Recommendations for Recruiters with Sentiment Detection

Ashish Lakhani
iCIMS
Fremont, CA, USA
Alakha@icims.com

ABSTRACT

Recommendation systems based on AI are widely used by technology companies, especially in the retail and entertainment domains. When you buy products on Amazon, for example, it recommends various other products based on what you have purchased in the past. Similarly, when you watch movies on Netflix, it recommends other movies based on your past viewing history.

Recommendation systems are a win-win for both consumers and service providers. For consumers, they are beneficial because they give them personalization based on their interests and affinities. For service providers, it creates an advertisement opportunity and brings their customers closer to the next items they may purchase on their platform.

While there are well-known methods for making product recommendations in domains like retail and entertainment, the recruiting industry has historically lacked well-known methods of providing similar recommendations for candidates and jobs.

Recruiters are looking to select the best candidates from a typically large pool of candidates. Many recruiting technology vendors offer a way to filter candidates based on their skills and experience, but even then, the pool of candidates is often unmanageably large. Previously, there was a gap in the hiring software market for narrowing down the applicant list by more relevant criteria, such as differentiating between ordinary candidates and highly motivated talent.

iCIMS identified the recruiters need and addressed it in the iCIMS Talent Cloud. In this paper, we explore how iCIMS could use intent and sentiment detection to build a recommendation system for recruiters.

RecSys in HR 2021, October 1, 2021, Amsterdam, the Netherlands 2021.

1 Basics about NLP, Sentiment, and Intent Detection

Natural language Processing (NLP) is a field of computer science — more specifically, a field of Artificial Intelligence—which deals with providing machines the ability to understand human language. Sentiment detection and intent detection are both Natural Language Processing tasks.

Sentiment analysis is a process of analyzing a piece of text and identifying the emotion behind it. Sentiment can be either positive, negative, or neutral.

It is also important to note that sentiment analysis is a domain-specific task [1], and what may be considered positive in one domain may not be positive in another domain. This is because sentiment expressions used in different domains are often unique [2].

Lets look at few examples. The words “boring” and “lengthy” are used to express negative sentiments in book domain; however, they are rarely used in the electronics domain [3].

It is also possible that the same words may express different sentiments in different domains. The word “easy” is used to express negative sentiments in movie domain, such as, “Movie was easy to guess.” However, the word “easy” is used to express positive sentiments in the electronics domain, e.g., “Camera was easy to use.” [2].

Thus, a sentiment classifier trained on one domain cannot be used to infer sentiments in another domain [4].

Intent detection is the methodology to understand the motive behind the user’s text. For example, in the sentence “I want to read the book,” the intent of the user is to read a book.

Let us consider a few sentiment examples in the recruiting domain. In talent acquisition, certain events—such as a candidate looking for a job, trying to schedule an interview, inquiring about job requirements, or asking for the next steps—are considered positive events. On the other hand, a candidate's lack of interest in the job, stalling the process, indecision on what they need to do, or canceling the interview are considered negative events.

Below is a table that outlines examples of sentiment and intent from candidate messages.

Examples	Sentiment	Intent
I have plenty of customer service experience and should easily fit this position.	Positive	The candidate is stating their experience
I already filled out my application. I just need to set up my interview.	Positive	The candidate is seeking an interview
Which holidays do your employees get off?	Neutral	The candidate wants to know the provided benefits
What are your hours of operation?	Neutral	The candidate is inquiring about normal hours of operation
I need to cancel my interview	Negative	The candidate wants to cancel the interview
Sorry, I have a job already	Negative	The candidate is not interested in the job

2 What is iCIMS?

iCIMS is the talent cloud company that empowers organizations to attract, engage, hire, and advance the right talent that builds a diverse, winning workforce. We accelerate transformation for a community of more than 4,000 customers, including 40 percent of the Fortune 100.

2.1 iCIMS Talent Cloud

The iCIMS Talent Cloud can help you stay engaged with top talent along every step of the talent acquisition lifecycle, deepen your talent pool, and fill positions faster by building strong connections with candidates at scale.
The iCIMS Talent Cloud helps recruiters to:

- Continually engage with candidates in their talent pipeline through robust talent pools.
- Send targeted, automated text and email campaigns that encourage candidates to apply faster.
- Host virtual career fairs to cut costs and expand their reach
- Get powerful recruitment marketing reporting and source analytics that helps measure ROI and efficiency

Using the iCIMS Talent Cloud, recruiters can make use of this recommendation system to engage with quality talent.

3 Training Data

At iCIMS we provide tools to recruiters that help them communicate with candidates via various channels, such as email, SMS, WhatsApp, Facebook, and more.

Training data examples described below occur between candidates and recruiters mostly via text messages. With the iCIMS platform recruiters can contact candidates via text messages and corresponding communication occurs between them via SMS.

After analyzing a vast number of conversations happening between recruiters and candidates on the iCIMS platform, we observed that conversations could be broadly classified into the below categories.

3.1 Extremely Positive Messages

These are messages from candidates who are very eager and highly motivated to find a job. Consider a few examples below:

- Ok, great, I'm excited! How soon can I start if I get hired?
- I have a construction background, and this is exactly what I have been looking for.
- My name is John, and I would love to be a delivery driver for ABC Pizza Company.
- I have already completed the application. I am a very dedicated employee, and I am interested in taking the next steps.

3.2 Positive Messages

These are messages from a candidate who is seen as somewhat interested in finding a job with the company. Consider these examples:

- I want to apply for the HR team member position.
- Can I apply online or turn in my resume in person?

- I am a dot net developer with more than 10 years of experience in web development.

3.3 Neutral Messages

These types of conversations are where candidates are either engaging in small talk with the Recruiters or asking basic exploratory questions about the company, its policies, benefits, etc. These candidates are still trying to figure out if there is a good match and whether they want to apply for job.

- Does your company offer good benefits?
- Who is part of your leadership team?
- Do you hire people with disabilities?

3.4 Negative Messages

These types of conversations are where a candidate is mildly frustrated at the service and/or is experiencing some difficulties in the application process.

- Why haven't you contacted me to follow up?
- Why are there so many steps involved in the application process?
- I need to cancel my interview.

3.5 Extremely Negative

These types of conversations are where a candidate is not interested in finding a job and/or is extremely dissatisfied with the hiring process.

- Sorry, not interested.
- I have decided to work at home. I'm not looking for a job now.
- I already have a job. Thank you for the opportunity.

You could add more granularity to the above sentiment classification scheme based on your needs. Instead of five levels of granularity for sentiment, you could have seven. iCIMS eventually used seven levels of granularity in the final model.

We annotate messages coming from candidates and identify the intent and sentiment behind those messages. Intent is helpful to the model because it can learn that some intents like "amountOfExperience" are associated positive sentiment while other intents like "cancelInterview" and "notInterested" are associated with negative sentiments.

Below are some examples of training data we generate:

Statements	Sentiment	Intent
I'm not having any interviews. This is the end of the conversation.	0 (very negative)	notInterested
I am trying to fill out an application, but it keeps redirecting me	1 (negative)	applicationProblems
I can't come to the interview	1 (negative)	cancelInterview
Can I bring my dog to work?	2 (neutral)	petToWork
Can you provide me with some information about the company?	2 (neutral)	companyInformation
I have 10 years of experience in software development.	3 (positive)	ammountOfExperience
I saw the job description. It is a perfect fit for my profile. How can I apply?	4 (very positive)	applicationProcess

iCIMS' models are trained on the joint features of intent and sentiment classification. We have identified dozens of intent categories and seven levels of sentiment in our training data. This way, sentiment classification, and intent classification tasks are helping each other to reduce the overall loss of the model. We are using transformers to generate embeddings of the incoming text and then use dense and softmax layers in the final layers of it to achieve classification. Although we are making use of transformers, this could also be also achieved with an LSTM or GRU-based RNN model.

Shown below is output of the model summary to give an idea of the joint model inference:

```
Model: "model"
__________________________________________________________________________________________________
Layer (type)                    Output Shape         Param #     Connected to
==================================================================================================
input_ids (InputLayer)          [(None, 55)]         0
__________________________________________________________________________________________________
attention_mask (InputLayer)     [(None, 55)]         0
__________________________________________________________________________________________________
tfxlm_roberta_model_1 (TFXLMRob TFBaseModelOutputWit 278043648   input_ids[0][0]
                                                                 attention_mask[0][0]
__________________________________________________________________________________________________
dropout_74 (Dropout)            (None, 768)          0           tfxlm_roberta_model_1[0][1]
__________________________________________________________________________________________________
dense (Dense)                   (None, 350)          269150      dropout_74[0][0]
__________________________________________________________________________________________________
dropout_75 (Dropout)            (None, 350)          0           dense[0][0]
__________________________________________________________________________________________________
output_1 (Dense)                (None, 73)           25623       dropout_75[0][0]
__________________________________________________________________________________________________
output_2 (Dense)                (None, 7)            2457        dropout_75[0][0]
==================================================================================================
Total params: 278,340,878
Trainable params: 278,340,878
Non-trainable params: 0
```

Figure 1: Output of the model summary

The model is then run on a test dataset to evaluate the results.

```
210/210 [==============================] -
20s 94ms/step - loss: 1.2001 - output_1_loss: 0.7019
- output_2_loss: 0.4982 - output_1_accuracy: 0.8789
- output_2_accuracy: 0.9098
[1.2000659704208374,
 0.7018905282020569,
 0.49817585945129395,
 0.8788713812828064,
 0.909767746925354]
```

As you will observe from the above output, the model has intent accuracy of 88% and sentiment accuracy of almost 91% on test dataset. Thus, we have a model which could predict a sentiment score for most text messages from a candidate.

4 Methodology

There are few other things we need to consider apart from the sentiment score described earlier.

4.1 Response Time

Response time is the median time taken by a candidate to respond to a recruiter's messages. A candidate who is engaged and interested in finding a job will often be very prompt in responding to recruiter messages and have a much shorter response time. On other hand, candidates who are not interested will likely have longer response times.

4.2 Response Rate

Response rate is the ratio of the number of messages which a recruiter sends to the candidate in relation to the number of messages the candidate sends to the recruiter. A high response rate is an indication that the candidate is very engaged and responsive.

Response Rate for Candidate = (total incoming messages to candidates / outgoing messages from the candidate)

4.3 Time Decay

We need to be mindful of the fact that candidates' recruiting needs can change over time. A candidate who was not interested in finding a job in the past may need to relocate ,or their employer may have cut down on their hours, so is, therefore, is back in the job market. Alternatively, there may be a candidate who has already found a job but is now happy with their current employer and is no longer interested in looking for a job. In those cases, we need to apply a time decay to the sentiment score. Time decay provides a way to give more weightage to recent messages from a candidate.

$$\text{Sentiment Score} = \sum_{0}^{n} \frac{\text{Sentiment Score}}{n+1}$$

Where n = number of months elapsed

Let's consider an example below:

Candidate Message to Recruiter (2 months ago)

Sorry, I am not interested in looking for a job. (very negative, sentiment score -2)

Candidate Message to Recruiter (1 month ago)

What kind of benefits do you offer? (Positive, sentiment score 1)

Candidate Message to Recruiter (this month)

I am interested in looking for a job at your company. (very Positive, sentiment score 2)

$$\text{Sentiment Score} = -2 * \frac{1}{2+1} + 1 * \frac{1}{1+1} + 2 * \frac{1}{0+1} = 1.86$$

The final score of a candidate is weighted more towards sentiment of their recent messages than past messages.

Similarly, time decay can be applied to response rate and response time as well. For example, if a candidate became responsive after a period of unresponsiveness, it could indicate that they are back in the job market.

We can calculate response time and response rate on a monthly basis, and then we could apply time decay.

4.4 Engagement Score

Finally, using the sentiment score, response rate, and response time, we can calculate an engagement score for the candidates.

Each of the factors, i.e., sentiment score, response rate, and response time, could be given some weights to calculate engagement score.

$$Engagement\ Score = (w1 \times Sentiment\ Score) + (w2 \times Response\ Rate) + (w3 \times \frac{1}{Response\ Time})$$

We could set these weights, w1, w2, and w3, manually based on what is important for the business goals. For example, we can give more weight to sentiment score and less weight to response rate and response time. We can even use machine learning to automatically compute those weights.

Let's illustrate using an example of Engagement Scores in surfacing candidates. Assume recruiters are looking for candidates with "marketing skills" and there are numerous candidates in their database with those skills. The recruiters need some means of prioritizing the candidates with which they want to first engage. They could use Engagement Score as the criteria to engage with the top 10 or 20 candidates and work down the list in that order. Engagement score bubbles the engaging candidates who are currently looking for a position at the top and those who are not interested to the bottom.

Note that no personal data is used in the calculation of engagement scores or their components. No automated actions are performed with these scores.

5 Limitations

There are some limitations in surfacing candidates to recruiters using the above methodology. This system suffers from a cold start problem which means that if the candidate is new, the system would not have sufficient data to derive any kind of sentiment for that candidate. So even if the candidate might be good, we would not be able to recommend his candidacy to recruiters until there are few messages exchanged.

Also, sometimes weak candidates might show high engagement or good candidates may not come across as very engaging in their initial conversations with recruiters.

6 Conclusion

At iCIMS, our focus is to help our customers find the most qualified talent. Spending time on candidates that are not that interested in the job can be a barrier to that goal because it takes time away from engaging with those that are invested. Plus, it's often time intensive for talent acquisition professionals to surface great-fit candidates without having to sift through hundreds of profiles, which ultimately increases time-to-fill. The iCIMS Talent Cloud provides solutions to both challenges.

To fill a gap in the recruiting software industry, iCIMS clients can use sentiment scores along with response time and response rate to come up with a highly accurate and valuable system for recruiters. Our clients can use engagement score and skills tags to filter quality candidates for their open roles. If a candidate is responsive and asking relevant questions, our system will give a higher engagement score to that candidate. On the other hand, if the candidate is not expressing much interest, stalling, or not truly searching for a job, our system will give a lower score to that candidate. That way, talent acquisition professionals can identify qualified talent more quickly than with the traditional candidate engagement methods.

By surfacing candidates that have relevant skills and are motivated to discuss their jobs, our clients can spend less time on back-and-forth communication and more time on strategic initiatives.

ACKNOWLEDGEMENTS

I would like to acknowledge Talent Cloud AI Team at iCIMS for providing their valuable feedback in writing this paper. I would also like to acknowledge the Marketing, Product, and Legal teams at iCIMS for helping me put all this together.

REFERENCES

[1] Glorot, Xavier and Bordes, Antoine and Bengio, Yoshua. Domain Adaptation for Large-Scale Sentiment Classification: A Deep Learning Approach. (2011).

[2] Wu, Fangzhao and Huang, Yongfeng. Sentiment Domain Adaptation with Multiple Sources. (2016), 301-310.

[3] Bollegala, Danushka and Weir, David and Carroll, John A. Using Multiple Sources to Construct a Sentiment Sensitive Thesaurus for Cross-Domain Sentiment Classification. (2011), 132-141.

[4] Pang, Bo & Lee, Lillian. Opinion Mining and Sentiment Analysis. Foundations and Trends in Information Retrieval, 2 (2008), 1-135.