A Stopping Criterion for Transductive Active Learning

Daniel Kottke[1 [0000−0002−7870−6033]] (✉), Christoph Sandrock[1], Georg Krempl[2 [0000−0002−4153−2594]], and Bernhard Sick[1 [0000−0001−9467−656X]]

[1] University of Kassel, Wilhelmshöher Allee 73, 34121 Kassel, Germany
{daniel.kottke,christoph.sandrock,bsick}@uni-kassel.de
[2] Utrecht University, Princetonplein 5, 3584 CC Utrecht, The Netherlands
g.m.krempl@uu.nl

Abstract. In transductive active learning, the goal is to determine the correct labels for an unlabeled, known dataset. Therefore, we can either ask an oracle to provide the right label at some cost or use the prediction of a classifier which we train on the labels acquired so far. In contrast, the commonly used (inductive) active learning aims to select instances for labeling out of the unlabeled set to create a generalized classifier, which will be deployed on unknown data. This article formally defines the transductive setting and shows that it requires new solutions. Additionally, we formalize the theoretically cost-optimal stopping point for the transductive scenario. Building upon the probabilistic active learning framework, we propose a new transductive selection strategy that includes a stopping criterion and show its superiority.

Keywords: Stopping Criteria · Active Learning · Transduction.

A Experimental Evaluation - Full Results

A.1 Hypothesis A

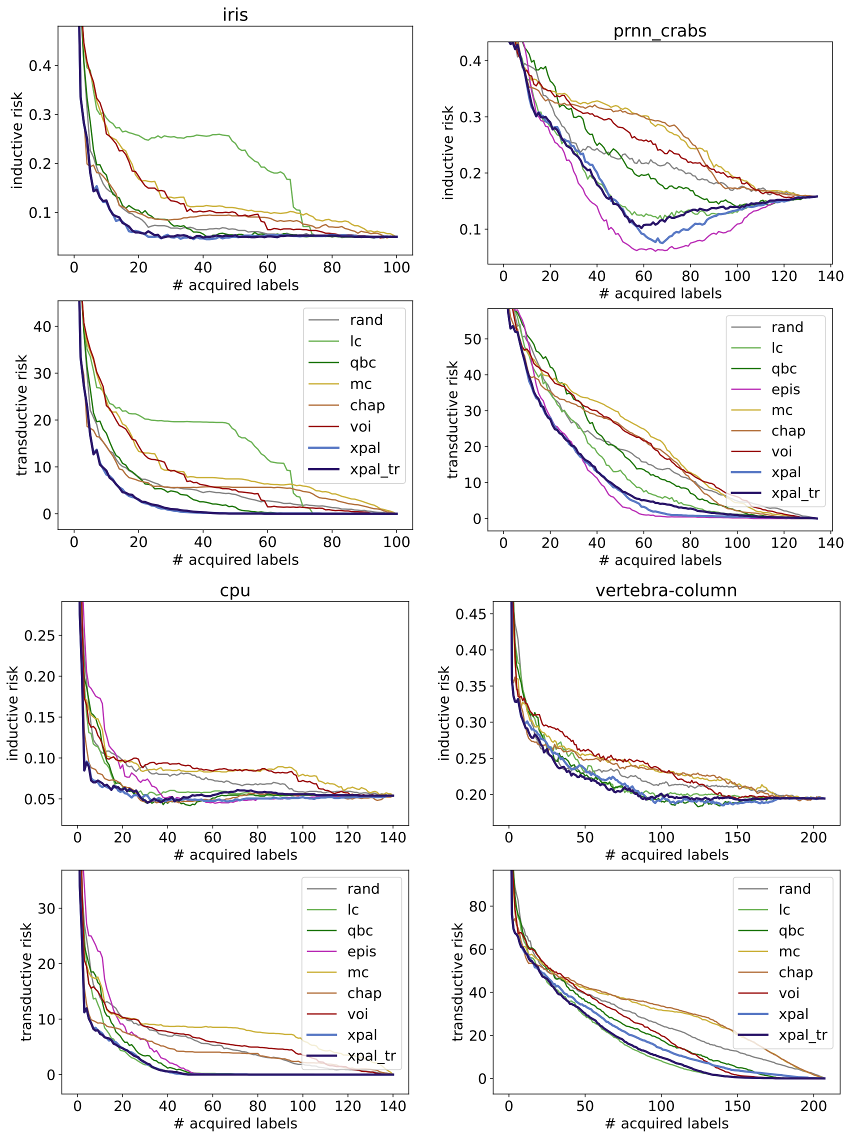

Fig. 1. Learning curves of selection strategies with respect to the inductive (upper) and the transductive (lower) risk.

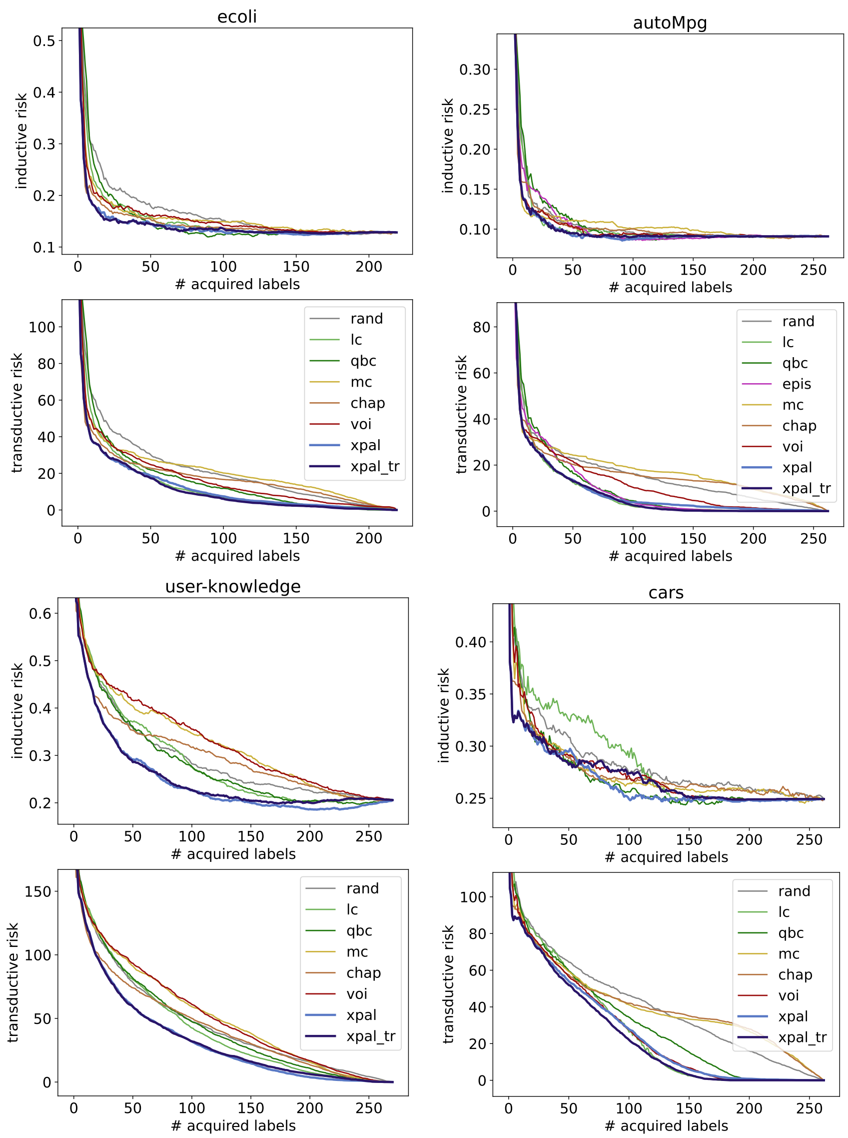

Fig. 2. Learning curves of selection strategies with respect to the inductive (upper) and the transductive (lower) risk.

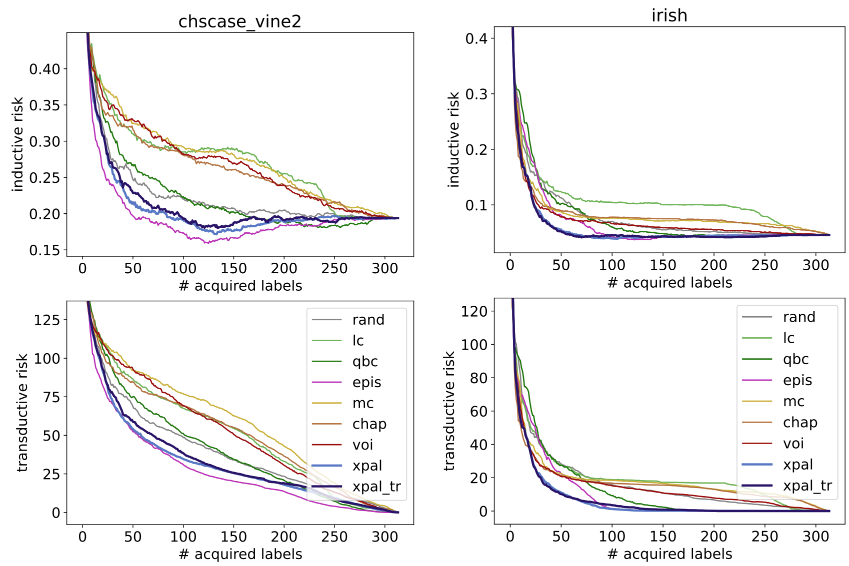

Fig. 3. Learning curves of selection strategies with respect to the inductive (upper) and the transductive (lower) risk.

	iris	prnn_crabs	cpu	vertebra-column	ecoli	autoMpg	user-knowledge	cars	chscase_vine2	irish	mean	win/tie/loss
rand	4.0↗	6.4↗	7.2↗	4.6↗	6.7↗	5.9↗	5.0↗	6.0↗	4.7↗	6.4↗	5.69	10 / 0 / 0
lc	6.0↗	3.1	5.5↗	4.4↗	4.9↗	4.9↗	4.1↗	6.2↗	6.9↗	6.0↗	5.19	9 / 1 / 0
qbc	4.7↗	5.9↗	4.9↗	3.5	4.9↗	6.4↗	3.9↗	3.6	4.6↗	5.8↗	4.81	8 / 2 / 0
epis	nan	1.6↘	5.5↗	nan	nan	5.0↗	nan	nan	1.2↘	4.3↗	3.52	3 / 0 / 2
mc	6.2↗	7.5↗	6.4↗	5.9↗	5.7↗	5.7↗	7.0↗	4.4↗	7.6↗	7.2↗	6.36	10 / 0 / 0
chap	4.5↗	7.8↗	3.2	5.6↗	4.0↗	5.4↗	5.4↗	4.4↗	7.0↗	6.7↗	5.40	9 / 1 / 0
voi	5.8↗	6.8↗	6.4↗	6.0↗	4.9↗	4.2	7.0↗	4.5↗	7.0↗	5.2↗	5.78	9 / 1 / 0
xpal	2.3	2.8	2.4	3.0	2.3	3.6	1.5	3.0	2.9	1.8	2.56	baseline
xpal_tr	2.6	3.2↗	3.4↗	3.0	2.7	3.9	2.0↗	3.9↗	3.2	1.7	2.95	4 / 6 / 0

Fig. 4. Ranking statistics for the area under the inductive risk.

	iris	prnn_crabs	cpu	vertebra-column	ecoli	autoMpg	user-knowledge	cars	chscase_vine2	irish	mean	win/tie/loss
rand	5.1↗	6.4↗	7.5↗	6.1↗	7.5↗	7.6↗	5.4↗	7.0↗	4.7↗	6.9↗	6.43	10 / 0 / 0
lc	5.6↗	4.2↗	3.2↗	1.4	3.0↗	2.2	4.0↗	2.9↗	6.5↗	5.9↗	3.89	8 / 2 / 0
qbc	4.3↗	6.4↗	4.4↗	4.4↗	4.6↗	4.3↗	5.0↗	4.8↗	4.9↗	5.0↗	4.81	10 / 0 / 0
epis	nan	2.2	5.6↗	nan	nan	4.1↗	nan	nan	1.3↘	3.6↗	3.36	3 / 1 / 1
mc	6.7↗	7.6↗	8.3↗	7.4↗	7.2↗	8.2↗	7.2↗	6.9↗	8.0↗	7.8↗	7.54	10 / 0 / 0
chap	5.4↗	6.6↗	6.0↗	7.4↗	5.6↗	7.9↗	4.4↗	7.0↗	7.2↗	6.9↗	6.44	10 / 0 / 0
voi	5.5↗	7.2↗	6.3↗	4.5↗	5.0↗	5.7↗	7.0↗	3.1↗	7.0↗	5.8↗	5.71	10 / 0 / 0
xpal	1.7	1.7↘	1.7	3.0↗	2.0↗	3.2↗	1.5	3.0↗	2.4	1.0↘	2.14	4 / 4 / 2
xpal_tr	1.7	2.7	1.9	1.7	1.1	1.8	1.6	1.3	3.0	2.0	1.87	baseline

Fig. 5. Ranking statistics for the area under the transductive risk.

A.2 Hypothesis B

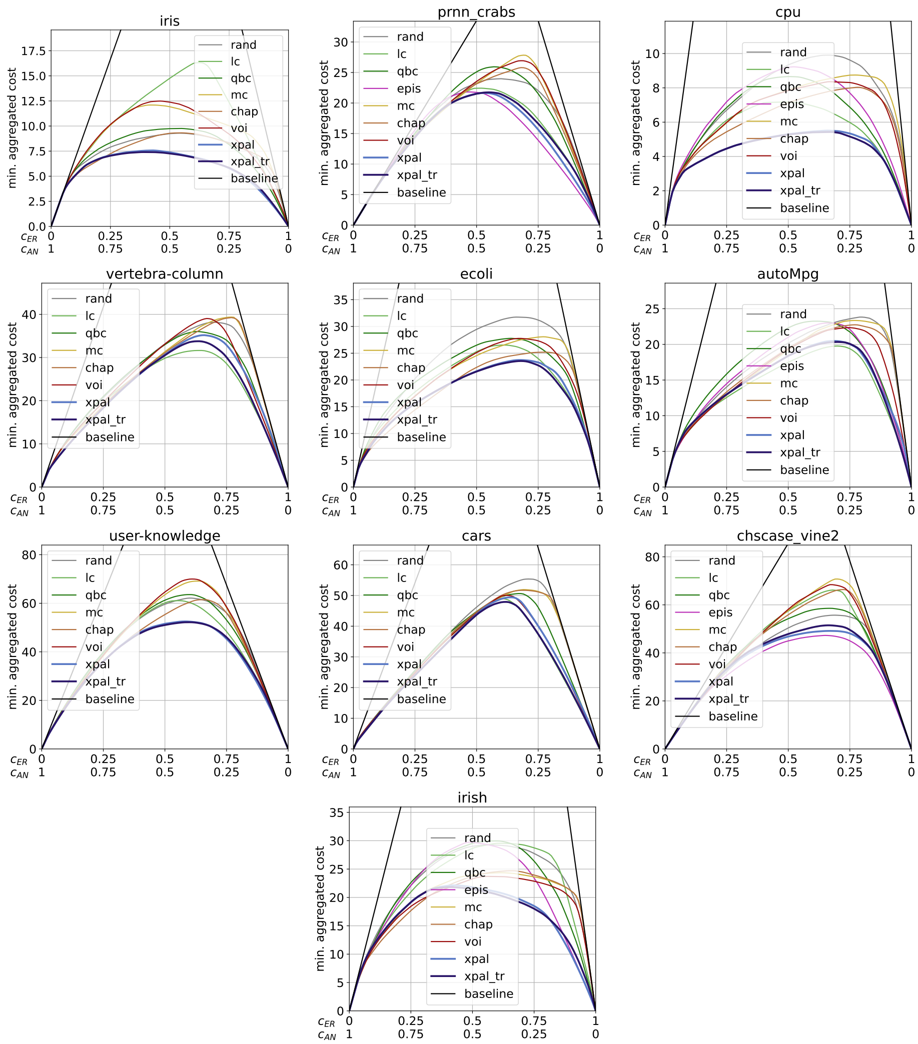

Fig. 6. Comparison of the minimum aggregated cost curves for all selection strategies.

	iris	prnn_crabs	cpu	vertebra-column	ecoli	autoMpg	user-knowledge	cars	chscase_vine2	irish	mean	win/tie/loss
rand	5.0↗	6.4↗	6.6↗	6.7↗	7.5↗	6.8↗	5.6↗	7.4↗	4.7↗	7.4↗	6.41	10 / 0 / 0
lc	5.5↗	4.0	4.2↗	2.0	4.4↗	2.6	4.5↗	2.6↗	6.6↗	6.4↗	4.28	7 / 3 / 0
qbc	4.6↗	6.5↗	5.4↗	5.4↗	5.9↗	6.6↗	5.5↗	5.0↗	5.8↗	6.8↗	5.75	10 / 0 / 0
epis	nan	2.2↘	6.4↗	nan	nan	5.5↗	nan	nan	1.2↘	6.2↗	4.30	3 / 0 / 2
mc	6.0↗	6.9↗	6.7↗	6.0↗	5.6↗	5.8↗	6.6↗	6.4↗	7.1↗	4.9↗	6.21	10 / 0 / 0
chap	4.9↗	6.2↗	5.3↗	5.4↗	3.3↗	6.1↗	3.8↗	6.4↗	6.7↗	4.7↗	5.30	10 / 0 / 0
voi	5.7↗	6.8↗	5.5↗	5.2↗	5.1↗	4.7↗	6.6↗	3.3↗	6.8↗	4.6↗	5.43	10 / 0 / 0
xpal	2.1	2.4↘	2.4	3.4↗	2.6↗	3.8↗	1.6	3.2↗	2.7↘	1.9	2.61	4 / 4 / 2
xpal_tr	2.2	3.5	2.6	1.9	1.7	3.1	1.8	1.6	3.3	2.1	2.37	baseline

Fig. 7. Ranking statistics for the area under minimum aggregated cost curves.

A.3 Hypothesis C

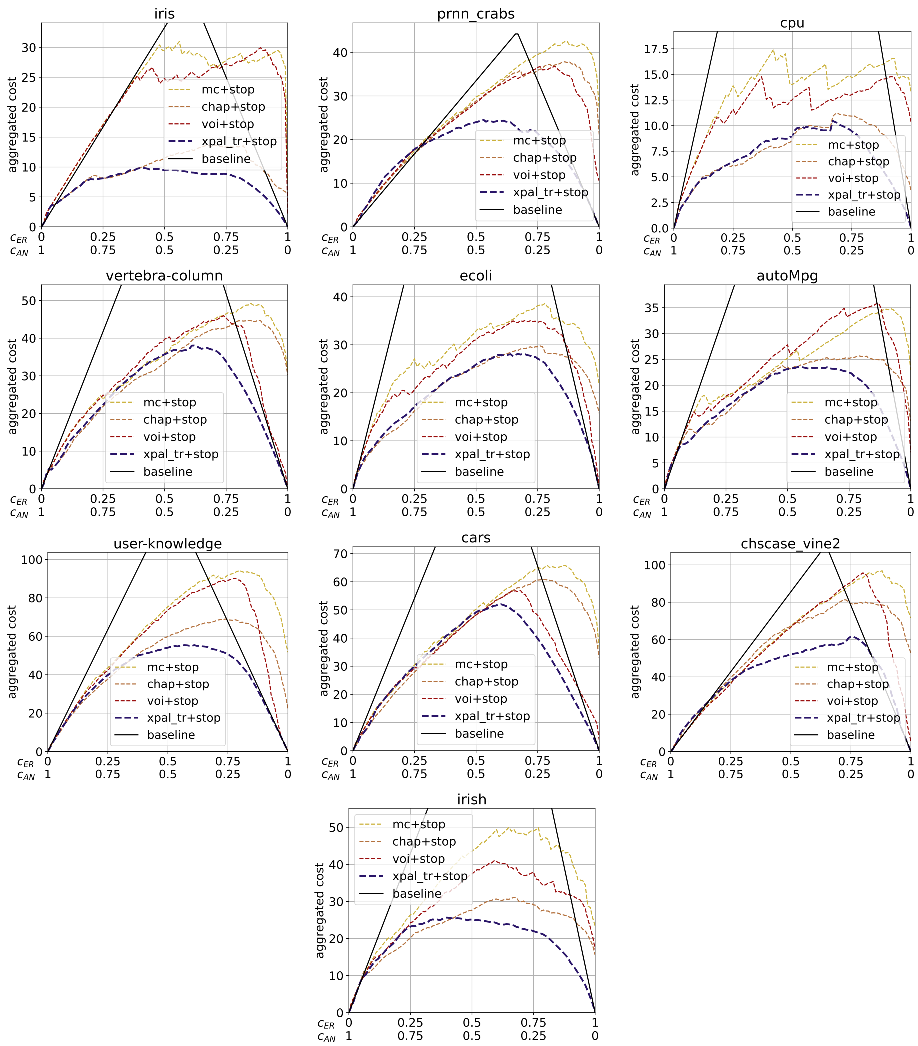

Fig. 8. Comparison of the aggregated cost curves for methods that implement their own stopping criterion.

	iris	prnn_crabs	cpu	vertebra-column	ecoli	autoMpg	user-knowledge	cars	chscase_vine2	irish	mean	win/tie/loss
mc+stop	3.9↗	3.6↗	3.6↗	3.8↗	3.8↗	3.4↗	4.0↗	3.7↗	3.5↗	3.6↗	3.67	10 / 0 / 0
chap+stop	2.0↗	2.9↗	2.0↗	2.9↗	2.2↗	2.6↗	2.2↗	3.2↗	3.0↗	2.2↗	2.53	10 / 0 / 0
voi+stop	3.1↗	2.5↗	2.9↗	2.3↗	2.7↗	2.5↗	2.8↗	1.7↗	2.5↗	3.1↗	2.61	10 / 0 / 0
xpal_tr+stop	1.0	1.0	1.5	1.1	1.2	1.4	1.0	1.4	1.0	1.1	1.18	baseline

Fig. 9. Ranking statistics for the area under aggregated cost curves for methods that implement their own stopping criterion.

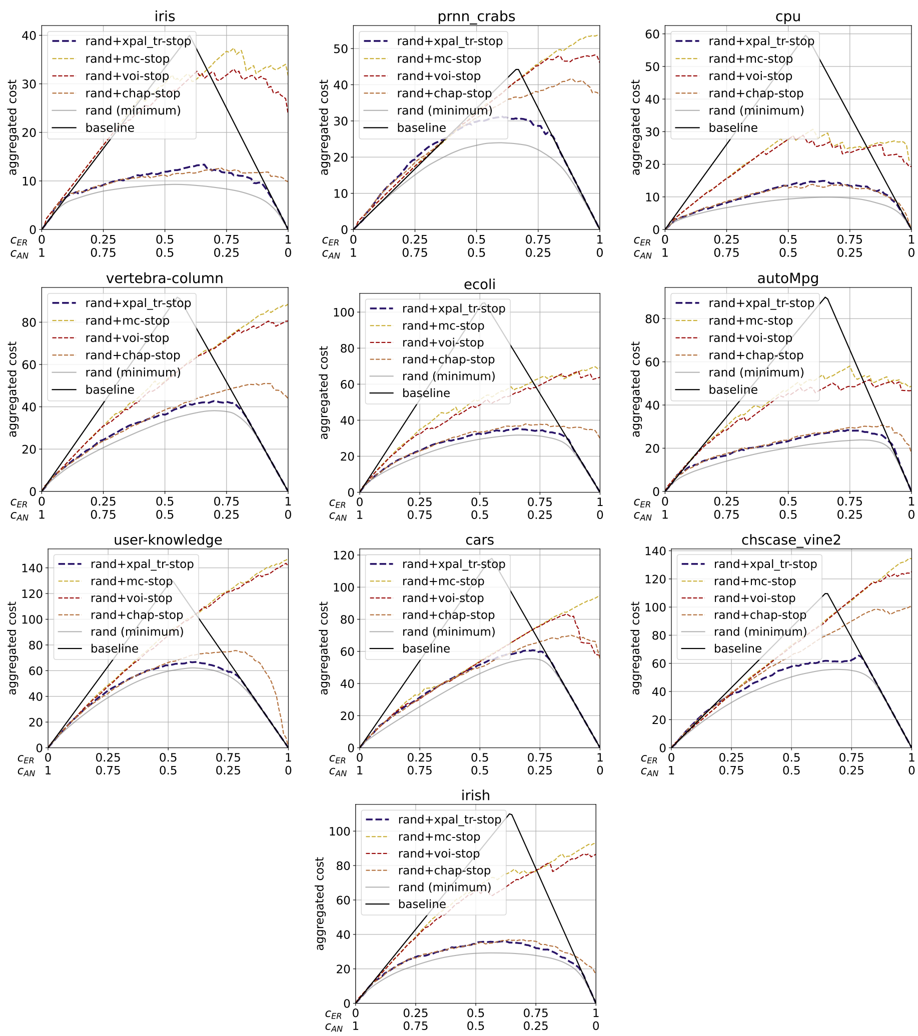

Fig. 10. Comparison of the aggregated cost curves for different stopping criteria with random selection.

	iris	prnn_crabs	cpu	vertebra-column	ecoli	autoMpg	user-knowledge	cars	chscase_vine2	irish	mean	win/tie/loss
rand+mc-stop	3.8↗	3.6↗	3.8↗	3.8↗	3.8↗	3.9↗	3.8↗	3.6↗	3.6↗	3.8↗	3.75	10 / 0 / 0
rand+chap-stop	1.7↗	2.2↗	1.4	2.1↗	2.1↗	2.1↗	2.0↗	2.5↗	2.2↗	1.9↗	2.02	9 / 1 / 0
rand+voi-stop	3.2↗	3.1↗	3.2↗	3.2↗	3.1↗	3.1↗	3.2↗	2.9↗	3.2↗	3.2↗	3.12	10 / 0 / 0
rand+xpal_tr-stop	1.3	1.0	1.6	1.0	1.0	1.0	1.0	1.0	1.0	1.1	1.10	baseline

Fig. 11. Ranking statistics for the area under the aggregated cost curves for different stopping criteria with random selection.

B Dataset Description

OpenML ID	name	instances	features	class distribution
61	iris	150	4	[0.33, 0.33, 0.33]
446	prnn_crabs	200	7	[0.5, 0.5]
796	cpu (v2)	209	7	[0.25, 0.75]
1523	vertebra-column	310	6	[0.19, 0.32, 0.48]
40671	ecoli (v3)	327	7	[0.44, 0.24, 0.11, 0.06, 0.16]
831	autoMpg (v2)	398	7	[0.47, 0.53]
1508	user-knowledge	403	5	[0.25, 0.32, 0.3, 0.06, 0.06]
455	cars	406	7	[0.63, 0.18, 0.19]
814	chscase_vine2 (v2)	468	2	[0.45, 0.55]
451	irish	500	5	[0.56, 0.44]