NASA/ESA/CXC/SSC/STScI

The rise of the machines
in our quest to understand the universe

Michelle Lochner explains how machine learning is used in astronomy and cosmology

In 1967 Jocelyn Bell, a young PhD student at the University of Cambridge, discovered something no one had ever seen before. She had been tasked with building a radio telescope, an instrument designed to pick up radio waves (as opposed to visible light) to study the universe. While looking at the data from her telescope, which in those days was in the form of long sheets of paper with lines on them, she observed something that looked like a pulse repeating with astonishing regularity from a distant corner of our galaxy. It was initially dismissed as noise, but Jocelyn persisted that she had discovered a true cosmic anomaly. Eventually, it was realised she had made one of the most important discoveries of the 20th century.

Jocelyn had discovered pulsars, incredibly dense stars that emit radio waves as they rapidly rotate. Pulsars are now some of the most well-studied objects in astronomy. The important thing to note is that it required a human being to look at the data to spot something weird.

Fast forward 52 years into the future: over 10 000 km away from Cambridge – in our backyard in the Karoo – another radio telescope is being built. With half a century of technological advancement, this telescope, the Square Kilometre Array (SKA), is significantly bigger and more powerful than Jocelyn's telescope. In fact, when finished the SKA will be the biggest telescope (of any kind) in the world.

The SKA is an array, meaning it will consist of hundreds of telescopes working together to make radio images of the sky. The SKA will be split in two, with one component being built in South Africa and the other in Australia. Currently, there is a precursor for the SKA in the Karoo, called MeerKAT. It is made up of 64 individual antennas, and is already one of the world's most sensitive telescopes. MeerKAT was proudly designed and built by South Africans, and is already in high demand in the scientific community.

Jocelyn Bell, 1967

Roger Haworth, CC BY-SA 2.0

As well as being much bigger and more sensitive than Jocelyn's telescope of 1967, the SKA will also produce dramatically more data – as much as 100 PB (petabytes; 1 PB = 1 million GB) a day even in its first phase of construction. If you could convert all that data into a song and listen to the universe singing to you, it would take approximately two million years to play back the data collected in a single day. The universe is a busy and complicated place!

Across the ocean in Chile, another marvellous telescope is under construction, called the Large Synoptic Survey Telescope (LSST). This big project is led by the United States, but several South African researchers are also involved. Unlike MeerKAT, LSST looks more like a traditional telescope that sees the universe in visible light. It is made up of huge mirrors as well as what will be the world's largest CCD camera (charge-coupled device, the type of sensor found in most digital cameras). LSST has clever optics that will allow it to observe almost the entire southern sky every three nights. What this means is that over its 10-year survey LSST will be able to make a movie of the universe.

When you look up at the sky at night, I'm sure you'll agree that it doesn't look like it changes much. Apart from the constellations appearing to move on their steady path throughout the night and the year, not much seems to happen in the sky. But if your eyes were bigger, like a telescope's eye, you'd be able to see fainter things – and then you'd realise that the universe is changing all the time.

Often the changing sky is associated with dramatic or explosive events. When a big star runs out of fuel, it blows up in a massive explosion called a supernova. Now the death of a star may be a rare event, but if you observe the whole sky you end up catching a lot of supernovae! There are other dramatic events like supermassive black holes devouring nearby stars, or when pairs of the same incredibly dense stars that form pulsars merge in an explosive fashion. LSST will find everything that changes in the sky, from nearby rocks hurtling through space to distant exploding stars.

Enter machine learning

Machine learning has revolutionised the modern world, and astronomy is no different. Not only can it automate tedious tasks once left to the astronomer, like detecting artefacts in images (think satellite trails and cosmic rays), but it can also be used to automate discovery and seek out unknown unknowns – things we don't know we don't know, being so unexpected that we would not consider or predict them.

With the imminent arrival of these two telescopes, the SKA and the LSST, astronomy is undergoing the same revolution as every other field. We are rapidly finding we have more data than we know how to handle! We have no choice but to automate many of the tasks that would traditionally be done by astronomers. Machine learning is becoming a key tool in smart automation.

Here's an example of the data deluge that we'll need machine learning to handle: LSST will detect 10 million transient events every night. That means 10 million things will have changed from the previous night. As much as half of these won't be real events, they will be things like problems with the CCD or aeroplanes flying past. But that still leaves millions of events that are real astrophysical sources. Most of these will be classes of objects we've seen before, like exploding stars (supernovae) or supermassive black holes swallowing stars (active galactic nuclei). How do we automatically classify these objects so that astronomers can just study the types of objects they're interested in? Or the bigger question, what if one of those 10 million transients is something we've never seen before, like pulsars in 1967? How do we find them when there just isn't enough time in the world to look through all the data with human eyes?

Our Data Science group, which is a mix of collaborative groups at the African Institute for Mathematical Sciences (AIMS) and the South African Radio Astronomy Observatory (SARAO), led by Professor Bruce Bassett, is engaged in a number of projects to make use of machine learning for astronomy.

We have teams working on classifying types of transients for LSST, MeerKAT and other telescopes. We have a large team focusing on automatically detecting radio frequency interference in data, which can be caused by satellites, aeroplanes and electronic devices near the MeerKAT telescope site. And Prof. Bassett and I lead a team focusing on anomaly detection in astronomy, one of the most exciting areas of machine learning.


MeerKAT, near Carnarvon in the Northern Cape, consists of 64 radio antennas.

SARAO

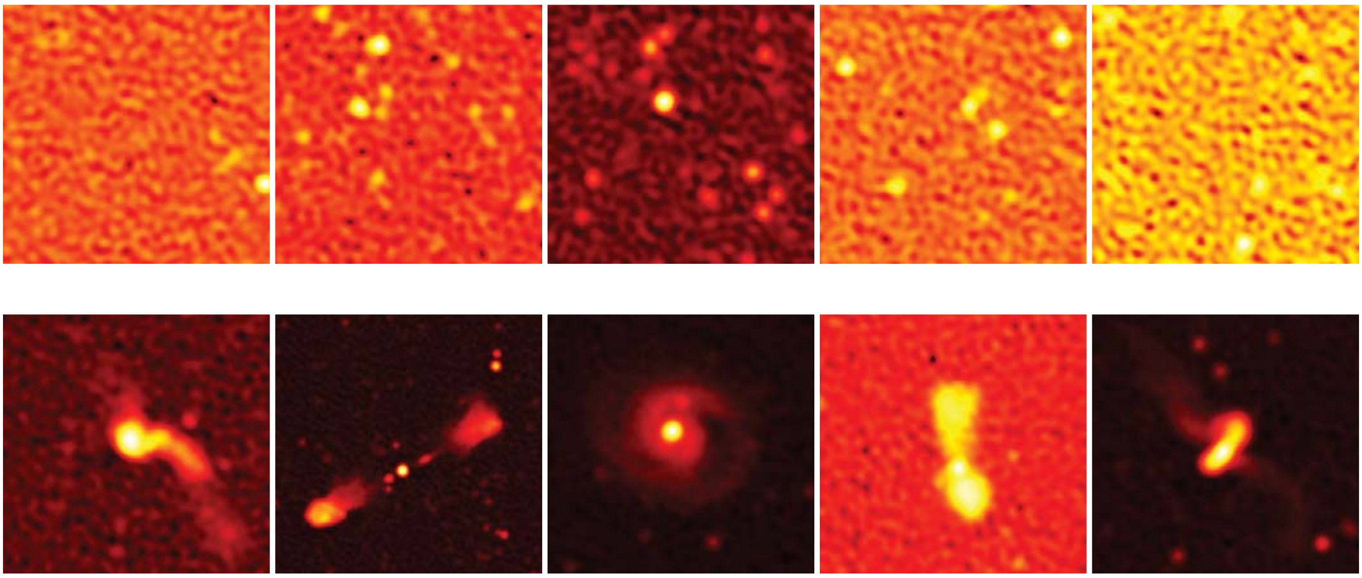

SARAO

Cutouts from a MeerKAT radio image showing (top row) point sources, representing radio galaxies too distant for MeerKAT to resolve, and (bottom row) anomalies, representing rare objects that might turn out to be new discoveries.

Anomaly detection

Anomaly detection is all about finding a needle in a haystack: finding that one rare or even never-before-seen object in a dataset of thousands of ordinary things. Let's illustrate anomalies with a real example. Above are some cutouts from a MeerKAT radio image. MeerKAT stared at a relatively empty patch of sky to make a large radio image that I chopped up into thousands of smaller cutouts (each cutout is a piece of sky 0.04 degrees across, which is around 10 times smaller than the moon). The sensitivity is turned up so that empty sky shows up as orange patchy noise, much like the white noise of a detuned television.

If we examine these cutouts, we find that most of them look like the ones in the top row: pretty boring, noisy images. We also see some cutouts with point sources – radio galaxies too distant for MeerKAT to resolve them clearly, so they just look like blobs. While lots of important science can be done with point sources, they're not anomalies. What we want to find automatically are things like those in the bottom row.

The problem with images is that they are actually very complicated for computers to understand. While even a three-year-old can tell the difference between an elephant and a mouse, it's taken decades of research and a revolution in machine learning to get a computer to be able to do the same thing. So we have to extract features from the images to be able to reduce them down to a simple set of numbers a computer can understand.

For this problem, we made use of an advanced school of machine learning called deep learning. Deep learning tries to emulate the way the human brain works by building a network of 'neurons' – a neural network – to map inputs (such as images) to outputs (such as the animal's name). We used a type of neural network called an autoencoder, which finds a small set of numbers that represents the image. We can then take a look at these numbers on a plot. We used a second stage of machine learning to do the actual anomaly detection. The algorithm we used works by trying to isolate points that are different from the others.

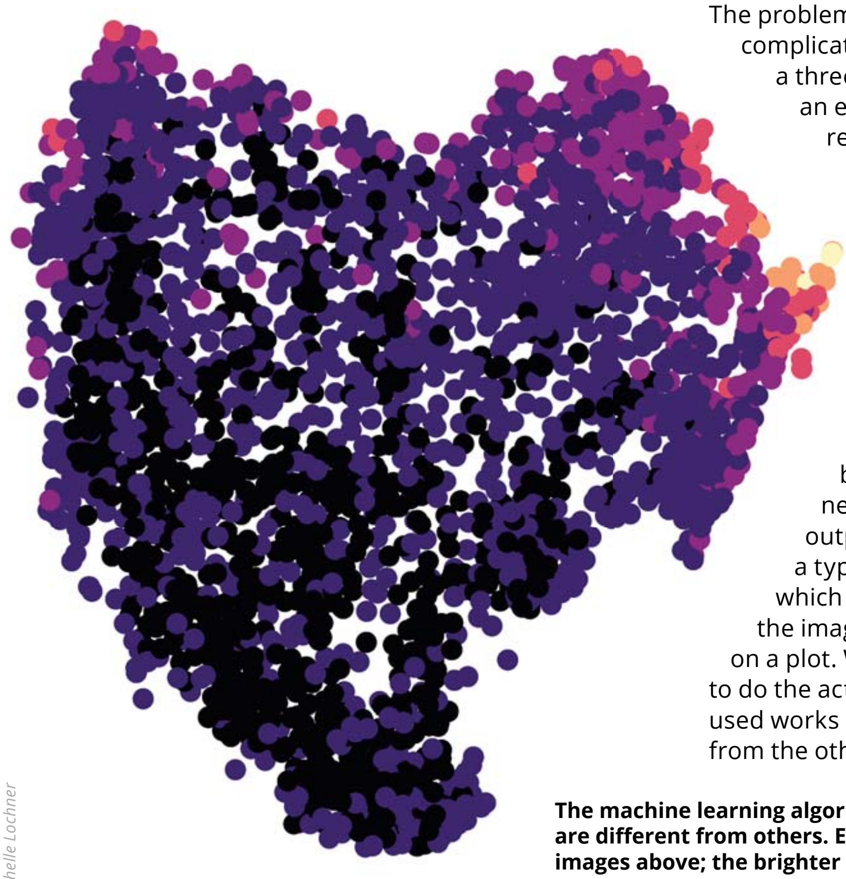

Michelle Lochner

The machine learning algorithm for anomaly detection isolates points that are different from others. Each dot is a representation of one of the cutout images above; the brighter the dot, the more anomalous the image.

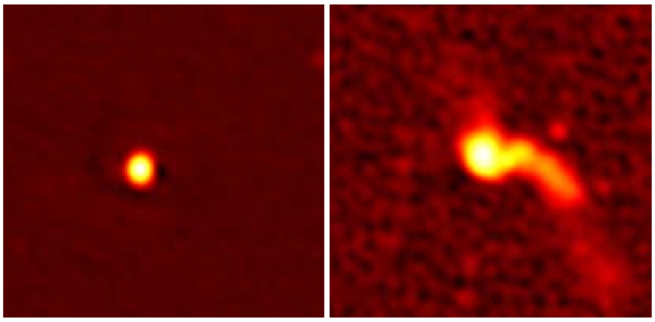

SARAO

One scientist may be curious about how this anomalous artefact was produced...

But another will be more interested in astronomical objects like these.

Each dot in the bottom left image is actually a representation of one of the cutouts seen earlier. The brighter the dot, the more anomalous the machine learning algorithm thinks the image is. In this way, we can turn quite a complicated thing (an image) into a simple set of numbers that a computer can understand. The entire set of 2 500 cutouts, which would take a human hours to go through looking for interesting objects, is processed in under two minutes on a normal desktop.

Finally, there is still a subtle distinction between what is anomalous and what is interesting. For instance, someone responsible for producing the left image above might be interested in this source. It looks a little strange because of the dark spot on its side and could be an artefact. However, a scientist looking for active galaxies would not be interested. So we build a final layer to turn our anomaly detection algorithm into an astronomical recommendation engine – an algorithm capable of learning users' interests and showing them more sources that it thinks they will be interested in.

This application of machine learning could be a massive time saver in astronomy. Indeed, it will be the only way to explore datasets that are too large to sift through manually. We hope our framework will be an entirely new way of working with data, and perhaps allow us to make the next big discovery, just like Jocelyn Bell did more than 50 years ago.

Some people are (quite understandably) afraid of machine learning and how it is impacting our society. Like any technological advancement, it can be used to help or harm. But by combining next-generation telescopes with machine learning tools, astronomers are finding new ways to unlock the mysteries of the universe.

Dr Michelle Lochner is a researcher in a position shared between the African Institute for Mathematical Sciences and South African Radio Astronomy Observatory. She did her undergraduate degree in Physics and Electronics at Rhodes University, before moving on to postgraduate studies at the University of Cape Town, where she obtained her PhD in Mathematics and Applied Mathematics.

Jocelyn Bell, now Dame Jocelyn Bell Burnell, was awarded the US$3 million Special Breakthrough Prize in Fundamental Physics in November 2018 in recognition of her 'fundamental contributions to the discovery of pulsars, and a lifetime of inspiring leadership in the scientific community'.

The award comes five decades after her discovery of pulsars, for which her PhD supervisor Antony Hewish and his collaborator Sir Martin Ryle were awarded the 1974 Nobel Prize in Physics. She has remained deeply engaged in astronomy, teaching at multiple research institutes and taking on leadership roles such as project manager of the James Clerk Maxwell Telescope in Hawaii. She has been President of the Royal Astronomical Society, the Institute of Physics and the Royal Society of Edinburgh, and is currently a Visiting Professor of Astrophysics at the University of Oxford and Chancellor of the University of Dundee. She received a CBE in 1999 and a DBE in 2007 for her services to astronomy.

Bell Burnell announced that she would donate the prize money to establish research studentships or scholarships for people from under-represented groups in physics. This is in keeping with her instrumental role in the creation of the Athena SWAN (Scientific Women's Academic Network) Charter and awards, established in 2005 to advance the careers of women in science, technology, engineering, maths and medicine.

Bell Burnell is the fourth recipient of the prize, which was previously awarded to Stephen Hawking, seven CERN scientists whose leadership led to the discovery of the Higgs boson, and the entire LIGO collaboration that detected gravitational waves.

For an entertaining account of Bell Burnell's life, see: https://physicsworld.com/a/look-happy-dear-youve-just-made-a-discovery/