



HHS Public Access

Author manuscript

Proc Int World Wide Web Conf. Author manuscript; available in PMC 2017 July 19.

Published in final edited form as:

Proc Int World Wide Web Conf. 2017 April ; 2017: 123–131. doi:10.1145/3038912.3052661.

Cataloguing Treatments Discussed and Used in Online Autism Communities

Shaodian Zhang[†], Tian Kang[†], Lin Qiu[‡], Weinan Zhang[‡], Yong Yu[‡], and Noémie Elhadad[†]

[†]Department of Biomedical Informatics, Columbia University, New York, NY, US

[‡]Apex Data and Knowledge Management Lab, Shanghai Jiao Tong University, Shanghai, China

Abstract

A large number of patients discuss treatments in online health communities (OHCs). One research question of interest to health researchers is whether treatments being discussed in OHCs are eventually used by community members in their real lives. In this paper, we rely on machine learning methods to automatically identify attributions of mentions of treatments from an online autism community. The context of our work is online autism communities, where parents exchange support for the care of their children with autism spectrum disorder. Our methods are able to distinguish discussions of treatments that are associated with patients, caregivers, and others, as well as identify whether a treatment is actually taken. We investigate treatments that are not just discussed but also used by patients according to two types of content analysis, cross-sectional and longitudinal. The treatments identified through our content analysis help create a catalogue of real-world treatments. This study results lay the foundation for future research to compare real-world drug usage with established clinical guidelines.

Keywords

Natural Language Processing; Autism; Online Health Community; Treatment; Conditional Random Fields

1. INTRODUCTION

Online health communities (OHCs) provide means of exchanging information and social support for a vast amount of patients, especially individuals with chronic or life-threatening diseases. A wide range of studies have been carried out for purposes such as creating social support interventions [28, 19, 31], understanding patient behaviors [18, 36], assisting community facilitators [24], and mining critical disease- or medication-specific information [30, 35, 38]. Information related to disease treatment, such as medications, therapeutic protocols, and surgeries, have been particularly prevalent in online health discussion. One specific research question has remained unexplored thus far is which treatments are actually taken or adopted by online health community members in their real lives, as opposed to

published under Creative Commons CC BY 4.0 License.

[†]sz2338@columbia.edu, tk2624@columbia.edu, noemie.elhadad@columbia.edu

[‡]lqiu@apex.sjtu.edu.cn, wnzhang@apex.sjtu.edu.cn, yyu@apex.sjtu.edu.cn

merely discussed within a community. This question is of interest to patients, clinical practitioners, and health researchers alike: patients want to know about the drugs and protocols their peers rely on, as means to inform their own decision making. Clinicians and health researchers can learn about adherence to various treatments. For health researchers, public online health communities provide massive cohorts of patients who are potential subjects of post-market research. For example, researchers may be interested in what treatments are actually consumed by patients, in contrast to what are suggested by established clinical guidelines. Traditionally, it is usually difficult to carry out such research analysis outside of clinical setting, which provides unprecedented opportunities to rely on content analysis of OHCs and social media. For example, to compare clinicians' and patients' perspectives on the symptomatic treatment of ALS, Nakamura et al. conducted a study by comparing data from a traditional survey study of clinicians with data from an OHC, PatientsLikeMe [27].

In most health communities, patients discuss and report on their treatments through posts, with free-text narratives. Critical to our task at hand is thus a method that can automatically identify signals of actual treatment usage from OHC content, especially OHC text. The task is not as trivial as simply extracting mentions of treatment from user posts. One particular challenge of connecting online content with reality should be tackled first: messages conveyed through content may not necessarily indicate the happening of corresponding events in real lives. In the case of treatment-related discussions, discussing a treatment online does not always indicate the action of adopting the treatment. Particularly, many of the mentions of treatments may not be attributed to the patients themselves. For example, users in OHCs may discuss related scientific findings about a treatment, in which a large number of treatment names may occur. Such mentions do not indicate any actual usage of the drugs, and therefore should be excluded when cataloguing patients' treatment usage.

The context for our study is Autism Spectrum Disorder (ASD), a range of developmental disorders originating in early childhood, and two ASD online communities. Unlike other disease-specific communities where participants are primarily patients or survivors, autism communities' members are mostly parents of children on the spectrum. In ASD forums, members primarily discuss their children's diagnoses and treatments and a wide range of topics related to their children's lives. They also rely on the community to exchange information and support about themselves as caregivers. As such, in order to build up catalogues of treatment for patients of interest (i.e., children with ASD), it is important to distinguish between treatments attributed to patients themselves and to the caregivers.

In this paper, we aim to build a catalogue of treatments for children with ASD through text mining of OHCs. Our approach consists of two steps: extracting mentions of treatment in OHC posts and classifying treatment attributions, i.e., attributing the treatment mention to a patient, a caregiver, someone else, or as a generic mention. We show that treatment mentions and their attributions could be identified jointly through an end-to-end classifier based on conditional random fields (CRF). Results of this study can help investigate further real-world use of ASD treatments as derived from caregivers' discussions: the catalogue of treatments, and even their established prevalence of use in a community, can be compared with established clinical guidelines, and help study the gap between clinical expectation and

patients' daily lives. In this paper, we will also take temporal information into consideration, investigating how members' perceptions and usages of treatments change throughout their community participation, thus investigating whether there is any influence from the community in the decision making process of community members when it comes to treatments.

2. BACKGROUND

2.1 Autism Spectrum Disorder (ASD)

Autism Spectrum Disorder (ASD or autism) is a neurodevelopment disorder characterized by deficits in social communication, social interaction, as well as the presence of restricted, repetitive behaviors [1]. Autism statistics from the U.S. Centers for Disease Control and Prevention (CDC) identify around 1 in 68 American children as on the autism spectrum, a ten-fold increase in prevalence in 40 years. ASD affects over 3 million individuals in the U.S. and tens of millions worldwide [6].

ASD is typically diagnosed on the basis of behavioral symptoms, without reference to etiology. However, considerable research has been devoted to investigations of etiological factors. To date, there is no accepted single cause of ASD, although there are numerous theories and studies suggest that autism results from different sets of causal factors, such as genetic, neurobiological, and environmental, that manifest in characteristic behavioral symptoms [5, 15]. The most obvious signs of autism and symptoms of autism tend to emerge between 2 and 3 years of age. Currently, for diagnosis of ASD, no established medical test was recognized, as the symptoms of autism vary [7]. Instead, specially trained physicians and psychologists administer autism-specific behavioral evaluations.

The treatment or intervention of ASD can involve behavioral treatments, medicines or both. Given the nature of autism and the needs of individuals with ASD, family involvement serves as a key role during the treatment [26]. Most comprehensive programs for individuals with ASD offer parents/caregiver training.

In fact, beyond the impact of ASD on the patients themselves, ASD impacts families as a whole. Caregivers as such have a with a wide range of emotional, instrumental, and informational needs. Online autism communities are a place for exchange of support for these three types of needs. Maybe because the etiology and mechanisms of ASD are still enigmatic, ASD communities contain much discussions and opinions amongst caregivers about causes and efficacy of various treatments, including misinformation (e.g., vaccines as a cause to autism) [11]. There is a critical need to understand to which extent caregivers' understanding of ASD and treatment decisions align with current clinical characterization of the disease.

2.2 Named Entity Recognition and Attribution Identification

In this paper, our primary task is to recognize treatments mentioned in OHCs and then distinguish their attributions. The first step, recognizing treatment names, is traditionally in the realm of named entity recognition (NER). One of the most important use case of clinical NER is mining patient related healthcare knowledge from Electronic Health Records (EHR),

which contain abundant information in free-text format. NER is one of the fundamental issues in information extraction of EHR [23]. Both supervised and unsupervised approaches are deployed in clinical NER systems. One of the most mature systems for clinical text processing is MedLEE [14], which uses heuristic rules and patterns for identifying clinical problems, treatment and events. Zhang et al. developed an unsupervised clinical NER system by generating seed terms from Unified Medical Language System (UMLS) [39]. Learning-based methods have been rapidly advanced in the past 5 years [16]. Representative ones include conditional random fields (CRF) and structured support vector machines (SVM) [34]. For example, the best system reported in the i2b2 NER challenge from deBruijn et al. used Semi-Markov (F score 85.2%) [10], followed by the system by Jiang et al. using CRF (F-score 83.9%) [20]. Mining important clinical entities and events in social media data for health related purpose is also a promising field, especially from the perspective of public health. Social networks have seen an unprecedented growth worldwide. A large population of patients are actively involved in sharing and posting health related information in social networks [32].

Pertaining to the second task of this paper, attribution identification, a recent survey [13] has revealed that 34% of caregivers and 20% of patients read or watch other's commentary or experience online. In addition, 11% of caregivers and 6% of patients share experience or post questions online. Particularly, in health related social platforms, users share and discuss their health-related experience with others facing similar problem, including use of drugs, side effects and treatments, which makes such social networks unique and robust source of information such as drugs and treatments [32]. Numerous studies have been published recently in this field, including studies on pharmacovigilance [17], identifying smoking cessation patterns [22], identifying user social circles with common experiences such as drug abuse [4], and tracking infectious disease spread [12, 2, 29].

3. METHODS

3.1 Dataset

The dataset of autism forums used in this paper were collected from two publicly available sources: autismweb.com and autism-pdd.net, which are primarily for ASD patients and caregivers. The forum from autism-pdd.net was officially closed in 2015 and could no longer be accessed. We crawled all content that was public available from these two forums in March 2015. These two forums are designed for the same audience and thus have similar functionalities, but the forum from autismweb.com is significantly larger than the one from autism-pdd.net. As such, we merged these two forums into one single dataset, with following information available: sub-forums, threads, posts, and authors. Detailed descriptive statistics of this dataset can be found in Table 1.

3.2 Annotation

Five types of attributions were considered in the manual annotation, with descriptions given in Table 2. In general, the labels were designed to reflect whom the treatment is tied to. We consider four types of attributions: mentions attributed to the patients (ASD children), the caregivers (usually parents of children, who actually participated in the discussions), others

(either other people in the community or other people in the real world), and no one (i.e., general discussions). Moreover, as we discussed previously, an entity of a treatment attributed to a patient does not necessarily indicate an actual history of usage. For instance, in “The doctor suggested to put my son on risperdal”, although the mention “risperdal” is associated with the patient (my son), it is not clear whether the drug is eventually prescribed by the doctor and hence taken by the patient. Therefore, in order to support subsequent user modeling in which we establish a treatment catalogue for each user, in the annotation schema we further distinguish mentions of treatments attributed to patients which do and do not indicate actual usage or usage history, i.e. *pt* and *pt-gen* in Table 2.

A randomly sampled 500 posts were extracted and split into two sub-sets, with 50 posts overlapping (i.e. first set from post 1 to post 275, second set from post 225 to post 500). Two annotators were asked to 1) identify mentions of treatments (entities) from text, and 2) annotate the attribution label according to Table 2 for each mention. It is noteworthy, however, that the annotators were asked to classify attributions locally, without considering context which may shift the attribution of a mention. For example, in “The doctor suggested to put my son on risperdal....My son tried risperdal and...”, the first mention of risperdal should be labeled as *pt-gen*, even if following context indicates an actual usage of the same drug. In the annotation for this task, we did not consider co-references, e.g., pronouns which refer to treatments. Since our purpose is to find treatment history of patients, ignoring co-references should not affect the final results significantly as long as all treatment names are properly identified.

The annotation started with each annotator coding the overlapping part of the two sets, on which we tracked inter-rater agreement. Our annotators reached a Kappa of 0.77 after 3 rounds of pre-annotation and disagreement resolving. The remaining parts of the two sets were coded by the two annotators independently. In total, 4,264 mentions of treatments were identified from the 500 posts. Among them, 434 were annotated as *pt-gen*, 1,830 as *pt*, 210 as *others*, 95 as *cg*, and 1,635 as *gen*.

3.3 Models and Experiments

We base on the popular sequence labeler, the conditional random fields (CRF), to jointly identify treatment mentions and their attributions. Three separate sets of evaluations were carried out for this task. Since the CRF model handles term identification and term classification jointly, it is necessary to evaluate these two separate steps in an explicit way. As such, the first set of evaluations is to examine how well the classifier can detect treatment entities, regardless of their attribution labels. The second set of evaluations takes attributions into consideration and evaluates the end-to-end performance of the method on the task. Finally, in our particular scenario of application in which we aim at building up treatment catalogue for each patient, we are more interested in one attribution label, the *pt* class. Therefore, one additional evaluation is also carried out in which only two attribution labels are considered, *pt* and *non-pt*. The *non-pt* class is simply the aggregation of all attribution labels other than *pt*. For each set of evaluations, we report the performance of CRF model with different sets of features, ranging from basic lexical ones to syntactic features and information from context posts. We also compare the system with baselines. The baseline

system for mention identification is based on a dictionary match using Consumer Health Vocabulary [33]. The baseline system for attribution classification is a logistic regression classifier with lasso regularization using the same set of features.

Features we used in this study are described as follows: *Lexical features* refer to words, lemmas, part-of-speech tags of the content. In this study, we rely on the existing open source tools, the OpenNLP [3], to extract these features from raw content. The feature also includes occurrences of non-semantic tokens, such as question marks, exclamation marks, and mentions of user names.

Syntactical features are ones relying on the parse tree of the sentences. In particular, subject (usually an NP) and predicate (usually a VP) of sentences, as well as the lengths of paths from current token to subject and predicate, are used as features. Sentences were parsed by the StanfordNLP toolkit [21] in our experiments.

Semantic features refer to those ones representing domain knowledge, relying on dictionary matching according to following lexicons: WordNet [25], UMLS metathesaurus [8], and Consumer health vocabulary [33].

4. RESULTS

4.1 Treatment Mention Identification

Table 3 lists performance measured by precision, recall, and F score for the treatment name identification, regardless of attribution types. All CRF-based systems, no matter what the features are, outperform the baseline (dictionary matching based on Consumer health vocabulary) significantly. However, syntactical features do not help the system performance. It seems to suggest that regardless of the treatment attributions, lexical features alone (including ones based on lexicons) are sufficient to identify the treatment mentions for CRF model.

4.2 End-to-end Evaluation

Performance of the end-to-end evaluation of joint treatment mention detection and attribution classification is given in Table 4. A true positive in this evaluation is a recognized treatment mention with both boundary and attribution correctly identified. As a result, it is a more challenging task since either an incomplete boundary or an incorrect attribution label will make the prediction counted as an error. The overall micro averaged F score is around 50 to 60, which varies by different feature sets. Syntactic features, which represent more global information from the whole sentence, are decisive in this task. Compared with the standalone evaluation of treatment identification, it seems to suggest that syntactic features are helpful for attribution classification, but not entity recognition.

The baseline systems use dictionary matching for term identification, and logistic regression with corresponding features for attribution classification. All baseline systems underperform their CRF counterparts significantly. However, the performance of baseline systems may just be compromised by the weak dictionary matching baseline for term identification. In order to compare CRF and logistic regression more fairly in attribution classification. We also

carried out a set of additional experiments, where CRF's results on term identification were given to logistic regression as input for attribution classification. The results are given as baseline+ in Table 4. It can be seen that systems based on CRF still outperform logistic regression, which suggests that attribution identification benefit from sequence learning.

Across the five attribution categories, our method is able to classify *gen* and *pt* better than the other three. This is primarily because of the distributions of these attributions in the training and test datasets - *pt* and *gen* are the most dominant attributions. Fortunately, in our downstream applications, building up treatment catalogues for patients, only treatment mentions with *pt* attribution will be used. To see if excluding other attributions from the dataset to make the classification as a binary choice (*pt* vs. *non-pt*) can help boost the accuracy of identifying mentions attributed to patients, we carried out an additional evaluation in which *gen*, *others*, *pt-gen*, and *cg* were merged into one class. The performance is given in Table 5. Compared with Table 4, accuracy of identifying Patient is boosted for around 4–5 percent, although the dataset, feature, and model keep exactly the same ones. The results suggest that properly formulating the task and setting up the target categories make significant difference in this type of tasks.

4.3 Creating Treatment Catalogues for Members

We applied the best system to the entire un-annotated dataset. In total, 164,335 mentions of 3,981 different treatment terms were identified in the entire ASD dataset. In average, around every three posts have one treatment mention. *pt*, which represents that a mention is attributed to patients of interest, is the most dominant attribution label, with 79,778 mentions of 3,552 treatment terms identified. Since some of the terms may refer to the same treatment (e.g. chelation and chelating), actual number of treatment identified may be less. 71,1102 mentions of 3,622 treatment terms, 7,783 mentions of 1,142 treatment terms, 5,297 mentions of 915 treatment terms, and 275 mentions of 176 treatment terms are identified for attribution *gen*, *pt-gen*, *others*, and *cg*, respectively.

The original top ten most frequent treatment terms with corresponding numbers of mentions for each attribution class are given in Table 6. The lists contain common treatment options for ASD patients, as well as alternative therapies. Prevalence of the same treatment in different attribution classes may differ. For instance, although chelation is the most prevalent treatment discussed in the forum, and is particularly popular when attributed to general discussions (See the class *gen* in Table 6), number of mentions of chelation which attributed to the users' actual usage is not that dominant. It is interesting that alternative therapies, such as probiotics and vitamins, are used by patients in the forum almost as frequently as conventional drugs such as Risperdal. Moreover, it is surprising that almost all the top ten terms identified for each attribution class are indeed either treatment options or nutritional supplements, with only one false positive appeared in the list of Caregiver (*cab*). Given the broad coverage of treatments identified, the high precision of the top term lists indicate a successful application of the our method to identify treatment terms. However, some of the terms identified in specific attribution classes are questionable. In particular, treatments attributed to caregivers in current result are mostly treatment options for ASD, which are likely to be caused by incorrect classifications.

After identifying attributions of treatment mentions, for each user we are able to create a treatment catalogue in which all treatments attributed to their ASD children are recorded. We obtain this by simply aggregating all treatment mentions whose attribution are *pt*, in all the posts of individual users in the forum. As such, we are able to create treatment catalogues for 3,635 members. Among them, 2,301 have tried more than one treatment according to the identification. Distributions of number of users, by number of used treatment, is given in Figure 1. Most of the members have tried multiple treatments, which is consistent with the fact that parents tried various treatments as well as supplements for their ASD children, since autism is complex and hardly be curable with standard conventional protocols.

Table 7 shows the top ten treatments most used by members in the autism communities. The difference between this table and the *pt* columns in Table 6 is that multiple mentions of a treatment of the same attribution posted by one user will only be counted once in this table. As such, numbers in Table 7 represents the true prevalence of treatment usage among autism forum users, rather than frequencies of mentions. Chelation, as a controversial therapy which lacks sufficient scientific evidence of effectiveness, attracts a lot of discussions in the forums, according to its overall frequency in Table 6. However, it only ranks 3rd as the mostly used treatment by patients. On the contrary, probiotics as a nutritional supplement, and speech therapy as a well-established psychosocial therapy for ASD children, are more popular in real practice.

4.4 Longitudinal analysis of treatment catalogues of members

We also investigate how frequencies of treatment mentions change through time, and how the patterns differ across attribution types. Specifically, treatment mentions of attribution type *pt* and ones of other types are considered separately. We illustrate how frequencies of mentions change through time in weeks and in days since members joining the community in Figure 2.

In general, no clear pattern can be identified for each treatment. Members do not necessarily focus on certain treatment at the beginning stage of participation. In the long run, members keep discussing treatment options throughout their participation, with no decline in frequencies of mentions of any terms significantly.

In terms of frequencies of mentions attributed to patients, it was expected that such mentions should occur more frequently at the initial stage of participation, when members join the community and introduce conditions and current treatment adoptions of their ASD children. However, such pattern is not found in our analysis. On the contrary, frequencies of mentions attributed to patients fluctuate with total frequencies, and maintain substantial percentages throughout members' participation. One possible explanation is that members try different treatment options for their children at different times, and keep updating about their effectiveness in the forums.

5. DISCUSSION AND CONCLUSION

The results suggest that abundant treatment options, ranging from conventional therapies to alternative ones, are discussed in the autism forums. Mentions of treatments are attributed to different stakeholders of autism care such as patients and caregivers. In the autism forums, most of the treatment discussions are attributed to ASD children of community members. Although not all mentions of treatment indicate actual adoption, around 90% of treatment mentions attributed to patients (ASD children) represent an ongoing treatment or a history of usage. Specifically, members keep updating status of their kids as they are treated, in which massive amount of treatment mentions occur. Members of the autism forums also discuss therapies frequently on issues like scientific evidence of effectiveness and information received from health professionals and online sources. A small proportion of treatment mentions are attributed to the caregivers themselves as well as other people in the community or in their real lives.

Our study also demonstrates that it is possible to rely on automated content analysis to overcome the issue by identifying attributions of extracted information and further filtering out information that is not associated with real actions. Similar issues also exist in studying other characteristics of OHC members (e.g. connecting sentiment expression to actual emotion), and should be the focus of future work.

We notice that some of the treatments, such as chelation, are discussed prevalently in the communities. However, they are not necessarily options that are mostly taken in practice. Within the top treatment list that represents actual usage, non-chemical psychosocial interventions such as speech therapy and special education are popular, although they are not necessarily the most popular ones under discussion. Our results provide a clear evidence that users' perceptions, and hence actual adoptions, of treatment may not be accurately reflected by popularities in discussions, not to mention merely frequencies of certain keywords. More broadly, the results remind us that when connecting online content to members' real life actions in a quantitative way, hidden information (e.g. attributions) of content must be taken into account to avoid mis-interpretation of results.

Longitudinally, we found that members discuss treatment therapies with quite constant frequency in the communities throughout their participation. No clear pattern could be identified in terms of how sustained participation affects frequencies of discussions of certain treatments. The finding is different from previous research findings where user sentiment and topics both show clear patterns longitudinally [40, 37]. Moreover, frequencies of mentions attributed to patients fluctuate with total frequencies, and maintain substantial percentages in all the mentions throughout members' participation, which is somewhat counter-intuitive. In the future work, it is therefore an interesting question to explore how and why members keep mentioning treatment attributed to their ASD children throughout their participation.

The most important building block of future work following this study is to compare the list of treatments discovered in OHCs automatically by the computational tool with established clinical guidelines. For example, while effectiveness of chelation is still under investigation

by researchers [9], it already becomes a rather popular choice among autism community members. It is therefore critical to further quantify how broad the gap is between established guideline and patients' actual practice. The future work will contribute to understanding how information support and consumption in OHC affect members' decision makings regarding disease management, and hence how OHC participation makes physical and psychological impact.

References

1. [Accessed: 2017-02-02] Autism spectrum disorder: Overview. <http://www.asha.org/Practice-Portal/Clinical-Topics/Autism/>
2. Medwatch: The fda safety information and adverse event reporting program.
3. Open NLP. <https://opennlp.apache.org/>
4. Ahmad SR. Adverse drug event monitoring at the food and drug administration. *Journal of general internal medicine*. 2003; 18(1):57–60. [PubMed: 12534765]
5. Bailey A, Le Couteur A, Gottesman I, Bolton P, Simonoff E, Yuzda E, Rutter M. Autism as a strongly genetic disorder: evidence from a british twin study. *Psychological medicine*. 1995; 25(01): 63–77. [PubMed: 7792363]
6. Baio, J. surveillance summaries. Vol. 61. Centers for Disease Control and Prevention; 2012. Prevalence of autism spectrum disorders: Autism and developmental disabilities monitoring network, 14 sites, united states, 2008. morbidity and mortality weekly report.
7. Baird G, Cass H, Slonims V. Diagnosis of autism. *Bmj*. 2003; 327(7413):488–493. [PubMed: 12946972]
8. Bodenreider O. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*. 2004; 32(suppl 1):D267–D270. [PubMed: 14681409]
9. Davis TN, Oâ Reilly M, Kang S, Lang R, Rispoli M, Sigafos J, Lancioni G, Copeland D, Attai S, Mulloy A. Chelation treatment for autism spectrum disorders: A systematic review. *Research in Autism Spectrum Disorders*. 2013; 7(1):49–55.
10. de Bruijn B, Cherry C, Kiritchenko S, Martin J, Zhu X. Machine-learned solutions for three stages of clinical information extraction: the state of the art at i2b2 2010. *Journal of the American Medical Informatics Association*. 2011; 18(5):557–562. [PubMed: 21565856]
11. Del Vicario M, Bessi A, Zollo F, Petroni F, Scala A, Caldarelli G, Stanley HE, Quattrocioni W. The spreading of misinformation online. *Proceedings of the National Academy of Sciences*. 2016; 113(3):554–559.
12. U. Food, D. Administration, et al. Fda adverse event reporting system (faers). 2014.
13. Fox, S., Jones, S. The social life of health information. pew research center; 2009.
14. Friedman C, Hripcsak G, Shagina L, Liu H. Representing information in patient reports using natural language processing and the extensible markup language. *Journal of the American Medical Informatics Association*. 1999; 6(1):76–87. [PubMed: 9925230]
15. Frith, U. Autism: Explaining the enigma. 1989.
16. Harpaz R, Callahan A, Tamang S, Low Y, Odgers D, Finlayson S, Jung K, LePendu P, Shah NH. Text mining for adverse drug events: the promise, challenges, and state of the art. *Drug safety*. 2014; 37(10):777–790. [PubMed: 25151493]
17. Harpaz R, DuMouchel W, Shah NH, Madigan D, Ryan P, Friedman C. Novel data-mining methodologies for adverse drug event discovery and analysis. *Clinical Pharmacology & Therapeutics*. 2012; 91(6):1010–1021. [PubMed: 22549283]
18. Hartzler A, Pratt W. Managing the personal side of health: how patient expertise differs from the expertise of clinicians. *Journal of medical Internet research*. 2011; 13(3):e62. [PubMed: 21846635]
19. Høybye MT, Johansen C, Tjørnhøj-Thomsen T. Online interaction. effects of storytelling in an internet breast cancer support group. *Psycho-Oncology*. 2005; 14(3):211–220. [PubMed: 15386774]

20. Jiang M, Chen Y, Liu M, Rosenbloom ST, Mani S, Denny JC, Xu H. A study of machine-learning-based approaches to extract clinical entities and their assertions from discharge summaries. *Journal of the American Medical Informatics Association*. 2011; 18(5):601–606. [PubMed: 21508414]
21. Klein, D., Manning, CD. Accurate unlexicalized parsing. *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*; Association for Computational Linguistics; 2003. p. 423-430.
22. Lazarou J, Pomeranz BH, Corey PN. Incidence of adverse drug reactions in hospitalized patients: a meta-analysis of prospective studies. *Jama*. 1998; 279(15):1200–1205. [PubMed: 9555760]
23. Leaman, R., Wojtulewicz, L., Sullivan, R., Skariah, A., Yang, J., Gonzalez, G. Towards internet-age pharmacovigilance: extracting adverse drug reactions from user posts to health-related social networks. *Proceedings of the 2010 workshop on biomedical natural language processing*; Association for Computational Linguistics; 2010. p. 117-125.
24. Mamykina, L., Nakikj, D., Elhadad, N. Collective sensemaking in online health forums. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*; ACM; 2015. p. 3217-3226.
25. Miller GA. Wordnet: a lexical database for english. *Communications of the ACM*. 1995; 38(11): 39–41.
26. Muhle R, Trentacoste SV, Rapin I. The genetics of autism. *Pediatrics*. 2004; 113(5):e472–e486. [PubMed: 15121991]
27. Nakamura C, Bromberg M, Bhargava S, Wicks P, Zeng-Treitler Q. Mining online social network data for biomedical research: a comparison of clinicians' and patients' perceptions about amyotrophic lateral sclerosis treatments. *Journal of medical Internet research*. 2012; 14(3):e90. [PubMed: 22721865]
28. Owen JE, Klapow JC, Roth DL, Shuster JL, Bellis J, Meredith R, Tucker DC. Randomized pilot of a self-guided internet coping group for women with early-stage breast cancer. *Annals of Behavioral Medicine*. 2005; 30(1):54–64. [PubMed: 16097906]
29. Paul MJ, Sarker A, Brownstein JS, Nikfarjam A, Scotch M, Smith KL, Gonzalez G. Social media mining for public health monitoring and surveillance. *Pacific Symposium on Biocomputing*. 2016; 21:468.
30. Portier K, Greer GE, Rokach L, Ofek N, Wang Y, Biyani P, Yu M, Banerjee S, Zhao K, Mitra P, et al. Understanding topics and sentiment in an online cancer survivor community. *JNCI Monographs*. 2013; 47:195–198.
31. Salzer MS, Palmer SC, Kaplan K, Brusilovskiy E, Ten Have T, Hampshire M, Metz J, Coyne JC. A randomized, controlled study of internet peer-to-peer interactions among women newly diagnosed with breast cancer. *Psycho-Oncology*. 2010; 19(4):441–446. [PubMed: 19484712]
32. Sarker A, Ginn R, Nikfarjam A, O'Connor K, Smith K, Jayaraman S, Upadhaya T, Gonzalez G. Utilizing social media data for pharmacovigilance: A review. *Journal of biomedical informatics*. 2015; 54:202–212. [PubMed: 25720841]
33. Smith, CA., Stavri, PZ. *Consumer Health Informatics*. Springer; 2005. Consumer health vocabulary; p. 122-128.
34. Tang B, Cao H, Wu Y, Jiang M, Xu H. Recognizing clinical entities in hospital discharge summaries using structural support vector machines with word representation features. *BMC medical informatics and decision making*. 2013; 13(1):1. [PubMed: 23289362]
35. Tuarob S, Tucker CS, Salathe M, Ram N. An ensemble heterogeneous classification methodology for discovering health-related knowledge in social media messages. *Journal of biomedical informatics*. 2014; 49:255–268. [PubMed: 24642081]
36. Wang Y-C, Kraut RE, Levine JM. Eliciting and receiving online support: using computer-aided content analysis to examine the dynamics of online social support. *Journal of medical Internet research*. 2015; 17(4):e99. [PubMed: 25896033]
37. Zhang, S., Bantum, E., Owen, J., Elhadad, N. Does sustained participation in an online health community affect sentiment?. *AMIA Annual Symposium Proceedings*; American Medical Informatics Association; 2014. p. 1970

38. Zhang S, Bantum EO, Owen J, Bakken S, Elhadad N. Online cancer communities as informatics intervention for social support: conceptualization, characterization, and impact. *Journal of the American Medical Informatics Association*. 2016:ocw093.
39. Zhang S, Elhadad N. Unsupervised biomedical named entity recognition: experiments with clinical and biological texts. *Journal of biomedical informatics*. Dec; 2013 46(6):1088–98. [PubMed: 23954592]
40. Zhang S, Grave E, Sklar E, Elhadad N. Longitudinal analysis of discussion topics in an online breast cancer community using convolutional neural networks. 2016 arXiv preprint arXiv: 1603.08458.




Figure 1. Distributions of number of users, by number of used treatment. The x axis is the number of used treatment identified, and the y axis is the number of users.

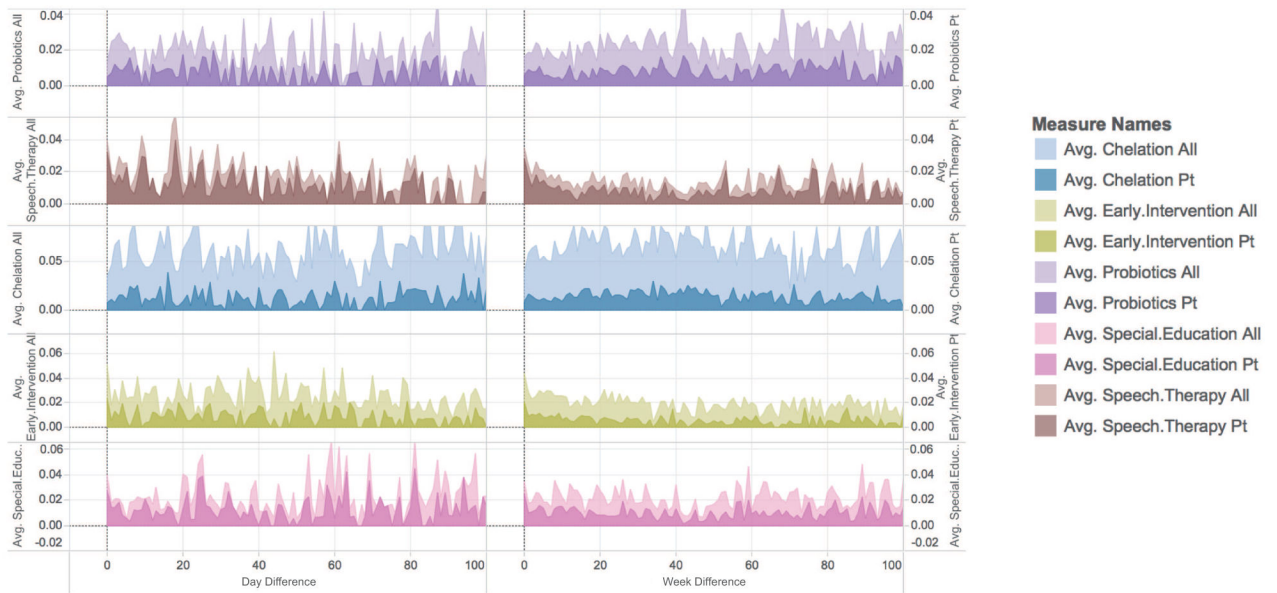


Figure 2.

Changes of frequencies (mention per post) of top five treatments in autism communities, since members joining the community. Two separate X-axes represent views in weeks (right) and in days (left), respectively. Variables (measure names) ending with “all” represent total frequencies of mentions of corresponding treatment, regardless of their attribution types. Variables ending with “pt” represent frequencies of mentions of attribution type *pt*.

Table 1

Descriptive statistics of the ASD dataset

Number of sub-forums	16
Number of threads	61,817
Number of posts	551,029
Number of authors	10,210

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2

Attribution labels for treatment mentions, their descriptions and their counts in dataset.

Attribution label	Description	Count
<i>pt</i>	Mention of treatment which indicates an actual usage or usage history of the patients of interest, usually children of the users in this particular study.	1830
<i>pt-gen</i>	Mention of treatment tied to the patient but does not indicate actual usage.	434
<i>cg</i>	Mention of treatment tied to the caregiver of the patient, usually the user herself.	95
<i>others</i>	Mention of treatment tied to specific individuals other than the caregiver or the patient. Can be other members in the community, or other people in the author's real life.	210
<i>gen</i>	Mention not tied to a specific individual.	1635

Table 3

System performance for binary treatment mention detection with different types of features.

Features	Precision	Recall	F
Baseline	78.2	56.5	65.6
lexical	82.1	83.6	82.9
lexical+semantic	81.4	83.8	82.6
lexical+semantic+syntactical	81.0	83.5	82.2

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 4

System performance (F score) for joint treatment detection and attribution classification with different types of features. The baseline systems use **dictionary matching** for term identification, and logistic regression with corresponding features for attribution classification. The baseline+ systems use **CRF** for term identification, and logistic regression with corresponding features for attribution classification.

Features	micro	cg	gen	others	pt	pt-gen
baseline (lexical)	33.4	9.8	31.2	29.1	45.0	13.4
baseline+ (lexical)	52.7	17.6	54.4	34.8	60.8	15.6
crf (lexical)	55.4	18.2	56.0	37.0	61.6	19.2
baseline (lexical+semantic)	34.1	10.2	31.1	29.0	45.6	14.5
baseline+ (lexical+semantic)	52.9	17.5	55.1	35.4	60.9	17.0
crf(lexical+semantic)	56.1	18.1	57.4	36.8	61.7	20.9
baseline (lexical+semantic+syntax)	35.0	14.6	29.9	31.2	44.1	15.4
baseline+ (lexical+semantic+syntax)	60.3	17.9	62.8	45.5	62.9	30.0
crf (lexical+semantic+syntax)	62.3	18.2	64.1	51.6	66.8	34.1

Table 5

System performance (CRF) for mentions with *pt* attribution with different types of features, when all other types of attributions are merged into one as *non-pt*.

Features	Prec. (pt)	Rec. (pt)	F (pt)
lexical	61.0	56.6	58.7
lexical+semantic	63.0	58.9	60.9
lexical+semantic+syntax	68.7	64.8	66.7

Table 6

Top 10 treatment with number of mentions for the five attribution classes, identified in the entire data set.

Term	Frequency	Term	Frequency
<i>pt</i>		<i>pt-gen</i>	
chelation	4935	chelation	1259
probiotics	2498	probiotics	389
zinc	2011	chelating	210
enzymes	1705	speech therapy	99
melatonin	1425	probiotic	98
special education	1287	activated charcoal	77
antibiotics	1283	nystatin	75
speech therapy	1245	melatonin	73
early intervention	1061	calcium	70
magnesium	889	early intervention	66
<i>cg</i>		<i>others</i>	
chelation	16	probiotics	424
progesterone	7	chelation	408
probiotics	7	probiotic	163
cod liver oil	5	chelating	150
chelator	4	melatonin	121
cab	4	enzymes	117
molybdenum glycinate chelate	4	zinc	114
sensory integration	4	risperdal	80
aloe vera	3	charcoal	77
pyridoxine hydrochloride	3	homeopathy	76
<i>gen</i>			
chelation	8341		
vitamin	1418		
early intervention	1268		
probiotics	1267		
special education	1153		
chelator	910		
vitamins	886		
melatonin	877		
homeopathy	862		
thimerosal	801		

Table 7

Top 10 treatment by number of users, identified in the ASD data set.

Term	Number of users
probiotics	819
speech therapy	565
chelation	520
early intervention	475
special education	395
melatonin	391
antibiotics	381
enzymes	352
zinc	332
vitamins	283

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript