

and energy. The previous reconstruction method was based on a computationally expensive set of Monte Carlo simulations, but the neural network approximates those simulations more efficiently. The new data set of 60 000 neutrino events is 20 times as large as previous sets. Francis Halzen, a theoretical physicist who is a member of the IceCube Collaboration, says that “given the new machine-learning techniques, it almost looked easy in retrospect, especially given our unsuccessful efforts to see our galaxy over the last decade.”

With those improvements, the collaboration found a signal of diffuse high-energy neutrinos from the galactic plane. Figure 2 shows the Milky Way as it’s typically been seen through radio, optical, and gamma-ray emissions. Those three wavelength regimes show a clear, bright galactic center flanked on either side by a thin plane of more diffuse emission. The Milky Way’s neutrino signal is about 10% of the total flux at 30 TeV, and although it may not be immediately clear,

the signal is consistent with the gamma-ray emission. In fact, the IceCube Collaboration found that the chance for the neutrino signal to arise randomly from background noise is less than 1 in 100 000, or a statistical confidence level of 4.5 standard deviations.

Guide to the galaxy

The neutrino picture of the Milky Way so far shows a diffuse pattern of emission that arises when cosmic rays interact with interstellar gas. But neutrinos could be coming from galactic point sources of cosmic rays. To figure that out, the team will need to review more data. Similar to how the atmospheric background overwhelms the astrophysical signal, the suspected point-source neutrino signal may be masked by the diffuse neutrino signal shown in figure 2.

Astronomers already have some ideas for where galactic cosmic rays, and by extension neutrinos, could come from. Supermassive black holes could produce them (see PHYSICS TODAY, August 2022,

page 14), but the one at the center of the Milky Way may not be active enough to generate such high-energy particles.

Neutrinos could also be sourced from active galactic nuclei (see “IceCube pinpoints an extragalactic neutrino source,” PHYSICS TODAY online, 12 July 2018), and some have already been spotted in an extragalactic supernova (see “A supernova for the ages, 30 years later,” PHYSICS TODAY online, 23 February 2017). In the Milky Way, says Halzen, “finding the sources are our next priority, and we are on it.”

Alex Lopatka

References

1. IceCube Collaboration, *Science* **380**, 1338 (2023).
2. M. Ackermann et al., *Astrophys. J.* **750**, 3 (2012).
3. IceCube Collaboration, *Science* **342**, 947 (2013).
4. M. Hünnefeld (on behalf of the IceCube Collaboration), *Proc. Sci.* (2017), <https://pos.sissa.it/301/1057/pdf>.

Toward faster, safer proton therapy

Most facilities for zapping tumors with protons are extremely inefficient. But perhaps they don’t have to be.

In the fight against cancer, the global medical community is placing an increasingly big bet on protons. Of the just over 100 proton-therapy facilities worldwide, more than half began operation in 2016 or later. Currently they treat some 50 000 patients per year, with a cumulative total of around 300 000.

The treatments are eye-wateringly expensive. A new facility can cost more than \$200 million just to build, not counting the cost of upkeep, operation, and doctors’ time. Because of the high price tag, plenty of critics are putting pressure on proton-therapy proponents to justify the expense with results. (See, for example, PHYSICS TODAY, October 2015, page 8.) For some applications—such as treating eye tumors, as shown in figure 1—proton therapy has shown clear advantages over other treatments. For others, achieving its potential is still a work in progress.

At the Paul Scherrer Institute in Switzerland, home to one of the oldest operating proton-therapy centers, a team of researchers led by Vivek Maradia and

his PhD adviser, Serena Psoroulas, is working to lessen a major source of inefficiency: In a typical treatment facility that accelerates its protons with a cyclotron, between 70% and 99.9% of the accelerated protons get thrown away.¹

Wasted protons don’t directly translate into wasted dollars. But getting protons more efficiently into patients could lead to faster treatments, happier patients, and eventually more economical facilities. And it could allow proton therapy to better benefit from a powerful but counterintuitive phenomenon called the FLASH effect: Delivering a lot of radiation to a tumor all at once could lessen the side effects on the surrounding healthy tissue.

Low transmission

Proton therapy’s appeal stems from the physics of how protons interact with matter. The goal of radiotherapy—and indeed of all forms of cancer treatment—is to kill cancer cells while sparing healthy ones. But a beam of radiation (of any

form) sent into the body encounters not just the tumor but also the healthy organs in front of and behind it. X rays, still the tool for most radiotherapy, deposit their energy everywhere along their path, in tumor and healthy cells alike.

Protons, on the other hand, have scattering cross sections that depend inversely on their kinetic energy. So a proton passing through the body slows down little by little, until it finally leaves most of its energy—and does most of its damage—right before it stops. By controlling the proton beam’s direction and initial kinetic energy, clinicians can position the radiation-affected region in all three dimensions.

In practical terms, however, controlling proton kinetic energies is not so easy. Clinicians need proton energies ranging from 60 MeV (for tumors close to the body surface, such as those in the eyes) to 230 MeV (for tumors tens of centimeters deep in the body). But a single cyclotron produces proton beams at just one energy. It would be far too expensive to have a separate cyclotron for every possible proton energy. A few facilities have opted for low-energy cyclotrons for

treating eye tumors, with the consequence that they can't treat anything else. Most proton-therapy centers, however, use cyclotrons at the top of the energy range, at either 230 MeV or 250 MeV. And that's where the inefficiency comes in.

One can turn high-energy protons into lower-energy protons by passing them through a chunk of solid material, usually carbon. But that energy-degradation process also turns a monoenergetic proton beam into one with a considerable energy spread—no longer suitable for clinical use, because the protons' localized depositions of energy, known as Bragg peaks, are no longer all in the same place. The standard approach is to use a dipole magnet to disperse the protons by energy and then pass them through a slit to select protons with as close to a single energy as possible. Most of the protons, as a result, are thrown away.

The waste is worst at low clinical energies. The more the cyclotron protons need to be degraded, the larger their energy spread, and the lower the fraction transmitted through the energy-selection slit. For target energies greater than 200 MeV, perhaps 10% or more of the initial protons can be salvaged. But for target energies less than 100 MeV, less than 1% can.

The low transmission makes it hard to treat eye tumors at facilities without low-energy cyclotrons. Delivering a radiation dose takes about a minute, which may not sound like much. But patients need to be kept from blinking or moving their eyes for that time, which is challenging and uncomfortable.

For tumors in parts of the body such as the lungs and abdomen, which inevitably move around despite the patient's best efforts, treatment can take even longer—up to 45 minutes—because clinicians need to continually rescan the patient's body to track the tumor's position. Again, the patient needs to stay as still as possible for the duration of the procedure.

Momentum cooling

Are the low transmission efficiencies and long treatment times an inherent limitation of cyclotron-based proton therapy? Much of the community thought so, says Maradia. "For years, it's been widely believed that there was no feasible way to enhance transmission," he says. "But Serena Psoroulas challenged that notion, and she conceived the idea for my PhD project."

For the first year of his PhD studies, Maradia tinkered with simulations of beamline ion optics, and he discovered

some new ways to wrangle more protons from the cyclotron to the patient.² In a nutshell, existing ion-optics setups treat the two dimensions perpendicular to the beam symmetrically, and they apply the same focusing and defocusing forces in both directions. But the dimensions aren't symmetrical—in part, because the protons are dispersed by energy in one direction but not the other. Maradia and colleagues predicted that by accounting for that asymmetry, they could improve transmission by up to a factor of six.

There remained the greatest source of inefficiency: the protons discarded at the energy-selection slit. The solution, it turned out, was deceptively simple. The protons were already dispersed by energy, and their momentum can be slowed by passing them through solid material. So Maradia proposed sticking a wedge into the beam, as shown in figure 2. The fastest protons pass through the thickest part of the wedge and are slowed most; the slowest pass through the thinnest part and are slowed least.

Maradia came up with the momentum-cooling idea on his own, but he noticed afterward that wedge-shaped absorbers had been used before in other areas of particle physics, such as muon experiments.³ They'd not been considered before for proton therapy, perhaps



FIGURE 1. EYE CANCER is rare, but it's extremely difficult to treat. At the OPTIS2 beamline at the Paul Scherrer Institute in Switzerland, where eye tumors are treated with proton-beam therapy, patients must keep their eyes open and immobile for the duration of the minute-long radiation delivery. The treatment time could be shortened by increasing the fraction of protons transmitted through the beamline. (Courtesy of the Paul Scherrer Institute.)

because when the protons scatter off the wedge, their momentum spread perpendicular to the beam increases. But Maradia and colleagues' improved ion optics were equipped to handle the increased spread.

Proposing and simulating improvements is one thing; actually implementing them can be quite another, especially in an active medical facility. "No one wanted to disrupt the ongoing clinical treatments," says Maradia. But with persistence, he eventually got permission to try out his wedge on the Paul Scherrer Institute's eye-treatment beamline.

The results were positive but modest: From an initial fraction of 0.27% of protons, the wedge increased transmission almost twofold, to 0.5%. Why such a small improvement? The beamline as a whole was designed on the basis of the assumption that only protons with one specific energy would ever make it through to the patient. After being dispersed by the dipole magnet, most of them crash into the beamline walls before they even reach the wedge.

If the dipole magnet deflected the protons at a shallower angle, the loss could be mitigated, and the researchers estimate that transmission at the lowest energies could be boosted to perhaps 7%. Making such a change to an existing beamline is probably not feasible. "However, it would be relatively easy to incorporate momentum cooling into future proton-therapy centers during their design and construction," says Maradia. Several dozen new proton-therapy facilities are currently in development around the world.

FLASH therapy

Increasing the fraction of protons that make it through the beamline has more implications than just reducing treatment times proportionally. For example, if proton treatment of a lung tumor could be sped up so much that the entire radiation dose is delivered while patients hold their breath, clinicians might no longer need to take elaborate steps—and employ expensive equipment—to track the motion of the tumor during treatment.

Alternatively, if the treatment times that are available today are considered acceptable, they could be achieved with much smaller and less powerful cyclotrons. A more modest cyclotron produces less radiation overall, so it requires less




FIGURE 2. WHEN PROTONS from a cyclotron are tuned to the energies required for proton-beam therapy, they end up with a momentum spread $p \pm \Delta p$ that's too wide for clinical use. The conventional approach is to use a dipole magnet (dark blue) to disperse the protons by momentum and then pass them through a slit to select out the protons with just one momentum value, p . But if, instead, a wedge-shaped absorber is used to cool all the protons' momentum to $p - \Delta p$, more protons can make it through the beamline and into the patient. (Adapted from ref. 1.)

concrete shielding and could be built at less cost.

But perhaps the most intriguing potential implication concerns the FLASH effect. Proton therapy's appeal is that most of the proton beam's energy is deposited at the Bragg peak. But most is not all, and proton-therapy clinicians have to go to great lengths to design treatments that avoid harming healthy tissues, especially when tumors lie close to critical organs or arteries. (See the article by Jeremy Polf and Katia Parodi, *PHYSICS TODAY*, October 2015, page 28.)

So when, in 2014, experiments started to show⁴ that if radiation is delivered very quickly, it does less harm to healthy tissues—despite being just as effective at killing the tumor—the radiotherapy community was captivated.

FLASH radiotherapy is still far from ready for clinical use, and much remains unknown. For example, researchers still don't know how the effect works—and not every experiment even agrees that it does. One popular hypothesis is that the fast delivery of radiation induces a temporary oxygen deficiency in healthy tissue, which protects it from damage because radiation works by creating oxygen radicals. The tumor, on the other hand, is already starved of oxygen, so it doesn't become more oxygen-deficient than it already is. But much more study is needed to see if that picture holds up.

It's also not known exactly how fast radiation must be delivered to produce the FLASH effect, but a rough consensus is that it needs to be several orders of magnitude faster than current treatments allow. That is, instead of lasting

minutes, delivery should take a fraction of a second.

The FLASH effect appears to be equally applicable to all forms of radiation: protons, x rays, electrons, and carbon ions. Out of all clinical radiation sources, proton-accelerating cyclotrons are the closest to being able to achieve FLASH intensities. But the catch is, they can do so only with the full-strength high-energy beam straight out of the cyclotron—which means forgoing all the advantages of the Bragg peak and its tunability.⁵

In their simulations, Maradia, Psoroulas, and colleagues estimate that with a beamline optimized for their momentum-cooling approach, they could reach FLASH intensities across the entire range of clinically relevant proton energies—as long as the beam is focused to a small enough spot. For tumors more than a few millimeters in diameter, however, the FLASH beam would need to be scanned over the tumor volume more rapidly than is currently possible.

Johanna Miller

References

1. V. Maradia et al., *Nat. Phys.* (2023), doi:10.1038/s41567-023-02115-2.
2. V. Maradia et al., *Med. Phys.* **48**, 7613 (2021); V. Maradia et al., *Med. Phys.* **49**, 2183 (2022).
3. M. Bogomilov et al. (MICE collaboration), *Nature* **578**, 53 (2020).
4. V. Favard et al., *Sci. Transl. Med.* **6**, 245ra93 (2014).
5. S. Jolly et al., *Phys. Medica* **78**, 71 (2020); M.-C. Vozenin, J. Bourhis, M. Durante, *Nat. Rev. Clin. Oncol.* **19**, 791 (2022). **PT**