

SHELL, KNOW THYSELF!

Back in simpler times, you used data that you were generally familiar with to perform calculations and perhaps as input for computer models to help you understand more-complex systems. Things have moved on. Given sufficient computer power, it is becoming possible to uncover relationships within large, diverse sets of data that were previously hidden. This is the result of advances in a variety of related areas, notably data analytics, artificial intelligence, machine learning and computational hardware. There are some, however, who question whether companies are really achieving maximum value from all this. That is certainly the feeling within Maana, a San Francisco-based company launched in 2012. With its unique emphasis on capturing the knowledge possessed by the human experts within most organisations, as well as in the increasing volumes of data now routinely available, Maana claims to be capable of generating levels of business value in excess of normal expectations.

Shell was quick to appreciate the direction Maana was taking and its fresh approach to data analytics. After following Maana for more than a year, Shell Technology Ventures (STV) decided to invest in the company in 2016. Carl Stjernfeldt, the STV lead for the Shell–Maana relationship, says, “The Maana technology stands out for being able

to extract useful information from very diverse data sources, notably including unstructured ones. It is then uniquely able to integrate, correlate and find patterns within the different kinds of information and to yield powerful insights that lead to higher-quality business decisions.”

Okay, so why invest in Maana and not just buy a set of Maana tools? The answer is simple according to Carl: “The technology can be applied to a range of challenges in various industries and, frankly, we wanted to encourage Maana to focus first on the oil, gas and energy sectors, and then to pay attention to some areas of particular interest to Shell. We have already identified more than 100 possible applications for the technology within the company.”

Since the STV investment, Maana has worked remarkably closely with Shell (see the Digital@ Work feature, page 30) and several successful pilot trials have been carried out using the Maana Knowledge Platform® all under the stewardship of the Technical and Competitive IT (TaCIT) division of Shell Projects & Technology.

To find out more about the technology’s status within Shell, *TechXplorer* interviewed Maana’s chief marketing officer, Azita Martin, and its customer solutions director, Michael Verkruse, together with TaCIT’s Freddie Darbyshire, who has supported the implementation of the Maana technology within Shell.

The search for true meaning

Maana is the word for “meaning” in several languages, including that of the company’s Turkish co-founder and chief executive officer Babur Ozden. The name reflects the fact that the company offers what it calls digital knowledge technology to extract meaningful insights from all kinds of data to optimise complex industrial operations.

Maana technology is poised to be used to deal with corrosion risks in the crude refining business.



In a technical sense, what is the most outstanding aspect of the Maana Knowledge Platform?

Azita: Our special sauce is the patented Maana Knowledge Graph®, which sits at the heart of the Maana Knowledge Platform (see boxed text, “Making connections”). Other companies offer knowledge graphs but they are typically designed with the consumer world and the general public in mind and they tend to be data-centric. Ours is specifically for industrial applications and is notably capable of capturing and containing not just data but also people’s expertise and their understanding of a particular business process or decision flow. The Maana Knowledge Graph is a way of digitalising human expertise, along with the data that are most relevant to the process that a company is seeking to optimise.

Freddie: Knowledge graphs are of great interest to Shell because they offer a way to represent, store and join data, and then to model them easily and quickly. Within Shell, we have masses of data presented in different formats and, perhaps more importantly, recorded within different business contexts. The Maana Knowledge Graph gets over the problem of comparing these data by separating the content of the data from their structure. You could say that it takes the data out of context, but in a smart way, which enables data from different sources to be integrated effectively and relationships to be readily explored.

Once a company has bought into the Maana technology, how does it get started?

Michael: Well, it starts with the people who know most about the chosen business area and who understand the challenges best. Maana works with these people, the subject matter experts, to create a picture of what a business wants to achieve from the decision-making perspective and then to break it down into discrete problem-questions, as we call them. These form the basis of the series of related business models, the Knowledge Models, that go to make up the Maana Knowledge Graph.

It is a crucial first step, as before you can start to come up with answers, you must ensure that everyone – business leaders, subject-matter experts, data scientists and business analysts – shares the same understanding of what is involved and what the goal is in terms of optimising a given operation or business process. It is only then that we can start to think about the data requirements and the algorithms that will get you to the answers you want.

Freddie: I think it is also important to stress that the Maana Knowledge Platform is new technology, and we are still at the point where Maana Knowledge Graphs can only be constructed and used effectively by Shell IT data scientists with good experience of coding. Having said this, I am hopeful that, as the technology develops further, the situation will change and business analysts, for example, will be able to use the technology to develop their own solutions, thereby democratising the power of the platform and increasing the value it can deliver to Shell.

Azita: In this context, it is worth mentioning our Knowledge Assistants. These are Maana’s artificial intelligence algorithms designed to help create Knowledge Models quickly. DocAssist is worth a particular mention, as it extracts unstructured information from PDF documents, for example, technical papers, reports or even emails. We are continuing to add to the list of available Knowledge Assistants, as we have found that they help customers to create models 3–10 times faster than when using any other technology. This is based on customer trials involving Maana and other vendors. Actually, it is customers, including Shell, who have been first to go public on this performance benefit.

Freddie: I can vouch for the Knowledge Assistants simply on the basis of the amount of time it can take Shell IT specialists to build the code necessary to carry out what these off-the-shelf tools can do. ▶▶▶

Maana Knowledge Platform and Maana Knowledge Graph are trademarks of Maana Inc.

INTERVIEW

Azita Martin, Michael Verkruyse & Freddie Darbyshire



So, you have captured the knowledge related to a particular asset or process in a Maana Knowledge Graph, but how do you ask questions and get the answers to improve the way you operate?

Michael: Users can start with Maana’s semantic search engine, which is included as part of the Maana Knowledge Platform. You simply type in questions and the software’s natural language algorithms select the data and knowledge most relevant to your problem from the Maana Knowledge Graph. Another way to interact with the platform is using tailored knowledge applications devised jointly by Maana and the customer. Each application is a specific solution template with a dedicated user interface. The applications search the accumulated knowledge within the platform and, powered by data and computational models, produce recommendations on which the experts can make reasoned operational decisions.

Freddie: I think it is important to stress that neither Shell nor Maana intends this technology to

replace expert decision making. The technology enables the collection of feedback about how people have worked with the models to come to their decisions. This can be used to improve the models, which leads to better recommendations. If you detect human bias in the decision-making process, for example, you can correct for it using suitable data analytics. But there is a limit to how far you can go. We are still a very long way from having an artificial intelligence manage even parts of our business.

How far is Shell from having Maana Knowledge Platforms up and running in the business?

Freddie: Shell has two projects close to fruition. The first is corrosion-related and concerns the impact of crude oil selection on refinery equipment integrity. We are confident that a newly developed knowledge application will help engineers to understand more fully the corrosion risks associated with specific crudes. Its ultimate aim is to reduce maintenance costs and unplanned downtime.

The second is a health, safety and environmental risk application designed to help Shell identify operational incidents that could have had a worse outcome than they did. The aim is to improve our understanding of what really went wrong in these cases and thus prevent recurrences. This has involved aggregating and analysing information about health, safety and environmental incidents, as well as mitigation measures, in the company’s primary health, safety and environmental database and creating an algorithm that accurately predicts incident outcomes. Analysing the sheer volume of fragmented data, both structured and unstructured,

has been a big challenge. Traditional statistical tools and techniques are unable to help us solve this kind of problem.

In addition, there are several other projects at an earlier stage of development. A particularly exciting and challenging one aims to improve decision making in the exploration area through better understanding of reservoir basin features. Automatically extracting knowledge from scientific papers will feature heavily here.

This is all about knowledge, clearly, but do you know where all this is going to lead?

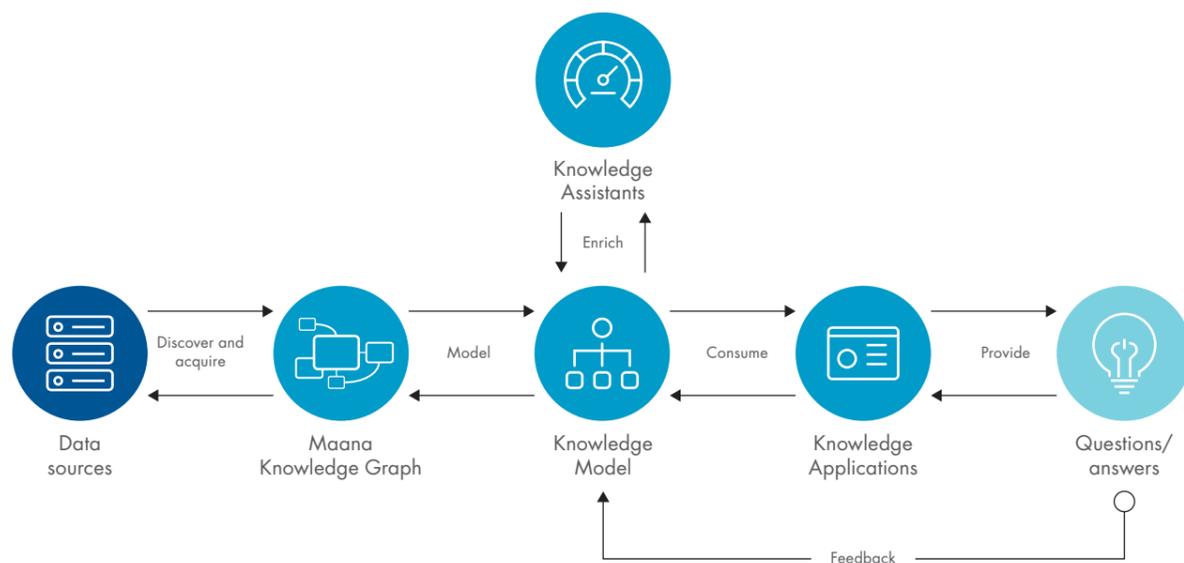
Freddie: The Maana Knowledge Platform is one of the enabling technologies that can support our vision of joining up the data from different parts of Shell in order to gain deeper insights into the broader enterprise. Decisions taken in one part of the business clearly affect others. A tool like this should help us to make improved business decisions by showing how any single decision affects not just that particular part of the business but also wider elements of Shell’s value chain, thus creating greater overall business value. Realising this vision will require the ability to encode even more kinds of data and information than is currently possible while also focusing on making the process of building knowledge models easier and faster.

Michael: Yes, the Maana vision is not to optimise just single activities but rather an entire enterprise – to transform whole organisations digitally. And I agree with Freddie that we have to make the technology more accessible to business users. Right now, progress is limited by the amount of IT technical expertise demanded at the front end of the process. We have to make that modelling part easier for business users. Training is a part of this, of course, and something into which we are putting more effort. The appetite is there and we just need to maintain it by getting a more user-friendly form of the technology into the hands of the real practitioners.

Making connections

If you think that a graph is a diagram that illustrates the relationships between different variables, then you are half way to understanding knowledge graphs. For these are conceptual constructs for capturing data, plus, crucially, their relationships and connections. The best systems are capable of making those connections according to the real meaning of the data rather than just a character-string match such as the “find” and “search” functions in Microsoft Word or Outlook. In other words, there is a strong semantic element to the process. Further, it is a dynamic process in which the connections are constantly enriched on the basis of the (semantic) searches being performed on the system. Google owns probably the best known (and the biggest) knowledge graph on the planet: it is what we turn to for information on just about anything these days.

What Maana is doing is extending these same principles to industrial and business operations where the scope of the data is obviously narrower but their complexity is greater; where the data exist in silos (and multiple formats) across an organisation; where the relationships are not always so obvious; and where much of the best subject matter and insights sit not in existing websites (as in the case of Google) but in people’s minds.



Within the Maana Knowledge Platform, the user creates a series of Knowledge Models with the help of the patented Maana Knowledge Graph and artificial intelligence algorithms (Knowledge Assistants). The models can be turned into Knowledge Applications that provide recommendations directly for day-to-day business operations.