Strategic trend forecasting: anticipating the future with artificial intelligence

Could machines do better than humans at spotting events so far considered unpredictable, such as the COVID-19 pandemic? **Jim Marshall** *and* **Stephan Mergenthaler** *have been using natural language-understanding technology to detect trends—and analyse them at depth. They write that this type of analysis is an important foundation for people to better anticipate the future and prepare for it.*

Nassim Taleb's best-selling 2007 book "The Black Swan: The Impact of the Highly Improbable" was meant to offer a blueprint for more resilient business strategies in a networked world. Instead, the title has become a shorthand for any event that caught us unprepared. But are these types of events quickly labelled as black swans really unpredictable, or is it our inability that keeps us from connecting the dots at scale to spot and prepare for them? Could machines do better than humans at spotting them, given their near infinite capacity to read, absorb, and assimilate? And could we help communities and stakeholders to "look around the corner" a little on the topics that have greatest potential to impact their organisations, to build the kind of resilience that Taleb called for?

As part of our work developing the World Economic Forum's strategic intelligence capabilities, we have been aggregating and surfacing high-quality content on over two hundred global issues to over three quarters of a million users for nearly four years now. This gives us a tremendous data set from which we can extrapolate trends that are shaping the world around us. What's more, our hand-picked network of content partners from around the world means that we automatically exclude much of the noisy clickbait, fake news, and poor quality content that plague the Internet at large. We work with hundreds of think tanks, universities, research institutions and independent publishers in all major regions of the world to provide a truly global perspective. As such, when examining trends, we can be reasonably confident that our data are well positioned when it comes to the intrinsic biases inherent to open text analysis on uncurated content from the Internet.

Our recent advancements in harnessing machine learning and natural language-understanding technology mean that we are able to understand these trends at depth. We can not only look at the technical factors, such as volume and velocity of discussions, but also the strategic and emotional aspects driving them as well. Based on these unique trend characteristics, we wondered whether we might be able to predict their direction of travel. Are they a "flash in the pan" that will quickly come to a conclusion, or are they part of a longer, slow-burning trend that will build momentum over a period of years? Could the trend present opportunity or risk, is it disruptive, and could it end with a crisis or breakthrough? Perhaps more importantly, should you act now, or wait?

To answer these questions, we decided to look back in time through our database of over one million high-quality articles. We wanted to know, in retrospect, if we had known what to pay attention to, if we could have better prepared ourselves for future events — for example, a pandemic.

Long before the onset of the COVID-19 crisis, there were very visible warnings about our level of preparedness for such kind of pandemic. For example, in April 2018, academic articles such as "Bats, Coronaviruses, and Deforestation: Toward the Emergence of Novel Infectious Diseases?" and similar articles published by Frontiers featured clear warnings about what was about to unfold almost two years later. While it stated that *"the exact time and nature of the emergence of a disease cannot be predicted,"* it warned that *"the increased probability [of such events] must be considered seriously"*. Throughout 2019 we surfaced articles like this from Science Daily, highlighting that *"not a single country in the world is fully prepared to handle an epidemic or pandemic"* and this one from Project Syndicate entitled "Preventing the next Pandemic".

These are just some data points representing a plethora of warnings from the wider expert community in the years leading up to the pandemic. Why is it then that we often revert to the black swan theory to exonerate ourselves from any degree of preparedness? It's a principle we see at play in many other areas as well: Global societal trends and emerging technologies interact and create effects that may be apparent and widely discussed within narrow expert communities, but often fail to gather broad-based visibility before a substantial disruption takes place. Their development is often characterised by exponentiality and complexity – dynamics that are hard to grasp for the human brain, but fundamental to the behaviour of networked societies.

In other words, the signals of disruption may be there, but for humans to pay attention to, and connect the dots between, possible interactions and repercussions across hundreds of different topic areas is overwhelming and easily leads to neglect or denial. Doing this kind of systemic horizon scanning is next to impossible without the filtering ability, throughput, and analytical powers of intelligent machines.

Figure 1. Automatically detecting and monitoring topics using machine learning and natural language understanding
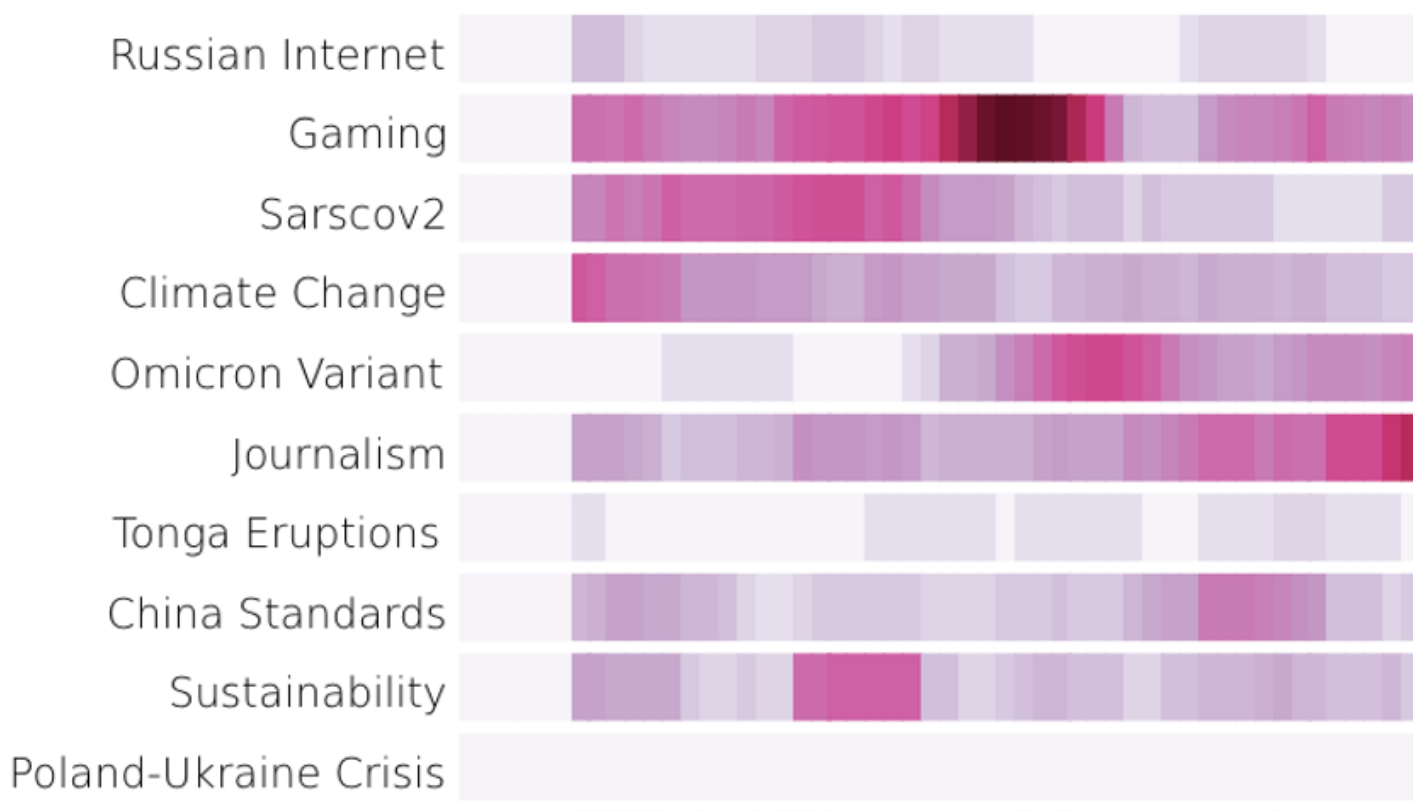
Building on top of our [previous work](#) using natural language understanding technology to detect important signals in content, we built an ensemble of advanced neural networks to identify, name, monitor and forecast trends. You can see a selection of the [COVID-19](#) related trends that our models have been tracking for us below. The height of the line corresponds to the relative volume of discussion about the topic and the colour of the line corresponds to the sentiment of the discussions (negative is red, neutral is yellow, positive is green). The shaded section on the right gives a glimpse of the future at the time of writing (mid-May 2022).

Figure 2. Trends detected by advanced neural networks

COVID-19

Travel restrictions

Vaccines

Shanghai Lockdowns
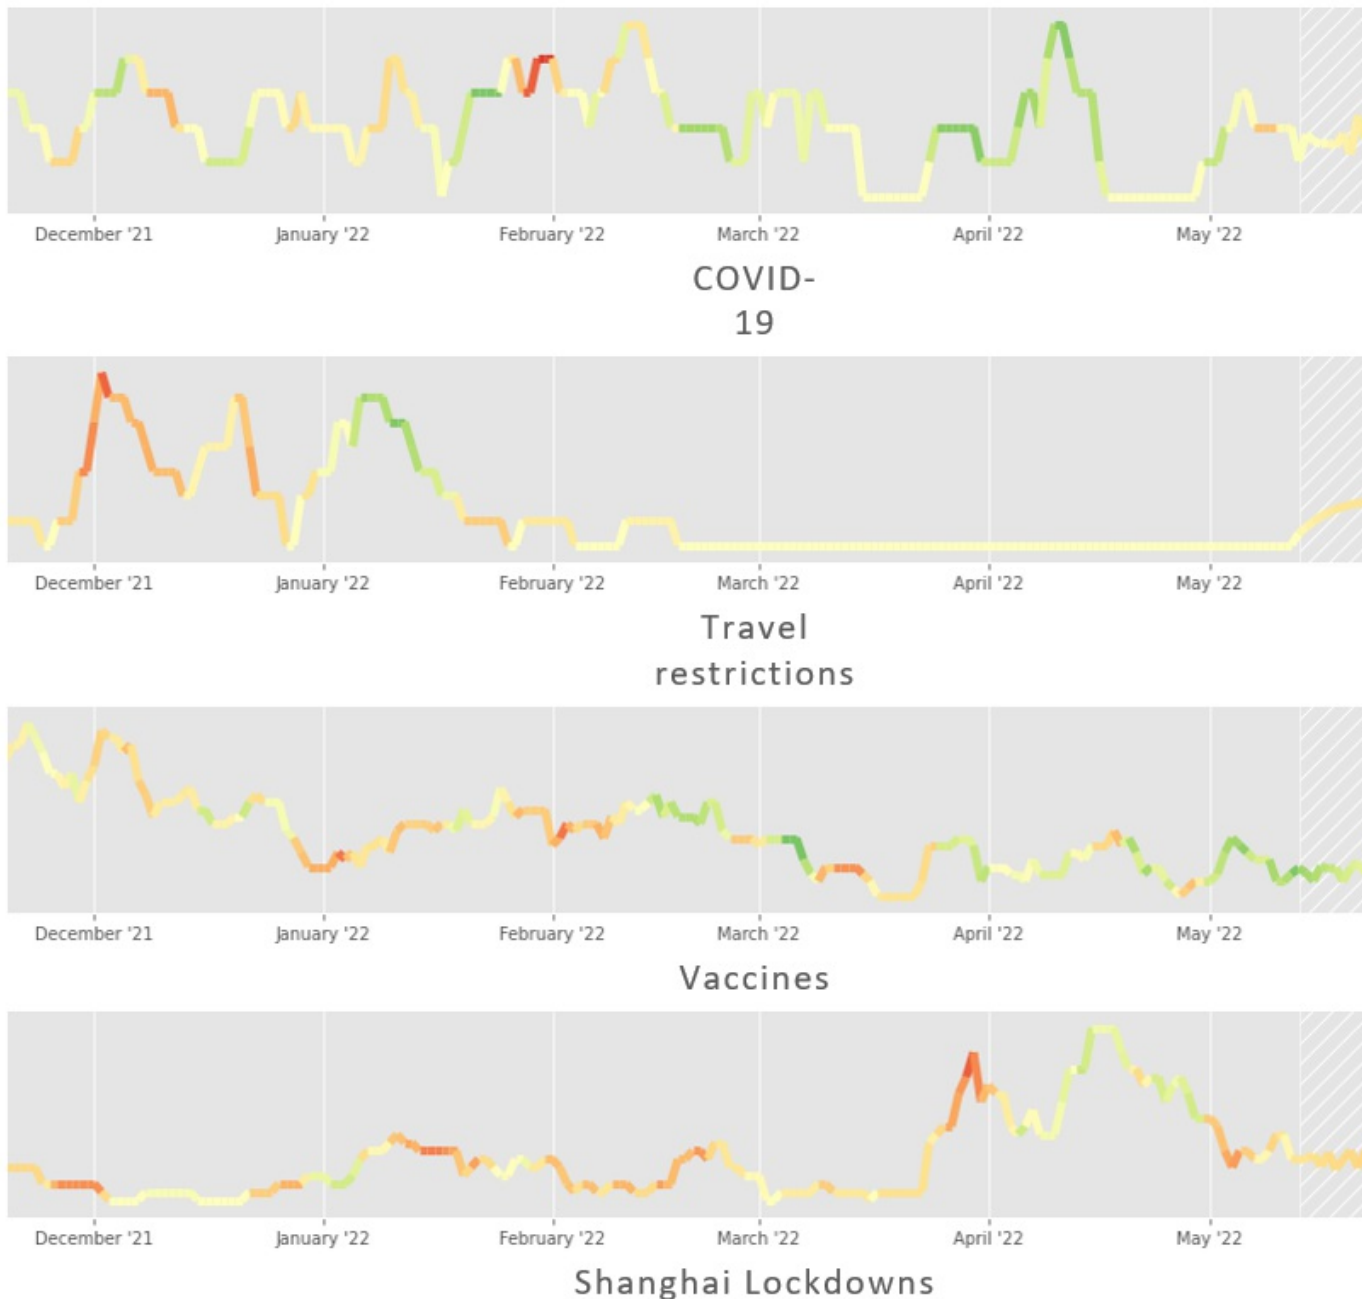
We are already piloting alerts based on the trends that our models predict will come to the fore in the next month. The beauty of this approach is that we can backtest our model over four years' worth of knowledge to give a general indication of its accuracy. So far, we are at around 70% accuracy within one standard deviation of the mean on forecasts up to one month in advance. As we add more signals into the mix, we expect to not only improve the accuracy but also the foresight timeframe.

Figure 3. Retrospectively comparing predicted vs. actual volume of discussions about the topic of stalkerware in April 2022

While much more work needs to be done to improve our ability to enact decisive collective actions when faced with signals of future disruption, we see this work as an important foundation for stakeholders to better anticipate and prepare for the future. In doing so, we will hopefully be in a better position to face the black swans of tomorrow.

♣♣♣

Notes:

- *This blog post represents the views of its author(s), not the position of LSE Business Review or the London School of Economics.*
- *Featured image by Brian McGowan on Unsplash*
- *When you leave a comment, you're agreeing to our Comment Policy*