nature

Learn from COVID before diving into a pandemic treaty

A treaty might help countries to prepare for the next pandemic – but first they must study what went wrong during this one.

Last week, 27 world leaders, including the president of the European Council and the director-general of the World Health Organization (WHO), called for the creation of a legally binding treaty so that the world is better prepared for the next pandemic. The European Union helped to initiate the idea, which is backed by countries in the African Union, Asia and South America. The United States and China are not in the founding group.

The leaders envisage a global agreement, organized through the WHO, that will enable nations to equitably share their expertise, equipment and knowledge in controlling and ending pandemics. "Together, we must be better prepared to predict, prevent, detect, assess and effectively respond to pandemics in a highly coordinated fashion," the leaders wrote in a series of opinion articles published simultaneously in newspapers around the world.

A pandemic treaty, they imply, would be the antidote to the current response, which is the opposite of cooperative. Countries that are doing vaccine research and manufacturing – such as those in the EU, the United States and India – are buying up most of the stock or restricting exports to other countries. Many of the richer nations are, for now, maintaining intellectual-property protection on important technologies. Earlier in the pandemic, countries were outbidding each other to find personal protective equipment and coronavirus testing kits.

Although the objective of a pandemic treaty is compelling, it's not clear whether, in practice, such an instrument would lead to a more unified and equitable response to a future pandemic. When the 2008 global financial crisis hit, world leaders, finance ministers and leaders of financial institutions worked in concert to stabilize economies. They didn't need a treaty. Nor should those working to navigate a path out of the pandemic.

It was – and remains – entirely possible for nations to come together and make a few key joint decisions, as many, including this journal, have been advocating. On vaccinations, for example, there could have been an agreement to share know-how and supplies so that all of the world's most vulnerable people could be protected first. Countries could similarly have cooperated by sharing research and experience to inform criteria for starting and ending lockdowns and for closing and opening borders.

Nonetheless, a pandemic treaty is not an unworkable

> **"It will need a broader constituency of support, and it will require patience, and give and take."**

idea, and multilateralism is infinitely preferable to nations acting alone. But it requires due diligence, alongside studies of the current pandemic response. It will also need a broader constituency of support beyond that of its founding leaders, and it will require patience, and give and take.

So what does a treaty need?

A pandemic treaty needs to meet at least four conditions. First, international treaties need the world to sign up. So far, only a small number of countries have done so. Different countries will have different reasons for taking part, which will need to be accommodated. A global treaty would also need participation from China and the United States, which is not certain at this time of renewed tensions in East–West relations. And treaty negotiation can take years, if not decades – so prepare for a long haul.

Second, a treaty needs researchers and non-governmental organizations (NGOs) to be actively involved. Historically, both have been crucial in creating treaties. Warnings from researchers helped the world to agree on international treaties on ozone depletion and climate change. NGOs have been instrumental in the creation of biodiversity treaties and the Treaty on the Prohibition of Nuclear Weapons, which entered into force in January. Their voices, along with those of the thousands of clinicians, epidemiologists, geneticists, public-health specialists, virologists and so many more who have been central to the pandemic response, have yet to be heard – and they need to be.

Third, before negotiating a new treaty, nations need to study why existing agreements are not working. Under the EU's pandemic plan, for example, procured vaccines are placed in a central pool and are released to countries according to priority categories. But this has proved controversial because it has prevented countries from purchasing their own supplies. Some EU countries have contravened the agreement and gone outside the pool to procure vaccines on their own.

The COVAX scheme has some elements of a treaty: countries and funders have promised to contribute funding and vaccines for the most vulnerable people in countries with the greatest need. This is an important initiative, and was designed to ensure that all of the world's most vulnerable people could be vaccinated first. However, it is clear that donor countries are waiting for vaccines to be rolled out to their own citizens before they release supplies to COVAX for recipient countries.

Clearly, although countries say they are committed to the principle of equitable allocation, this is not actually happening. That means there is a risk that, when the next pandemic strikes, countries will return to competing with each other for supplies, even with a treaty committing them to equitable access.

Fourth, if a new treaty is to be administered by the WHO, nations need to have an honest conversation about whether they will give the agency more powers. The WHO has been offering constant guidance since before it declared a Public Health Emergency of International Concern more than a year ago. Some countries were better able to follow its advice and have had more success at eliminating the virus,

nature

others less so. To succeed, a treaty that is administered by the WHO will need every country to respect its instructions.

After the 2008 global financial crisis, world leaders realized that parts of the architecture of international finance needed to be mended. But you cannot fix a broken system in the middle of a crisis. A treaty to fix today's ills has the potential to be a powerful instrument in a future pandemic, but, with countries still navigating their way out of this one, it's important to remember that people don't need an international law to pick up the phone and talk.

Rise of AI debaters highlights need for transparency

With artificial intelligence starting to take part in debates with humans, more oversight is needed to avoid manipulation and harm.

> **"Researchers are starting to incorporate elements of a theory of mind into AI models."**

Can a machine powered by artificial intelligence (AI) successfully persuade an audience in debate with a human? Researchers at IBM Research in Haifa, Israel, think so.

They describe the results of an experiment in which a machine engaged in live debate with a person. Audiences rated the quality of the speeches they heard, and ranked the automated debater's performance as being very close to that of humans. Such an achievement is a striking demonstration of how far AI has come in mimicking human-level language use (N. Slonim *et al. Nature* **591**, 379–384; 2021). As this research develops, it's also a reminder of the urgent need for guidelines, if not regulations, on transparency in AI — at the very least, so that people know whether they are interacting with a human or a machine. AI debaters might one day develop manipulative skills, further strengthening the need for oversight.

The IBM AI system is called Project Debater. The debate format consisted of a 4-minute opening statement from each side, followed by a sequence of responses, then a summing-up. Although Project Debater was able to match its human opponents in the opening statements, it didn't always match the coherence and fluency of human speech. This is partly because Project Debater is a machine-learning algorithm, meaning that it is trained on existing data. It first extracts information from a database of 400 million newspaper articles, combing them for text that is semantically related to the topic at hand, before compiling relevant material from those sources into arguments that can be used in debate.

Systems such as this, that rely on a version of machine learning called deep learning, are taking great strides in the interpretation and generation of language. But because training data are drawn from human output, AI systems can end up repeating human biases, such as racism and sexism. Researchers are aware of this, and although some are making efforts to account for such biases, it cannot be taken for granted that corporations will do so.

As AI systems become better at framing persuasive arguments, should it always be made clear whether one is engaging in discourse with a human or a machine? AI specialist Stuart Russell at the University of California, Berkeley, told *Nature* that humans should always have the right to know whether they are interacting with a machine — which would surely include the right to know whether a machine is seeking to persuade them. It is equally important to make sure that the person or organization behind the machine can be traced and held responsible in the event that people are harmed. Project Debater's principal investigator, Noam Slonim, says that IBM implements a policy of transparency for its AI research, for example making the training data and algorithms openly available.

Right now, it's hard to imagine systems such as Project Debater having a big impact on people's judgements and decisions, but the possibility looms as AI systems begin to incorporate features based on those of the human mind. Unlike a machine-learning approach to debate, human discourse is guided by implicit assumptions that a speaker makes about how their audience reasons and interprets, as well as what is likely to persuade them — what psychologists call a theory of mind.

Nothing like that can simply be mined from training data. But researchers are starting to incorporate some elements of a theory of mind into their AI models (L. Cominelli *et al. Front. Robot. AI* https://doi.org/ghmq5q; 2018) — with the implication that the algorithms could become more explicitly manipulative (A. F. T. Winfield *Front. Robot. AI* https://doi.org/ggvhvt; 2018). Given such capabilities, it's possible that a computer might one day create persuasive language with stronger oratorical ability and recourse to emotive appeals — both of which are known to be more effective than facts and logic in gaining attention and winning converts, especially for false claims (C. Martel *et al. Cogn. Res.* https://doi.org/ghhwn7 (2020); S. Vosoughi *et al. Science* **359**, 1146–1151; 2018).

As former US president Donald Trump has demonstrated, effective orators need not be truthful to succeed in persuading people to follow them. Although machines might not yet be able to replicate this, it would be wise to propose regulatory oversight that anticipates harm, rather than waiting for problems to arise. Equally, AI will surely look attractive to those companies looking to persuade people to buy their products. This is another reason to find a way, through regulation if necessary, to ensure transparency and reduce potential harms. AI algorithms could also be required to undergo trials akin to those required for new drugs, before they can be approved for public use.

Government is already undermined when politicians resort to compelling but dishonest arguments. It could be worse still if victory at the polls is influenced by who has the best algorithm.