

CART for letter-to-sound

worked example for the letter “a”

predictors							predictee
ppp	pp	p		n	nn	nnn	
o	g	w	a	s	h	-	aa
a	r	w	a	s	h	-	aa
-	-	w	a	r	r	a	ao
-	-	w	a	r	r	a	ao
r	d	w	a	r	e	-	eh
f	t	w	a	r	e	-	eh
-	-	w	a	r	e	s	eh

Entropy of distribution at the root

aa	ao	eh

$$H = - \sum_x p(x) \log(p(x))$$

$$p(\text{aa}) = 2 / 7 \approx 0.286$$

$$p(\text{ao}) = 2 / 7 \approx 0.286$$

$$p(\text{eh}) = 3 / 7 \approx 0.428$$

$$\log_2(2 / 7) = \log_2(2) - \log_2(7)$$

$$\approx 1.00 - 2.81 \approx -1.81$$

$$\log_2(3 / 7) = \log_2(3) - \log_2(7)$$

$$\approx 1.58 - 2.81 \approx -1.23$$

Entropy of distribution at the root

$$H = - \sum_x p(x) \log(p(x))$$

aa	ao	eh

$$\begin{aligned} H &\approx - (0.286 \times -1.81 \\ &\quad + 0.286 \times -1.81 \\ &\quad + 0.428 \times -1.23) \\ &\approx - (-0.518 - 0.518 - 0.526) \end{aligned}$$

$$H \approx 1.56 \text{ bits}$$

Example split: *Does letter n = "r" ?*

predictors							predictee	<i>Answer to question</i>
ppp	pp	p		n	nn	nnn		
o	g	w	a	s	h	-	aa	N
a	r	w	a	s	h	-	aa	N
-	-	w	a	r	r	a	ao	Y
-	-	w	a	r	r	a	ao	Y
r	d	w	a	r	e	-	eh	Y
f	t	w	a	r	e	-	eh	Y
-	-	w	a	r	e	s	eh	Y

Example split: *Does letter n = "r" ?*

predictors							predictee	<i>Answer to question</i>
ppp	pp	p	n	nn	nnn			
o	g	w	a	s	h	-	aa	N
a	r	w	a	s	h	-	aa	N

predictors							predictee	<i>Answer to question</i>
ppp	pp	p	n	nn	nnn			
-	-	w	a	r	r	a	ao	Y
-	-	w	a	r	r	a	ao	Y
r	d	w	a	r	e	-	eh	Y
f	t	w	a	r	e	-	eh	Y
-	-	w	a	r	e	s	eh	Y

predictors							predictee	<i>Answer to question</i>
ppp	pp	p		n	nn	nnn		
o	g	w	a	s	h	-	aa	N
a	r	w	a	s	h	-	aa	N

aa	ao	eh

$$H = - \sum_x p(x) \log(p(x))$$

$$H = - (1 \times 0 + 0 + 0)$$

$$H = 0 \text{ bits}$$

predictors							predictee	<i>Answer to question</i>
ppp	pp	p	n	nn	nnn			
-	-	w	a	r	r	a	ao	Y
-	-	w	a	r	r	a	ao	Y
r	d	w	a	r	e	-	eh	Y
f	t	w	a	r	e	-	eh	Y
-	-	w	a	r	e	s	eh	Y

aa	ao	eh

$$p(aa) = 0 / 5 = 0$$

$$p(ao) = 2 / 5 = 0.4$$

$$p(eh) = 3 / 5 \approx 0.6$$

$$\begin{aligned} \log_2(2 / 5) &= \log_2(2) - \log_2(5) \\ &\approx 1.00 - 2.32 \approx -1.32 \end{aligned}$$

$$\begin{aligned} \log_2(3 / 7) &= \log_2(3) - \log_2(7) \\ &\approx 1.58 - 2.81 \approx -1.23 \end{aligned}$$

predictors							predictee	<i>Answer to question</i>
ppp	pp	p	n	nn	nnn			
-	-	w	a	r	r	a	ao	Y
-	-	w	a	r	r	a	ao	Y
r	d	w	a	r	e	-	eh	Y
f	t	w	a	r	e	-	eh	Y
-	-	w	a	r	e	s	eh	Y

aa	ao	eh

$$H = - \sum_x p(x) \log(p(x))$$

$$\begin{aligned}
 H &\approx - (0 \\
 &\quad + 0.4 \times -1.32 \\
 &\quad + 0.6 \times -0.74) \\
 &\approx - (0 - 0.528 - 0.444)
 \end{aligned}$$

$$H \approx 0.972 \text{ bits}$$

Total entropy before split

predictors							predictee
ppp	pp	p	n	nn	nnn		
o	g	w	a	s	h	-	aa
a	r	w	a	s	h	-	aa
-	-	w	a	r	r	a	ao
-	-	w	a	r	r	a	ao
r	d	w	a	r	e	-	eh
f	t	w	a	r	e	-	eh
-	-	w	a	r	e	s	eh

$H \approx 1.56$ bits

total $H \approx 1.56$ bits

Entropy after split *Does letter n = "r" ?*

predictors							predictee	<i>Answer to question</i>
ppp	pp	p	n	nn	nnn			
o	g	w	a	s	h	-	aa	N
a	r	w	a	s	h	-	aa	N

H = 0 bits

predictors							predictee	<i>Answer to question</i>
ppp	pp	p	n	nn	nnn			
-	-	w	a	r	r	a	ao	Y
-	-	w	a	r	r	a	ao	Y
r	d	w	a	r	e	-	eh	Y
f	t	w	a	r	e	-	eh	Y
-	-	w	a	r	e	s	eh	Y

H ≈ 0.972 bits

Information gain

- We now compute how much the entropy has been decreased by this split
- This is the **weighted** sum of the entropies in each partition
 - the weights are based on the number of data points in each partition
 - N branch: $2/7 \times 0$ bits
 - Y branch: $5/7 \times 0.972$ bits = 0.694 bits
 - total = $0 + 0.694 = 0.694$ bits
- Before split: 1.56 bits
- After split: 0.694 bits
- Information gain = 0.866 bits

in the lecture we simplified this and used the unweighted sum

Choosing the best question

- Just try them all, and pick the one with the highest information gain
- Try working through some of the other questions yourself, then make a table like this:

question	information gain (bits)
n = "a" ?	
n = "b" ?	
....	
n = "r" ?	0.866
n = "s" ?	
...	
p = "a" ?	
p = "b" ?	
...	