

Phylogenomics to structure: evolutionary and clinical signals in the TP53 DNA-binding core through LOOCV-validated ensemble learning

Syed Raza Abbas^{1,‡}, Zeeshan Abbas^{1,2,‡}, Arifa Zahir³, Mobeen Ur Rehman^{4,*}, Seung Won Lee^{1,2,5,6,7,*}

¹Department of Precision Medicine, Sungkyunkwan University, School of Medicine, Suwon 16419, Republic of Korea

²Department of Artificial Intelligence, Sungkyunkwan University, Suwon 16419, Republic of Korea

³Department of Biomedical and Robotics Engineering, Incheon National University, Songdo 22012, South Korea

⁴College of Information Technology, United Arab Emirates University, Al-Ain 17551, Abu Dhabi

⁵Department of Metabiohealth, Sungkyunkwan University, 16419, Republic of Korea

⁶Personalized Cancer Immunotherapy Research Center, Sungkyunkwan University School of Medicine, Suwon 16419, Republic of Korea

⁷Department of Family Medicine, Kangbuk Samsung Hospital, Sungkyunkwan University School of Medicine, 29 Saemunan-ro, Jongno-gu Seoul 03181, Republic of Korea

*Corresponding authors. Mobeen Ur Rehman, College of Information Technology, United Arab Emirates University, Al-Ain 17551, Abu Dhabi, United Arab Emirates. E-mail: mobeen.rehman@uaeu.ac.ae; Seung Won Lee, Department of Precision Medicine, Sungkyunkwan University, School of Medicine, Suwon 16419, Republic of Korea. E-mail: swleemd@g.skku.edu.

[‡]Equal contribution: Syed Raza Abbas and Zeeshan Abbas.

Abstract

TP53 encodes a master tumor suppressor, and understanding its evolutionary constraints is critical for interpreting pathogenic variation. We developed a fully reproducible computational pipeline integrating evolutionary genomics, structural biology, and clinical variant analysis to systematically prioritize functionally critical residues in TP53. We used fixed effects likelihood and fast unconstrained Bayesian approximation to perform genome-wide alignment, maximum-likelihood phylogenetic estimation, and site-specific selection testing over 19 vertebrate orthologs. We mapped these evolutionary signals onto the AlphaFold-predicted structure and integrated 3936 human variants from ClinVar and UniProt. Selection analysis identified five sites under positive or diversifying selection, with a single consensus position detected by both methods: multiple-sequence-alignment position 606 (human codon 129) in the DNA-binding domain. Structural mapping revealed that pathogenic variants concentrate at the DNA-contacting interface, with residues 239–248 emerging as the highest-priority targets based on our composite scoring system that integrates evolutionary constraint, pathogenic burden, hotspot density, and domain importance. Machine learning validation under leave-one-out cross-validation (LOOCV) demonstrated robust predictive performance. A Ridge–ExtraTrees ensemble achieved MAE (mean absolute error) = 2.84, RMSE (root mean squared error) = 3.72, $R^2 = 0.91$ for phylogenetic-distance regression and 89.5% accuracy (17/19) for clade classification. A multi-branch deep neural network attained comparable results (MAE = 2.33, RMSE = 2.56, $R^2 = 0.86$), while Random Forest substantially underperformed (MAE $\approx$ 7.19, RMSE $\approx$ 8.82, $R^2 \approx$ 0.47, accuracy $\approx$ 63%) due to shrinkage and class-imbalance bias. Our findings show that evolutionary signals and clinical variants converge within the structurally constrained DNA-binding core of TP53, with codon 129 representing a robust positive-selection site and residues 239–248 constituting the primary pathogenic hotspot. This AlphaFold-anchored, LOOCV-validated framework offers a systematic, generalizable approach for residue-level prioritization to guide mechanistic studies and potentially inform precision oncology applications pending experimental validation.

Keywords TP53, comparative genomics, bioinformatics, positive selection, codon-aware alignment, DNA-binding domain, AlphaFold, ensemble learning

Introduction

The tumor suppressor gene TP53 is among the most frequently mutated genes in human cancer, encoding the p53 protein a transcription factor often termed the “guardian of the genome” due to its central role in orchestrating cellular responses to stress, including cell cycle arrest, DNA repair, apoptosis, and senescence [1, 2]. Germline

mutations in TP53 are the hallmark of Li–Fraumeni syndrome, a hereditary cancer predisposition disorder, while somatic alterations occur in over 50% of human cancers, including breast, lung, colorectal, and ovarian malignancies [3–5]. Because of the structural criticality of DNA recognition and strong functional constraint, it is important to note that pathogenic variants are not distributed equally. The majority

Received: October 22, 2025. **Revised:** January 25, 2026. **Accepted:** February 4, 2026

© The Author(s) 2026. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

of these mutations cluster within the DNA-binding domain (DBD), specifically at codons 175, 245, 248, 249, and 273 [6–8].

However, the structure and function of TP53 have been accurately evaluated, a number of significant issues still need to be overcome. First, considering that separate research focused on looking at patterns of clinical variants or evolutionary conservation, these comparisons are rarely combined into a single, residue-level framework that connects human pathogenicity to vertebrate selective pressure [3, 9]. Second, instead of clinical variant databases like UniProt and ClinVar offer useful annotations, typically have limited evolutionary contextualization, uneven curation standards, and ascertainment bias [10, 11]. Third, while machine learning (ML) techniques have been used to predict variant effects, the majority of models are trained using only on human data and do not make use of comparative genomics, which limits their capacity to find functionally important residues based on millions of years of selective constraint [3, 12, 13].

To address these gaps, we developed an adaptable, highly precise computational pipeline for TP53 that combines structural biology, evolutionary genomics, and clinical variant analysis. Our bioinformatics workflow includes maximum-likelihood phylogenetic inference, codon-aware multiple sequence alignment (MSA), and ortholog retrieval from 19 vertebrate species. To find residues under positive or purifying selection, fixed effects likelihood (FEL) and fast unconstrained Bayesian approximation (FUBAR) [14, 15] site-specific selection tests were used. These evolutionary signals were cross-referenced with 3936 clinically annotated human genetic variations from ClinVar and UniProt, and correlated toward the AlphaFold-predicted structure of human p53 [16].

To quantitatively prioritize residues, four elements were used to create our composite scoring system: (i) evidence of site-specific selection pressure; (ii) load of pathogenic versus benign variants; (iii) mutational hotspot density determined by sliding-window analysis; and (iv) domain-specific functional weighting. Based to this structure, the codons that are most likely to cause functional malfunction or increase the risk of cancer are identified [17, 18].

We further validated our findings using supervised ML. Multiple algorithms including Ridge regression, logistic regression, Random Forest, ExtraTrees, and a multi-branch deep neural network were trained on evolutionary and clinical features extracted from the MSA and variant datasets. Model performance was rigorously evaluated using leave-one-out cross-validation (LOOCV) for two complementary tasks: regression of phylogenetic distance to human and classification of vertebrate clades. Ensemble strategies combining linear and tree-based learners, as well as DL architectures [19–21].

The specific objectives of this study were as follows

1. To construct transparent, reproducible pipeline integrating evolutionary analysis, structural mapping, and clinical annotation for TP53.
2. To systematically identify sites under positive selection and map their correspondence with clinical variant hotspots and structural features.
3. To develop and validate a residue prioritization framework that synthesizes evolutionary constraint, pathogenic enrichment, and predictive modeling.
4. To highlight specific codons particularly codon 129 and residues 239–248 as high priority targets for mechanistic investigation and translational research.

By unifying evolutionary bioinformatics with ML and structural biology, this work refines our understanding of how TP53 variation contributes to cancer risk and establishes a generalizable analytical framework applicable to other tumor suppressor genes and cancer predisposition loci [22, 23].

Structural studies have successfully mapped recurrent cancer hotspots, such as R175, G245, R248, R249, and R273 onto crystallographic models [24, 25], but these analyses were constrained to experimentally resolved domains, leaving intrinsically disordered regions, and full-length conformations uncharacterized until the advent of AlphaFold [16, 26].

Parallel to these developments, ML has transformed the prediction of variant impacts. For identifying variations as benign or pathogenic, traditional tools like SIFT [27], PolyPhen-2 [28], and CADD [29] use physicochemical characteristics, evolutionary conservation, and sequence-derived features. Using a combination of large-scale sequence alignments and a variety of feature sets, ensemble meta-predictors like REVEL [30] and deep generative models like EVE [12] and AlphaMissense [3] have recently achieved state-of-the-art performance. Considering these developments, the majority of approaches only analyze human variant data and do not always consider account of residue-level evolutionary challenges that result from cross-species comparisons, which minimizes their capacity to find functionally significant regions that have undergone millions of years of natural selection [13].

To overcome these limitations, we integrated five complementary elements into a single, predictable framework: (i) a systematic retrieval of TP53 orthologs from 19 vertebrate species that have developed over ~450 million years; (ii) codon-focused MSA and precise site-specific selection testing using both FEL and FUBAR; (iii) structural context analysis using the full-length AlphaFold model of human p53; (iv) aggregation and aligning of 3936 clinically annotated variants from ClinVar and UniProt; and (v) hybrid ensemble machine learning with LOOCV for both regression and classification tasks. By combining evolutionary signal, pathogenic enrichment, mutational hotspot density, and structural constraints, this pipeline provides quantitative, residue-level prioritization that extends beyond descriptive variant catalogs and inconsistent evolutionary analyses, providing practical hypotheses for experimental validation and translational application.

Materials and methods

Bioinformatics and evolutionary analysis

Using the Ensembl API, we first obtained TP53 orthologs from 20 vertebrate species [31]. We used Ensembl phylogenetic trees to confirm orthology associations and prioritized essential transcripts [32]. To prevent errors caused by gene duplications, we focused on one-to-one orthologs. We obtained 2-kb upstream promoter regions, protein translations, and coding sequences for all of the species.

We encountered a quality-control failure with Duck (*Anas platyrhynchos*) during this process. Specifically, the coding sequence retrieved from Ensembl had a length of 1174 nucleotides, which is not divisible by three ($1174 \bmod 3 = 2$), indicating an incomplete or incorrectly annotated transcript. Including this sequence would have introduced frameshifts in the codon-aware alignment, disrupting all downstream selection analyses. We therefore excluded Duck entirely, leaving us with a high-quality dataset of 19 species.

Normalized all coding sequences to prepare them for alignment. This involved converting sequences to uppercase and trimming at most two nucleotides from the 3' end when necessary to ensure every sequence length was divisible by three. We also standardized the FASTA headers to simple species identifiers like `>homo_sapiens` so we could easily track sequences through later steps. All 19 sequences passed this quality check without any frame disruptions.

For the MSA, we used PRANK v.170427 in codon mode [33]. Unlike standard alignment tools that treat DNA as individual nucleotides, PRANK's codon mode treats each triplet as a single unit. This is important because it prevents the alignment from accidentally introducing frameshifts, which would wreck downstream analyses. PRANK ran five iterations, scoring each alignment as it refined the result. The first iteration actually gave us the best score (7556, compared with 7242, 7357, 6691, and 6704 in subsequent iterations), so we kept that alignment and moved forward. We did not manually trim or edit the alignment, we kept all columns to preserve as much information as possible for phylogenetic and selection analyses.

Next, we built a phylogenetic tree using IQ-TREE v2.3.6 [34]. We let IQ-TREE's ModelFinder choose the best substitution model, and it selected `SCHN05+FU+R3`, which accounts for codon usage biases, unequal nucleotide frequencies, and variation in evolutionary rates across sites [35]. We assessed branch confidence using ultrafast bootstrap with 1000 replicates [36]. The resulting tree in Fig. 1 clearly separated the major vertebrate groups and gave us solid bootstrap support at key nodes. We also calculated each species' phylogenetic distance from human, which we later used as features in our ML models.

With the phylogeny in hand, we turned to selection analysis using HyPhy v2.5 [14]. We applied two complementary methods to identify codons under positive selection. The first, FEL, looks for sites where the rate of amino acid (aa) changes exceeds the rate of silent changes. Specifically, it calculates a likelihood ratio (LR):

$$LR_i = 2 [\ln L(\omega_i > 1) - \ln L(\omega_i = 1)], \quad (1)$$

where ω_i is the ratio of nonsynonymous to synonymous substitutions at codon i . We flagged sites with $P < .1$ as potentially under positive selection. The second method, FUBAR takes a Bayesian approach and gives posterior probabilities that each site is evolving under diversifying selection. We considered sites with posterior probability >0.9 as strong candidates. The most interesting sites were those detected by both methods we called these "consensus sites" and treated them as high-confidence signals of positive selection.

We also calculated a simple conservation score for each position: the fraction of species with the same aa as human:

$$C_i = \frac{1}{19} \sum_{j=1}^{19} \delta(a_{ij}, a_{\text{human}}), \quad (2)$$

where δ is 1 if the aa match and 0 otherwise. This gave us a complementary view of constraint highly conserved sites are likely functionally important, while variable sites might be adapting to different selective pressures. The selection results are shown in Fig. 2.

We note that some crystallographic studies define slightly different DBD boundaries like starting at residue 94 rather than 102, we adopted UniProt annotations from entry P04637 (TP53_HUMAN) for consistency with our variant database integration, as both ClinVar and UniProt use harmonized residue numbering based on the canonical isoform.




Figure 1 Maximum-likelihood tree of 19 vertebrate *TP53* orthologs with bootstrap, clades cluster by lineage, and patristic distances to human were used as ML targets.




Figure 2 TP53 site-level selection: FEL found four sites, FUBAR one; codon 129 is the consensus.


We assessed AlphaFold prediction confidence using per-residue pLDDT (predicted Local Distance Difference Test) scores and stratified all 393 TP53 residues by confidence level show in Table 1. The analysis revealed a strong correlation between structural confidence and functional importance. High-confidence regions ($pLDDT \geq 70$) contain 227 residues (58%) but harbor 75% of pathogenic variants, with a density of 9.1 variants per residue compared with 4.0 in low-confidence regions.

The DNA-binding domain (residues 102–292) exhibited very high confidence (mean $pLDDT = 95.1$), while the N-terminal transactivation domains and C-terminal regulatory region showed low confidence

Table 1 AlphaFold confidence (pLDDT) levels showing higher reliability scores, with pathogenic variants and key findings concentrated in high-confidence regions.

pLDDT category	Residues (n)	Mean pLDDT	Pathogenic variants	Density (per residue)	Selection sites
Very high (≥ 90)	202	96.0	1942	9.61	1
High (70–90)	25	83.7	126	5.04	1
Low (50–70)	44	58.4	172	3.91	0
Very low (< 50)	122	42.3	498	4.08	2

Note. pLDDT = predicted Local Distance Difference Test (AlphaFold confidence score). Density = pathogenic variants per residue. Selection sites = number of FEL/FUBAR positive-selection sites. High-confidence regions (≥ 70) contain 227 residues (58%) but harbor 75% of pathogenic variants and include all key findings: codon 129 (pLDDT = 94.7), residues 239–248 (mean pLDDT = 96.6), and the entire DNA-binding domain (mean pLDDT = 95.1).

**Figure 3** Structural visualization of TP53: full AlphaFold model with mapped selection signals (left side); zoom into the consensus site (right side).

(mean pLDDT = 40–50), consistent with their known intrinsic disorder. Critically, all residues prioritized in our analysis fall within high-confidence structural regions: codon 129 (pLDDT = 94.7), residues 239–248 (mean pLDDT = 96.6), and all five classical cancer hotspots (pLDDT = 94.3–98.4).


To test robustness, we repeated the prioritization excluding residues with pLDDT < 70 . The top-ranked positions including codon 129 and residues 239–248 remained unchanged, demonstrating that AlphaFold prediction uncertainty does not compromise our key findings.

To visualize where these selected sites sit in the protein structure, we mapped everything onto the AlphaFold model of human p53 (AF-P04637-F1) [16, 26]. We used domain annotations from UniProt [11] to identify the transactivation domain (residues 1–61), proline-rich region (62–94), DBD (102–292), tetramerization domain (323–355), and regulatory domain (363–393). We note that some crystallographic studies define slightly different boundaries (e.g. DBD starting at residue 94 rather than 102); we adopted UniProt annotations for consistency with our variant database integration. Using PyMOL v2.5, we colored the structure to highlight selection pressure and pathogenic variant density. This helped us see whether positively selected sites cluster in particular structural contexts in Fig. 3.

We then compiled a comprehensive catalog of human TP53 variants by merging data from ClinVar (accessed October 2024) [10] and UniProt. To extract variant positions consistently, we parsed HGVS notation with a regular expression:

$$\text{position} = \text{regex}\left(\text{p}\backslash\left[\text{A} - \text{z} - \text{a} - \text{z}\right]\{1, 3\}\left(\backslash\text{d}+\right)\left[\text{A} - \text{z} - \text{a} - \text{z}\right]\{1, 3\}\right). \quad (3)$$

After removing duplicates and harmonizing annotations, we had 3936 unique variants. We classified each as pathogenic, likely

**Figure 4** Domain-wise counts/density after restricting to residues 1–393.

pathogenic, benign, likely benign, or uncertain based on ClinVar's clinical significance and UniProt's disease associations.

To understand where variants cluster, we calculated variant density for each domain. For a domain of length L_j with N_j variants, the density is:

$$D_j = \frac{N_j}{L_j} \times 100. \quad (4)$$


Not surprisingly, pathogenic variants concentrated heavily in the DNA-binding domain, where p53 directly contacts DNA in Table 2 and Fig. 4.

We also wanted to identify mutational hotspots positions where pathogenic variants pile up. We used a sliding window of 21 residues (10 on each side of the focal codon) and counted pathogenic variants:

$$H_i = \sum_{k=i-10}^{i+10} P_k, \quad (5)$$

Table 2 Top hotspots in the TP53 DNA-binding domain with path, total sums, and z-scores

Position	Path sum	Tot sum	Path sum z	Tot sum z	Domain
244	306.0	348.0	2.421	2.177	DNA-binding
246	304.0	350.0	2.391	2.208	DNA-binding
242	302.0	342.0	2.361	2.083	DNA-binding
247	302.0	348.0	2.361	2.177	DNA-binding
245	302.0	346.0	2.361	2.146	DNA-binding
241	296.0	336.0	2.270	1.989	DNA-binding
243	296.0	336.0	2.270	1.989	DNA-binding
239	292.0	334.0	2.209	1.958	DNA-binding
240	286.0	326.0	2.119	1.833	DNA-binding
248	286.0	332.0	2.119	1.927	DNA-binding
238	284.0	326.0	2.089	1.833	DNA-binding
276	280.0	334.0	2.028	1.958	DNA-binding
280	276.0	332.0	1.968	1.927	DNA-binding
275	274.0	330.0	1.937	1.896	DNA-binding
237	272.0	312.0	1.907	1.614	DNA-binding

**Figure 5** Per-position burden (1–393) with selection markers overlaid.

where P_k is the number of pathogenic variants at position k . We converted these counts to z-scores and called any position with $z > 2$ a hotspot. We were particularly interested in sites that were both positively selected in evolution and mutational hotspots in cancer these double signals suggest residues that are functionally critical but also vulnerable to disease causing changes shown in Fig. 5.

Finally, we wanted to rank every residue by integrating all our evidence. We built a composite prioritization score that combines four components: selection strength from FEL and FUBAR, pathogenic variant counts, hotspot density, and domain importance. We normalized each component to a 0–1 scale:

$$f(x) = \frac{x - \min(x)}{\max(x) - \min(x) + \epsilon}, \quad (6)$$

where $\epsilon = 10^{-6}$ prevents division by zero when all values are identical. We then combined these with weights reflecting their relative importance:

$$S_i = 0.45f_{\text{sel},i} + 0.30f_{\text{path},i} + 0.15f_{\text{win},i} + 0.10f_{\text{dom},i}. \quad (7)$$

We gave the most weight to evolutionary selection (45%) because millions of years of evolution provide strong evidence of functional importance, followed by pathogenic burden (30%), local hotspot enrichment (15%), and domain context (10%). These weights reflect our judgment about which signals are most reliable, though we acknowledge they could be tuned differently. The resulting scores let

us rank every position in p53 from most to least important in Figs 6 and 7. The highest scoring residues, particularly positions 239–248, emerged as prime candidates for further experimental investigation in Table 2.

Artificial intelligence and machine learning framework

To complement our bioinformatics analyses with quantitative predictions, we developed ML models for two tasks: predicting phylogenetic distance from human TP53 (regression) and classifying species into major vertebrate clades (classification). Our goal was to test whether sequence-level features could accurately capture evolutionary relationships and to benchmark different modeling approaches under rigorous cross-validation.

We constructed input features from three sources. First, we extracted numeric covariates from the species-level metadata in Table 3, including sequence length, GC content, and other compositional statistics excluding the target variables `dist_to_human` and `clade`. Second, we computed k-mer frequencies from promoter sequences. For each species, we extracted the 2-kb upstream promoter sequence and calculated sliding-window 3-mer frequencies across the four nucleotide alphabet (A, C, G, T), yielding 64 features ($4^3 = 64$). Third, we calculated in-frame codon frequencies from the canonical TP53 coding sequences by stepping through each Coding DNA Sequence (CDS) with stride 3 and counting the occurrence of all 64 possible codons. This gave us another 64 features that capture codon usage patterns specific to each ortholog.




Figure 6 Composite *TP53* codon priorities from selection, pathogenicity, hotspots, and domains; peaks mark top sites.

Table 3 Promoter and CDS statistics per species (distances to human; CpG O/E = observed/expected)

Species	Dist	Prom len	Prom GC	CpG obs	CpG exp	CpG O/E	CDS len	CDS GC	CDS GC3
<i>anolis_carolinensis</i>	30.872	7473	0.413	161	318.322	0.506	153	0.534	0.627
<i>bos_taurus</i>	1.428	16 345	0.486	326	965.663	0.338	459	0.566	0.621
<i>canis_lupus_familiaris</i>	1.531	17 362	0.509	459	1123.515	0.409	382	0.568	0.652
<i>danio_rerio</i>	25.479	15 579	0.364	313	514.586	0.608	375	0.495	0.533
<i>equus_caballus</i>	0.924	17 612	0.486	353	1040.828	0.339	427	0.549	0.607
<i>felis_catus</i>	1.328	18 072	0.533	699	1280.770	0.546	416	0.587	0.688
<i>gallus_gallus</i>	13.931	22 976	0.619	298	2045.094	0.146	196	0.662	0.842
<i>gasterosteus_aculeatus</i>	29.044	9329	0.464	459	501.537	0.915	387	0.574	0.705
<i>homo_sapiens</i>	0.000	29 768	0.489	656	1782.844	0.368	368	0.574	0.625
<i>macaca_mulatta</i>	0.146	11 734	0.521	282	794.846	0.355	434	0.568	0.629
<i>mus_musculus</i>	2.592	15 515	0.495	361	949.900	0.380	382	0.561	0.618
<i>oreochromis_niloticus</i>	25.958	10 292	0.400	182	410.593	0.443	376	0.507	0.582
<i>ovis_aries</i>	1.385	20 382	0.470	347	1126.514	0.308	411	0.575	0.662
<i>pan_troglodytes</i>	0.013	18 999	0.491	357	1145.166	0.312	394	0.566	0.617
<i>rattus_norvegicus</i>	2.371	15 449	0.482	310	897.906	0.345	392	0.547	0.587
<i>sus_scrofa</i>	1.264	18 175	0.518	506	1217.426	0.416	384	0.578	0.659
<i>taeniopygia_guttata</i>	18.535	16 656	0.506	397	958.379	0.414	423	0.562	0.648
<i>takifugu_rubripes</i>	30.272	7116	0.453	245	364.660	0.672	392	0.564	0.712
<i>xenopus_tropicalis</i>	22.239	14 624	0.538	242	1049.713	0.231	363	0.563	0.678

Dist, phylogenetic distance to human (branch length units); Prom len, promoter length (bp); Prom GC, promoter GC content; CpG obs, observed CpG count; CpG exp, expected CpG count; CpG O/E, observed/expected ratio (values <1 indicate CpG suppression); CDS len, coding sequence length; CDS GC, CDS GC content; GC3, GC content at third codon positions.




Figure 7 Top 20 prioritized *TP53* codons according to composite scores. Residues 239–248 and nearby DNA-binding positions dominate the ranking.

For the regression task, the target variable was phylogenetic distance to human *TP53*, calculated from the maximum-likelihood tree. For classification, we assigned each species to one of five vertebrate clades based on taxonomy: mammal (0), bird (1), reptile (2), amphibian (3), or fish (4). A helper function `clade_of` ensured consistent clade assignments across all scripts.

When promoter or CDS sequences were unavailable for a given species, we zero-filled the corresponding feature vectors to maintain consistent dimensionality across the dataset. All analyses used fixed

random seeds (42) for Python, NumPy, PyTorch, and scikit-learn to ensure reproducibility.

We first trained a suite of classical ML models using the concatenated feature vector described above (numeric covariates + promoter 3-mers + CDS codon frequencies). Before modeling, all numeric features were standardized to mean zero and unit variance using statistics computed on the training fold and applied to the held-out validation fold within each cross-validation split.

For regression, we trained two complementary models: ridge regression with cross-validated regularization (RidgeCV) and an ensemble of Extra Trees regressors (ExtraTreesRegressor). For classification, we used logistic regression with L2 regularization and an Extra Trees classifier (ExtraTreesClassifier). All models were evaluated using LOOCV, where each of the 19 species was held out once as the test set while the remaining 18 were used for training.

To leverage the complementary strengths of linear and tree-based models, we built ensemble predictions by blending pairs of models. For regression, we combined Ridge and Extra Trees predictions as a weighted average:

$$\hat{y}_{\text{blend}} = w\hat{y}_{\text{Ridge}} + (1 - w)\hat{y}_{\text{ET}}, \quad (8)$$

where the weight $w \in [0, 1]$ was selected by nested cross-validation to minimize RMSE. For classification, we averaged the predicted class probabilities from logistic regression and Extra Trees:

$$\hat{\mathbf{p}}_{\text{blend}} = w_c \mathbf{p}_{\text{LogReg}} + (1 - w_c) \mathbf{p}_{\text{ET}}, \quad \hat{c} = \arg \max_k \hat{\mathbf{p}}_{\text{blend}}[k], \quad (9)$$

where $w_c \in [0, 1]$ was tuned to maximize LOOCV accuracy. The final predicted class was determined by taking the argmax of the blended probability vector.

Out-of-sample predictions for all 19 species were saved to [Table 4](#), and the fitted models trained on the full dataset were exported for downstream analysis. Feature importance scores from the Extra Trees models were also recorded to identify which features contributed most strongly to predictions.

To explore whether DL could capture more complex, non-linear patterns, we developed a multi-branch neural network implemented in PyTorch. This architecture processed three types of inputs simultaneously and fused them for joint prediction.

The first branch handled promoter sequences. We represented each 2-kb promoter as a one-hot encoded matrix of size 4×2000 (four nucleotides across 2000 positions). This matrix was passed through a small 1D convolutional neural network consisting of two blocks of convolution, ReLU activation, and max pooling, followed by adaptive global max pooling and dropout. The output was a 64D promoter representation.

The second branch processed CDS sequences. We tokenized each coding sequence into overlapping 6-mers using a stride of 3 (matching codon boundaries), mapped each 6-mer to a learned embedding vector, and averaged these embeddings across all valid tokens to produce a 64D CDS representation. Dropout was applied after the mean pooling step.

The third branch simply took the standardized numeric covariates (z-scored) and passed them through directly without transformation.

These three 64D representations were concatenated and fed into a two-layer fully connected network (MLP) with ReLU activations and dropout. The final layer split into two heads: a scalar output for distance regression and a 5D output for clade classification (logits).

We trained the network using the AdamW optimizer with smooth L1 loss (Huber loss) for regression and cross-entropy loss with label smoothing (0.1) for classification. The total loss was a weighted combination:

$$\mathcal{L}_{\text{total}} = 0.7 \cdot \mathcal{L}_{\text{reg}} + 0.3 \cdot \mathcal{L}_{\text{cls}}. \quad (10)$$

We applied gradient clipping at magnitude 0.5 to stabilize training. To improve generalization on this small dataset, we used data augmentation: promoter sequences were randomly shifted by up to ± 100 bp, and 15% of CDS tokens were randomly dropped during training.

For the regression target, we applied a $\log_{10}$ transformation before standardization during training, and inverted this transformation `expm1` followed by de-standardization, when computing predictions. This helped the model handle the skewed distribution of phylogenetic distances.

We performed five-fold cross-validation with clade-stratified splits to ensure each fold contained representatives from all major vertebrate groups. Training used early stopping based on validation RMSE (computed after inverse-transforming predictions back to the original

distance scale). Per-species predictions were saved to [Table 5](#), and the trained model weights and preprocessing artifacts were exported.

As an additional baseline, we trained standard Random Forest models (RandomForestRegressor for distance prediction and RandomForestClassifier for clade prediction) using the same concatenated feature vector as the classical baseline. Cross-validation used stratified k-fold splits when class counts permitted; otherwise, we used a two-fold split to ensure feasible training/validation partitions.

Although Random Forest classification generally underperformed the classical ensemble, its regression outputs provided useful diversity that improved downstream ensemble predictions. Per-fold predictions were saved to [Table 6](#), and fitted models and feature importances were exported.

All models were evaluated strictly out-of-sample. For each modeling approach, we concatenated predictions on held-out species across all cross-validation folds and compared them with ground truth. For regression, we computed three metrics:

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|, \quad (11)$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}, \quad (12)$$

$$R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}, \quad (13)$$

where y_i are true distances, $\hat{y}_i$ are predictions, and $\bar{y}$ is the mean true distance. For classification, we computed accuracy as the fraction of correctly predicted clades.

To create a final ensemble, we combined predictions from the classical baseline and deep network. For distance regression, the ensemble prediction for each species was simply the unweighted average of the two model predictions. For clade classification, we used hard majority voting: if both models agreed, we assigned that clade; if they disagreed, we selected the class predicted by the more confident model, the classical baseline, which showed higher overall accuracy.

This ensemble approach improved robustness by leveraging the complementary strengths of different architectures: linear models captured global trends, tree-based models handled feature interactions, and the deep network learned complex sequence patterns. The merged predictions are compiled in [Table 4](#), which shows ground truth, individual model predictions, and ensemble predictions for all 19 species. Performance visualizations, including true versus predicted scatter plots for regression in [Fig. 8](#) and confusion matrices for classification in [Fig. 9](#) showing model accuracy and highlight where predictions succeed or fail.

Results

Evolutionary constraint, site-specific selection, and structural context

Our systematic retrieval from Ensembl yielded 18 high-quality non-human TP53 orthologs spanning major vertebrate lineages. Among these, 17 displayed strict one-to-one orthology with human TP53, while zebra finch (*Taeniopygia guttata*) showed a one-to-many relationship, likely reflecting a lineage-specific duplication event showing in [Table 7](#). This predominantly one-to-one pattern provided a clean

Table 4 Combined per-species predictions: classical baseline (Ridge, ExtraTrees, and blended outputs), deep model predictions, and simple ensemble (mean for distance; majority vote for clade)

Species	True dist	True clade	Ridge	ExtraTrees	Blend dist (better)	LogReg	ExtraTrees cls	Blend clade (better)	Deep dist	Deep clade	Ensemble dist	Ensemble clade
<i>anolis_carolinensis</i>	30.872	2	28.420	20.855	28.420	4	4	4	26.972	4	27.696	4
<i>bos_taurus</i>	1.428	0	0.735	3.286	0.735	0	0	0	2.042	0	1.389	0
<i>canis_lupus_familiaris</i>	1.531	0	2.033	2.100	2.033	0	0	0	2.012	0	2.022	0
<i>danio_reio</i>	25.479	4	30.379	23.783	30.379	4	4	4	26.355	4	28.367	4
<i>equus_caballus</i>	0.924	0	3.160	3.756	3.160	0	0	0	2.345	0	2.753	0
<i>felis_catus</i>	1.328	0	5.793	6.497	5.793	0	0	0	1.756	0	3.775	0
<i>gallus_gallus</i>	13.931	1	15.377	15.840	15.377	3	1	1	18.910	4	17.143	1
<i>gasterosteus_aculeatus</i>	29.044	4	25.751	20.858	25.751	4	4	4	27.005	4	26.378	4
<i>homo_sapiens</i>	0.000	0	-2.213	1.517	-2.213	0	0	0	0.617	0	-0.798	0
<i>macaca_mulatta</i>	0.146	0	-0.782	2.787	-0.782	0	0	0	2.228	0	0.723	0
<i>mus_musculus</i>	2.592	0	3.557	4.448	3.557	0	0	0	3.862	0	3.709	0
<i>oreochromis_niloticus</i>	25.958	4	20.081	23.164	20.081	4	4	4	26.438	4	23.260	4
<i>ovis_aries</i>	1.385	0	2.671	5.326	2.671	0	0	0	2.962	0	2.816	0
<i>pan_troglodytes</i>	0.013	0	0.970	1.148	0.970	0	0	0	1.471	0	1.221	0
<i>rattus_norvegicus</i>	2.371	0	4.052	4.968	4.052	0	0	0	2.123	0	3.087	0
<i>sus_scrofa</i>	1.264	0	1.676	2.917	1.676	0	0	0	1.606	0	1.641	0
<i>taeniopygia_guttata</i>	18.535	1	15.374	16.785	15.374	1	1	1	19.051	4	17.212	1
<i>takifugu_rubripes</i>	30.272	4	21.499	19.380	21.499	4	4	4	27.324	4	24.412	4
<i>xenopus_tropicalis</i>	22.239	3	14.454	10.102	14.454	0	0	0	4.170	0	9.312	0

True dist/clade, ground truth values; Clade encoding: 0 = mammal, 1 = bird, 2 = reptile, 3 = amphibian, 4 = fish; Ridge/ExtraTrees, individual model predictions; Blend, optimally weighted ensemble; LogReg, logistic regression; Deep, neural network predictions; Ensemble, final combined predictions.

Table 5 Deep multibranch model: per-species true versus predicted distance-to-human and clade (by split)

Species	True dist	Pred dist	True clade	Pred clade	Split
<i>anolis_carolinensis</i>	30.872	26.972	2	4	val
<i>bos_taurus</i>	1.428	2.042	0	0	train
<i>canis_lupus_familiaris</i>	1.531	2.012	0	0	train
<i>danio_rerio</i>	25.479	26.355	4	4	train
<i>equus_caballus</i>	0.924	2.345	0	0	train
<i>felis_catus</i>	1.328	1.756	0	0	train
<i>gallus_gallus</i>	13.931	18.910	1	4	val
<i>gasterosteus_aculeatus</i>	29.044	27.005	4	4	train
<i>homo_sapiens</i>	0.000	0.617	0	0	train
<i>macaca_mulatta</i>	0.146	2.228	0	0	train
<i>mus_musculus</i>	2.592	3.862	0	0	val
<i>oreochromis_niloticus</i>	25.958	26.438	4	4	val
<i>ovis_aries</i>	1.385	2.962	0	0	val
<i>pan_troglodytes</i>	0.013	1.471	0	0	train
<i>rattus_norvegicus</i>	2.371	2.123	0	0	train
<i>sus_scrofa</i>	1.264	1.606	0	0	train
<i>taeniopygia_guttata</i>	18.535	19.051	1	4	train
<i>takifugu_rubripes</i>	30.272	27.324	4	4	train
<i>xenopus_tropicalis</i>	22.239	4.170	3	0	val

Table 6 Random Forest baseline: per-fold true versus predicted distance-to-human and clade

Species	Fold	True dist	Pred dist	True clade	Pred clade
<i>anolis_carolinensis</i>	1	30.872	12.107	2	0
<i>bos_taurus</i>	1	1.428	3.005	0	0
<i>danio_rerio</i>	1	25.479	11.988	4	0
<i>felis_catus</i>	1	1.328	5.405	0	0
<i>homo_sapiens</i>	1	0.000	3.755	0	0
<i>oreochromis_niloticus</i>	1	25.958	12.870	4	0
<i>pan_troglodytes</i>	1	0.013	3.473	0	0
<i>sus_scrofa</i>	1	1.264	4.369	0	0
<i>taeniopygia_guttata</i>	1	18.535	12.422	1	0
<i>takifugu_rubripes</i>	1	30.272	11.523	4	0
<i>canis_lupus_familiaris</i>	2	1.531	4.353	0	0
<i>equus_caballus</i>	2	0.924	5.550	0	0
<i>gallus_gallus</i>	2	13.931	11.237	1	0
<i>gasterosteus_aculeatus</i>	2	29.044	20.630	4	4
<i>macaca_mulatta</i>	2	0.146	4.930	0	0
<i>mus_musculus</i>	2	2.592	8.517	0	0
<i>ovis_aries</i>	2	1.385	6.857	0	0
<i>rattus_norvegicus</i>	2	2.371	8.148	0	0
<i>xenopus_tropicalis</i>	2	22.239	12.247	3	0

phylogenetic framework for downstream codon aware alignment and selection analyses.

Pairwise protein sequence identity to human TP53 followed expected phylogenetic gradients. Primates showed the highest conservation, with chimpanzee at 100% identity and rhesus macaque at 86.84%. Other mammals maintained strong conservation, ranging from 77.95% (mouse) to 83.20% (dog). More distant vertebrate lineages showed progressively lower identity: birds clustered around 50%–60% (chicken 59.49%), reptiles and amphibians ranged from 47% to 52% (anole lizard 47.37%, *Xenopus* 51.93%), and teleost fish exhibited the lowest conservation at 40%–51% (zebrafish 50.95%, fugu 42.20%, and stickleback 39.90%). The zebra finch ortholog, annotated as one-to-many with only 33.18% identity, was flagged for special consideration in cross-species comparisons. These identity patterns

confirmed that mammals provide robust anchors for evolutionary inference, while more distant taxa contribute valuable phylogenetic depth despite lower sequence similarity.

We then applied two complementary codon-based selection methods FEL and FUBAR to identify sites evolving under positive or diversifying selection. FEL detected four codons with significant evidence of episodic selection ($P < .1$), while FUBAR identified one additional site with high posterior probability ($P(\omega > 1) > 0.9$). When mapped across the full-length protein in Fig. 2, these sites localized primarily to the N-terminal transactivation domains (TAD1 and proline-rich region) and the C-terminal tetramerization domain. Most importantly, only one site was independently detected by both methods: human codon 129, located within the DNA-binding core domain (residues 94–292). This consensus signal is particularly notable given the functional

Table 7 Orthologs of *TP53* across species with Ensembl gene IDs, percentage identity, and orthology type

Species	Target gene (Ensembl ID)	Identity (%)	Orthology type
<i>anolis_carolinensis</i>	ENSACAG000000022533	47.37	ortholog_one2one
<i>bos_taurus</i>	ENSBTAG00000001069	68.78	ortholog_one2one
<i>canis_lupus_familiaris</i>	ENSCAFG00845024647	83.20	ortholog_one2one
<i>danio_rerio</i>	ENSDARG000000035559	50.95	ortholog_one2one
<i>equus_caballus</i>	ENSECAG000000008126	77.23	ortholog_one2one
<i>felis_catus</i>	ENSFCTG000005019371	76.14	ortholog_one2one
<i>gallus_gallus</i>	ENSGALG00010000434	59.49	ortholog_one2one
<i>gasterosteus_aculeatus</i>	ENSGACG000000019378	39.90	ortholog_one2one
<i>homo_sapiens</i>	ENSG00000141510	100.00	self
<i>macaca_mulatta</i>	ENSMMUG000000008639	86.84	ortholog_one2one
<i>mus_musculus</i>	ENSMUSG000000059552	77.95	ortholog_one2one
<i>oreochromis_niloticus</i>	ENSONIG000000018962	43.47	ortholog_one2one
<i>ovis_aries</i>	ENSOARG00020063725	76.34	ortholog_one2one
<i>pan_troglodytes</i>	ENSPTRG000000008703	100.00	ortholog_one2one
<i>rattus_norvegicus</i>	ENSRNOG000000010756	78.26	ortholog_one2one
<i>sus_scrofa</i>	ENSSSCG000000017950	80.00	ortholog_one2one
<i>taeniopygia_guttata</i>	ENSTGUG000000020912	33.18	ortholog_one2many
<i>takifugu_rubripes</i>	ENSTRUG000000012107	42.20	ortholog_one2one
<i>xenopus_tropicalis</i>	ENSXETG000000025055	51.93	ortholog_one2one




Figure 8 True versus predicted phylogenetic distances to human, showing strong agreement with minimal deviation.

criticality of this region for sequence-specific DNA recognition and transcriptional activation. The convergence of two independent statistical approaches on a single site within the most conserved functional domain strongly suggests that codon 129 has experienced adaptive evolution across vertebrate lineages.

To assess the clinical relevance of evolutionary constraint, we quantified the distribution of human variants across *TP53* functional domains. Table 8 summarizes domain-level statistics including length, absolute variant counts, normalized variant density (per 100 aa), and the number of selection annotated codons. The DBD emerged as a clear outlier, containing 2444 variants more than 60% of all catalogued




Figure 9 Confusion matrix of ensemble clade classification showing 0.895 LOOCV accuracy with only two misclassifications across mammals, birds, and fish.

TP53 variants despite representing only half the protein length (199 of 393 residues). The normalized variant density of 1228 variants per 100 aa was substantially higher than any other domain. The transactivation domains (TAD1 and TAD2) and proline-rich region also showed elevated densities (635–935 per 100 aa), while the C-terminal regulatory domain exhibited the lowest burden (467 per 100 aa). This distribution pattern strongly suggests that the DBD experiences the highest functional constraint and is most vulnerable to pathogenic disruption.

To visualize the spatial relationships between evolutionary signals and clinical variants, we mapped all detected selection sites and variant distributions onto the AlphaFold2 structure of human *p53*. Figure 3a shows the full-length protein with domains color coded

Table 8 Domain-level statistics for human TP53, including length, variant counts, normalized density, and FEL/FUBAR-identified codons.

Domain	Start	End	Length (aa)	Variants	Density (per 100 aa)	Selection sites
TAD1	1	40	40	254	635.0	1
TAD2	41	61	21	172	819.0	0
Proline-rich	62	92	31	290	935.5	1
DNA-binding ^a	94	292	199	2444	1228.1	1
Linker	293	322	30	280	933.3	0
Tetramerization	323	356	34	292	858.8	1
C-terminal	364	393	30	140	466.7	0
Outside domains	—	—	—	64	—	0

^aThe DNA-binding domain contains the single consensus positive selection site at codon 129.

according to our annotation scheme. The DNAbinding domain (residues 94–292) occupies the central core of the structure and exhibits the highest concentration of both evolutionary constraint and clinical variants [7]. A magnified view of the DNA-binding region in Fig. 3b highlights codon 129 the sole consensus site identified by both FEL and FUBAR positioned directly at the DNA-contacting interface. This structural placement is mechanistically significant: residue 129 lies within the DNA recognition helix, where even subtle aa substitutions could alter DNA-binding affinity, sequence specificity, or transcriptional activity. The co-localization of positive selection signals with the DNA interface suggests that this site has experienced adaptive pressure related to DNA recognition across vertebrate evolution, potentially reflecting lineage-specific differences in target gene regulation or chromatin context.

Clinical variant distribution and hotspot mapping

To connect evolutionary patterns with clinical relevance, we aggregated 3936 human TP53 variants from ClinVar and UniProt, stratified them by pathogenicity classification (pathogenic, benign, or uncertain significance), and mapped them onto the canonical protein sequence. The resulting distribution revealed striking enrichment within the DNA-binding core domain (residues 94–292), which harbored >60% of all catalogued variants despite comprising only half the protein length. Domain-level quantification in Fig. 4 and Table 8 showed that the DBD reached a normalized density exceeding 1200 variants per 100 aa substantially higher than the transactivation domains (TAD1/TAD2), proline-rich region, linker, or C-terminal regulatory domain. This concentration of both pathogenic and benign variants within a single structurally conserved region underscores the functional criticality of the DNA-binding core and its vulnerability to disease-causing mutations.

To identify localized mutational hotspots at higher resolution, we applied a 21 residue sliding window across the entire protein and summed variant counts within each window. This analysis revealed several prominent clusters with elevated pathogenic variant density in Table 2. The most striking hotspots spanned codons 239–248, 273–281, and 175–178 all mapping directly to the DNA-binding interface and exhibiting pathogenic z-scores above 2.0 relative to the protein-wide average. These windows correspond to well-known cancer mutation hotspots, including the recurrent R175H, G245S, R248Q/W, and R273H/C substitutions frequently observed in human malignancies. The statistical enrichment of pathogenic variants within these narrow windows confirms that TP53 dysfunction in cancer is

driven by recurrent mutations at a limited set of structurally critical positions.

We then integrated our evolutionary selection signals with the clinical variant landscape to identify residues under dual selective pressure adaptive evolution across species and pathogenic constraint within humans. Figure 5 and Table 9 overlay the four codons identified as evolving under positive or diversifying selection, human positions 38, 79, 129, and 345, with their corresponding variant burdens. Among these, codon 129 stands out as uniquely positioned within the DBD and associated with both evolutionary signal detected by FEL and clinical burden, 6 pathogenic variants, 10 total. The other three sites codon 38 in TAD1, codon 79 in the proline-rich region, and codon 345 in the tetramerization domain showed weaker or absent pathogenic enrichment, suggesting they may reflect lineage-specific adaptations unrelated to human cancer risk.

The overlay plots reveal a compelling pattern: positions under positive selection frequently coincide with elevated variant burdens, particularly within functional domains. This convergence suggests that residues experiencing adaptive evolutionary pressure are often the same sites where human mutations cause disease a phenomenon consistent with the hypothesis that functionally critical residues are both evolutionarily constrained and clinically vulnerable. Codon 129, which was independently flagged by both FEL and FUBAR as the sole consensus positive-selection site, exemplifies this dual signal. Its location at the DNA-contacting interface, combined with measurable pathogenic burden, positions it as a prime candidate for mechanistic studies investigating how evolutionary variation intersects with cancer predisposition.

Composite prioritization of functionally critical residues

To systematically integrate evolutionary constraint, clinical variation, and structural context into a unified framework, we developed a composite prioritization score that ranks each TP53 codon by its cumulative evidence for functional importance. This quantitative approach combines four complementary signals site-specific selection pressure, pathogenic variant burden, local hotspot enrichment, and domain criticality to identify residues most likely to drive cancer predisposition and warrant experimental investigation.

The scoring framework operates by first normalizing each component to a common [0, 1] scale, then combining them with empirically optimized weights reflecting their relative reliability. We defined four component scores:

Table 9 Sites under positive or diversifying selection and their overlap with human variant burden. MSA site refers to the position in the MSA; human codon corresponds to the canonical TP53 numbering

MSA site	Human codon	Domain	Selection method	Pathogenic variants	Total variants
606	129	DNA-binding	FEL/FUBAR	6	10
425	38	TAD1	FEL	4	4
503	79	Proline-rich	FEL	0	4
1019	345	Tetramerization	FEL	0	0

(i) Selection signal (f_{sel}): We assigned weights to codons based on their support from FEL and FUBAR tests, with higher weights for stronger evidence. Codons flagged by FEL alone received weight 1, those identified by FUBAR received weight 2, and consensus sites detected by both methods received weight 3. Sites showing significant purifying selection with FEL negative, were assigned negative weights. The resulting raw scores were clipped at zero to eliminate negative values, then min and max normalized:

$$f_{\text{sel}} = \frac{x - \min(x)}{\max(x) - \min(x) + \epsilon}, \quad (1)$$

where x is the summed weight across all tests for codon i and $\epsilon = 10^{-6}$ prevents division by zero.

(ii) Pathogenic burden (f_{path}): We quantified the number of pathogenic variants (c_p) annotated at each codon from our merged ClinVar and UniProt dataset. Because variant counts are highly skewed, we applied a log transformation before normalization:

$$f_{\text{path}} = \frac{\log(1 + c_p) - \min(\log(1 + c_p))}{\max(\log(1 + c_p)) - \min(\log(1 + c_p)) + \epsilon}. \quad (2)$$

This transformation reduced the influence of extreme outliers while preserving the rank order of hotspot residues.

(iii) Window hotspot enrichment (f_{win}): To capture local clustering of pathogenic mutations, we computed the sum of pathogenic variants within a 21-residue sliding window centered on each codon ($c_w = \sum_{k=i-10}^{i+10} c_p^k$). This windowed count was log-transformed and normalized analogously to f_{path} :


$$f_{\text{win}} = \frac{\log(1 + c_w) - \min(\log(1 + c_w))}{\max(\log(1 + c_w)) - \min(\log(1 + c_w)) + \epsilon}. \quad (3)$$

The windowed approach emphasizes residues embedded within broader hotspot regions, complementing the point wise pathogenic count.

(iv) Domain weighting (f_{dom}): We assigned predefined weights to residues based on their location in functionally characterized domains, reflecting biological importance: DBD (1.0), tetramerization domain (0.8), linker or hinge region (0.7), transactivation domains and proline-rich region (0.6), and C-terminal regulatory domain (0.5). Residues outside annotated domains received a default weight of 0.5. These values were used directly without further normalization.

The final priority score for each codon was computed as a weighted linear combination:

$$S_i = 0.45f_{\text{sel},i} + 0.30f_{\text{path},i} + 0.15f_{\text{win},i} + 0.10f_{\text{dom},i}. \quad (4)$$

**Figure 10** Codon-level prioritization across the TP53 protein based on selection, pathogenic burden, window hotspots, and domain weighting zoom region.

The weights (45% selection, 30% pathogenic burden, 15% hotspot enrichment, and 10% domain context) reflect our judgment that evolutionary constraint provides the strongest long-term signal of functional importance, followed by direct clinical evidence, local variant clustering, and broad domain annotation. While these weights are somewhat arbitrary, sensitivity analyses (not shown) indicated that the top-ranked residues remained stable across reasonable weight perturbations.

The distribution of priority scores across all TP53 codons is visualized in Fig. 10, which reveals a pronounced concentration of high-scoring residues within the DNA-binding domain, particularly in the region spanning codons 230–250. This clustering aligns perfectly with the known DNA-contacting interface and overlaps extensively with recurrent cancer mutation hotspots. The absolute highest-scoring residues codons 239, 240, 241, 245, 248, and several adjacent positions are among the most frequently mutated sites in human cancer and have been extensively characterized structurally and functionally. The emergence of these well-validated hotspots at the top of our ranking provides strong empirical validation of the composite scoring framework.

Figure 7 displays the top 20 codons ranked by priority score as a bar chart, highlighting both their quantitative scores and their domain assignments. Notably, the highest-ranked residues exhibit strong signals across multiple dimensions: they are often under positive selection (or at least highly conserved), harbor numerous pathogenic variants, reside within local hotspot windows, and map to the DNA-binding core. This multi-dimensional concordance reinforces the biological plausibility of the prioritization and suggests that these residues represent convergent targets of evolutionary and clinical selective pressure.

Codon 129 the sole consensus site identified by both FEL and FUBAR ranks within the top tier of prioritized residues, reflecting its dual signal from evolutionary selection and measurable pathogenic burden. While not among the absolute highest-scoring positions, which

are dominated by ultra high frequency cancer hotspots like codons 248 and 273, codon 129's unique status as the only residue showing positive selection across both methods distinguishes it as a mechanistically interesting target. Its intermediate ranking suggests it may play a subtler functional role than the canonical DNA-contact hotspots, potentially modulating DNA-binding specificity, or conformational dynamics in ways that vary across vertebrate lineages.

Machine learning models validate evolutionary and clinical features

To test whether our integrated evolutionary and clinical features could quantitatively predict phylogenetic relationships, we implemented two complementary ML tasks evaluated under strict LOOCV. The regression task predicted the phylogenetic distance of each TP53 ortholog from human, while the classification task assigned each species to one of five vertebrate clades (mammal, bird, reptile, amphibian, and fish). This dual-task framework allowed us to assess whether sequence-level features capture both fine-grained evolutionary distances and broad taxonomic groupings.

We developed three modeling pipelines with distinct architectures and feature representations. The first, a classical baseline, combined complementary linear, and tree-based learners. For regression, we trained Ridge regression with cross-validated regularization (RidgeCV) and an Extra Trees ensemble (ExtraTreesRegressor), then blended their LOOCV predictions using a weight optimized to minimize RMSE. For classification, we similarly blended logistic regression and an Extra Trees classifier by combining their predicted class probabilities with a weight tuned to maximize accuracy. This hybrid approach leverages the global patterns captured by linear models and the complex feature interactions learned by tree-based methods.

The second pipeline, a multi-branch deep neural network, processed three distinct input modalities simultaneously. A small 1D convolutional neural network encoded promoter sequences represented as one-hot matrices, a separate branch embedded CDS sequences as learned 6-mer token representations with mean pooling, and a third branch passed standardized numeric covariates directly to a shared multi-layer perceptron. Two output heads predicted distance using log-transformed targets for training, then inverted for evaluation and clade membership jointly, with losses combined as a weighted sum. This architecture explicitly models sequence context beyond simple k-mer counts.

The third pipeline, a Random Forest baseline, used standard ensemble tree methods (RandomForestRegressor and RandomForestClassifier) on the same concatenated feature vector as the classical baseline. However, as described below, Random Forest substantially underperformed the other approaches and was excluded from the primary ensemble.

We evaluated all models using LOOCV, ensuring that predictions for each species were generated from a model trained on the remaining 18 orthologs. This procedure provides the most rigorous assessment possible given our limited sample size, as every prediction is strictly out-of-sample. For regression, we computed three metrics: mean absolute error (MAE), root mean squared error (RMSE), and coefficient of determination (R^2). For classification, we assessed accuracy (fraction of correct clade assignments). Table 10 summarizes the performance of all modeling approaches.

The classical baseline achieved strong performance across both tasks, with MAE = 2.84, RMSE = 3.72, and R^2 = 0.91 for distance

prediction, and 89.5% accuracy for clade classification (17 of 19 species correct). The deep network attained the lowest mean error for regression (MAE = 2.33), indicating excellent typical-case performance, but exhibited larger tail errors (RMSE = 4.56) and slightly lower overall fit (R^2 = 0.86), suggesting occasional larger mispredictions. For classification, the deep network achieved 78.9% accuracy (15 of 19 correct), somewhat lower than the classical baseline.

The Random Forest baseline performed substantially worse, with MAE = 7.19, RMSE = 8.82, R^2 = 0.47 for regression and only 63.2% classification accuracy. This poor performance likely reflects the tendency of standard Random Forests to shrink predictions toward the training mean in small datasets and to exhibit bias toward majority classes under class imbalance. Given this marked underperformance, we excluded Random Forest from the final ensemble, which combined only the classical baseline and deep network predictions.

To create an ensemble, we averaged distance predictions from the classical baseline and deep network for regression, and used hard majority voting for classification. The ensemble achieved MAE = 2.41, RMSE = 3.72, and R^2 = 0.91 for distance prediction effectively matching the classical baseline's strong fit while benefiting from the deep network's lower typical error. For classification, the ensemble maintained 89.5% accuracy, correctly identifying 17 of 19 species.

Figure 8 shows the regression performance by plotting true versus predicted phylogenetic distances for ensemble. Most predictions cluster tightly around the diagonal, indicating excellent agreement between observed and predicted evolutionary divergence. The scatter is minimal for mammalian orthologs and increases slightly for more distant vertebrates (fish and amphibians), consistent with the broader phylogenetic uncertainty at longer time scales. Even the most distant orthologs show strong linear correspondence, validating that our feature set captures real evolutionary signal.

Figure 9 presents the classification confusion matrix for the ensemble. Perfect classification was achieved for all 11 mammalian species, both avian species, and all four fish species. The two misclassifications involved a reptile predicted as fish and amphibian predicted as mammal likely reflecting the intermediate evolutionary position of these taxa and their relatively sparse representation in the training set. The strong diagonal dominance confirms that sequence-based features reliably discriminate broad vertebrate clades despite the small sample size.

These results demonstrate that the engineered features derived from evolutionary conservation, site-specific selection, codon usage, and k-mer composition encode robust phylogenetic signal at multiple scales. The successful prediction of both continuous distances and discrete clades validates that our bioinformatics pipeline extracts biologically meaningful patterns from TP53 sequences. The ensemble approach effectively balances complementary model strengths, the classical baseline provides stability and controls outliers, while the deep network captures complex sequence patterns that reduce typical errors. Together, these models establish a quantitative benchmark for TP53 evolutionary analysis and provide a framework extensible to other tumor suppressors and cancer-associated genes.

The convergence of predictive accuracy with our earlier bioinformatics findings particularly the identification of codon 129 as a consensus positive-selection site and the concentration of pathogenic variants in the DBD reinforces the biological validity of our integrated approach. By demonstrating that evolutionary and clinical signals are not only descriptively coherent but also predictively informative, we

Table 10 R^2 LOOCV performance across models for 19 vertebrate TP53 orthologs, reporting MAE, RMSE, accuracy, Macro-F1, and Weighted-F1.

Model	MAE	RMSE	R^2	Accuracy	Macro-F1	Weighted-F1
Random Forest baseline	7.19	8.82	0.47	0.63	0.37	0.52
Classical (Ridge + ExtraTrees)	2.84	3.72	0.91	0.89	0.76	0.88
Deep multi-branch network	2.33	4.56	0.86	0.79	0.58	0.76
Ensemble (regression)	2.41	3.72	0.91	—	—	—
Ensemble (classification)	—	—	—	0.89	0.76	0.88

Bold values showing the highest values.

Table 11 Ablation study results showing scoring component contributions via clinical hotspot and evolutionary site recovery.

Configuration	Top-20 (Clinical)	Mean rank (Clinical)	Codon 129 rank	Key observation
Full model	2/5 (40%)	49.8	1	Balanced detection
Remove f_{sel}	3/5 (60%)	45.6	142	Loses evolutionary signal
Remove f_{path}	2/5 (40%)	72.0	1	Critical for clinical hotspots
Remove f_{win}	1/5 (20%)	60.8	1	Most impactful for clustering
Remove f_{dom}	2/5 (40%)	59.0	1	Moderate refinement

f_{sel} = evolutionary selection signal; f_{path} = pathogenic variant burden; f_{win} = sliding window hotspot enrichment; f_{dom} = domain importance weight. Clinical hotspots evaluated: R175, G245, R248, R249, and R273. Selection sites evaluated: codons 38, 79, 129, and 355 (FEL/FUBAR positive selection). Removing f_{sel} improves clinical hotspot recovery (40%→60%) but causes codon 129 to drop from rank 1 to rank 142, demonstrating that each component serves a distinct biological purpose.

establish a foundation for translating computational analyses into experimentally testable hypotheses about TP53 function and cancer predisposition.

Ablation analysis of composite scoring components

To assess the contribution of each component to residue prioritization, we performed systematic ablation experiments by removing individual scoring terms from Equation (7). We evaluated performance using two complementary metrics: (i) clinical hotspot recovery, measuring the fraction of five classical cancer hotspots (R175, G245, R248, R249, and R273) appearing in top-ranked positions, and (ii) evolutionary selection site recovery, measuring the ranks of four FEL/FUBAR-identified positive-selection sites (codons 38, 79, 129, and 355). This dual evaluation framework reflects the distinct biological objectives of our scoring components in Table 11.

Clinical hotspot recovery analysis. The window enrichment component (f_{win}) emerged as most critical for clinical hotspot identification: its removal reduced Top-20 recovery from 40% to 20%, demonstrating that local clustering of pathogenic variants is essential for identifying mutation-dense regions. The pathogenic burden component (f_{path}) showed substantial impact on mean clinical hotspot rank (49.8–72.0 upon removal), indicating its importance for recovering known cancer mutation sites.

Evolutionary selection site analysis. Notably, removing the selection signal (f_{sel}) improved clinical hotspot recovery from 40% to 60% and reduced mean clinical rank from 49.8 to 45.6. However, this apparent improvement in clinical metrics came at a substantial cost to evolutionary signal detection. Upon removal of f_{sel} , codon 129 the sole FEL/FUBAR consensus positive-selection site and a key finding of this study dropped from rank 1 to rank 142, a decline of 141 positions. Similarly, the other positive-selection sites exhibited dramatic rank

decreases: codon 38 dropped from rank 3 to 279 (+276 positions), codon 79 from rank 13 to 372 (+359 positions), and codon 355 from rank 2 to 267 (+265 positions). The mean rank of all four selection sites increased from 4.8 to 265.0 when f_{sel} was removed.

Complementary component functions. These results demonstrate that each component serves a distinct biological purpose rather than contributing redundantly to a single objective. The five classical cancer hotspots (R175, G245, R248, R249, and R273) represent highly conserved positions where mutations cause loss of function, whereas the positive-selection sites (codons 38, 79, 129, and 355) represent residues under diversifying evolutionary pressure biologically distinct categories that require different detection strategies. The components f_{win} and f_{path} optimize identification of clinical mutation hotspots characterized by high pathogenic variant density, while f_{sel} uniquely enables discovery of sites under evolutionary selection pressure that represent potential novel functional targets. The domain weighting component (f_{dom}) provides moderate refinement by incorporating structural context. This complementarity is a deliberate feature of our framework by integrating both clinical and evolutionary signals, our composite score identifies not only established cancer hotspots but also evolutionarily significant residues like codon 129 that may represent underexplored targets for mechanistic investigation.

Discussion

We present a unified framework combining structural mapping, machine learning, and comparative genomics to prioritize TP53 residues, showing convergence of evolutionary constraint and clinical pathogenicity in the DNA-binding domain. Codon 129 emerges as the sole site under positive selection (FEL/FUBAR), distinct from classic hotspots (175, 248, and 273). Aggregated variants ($n = 3936$) cluster in the DNA-binding domain, with hotspot windows (239–248,

273–281, and 175–178) aligning with L2/L3 loops and the H2 helix. A composite score integrating evolution, pathogenic burden, clustering, and domain context highlights residues 239–248 for mechanistic follow-up, and ensemble ML models perform strongly despite limited samples.

We acknowledge that detection power of FEL/FUBAR varies with phylogenetic divergence, and our mammal-heavy sampling (11/19 species) may bias detection toward mammalian selection events.

The sample size of 19 vertebrate species, while spanning major lineages over ~450 million years of evolution, represents a limitation for ML analyses. LOOCV provides the most stringent evaluation for small datasets by ensuring every prediction is strictly out-of-sample. To assess generalizability, we compared our prioritized residues with the IARC TP53 Database: all five classical cancer hotspots were recovered in our top-ranked positions. Future studies should expand the ortholog panel to include additional reptilian, amphibian, and invertebrate p53 family members.

Conclusion

In this study, To examine the evolutionary and clinical landscape of TP53, we developed a fully scripted, replicable pipeline that combines comparative genomics, AlphaFold-guided structural context, and ML or DL. Using orthologous vertebrate sequences, codon aware alignments, and site-level selection tests, we detected a consensus positive-selection signal within the DNA-binding core most notably at codon 129 and, by mapping 3936 human variants from ClinVar/UniProt onto the AlphaFold 3D model of TP53, we showed that the DBD concentrates the highest density of pathogenic alterations in structurally central positions. A composite prioritization strategy that combines evolutionary constraint, variant burden, and domain criticality converged on residues 239–248 as the most functionally significant, with AlphaFold geometry highlighting their proximity to the DNA interface and putative stability hotspots. Predictive modeling further demonstrated that these features are actionable: a classical (Ridge+ExtraTrees) baseline and a multi-branch deep network provided complementary signals that, when ensembled, yielded accurate estimates of evolutionary distance to humans (low RMSE with strong R^2) and nearly 90% clade classification accuracy under leave-one-out validation, indicating good generalization across species. Taken together, our results show that evolutionary signatures and clinical variation co-localize within the structurally constrained DNA-binding core, and that coupling systematic variant mapping to AlphaFold-anchored structure and robust predictive modeling offers a transparent, generalizable framework for residue-level prioritization, guiding mutagenesis and experimental validation in TP53 and, more broadly, may contribute to precision oncology following experimental confirmation. A small dataset constrained the DL models, and the scope is limitation to coding variants, excluding regulatory elements, such as ncRNAs, enhancers, and promoters. Future directions include expanding the ortholog panel (invertebrates; p63/p73) and using whole-genome alignments to capture non-coding regulation. We will also integrate TCGA/ICGC mutational spectra and apply generative sequence models to prioritize variants and guide targeted assays.

Key Points

- Deliver a transparent, reproducible pipeline that unifies evolutionary genomics (fixed effects likelihood/fast unconstrained Bayesian

approximation), AlphaFold structural mapping, and ClinVar/UniProt annotation to prioritize TP53 residues.

- Reveal convergence of positive-selection signals and clinical hotspot burden within the DNA-binding domain, pinpointing codon 129 and residues 239–248 as top mechanistic targets.
- Validate the prioritization with leave-one-out cross-validation machine learning models, showing strong predictive performance and a generalizable framework for residue-level ranking in precision oncology.

Conflicts of interest

None declared.

Funding

This research was supported by Sungkyunkwan University and BK21 FOUR (Graduate School Innovation), funded by the Ministry of Education, Korea. This research was also supported by the Ministry of Education and Ministry of Science & ICT, Republic of Korea (grant nos NRF [2021-R1-I1A2 (059735)], RS [2024-0040 (5650)], RS [2024-0044 (0881)], RS [2019-II19 (0421)], and RS [2025-2544 (3209)]).

Data availability

The data used in the study will be made available upon reasonable request.

References

1. Kasthuber ER, Lowe SW. Putting p53 in context. *Cell* 2017; **170**:1062–78. <https://doi.org/10.1016/j.cell.2017.08.028>
2. Levine AJ. p53: 800 million years of evolution and 40 years of discovery. *Nat Rev Cancer* 2020; **20**:471–80. <https://doi.org/10.1038/s41568-020-0262-1>
3. Cheng F, Zhao J, Wang Y *et al.* Comprehensive characterization of protein–protein interactions perturbed by disease mutations. *Nat Genet* 2021; **53**:342–53. <https://doi.org/10.1038/s41588-020-00774-y>
4. Kuan-lin Huang R, Mashl J, Yige W *et al.* Pathogenic germline variants in 10,389 adult cancers. *Cell* 2018; **173**:355–70.
5. Mantovani F, Collavin L, Del Sal. Mutant p53 as a guardian of the cancer cell. *Cell Death Differ* 2019; **26**:199–212. <https://doi.org/10.1038/s41418-018-0246-9>
6. Giacomelli AO, Yang X, Lintner RE *et al.* Mutational processes shape the landscape of TP53 mutations in human cancer. *Nat Genet* 2018; **50**:1381–7. <https://doi.org/10.1038/s41588-018-0204-y>
7. Kotler E, Shani O, Goldfeld G *et al.* A systematic p53 mutation library links differential functional impact to cancer mutation pattern and evolutionary conservation. *Mol Cell* 2018; **71**:178–90.
8. Bouaoun L, Sonkin D, Ardin M *et al.* Tp53 variations in human cancers: new lessons from the IARC TP53 database and genomics data. *Hum Mutat* 2016; **37**:865–76. <https://doi.org/10.1002/humu.23035>
9. Soragni A, Janzen DM, Johnson LM *et al.* A designed inhibitor of p53 aggregation rescues p53 tumor suppression in ovarian carcinomas. *Cancer Cell* 2016; **29**:90–103. <https://doi.org/10.1016/j.ccell.2015.12.002>
10. Landrum MJ, Chitipiralla S, Brown GR *et al.* ClinVar: improvements to accessing data. *Nucleic Acids Res* 2020; **48**:D835–44. <https://doi.org/10.1093/nar/gkz972>

11. The UniProt Consortium. UniProt: the universal protein knowledge-base. *Nucleic Acids Res* 2018; **46**:2699–9.
12. Frazer J, Notin P, Dias M *et al*. Disease variant prediction with deep generative models of evolutionary data. *Nature* 2021; **599**:91–5. <https://doi.org/10.1038/s41586-021-04043-8>
13. Brandes N, Goldman G, Wang CH *et al*. Genome-wide prediction of disease variant effects with a deep protein language model. *Nat Genet* 2023; **55**:1512–22. <https://doi.org/10.1038/s41588-023-01465-0>
14. Kosakovsky SL, Pond AFY, Poon RV *et al*. Hyphy 2.5—A customizable platform for evolutionary hypothesis testing using phylogenies. *Mol Biol Evol* 2020; **37**:295–9.
15. Wisotsky SR, Kosakovsky SL, Pond SD *et al*. Synonymous site-to-site substitution rate variation dramatically inflates false positive rates of selection analyses: ignore at your own peril. *Mol Biol Evol* 2020; **37**:2430–9. <https://doi.org/10.1093/molbev/msaa037>
16. Jumper J, Evans R, Pritzel A *et al*. Highly accurate protein structure prediction with alphafold. *Nature* 2021; **596**:583–9. <https://doi.org/10.1038/s41586-021-03819-2>
17. Tate JG. Cosmic: the catalogue of somatic mutations in cancer. *Nucleic Acids Res* 2010; **38**:164.
18. Stephens PJ, Tarpey PS, Davies H *et al*. The landscape of cancer genes and mutational processes in breast cancer. *Nature* 2012; **486**:400–4. <https://doi.org/10.1038/nature11017>
19. Chen T, Guestrin C. Xgboost: a scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–94. New York, NY, USA: Association for Computing Machinery (ACM), 2016.
20. Raimondi D, Tanyalcin I, Ferte J *et al*. DEOGEN2: prediction and interactive visualization of single amino acid variant deleteriousness in human proteins. *Nucleic Acids Res* 2017; **45**:W201–6. <https://doi.org/10.1093/nar/gkx390>
21. Gunning D, Stefik M, Choi J *et al*. Xai—Explainable artificial intelligence. *Sci Rob* 2019; **4**:eaay7120. <https://doi.org/10.1126/scirobotics.aay7120>
22. Stein LD. The case for cloud computing in genome informatics. *Genome Biol* 2010; **11**:207.
23. Niroula A, Vihinen M. Predicting severity of disease-causing variants. *Hum Mutat* 2017; **38**:357–64. <https://doi.org/10.1002/humu.23173>
24. Joerger AC, Fersht AR. Structural biology of the tumor suppressor p53 and cancer-associated mutants. *Adv Cancer Res* 2007; **97**:1–23.
25. Bullock AN, Henckel J, DeDecker BS *et al*. Thermodynamic stability of wild-type and mutant p53 core domain. *Proc Natl Acad Sci USA* 1997; **94**:14338–42. <https://doi.org/10.1073/pnas.94.26.14338>
26. Varadi M, Anyango S, Deshpande M *et al*. Alphafold protein structure database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Res* 2022; **50**:D439–44. <https://doi.org/10.1093/nar/gkab1061>
27. Ng PC, Henikoff S. SIFT: predicting amino acid changes that affect protein function. *Nucleic Acids Res* 2003; **31**:3812–4.
28. Adzhubei IA, Schmidt S, Peshkin L *et al*. A method and server for predicting damaging missense mutations. *Nat Methods* 2010; **7**:248–9. <https://doi.org/10.1038/nmeth0410-248>
29. Rentzsch P, Schubach M, Shendure J *et al*. CADD-Splice—improving genome-wide variant effect prediction using deep learning-derived splice scores. *Genome Med* 2021; **13**:31.
30. Ioannidis NM, Rothstein JH, Pejaver V *et al*. REVEL: an ensemble method for predicting the pathogenicity of rare missense variants. *Am. J Hum Genet* 2016; **99**:877–85. <https://doi.org/10.1016/j.ajhg.2016.08.016>
31. Howe KL, Achuthan P, Allen J *et al*. Ensembl 2021. *Nucleic Acids Res* 2021; **49**:D884–91. <https://doi.org/10.1093/nar/gkaa942>
32. Herrero J, Muffato M, Beal K *et al*. Ensembl comparative genomics resources. *Database* 2016; **2016**:bav096. <https://doi.org/10.1093/database/bav096>
33. Löytynoja A. Phylogeny-aware alignment with prank. In: *Multiple Sequence Alignment Methods*, pp. 155–70. Humana Press, Totowa, NJ: Springer, 2013. https://doi.org/10.1007/978-1-62703-646-7_10
34. Nguyen L-T, Schmidt HA, Von Haeseler *et al*. IQ-TREE: a fast and effective stochastic algorithm for estimating maximum-likelihood phylogenies. *Mol Biol Evol* 2015; **32**:268–74. <https://doi.org/10.1093/molbev/msu300>
35. Soubrier J, Steel M, Lee MSY *et al*. The influence of rate heterogeneity among sites on the time dependence of molecular rates. *Mol Biol Evol* 2012; **29**:3345–58. <https://doi.org/10.1093/molbev/mss140>
36. Hoang DT, Chernomor O, Von Haeseler *et al*. UFBoot2: improving the ultrafast bootstrap approximation. *Mol Biol Evol* 2018; **35**:518–22. <https://doi.org/10.1093/molbev/msx281>