TECHNOLOGY FEATURE

CRYO-ELECTRON MICROSCOPY SHAPES UP

As the imaging technique produces ever-sharper protein structures, researchers are racing to develop tools to assess how accurate they are.

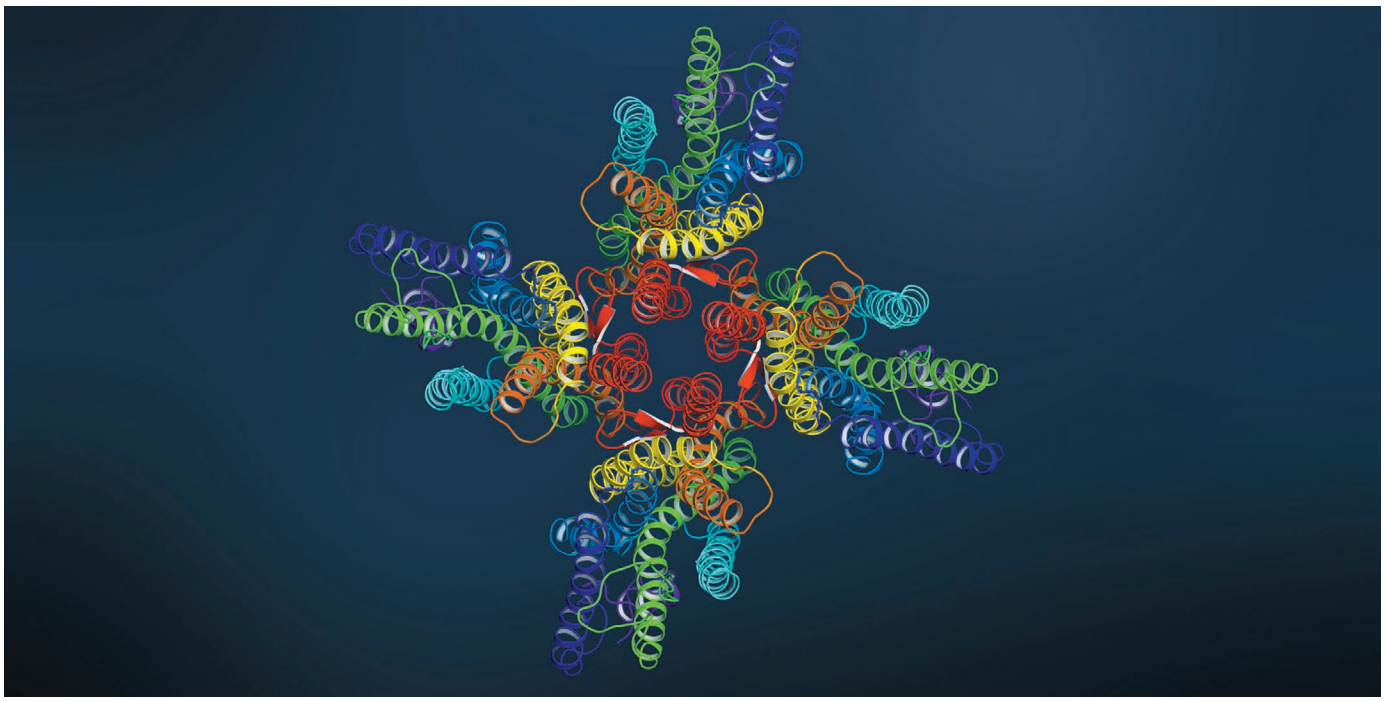

LABORATORY OF NEUROPHYSIOLOGY AND BEHAVIOR/THE ROCKEFELLER UNIVERSITY

Structures of ion channels, such as this insect smell receptor, are tough targets for X-ray crystallography. Cryo-electron microscopy has revealed this and more.

BY MONYA BAKER

Once derided as 'blobology' for its blurry images, cryo-electron microscopy (cryo-EM) is now churning out high-resolution structures of everything from virus particles to enzymes. The number of cryo-EM images uploaded to the Electron Microscopy Data Bank (EMDB) has boomed from just 8 in 2002 to 1,106 last year — the same year the technique won its developers the Nobel Prize in Chemistry.

The quality of cryo-EM images now rivals that of X-ray crystallography, long the dominant technique for solving protein structures. The technique has also succeeded where crystallography has struggled: showing, for instance, how temperature-sensitive ion channels work, characterizing pathological proteins in neurodegenerative disease and detailing how viruses can interact with antibodies[1]. Consequently, many veteran crystallographers are giving up on crystals and freezing proteins for cryo-EM instead.

Publications of cryo-EM structures are coming in fast (see 'Widening the bottleneck'). But, some researchers worry, not everyone knows how to evaluate them, and some are calling for new practices and tools to help them do so.

CRITERIA CREATION

The surge in cryo-EM is largely a result of better electron detectors and image-processing techniques, says Richard Henderson of the Medical Research Council Laboratory of Molecular Biology in Cambridge, UK, who shared Nobel prize. But, he says, the field still lacks the kind of standardized tools for producing robust structural models that crystallographers developed as their field matured. "This has led to a lot of sloppiness," he says. "What is needed now are better criteria to encourage researchers to put more work into their model-building."

Instead, the race is on to publish structures with ever-better resolution, and that is discouraging careful work, says Holger Stark, an electron microscopist at the Max Planck Institute for Biophysical Chemistry in Göttingen, Germany. Some published structures depict atom-level precision without acknowledging that certain regions of the structure are "fantasy", with scant data to back up any particular interpretation, he says. "It's just noise in areas where people have put in atomic coordinates."

There is no question that cryo-EM has enabled fantastic discoveries and that many structures are solid, says Gabriel Lander, a structural biologist at the Scripps Research Institute in La Jolla, California. But he cautions that many researchers are too quick to assume that all the details in the structure are correct. As a result, someone who uses a structure to design mutant versions of a protein to understand its mechanism, or who sees a ligand binding in a poorly defined spot, could end up doing ▶

months of failed experiments, he explains. "I don't want the reputation of cryo-EM sullied by over-interpretation."

RESOLUTION FIXATION

Protein structures are often judged by a single factor: resolution, the level of detail a structure shows. That metric is straightforward to ascertain in crystallography, but not in cryo-EM.

In crystallography, a highly ordered lattice of tightly packed molecules is rotated through an X-ray beam, and the resolution of the resulting image can be calculated directly from the diffraction patterns made by the deflected photons. Those patterns are then transformed into 'maps' of electron density, which researchers combine with the protein sequence to build a model. The model represents how specific chemical building blocks of a protein fold into sheets and helixes (visible at a resolution of around 5 ångstroms), and how side chains of amino acids are positioned (which start to become visible around 3.5 Å). Big, floppy objects tend not to form ordered crystals, so as a rule, the smaller and more rigid the protein, the more amenable it is to crystallography.

In cryo-EM, proteins and other macromolecular complexes are flash-frozen in a thin layer of water, ideally not much thicker than the protein itself. Irradiating that layer with low-energy electrons produces 2D images of individual particles on the detector — fuzzy shadows cast from scattered electrons (see 'Modelling in ice'). Thousands or even hundreds of thousands of these noisy images are then computationally sorted and reconstructed to create a 3D map. Finally, other types of software fit the protein sequence into the map to create a model. The smaller the object, the noisier the images, so cryo-EM tends to work best for larger structures.

To avoid mistaking noise for signal, researchers typically split particles into two subsets and build 'half maps' from each. The correlation between those two maps is used to calculate resolution — but it's an imperfect proxy, says Edward Egelman, a structural biologist at the University of Virginia in Charlottesville. "That's not measuring resolution, per se, it's measuring consistency." And the resulting values, he says, must be taken with a grain of salt. Indeed, he says that the race to claim high resolution has sometimes led researchers to "silliness" — such as reporting resolution to a hundredth or even a thousandth of an ångstrom, a level of precision that makes no sense with cryo-EM.

Also, not all false signals are random noise. Egelman has demonstrated that systematic artefacts (such as computationally adding non-existent cylinders into both half maps) can drastically (and erroneously) improve the apparent resolution of a structure[2].

Sometimes researchers actually back-compute an electron-density map from the structural model that created it, and then revisit their data to select particles that are most likely to confirm the model. "It's a kind of bias," says crystallographer Piotr Neumann at the University of Göttingen in Germany. "This kind of cheating is not acceptable, but it's okayish." Another, more common, technique is to create a 'mask' of the expected overall shape of the protein and use that to exclude portions of images. Done judiciously, this boosts the signal-to-noise ratio; done aggressively, it shoehorns or 'overfits' data.

TWEAKED TO FIT

Structural biologists joke that there are many more structures published with resolutions of 2.9 Å than of 3.0 Å — an apparent symptom of over-aggressive analyses. But even without gaming, describing a protein with a single number is problematic, says Lander. It obscures the fact that the quality of a cryo-EM map varies dramatically, with the poorest-quality fit often occurring in the most flexible and biologically interesting areas of the protein. "There is no one metric that is good," says Neumann. "All metrics can be biased or not fully reliable. So, we need to use many simultaneously."

Earlier this year, Neumann and his colleagues set out to document how well protein structure models in the Protein Data Bank fit the corresponding maps in the EMDB. They found only low or moderate agreement for more than three-quarters of the 565 structures examined, suggesting that large swathes of the models should be viewed with scepticism[3].

Some drug developers, at least, are approaching the models with caution. Christian Wiesmann, head of the cryo-EM team at the Novartis Institutes for Biomedical Research in Basel, Switzerland, says that when looking at models of proteins bound to small molecules, he typically downloads maps from the EMDB, assesses how other researchers nestled the compounds into the protein and then uses his own judgement. More than once, Wiesmann says, he would have made different calls — differences that could affect drug design.

Not every researcher possesses that level of structural sophistication. But even if they did, maps can be hard to vet. Authors must deposit their maps in the EMDB when publishing papers, but these deposits are often insufficiently annotated, says Alex Wlodawar, a structural biologist at the US National Cancer Institute in Frederick, Maryland, who has compared crystal and cryo-EM structures at high resolutions and found that the latter are often "optimistic"[4]. Researchers might deposit the raw map without the refined or 'sharpened'

MODELLING IN ICE

In cryo-EM, thousands of raw electron-microscopy images are collected and computationally analysed to build up a density map that reflects the shape of the protein.

RAW IMAGE

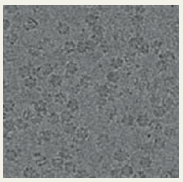

Where to share data
Electron Microscopy Public Image Archive (EMPIAR)

MAP

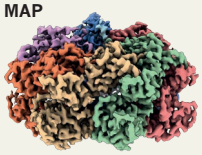

Electron Microscopy Data Bank (EMDB)

This is then combined with the known protein sequence to create a final model showing the placement of atomic groups.

MODEL

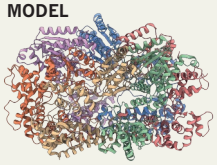

Protein Data Bank (PDB)

SOURCE: C. PUCHADES *ET AL. SCIENCE* **358**, EAAO0464 (2017)

Widening the bottleneck

Better validation tools and practices are not the only thing holding back cryo-electron microscopy; many researchers are unable to produce images of sufficient quality to even start the process.

"The real bottleneck is specimen prep," says Bridget Carragher, electron microscopy co-director at the New York Structural Biology Center. But a technique she has co-developed might substantially broaden that bottleneck.

Carragher and her co-director Clint Potter developed robots that have droplet-sensing cameras and piezoelectric devices similar to those found in an inkjet printer. The robots can apply small volumes of sample onto a grid just as another robot plunges the sample swiftly and evenly into cryogen, creating a thin, uniform, frozen layer that is perfect for imaging. The grid itself is covered in nanowires that wick away excess solution in a more controlled way than the blotting paper typically used. That also reduces the chance of proteins getting stuck at the air–water interface, which can cause them to denature or adopt preferred orientations rather than the random orientations that are necessary for them to be seen from every angle.

The system, called Spotiton, allows users to prepare more samples more quickly and using less protein, and ensures that a greater fraction will be usable. The pair have licensed Spotiton to TTP Labtech in Melbourn, UK, which plans to commercialize the system in the next year or so under the name Chameleon. M.B.

map used to build the model, or without reporting whether they used a mask in building it. And very few deposit the half maps used to validate their analysis.

MAPPING THE FUTURE

Like models, maps are highly variable in quality, says Ardan Patwardhan, who manages the EMDB. Suites of automated and semi-automated tools have been created to help researchers turn 2D cryo-EM images into 3D maps. To help assess these workflows, the EMDB has run several validation competitions. It found that the greatest variability came not from the software packages, but from the experience level of the users. Less-experienced groups used default parameters; the best teams tailored settings to the data they had. That can make the difference between clearly visible side chains and blurry secondary structures, even when starting from the same raw images[5].

Today, researchers are calling for better methods for validating cryo-EM maps and models[6] — and raw image data could help. In 2014, Patwardhan and his colleagues at the European Bioinformatics Institute (EBI) in Cambridge, UK, created the Electron Microscopy Public Image Archive. The largest of the current 175 deposits of raw image data is more than 12 terabytes, which takes about 5 days to download.

Better methods for representing uncertainty could also help. Lander has proposed that researchers provide a spectrum of models[7] to better illustrate the range of structures that might fit the data. Maya Topf, a computational structural biologist at Birkbeck, University of London, has helped to create software called TEMPy that measures the quality of the model at the scale of amino acids rather than of the entire structure. Although this is not yet mandatory, the research community is starting to expect these kinds of evaluations, she says[8]. "The awareness is growing. More and more people are reporting in papers the local resolution."

Still, cryo-EM has a long way to go to match practices of crystallography. "The fact that data and models need to be validated has to become ingrained in people's minds, especially as the field attracts many new practitioners who don't have decades of experience," says Gerard Kleywegt, a structural biologist at the EBI. And, of course, some things are fundamentally different: crystallography captures proteins in rigid conformations, whereas cryo-EM can show more natural, and naturally ambiguous, conformations for which people are still developing the language to describe. Improvements will require better methods, greater consensus and better practices — all of which take time to develop. A validation task force met in September 2010 to develop recommendations, Kleywegt notes. "The field has evolved so rapidly since then that a follow-up meeting is overdue." Planning for a 2019 meeting is already under way. ■

Monya Baker *is an editor at* Nature.

1. Renaud, J. P. *et al. Nature Rev. Drug Discov.* **17**, 471–492 (2018).
2. Subramaniam, S., Earl, L. A., Falconieri, V., Milne, J. L. S. & Egelman, E. H. *Curr. Opin. Struct. Biol.* **41**, 194–202 (2016).
3. Neumann, P., Dickmanns, A. & Ficner, R. *Structure* **26**, 785–795.e4 (2018).
4. Wlodawer, A., Li, M. & Dauter, Z. *Structure* **25**, 1589–1597.e1 (2017).
5. Heymann, J. B. *et al. J. Struct. Biol.* https://doi.org/10.1016/j.jsb.2018.08.010 (2018).
6. Afonine, P. V. *et al. Acta Cryst. D* **74**, 814–840 (2018).
7. Herzik, M. A. Jr, Fraser, J. & Lander, G. C. Preprint at bioRxiv https://doi.org/10.1101/128561 (2017).
8. Joseph, A. P., Lagerstedt, I., Patwardhan, A., Topf, M. & Winn, M. *J. Struct. Biol.* **199**, 12–26 (2017).

Big data fresh from the sea

Machine learning helps marine biologists to churn through millions of plankton images.

BY JEFFREY M. PERKEL

When they think about big data, most researchers probably imagine genomics, neuroscience or particle physics. Kelly Robinson's data challenge involves plankton.

"A lot of things that we enjoy seafood-wise — from fish to oysters to mussels to shrimp — almost everything starts their lives as plankton," says Robinson, who studies marine ecosystems at the University of Louisiana at Lafayette. In photographs, they look like floating specks of dust, and her research involves quantifying and mapping their distribution and predator–prey interactions. The problem is, she must do so in millions upon millions of images.

Robinson collects data by towing a remote-camera platform called ISIIS — the *In Situ* Ichthyoplankton Imaging System — behind a boat. ISIIS captures about 80 photos per second, or 288,000 images (660 gigabytes) per hour. For one project in the Straits of Florida, when Robinson was a postdoc, she generated 340 million pictures; a colleague working in the Gulf of Mexico generated billions.

"You start to learn about things that you never thought you would learn," Robinson says, "like the number of files that you can store on an individual computer. It's 30 million, by the way, on your regular PC." On her most recent cruise, Robinson sailed with 52 2-terabyte hard drives, which a student had to monitor and replace as they filled up. Someone then must get that collection to the university, convert the files to Linux formatting, and upload them to a server — a process that takes 24 hours per drive.

Kelly Robinson and her team scrutinize images.

The team uses machine-learning software to automatically pick out and identify objects in the images. But the algorithms must be taught what to look for — this is a starfish, that is a prawn. Such features are relatively rare in the water, so finding pictures for the training set takes time. Over two months, Robinson and her team manually sorted through 2 million images to find enough to feed the algorithm. "It's a little mind-numbing, but if you're under the gun you can do it," she says.

Naturally, the team is looking to optimize the process. Working with colleagues at Oregon State University in Corvallis, where she was a postdoc, Robinson is testing whether she could accelerate her work by processing the images on multiple video card graphical processing units (GPUs) running in parallel. She is also looking into cloud computing as an alternative to Earth-bound clusters.

But infrastructure goes only so far; what the team really needs, she says, is more people to crunch the numbers. Unfortunately, data scientists are in high demand, and industry jobs are lucrative. "We have a lot of turnover," she says. ■

Jeffrey M. Perkel *is technology editor at* Nature.

CORRECTION

The referencing in the Technology feature 'How to teach an old sequencer new tricks' (*Nature* **559**, 643–645; 2018) was incorrect. The correct version can be found online at at go.nature.com/2wmnhgc.