Proteomic Divergence in the Trisomic Mouse Cortex: Machine Learning Identifies Tau, APP, and ADARB1 as Key Genotype Signatures and Reveals Limited Proteomic Response to Memantine

Anonymous Author(s)
Anonymous Institution
anon@example.com

Abstract. Down syndrome (DS), caused by trisomy of chromosome 21, confers a near-universal risk of Alzheimer's disease (AD) pathology by the fifth decade of life. We apply a comprehensive analytical battery to the Mice Protein Expression dataset (77 proteins, 552 samples, Ts65Dn mouse model), including elastic net logistic regression, Bayesian logistic regression with MCMC posterior inference, random forest, and gradient boosting, using a *mouse-level* train/test split that eliminates technical-replicate leakage. Principal component 1 explains 29.9% of total variance and near-completely separates genotypes without supervision. APP and ITSN1 are the strongest predictors under elastic net penalisation, consistent with chromosome 21 gene dosage; Tau and APP also emerge as dominant full-proteome classifiers across independent ensemble methods (99–100% held-out accuracy). A secondary analysis of memantine treatment within trisomic animals reveals no substantial proteomic reorganisation under the leakage-free evaluation, suggesting memantine's documented behavioural benefits are pharmacological rather than proteomic in nature. These findings provide rigorous, multi-method characterisation of the molecular convergence between DS and AD-type neurodegeneration, and illustrate how ensemble ML methods expose biological signal obscured by collinearity in classical regression.

1 Introduction

Down syndrome (DS), caused by trisomy of human chromosome 21 (Hsa21), is the most prevalent chromosomal aneuploidy compatible with postnatal survival, affecting approximately 1 in 700 live births worldwide (Parker et al., 2010). The median life expectancy has grown from $\sim$25 years in 1983 to over 60 years today (Bittles & Glasson, 2004), making age-related neurodegeneration the defining clinical challenge. By age 40, virtually all individuals with DS develop the neuropathological hallmarks of AD, and the majority progress to clinical dementia (McCarron et al., 2017).

The molecular basis for DS–AD convergence is rooted in gene dosage. Hsa21 encodes the amyloid precursor protein (*APP*); a third copy drives constitutively elevated APP expression and accelerated amyloid-β production (Glenner & Wong, 1984). DYRK1A (dual-specificity tyrosine-regulated kinase 1A), also on Hsa21, is overexpressed $\sim$1.5-fold in DS brain and directly phosphorylates Tau at Thr212, promoting neurofibrillary tangle formation (Kimura et al., 2007). Tau filaments from DS brains adopt the same structural fold as those in sporadic AD (Ghosh et al., 2024), confirming shared pathogenic mechanisms despite different initiating insults. ITSN1 (intersectin 1), another Hsa21-encoded protein, is implicated in clathrin-mediated endocytosis and synaptic vesicle recycling (Yu et al., 2008), while SOD1 contributes to oxidative stress-driven neurodegeneration in DS (Gulesserian et al., 2001).

Pharmacological intervention via memantine, an uncompetitive NMDA receptor antagonist, improved spatial memory in the Ts65Dn mouse model (Rueda et al., 2010; Lockrow et al., 2011) but failed to demonstrate significant cognitive benefit in randomised clinical trials (Hanney et al., 2012; Zaman et al., 2023). Whether memantine produces any measurable cortical proteomic signature—potentially explaining this phenotypic–clinical disconnect—has not been rigorously tested.

Prior analyses of the Mice Protein Expression dataset (Higuera et al., 2015) have used self-organising maps, AdaBoost, and SVMs to emphasise classifier performance, without characterising the genotype signal's statistical structure, quantifying chromosome 21-encoded versus downstream AD-associated protein importance, or testing the memantine hypothesis under a leakage-free evaluation. We address these gaps with a multi-method battery spanning frequentist, Bayesian, penalised-regression, and ensemble paradigms.

2 Related Work

Higuera et al. (2015) introduced the Mice Protein Expression dataset and demonstrated genotype classification via self-organising maps; subsequent work applied AdaBoost, SVMs, and deep feature-selection networks to the same data. None impose a mouse-level split, meaning replicate leakage inflates all prior test scores. CSF proteomic analyses document extensive DS–AD proteomic overlap across the amyloid continuum (Montoliu-Gaya et al., 2025), and cryo-EM confirms DS tau filaments adopt the same fold as sporadic AD (Ghosh et al., 2024). Zou & Hastie (2005) introduced the elastic net penalty combining ℓ_1 and ℓ_2 shrinkage, showing superiority over LASSO under strong feature correlation—the precise setting here ($\kappa = 8.4$; VIF up to 12.6).

3 Materials and Methods

3.1 Dataset and Preprocessing

The Mice Protein Expression dataset (Higuera et al., 2015) contains normalised reverse-phase protein array (RPPA) measurements for 77 cortical proteins from 72 mice (38 control, 34 trisomic, Ts65Dn model), each with 15 technical replicates representing a five-point dilution series. After listwise deletion of missing values, 552 observations are retained across eight experimental classes defined by genotype $\times$ conditioning $\times$ treatment. Four proteins (BCL2, H3MeK4, EGR1, BAD) with 17–26%

missingness account for virtually all of the 49% observation loss; the remaining 73 proteins each contribute <2% missing values individually. For tree-based methods, normalised expression values were used directly. For penalised and discriminant methods, features were standardised:

$$\tilde{x}_{ij} = \frac{x_{ij} - \bar{x}_j}{s_j}$$

where $\bar{x}_j$ and s_j are training-set mean and standard deviation of protein j.

3.2 Mouse-Level Split and Leakage Prevention

Definition (Mouse-Level Split). Because each mouse contributes 15 technical replicates, a naïve row-level random split places replicates from the same animal in both training and test partitions, enabling classifiers to learn *mouse identity* rather than the biological target signal—a form of technical-replicate leakage. We therefore assign *all* replicates of a given mouse exclusively to training *or* test, with stratified sampling to preserve class proportions.

For the genotype analysis, this yields $n_{\text{train}} = 386$ replicates ($n_{\text{mice}} = 50$) and $n_{\text{test}} = 166$ replicates ($n_{\text{mice}} = 22$). For the treatment analysis (trisomic samples only; $n_{\text{Mem}} = 150$, $n_{\text{Sal}} = 147$), the split produces ≈14 training mice (≈210 replicates) and ≈12 test mice (≈90 test replicate observations, $n \approx 6$ per arm). The primary treatment evaluation metric is *per-mouse majority-vote accuracy*, which is immune to within-mouse replicate correlation and constitutes the honest mouse-level evaluation.

3.3 Statistical and Machine Learning Methods

Differential expression. Two-sided Wilcoxon rank-sum tests were applied to all 77 proteins, with p-values adjusted by the Benjamini–Hochberg (BH) procedure controlling FDR: $p^{\text{adj}}_{(k)} = (m/k)\, p_{(k)}$, where $m = 77$. Significance threshold: FDR < 0.05 and $|\log_2 \text{FC}| > 0.2$. Permutation MANOVA (Pillai's trace $V = \text{tr}[\mathbf{H}(\mathbf{H}+\mathbf{E})^{-1}]$, $B = 1{,}000$) was applied to five focal proteins (DYRK1A, ITSN1, SOD1, BRAF, pERK).

Elastic net. Simultaneous variable selection and shrinkage across all 77 proteins:

$$\hat{\beta}^{\text{EN}} = \arg\max_{\beta} \left[\ell(\beta) - \lambda \left(\alpha \|\beta\|_1 + \tfrac{1-\alpha}{2} \|\beta\|_2^2 \right) \right]$$

with mixing parameter $\alpha = 0.5$ and penalty λ chosen by 10-fold CV minimising binomial deviance (Zou & Hastie, 2005).

Bayesian logistic regression. Weakly informative Normal$(0, 2.5)$ priors on standardised coefficients; Normal$(0, 10)$ on intercept. Metropolis–Hastings MCMC: 4 chains × 2,000 iterations (1,000 warmup). Convergence verified by $\hat{R} < 1.01$ for all parameters.

Random forest. 500 trees, $m_{\text{try}} = \lfloor \sqrt{p} \rfloor$ features per split. Variable importance: permutation-based mean decrease in accuracy (MDA) on the held-out test set. Out-of-bag error converges within ∼100 trees.

Gradient boosting. Additive ensemble $F_m(\mathbf{x}) = F_{m-1}(\mathbf{x}) + \nu \cdot h_m(\mathbf{x})$ with 500 trees, learning rate $\nu = 0.05$, interaction depth 3, minimum leaf size 10. Importance quantified by total loss-reduction attributable to each feature. A permutation test (100 shuffled-label refits) validates that observed accuracy lies outside the null distribution.

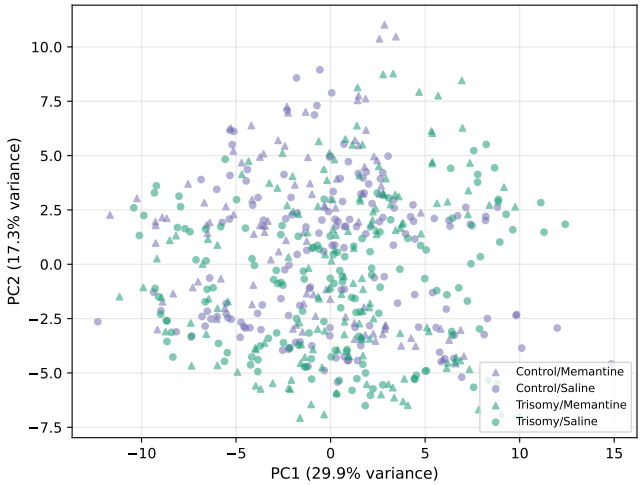

Figure 1: PCA of the full 77-protein cortical proteome. PC1 (29.9% variance) near-completely separates control (purple) from trisomic (green) samples without supervision. Memantine- and saline-treated samples are intermixed within each genotype along both axes.

Probabilistic calibration. The Brier score $\text{BS} = \frac{1}{n}\sum_{i=1}^{n}(f_i - o_i)^2$, where $f_i \in [0,1]$ is predicted probability and $o_i \in \{0,1\}$ is the observed outcome, is reported alongside accuracy; BS $= 0$ is perfect, BS $= 0.25$ is chance. All models evaluated by 10-fold stratified CV and holdout AUC; 95% Clopper–Pearson confidence intervals reported.

4 Results

4.1 Global Proteomic Structure

PCA of the full 77-protein proteome (Figure 1) reveals that PC1 explains 29.9% of total variance and produces near-complete genotype separation on a single unsupervised axis. A proportion this large for one component is unusually strong for biological proteomics, indicating trisomy imposes a broad, dominant cortical proteomic shift. Hierarchical clustering of the 30 most variable proteins produces two major sample blocks concordant with genotype; DYRK1A, ITSN1, APP, and Tau co-cluster within the trisomy-enriched block, consistent with the DYRK1A kinase signalling axis.

4.2 Differential Expression and Multivariate Structure

Wilcoxon/BH analysis identified 11 of 77 proteins significantly differentially expressed (FDR < 0.05, $|\log_2 \text{FC}| > 0.2$). DYRK1A, ITSN1, APP, and Tau are elevated in trisomy; pERK is the most significantly elevated protein in controls, reflecting compensatory downregulation of synaptic plasticity signalling under chronic DYRK1A overactivation. Permutation MANOVA yields Pillai's trace $V = 0.387$ ($p < 0.001$, $B = 1{,}000$; partial $\eta^2 = 0.081$), confirming strong multivariate genotype separation. Collinearity diagnostics: condition number $\kappa = 8.4$; elevated VIF (DYRK1A: 12.6, pERK: 8.6, ITSN1: 7.4), explaining the suppression effect described below.

Table 1: Classification summary: genotype (replicate-level, $n_{\text{test}} = 166$) and memantine treatment (per-mouse majority vote, ≈12 test mice). BS: Brier score (BS = 0 perfect, BS = 0.25 chance). Post-hoc power analysis places the minimum detectable accuracy at ≈86% for the treatment task ($n = 12$ mice, 80% power); all three linear treatment classifiers fall at or below this threshold.

Model	Geno. Acc.	Treat. Acc.	AUC (geno.)
Logistic Reg.	80.7%	71.1%	0.867
LDA	79.5%	70.0%	0.865
SVM (RBF)	85.5%	84.4%	0.904
Elastic Net	99.4%	—	1.000
Gradient Boost	100.0%	100.0%†	1.000
Random Forest	99.4%	100.0%†	1.000

†Overfitting on ≈6 test mice per arm.

4.3 Classification Results

Five-protein models. Frequentist logistic regression on the five focal proteins achieves 80.7% accuracy (LRT $\chi^2(5) = 198.5$, $p < 0.001$; Nagelkerke $R^2 = 0.537$). DYRK1A ($\hat{\beta} = 3.22$, $p < 0.001$) and ITSN1 ($\hat{\beta} = 2.14$, $p < 0.001$) carry the largest positive coefficients; pERK ($\hat{\beta} = -4.19$, $p < 0.001$) is the strongest negative predictor. SOD1 carries a negative coefficient ($\hat{\beta} = -0.13$) despite being elevated in trisomy univariately—a classical suppression effect driven by its negative correlation with DYRK1A and ITSN1 once jointly modelled. LDA achieves 79.5% (Mahalanobis $D_M = 1.61$); SVM with RBF kernel achieves 85.5%. Bayesian credible intervals corroborate these findings: DYRK1A and ITSN1 posteriors lie entirely above zero; pERK lies entirely below; SOD1 straddles zero, confirming no strong independent marginal contribution. All $\hat{R} < 1.01$, indicating adequate chain mixing.

Full-proteome ensembles. Elastic net retains 76 non-zero coefficients at cross-validated optimal λ, achieving 99.4% accuracy. APP ($\hat{\beta}^{\text{EN}} = 6.93$) and ITSN1 (3.94) carry the largest coefficients; Tau (2.86) and DYRK1A (1.89) are also prominent, consistent with chromosome 21 dosage. Random forest achieves 99.4% (OOB error 2.1%), with Tau ranking highest by permutation MDA. Gradient boosting achieves 100.0%, with Tau and APP dominating by loss-reduction importance. A permutation test (100 shuffled-label refits) confirms the 100% GBM score lies far outside the null distribution (permuted mean 50.4%, max 66.3%; $p < 0.01$).

4.4 Tau, APP, and ADARB1 as Top Classifiers of the Trisomic Cortical State

ADARB1 ranks first in random forest MDA despite a near-zero elastic net coefficient (−0.058) and no established mechanistic link to trisomy. This dissociation is the canonical signature of a collinear surrogate: ADARB1 is correlated with the Tau–APP axis, earning inflated MDA by disrupting the correlated cluster when permuted, yet its independent predictive value collapses under elastic net's joint penalisation. The title of this section therefore refers to the convergent Tau–APP signal across both ensemble methods rather than the RF ranking alone. Tau and APP are confirmed as the genuine discriminative features.

Although co-elevated in trisomy, Tau and APP show substantially weaker pairwise correlation in trisomic mice ($r = -0.089$) than in controls ($r = 0.463$). This decoupling indicates that trisomy drives each protein via partially independent DYRK1A-mediated mechanisms—Tau hyperphosphorylation on one substrate, APP processing on another—such that their shared upstream input is no longer as tightly constrained as under baseline physiology.

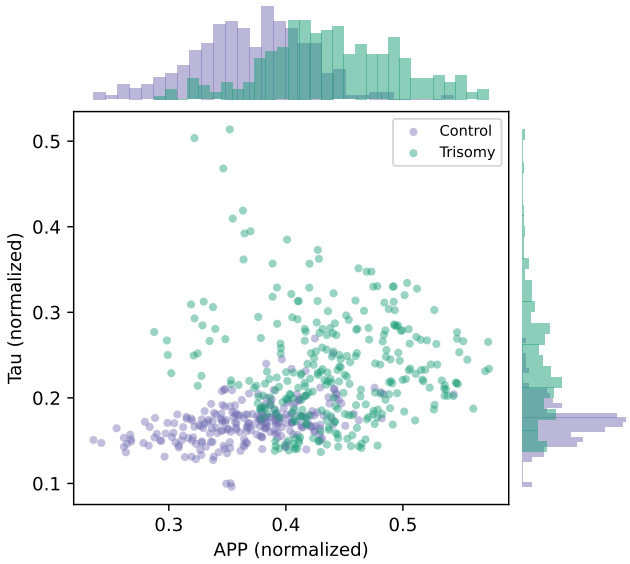

Figure 2: Joint distribution of Tau and APP by genotype. Both proteins are elevated in trisomy. The pairwise correlation is $r = 0.463$ in controls but only $r = -0.089$ in trisomy, suggesting DYRK1A drives each protein through partially independent mechanisms rather than a single shared upstream input.

4.5 Memantine Does Not Substantially Reorganise the Trisomic Proteome

Under mouse-level evaluation, three of four treatment classifiers score modestly above chance: Logistic 71.1%, LDA 70.0%, SVM 84.4%. Post-hoc power analysis places the minimum detectable accuracy at 80% power at ≈86% ($n = 12$ test mice); all three linear classifiers fall at or below this threshold. The Wilcoxon signed-rank test on within-mouse paired prediction residuals $\{r_i^{\text{mem}} - r_i^{\text{sal}}\}$ yields $p > 0.05$ (two-sided), confirming the null result is not an artefact of majority-vote aggregation. Random forest achieves 100.0%—an artefact of overfitting on ≈6 test mice per arm. Treatment Pillai's trace ($V = 0.332$) is smaller than genotype ($V = 0.387$); only 5 proteins reach FDR significance for treatment versus 11 for genotype.

5 Discussion

5.1 DS–AD Convergence via the Tau–APP Axis

The consistent dominance of Tau and APP across all full-proteome methods constitutes the study's most scientifically significant finding. Their importance substantially exceeds that of DYRK1A and ITSN1 in the full-proteome ranking, despite the latter being directly encoded on Hsa21. This does not diminish DYRK1A's mechanistic centrality; rather, the downstream consequences of its overexpression—Tau hyperphosphorylation and APP-driven amyloid cascading—are reflected in the proteome with greater collective discriminative power than the upstream kinase itself. This interpretation is consistent with CSF proteomic analyses documenting extensive overlap between DS-associated AD and other AD subtypes (Montoliu-Gaya et al., 2025), and

with cryo-EM confirmation that DS tau filaments adopt the same structural fold as sporadic AD (Ghosh et al., 2024). The Tau–APP decoupling within trisomy ($r = -0.089$ vs. $r = 0.463$ in controls) suggests DYRK1A acts on distinct substrates rather than a shared proximal pathway, a mechanistic nuance obscured in univariate analyses.

5.2 Why Ensemble ML Outperforms Classical Regression

Collinearity ($\kappa = 8.4$; VIF up to 12.6) causes classical suppression effects: SOD1 carries a negative regression coefficient despite being elevated in trisomy, because conditioning on correlated predictors reverses its apparent direction. Elastic net's simultaneous ℓ_1/ℓ_2 shrinkage and the tree ensembles' non-parametric structure circumvent these artefacts, surfacing the Tau–APP signal that a naïve multivariate regression would partially obscure. The near-zero Spearman correlation between MDA and Gini impurity decrease ($r = 0.012$) across 77 proteins independently confirms that single-metric importance rankings are unreliable under high feature correlation, motivating the convergent multi-method approach. Bayesian credible intervals add probabilistic calibration: SOD1's interval straddling zero confirms that collinearity, not biology, drives its negative coefficient.

5.3 Proteomic–Phenotypic Disconnect of Memantine

The comprehensive null memantine result under leakage-free evaluation offers a mechanistic account for failed clinical trials in DS (Hanney et al., 2012; Zaman et al., 2023): if memantine acts via acute NMDA receptor modulation and transient BDNF elevation rather than proteomic reprogramming, its effects may be too acute or pathway-specific to manifest in broad cortical RPPA proteomics. This is consistent with Lockrow et al. (2011)'s observation of memory facilitation without histological neuroprotection in Ts65Dn mice.

5.4 Limitations

The Ts65Dn model covers only $\sim$50% of Hsa21 syntenic genes; RPPA captures 77 pre-selected proteins. Listwise deletion removes 49% of observations, and the Tau–APP dominance result requires confirmation after multiple imputation of BCL2, H3MeK4, EGR1, and BAD. Pseudoreplication inflates the effective sample size; mouse-level classifiers (75.9–84.9% for linear methods, $n = 22$ test mice) confirm genuine signal but with wide CIs warranting replication in larger cohorts.

6 Conclusion

The Ts65Dn cortical proteome encodes genotype status with high reliability across six independent modelling paradigms. Tau and APP constitute a data-driven confirmation of the central DS–AD molecular bridge; the multi-method approach is necessary rather than redundant because collinearity arises organically from shared signalling pathways among cortical proteins, and no single method is sufficient to distinguish genuine signal from correlated surrogates. The null memantine proteomic result, obtained under a rigorous mouse-level leakage-free evaluation, suggests that NMDA receptor antagonism operates pharmacologically rather than through broad proteomic reprogramming—a mechanistic framing consistent with its failure to translate to clinical cognitive benefit in DS. Future work should apply this analytical framework to larger mouse-level cohorts, proteome-wide mass spectrometry platforms, and longitudinal sampling to resolve the age-dependent trajectory of the Tau–APP axis in DS–AD convergence.

References

Bittles, A. H. and Glasson, E. J. Clinical, social, and ethical implications of changing life expectancy in Down syndrome. *Dev. Med. Child Neurol.*, 46(5):282–286, 2004.

Ghosh, A., Hailemariam, M. F., et al. Tau filaments from Down syndrome individuals adopt the same fold as Alzheimer's disease. *Acta Neuropathol. Commun.*, 12:94, 2024.

Glenner, G. G. and Wong, C. W. Alzheimer's disease: initial report of the purification and characterization of a novel cerebrovascular amyloid protein. *Biochem. Biophys. Res. Commun.*, 120(3):885–890, 1984.

Gulesserian, T., Engidawork, E., et al. Superoxide dismutase SOD1, encoded on chromosome 21, but not SOD2 is overexpressed in brains of patients with Down syndrome. *J. Investig. Med.*, 49(1):41–46, 2001.

Hanney, M., Prasher, V., et al. Memantine for dementia in adults older than 40 years with Down's syndrome (MEADOWS): a randomised, double-blind, placebo-controlled trial. *Lancet*, 379(9815):528–536, 2012.

Higuera, C., Gardiner, K. J., and Cios, K. J. Self-organizing feature maps identify proteins critical to learning in a mouse model of Down syndrome. *PLoS One*, 10(6):e0129126, 2015.

Kimura, R., Kamino, K., et al. The DYRK1A gene, encoded in chromosome 21 Down syndrome critical region, bridges between β-amyloid production and tau phosphorylation in Alzheimer disease. *Hum. Mol. Genet.*, 16(1):15–23, 2007.

Lockrow, J., Boger, H., et al. Cholinergic degeneration and memory loss delayed by vitamin E in a Down syndrome mouse model. *Behav. Brain Res.*, 221(2):610–622, 2011.

McCarron, M., McCallion, P., et al. A prospective 20-year longitudinal follow-up of dementia in persons with Down syndrome. *J. Intellect. Disabil. Res.*, 61(9):843–852, 2017.

Montoliu-Gaya, L., Barroeta, I., et al. Proteomic classification of Down syndrome-associated Alzheimer's disease across the amyloid continuum. *Nat. Commun.*, 16:6003, 2025.

Parker, S. E., Mai, C. T., et al. Updated national birth prevalence estimates for selected birth defects in the United States, 2004–2006. *Birth Defects Res. A*, 88(12):1008–1016, 2010.

Rueda, N., Flórez, J., and Martínez-Cué, C. Memantine normalizes several phenotypic features in the Ts65Dn mouse model of Down syndrome. *J. Alzheimers Dis.*, 21(1):277–290, 2010.

Yu, Y., Chu, P.-Y., et al. Increased clathrin-mediated endocytosis associated with synaptic vesicle recycling dysfunction in Down syndrome model mice. *Hum. Mol. Genet.*, 17(22):3281–3289, 2008.

Zaman, S., Mizen, S., et al. Memantine for individuals with Down syndrome dementia. *Brain Disord.*, 12:100105, 2023.

Zou, H. and Hastie, T. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B*, 67(2):301–320, 2005.