

September 18, 2018

Deanna Zhang | Associate | dezhang@tphco.com | 713.333.5424

TPH Energy Tech Research Company Spotlight

The latest and greatest developments in e-tech

In this series, we will keep a running tab of e-tech-related companies and their newest developments - from milestones reached to technology deployed to sales leads acquired, and everything in-between.

ThoughtTrace

Documentation Nation

Natural language processing (NLP) has long been a point of fascination for machine learning researchers and AI enthusiasts. Having the ability to teach a machine how to understand human language has enormous implications for what types of data can be processed. Think about how much of your time at your job is spent reading and understanding at work - and now think about how that would change if a machine could help you do it. One of the bigger obstacles to this is that language can vary tremendously. Many of the problems with modern NLP revolve around how to deal with aberrations from a standard set of rules - verbal phenomena like slang, idioms, ironic expressions, and grammatical errors add many, many more dimensions to the already complex problem. Because of this, the NLP engines that back Siri, Alexa, and Google Assistant are still far from being able to handle the majority of potential inquiries (beyond the “I don’t understand the question, try searching on the web” cop-out). Fortunately, there is still a subset of modern language that is governed by strict rules with limited deviation, and you can find its presence in every “official” transaction between two or more parties. Contract language is formal, precise, and consequently very good fodder for machine learning. Houston-based ThoughtTrace has recognized this potential and tackled its application in oil and gas with extraordinary success.

ThoughtTrace, previously called Agile Upstream before a rebranding earlier this month, is focused on streamlining the process of organizing, reviewing, and analyzing oil and gas-related contracts. Their current flagship product, ALI, can take hundreds of leases uploaded by the user, OCR those leases for analysis, extract information, organize and cluster (*Cont'd on next page*)

Company Spotlight: ThoughtTrace (Cont'd)

that information, and present certain metrics for analysis. Anyone who has dealt with leases en masse knows that extricating even basic data like parties involved, date, term, and tracts involved can be a time-consuming, tedious task if done by hand. ThoughtTrace's oil and gas lease tool can automatically extract 240+ different lease provisions and data elements, which, in addition to reducing or even eliminating the need for hiring a third party to process the files manually, allows the associated A&D and land teams to dig in almost immediately on their lease package, decreasing the turnaround time for transactional due diligence and day-to-day review.

It's not only leases that ThoughtTrace is equipped to handle. The company recently released a tool for analyzing Surface Use and Right of Way agreements. This tool can extract provisions such as rights associated with laying multiple pipelines, product specifications, lease termination, etc. The same can be done for Master Service Agreements, another tool under development. The MSA tool will target provisions like indemnification, insurance requirements, definitions of contractual parties/groups and terms, and payment terms and obligations. The combined analytical package of leases, RoW/Surface agreements, and MSAs covers nearly all of the major contractual agreements in oil and gas. Mass-deployment of such a system could have far-reaching effects. If the A&D due diligence process can be shortened and made easier, would that increase the number of average participants in a process and make packages more competitive? Would a shorter evaluation process increase the frequency with which packages are marketed? With greater insight into contract provisions, would disputes between service providers and their customers be resolved more quickly? With instant reference back to the whole portfolio of contracts and their provisions, would that standardize negotiation terms or even level pricing across contracts?

As an industry with an abundance of underutilized structured data, there is often little screentime given to relatively unstructured data like text in contracts. But there is clearly massive potential for machine learning to create value in this area - and we are only at the beginning of such a revolution. In the last 6 months, ThoughtTrace has grown the volume of documents it processes from 190k to >1.8mm. This number will grow exponentially as new tools are added and more customers are signed up. We look forward to seeing what happens then.