

An aerial photograph of an excavator working in a dark, sandy environment. The excavator's tracks and mechanical components are visible. Three colored dots are overlaid on the image: a yellow dot on the upper arm, a blue dot on the lower arm, and a green dot on the ground below. A large red circle is partially visible on the left side of the image, containing the text.

Machine learning:
Asking the
right questions

Machine learning is promising to uncover new insight in the mining industry. One of the challenges to applying a successful machine learning initiative is to have a good understanding of the types of questions machine learning can answer.

Ask specific questions

One of the things to think about is how you ask the question. For instance, suitable machine learning questions are precise and will usually look to have a target number or word that describes the outcome.

Examples of good questions are:

- Will the recovery calculation from the mineral process be within the range 80-85%?
- What is the expected recovery going to be for this current shift in the tailings circuit?
- What is the chance that the recovery will be higher than last week's maximum recovery?
- What is the relationship between my incoming head grade and the expected recovery from the process?
- What is the expected commodity price for copper next month?

Examples of poor questions are:

- Will I get good recovery today?
- What is causing problems today on my site?
- Why am I getting poor recovery today?

Classification questions

Can you structure the question only to have a list of possible answers? Typical examples will just have two possible answers and will be a two-class classification question. These could be selective answers (A or B), logical (yes or no) or specific to a particular problem (assigned or not).

Some examples of two-class classification are:

- Will this tyre explode in the next 1000 km?

- Will this component fail in the next three days?
- Is this image taken of ore or waste?

If there are more than two alternative answers, then it could be a multi-class classification.

Some examples include:

- Which component will fail in the next seven days: gearbox, tyres, engine or hydraulics?
- What is the rock type shown in this image?
- Will it rain, snow or be fine tomorrow?

Anomaly detection

Maybe your problem is trying to determine if the data is normal or not. This looks like the two-class classification, but is asking if the data is weird or abnormal. Some anomaly detection algorithms can detect for abnormalities in the data even when there are no examples in the available training data set.

Some examples of anomaly detection questions are:

- Is this tyre pressure reading normal?
- Is the combination of readings typical for this scenario?
- Are the power readings normal for this time of year?
- Is this combination of process readings unusual?

Regression

If the purpose of the question is to get a number rather than a category or a class, then the question can be a regression question. Regression results will usually be a real number that can sometimes be negative or have lots of decimal points. These results may need to be interpreted to get the outcome required. Some interpretation examples: are rounding to the nearest whole number and assuming that negative numbers indicate a zero result.



Successful machine learning initiative is to have a good understanding of the types of questions machine learning can answer.

Examples of regression questions include:

- What will be the expected temperature tomorrow?
- How much power will be required for my mill today given my current hardness of ore?
- What percentage of these tyres are expected to be still operational at 20,000km?

Multi-class classification questions as regression

A regression approach to multi-class classification questions can also be useful. For example, “Which component will fail in the next seven days: engine, gearbox, tyres or hydraulics?” seems to require a classification or a single component that will fail. Taking a regression approach would reformulate the question to, “How likely is each component (engine, gearbox, tyres, hydraulics) to fail in the next seven days?” and would provide a numerical failure score for each component. The result would then be the highest scoring component.

Another example of restructuring a multi-class classification to a regression could be:

- “Which truck in my fleet needs servicing the most?” can be rephrased as, “How urgently does each truck in my fleet need servicing?”

Two-class classification as regression

Sometimes it is beneficial to reformulate two-class classification questions as regression questions. The regression version of the questions provides two scores that provide a ‘yes’ score and a ‘no’ score. The highest score can still be interpreted as either ‘yes’ or ‘no’, but can also handle the situation where there is ‘partly yes’ and ‘partly no’. Each of the scores for ‘yes’ and ‘no’ can be partial or complete scores and may provide more information than just a ‘yes’ or ‘no’.

Questions of this type often begin ‘how likely...’ or ‘what fraction...’

- How likely will it rain tomorrow?
- What fraction of components will fail in 3 months?
- What is the likelihood that this picture is of waste material?

Clustering questions

Clustering questions look to understand the structure of the data and tries to separate data into natural ‘clumps’ that a human can easily interpret.

Some examples of clustering questions include:

- Which pieces of equipment fail the same way?
- What is the natural way to group different operator behaviors?
- What process conditions cause similar outputs in the process?
- What is a natural way to group these safety incidents?

“What should I do next?” questions

Reinforcement learning algorithms allow a more advanced type of question to be asked. A question that can be linked to an action. What should I do next?

These questions are model-based and are rewarded when they make a ‘good decision’ and are penalized when they make a poor decision.

Examples of questions that are well suited to reinforcement learning include:

- Should the speed of the conveyor increase, decrease or stay the same?
- When should I order more consumables to ensure I minimise my inventory, but still meet my service level agreements?
- In what direction should the robot move given the current environment?

Conclusion

By understanding the different forms of machine learning questions and different algorithms, the mining industry will be able to initiate a successful machine learning initiative. While having the right data will still be essential, an insight into how to ask the question and the implications may uncover a new way to think about how to leverage the available data.

References:

<https://blogs.technet.microsoft.com/machinelearning/2015/08/26/what-can-data-science-do-for-me/>

<https://blogs.technet.microsoft.com/machinelearning/2015/08/27/what-types-of-questions-can-data-science-answer/>

About the author

Joseph Plant

**Natural Resources, Global Practice Lead –
Operations Improvement Energy,
Natural Resources & Utilities, Wipro Ltd.**

Joseph leads Wipro’s Global Operations Improvement Practice as part of the Natural Resources vertical and has experience in creating, designing and implementing operational improvement solutions for mining and metals customers.

He has a background in industrial automation and software development. He has deep experience in operational systems such as MES (Manufacturing Execution Systems), and leading data science initiatives and analytics for operational systems.

Joseph has been involved in creating and delivering numerous energy management initiatives for the mining and metals industry.

He has global experience and has worked on projects and opportunities across the globe including Australia, South Africa, China, Brazil, India, USA, UK, France, Belgium, Finland and Canada.

Joseph can be contacted on LinkedIn:

<https://au.linkedin.com/in/josephplant>



Wipro Limited

Doddakannelli, Sarjapur Road,
Bangalore-560 035,
India

Tel: +91 (80) 2844 0011

Fax: +91 (80) 2844 0256

wipro.com

Wipro Limited (NYSE: WIT, BSE: 507685, NSE: WIPRO) is a leading global information technology, consulting and business process services company. We harness the power of cognitive computing, hyper-automation, robotics, cloud, analytics and emerging technologies to help our clients adapt to the digital world and make them successful. A company recognized globally for its comprehensive portfolio of services, strong commitment to sustainability and good corporate citizenship, we have over 160,000 dedicated employees serving clients across six continents. Together, we discover ideas and connect the dots to build a better and a bold new future.

For more information,
please write to us at
info@wipro.com

