CS230

Automated medical triage using deep reinforcement learning

Massimiliano Gori*
Graduate School of Business
Stanford University
mgori@stanford.edu

Ciro di Lanno
Graduate School of Business
Stanford University
cirodil@stanford.edu

Abstract

Prior to seeking professional medical care, it is increasingly common for patients to use online resources such as auto- mated symptom checkers or Google. This work proposes an automated triage symptom checker built using deep reinforcement learning that, given a patient gender, age and initial symptom, is able to ask questions about related symptoms and eventually propose a probable disease based on the information collected. Our work is an evolution of traditional symptoms checker applications, like WebMD, and it aims to solve the usability issues that such system present, guiding the user through the process of listing all its symptoms.

1 Introduction

In this paper, we propose neural symptom checking, which learns to inquire and diagnose based on limited patient data. Unlike existing systems which use approximation schemes to select symptoms, we adopt a reinforcement learning framework and formulate inquiry and diagnosis policies as Markov decision processes. The optimization objective directly optimizes a policy function that can be used to select symptoms to inquire patients. At the start, our symptom checker instructs a user to input his age, sex and main symptom (e.g., abdominal pain or headache) so that the model can ask further questions about related symptoms and present a small list of symptoms to choose from. After every question the RL model decides whether to make a disease prediction of suggest additional symptoms to the user. The benefits of this approach are that not only does it improve model accuracy, but it also provides better user experience, guiding users step by step instead of providing an infinite list of choices. Compared to similar commercially available systems, our disease prediction offers a superior user experience, with a marginally lower prediction, that can be improved through further tweaks to the model parameters and training set.

2 Related work

While there has been no relevant academic deeplearning work in the space we have used the findings in the following papers to guide us in the definition of the model and the evaluation of its performance:

- R. Kohavi. Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. In Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96), Portland, Oregon, USA, pages 202–207, 1996.

*Use footnote for providing further information about author (webpage, alternative address)—*not* for acknowledging funding agencies.

CS230: Deep Learning, Winter 2018, Stanford University, CA. (LateX template borrowed from NIPS 2017.)

- I. Kononenko. Inductive and bayesian learning in medical diagnosis. Applied Artificial Intelligence, 7(4):317–337, 1993.
- I. Kononenko. Machine learning for medical diagnosis: history, state of the art and perspective. Artificial Intelligence in Medicine, 23(1):89–109, 2001.
- R. Ledley and L. Lusted. Reasoning foundations of medical diagnosis symbolic logic, probability, and value theory aid our understanding of how physicians reason. Science, 130(3366):9–21, 1959.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. A. Riedmiller. Playing atari with deep reinforcement learning. CoRR, abs/1312.5602, 2013.
- H. L. Semigran, J. A. Linder, C. Gidengil, and A. Mehrotra. Evaluation of symptom checkers for self diagnosis and triage: audit study. BMJ, 351, 2015.
- R. Sutton and A. Barto. Reinforcement learning: An introduction, volume 116. Cambridge Univ Press, 1998.

3 Dataset and Features

Due to privacy laws (e.g., the Health Insurance Portability and Accountability Act; HIPAA) and concerns, real clinical data may not be publicly available, and even anonymized clinical data cannot be shared among researchers. Furthermore, clinical data obtained from hospitals will be biased as it does not represent the real statistical distribution of a symptom across the US population of a certain age and gender. To bridge the gap between limited available data and data-driven methodologies, we propose an approach to generate synthetic clinical data. We first composed the disease set as follows: At the start, we chose SymCat's symptom-disease database as our target since it contains 801 diseases, and each disease is annotated with its symptom distribution across age and gender categories. Then we performed two pre-processing steps to rule out the most extreme/unlikely diseases from the SymCat database. First, we removed the diseases that are not contained in the Centers for Disease Control and Prevention (CDC) database. Second, we observed that SymCat's diseases contain several parent-child relationships. We thus identified all these relationships by querying the UMLS medical database and removed all parent diseases to provide fine-grained disease predictions. For example, skin disorder, atrophic skin condition, and psoriasis are contained in the SymCat database. Since skin disorder is a collective name and more generic than the other two, we removed skin disorder. We then selected the most frequent 220 diseases and 358 associated symptoms in order to have a training set that was large enough to be representative of the complexity of the task, but still perform rapid iterations on the model. Once we created our labelled database we used the probabilities to generate a synthetic list of 1 million patients, with associated age, gender, symptoms and disease, which was organised csv file organized as follows:

Disease	Age group	Gender	Symptom1	Symptom2	...	Symptom10
Common cold	20 – 30	Male	124	102	...	45

The symptoms were translated to a numerical dictionary and the corresponding number were added in the spreadsheet. Given a disease the number of symptoms depended on their probability and their sequence was casual to better simulate a real case. The resulting csv file was used for the training of all the 3 models discussed in this paper.

4 Methods

Given the relatively high size of the output space compared to the number of inputs using an end to end deep reinforcement learning algorithm was not possible. The approach taken was to build disease prediction model and a related symptom prediction model, train them and then connect them to a deep reinforcement learning model that at each state was deciding to redirect the user to either of the 2.

4.1 Disease prediction

This is a multi-label classification model that outputs a disease prediction given information on a person age, gender and symptoms. The model architecture is represented below: (layer size in parenthesis)

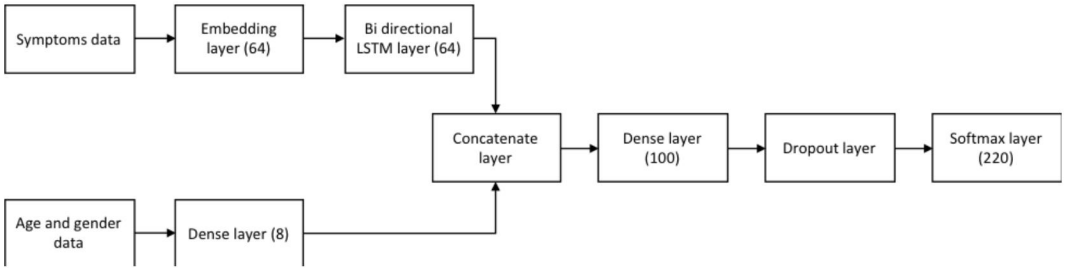

The first step was to split the gender/age data and the symptoms data in 2 different vectors. The age and gender data were one-hot encoded to Due to the high number of symptoms and the relative variance that can occur in the symptoms column we needed to split In order to amplify the signal, the symptoms data is first processed by a 64-dimensional embedding layer with masking enabled to reduce the computable cost associated with sparse input vectors. The resulting output is then processes by a bidirectional LSTM layer, where the output of the forward pass is averaged with the one of the backward pass to eliminate any bias associated to the order in which the symptoms are presented by the user. The resulting output is then merged with the one hot encoded information about sex and gender and then fed to a set of dense layers and eventually the soft ax classifier. Several iterations on the data have shown that the performance of a model with a single intermediate dense layer is comparable to the one with 2 or 3, however the model ability to generalize on new data decreases as the number of dense layers increase. Training iterations have proven that the optimization algorithm yielding the highest accuracy was Adam with Nesterov momentum. This enables a very high learning rate on the LSTM layer, while keeping noise to a minimum. The algorithm reached a maximum training and test accuracy of 77, which is very close to the one of commercially available similar solution. The resulting output of the embedding layer is of particular significance as it correctly shows that related symptoms have the shortest Euclidean distance between them. The network was also able to abstract more nuanced relationships between symptoms, which can only be captured by people with medical training.

4.2 Symptoms prediction

This is a multi-label classification model that outputs the probabilities of related symptoms given information on a person age, gender and symptoms. The model architecture is represented below: (layer size in parenthesis)

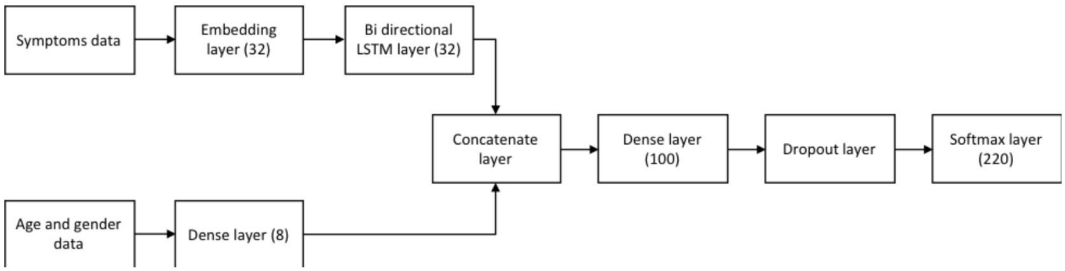

As shown by the above diagram the related symptoms prediction algorithm is topologically similar to the disease prediction, however in this case the algorithm has to predict the 4 symptoms that a user is most likely to have given his, age and a set of symptoms. In order to generate the output, we have followed the scheme described in the below picture: after adding a seize 2 padding to the left of the symptoms and then used a sliding window of size 3 to create the symptoms vector and the adjacent label Y. Another difference from the previous model is that the last layer has a sigmoid activation function as the desired output were the 4 diseases with the highest probabilities of occurrence given a certain input. The model training was done using the same patient list used in the previous case and the final accuracy was 73 percent.

4.3 Inference engine

This is a deep reinforcement learning model that redirects the input from the user to either of the above models given information on a person age, gender and symptoms, current and past states. The model architecture is represented below: (layer size in parenthesis)

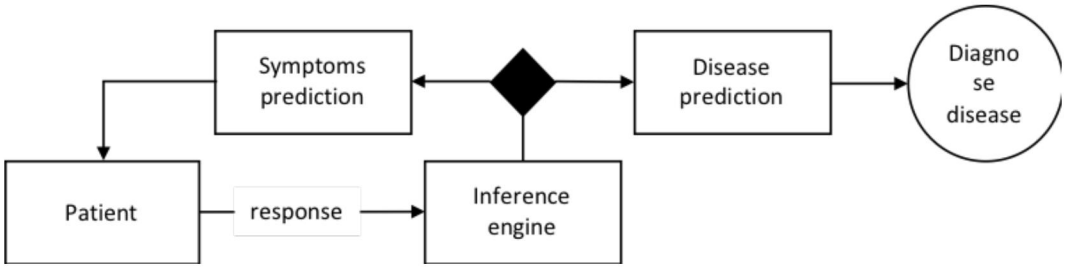

As shown in the above diagram, the inference engine uses a similar network architecture as the one of the disease prediction model as the input structure is exactly the same. This architecture also allowed us to reuse the symptoms embedding matrix we learned in the disease prediction model. The main difference is that the last layer has 2 nodes with a linear activation function that are responsible for passing the input either to the symptom or the disease prediction models. In the first state the user provides information his gender, age and first disease. The model then can either decide to ask for other symptoms or make a prediction about the disease the user has. If the model feeds the information to the symptoms predictor then the user is presented with a list of 4 symptoms he is likely to have and from which he has to select the ones he has and the ones he doesn't. After the selection is complete a new state begins, and the information are fed to the inference model, which can again redirect the input either to the symptom or the disease predictor algorithm. When the input is redirected to the disease prediction model the game ends and the model gets a reward or is penalized depending on whether the disease predicted was the correct one or not. In order to facilitate convergence, we have set the maximum number of steps, meaning that if we get to step 10 the model will make a prediction. More specifically, we use the DQN training algorithm [5] proposed by Mnih et al. The loss function is defined as Lj (Tetaj) = Es,a,r,s [(yj - Q(s, a; Tetaj)2], where target yj =r+gamma max a Q(s,a;Teta-) is evaluated by a separate target network Q(s,a;Teta-)with parameters Teta. The variable j is the index of training iteration. To improve training stability and convergence, the target network is fixed for a number of training iterations. The parameters Teta can be updated by the standard backward propagation algorithm. In order to train the model we have built a custom program that was automatically sampling a patient form the list created in the initial step, selecting a random symptom to supply to the model and then was supplying binary answer to the related symptoms inferences oprovided by the symptoms checker model. We have observed that the overall accuracy of the system is 70. As expected the model performs very well in the case of common diseases, but its performance declines when it is faced with extremely infrequent diseases. It has to be noted that the overall accuracy is very high in relation to the accuracy of the individual downstream symptoms and disease predictor models, as they represent the model performance cap.

5 Limitations

Although the trained model can handle some common scenarios of human practice, there is an essential limitation of the trained model: it only considers symptoms. Some similar diseases require differential diagnosis taking the results of physical or laboratory examinations into consideration to accurately distinguish them. The use of attention, coupled with a starched LSTM layer could be used in order to improve the accuracy of both models. Furthermore, the current model does not incorporate negations, meaning that

6 Conclusion/Future Work

We have shown that the proposed neural symptom checker can imitate the behaviour of inquiry and diagnosis process performed by doctors. One direction of our future work is to develop methods that can recommend more complex diagnosis tasks including physical and laboratory examinations.

Another direction is to cover more diseases (addressing the issue of scalability) without degrading the diagnosis accuracy.

7 Contribution

Ciro worked on the generation of the medical database, the training sample and the deep reinforcement learning model. Massi worked on the symptoms and disease prediction models, as well as the deep reinforcement learning agent.

8 References

[1] Alexander, J.A. & Mozer, M.C. (1995) Template-based algorithms for connectionist rule extraction. In G. Tesauro, D.S. Touretzky and T.K. Leen (eds.), *Advances in Neural Information Processing Systems 7*, pp. 609–616. Cambridge, MA: MIT Press.

[2] Bower, J.M. & Beeman, D. (1995) *The Book of GENESIS: Exploring Realistic Neural Models with the GEneral NEural SImulation System.* New York: TELOS/Springer–Verlag.

[3] Hasselmo, M.E., Schnell, E. & Barkai, E. (1995) Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. *Journal of Neuroscience* **15**(7):5249-5262.