CPSC 322
Introduction to Artificial Intelligence

December 1, 2004

Things...

Welcome to December!

Kaili's practice homework assignment has been posted.

Is there some law prohibiting the sale of Christmas trees in British Columbia?

Example: Supervised learning of concept

Negative examples help the learning procedure specialize. If the model of an arch is too general (too inclusive), the negative examples will tell the procedure in what ways to make the model more specific.

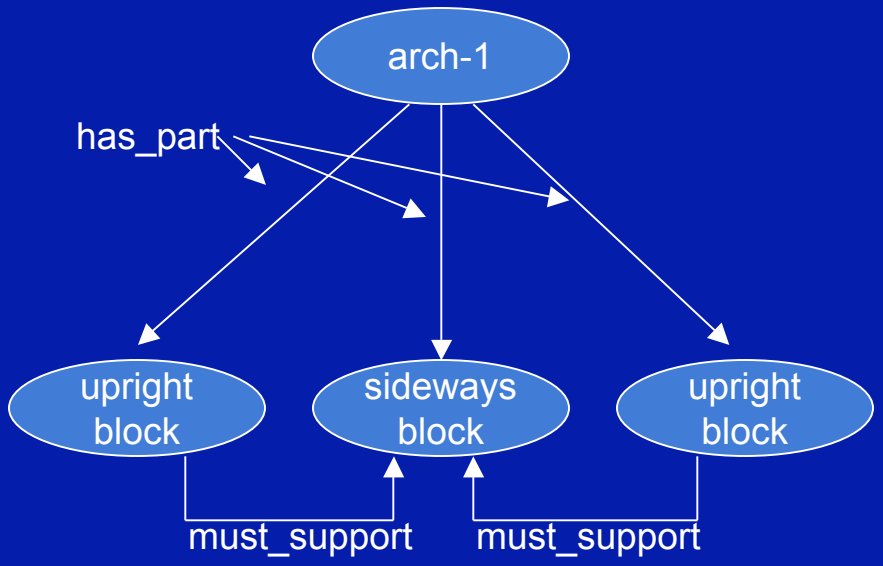

Example: Supervised learning of concept

Here comes another near miss. What's the difference between the near miss and the current concept of an arch?

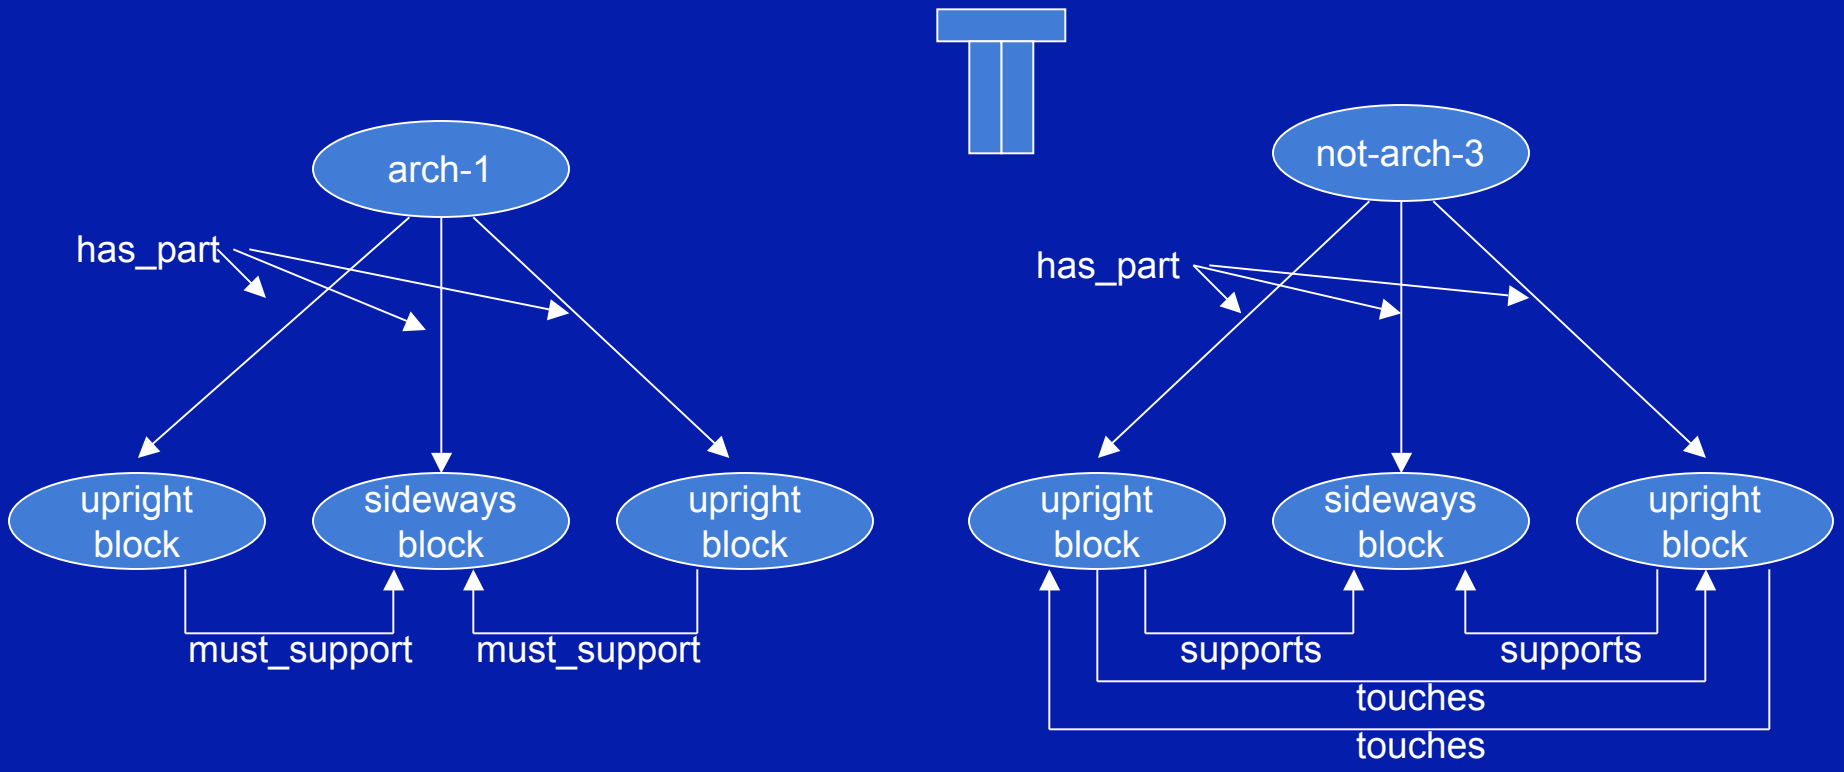

Example: Supervised learning of concept

The difference is the existence of the touches links in the near miss. That is, there's no gap between the upright blocks. Since that's the only difference, then the supporting blocks in an arch must not touch.

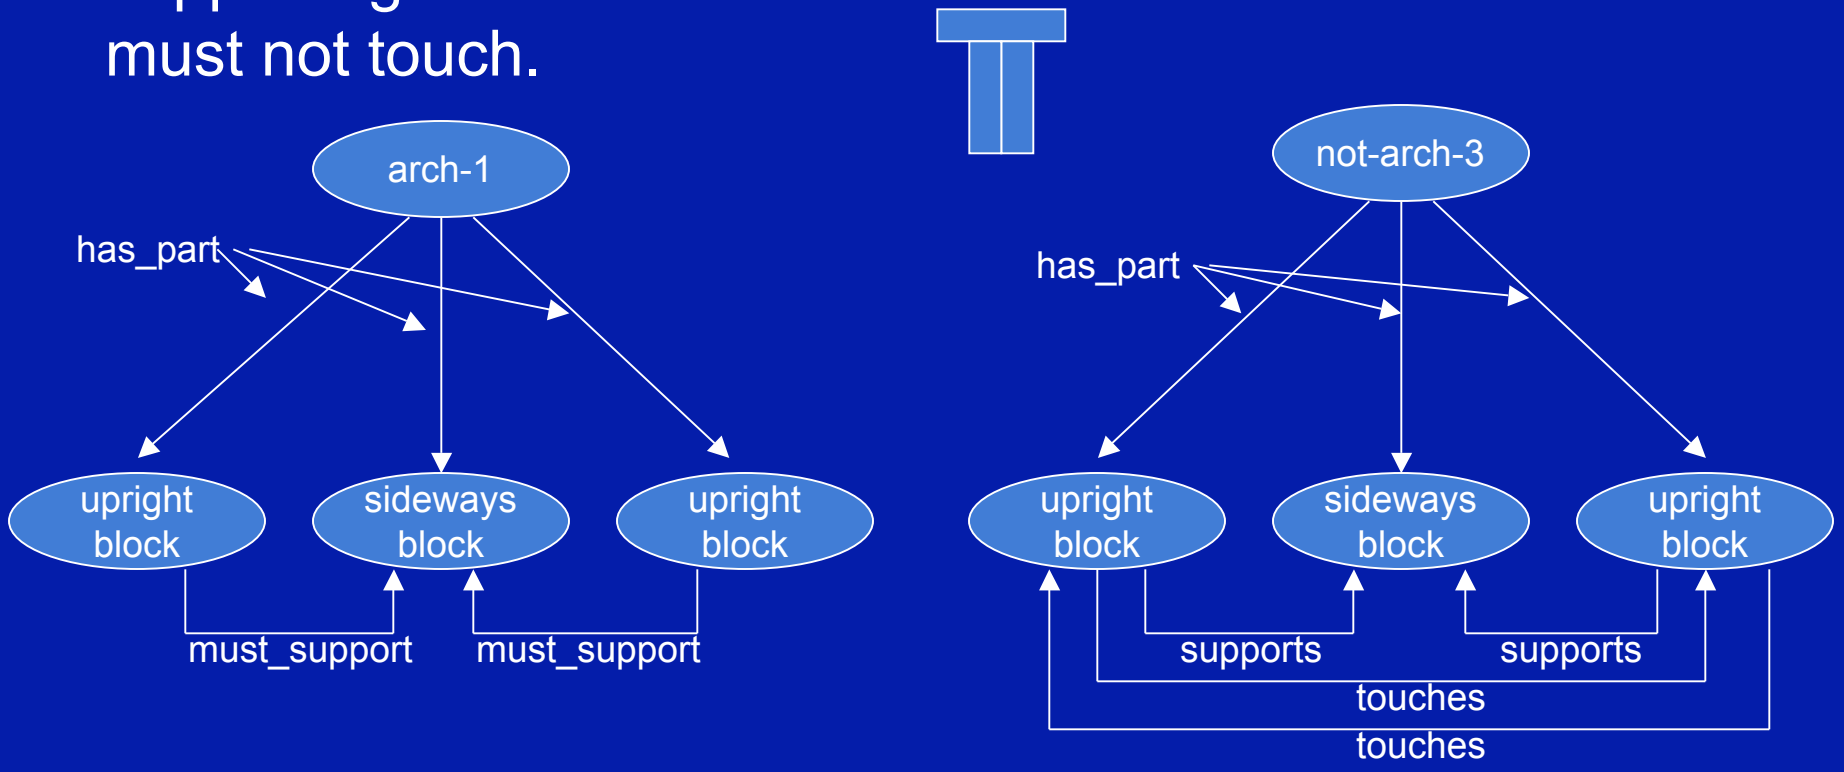

Example: Supervised learning of concept

The program updates its representation to reflect that the touches links between the upright blocks are forbidden.

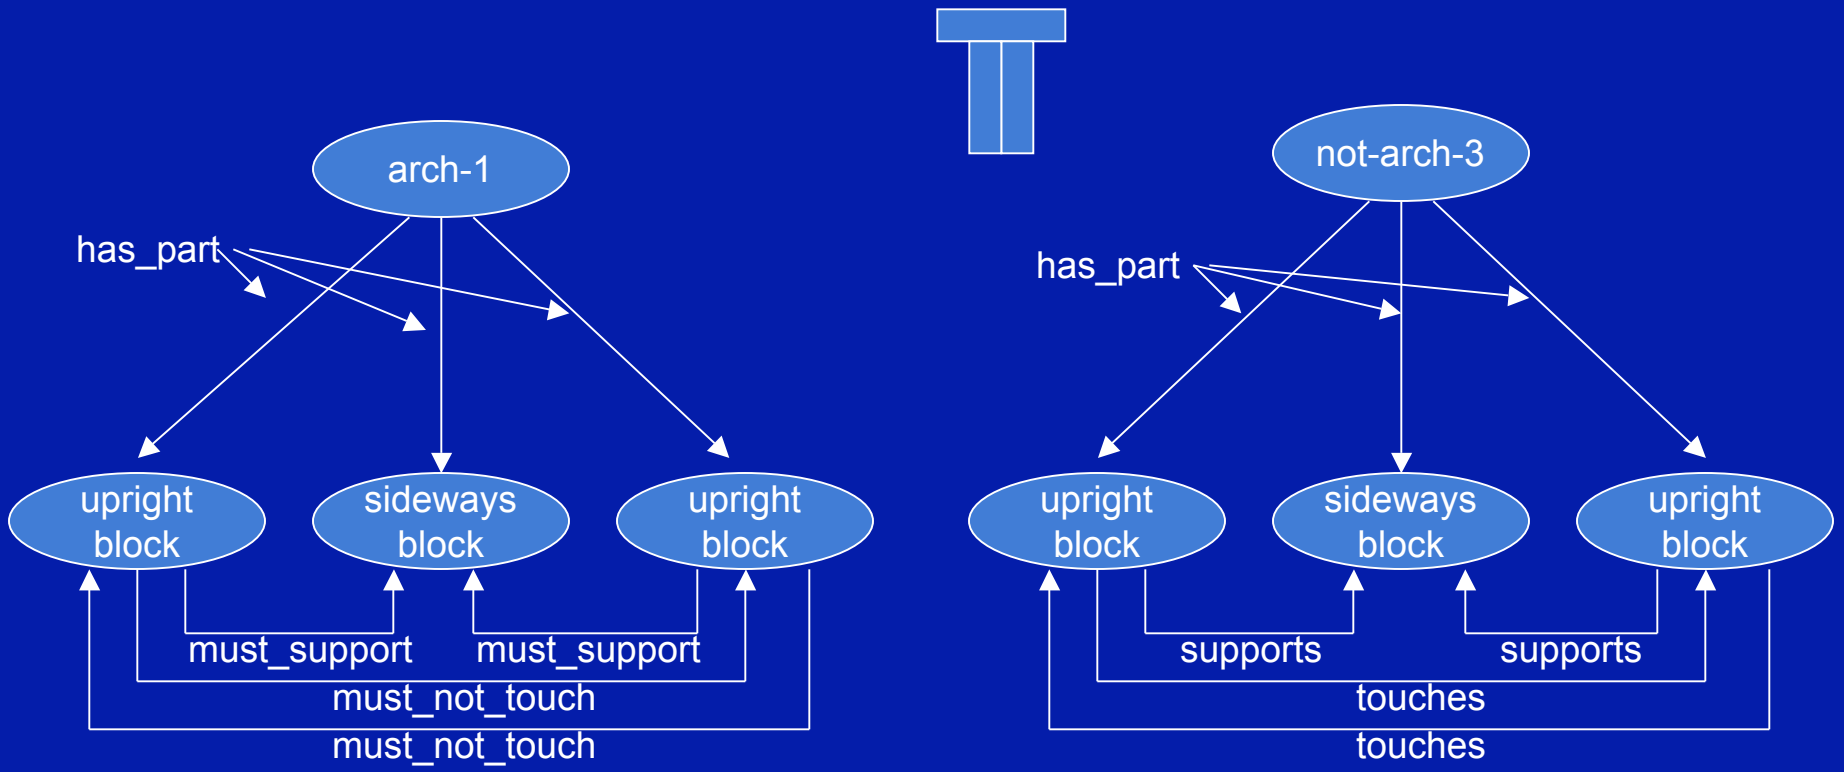

Example: Supervised learning of concept

Because of the second negative example, the concept of the arch is even more specific than before.

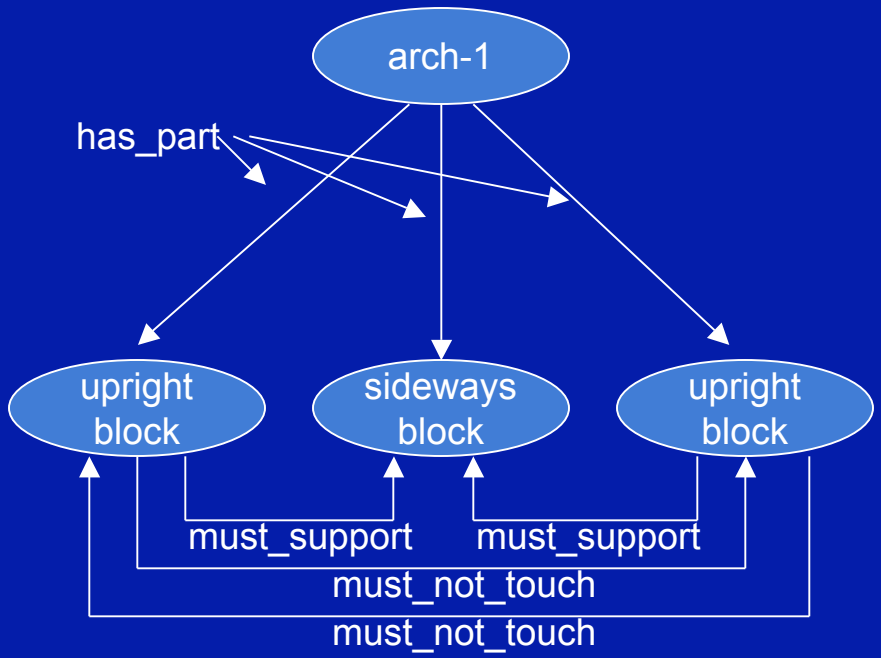

Example: Supervised learning of concept

Here's yet another training example, but this time it's a positive example. What's the difference between the new positive example and the current concept of an arch?

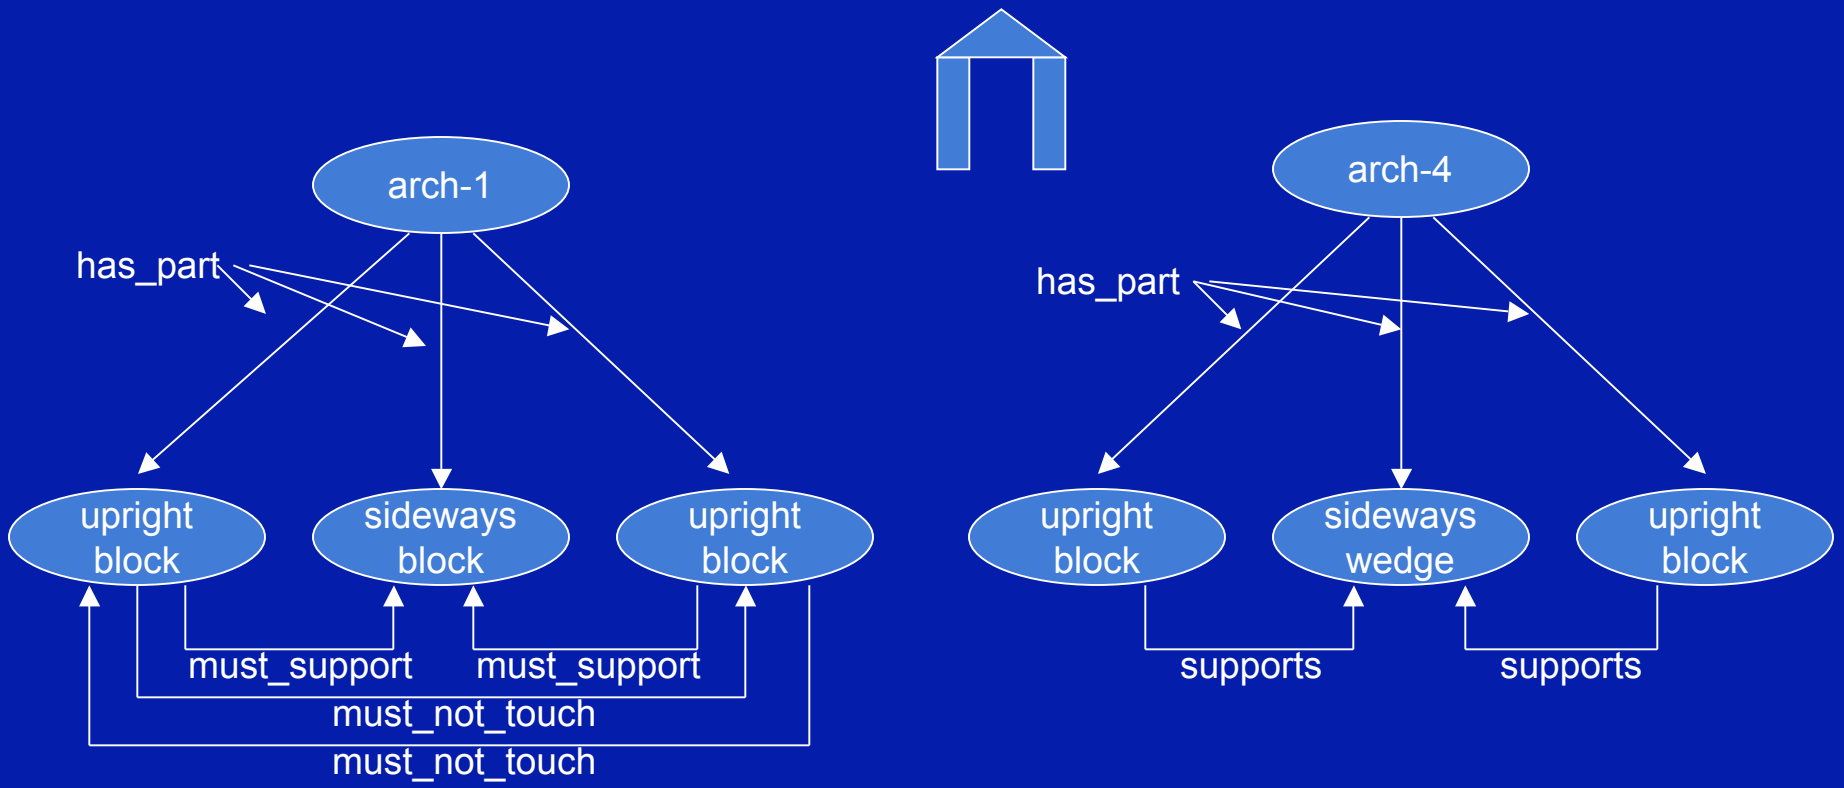

Example: Supervised learning of concept

The difference is that the block being supported has a different shape: it's a wedge. So the block being supported can be either a rectangular block or a wedge. The model of the arch is updated accordingly.

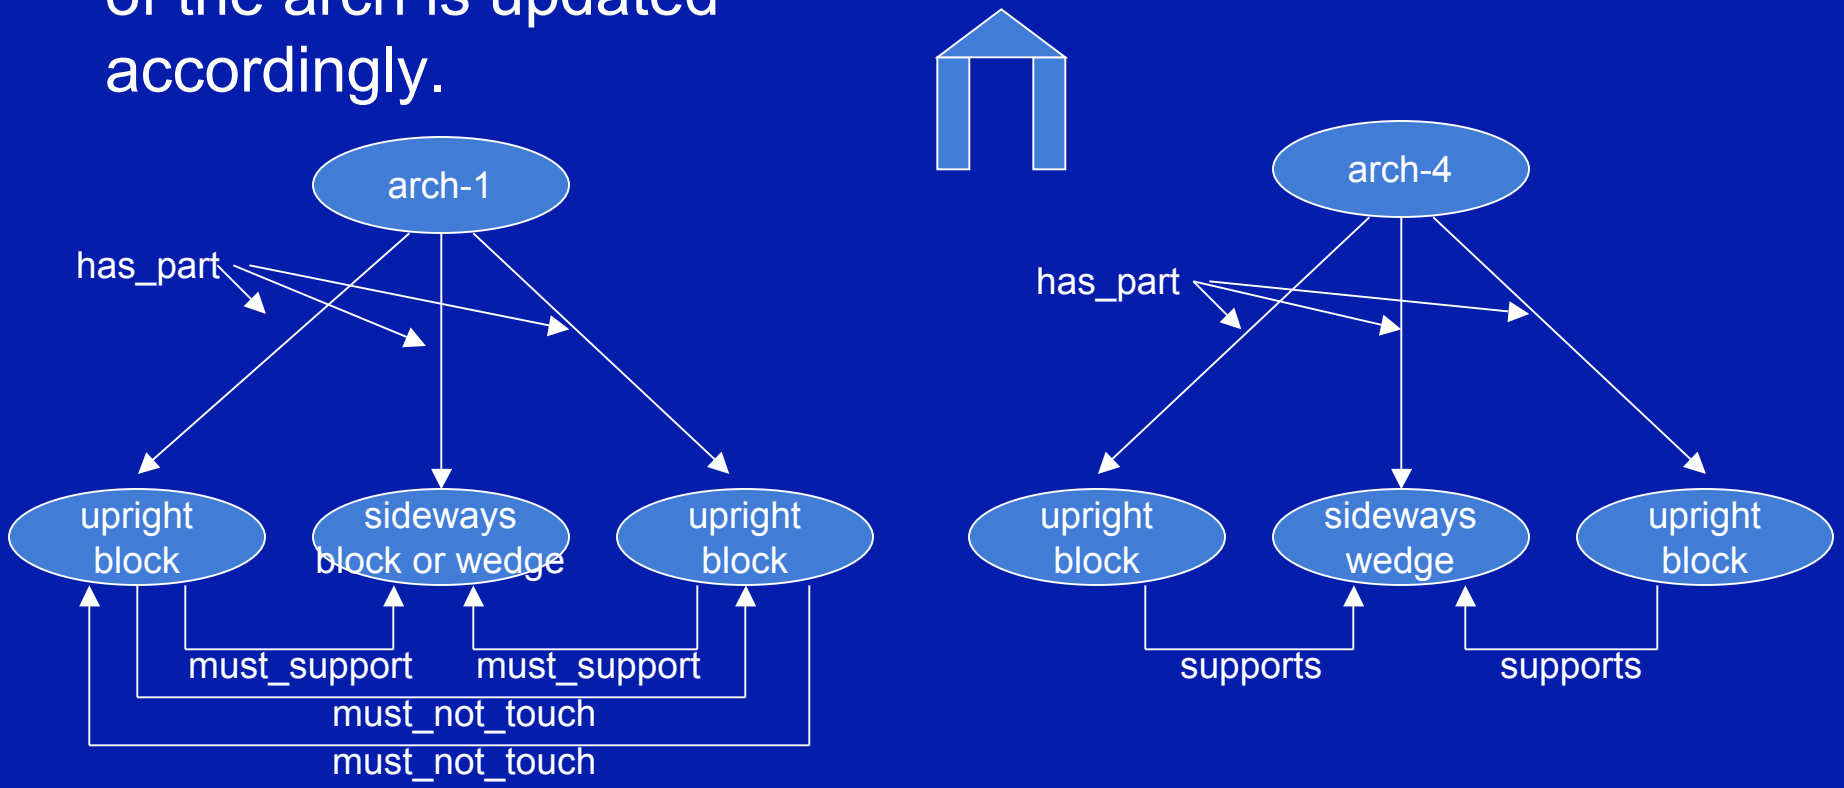

Example: Supervised learning of concept

Positive examples tell the learning procedure how to make its model more general, to cover more instances with the model.

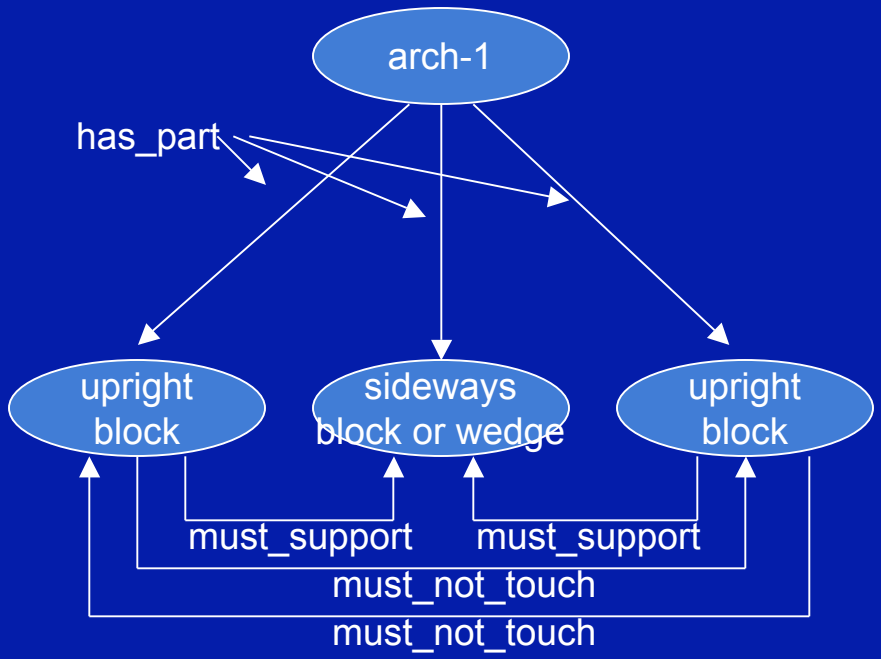

Example: Supervised learning of concept

If we take the program out of learning mode and ask it to classify a new input, what happens?

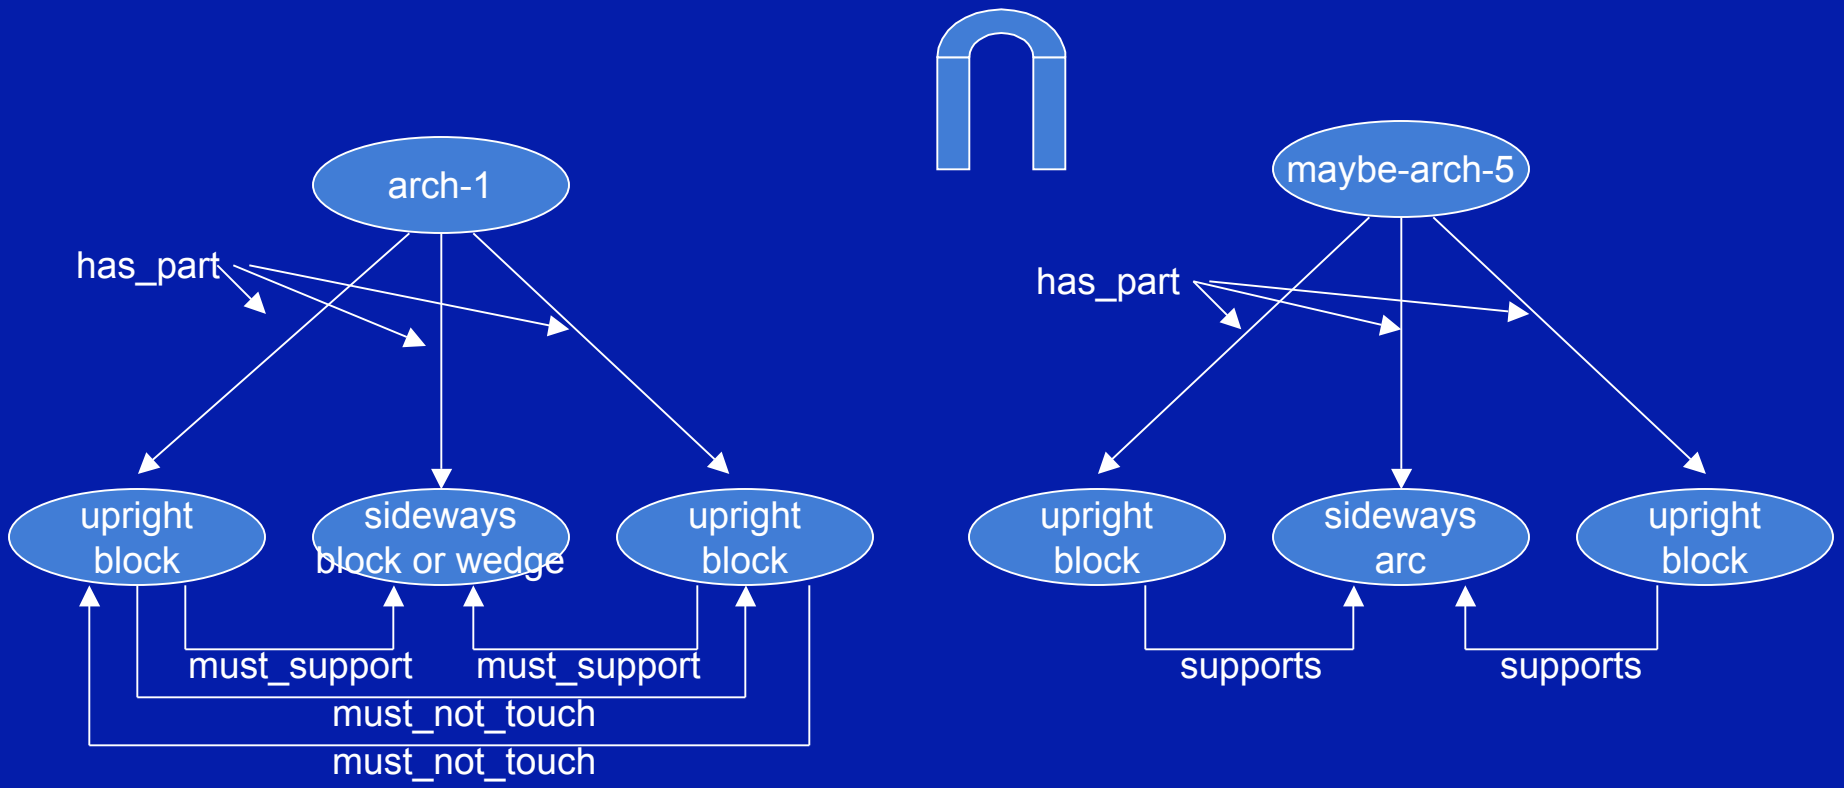

Example: Supervised learning of concept

Warning: Do not learn the wrong things from this example. It is not the case that negative examples are only about links and positive examples are only about nodes!

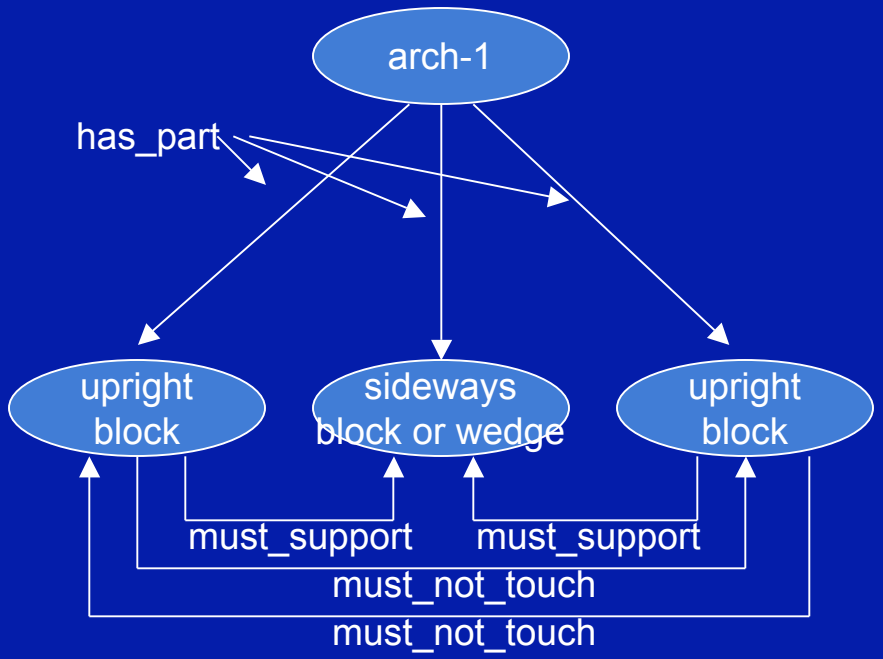

Example: Supervised learning of concept

This is a very simple model of one specific kind of learning, but it's easy to understand and easy to implement. That's one reason it's presented in just about every introductory AI course. But it also presents many of the issues that are common to all sorts of approaches to learning.

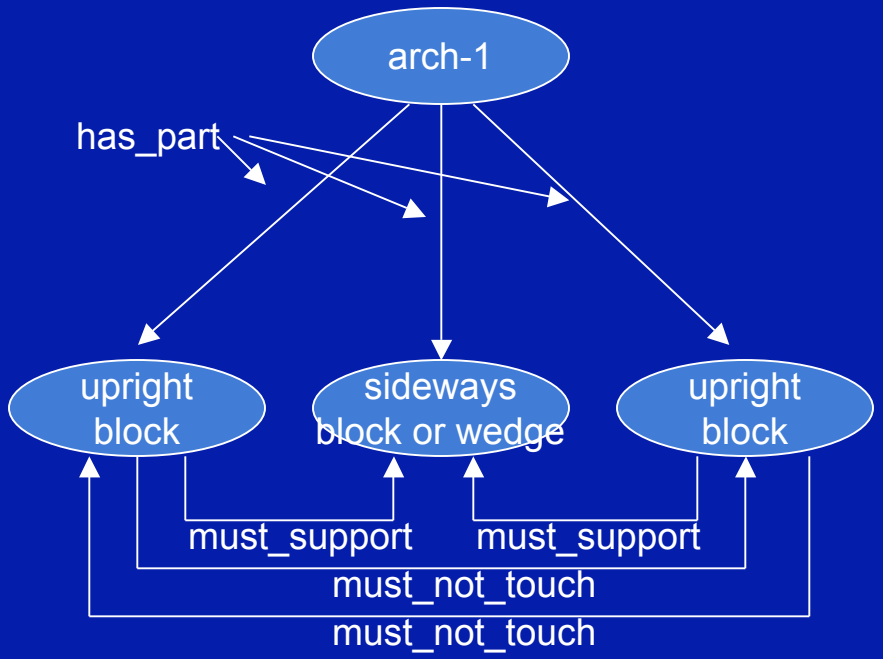

What's Minsky teaching to his class?

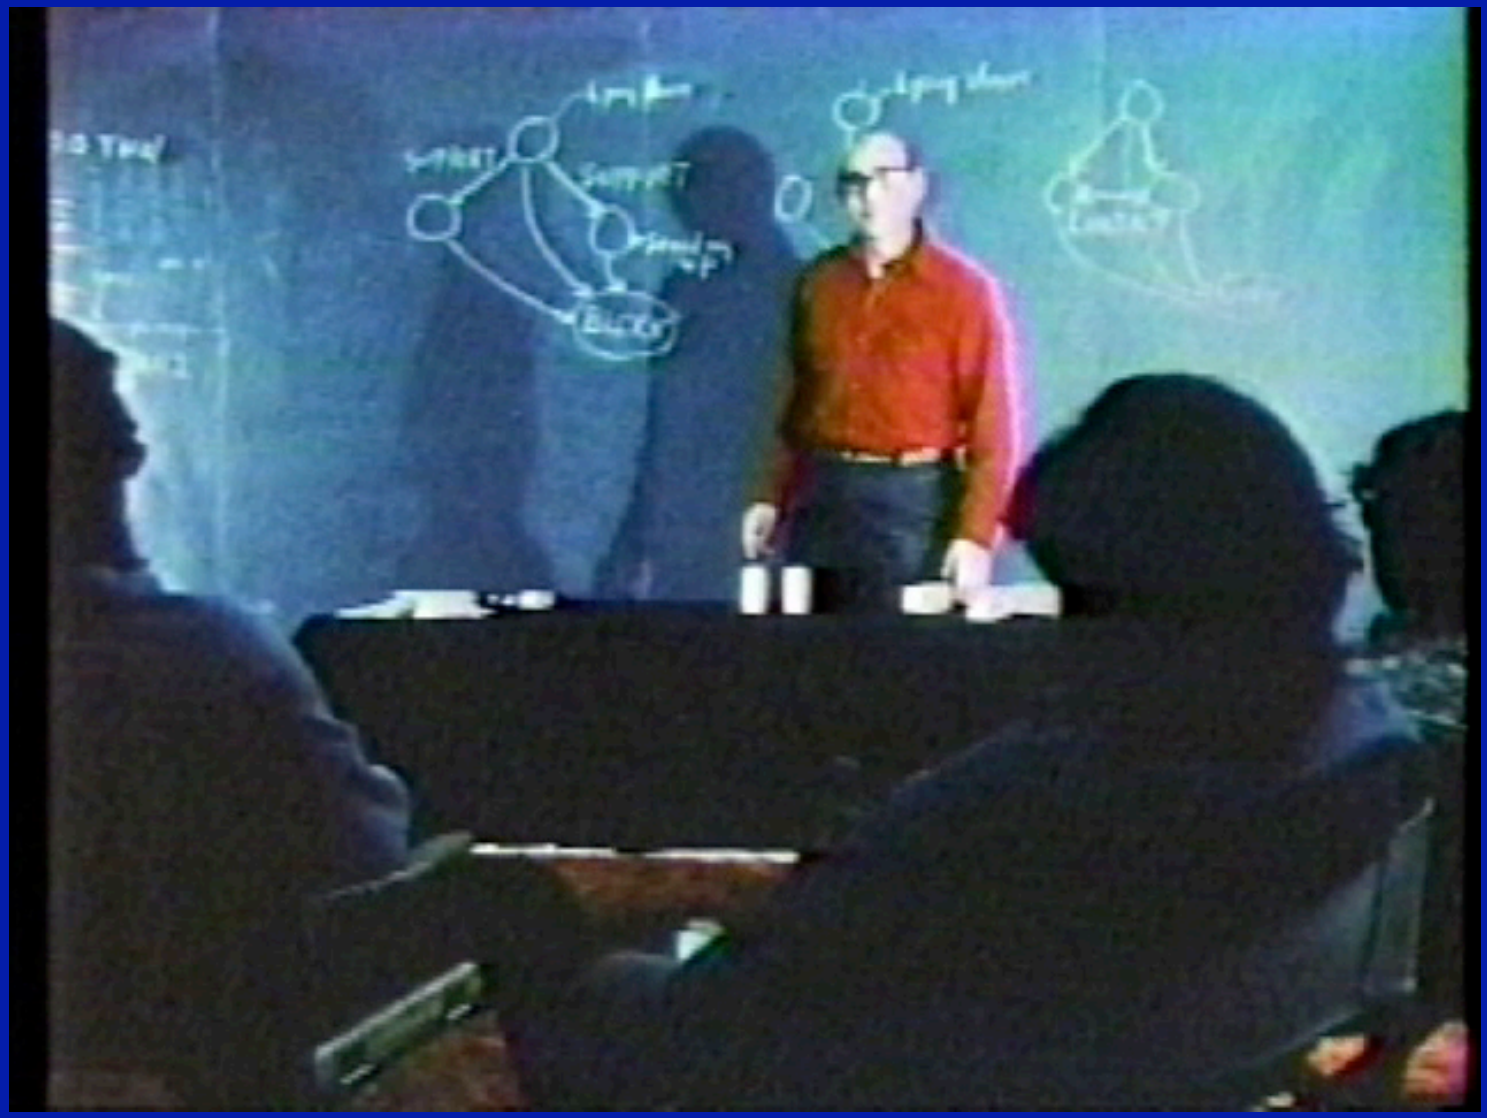

Example: Supervised learning of concept

Another way of looking at this is that the system is trying to gain the ability to make predictions beyond the given data. We've talked about this before, when we discussed inductive inference -- making generalizations from lots of examples.

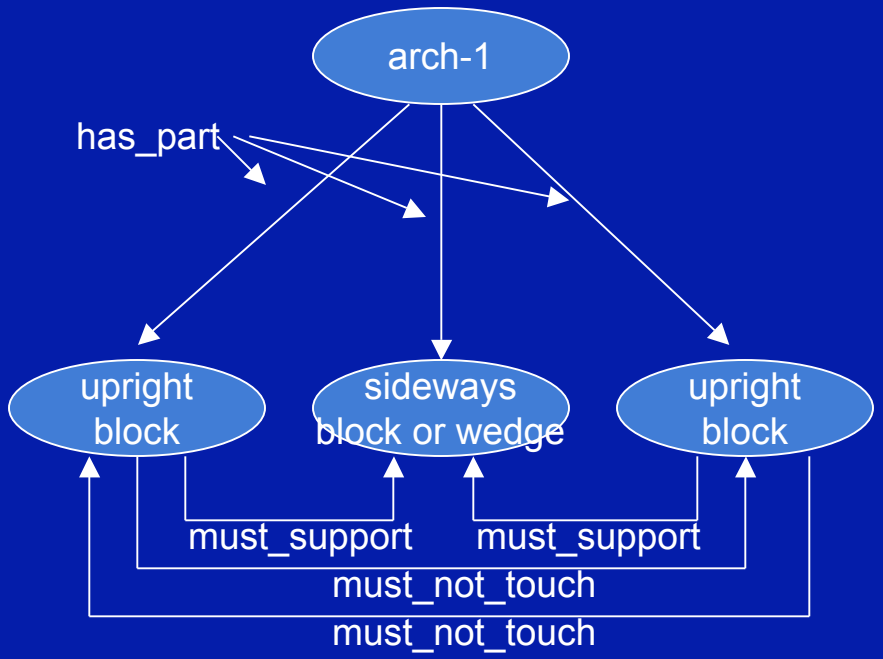

Inductive inference and learning by example

If you were a robot trying to learn when to cross the street and had seen lots of successful and unsuccessful examples of street crossings, how would you know what to pay attention to?

- width of street
- number of cars
- speed of cars
- weather
- daytime/nighttime?
- color of cars
- driver attributes
- type of cars
- trees along the street
- gas station on corner
- and countless others

Inductive inference and learning by example

While computers so far tend to be bad at figuring out for themselves what the salient features are, people are magically great at this sort of thing. It's a survival thing.

width of street
number of cars
speed of cars
weather
daytime/nighttime?
color of cars

driver attributes
type of cars
trees along the street
gas station on corner
and countless others

Learning is about choosing the best representation

That's certainly true in a logic-based AI world.

Our arch learner started with some internal representation of an arch. As examples were presented, the arch learner modified its internal representation to either make the representation accommodate positive examples (generalization) or exclude negative examples (specialization).

There's really nothing else the learner could modify... the reasoning system is what it is. So any learning problem can be mapped onto one of choosing the best representation...

Learning is about search

...but wait, there's more!

By now, you've figured out that the arch learner was doing nothing more than searching the space of possible representations, right?

So learning, like everything else, boils down to search.

If that wasn't obvious, you probably will want to do a little extra preparation for the final exam....

Same problem - different representation

The arch learner could have represented the arch concept as a decision tree if we wanted

Same problem - different representation

The arch learner could have represented the arch concept as a decision tree if we wanted

Same problem - different representation

The arch learner could have represented the arch concept as a decision tree if we wanted

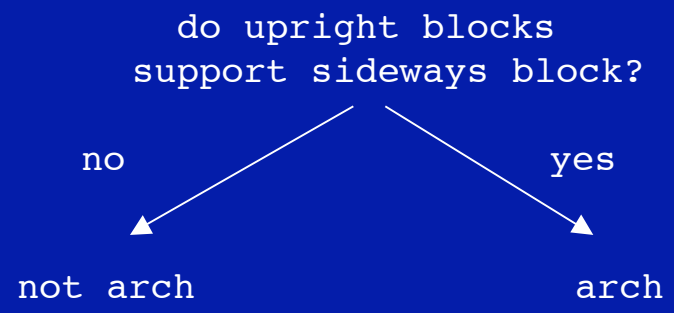

Same problem - different representation

The arch learner could have represented the arch concept as a decision tree if we wanted

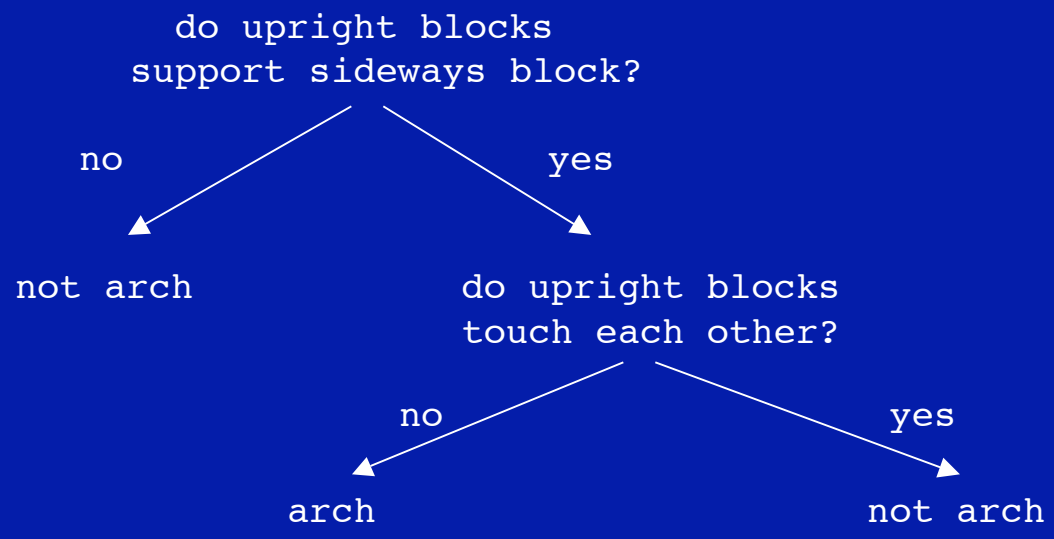

Same problem - different representation

The arch learner could have represented the arch concept as a decision tree if we wanted

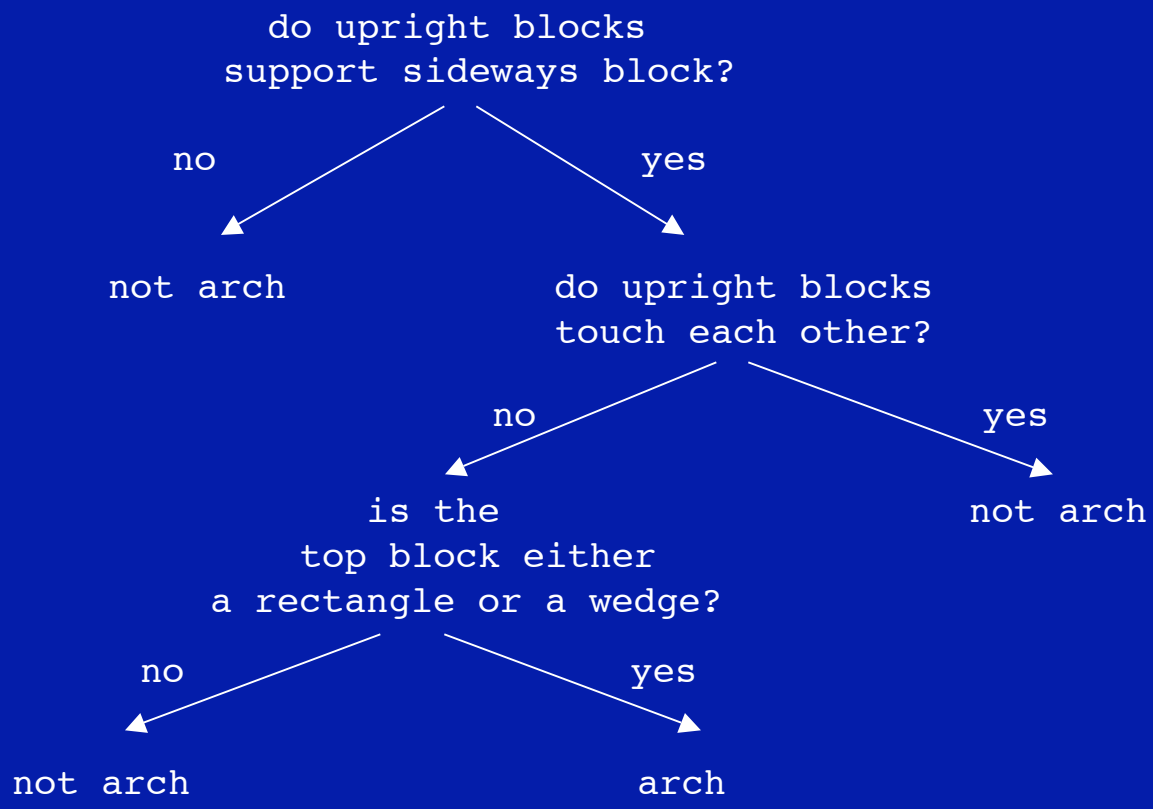

Other issues with learning by example

The learning process requires that there is someone to say which examples are positive and which are negative.

This approach must start with a positive example to specialize or generalize from.

Learning by example is sensitive to the order in which examples are presented.

Learning by example doesn’t work well with noisy, randomly erroneous data.