An auditable structural pilot of human EEF1A2 D91N

Matched EEF1B2-GDP-Mg predictions and a hypothesis about productive nucleotide exchange

Public working report | eef1a2-d91n-seed11-v1.0.0 | 12 November 2026 UTC

Filippo D. Galbo

Independent researcher and ExoMBridge developer

This is an exploratory computational study and a new public derivative of the preserved seed-11 evidence. The original private report was dated 12 November 2026 UTC; this release adds sole-human-author attribution, public methods and assistance disclosures, privacy edits, and navigable data access. The preserved source models and numerical results are unchanged. It is not peer reviewed, a clinical interpretation, a preprint submission, or a demonstration of a D91N mechanism.

The results and progress figures are a historical seed-11 snapshot: three retained predictions, one per condition, from five attempts. A separate seed-23 successor had begun execution after the original plan; its results are outside this version. No new prediction was launched to prepare this release.

1. Results summary

Abstract

Objective. To establish a traceable structural comparison of human EEF1A2 WT, D91N and D252H in a limited complex with the nucleotide-exchange partner EEF1B2, and use the comparison to formulate experimentally distinguishable hypotheses.

Approach. Canonical full-length proteins were modeled with one GDP and one Mg using the existing Boltz-2 backend. Alignments and inference settings were controlled across variants. Of nine planned condition-seed positions, three seed-11 positions completed, yielding one retained model per condition. Experimental coordinates were audited separately, and molecular identities, confidence artifacts, local distances and protein-interface proximity were checked. Failed attempts, timing scopes and source hashes were retained.

Observations. WT, D91N and D252H contained 461, 365 and 454 protein-interface heavy-atom pairs within 4.0 angstrom, respectively. D91N differed from WT in common-backbone local geometry, while residue 91 remained proximal to EEF1B2 residue 224. The WT/D91N backbone O91-N224 distance changed from 2.661 to 3.869 angstrom, although the residue-level minimum remained 2.661 versus 2.721 angstrom because different atoms participated. Interchain uncertainty was substantial, and single predictions cannot resolve ordinary predictor variation.

Interpretation. The observations motivate a post hoc hypothesis that D91N alters coupling between exchange-partner engagement and nucleotide release. They do not establish an exchange-rate defect, reduced affinity, instability or pathogenicity. Abundance/biogenesis and intrinsic nucleotide-state effects remain competing explanations. The most discriminating follow-up is measurement of intrinsic and

EEF1B2-stimulated GDP release at matched soluble, nucleotide-competent EEF1A2 abundance, with association measured separately.

1.1 What actually ran

Retained condition	Seed	Models	Confidence artifacts	Queue job elapsed	Container elapsed
WT	11	1	4	3 h 52 m 58.489 s	3 h 51 m 59.659 s
D91N	11	1	4	4 h 16 m 13.864 s	4 h 15 m 17.981 s
D252H, successful retry	11	1	4	7 h 54 m 46.151 s	7 h 53 m 53.797 s

At the preserved report cutoff, three analytical positions were complete and six positions at seeds 23 and 47 were unstarted. This is historical accounting, not current queue status. Two additional historical attempts produced no model: an initial one-hour WT timeout and an interrupted D252H attempt during Windows servicing. Thus five execution attempts yielded three retained models. Failed attempts are not computational replicates and processed input files are not models. No new prediction was run while preparing this report.

Execution and analysis chronology

n = 1 retained prediction per condition, seed 11; not measured biology

A Actual container intervals (UTC; November 2026)

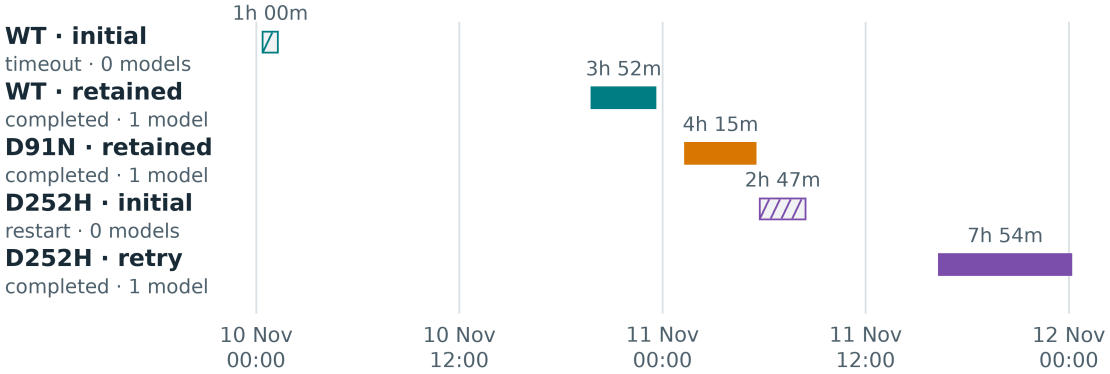

Solid: retained model. Hatched: stopped without a model.

B Analysis milestones after the final run

12 Nov 2026, UTC

00:11:27	• Final prediction container stopped
00:18:06	• D252H coordinate analysis completed
00:19:32	• Matched three-condition comparison completed
00:24:47	• Independent numeric/source audit recorded

C Nine-position analytical panel

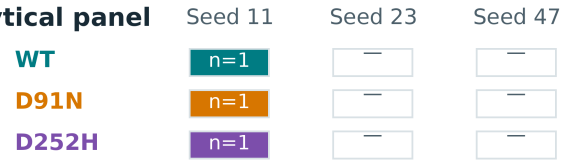

Three completed seed-11 positions; six seed-23/47 positions not started.

Earlier unsuccessful attempts are separate. Container time is not a profiled inference stage.

Figure 1. Historical seed-11 study progress. The figure label "Independent numeric/source audit" refers to an internal automated-agent check, not independent human verification or external peer review. The graphic distinguishes execution attempts, post-run analysis milestones and retained versus unstarted analytical subjects. Completion describes runtime/output acceptance, not demonstrated biological accuracy. Unstarted positions in this historical snapshot are not negative outcomes; subsequent execution is outside this release.

1.2 Measurements of retained predictions

Measurement	WT	D91N	D252H
Interface heavy-atom pairs, distance <=4.0 angstrom	461	365	454
Interface residue pairs, same cutoff	125	106	138
Mean PAE: EEF1A2 frame to EEF1B2 position, angstrom	13.870	16.461	14.604
Mean PAE: EEF1B2 frame to EEF1A2 position, angstrom	15.036	17.660	15.789
Native ipTM	0.85510	0.82933	0.81790
Residue-91 token pLDDT, scale 0-100	72.935	73.789	74.329
EEF1A2 C-alpha20 to C-alpha91, angstrom	6.723	7.479	6.428
EEF1A2 backbone O91 to EEF1B2 backbone N224, angstrom	2.661	3.869	2.696
Residue91 to EEF1B2 residue224 minimum, angstrom	2.661	2.721	2.635
Residue91 to GDP minimum, angstrom	4.776	6.084	4.654
Residue91 to Mg minimum, angstrom	3.904	5.519	4.028
GDP to Mg minimum, angstrom	1.857	2.077	1.911

All values summarize existing predicted coordinates or native confidence outputs. A residue minimum can involve different atoms in different variants. In particular, Asp and Asn side-chain minima are not chemically equivalent correspondences. PAE is directional; the two means are not interchanged or averaged into a binding score. Counts and confidence values were not converted into affinity, exchange rate, stability or clinical effect.

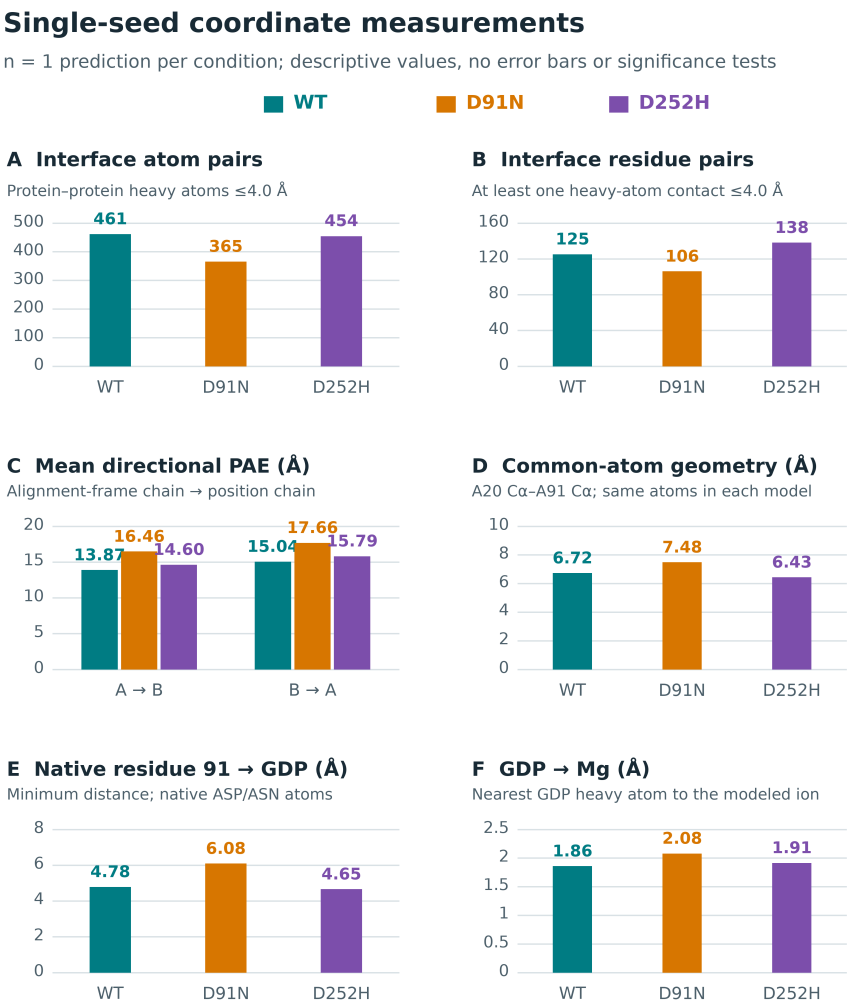

A: EEF1A2; B: EEF1B2. PAE is predictor uncertainty; contact count is not binding affinity.

Panel D uses common Cα atoms. Panel E uses native inventories, not a chemical atom mapping.

Figure 2. Descriptive measurements, one prediction per condition. Distances use the original coordinates and the previously specified atom identities or residue-level minima. No error bars, biological replicate statistics or significance claims are appropriate. Global PAE summaries describe predictor uncertainty, not experimental error bars for individual distances.

1.3 Fitted comparison and technical suitability

Each mutant was fitted to WT once using the same provisional 120-position EEF1A2 C-alpha core. The original fitted transformation was then applied to all evaluation regions without refitting. D91N and D252H core-fit RMSDs were 0.928 and 0.883 angstrom; the corresponding residues 85-110 RMSDs were 1.561 and 1.274 angstrom. Position-91 C-alpha displacements were 1.026 and 0.454 angstrom. These measurements indicate local differences within similarly fitted cores, not thermodynamic stability or folding efficiency.

Entire EEF1B2 RMSDs under that EEF1A2 fit were 21.851 and 32.240 angstrom. These combine partner placement and internal variation, including regions remote from the interface; they must not be reported as interface-only movement or dissociation. The fitting core was chosen as a static correspondence and has not been validated as dynamically invariant.

All 688 protein residues were represented. Identity and finite-coordinate checks passed; the models had 717 confidence tokens and the expected GDP/Mg entities. No cross-entity heavy-atom approaches below 1.5 angstrom were detected under the recorded rule. This limited check does not establish chemically correct ligand geometry or eliminate all clashes. Waters, protonation-dependent interactions and coordination chemistry were not validated.

2. Discussion, literature review and hypothesis development

2.1 Scope of the literature review

This is a focused narrative review, not a systematic review or exhaustive variant catalog. Primary structural, mechanistic and variant-function studies were selected for direct relevance to exchange-factor engagement, residue-91 context, D252H interpretation and protein availability. Authoritative bibliographic records and available primary full text were checked on 12 November 2026 UTC. The source ledger records species, assays, locations and inference boundaries. Literature descriptions below are original summaries, not reproduced figures or text.

eEF1A participates in the nucleotide-dependent translation cycle; exchange-factor action helps restore a GTP-competent state after GDP release. Yeast structural studies of an exchange-factor fragment show that partner engagement can reorganize the nucleotide-binding region 1. Nucleotide-containing yeast complexes provide additional state snapshots, but are not direct human D91N experiments 2. Our binary model also omits other exchange-complex components, the ribosome and tRNA.

2.2 Residue 91 provides a plausible mechanistic connection

Rabbit eEF1A2-GDP structural work proposed that Asp91 participates in a state-dependent rearrangement between a network involving His95, Asn101 and Ser107 and the Lys20/Ser21 region during exchange. The inference used comparisons across rabbit and yeast structures 3. The replacement of a usually charged Asp carboxylate by an uncharged Asn amide therefore motivates examining interaction chemistry as well as overall shape. It does not specify the direction of an exchange-rate effect.

Importantly, our WT complex does not demonstrate the nearby His95/Asn101/Ser107 arrangement observed in the isolated rabbit GDP reference. Nor does it demonstrate a Lys20-Asp91 salt bridge: WT Lys20 NZ is 5.770 and 7.009 angstrom from Asp91 OD1 and OD2. The shorter 3.401-angstrom side-chain minimum is a CB-OD1 pair. A claim that D91N broke a salt bridge or a particular hydrogen-bond network would therefore exceed our evidence.

The literature also requires restraint concerning Mg. Yeast exchange-factor structures and mutational analysis emphasized Mg-related mechanisms 2, whereas the rabbit study found modest Mg effects and questioned a direct bacterial-style Mg mechanism in the mammalian context 3. Species, state and assay differences matter. Our residue91-Mg distances are descriptive and do not establish direct coordination or its loss.

2.3 What D252H contributes, and what it does not

Davies and colleagues reported impaired recovery of the eEF1B complex with D252H in cellular association experiments and examined a corresponding mouse model 4. Mohamed and Klann subsequently studied G70S, E122K and D252H, including complex association and distinct translation/tRNA/actin endpoints; they did not study D91N 5. Cellular co-immunoprecipitation is not an isolated EEF1B2 dissociation constant or a GDP-release assay. Findings for those variants cannot simply be assigned to D91N.

Our D252H model retains approximately the WT number of atom-pair proximities and more residue-pair proximities. The pilot therefore does not recover a straightforward geometric proxy for the published association phenotype. D252H is a biologically motivated comparator, not a validated positive control for our contact count, ipTM or PAE. This weakens any attempt to rank D91N by predicted binding impairment.

2.4 Protein availability is a substantive competing explanation

Minoia and colleagues tested D91N in yeast V5-TEF2 and reported reduced protein abundance and reduced Chp1 co-immunoprecipitation recovery, with lower levels in the absence of Chp1. Detailed synthesis/turnover follow-up concerned G70S and F98C rather than D91N 6. The result supports investigating biogenesis/availability, but does not establish human D91N turnover, human chaperone dependence or a binding constant. Normal-looking predicted folds cannot exclude inefficient production, aggregation or loss of competent protein.

WT paralog work combining molecular dynamics, HDX-MS and SAXS used rabbit-derived proteins and human-sequence computational models, illustrating why local nucleotide-dependent behavior can matter without a gross fold change. That study did not test D91N or D252H 7. It motivates dynamics-sensitive experiments rather than validating our mutant mechanism.

2.5 How the current hypothesis emerged

The initial question was whether a matched structural comparison involving EEF1B2 could provide useful D91N hypotheses. Sequence and reference audits came first. A resource-limited WT attempt timed out; a longer WT attempt established execution feasibility. Before mutant inference, a native-parser audit exposed a query-dependent alignment-deduplication difference. A normalization repair and downstream checks established the scoped effective-MSA comparability described in Methods. Matched D91N completed, and an interrupted D252H run was preserved before a separate retry produced the third model. Technical corrections, admission checks and receipts are recorded separately from biological interpretation.

The first descriptive result was fewer interface proximity pairs in D91N. Atom-level review then showed that residue 91 remained close to EEF1B2 residues 223 and 224 despite altered participating groups. The common backbone O91-N224 approach lengthened, whereas an Asn91 ND2-to-EEF1B2 backbone-O approach was short. Those observations shifted the useful question from a contact-count ranking to the functional consequences of a changed local interaction pattern. No hydrogen-bond classification was performed.

The D252H result limited the association interpretation, and the yeast D91N abundance evidence introduced a distinct alternative. Together, these considerations led to the following hypothesis, formulated after viewing the seed-11 results:

D91N may alter the local interaction network coupling EEF1A2 engagement with EEF1B2 to productive GDP release, while retaining partner proximity.

This is a post hoc, falsifiable hypothesis, not a preregistered prediction confirmed by the data. Slower stimulated exchange is one possible branch; faster intrinsic release, unchanged exchange, altered occupancy and reduced competent-protein abundance remain plausible. Multiple mechanisms could coexist. A failure to resolve a difference would weaken this particular explanation without establishing that the variant is functionally normal.

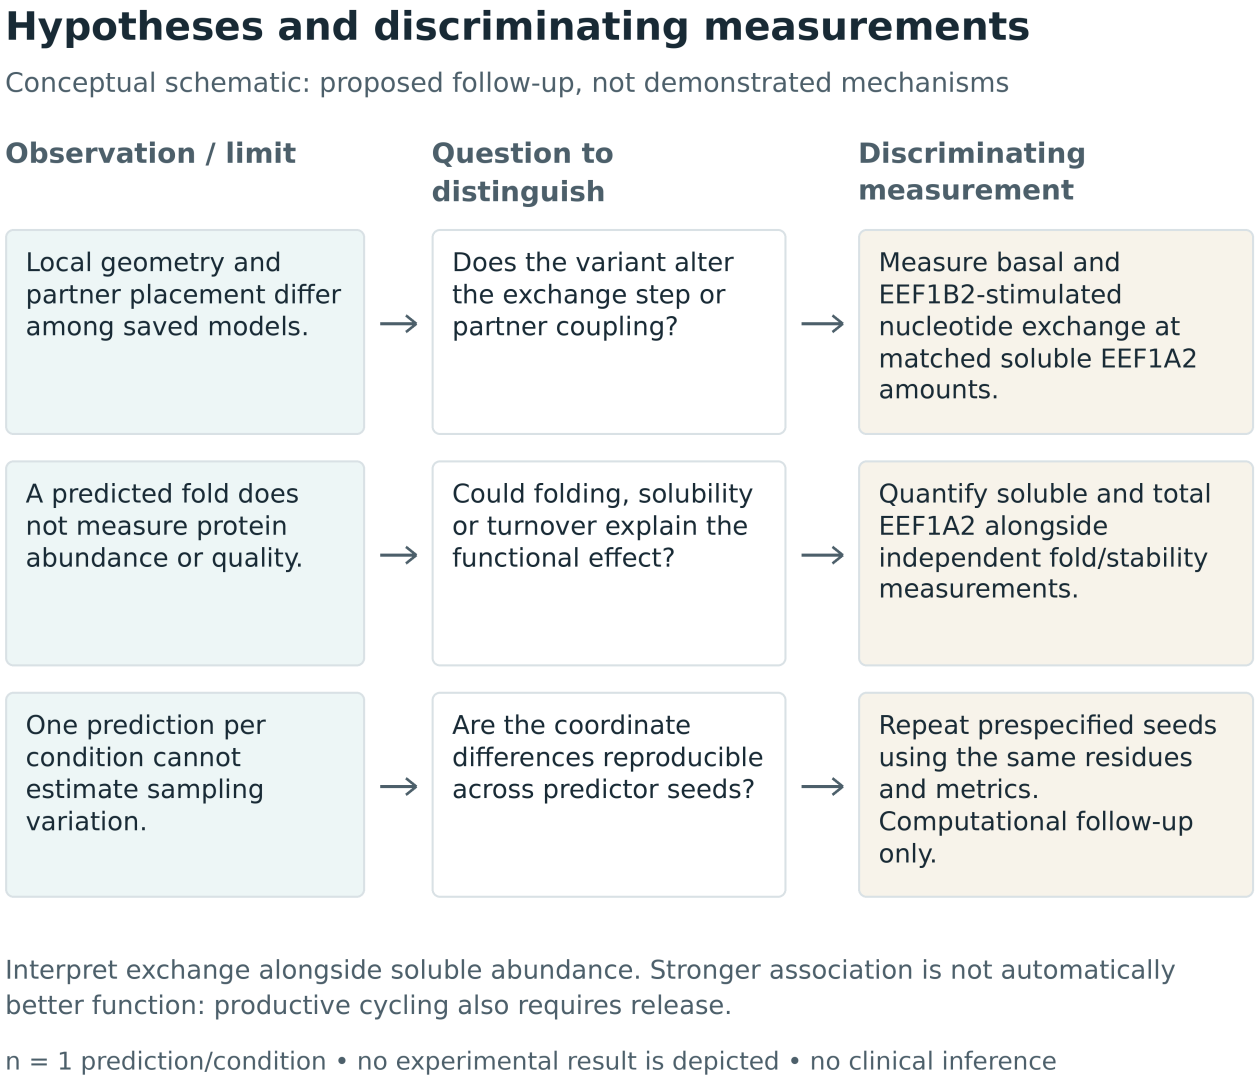

Figure 3. Conceptual hypothesis diagram, not a molecular rendering or measured kinetic pathway. The proposed relationship between local interaction changes and exchange behavior requires testing. The alternatives deliberately separate association, nucleotide handling and competent-protein availability.

2.6 Prioritized follow-up experiments

The primary question is: **Does D91N change EEF1B2-dependent acceleration of GDP release after controlling for soluble, nucleotide-competent EEF1A2 abundance?** The designs below are proposals; none was performed in this pilot.

Priority	Proposed measurement	Essential comparison/control	Interpretation boundary
1	Intrinsic and EEF1B2-stimulated nucleotide release/exchange time courses	WT/D91N/D252H, matched competent protein and nucleotide loading; no-GEF baseline; EEF1B2 concentration series	Distinguishes stimulation from basal behavior; does not alone identify a physical interaction
2	Partner association under matched nucleotide conditions	Independent concentration-dependent assay; expression/input controls for cellular co-IP	Separates occupancy from productive exchange; co-IP is not a Kd
3	Soluble abundance, competence and biogenesis	Matched expression/preparation; total versus soluble protein; nucleotide-binding competence	Tests whether fewer usable molecules explain a functional difference
4	Matched dynamics and orthogonal structural probing	Defined nucleotide/Mg/protonation systems; multiple independent trajectories; reference-state controls	Tests persistence of local differences; simulations alone do not measure biological exchange
5	Frozen computational seed expansion	Seeds 23 and 47 for every condition, identical backend/settings and prespecified readouts	Tests prediction sensitivity, not biological replication or clinical validity

For the exchange experiment, a validated nucleotide-release or exchange assay should be selected and validated for the experimental system. Record complete time courses, nucleotide occupancy and competent protein concentration. If a fluorescent nucleotide analog is used, check reporter dependence with an orthogonal readout where feasible. Keep buffer, temperature, Mg conditions, construct identity, purification and nucleotide-loading history matched. Do not infer an exchange rate from an endpoint signal alone. Fit a justified kinetic model and report fit residuals and uncertainty; avoid imposing a single-exponential model merely for convenience.

A useful primary endpoint is the change in EEF1B2 stimulation relative to intrinsic release at matched competent-protein concentration. Report both the absolute basal/stimulated rates and the concentration response; a fold-change alone can conceal an altered baseline. Similar association with altered stimulation favors coupling; a stimulation change explained by reduced occupancy favors association. Altered intrinsic release points more directly to nucleotide-state behavior. Similar rates per competent molecule accompanied by reduced availability favors biogenesis. These are distinguishable tendencies, not exclusive diagnoses.

Use independently prepared protein batches as experimental units, with technical measurements nested within them. Randomize assay order, blind analysis labels where practical, and select replication based on preliminary assay variability and a prespecified effect of interest. Three predictor seeds are not a justification for three biological replicates. Publish or retain all curves, exclusions and failed preparations, with reasons defined before comparing conditions.

For a cellular follow-up, distinguish total abundance, soluble abundance and functional competence before attributing translation changes to exchange. If later assays show a phenotype, a controlled WT/mutant mixture could address whether behavior in a mixed-protein context differs from the isolated

variants. Such a mixture would be a new experiment, not a modeled feature of this pilot. Chaperone observations from yeast require independent validation in the chosen mammalian system.

For a future MD study, first define exact constructs, nucleotide state, metal and water treatment, protonation, force fields, restraints, preparation history and trajectory protocols. Compare distributions of common-backbone distances, chemistry-specific side-chain contacts and hydration rather than selecting illustrative frames. The His95/Asn101/Ser107 region, Lys20/Ser21 region and residue91-EEF1B2 terminal interface are candidate observables, not assumed persistent bonds. Independent starting conformations and trajectory replicas should distinguish initial-model dependence from reproducible behavior. No MD trajectory was generated by this pilot.

Before additional predictions, freeze the readout list and interpretation criteria: common-atom distances, native-chemistry minima kept separate, original core-fit displacements, local confidence and directional interchain uncertainty. Analyze all conditions and samples, including failures. A consistent difference across seeds is computational support for follow-up, not proof of exchange kinetics. No new prediction was launched to prepare this report.

2.7 Limitations and scientific claim boundary

There is one retained prediction per condition, no experimental D91N exchange data from this project, no measured affinity and no MD trajectory generated here. Predictor training-data independence is not established by omitting explicit templates. The binary full-length EEF1B2 complex does not capture all exchange machinery, nucleotide-cycle states, modifications, solvent chemistry or cellular context. Whole-partner alignment measurements are influenced by uncertain regions. The provisional core selection and local readouts require sensitivity analysis.

Technical checks, deterministic exports, hash preservation and same-coordinate replay demonstrate traceability and calculation reproducibility within their tested boundaries. They do not validate molecular accuracy, causal interpretation or clinical relevance. The previously inspected reference structures are homologous experimental context, not direct human-mutant controls. The central deliverable is a reproducible question with explicit alternatives.

3. Methods

3.1 Molecular identities and provenance

Human EEF1A2 was based on canonical UniProt Q05639, 463 residues, with Asp at positions 91 and 252, cross-checked against RefSeq NP_001949.1. Human EEF1B2 was canonical P24534, 225 residues. Separate EEF1A2 inputs contained WT, D91N only or D252H only; no double mutant, truncation or affinity tag was introduced. Public records were retrieved on 9 November 2026 at approximately 23:54:46 UTC; exact offset-bearing timestamps and byte hashes were retained. UniProt is the sequence-source resource, not evidence of a functional effect 8.

The modeled chain mapping was A, EEF1A2; B, EEF1B2; C, GDP; D, Mg. One GDP and one Mg per binary protein complex were a preliminary matched modeling assumption. No covalent Mg bond or preferred coordination geometry was imposed. The GDP entity contained 28 heavy atoms and the Mg entity one atom. All three conditions preserved identical partner and ligand composition.

3.2 Frozen alignments and inference

Raw alignment acquisition used the installed Boltz 2.2.1 compute_msa workflow with public canonical WT proteins and ColabFold-server greedy pairing 9. The returned files were frozen on 10 November 2026 at approximately 00:08 UTC. The server database revision was not exposed; exact returned raw files are therefore the reproducibility anchor, rather than a claim that a future server query will return identical alignments. No new MSA request was made for this report.

The successful WT used the original raw alignment. Before mutant runs, an audit found that query-only mutation did not guarantee identical effective homolog evidence: native deduplication retained a WT-identical homolog for mutant queries, affecting which row survived the native limit. A repair removed the same six query-colliding nonquery rows from all three derived EEF1A2 raw alignments (indices 472, 478, 488, 522, 1027 and 11921, query row zero and header excluded). Retained homolog bytes, order and pairing keys were preserved. EEF1B2 raw alignment bytes were unchanged. Mutants used these normalized inputs; WT was not rerun.

Native parsing produced 8,192 rows per protein. Normalized WT parsed arrays were exactly equal to the successful WT job's saved arrays. Reconstruction of downstream processing yielded 11,067 effective rows across 717 tokens, with 19 token features and seven MSA features equal for WT. Mutant changes were confined to the intended residue token and source-traced query-derived MSA/profile cells. Pairing seeds 42, 11, 23 and 47 were checked. The native pairing contains query-derived rows at both 0 and 1; an earlier audit's row-zero-only assumption failed and was corrected, with the unsuccessful receipt retained.

These checks support matched effective homolog evidence within the audited feature scope, not byte-identical WT/mutant raw inputs. They reconstruct scoped features from saved processed inputs and byte-identical Boltz Python source; the original inference did not save every feature tensor, and compiled-library equivalence is not asserted. Chemically different residues necessarily differ in molecular features. Processed NPZ files were never substituted for raw MSA inputs. Exact files, hashes, query checks and provenance are linked in the methods audit and study manifest.

Boltz 2.2.1 was the installed prediction backend. The Boltz-2 method publication is cited for the model, not as validation of protein-protein affinity in this use case 10. Each retained job used seed 11, three recycling steps, 200 sampling steps, one diffusion sample, maximum parallel samples one, inference potentials enabled, mmCIF output and the existing no_kernels configuration. PAE and PDE outputs were retained. No explicit structural templates, imposed interface/pocket restraints or affinity predictions were requested. The report does not mix AF3 predictions with Boltz predictions, even though inherited job/container names contain AF3 identifiers.

Runs used the existing local GTX 1080 workflow with four CPUs, a 14-GiB container-memory limit and the recorded 12-hour total-job deadline for the longer attempts. The original WT one-hour deadline belongs to the unsuccessful feasibility attempt. Five-second sampled GPU supervision and the recorded temperature/staleness policy were retained. Device telemetry is sampled whole-device telemetry, not continuous protection or process-specific peak memory. Runtime/image/checkpoint identity is preserved in the source receipts and methods ledger.

The installed runtime used bf16-mixed precision; the GTX 1080 lacks native bfloat16 support and Torch reported emulation support. No precision override, GPU change or host fallback was applied. Native MSA settings included max_msa_seqs=8192, subsample_msa=false and num_subsampled_msa=1024 (the latter is not an assertion that subsampling was enabled). The isolated study admitted up to 700 polymer residues; the normal production 128-residue guardrail remained unchanged. Launch required temperature

<=75 C with telemetry no older than 15 seconds; stale supervision or sampled temperature >=93 C invoked the existing cancellation policy.

The runtime source revision was cb04aeccdd480fd4db707f0bbafde538397fa2ac. The image digest was sha256:d4f637900d8608dd84f38379aa9ffb8b33f82824139888d312c50bbe3e48704c. The structure/confidence checkpoint was boltz2_conf.ckpt (2,286,561,469 bytes), SHA-256 090e82ac8c92f5e943fa1b39e7410a44027bea7243c0bbb3caa67a77fc1428e1. Per-study configuration fingerprints differ; the matching claim concerns the verified engine image, checkpoint and effective scientific settings, not identical whole-study configurations.

3.3 Reference structures

The preserved experimental references were 4C0S (rabbit eEF1A2 with nucleotide/metal context), 1F60 (yeast eEF1A with an exchange-factor fragment, nucleotide free), and 1IJE/1IJF (yeast exchange-factor complexes with partially modeled GDP, with Mg in 1IJF) 3 1 2. References were downloaded and audited separately from predictions. Actual ligand atoms, modifications, missing coordinates and sequence correspondence were inspected; an entry title was not treated as proof of complete GDP geometry.

The canonical rabbit sequence matched canonical human EEF1A2, but deposited modifications and missing coordinates still mattered. Yeast residue mapping was retained explicitly. Both nucleotide-containing yeast references lacked beta-phosphate atoms in the deposited coordinate audit, limiting direct comparisons to the modeled full GDP. Experimental structures and predicted complexes were never counted together as a sample ensemble.

3.4 Geometry, alignment and uncertainty

The saved analysis used explicit chain/residue/atom identities and original Cartesian coordinates. Contacts were protein-protein heavy-atom distances <=4.0 angstrom, inclusive. Atom-pair and residue-pair counts were reported separately. A contact means spatial proximity, not a bond or favorable energetic interaction. The minimal cross-entity approach screen used <1.5 angstrom. Neither criterion is a complete stereochemical assessment.

Mutant-to-WT fitting used one rigid transformation based on the same saved 120-position C-alpha correspondence. Evaluation of other residues and EEF1B2 used that transformation without refitting. Original model files were unchanged. Exact core positions, method source, transformations and per-residue values are supplied in the companion evidence. Missing or ambiguous data were not replaced with invented contacts.

The core mask was the intersection of deposited helix/sheet positions and observed C-alpha coordinates in both 4C0S protein copies within human positions 4-234, excluding 15-25 and 85-110. The 120 positions in compact range notation are 7-13, 26-31, 35-47, 57-66, 75-77, 79-81, 111-116, 121-127, 132-143, 147-152, 154-158, 163-182, 184-188, 190-192, 214-216, 223-225 and 227-234. Position 163 is modified in the experimental source and was retained for backbone correspondence; no side-chain equivalence is implied. The existing unweighted Gemmi 0.6.5 right-to-left fit and NumPy 1.26.4 analysis were reused.

Native pLDDT, pTM/ipTM, PAE and PDE were preserved. Local token values and direction-specific interchain PAE summaries were separated from coordinate distances. The confidence values are model outputs, not measurements of affinity, catalytic activity, pathogenicity or experimental uncertainty. No significance tests, biological confidence intervals, confidence-weighted affinity scores or automatic best-binder rankings were produced.

3.5 Verification and reproduction

Internal AI-assisted reviews checked molecular identity, source integrity, exported values and atom-level interpretation. The existing numerical audit recomputed 1,376 C-alpha displacements from the original models using saved transformations, checked 926 per-residue rows, 51 directed-PAE rows, 231 local-geometry rows, 11 Mg-neighbor rows, 369 residue pairs and 50 native mutation-atom records. It did not independently recompute complete contact enumeration or all raw confidence matrices; those limits remain explicit.

A preserved same-coordinate replay reproduced seven scientific CSVs, totaling 1,641 data rows, and the saved fit JSON byte-for-byte. Runtime metadata differed as expected. This establishes reproducibility of those calculations on the same coordinates, not independent inference reproducibility. Wrapper, routing and archive tests are historical checks with exact receipts; they were not falsely rerun during manuscript preparation. No production application code was changed for this report.

4. Detailed evidence and reproducibility appendix

4.1 Evidence hierarchy

Evidence class	Present in this work	What it establishes
Authoritative molecular inputs	Public sequence records, raw alignments, exact submitted molecules	Identity and traceability
Experimental structural references	Audited deposited coordinates	Existing-coordinate context, with species/state limitations
New predictions	Three retained seed-11 models and confidence artifacts	Executed model outputs, not biological truth
Derived geometry	CSV/JSON, saved transforms and unchanged-source checks	Reproducible measurements of these coordinates
Hypothesis	Post hoc interaction-coupling proposal with alternatives	A testable interpretation, not a demonstrated mechanism
Proposed experiments	Designs and discriminating outcomes	Future work only; no invented results

4.2 Public evidence and analysis

The release provides all three model coordinates and their twelve native confidence artifacts in 02_models/; measurements and saved transformations in 03_results/; molecular inputs, frozen alignments and public sequence/structure records in 05_inputs/; and a sanitized scientific methods record in 06_methods/. Source date fields retain their original timestamps. Job identifiers and machine paths have been removed from public derivatives.

The saved coordinate analyses are repeatable independently of the private workbench. recompute_anchors.py reads the released mmCIF and confidence arrays and checks the five headline readouts for every condition. It launches no inference and writes no model files. Historical replay evidence in 06_methods/historical-analysis-replay.json separately records seven scientific CSVs and saved fits

reproduced from the same coordinates. Repeating an analysis of saved coordinates is different from reproducing the original prediction runtime.

Queue-job, container, native progress, runner and analysis durations are different scopes. The historical D252H telemetry epoch and monotonic spans differed by 240.695 seconds for an unestablished reason; timing differences are not inferred stage durations. These operational uncertainties do not change the saved coordinates or the measurements reported here.

4.3 Authorship and contributions

Filippo D. Galbo is the sole human author of this public working report and dataset, and the independent researcher and developer responsible for the ExoMBridge project and release decisions. His contribution connects the biological question, controlled molecular inputs, workflow development, evidence curation, interpretation boundaries and public research communication. Software development, analysis implementation and drafting used AI assistance; this statement does not imply that every line or calculation was performed manually or independently verified by a human.

Boltz-2 produced the molecular predictions. ExoMBridge supplied the documented local orchestration, input/output records, provenance and structural-analysis workflow. Filippo did not create or train the Boltz foundation model, and the experimental reference structures are credited to their public depositors and cited scientific authors. There is no claimed institutional or employer sponsorship, research team, clinical validation, or commercial service.

4.4 AI assistance, tools, reuse and citation

Recorded automated-agent assistance included code and analysis implementation, execution oversight, data checking, literature retrieval, hypothesis exploration, figure generation and drafting. These checks are internal quality controls, not external peer review. Historical records do not establish every model identity or service revision. The public tools ledger separates established software versions from incomplete historical service evidence; no blanket compliance certification or provider endorsement is claimed. Scientific interpretation and suitability for new research require substantive review beyond these automated checks.

See the tools and methods ledger, also tools-ledger.md in the package root, for scientific citations, versions, applicable license and service documentation, and what was submitted or redistributed. Third-party source notices accompany the sequence, alignment and structure data. Public availability is not an unrestricted license for the author's new report, dataset assembly or private workbench. No new open-source license is assigned to those materials by this release; contact the author through the public website for reuse questions.

Citation: Galbo FD. An auditable structural pilot of human EEF1A2 D91N: matched EEF1B2-GDP-Mg predictions and a hypothesis about productive nucleotide exchange. Public working report and dataset, eef1a2-d91n-seed11-v1.0.0. 2026-11-12. https://www.filippogalbo.com/research/eef1a2-d91n . Exploratory computational study; not peer reviewed.

References

1. G. R. Andersen; L. Pedersen; L. Valente; I. Chatterjee; T. G. Kinzy; M. Kjeldgaard; J. Nyborg. Structural basis for nucleotide exchange and competition with tRNA in the yeast elongation factor complex eEF1A:eEF1Balpha. *Molecular Cell*. 2000;6(5):1261-1266. doi:10.1016/S1097-2765(00)00122-2.

2. G. R. Andersen; L. Valente; L. Pedersen; T. G. Kinzy; J. Nyborg. Crystal structures of nucleotide exchange intermediates in the eEF1A-eEF1Balpha complex. *Nature Structural Biology*. 2001;8(6):531-534. doi:10.1038/88598.

3. Thibaut Crepin; Vyacheslav F. Shalak; Anna D. Yaremchuk; Dmytro O. Vlasenko; Andrew McCarthy; Boris S. Negrutskii; Michail A. Tukalo; Anna V. El'skaya. Mammalian translation elongation factor eEF1A2: X-ray structure and new features of GDP/GTP exchange mechanism in higher eukaryotes. *Nucleic Acids Research*. 2014;42(20):12939-12948. doi:10.1093/nar/gku974.

4. Faith C. J. Davies; Jilly E. Hope; Fiona McLachlan; Grant F. Marshall; Laura Kaminioti-Dumont; Vesa Qarkaxhija; Francis Nunez; Owen Dando; Colin Smith; Emma Wood; Josephine MacDonald; Oliver Hardt; Catherine M. Abbott. Recapitulation of the EEF1A2 D252H neurodevelopmental disorder-causing missense mutation in mice reveals a toxic gain of function. *Human Molecular Genetics*. 2020;29(10):1592-1606. doi:10.1093/hmg/ddaa042.

5. Muhaned S. Mohamed; Eric Klann. Autism- and epilepsy-associated EEF1A2 mutations lead to translational dysfunction and altered actin bundling. *Proceedings of the National Academy of Sciences of the United States of America*. 2023;120(38):e2307704120. doi:10.1073/pnas.2307704120.

6. Melania Minoia; Jany Quintana-Cordero; Katharina Jetzinger; Ilgin Eser Kotan; Kathryn Jane Turnbull; Michela Ciccarelli; Anna E. Masser; Dorina Liebers; Eloïse Gouarin; Marius Czech; Vasili Hauryliuk; Bernd Bukau; Günter Kramer; Claes Andréasson. Chp1 is a dedicated chaperone at the ribosome that safeguards eEF1A biogenesis. *Nature Communications*. 2024;15(1):1382. doi:10.1038/s41467-024-45645-w.

7. Oleksandra Novosylna; Vyacheslav Shalak; Katarzyna Dąbrowska; Ilias Patmanidis; Dmytro Lozhko; Tetiana Bondarchuk; Birgit Schiøtt; Jan S. Pedersen; Charlotte R. Knudsen; Poul Nissen; Michał Dadlez; Boris Negrutskii. Dynamics and structural features of the eEF1A1 and eEF1A2 paralogs. *Nucleic Acids Research*. 2025;53(21):gkaf1177. doi:10.1093/nar/gkaf1177.

8. The UniProt Consortium. UniProt: the Universal Protein Knowledgebase in 2025. *Nucleic Acids Research*. 2025;53(D1):D609-D617. doi:10.1093/nar/gkae1010.

9. Milot Mirdita; Konstantin Schütze; Yoshitaka Moriwaki; Lim Heo; Sergey Ovchinnikov; Martin Steinegger. ColabFold: making protein folding accessible to all. *Nature Methods*. 2022;19(6):679-682. doi:10.1038/s41592-022-01488-1.

10. Saro Passaro; Gabriele Corso; Jeremy Wohlwend; Mateo Reveiz; Stephan Thaler; Vignesh Ram Somnath; Noah Getz; Tally Portnoi; Julien Roy; Hannes Stark; David Kwabi-Addo; Dominique Beaini; Tommi Jaakkola; Regina Barzilay. Boltz-2: Towards Accurate and Efficient Binding Affinity Prediction. *bioRxiv*. 2025;2025.06.14.659707. doi:10.1101/2025.06.14.659707. Cited version: bioRxiv v1 preprint; not journal-peer-reviewed.