

Computer Science

Autonomous balloons take flight with AI

Scott M. Osprey

An artificially intelligent controller can station a stratospheric balloon for weeks at a time without full knowledge of surrounding winds, opening the prospect of unsupervised environmental monitoring. See page XXX

The goal of an autonomous machine is to achieve an objective by making decisions while negotiating a dynamic environment. Given complete knowledge of a system's current state, artificial intelligence and machine learning can excel at this, and even outperform humans at certain tasks — for example, when playing arcade and turn-based board games¹. But beyond the idealized world of games, real-world deployment of automated machines is hampered by environments that are sometimes noisy and chaotic, and which are not adequately observed. The difficulty of devising long-term strategies from incomplete data can also hinder the operation of independent AI agents in real-world challenges. On page XXX, Bellemare *et al.*² describe a way forward by demonstrating that stratospheric balloons, guided by AI, can pursue a long-term strategy for positioning themselves about a location on the equator, even when precise knowledge of buffeting winds is not known.

Fixed-volume balloons, known as super-pressure balloons, are often used to carry out unmanned experiments in the upper atmosphere. Station-keeping is the act of maintaining the position of such a balloon within a certain horizontal distance of a ground location (the

station). This involves changing the balloon's height to move it between regions in which winds blow in different directions — when the balloon is driven away from its station by winds at one height, it moves to a different height where the winds can blow it back again (Fig. 1).

Self-navigating balloons do one of two things to stay within range of their stations. When a balloon is outside its range, the controller seeks winds pointing to within a small angle of the station. However, the balloon preferentially seeks out lighter winds when inside the target range and close to the station. Balloons that are more active in exploring winds above and below are more likely to find suitable winds to help achieve station keeping, but this comes at the expense of using battery power that might be needed for other tasks, such as relaying telecommunications or environmental monitoring. These competing factors need to be weighed up carefully.

A type of machine learning known as reinforcement learning can be used to train an artificial agent to make an optimal sequence of decisions. In the case of a super-pressure balloon, the decisions are whether to rise, fall or do nothing, and are based on an historical record of global winds³, local observed and forecast winds, and projected future flight paths. Crucially, the available wind data are sparse and do not fully constrain the flight controller's decision-making.

In their system, Bellemare *et al.* filled in the gaps by adding randomly generated 'noise' to the wind data, to better map out the range of winds that could plausibly occur, and to improve assessments of the variety of paths the balloon might take in the future. The resulting wind information and its statistical uncertainty, together with a small number of balloon-relevant

parameters, were used to train a machine-learning system known as an artificial neural network, and ultimately improved decision-making time during flights compared with previously used control systems using similar power.

Earlier applications of reinforcement learning, which included playing classic board and arcade games, were trained using complete information sets¹ – the same information that is available to human controllers⁴. These allowed like-for-like performance comparisons between humans and AI players. However, the challenge confronting Bellemare and colleagues was that incomplete knowledge of environmental winds not only makes it difficult to judge the optimal actions to take, but also makes forecasts of future states following these actions uncertain. These problems are further compounded by other practical uncertainties that don't affect game controllers, such as those associated with internal balloon motions, power management and battery health. Bellemare and co-workers' success therefore represents a big advance in the use of reinforcement learning for real-world applications.

Station-keeping performance is ultimately limited by the range of wind speeds and directions in the region surrounding the balloons (at heights of 15 to 20 kilometres, for the present study). The winds must also switch direction so that balloons can adjust their trajectory to stay within range of the station. These special conditions only persist for months at a time within the equatorial stratosphere, where Bellemare and colleagues' study was carried out – and where a slow procession of opposing winds peak in strength near 30 km, before descending and dissipating near 15km, switching direction every 14 months or so⁵. Such wind diversity also occurs elsewhere, but is less reliable and generally occurs beyond the range of heights at which a single super-pressure balloon can operate. During the flight campaign

described in the present study, larger wind disturbances originating from high latitudes occurred in the tropical stratosphere, and probably assisted station-keeping (Fig. 1). Bellemare and colleagues' system might therefore struggle to achieve the same success at other locations. However, smaller more rapid wind variations can also occur, including atmospheric waves of various types⁶, which a skilful controller could navigate to its advantage.

The advent of effective autonomous super-pressure balloons would open up a range of commercial and scientific applications for probing Earth's atmosphere and that of other planets. Such balloons are already used to study small and large-scale waves in the tropical stratosphere⁷, and to detect low-frequency sounds produced by the ocean⁸, lightning⁹ and earthquakes¹⁰. They have also been proposed for use in future explorations of Venus's upper atmosphere¹¹, to search for signs of active volcanism and chemical signatures of life¹². Moreover, the ability to fix a balloon's geographic position is crucial if they are to be used to build an aerial wireless network for telecommunications – an early objective of Project Loon, the owners of the balloons used in Bellemare and colleagues' study.

Fixing a balloon's location for months at a time would allow long-term environmental monitoring, for example, of air quality over cities, of carbon fluxes from heat-stressed forests and of regions of thawing permafrost. Other applications include monitoring animal migration routes and illicit trafficking of goods and people across borders. These applications will become increasingly relevant as the effects of climate change become more pronounced, as restrictions on movement are imposed by global events such as

COVID-19, and as long-term climate-change mitigation involving aviation prompt the search for alternative platforms for making aerial observations.

Scott M. Osprey is from the National Centre for Atmospheric Science based in the Department of Atmospheric, Oceanic & Planetary Physics, University of Oxford, Oxford OX1 3PU, UK.

e-mail: scott.osprey@physics.ox.ac.uk

References

- 1 Silver, D. *et al.* Mastering the game of Go without human knowledge. *Nature* **550**, 354–359 (2017) doi:10.1038/nature24270
- 2 Bellemare et al., *Nature* XXX, XXX-XXX (2020).
- 3 Hersbach, H. *et al.* The ERA5 global reanalysis. *Q. J. R. Meteorol. Soc.* qj.3803 (2020) doi:10.1002/qj.3803.
- 4 Mnih, V. *et al.* Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015) doi:10.1038/nature14236
- 5 Baldwin, M. P. *et al.* The quasi-biennial oscillation. *Rev. Geophys.* **39**, 179–229 (2001) doi:10.1029/1999RG000073
- 6 Schoeberl, M. R. *et al.* Gravity wave spectra in the lower stratosphere diagnosed from project loon balloon trajectories. *J. Geophys. Res. Atmos.* **122**, 8517–8524 (2017) doi:10.1002/2017JD026471
- 7 Haase, J. *et al.* Around the World in 84 Days. *Eos (Washington. DC)*. **99**, (2018) doi:10.1029/2018EO091907.

- 8 Bowman, D. C. & Albert, S. A. Acoustic event location and background noise characterization on a free flying infrasound sensor network in the stratosphere. *Geophys. J. Int.* **213**, 1524–1535 (2018) doi:10.1093/gji/ggy069
- 9 Lamb, O. D., Lees, J. M. & Bowman, D. C. Detecting Lightning Infrasound Using a High-Altitude Balloon. *Geophys. Res. Lett.* **45**, 7176–7183 (2018) doi:10.1029/2018GL078401
- 10 Krishnamoorthy, S. *et al.* Detection of Artificially Generated Seismic Signals Using Balloon-Borne Infrasound Sensors. *Geophys. Res. Lett.* **45**, 3393–3403 (2018) doi: 10.1002/2018GL077481
- 11 Glaze, L. S., Wilson, C. F., Zasova, L. V., Nakamura, M. & Limaye, S. Future of Venus Research and Exploration. *Space Sci. Rev.* **214**, 89 (2018).
- 12 Greaves, J. S. *et al.* Phosphine gas in the cloud decks of Venus. *Nat. Astron.* (2020) doi:10.1038/s41550-020-1174-4.




Figure 1 | Station-keeping for stratospheric balloons. Unmanned balloons, known as super-pressure balloons, are used to carry out experiments in the upper atmosphere. Station-keeping is the act of maintaining the position of a balloon within a certain range of a specific position on the ground. The balloon's height is altered to move it between altitudes at which winds blow in different directions — when the balloon is driven away from its station by winds at one height, it moves to a different height at which the winds can blow it back again. Bellemare *et al.*² report that a type of machine learning, known as reinforcement learning, can be used to train an autonomous control system for station-keeping, which outperforms previously used control systems.