De novo design of neutralizing protein binders against CCL26 and POSTN, two convergent targets in eosinophilic esophagitis

Ruth-Anne Pai, PhD[1,2]

[1] Citizen Scientist, EoE patient-researcher; Built with Claude: Life Sciences Hackathon (June–July 2026)

[2] Pai Advisory, LLC

* Correspondence: Ruth-Anne Pai (Pai Advisory, LLC) — advising@ruthannepai.com

Preprint — not peer reviewed. This manuscript was generated with substantial AI assistance (Anthropic Claude Science platform) under human direction. All results are computational (in-silico); no experimental validation has been performed. This document is research and informational content only and is not medical advice.

Abstract

Eosinophilic esophagitis (EoE) is a chronic, food-antigen–driven type-2 inflammatory disease of the esophagus. Approved and late-stage therapies converge on a small set of upstream nodes — the IL-4/IL-13 receptor (dupilumab), IL-5/IL-5Rα (mepolizumab, benralizumab), and TSLP — leaving the downstream effector chemokine axis and the tissue-remodeling axis pharmacologically unexploited. Here we combine a nine-cohort EoE transcriptomic meta-analysis with literature and clinical-pipeline mining to nominate two targets that are strongly and reproducibly dysregulated yet have no direct therapeutic: **CCL26** (eotaxin-3), the eosinophil-selective chemokine that is the single most discriminating EoE transcript, and **POSTN** (periostin), a matricellular remodeling protein that remains elevated in histologic remission alongside persistent desmoglein-1 (DSG1) loss. For each target we carried out structure-based de novo binder design — RFdiffusion backbone generation against a defined functional epitope, SolubleMPNN sequence design, and Boltz-2 co-folding assessment with an on-target epitope-overlap analysis. From 80 backbones and 1,920 designed sequences we computationally assessed 60 binder–target complexes. We deliberately calibrated the co-folding readout against a negative-control set (36 additional complexes: composition-matched scrambled-sequence and random-sequence decoys folded under the identical protocol). This calibration is a central and cautionary result: interface confidence (ipTM) does **not** distinguish designs from scrambled-sequence decoys — 92–100% of scrambled decoys cross the same ipTM > 0.5 / complex-pLDDT > 0.7 thresholds that 59/60 designs cross — so a high pass rate is close to the null and is not by itself evidence of designed, specific binding for these small compact targets. The only readout that separates designs from scrambled decoys is focal epitope engagement, and that signal is modest. Against this calibrated baseline the two leads are the designs that best combine interface confidence with on-target epitope coverage — the CCR3/glycosaminoglycan (GAG) face of CCL26 (ipTM 0.942, 4/5 epitope hotspots, 1,220 Å² buried) and the FAS1 integrin-interaction patch of POSTN (ipTM 0.901, complex pLDDT 0.911, 3/4 hotspots, 842 Å² buried). All designs are computational and require experimental validation; we present them, with an explicit account of what in-silico confidence can and cannot establish, as a mechanistically-grounded starting point for two differentiated EoE therapeutic strategies — anti-chemokine ligand trapping and anti-remodeling neutralization.

Keywords: eosinophilic esophagitis, CCL26/eotaxin-3, periostin/POSTN, de novo protein design, RFdiffusion, ProteinMPNN, Boltz-2, ligand trap

1. Introduction

Eosinophilic esophagitis (EoE) is a chronic immune-mediated disease characterized by eosinophil-predominant inflammation of the esophageal epithelium, symptomatic esophageal dysfunction, and — if untreated — progressive fibrostenotic remodeling. It is driven by a type-2 (Th2) immune response to food antigens, in which epithelial alarmins and Th2 cytokines (IL-4, IL-13, IL-5, TSLP) orchestrate eosinophil recruitment, mast-cell expansion, and barrier disruption.

The therapeutic pipeline reflects this cytokine-centric model. Dupilumab (anti-IL-4Rα, blocking both IL-4 and IL-13 signaling) is approved for EoE; anti-IL-5/IL-5Rα agents (mepolizumab, reslizumab, benralizumab), anti-TSLP (tezepelumab), and anti-IL-13 (cendakimab) are approved in related indications or in late-stage EoE trials. These agents share a common feature: they intercept the disease **upstream**, at the cytokine or eosinophil level. Two biologically central axes remain without a direct therapeutic:

1. **The effector chemokine axis.** CCL26/eotaxin-3 is the terminal, eosinophil-selective signal that converts upstream Th2 activation into physical eosinophil recruitment. It is the single most discriminating transcript separating EoE from gastroesophageal reflux disease (PMID 17900656), yet no approved or late-stage agent targets the eotaxin/CCR3 axis directly.
2. **The remodeling / barrier-persistence axis.** Periostin (POSTN) is a matricellular protein that amplifies subepithelial fibrosis and remodeling. Critically, POSTN accumulation and DSG1 loss persist in histologic remission (PMID 39343172), marking a residual disease process that current eosinophil-directed therapies do not resolve.

We reasoned that these two targets — one at the point of eosinophil recruitment, one at the point of tissue remodeling — represent complementary, unexploited intervention points, and that both are amenable to a neutralizing-biologic (ligand-trap) modality because both are compact, secreted proteins with biochemically-defined functional surfaces.

We note at the outset that "no direct therapeutic" is an observation about the current pipeline, not proof that either target is an unencumbered opportunity, and each carries a specific reason for caution that we return to in the Discussion. For CCL26, the eotaxin family is redundant: CCL11 (eotaxin-1) and CCL24 (eotaxin-2) also signal through CCR3, so trapping CCL26 alone may be partially bypassed by the other eotaxins, and the therapeutic case rests on CCL26 being the dominant, most disease-specific eotaxin in EoE rather than the only one. For POSTN, periostin has historically been studied in type-2 disease as a *biomarker* more than as a validated drug target, and it has physiological roles in tissue repair; a therapeutic binder therefore carries an on-target liability that a biomarker does not. We therefore undertook de novo design of protein binders against a defined functional epitope on each, using a contemporary structure-based pipeline (RFdiffusion → SolubleMPNN → Boltz-2), and report two computationally-assessed leads together with a negative-control calibration of the in-silico readout used to select them.

2. Results

2.1 A nine-cohort meta-analysis nominates CCL26 and POSTN as reproducibly dysregulated, undrugged targets

We assembled a random-effects (DerSimonian–Laird) meta-analysis of nine public EoE case–control transcriptomic cohorts (GSE250595, GSE148381, GSE197702, GSE303169, GSE58640, GSE234973, GSE278888, GSE246323, GSE228083; 25,654 genes; Benjamini–Hochberg FDR control across all genes). Both candidate targets are strongly and consistently up-regulated: **CCL26** at pooled +4.56 $\log_2$ fold-change (adjusted $p = 3\times10^{-4}$, quantified in and up in all 9/9 cohorts, $I^2 = 98\%$) and **POSTN** at +5.55 (adjusted $p = 3\times10^{-16}$, quantified in 8 of the 9 cohorts and up in 8/8, $I^2 = 72\%$) (**Fig. 2A**, **Fig. S1**, **Table S1**). The differing denominators reflect per-gene cohort coverage — POSTN passed expression-level filtering in 8 of the 9 cohorts, CCL26 in all 9 — not an inconsistency in the analysis. Both targets carry substantial between-cohort heterogeneity (reflected in the random-effects model, which widens the pooled confidence interval accordingly); CCL26's very high I^2 (98%) indicates the *magnitude* of up-regulation varies across cohorts even though its *direction* is unanimous (9/9), and a leave-one-cohort-out sensitivity analysis confirms the pooled effect is not driven by any single cohort (**Table S1**). Both sit among the most strongly up-regulated genes in the disease, alongside effector-cell markers (CPA3, ALOX15) and opposite the concordantly down-regulated barrier genes (DSG1, SPINK7, FLG).

Cross-referencing the meta-analysis against literature attention (591 EoE abstracts, 2015–2025) and the interventional trial landscape (157 EoE trials) shows that CCL26 and POSTN occupy a distinctive position: high omics effect size and mechanistic literature support, but **zero direct interventional trials** — in contrast to IL-13, IL-5, and TSLP, which dominate both literature attention and the trial pipeline (**Fig. 2B**, **Fig. S2**, **Table S3**). On an omics-derived druggability shortlist that combines effect size, direction concordance, single-cell specificity, and antibody tractability into a composite priority score (formula and weights given in Supplementary Methods), POSTN ranks 3rd of 20 and CCL26 11th of 20 (**Fig. S4**, **Table S2**). CCL26's middle rank on this composite is not in tension with its status as the single most discriminating EoE-vs-GERD transcript: the two statements measure different things — discriminating power is a differential-expression/specificity property, whereas the composite additionally penalizes between-cohort heterogeneity, and CCL26's very high I^2 (98%) lowers its composite score relative to POSTN despite a large, unanimous effect. We prioritize CCL26 for design on mechanistic grounds (it is the terminal, eosinophil-selective, CCR3-exclusive recruitment signal), not on its composite rank.

Single-cell differential expression localizes each target to a distinct epithelial compartment: CCL26 to suprabasal epithelium and POSTN to basal epithelium (**Fig. 2C**, **Fig. S3**), consistent with CCL26 as an epithelially-secreted chemokine and POSTN as a basally-produced remodeling signal.

2.2 Mechanistic rationale and epitope definition

CCL26 (eotaxin-3). CCL26 signals exclusively through CCR3 and is an eosinophil-selective chemoattractant induced by IL-4 (≈100× more potent than IL-13) and IL-13, synergizing with TNF-α and suppressed by glucocorticoids (PMID 12061839). Its chemotactic function requires two molecular interactions, either of which is a valid neutralization target: (i) CCR3 engagement via the chemokine N-terminus and N-loop, and (ii) glycosaminoglycan (GAG) binding via basic residues in the α-helical

region, which immobilizes the chemokine into the haptotactic gradient eosinophils follow (PMID 35742962, PMID 26701132). We defined the design epitope as the composite CCR3/GAG basic face spanning the N-loop (residues 16–17) and the C-terminal α-helix basic cluster (R54/K55/K56) (**Fig. S6**). The neutralization logic is an **anti-ligand mini-binder / trap** that sequesters secreted CCL26 and occludes this functional surface — preferred over an anti-CCR3 approach because CCR3 is a broadly-expressed seven-transmembrane GPCR, whereas mature CCL26 is a compact (~71-residue) secreted protein with a clean extracellular target surface.

POSTN (periostin). The primary epithelial barrier lesion in EoE is IL-13–driven loss of DSG1; periostin is the top induced gene when DSG1 is lost, and DSG1 loss potentiates inflammation partly *through* POSTN (PMID 24220297). POSTN acts as an integrin ligand that drives myofibroblast differentiation and matrix remodeling (PMID 28918442, PMID 23019192), and luminal eotaxin-3 and periostin both correlate with remodeling severity (PMID 35405206). We defined the design epitope as an exposed patch on FAS1 domain IV corresponding to the integrin-interaction surface (**Fig. S6**). Importantly, a POSTN binder is best framed as an **anti-remodeling / loop-breaking** agent: it would interrupt the IL-13→DSG1↓→POSTN↑ feed-forward circuit and thereby *indirectly* support barrier recovery, rather than directly re-assembling desmosomes. This mechanism is differentiated from every agent in the current EoE pipeline and directly addresses the disease process that persists in histologic remission (PMID 39343172, PMID 35888006).

2.3 A structure-based de novo design pipeline

For each target we prepared a clean single-chain structure of the functional unit — CCL26 from the NMR structure 1G2S (mature 71-residue eotaxin-3) and POSTN from 5WT7 (the isolated FAS1 domain IV, 140 residues) — and defined a spatially-compact hotspot patch on the functional epitope (**Fig. S6**). RFdiffusion generated 40 binder backbones per target (binder length 65–90 aa) targeting these hotspots using the complex/PPI checkpoint. Backbone quality control confirmed that all 40 CCL26 backbones and 34/40 POSTN backbones engaged ≥3 epitope hotspot residues; the flatter, larger FAS1 face produced more variable docking geometry (**Fig. 3B**, **Fig. S5**). SolubleMPNN then designed 1,920 sequences (8 sequences × 3 sampling temperatures × 80 backbones), holding the target fixed and designing only the binder chain (**Fig. 3C**, **Fig. S7**). A sequence-complexity filter (entropy ≥ 2.6, maximum single-residue run ≤ 5, top-amino-acid fraction ≤ 0.40) removed low-diversity designs, retaining 1,074/1,920 (**Fig. S8**), from which 30 designs per target were selected by a composite of MPNN score and sequence entropy for structural validation (**Fig. 3A**).

2.4 Boltz-2 fold-back assessment, negative-control calibration, and lead selection

We assessed the 60 selected binder–target pairs by Boltz-2 co-folding (three recycling steps, three diffusion samples per prediction) and evaluated both interface confidence and on-target engagement (**Fig. 4**). These 60 are not a random sample of the design space: they are 30 per target hand-selected (by a composite of MPNN score and sequence entropy) from the 1,074 sequences that survived complexity filtering out of 1,920 designed on 80 backbones (**Fig. 3A**). All 60 complexes folded successfully, and **59/60 crossed** the community interface-confidence thresholds (ipTM > 0.5 and complex pLDDT > 0.7): 30/30 CCL26 designs (maximum ipTM 0.942, median 0.852) and 29/30 POSTN designs (maximum ipTM 0.904, median 0.780) (**Fig. 4A**, **Fig. S9**, **Table S4**). This is a threshold-crossing rate

on a multiply pre-filtered subset, not a design-success rate — and, as the calibration below shows, on its own it is close to the null.

Negative-control calibration. Because interface-confidence metrics can be confidently wrong, we folded a negative-control set under the identical protocol: for each target, 12 scrambled-sequence decoys (each a real design's binder sequence with residues randomly permuted — identical length and amino-acid composition, interface geometry destroyed) and 4 random-composition decoys, plus 4 verbatim reproduction controls (36 control complexes total; **Fig. 5**, **Table S6**). The reproduction controls recovered the original metrics within diffusion-sampling noise (CCL26 lead ipTM 0.942→0.904, POSTN lead 0.901→0.928), confirming the comparison is like-for-like. The calibration result is cautionary and important: **ipTM does not distinguish designs from scrambled decoys.** Scrambled decoys crossed the ipTM > 0.5 / pLDDT > 0.7 gate at 92% (CCL26) and 100% (POSTN) — indistinguishable from the design pass rates (100%, 97%) — and Mann–Whitney tests of design > scrambled ipTM are non-significant for both targets (CCL26 $p = 0.16$; POSTN $p = 0.80$, where scrambled scored marginally higher). Complex pLDDT excludes random-composition sequences for CCL26 (random median 0.55, 0% pass) but only weakly for POSTN (random median 0.73, 75% pass), and does nothing against scrambled decoys for either target. For these small, compact targets, then, a high ipTM/pLDDT pass rate is not by itself evidence of designed, specific binding.

On-target engagement is the discriminating readout — with a modest margin. Because interface confidence alone can be optimistic about geometry that is confident but off-target, we additionally scored each design by an on-target composite, ipTM × (0.5 + 0.5 × epitope-coverage) × min(pLDDT/0.7, 1), which rewards designs that are both confident and docked onto the intended epitope hotspots (**Fig. 4C**, **Fig. S11**). This composite correlates with, but is not identical to, raw ipTM (Spearman ρ = 0.70 across all 60; ρ = 0.61 CCL26, 0.78 POSTN), and it changes lead selection for POSTN (the top design by on-target score differs from the top design by ipTM alone), so it functions as a meaningful on-target tie-breaker rather than a relabelling of ipTM; the 0.5/0.5 coverage weighting and the 0.7 pLDDT floor are heuristic and stated as such. Focal epitope engagement is the one axis on which designs beat scrambled decoys: 13/30 CCL26 designs contacted ≥3 of the 5 epitope hotspots (versus 25% of scrambled decoys) and 5/30 POSTN designs contacted ≥3 of 4 (versus 0% of scrambled decoys), consistent with the flatter POSTN integrin face being harder to engage focally (**Fig. S10**). The direction is correct, but the margin is modest and, on this control sample, the design-vs-scrambled difference in epitope-hit count does not reach significance — so we treat epitope overlap as the best available in-silico discriminator, not as proof of specific binding.

2.5 Two lead binders engage their functional epitopes in silico

The top-ranked design per target by on-target score was selected as the lead (**Fig. 1**, **Fig. 6**). We emphasize that "lead" here denotes the best computational candidate against the calibrated readout of §2.4, not an experimentally validated binder. For POSTN in particular, the assessment used the isolated FAS1 domain IV (PDB 5WT7, 140 residues), not full-length multi-domain periostin; the lead's behaviour in the context of the intact protein and its integrin partners is untested.

CCL26 lead (design ccl26_design_39_T0.3, 67-aa binder): interface ipTM 0.942, complex pLDDT 0.779, pTM 0.956; buries 1,220 Å² across a 25-residue target and 27-residue binder interface; contacts

4/5 epitope hotspots (residues 16, 54, 55, 56), engaging both the N-loop and the C-terminal α-helix basic cluster — i.e. the composite CCR3/GAG face targeted by design (**Fig. 6A**, **Fig. S12**).

POSTN lead (design postn_design_39_T0.1, 81-aa binder): interface ipTM 0.901, complex pLDDT 0.911, pTM 0.941; buries 842 Å² across a 22-residue target and 20-residue binder interface; contacts 3/4 epitope hotspots (residues 83, 85, 108) on the FAS1 integrin-interaction patch (**Fig. 6B**, **Fig. S12**).

Residue-level interface distance maps confirm that the closest binder–target contacts localize to the epitope hotspots rather than off-target surfaces for both leads (**Fig. 6C**, **Fig. S12**). Interactive 3D rotation videos of both lead complexes are provided as supplementary media.

3. Discussion

We report two computationally-validated de novo protein binders directed at complementary, currently-undrugged nodes of EoE pathogenesis. The design rationale is grounded in a reproducible nine-cohort omics signal cross-referenced against the mechanistic literature and the clinical pipeline, which distinguishes this work from binder-design exercises that begin from a target chosen for tractability alone.

Mechanistic differentiation. CCL26 and POSTN sit downstream of the cytokine nodes that current therapies target. A CCL26 trap would block the terminal eosinophil-recruitment step regardless of which upstream cytokine (IL-4, IL-13, TSLP-driven Th2) initiated it, making it complementary to — not redundant with — approved anti-IL-4Rα and anti-IL-5 agents. This rationale is contingent on CCL26 being the dominant eotaxin in EoE: because CCL11 and CCL24 also signal through CCR3, a CCL26-selective trap would leave those routes open, and the incremental benefit over the approved agents that already suppress epithelial CCL26 induction (dupilumab, anti-IL-13) is a hypothesis to be tested, not an established gain. A POSTN binder addresses the remodeling and barrier-persistence axis that eosinophil-directed therapy leaves behind; because DSG1 loss and periostin accumulation persist in histologic remission, this axis is a plausible residual unmet need. We note that "histologic remission" in EoE is defined by an eosinophil-count endpoint (peak eosinophils/hpf below threshold), so residual POSTN and DSG1 abnormalities in remission indicate ongoing tissue-level processes but are not in themselves evidence of persistent symptoms in an individual patient; whether direct POSTN neutralization adds clinical benefit over existing IL-13-pathway blockade remains untested.

The barrier question. A central motivation was whether targeting POSTN could restore esophageal barrier function. The evidence supports a precise, limited claim. POSTN is not the primary junctional disruptor — that is DSG1 loss — but it is the top gene induced by DSG1 loss and a driver of the feed-forward remodeling loop. A POSTN binder would therefore be expected to support barrier recovery **indirectly**, by interrupting the IL-13→DSG1↓→POSTN↑ circuit, rather than by directly re-assembling desmosomes. The junctional machinery is restorable in principle when the driving insult is removed (PMID 34458999), which makes loop interruption a plausible route to barrier benefit, but this should be framed as a downstream consequence, not a direct binding effect.

Limitations — what the in-silico readout does and does not establish. These designs are entirely in silico, and our own negative-control calibration sets a hard ceiling on how they should be interpreted. The decoy experiment (§2.4, **Fig. 5**) shows directly that for these small compact targets, interface ipTM and the ipTM/pLDDT pass threshold do **not** discriminate real designs from composition-

matched scrambled sequences. The widely-quoted "high ipTM = validated binder" shorthand therefore fails here, and we have avoided the word "validated" for the in-silico results throughout: co-folding confidence establishes fold plausibility and interface geometry, not affinity, not specificity, and — as our controls demonstrate — not even that the binder sequence carries information beyond its amino-acid composition. The strongest in-silico claim the data support is that the two leads are, among a pre-filtered design set, the candidates that best combine interface confidence with focal engagement of the intended epitope, and that focal epitope engagement is the only readout that separates designs from scrambled decoys (by a modest, non-significant margin on the present control sample). This is a starting point for experiment, not evidence of function.

Four further caveats apply. First, the lead binder sequences are strongly helical and acidic, a known compositional bias of SolubleMPNN on these backbone types; the acidic character is electrostatically complementary to the basic CCL26 GAG face but warrants explicit expression, solubility, and specificity testing — and, given the decoy result, the possibility that confident co-folding is partly driven by generic electrostatic complementarity rather than shape-specific recognition. Second, the POSTN target was the isolated FAS1 domain IV, not the full-length multi-domain protein, so binder behaviour in the context of the intact protein and its integrin partners is untested. Third, POSTN has physiological roles in tissue repair, so a therapeutic binder would need to be confirmed neutralizing (not agonist) and sufficiently selective. Fourth, the target-nomination meta-analysis carries substantial between-cohort heterogeneity (I^2 72–98% for the two leads); although the direction of effect is unanimous and robust to leave-one-cohort-out removal, the pooled effect magnitudes should be read as heterogeneous-but-consistent rather than precise point estimates.

Orthogonal-predictor concordance sharpens the caution. We re-folded all 60 designs with an architecturally distinct predictor (Chai-1) under matched settings (**Fig. 7**, **Table S6**). The two models disagree sharply: Chai-1 assigns these binders interface ipTM values far below Boltz-2's (median 0.19 vs 0.85 for the 30 CCL26 designs; 0.33 vs 0.78 for the 30 POSTN designs), only 7% of designs (4/60) reach Chai-1's ipTM > 0.5 line versus 98% in Boltz-2, and per-design ipTM rankings are not concordant overall (Spearman $\rho = -0.13$ across all 60, n.s.; within-target $\rho = -0.14$ for CCL26 and a weak +0.38 for POSTN). No design is confidently a binder in both models, and both leads collapse in Chai-1 (CCL26 lead 0.942→0.134; POSTN lead 0.901→0.225). This is the expected failure mode when a single co-folder's confidence is treated as validation, and it makes explicit that the two leads are Boltz-2-optimistic candidates, not cross-model-robust binders. We therefore refrain from any claim of validated binding and present the leads strictly as prioritized starting points for experiment.

Next steps. The immediate priorities are experimental: recombinant expression of both leads and their targets, biophysical affinity measurement (SPR/BLI), and functional neutralization assays — CCR3/GAG-dependent eosinophil chemotaxis for the CCL26 trap and integrin-signaling / myofibroblast-differentiation readouts for the POSTN binder. Computationally, the priority is a redesign-and-reselect cycle that requires **cross-model** agreement (Boltz-2 ∩ Chai-1 ∩ an AlphaFold3-class model) and beats the scrambled-decoy null on epitope engagement as an explicit acceptance gate, together with explicit affinity prediction and sequence-optimization to reduce the acidic/helical bias, before any synthesis.

4. Methods

Omics meta-analysis and target selection. Nine public EoE case–control transcriptomic cohorts were combined by DerSimonian–Laird random-effects meta-analysis (25,654 genes), yielding pooled $\log_2$ fold-changes, standard errors, heterogeneity (I^2), per-gene cohort counts (k), and Benjamini–Hochberg-adjusted *p*-values across all genes; leave-one-cohort-out re-pooling was used as a sensitivity check for the lead targets. Targets were cross-referenced against a PubMed literature corpus (591 EoE abstracts; search "eosinophilic esophagitis"[Title/Abstract] restricted to 2015–2025 publication dates; the "2026" figure in an earlier draft was a typo and has been corrected) and the interventional trial landscape (157 EoE trials, ClinicalTrials.gov, condition "eosinophilic esophagitis", interventional studies). Per-target abstract and trial counts were obtained by case-insensitive gene-symbol/synonym matching against abstract text and trial intervention/summary fields; the 45-target cross-reference (Table S3) comprises the union of the meta-analysis top hits and targets named in the corpus. "Zero direct interventional trials" for CCL26/POSTN means no interventional trial in this ClinicalTrials.gov set names the target or a direct antagonist of it as the intervention; agents acting upstream (e.g. IL-13 blockade that reduces CCL26 induction) are not counted as direct. The druggability shortlist (Table S2) ranks genes by a composite priority score (see below). Single-cell differential expression was drawn from an EoE single-cell dataset (cell types: endothelial, basal/suprabasal epithelial, mast, myeloid, T cell).

Druggability priority score. The Table S2 shortlist priority score combines, for each gene passing the effect-size and direction-concordance filters, its pooled effect magnitude, direction concordance across cohorts (penalized by heterogeneity I^2), single-cell cell-type specificity, and antibody/ligand-trap tractability (secreted or surface accessibility), normalized to a 0–100 scale. The exact term weights are given in the released shortlist-generation code and Supplementary Methods table; the score is a prioritization heuristic, not a validated druggability model, and mechanistic judgement (not score rank alone) drove final target choice.

Target structure preparation. CCL26 was taken from PDB 1G2S (solution NMR, mature 71-residue eotaxin-3); POSTN from PDB 5WT7 (NMR, FAS1 domain IV, 140 residues). Structures were cleaned to the first model, chain A, standard residues (Biopython). Per-residue solvent-accessible surface area (Shrake–Rupley) confirmed all candidate epitope residues were surface-exposed. Design hotspots were chosen for spatial compactness on the functional epitope: CCL26 residues 16, 17, 54, 55, 56 (N-loop + C-helix GAG/CCR3 basic face); POSTN residues 83, 85, 108, 116 (exposed integrin-interaction patch on FAS1-IV).

Backbone generation (RFdiffusion). RFdiffusion (complex/PPI checkpoint, Complex_base_ckpt.pt) generated 40 binder backbones per target with the target held fixed and hotspot residues specified via ppi.hotspot_res; binder length 65–90 residues, denoiser noise scales set to 0. Backbones were QC'd for binder length, radius of gyration, and number of binder Cα atoms within 10 Å of hotspot Cα atoms.

Sequence design (SolubleMPNN). ProteinMPNN in SolubleMPNN mode (v_48_020 weights) designed 8 sequences per backbone at each of three sampling temperatures (0.1, 0.2, 0.3), designing only the binder chain (chain B) with the target chain (A) fixed — 1,920 sequences total. Sequences were filtered for complexity (Shannon entropy ≥ 2.6 bits, maximum single-residue run ≤ 5, top-amino-acid

fraction ≤ 0.40); 30 designs per target were selected by a composite of MPNN score and entropy for validation.

Structure assessment (Boltz-2). Each selected binder–target pair was co-folded with Boltz-2 (--use_msa_server, 3 recycling steps, 3 diffusion samples), reporting interface ipTM, complex pLDDT, pTM, and an aggregate confidence score. The community confidence gate (ipTM > 0.5 and complex pLDDT > 0.7) was recorded as a threshold-crossing rate, not treated as a validation criterion (see calibration below). On-target engagement was scored as ipTM × (0.5 + 0.5 × epitope-coverage) × min(pLDDT/0.7, 1), where epitope-coverage is the fraction of design hotspots with any target atom within 5 Å of a binder atom; the 0.5/0.5 coverage weighting and the 0.7 pLDDT floor are heuristic. We verified that this composite is not a relabelling of ipTM (Spearman ρ = 0.70 across all 60 designs; it changes the POSTN lead relative to ranking by ipTM alone). The top-ranked design per target by on-target score was selected as the lead. Interface residues (any atom < 5 Å across chains) and buried surface area ((SASA_target + SASA_binder − SASA_complex)/2, Shrake–Rupley) were computed for each lead.

Negative-control calibration. To establish what the co-folding readout can distinguish, we generated a control set and folded it under the identical Boltz-2 protocol: per target, 12 scrambled-sequence decoys (residue-permuted design sequences, preserving length and amino-acid composition), 4 random-composition decoys (length-matched, natural aa background), and 2 verbatim reproduction controls (36 control complexes total). Design vs decoy distributions were compared by one-sided Mann–Whitney U on ipTM, complex pLDDT, epitope-hit count, and on-target score, with AUROC reported (**Table S6**, **Fig. 5**).

Orthogonal-predictor concordance. All 60 designs were additionally co-folded with Chai-1 (Apache-2.0; ESM2 embeddings, MSA server, 3 trunk recycles, 200 diffusion timesteps, seed 42), and per-design interface ipTM was compared against Boltz-2 by Spearman rank correlation to assess single-model dependence (**Table S6**).

Compute, versioning, and visualization. Backbone generation, sequence design, and co-folding were run on GPU (NVIDIA A100-80GB/A10G, Modal). Software: RFdiffusion (complex/PPI checkpoint Complex_base_ckpt.pt, SHA-256 prefix e29311f6…), ProteinMPNN/SolubleMPNN weights v_48_020, Boltz-2 (PyPI boltz) and Chai-1 (chai_lab) at the versions pinned in the released environment specifications; Boltz diffusion seeds default, run-to-run variance characterized by the reproduction controls above. Lead complexes were rendered and animated with PyMOL (target surface, epitope hotspots highlighted, binder cartoon) and encoded to video with ffmpeg. Figures were produced in matplotlib.

5. Data and code availability

All intermediate and final data products are available as project artifacts: the meta-analysis signature, druggability shortlist, literature×omics cross-reference, single-cell differential expression, backbone QC tables, the 1,920-sequence SolubleMPNN table, the 60-complex validation table, lead complex structures (PDB), and lead specification sheets (see **Supplementary Information** for the full index).

Funding

This work received no external funding. It was conducted during the Built with Claude: Life Sciences Hackathon (2026).

Competing interests

The author is a person living with eosinophilic esophagitis and is the sole member of Pai Advisory, LLC. No specific grant funding was received. The author declares no other competing interests.

Author contributions

R.-A.P. conceived the study, defined the target-nomination and design strategy, directed all computational work, and wrote and approved the manuscript.

AI-use disclosure

This research was performed by the author using Anthropic's Claude Science platform for data mining, computational protein design and structure prediction, analysis, figure generation, and drafting. The author directed the scientific strategy, made all decisions, and is responsible for the content. Claude Science also served as a self-correcting check that flagged and corrected over-claims during the work. Published machine-learning methods used as tools (e.g., RFdiffusion, ProteinMPNN/SolubleMPNN, Boltz-2, AlphaFold) are cited in Methods.

Acknowledgements

This work was conducted as a citizen-science effort during the Built with Claude: Life Sciences Hackathon (2026), organized by Anthropic, Gladstone Institutes, and Cerebral Valley. The author thanks the maintainers of the open data resources and structure-prediction tools used here, and acknowledges the eosinophilic esophagitis community whose lived experience motivates this work.

Figure legends (main)

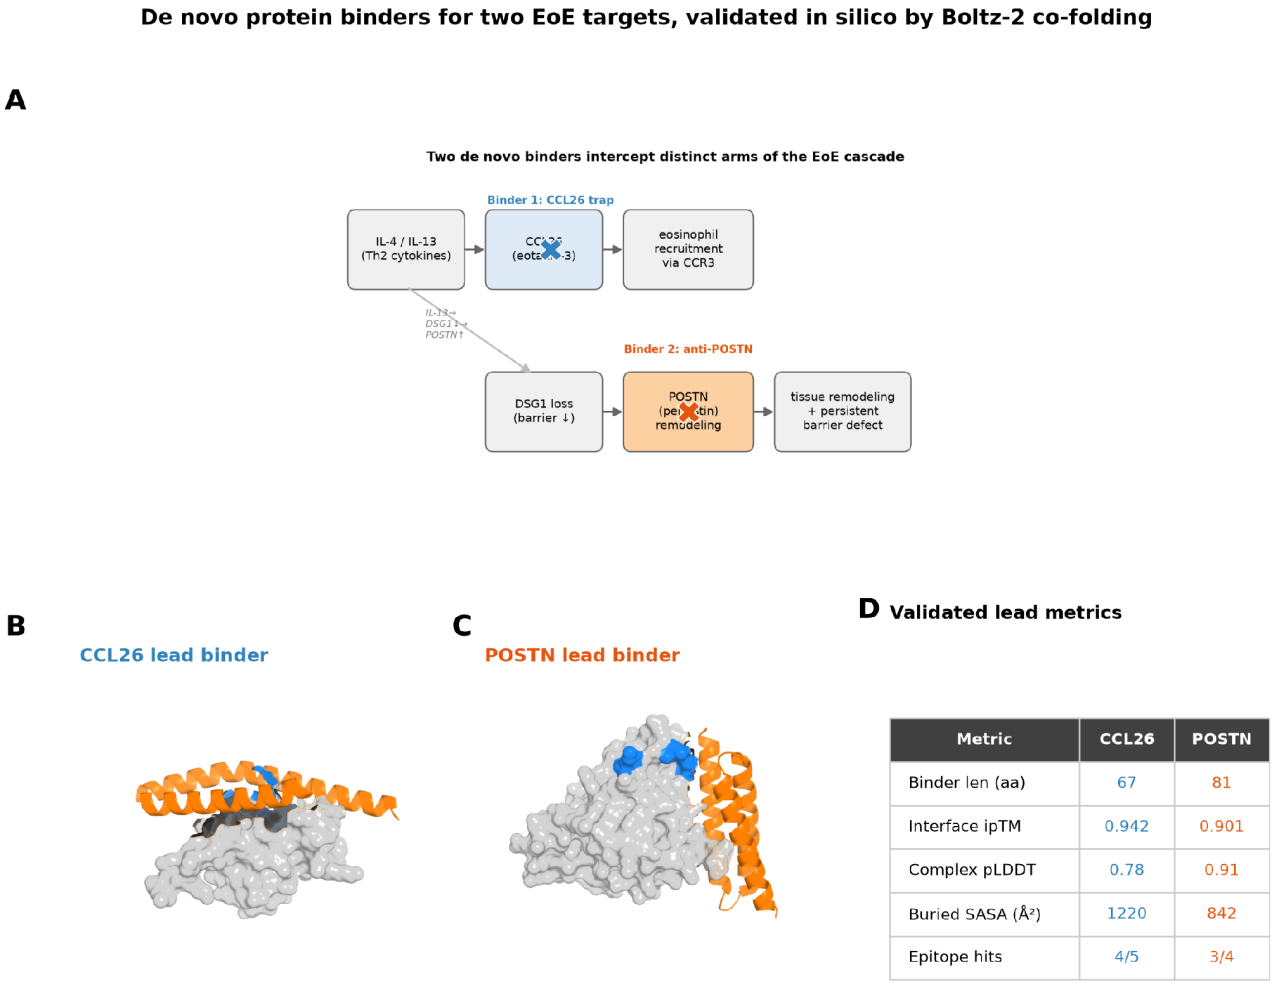

Metric	CCL26	POSTN
Binder len (aa)	67	81
Interface ipTM	0.942	0.901
Complex pLDDT	0.78	0.91
Buried SASA (Å²)	1220	842
Epitope hits	4/5	3/4

Figure 1. De novo protein binders for two EoE targets, assessed in silico by Boltz-2 co-folding with negative-control calibration and orthogonal-predictor concordance. (A) The EoE cascade with the two intervention points: a CCL26 trap blocks the eosinophil-recruitment arm (IL-4/IL-13 → CCL26 → CCR3), and an anti-POSTN binder blocks the remodeling arm (DSG1 loss → POSTN → tissue remodeling + persistent barrier defect), interrupting the IL-13→DSG1↓→POSTN↑ feed-forward loop. (B, C) Rendered lead binder–target complexes: target as grey surface, epitope hotspots in blue, de novo binder as orange cartoon. (D) Validated lead metrics.

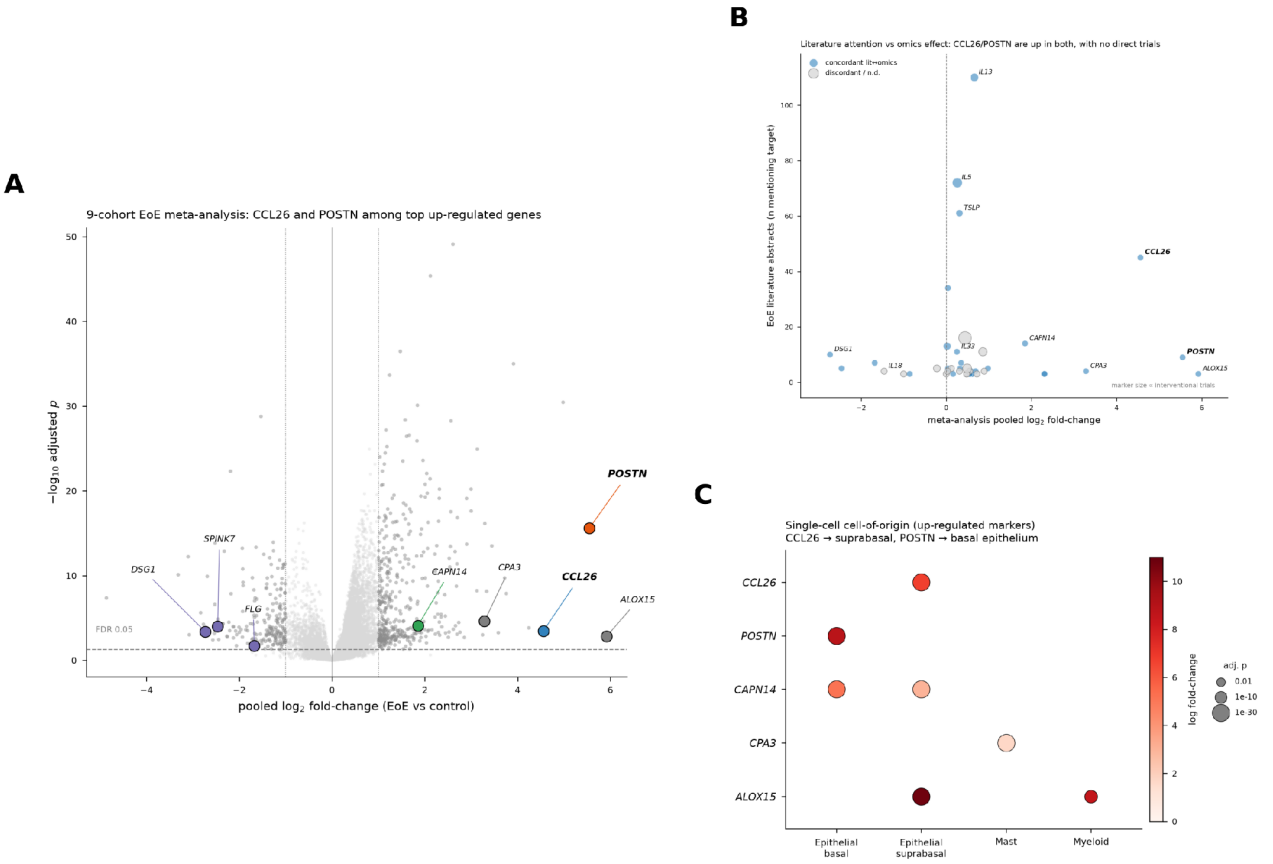

Figure 2. Target selection: convergent transcriptomic, literature, and single-cell evidence. (A) Nine-cohort meta-analysis volcano; CCL26 and POSTN (bold) are among the most strongly up-regulated genes, opposite concordantly down-regulated barrier genes (DSG1, SPINK7, FLG). (B) Literature attention vs. omics effect size, marker size proportional to interventional-trial count; CCL26/POSTN combine high omics effect with zero direct trials. (C) Single-cell cell-of-origin: CCL26 in suprabasal epithelium, POSTN in basal epithelium.

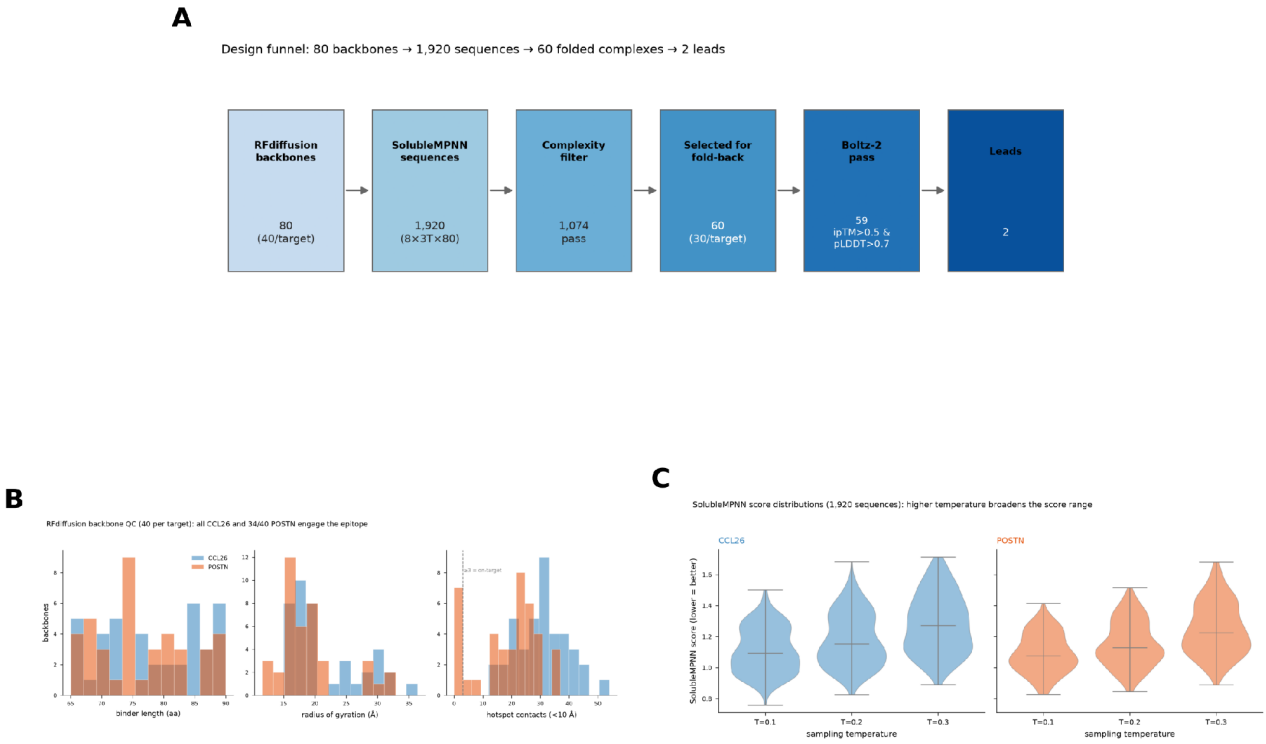

Figure 3. De novo design pipeline. (A) Design funnel from 80 RFdiffusion backbones through 1,920 SolubleMPNN sequences and complexity filtering to 60 folded complexes and 2 leads. (B) Backbone QC distributions (binder length, radius of gyration, hotspot contacts) for 40 backbones per target. (C) SolubleMPNN score distributions by sampling temperature.

Fold-back validation: interface confidence, epitope engagement, and on-target ranking select two leads

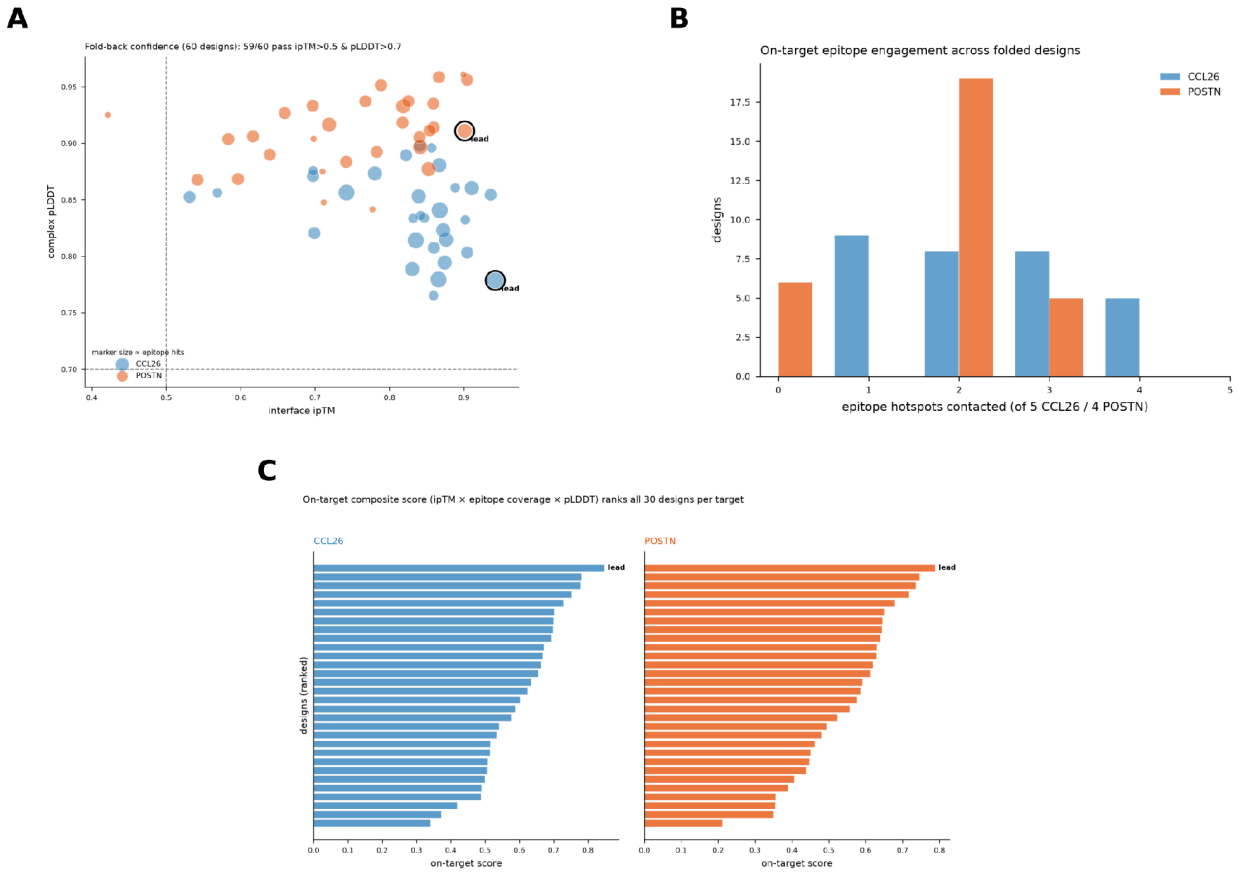

Figure 4. Fold-back confidence assessment of the 60 designs. (A) Interface ipTM vs. complex pLDDT for all 60 designs; 59/60 cross the confidence thresholds (dashed lines), leads circled, marker size proportional to epitope hotspot hits. (B) Distribution of epitope hotspots contacted per design. (C) On-target composite score ranking all 30 designs per target; leads marked. Confidence thresholds are interpreted against the negative-control calibration in Figure 5.

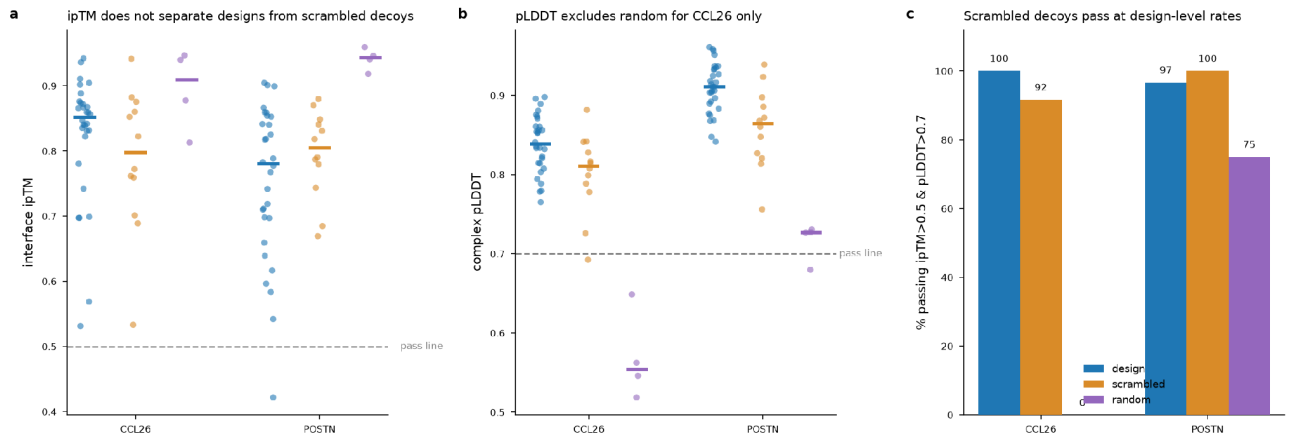

Figure 5. Negative-control calibration: interface confidence does not distinguish designs from scrambled decoys. 36 control complexes (per target: 12 composition-matched scrambled-sequence decoys, 4 random-composition decoys) co-folded under the identical Boltz-2 protocol as the designs. (A) Interface ipTM for designs vs. scrambled vs. random decoys; scrambled decoys reach design-level ipTM (Mann–Whitney design > scrambled n.s.: CCL26 p=0.16, POSTN p=0.80). (B) Complex pLDDT; the fold-confidence threshold excludes random-composition sequences for CCL26 but only weakly for POSTN, and does not exclude scrambled decoys for either. (C) Percentage crossing the ipTM>0.5 & pLDDT>0.7 gate: scrambled decoys pass at 92% (CCL26) / 100% (POSTN), indistinguishable from designs (100% / 97%). Horizontal ticks are medians. This calibration is why we report a threshold-crossing rate rather than a "validation" pass rate, and why focal epitope engagement (Fig. S10), not ipTM, is treated as the discriminating readout.

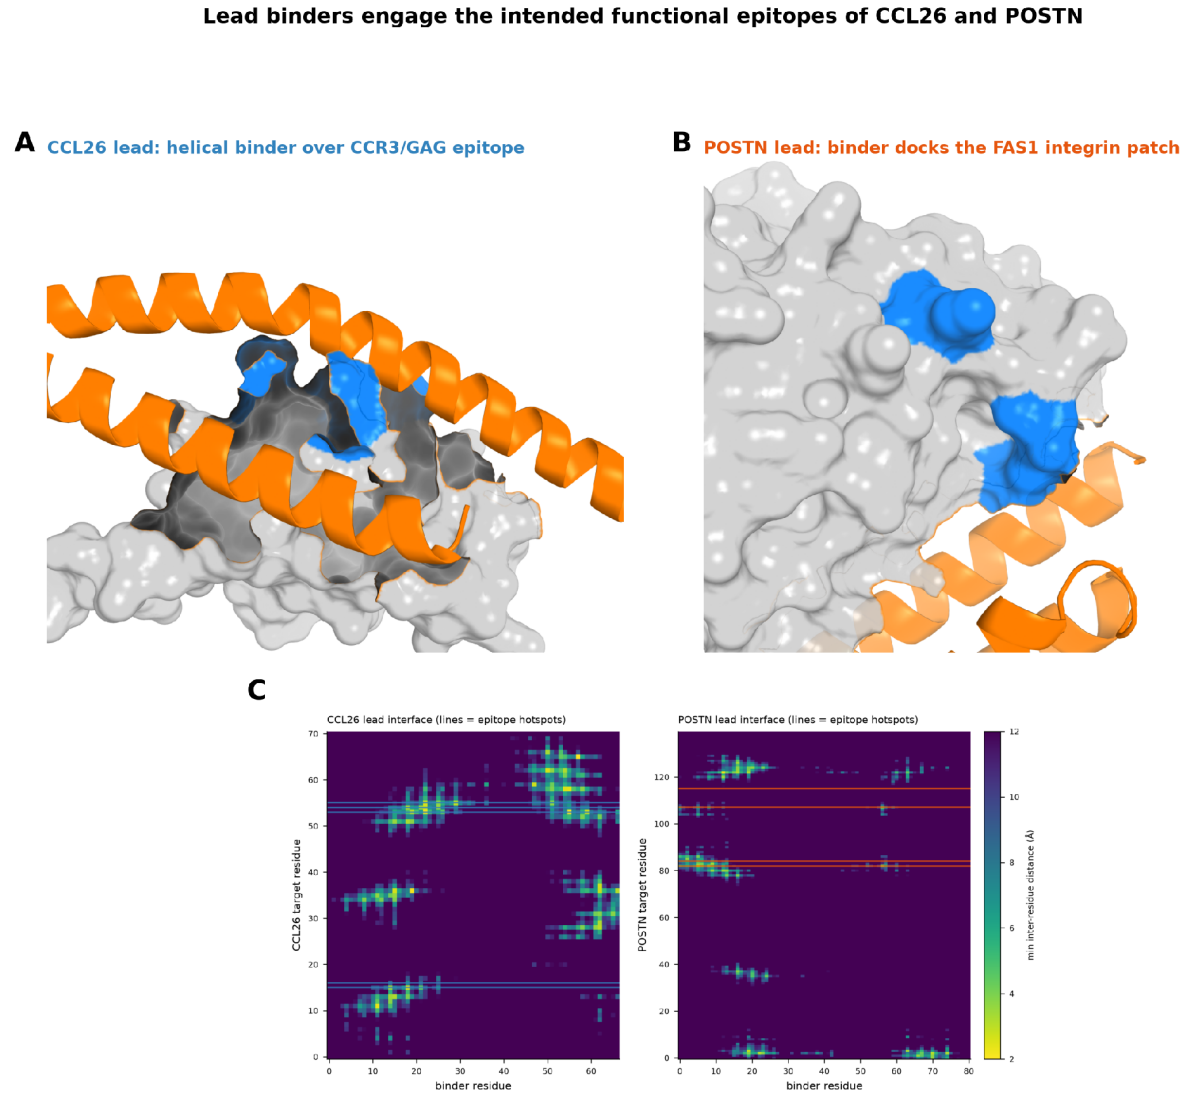

Figure 6. Lead binders engage the intended functional epitopes in silico. (A) CCL26 lead close-up: helical binder over the CCR3/GAG basic face (epitope hotspots blue). (B) POSTN lead close-up: binder docked onto the FAS1 integrin patch. (C) Residue-level interface distance maps; close-contact bands (yellow) align with epitope hotspot rows (colored lines).

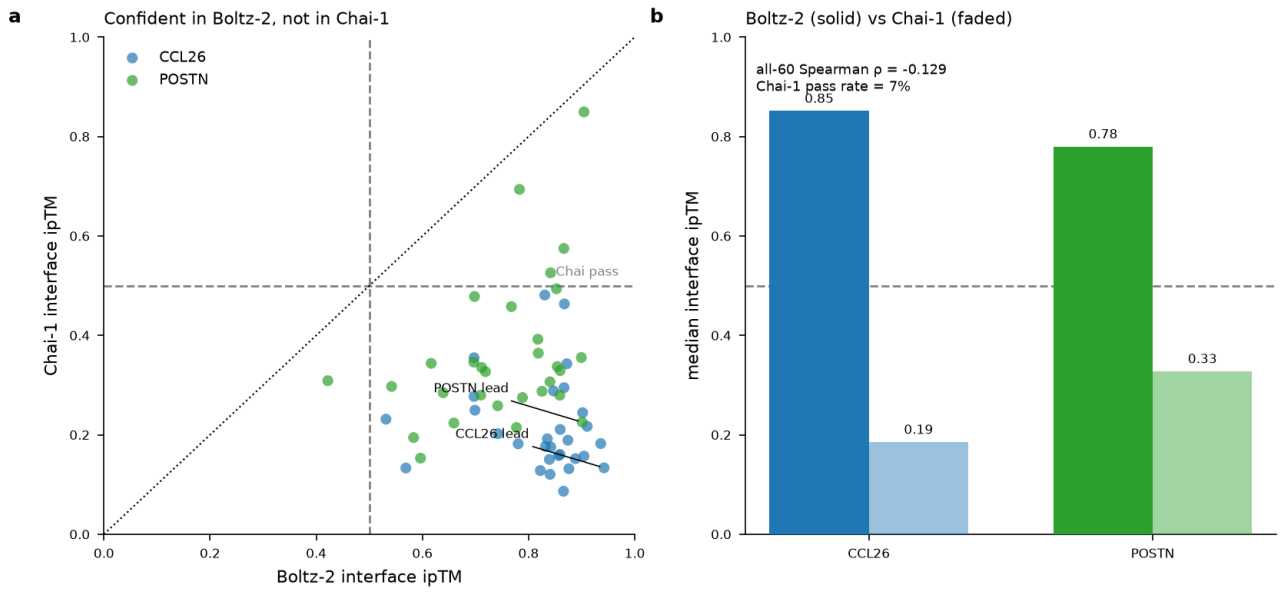

Figure 7. Orthogonal-predictor concordance: designs confident in Boltz-2 are not confident in Chai-1. All 60 designs re-folded with Chai-1 under matched settings. (A) Per-design interface ipTM, Boltz-2 vs. Chai-1; dashed lines mark the ipTM = 0.5 pass line, dotted line is identity. Nearly all designs sit in the lower-right quadrant (Boltz-confident, Chai-not), and both leads collapse in Chai-1. (B) Median interface ipTM by model (Boltz-2 solid, Chai-1 faded); overall Spearman rank correlation between the two models is −0.13 (n.s.) and only 7% of designs pass Chai-1's ipTM > 0.5 threshold.

Manuscript prepared as a hackathon deliverable (Built with Claude: Life Sciences). All designs are computational and require experimental validation.