

1 We thank reviewers for their time and valuable comments. We will try to answer the concerns raised by each reviewer.

2 We thank R1 for their positive feedback. F3 and T2 are indeed results on *small* instances. We will clarify the captions.

3 We thank R2 for their valuable comments. 1) We respectfully disagree
4 on 1., as we do not consider this as a weakness of the paper. As we
5 explained in 1.52-60, state-of-the-art MILP solvers are CPU-based, and
6 performance of these solvers is benchmarked in terms of solving time.
7 As such, fast inference is crucial for any ML-based branching strategy
8 to be competitive w.r.t. this metric. Devising an effective model that can
9 learn how to branch accurately is one thing, and making such a model
10 practical, in realistic scenarios where practitioners do not have access to
11 high-end GPU cards, is another. The high inference cost of GNNs was

12 one of the main limitations for integration within MILP solvers, which we believe we addressed in the submission. Our
13 proposed hybrid model offers a better computational trade-off, while retaining most of the accuracy and generalization
14 ability of GNNs. We propose to make a better case of this question, which we will discuss more thoroughly in a new
15 sub-section “The accuracy/complexity trade-off”. 2) We also respectfully disagree on 2. Taken individually, it is true
16 that each part of our proposed methodology has been proposed in a different ML context. However the core idea of
17 running an expressive, but expensive, model once at the root node, and using the resulting transformation to modulate a
18 simpler model everywhere else in the tree, is novel. We believe it is a natural idea, specific to the problem of deploying
19 ML models in time-sensitive B&B environments, which the ML/OR community will find interesting. The paper shows
20 that this intuitive idea is practical, and promising for the field. 3) We agree on 3., as our training protocol improvements
21 have only a minor impact on the final performance of the model for branching. We conducted additional preliminary
22 experiments (Table A1), and it seems like KD has an effect on generalization which is visible only on hard instances,
23 while AT has a globally positive effect (except for one scenario in Table A1). We propose to lower our claim from
24 “substantial” to “minor”, and we will conduct and report the complete experiment in the final version of the paper,
25 including the small instances and the other 2 problem families. However, we would like to point out that the training
26 protocol improvements do not constitute the main contribution of the paper. 4) Finally, we thank R2 for the suggestion
27 regarding 1.255, which more accurately reflects the behaviour of hybrid-SVM models.

28 We thank R3 for their constructive comments. 1) We agree that our training protocol
29 tweaks do have a minor impact, and we will make this more clear in the text (see
30 response 3 to R2). 2) Regarding the trade-off between model accuracy and complex-
31 ity, we agree that our discussion in 1.52-60 is too brief, and will put more emphasis
32 on this discussion in a new sub-section (see answer 1 to R2). 3) Regarding KD, we
33 will report additional experiments to better highlight its contribution (see answer 3 to
34 R2). 4) We agree that a separate experiment with only AT would bring value to the
35 paper. We will conduct it and report the results in the supplementary materials, or
36 the main paper if space permits. 5) We will fix the two abstract misphrasings, thank
37 you. 6) We did measure the suggested additional evaluation metric, optimality gap at
38 time out (see Table A2). The results are again in favor of our proposed approach, we
39 will include them in the paper. Thank you for the suggestion.

40 We thank R4 for their insightful comments. 1) The proposed method is indeed specifically trained on one problem type,
41 which is a very valid point that we do not discuss in the paper. We follow the same setting as Gasse et al., who consider
42 this scenario as relevant because most practitioners usually only care about solving very specific problem types. We will
43 add a reference to that argument to clarify this point. 2) Regarding “different problem types require different models”,
44 this might be true, but here we provide evidence that goes in the other direction. Indeed, the same model seems to
45 work reasonably well on a variety of problem types, without the need for specific adaptations (except for re-training).
46 3) We agree that our model does not yield a general MILP branching strategy. Attaining this goal with ML would be a
47 substantial achievement, and indeed is not an easy task. We will include this remark in the concluding section. 4) Our
48 comparison to general MILP branching strategies like RPB is not completely fair. It is a very good point. However,
49 developing specialized versions of such strategies is a challenge in itself, and would require a dedicated paper. We would
50 gladly compare to such approaches, if available. Thank you for raising this point, and we will include this discussion as
51 well. 5) Regarding online learning, we do not share the same view here. Those branching rules exploit non-local LP
52 information to improve future decisions based on past branching decisions. But our model exploits such information as
53 well, as pseudo-cost is also used as an input feature, at every node. Online learning would imply changing the branching
54 strategy behaviour given the exact same features, which, in our knowledge, none of the branching strategies do.

55 Finally, both R3 and R4 do not consider our work reproducible, despite the source code being provided. We take note,
56 and will provide more exhaustive implementation details in the appendix of the final paper.

Table A1: Effect of training protocols on B&B tree size for the FiLM model (geometric mean over commonly solved instances).

	cautions		facilities	
	medium	big	medium	big
e2e	702	8939	340	339
e2e & KD	720	8846	343	325
e2e & KD & AT	686	8711	336	345

Table A2: Mean optimality gap (lower the better) of commonly unsolved “big” instances (number of such instances brackets).

	setcover (33)	indset (39)
FSB	0.170944	0.075543
PB	0.071286	0.029791
RPB	0.062841	0.025186
COMP	0.074041	0.025214
GNN	0.103906	0.034104
FiLM	0.059692	0.018685