

Enriching an Online Suicidal Dataset with Active Machine Learning

Tianrui Liu
Beijing Institute of Technology
Beijing, China
1120192018@bit.edu.cn

Yuchen Zhou
Tianjin University
Tianjing, China
zhouyuchentju2021@163.com

Zhiwei Zheng
Huazhong University of Science and Technology
Wuhan, China
zzwalala@hust.edu.cn

Yiyi Yang, Yang Song
University of North Carolina Wilmington
Wilmington, North Carolina, USA
{yangy, songy}@uncw.edu

ABSTRACT

The scarce and often small size of annotated suicide-related datasets is one of the main obstacles toward automating the process of identifying the online users of high suicide risk. In this paper, we present a framework to annotate a mental-health-related textual dataset with suicide attempts and suicide ideations in the posts and comments using active machine learning method. This approach starts from a relatively small annotated dataset, and learns from a domain expert by obtaining the expert's judgments on the most contradictory samples when the active machine learning model is not able to make the judgments. Meanwhile, the model annotates new samples without asking for the expert's input when it is confident enough about the new samples. The active machine learning models were evaluated and updated when a batch of new samples was annotated, including parameter tuning, replacing models, or changing the representations of samples. The dataset we used is from the SuicideWatch Reddit channel. We expanded the dataset from initially 200 manually annotated samples to 1000 ones.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning.**

KEYWORDS

Active Machine Learning, Suicide Prevention, Suicide Ideation, Suicide Attempts

ACM Reference Format:

Tianrui Liu, Zhiwei Zheng, Yuchen Zhou, Yiyi Yang, and Yang Song. 2022. Enriching an Online Suicidal Dataset with Active Machine Learning. In *2022 ACM Southeast Conference (ACMSE 2022)*, April 18–20, 2022, Virtual Event, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3476883.3520213>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ACMSE 2022, April 18–20, 2022, Virtual Event, USA

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-8697-5/22/04...\$15.00
<https://doi.org/10.1145/3476883.3520213>

1 INTRODUCTION

Suicides take a high toll. There are indications that for each adult who dies of suicide there may be more than 20 others attempting suicide [22]. The impact on families, friends, and communities is devastating and far-reaching. Social, psychological, cultural, and other factors can interact, leading a person to suicidal behavior, and the stigma attached to suicide makes many people unwilling to seek help. Most suicides occur in low- and middle-income countries where resources and services, if they do exist, are often scarce and limited for early identification, treatment, and support of people in need [22]. These striking facts and the lack of implemented timely interventions make suicide a serious global public health problem that needs to be tackled urgently.

Suicide is shrouded in stigma, shame and misunderstanding. This means that people often do not or cannot seek adequate help [10, 18, 21]. Prevention of suicide cannot be accomplished by one person, organization or institution alone; it requires support from the whole community. Communities can play a critical role in suicide prevention. They can provide social support to vulnerable individuals and engage in follow-up care, fight stigma and support those who have suicide ideation. A community can give individuals a sense of belonging and a feeling of connection by being part of it. Lastly, communities can also implement specific suicide prevention strategies relevant to their situation [21], namely, individual needs and struggles [5, 13].

However, as mobile Internet technologies and online social networks prevail, people tend to talk more about their suicide intentions in online communities, where they can seek help from a large number of specific forums, newsgroups, or virtual self-help groups. Some people, especially adolescents, choose to post their suicidal thoughts online, asking how to commit suicide, or enter into online suicide pacts [25]. This online user-generated content could be helpful for detecting individuals' suicidal ideations and suicidal attempts, thus providing another possible angle for early suicide detection and prevention [18]. Therefore, online communities, especially suicide peer-support communities, can play an important role in suicide prevention.

Online communities have an ability to transcend geographical and temporal constraints, which is much better for virtual support groups. There is less risk in disclosure of health information and there is greater access to diverse sources of health information than face-to-face support groups. Due to reduced social status cues,

more heterogeneous supportive relationships are possible [19]. In addition, writing about health problems and formulating disclosure of personal concerns has a therapeutic value by itself [30] (e.g. if one described his/her struggles in words, this person may have a thorough review and a deeper understanding of his/her existing concerns). Besides, it not only makes it easier for early detection but also allows members to support one another. Researchers have conducted an experiment to simulate environments where youth help-seekers wrote about their purposes for seeking help while trained crisis-intervention moderators provided social support and referrals to offline services. Over time, help-seekers could begin to provide support to other help-seekers, developing a community characterized by reciprocal help-seeking and help-providing [9]. Considering it is possible to “cultivate” a peer-supporting atmosphere in an experimental setting, researchers also expect in the real-world online communities, this type of mutual-supportiveness also exists, and they are producing a rich dataset for researchers to better understand the communication of suicide ideations.

The big dataset that the online community creates new challenges for suicide detection and prevention because 1) the volume of online content generated each day is large, and 2) processing the data for suicide prevention should be done in real-time. To train a predictive model for identifying the users at-risk, an annotated dataset is required. So far, there are only a handful of validated datasets that are related to suicide [8, 11] and the sample sizes are usually small (dozens to a few hundreds), which is not enough to train a reliable machine learning model to automate the process of identifying the online users who are likely to attempt suicide.

In this paper, we present a framework to annotate suicide ideation and suicide attempts that are discussed in posts and comments from an online suicide peer-support community. An active machine learning model is used to learn from the domain expert so that the annotation process can be much faster than being done manually. The active machine learning model can learn from the expert, and the model can be updated periodically to better capture the hidden features from the expert’s annotation. We started from 201 samples that are annotated manually, and the experiment demonstrates how we expand the dataset to more than 1001 in a few iterations.

The rest of the paper is organized as follows: Section 2 introduces the related work on active machine learning, suicide ideation detection using NLP (natural language processing), and machine learning; Section 3 introduces our framework of annotation using active machine learning; Section 4 presents our experiment process of iterative annotation to expand a small annotated dataset to more than 1000; Section 5 presents our conclusion of this project and the future work.

We hope that our annotation models will be used in many communities and contexts to help research communities enrich their datasets related to suicide communication. Together, we are working towards the ultimate goal of reducing suicide.

2 RELATED WORK

Active learning is one kind of machine learning, which was first introduced nearly 30 years ago [15, 31]. Nowadays, with the development of the Internet and its virtual connection with humans, an

increasingly amount of information could be gathered and used, and thereby, many studies have been carried out. Nonetheless, in the field of supervised machine learning, one big problem that follows is how to label a large amount of data. Active learning is one of the most effective keys to this problem. The main idea is to select a small subset of unlabeled samples, instead of the whole, then annotate them by human experts. In other words, when we try to annotate datasets, since annotating the whole dataset manually is so costly which often is far beyond the budget, we should start by annotating a small part.

An active machine learning model has a set of independent classifiers. This type of independent model fusion is known as a committee (ensemble) of learners. In annotation work, active learning would send the most controversial samples to an oracle, typically a person with extensive knowledge of the domain. By learning from the domain experts’ judgments on the most controversial samples, the active machine learning model increases its effectiveness, and thereby, reduces the number of data that needs to be labeled manually. Active learning has already been used in many fields, such as clinical concept extraction [12], medical image analysis [3], drug discovery [23], image/video annotation [17]. However, in the light of our research, suicide ideation/attempt detection, and the related research, though a great number of relative work have been done [20, 24, 26, 27], the implementation of active learning in annotation is less explored. Speaking of how previous work annotate the data, either the number of collected data is so limited that the data could be manually annotated [20], or the data is labelled by trained annotators or experts from enterprises[11], which would potentially cost a lot if the sample size is large. In addition to the advantage mentioned above, active learning can annotate a dataset with high quality since the more ambiguous samples are labeled by an oracle.

In the general process of active machine learning, the algorithm will first take a small set of already labeled samples to train a committee of classifiers. After that, a large number of unlabeled samples will be sent to the classifier. The classifiers will then try to annotate these samples independently. Once the classifier is not sure about one sample, it will be pushed to an oracle for a judgment. The label from the oracle will be regarded as the ground truth, and both the sample and the label will be appended to the training dataset to retrain the classifiers. The overall idea is to minimize the human effort during the process of annotation by querying the human experts as few samples as possible.

There are two general ways to quantify how uncertain a model is on a given sample. One is called Uncertainty Sampling, and it could be further divided into two standards. One standard is to select samples with the lowest confidence generated along with the annotation result, like the entropy of the class prediction [14]. The second standard is that the model will take the lowest margin into consideration [6], sometimes as well as the confidence [28].

The other one is called Query by the Committee. In this context, many classifiers will first be grouped into one committee. The committee will be trained before each classifier their judgment independently based on their standards, and the committee will vote/give the label according to the majority. However, if the voting result is close, which means the committee is not sure, then the sample will be sent to an oracle [7, 16].

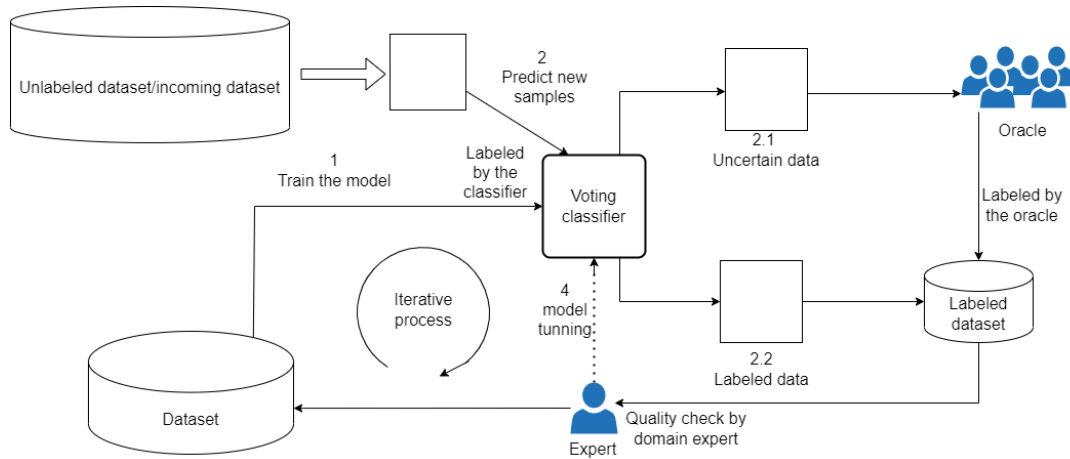


Figure 1: Workflow of the Active Machine Learning Model

In a machine learning project, to train the models, researchers first need enough annotated data. In the last several years, many datasets have been derived from social media data, mainly from Twitter. In one project focusing on the analysis of six common emotions, 305,310 tweets were first downloaded and then 8,150 tweets were retrieved among them, based on a list of seed words for the six emotions [1]. All these emotions were then labeled by five people who “received no training”. Instead of manually labeling all messages, in another project of finding out how the media portray of public transit services can influence the way voters think about the future government investments in transit, the author collected 65,000 tweets from celebrities, public parks, etc. [24]. Then the researchers used a text-mining algorithm to score some specific words in these tweets separately based on the lexicon derived from Hu and Liu [17], after which all scores were summed for each tweet. Apart from some English social media, researchers also used data from Weibo, one of the Chinese most populated social media. A research team downloaded posts from Weibo to assess the suicide risk and emotional distress of Weibo users [4]. All posts were then parsed and fitted into Simplified Chinese-Linguistic Inquiry and Word Count (SC-LIWC) categories to build a dataset. The features collected from the SC-LIWC were then associated with the five 5 suicide risk factors; then the authors used machine learning to classify whether a Weibo user presents those risk factors.

3 OUR APPROACH

3.1 Dataset

All the dataset used in our work was collected from the Suicide-watch subreddit. Reddit is an American social news aggregation, web content rating, and discussion website [29]. Table 2 shows a small sample of our dataset. Registered members submit content to the site, which is then voted up or down by other members, or post their own thoughts and comments below. The subreddit SuicideWatch was created in 2008 to provide peer support for people struggling with suicidal thoughts. It’s public, so we collected posts and comments using the Pushshift API [2], which was created by the Subreddit Datasets Mod Team. The data was broken into months,

and the features collected were: Id, author, date, score, number of comments, title, and self text. The submissions and comments were collected from 2019 to July 2021.

In the initial stage of our research, we have trained two students to label the data with suicide attempts and ideations using a randomly selected sample ($n = 201$). We labeled whether the data included suicide attempts and suicidal ideations, as well as statements supporting the label.

3.2 Methodology

In this project, we used the Query by the Committee approach in building our active machine learning model. At first, we chose K Nearest Neighbors (KNN), Decision Tree (DT), Random Forest (RF), Logistic Regression (LR), Stochastic Gradient Descent (SGD) Classifier, Naive Bayes (NB), and Support Vector Machine (SVM with the linear kernel) to build the voting classifier. We used the initial small annotated dataset to train the voting classifier, but the accuracy is very low, around 50%. We then deployed grid search to do the parameters tuning and selected the best five classifiers, which are KNN ($K=4$), DT, LR, SVM (with the linear kernel), and SGD. Overall, the accuracy of individual classifiers ranges between 68% to 73% on the initial dataset (201 samples).

In our dataset, we label a post presenting suicidal ideation/attempt as 1, otherwise, the sample is labeled as 0. The number of 1 and 0 is unbalanced because the dataset is small and the people who have suicidal attempts are the minority among the users on the SuicideWatch subreddit. As a result, it is difficult to improve the model using the limited initial dataset, instead, we started with this group of “weak classifiers” as the committee of our active machine learning model to annotate more samples.

To help train the active machine learning model, we have a human expert who is a researcher in suicide communication. This researcher interacted with our model as the oracle/human expert. With more training and domain knowledge, the human expert can label the data reliably. Therefore, when the active machine learning model is unsure about an unlabeled sample, it can “query” by *pushing* it to the expert, who “answers” the query by labeling

Table 1: Sample Size and Model Performances on Identifying Suicide Ideation and Attempts

Round	Sample size (positive, negative)		Improve model performance-ideation			Improve model performance-attempt		
	Suicide ideation	Suicide attempt	Accuracy	Sensitivity	Sepecificity	Accuracy	Sensitivity	Sepecificity
Initial annotation	201(125,76)	201(71,130)	n/a	n/a	n/a	n/a	n/a	n/a
Active learning round 1	100(71,29)	100(5,95)	0.75	0.85	0.59	0.77	0.36	0.91
Active learning round 2	200(144,56)	200(31,169)	0.79	0.86	0.65	0.84	0.51	0.93
Active learning round 3	150(135,20)	150(52,98)	0.82	0.89	0.69	0.86	0.53	0.94
Active learning round 4	240(174,66)	240(42,198)	0.87	0.92	0.76	0.87	0.71	0.91
Active learning round 5	110(81,29)	110(14,96)	0.86	0.93	0.77	0.87	0.68	0.91

the data. The model updates and the whole process will be executed iteratively. In our workflow design, after each batch of new data is annotated by the human expert, the expert also validates the annotation quality in the current round, and makes corrections if necessary. After each round, we update the committee of learners and tune the parameters again.

In our active machine learning algorithm, the committee records the results of each classifier’s annotation to the data and adds the numbers annotated by the five classifiers/committee members. If more than 75% of the classifiers labeled the data as 1 (in our current algorithm, it means the sum is greater than or equal to 4), the data should be labeled as 1; if more than 75% of the classifiers labeled the data as 0 (in our current algorithm, it means the sum is less than or equal to 1), the data should be labeled as 0; otherwise (the sum equals 2 or 3), the committee is unsure about the data, and it pushes the data to the expert asking for a label. Figure 1 shows the general workflow of our active machine learning model.

4 EXPERIMENT RESULTS

We have created an active machine learning model that contains a set of independent classifiers to analyze suicide ideation and attempts. The suicide ideation and suicide attempts are treated as two independent classification tasks. Therefore, the committee members of learners are not the same, and the parameters are tuned separately as well. The model has been evaluated and its performance is promising. Table 1 presents the sample size generated from each round, as well as the test results on the model’s sensitivity, specificity, and overall accuracy.

From Table 1, we found that the overall accuracy of predicting both suicide ideation and suicide attempt has been improving as we collect more data and iteratively improve/update the committee of learners. However, we also notice that the data we have is imbalanced on the positive samples and negative samples: there are generally more samples that talked about suicide ideation, but among them, there are fewer that really talked about suicide attempts. We find that the active learning model has a higher sensitivity on identifying suicide ideation, but a higher specificity on identifying suicide attempts, and this is likely to be caused by the skewed distribution of the data.

The model is being developed and will continue to be improved after each round of new samples is annotated and validated by the human expert. When new samples are annotated, we record the predictions from each learner on each sample. When a new batch of samples is annotated, we evaluate the performance of

each individual learner. We decide whether to discard a learner given its poor performance, or to tune the parameters to improve its performance. Table 3 records the improvements we have done on the active machine learning models after each round of active learning. We have done a total of 5 rounds. To report the human expert’s involvement in the annotation process, we use round 5 as an example: for suicide attempt and suicide ideation respectively there were 12.5% and 7.08% of samples pushed to the oracle; the oracle also manually fixed 7.08% and 4.17% of samples during the quality check phase, too.

5 CONCLUSION

In this paper, we present our project that uses active machine learning (the Query by the Committee approach) to annotate an online dataset with suicide attempts and suicide ideations in the posts and comments. This approach starts from a relatively small annotated dataset (201), and learns from a domain expert by pushing the most contradictory samples to the expert and obtaining the expert’s judgments. For the samples that the active machine learning model is confident enough, the model will annotate them without asking for the expert’s input. The active machine learning models were evaluated and updated by us when a batch of new samples (usually 100-200) were annotated. This process may include tuning the models, replacing models, or changing the representations of samples.

The dataset we used is from the SuicideWatch Reddit channel. In our experiments, we did 5 rounds of active machine learning with a domain expert’s input. In each round, the active machine learning is updated to improve the performance. We were able to use this approach to expand our dataset to 1001 samples annotated with both suicide ideations and suicide attempts presented.

The main goal of this project is to present a framework to annotate a mental-health-related textual dataset that does not require domain experts to read and annotate all the samples manually. In addition, we will open this dataset to the research community to build better models that can help identify the users who are likely to have suicidal ideations/behaviors. Our effort in expanding this dataset will continue and thereby make it more useful for the machine learning task.

ACKNOWLEDGMENTS

Firstly, we believe bias exists in the dataset against minority communities, regarding ethnicities, sexual orientations, etc. More samples should be collected to ensure there are enough samples for

Table 2: Raw Data

Index	id	author	created_utc	score	num_comments	url	title	selftext
18	bjcbgh	iei91ojw9qiw29o22i	4/30/2019 10:53:53 PM	0	7	https://www.reddit.com/r/SuicideWatch/comments/bjcbgh/just_kill_your_self/	Just kill your self	No one loves you
35	bjbqh5	Dipshit_Kris	4/30/2019 9:49:54 PM	1	0	https://www.reddit.com/r/SuicideWatch/comments/bjbqh5/what_do_we_say_to_the_god_of_death/	What do we say to the god of death?	Not today.

Table 3: Model Improvements After Each Batch of Data From Active Learning

Round	Improvement
1	<ul style="list-style-type: none"> Removed KNN for suicide attempts, replace with SVM (with the polynomial kernel) Tuned parameters for suicide ideation: SVM (with the linear kernel), KNN, DT, SGD Tuned parameters for suicide attempt: SVM (with the linear kernel), RF, DT, SGD
2	<ul style="list-style-type: none"> Tuned parameters for suicide attempt: DT Tuned parameters for suicide ideation: KNN, LR, SGD Introduced downsampling for the data-balancing purpose on suicide attempt prediction.
3	<ul style="list-style-type: none"> Removed KNN for suicide ideation prediction, replaced with SVM with the polynomial kernel Tuned parameters for suicide attempt prediction: SVM(with the linear kernel), DT, SGD
4	<ul style="list-style-type: none"> Added Bibtex to represent features Tuned parameters for suicide ideation prediction: SVM(with the linear kernel), SGD
5	<ul style="list-style-type: none"> Not changed

each minority population’s struggle to be seen and captured by the active machine learning model. Secondly, further intervention is harder than expected since the data is pulled directly from on-line communities. However, the authors are collaborating with the social-media industry to integrate our model into their server-side so that timely intervention can be deployed once our models find users of high suicide risk. This project is supported by the UNCW research momentum grant in the summer and fall of 2021.

REFERENCES

- [1] Rakesh C. Balabantaray, Mudasir Mohammad, and Nibha Sharma. Multi-class Twitter Emotion Classification: A New Approach. *International Journal of Applied Information Systems*, 4(1):48–53, 2012.
- [2] Jason Michael Baumgartner. Pushshift API. <https://github.com/pushshift/api>.
- [3] Samuel Budd, Emma C. Robinson, and Bernhard Kainz. A Survey on Active Learning and Human-in-the-loop Deep Learning for Medical Image Analysis. *Medical Image Analysis*, page 102062, 2021.
- [4] Qijin Cheng, Tim MH Li, Chi-Leung Kwok, Tingshao Zhu, and Paul SF Yip. Assessing Suicide Risk And Emotional Distress in Chinese Social Media: A Text Mining and Machine Learning Study. *Journal of Medical Internet Research*, 19(7):e243, 2017.
- [5] Evelien Coppens, Chantal Van Audenhove, Samuel Iddi, Ella Arensman, Katrin Gottlebe, Nicole Koburger, Claire Coffey, Ricardo Gusmão, Sónia Quintão, Susana Costa, et al. Effectiveness of Community Facilitator Training in Improving Knowledge, Attitudes, and Confidence in Relation to Depression and Suicidal Behavior: Results of the OSPI-Europe Intervention in Four European Countries. *Journal of Affective Disorders*, 165:142–150, 2014.
- [6] Mark Dredze and Koby Crammer. Active Learning with Confidence. In *Proceedings of ACL-08: HLT, Short Papers*, pages 233–236, Columbus, USA, 2008.
- [7] Yoav Freund, H. Sebastian Seung, Eli Shamir, and Naftali Tishby. Selective Sampling Using the Query By Committee Algorithm. *Machine Learning*, 28(2):133–168, 1997.
- [8] Soumitra Ghosh, Asif Ekbal, and Pushpak Bhattacharyya. Cease, A Corpus of Emotion Annotated Suicide Notes in English. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1618–1626, 2020.
- [9] Elaine Greidanus and Robin D. Everall. Helper Therapy in an Online Suicide Prevention Community. *British Journal of Guidance & Counselling*, 38(2):191–204, 2010.
- [10] Mark E. Hastings, Lisa M. Northman, and June P. Tangney. *Shame, Guilt, and Suicide*, pages 67–79. Springer US, Boston, MA, 2002.
- [11] Yang Hui, Alistair Willis, Anne De Roeck, and Bashar Nuseibeh. A Hybrid Model for Automatic Emotion Recognition in Suicide Notes. *Biomedical Informatics Insights*, 5:BI1–S8948, 2012.
- [12] Mahnoosh Kholghi, Laurianne Sitbon, Guido Zuccon, and Anthony Nguyen. Active Learning Reduces Annotation Time for Clinical Concept Extraction. *International Journal of Medical Informatics*, 106:25–31, 2017.
- [13] Michael J. Kral, Patricia K. Wiebe, Kari Nisbet, Catherine Dallas, Looee Okalik, Nubiya Enuaraq, and James Cinotta. Canadian Inuit Community Engagement in Suicide Prevention. *International Journal of Circumpolar Health*, 68(3):292–308, 2009.
- [14] David Lewis and William Gale. A Sequential Algorithm for Training Text Classifiers. *Proceedings of the 17th Annual International Acm Sigir Conference on Research and Development in Information Retrieval*, 29, 08 2001.
- [15] David JC MacKay. Information-based Objective Functions for Active Data Selection. *Neural computation*, 4(4):590–604, 1992.
- [16] Andrew McCallum and Kamal Nigam. Employing EM and Pool-Based Active Learning for Text Classification. In *ICML*, San Francisco, United States, 1998.
- [17] Wang Meng and Hua Xian-Sheng. Active Learning in Multimedia Annotation and Retrieval: A Survey. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(2):1–21, 2011.
- [18] Christine Yu Moutier, Michael F. Myers, Jennifer Breen Feist, J. Corey Feist, and Sidney Zisook. Preventing Clinician Suicide: A Call to Action During the COVID-19 Pandemic and Beyond. *Academic Medicine*, 96(5):624–628, 2021.
- [19] Monica Murero and Ronald E Rice. *The Internet and Health Care: Theory, Research, and Practice*. Routledge Communication Series. Taylor & Francis, New York, 2013.
- [20] Alicia L Nobles, Jeffrey J Glenn, Kamran Kowsari, Bethany A Teachman, and Laura E Barnes. Identification of Imminent Suicide Risk Among Young Adults Using Text Messages. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–11, Montréal, Canada, 2018.
- [21] World Health Organization et al. Integrated Care for Older People: Guidelines on Community-level Interventions to Manage Declines in Intrinsic Capacity. 2017.
- [22] World Health Organization et al. Preventing Suicide: A Community Engagement Toolkit. 2018.
- [23] Daniel Reker and Gisbert Schneider. Active-learning Strategies in Computer-assisted Drug Discovery. *Drug Discovery Today*, 20(4):458–465, 2015.
- [24] Lisa Schweitzer. Planning and Social Media: A Case Study of Public Transit and Stigma on Twitter. *Journal of the American Planning Association*, 80(3):218–238, 2014.
- [25] Ji Shaoxiong, Yu Celina Ping, Fung Sai-fu, Pan Shirui, and Long Guodong. Supervised Learning for Suicidal Ideation Detection in Online User Content. *Complexity*, 2018, 2018.
- [26] Nadiya Straton, Hyeju Jang, and Raymond Ng. Stigma Annotation Scheme and Stigmatized Language Detection in Health-Care Discussions on Social Media. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1178–1190, Marseille, France, 2020.
- [27] Farah Tokmic, Mirsad Hadzikadic, James R. Cook, and Oleg V. Tcheremissine. Development of a Behavioral Health Stigma Measure and Application of Machine Learning for Classification. *Innovations in Clinical Neuroscience*, 15(5-6):34–42, June 2018.
- [28] Simon Tong and Daphne Koller. Support Vector Machine Active Learning with Applications to Text Classification. *Journal of Machine Learning Research*, 2(Nov):45–66, 2001.
- [29] Wikipedia. Reddit. <https://en.wikipedia.org/wiki/Reddit>.
- [30] Kevin Wright and Sally Bell. Health-related Support Groups on the Internet: Linking Empirical Findings to Social Support and Computer-mediated Communication Theory. *Journal of Health Psychology*, 8(1):39–54, 2003.
- [31] Byoung-Tak Zhang and Gerd Veenker. Neural Networks that Teach Themselves Through Genetic Discovery of Novel Examples. In *Proceedings 1991 IEEE International Joint Conference on Neural Networks*, pages 690–695 Vol.1, Seattle, USA, 1991.