can reduce the emission quality.

The electronic structure of lead halide perovskites makes photon emission in these materials much more resilient to the presence of crystal and surface defects than in III–V semiconductors, and might allow quantum-dot devices to be made without the need for rigorous surface-passivation methods. Highly efficient emission from perovskite quantum dots in solution has been demonstrated, including in the notoriously challenging blue spectral region[5], but the fabrication of blue-emitting LEDs has been challenging. This is because ultra-small, quantum-confined perovskite quantum dots are not stable enough to withstand the elaborate procedures used to deposit them from solution to form solids and to exchange ligands at their surfaces, as is needed to enable effective injection of charge carriers into the quantum dots in devices[6].

Jiang and colleagues now report a conceptually novel method for forming quantum-confined, blue-emitting quantum dots directly on the substrate of a device, circumventing the need for deposition and ligand exchange. Their method is remarkably easy to carry out: chemical precursors of perovskites are simply mixed with ligands on the surface of a substrate, directly yielding a layer of the desired quantum dots.

The key to success lies in the rational design of the ligands (Fig. 1). First, the authors' choice of a small, electrically conductive tail allows injection of electrons and holes into the inorganic quantum dots in the emissive layer of the devices. Second, the use of a bulky head group stops the ligand molecules from binding in the voids in the perovskite crystal structure, thus preventing the formation of a phase known as a layered perovskite, which has undesirable properties[7]. Finally, the incorporation of an electron-attracting halogen atom (such as a bromine atom) into the tail increases the strength of binding of the head group to the perovskites. Overall, these ligand effects thermodynamically favour the formation of stable perovskite quantum dots that are small enough to have the quantum confinement needed for blue emission.

The authors' method provides highly emissive, colour-pure solids of quantum dots, meeting the stringent requirements of energy-efficient blue-emitting perovskite LEDs. Indeed, the authors used their method to make LEDs that surpass the best carrier-to-photon conversion efficiencies of previously reported perovskite devices by a factor of up to two[8].

The reported method for making solids of perovskite quantum dots might lead to further advances in the development of perovskite LEDs. It should be noted that the stability of perovskite quantum-dot LEDs is not yet competitive with that of equivalent devices that use III–V semiconductor quantum dots, but there is likely to be room for improvement as the chemistry of lead halide perovskites becomes better understood and controlled. Advances such as those of Jiang *et al.* might also present opportunities for making solids composed of strongly coupled quantum dots, for example as non-classical light sources that exhibit exotic quantum behaviour[9].

Hendrik Utzat is at the College of Chemistry, University of California, Berkeley, Berkeley, California 94720, USA. **Maria Ibáñez** is at the Institute of Science and Technology of Austria, Klosterneuburg 3400, Austria.
e-mails: hutzat@berkeley.edu; mibanez@ist.ac.at

1. Fakharuddin, A. *et al. Nature Electron.* **5**, 203–216 (2022).
2. Jiang, Y. *et al. Nature* **612**, 679–684 (2022).
3. Alivisatos, A. P. *J. Phys. Chem.* **100**, 13226–13239 (1996).
4. Won, Y.-H. *et al. Nature* **575**, 634–638 (2019).
5. Protesescu, L. *et al. Nano Lett.* **15**, 3692–3696 (2015).
6. Shamsi, J. *et al. Nature Nanotechnol.* **16**, 1164–1168 (2021).
7. Li, X., Hoffman, J. M. & Kanatzidis, M. G. *Chem. Rev.* **121**, 2230–2291 (2021).
8. Dong, Y. *et al. Nature Nanotechnol.* **15**, 668–674 (2020).
9. Rainó, G. *et al. Nature* **563**, 671–675 (2018).

The authors declare no competing interests.

Medical research

A refined use of mutations to guide immunotherapy

Chao Cheng & Christopher I. Amos

Assessment of a tumour's mutational profile offers a way of predicting a person's response to anticancer therapies called immune-checkpoint inhibitors. It seems that such approaches might fall short for people who are not of European ancestry.

Immunotherapy harnesses the immune system to target tumour cells, and is used to treat several types of cancer. The therapy targets 'immune checkpoint' proteins, such as the protein PD-1, that can dampen immune responses. However, only a subset of people respond to treatment with what are called immune-checkpoint inhibitors (ICIs)[1]. Writing in *Cancer Cell*, Nassar *et al.*[2] provide insight into how decisions to use such treatments might be improved.

A measurement called the tumour mutation burden (TMB) is defined as the number of somatic mutations in the protein-coding region of the genome (the exome) per million bases of DNA. (Somatic mutations are those that are not inherited (germline) and that instead arise after birth.) The TMB correlates with the overall response to ICIs: in general, a better response is observed in certain types of cancer, such as melanoma and lung cancer, that have a high TMB. For a specific type of cancer, tumours with a high TMB are thus more likely to respond to ICIs than are those with a low TMB. A high TMB could mean that many abnormal proteins have been generated that might be recognized by the immune system and trigger an anticancer response.

"Reference genomes tend to have a European-descent bias in terms of the genetic variation represented."

A clinical study of the anti-PD-1 inhibitor pembrolizumab reported an objective response rate (corresponding to a reduction in tumour size) of 31.4% in people with a TMB of at least 175 mutations across the exome, compared with 9.5% for individuals with a TMB of fewer than 175 mutations across the same region[3]. The predictive value of TMB for a person's response was independent of other potentially relevant factors, such as expression of the protein PD-L1 (PD-1's binding partner); the tumour type; or the status of a genomic feature called microsatellite instability, which is associated with a high rate of a specific type of mutation.

Nassar and colleagues investigated how an individual's genetic ancestry affected TMB estimates obtained from sequencing specific genomic regions (Fig. 1). For this, the authors used sequencing panels, which focus on a selected set of genes that have variants known to be associated with cancer. Tissues from two groups of patients were assessed: one from the Dana–Farber Cancer Institute (DFCI) and the other from the Memorial Sloan Kettering Cancer Center (MSKCC), both in the United States.

People in the DFCI group had one of seven types of solid tumour, and their cancer had spread to other parts of their bodies. They had been treated with antibodies targeting PD-1, PD-L1 or another checkpoint protein called CTLA-4. The TMB of samples from these patients was determined using Oncopanel, a

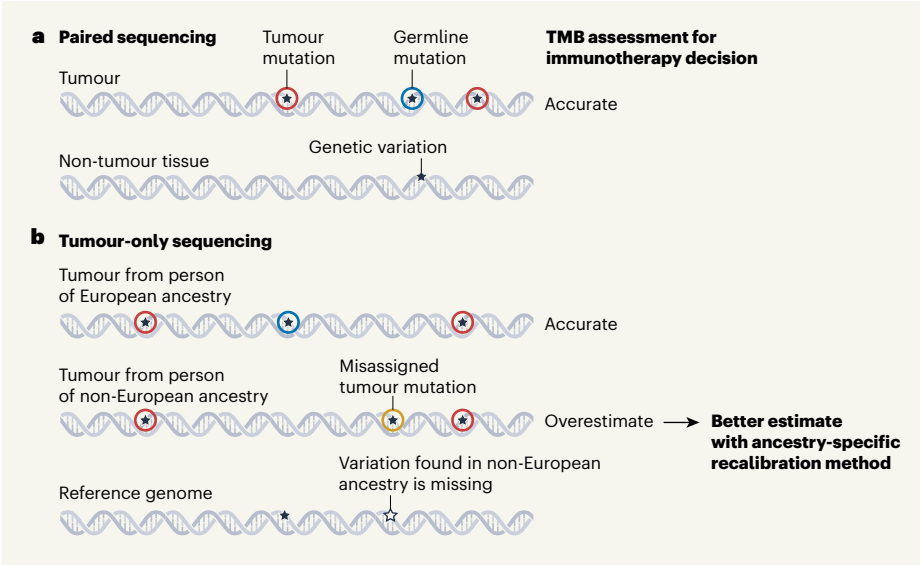

Figure 1 | Identifying tumour mutations to guide a decision about treatment. A type of anticancer immunotherapy is recommended only if the number of tumour mutations, termed the tumour mutation burden (TMB), is sufficiently high. Nassar *et al.*[2] assessed the accuracy of approaches used to estimate TMB. **a**, In a paired-sequencing approach, DNA is sequenced from samples of a person's tumour and of their non-tumour tissue. By comparing the results, sites of genetic variation are assigned as either inherited mutations (germline mutations) or tumour mutations. This method accurately assesses the TMB. **b**, In a tumour-only sequencing approach, the tumour-sequencing results are compared with the use of non-tumour sequences obtained from reference genomes. This approach is suitable for people of European ancestry, but tends to overestimate the TMB for people of non-European ancestry. This is because genetic variations in individuals of non-European ancestry tend to be under-represented in reference sequences. The authors developed a method to recalibrate tumour-only approaches so that ancestry is taken into account when estimating the TMB.

panel-based targeted-sequencing assay. The MSKCC group consisted of 1,898 individuals with non-small cell lung cancer (NSCLC) who had also been treated with ICIs. The TMB of these samples was determined by a different approach that involved tumour sequencing along with paired sequencing of non-tumour tissue from the individuals.

In addition, a separate group of 466 individuals was analysed; these included 327 people with NSCLC (234 of whom had been treated with ICIs) and 139 with other cancer types. Importantly, DNA-sequencing data for these samples were obtained from paired tumour/non-tumour assays and from tumour-only assays. For all of the studies assessed, Nassar and colleagues used ancestry information that was predicted on the basis of genetic analysis.

Nassar *et al.* found that tumour-only assays substantially overestimate TMB compared with the gold-standard assay results based on paired tumour/non-tumour sequencing. This overestimation is particularly high in people of non-European ancestry (African American or Asian in this study) because of incorrect assignment of germline mutations as tumour mutations.

DNA sequencing identifies both somatic and germline mutations in the genome of tumour samples. In a paired tumour/non-tumour sequencing panel, germline mutations can be excluded by removing the non-tumour mutational variation identified in the matched non-tumour samples. However, such information is not available from a tumour-only panel. For mutations identified from tumour-only sequencing panels, databases of genetic variation can be consulted to filter out germline mutations. However, this filtering step works best for samples from individuals of European ancestry. This is because reference genomes tend to have a European-descent bias in terms of the genetic variation represented.

The authors found that tumours from people of non-European ancestry tended to have a greater number of genetic variants seen at a relatively high frequency than did tumours from individuals of European descent. When TMB estimates from tumour-only and paired tumour/non-tumour panels were compared, people of non-European ancestry had a significantly higher TMB difference than did those of European descent.

To address this issue, Nassar and colleagues developed an approach called a regression-based analysis to recalibrate TMBs derived from tumour-only assays in an ancestry-specific manner. Ancestry-specific calibration measurements for European and non-European individuals were estimated separately by modelling the relationship between the TMBs from tumour-only assays and those from paired tumour/non-tumour assays. For tumour/non-tumour pairs, the authors used a data resource for sequenced protein-coding regions of the human genome (the TCGA whole-exome sequencing data). This enabled Nassar *et al.* to simulate the TMBs from multigene sequencing panels with or without matched non-tumour samples. The resulting recalibrated TMB was called the TMB-c.

TMB recalibration had a notable effect on decisions that might have been made concerning ICI-based immunotherapy. When a tumour-only panel was used, the authors' analysis indicates that 43.6% of African American and 37% of Asian individuals in the DCFI group had their tumours erroneously classified as TMB-high and thus (incorrectly) eligible for treatment. This is significantly higher than the corresponding erroneous value for people of European ancestry (21%). After recalibration, the authors estimate that the proportion of African American and Asian individuals falsely classified as having high-TMB tumours would fall to 23% and 17%, respectively.

Importantly, TMB-c provides a better overall prediction of survival time for people treated with ICIs than do non-calibrated TMBs. Specifically, individuals with NSCLC in the MSKCC group were separated into three groups: true TMB-low (TMB less than 10, TMB-c less than 10), true TMB-high (TMB more than 10, TMB-c more than 10) and false TMB-high (TMB more than 10, TMB-c less than 10). Individuals in the true TMB-high group had a significantly longer survival time after ICI treatment than did those in the other two groups, whereas people in the false TMB-high and the true TMB-low groups had a similarly worse prognosis. However, in both the DFCI and MSKCC groups, a higher TMB-c was associated with a longer overall survival time and time to ICI failure in individuals of European than non-European ancestry, after adjusting for factors such as previous therapy, type of ICI and cancer classification by histological subtype. Neither TMB-c nor TMB was significantly associated with survival outcome in people of Asian or African ancestry, presumably as a result of the relatively small number of samples for these two populations.

Nassar and colleagues' study shows that tumour-only sequencing panels tend to overestimate TMBs, especially those in tumours of people of non-European ancestry. To obtain a TMB estimate that would reliably enable a decision about an individual's eligibility for ICI treatment, a paired tumour/non-tumour sequencing panel would be recommended. Alternatively, when a tumour-only assay must be used, the resultant TMB estimate should be recalibrated in an ancestry-specific manner. The authors' findings indicate the need for improved approaches to identifying genetic variants[4], so that more-diverse reference panels of genomic information can be combined with more-comprehensive assessments of germline variants for people of non-European ancestry.

Chao Cheng and **Christopher I. Amos** are in the Department of Medicine, Baylor College

of Medicine, Houston, Texas 77030, USA, and at the Dan L. Duncan Comprehensive Cancer Center and the Institute for Clinical and Translational Research, Baylor College of Medicine, Houston.
e-mails: chao.cheng@bcm.edu; chris.amos@bcm.edu

1. Haslam, A. & Prasad. V. *JAMA Netw. Open* **2**, e192535 (2019).
2. Nassar, A. H. *et al. Cancer Cell* **40**, 1161–1172 (2022).
3. Cristescu, R. *et al. J. Immunother Cancer* **10**, e003091 (2022).
4. Aganezov, S. *et al. Science* **376**, eab13533 (2022).

The authors declare no competing interests.

Astronomy

Magnetic field lights up a stellar graveyard

Samar Safi-Harb

An X-ray imaging mission has unveiled the magnetic field in the environment of a dead star. The order and symmetry of the field will reshape our understanding of how it accelerates particles to ultra-high energies. **See p.658**

When massive stars die in dramatic supernovae, the beautiful stellar debris lingers in space for tens of thousands of years. The core of the dead star becomes a pulsar – a rapidly rotating, highly magnetized pulsating star the size of a city, which spews out a wind of fast-moving, electrically charged particles. This pulsar wind carries a strong magnetic field, typically shaped like a doughnut, that generates a gas cloud called a pulsar wind nebula as it slams into the surrounding material[1], accelerating the particles to energies much higher than can be achieved on Earth. On page 658, Xie *et al.*[2] report imaging of the magnetic field in the acceleration zone of the Vela pulsar wind nebula, revealing a surprisingly highly ordered and symmetrical structure. This could help to shed light on why these nebulae are among the most powerful particle accelerators in the Universe[3].

Astronomers have been exploring pulsars ever since their discovery 55 years ago[4]. In the past two decades, NASA's Chandra X-ray Observatory has enabled close-up views of their nebulae[5], but direct imaging of the magnetic fields responsible for their high-energy radiation has not been possible. This is largely because magnetic fields are invisible – to visualize them, astronomers use the properties of light emitted when particles are accelerated, which occurs through a process known as synchrotron emission.

This type of radiation is generated when magnetic fields bend the paths of particles that are relativistic (meaning they are travelling at speeds close to that of light), and it can be emitted with wavelengths ranging from radio waves to X-rays. The electromagnetic field of the photons emitted by synchrotron radiation is expected to vibrate in one direction, which is perpendicular to that of the nebula's magnetic field. So, measuring the direction of this vibration – the 'polarization' of the emission – can reveal which way the magnetic field is pointing (see go.nature.com/3ykafb7).

The onset of a pulsar wind nebula is referred to as the termination shock, and it resembles an arc of X-rays, which typically appears a short distance away from the pulsar (Fig. 1). Astronomers think that the particle energies get boosted precisely at the termination shock, although the mechanism through which this occurs is a topic of intense debate. Imaging the magnetic field driving the acceleration is crucial for solving this puzzle.

Xie *et al.* detected X-ray polarization from the Vela pulsar wind nebula using data from NASA's Imaging X-Ray Polarimetry Explorer (IXPE) mission, which was launched in December 2021. Astronomers had previously imaged magnetic fields in astronomical sources using radio and optical telescopes, but these types of radiation are subject to rotation effects and absorption by dust, respectively, both of which hinder observation. By contrast, X-ray polarization offers a direct probe of the geometry of the magnetic field in the acceleration zone, where the pulsar dumps its particle wind and emits high-energy radiation.

Before the IXPE mission, the only X-ray polarization measurement was that of the Crab nebula, first reported in 1976 (ref. 6). This nebula is perhaps the most iconic cosmic object, the supernova of which was witnessed from Earth in the year 1054. The Vela nebula is much older than the Crab, but looks like its identical twin when viewed with Chandra's sharp X-ray eyes. At a distance of around 300 parsecs away[7], Vela is also much closer to Earth than is the Crab, which makes it possible to obtain a close-up view. Vela formed when a massive star exploded sometime between 10,000 and 20,000 years ago, leaving behind a supernova remnant spanning around 16 full Moons across Earth's sky.

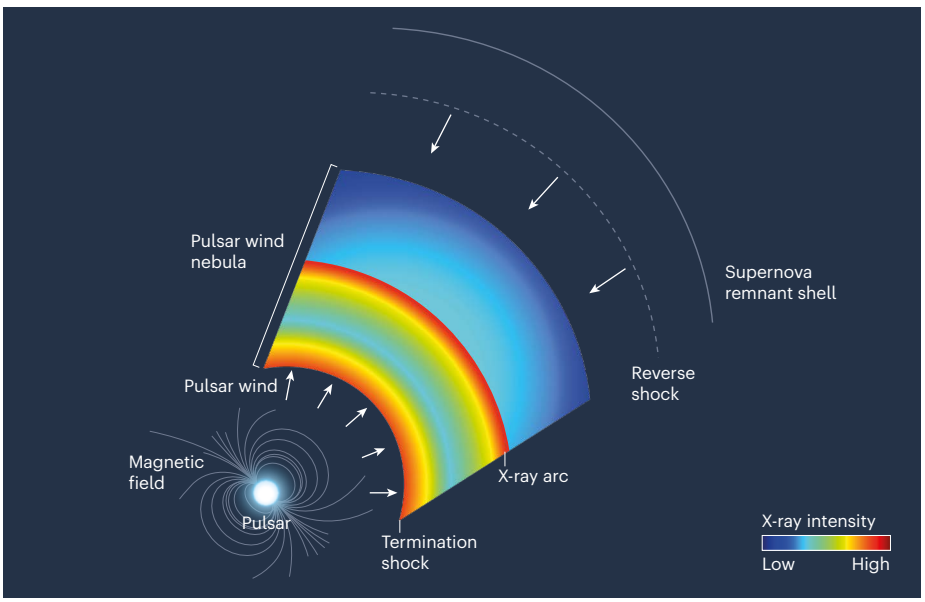

Figure 1 | The structure of a pulsar wind nebula in a supernova remnant. In the remnant of a supernova, the core of the dead star (a pulsar) emits a particle wind that gyrates within a magnetic field that is typically doughnut-shaped, and accelerates the particles to ultra-high energies, generating a gas cloud called a pulsar wind nebula. The onset of the pulsar wind nebula is known as the termination shock, and the nebula is also subjected to a reverse shock in the opposite direction as it ages. The magnetic field causes particles to emit radiation, and emission with X-ray wavelengths is visible as arcs that can indirectly reveal the shape of the field. Xie *et al.*[2] used X-ray emission data from the Vela pulsar wind nebula to infer the geometry of its magnetic field, which is a symmetrical and surprisingly highly ordered structure. Graphic not to scale.