Toy Example

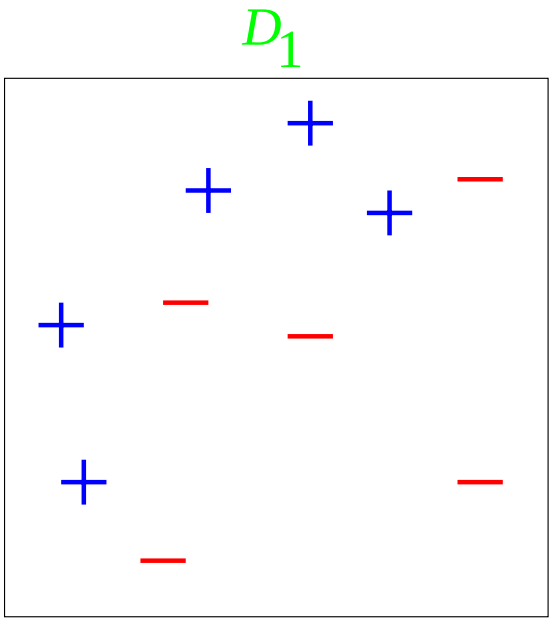

weak classifiers = vertical or horizontal half-planes

Round 1

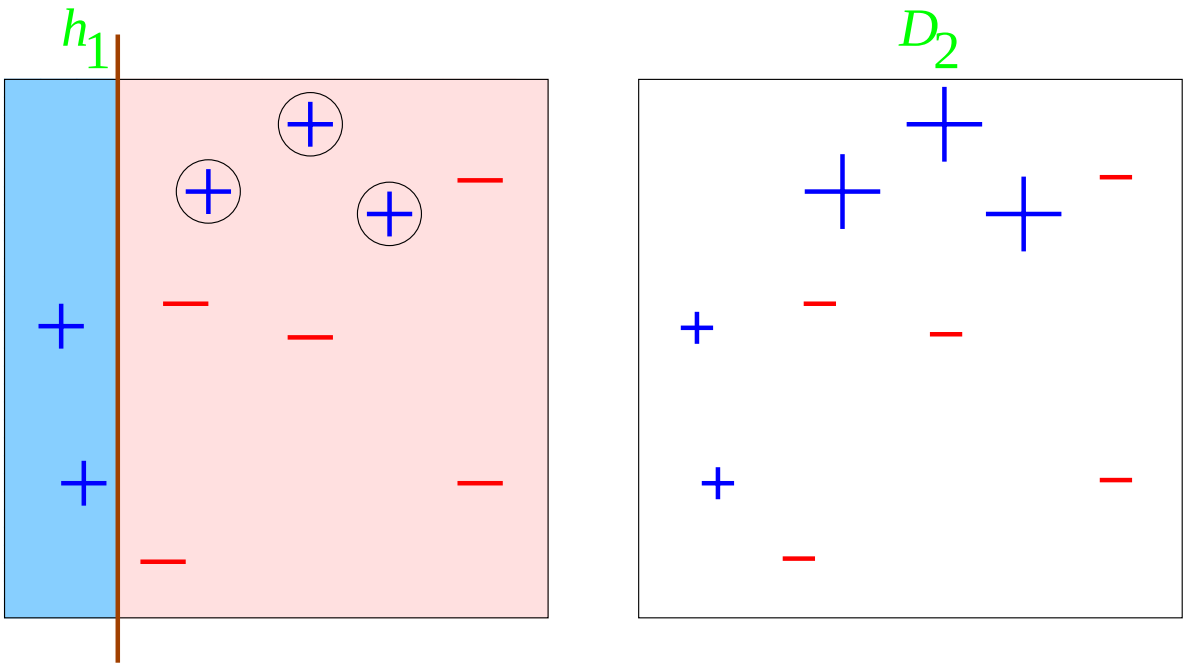

$\varepsilon_1 = 0.30$

$\alpha_1 = 0.42$

Round 2

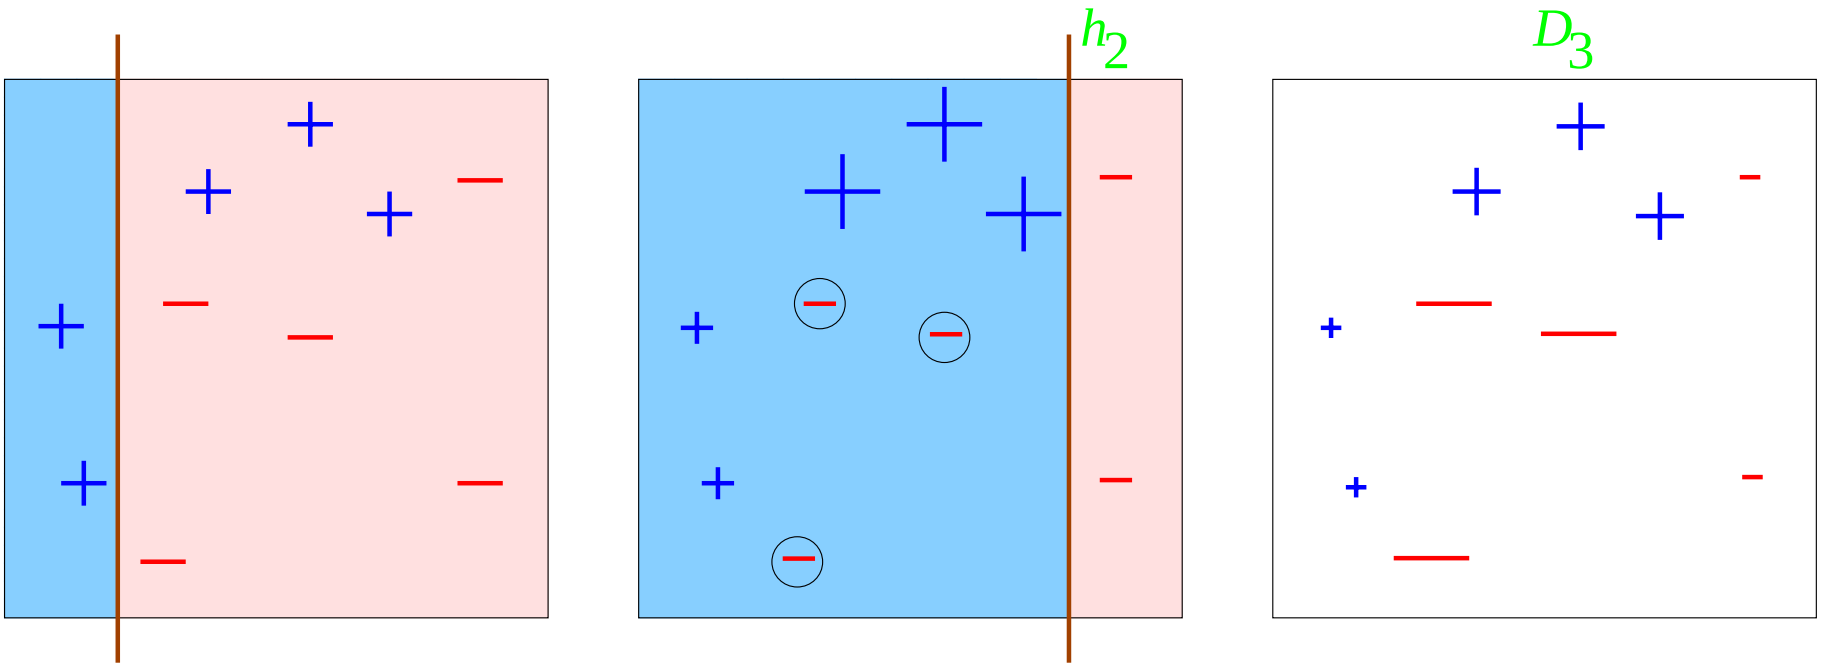

$\varepsilon_2 = 0.21$

$\alpha_2 = 0.65$

Round 3

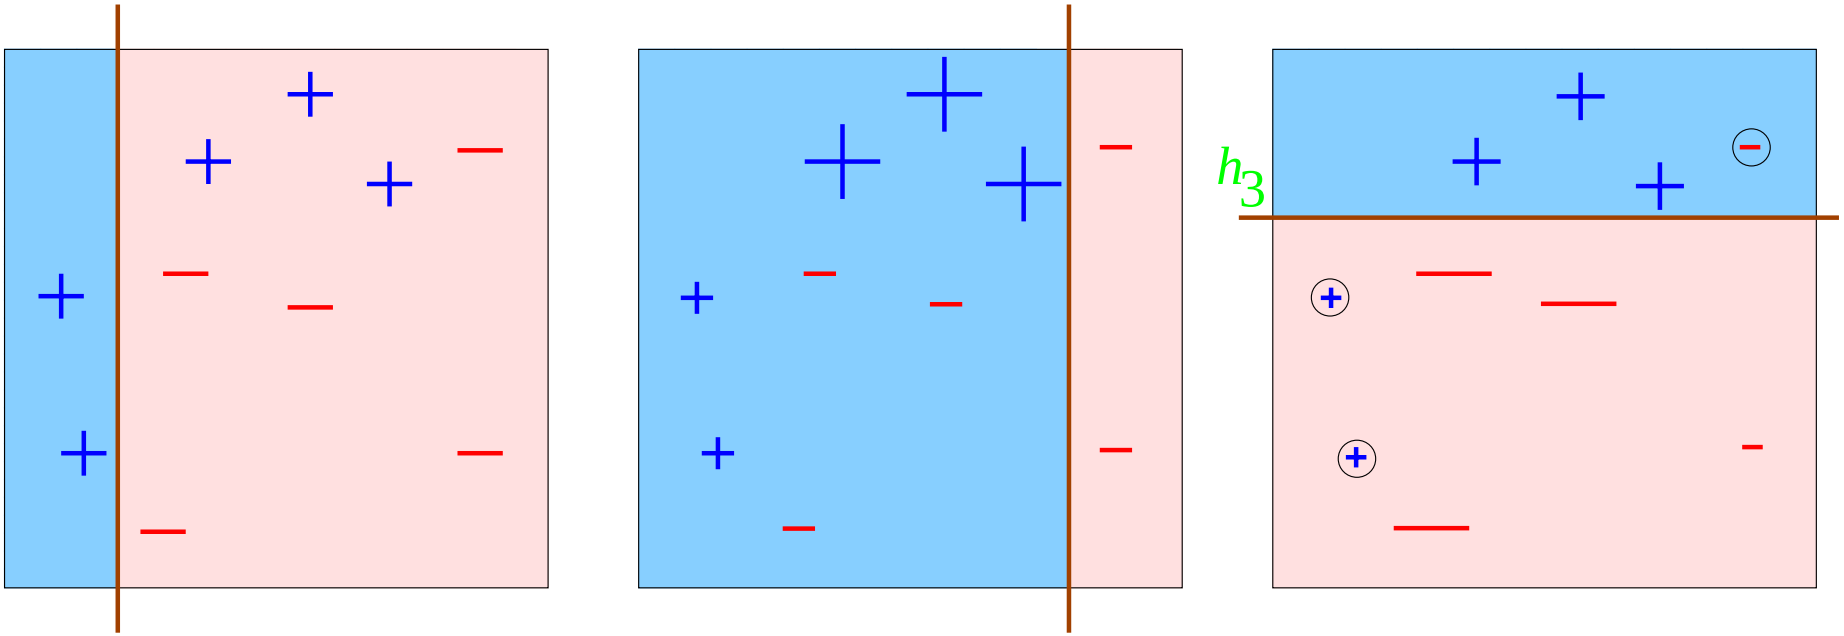

$\varepsilon_3 = 0.14$

$\alpha_3 = 0.92$

Final Classifier

$H_{\text{final}} = \text{sign}\left(0.42 \right.$ 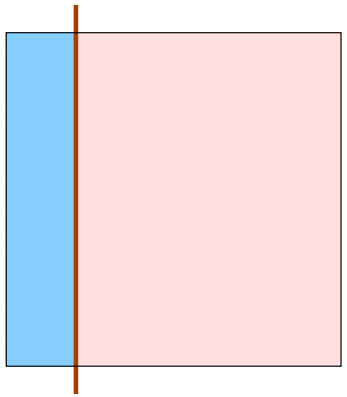 $+ 0.65$ 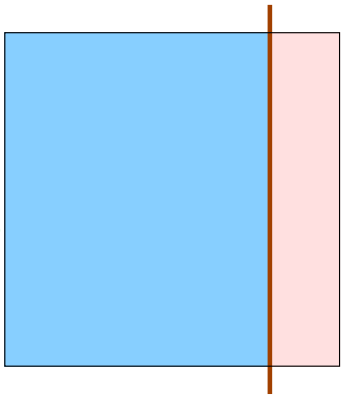 $+ 0.92$ 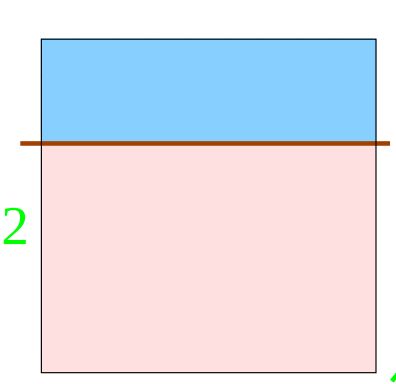 $\left.\right)$

$=$ 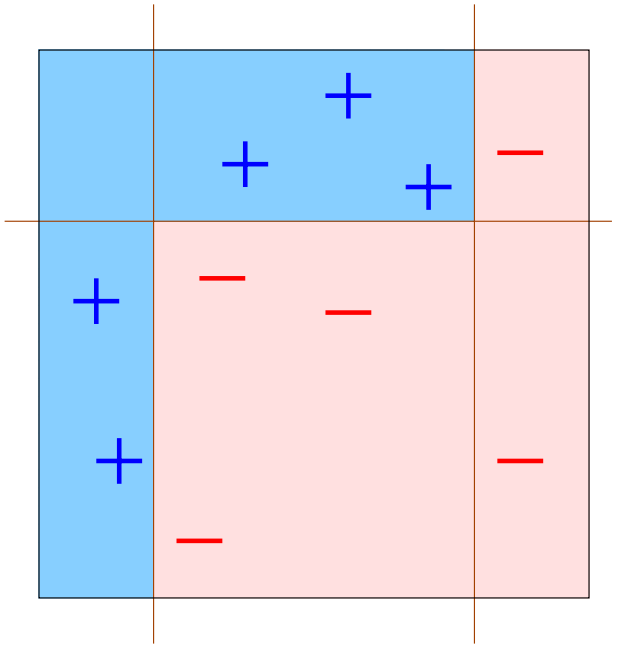

Actual Typical Run

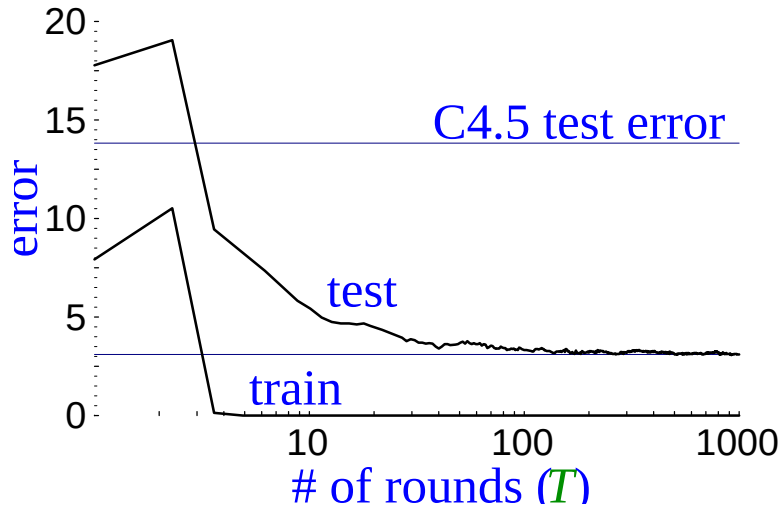

(boosting C4.5 on "letter" dataset)

- test error does not increase, even after 1000 rounds
  - (total size $>$ 2,000,000 nodes)
- test error continues to drop even after training error is zero!

	# rounds		
	5	100	1000
train error	0.0	0.0	0.0
test error	8.4	3.3	3.1

- Occam's razor wrongly predicts "simpler" rule is better

Empirical Evidence: The Margin Distribution

- margin distribution
  = cumulative distribution of margins of training examples

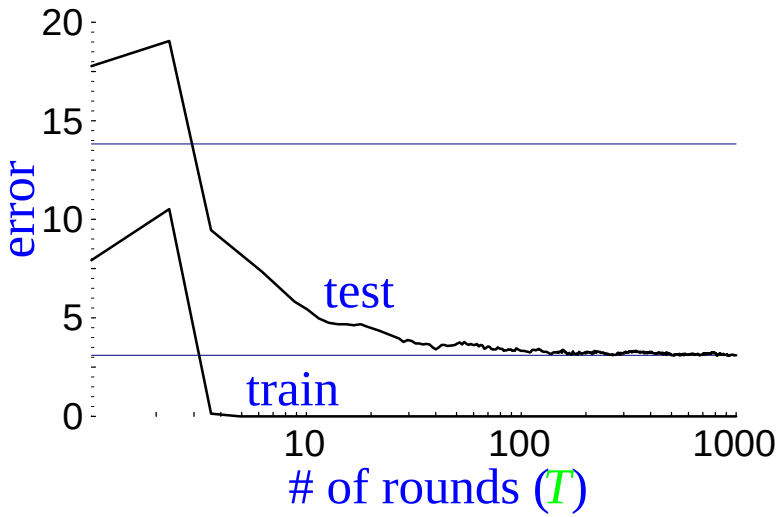

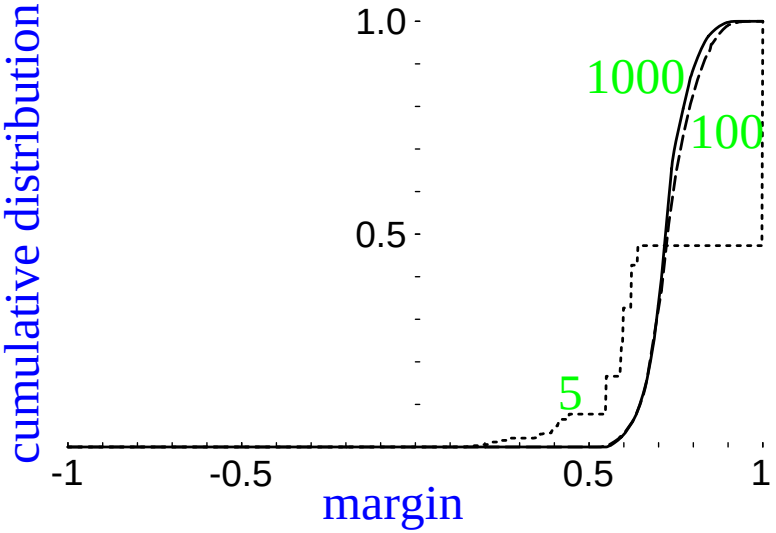

	# rounds 5	100	1000
train error	0.0	0.0	0.0
test error	8.4	3.3	3.1
% margins ≤ 0.5	7.7	0.0	0.0
minimum margin	0.14	0.52	0.55