

Why natural language processing is AI's jewel in the crown

Liu Qun, Chief Scientist of Speech and Language Computing at Huawei's Noah's Ark Lab, believes that as humanity enters the intelligent era, smart devices and data volume will rise at unprecedented speeds, as will the demand for human-computer interaction through speech and language. Significant cutting-edge research and technological innovations will emerge from the fields of speech and natural language processing.

By Linda Xu, Xue Hua

Natural language processing (NLP) is a key area of research in the fields of artificial intelligence (AI) and computer science that aims to study the theories and methodologies for effective communications between humans and machines using natural language. According to Gartner's 2018 World AI Industry Development Blue Book, the global NLP market will be worth US\$16 billion by 2021.

Three types of intelligence

WinWin: From the scientific perspective, what is the significance of NLP?

Liu Qun: Linguistics is an age-old subject. Why do we have language? Why haven't animals developed language as complex and advanced as humans have? Does the human brain have an innate language mechanism or is language acquired like other abilities? How is language formed and developed? What rules does it follow? These are some of the mysteries that scientists have yet to answer.

Computational linguistics, or NLP, is a science as





If we want to stay ahead we need to adopt an AI mindset – use AI concepts and technologies to tackle both existing and future problems.



well as an application technology. From a scientific perspective, like other computer sciences, it's a discipline that involves the study of language from a simulated perspective. NLP isn't directly concerned with the study of the mechanisms of human language; instead, it's the attempt to make machines simulate human language abilities. For a computer to have human-like language ability would indicate, to some extent, that we have an understanding of human language mechanisms. Since understanding natural language requires extensive knowledge of the external world and the ability to apply and manipulate this knowledge, NLP is an AI-complete issue and is considered one of the core issues of AI.

WinWin: Some people think that NLP is the key to achieving inclusive artificial intelligence. What's your view?

Liu: There's a certain amount of truth to that. Some people divide human intelligence into three main categories: perceptual intelligence, motor intelligence, and cognitive intelligence.

First, perceptual intelligence includes hearing, vision,

touch, and so on. In the last two years, the use of deep learning has significantly improved speech and image recognition rates. Computers have therefore done quite well at the perceptual intelligence level, in some classic tests reaching or exceeding the average level of human beings.

Second, motor intelligence refers to the ability to move about freely in complex environments. Motor intelligence is one of the core research areas of robotics.

Third, cognitive intelligence is the most advanced of intelligent activities. Animals have perceptual and motor intelligence, but their cognitive intelligence is far inferior to ours. Cognitive intelligence involves the ability to understand and use language; master and apply knowledge; and infer, plan, and make decisions based on language and knowledge. The basic and important aspect of cognitive intelligence is language intelligence – and NLP is the study of that.

The object of NLP study is human language, including words, phrases, sentences, and chapters. By analyzing these language units, we hope to understand not just



We've achieved a great deal of success with AI and machine learning technologies in the area of image recognition, but NLP is still in its infancy.



the literal meaning expressed by the language, but also the emotions expressed by the speaker and the intentions conveyed by the speaker through language. True cognitive intelligence isn't possible without successful NLP.

Natural language understanding and processing are also the most difficult for AI. If, for example, you alter a few pixels or a part of an image, it doesn't have much effect on the content of the image as a whole. But with text, it's different. Changing one word in a sentence in many cases would completely change the meaning.

We've achieved a great deal of success with AI and machine learning technologies in the area of image recognition, but NLP is still in its infancy. However, with style generation applied to an image we can easily replicate the style of Van Gogh, but we still don't have the technological capability to accurately replicate a passage of text into the style of Shakespeare.

This is why many experts call NLP 'the jewel in the crown of AI'. It makes a lot of sense and I completely agree with this description.

Latest developments and

challenges in NLP

WinWin: What are the biggest problems facing NLP?

Liu: I think there will be two main problems facing NLP in the future.

The first is semantic understanding, that is to say the problem of learning knowledge or common sense. This problem is about how NLP technology can get "deeper". Although humans don't have any problem understanding common sense, it's very difficult to teach this to machines. For example, you can tell a mobile assistant to "find nearby restaurants" and your phone will display the location of nearby restaurants on a map. But if you say "I'm hungry", the mobile assistant won't give you any results because it lacks the logical connection that if you're hungry, you need to eat, unless the phone designer programs this into the system. But a lot of this kind of common sense is buried in the depths of our consciousness, and it's practically impossible for AI system designers to summarize all of this common sense and program it into a system.

The second problem is a lack of resources. Technologies such as unsupervised learning, zero-shot learning, few-shot learning, meta-learning, and

migration learning are all essentially attempts to solve the low-resource problem. NLP is unable to effectively deal with the lack of labelled data that may exist in the machine translation of minority languages, dialogue systems for specific domains, customer service systems, Q&A systems, and so on. These kinds of problems are collectively referred to as low-resource NLP. To address these problems and enhance data capabilities, in addition to trying to introduce domain knowledge such as dictionaries and rules, we can also use active learning methodologies to add manually labelled data and unsupervised and semi-supervised methodologies to utilize unlabeled data, or multitask learning to utilize data from other tasks or even from other languages. And you can also apply migration learning to leverage other models. This is how NLP technology can become broader.

WinWin: What research has had the most profound impact in NLP over the last decade?

Liu: Deep learning, without a doubt. Deep learning based on deep neural networks has altered NLP technology on a basic level, moving us from discrete symbols to defining and solving NLP problems and using continuous numerical fields. This has led to a complete change in the way the whole problem is defined and the mathematical tools we use, and has driven huge developments in NLP research.

Before the application of deep learning techniques in NLP, the mathematical tools used were completely different to the ones adopted for speech, image, and video processing, creating a huge barrier to the flow of information between these different modes. But using deep learning in NLP means that the same mathematical tools are used. This has removed the barrier between different modes of information, making multi-modal information processing and fusion possible.

The application of deep learning has led NLP to an unprecedented level and greatly expanded the scope of NLP applications. The spring of NLP has come, as it were.

WinWin: What unique demands does NLP present for frameworks and hardware?

Liu: AI research has endless demands on hardware, but inadequate hardware can limit models – scientists require better hardware to experiment with more complex models and develop newer and better methodologies.

I don't think NLP has unique demands on frameworks or hardware, and they're similar to those in other areas of AI research. You always need more memory, higher bandwidth, more parallel computing power, and higher speeds. So, optimizations for specific NLP scenarios are no particular barrier.

WinWin: What new methodologies or trends impacted NLP applications in 2018?

Liu: In 2018, the most striking achievement in NLP research was pre-training language models, including RNN-based ELMo, Transformer-based GPT, and BERT. The success of pre-training language models proves that we can learn much from massive volumes of unlabeled text without having to label large amounts of data for every NLP task.

In terms of applications, Google's Duplex was something we'd never seen before. Several Chinese companies have also developed very impressive simultaneous interpretation technology. Although it still makes many mistakes in simultaneous interpretation and is still a long way off being as good as simultaneous interpretation by humans, it's undoubtedly very useful. It was hard to imagine this technology actually getting used a few years ago, so



Over the past five years, we've submitted one of top 50 papers cited by NIPS. We have also submitted one paper in the top 20 and three in the top 30 papers cited by ACL.



it's completely unexpected to have reached a level of preliminary practical application in such a short time.

Huawei's work to advance NLP

WinWin: What is the focus of NLP research in Huawei Noah's Ark Lab? And what progress has been made?

Liu: Huawei Noah's Ark Lab is carrying out three main areas of NLP research: voice technology, machine translation, and dialogue systems.

We've already started to apply Noah's Ark's NLP in a wide range of Huawei products and services. For example, Huawei's mobile phone voice assistant integrates Noah's Ark's voice recognition and dialogue technology. Noah's Ark's machine translation technology supports the translation of massive technical documents within Huawei. Noah's Ark's Q&A technology based on knowledge graphs enables Huawei's Global Technical Support (GTS) to quickly and accurately answer complex technical questions.

Noah's Ark Lab has also achieved outstanding achievements in NLP research. Our research results in natural language text matching, dialogue generation,

and neural network machine translation have been widely cited by researchers. Over the past five years, we've submitted one of top 50 papers cited by NIPS. We have also submitted one paper in the top 20 and three in the top 30 papers cited by ACL.

WinWin: Is multimodality that combines hearing and vision a promising future direction of research in NLP? Is Huawei working on that?

Liu: Yes, and we've started research in this area.

NLP use cases and industry adoption

WinWin: How can NLP be applied in vertical industries such as finance, law, and healthcare?

Liu: NLP has been an integral part of our day-to-day lives for a long time now. Many people don't realize that we enjoy the convenience of NLP technology every day, including the Pinyin input method in China. Two decades ago, the Wubi input method was popular, but has since been almost completely replaced by the Pinyin input method.

Pinyin input methods did actually exist when Wubi was



When you call customer services, you won't have to choose from a whole bunch of voice menu options. Voice assistants will understand your requirements and help you with a range of day-to-day tasks.



popular, but at the time had very limited intelligence. Users had to select the correct Chinese characters from a large number of homophones. Word prediction functionality was also very weak and input was very slow. It was only after we made a lot of progress in NLP technology and adopted statistical language models to enable the automatic selection of the most likely sequence of Chinese characters from long strings of pinyin, that pinyin input methods were able to overtake Wubi as the main Chinese character input method.

Today's search engines rely heavily on NLP technology. If you search for "the population of Sichuan", for example, search engines will give you a specific answer by using natural language Q&A technology, as well as listing a series of related web pages.

In fields like finance, law, and healthcare, NLP technology is also gaining traction. In finance, NLP can provide analytical data for investing in stocks, such as identifying trends, analyzing public opinion, analyzing financial risks, and identifying fraud. In law, NLP can help with case searches, judgment predictions, the automatic generation of legal documents, the translation of legal text, intelligent Q&A, and more. And

in healthcare, NLP has a broad avenue of application, for example, assisting medical record entry, retrieving and analyzing medical materials, and assisting medical diagnoses. There are massive modern medical materials and new medical methods and approaches are developing rapidly. No single doctor or expert can be expert at all the latest medical developments. NLP can help doctors quickly and accurately find the latest research results for various difficult diseases, so that patients can benefit from advancements in medical technology more quickly.

WinWin: How will advances in NLP benefit people?

Liu: It will make life more convenient for everyone. For example, when you call customer services, you won't have to choose from a whole bunch of voice menu options. Voice assistants will understand your requirements and help you with a range of day-to-day tasks. Machines will be able to help you write reports and even poems. Advances in technology will also have a disruptive effect. In terms of employment, for example, machines will replace manual labor and thus lead to job losses. However, while the application of new technologies will cause some jobs to disappear, it will also create a large number of new employment opportunities. 