

Collecting Better Training Data using Biased Agent Policies in Negotiation Dialogues

Vasily Konovalov[†] Oren Melamud[†] Ron Artstein[‡] Ido Dagan[†]

[†]Computer Science Department, Bar-Ilan University, Ramat-Gan, Israel
vaskonov@yahoo.com, {melamuo, dagan}@cs.biu.ac.il

[‡]USC Institute for Creative Technologies, USA
artstein@ict.usc.edu

Abstract. When naturally occurring data is characterized by a highly skewed class distribution, supervised learning often benefits from reducing this skew. Human-agent dialogue data is commonly highly skewed when using standard agent policies. Hence, we suggest that agent policies need to be reconsidered in the context of training data collection. Specifically, in this work we implemented biased agent policies that are optimized for data collection in the negotiation domain. Empirical evaluations show that our method is successful in collecting a reasonably balanced corpus in the highly skewed Job-Candidate domain. Furthermore, using this balanced corpus to train a negotiation intent classifier yields notable performance improvements relative to naturally distributed data.

Keywords: Dialogue System, Crowd-sourcing, Natural Language Understanding

1 Introduction

Various real-world datasets have highly skewed class distributions. Such skew presents a major challenge to supervised learning systems [2]. For example, consider an outlier class detection task, as in the detection of oil spills in satellite radar images [13], or the detection of fraudulent telephone calls [9]. In such tasks, more than 99% of the data could be labeled as ‘normal’, and hence a classifier trained to optimize prediction accuracy on this data, may achieve 99% accuracy by simply predicting ‘normal’ at all times, yet fail to achieve its real goal of detecting the outlier instances.

To better reflect performance on skewed datasets, alternative measures to accuracy are used, such as precision, recall and F1, and various strategies were developed to optimize them. While the natural distribution is considered best for optimizing accuracy, a balanced class distribution is considered better for optimizing precision and recall measures [24].

Indeed, class balancing is a prominent approach in supervised machine learning for addressing skewed datasets. One popular method is to sample the original training data set in a way that generates a more balanced sampled dataset. This is done either by *under-sampling* or *over-sampling*.

Another method is Active Class Selection, which actively selects the class labels from which training instances are to be generated.¹

Dialogue corpora typically suffer from a skewed distribution of the classes representing the intents expressed in each utterance. This is reported, for example, as a major cause of misclassification in the ATIS corpus [23]. The skew typically hurts classification performance on low frequency intents, whose correct detection by automated agents is nevertheless critical for the dialogue’s successful progress.

In this work, we propose to bias the policies of the automated agent when collecting a dialogue training set, in order to obtain a more balanced intent distribution. We apply our approach to a chat-based human-agent negotiation dialogue system in the *Job-Candidate* domain. To make the collection process as similar as possible to the expected human-agent interaction, we collect it using this same system, following the Wizard-of-Oz (WOZ) approach. Under this approach, a human ‘wizard’ maps the natural language utterances of the human user into the semantic language of the automated agent, thereby facilitating the online collection of the corpus without requiring an automated natural language understanding (NLU) component, which is yet to be trained [8]. While highly effective, this data collection procedure is also unfortunately very expensive and time-consuming to perform. Therefore, we focus on exploiting a typical relatively small train set size.

To solicit the low-frequency intents from the human user, while keeping the dialogue flow as natural as possible, we employ two techniques: (1) During dialogue, we dynamically modify the agent’s policies to increase the probability of low-frequency response types from the human user (e.g. make a ‘good’ offer in order to solicit an ‘Accept’ response); and (2) We ask the human user to rephrase their last utterance of a low-frequency intent. Our empirical evaluations show that a standard SVM classifier trained over our biased corpus notably outperforms a naturally distributed baseline of the same size, as well as biased corpora generated by using standard under-sampling and over-sampling techniques.²

2 Negotiation Dialogue System

Dialogue System Architecture Our chat-based negotiation dialogue system is based on the NCAgent negotiation agent [18], as described in Figure 1. It includes the following components: (1) Natural Language Understander (NLU) - maps natural language human utterances to a labeled semantic representation. (2) Dialogue Manager (DM) and Agent - the DM keeps track of the dialogue state and maintains conversation flow, while the Agent makes the negotiation decisions. (3) Natural Language Generator (NLG) - translates Agent statements to natural language utterances.

¹ Active Class Selection is not to be confused with Active Learning [22]. The latter chooses the training instances that should be manually labeled, while the former selects the class labels from which training instances should be generated [16].

² Our code is freely available at: <https://github.com/vaskonov>

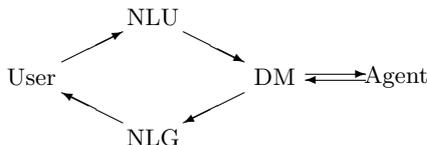


Fig. 1: Dialogue system architecture.

Negotiation Domain While our approach should be broadly applicable to the negotiation domain, in this work we applied it specifically to the *Job-Candidate* sub-domain, which was used in several previous works [15,17,12].

This domain is compatible with Genius³, a general platform for negotiation agents, which was also adapted by the Automated Negotiating Agents Competition (ANAC) as the official competition platform [14,1]. Our research concentrates on the technique that will allow anyone to collect a small and a well-balanced training set for any Genius-compatible domain with minimal efforts.

The negotiation takes place between an *employer* and a *candidate*. In our experiments the agent is always the candidate, while the human is the employer. The goal of both sides is to reach a consensus on several issues in order to sign a hiring agreement, while optimizing their own score objective.

The list of issues, denoted as attributes, with their corresponding values can be found in Table 1. Each party to the negotiation is given a scorecard representing their subjective utilities for each attribute. The negotiation dialogue comprises the exchange of one or more of the following intents in each utterance: *Offer*, *Accept*, *Reject*, *Query*, *Greet*, *Quit*, where all but the last two are considered crucial for maintaining the dialogue flow. This list of the possible intents was defined in [19].

Attributes	Values
<i>Salary</i>	60.000 USD, 90.000 USD, 120.000 USD
<i>Working hours</i>	8 hours, 9 hours, 10 hours
<i>Leased car</i>	With leased car, Without leased car, No agreement
<i>Pension fund</i>	0%, 10%, 20%, No agreement
<i>Promotion possibilities</i>	Fast promotion track, Slow promotion track, No agreement
<i>Job position</i>	QA, Programmer, Team manager, Project manager

Table 1: The list of *JobCandidate*'s attributes and their possible values

The NegoChat Corpus The NegoChat negotiation corpus was collected in the Job-Candidate domain using the NCAgent [12]. The NegoChat corpus comprises of 100 dialogues with 1409 human utterances. Each utterance in this corpus is labeled

³ <http://ii.tudelft.nl/genius/>

with one or more intent and an associated issue and value. For example, the utterance “I offer you a pension of 10%” is labeled as (*intent=Offer, issue=Pension fund, value=10%*).

Some utterances can be labeled by intents and issues only, for example, “I accept the suggested salary” is labeled as (*intent=Accept, issue=Salary*), while some utterances can be labeled by intents only, “I reject your offer” (*intent=Reject*). *Query* is a special intent that can be used to ask an opponent for a preferred value of a specific issue or an entire offer. For example, the utterance “What position are you looking for?” is labeled as (*intent=Query, issue=Job position*) and the utterance “What is your offer?” is labeled as (*intent=Query*).

The intents are considered domain-independent, while the attributes and values are specific to the particular type of a negotiation domain.

Despite the limited number of intents, the suggested semantic language covers 93% of the human utterances, collected in previous experiments [12]. Among the uncovered utterances there are emotional expressions like “Welcome to the company”, and utterances without any negotiation content “Here’s an offer, take a look”. Possible extension to the semantic language can be found in [12].

The ‘natural’ column in Table 2 shows the highly skewed distribution of the recorded intents in the NegoChat corpus.

	natural		biased with rephrase		biased no rephrase	
	count	%	count	%	count	%
<i>Offer</i>	840	59	163	33	137	37
<i>Accept</i>	237	16	200	40	130	35
<i>Reject</i>	241	17	102	20	78	21
<i>Query</i>	74	5	26	5	14	5
<i>Greet</i>	29	2	6	1	6	1
<i>Quit</i>	18	1	1	1	1	1

Table 2: The distribution of intents.

3 Dialogue Biasing Strategies

We propose the following strategies for biasing the course of a negotiation dialogue with the goal of collecting a more balanced training corpus.

3.1 Manipulating Agent Policies

‘Offer’ is the most common intent in our domain by a large margin, while the ‘reactive’ intents, namely ‘Accept’ and ‘Reject’, are comparatively rare. This is mainly because users do not tend to accept the first offer in a negotiation, and

frequently respond to an offer with a counteroffer instead of explicitly rejecting it. To increase the proportion of the reactive intents relative to ‘Offer’ in the human utterances, we perform the following manipulations in the automated agent policies. First, the agent does not respond with an ‘Accept’ to an offer made by the user, but rather makes an offer regarding a different issue, as long as there are other issues that have not yet been discussed. With this strategy the agent takes more of an initiating role, pushing the human into a reactive position. Second, when making an offer, the agent makes a decision whether to solicit either an ‘Accept’ or ‘Reject’ response, by sampling a Bernoulli distribution, with $p(\text{Accept}) = \frac{N_{\text{Reject}}}{N_{\text{Reject}} + N_{\text{Accept}}}$, where N_x is the number of responses of type x , already recorded in the current dialogue.

When soliciting a ‘Reject’ response, the agent makes an offer according to its ‘regular’ policies, but when soliciting an ‘Accept’ response, it chooses a value which optimizes the user’s objectives, or in other words it makes a ‘good’ offer. This aims to balance the distribution between the ‘Accept’ and ‘Reject’ utterances.

3.2 Rephrasing Requests

As an additional simple strategy to solicit less-frequent intents, we used rephrasing requests. At each point in the dialogue where the user’s utterance includes any intent but the most frequent one up to that point, the agent asks for a rephrase by uttering something like: “*I didn’t understand your last sentence, please rephrase it in other words*”.

The rephrasing strategy has a major limitation, since for asking a human to rephrase an utterance of an underrepresented intent, the human should first proactively express this intent. Therefore, the rephrasing component is used along with the previous strategy of manipulating the agent’s policies.

Yet, we cannot use the rephrasing strategy without any limit, since at some point the human participant will get annoyed by an agent repeatedly asking to rephrase every utterance. Therefore, the agent refrains from asking consecutively for rephrases of the same intent type. The rephrasing component is thus used as an additional opportunity to achieve more utterances of the under-represented intents. The contribution of the rephrasing component to the intent classification is explored in Section 4.2. The success rate of soliciting rephrases for intent type is presented in Section 5.

4 Evaluation and Results

We used several methods to generate different training datasets, then used them to train respective supervised intent classifiers and finally compared the performance of the different classifiers using the very same test corpus. We note that while the generated training datasets may (intentionally) not reflect the natural course of dialogues, the shared test dataset was generated in the ‘natural’ setting. We report learning curve performance, where we plot F1-measure results as a function of an incrementally growing train set size and focus on small train set sizes.

4.1 Generated Training Corpora

The following training datasets were compared:

Natural The 100 NegoChat dialogues collected in [12] were used as the naturally distributed dataset.

Over/Under-sampled We used over/under sampling to generate a balanced training dataset from the natural dataset. The over-sampling method randomly re-samples the utterances of all but the most frequent intent, until obtaining a fully balanced training set. The under-sampling method, instead, randomly discards utterances of all but the least frequent intent to reach a balanced distribution [10]. We re-sampled the naturally distributed train set at every point in the learning curve, while reporting the original number of the utterances at every point.

Biased To collect biased dialogues according to our method, we recruited 30 English speakers on Amazon Mechanical Turk (AMT) [5] and collected 30 biased dialogues comprising of 356 human and 392 agent utterances. To estimate the contribution of the rephrasing component, we performed an ablation test, where we compared the performance between the biased set with rephrases and the biased set without rephrases.

An example of an annotated dialogue appears in the Appendix.

The intent distributions of the datasets generated by the natural and biased methods appears in Table 2 (the over/under-sampling methods are fully balanced). As an initial result it can be seen that our biasing method is indeed effective in largely balancing the skewed intent distribution even without using the rephrase component.

4.2 Intent Classification

We used an SVM [20] for intent classifier [23] to evaluate the effectiveness of the compared train sets.⁴ The classifier converts the text of an utterance into a set of input features. As input features to the classifier we used unigrams and context features. The context features indicate how the issues and the values, found in the human’s utterance, were used in the previous agent’s utterance. The insight behind the context features is that in order for a human to ‘Accept’ some issue, the issue should be ‘Offer’ed by an agent in previous turns, and similarly for ‘Reject’. Contextual intent prediction was explored in detail in [3]. At the end of the preprocessing, we discard the parts of the text that describe the issues and their values (as identified by the rule-based component), in order to let the classifier focus more easily on the words that express intents.

We compared the performance of the main⁵ intents of the above train sets with a 10-folds cross validation experiment, as follows. The natural dataset was divided into 10 folds with 10 dialogues in each. Each time, a single fold was

⁴ we used the LibSVM implementation [7].

⁵ ‘Greet’ and ‘Quit’ intents were not included as they are rare and often completely missing in a dialogue.

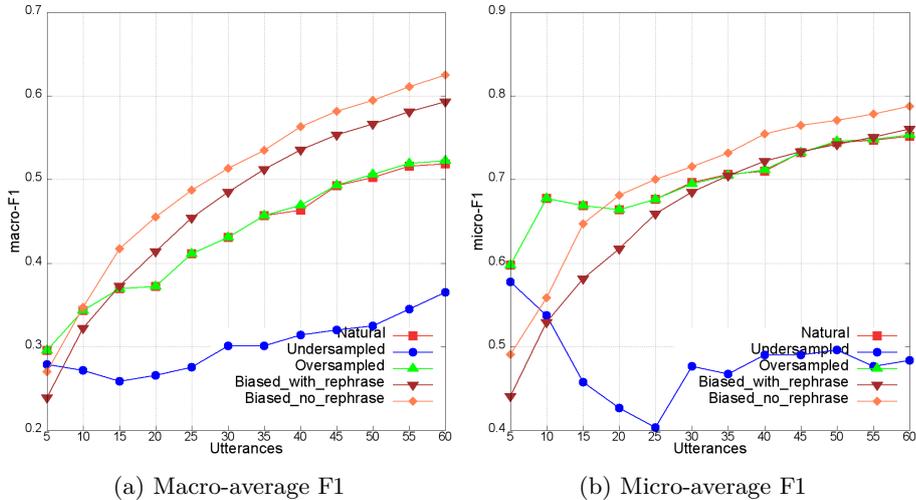


Fig. 2: F1 performance measures of intent classification.

designated as the train set and the remaining nine folds as the test set. To make a fair comparison with the other methods, we kept the exact same nine folds of natural data as test set, while replacing only the training data per each method. For the over and under sampling methods, we applied these methods to balance the original natural train fold. For our biased method, we randomly sampled each time 10 dialogues from our 30 biased dialogues as the train set. We repeated this procedure 10 times and report the mean F1 results.

Figure 2a shows the macro average F1 results as a function of the number of manually labeled utterances used to train the classifier. As seen, the classifier trained on our biased dialogues notably outperforms all other baselines, and it seems that this advantage is gained even without the rephrase component. On the other hand, the over/under-sampling methods do not seem effective in improving performance over the natural baseline.

Macro-average measures give equal importance to classes of different frequencies (in the test set). For completeness, in Figure 2b we also report performance using the micro-average measure, which gives more weight to frequent classes [21]. As seen, the micro-average performance of our biasing method is almost identical to that of the natural baseline. In both cases, the performance of the natural distribution is almost identical to the one of the oversampled distribution.

By definition, the oversampled set of utterances contains the entire set of the available utterances plus some duplicates of the rare classes utterances. The results seem to indicate that in our case, where the training set is small, adding duplicated instances is less likely to influence class separability, therefore the performance of the oversampled baseline is almost identical to the performance of the natural distribution.

On the other hand, the size of the under-sampled set of utterances is always smaller than the size of the original natural set, when it is unbalanced. Specifically, the size of the under-sampled set equals the number of classes times the number of the instances of the least frequent class. In our case, almost every fold of the under-sampled set ends up with few instances of each class, even though the original naturally-distributed set has tens of utterances. Interestingly, while using under-sampling and computing learning curves, we observed some instability of the learning behavior. In our case, the least frequent intent is ‘Query’. When the first ‘Query’ utterance occurs at the late steps of the learning curves it resets the counters of the other intents to one, which explains the unstable form of the under-sampling learning curve. It was shown previously that under-sampling is not a very useful scheme in case of small training sets, when it omits a lot of useful training instances [11].

Figure 3 shows F1 results for the individual intents. The biased dialogues attained better performance on all intents.

5 Analysis

In Table 3 we analyse the effectiveness of our method to solicit an ‘Accept’/‘Reject’ response as described in Section 3.1. As expected, soliciting an ‘Accept’ response by having the agent adopt the human objectives is mostly successful. However, in the opposite case, rather than responding with a ‘Reject’ the human user mostly reacts with another ‘Offer’, mostly a counteroffer.

	solicit ‘Accept’	solicit ‘Reject’
<i>Offer</i>	11	58
<i>Accept</i>	88	23
<i>Reject</i>	0	18
<i>Query</i>	1	1

Table 3: The distribution of user intents in response to our ‘Accept’/‘Reject’ solicitation (%).

The success ratio in soliciting a rephrase of the intent in the last human utterance is shown in Table 4, where columns show desired intents and rows show actual responses. Interestingly, asking a rephrase for a ‘Reject’ is not very successful. A closer look at the data reveals that this is because in such cases the human chooses to respond with an ‘Offer’ instead.

In order to measure the degree of balancing achieved through both biasing and rephrasing vs. biasing only, we calculated the *total variation distance* from the uniform ‘balanced’ distribution. The total variation distance between two probability distributions P and Q is defined as [6]: $\delta(P, Q) = \frac{1}{2} \sum_x (|P(x) - Q(x)|)$. We found that, the distance between the uniform distribution and the distribution of the biased set without rephrases is slightly smaller than the distance for the

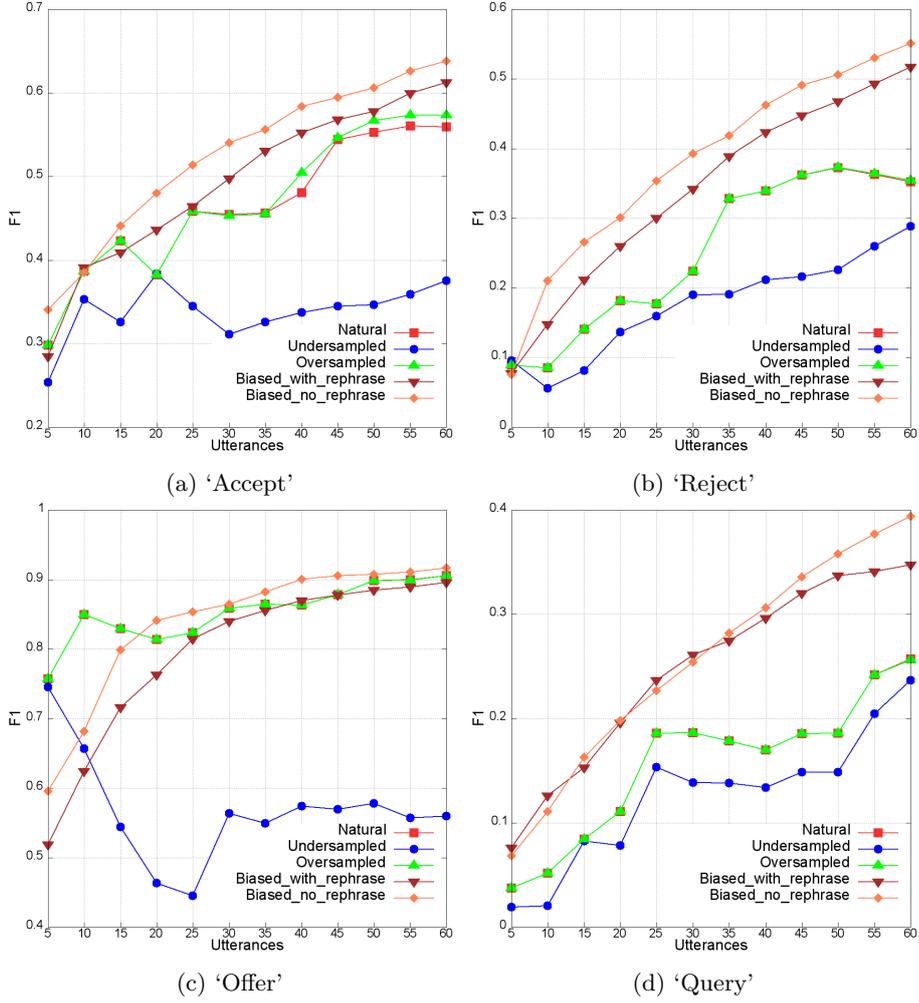


Fig. 3: F1 performance of intent classification.

	<i>Offer</i>	<i>Accept</i>	<i>Reject</i>	<i>Query</i>
<i>Offer</i>	100	3	46	11
<i>Accept</i>	0	94	0	0
<i>Reject</i>	0	0	48	0
<i>Query</i>	0	3	6	89

Table 4: The success rate in soliciting a rephrase for a target intent (%)

biased dataset with rephrases. This means that the rephrasing strategy does not contribute to the degree of balancing obtained by the biased distribution.

6 Conclusion

The skewed class distribution in dialogue corpora and the high cost associated with collecting them are critical challenges in successfully training supervised NLUs for dialogue agents. In this work we proposed to bias the policies of the agent when collecting training data in order to reduce this skew. Evaluation in our negotiation domain yielded notable improvement in supervised intent classification performance. To the best of our knowledge we are the first to introduce the concept of biasing the course of a dialogue in the interest of collecting a more effective training corpus. We believe that there is a good potential for further research on this direction, and make our biased corpus and code publicly available to promote such research. We observed that the agent with the proposed biasing approach cannot completely gain initiative in the course of the negotiation dialogues, since in most cases a human takes initiative by providing counteroffers. Therefore, future work should explore more advanced methods for gaining initiative in negotiation dialogues.

Further, even though in our work the rephrasing strategy did not improve performance, we believe that by exploring different rephrasing strategies [4] we may achieve better performance.

7 Acknowledgments

We thank Sarit Kraus, Avi Rosenfeld, Erel Segal-Halevi, Osnat Drein and Inon Zuckerman for their assistance and contribution. This work was partly supported by ERC Grant #267523.

References

1. Baarslag, T., Fujita, K., Gerding, E.H., Hindriks, K., Ito, T., Jennings, N.R., Jonker, C., Kraus, S., Lin, R., Robu, V., et al.: Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. vol. 198, pp. 73–103. Elsevier (2013)
2. Batista, G.E., Prati, R.C., Monard, M.C.: A study of the behavior of several methods for balancing machine learning training data. *ACM Sigkdd explorations newsletter* 6(1), 20–29 (2004)
3. Bhargava, A., Celikyilmaz, A., Hakkani-Tür, D., Sarikaya, R.: Easy contextual intent prediction and slot detection. In: 2013 IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 8337–8341. IEEE (2013)
4. Bohus, D., Rudnicky, A.I.: Sorry, I didn't catch that! An investigation of non-understanding errors and recovery strategies. In: 6th SIGdial workshop on discourse and dialogue (2005)
5. Buhrmester, M., Kwang, T., Gosling, S.D.: Amazon's mechanical turk a new source of inexpensive, yet high-quality, data? *Perspectives on psychological science* 6(1), 3–5 (2011)
6. Chambolle, A.: An algorithm for total variation minimization and applications. *Journal of Mathematical imaging and vision* 20(1-2), 89–97 (2004)

7. Chang, C.C., Lin, C.J.: Libsvm: a library for support vector machines. In: *ACM Transactions on Intelligent Systems and Technology*. vol. 2, pp. 27:1–27:27 (2011)
8. Dahlbäck, N., Jönsson, A., Ahrenberg, L.: Wizard of Oz studies: why and how. In: *Proceedings of the 1st international conference on intelligent user interfaces*. pp. 193–200. ACM (1993)
9. Fawcett, T., Provost, F.: Adaptive fraud detection. *Data mining and knowledge discovery* 1(3), 291–316 (1997)
10. Japkowicz, N.: Learning from imbalanced data sets: a comparison of various strategies. In: *AAAI workshop on learning from imbalanced data sets*. vol. 68, pp. 10–15. Menlo Park, CA (2000)
11. Japkowicz, N., Stephen, S.: The class imbalance problem: a systematic study. *Intelligent data analysis* 6(5), 429–449 (2002)
12. Kononov, V., Artstein, R., Melamud, O., Dagan, I.: The Negochat corpus of human-agent negotiation dialogues. In: *LREC* (2016)
13. Kubat, M., Matwin, S., et al.: Addressing the curse of imbalanced training sets: one-sided selection. In: *ICML*. vol. 97, pp. 179–186. Nashville, USA (1997)
14. Lin, R., Kraus, S., Baarslag, T., Tykhonov, D., Hindriks, K., Jonker, C.M.: Genius: an integrated environment for supporting the design of generic automated negotiators. In: *Proceedings of ANAC*. Wiley Online Library (2012)
15. Lin, R., Kraus, S., Wilkenfeld, J., Barry, J.: Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence* 172(6), 823–851 (2008)
16. Lomasky, R., Brodley, C.E., Aernecke, M., Walt, D., Friedl, M.: Active class selection. In: *Machine learning: ECML 2007*, pp. 640–647. Springer (2007)
17. Oshrat, Y., Lin, R., Kraus, S.: Facing the challenge of human-agent negotiations via effective general opponent modeling. In: *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. pp. 377–384. International Foundation for Autonomous Agents and Multiagent Systems (2009)
18. Rosenfeld, A., Zuckerman, I., Segal-Halevi, E., Drein, O., Kraus, S.: Negochat a chat-based negotiation agent. In: *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. pp. 525–532. International Foundation for Autonomous Agents and Multiagent Systems (2014)
19. Rosenfeld, A., Zuckerman, I., Segal-Halevi, E., Drein, O., Kraus, S.: Negochat: a chat-based negotiation agent with bounded rationality. *Autonomous Agents and Multi-Agent Systems* 30(1), 60–81 (2016)
20. Schölkopf, B., Smola, A.J., Williamson, R.C., Bartlett, P.L.: New support vector algorithms. *Neural computation* 12(5), 1207–1245 (2000)
21. Schütze, H.: Introduction to information retrieval. In: *Proceedings of the international communication of association for computing machinery conference* (2008)
22. Settles, B.: Active learning literature survey. *University of Wisconsin, Madison* 52(55-66), 11 (2010)
23. Tur, G., Hakkani-Tur, D., Heck, L.: What is left to be understood in ATIS? In: *Spoken Language Technology Workshop (SLT), 2010 IEEE*. pp. 19–24. IEEE (2010)
24. Weiss, G.M., Provost, F.: Learning when training data are costly: the effect of class distribution on tree induction. *Journal of Artificial Intelligence Research* pp. 315–354 (2003)

Appendix: An example of an annotated dialogue

```
[
  {
    "role": "Agent",
    "data": "Hello. I would like to discuss the issues of my contract."},
  {
    "role": "Agent",
    "desired intent": "Reject",
    "input": "I am willing to have a position of project manager",
    "output": { "Offer": { "Job Description": "Project Manager" } }},
  {
    "role": "Human",
    "input": { "text": "I agree" },
    "output": { "Accept": true }},
  {
    "role": "Agent",
    "data": "I didn't understand your last sentence, please rephrase it in other words"},
  {
    "role": "Human",
    "input": { "text": "I accept this position" },
    "output": { "Accept": "Job Description" }},
  {
    "role": "Agent",
    "desired intent": "Reject",
    "input": "I am expecting a salary of 120,000",
    "output": { "Offer": { "Salary": "120,000 USD" } }},
  {
    "role": "Human",
    "input": { "text": "I was thinking more like 90,000" },
    "output": { "Offer": { "Salary": "90,000 USD" } }},
  {
    "role": "Agent",
    "desired intent": "Reject",
    "input": "I am expecting a pension of 20%",
    "output": { "Offer": { "Pension Fund": "20%" } }},
  {
    "role": "Human",
    "input": { "text": "pension at 20% is fine" },
    "output": { "Accept": { "Pension Fund": "20%" } }},
  {
    "role": "Agent",
    "data": "I didn't understand your last sentence, please rephrase it in other words"},
  {
    "role": "Human",
    "input": { "text": "You can get a pension of 20%" },
    "output": { "Accept": { "Pension Fund": "20%" } }},
  {
    "role": "Agent",
    "desired intent": "Accept",
    "input": "I would like a workday of 10 hours",
    "output": { "Offer": { "Working Hours": "10 hours" } }},
  {
    "role": "Human",
    "input": { "text": "10 Hour work day is great, i agree" },
    "output": { "Accept": { "Working Hours": "10 hours" } }},
  {
    "role": "Agent",
    "desired intent": "Reject",
    "input": "I would like a company car",
    "output": { "Offer": { "Leased Car": "With leased car" } }}
]
```