So can a CART model actually predict Supreme Court case outcomes better than a group of experts?

Martin and his colleagues used 628 previous Supreme Court cases between 1994 and 2001 to build their model.

They made predictions for the 68 cases that would be decided in October, 2002, before the term started.

Their model had two stages of CART trees.

The first stage involved making predictions using two CART trees.

One to predict a unanimous liberal decision and one to predict a unanimous conservative decision.

If the trees gave conflicting responses or both predicted no, then they moved on to the next stage.

It turns out that about 50% of Supreme Court cases result in a unanimous decision.

So this was a nice first step to detect the easier cases.

The second stage consisted of predicting the decision of each individual justice.

And then using the majority decision of all nine justices as a final prediction for the case.

In this lecture, we constructed the CART tree for Justice Stevens.

Here's a different tree, this one for Justice O'Connor.

The first split is whether or not the lower court decision is liberal.

If yes, then we predict that she will reverse the case.

This makes sense because Justice O'Connor is generally viewed as a conservative judge.

On the other hand, if the lower court decision is conservative, we check for the circuit court of origin.

If it's the second, third, DC, or federal court, we predict that she will affirm the case.

If it's not one of these courts, we move on to the next split.

The remaining two splits are for the respondent and the primary issue.

Here's another tree, this one for Justice Souter.

This shows an unusual property of the CART trees that Martin and his colleagues developed.

They use predictions for some trees as independent variables for other trees.

In this tree, the first split is whether or not Justice Ginsburg predicted decision is liberal.

So we have to run Justice Ginsburg's CART tree first.

See what the prediction is.

And then use that as input for Justice Souter's tree.

If Justice Ginsburg predicted decision is liberal and the lower court decision is liberal, then we predict that Justice Souter will affirm the case.

But if the lower court decision is conservative, then we predict that Justice Souter will reverse the case.

On the other side of the tree, if Justice Ginsburg predicted decision is conservative, but the lower court decision is liberal, then we predict that Justice Souter will reverse the case.

But if the lower court decision is conservative, then we predict that Justice Souter will affirm the case.

In summary, if we predict that Justice Ginsburg will make a liberal decision, then Justice Souter will probably make a liberal decision too.

But if we predict that Justice Ginsburg will make a conservative decision, then we predict that Justice Souter will probably make a conservative decision too.

Martin and his colleagues also recruited 83 legal experts, 71 academics, and 12 attorneys.

38 had previously clerked for a Supreme Court Justice, 33 were chaired professors, and five were current or former law school deans.

So this was really a dream team of experts.

Additionally, the experts were only asked to predict within their area of expertise.

So not all experts predicted all cases.

But there was more than one expert making predictions for each case.

When making their predictions, the experts were allowed to consider any source of information.

But they were not allowed to communicate with each other regarding the predictions.

For the 68 cases in October, 2002, the predictions were made, and at the end of the month the results were computed.

For predicting the overall decision that was made by the Supreme Court, the models had an accuracy of 75%, while the experts only had an accuracy of 59%.

So the models had a significant edge over the experts in predicting the overall case outcomes.

However, when the predictions were run for individual justices, the model and the experts performed very similarly.

For some justices, the model performed better.

And for some justices, the experts performed better.

Being able to predict Supreme Court decisions is very valuable to firms, politicians, and nongovernmental organizations.

We saw in this lecture that a model that predicts overall Supreme Court decisions is both more accurate than experts and can be run much faster than experts can make their predictions.

The CART models that we built were based on very high level components of the cases, compared to the experts who can process much more detailed and complex information.

This example really shows the edge that analytics can provide in traditionally qualitative applications.