How Natural Language Generation (NLG, branch of AI) is helping clinicians improve patient outcomes?

Authors:
Mr. Atri Vivek Sharma, Technical Project Manager, Pangaea Data
Mr. Jingqing Zhang, Head of AI, Pangaea Data
Dr. Julia Ive, Head of NLP Research, Pangaea Data
Dr. Vibhor Gupta, Founder, Pangaea Data

Imagine a clinician is examining a chronically ill patient, with a long history of organ failures and medical procedures, and drug allergies. This historical information is critical for making a diagnosis; however, it is vast and is contained inside a large electronic health record (EHR) with lab test results, doctors' notes, past procedures and imaging reports. Ordinarily, a clinician would have to painstakingly go through each of those documents manually to understand the patient's medical history.

Now imagine that using artificial intelligence (AI), the vast amount of information contained in the electronic health record can be automatically summarised in an instant, such that the clinician will be presented with all the relevant information in a concise and comprehensible format. The time that the clinician would have previously spent on thoroughly scanning patient information can now be spent on what truly matters – care for the patient.

Artificial Intelligence that enables electronic health summarisation is known as Natural Language Generation, or NLG. NLG comes under the umbrella of NLP (Natural Language Processing), which essentially encompasses all techniques that allow computers to understand as well as generate human language.

At a very high level, NLP creates "vectors" that represent words and sentences – essentially converting the written words into a series of numbers that represent the "essence" of the sentence. Natural Language Generation is when these vectors are converted back to words arranged in a meaningful and grammatically correct sentence.

NLG technology surrounds us already. From the word suggestions that pop up in your smartphone keyboard, to sentence completion suggestions while composing an email, to the chatbots that you can interact with in customer service, are all prime examples of where NLG has had a massive impact on consumers and organisations. We at Pangaea are working towards employing these techniques in a medical setting, such that they can assist clinicians and scientists in improving patient outcomes, building on the pioneering work of our lead researchers who developed a state-of-the-art NLG framework, PEGASUS.[1]

Medical Summarisation

[1] https://arxiv.org/abs/1912.08777

As alluded to previously, the increased adoption of EHRs has led to an overwhelming amount of patient data which has potentially negative consequences on clinical work – such as errors of omission and increased waiting time for patients. Chronically ill patients, the ones who require the maximum amount of care and attention, often have the largest datasets which makes their health records difficult and time consuming to understand. As an example, in one institution, EHRs of patients with chronic kidney disease had on average 338 notes captured over 14 years, with several patients' records containing over 4000 notes [2]. NLG techniques can be used to generate a comprehensive summary of such long EHRs based on information from discreet values (structured data) such as lab results; and textual (unstructured) data such as clinical notes and pathology or radiology reports, which makes it easier and effective for clinicians to understand the patient journey.

Our lead researchers at Pangaea have published a high impact paper[3] where we demonstrated the application of NLG to successfully generate synthetic EHRs with high accuracy after reading a patient's original health record. This is a key development in medical summarisation, as we have shown that EHRs can be read and summaries from them can be created with high accuracy in a privacy compliant manner.

Generation of Clinical Study Reports

Natural Language Generation can also be extended to the drug development domain. As anybody who has ever been involved in clinical studies can attest, report writing is an extremely tedious task, where strict guidelines and writing styles must be followed. Clinical study reports routinely take experts up to 10-25 days [4] to draft, and even longer to finalise. This repetitive task of converting tabular clinical data into a textual reports is an ideal use case for NLG as it would accelerate report writing, allowing medical experts to better utilise their precious time. Developments in NLG models are enabling them to automatically generate text with ever increasing consistency [5], which would ensure that the generated reports would meet the strict guidelines and allows scientists to focus their precious time on inference and analysis of the contents of the reports as opposed to manually write the textual narratives for the report.

At Pangaea, we have successfully developed PIES (Pangaea's Intelligence Extraction and Summarisation), a neural architecture that can automatically generate regulatory grade reports for scientists and clinicians in the Biopharmaceutical industry. PIES is a generalisable architecture that can use various types of input data to automatically generate

[2] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4986665/
[3] https://www.nature.com/articles/s41746-020-0267-x
[4] https://www.samhamiltonmwservices.co.uk/pdf/hamiltontws-2008_173.pdf
[5] https://www.nature.com/articles/s41746-020-0267-x

assay validation, PK, Toxicology, clinical study, adverse event reports and alike, which are critical across pre-clinical, clinical and pharmacovigilance domains of the industry.

Dialogue Generation for Chatbots

Although chatbots have been around for quite some time, the constant evolution of NLP and NLG technology continues to make their interaction a lot more natural and human-like. They are very widely used in several areas, most notably in customer service. In this use-case, NLG is used to generate text responses to the messages sent by an individual, adjusting the content and tone based on the context of the initial message.

In the medical domain, chatbots are currently being used for providing diagnostic information, health tips, and information about local healthcare providers. They are also being utilised for combating mental health issues – a prime example is Woebot, designed by Stanford University. This chatbot provides mental health assistance using cognitive behavioural therapy techniques [6] . Using novel NLP and NLG techniques, researchers have developed a therapeutic relational agent that can form a bond with its users and deliver human-like therapeutic encounters. Their clinical trials have demonstrated a 30% decrease in substance abuse, and a 23% decrease in symptoms of depression and anxiety amongst its users. Such chatbots could be truly revolutionary in areas with limited medical staff – as they would allow for much greater access to medical information and treatment of mild symptoms.

Ethical Considerations and Validation by Human Experts

The increased adoption of artificial intelligence technology in our daily lives continues to raise concerns regarding its ethical considerations, and rightfully so. However, it is important to bear in mind that AI in the medical domain cannot and will not replace human clinicians and experts anytime in the future. It simply acts as an assistant that allow experts to complete tasks faster and focus on patient care or drug development. Assisting human experts in their work has been, and always will be the primary focus for the team at Pangaea.

Furthermore, at Pangaea we place a heavy emphasis on including the “Human (Expert) in the loop” when building our AI models.[7] This means that we incorporate human inputs for both training and inference purposes, which boosts the confidence of the outputs and improves the performance of our AI. The model we have developed, PIES, has

[6] https://woebothealth.com

[7] https://www.forbes.com/sites/edstacey/2021/04/09/what-humans-can-learn-from-human-in-the-loop-learning/

implemented a unique human-in-the-loop protocol that drastically reduces (by 60%) the human effort required to improve the performance of such models.

Challenges of NLG development

There are several challenges facing the development of NLG techniques specifically in the medical domain. The biggest, by far, is the sparsity of data, which refers to the AI model having too little data to learn from.

In simple terms, any AI or machine learning model requires a large amount of real-world data, which it uses for training. For example, in the context of NLG, the model requires a large amount of grammatically correct sentences, which would collectively contain all the possible words, phrases and terminology. However, in the medical domain, there are several diseases and medical terminology that are very rare, and thus have few sentences associated with them. Thus, the NLG model would simply not know how to generate sentences that contain references to those rarely occurring diseases or terminologies. To deal with this data sparsity issue, our PIES architecture uses a unique data augmentation procedure to generate additional training data artificially. This allows PIES to generate textual narratives for regulatory reports after being trained on as few as 200 text examples.

The Next Steps for NLG

Thus, we have seen how Natural Language Generation technology has the potential to improve patient outcomes. By offloading the repetitive and time-consuming tasks of clinicians and medical experts, NLG techniques allow experts to spend more time on more critical tasks such as patient care or drug development. Natural Language Generation has seen major advancements in recent times, which can immensely benefit the medical domain.

Natural Language Generation has been successfully implemented in a variety of other industries as well, to automatically generate big-data driven financial reports in the finance sector[8], produce automated weather/news reports in journalism[9], deliver chatbots for customer service, generate product descriptions for e-commerce[10], and many, many more.

We have barely scratched the surface of the potential benefits that NLG advancements can offer to the life sciences community. We continue to learn of other use cases and areas where such novel techniques can be applied to improve patient outcomes, and therefore would love to hear your thoughts and engage in a discussion! We are reachable on vgupta@pangaeadata.ai and through www.pangaeadata.ai

[8] https://link.springer.com/chapter/10.1007%2F978-3-030-67664-3_40
[9] https://www.aclweb.org/anthology/D18-1206.pdf
[10] https://dl.acm.org/doi/pdf/10.1145/3308558.3313407