

Considerations of Artificial Intelligence for Monitoring, Control, and Surveillance Officers

What is Artificial Intelligence?

Computers performing human tasks (for example, playing games, writing, listening), while also potentially learning and improving

Machine learning (ML) is a part of artificial intelligence (AI)

Goal today:

- Introduce some terms that AI engineers will use

- Show potential pitfalls of using AI/ML to make decisions

- Highlight the value of ML for the right tasks

Machine Learning for Monitoring, Control, and Surveillance

Problem:

We need to monitor the ocean for dark vessels

Machine Learning for Monitoring, Control, and Surveillance

Problem:

We need to monitor the ocean for dark vessels

Solution:

Public satellite imagery is available!



Machine Learning for Monitoring, Control, and Surveillance

Problem:

We need to monitor the ocean for dark vessels

Solution:

Public satellite imagery is available!

...but there is a lot of it, and the ocean area to search is large

Machine Learning for Monitoring, Control, and Surveillance

Problem:

We need to monitor the ocean for dark vessels

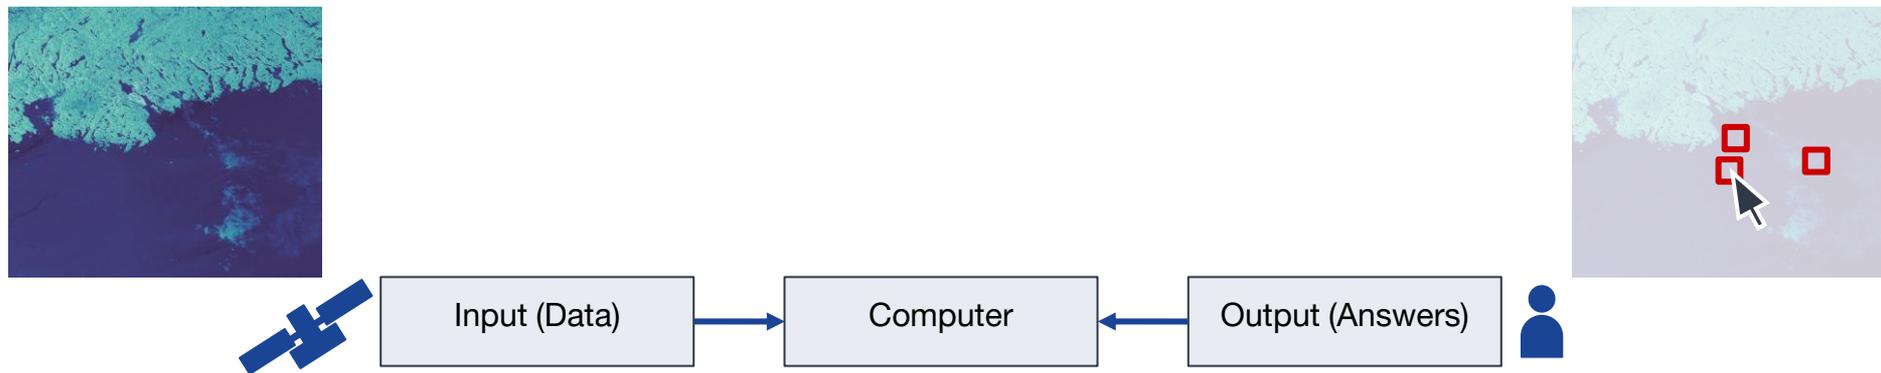
Solution:

Public satellite imagery is available!

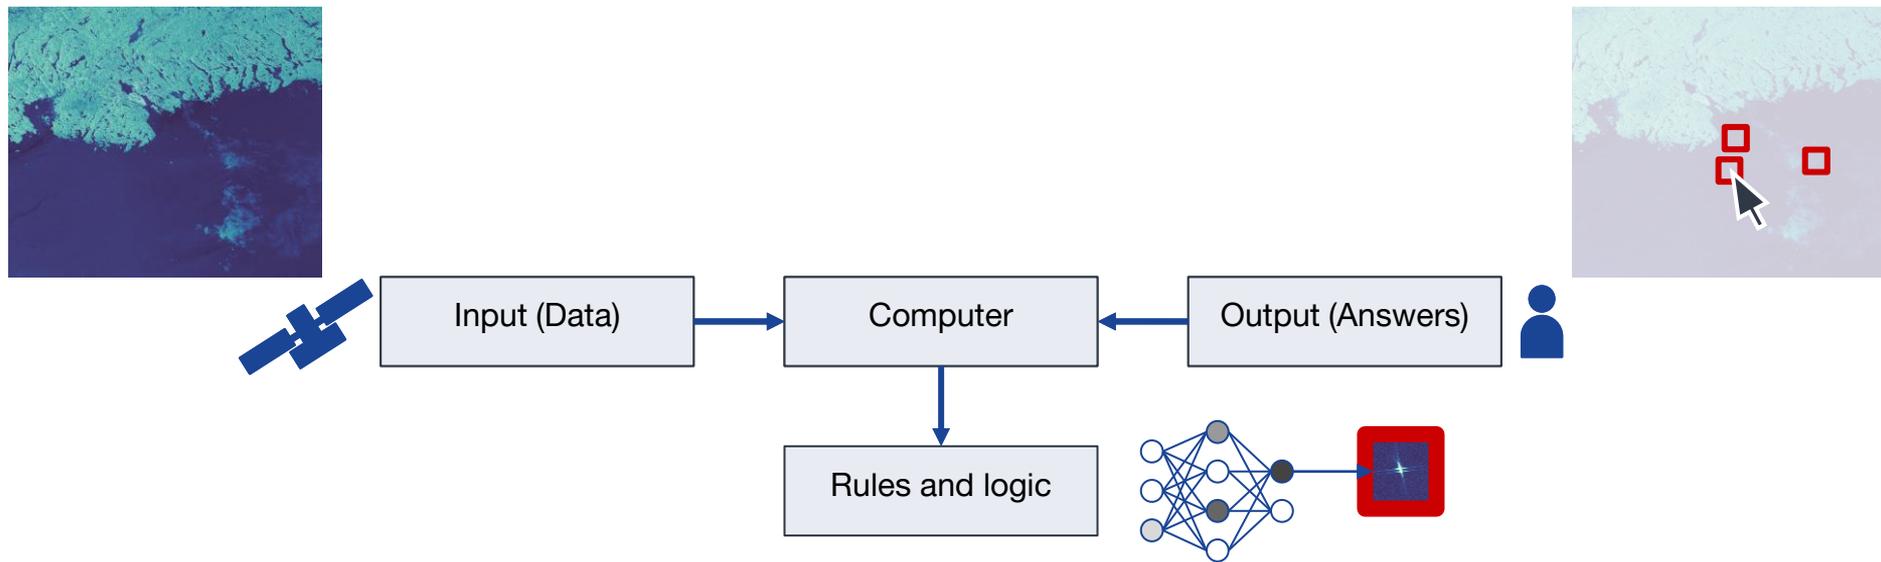
...but there is a lot of it, and the ocean area to search is large

...so let's use machine learning to automate the job of looking through lots of satellite imagery for vessels

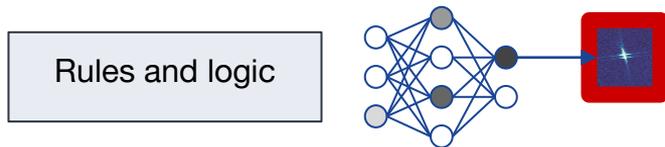
Machine Learning for Monitoring, Control, and Surveillance



Machine Learning for Monitoring, Control, and Surveillance

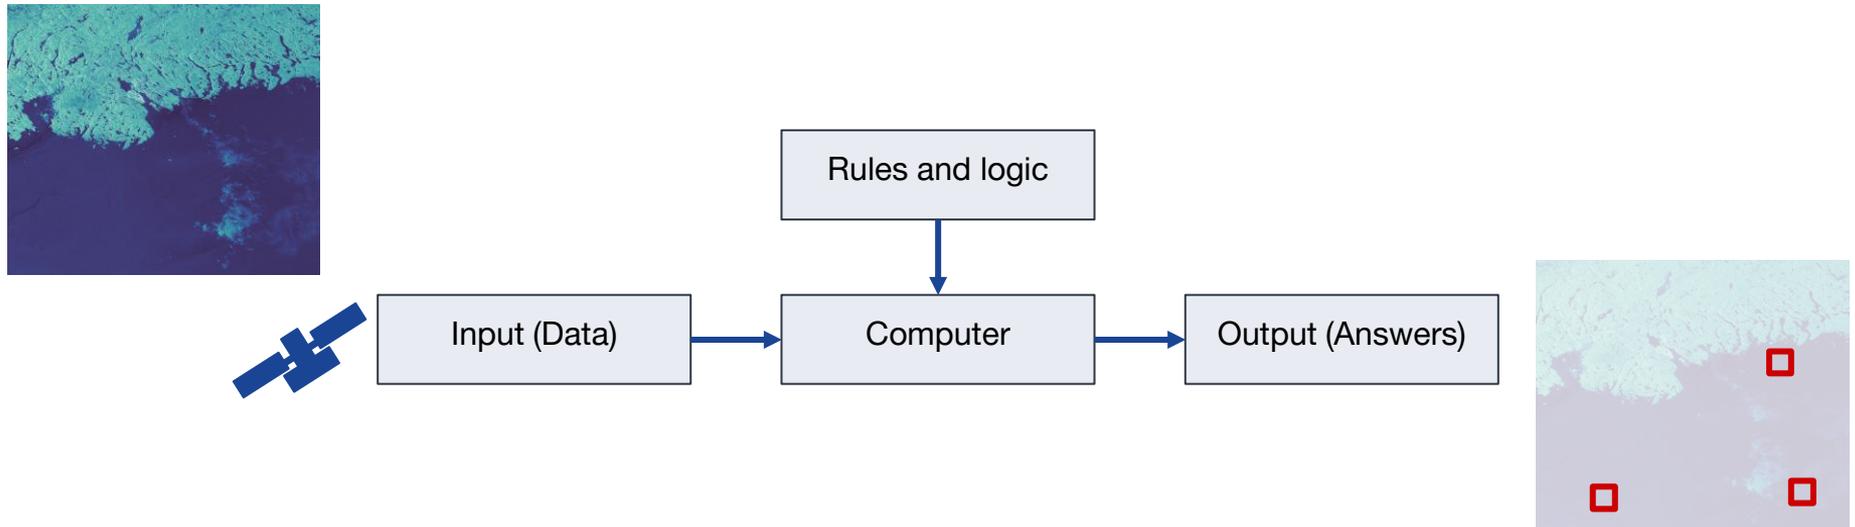


Machine Learning for Monitoring, Control, and Surveillance



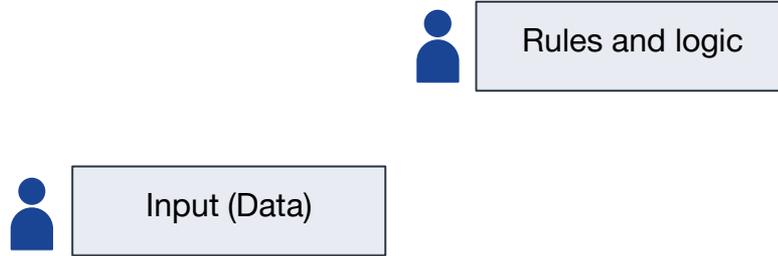
Machine Learning for Monitoring, Control, and Surveillance

Result: timely vessel detections without a need for human analysis



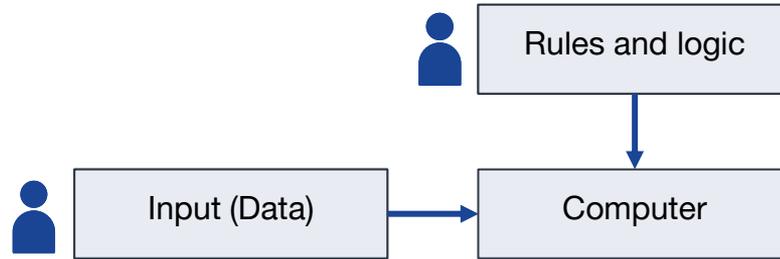
What is Machine Learning?

Traditional
Programming



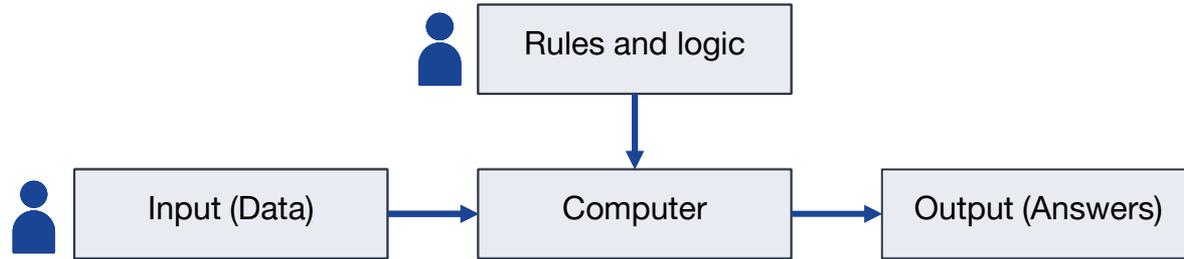
What is Machine Learning?

Traditional
Programming



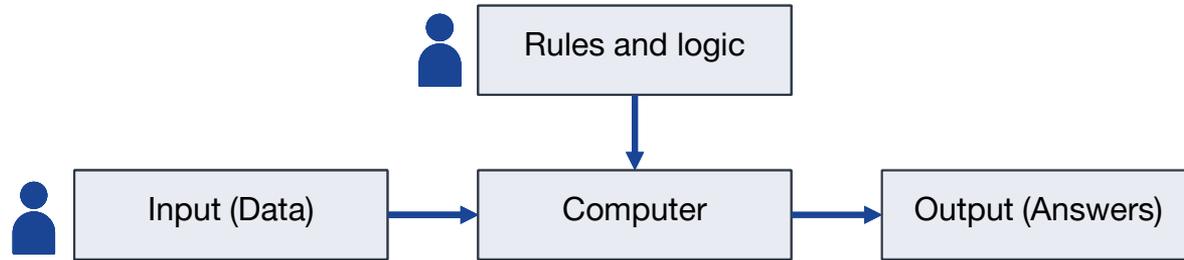
What is Machine Learning?

Traditional
Programming



What is Machine Learning?

Traditional
Programming

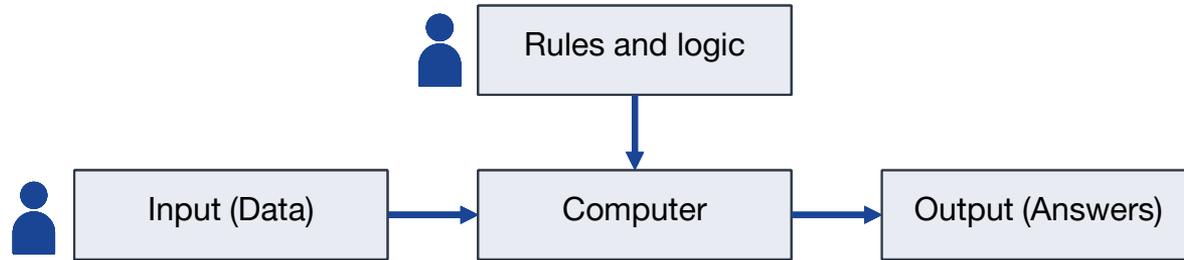


Machine
Learning



What is Machine Learning?

Traditional
Programming

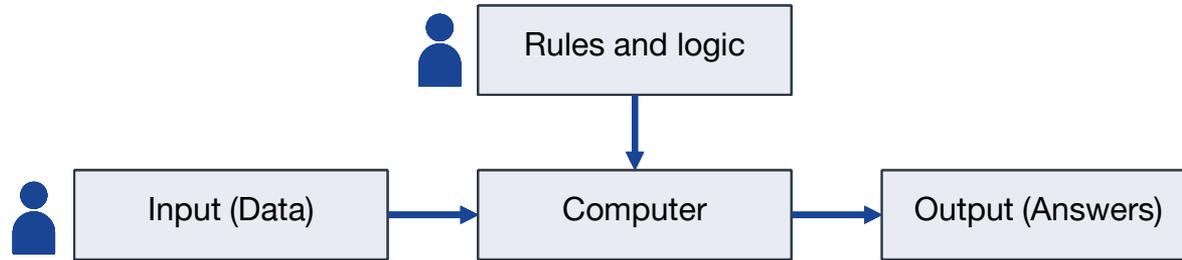


Machine
Learning

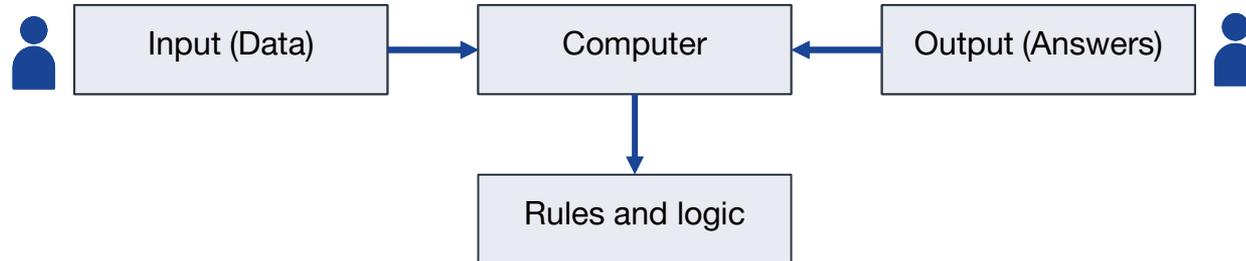


What is Machine Learning?

Traditional
Programming

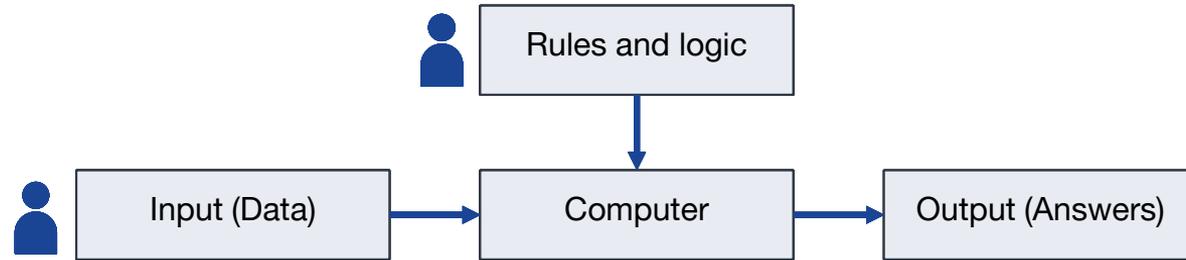


Machine
Learning

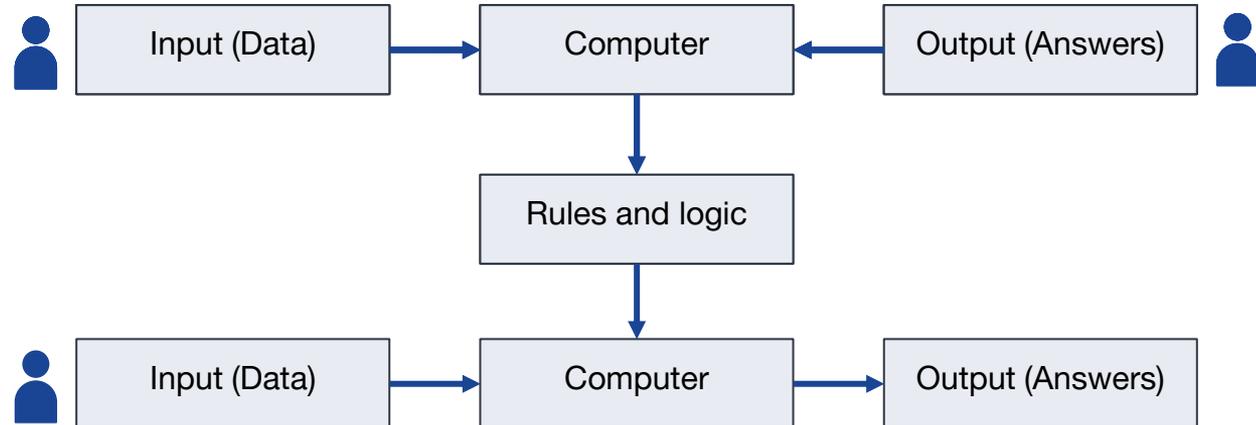


What is Machine Learning?

Traditional
Programming



Machine
Learning



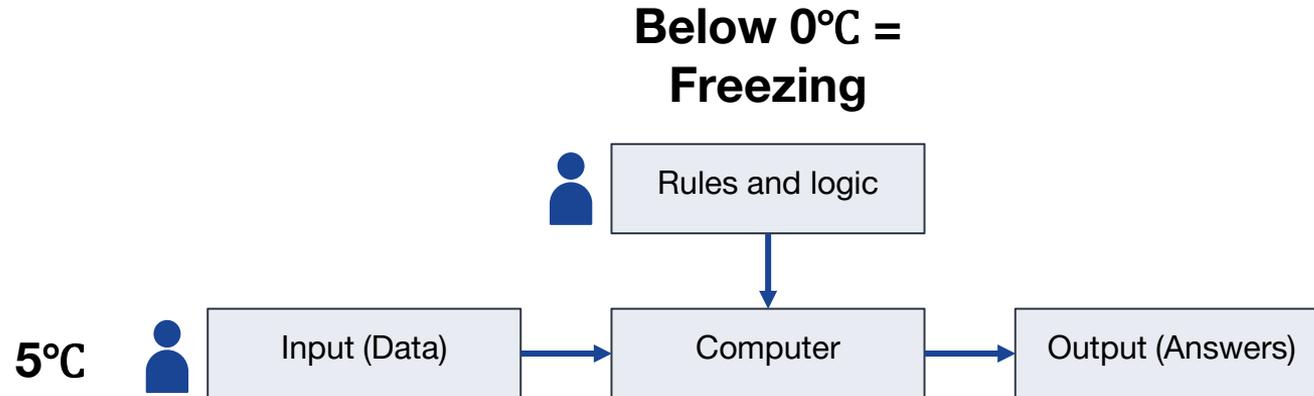
Then why machine learning?

Then why machine learning?

It can help us find more complicated rules, or unknown relationships

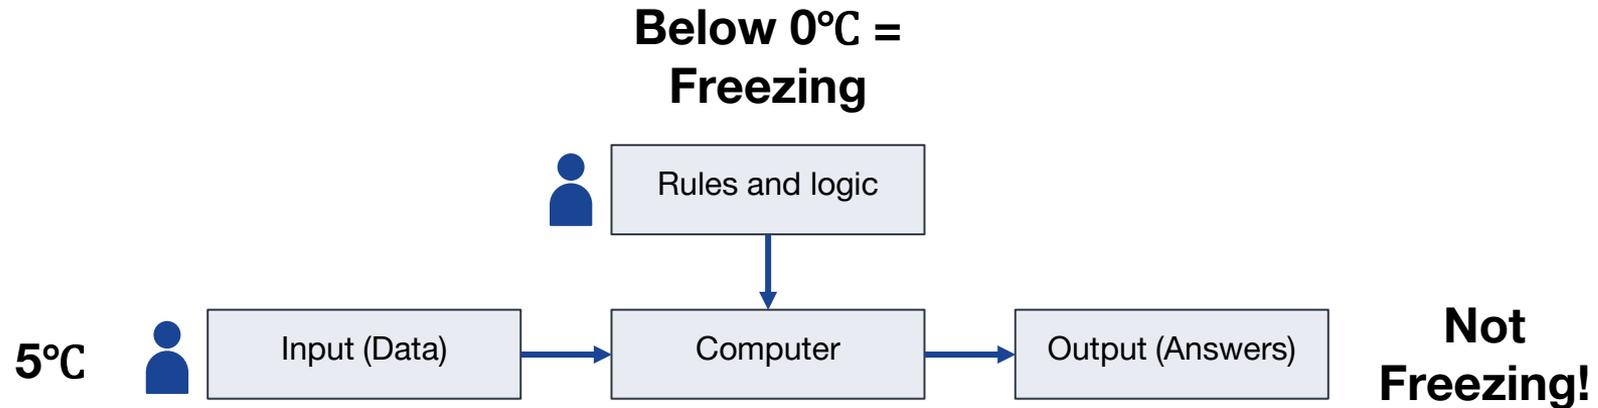
Then why machine learning?

Traditional programming answering, “Is it freezing?”



Then why machine learning?

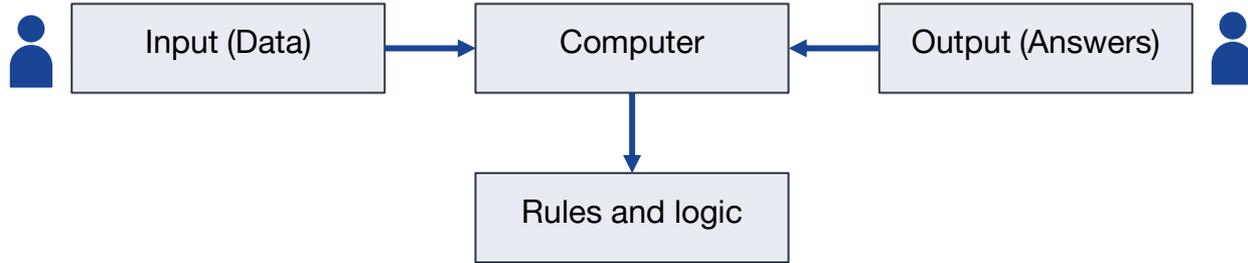
Traditional programming answering, “Is it freezing?”



Then why machine learning?

Machine learning answering, “What do people consider cold?”

12°C
5°C
24°C
17°C
21°C
...

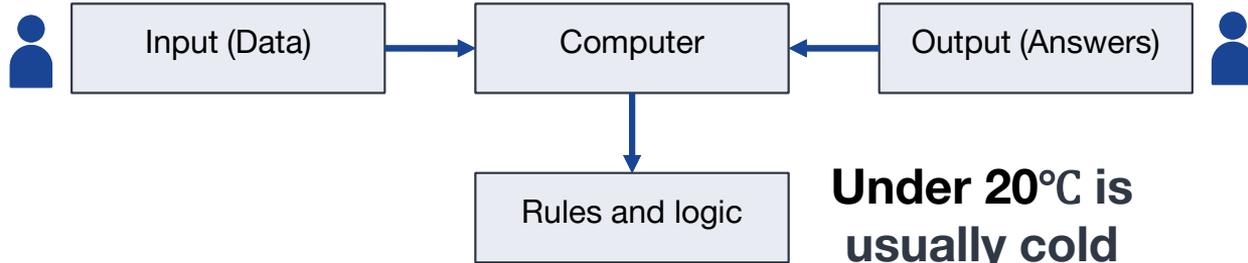


Warm!
Cold!
Warm!
Cold!
Warm!
...

Then why machine learning?

Machine learning answering, “What do people consider cold?”

12°C
5°C
24°C
17°C
21°C
...



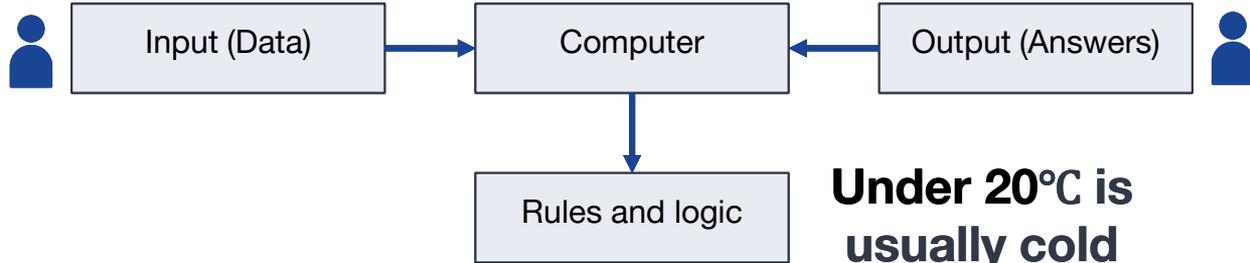
**Under 20°C is
usually cold**

**Warm!
Cold!
Warm!
Cold!
Warm!
...**

Then why machine learning?

This is called “training” in machine learning

12°C
5°C
24°C
17°C
21°C
...

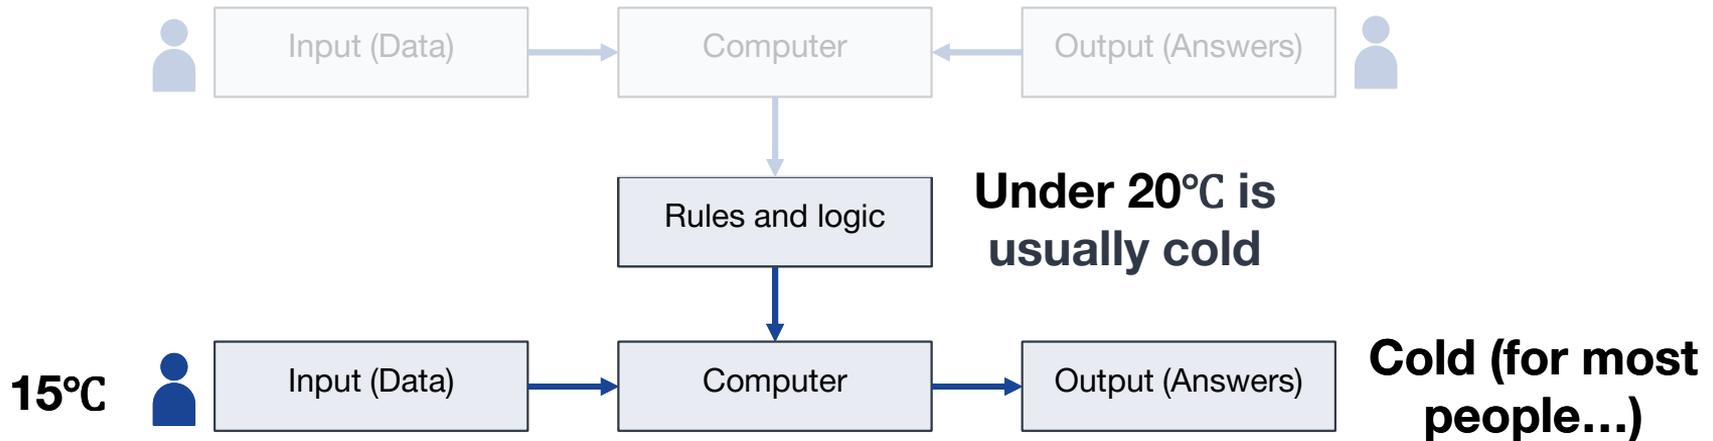


**Under 20°C is
usually cold**

Warm!
Cold!
Warm!
Cold!
Warm!
...

Then why machine learning?

Now we can use the rule we've learned to give answers:



Potential problems?

The model will learn the data we give it...

...including any flaws and biases

These results don't work for the person who thought 12 degrees was warm

These results might depend on where people are accustomed to living

Doing ML the Right Way

Problems in machine learning can be addressed by careful design:

- Collect a lot of data from diverse sources
- Choose the right type of model for the desired outcome
- Focus on a specific task or problem

...and carefully examining the data:

- Are there human prejudices captured in the data?
- Are we only observing certain things, in certain places, at certain times?
- Do we have an equal amount of examples of everything we expect to see in the real world?

Doing ML the Right Way

Machine learning does not immediately have the correct answer, but must find it over time

A trial and error process, training a machine learning system needs:

- a lot of high quality data (tens to hundreds of thousands of examples!)
- a lot of computing power (thousands of processors!)
- a lot of time to try different options (thousands of attempts to improve!)

Doing ML the Right Way

We decide on:

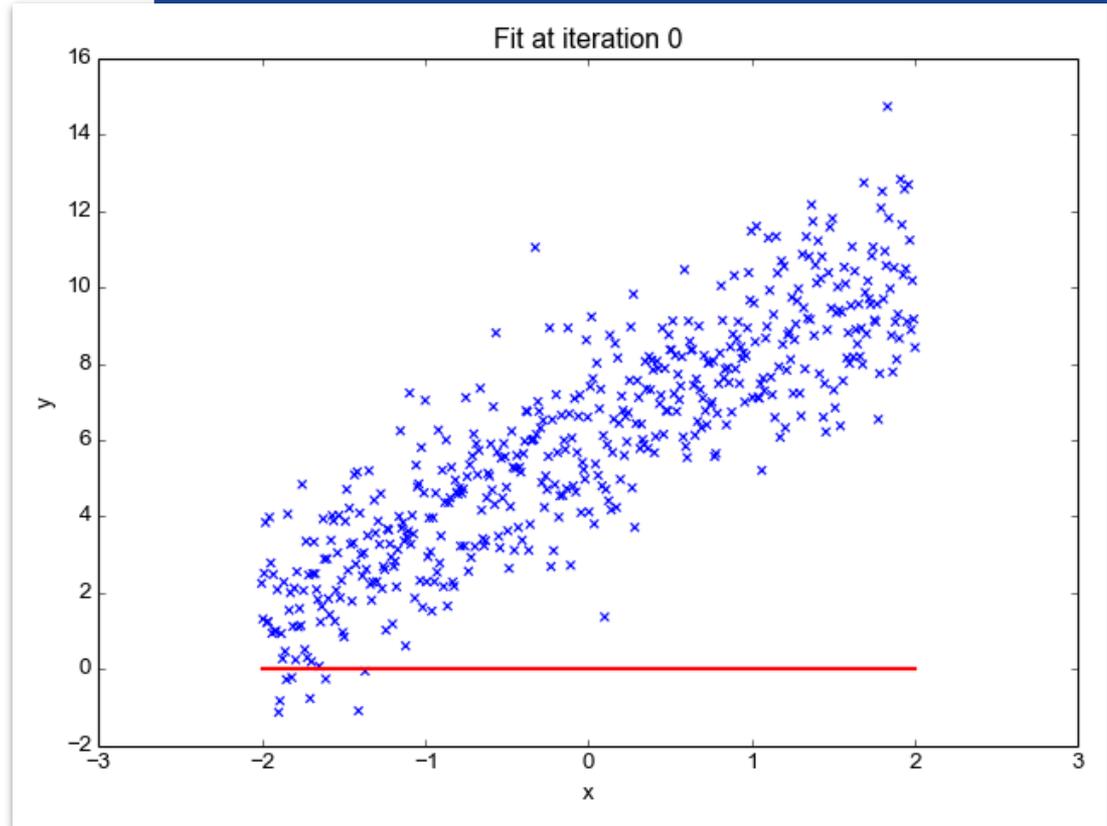
- A model (how have similar problems been solved in the past?)
- A goal:
 - Do we want to be highly accurate at the risk of missing things? This is called “precision”
 - Do we want to catch more at the risk of getting some things wrong? This is called “recall”
 - You may hear “F1 score” which is a way to capture these ideas together
- And how we will score or grade the model:
 - We can use a scoring method called a “loss function” to tell the model how incorrect it was
 - The model then attempts to improve its rules and tries again, hopefully getting closer to a best possible solution, often using something called “gradient descent”
 - “Attention maps” tell us what a model thinks is important
 - “Confusion matrices” tell us what a model cannot tell apart (think: fixed versus drifting gear)

What ML models can do: score (regression)

We can find a “line of best fit” that gives us a general idea of the data

We can use this line to put in an “x” and get back “y” as an answer or prediction, maybe fishing yield based on vessels at sea

...but depending on the data, there will probably be some amount of error

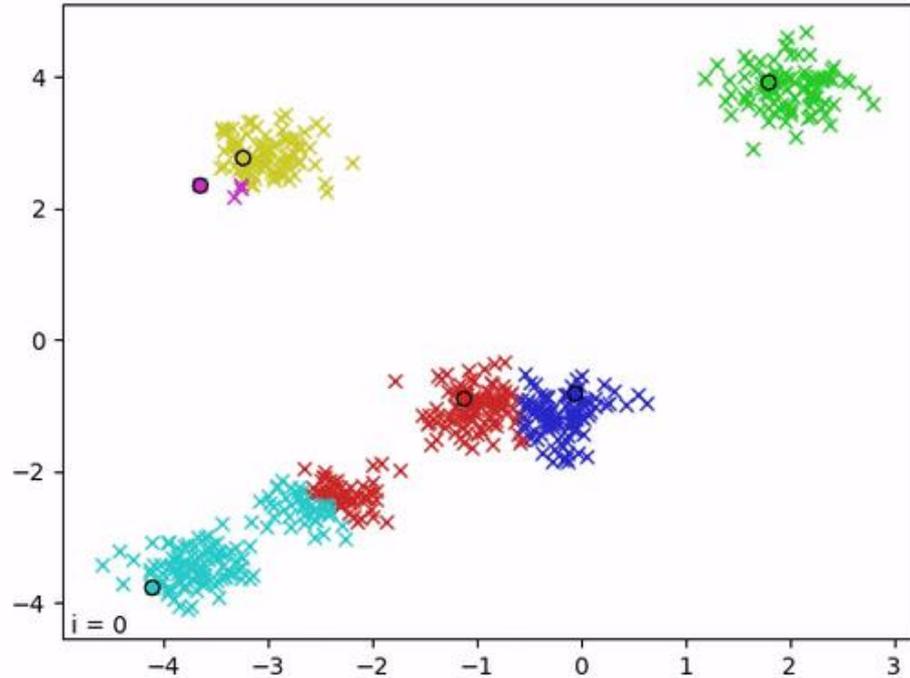


What ML models can do: separate (classify)

“K-Means Clustering” separates groups solely based on what they have in common without the groups being identified in advance

This is called “unsupervised” learning because we don’t know the answers

Example: finding groups of similar vessels and activity

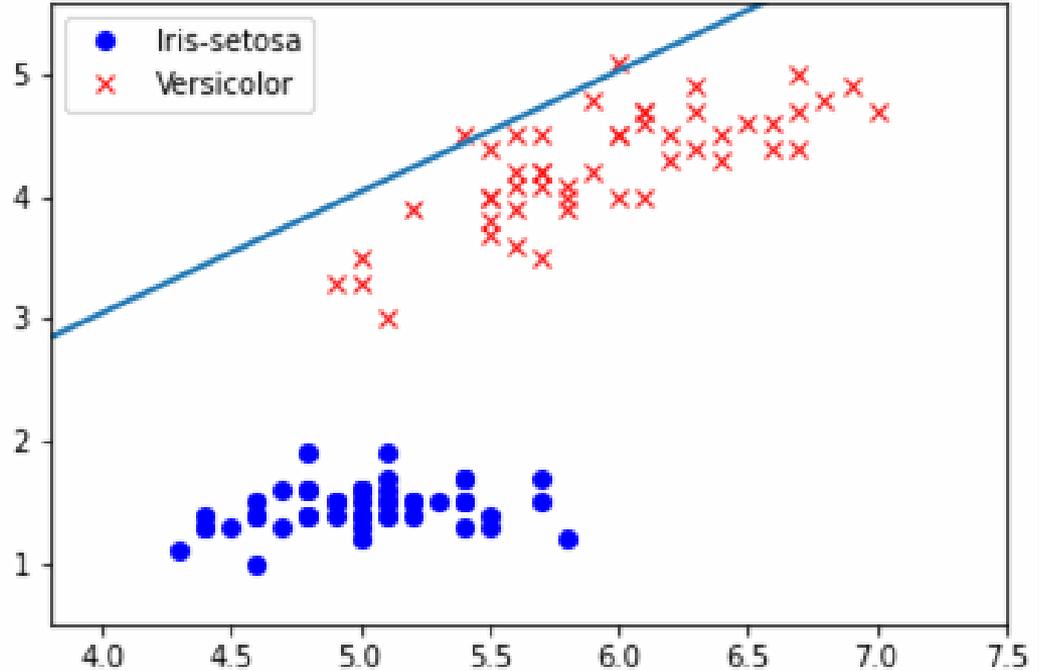


What ML models can do: separate (classify)

“Perceptrons,” invented in 1957, find a boundary line to separate two known groups, so it can predict to which group a new unknown point belongs

It corrects itself based on how wrong it is about a point

It finds a dividing line that works, and then stops

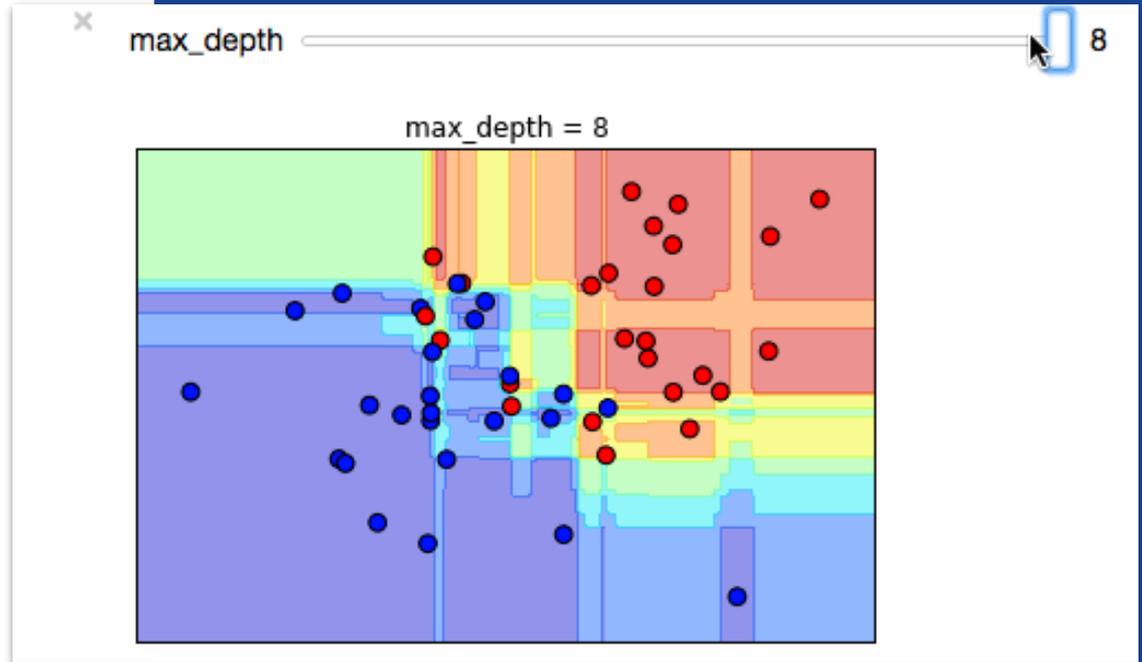


What ML models can do: separate (classify)

A “random forest” creates smaller and smaller regions until red dots are in red regions, and blue dots are in blue regions (with some room for uncertainty)

These red and blue areas can be used to predict if an unknown dot is red or blue

This is called “supervised” learning

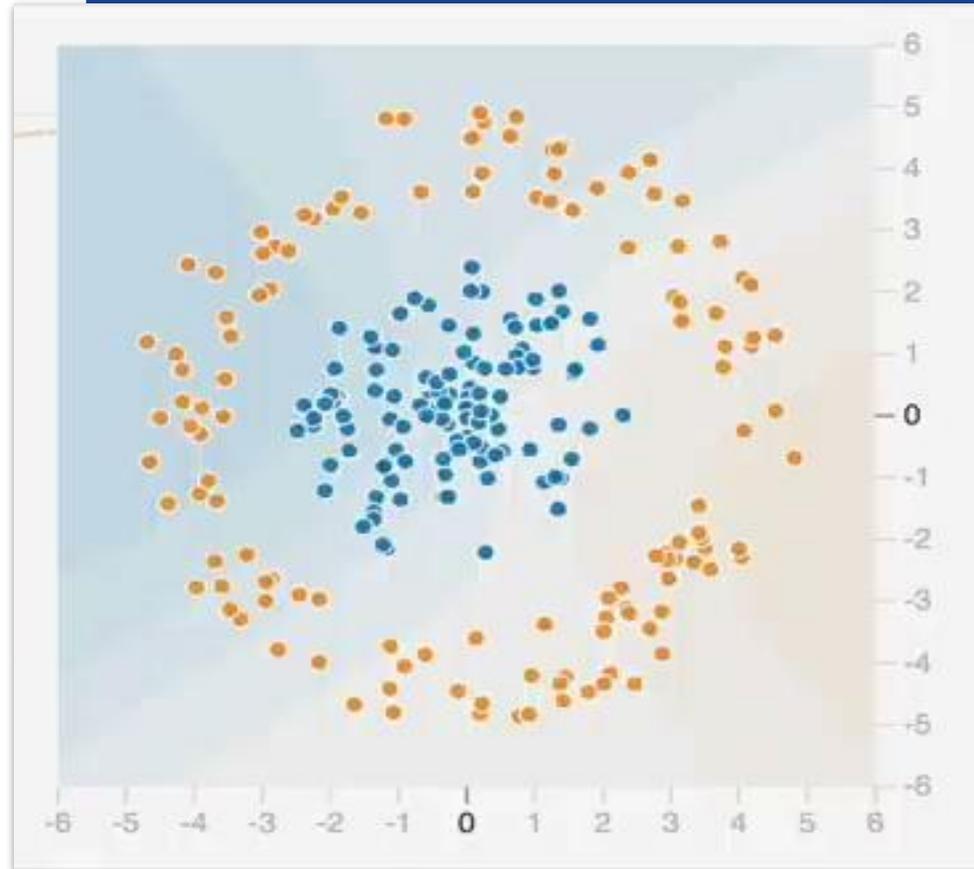


What ML models can do: separate (classify)

A “neural net” can also separate data, and more!

Neural nets are the basis for many of the AI breakthroughs you may hear about in the news lately

Data flows through a neural net from the inputs, being redirected along the way, until it reaches the right output (again figured out through a lot of trial and error!)



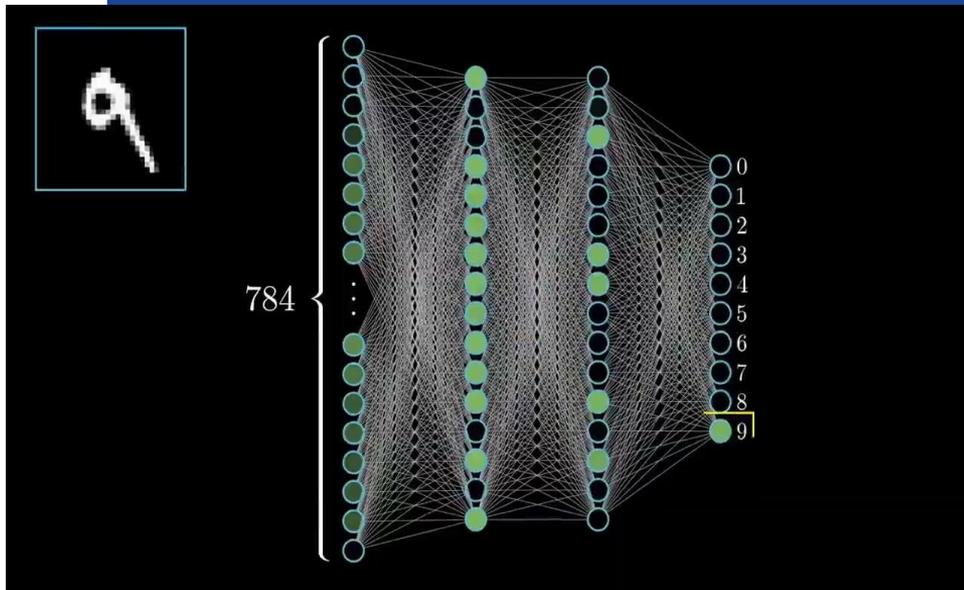
Let's talk about “large language models”...

Generative Pre-trained Transformer (GPT) are very large neural networks

They have close to as many “neurons” as the human brain

They are trained on massive amounts of data (the entire internet)

They can cost hundreds of millions of dollars to train



What does the track of a longliner fishing vessel look like when it is setting, soaking, and hauling back fishing gear?

The track of a longliner fishing vessel when it is setting, soaking, and hauling back fishing gear will typically follow a specific pattern, depending on the fishing method and the species being targeted. Longliner vessels are designed to deploy long lines with baited hooks to catch fish, so their track will reflect the process of laying out the fishing gear, letting it soak in the water, and retrieving it.

1. **Setting:** During the setting phase, the longliner vessel moves in a straight line or a gentle curve while deploying the fishing gear. The fishing gear consists of a mainline that can stretch for many miles, with smaller lines or branch lines attached at regular intervals. The baited hooks are usually spaced along the branch lines. As the vessel moves forward, crew members carefully pay out the gear to ensure proper deployment.
2. **Soaking:** Once the fishing gear is fully deployed, the vessel will typically come to a stop or maintain slow forward motion. The fishing gear will be left to "soak" in the water for a certain period, allowing the baited hooks to attract and catch fish. The duration of soaking varies depending on the target species and fishing regulations.
3. **Hauling back:** After the soaking period is over, the vessel will start hauling back the fishing gear. This phase involves winding in the longline and bringing the catch on board. The vessel's track during hauling back will often be the reverse of the setting track. Crew members will retrieve the fishing gear, unhook the catch, and prepare the gear for the next setting.

The track of a longliner fishing vessel can be recorded using GPS systems and chartplotters on the vessel. This data can help optimize fishing operations, identify productive fishing grounds, and adhere to sustainable fishing practices.

It's important to note that the specific track and fishing methods can vary depending on the type of longline fishing (e.g., pelagic longline, demersal longline) and the target species (e.g., tuna, swordfish, cod). Additionally, responsible fishing practices and regulations may require vessels to follow specific guidelines, such as avoiding protected areas or using specific gear configurations to reduce bycatch of non-target species.

Regenerate

Send a message

Areas

Events (2)

Filters (1) Clear

Save as Alert



F/V LADY CHRIS 5

Fishing

Event Date/Time

2023-07-24 T 0:05:10 GMT

Lat, Lon: -15.6055, -153.4752

F/V LADY CHRIS 5

French Polynesia

MMSI 546012900

Vessel Type Fishing

IMO -

Length 21 M

Event History 312 Fishing Events In

Our Areas

