Supplementary Material

Appendix A Results Tables

Table A1 Performance comparison of ProtoSleepNet and SeqSleepNet on the clean test set across all evaluated datasets. Results are reported in % for accuracy (Acc.), macro-averaged F1-score (F1), recall (RC), and precision (PR), whereas Cohen's κ (CK) is reported as a unitless agreement coefficient. For the Alzheimer Sleep Dataset (ASD) and the KU Leuven Parkinson dataset (KPD), the label pairs denote the clinical cohort used to evaluate the sleep-staging model: AD–AD (respectively PD–PD) refers to the Alzheimer's disease (respectively Parkinson's disease) cohort, while AD–HC and PD–HOA refer to the corresponding control cohorts, namely healthy controls (HC) and healthy older adults (HOA). Overall, ProtoSleepNet with 15 prototypes achieves performance comparable to SeqSleepNet across most datasets. Increasing the number of prototypes to 65 yields consistent gains in Acc., CK, and F1, and results in ProtoSleepNet65 outperforming SeqSleepNet on the majority of datasets, suggesting that a richer prototype set improves representational capacity and downstream generalization on the clean test condition.

Model	Dataset	Acc.	CK	F1	RC	PR
ProtoSleepNet15	ASD - AD	64.08	0.51	62.62	64.08	71.48
	ASD - HC	82.32	0.76	82.14	82.32	86.44
	KPD - HOA	85.06	0.78	84.56	85.06	85.48
	KPD - PD	79.61	0.69	78.79	79.61	82.49
	DCSM	89.09	0.83	88.73	89.09	89.82
	HMC	76.89	0.68	75.85	76.89	77.96
	MASS	84.00	0.76	83.01	84.00	83.89
	MESA	84.81	0.77	84.06	84.81	85.70
	MROS	88.53	0.82	88.10	88.53	88.93
	SHHS	86.50	0.80	86.14	86.50	87.79
	SleepEDF	81.54	0.74	81.01	81.54	82.88
	WSC	83.65	0.72	81.59	83.65	82.77
ProtoSleepNet65	ASD - AD	67.50	0.53	66.94	67.50	73.20
	ASD - HC	83.79	0.78	83.11	83.79	86.81
	KPD - HOA	88.20	0.82	87.57	88.20	88.18
	KPD - PD	81.56	0.72	80.73	81.56	83.29
	DCSM	89.67	0.84	89.43	89.67	90.59
	HMC	80.21	0.73	79.00	80.21	80.85
	MASS	85.79	0.79	85.36	85.79	86.02
	MESA	86.28	0.80	85.69	86.28	87.38
	MROS	89.34	0.83	88.89	89.34	89.67
	SHHS	87.71	0.82	87.47	87.71	89.03
	SleepEDF	83.66	0.77	83.26	83.66	84.95
	WSC	86.49	0.77	85.10	86.49	85.99
SeqSleepNet	ASD - AD	62.53	0.49	60.69	62.53	65.21
	ASD - HC	80.91	0.74	80.74	80.91	84.35
	KPD - HOA	88.05	0.82	87.46	88.05	88.17
	KPD - PD	78.70	0.66	74.88	78.70	76.36
	DCSM	89.36	0.84	88.90	89.36	89.91
	HMC	75.17	0.67	73.75	75.17	76.42
	MASS	86.06	0.79	85.35	86.06	86.33
	MESA	84.06	0.76	83.84	84.06	86.38
	MROS	87.56	0.81	86.95	87.56	88.32
	SHHS	86.46	0.80	86.30	86.46	88.29
	SleepEDF	82.14	0.75	81.87	82.14	84.48
	WSC	87.50	0.78	86.11	87.50	87.20

Table A2 Robustness to missing-sensor information assessed via a channel-occlusion stress test. For each dataset, we randomly occlude a fixed fraction of input channels at inference time, using two occlusion probability for each epoch (25% and 50%), and report the resulting metrics as $Acc_{25/50}$, $F1_{25/50}$, $RC_{25/50}$, and $PR_{25/50}$ in %, while $CK_{25/50}$ denotes Cohen's κ (unitless, i.e., not in %) measuring agreement beyond chance [**?**]. Across datasets, ProtoSleepNet exhibits substantially smaller performance degradation than SeqSleepNet as the occlusion rate increases from 25% to 50%, indicating higher resilience to partial channel loss. In particular, SeqSleepNet shows pronounced drops at 50% occlusion (e.g., on ASD–HC, $Acc_{25} = 74.45\%$ to $Acc_{50} = 60.56\%$; on MESA, $Acc_{25} = 64.95\%$ to $Acc_{50} = 47.81\%$), whereas ProtoSleepNet15 and ProtoSleepNet65 remain comparatively stable over the same perturbation (e.g., ASD–HC: 82.27% to 81.52% and 82.57% to 81.54%, respectively; MESA: 84.05% to 82.90% and 85.52% to 84.43%). Moreover, increasing the number of prototypes from 15 to 65 systematically improves robustness under occlusion, yielding the strongest overall performance at both occlusion levels (notably on KPD–HOA and on multiple population cohorts such as SHHS and WSC).

Model	Dataset	Acc_{25}	Acc_{50}	CK_{25}	CK_{50}	$F1_{25}$	$F1_{50}$	RC_{25}	RC_{50}	PR_{25}	PR_{50}
ProtoSleepNet15	ASD - AD	62.69	62.66	0.49	0.48	61.22	60.79	62.69	62.66	70.39	69.11
	ASD - HC	82.27	81.52	0.76	0.75	81.91	81.09	82.27	81.52	86.27	85.09
	KPD - HOA	84.81	84.17	0.77	0.76	84.09	83.42	84.81	84.17	84.74	84.07
	KPD - PD	78.96	78.05	0.68	0.66	77.65	76.06	78.96	78.05	81.00	79.83
	DCSM	88.72	88.00	0.83	0.81	88.11	87.13	88.72	88.00	89.11	88.13
	HMC	76.18	73.82	0.67	0.64	74.86	72.16	76.18	73.82	77.00	75.19
	MASS	83.33	81.74	0.75	0.72	82.27	80.59	83.33	81.74	83.30	81.83
	MESA	84.05	82.90	0.76	0.74	83.28	82.18	84.05	82.90	84.89	83.73
	MROS	87.87	86.90	0.81	0.80	87.40	86.34	87.87	86.90	88.24	87.23
	SHHS	85.73	84.38	0.79	0.77	85.34	83.91	85.73	84.38	86.96	85.63
	SleepEDF	81.31	80.45	0.73	0.72	80.54	79.38	81.31	80.45	82.32	81.16
	WSC	83.24	82.08	0.71	0.69	81.40	80.54	83.24	82.08	82.49	81.53
ProtoSleepNet65	ASD - AD	66.29	64.83	0.52	0.49	65.29	62.91	66.29	64.83	72.58	71.70
	ASD - HC	82.57	81.54	0.76	0.74	82.05	80.89	82.57	81.54	85.93	85.17
	KPD - HOA	87.44	86.31	0.81	0.79	86.83	85.45	87.44	86.31	87.45	86.02
	KPD - PD	80.92	79.13	0.71	0.68	79.60	77.66	80.92	79.13	82.74	79.49
	DCSM	89.25	88.40	0.84	0.82	88.88	87.76	89.25	88.40	89.92	88.95
	HMC	78.62	74.36	0.70	0.64	77.43	72.75	78.62	74.36	79.76	77.16
	MASS	84.92	83.45	0.77	0.75	84.41	82.78	84.92	83.45	85.04	83.45
	MESA	85.52	84.43	0.78	0.77	84.82	83.51	85.52	84.43	86.48	85.32
	MROS	88.85	87.92	0.83	0.81	88.22	86.97	88.85	87.92	89.07	88.00
	SHHS	87.03	85.81	0.81	0.79	86.65	85.14	87.03	85.81	88.23	86.95
	SleepEDF	83.44	82.43	0.76	0.74	82.75	81.20	83.44	82.43	84.36	82.95
	WSC	85.87	84.91	0.76	0.74	84.40	83.35	85.87	84.91	85.32	84.38
SeqSleepNet	ASD - AD	63.42	55.67	0.43	0.27	60.23	48.82	63.42	55.67	68.26	61.46
	ASD - HC	74.45	60.56	0.65	0.45	73.37	55.21	74.45	60.56	82.29	68.78
	KPD - HOA	82.29	69.66	0.73	0.51	80.79	65.21	82.29	69.66	82.65	72.36
	KPD - PD	67.17	55.82	0.50	0.35	61.92	48.56	67.17	55.82	68.55	57.04
	DCSM	84.85	76.05	0.76	0.61	83.62	72.41	84.85	76.05	85.74	78.78
	HMC	66.67	49.62	0.54	0.31	64.27	43.38	66.67	49.62	72.65	63.56
	MASS	76.30	56.04	0.63	0.34	74.46	51.09	76.30	56.04	79.86	68.08
	MESA	64.95	47.81	0.45	0.15	60.53	37.35	64.95	47.81	70.91	49.90
	MROS	81.75	70.80	0.71	0.52	79.99	66.18	81.75	70.80	82.82	73.99
	SHHS	73.28	52.43	0.60	0.30	71.72	46.68	73.28	52.43	78.92	62.26
	SleepEDF	77.34	62.97	0.67	0.46	75.06	58.07	77.34	62.97	79.30	69.55
	WSC	81.65	69.00	0.68	0.48	79.35	65.57	81.65	69.00	81.99	74.27

Table A3 Modality-specific occlusion analysis separating EEG from EOG-EMG channels. Condition (1) ($\cdot_1$) corresponds to occluding the EEG channels (thereby forcing the models to rely primarily on EOG/EMG information), whereas condition (2) ($\cdot_2$) corresponds to occluding the EOG–EMG channels (thereby forcing reliance on EEG). This ablation is motivated by the complementary role of EEG, EOG and EMG in standard PSG-based sleep staging, where EEG provides core cortical rhythms while EOG and EMG contribute distinctive cues (notably for REM and wake) [**?**]. We report $Acc_{1/2}$, $F1_{1/2}$, $RC_{1/2}$, and $PR_{1/2}$ in %, and Cohen's κ ($CK_{1/2}$) as a unitless chance-corrected agreement measure [**?**]. Across datasets, occluding EEG (condition 1) is generally more detrimental than occluding EOG–EMG (condition 2), especially for SeqSleepNet (e.g., DCSM $Acc_1 = 83.14\%$ vs. $Acc_2 = 67.78\%$; HMC 67.00% vs. 42.19%; MASS 66.84% vs. 45.99%), indicating a strong dependence on EEG features for temporal modeling in that architecture. In contrast, ProtoSleepNet remains markedly more robust under both modality dropouts, with smaller gaps between $\cdot_1$ and $\cdot_2$ on most datasets and consistently higher κ and F_1 under severe single-modality conditions, suggesting that prototype-based representations improve stability when one physiological modality is unavailable. Notably, ProtoSleepNet65 achieves the strongest overall resilience, often maintaining high performance even when either EEG or EOG–EMG is fully occluded (e.g., SHHS $Acc_1 = 82.80\%$, $Acc_2 = 84.54\%$; WSC 82.94%, 84.86%).

Model	Dataset	Acc_1	Acc_2	CK_1	CK_2	$F1_1$	$F1_2$	RC_1	RC_2	PR_1	PR_2
ProtoSleepNet15	ASD - AD	59.82	60.46	0.45	0.44	58.02	57.96	59.82	60.46	64.26	64.29
	ASD - HC	75.39	81.63	0.68	0.74	73.99	80.53	75.39	81.63	84.49	85.63
	KPD - HOA	82.06	79.27	0.73	0.69	81.34	76.35	82.06	79.27	82.46	78.36
	KPD - PD	76.41	75.40	0.64	0.62	74.28	72.59	76.41	75.40	76.22	74.89
	DCSM	86.62	87.08	0.80	0.80	86.20	85.70	86.62	87.08	87.98	87.34
	HMC	65.46	72.43	0.52	0.62	63.00	69.91	65.46	72.43	68.23	73.97
	MASS	79.36	77.80	0.69	0.66	78.09	76.18	79.36	77.80	81.24	79.00
	MESA	82.13	79.35	0.73	0.69	81.03	78.30	82.13	79.35	83.09	81.32
	MROS	85.83	83.85	0.78	0.75	85.32	82.44	85.83	83.85	86.86	84.62
	SHHS	81.77	82.00	0.73	0.73	81.31	80.59	81.77	82.00	84.13	83.42
	SleepEDF	77.49	74.80	0.68	0.62	76.40	71.40	77.49	74.80	79.12	77.47
	WSC	80.52	79.33	0.67	0.65	78.64	78.13	80.52	79.33	80.36	80.41
ProtoSleepNet65	ASD - AD	56.40	62.16	0.40	0.42	54.07	58.95	56.40	62.16	63.61	67.67
	ASD - HC	84.00	78.80	0.78	0.71	83.52	77.22	84.00	78.80	84.44	79.79
	KPD - HOA	83.61	83.24	0.75	0.75	82.45	82.24	83.61	83.24	83.72	83.64
	KPD - PD	79.61	74.10	0.68	0.61	77.18	72.73	79.61	74.10	77.17	78.24
	DCSM	86.98	88.19	0.80	0.82	86.74	87.18	86.98	88.19	88.77	88.65
	HMC	67.86	70.46	0.55	0.59	65.49	67.76	67.86	70.46	71.18	75.84
	MASS	80.95	80.44	0.71	0.70	80.51	79.67	80.95	80.44	82.74	81.92
	MESA	83.32	82.15	0.75	0.73	82.70	81.07	83.32	82.15	84.86	83.67
	MROS	86.53	86.08	0.79	0.78	85.48	85.01	86.53	86.08	86.80	86.83
	SHHS	82.80	84.54	0.74	0.77	81.76	83.89	82.80	84.54	84.31	86.27
	SleepEDF	79.63	78.95	0.71	0.69	78.84	77.82	79.63	78.95	80.72	81.45
	WSC	82.94	84.86	0.71	0.74	80.94	83.64	82.94	84.86	82.29	85.07
SeqSleepNet	ASD - AD	40.88	52.77	0.18	0.19	36.14	42.49	40.88	52.77	45.43	45.47
	ASD - HC	71.55	54.88	0.61	0.36	70.63	42.89	71.55	54.88	76.73	37.84
	KPD - HOA	63.87	67.99	0.47	0.49	60.01	60.46	63.87	67.99	72.22	60.24
	KPD - PD	62.21	51.90	0.42	0.31	56.24	42.53	62.21	51.90	57.89	42.40
	DCSM	83.14	67.78	0.73	0.47	81.46	60.44	83.14	67.78	84.17	65.68
	HMC	67.00	42.19	0.54	0.20	63.81	30.74	67.00	42.19	68.51	36.09
	MASS	66.84	45.99	0.43	0.21	61.36	38.69	66.84	45.99	70.04	46.97
	MESA	53.17	46.98	0.24	0.12	44.54	34.64	53.17	46.98	53.42	38.62
	MROS	66.50	70.81	0.44	0.52	60.06	64.11	66.50	70.81	65.30	65.79
	SHHS	52.37	56.65	0.23	0.35	44.12	48.13	52.37	56.65	57.45	53.19
	SleepEDF	59.99	59.17	0.39	0.40	52.15	51.73	59.99	59.17	64.25	60.15
	WSC	69.13	65.37	0.39	0.40	62.12	58.54	69.13	65.37	66.41	60.65

Table A4 Comparison between *post-hoc* prototype-learning constraints applied to a pre-trained SeqSleepNet and the *built-in* prototype-learning mechanism of ProtoSleepNet. ProtoNet and SimVQ are evaluated by extracting prototypes *after* training and, at evaluation time, forcing SeqSleepNet's decision process to be mediated by these learned prototypes (i.e., predictions are constrained to follow the space quantization determined by the learned protypes). The reported scores (Acc., CK, F1, PR, RC) quantify performance under this forced prototypical decision-making; the columns marked Δ indicate the performance change induced by prototyping relative to the unconstrained SeqSleepNet baseline evaluated in its original continuous regime (negative Δ denotes degradation, positive Δ denotes improvement). For ProtoSleepNet, Δ is instead defined as the difference with respect to SeqSleepNet (ProtoSleepNet minus SeqSleepNet), reflecting the net gain of the proposed *end-to-end* prototypical architecture over the non-prototypical baseline; note that CK is Cohen's κ, a unitless chance-corrected agreement coefficient rather than a percentage. Overall, enforcing post-hoc prototyping on SeqSleepNet yields large and systematic drops across most datasets (particularly for ProtoNet, and to a slightly lesser extent for SimVQ), whereas ProtoSleepNet achieves neutral or positive Δ on many cohorts and attains the strongest performance among the prototype-based mechanisms considered. These results support the interpretation that, when representations are learned to optimize a continuous decision head, the intermediate feature spaces can be highly fragmented and not naturally amenable to a faithful post-hoc partition into a small set of prototypes; conversely, ProtoSleepNet learns prototypes as part of the predictive pathway, encouraging a latent organization that remains both performant and compatible with prototype-based decision making.

Model	Dataset	Acc.	Acc. Δ	CK	CK Δ	F1	F1 Δ	PR	PR Δ	RC	RC Δ
ProtoNet	ASD - AD	57.21	-5.32	0.33	-0.15	51.72	-8.97	65.92	0.71	57.21	-5.32
	ASD - HC	60.16	-20.75	0.46	-0.29	53.52	-27.22	65.25	-19.10	60.16	-20.75
	DCSM	77.38	-11.98	0.64	-0.20	72.95	-15.95	76.47	-13.44	77.38	-11.98
	HMC	51.21	-23.96	0.31	-0.35	44.05	-29.70	57.48	-18.94	51.21	-23.96
	MASS	68.19	-17.87	0.53	-0.26	64.06	-21.29	71.87	-14.46	68.19	-17.87
	MESA	76.08	-7.98	0.64	-0.12	75.25	-8.59	78.82	-7.56	76.08	-7.98
	MROS	84.65	-2.92	0.76	-0.05	83.61	-3.34	85.41	-2.91	84.65	-2.92
	KPD - HOA	75.39	-12.66	0.61	-0.22	70.55	-16.91	75.46	-12.71	75.39	-12.66
	KPD - PD	61.14	-17.57	0.46	-0.20	54.87	-20.01	69.25	-7.11	61.14	-17.57
	SHHS	84.02	-2.43	0.76	-0.04	83.21	-3.08	85.41	-2.88	84.02	-2.43
	SleepEDF	72.22	-9.92	0.59	-0.16	67.79	-14.09	74.25	-10.24	72.22	-9.92
	WSC	80.83	-6.67	0.66	-0.13	77.87	-8.24	79.09	-8.11	80.83	-6.67
SimVQ	ASD - AD	57.56	-4.97	0.35	-0.13	53.07	-7.62	69.75	4.54	57.56	-4.97
	ASD - HC	62.43	-18.49	0.49	-0.25	57.65	-23.09	71.93	-12.42	62.43	-18.49
	DCSM	76.91	-12.45	0.63	-0.21	72.16	-16.74	74.99	-14.91	76.91	-12.45
	HMC	52.30	-22.88	0.33	-0.34	45.42	-28.33	60.26	-16.16	52.30	-22.88
	MASS	70.76	-15.30	0.58	-0.21	68.62	-16.73	76.34	-9.99	70.76	-15.30
	MESA	77.02	-7.04	0.65	-0.11	75.66	-8.18	78.90	-7.48	77.02	-7.04
	MROS	84.90	-2.67	0.76	-0.04	83.99	-2.96	85.72	-2.61	84.90	-2.67
	KPD - HOA	79.39	-8.65	0.68	-0.14	76.22	-11.24	80.46	-7.71	79.39	-8.65
	KPD - PD	58.82	-19.88	0.42	-0.23	52.78	-22.10	66.16	-10.21	58.82	-19.88
	SHHS	83.58	-2.88	0.75	-0.05	82.78	-3.51	84.91	-3.38	83.58	-2.88
	SleepEDF	74.66	-7.48	0.63	-0.12	70.85	-11.02	77.66	-6.83	74.66	-7.48
	WSC	80.22	-7.28	0.65	-0.14	77.48	-8.63	79.19	-8.00	80.22	-7.28
ProtoSleepNet15	ASD - AD	64.08	1.55	0.51	0.02	62.62	1.93	71.48	6.27	64.08	1.55
	ASD - HC	82.32	1.40	0.76	0.02	82.14	1.40	86.44	2.09	82.32	1.40
	DCSM	89.09	-0.26	0.83	-0.00	88.73	-0.17	89.82	-0.09	89.09	-0.26
	HMC	76.89	1.72	0.68	0.02	75.85	2.10	77.96	1.54	76.89	1.72
	MASS	84.00	-2.06	0.76	-0.03	83.01	-2.34	83.89	-2.44	84.00	-2.06
	MESA	84.81	0.74	0.77	0.01	84.06	0.21	85.70	-0.68	84.81	0.74
	MROS	88.53	0.97	0.82	0.02	88.10	1.15	88.93	0.61	88.53	0.97
	KPD - HOA	85.06	-2.99	0.78	-0.05	84.56	-2.89	85.48	-2.69	85.06	-2.99
	KPD - PD	79.61	0.91	0.69	0.03	78.79	3.91	82.49	6.13	79.61	0.91
	SHHS	86.50	0.05	0.80	-0.00	86.14	-0.16	87.79	-0.51	86.50	0.05
	SleepEDF	81.54	-0.59	0.74	-0.01	81.01	-0.86	82.88	-1.61	81.54	-0.59
	WSC	83.65	-3.85	0.72	-0.07	81.59	-4.53	82.77	-4.43	83.65	-3.85
ProtoSleepNet65	ASD - AD	67.50	4.97	0.53	0.05	66.94	6.25	73.20	7.99	67.50	4.97
	ASD - HC	83.79	2.87	0.78	0.03	83.11	2.38	86.81	2.46	83.79	2.87
	DCSM	89.67	0.32	0.84	0.01	89.43	0.53	90.59	0.68	89.67	0.32
	HMC	80.21	5.03	0.73	0.06	79.00	5.25	80.85	4.43	80.21	5.03
	MASS	85.79	-0.27	0.79	-0.00	85.36	0.01	86.02	-0.32	85.79	-0.27
	MESA	86.28	2.21	0.80	0.03	85.69	1.85	87.38	1.00	86.28	2.21
	MROS	89.34	1.77	0.83	0.03	88.89	1.94	89.67	1.34	89.34	1.77
	KPD - HOA	88.20	0.15	0.82	0.00	87.57	0.12	88.18	0.01	88.20	0.15
	KPD - PD	81.56	2.86	0.72	0.06	80.73	5.85	83.29	6.93	81.56	2.86
	SHHS	87.71	1.26	0.82	0.02	87.47	1.17	89.03	0.74	87.71	1.26
	SleepEDF	83.66	1.52	0.77	0.02	83.26	1.39	84.95	0.47	83.66	1.52
	WSC	86.49	-1.01	0.77	-0.02	85.10	-1.01	85.99	-1.20	86.49	-1.01

Appendix B Local-instance to prototype explanation framework

Let $x \in \mathbb{R}^{C\times T\times F}$ denote a multichannel time–frequency representation with C channels, T time steps, and F frequency bins. The matching between a local instance and a prototype is defined in the latent embedding space learned by the model. Let $\phi(\cdot)$ be the projection implemented by the ProtoSleepNet front-end and temporal encoder. Concretely, the pipeline reshapes and filters the input through the fixed filterbank, then applies a sequence model and time masking, aggregates across channels, and yields a temporal embedding:

$$\phi(x) = \mathrm{mean}_{c=1}^{C}\Big(h_c(x) + \mathrm{ChannelMix}(h_c(x))\Big),$$

where $h_c(x)$ denotes the per-channel latent representation after filterbank and recurrent processing, and $\mathrm{ChannelMix}(\cdot)$ is the learned channel mixing operator. The resulting $\phi(x) \in \mathbb{R}^d$ is compared to each prototype $p_k \in \mathbb{R}^d$ from the learned codebook.

The similarity between an input instance and a prototype is measured by cosine similarity:

$$s_k(x) = \cos\big(\phi(x), p_k\big) = \frac{\phi(x)^\top p_k}{\|\phi(x)\|_2\,\|p_k\|_2}.$$

For a given instance, the assigned prototype is the maximizer $k^\star = \arg\max_k s_k(x)$.

To explain why a particular instance matches a prototype, we define a scalar function

$$f_k(x) = \cos\big(\phi(x), p_k\big),$$

and compute input attributions using Integrated Gradients with a zero baseline $x_0 = \mathbf{0}$. Since our data is standardized before being fitted to the network this corresponds to integrate the gradients from the input to the mean of the distribution:

$$\mathrm{IG}_k(x) = (x - x_0)\int_0^1 \nabla_x f_k\big(x_0 + \alpha(x - x_0)\big)\,d\alpha.$$

This yields a relevance tensor in the same shape as x, indicating which time–frequency bins and channels increase similarity to the prototype.

Example.

To provide a practical and simple example of our local-explanation framework we trained ProtoSleepNet on the MASS dataset and inspect one of its prototypes, which the model associates predominantly with stage N3 (confidence 97.3%).

To select a local instance to inspect we used the data-driven reconstruction of the prototype, obtained by scanning the data and selecting the local PSG segment whose embedding maximizes cosine similarity to the prototype. Starting from this representative instance, we compute prototype-conditioned relevance using Integrated Gradients, and visualize the resulting frequency-wise relevance maps alongside the reconstructed power spectra (Figure B1).

The plot highlights a pronounced low-frequency dominance in the EEG and EOG channels, with relevance concentrated in the slow-wave band, while EMG shows comparatively low and more diffuse relevance. This pattern is coherent with PSG-based sleep staging: N3 is characterized by high-amplitude slow waves in EEG and reduced muscle activity, hence a relevance peak at low frequencies for EEG/EOG and a subdued EMG contribution. The figure therefore provides a mechanistic explanation linking the prototype match to physiologically meaningful spectral features in the PSG, reinforcing the interpretability of prototype assignments in the clinically relevant N3 regime.

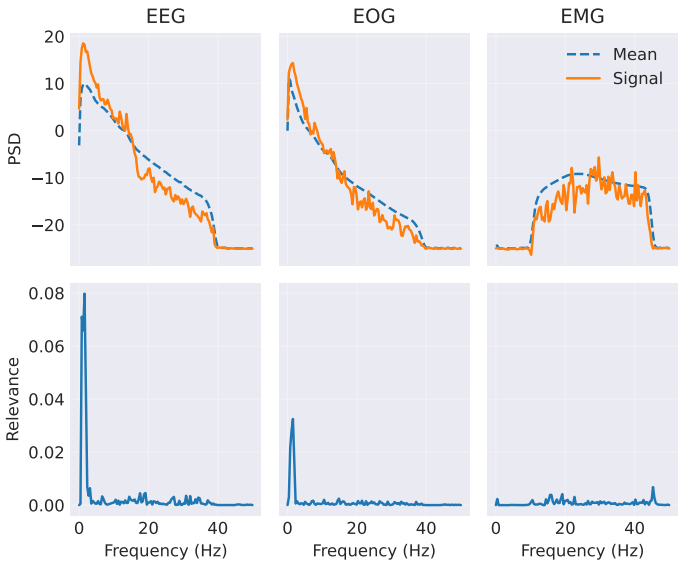

Fig. B1 Local N3-associated prototype explanation on MASS. Top row: power spectral density of the data-driven reconstruction compared to the dataset mean for EEG, EOG, and EMG. Bottom row: frequency-wise relevance obtained with Integrated Gradients for the same instance. Relevance concentrates at low frequencies in EEG/EOG and remains low and diffuse in EMG, consistent with PSG-based N3 staging dominated by slow-wave activity and reduced muscle tone.

Appendix C Global prototype-based explanation framework

To provide interpretable and representative examples for each learned prototype in the model's latent space, we employ a hybrid optimization approach. For each prototype vector $\mathbf{p}_k$, we first identify a real data instance $\mathbf{x}_k^*$ from the dataset that is both assigned to prototype k and exhibits maximal similarity (or minimal distance) to $\mathbf{p}_k$ in the embedding space. This selection ensures that the starting point for further optimization is a plausible, data-driven example.

Subsequently, we perform a gradient-based optimization in the input space, initializing from $\mathbf{x}_k^*$, to maximize the similarity between the projected input $\phi(\mathbf{x})$ and the prototype $\mathbf{p}_k$. Formally, we solve

$$\tilde{\mathbf{x}}_k = \arg\max_{\mathbf{x}} \ \mathrm{sim}(\phi(\mathbf{x}), \mathbf{p}_k) - \lambda \mathcal{R}(\mathbf{x}, \mathbf{x}_k^*)$$

where $\mathrm{sim}(\cdot, \cdot)$ denotes the cosine similarity, ϕ is the model's projection function, and $\mathcal{R}$ is a regularization term (e.g., an L_2 penalty) that encourages the optimized instance $\tilde{\mathbf{x}}_k$ to remain close to the original data instance $\mathbf{x}_k^*$. This hybrid approach yields prototype instances that are both highly representative of the prototype and remain within the data manifold, thus facilitating interpretation.

To quantify the contribution of each input feature or channel to the activation of a given prototype, we employ a relevance attribution mechanism based on feature perturbation. For each prototype instance $\tilde{\mathbf{x}}_k$, we systematically perturb subsets of input features (e.g., frequency bands, channels) by replacing them with their empirical mean values, and measure the resulting drop in similarity to the prototype:

$$\mathrm{rel}_{k,S} = \frac{\mathrm{sim}(\phi(\tilde{\mathbf{x}}_k), \mathbf{p}_k) - \mathrm{sim}(\phi(\tilde{\mathbf{x}}_k^{(S)}), \mathbf{p}_k)}{\mathrm{sim}(\phi(\tilde{\mathbf{x}}_k), \mathbf{p}_k) - \mathrm{sim}(\phi(\tilde{\mathbf{x}}_k^{(\mathrm{all})}), \mathbf{p}_k)}$$

where S denotes the subset of features being swapped, $\tilde{\mathbf{x}}_k^{(S)}$ is the instance with features S replaced by their mean, and $\tilde{\mathbf{x}}_k^{(\mathrm{all})}$ is the instance with all features replaced. This normalized relevance score reflects the relative importance of each feature or channel in supporting the prototype's activation, enabling a fine-grained interpretation of the model's internal representations.

The results of prototype instance generation and relevance attribution are visualized through a set of informative plots. For each prototype, we display: (i) the power spectral density (PSD) of the prototype instance compared to the dataset mean and the original data instance, (ii) bar plots summarizing the relevance scores for each channel or feature, and (iii) frequency-resolved profiles of feature relevance. These visualizations provide a comprehensive and interpretable summary of the prototypical patterns learned by the model, the features that drive their activation, and their correspondence to real data.

Figure C2 provides a comprehensive overview of the 15 prototypical reconstructions learned by the ProtoSleepNet model trained with $K = 15$ prototypes (PSN15). Each panel in the figure corresponds to a distinct prototype and displays the optimized instance generated through the hybrid reconstruction process described above.

For each prototype, the figure shows the power spectral density (PSD) profiles for the three input channels (EEG, EOG, EMG), averaged over time and across the optimized instance. The solid lines represent the PSD of the prototype instance, while the dashed lines indicate the mean PSD computed over the entire training dataset, serving as a baseline for comparison. This visualization allows for a direct assessment of the spectral characteristics that distinguish each prototype from the population mean.

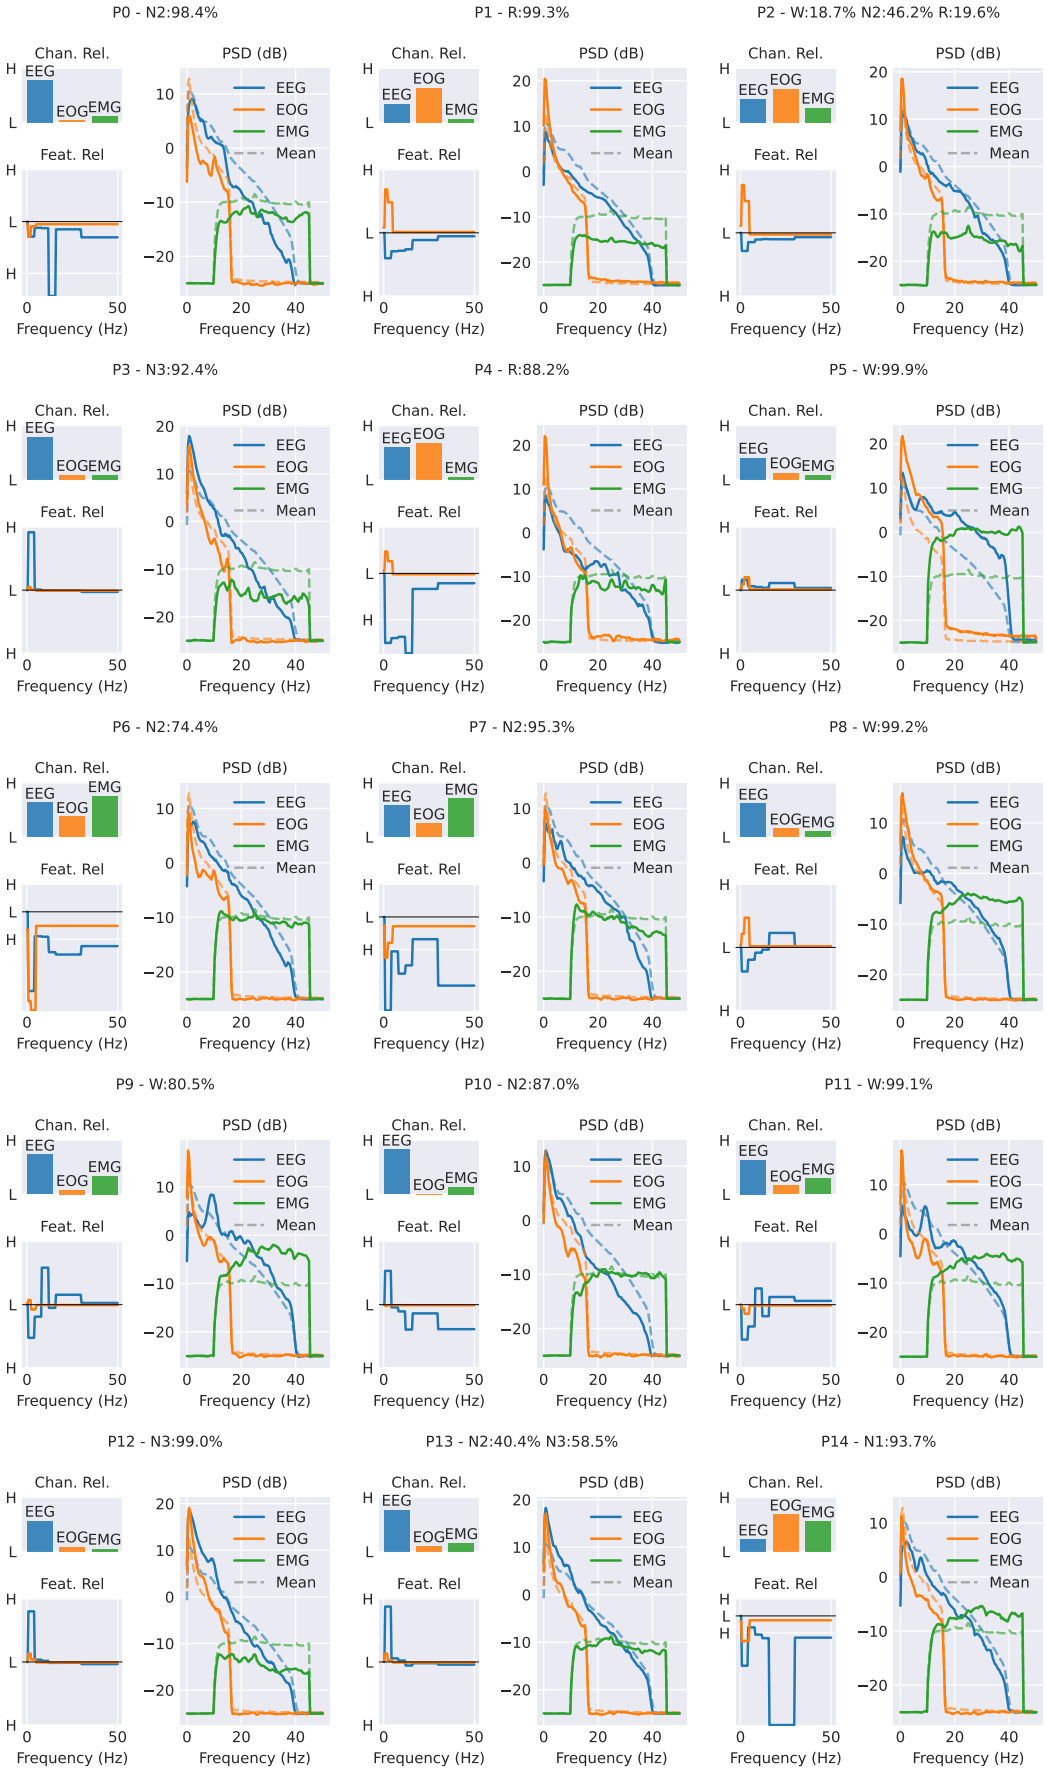

Fig. C2 Optimized prototype reconstructions learned by ProtoSleepNet with $K = 15$ prototypes (PSN15). Each panel displays the power spectral density (PSD) profiles of the optimized instance associated with a distinct prototype, for the EEG, EOG, and EMG channels. Solid lines represent the PSD of the prototype instance, while dashed lines indicate the mean PSD over the training dataset. The diversity of spectral patterns across prototypes reflects the model's ability to capture a wide range of physiologically meaningful features, supporting both interpretability and stage-specific specialization.

Appendix D Rule-Generation from Prototype instances

To further enhance interpretability, we derive explicit rules that characterize the prototypical patterns learned by the model. For each prototype, the rule-generation process synthesizes a concise, human-readable description based on the relevance scores attributed to channels and features.

The procedure begins by filtering channels whose relevance exceeds a predefined threshold, ensuring that only the most influential channels are considered. For each relevant channel, we identify the top features (e.g., frequency bands for EEG, categories for EOG, or tonic/atonic state for EMG) that contribute most strongly to the prototype's activation. Each feature is associated with its mean reconstruction value, its normalized relevance score, and the direction of its effect (increase or decrease in similarity), as determined by the perturbation analysis.

The resulting rule for prototype k takes the form:

Prototype k represents epochs showing [channel: feature, value, effect] ... and is associated with sleep stages [stage, probability].

where the channel-feature-value-effect tuples are selected based on their relevance and direction, and the stage probabilities are derived from the model's softmax predictions for the prototype.

Mathematically, the rule is constructed by ranking features according to their relevance scores $\text{rel}_{k,f}$, and including those for which $\text{rel}_{k,f} > \tau$, where τ is the threshold. The direction of effect is determined by the sign of the mean reconstruction value relative to the dataset mean, and is encoded using symbolic arrows (e.g., ↑ for increase, ↓ for decrease).

Tab. D5 shows the rules extracted from the 15 different prototypes learned in the PSN15 architectures and reconstructed using an hybrid strategy.

Table D5 Extracted rules from PSN15 architecture.

P.ID	Rule	Stages
0	represents epochs showing EEG: γ(30-50Hz) -22.27 dB ↓, σ(12-16Hz) -1.85 dB ↓, β(16-30Hz) -10.09 dB ↓; EMG: EMG atonic -12.65 dB ↓.	N2 98.4%
1	represents epochs showing EEG: β(16-30Hz) -7.30 dB ↓, δ(0.5-4Hz) +5.96 dB ↓, θ(4-8Hz) +1.24 dB ↓; EOG: REM(2-5Hz) +5.72 dB ↑, blink(0.5-2Hz) +15.52 dB ↑, (>5Hz) -19.08 dB ↑.	R 99.3%
2	represents epochs showing EEG: β(16-30Hz) -6.87 dB ↓, γ(30-50Hz) -20.99 dB ↓, δ(0.5-4Hz) +9.12 dB ↓; EOG: REM(2-5Hz) +5.51 dB ↑, blink(0.5-2Hz) +13.92 dB ↑, (>5Hz) -19.22 dB ↓; EMG: EMG atonic -15.33 dB ↓.	W 18.7%, N2 46.2%, R 19.6%
3	represents epochs showing EEG: δ(0.5-4Hz) +14.40 dB ↑, γ(30-50Hz) -22.10 dB ↓; EOG: (>5Hz) -20.20 dB ↓; EMG: EMG atonic -15.45 dB ↓.	N3 92.4%
4	represents epochs showing EEG: σ(12-16Hz) -7.10 dB ↓, θ(4-8Hz) -1.07 dB ↓, α(8-12Hz) -4.62 dB ↓.	R 88.2%
5	represents epochs showing EEG: β(16-30Hz) +2.23 dB ↑.	W 99.9%
6	represents epochs showing EEG: γ(30-50Hz) -21.22 dB ↓, β(16-30Hz) -7.64 dB ↓, δ(0.5-4Hz) +6.82 dB ↓; EOG: (>5Hz) -19.83 dB ↓, REM(2-5Hz) -0.76 dB ↓, blink(0.5-2Hz) +5.81 dB ↓; EMG: EMG atonic -10.70 dB ↓.	N2 74.4%
7	represents epochs showing EEG: γ(30-50Hz) -21.22 dB ↓, δ(0.5-4Hz) +5.77 dB ↓, β(16-30Hz) -6.94 dB ↓; EOG: (>5Hz) -19.94 dB ↓, REM(2-5Hz) +2.02 dB ↓, blink(0.5-2Hz) +7.99 dB ↓; EMG: EMG atonic -11.10 dB ↓.	N2 95.3%
8	represents epochs showing EEG: β(16-30Hz) -4.20 dB ↑, δ(0.5-4Hz) +4.11 dB ↓; EOG: REM(2-5Hz) +6.58 dB ↑; EMG: EMG tonic -5.84 dB ↑.	W 99.2%
9	represents epochs showing EEG: α(8-12Hz) +6.45 dB ↑, β(16-30Hz) -2.44 dB ↑, δ(0.5-4Hz) +3.91 dB ↓; EMG: EMG tonic -4.92 dB ↑.	W 80.5%
10	represents epochs showing EEG: γ(30-50Hz) -21.90 dB ↓, β(16-30Hz) -10.54 dB ↓, δ(0.5-4Hz) +10.09 dB ↑; EMG: EMG atonic -10.71 dB ↓.	N2 87.0%
11	represents epochs showing EEG: δ(0.5-4Hz) +2.79 dB ↓, β(16-30Hz) -3.86 dB ↑; EMG: EMG tonic -6.02 dB ↑.	W 99.1%
12	represents epochs showing EEG: δ(0.5-4Hz) +15.77 dB ↑, γ(30-50Hz) -21.26 dB ↓.	N3 99.0%
13	represents epochs showing EEG: δ(0.5-4Hz) +13.78 dB ↑, γ(30-50Hz) -20.34 dB ↓, β(16-30Hz) -6.97 dB ↓; EOG: (>5Hz) -19.51 dB ↓; EMG: EMG atonic -11.18 dB ↓.	N2 40.4%, N3 58.5%
14	represents epochs showing EEG: β(16-30Hz) -7.35 dB ↓, γ(30-50Hz) -20.70 dB ↓, δ(0.5-4Hz) +5.82 dB ↓; EOG: (>5Hz) -20.21 dB ↓, REM(2-5Hz) +0.18 dB ↓, blink(0.5-2Hz) +6.60 dB ↓; EMG: EMG tonic -7.66 dB ↑.	N1 93.7%

The 15 prototypes capture distinct physiological patterns across sleep stages with strong stage-specificity (prior to sequential context). Some prototypes rely on selective single-channel markers (e.g., P3, P12 dominated by delta), while others integrate multi-channel evidence to disambiguate overlapping stages. The learned rules align with established AASM staging criteria [**?**].

Wakefulness (P5, P8, P9, P11) prototypes are marked by beta activity (16–30 Hz), elevated EMG tone, and REMs P8, likely reflecting high-frequency blinks. P5 and P11 show slightly elevated beta power with, suppressed REMs, and high EMG, clearly distinguishing wakefulness from REM sleep. P8 and P9 use simpler signatures, predominant REMs

(P8) or isolated beta elevation with high EMG (P9), illustrating the flexibility of prototypes in capturing multiple valid physiological patterns for the same stage.

The N1 prototype (P14) is identified by suppressed beta, gamma, and delta bands, consistent with its transitional nature between wakefulness and deeper sleep.

N2 prototypes (P0, P6, P7, P10, P13) are characterized by elevated sigma (12–16 Hz) and delta power (0.5–4 Hz) with reduced gamma activity. P13 shows intermediate specificity (N2 40.4%, N3 58.5%), reflecting the continuous delta transition between N2 and N3. Blink features appear suppressed across N2 prototypes, consistent with reduced eyelid activity during deeper sleep.

N3 prototypes (P3, P12) are dominated by delta-band activity achieving high N3 confidence (92.4–99.0%) using elevated delta power as the sole discriminative marker, reflecting the overwhelming dominance of slow waves in deep sleep.

R prototypes (P1, P4) combine EOG REMs activity (2–5 Hz) with blink features (0.5–2 Hz). While these bands partially overlap, blinks are larger-amplitude, stereotyped vertical deflections, whereas REMs are rapid, lower-amplitude oscillations. PSN successfully distinguishes these features as complementary REM markers.

Appendix E Prototype-Based Disease Characterization

To investigate group-specific neurophysiological patterns, we employed a multi-stage analysis pipeline leveraging the ProtoSleepNet (PSN15) architecture with 15 prototypes. This approach enabled the extraction and characterization of prototypical sleep patterns across healthy controls, Parkinson's disease (PD), and Alzheimer's disease (AD) patients.

The analysis proceeded through the following stages:

1. **Prototypogram extraction:** Using the PSN15 architecture, we extracted the prototypogram—the sequential flow of prototypes throughout an entire night—for each subject across all study groups (healthy controls, PD, and AD patients). This representation captures the temporal dynamics of sleep architecture at the prototype level.
2. **Statistical prototype selection:** We performed Mann-Whitney U-tests to identify prototypes exhibiting significant differences in average time spent between groups. This analysis revealed prototypes P1, P2, P4, P5, P6, P7, P8, P9, P11, and P14 as displaying statistically significant group differences in their nocturnal prevalence.
3. **Subject-specific reconstruction:** For each selected prototype and each subject, we identified the specific PSG epoch that the PSN15 network projected closest to the prototype in the latent space. This reconstruction process yielded subject-specific instantiations of each prototypical pattern.
4. **Feature extraction with sequential context:** From each reconstructed epoch, we extracted neurophysiological features including power spectral density across canonical frequency bands (δ: 0.5-4Hz, θ: 4-8Hz, α: 8-12Hz, σ: 12-16Hz, β: 16-30Hz, γ: 30-50Hz), ocular artifacts (slow eye movements: 0-0.5Hz, blinks: 0.5-2Hz, REM: 2-5Hz), high-frequency activity (>5Hz), electromyographic activity (EMG), and sleep stage proportions. Critically, we computed the network's confidence for each prototype assignment after sequential context integration, thereby accounting for the prototype's specific temporal position within the night's architecture.
5. **Statistical testing:** Mann-Whitney U-tests were conducted to assess differences between healthy controls and each disease group (AD and PD) for each feature within each prototype. Multiple comparison correction was applied to control the family-wise error rate.

Tables E6- E7 present the extracted features organized by prototype pairs (P1-P2, P4-P5, P6-P7, P8-P9, P11-P14), with each row corresponding to a specific neurophysiological feature and columns representing the three study groups. All spectral power values are expressed in decibels (dB).

The prototype-specific patterns reveal differences between groups that are in line with known alterations of sleep in Alzheimer's disease (AD) and Parkinson's disease (PD). In the tables, each prototype highlights a particular combination of frequency bands, eye movements, muscle tone, and sleep stages, allowing a compact comparison of how these elements change across healthy controls, AD, and PD.

Slow-wave activity prototypes (P1, P2, P4). In several prototypes dominated by slow-wave (delta) activity, healthy participants show higher power than patients with AD, indicating that deep sleep is less prominent in AD in these patterns. This agrees with studies reporting that people with AD spend less time in slow-wave sleep than healthy older adults. PD patients also show changes in slow-wave activity and deep sleep in the literature, but in these prototypes they often fall between healthy controls and AD, suggesting a different alteration pattern rather than a simple loss of slow waves.

Sleep spindle-related prototypes (P4, P6, P7). Prototypes with strong activity in the sigma band (classically linked to sleep spindles) tend to show lower values in AD than in healthy controls in our tables. This is consistent with work showing that sleep spindles and related sigma power are reduced in AD and are associated with markers of neurodegeneration and cognitive decline. PD patients sometimes show intermediate values in these spindle-rich prototypes, which may indicate that thalamocortical sleep mechanisms are affected in both conditions, although probably through partly different pathways.

Alpha and theta modulation (P5, P8, P9, P11). Several prototypes emphasize alpha and theta activity, which reflect more relaxed or light-sleep brain rhythms. In these prototypes, AD patients tend to show lower or more altered alpha–theta power than healthy participants, in line with the broader picture of disturbed cortical rhythms in AD. PD patients generally appear closer to the healthy group in these patterns in our data, suggesting that, at least for these prototypes, large-scale cortical rhythmicity may be relatively less affected than in AD, even though both diseases are known to alter EEG activity.

REM-associated prototypes (P1, P4). Some prototypes correspond to segments with a high proportion of REM sleep. In these REM-rich prototypes, AD patients show reduced REM content compared with healthy controls, which is compatible with reports of reduced REM sleep and altered REM regulation in AD. PD is also associated with reduced REM sleep and fragmented REM periods, but in our prototypes PD often shows values between healthy controls and AD, pointing to a different balance between REM reduction and sleep fragmentation.

Beta and gamma activity (multiple prototypes). Across many prototypes, higher-frequency activity in the beta and gamma ranges appears lower in the patient groups than in healthy controls. Reviews of sleep oscillations in AD describe broad changes in these faster rhythms together with alterations in slower activity, suggesting a general disruption of coordinated network activity during sleep. PD is likewise associated with abnormal beta activity and changes in sleep-related oscillations, although most work has focused on waking motor symptoms and subcortical recordings rather than detailed scalp sleep spectra.

Wake-related prototypes (P5, P8, P9, P11). Some prototypes represent states with a high proportion of wakefulness during the night. Both AD and PD are known to show more time awake after sleep onset and more fragmented sleep compared with healthy older adults. In our tables, these wake-related prototypes often show that patient groups either spend more time in such patterns or have more irregular values, consistent with the idea that neurodegeneration destabilizes sleep–wake regulation and makes sleep more fragmented.

Ocular and muscular activity. The tables also summarize eye-movement-related activity (such as slow eye movements, blinks, and REM-related bands) and muscle tone (EMG) for each prototype. Clinically, PD is characterized by abnormal motor control during sleep (for example, REM sleep behavior disorder), and AD is associated with changes in eye movements and overall sleep stability. In our prototypes, lower EMG levels and altered blink- and REM-related activity in the patient groups complement the spectral and sleep-stage findings, suggesting that changes in motor and ocular control during sleep accompany the more global alterations of sleep architecture.

Taken together, these prototype-based patterns show that the PSN15-derived representations capture sleep changes that are coherent with established clinical knowledge: AD is strongly linked to losses in deep sleep, spindles, and REM organization, while PD shows a combination of reduced restorative sleep and pronounced fragmentation, with partially overlapping but not identical alterations in sleep-related brain rhythms.

Table E6 Features extracted from per-subject prototypical instances of PSN15 across patient groups. Values represent power spectral density (dB) for frequency bands (δ: 0.5-4Hz, θ: 4-8Hz, α: 8-12Hz, σ: 12-16Hz, β: 16-30Hz, γ: 30-50Hz), ocular artifacts (SEM: 0-0.5Hz, blink: 0.5-2Hz, REM: 2-5Hz), high-frequency activity (>5Hz), electromyographic activity (EMG), and sleep stage confidence posterior to the sequential context (W: wake, N1-N3: non-REM stages, R: rapid eye movement). Each table presents paired prototypes (P1-P2, P4-P5, P6-P7, P8-P9, P11-P14) with measurements from Healthy controls, Alzheimer's disease, and Parkinson's disease patients. Asterisks denote statistical significance levels from Mann-Whitney U-test with correction for multiple comparisons: $*p < 0.05, **p < 0.01, ***p < 0.001$.

PID	Feat	Healthy	Alzheimers	Parkinsons	PID	Healthy	Alzheimers	Parkinsons
P1	α(8-12Hz)	-0.22	-5.70***	0.0917	P2	-0.09	-4.79***	0.0355
	β(16-30Hz)	-6.91	-11.24***	-7.7397		-7.90	-11.39***	-8.0035
	δ(0.5-4Hz)	6.05	1.39***	6.0875		7.48	3.22***	6.5106
	γ(30-50Hz)	-12.77	-15.43**	-15.6013		-14.39	-15.9989	-14.9141
	σ(12-16Hz)	-3.24	-7.97***	-3.1983		-3.32	-7.50***	-3.2021
	θ(4-8Hz)	2.75	-2.66***	2.9414		3.57	-1.02***	3.5681
	(>5Hz)	-6.24	-10.44***	-10.3484		-7.06	-11.41***	-10.11***
	EMG	-15.00	-14.4920	-12.9057		-16.71	-16.3677	-14.3495
	N1	0.03	0.04***	0.0363		0.05	0.0548	0.0442
	N2	0.05	0.07***	0.0815		0.24	0.37**	0.2394
	N3	0.00	0.02***	0.0177		0.02	0.09***	0.0346
	REM	0.72	0.7491	0.7677		0.59	0.44**	0.5325
	REM(2-5Hz)	5.48	5.2119	3.4984		5.86	4.3289	2.85***
	SEM(0-0.5Hz)	10.74	11.8904	9.2484		10.50	7.37*	6.50***
	W	0.07	0.09***	0.0317		0.02	0.04***	0.0188
	blink(0.5-2Hz)	10.99	13.69**	10.5569		11.38	11.1220	8.30***

PID	Feat	Healthy	Alzheimers	Parkinsons	PID	Healthy	Alzheimers	Parkinsons
P4	α(8-12Hz)	-0.96	-5.79***	-0.6237	P5	6.20	-1.36***	3.2754
	β(16-30Hz)	-8.18	-10.68***	-8.0609		3.21	-4.54***	0.5968
	δ(0.5-4Hz)	6.46	1.04***	5.10*		12.71	4.75***	5.3204
	γ(30-50Hz)	-14.89	-14.7328	-16.0454		0.20	-6.75***	-4.8699
	σ(12-16Hz)	-4.47	-8.11***	-4.0506		4.72	-3.04***	1.9172
	θ(4-8Hz)	2.27	-2.79***	1.8168		7.72	1.15***	2.5468
	(>5Hz)	-7.54	-9.69**	-10.31***		3.52	-2.23***	-6.0689
	EMG	-17.41	-13.02**	-14.8565		-1.32	1.9272	-1.9375
	N1	0.06	0.0710	0.0511		0.02	0.04***	0.0266
	N2	0.08	0.1139	0.0904		0.03	0.12***	0.0640
	N3	0.00	0.02***	0.0029		0.01	0.02***	0.0130
	REM	0.82	0.64***	0.7588		0.00	0.06***	0.0100
	REM(2-5Hz)	6.78	5.9358	4.79**		13.67	12.1328	9.6624
	SEM(0-0.5Hz)	14.15	12.1136	11.7320		17.86	15.7851	12.9757
	W	0.03	0.09***	0.0315		0.85	0.75***	0.8864
	blink(0.5-2Hz)	13.86	14.0197	12.5692		19.64	19.9468	15.6120

PID	Feat	Healthy	Alzheimers	Parkinsons	PID	Healthy	Alzheimers	Parkinsons
P6	α(8-12Hz)	0.99	-4.30***	1.7330	P7	1.16	-3.99***	1.6585
	β(16-30Hz)	-5.72	-10.61***	-8.31**		-8.94	-11.13***	-8.5514
	δ(0.5-4Hz)	5.75	2.98***	8.04**		9.06	3.60***	8.6226
	γ(30-50Hz)	-10.36	-14.93***	-15.35**		-15.31	-15.3971	-15.7879
	σ(12-16Hz)	-0.95	-6.78***	-2.03**		-1.64	-6.51***	-2.2959
	θ(4-8Hz)	2.96	-0.59***	5.06***		4.31	-0.28***	5.0391
	(>5Hz)	-5.71	-10.05***	-11.16***		-7.81	-10.62**	-11.4752
	EMG	-8.68	-9.6430	-12.75**		-13.58	-10.3524	-13.0206
	N1	0.06	0.15***	0.08***		0.04	0.06***	0.0413
	N2	0.60	0.6082	0.8438		0.91	0.77***	0.8676
	N3	0.01	0.02***	0.02***		0.02	0.05***	0.0267
	REM	0.03	0.04***	0.02***		0.03	0.0393	0.0350
	REM(2-5Hz)	3.47	2.9665	1.8995		5.87	3.23***	2.3528
	SEM(0-0.5Hz)	3.77	3.8976	1.19*		6.79	3.51*	1.0187
	W	0.02	0.13***	0.02***		0.01	0.05***	0.0076
	blink(0.5-2Hz)	6.56	7.9953	6.0043		10.75	8.30*	6.3524

Table E7 Continues from E6

PID	Feat	Healthy	Alzheimers	Parkinsons	PID	Healthy	Alzheimers	Parkinsons
P8	α(8-12Hz)	2.70	-4.03***	3.2230	P9	4.06	-3.03***	4.3024
	β(16-30Hz)	-2.39	-8.31***	-2.6792		-0.85	-7.51***	-2.7195
	δ(0.5-4Hz)	9.24	-0.22***	5.4433		7.83	0.61***	4.90*
	γ(30-50Hz)	-7.16	-12.02***	-9.1602		-5.22	-11.91***	-9.49**
	σ(12-16Hz)	0.21	-5.97***	0.1975		1.51	-4.86***	0.4277
	θ(4-8Hz)	4.62	-2.05***	3.1650		4.60	-1.65***	4.2259
	(>5Hz)	-1.17	-5.49***	-6.3002		-0.33	-5.53***	-7.29***
	EMG	-6.70	-4.8385	-5.3919		-4.62	-0.49*	-6.0917
	N1	0.02	0.02**	0.0233		0.08	0.09*	0.12*
	N2	0.01	0.04***	0.0248		0.14	0.18**	0.23*
	N3	0.00	0.02***	0.0038		0.02	0.03***	0.0157
	REM	0.01	0.04***	0.0300		0.01	0.03***	0.02***
	REM(2-5Hz)	10.52	8.6064	10.2574		7.13	6.1057	5.3077
	SEM(0-0.5Hz)	14.12	12.5353	13.8729		12.73	10.5664	9.05*
	W	0.92	0.88***	0.9182		0.68	0.6706	0.6189
	blink(0.5-2Hz)	16.12	16.0759	16.5890		13.23	13.8240	11.6758

PID	Feat	Healthy	Alzheimers	Parkinsons	PID	Healthy	Alzheimers	Parkinsons
P11	α(8-12Hz)	4.24	-2.78***	4.6788	P14	0.49	-4.76***	1.3732
	β(16-30Hz)	-1.33	-7.47***	-1.6659		-7.20	-10.54***	-7.6682
	δ(0.5-4Hz)	7.12	0.71***	4.5717		6.83	2.60***	6.4243
	γ(30-50Hz)	-6.13	-11.66***	-8.4589		-13.74	-15.07**	-15.6571
	σ(12-16Hz)	0.78	-5.80***	0.7442		-2.38	-7.02***	-2.3619
	θ(4-8Hz)	3.70	-1.52***	3.8155		3.42	-0.97***	3.9982
	(>5Hz)	-0.77	-5.51***	-7.2094		-6.82	-9.83***	-10.86***
	EMG	-6.04	-0.75**	-4.3865		-11.50	-8.6691	-10.9797
	N1	0.07	0.0922	0.1036		0.24	0.16***	0.2139
	N2	0.06	0.10***	0.0768		0.40	0.4062	0.46**
	N3	0.00	0.01**	0.0025		0.01	0.04***	0.0186
	REM	0.06	0.0831	0.0432		0.14	0.05***	0.09**
	REM(2-5Hz)	6.73	5.9334	3.5039		4.47	3.1508	0.90***
	SEM(0-0.5Hz)	15.03	10.71*	9.2309		9.58	6.17*	4.24***
	W	0.80	0.71**	0.7522		0.21	0.34***	0.2134
	blink(0.5-2Hz)	14.25	13.7645	10.3173		10.15	9.3544	6.43***