

The New York Times

INTERNATIONAL EDITION | TUESDAY, AUGUST 20, 2019



PHOTOGRAPHS BY REBECCA CONWAY FOR THE NEW YORK TIMES

Workers at the iMerit office in Kolkata, India, learn to spot and label signs of disease in a medical scan or annotate images of stop signs in videos for use in artificial intelligence.

The tedium of teaching A.I.

BHUBANESWAR, INDIA

Making algorithms smarter is a labor-intensive task for thousands of workers

BY CADE METZ

Namita Pradhan sat at a desk in downtown Bhubaneswar, India, about 40 miles from the Bay of Bengal, staring at a video recorded in a hospital on the other side of the world.

The video showed the inside of someone's colon. Ms. Pradhan was looking for polyps, small growths in the large intestine that could lead to cancer. When she found one — they look a bit like a slimy, angry pimple — she marked it with her computer mouse and keyboard, drawing a digital circle around the tiny bulge.

She was not trained as a doctor, but she was helping to teach an artificial intelligence system that could eventually do the work of a doctor.

Ms. Pradhan was one of dozens of young Indian women and men lined up at desks on the fourth floor of a small office building. They were trained to annotate all kinds of digital images, pinpointing things like stop signs and pedestri-



Prasenjit Baidya, a manager at iMerit, and his wife, Barnali Paik, who also works at the company. Their jobs have enabled the couple to afford an apartment near the office.

ans in street scenes and factories and oil tankers in satellite photos.

Artificial intelligence, most people in the tech world would tell you, is the future of their industry, and it is improving fast, thanks to something called machine learning. But tech executives rarely discuss the labor-intensive

process that goes into its creation. A.I. is learning from humans. Lots and lots of humans.

Before an A.I. system can learn, someone has to label the data supplied to it. Humans, for example, must pinpoint the polyps. The work is vital to the creation of artificial intelligence in appli-

cations like self-driving cars, surveillance systems and automated health care.

Tech companies keep quiet about this work. And they face growing concerns from privacy activists over the large amounts of personal data they are storing and sharing with outside businesses.

Earlier this year, I negotiated a look behind the curtain that Silicon Valley's wizards rarely grant. I made a meandering trip across India and stopped at a facility across the street from the Superdome in downtown New Orleans. In all, I visited five offices where people are doing the endlessly repetitive work needed to teach A.I. systems, all run by a company called iMerit.

There were intestine surveyors like Ms. Pradhan and specialists in telling a good cough from a bad cough. There were language specialists and street scene identifiers. What is a pedestrian? Is that a double yellow line or a dotted white line? One day, a robotic car will need to know the difference.

What I saw didn't look very much like the future — or at least the automated one you might imagine. The offices could have been call centers or payment processing centers. One was a timeworn former apartment building in the middle of a low-income residential neighbor-

A.I. is learning from humans. Many humans.

A.I., FROM PAGE 1

hood in western Kolkata that teemed with pedestrians, auto rickshaws and street vendors.

In facilities like the one in Bhubaneswar and in other cities in India, China, Nepal, the Philippines, East Africa and the United States, tens of thousands of workers are punching a clock while they teach the machines.

Tens of thousands more workers, independent contractors usually working in their homes, also annotate data through crowdsourcing services like Amazon Mechanical Turk, which lets anyone distribute digital tasks to independent workers in the United States and other countries. The workers earn a few pennies for each label.

Based in India, iMerit labels data for many of the biggest names in the technology and automobile industries. It declined to name these clients publicly, citing confidentiality agreements. But it recently revealed that its more than 2,000 workers in nine offices around the world are contributing to an online data-labeling service from Amazon called SageMaker Ground Truth. Previously, it listed Microsoft as a client.

One day, who knows when, artificial intelligence could hollow out the job market.

But for now, it is generating relatively low-paying jobs. The market for data labeling passed \$500 million in 2018 and it will reach \$1.2 billion by 2023, according to the research firm Cognilytica. This kind of work, the study showed, accounted for 80 percent of the time spent building A.I. technology.

Is the work exploitative? It depends on where you live and what you're working on. In India, it is a ticket to the middle class. In New Orleans, it's a decent enough job. For someone working as an independent contractor, it is often a dead end.

There are skills that must be learned — like spotting signs of a disease in a video or medical scan or keeping a steady hand when drawing a digital lasso around the image of a car or a tree. In some cases, when the task involves medical videos, pornography or violent images, the work turns grisly.

"When you first see these things, it is deeply disturbing. You don't want to go back to the work. You might not go back to the work," said Kristy Milland, who spent years doing data-labeling work on Amazon Mechanical Turk and has become a labor activist on behalf of workers on the service.

"But for those of us who cannot afford to not go back to the work, you just do it," Ms. Milland said.

Before traveling to India, I tried labeling images on a crowdsourcing service, drawing digital boxes around Nike logos and identifying "not safe for work" images. I was painfully inept.

I had to pass a test before starting the work. Even that was disheartening. The first three times, I failed. Labeling images so people could instantly search a website for retail goods wasn't exactly inspiring.

A.I. researchers hope they can build systems that can learn from smaller amounts of data. But for the foreseeable future, human labor is essential.

"This is an expanding world, hidden beneath the technology," said Mary Gray, an anthropologist at Microsoft and the co-author of the book "Ghost Work," which explores the data labeling market. "It is hard to take humans out of the loop."

CITY OF TEMPLES

Bhubaneswar is called the City of Temples. Ancient Hindu shrines rise over roadside markets at the southwestern end of the city — giant towers of stacked stone that date to the first millennium. In the city center, many streets are unpaved. Cows and feral dogs meander among the mopeds, cars and trucks.

The city of 830,000 people is also a rapidly growing hub for online labor. About a 15-minute drive from the temples, on a (paved) road near the city center, a white, four-story building sits behind a stone wall. Inside, there are three rooms filled with long rows of desks, each with its own wide-screen computer display. This was where Namita Pradhan spent her days labeling videos when I met her.

Ms. Pradhan grew up just outside the city and earned a degree from a local college, where she studied biology and other subjects before taking the job with iMerit. It was recommended by her brother, who was already working for the company. She lived at a hostel near her office during the week and took the bus back to her family home each weekend.

I visited the office on a temperate January day. Some of the women sitting at the long rows of desks were traditionally dressed — bright red saris, long gold earrings. Ms. Pradhan wore a green long-sleeve shirt, black pants, and white lace-up shoes as she annotated videos for a client in the United States.

Over the course of what was a typical eight-hour day, the shy 24-year-old watched about a dozen colonoscopy videos, constantly reversing the video for a closer look at individual frames.

Every so often, she would find what she was looking for.

She would lasso it with a digital "bounding box." She drew hundreds of these bounding boxes, labeling the polyps and other signs of illness, like blood clots and inflammation.

Her client, a company in the United States that iMerit is not allowed to name, will eventually feed her work into an A.I. system so it can learn to identify medical conditions on its own. The colonoscopy patient is not necessarily aware the video exists. Ms. Pradhan doesn't know where the images came from. Neither does iMerit.

Ms. Pradhan learned the task during seven days of online video calls with a nonpracticing doctor, based in Oakland, Calif., who helps train workers at many iMerit offices. But some ask whether experienced doctors and medical students should do this labeling themselves.

This work requires people "who have a medical background, and the relevant knowledge in anatomy and pathology," said Dr. George Shih, a radiologist at Weill Cornell Medicine and NewYork-Presbyterian and the co-founder of the start-up MD.ai, which helps organizations build artificial intelligence for health care.

When we chatted about her work, Ms. Pradhan called it "quite interesting," but tiring. As for the graphic nature of the videos? "It was disgusting at first, but then you get used to it."

The images she labeled were grisly, but not as grisly as others handled at iMerit. Their clients are also building artificial intelligence that can identify and remove unwanted images on social networks and other online services. That means labels for pornography, graphic violence and other noxious images.

This work can be so upsetting to work-

ers, iMerit tries to limit how much of it they see. Pornography and violence are mixed with more innocuous images, and those labeling the grisly images are sequestered in separate rooms to shield other workers, said Liz O'Sullivan, who oversaw data annotation at an A.I. start-up called Clarifai and has worked closely with iMerit on such projects.

Other labeling companies will have workers annotate unlimited numbers of these images, Ms. O'Sullivan said.

"I would not be surprised if this causes post-traumatic stress disorder — or worse. It is hard to find a company that is not ethically deplorable that will take this on," she said. "You have to pad the porn and violence with other work, so the workers don't have to look at porn, porn, porn, beheading, beheading, beheading."

iMerit said in a statement it does not compel workers to look at pornography or other offensive material and only takes on the work when it can help improve monitoring systems.

Ms. Pradhan and her fellow labelers earn between \$150 and \$200 a month, which pulls in between \$800 and \$1,000 of revenue for iMerit, according to one company executive.

By United States standards, Ms. Pradhan's salary is indecently low. But for her and many others in these offices, it is about an average salary for a data-entry job.

TEDIOUS WORK, BUT IT PAYS

Prasenjit Baidya grew up on a farm about 30 miles from Kolkata, the largest

city in West Bengal, on the east coast of India. His parents and extended family still live in his childhood home, a cluster of brick buildings built at the turn of the 19th century. They grow rice and sunflowers in the surrounding fields and dry the seeds on rugs spread across the rooftops.

He was the first in his family to get a college education, which included a computer class. But the class didn't teach him all that much. The room offered only one computer for every 25 students. He learned his computer skills after college, when he enrolled in a training course run by a nonprofit called Anudip. It was recommended by a friend, and it cost the equivalent of \$5 a month.

Anudip runs English and computer courses across India, training about 22,000 people a year. It feeds students directly into iMerit, which its founders set up as a sister operation in 2013. Through Anudip, Mr. Baidya landed a job at an iMerit office in Kolkata, and so did his wife, Barnali Paik, who grew up in a nearby village.

Over the last six years, iMerit has hired more than 1,600 students from Anudip. It now employs about 2,500 people in total. More than 80 percent come from families with incomes below \$150 a month.

Founded in 2012 and still a private company, iMerit has its employees perform digital tasks like transcribing audio files or identifying objects in photos. Businesses across the globe pay the company to use its workers, and in-

creasingly, they assist work on artificial intelligence.

"We want to bring people from low-income backgrounds into technology — and technology jobs," said Radha Basu, who founded Anudip and iMerit with her husband, Dipak, after long careers in Silicon Valley with the tech giants Cisco Systems and HP.

The average age of these workers is 24. Like Mr. Baidya, most of them come from rural villages.

At first, iMerit focused on simple tasks — sorting product listings for online retail sites, vetting posts on social media. But it has shifted into work that feeds artificial intelligence.

The growth of iMerit and similar companies represents a shift away from crowdsourcing services like Mechanical Turk. iMerit and its clients have greater control over how workers are trained and how the work is done.

Mr. Baidya, now a manager at iMerit, oversees an effort to label street scenes used in training driverless cars for a major company in the United States. His team analyzes and labels digital photos as well as three-dimensional images captured by Lidar, devices that measure distances using pulses of light. They spend their days drawing bounding boxes around cars, pedestrians, stop signs and power lines.

He said the work could be tedious, but it had given him a life he might not have otherwise had. He and his wife recently bought an apartment in Kolkata, within walking distance of the iMerit office where she works.

"The changes in my life — in terms of my financial situation, my experiences, my skills in English — have been a dream," he said. "I got a chance."

LISTENING TO PEOPLE COUGH

A few weeks after my trip to India, I took an Uber through downtown New Orleans. About 18 months ago, iMerit moved into one of the buildings across the street from the Superdome.

A major American tech company needed a way of labeling data for a Spanish-language version of its home digital assistant. So it sent the data to the new iMerit office in New Orleans.

After Hurricane Katrina in 2005, hundreds of construction workers and their families moved into New Orleans to help rebuild the city. Many stayed. A number of Spanish speakers came with that new work force, and the company began hiring them.

Oscar Cabezas, 23, moved with his mother to New Orleans from Colombia. His stepfather found work in construction, and after college Mr. Cabezas joined iMerit as it began working on the Spanish-language digital assistant.

He annotated everything from tweets to restaurant reviews, identifying people and places and pinpointing ambiguities. In Guatemala, for instance, "pisto" means money, but in Mexico, it means beer. "Every day was a new project," he said.

The office has expanded into other work, serving businesses that want to keep their data within the United States. Some projects must remain in America for legal and security purposes.

Glenda Hernandez, 42, who was born in Guatemala, said she missed her old work on the digital assistant project. She loved to read. She reviewed books online for big publishing companies so she could get free copies, and she relished the opportunity of getting paid to read in Spanish.

"That was my baby," she said of the project.

She was less interested in image tagging or projects like the one that involved annotating recordings of people coughing; it was a way to build A.I. that identifies disease symptoms of illness over the phone.

"Listening to coughs all day is kind of disgusting," she said.

The work is easily misunderstood, said Ms. Gray, the Microsoft anthropologist. Listening to people cough all day may be disgusting, but that is also how doctors spend their days. "We don't think of that as drudgery," she said.

Ms. Hernandez's work is intended to help doctors do their jobs or maybe, one day, replace them. She takes pride in that. Moments after complaining about the project, she pointed to her colleagues across the office.

"We were the cough masters," she said.

NOT ENOUGH TO LIVE ON

In 2005, Kristy Milland signed up for her first job on Amazon Mechanical Turk. She was 26 and living in Toronto with her husband, who managed a local warehouse. Mechanical Turk was a way of making a little extra money.

The first project was for Amazon itself. Three photos of a storefront would pop up on her laptop, and she would choose the one that showed the front door. Amazon was building an online service similar to Google Street View, and the company needed help picking the best photos.

She made three cents for each click, or about 18 cents a minute. In 2010, her husband lost his job, and "MTurk" became a full-time gig. For two years, she worked six or seven days a week, sometimes as much as 17 hours a day. She made about \$50,000 a year.

"It was enough to live on then. It wouldn't be now," Ms. Milland said.

The work at that time didn't really involve A.I. For another project, she would pull information out of mortgage documents or retype names and addresses from photos of business cards, sometimes for as little as a dollar an hour.

Around 2010, she started labeling for A.I. projects. Ms. Milland tagged all sorts of data, like gory images that showed up on Twitter (which helps build A.I. that can help remove gory images from the social network) or aerial footage likely taken somewhere in the Middle East (presumably for A.I. that the military and its partners are building to identify drone targets).

Projects from American tech giants, Ms. Milland said, typically paid more than the average job — about \$15 an hour. But the job didn't come with health care or paid vacation, and the work could be mind-numbing — or downright disturbing. She called it "horribly exploitative." Amazon declined to comment.


Since 2012, Ms. Milland, now 40, has been part of an organization called TurkerNation, which aims to improve conditions for thousands of people who do this work. In April, after 14 years on the service, she quit.

She is in law school, and her husband makes \$600 less than they pay in rent each month, which does not include utilities. So, she said, they are preparing to go into debt. But she will not go back to labeling data.

"This is a dystopian future," she said. "And I am done."



REBECCA CONWAY FOR THE NEW YORK TIMES



BRYAN TARNOWSKI FOR THE NEW YORK TIMES

Top, Namita Pradhan, second from right, identifying anomalies in medical scans. Above, Oscar Cabezas helps with Spanish-language tasks at iMerit in New Orleans.