



Hateful Messages: A Conversational Data Set of Hate Speech Produced by Adolescents on Discord

Jan Fillies, Silvio Peikert, and Adrian Paschke

Abstract With the rise of social media, an increase of hateful content online can be observed. Even though the understanding and definitions of hate speech vary, platforms, communities, and legislature all acknowledge the challenge. Adolescents are a new and active group of social media users. The majority of adolescents experience or witness online hate speech. Research in the field of automated hate speech classification has been on the rise and focuses on aspects such as bias, generalizability, and performance. To increase generalizability and performance, it is important to understand biases within the data. This research addresses the bias of youth language within hate speech classification and contributes by providing a modern and anonymized hate speech youth language data set consisting of 88.395 annotated chat messages. The data set consists of publicly available online messages from the chat platform Discord. For 35.553 messages, the user profiles provided age annotations, setting the average author age to under 20 years old. 6,4% of the total messages were classified as hate speech using the annotation schema, which was adapted for this data set.

Keywords Hate speech · Youth language · Bias · Data set · NLP

1 Introduction

Research shows that there are differences within the language used by age groups online [16]. Most teenagers within the United States use social media [18]. Between January 2020 and March 2020 Facebook removed 9.6 million posts containing hate speech.¹ As of today, it is clear that social media is used often and frequently by adolescents. Hate speech and its algorithmic detection has had an increasing interest in social media platforms such as Facebook.² This development is especially supported by the harmful effects hate speech has on its recipients [15].

Based on the research that identifies a difference in language and topic in conversations between adolescents and adults [16], it is necessary to build a database of youth language to explore the impact the language has upon algorithmic hate speech detection. This research lays the groundwork to close the gap by introducing an annotated hate speech data set focusing on youth language. The data set was collected in a real-world environment between March 2021 and June 2022. It provides the scientific community a modern corpus that can be used to evaluate the bias in existing classification algorithms for hate speech and further train domain specific algorithms to the setting of hate speech within the online chat conversations of adolescents. This modern real-world data set overcomes the status quo of identifying hate speech connected to geolocation and introduces the view that hate speech is also unique to international group conversations on the internet. It provides the field not just an age annotated data set, but introduces data collected from the chat platform discord in connection with an unseen real-world

¹Forbes Media, accessed on 23.03.2023, <https://www.forbes.com/sites/niallmccarthy/2020/05/13/facebook-removes-record-number-of-hate-speech-posts-infographic/?sh=20c0ef983035>.

²Meta AI, accessed on 23.03.2023, <https://ai.facebook.com/blog/how-facebook-uses-super-efficient-ai-models-to-detect-hate-speech>.

J. Fillies (✉) · A. Paschke
Institut für Angewandte Informatik, Goedelerring 9, 04109 Leipzig, Germany
e-mail: fillies@infai.org

Freie Universität Berlin, Kaiserswerther Str. 16-18, 14195 Berlin, Germany

S. Peikert · A. Paschke
Fraunhofer-Institut für Offene Kommunikationssysteme FOKUS, Kaiserin-Augusta-Allee 31, 10589 Berlin, Germany

chat conversation spreading over a time period of 15 months. Contrary to similar data sets, this research does not focus on filtered Tweets or comments. It is available, on request, for further research at Zenodo.org.³

2 Related Literature

Annotated data sets in the field of hate speech detection are available (e.g. [7, 20]) though there are fewer multilingual data sets with fitting annotations available [4]. Hate speech data sets have many annotation schemes [4], from binary to multi-class hierarchies. Other universal annotation schemes exist [11] and are deployed in hate speech annotation or similar contexts, such as cyberbullying [17]. It is difficult to obtain hate speech data sets and hate speech information within adolescents. Research focusing on cyberbullying in pre-teens can be found in the research of Sprugnoli et al. [17]. They created a data set containing annotated hate speech chat conversations between Italian high school students. The data set was created in an experimental setting to foster a safe environment, moderated by the researcher. In 2019, Menini et al. [12] presented a monitoring system for cyberbullying. They identified a network of multiple high schools, their students and their friends in the United Kingdom’s Instagram community. In 2020, Wijesiriwardene et al. [21] published a multimodal data set containing Tweets labeled for toxic social media interactions. The data set was created focusing on American high school students. In 2011, Bayzick [1] created a data set consisting of messages from MySpace.com. They organized the messages into groups of ten and annotated the messages, some of which contained cyberbullying. The data set includes self-provided information about the age of the author. Dadvar et al. [5] showed that user context, including attributes such as age, gender, and cyberbullying history of the user, improves the detection of cyberbullying. Chen et al. [3] analyzes the personal writing style and other user specific attributes to identify the potential of the user spreading hate speech. To give closer understanding of the mentioned data sets, the key attributes are displayed in Table 1.

Natural language processing is required to be algorithmically fair and fitted to many social groups [2]. Classification algorithms can be biased towards many minority groups of people, for example, bias by gender [10] or race [9]. Even though age is a known source for bias in data [8], it is not widely analyzed in pretrained networks. To counter these biases, there are different approaches. Some focus on single domains, or tasks, via fine-tuning using new data [14].

As shown, there are numerous publicly available hate speech data sets, some of which address the adolescent audience and are annotated for cyberbullying or hate speech. However, there are three missing fields. Firstly, our research focuses on online conversations, not comments under posts. Secondly, the introduced dataset is drawn out of a real-world setting and not created in an artificial experimental setting. Thirdly, this data focuses on an international online English-speaking community, not a regional community.

The aspect of time needs to be considered, as it is necessary to collect and analyze the data sets from a recent time frame, considering the shifts in topic and language.

Within the last five years, no real-world hate speech data set containing online conversations of adolescents could be found.

3 Hate Speech Data Set

3.1 Methodology

Vidgen and Derczynski [19] recommend addressing the following points when creating an abusive content data set: purpose, explore new source, clear taxonomy, develop guidelines iteratively with your annotators, and data statement.

In the context of this research, the points are addressed as follows:

- Purpose: the purpose of this data set is to build a base for validation and improve hate speech detection within youth language, further explained in Sect. 1 (“Introduction”).
- Explore new source: no hate speech Discord data set has been discovered so far.
- Clear taxonomy: the used taxonomy is based on Paasch-Colberg et al. [13] and is described in the Sect. 3.3 (“Annotation Guidelines”).
- Develop guidelines iteratively with your annotators: this has been done and is described in the Sect. 3.4 (“Annotations Procedure”).
- Data Statement: a data statement is provided using the format suggested by Gebru et al. [6].

³Zenodo.org, accessed on 13.04.2023, <https://doi.org/10.5281/zenodo.7824768>.

Table 1 Table of the mentioned existing data sets and the new introduced data set

Title	Author	Topic	Timeframe	Language	Platform	Number of entries	Percentage of hate speech (%)	Number of participants	Age of authors available
Hate speech dataset from a white supremacy forum	Gibert et al. [7]	Hate speech	2002–2017	English	Stormfront	9,916	11,29	Unknown	No
Creating a whatsapp dataset to study pre-teen cyberbullying	Sprugnoli et al. [17]	Cyberbullying	2018 or earlier	Italian	WhatsApp	ca. 2927 (14,600 tokens)	41,1	70	Yes
Alone: A dataset for toxic behavior among adolescents on twitter	Wijesirwardene et al. [21]	Toxic Behavior	up to May 2018	English	Twitter	688 Interactions	17,15	Unknown	No
Detecting the presence of cyberbullying using computer software	Bayzick et al. [1]	Cyberbullying	2011	English	Myspace	4,813	21,0	Unknown	Yes
Hateful Symbols or Hateful People? Predictive Features for Hate Speech Detection on Twitter	Wasceem and Hovy [20]	Cyberbullying	before 2012	English	Youtube	4,626	9,7	3,858	Yes
Hateful Messages: A Conversational Data Set of Hate Speech produced by Adolescents on Discord	Removed for Review	Hate Speech	3.2021–06.2022	English	Discord	88,395	6,42	249	Yes

3.2 Data Identification

Discord is a chat platform that provides spaces for communication between users. These servers are publicly available and, if configured, could be joined by anybody interested. There are public lists available for existing chat servers, filterable by language, name, and topics. The research project pre-selected a list of servers based on their names and descriptions. The pre-selected servers were further evaluated following these five criteria: Firstly, conversation language in English. Secondly, high appearance of general derogatory terms through a simple key word search. Thirdly, amount of active users. Fourthly, amount of messages sent in the group chat. And lastly, available information on the age of the users. The chosen server fulfills these criteria and was exported for the purpose of furthering research within the topic.

3.3 Annotation Guidelines

The annotation guidelines were developed iteratively with and by the annotators, ensuring a high understanding of the process and definitions. A common definition of hate speech was established as follows: a statement is viewed as hate speech if it is directed towards a group or an individual of a group with the characteristic of excluding or stigmatizing. A statement is further considered hate speech if it is hostile or implies the desire to harm or incite violence. Based on Paasch-Colberg et al. [13] a new annotation schema was defined including descriptions and examples. All nine categories of the schema are explained in the following Table 2.

3.4 Annotations Procedure

Five annotators have annotated the data set. The team of annotators consisted of Bachelor and Master computer science students, and the average age of the members was 29 years. The ages varied from 21 to 58 years. The team consisted of two female and three male annotators. For four out of five members, the ethnic background and mother tongue was German. One

Table 2 Table of annotated classes

Label	Definition	Example
No Hate	Positive and neutral conversations, but also criticism, rejection and disliking	“I don’t like the new chairs”
Negative Stereotyping	Generalizations in which hurtful intent is a central motivator	“All blondes are stupid”
Dehumanization	Non-human character traits are attributed to humans, people, or groups. They are paired with elements not belonging to the human species	“Asian rats”
Violence and Killing	Endorse, glorify, or fantasize about violence or killing, the explicit call to violence and murder	“The only thing that [...] helps is pure violence”
Equation	Associates’ people and groups of people with negative characteristics	“Poor = Africa”
Norm. of Exi. Dis.	Existing discrimination is downplayed and or manifested	“No wonder blacks are treated this way.”
Disguise as Irony	Disguises negative statements as irony and downplays them as humor	“In my next life I’ll be a social welfare recipient, there I can chill”
Harmful Slander	All other forms of insults and hurtful statements that cannot be classified into the previous labels	“I don’t know any “normal” Jew”
Skip	Comments that cannot be understood and not assigned to labels above, due to a linguistic or symbolic barrier	Examples of this are emojis or Asian characters

Table 3 The data statement

Characteristic	Description
CURATION RATIONALE	The data set consists of Discord chat messages. The chat room was selected due to the high level of hate speech, the age of the authors and the platform’s young user base
LANGUAGE VARIETY	The messages are online, written in English, and provided by a multinational setting, predominantly Europe, United Kingdom and America. It is partly youth language
SPEAKER DEMOGRAPHIC	There are 249 unique author IDs. In the data set and, the average age of the user who provided information was under 20. Out of the users who provided information, 19 are female and 22 male. Out of the users who provided information, 15 were from the UK, 18 from the USA, 17 from Europe, and 13 from other countries. Disordered speech is present
ANNOTATOR DEMOGRAPHIC	Five annotators were used. They are full time students with an age range between 21–58, average age 29. The group consisted of three males, two females, four native German speakers, one native Albanian speaker. Four members were German natives and one member was an native Albanian. Lastly, one member holds a degree as a translator
SPEECH SITUATION	The data set was collected between 26.03.21 and 15.06.2022. It consists of written unscripted asynchronous messages. And the intended audience were the other participants in the chat
TEXT CHARACTERISTICS	It is an everyday conversational setting which guarantees its members free speech

annotator’s mother tongue and ethnic background was Albanian. One group member brought domain-specific knowledge through a degree as a translator.

The data set was divided into equal parts so that simultaneous annotation was possible. An annotation tool was used. Messages that were uncertain or not clear for the annotator were jointly annotated in the peer review process.

3.5 Data Statement

The statement is provided in Table 3 and based on Gebru et al. [6]. The classes “RECORDING QUALITY”, “OTHER”, and “PROVENANCE APPENDIX” were not available or applicable for the data set.

4 Data Evaluation

The evaluation is oriented on the data evaluation performed by de Gibert et al. [7]. The hate speech data set contains online conversations of adolescents on Discord, written in English. All messages have time stamps and author id’s attached. The users are from different countries, mainly the USA, the EU, and GB. The data set consists of 88.395 messages. Out of these, 35.553 have an age annotation available and 52.895 do not. Table 4 shows the distribution of messages over all nine annotated categories. It is visible that the classes are not balanced, with most classes having less than 1.000 messages assigned and the non hate speech class dominates with over 87% of the data set. The whole data set contains 6,41% hateful messages and the age annotated subset contains 5,07% hateful messages. There are 9 members in the age group 14–17, 19 members in the age group 18–25, and 4 members in the age group 26+.

The distribution of the comments in relation to the 249 registered users shows directly that 90% of all messages were written by 30 users. On average, a user posts 2662 messages to the chat room, and 90% of the hate speech was produced by 85 users with an average user posting 60 hateful messages. It was discovered that one highly active user wrote 33.372 messages, accounting for 2488 or 43,87% of the hateful messages. This user is not classified as a chatbot and did not provide data regarding their age, therefore is not influencing the age annotated data sub set. In the age annotated sub set, 90% of the messages are attributed to 10 users, with 35 users providing 90% of the hateful messages in the sub set, sending an average of 54 hateful messages per user.

Based on de Gibert et al. [7] a hate score (HS) for each word (w) has been calculated as a simple way to create insight into the context in which a word appears. For this, all hateful classes (hate) have been combined into one. A Pointwise Mutual Information (PMI) score has been calculated between each word, the hateful class, and the non-hateful (nohate) class. The PMI score considers the relationship between the joint probability (P) and the product of the individual probabilities of two instances. Based on this, the hate score of each word was calculated by subtracting the non-hateful PMI from the hateful PMI.

Table 4 Table with distribution of labels

Label	All messages	Percentage of all messages	Messages with age attributes	Percentage of messages with age attributes
No Hate	77.034	87,15	31.037	87,30
Negative Stereotyping	768	0,9	232	0,7
Dehumanization	499	0,6	165	0,5
Violence and Killing	651	0,7	205	0,6
Equation	124	0,1	27	0,08
Norm. of Exi. Dis.	145	0,2	40	0,1
Disguise as Irony	181	0,2	60	0,2
Harmful Slander	3.303	3,7	1.075	3,0
Skip	5.689	6,4	2.712	7,6
Total	88.395	100	35.553	100

Table 5 Table with the most positive (hateful) and most negative (least hateful) HS

Hateful terms	Hate score	Non-hateful terms	Hate score
ni**er	11,0493	👁👁	-4.4188
b*tch	9,6315	fair	-4.0526
fa*got	9,2531	🤔	-3.6197
h*e	9,1535	plus	-3.4836
kys	7,4165	huh	-3.4132

$$PMI(w, hate) = \log_2 \frac{P(w, hate)}{P(w)P(hate)} \quad (1)$$

$$HS(w) = PMI(w, hate) - PMI(w, nohate) \quad (2)$$

The Table 5 displays the five words with the highest and lowest correlation to the hateful class. The hateful terms were modified to reduce their impact on the reader. The hateful words are strong, well known slurs and defamations. The least correlated term is an emoticon followed by general words including the word “fair”. In these ten words, a youthful character can be identified, for example, by the extensive emoticon use and the usage of hateful modern abbreviation such as “kys”. Overall, there are 3542 words with a negative Hate Score and 1608 words with a positive Hate Score.

5 Discussion

It is important to understand that online chat rooms, like the one evaluated, are an ecosystem, meaning the users influence each other in language and topic. Therefore, this youth language corpus might be fundamentally different from other youth language corpora. The approach used in this paper is an important contribution to the world of youth language data sets due to the use of modern language with the provided self-identification, putting the general discussion in an age range under 20.

There is no way to guarantee that the given age ranges are truthful. This limitation cannot be easily circumvented in a non-experimental setting if a real-world data set guaranteeing data protection is wanted.

During the creation of the data set, which followed the official recommendation, the classes of the annotation schema were developed in communication with the annotators. This arguably led to an improvable annotation schema. Similarly, it is arguable that the used form of communal decision on uncertain classifications is not transparent and missing inter-annotator agreements. The work is open to updates and changes in the class definitions or reannotation.

The data set is heavily unbalanced in regard to authorship of the messages and the labeled hate speech classes. Furthermore, it is clear that the number of authors in general is small. These factors are due to the real-life character of the data set and are common problems in the field of hate speech.

It is important to start collecting and publishing subdomain data sets to understand the difference and uniqueness of languages in these groups and to best identify performing hate speech classification algorithms.

6 Conclusion and Future Work

This paper collected and annotated a youth language data set containing 88,395 online chat messages. It provides an unseen amount of unfiltered annotated conversational data between multiple international authors, novel to the domain of hate speech detection.

Of the 249 unique users, 31 provided information about their age, averaging to under 20 years. The data set is labeled into nine classes in the field of hate speech. A data analysis has been conducted and influential terms for the “Hate” and “No Hate” classes have been established. A data statement is provided. The data set is available for scientific research.

This research is the ground for further work in the field of hate speech detection within youth language. The next step is to identify a non youth language online chat conversation and annotate it for hate speech, comparing the differences in language and use of hateful terms. Overall, the research can be used to train youth language specific hate speech classifiers and identify the influence of youth language on their performance.

This research opens up the possibility to analyze the bias youth language introduces into existing pretrained hate speech detection models. Furthermore, the generalizability of existing prediction models can be tested and increased by using this new data set. Lastly, topics of interest within groups of adolescent can be established and compared to other communities or research results.

Acknowledgements This research was supported by the Citizens, Equality, Rights and Values (CERV) Programme under Grand Agreement No. 101049342.

References

1. Bayzick, J., Kontostathis, A., Edwards, L.: Detecting the presence of cyberbullying using computer software (2011)
2. Blodgett, S.L., O’Connor, B.: Racial Disparity in Natural Language Processing: A Case Study of Social Media African-American English. arXiv e-prints [arXiv:1707.00061](https://arxiv.org/abs/1707.00061) (Jun 2017). <https://doi.org/10.48550/arXiv.1707.00061>
3. Chen, Y., Zhou, Y., Zhu, S., Xu, H.: Detecting offensive language in social media to protect adolescent online safety. In: 2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing, pp. 71–80 (2012). <https://doi.org/10.1109/SocialCom-PASSAT.2012.55>
4. Chung, Y.L., Kuzmenko, E., Tekiroglu, S.S., Guerini, M.: CONAN – COUNTER NARRATIVES THROUGH NICHE SOURCING: A MULTILINGUAL DATASET OF RESPONSES TO FIGHT ONLINE HATE SPEECH. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 2819–2829. Association for Computational Linguistics, Florence, Italy (Jul 2019). <https://doi.org/10.18653/v1/P19-1271>, <https://aclanthology.org/P19-1271>
5. Dadvar, M., Trieschnigg, D., Ordelman, R., de Jong, F.: Improving cyberbullying detection with user context, pp. 693–696 (1 2013). https://doi.org/10.1007/978-3-642-36973-5_62
6. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J.W., Wallach, H., III, H.D., Crawford, K.: Datasheets for datasets. *Commun. ACM* **64**(12), 86–92 (2021). <https://doi.org/10.1145/3458723>, <https://doi.org/10.1145/3458723>
7. de Gibert, O., Perez, N., García-Pablos, A., Cuadros, M.: Hate speech dataset from a white supremacy forum. In: Proceedings of the 2nd Workshop on Abusive Language Online (ALW2), pp. 11–20. Association for Computational Linguistics, Brussels, Belgium (Oct 2018). <https://doi.org/10.18653/v1/W18-5102>, <https://aclanthology.org/W18-5102>
8. Hovy, D., Prabhumoye, S.: Five sources of bias in natural language processing. *Lang. Linguist. Compass* **15**(8), e12432 (2021). <https://doi.org/10.1111/lnc3.12432>, <https://compass.onlinelibrary.wiley.com/doi/abs/10.1111/lnc3.12432>
9. Kennedy, B., Jin, X., Mostafazadeh Davani, A., Dehghani, M., Ren, X.: Contextualizing hate speech classifiers with post-hoc explanation. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 5435–5442. Association for Computational Linguistics, Online (Jul 2020). <https://doi.org/10.18653/v1/2020.acl-main.483>, <https://aclanthology.org/2020.acl-main.483>
10. Kurita, K., Vyas, N., Pareek, A., Black, A.W., Tsvetkov, Y.: Measuring bias in contextualized word representations. In: Proceedings of the First Workshop on Gender Bias in Natural Language Processing, pp. 166–172. Association for Computational Linguistics, Florence, Italy (Aug 2019). <https://doi.org/10.18653/v1/W19-3823>, <https://aclanthology.org/W19-3823>
11. Lenzi, V.B., Moretti, G., Sprugnoli, R.: Cat: the celct annotation tool. In: Chair, N.C.C., Choukri, K., Declerck, T., Doğan, M.U., Maegaard, B., Mariani, J., Moreno, A., Odijk, J., Piperidis, S. (eds.) Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12). European Language Resources Association (ELRA), Istanbul, Turkey (May 2012)
12. Menini, S., Moretti, G., Corazza, M., Cabrio, E., Tonelli, S., Villata, S.: A system to monitor cyberbullying based on message classification and social network analysis. In: Proceedings of the Third Workshop on Abusive Language Online, pp. 105–110. Association for Computational Linguistics, Florence, Italy (Aug 2019). <https://doi.org/10.18653/v1/W19-3511>, <https://aclanthology.org/W19-3511>
13. Paasch-Colberg, S., Strippel, C., Trebbe, J., Emmer, M.: From insult to hate speech: mapping offensive language in German user comments on immigration. *Media Commun.* **9**(1), 171–180 (2021). <https://doi.org/10.17645/mac.v9i1.3399>

14. Park, J.H., Shin, J., Fung, P.: Reducing gender bias in abusive language detection. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 2799–2804. Association for Computational Linguistics, Brussels, Belgium (Oct–Nov 2018). <https://doi.org/10.18653/v1/D18-1302>
15. Saha, K., Chandrasekharan, E., De Choudhury, M.: Prevalence and psychological effects of hateful speech in online college communities. In: Proceedings of the 10th ACM Conference on Web Science, pp. 255–264. WebSci '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3292522.3326032>
16. Schwartz, H.A., Eichstaedt, J.C., Kern, M.L., Dziurzynski, L., Ramones, S.M., Agrawal, M., Shah, A., Kosinski, M., Stillwell, D., Seligman, M.E.P., Ungar, L.H.: Personality, gender, and age in the language of social media: the open-vocabulary approach. PLOS ONE **8**(9), 1–16 (09 2013). <https://doi.org/10.1371/journal.pone.0073791>
17. Sprugnoli, R., Menini, S., Tonelli, S., Oncini, F., Piras, E.: Creating a WhatsApp dataset to study pre-teen cyberbullying. In: Proceedings of the 2nd Workshop on Abusive Language Online (ALW2), pp. 51–59. Association for Computational Linguistics, Brussels, Belgium (Oct 2018). <https://doi.org/10.18653/v1/W18-5107>, <https://aclanthology.org/W18-5107>
18. Thapa, R., Subedi, S.: Social media and depression. J. Psychiatr. Assoc. Nepal **7**(2), 1–4 (2018). <https://doi.org/10.3126/jpan.v7i2.24607>, <https://www.nepjol.info/index.php/JPAN/article/view/24607>
19. Vidgen, B., Derczynski, L.: Directions in abusive language training data, a systematic review: Garbage in, garbage out. PLOS ONE **15**(12), 1–32 (12 2021). <https://doi.org/10.1371/journal.pone.0243300>
20. Waseem, Z., Hovy, D.: Hateful symbols or hateful people? predictive features for hate speech detection on Twitter. In: Proceedings of the NAACL Student Research Workshop, pp. 88–93. Association for Computational Linguistics, San Diego, California (Jun 2016). <https://doi.org/10.18653/v1/N16-2013>, <https://aclanthology.org/N16-2013>
21. Wijesiriwardene, T., Inan, H., Kursuncu, U., Gaur, M., Shalin, V.L., Thirunarayan, K., Sheth, A., Arpinar, I.B.: Alone: a dataset for toxic behavior among adolescents on twitter. In: Social Informatics: 12th International Conference, SocInfo 2020, Pisa, Italy, Oct 6–9, 2020, Proceedings, pp. 427–439. Springer, Berlin, Heidelberg (2020). https://doi.org/10.1007/978-3-030-60975-7_31