From first lines of code to peer-reviewed science: what high schoolers can build with AI research

Ricardo A. Gonzales

When I began teaching at Veritas AI, I first worked in the AI Scholars program, helping students build a foundation in Python, data science, and core machine learning. Later, I designed and taught AI + Medicine, which moves quickly from those fundamentals into clinical datasets, medical imaging, and model evaluation. These courses give students shared tools and a common starting point. From there, some students choose to continue their work in a more individualized research mentorship. In that setting, the student takes ownership of the question, and we work together to develop the methods, run experiments, and understand what the results mean. It is simply a different mode of learning: more open-ended, more iterative, and shaped by the student's curiosity.

In the fellowship program, students learn by doing. They explore datasets that are noisy and imperfect. They test baselines, refine architectures, troubleshoot errors that make no sense at first, and get comfortable asking, *"Why did the model behave that way?"* Research becomes less about performing correctness and more about learning how to think. The reward is not just a result: it is the moment the student realizes that they are capable of building knowledge, not just consuming it.

Over the past two years, several of my mentees have taken their research from idea to publication. Five of these projects were published in the Journal of Emerging Investigators, each one shaped by a student's personal curiosity and their persistence in solving a problem that mattered to them. What stands out is how different these problems are: ocean sustainability, global health, oncology. The diversity reflects the students, not any instruction to pursue a particular theme.

Below are the projects that grew from that shared work.

1. Victoria Wahlig — Underwater trash identification for ocean conservation

Victoria grew up sailing and saw ocean conservation as deeply personal. She began with a simple ROV (remote operating vehicle) image classifier, then expanded her project into multi-class segmentation to differentiate marine life, vegetation, and specific debris [1]. *"Merging my love for the ocean with my coding skills,"* she wrote, *"I was able to create an artificial intelligence system to do just that."* Her model offers a realistic path for ROV-based environmental stewardship.

2. Yilou Ma — Pneumonia detection from pediatric chest X-rays

Yilou compared an MLP, a standard CNN, and an adapted convolutional model to classify pneumonia [2]. The adapted model performed the best and demonstrated how controlling capacity and regularization directly impacts clinical reliability. She wrote of how comparing models side by side showed her *"there are always multiple ways to solve the problem,"* and that learning which model fits the data is a skill earned through practice.

3. Anusuiya Bhorkar — Tuberculosis diagnosis and data diversity

Anusuiya was motivated by witnessing how delayed or uncertain diagnosis can profoundly affect patients in communities with limited healthcare access. She trained multiple CNNs on several publicly available chest X-ray datasets and found that the models trained on more varied and representative data generalized more reliably to new patient groups [3]. The insight was simple and powerful: medical AI must reflect the diversity of the people it is meant to serve. Her school highlighted[1] her research and the meaningful growth she experienced along the way.

4. Amaan Mohammed — Glioma tumor segmentation using MRI

Amaan trained two different neural network architectures to segment brain tumors in MRI scans and compared how well they captured the tumor boundaries [4]. U-Net performed better because its skip connections help retain fine details that matter in medical imaging. When he began, Amaan was still getting comfortable with Python, and the debugging could be slow and cumbersome. But step by step, he built a model that worked. The project reflects not just technical skill, but patience and steady problem-solving.

5. Arjun Ulag — Multimodal early lung cancer screening

Arjun combined imaging data with clinical risk factors to improve classification accuracy compared to image-only models [5]. His interest came from prior genomics and oncology work, and his research emphasizes early detection and precision medicine. He described how integrating modalities changed the way he thought about diagnosis, noting that medicine is rarely one-dimensional.

[1] `https://www.dominicanacademy.org/about/news/news-details/~board/news-post/post/da-student-publishes-research-on-using-ai-to-diagnose-tuberculosis`

None of these projects happened overnight. Each student encountered dead-ends and confusion, and each discovered that improvement comes from iteration rather than inspiration. Research is not about producing perfect models. It is about learning to sit with uncertainty, adjust, try again, and stay curious long enough for insight to appear. The pace of AI can seem fast from the outside, but progress at the student level is grounded in something steady: learning the fundamentals and applying them with care. Along the way, students develop habits that matter in any field. They learn to build a simple, clean baseline before chasing complexity. They learn to evaluate results honestly on data the model has not seen. They learn that every model has consequences in the real world, and that accuracy is only one part of responsibility. They learn to write clearly, explain what happened, and defend their reasoning.

Those are durable skills. Whether a student eventually becomes an engineer, a doctor, a researcher, or something entirely different, they carry forward the ability to frame a question, experiment with ideas, and trust evidence over assumption. What these projects show is that when students are given structure, mentorship, and the space to pursue a question they care about, they are capable of producing work that is thoughtful, rigorous, and meaningful. The work speaks in the models themselves, in the tables of results, in the segmentation masks aligned to clinical boundaries, and in the reflections students write when they realize how far they have come. The role of a mentor is simply to make room for that growth and to help students recognize that they are already capable of real scientific thinking. The students do the rest. And that is exactly how it should be.

References

1. Wahlig V and Gonzales RA. Enhancing marine debris identification with convolutional neural networks. Journal of Emerging Investigators 2024; 7. DOI: `10.59720/23-150`
2. Ma Y and Gonzales RA. Convolutional neural network-based analysis of pediatric chest X-ray images for pneumonia detection. Journal of Emerging Investigators 2024; 7. DOI: `10.59720/24-031`
3. Bhorkar A and Gonzales RA. Effects of data amount and variation in deep learning-based tuberculosis diagnosis in chest X-ray scans. Journal of Emerging Investigators 2025; 8. DOI: `10.59720/24-278`
4. Mohammed AA and Gonzales RA. Advancements in glioma segmentation: comparing the U-Net and DeconvNet models. Journal of Emerging Investigators 2025; 8. DOI: `10.59720/24-022`
5. Ulag AS and Gonzales RA. Lung cancer AI-based diagnosis through multi-modal integration of clinical and imaging data. Journal of Emerging Investigators 2025; 8. DOI: `10.59720/24-190`

Ricardo A. Gonzales is a Postdoctoral Research Fellow at Harvard Medical School, where he works on deep learning methods for neuroimaging. He earned his doctorate in Medical Sciences from the University of Oxford and has conducted research on AI for cardiac MRI at Oxford, Yale, and Lund University. At Veritas AI, Ricardo taught AI Scholars and led the development of the AI + Medicine curriculum while mentoring students through independent research projects. He is passionate about helping students learn by doing and supporting them as they turn their interests into rigorous, meaningful work at the intersection of AI and real-world impact.