

APPENDIX

A RELATED WORKS

Continual learning. Continual learning methodologies (Parisi et al., 2019; De Lange et al., 2021; Masana et al., 2022) can be broadly classified into three categories: *regularization-based*, *replay-based*, and *parameter-isolation* methods. *Regularization-based* approaches typically introduce a regularization term in the loss function to constrain changes to parameters relevant to prior tasks. These can further be categorized as data-focused (Li & Hoiem, 2017; Kim et al., 2023), leveraging knowledge distillation from previously trained models, or prior-focused (Kirkpatrick et al., 2017; Zenke et al., 2017; Aljundi et al., 2018), estimating parameter importance as a prior for the new model. Recent research proposed enforcing weight updates within the null space of feature covariance (Wang et al., 2021; Tang et al., 2021). *Replay-based* methods rely on memory and rehearsal mechanisms to recall episodic memories of past tasks during training, thereby keeping the loss low in those tasks. Two main strategies are: exemplar replay - which stores selected training samples (Riemer et al., 2018; Buzzega et al., 2020; Chaudhry et al., 2018; Prabhu et al., 2020; Chaudhry et al., 2019; Liang & Li, 2024) and generative replay - with models that synthesize previous data with generative models (Shin et al., 2017; Wu et al., 2018). *Parameter isolation* methods aim to learn task-specific sub-networks within a shared network. Various techniques, such as Piggyback (Mallya et al., 2018), PackNet (Mallya & Lazebnik, 2018), SupSup (Wortsman et al., 2020), HAT (Serra et al., 2018), and Progressive Neural Network (Rusu et al., 2016), allocate and combine parameters for individual tasks. While effective in task-aware settings, these methods are most suited for scenarios with a known task sequence or oracle.

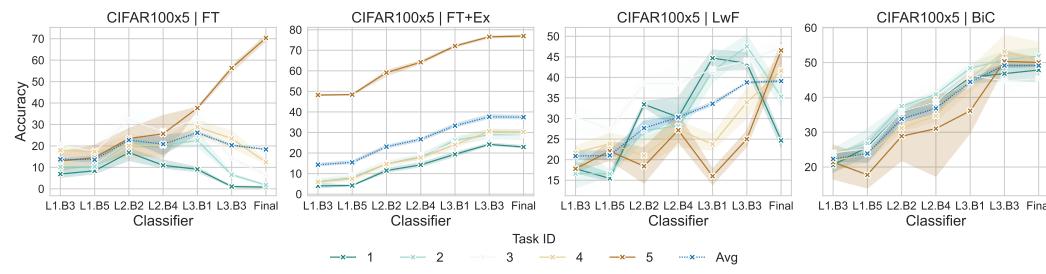
Representations in neural networks. Understanding and comparing the representations at different layers in deep neural networks is an active area of research, and tools such as CKA (Kornblith et al., 2019) emerged to measure the similarity of representations across layers. Several works have investigated how network representations behave during continual learning: Ramasesh et al. (2020) noticed that the most forgetting occurs in the deeper network layers, Zhao et al. (2023) likewise demonstrated that only a subset of modules is sensitive to the changes during continual learning, and recent work of Masarczyk et al. (2023) showed that neural networks split into parts that build different representations. Similar insights motivated continual learning methods that enforce stability through replay (Liu et al., 2020a; Pawlak et al., 2022) or regularization (Douillard et al., 2020) at the level of intermediate network layers. In representation learning, several works demonstrated that probing classifiers learned on top of intermediate network representations tend to perform relatively well (Davari et al., 2022), although usually worse than the final classifier. Early-exit techniques (Panda et al., 2016; Teerapittayanan et al., 2016; Kaya et al., 2019) use intermediate representations to reduce the inference cost through dynamic inference that enables skipping later model layers. Several works on early-exits also propose more advanced strategies (Liao et al., 2021; Sun et al., 2021; Han et al., 2022; Wójcik et al., 2023) that improve the effectiveness. Use of multiple classifiers in continual learning has been explored in an ensembling-like manner (Liu et al., 2020b), and (Yan et al., 2024) utilized intermediate classifiers for online continual learning and different motivations. Our method is dedicated to offline continual learning.

810 B ANALYSIS RESULTS FOR 5-TASK SPLITS 811

812 In this section, we include results with 5-tasks split corresponding to the experiments from Sec-
813 tions 2.1 to 2.3. Figures 7 to 9 demonstrate that our previous insights also hold for different task
814 split.
815

	CIFAR100x5 FT				CIFAR100x5 FT+Ex				CIFAR100x5 LwF				CIFAR100x5 BiC				
	L1.B3	0.95	0.94	0.90	0.87	0.95	0.95	0.92	0.91	0.99	0.99	0.99	0.98	0.99	0.98	0.96	0.94
	L1.B5	0.92	0.90	0.85	0.81	0.93	0.91	0.87	0.84	0.98	0.98	0.97	0.97	0.98	0.96	0.94	0.92
	L2.B2	0.93	0.90	0.86	0.84	0.94	0.91	0.89	0.87	0.99	0.98	0.98	0.98	0.98	0.96	0.94	0.92
	L2.B4	0.91	0.86	0.82	0.79	0.92	0.87	0.84	0.83	0.99	0.98	0.97	0.96	0.97	0.94	0.92	0.89
	L3.B1	0.86	0.81	0.78	0.77	0.88	0.85	0.82	0.81	0.97	0.96	0.96	0.95	0.95	0.92	0.89	0.87
	L3.B3	0.74	0.69	0.66	0.65	0.79	0.74	0.72	0.70	0.94	0.92	0.91	0.90	0.91	0.85	0.81	0.78
	L3.B5	0.57	0.51	0.51	0.51	0.73	0.66	0.63	0.61	0.91	0.88	0.86	0.85	0.86	0.79	0.74	0.71
	2	3	4	5	2	3	4	5	2	3	4	5	2	3	4	5	
	Task index				Task index				Task index				Task index				

825 Figure 7: CKA of the first task representations across different ResNet32 layers (L1.B3-L3.B5)
826 through continual learning on CIFAR100 split into 5 tasks.
827



838 Figure 8: Per-task final accuracy of the auxiliary classifiers trained with linear probing on top of
839 several network layers and final network classifier on CIFAR100 split into 5 tasks.
840

	CIFAR100x5 FT					CIFAR100x5 FT+Ex					CIFAR100x5 LwF					CIFAR100x5 BiC									
	L1.B3	2.68	2.13	2.48	1.95	0.12	1.87	0.25	0.37	0.15	0.30	0.42	0.30	0.78	1.10	1.50	1.57	0.85	1.16	0.52	0.98	0.77	0.63	1.32	0.84
	L1.B5	3.32	1.92	1.97	1.37	0.35	1.78	0.27	0.57	0.37	0.30	0.57	0.41	0.93	0.97	1.08	1.95	1.92	1.37	0.72	0.78	0.92	1.00	0.86	0.85
	L2.B2	6.52	3.88	4.18	1.65	0.20	3.29	0.55	0.63	0.30	0.47	0.23	0.44	1.72	0.82	1.13	0.45	0.30	0.88	0.73	0.75	1.02	0.77	0.53	0.76
	L2.B4	3.10	5.05	2.30	1.73	0.67	2.57	0.78	0.88	0.90	0.93	0.42	0.78	0.88	0.80	1.12	1.07	1.40	1.05	0.95	0.72	1.00	0.82	0.73	0.84
	L3.B1	3.05	7.43	6.22	4.95	0.75	4.48	2.00	2.47	1.92	1.83	0.93	1.83	3.27	2.75	1.35	0.62	0.28	1.65	1.63	1.77	2.20	1.45	0.78	1.57
	L3.B3	0.43	2.05	2.47	4.90	1.48	2.27	3.03	2.48	2.15	3.10	1.23	2.40	4.15	4.92	1.60	2.05	0.65	2.67	1.53	1.97	1.67	3.33	2.85	2.27
	Final	0.30	0.43	1.00	2.28	10.88	2.98	3.63	3.98	4.70	5.27	2.37	3.99	0.33	1.93	4.12	6.17	11.32	4.77	3.60	3.97	3.53	3.88	4.40	3.84
	1	2	3	4	5	Avg	1	2	3	4	5	Avg	1	2	3	4	5	Avg	1	2	3	4	5	Avg	
	Task					Task					Task					Task					Task				

850 Figure 9: Unique accuracy (a subset of samples that a single given classifier classifies correctly)
851 of auxiliary classifiers and final network classifier for different task data on CIFAR100 split into 5
852 tasks.
853
854
855
856
857
858
859
860
861
862
863

864 C ANALYSIS FOR METHODS WITH ENABLED GRADIENT PROPAGATION 865

866 We also replicate the analysis from Section 2 for methods with gradient propagation enabled through
867 the network, using both 5 and 10 task split. In Figures 10 to 12 we observe similar patterns as in
868 case of networks without gradient propagation, which validates our idea to train the network together
869 with the classifiers.

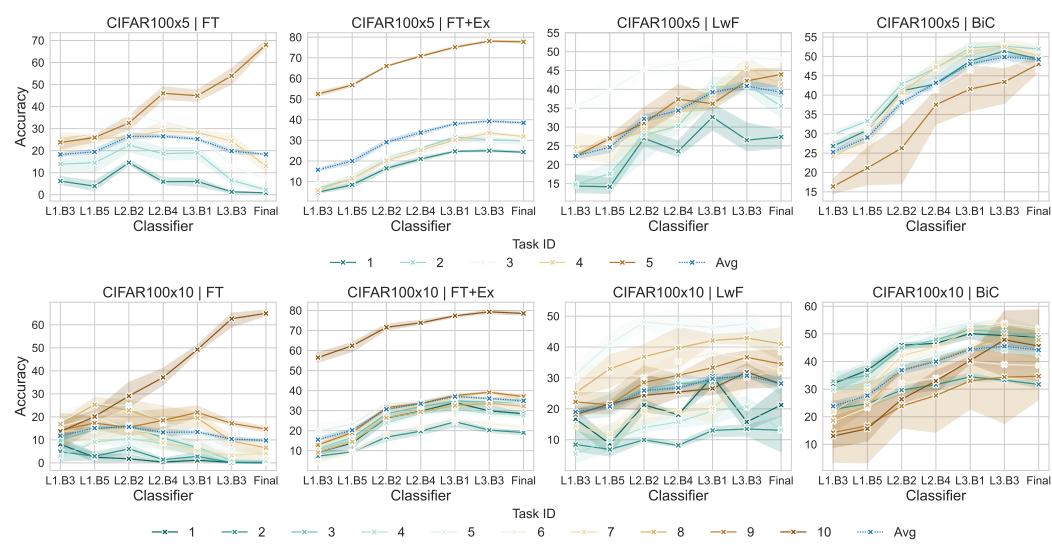
871 C.1 CKA 872

CIFAR100x5 FT					CIFAR100x5 FT+Ex					CIFAR100x5 LwF					CIFAR100x5 BiC				
L1.B3	0.95	0.91	0.86	0.83	0.96	0.92	0.89	0.86	1.00	0.99	0.99	0.99	0.98	0.97	0.95	0.94			
L1.B5	0.92	0.88	0.84	0.81	0.94	0.89	0.86	0.83	0.99	0.99	0.99	0.99	0.98	0.96	0.93	0.91			
L2.B2	0.89	0.84	0.81	0.80	0.91	0.87	0.85	0.83	0.99	0.98	0.98	0.98	0.97	0.94	0.92	0.91			
L2.B4	0.85	0.80	0.77	0.75	0.88	0.84	0.81	0.80	0.98	0.97	0.97	0.96	0.96	0.92	0.90	0.88			
L3.B1	0.80	0.77	0.73	0.71	0.85	0.81	0.78	0.77	0.97	0.96	0.95	0.95	0.94	0.90	0.87	0.85			
L3.B3	0.67	0.63	0.60	0.58	0.77	0.71	0.68	0.67	0.94	0.92	0.91	0.90	0.89	0.83	0.78	0.76			
L3.B5	0.53	0.47	0.49	0.49	0.72	0.65	0.63	0.61	0.91	0.88	0.87	0.86	0.86	0.79	0.74	0.71			
Task index					Task index					Task index					Task index				

CIFAR100x10 FT											CIFAR100x10 FT+Ex											CIFAR100x10 LwF											CIFAR100x10 BiC										
L1.B3	0.95	0.92	0.88	0.85	0.82	0.75	0.80	0.77	0.76	0.95	0.91	0.88	0.86	0.84	0.81	0.82	0.79	0.77	0.99	0.99	0.99	0.99	0.98	0.98	0.98	0.98	0.97	0.99	0.95	0.93	0.93	0.91	0.90	0.88	0.86								
L1.B5	0.92	0.89	0.86	0.83	0.83	0.77	0.81	0.78	0.77	0.93	0.87	0.85	0.82	0.82	0.80	0.80	0.78	0.77	0.99	0.99	0.98	0.98	0.98	0.98	0.98	0.98	0.97	0.98	0.96	0.94	0.92	0.91	0.90	0.89	0.88	0.86							
L2.B2	0.87	0.83	0.83	0.80	0.78	0.72	0.75	0.72	0.73	0.90	0.84	0.82	0.81	0.80	0.79	0.79	0.78	0.77	0.99	0.98	0.97	0.97	0.97	0.96	0.96	0.96	0.97	0.94	0.92	0.90	0.89	0.88	0.87	0.86	0.85								
L2.B4	0.81	0.77	0.76	0.75	0.73	0.67	0.71	0.67	0.69	0.88	0.82	0.81	0.80	0.79	0.78	0.77	0.76	0.75	0.98	0.97	0.96	0.96	0.95	0.95	0.94	0.93	0.92	0.91	0.90	0.89	0.88	0.87	0.86	0.85									
L3.B1	0.72	0.69	0.69	0.69	0.67	0.62	0.66	0.63	0.62	0.86	0.81	0.79	0.78	0.76	0.76	0.76	0.75	0.74	0.96	0.95	0.95	0.94	0.93	0.93	0.92	0.91	0.90	0.89	0.88	0.87	0.86	0.85	0.84	0.83	0.82	0.81	0.80	0.79					
L3.B3	0.59	0.55	0.56	0.58	0.55	0.47	0.53	0.51	0.52	0.81	0.76	0.72	0.72	0.70	0.69	0.68	0.67	0.66	0.94	0.92	0.91	0.90	0.89	0.89	0.88	0.87	0.86	0.92	0.86	0.83	0.80	0.77	0.76	0.74	0.73	0.72							
L3.B5	0.41	0.38	0.41	0.43	0.43	0.30	0.42	0.36	0.41	0.80	0.74	0.69	0.68	0.65	0.65	0.62	0.61	0.61	0.92	0.90	0.88	0.87	0.86	0.86	0.85	0.85	0.84	0.90	0.83	0.79	0.76	0.73	0.71	0.68	0.67	0.66							
Task index											Task index											Task index											Task index										

891 Figure 10: CKA of the first task representations across different ResNet32 layers (L1.B3-L3.B5)
892 through continual learning on CIFAR100, **with enabled gradient propagation**.

893 C.2 INDIVIDUAL CLASSIFIER ACCURACY 894



915 Figure 11: Per-task final accuracy of the auxiliary classifiers trained with linear probing on top of
916 several network layers and final network classifier on CIFAR100, **with enabled gradient propaga-**
917 **tion**.

918
919

C.3 UNIQUE ACCURACY

920

	CIFAR100x5 FT					CIFAR100x5 FT+Ex					CIFAR100x5 LwF					CIFAR100x5 BiC								
	1	2	3	4	5	Avg	1	2	3	4	5	Avg	1	2	3	4	5	Avg	1	2	3	4	5	Avg
L1.B3	2.48	3.25	2.07	2.47	0.50	2.15	0.18	0.18	0.27	0.08	0.42	0.23	0.98	0.72	1.47	1.20	0.72	1.02	0.65	0.68	1.33	0.35	0.47	0.70
L1.B5	0.78	2.25	3.15	2.58	0.45	1.84	0.42	0.60	0.18	0.58	0.47	0.45	0.53	0.85	1.30	0.75	0.70	0.83	0.52	0.60	0.73	0.60	0.83	0.66
L2.B2	7.77	5.68	4.40	2.08	0.60	4.11	0.85	0.83	1.08	0.58	0.42	0.75	1.77	1.33	0.75	0.63	0.55	1.01	0.88	0.83	0.98	1.05	0.58	0.87
L2.B4	1.90	3.20	3.33	3.07	1.00	2.50	1.40	1.48	1.73	1.22	0.45	1.26	0.48	1.18	0.82	0.97	1.62	1.01	0.77	1.02	0.82	1.45	1.52	1.11
L3.B1	2.53	5.18	3.57	3.00	0.68	2.99	2.37	2.63	2.07	1.85	0.90	1.96	2.58	3.10	1.28	1.15	0.97	1.82	1.10	1.53	1.37	1.43	1.75	1.44
L3.B3	0.47	1.22	1.45	3.25	0.58	1.39	1.80	1.60	1.90	2.12	0.95	1.67	0.92	2.63	1.13	2.57	1.62	1.77	1.55	1.42	1.83	1.93	1.30	1.61
Final	0.38	0.30	0.90	1.00	6.98	1.91	2.48	2.68	2.85	2.83	1.40	2.45	2.10	1.97	2.15	2.52	3.82	2.51	1.95	2.75	2.42	2.05	4.78	2.79
	1	2	3	4	5	Avg	1	2	3	4	5	Avg	1	2	3	4	5	Avg	1	2	3	4	5	Avg

	CIFAR100x10 FT										Avg	CIFAR100x10 FT+Ex										Avg	CIFAR100x10 LwF										Avg	CIFAR100x10 BiC										Avg
	1	2	3	4	5	6	7	8	9	10	Avg	1	2	3	4	5	6	7	8	9	10	Avg	1	2	3	4	5	6	7	8	9	10	Avg											
L1.B3	7.50	3.93	3.27	1.83	3.07	4.83	3.67	1.93	2.83	0.27	3.31	0.37	0.60	0.30	0.27	0.77	0.23	0.60	0.20	0.40	0.70	0.44	3.20	1.97	2.20	0.57	0.93	1.93	1.60	0.87	1.73	0.80	1.58	0.70	1.07	0.87	1.30	0.73	0.50	0.27	0.73	0.70	0.60	0.74
L1.B5	2.07	1.03	5.50	5.23	6.23	6.10	4.23	4.33	3.97	0.27	3.90	0.63	0.50	0.63	1.17	0.60	0.90	0.30	0.60	0.57	0.43	0.63	0.37	0.83	1.33	1.47	1.57	1.90	0.67	1.70	0.80	1.03	1.17	0.70	0.73	0.87	0.93	1.00	0.67	0.67	0.83	0.50	0.76	
L2.B2	1.37	4.20	5.63	5.10	5.13	2.96	5.23	3.80	2.20	0.50	3.61	1.67	1.17	1.40	1.60	1.73	1.57	1.70	1.63	1.70	0.57	1.47	1.77	0.73	1.10	1.23	1.20	0.60	0.87	1.13	0.70	1.13	1.05	1.46	1.03	1.03	1.23	0.97	1.17	1.03	0.93	1.03	0.73	1.06
L2.B4	0.77	0.43	3.10	2.97	2.80	2.87	2.23	2.43	2.20	0.45	1.96	1.57	1.73	1.70	1.73	1.23	1.67	1.40	1.40	1.07	0.57	1.41	0.50	0.50	1.20	1.23	0.97	0.83	0.83	0.80	0.40	0.85	0.97	1.23	1.07	1.07	1.17	1.27	0.97	1.57	1.07	1.43	1.18	
L3.B1	0.70	1.60	2.13	1.83	1.77	1.87	1.07	3.40	3.27	0.87	1.85	2.83	3.40	2.53	2.57	1.70	3.03	1.97	1.93	1.63	1.07	2.27	4.77	1.53	1.33	2.60	1.30	1.70	0.53	1.37	0.80	0.83	1.68	1.37	1.87	1.47	1.70	1.43	1.87	1.90	1.67	1.93	1.87	1.71
L3.B3	0.13	0.10	0.43	0.20	1.50	0.63	0.70	0.99	1.53	2.10	0.83	1.30	1.23	1.23	1.13	1.10	2.10	1.03	1.57	1.67	0.73	1.31	0.47	1.57	2.13	2.60	1.53	1.97	2.00	1.40	2.00	2.23	1.79	1.10	1.33	1.20	1.10	1.60	1.60	1.47	2.30	1.90	2.30	1.59
Final	0.07	0.10	0.33	0.20	0.33	0.90	0.27	0.70	1.27	3.87	0.90	1.40	1.43	2.37	1.70	1.83	2.67	2.83	2.50	2.43	0.97	2.01	2.17	1.97	2.50	1.17	0.93	2.00	4.23	2.23	2.90	2.13	2.22	1.50	1.87	1.83	2.10	1.90	2.53	2.20	3.13	4.20	3.03	2.43
	1	2	3	4	5	6	7	8	9	10	Avg	1	2	3	4	5	6	7	8	9	10	Avg	1	2	3	4	5	6	7	8	9	10	Avg											

Figure 12: Unique accuracy (a subset of samples that a single given classifier classifies correctly) of auxiliary classifiers and final network classifier for different task data on CIFAR100, **with enabled gradient propagation**.938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

972 D MULTI-CLASSIFIER PERFORMANCE UPPER BOUND ANALYSIS
973

974 To quantify the potential of the auxiliary classifiers, we consider an oracle multi-classifier network
975 as an upper bound for our method. When evaluating such an oracle, we obtain predictions from all
976 of its classifiers and consider a prediction correct when at least one classifier (auxiliary classifier or
977 the original network classifier) is correct. Therefore, the oracle has an ideal algorithm for combining
978 classifier predictions and always returns the ‘best case’ prediction from all the classifiers. We mea-
979 sure the difference between the accuracy of such an oracle network and the accuracy of a standard
980 single-classifier network and present the results for first task data, last task data, and average over
981 all the tasks in Table 8 and Table 9 for both linear probing and ACs. As in our previous analysis in
982 Section 2, exemplar-free methods show more variance in the performance across tasks. However,
983 the average difference across all tasks is also significant for exemplar-based methods, with the oracle
984 for the best-performing method - BiC - achieving approximately a 30-40% relative increase in over-
985 all accuracy. Those results indicate that, while our simple setup achieves consistent improvement, it
986 is still far from the best-case utilization of multi-classifier networks and still leaves room for future
987 work.
988

989 Table 8: Upper bound on accuracy improvement on 5 tasks of CIFAR100 when using oracle multi-
990 classifier network, trained with linear probing and auxiliary classifiers.
991

992

	Task 1	Task 2	Task 3	Task 4	Task 5	Avg
CIFAR100x5, Linear Probing						
FT	31.82±2.23	46.90±1.60	54.22±0.59	43.30±1.24	6.57±0.96	36.56±0.52
FT+Ex	12.28±0.10	14.58±0.33	12.72±1.80	13.87±0.93	11.37±0.71	12.96±0.62
LwF	35.07±2.03	28.88±2.51	20.92±2.01	15.25±1.13	9.55±1.47	21.93±0.20
BiC	15.15±0.91	17.00±2.33	18.73±1.50	17.93±2.26	16.03±2.39	16.97±0.50
CIFAR100x5, Auxiliary Classifiers						
FT	24.65±2.82	45.07±5.77	56.87±3.68	48.82±4.36	9.05±0.65	36.89±1.33
FT+Ex	14.42±0.60	17.42±1.12	15.70±1.58	15.30±1.59	10.88±0.40	14.74±0.66
LwF	18.52±1.89	22.57±2.84	21.57±0.93	18.57±0.95	15.07±2.08	19.26±0.61
BiC	15.88±0.96	18.42±0.46	18.83±2.29	18.72±0.06	15.43±3.06	17.46±0.61

993 Table 9: Upper bound on accuracy improvement on 10 tasks of CIFAR100 when using oracle multi-
994 classifier network, trained with linear probing and auxiliary classifiers.
995

996

	Task 1	Task 2	Task 3	Task 4	Task 5	Task 6	Task 7	Task 8	Task 9	Task 10	Avg
CIFAR100x10, Linear Probing											
FT	30.00±7.52	16.10±3.90	40.47±6.03	23.37±4.57	48.70±0.72	41.43±5.52	32.80±1.78	44.40±4.83	28.87±0.97	9.03±1.72	31.52±0.64
FT+Ex	17.17±1.24	16.63±0.42	18.40±2.72	18.10±1.44	17.67±0.78	16.30±1.31	18.90±0.79	15.77±1.30	17.00±0.44	11.90±1.32	16.78±0.38
LwF	48.93±6.24	30.30±3.00	33.63±3.35	28.93±5.05	25.60±2.13	21.50±5.50	13.30±2.17	14.47±1.50	11.33±1.24	8.40±1.54	23.64±1.29
BiC	19.03±1.10	20.73±2.04	20.50±2.74	20.83±2.37	17.83±1.12	19.17±0.67	16.83±3.16	18.90±0.53	20.77±4.80	17.73±4.31	19.23±0.70
CIFAR100x10, Auxiliary Classifiers											
FT	12.93±0.38	14.30±4.69	35.07±0.21	25.70±4.90	46.27±4.57	38.70±2.77	32.50±2.17	43.77±0.93	34.73±3.46	9.20±0.62	29.32±0.98
FT+Ex	20.03±1.56	17.77±1.47	18.87±0.60	20.33±1.70	18.13±2.25	20.50±0.78	18.97±1.95	17.47±1.36	19.13±2.48	12.30±0.80	18.35±0.01
LwF	22.80±3.46	12.60±4.18	22.93±1.65	21.50±2.91	28.30±1.80	25.33±4.65	12.70±1.49	21.53±3.31	18.00±5.12	15.27±2.77	20.10±0.74
BiC	18.87±2.76	20.60±0.46	21.53±2.74	22.70±2.52	18.60±1.51	19.87±2.01	17.17±1.88	19.77±0.57	18.77±8.36	17.73±7.03	19.56±0.78

1000
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

1026 **E DYNAMIC INFERENCE RULE ABLATION**
 1027

1028 In Table 10 we demonstrate the accuracy of two variants of dynamic inference for different confi-
 1029 dence thresholds λ for CIFAR100. We compare the standard, early-exit paradigm, where the net-
 1030 work returns a final classifier prediction in case no classifier meets the exit rule, and the paradigm
 1031 used in our experiments where the network defaults to the most confident prediction. Using the most
 1032 confident prediction outperforms the standard early-exit rule, which is consistent with our analysis
 1033 that showed that the last classifier is not always the best in continual learning and that the early
 1034 classifiers exhibit lower forgetting for earlier task data.

1035 Table 10: Comparison between dynamic inference performance with different confidence thresholds
 1036 λ when using maximum confidence prediction (MC) and final classifier prediction (Last) as the de-
 1037 fault output for multi-classifier networks trained with linear probing (LP) or jointly with the network
 1038 with gradient propagation (AC). Using max confidence prediction yields better accuracy.
 1039

	λ	FT (LP)	FT+Ex (LP)	LwF (LP)	BiC (LP)	FT (AC)	FT+Ex (AC)	LwF (AC)	BiC (AC)
CIFAR100x5									
Last	0.5	24.14 \pm 1.35	28.43 \pm 0.87	37.34 \pm 0.09	48.35 \pm 0.23	27.64 \pm 0.97	28.87 \pm 0.32	38.08 \pm 0.87	48.15 \pm 0.40
MC		24.73 \pm 1.48	28.33 \pm 0.96	36.95 \pm 0.20	48.48 \pm 0.40	28.10 \pm 1.03	28.85 \pm 0.31	38.36 \pm 0.59	48.37 \pm 0.31
Last	0.75	23.67 \pm 1.32	35.60 \pm 1.26	40.21 \pm 0.08	49.25 \pm 0.33	26.00 \pm 0.67	36.27 \pm 0.28	39.41 \pm 1.01	49.45 \pm 0.73
MC		26.66 \pm 1.70	35.09 \pm 1.38	39.70 \pm 0.10	49.74 \pm 0.48	28.91 \pm 1.07	36.18 \pm 0.17	40.33 \pm 0.76	50.19 \pm 0.63
Last	0.9	20.98 \pm 0.99	37.27 \pm 1.10	39.97 \pm 0.27	49.18 \pm 0.36	22.38 \pm 0.33	38.22 \pm 0.19	39.28 \pm 1.28	49.35 \pm 0.65
MC		26.70 \pm 1.55	36.57 \pm 1.33	40.07 \pm 0.10	49.89 \pm 0.52	28.44 \pm 1.04	38.27 \pm 0.13	40.50 \pm 0.91	50.39 \pm 0.65
Last	0.95	19.91 \pm 0.76	37.48 \pm 0.98	39.62 \pm 0.23	49.15 \pm 0.31	20.69 \pm 0.24	38.50 \pm 0.17	39.25 \pm 1.36	49.24 \pm 0.63
MC		26.74 \pm 1.40	36.77 \pm 1.29	40.07 \pm 0.08	49.89 \pm 0.52	28.25 \pm 1.05	38.62 \pm 0.21	40.53 \pm 0.94	50.40 \pm 0.66
Last	0.98	19.19 \pm 0.27	37.46 \pm 0.92	39.40 \pm 0.21	49.14 \pm 0.32	19.49 \pm 0.18	38.55 \pm 0.38	39.22 \pm 1.34	49.24 \pm 0.65
MC		26.83 \pm 1.23	36.81 \pm 1.27	40.10 \pm 0.07	49.89 \pm 0.52	28.19 \pm 1.08	38.72 \pm 0.24	40.55 \pm 0.95	50.40 \pm 0.68
Last	0.99	18.94 \pm 0.15	37.44 \pm 0.88	39.32 \pm 0.22	49.14 \pm 0.32	19.08 \pm 0.20	38.55 \pm 0.44	39.22 \pm 1.35	49.23 \pm 0.64
MC		26.82 \pm 1.22	36.83 \pm 1.27	40.10 \pm 0.07	49.89 \pm 0.52	28.20 \pm 1.09	38.75 \pm 0.27	40.55 \pm 0.95	50.40 \pm 0.68
Last	1.0	18.39 \pm 0.08	37.43 \pm 0.85	39.11 \pm 0.26	49.14 \pm 0.32	18.35 \pm 0.30	38.51 \pm 0.43	39.22 \pm 1.34	49.22 \pm 0.65
MC		26.82 \pm 1.19	36.83 \pm 1.27	40.10 \pm 0.07	49.89 \pm 0.52	28.18 \pm 1.07	38.75 \pm 0.26	40.55 \pm 0.95	50.40 \pm 0.68
CIFAR100x10									
Last	0.5	15.22 \pm 1.64	27.03 \pm 0.90	28.69 \pm 0.79	42.37 \pm 1.60	15.85 \pm 1.10	27.68 \pm 0.42	28.96 \pm 1.29	42.40 \pm 1.17
MC		16.05 \pm 1.95	27.04 \pm 0.87	28.06 \pm 0.93	42.62 \pm 1.75	16.79 \pm 1.34	27.72 \pm 0.37	29.09 \pm 1.10	42.67 \pm 1.44
Last	0.75	14.12 \pm 0.98	33.69 \pm 1.20	31.10 \pm 0.87	43.95 \pm 1.69	13.31 \pm 1.17	34.40 \pm 0.47	28.65 \pm 1.72	44.93 \pm 0.93
MC		17.47 \pm 1.51	33.84 \pm 1.03	30.17 \pm 0.77	44.59 \pm 1.75	16.77 \pm 1.22	34.64 \pm 0.41	29.72 \pm 1.19	45.83 \pm 1.51
Last	0.9	12.19 \pm 0.51	34.61 \pm 1.05	31.03 \pm 0.92	43.87 \pm 1.63	11.25 \pm 1.04	35.83 \pm 0.51	28.30 \pm 1.83	44.55 \pm 0.62
MC		17.71 \pm 1.33	35.41 \pm 0.87	30.54 \pm 0.76	44.82 \pm 1.71	16.82 \pm 1.14	36.59 \pm 0.52	29.79 \pm 1.21	46.14 \pm 1.46
Last	0.95	11.25 \pm 0.55	34.49 \pm 1.10	30.63 \pm 1.10	43.79 \pm 1.72	10.62 \pm 0.83	35.63 \pm 0.43	28.27 \pm 1.79	44.33 \pm 0.53
MC		17.74 \pm 1.31	35.58 \pm 0.92	30.59 \pm 0.67	44.84 \pm 1.72	16.89 \pm 1.11	36.92 \pm 0.39	29.79 \pm 1.21	46.17 \pm 1.45
Last	0.98	10.51 \pm 0.35	34.27 \pm 1.13	30.17 \pm 1.19	43.77 \pm 1.72	10.09 \pm 0.73	35.29 \pm 0.39	28.22 \pm 1.81	44.23 \pm 0.54
MC		17.77 \pm 1.30	35.61 \pm 0.90	30.60 \pm 0.68	44.84 \pm 1.73	16.88 \pm 1.08	36.97 \pm 0.40	29.79 \pm 1.21	46.19 \pm 1.47
Last	0.99	10.19 \pm 0.37	34.26 \pm 1.10	30.00 \pm 1.20	43.76 \pm 1.72	9.92 \pm 0.69	35.10 \pm 0.45	28.20 \pm 1.83	44.21 \pm 0.53
MC		17.77 \pm 1.30	35.61 \pm 0.89	30.60 \pm 0.68	44.84 \pm 1.73	16.88 \pm 1.08	36.98 \pm 0.40	29.79 \pm 1.21	46.19 \pm 1.47
Last	1.0	9.79 \pm 0.33	34.22 \pm 1.08	29.65 \pm 1.19	43.75 \pm 1.72	9.76 \pm 0.63	34.93 \pm 0.40	28.19 \pm 1.84	44.19 \pm 0.51
MC		17.77 \pm 1.30	35.62 \pm 0.89	30.60 \pm 0.69	44.84 \pm 1.73	16.88 \pm 1.08	36.97 \pm 0.39	29.79 \pm 1.21	46.19 \pm 1.47

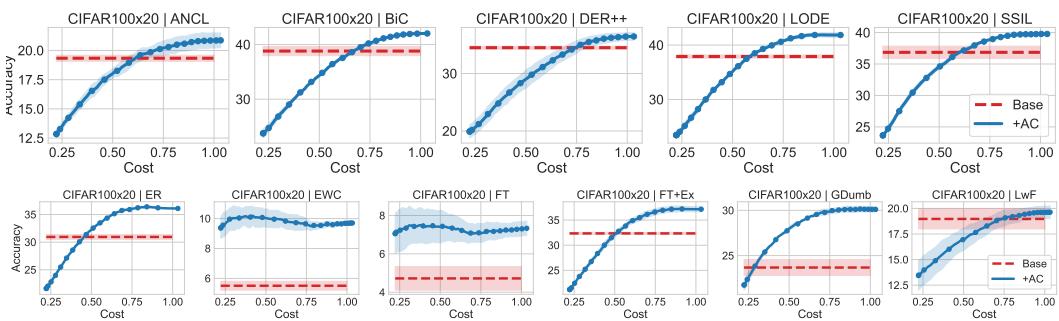
1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

1080 F AC-ENHANCED METHODS ON LONGER TASK SEQUENCES 1081

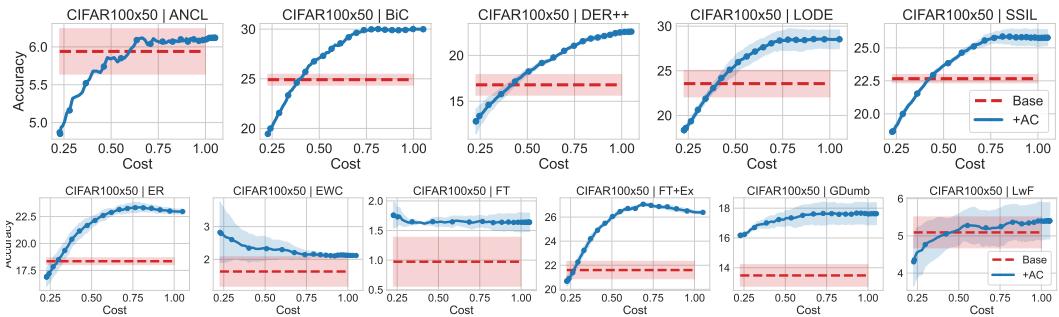
1082 In Table 11 and Figures 13 and 14 we present results for experiments with 20 and 50 task split,
1083 following the setup from Section 4. For the 50-task split, we use a growing memory of 20 exemplars
1084 instead of a constant memory of 2000 due to the task size. Non-replay-based methods perform
1085 poorly on longer sequences of tasks, especially on 50 tasks, but ACs robustly enhance the method
1086 performance in all tested scenarios.

1087 **Table 11:** Additional results for AC-enhanced methods with longer sequences of tasks on CIFAR100
1088

Method	FT	FT+Ex	GDumb	ANCL	BiC	DER++	ER	EWC	LwF	LODE	SSIL	Avg
CIFAR100x20												
Base	4.72 \pm 0.75	32.35 \pm 0.26	23.68 \pm 1.08	19.34 \pm 0.32	38.81 \pm 1.02	34.50 \pm 0.33	30.94 \pm 0.53	5.51 \pm 0.36	18.97 \pm 1.20	37.90 \pm 0.37	36.86 \pm 1.21	25.78 \pm 0.32
+AC	7.33\pm0.45	37.16\pm0.63	30.11\pm0.36	20.87\pm0.78	42.03\pm0.18	36.45\pm0.81	36.10\pm0.07	9.70\pm0.18	19.63\pm0.75	41.85\pm0.49	39.78\pm0.24	29.18\pm0.10
Δ	+2.61 \pm 0.61	+4.81 \pm 0.86	+6.43 \pm 0.98	+1.53 \pm 0.91	+3.22 \pm 1.01	+1.95 \pm 1.05	+5.16 \pm 0.60	+4.19 \pm 0.40	+0.66 \pm 1.69	+3.95 \pm 0.65	+2.92 \pm 1.03	+3.40 \pm 0.22
CIFAR100x50												
Base	0.97 \pm 0.51	21.60 \pm 0.88	13.48 \pm 0.89	5.94 \pm 0.37	24.91 \pm 0.67	16.79 \pm 1.38	18.35 \pm 0.38	1.63 \pm 0.55	5.09 \pm 0.51	23.56 \pm 1.82	22.66 \pm 0.39	14.09 \pm 0.28
+AC	1.64\pm0.19	26.39\pm0.16	17.63\pm0.90	6.41\pm0.59	30.10\pm0.36	22.59\pm0.39	22.96\pm0.29	2.12\pm0.02	5.40\pm0.60	28.53\pm1.29	25.77\pm0.74	17.23\pm0.18
Δ	+0.67 \pm 0.67	+4.79 \pm 0.88	+4.15 \pm 0.60	+0.47 \pm 0.22	+5.19 \pm 0.70	+5.80 \pm 1.01	+4.61 \pm 0.52	+0.49 \pm 0.57	+0.31 \pm 0.88	+4.97 \pm 2.06	+3.11 \pm 1.12	+3.14 \pm 0.44



1097 **Figure 13:** Dynamic inference plots for 20 task split of CIFAR100.
1098



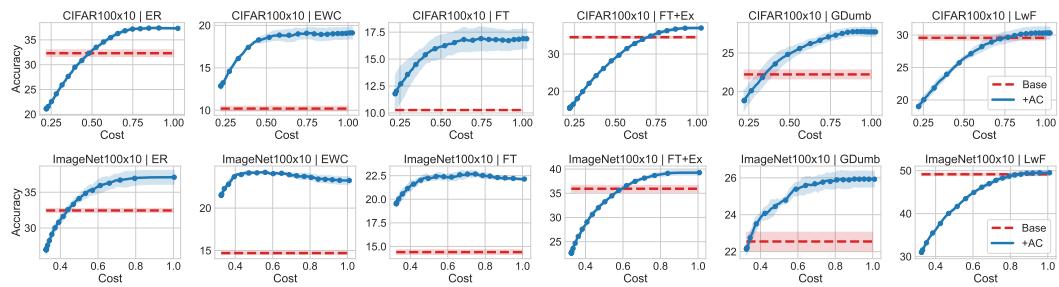
1108 **Figure 14:** Dynamic inference plots for 50 task split of CIFAR100.
1109

1134 G ADDITIONAL DYNAMIC INFERENCE PLOTS

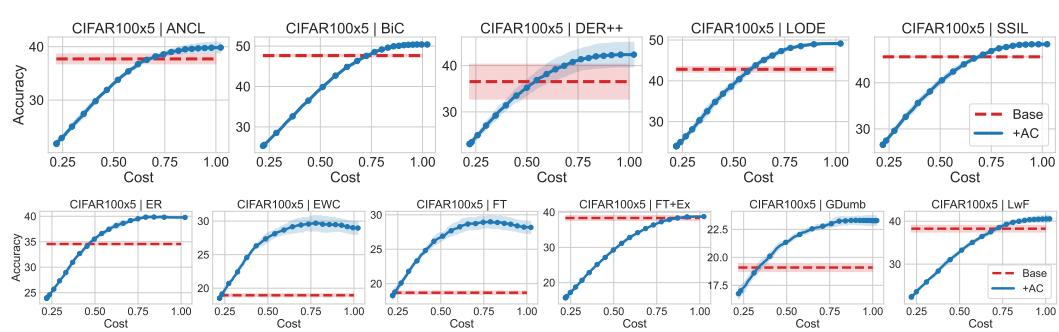
1135
 1136 In this section, we present dynamic inference plots obtained for experiments performed in Section 4
 1137 or complementary to those experiments.
 1138

1139 G.1 STANDARD BENCHMARKS

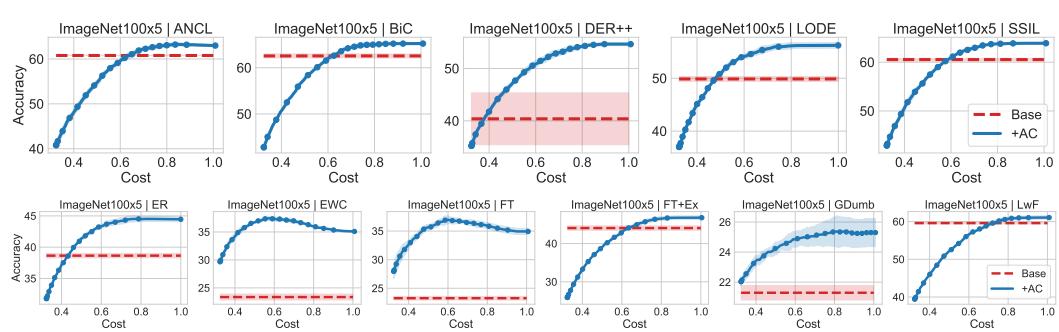
1140
 1141 In Figure 5, we present dynamic inference results for 10 task splits of CIFAR100 and ImageNet100
 1142 using ANCL, BiC, ER, LODE, and SSIL. In this section, we provide corresponding results for the
 1143 10-task split and the rest of the analyzed methods: EWC, FT, FT+Ex, GDumb, and LwF.
 1144



1145
 1146 Figure 15: Dynamic inference plots as in Figure 5 for additional continual learning methods ex-
 1147 tended with auxiliary classifiers on CIFAR100 and ImageNet100 split into 10 tasks.
 1148



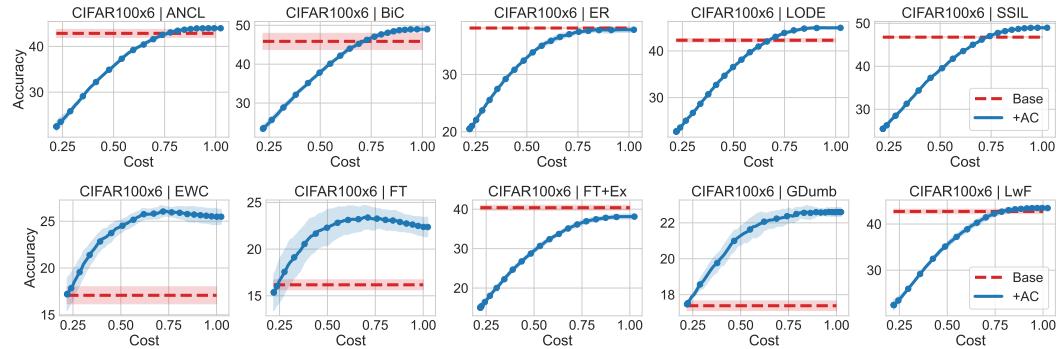
1149
 1150 Figure 16: Dynamic inference plots for continual learning methods extended with auxiliary classi-
 1151 fiers on CIFAR100 split into 5 tasks.
 1152



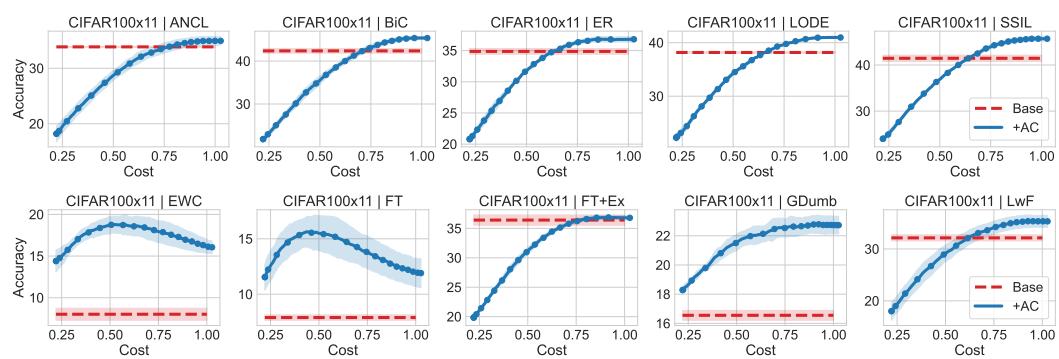
1153
 1154 Figure 17: Dynamic inference plots for continual learning methods extended with auxiliary classi-
 1155 fiers on ImageNet100 split into 5 tasks.
 1156

1188 **G.2 DYNAMIC INFERENCE PLOTS FOR WARM START SETTING**
 1189

1190 In this section, we provide dynamic accuracy plots for the warm start setting described in Section 4.
 1191



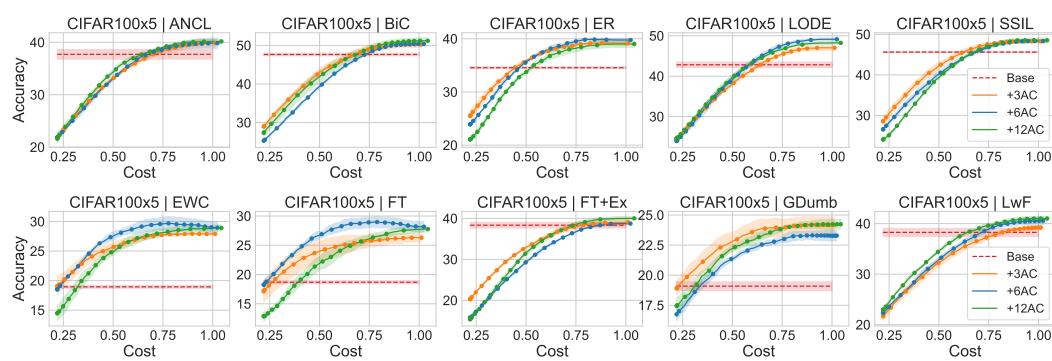
1204 Figure 18: Dynamic inference plots for continual learning methods extended with auxiliary
 1205 classifiers compared with the baselines corresponding to the results from the Table 2 for CIFAR100
 1206 starting from 50 classes and then training on the sequence of 5 tasks with 10 classes each.



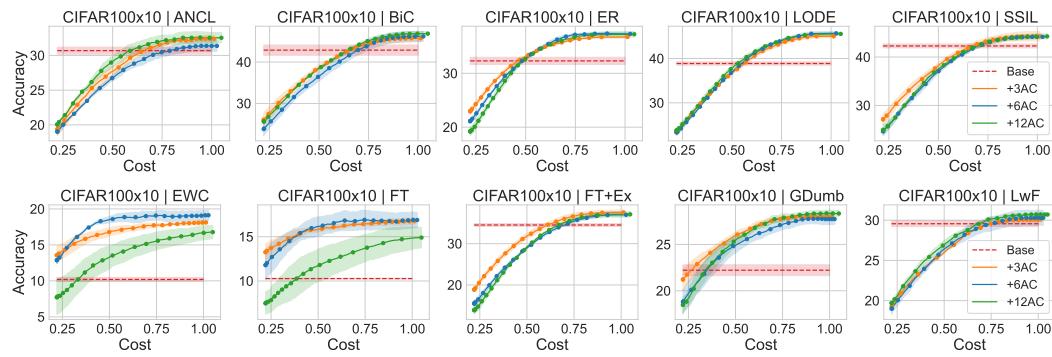
1220 Figure 19: Dynamic inference plots for continual learning methods extended with auxiliary classifiers
 1221 compared with the baselines corresponding to the results from the Table 2 for CIFAR100
 1222 starting from 50 classes and then training on the sequence of 10 tasks with 5 classes each.

1242 **G.3 DYNAMIC ACCURACY WITH DIFFERENT NUMBER OF ACs**
 1243

1244 In this section, we provide dynamic accuracy plots for experiments performed with different num-
 1245 bers of ACs in Section 4.



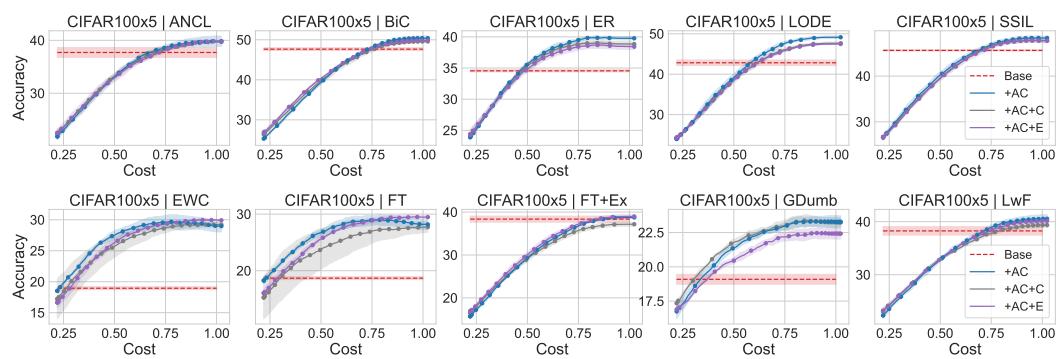
1258 Figure 20: Dynamic inference plots for continual learning methods extended using a different num-
 1259 ber of auxiliary classifiers on CIFAR100 split in 5 tasks.
 1260



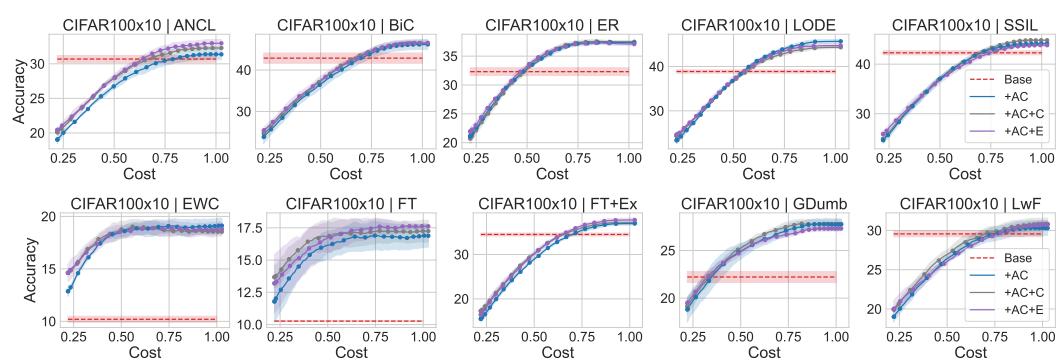
1273 Figure 21: Dynamic inference plots for continual learning methods extended using a different num-
 1274 ber of auxiliary classifiers on CIFAR100 split in 10 tasks.
 1275

1296 **G.4 DYNAMIC ACCURACY FOR DIFFERENT CLASSIFIER ARCHITECTURES**
 1297

1298 In this section, we provide dynamic accuracy plots for experiments performed with different classifier
 1299 architectures in Section 4.

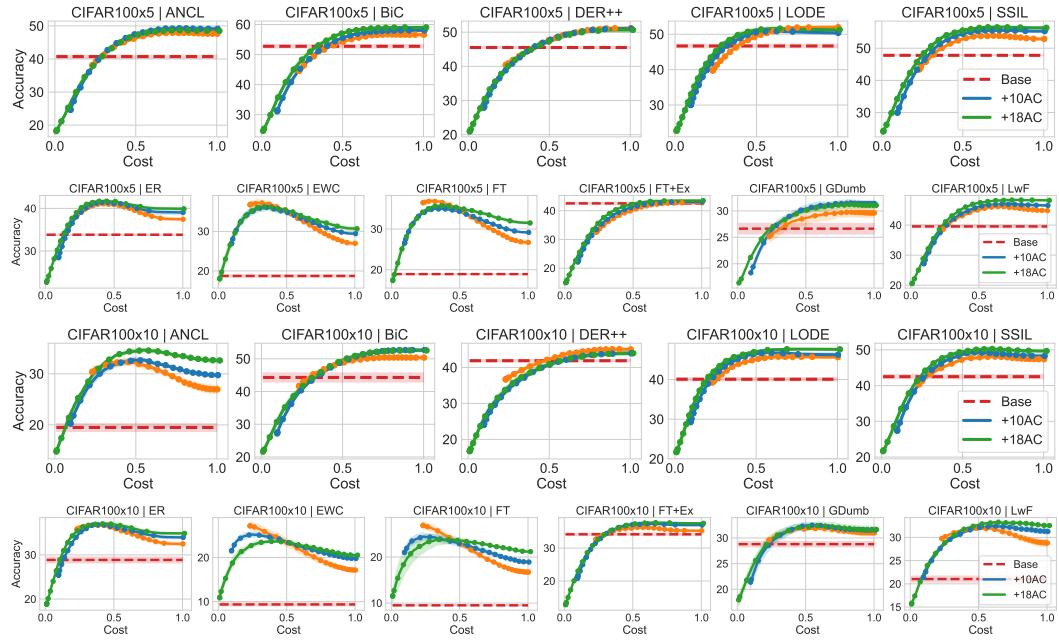


1300
 1301 Figure 22: Dynamic inference plots for continual learning methods extended with different auxiliary
 1302 classifier architecture on CIFAR100 split in 5 tasks.
 1303
 1304



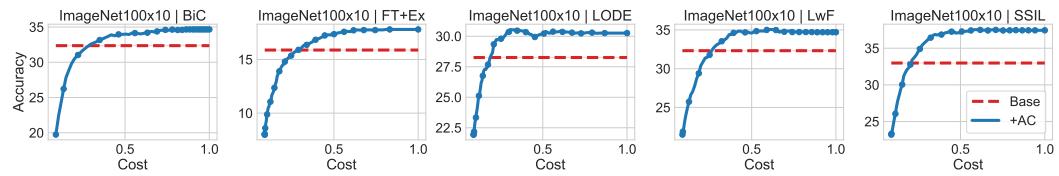
1305
 1306 Figure 23: Dynamic inference plots for continual learning methods extended with different auxiliary
 1307 classifier architecture on CIFAR100 split in 10 tasks.
 1308
 1309
 1310
 1311

1350
 1351 **G.5 VGG19 NETWORKS**
 1352



1373 Figure 24: Dynamic inference experiments on CIFAR100 with VGG19 network corresponding to
 1374 the ones from Section 4. As VGG networks are deeper, we attach either 6, 10, or 18 ACs. AC-
 1375 enhanced methods perform even better than in the case of ResNet, matching the accuracy of the
 1376 base method at a small fraction of its compute and outperforming the baseline significantly at the
 1377 full computational budget.

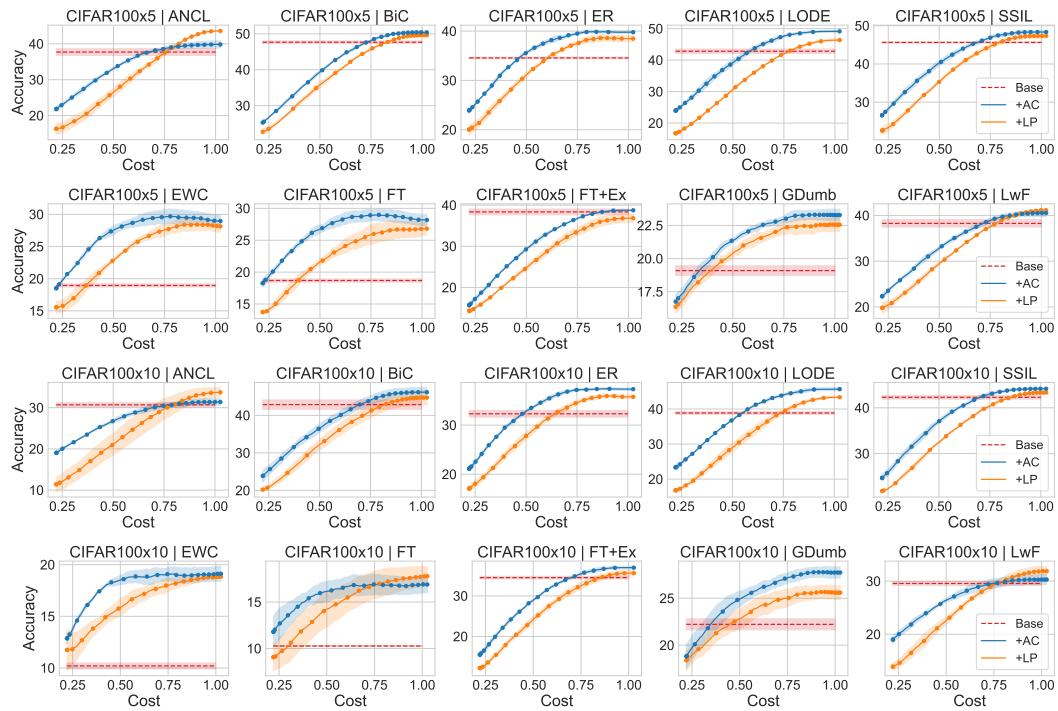
1378
 1379 **G.6 VISION TRANSFORMERS**
 1380



1388 Figure 25: Dynamic inference plots for several continual learning methods extended with auxil-
 1389 iary classifiers compared with the baselines using Vision Transformer trained from scratch on Im-
 1390 ageNet100 split into 5 tasks.

1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403

1404
 1405 **H EXPERIMENTAL COMPARISON BETWEEN OUR METHOD AND LINEAR
 1406 PROBING**



1430
 1431 Figure 26: Dynamic inference plots for several continual learning methods extended with auxiliary
 1432 classifiers when using auxiliary classifiers with enabled gradient propagation (AC) or without (LP).
 1433 Enabled propagation generally improves the results at low to mid-computation budgets, with ANCL
 1434 and LwF being the only outliers at high computational budgets. This variation could be explained
 1435 by the variance in intermediate classifier predictions shown in Section 2.2.

1436 In Section 2.4, we advocate for training the network and ACs jointly with enabled gradient propa-
 1437 gation, as it leads to better performance of individual classifiers. In this section, we investigate the
 1438 final performance of linear probing classifiers in comparison with jointly trained ACs on CIFAR100.
 1439 For ACs use the same setup as in Section 4, and in the case of linear probing the only difference is
 1440 that the classifiers are trained without gradient propagation. We show the final performance of both
 1441 settings in Table 12, and also demonstrate their cost-accuracy characteristics in Figure 26. Aside
 1442 from distillation-based exemplar-free methods, ACs outperform probing accuracy, and in all cases,
 1443 networks learned with probing still achieve an improvement upon the baseline. However, the dy-
 1444 namic accuracy curves highlight that enabled gradient propagation allows AC methods to achieve
 1445 greater accuracy at lower computational costs due to the ability to learn better early classifiers.

1446
 1447 **I AC ARCHITECTURE AND TRAINING DETAILS**
 1448

1449 In our main experiments, we follow the insights from Kaya et al. (2019) and Wójcik et al. (2023)
 1450 in our design of multi-classifier networks. We place the ACs after layers that perform roughly 15%,
 1451 30%, 45%, 60%, 75%, 90% of the computations of the full network. For convolutional networks,
 1452 ACs are composed of pooling layers to reduce the input size for the fully connected networks that
 1453 produce the predictions. For experiments on ViT, we apply a fully connected classifier on top of the
 1454 LayerNorm layer on the first token. All the classifiers in our model are composed of heads for each
 1455 task, and we add a new head upon encountering a new task.

1456 Our main objective used for training the network on any given task is a weighted sum of losses for
 1457 each classifier. For continual learning methods, we use the additional losses alongside the cross-
 1458 entropy and weigh the total loss. We train the model for each task jointly with all the ACs, updating

1458
 1459 Table 12: Comparison between final results when using intermediate classifiers trained together
 1460 with the network (AC) or trained with linear probing (LP). Training classifiers together generally
 1461 yields better performance, with the only noticeable exception being exemplar-free distillation-based
 1462 methods (ANCL and LwF), which could be caused by significant variance in per-task accuracy of
 1463 intermediate classifiers as shown in Section 2.2.

Method	FT	FT+Ex	GDumb	ANCL	BiC	ER	EWC	LwF	LODE	SSIL	Avg
CIFAR100x5											
Base	18.68 \pm 0.31	38.35 \pm 0.86	19.09 \pm 0.44	37.71 \pm 1.14	47.66 \pm 0.43	34.55 \pm 0.21	18.95 \pm 0.29	38.26 \pm 0.98	42.82 \pm 0.84	45.62 \pm 0.16	34.17 \pm 0.27
+LP	26.82 \pm 1.19	36.83 \pm 1.27	22.56 \pm 0.63	43.60\pm0.19	49.62 \pm 0.07	38.47 \pm 0.78	28.13 \pm 1.11	41.13\pm0.33	46.35 \pm 0.45	47.33 \pm 0.60	38.09 \pm 0.45
+AC	28.18\pm1.07	38.75\pm0.26	23.29\pm0.54	39.83 \pm 1.22	50.40\pm0.68	39.77\pm0.32	28.96\pm1.13	40.55 \pm 0.95	49.13\pm0.35	48.35\pm0.50	38.72\pm0.61
CIFAR100x5											
Base	10.27 \pm 0.05	34.51 \pm 0.40	22.22 \pm 0.72	30.69 \pm 0.62	42.87 \pm 1.51	32.31 \pm 0.82	10.20 \pm 0.35	29.56 \pm 0.44	38.87 \pm 0.45	42.29 \pm 0.49	29.38 \pm 0.26
+LP	17.77\pm1.30	35.62 \pm 0.89	25.60 \pm 0.91	33.72\pm1.38	44.74 \pm 2.31	35.78 \pm 0.46	18.84 \pm 0.19	31.88\pm1.11	43.37 \pm 0.31	43.33 \pm 0.08	33.06 \pm 0.17
+AC	16.88 \pm 1.08	36.97\pm0.39	27.74\pm0.73	31.37 \pm 0.94	46.19\pm1.47	37.32\pm0.28	19.12\pm0.88	30.31 \pm 1.14	45.67\pm0.52	44.17\pm0.28	33.57\pm0.22

1471
 1472 all the parameters of the network. We follow the weight scheduler from Kaya et al. (2019) and
 1473 progressively increase loss weights for different ACs over the training phase to the values matching
 1474 their computational cost (e.g. the weight for the first classifier for ResNet32 would increase up to
 1475 0.15, for the second classifier to 0.30, and so on). For ResNet18 we use 6 ACs and set weights
 1476 to [0.3, 0.4, 0.55, 0.65, 0.8, 0.9], as the network contains only 8 blocks whose computational cost
 1477 distributes approximately like that. For experiments with 12 ACs, we attach classifiers to all blocks
 1478 L1.B3-L3.B4 and interpolate the weights from the standard setting. For 3 ACs, we use blocks
 1479 L1.B3, L2.B2, and L3.B1 with weights [0.15, 0.45, 0.75]. When training ViT or VGG networks, for
 1480 each model block we use multiplies of a given base weight (e.g. 0.08 for ViT-base, 0.05 or 0.09
 1481 for 18 and 10 AC setup for VGG19). For example, we set the AC weights for 11 ACs in ViT as
 1482 [0.08, 0.16, ..., 0.80, 0.88]. Different loss weights for each AC serve to stabilize the training and
 1483 mitigate overfitting in the earlier layers, which may have lower learning capacity.

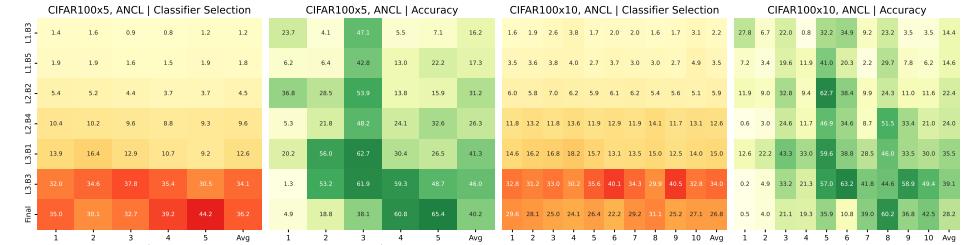
1484 We train the ResNet32 models on CIFAR100 for 200 epochs on each task, using SGD optimizer
 1485 with a batch size of 128 with a learning rate initialized to 0.1 and decayed by a rate of 0.1 at the
 1486 60th, 120th, and 160th epochs. For training ResNet18 on ImageNet100, we change the scheduler
 1487 to cosine with a linear warmup and train for 100 epochs with 5 epochs of warmup, as we find it to
 1488 converge to similar results in a shorter time. For ViT, we use AdamW and train each task for 100
 1489 epochs with a learning rate of 0.01 and batch size of 64. We also use a cosine scheduler with a linear
 1490 warmup for 5 epochs. We use a fixed memory of 2000 exemplars selected with herding (Rebuffi
 1491 et al., 2017). For ER each batch is balanced between old and new data, and for SSIL we use a
 1492 4:1 ratio of new to old data. Otherwise, for other exemplar-based methods, we follow the standard
 1493 FACIL procedure for exemplars and just add them to the training data without any balancing.

1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511

1512
1513**J CLASSIFIER SELECTION AND ACCURACY DURING INFERENCE**1514
1515
1516
1517
1518
1519

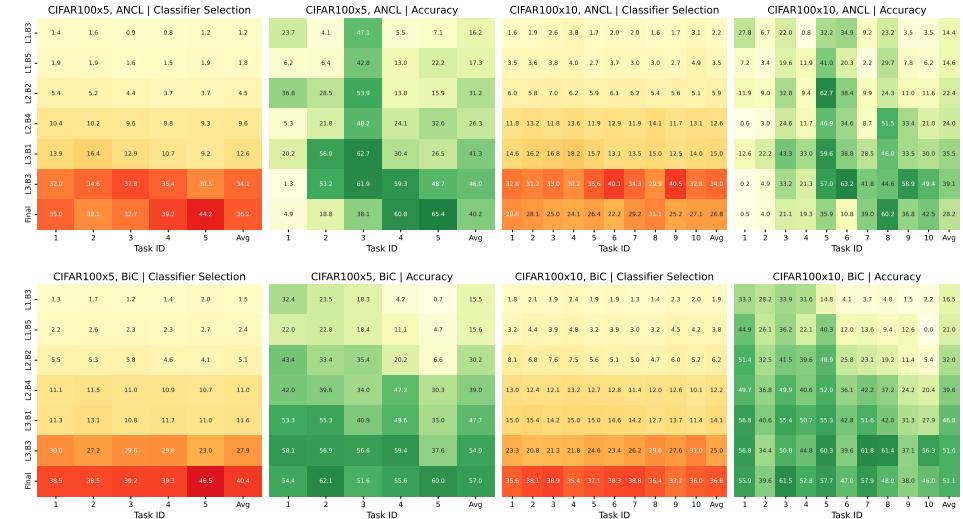
In Figures 27 and 28 we present the final distribution of the selected classifiers per task and on average across the test dataset for the CIFAR100 experiments presented in Table 1. We also present the corresponding accuracy of each classifier, assuming the classifier was selected for static inference as described in Section 3. We observe that early classifiers are not selected that often. However, when selected, intermediate classifiers (L2.B2-L3.B3) usually exhibit accuracy on par or better than the final classifier. This hints at the cause of the improvement observed from adding the ACs.

1520

1521
1522

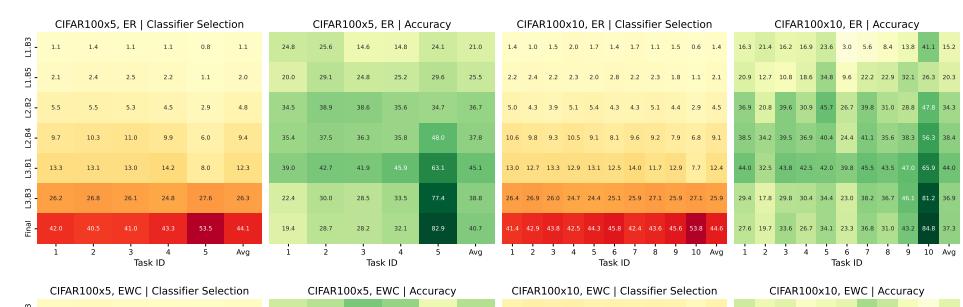
1528

1529

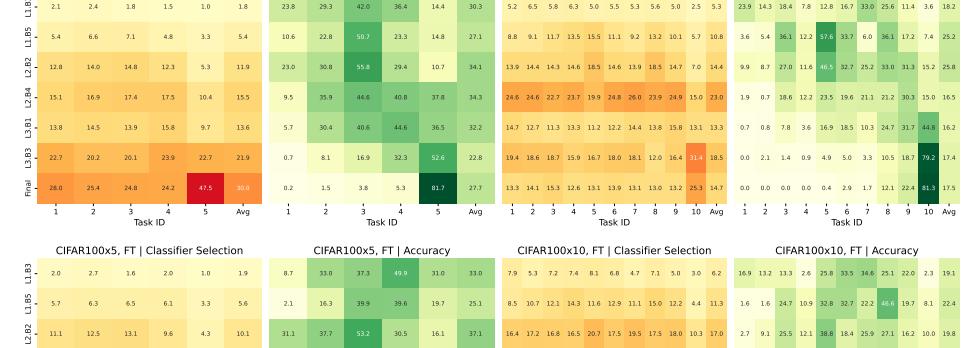


1537

1538



1545



1566

1567

1568

1569

1570

1571

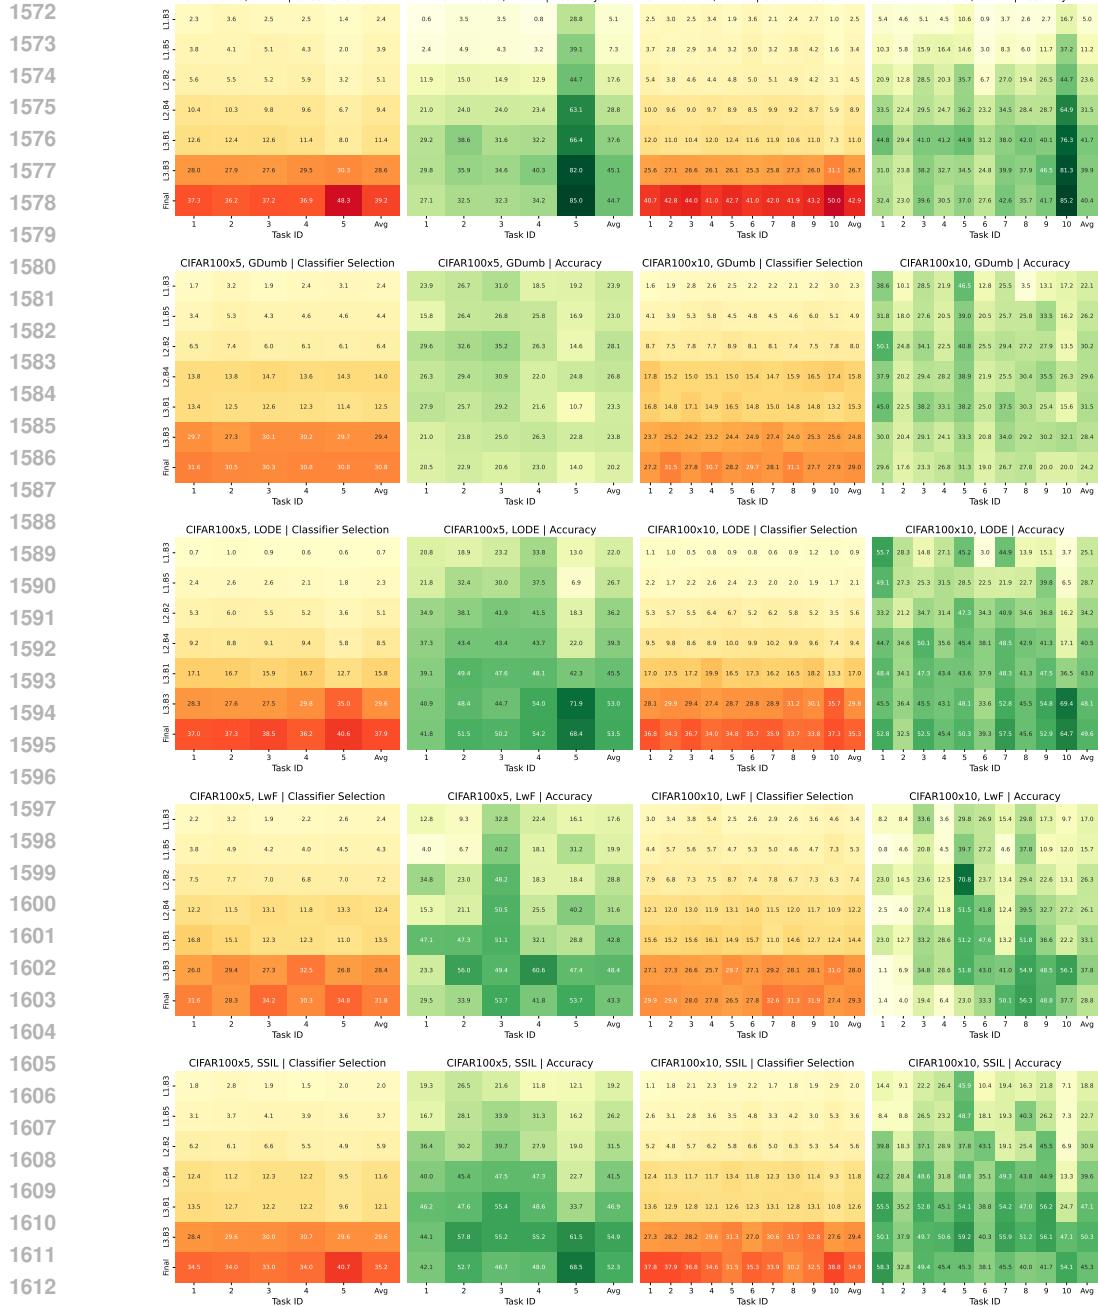


Figure 28: Distribution and accuracy of classifiers for FT+Ex, GDumb, LODE, LwF and SSIL.

1615

1616

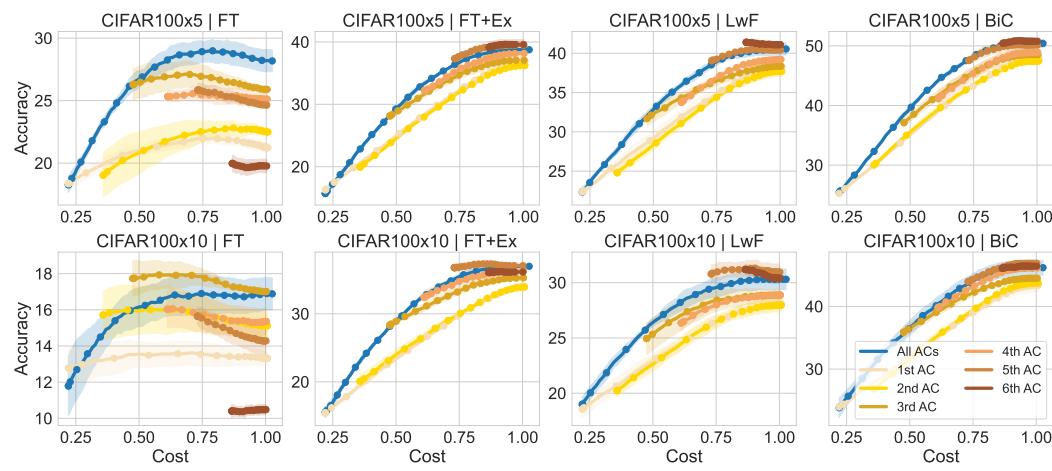
1617

1618

1619

1620 **K SINGLE AC ABLATION**
 1621

1622 In this section, we perform leave-one-out ablation of the setting explored in Section 4 for CIFAR100
 1623 on methods explored in Section 2. Namely, we train the model with only one auxiliary classifier out
 1624 of the original six, with the classifier weight equal to 1. during training and present dynamic infer-
 1625 ence results in Figure 29. For non-naive methods, we observe that 5th or 6th AC achieve the best
 1626 performance. Interestingly for finetuning, later AC yield lower performance, which is consistent
 1627 with our observations on more native stability in early layers. All tested AC setups achieve com-
 1628 parable performance at the full computational budget, but compared to our based setup of using all 6
 1629 ACs they tend to underperform at lower compute budget. Slightly better performance of single AC
 1630 setup for FT+Ex and LwF hints that AC placement in our work could be further optimized. How-
 1631 ever, overall similar performance across all the tested scenarios prove the robustness of our idea.
 1632



1648 Figure 29: Leave one out AC ablation for FT, FT+Ex, LwF and BiC.
 1649

1650
 1651
 1652
 1653
 1654
 1655
 1656
 1657
 1658
 1659
 1660
 1661
 1662
 1663
 1664
 1665
 1666
 1667
 1668
 1669
 1670
 1671
 1672
 1673