. As mentioned in section II, AI- and Data Readiness have interdependencies that also need to be implemented.

When using Kaufmann et al.'s BDMC planning method for Big Data projects, the following work packages for implementing this AI emerge [31]: Firstly, the integration of the AI in business processes must be modeled (*effectuation*). Scrum user stories along with process modeling methods are appropriate tools for this task. Secondly, in the field of applied data science, appropriate classification algorithms must be selected (*analysis*). Additionally, conversation ruled need to be modeled in order to enable the conversational AI. Thirdly, necessary training data to create a baseline system must be identified and gathered (*data integration* and *datafication*). Given the amount of classes that both parts of the AI need to implement, this is a huge project requiring data from multiple innovation projects in order to achieve reasonable results. This task is made even more herculean as specifics differ from region to region as different languages are spoken and authorities require different forms. This underlines the importance of transfer learning and bootstrapping techniques for this endeavor.

Besides these four more business oriented work packages, respective engineering work packages have to be implemented accordingly. This means, that data storage and integration systems need to be developed or procured before the classifiers and conversational AI features can be implemented. Additionally, the chat bot interface needs to be implemented. There is also the decision if there is a specific automated innovation coach chat client that innovators are using or if the system should integrate into existing chat software.

Each of these activities will raise the technology- data- and AI readiness of the automated innovation coach. Following Hasenauer et al.'s method, adequate activities to raise market readiness also need to be performed in order to create a successful innovation [2].

V. Conclusions

This paper provides the reader with summaries of the state of the art in two highly important fields of research: Innovation readiness assessment and artificial intelligence. As shown, innovation assessment software using a model that is proven to increase success rates exists in form of the READINESSnavigator. It has been tested with multiple innovation projects. The overall feedback is that users require a consultant or coach to use it. Therefore we propose an automated innovation coach.

Our contributions are threefold: Firstly, we lined out the high level architecture of an automated innovation coach. Secondly, we proposed specific components to implement this AI while thirdly; we assessed the different sub projects necessary to implement this system using a state of the art planning method.

We find, that the proposed AI is definitively feasible but will require a great amount of regionalization effort as well as industry specification. For instance, a system that does well in the domain of software engineering can generate poor results when used in the pharmaceutical industry.

References

[1] J. Schumpeter, Business cycles. A theoretical, historical and statistical analysis of the capitalist process. New York/London, 1939.

[2] R. Hasenauer, A. Gschöpf, C. Weber, "Technology Readiness, Market Readiness and the Triple Bottom Line: An Empirical Analysis of Innovating Startups in an Incubator," Proceedings of PICMET 2016, pp. 1387-1428

[3] R. Hasenauer, C. M. Weber, P. Filo, J. Orgonas, "Managing Technology Push through Marketing Testbeds: The Case of the Hi-Tech Center in Vienna, Austria," MANAGEMENT OF THE TECHNOLOGY AGE Proceedings of PICMET 2015, pp. 99 – 127

[4] Ontec AG, The READINESSnavigator. Retrieved April 10, 2019 from https://ontec.at/the-readinessnavigator/

[5] T. Eljasik-Swoboda, C. Rathgeber, R. Hasenauer, "Assessing technology readiness for artificial intelligence and machine learning based innovations," Proceedings of the 8th international conference on data science, technology and applications (DATA 2019), in press.

[6] S. R. Sadin, F. P. Povinelli, R. Rosen, "The NASA technology push towards future space mission systems," Acta Astronautica, 1989, pp. 73-77.

[7] D. Dent, B. Pettit, "Technology and market readiness levels", white paper Dent associates, 2011

[8] Intel, "The AI readiness model, judging an organization's Ability to generate business value from artifical intelligence." Retrieved March 28, 2019 from https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/ai-readiness-model-whitepaper.pdf

[9] D. Tinholt, E. van Niel, C. van Kraaji, M. Knödler, "Artificial intelligence benchmark", Retrieved March 29, 2019 from https://www.capgemini.com/wp-content/uploads/2018/07/AI-Readiness-Benchmark-POV.pdf

[10] W. Zhang, "From equilibrium-based business intelligence to information conservational quantum-fuzzy cryptography - A cellular transformation of bipolar fuzzy sets to quantum intelligence machinery," IEEE Trans. on Fuzzy Systems. IEEE Explore, March 24, 2017. 10.1109/TFUZZ.2017.2687408.

[11] S. J. Russel, P. Norvig, Artificial intelligence, a modern approach, 3rd edition, Upper Saddle River, New Jersey: Prentice Hall, ISBN 978-0-13-604259-4, 2009

[12] J. R. Searle, "Minds, brains and programs," Behaviour and Brain Sciences 3, 1980, pp. 417-424

[13] O. Vinyals et al., "AlphaStar: Mastering the real-time strategy game StarCraft II," 2019, Retrieved May 26, 2019 from https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/

[14] S. J. Pan, Q. Yang, "A survey on transfer learning," IEEE Transactions of Knowledge and Data Engineering, (vol. 22, issue: 10),October 2009, pp. 1345 – 1359

[15] K. Clark, M. Luong, C. D. Manning, Q. V. Le, "Semi-supervised sequence modeling with cross-view training," Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018, pp. 1914–1925

[16] J. Moor, "T he Dartmouth College Artificial Intelligence Conference: The Next Fifty years, " AI Magazine, vol 27, no. 4, 2006, pp. 87-91

[17] Bibliotheke 1.9.26

[18] J. Haugenland, Artificial Intelligence, the Very Idea, Cambridge, Mass: MIT Press, ISBN 0-262-08153-9, 1985

[19] M. Mohri, A. Rostamizadeh, A. Talwalkar, Foundations of Machine Learning, MIT Press, Cambridge, Massachusetts, USA, 2012

[20] S. Bader, P. Hitzler, Dimensions of neural-symbolic integration - a structured survey. arXiv:cs/0511042 [cs.AI] 2005.

[21] W. S. McCulloch, W. Pitts, "A logical calculus of the ideas immanent in nervous activity," Bulletin of Mathematical Biophysics, 5, 1943, pp. 115-133

[22] H.Helbig, A. Scherer, Kurs 1830: Neuronale Netze, University of Hagen, Germany, 2011

[23] V. N. Vapnik, A. Ya. Chervonenkis, On a class of algorithms of learning pattern recognition, Framework of the Generalised Portrait Method, 1964, Об одном классе алгоритмов обучения распознаванию образов, Автоматика и телемеханика

[24] I. Arel, D. C. Rose, T. P. Karnowski, "Deep machine learning – A new frontier in artificial intelligence research," IEEE Computational Intelligence Magazine, USA, November 2010 pp. 13 – 18,

[25] J. Hawkins, S.Blakelslee, On Intelligence, Times Books, USA, 2004

[26] F. Sebastiani, “Machine learning in automated text categorization,” ACM Computing Surveys, Vol 34, pp. 1-47, 2002

[27] C. Nawroth, F. Engel, T. Eljasik-Swoboda, M. Hemmje, “Towards enabling emerging named entity recognition as a clinical information and argumentation support,” Proceedings of the 7th International conference on Data Science, Technology and Applications (DATA 2018), ISBN: 978-989-758-318-6, 2018, pp. 47-55

[28] T. Eljasik-Swoboda, F. Engel, M. Hemmje, “Using topic specific features for argument stance recognition,” Proceedings of the 8th international conference on data science, technology and applications (DATA 2019), in press.

[29] T. Mikolov, K. Chen, G. Corrado, J. Dean, “Efficient estimation of word representation in vector space” Proceedings of Workshop at ICLR 2013

[30] S. Mohammad, S. Kiritchenko, P. Sobhani, X.Zhu, C. Cherry, “Semeval-2016 task 6: Detecting stance in tweets,” Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016) pp. 31-41

[31] M. Kaufmann, T. Eljasik-Swoboda, C. Nawroth, K. Berwind, M. Bornschlegl, M. Hemmje The Big Data Management Canvas Method, In: *Proceedings of the 6th international conference on data science, technology and applications (DATA 2017)*, DOI: 10.5220/0006397101490156, 2017, pp. 149-156