

IT'S ALL ABOUT THE GROUND TRUTH

THE ART AND SCIENCE OF DIAGNOSTICS IS PERHAPS THE MOST DIFFICULT CHALLENGE FOR AI AS EACH CLINICIAN'S DECISION IS OFTEN INFORMED BY A LIFETIME OF EXPERIENCE AND STUDY

As the AI revolution continues to grow, companies increasingly recognize that the magic wand of automation needs the hand of human intelligence to wave it. In industry parlance, this human intelligence is conveyed to AI models in the form of what is termed "ground truth" training data. The algorithm learns iteratively by looking a large number of data points which have been annotated (key portions of the data segmented) and labeled (classified) with the desired answer. Over a period of time the algorithm becomes better at matching the answer and eventually can generalise its operation to previously unseen data. ▶

“ A KPMG report released in 2018 showed that investment in AI for healthcare had reached almost \$1.3 billion across 107 deals in 2017.

In practical terms, an AI algorithm can only succeed if it has access to well structured training data which has been labeled at high accuracy by a human workforce. The old adage of “garbage in garbage out” applies even more in the world of machine learning. Herein lie both the challenges and opportunities of building large scale accurate datasets using human judgements and decisions. In the booming healthcare space, this problem is further intensified by the needs of deep domain expertise, regulatory barriers, and complexities surrounding data management.

A KPMG report released in 2018 showed that investment in AI for healthcare had reached almost \$1.3 billion across 107 deals in 2017. The report also noted that access to the right data at scale is a critical enabler of artificial intelligence in this space. While algorithms like TensorFlow and large scale computation resources like the cloud, have become widely accessible to big and small players, the data factor remains a crucial differentiator in the success of these investments. A year later, investment and interest in the medical AI space has grown even further as AI-powered technology begins to emerge in the marketplace to tackle everything from patient care to back-end operations.

The art and science of diagnostics is perhaps the most difficult challenge for AI as each clinician's decision is often informed by a lifetime of experience and study. It is also the part of the medical pipeline that offers the greatest opportunity for machine learning to produce a significant impact. An accurate image analysis algorithm, for instance, can free up the bottleneck caused by the low specialist to patient ratio in most geographies around the world. But for AI to be successful, three facets of the data challenge - scale, accuracy and cost - have to be resolved to train models to reach their full potential.

The sheer volume of medical data presents a challenge to scaling. A single MRI, CT, PET, or ultrasound scan generates thousands of images, and up to 80% of a radiologist's time is spent going through each image one at a time and organizing the findings. However if experts like radiologists, cardiologists, or pathologists have to work on segmentation and labeling of datasets with hundreds of thousands of images for the purpose of teaching algorithms, the same bottleneck becomes further aggravated. Using medical experts to label the data also quickly escalates the expense of the operation and the cost of building a viable product.

In recent years, micro-task based crowdsourcing has been used as a large scale, low-cost solution to this problem. As the online gig economy has picked up in popularity, a digital-savvy workforce that can be assigned repetitive labeling and annotation tasks has been created. Complex tasks are disassembled and assigned to online labelers, and then reassembled to create the completed package. While this addresses the scaling issue and provides access to a 24/7 workforce, accuracy and consistency can be stumbling blocks. Consistent quality control in particular is crucial as inconsistent interpretation of the medical data may have very real patient outcome consequences.

In leading the Medical Division at iMerit, this is where I've witnessed robust training and knowledge sharing models make a significant impact. Through specialized training, the expertise gained by medical professionals through years of practise can be transferred to a dynamic workforce with the right skillset, wherein pattern recognition and memorization abilities are leveraged to handle medical data at scale. In essence, this creates a new field of experts in specialized medical data labeling that is defined by narrow scope, specific context, and deep subject understanding. Such a model equips data labeling experts to contribute to massive medical data labeling projects. This provides a sustainable skilling model that can handle data at scale, helping game-changing medical applications go from the R&D stage to commercial use as well as providing life-changing opportunities for a work-force.



To tackle the skilling challenge, our first step is to develop a curriculum based on understanding medical lexicon, pathology, spatial orientation, and data manipulation. Just as in medical school, we also get a sense of who is curious about learning about the body and who is turned off from exposure to graphic medical images. Project-specific training includes live-demos, videos, models, and instruction guides.

There are several advantages to partnering with a dedicated inhouse workforce. The secure environment helps customers manage their data pipeline with confidence. Once the project is underway, a continuous and iterative review process takes place, particularly while dealing with tricky edge cases. Large volume data

THE TECH PERSPECTIVE

exposure combined with expert supervision then provides insights to long-tail problems and obstacles. Lastly, a multipass workflow ensures the degree of precision and accuracy necessary for medical applications. Labelers gain further expertise by expanding their understanding to new modalities and pathologies over time.

Equipped with specialized knowledge and a trained eye, a medical data annotation expert team is able to focus its attention entirely on identifying abnormalities in images, slides, videos, and other datasets. A day's work is spent dealing with thousands of anatomical images with extreme concentration, but in talking to the team, I've found this motivates the team as they are compelled by the valuable use-case and the endless variations in the data.

Data labeling experts like Namita Pradhan and Chinmayee Swain, who were initially recruited for their eye for detail and accuracy, have now become deeply interested in the subject matter. Namita talks about reading medical material on her own time to prepare for potential use cases where her data labeling skills might be leveraged.

Another contributor, Chinmayee talks about the pride in her work and connects her work as some day helping her own community meet its healthcare needs. Being fairly compensated for their skilled work and thriving in a collaborative work environment also increases the investment of each employee while undertaking projects of this complexity.

The human intelligence that we need for medical AI to work is distributed equally across the world but as we all know, opportunity is not. An enterprise skilling model bridges the gaps, both with expertise and access, and enables a global workforce to participate in the development of cutting edge technologies. Specialised data workers can ease the burden of doctors just like paramedics and nurses do in their areas of expertise.

BIOGRAPHY



DR. SINA BARI is a Stanford-trained surgeon and Medical Director for iMerit where we help companies create high quality ground truth data to train AI models.