The TAR2 Treatment Learner

Tim Menzies, Ying Hu
University of British Columbia
tim@menzies.com, yingh@ece.ubc.ca

September 7, 2001

1 Treatment Learning

Classical machine learners like C4.5 aim at discovering classification rules: i.e. given a classified training set, they output rules that are predictive of the class variable. TAR2 differs from those learners in that:

1. TAR2 assumes the classes are ordered by their *score* (some domain-specific measure). Highly scored classes are preferable to lower scoring classes. Further, one class is more desirable than all others, which is called the *best* class.

2. Rather than finding the classification rules, TAR2 finds rules that predict both increased frequency of the best class and decreased frequency of the worst class. That is, TAR2 finds discriminate rules that drive the system away from the worst class to the best class.

3. TAR2 output treatments rather than classifications. A treatment is one or a conjunction of attribute value ranges. It is a constraint on future control inputs of the system.

2 Golf Example

2.1 Input

TAR2 can best be introduced via example. Consider the log golf playing behavior shown in Figure 1. This log contains 4 attributes and 3 classes. Recall that TAR2 accesses a *score* for each class. For a golfer, the classes in Figure 1 could be scored as *none=2* (i.e. worst), *some=4*,

#	*outlook*	*temp(^oF)*	*humidity*	*windy?*	*class*
1	*sunny*	*85*	*86*	*false*	*none*
2	*sunny*	*80*	*90*	*true*	*none*
3	*sunny*	*72*	*95*	*false*	*none*
4	*rain*	*65*	*70*	*true*	*none*
5	*rain*	*71*	*96*	*true*	*none*
6	*rain*	*70*	*96*	*false*	*some*
7	*rain*	*68*	*80*	*false*	*some*
8	*rain*	*75*	*80*	*false*	*some*
9	*sunny*	*69*	*70*	*false*	*lots*
10	*sunny*	*75*	*70*	*true*	*lots*
11	*overcast*	*83*	*88*	*false*	*lots*
12	*overcast*	*64*	*65*	*true*	*lots*
13	*overcast*	*72*	*90*	*true*	*lots*
14	*overcast*	*81*	*75*	*false*	*lots*

Figure 1: A log of some golf-playing behavior.

lots=8 (i.e. best). Note that the preferred classes score exponentially higher than the non-preferred classes. As we shall see, this disproportionate weighting scheme strongly encourages TAR2 to chase the better classes.

2.2 The Mining Algorithm

TAR2 seeks attribute ranges that occur more frequently in the highly scored classes than in the lower scored classes. Let $a.r$ be some attribute range e.g. *outlook=overcast*) $\Delta_{a.r}$ is a heuristic measure of the worth of $a.r$ to improve the frequency of the $best$ class. $\Delta_{a.r}$ uses the following definitions:

$X(a.r)$**:** is the number of occurrences of that attribute range in class X; e.g. *lots(outlook.overcast)=4*.

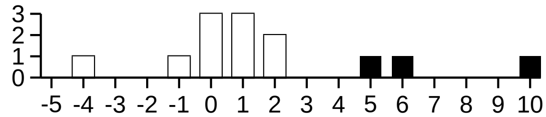

Figure 2: Δ distribution seen in golf data sets. Outstandingly high Δ values shown in black. Y-axis is the number of attribute ranges that have a particular Δ.

$all(a.r)$**:** is the total number of occurrences of that attribute range in all classes; e.g. *all(outlook.overcast)=4*.

$best$**:** the highest scoring class; e.g. $best = lots$;

$rest$**:** the non-best class; e.g. $rest = \{none, some\}$;

$score$**:** The score of a class X is $\$X$;.

$\Delta_{a.r}$ is calculated as follows:

$$\Delta_{a.r} = \frac{\sum_{X \in rest} (\$best - \$X) * (best(a.r) - X(a.r))}{all(a.r)}$$

When *a.r* is *outlook.overcast*, then $\Delta_{outlook.overcast}$ is calculated as follows:

$$\frac{\overbrace{((8-2)*(4-0))}^{lots \rightarrow none} + \overbrace{((8-4)*(4-0))}^{lots \rightarrow some}}{4+0+0} = \frac{40}{4} = 10$$

The attribute ranges in our golf example generate the Δ histogram shown in Figure 2. Note that *outlook=overcast*'s Δ is the highest, potentially most effective, attribute range.

2.3 Treatment

A *treatment* is a subset of the attribute ranges with an *outstanding* $\Delta_{a=r}$ value. For our golf example, such attributes can be seen in Figure 2: they are the outliers with outstandingly large Δs on the right-hand-side. (These outliers include *outlook=overcast*.

To *apply* a treatment, TAR2 rejects all example entries that contradict the conjunction of the attribute rages in the treatment. The ratio of classes in the remaining examples is compared to the ratio of classes in the original example set. The *best treatment* is the one that most increases the relative percentage of preferred classes. In our golf example, the best treatment is *outlook=overcast*; Figure 3 shows the class distribution before and after that treatment. i.e. if we bribe disc jockeys to always forecast overcast weather, then in 100% of cases, we should be playing lots of golf, all the time.

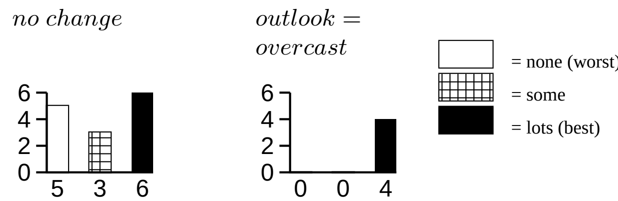

Figure 3: Finding treatments that can improve golf playing behavior. With no treatments, we only play golf lots of times in $\frac{6}{5+3+6} = 57\%$ of cases. With the restriction that *outlook=overcast*, then we play golf lots of times in 100% of cases.

3 Validation

The treatments learnt from TAR2 can be assessed using at least the following methods:

N-way cross-validation In this procedure, the available data is divided into N blocks so as to make each block's number of cases and class distribution as similar as possible. Next, N times, learning is performed on $\frac{N-1}{N}$ of the data and tested on the remaining $\frac{1}{N}$th of the data.

Model Feedback Ideally, the results can be applied to the model that generated the data. This is a more robust validation because the treatments are assessed by an outside device, avoided the effect of the training data.