

De novo design of cysteine proteases

Corresponding Author: Dr David Baker

Any redactions in this file are there to maintain patient confidentiality, the confidentiality of unpublished data, or to remove third-party material.

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Referee #1

(Remarks to the Author)

In this study, the authors use two-sided design to generate de novo cysteine proteases and their peptide substrates using the RFD2-MI diffusion framework. As the authors justly claim, protease activity is challenging due to the high energy barrier (stability of the peptide bond). Among 69 designed protease–substrate pairs, 13 exhibited detectable proteolytic activity in the tethered (cis) configuration, originating from three distinct structural backbones. The best-performing variants displayed modest kinetic parameters (k_{cat}/K_M of $12 \text{ M}^{-1} \text{ s}^{-1}$, k_{cat} 0.0011 s^{-1} , TTN 19). Analysis of the catalytic mechanisms (pK_a) of the active designs indicates that they do not generate the chemical environment conducive for efficient catalysis, with predicted pK_a shifts of the catalytic residues deviating by 2–4 pK_a units from those observed in natural cysteine proteases. Crystal structures of the designed enzymes reveal an overall accurate reproduction of the intended backbone. However, there are backbone deviations near the catalytic residues and substantial rotameric rearrangements. These structural differences may lead to the low activities.

Major comments

As in the companion paper, the kinetic parameters of the best designs are extremely low, with the additional concern regarding the total turnover number (TTN). In fact, the observed activities are significantly lower than those of catalytic (proteolytic) antibodies. The work would greatly benefit from a mechanistic analysis aimed at identifying the reasons behind these low rates. A plausible explanation could involve product inhibition or oxidation of the catalytic cysteine, as observed in a crystal structure.

For Gawi_09, the presence of a zinc ion in the crystal structure raises questions, given that zinc can act as a cofactor for protease activity. Perhaps conducting the functional assays in the presence of EDTA would refute this possibility. In the case of the best-performing design, Dokki_15, the crystallographic structure reveals a non-reactive conformation of the catalytic histidine, a deviation substantially greater than that observed in Extended Data Figure 12 for natural ULP1. Therefore, the statement in lines 211–212—“In our design structures, slight mispositioning of histidine may contribute to reduced catalytic efficiency”—understates the structural evidence. Taken together, the crystallographic data and the very low catalytic activity suggest insufficient accuracy or preorganization of the active site (although, please note Minor Comment #1. There may be confusion due to the naming schemes). As mentioned in the review of the companion paper, the low turnover numbers are suggestive of a scenario similar to the observations on catalytic antibodies, and structural validation of the active site is essential for the claims made in the paper on the efficacy of the design process.

Furthermore, there is insufficient data to assess the substrate selectivity of the designed proteases. Although 69 peptide substrates were generated, the selectivity assays evaluated only two (lines 161–162). As a result, the selectivity of the designed enzymes remains largely unknown and is crucial for understanding the efficacy of the design strategy.

Minor Comments

The naming scheme is confusing (HC vs Dokki/Gawi/Kal). In Fig. 3, it is unclear which crystal structures correspond to which scaffolds, and the design names HC34 and Dokki_15 appear to be used interchangeably across figures and text. Graphs lack explanations of bars and the number of replicates. For all gels showing cleavage, indicate reaction times and conditions; this information should also be added to lines 105–106 and the Methods.

Extended Data Figures 8 and 9 appear to show identical Michaelis–Menten analyses—please clarify or merge.

Provide a table summarizing all kinetic parameters; currently, only Gawi_09, Dokki_15, and their precursors are reported.

As in the companion paper, report design success relative to the trans assay using soluble peptide substrates as opposed to the fused version of the experiment.

Lines 121–122: specify which designs were incubated with the synthetic substrate—all 69, the 13 active in cis, or only Gawi and Dokki?

Referee #2

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports.

Referee #3

(Remarks to the Author)

This manuscript describes the computational two-sided design of cysteine proteases and their substrates through the use RFD2-MI, a recently introduced extension of RFdiffusion 2, reported in reference 2, with decent activities out of the box. The computational design methodology and the experimental validation approach is exactly the same as in the other manuscript I reviewed, about “de novo designed metalloproteases”. The manuscript is also structured in more or less the same way and also the discussion about the challenges and benefits of designed proteases is similar.

In the present work the catalytic machinery of the cysteine protease ulp1, i.e. catalytic triad and oxyanion hole, was used as input for the design. The success rate of active designs in this case was ~19%. Three enzymes, called Gawi, Kal, and Dokki, were selected and further investigated. Two of these were then subjected to a round of redesign because their initial kinetics were somewhat disappointing. This resulted in enzymes that showed up to 10^7 fold rate enhancement compared to the background reaction. This is impressive, but still $\sim 10^4$ fold less than for natural cysteine proteases and an order of magnitude less than the aforementioned designed zinc proteases from the other manuscript.

A bell shaped pH curve was obtained and it was found that the pKas of the active site cysteine and histidine were shifted, albeit to a different extent than in natural cysteine proteases.

The Dokki design showed no cross reactivity with the substrate of the Gawi design. However, the Gawi did show some preference for its designed substrate, it also showed high activity for the Dokki substrate.

Mutation experiments show that all residues of the catalytic machinery are required. But the decrease of activity of the residues from the oxyanion hole is only moderate, which may suggest the design was not optimal in this regard. Maybe the authors could discuss this point a bit more.

Structural characterization showed the accuracy of the design, albeit that there were some slight deviations observed in the positioning of the active site residues.

The overall methodologies and reporting are very high quality.

Overall: this is an impressive example of enzyme design, showing how computational design approaches can give completely novel designs of enzymes for challenging reactions such as peptide bond hydrolysis. But I already noted the similarity to the metalloprotease manuscript by the same groups. The metalloprotease manuscript seems a lot more fleshed out than this cysteine protease manuscript, and outperforms it with both scope and attention to detail, while the methods are the same or similar. For example, there is a less detailed investigation of substrate selectivity and no application to more biologically relevant reactions, although this prospect is raised in the discussion.

Finally, since both works and manuscripts are so similar, the general comments I made on the metalloproteases work also apply to the work described in this manuscript.

Referee #4

(Remarks to the Author)

In this study, Choi and colleagues apply a previously developed deep learning method, RFD2-MI, to design enzymes that use an activated cysteine nucleophile to cleave peptide bonds. The resulting cysteine proteases are built on entirely novel protein folds and exhibit rate enhancements of up to 10^7 relative to the uncatalyzed reaction. The crystal structures align closely with the computational models, underscoring the high accuracy of the structural design. However, the catalytic parameters, particularly k_{cat} , remain quite low, indicating that turnover rates are much below those of natural cysteine proteases. This highlights that, despite strong structural control, the method still lacks comparable precision in functional

design. Given that the computational approach was introduced previously (the paper in which RFD2-MI is developed is included as SI) and that the catalytic improvements achieved here are modest, we do not consider the results to meet the level of novelty and impact expected for publication in Nature.

Major points:

1. The authors tested the importance of the designed cysteine nucleophile, histidine base, and glutamine/asparagine oxyanion contact by mutating each residue individually to alanine. The authors mention that mutation of Cys or His to Ala abolished activity, however, the plots shown in Figure 2c (top panel) indicates that product formation is still observed for Gawi_09 (for Dokki_15 the scale is different, so it is much difficult to visually assess). The essential role of Cys/His is not confirmed in this Gawi_09 design.
2. The X-ray structures confirmed the structure of the designs with RMSD below 1.2 Å. However, the structural panels shown in Figure 2 and 3 should be improved by adding more information regarding key catalytic distances.
3. In some of the structures displacement of His is observed, which the authors assume it is caused by the presence of a Zn²⁺ ion present in the crystallization solution. However, in their discussion, when justifying the low catalytic constants of the designs, they hypothesized that this displacement of His is the major factor contributing to reduced catalytic efficiency, which they want to improve in future design efforts using alternative scaffolds. In our opinion, this is mandatory for publishing this paper in Nature.

The references are appropriate and the whole manuscript is well-written and easy to read.

Referee #5

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports.

Referee #6

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports.

Version 1:

Reviewer comments:

Referee #1

(Remarks to the Author)

Given the similarities between the papers, below we assess both together.

The authors have considerably revised the manuscripts, adding data on structures that exhibit high accuracy, mutations that improve catalytic rates, and molecular dynamics simulations that help explain the partial improvement in catalytic rates in mutants based on active-site preorganization. These are important additions.

Yet, even after optimization by systematic mutagenesis, the designed enzymes remain 2-3 orders of magnitude less active than the natural counterpart from which the active site was derived. This is the most relevant comparison, particularly because several recent studies have already demonstrated that RFDiffusion can produce de novo enzymes with low k_{cat}. We agree that the catalytic parameters of natural enzymes vary, but it is important to remember that much of this variation is due to trade-offs: TEV trades off rate for very high specificity, and the specificity tests carried out by the authors are far from sufficient to suggest they have achieved similar levels. We do not doubt that hydrolysing amides presents a greater challenge in terms of activation barriers than previous studies on hydrolases. But the methods are very similar to those of the previous studies and exactly the same in the two companion papers, limiting novelty, and questioning why the authors did not integrate them into one paper.

Furthermore, the gap from the natural enzymes from which the active site was derived limits confidence in the authors' conclusions that low rates in the final designs are due to low preorganization. Given that all recent de novo hydrolases have shown low rates, there may be a ceiling to the rates that are feasible using the current generation of de novo fold design methods. Identifying such missing factors may explain what makes natural enzymes so proficient, beyond the accuracy of the active site, and suggest ways to break the rate ceiling.

Regarding generality, except for one designed enzyme, all were generated to target custom-designed substrates (two-sided designs), and the design process relied on the simplifying assumption that the substrate associates with the enzyme through β -sheet interactions. It is not clear from the writing whether the targeted human protein (one-sided design) was chosen because of similarity to some of the previously designed substrates, adding another limitation to generality.

Given the concerns on rates and generality, the authors' assertions that the work opens the way to therapeutic and industrial applications are not tenable.

Referee #2

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports.

Referee #3

(Remarks to the Author)

Additional experimental work has been performed and included. A mutagenesis study was performed for 20 residues around the active site, which led to the discovery of a variant "Dokki-v2" that has four additional mutations compared to the designed Dokki-15 enzyme. The Dokki-v2 variant showed a 1-2 orders of magnitude increased k_{cat} , attributed to better preorganization of active site residues. The MM parameters of this improved variant are only 1 order of magnitude below that of the natural TEV protease, which is an impressive result. These results and their discussion was included in the manuscript.

I have the feeling the effect of the His rotamers is overemphasized in the manuscript. If in the Dokki-15 only 6-10% of the His occupies the non-productive rotameric states, i.e. >90% is in the correct "active" rotameric state, then it is hard to understand the more than an order of magnitude improvement of activity observed for Dokki-v2, which has 100% occupancy of His in the active rotameric state. I am convinced that the controlling the His rotameric states is an important contributor to the increased activity of Dokki-v2, but in combination with other structural effects. The authors do note this briefly in the manuscript, by mentioning "Gln65 oxyanion hole contacts with shorter, less variable 233 H-bond distances" as possible additional contributor. But the impression given in the manuscript is still very much that His rotameric states are the main cause of the increased activity, whereas in their rebuttal to reviewer 1 they note that "it does seem like at least part of the gap in activity was a result of imprecise positioning of the catalytic triad in Dokki-v2." I recommend to make the statements in the manuscript similarly more nuanced.

With regard to the response to my comments:

In my review, I noted the similarity of this manuscript to the companion manuscript on metalloprotease design: "The companion metalloprotease manuscript seemed a lot more fleshed out than this cysteine protease manuscript, and outperforms it with both scope and attention to detail, while the methods are the same or similar. For example, there is a less detailed investigation of substrate selectivity and no application to more biologically relevant reactions, although this prospect is raised in the discussion."

In their rebuttal, the authors do not really address this issue. They acknowledge the similarity in methods used, but emphasize the challenge of sub-angstrom design of the Cys-His-Asp catalytic triad that is addressed in this manuscript. While this is certainly true, the precise design of catalytic triads has been reported before by the Baker lab, e.g. for the design of Serine hydrolases (Science 2025, 10.1126/science.adu2454), using another approach. The current work, while impressive with regard to the results, does feel a bit as another, indeed "less fleshed out", example of the application of the same methodology as used in the companion metalloprotease study – even though the quality and impact of the work itself, in isolation, is undisputed.

With regard to the response to my comment 3, about why not all transition states in the design process are included as reported before for the serine hydrolases (Science 2025, 10.1126/science.adu2454): I understand the reasoning provided, i.e. to focus on the first tetrahedral intermediate because that is "considered to approximate the rate-limiting transition state for peptide bond hydrolysis", and the results show that this was a good choice. But it still feels a little bit like a step back. In the Science 2025 paper the importance of including all transition states in the design process was emphasized as key to achieving the most active designed enzymes and here this is only mentioned as a possibility to further improve the design. In itself this is fine, but it feels a bit unsatisfying.

A final note to the authors: a version of the manuscript with the changes highlighted, or at least a clear indication of where in the manuscript the changes were made would have been helpful.

Referee #4

(Remarks to the Author)

In this revised version of the manuscript, the authors provide computational evidence suggesting that incomplete active-site preorganization in the computational variants limits the activity of the first-generation designs. They further suggest that introducing mutations within the active-site pocket can improve the positioning of catalytic rotamers (see below) and thereby enhance the activity of the computational designs. However, as also noted for the related metalloprotease study, much of this additional improvement appears to have been achieved experimentally through a brute-force approach, involving the testing of 185 single-site variants at a selected set of positions within the active-site pocket (Tyr41, Ser44, Pro67, and Thr99).

A particularly concerning aspect is that none of these positions were sampled in the original design calculations because ProteinMPNN assigned them low probabilities. This raises questions about the reliability of the computational pipeline, since positions that ultimately proved important for improving active-site preorganization were not prioritized by

ProteinMPNN. Another important concern is that the mechanistic basis by which these mutations enhance active-site preorganization remains inconclusive. Much of the discussion focuses on the different side-chain rotamers of the catalytic histidine as predicted by PLACER, ranging from 100% in Ulp1 to 90% in Dokki, 94% in Dokki-15—where, notably, two conformations are also predicted for the catalytic Asp—and 100% in Dokki-v2. However, analysis of the catalytic Cys-His distance shows that this distance progressively increases along the Ulp1 (3.85 Å) < Dokki < Dokki-15 < Dokki-v2 (4.2 Å) series, which does not support the conclusion of improved preorganization of the catalytic Cys-His dyad. Likewise, the discussion of how the Y41T and S44V mutations increase activity remains largely speculative, as does the proposed effect of P67V and T99F on oxyanion-hole preorganization.

Overall, although this revised version represents an improvement over the original submission, we do not believe that the manuscript is suitable for publication in Nature.

Referee #5

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports.

Referee #6

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports.

Version 3:

Reviewer comments:

Referee #1

(Remarks to the Author)

The current revision adds some computational analyses but not substantial new data to counter the criticisms of the previous reviews. The authors' assertion that the low K_{cat} is true also for TEV protease is, as noted in a previous review cycle, not justifiable, as TEV protease exhibits extremely high selectivity, which is why it is so useful in molecular biology. Such levels of specificity were not demonstrated for the current work.

The authors' computational analyses suggest that the performance gap relative to comparable natural enzymes stems from insufficient preorganization and incomplete substrate positioning. This is an almost self-evident conclusion that appears in many enzyme design studies. The critical test is whether the study can address these deficiencies and improve the enzymes, but this is not shown here.

We previously noted that the methods are not sufficiently distinct from previous studies by the Baker lab. The authors explain that they implemented a different theozyme in this case, but this is not a substantial methodological difference. Given that in this and previous studies, the catalytic rates were low, and this study does not address this fundamental problem nor propose a substantially new approach, we do not see sufficient merit for publication in Nature.

Referee #7

(Remarks to the Author)

- HC12: Overall the details are right and well explained in the manuscript. A minor detail: Gly91 adjacent to Cys94 (catalytic residue) has an unmodeled peptide flip; this may be related to the small difference density accumulating on Cys94, which may be indicative of the start of radiation damage. The related structure Kal C103A complex_D_1000301017_model-annotate_P1 was supplied but not mentioned in main text the manuscript. It is the same complex with the Cys mutated and indeed has the peptide flip as the main/only conformation.

- Gawi_09: Overall I agree with the interpretation. Only one chain has alt conf for catalytic Cys, justified by slightly more pronounced density than the other chain. Overall the active site density features are very similar between the two chains, further supporting the interpretation and the Zn displacement hypothesis.

- HC34/Dokki WT apo_D_1000301020_model-annotate_P1 looks more like Dokki_15 C100A complex_D_1000300811_model-annotate_P1. This despite the cysteine sulfonic acid on HC34. This suggests that the larger His shift of the Dokki_15_WT structure may be due to the absence of substrate, not just the damaged Cys. Furthermore, the absence of the substrate and the Cys sulfonic are accompanied by not just a side chain shift. Rather, the whole beta hairpin His45 sits on shifts. None of this materially changes any conclusions of the paper or even this section; I would just would reword as I think it is a bit more complicated than what they are explaining here.

- o For the Cys sulfonic, the authors may want to speculate whether this is caused by radiation damage or accumulates during the purification/crystallization. There may be ways to salvage the dataset and resolve this if it is radiation damage, which can address some of the issues they are seeing. For example, reprocessing the first half of the data, etc.

James Fraser (UCSF)

Open Access This Peer Review File is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source.

The images or other third party material in this Peer Review File are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

Responses are unitalicized below and interspersed with the reviewer comments that are italicized and underlined.

Reviewer 1:

In this study, the authors use two-sided design to generate de novo cysteine proteases and their peptide substrates using the RFD2-MI diffusion framework. As the authors justly claim, protease activity is challenging due to the high energy barrier (stability of the peptide bond). Among 69 designed protease–substrate pairs, 13 exhibited detectable proteolytic activity in the tethered (cis) configuration, originating from three distinct structural backbones. The best-performing variants displayed modest kinetic parameters (k_{cat}/K_M of $12 \text{ M}^{-1} \text{ s}^{-1}$, k_{cat} 0.0011 s^{-1} , TTN 19). Analysis of the catalytic mechanisms (pK_a) of the active designs indicates that they do not generate the chemical environment conducive for efficient catalysis, with predicted pK_a shifts of the catalytic residues deviating by 2–4 pK_a units from those observed in natural cysteine proteases. Crystal structures of the designed enzymes reveal an overall accurate reproduction of the intended backbone. However, there are backbone deviations near the catalytic residues and substantial rotameric rearrangements. These structural differences may lead to the low activities.

As in the companion paper, the kinetic parameters of the best designs are extremely low, with the additional concern regarding the total turnover number (TTN). In fact, the observed activities are significantly lower than those of catalytic (proteolytic) antibodies. The work would greatly benefit from a mechanistic analysis aimed at identifying the reasons behind these low rates. A plausible explanation could involve product inhibition or oxidation of the catalytic cysteine, as observed in a crystal structure.

The referee raises important concerns regarding the kinetic parameters and total turnover number (TTN) of the first-generation designs. In the revised manuscript, we have substantially strengthened both the mechanistic and functional analyses. Most notably, we report an improved variant, Dokki-v2, that achieves a catalytic efficiency (k_{cat}/K_M) of $413 \pm 50 \text{ M}^{-1} \text{ s}^{-1}$ ($k_{\text{cat}} = 0.064 \pm 0.005 \text{ s}^{-1}$; $K_M = 154 \pm 14 \text{ }\mu\text{M}$) and a TTN > 130. This performance approaches that of native cysteine proteases, within ~2.5-fold of the widely used TEV protease in turnover number and within one order of magnitude in catalytic efficiency. We further provide computational evidence that suggests that incomplete active-site preorganization limits activity in the first-generation designs and that beneficial substitutions surrounding the active site improve positioning of the catalytic rotamers.

As noted in our previous rebuttal, many reported catalytic antibodies either catalyze chemically distinct and more facile reactions, rely on activated substrates, or lack comprehensive biochemical and structural characterization. In contrast, our designs employ a fully engineered Cys–His–Asp catalytic triad embedded within a de novo protein scaffold and catalyze unactivated peptide bond hydrolysis.

For Gawi_09, the presence of a zinc ion in the crystal structure raises questions, given that zinc can act as a cofactor for protease activity. Perhaps conducting the functional assays in the presence of EDTA would refute this possibility.

The presence of a zinc ion in the crystal structure reflects the high zinc acetate concentration (50 mM) in the crystallization buffer.

In the case of the best-performing design, Dokki_15, the crystallographic structure reveals a non-reactive conformation of the catalytic histidine, a deviation substantially greater than that observed in Extended Data Figure 12 for natural ULP1. Therefore, the statement in lines 211–212—“In our design structures, slight mispositioning of histidine may contribute to reduced catalytic efficiency”—understates the structural evidence.

Taken together, the crystallographic data and the very low catalytic activity suggest insufficient accuracy or preorganization of the active site (although, please note Minor Comment #1. There may be confusion due to the naming schemes). As mentioned in the review of the companion paper, the low turnover numbers are suggestive of a scenario similar to the observations on catalytic antibodies, and structural validation of the active site is essential for the claims made in the paper on the efficacy of the design process.

As noted by the reviewer, the crystallographic structure reveals a non-reactive conformation of the catalytic histidine that deviates substantially from the arrangement observed in natural ULP1. In the revised manuscript, we address the structural origins of the reduced catalytic efficiency and provide experimental optimization and computational analysis to understand the activity gap between designed/native. We investigated substitutions surrounding the active site and identified an improved variant, Dokki-v2 (Y41T:S44V:P67V:T99F). We hypothesize that the mutations function to improve the accuracy and preorganization of the catalytic residues since the substitutions Y41T and S44V are both β -branched amino acids located on β -strands that flank the catalytic histidine, likely helping to eliminate the non-productive rotamers observed in the crystal structures of Dokki-15.

To test this hypothesis, we analyzed active-site ensembles using PLACER, focusing on the first tetrahedral intermediate. In contrast to Dokki and Dokki-15, which populate off-target histidine rotamers, PLACER ensembles of Dokki-v2 exhibit complete preorganization of the catalytic His and Asp residues in their functional conformations and are comparable to native Ulp1. The beneficial substitutions Y41T and S44V do appear to eliminate off-target histidine conformations in Dokki and Dokki-15, enforcing the productive catalytic rotamer in Dokki-v2 (Fig. 3c,d and 4f; Extended Data Fig. 9d). In addition, the oxyanion contact (Gln65) becomes more tightly organized and less heterogeneous in response to the P67V and T99F mutations. Taken together, it does seem like at least part of the gap in activity was a result of imprecise positioning of the catalytic triad in Dokki-v2.

Furthermore, there is insufficient data to assess the substrate selectivity of the designed proteases. Although 69 peptide substrates were generated, the selectivity assays evaluated only two (lines 161–162). As a result, the selectivity of the designed enzymes remains largely unknown and is crucial for understanding the efficacy of the design strategy.

In this study, we co-designed the substrate together with each protease scaffold, and from the third round of the pipeline onward the substrate sequence was fixed for each backbone to simplify and standardize functional testing. As a result, a total of 12 distinct peptide substrates were generated and associated with specific scaffolds.

In the revised manuscript, we further examined the specificity of Dokki-15 using a panel of eGFP–peptide–mScarlet fusion substrates comprising peptides derived from multiple scaffold designs. In total, 6 non-cognate peptide substrates with its cognate peptide substrate were tested. These peptides share a similar amphiphilic sequence pattern but each contains a distinct sequence at the cleavage site. Dokki-15 exhibited high selectivity for its cognate substrate (VPTLTGG/G) and showed no detectable cleavage of the seven closely related non-cognate substrates (Fig. 2h). These non-cognate peptides are designed substrates from other scaffolds and contain hydrophobic residues in positions corresponding to the P4,P6 residues. Despite the shared amphiphilic patterning, Dokki-15 did not cleave these sequences, demonstrating that substrate recognition extends beyond general physicochemical patterning and reflects sequence-specific interactions.

Minor Comments

The naming scheme is confusing (HC vs Dokki/Gawi/Kal). In Fig. 3, it is unclear which crystal structures correspond to which scaffolds, and the design names HC34 and Dokki_15 appear to be used interchangeably across figures and text.

We apologize for the confusion caused by the naming scheme. The HC identifiers (hc04, hc12, hc34) refer to the original design IDs generated during the initial computational campaign. Following experimental screening, the most active design from each successful backbone was selected and assigned a distinct name: hc04 → Gawi, hc12 → Kal, and hc34 → Dokki. We have clarified this in the revised manuscript.

Graphs lack explanations of bars and the number of replicates. For all gels showing cleavage, indicate reaction times and conditions; this information should also be added to lines 105–106 and the Methods.

All FRET assays were performed in three independent replicates, with bars in the graphs indicating the mean $\pm$ s.d. For SDS–PAGE gels depicting cis-cleavage, reactions correspond to protease expression for 16 hours under the indicated conditions. We now explicitly describe the number of replicates, bar definitions, reaction times, and conditions to the figure legends, and have expanded the Methods section accordingly.

Extended Data Figures 8 and 9 appear to show identical Michaelis–Menten analyses—please clarify or merge.

To eliminate confusion, we have removed the mistakenly duplicated Michaelis–Menten plots from Extended Data Figures 8 and 9 (now Extended Data Figures 6), as the full Michaelis–Menten analysis is already presented in Figure 2. The Extended Data figures now retain only the corresponding linear representations.

In addition, we have streamlined the Extended Data section by consolidating figures as follows: previous Extended Data Figures 2 and 3 are now merged as Extended Data Figure 2; Figures 4 and 5 are merged as Extended Data Figure 3; Figures 6 and 11 are merged as Extended Data Figure 4; Figures 7 and 8 are merged as Extended Data Figure 5; Figures 9 and 10 are merged as Extended Data Figure 6; and Figures 12 and 13 are merged as Extended Data Figure 10.

Provide a table summarizing all kinetic parameters; currently, only Gawi_09, Dokki_15, and their precursors are reported.

We have added Supplementary Table 5 summarizing the kinetic parameters for all characterized designs, namely Gawi_09, Dokki, Dokki_15, and Dokki-v2. As noted in the manuscript and above, Kal could not be reliably characterized owing to poor substrate solubility. We measured kinetics for the most active variant identified for each structural scaffold.

As in the companion paper, report design success relative to the trans assay using soluble peptide substrates as opposed to the fused version of the experiment. Lines 121–122: specify which designs were incubated with the synthetic substrate—all 69, the 13 active in cis, or only Gawi and Dokki?

Only the 13 designs that exhibited detectable cis-cleavage activity were subsequently evaluated for trans-cleavage using the cognate synthetic peptide substrates. Among these, the hc10-12 could not be reliably assessed owing to limited substrate solubility; therefore, only hc4-10 and hc32-38 were characterized for trans-cleavage activity (Extended Data Fig. 4a). We have clarified this in the revised manuscript.

Reviewer 3:

This manuscript describes the computational two-sided design of cysteine proteases and their substrates through the use RFD2-MI, a recently introduced extension of RFdiffusion 2, reported in reference 2, with decent activities out of the box. The computational design methodology and the experimental validation approach is exactly the same as in the other manuscript I reviewed, about “de novo designed metalloproteases”. The manuscript is also structured in more or less the same way and also the discussion about the challenges and benefits of designed proteases is similar.

In the present work the catalytic machinery of the cysteine protease ulp1, i.e. catalytic triad and oxyanion hole, was used as input for the design. The success rate of active designs in this case was ~19%. Three enzymes, called Gawi, Kal, and Dokki, were selected and further investigated. Two of these were then subjected to a round of redesign because their initial kinetics were somewhat disappointing. This resulted in enzymes that showed up to 10^7 fold rate enhancement compared to the background reaction. This is impressive, but still $\sim 10^4$ fold less than for natural cysteine proteases and an order of magnitude less than the aforementioned designed zinc proteases from the other manuscript.

As detailed in our response to Reviewer 1, the revised manuscript now includes additional data describing an improved variant, Dokki-v2, that approaches the activity of the native cysteine protease TEV, reaching within 2.5-fold in turnover number and within one order of magnitude in catalytic efficiency.

A bell shaped pH curve was obtained and it was found that the pKas of the active site cysteine and histidine were shifted, albeit to a different extent than in natural cysteine proteases.

In many cysteine proteases the resting state is often described as a thiolate–imidazolium pair (deprotonated Cys / protonated His), and multiple studies support substantial perturbation of dyad protonation relative to free amino acids. For example, constant-pH MD across cysteine proteases has predicted low Cys pKa and high His pKa consistent with stabilization of the ion pair, and NMR/mutational analyses in papain-family enzymes demonstrate strong coupling between the catalytic Cys and His ionizations. At the same time, pH-rate analyses in other systems (e.g., cruzain) have yielded different apparent assignments depending on enzyme state and assay readout, emphasizing that dyad protonation is context-dependent (Zhai *et al.*, 2018).

The Dokki design showed no cross reactivity with the substrate of the Gawi design. However, the Gawi did show some preference for its designed substrate, it also showed high activity for the Dokki substrate.

We further expanded the substrate specificity analysis using GFP–mScarlet paired substrates, as synthesis of additional DABCYL–EDANS FRET peptides is time-intensive. However, the lower FRET efficiency of the GFP–mScarlet substrate pair relative to the DABCYL–EDANS peptide substrates, combined with the relatively low activity of Gawi, reduced assay sensitivity and made accurate quantification of Gawi’s specificity challenging. Accordingly, we focused our detailed selectivity analysis on Dokki-15, as described in our response to Reviewer 1.

Mutation experiments show that all residues of the catalytic machinery are required. But the decrease of activity of the residues from the oxyanion hole is only moderate, which may suggest the design was not optimal in this regard. Maybe the authors could discuss this point a bit more.

We added to the text a discussion on potential shortcomings of the oxyanion hole geometry, but note that similar trends are observed in natural enzymes. For example, the impact of mutating the sidechain asparagine oxyanion contact in subtilisin is much lower than that of residues in the catalytic triad. (Bryan *et al.*, 1986). This reduced sensitivity is thought to arise because oxyanion stabilization can be partially compensated by backbone amides and/or structured water molecules in the active site. In addition to this, we have more extensively characterized Dokki by site-saturation mutagenesis and found mutations flanking the oxyanion hole and catalytic histidine that appear to improve activity by improving the positioning and/or preorganization of these key functional motifs, and we discuss this in detail in the main text.

Structural characterization showed the accuracy of the design, albeit that there were some slight deviations observed in the positioning of the active site residues. The overall methodologies and reporting are very high quality.

Overall: this is an impressive example of enzyme design, showing how computational design approaches can give completely novel designs of enzymes for challenging reactions such as peptide bond hydrolysis. But I already noted the similarity to the metalloprotease manuscript by the same groups. The metalloprotease manuscript seems a lot more fleshed out than this cysteine protease manuscript, and outperforms it with both scope and attention to detail, while the methods are the same or similar. For example, there is a less detailed investigation of substrate selectivity and no application to more biologically relevant reactions, although this prospect is raised in the discussion.

We appreciate the reviewer’s acknowledgement that the protease design work described here is impressive. Although both studies employ related computational frameworks, the catalytic problems addressed are fundamentally distinct. Metalloproteases benefit from a preorganized

Lewis acidic metal center that directly polarizes the substrate and stabilizes the tetrahedral intermediate and proceeds through a single-step mechanism. In contrast, cysteine proteases rely on a precisely arranged Cys–His–Asp catalytic triad that must coordinate nucleophile activation, proton transfer, and oxyanion stabilization through a multi-step proton-coupled mechanism that is highly sensitive to sub-angstrom geometric deviations, oxidation state, and active-site electrostatics, all without the assistance of a reactive metal ion.

Furthermore, while natural metalloproteases often exhibit very high turnover numbers, many natural cysteine proteases display more modest turnover numbers and depend critically on precise active-site preorganization. In this context, bringing a fully designed cysteine protease within ~2.5-fold of the turnover number of a widely used native enzyme (TEV) represents a substantial achievement. As noted above, we have also expanded the substrate selectivity analysis in the revised manuscript.

Finally, since both works and manuscripts are so similar, the general comments I made on the metalloproteases work also apply to the work described in this manuscript.

1. The computational design techniques used, RFD2-MI, as well as the design of a zinc hydrolase were established by the same group in recent work. So, the main novelty here is the application towards the creation of metalloproteases and achieving peptide bond hydrolysis.

Although RFD2-MI was previously introduced and applied to post-translational modification (PTM) binder design, it had not been used for enzyme creation, despite being trained on enzyme interactions. In contrast, the present study establishes a substantially expanded capability on enzymes: the de novo design of cysteine proteases that catalyze hydrolysis of unactivated peptide bonds—one of the most chemically demanding reactions in biology.

2. This is maybe considered semantics, but I am not sure if the claim of “de novo design”, i.e. design from first principles, of metalloproteases is warranted here. The active site structure, i.e. the ligation of the catalytic metal, as well as the positioning of the oxy anion stabilizing residues was not designed, but taken from an existing metalloproteases, aminopeptidase N and astacin, and precisely reproduced in the designed proteins. So, one could argue that while the protein environment was designed de novo, the part of the structure that makes it an enzyme, the catalytic site, was not and, thus, that this work represents a (very impressive) redesign of aminopeptidase N and/or astacin. I think calling it a “designer enzyme” or “designer metalloprotease” is probably better justified than calling it a “de novo designed enzyme/metalloprotease”.

While the catalytic motif was taken from a natural protease, the surrounding protein scaffold was generated entirely de novo and does not resemble any known natural fold (TM-score < 0.5). Thus,

although the catalytic motif was transferred, the global architecture and structural context in which it operates are novel.

In typical enzyme redesign studies, an existing natural scaffold is retained and modified through sequence optimization or localized structural adjustments. In contrast, our approach embeds a functional catalytic motif into a completely new protein framework that lacks structural homology to known enzymes. We therefore believe that the term “de novo designed protease” accurately reflects the novelty of the scaffold and its integration with the catalytic site.

3. The design was modeled in three states: apo, substrate bound and transition state. This kind of multistate design is clearly important to achieve proficient enzymes. However, only one transition state (TS) was modeled, whereas there are more along the reaction coordinate. The Baker lab has shown in recent work (Science 2025, <https://doi.org/10.1126/science.adu2454>) that including all TS's along the reaction coordinate in the design improves the quality of the outcome. The authors should comment on why this approach was not used here; it might have delivered enzymes that are even closer to the catalytic activity observed for natural metalloproteases.

In the present work, we focused on the first tetrahedral intermediate, which is widely considered to approximate the rate-limiting transition state for peptide bond hydrolysis. While inclusion of additional transition states may further refine catalytic efficiency, our PLACER ensemble analysis centered on this mechanistically critical state. As discussed above, the revised manuscript demonstrates improved geometric preorganization of key catalytic residues at this dominant rate-limiting step, providing a mechanistic explanation for the substantial rate enhancement observed in Dokki-v2 relative to the first-generation designs.

Referee 4:

In this study, Choi and colleagues apply a previously developed deep learning method, RFD2-MI, to design enzymes that use an activated cysteine nucleophile to cleave peptide bonds. The resulting cysteine proteases are built on entirely novel protein folds and exhibit rate enhancements of up to 10^7 relative to the uncatalyzed reaction. The crystal structures align closely with the computational models, underscoring the high accuracy of the structural design. However, the catalytic parameters, particularly k_{cat} , remain quite low, indicating that turnover rates are much below those of natural cysteine proteases. This highlights that, despite strong structural control, the method still lacks comparable precision in functional design. Given that the computational approach was introduced previously (the paper in which RFD2-MI is developed is included as SI) and that the catalytic improvements achieved here are modest, we do not consider the results to meet the level of novelty and impact expected for publication in Nature.

As noted above in response to Reviewer 3, although RFD2-MI was previously applied to PTM binder design, it had not been used for enzyme creation despite being trained on enzyme interactions. Here, we demonstrate its first use for the de novo design of cysteine proteases that catalyze unactivated peptide bond hydrolysis—one of the most challenging reactions in biology. And as detailed above in response to Reviewer 1, the revised manuscript improves the functional limitations of the first-generation designs by identifying incomplete active-site preorganization as the primary constraint and demonstrating that simple substitutions improve catalytic geometry and substantially enhance turnover, bringing catalytic rates within ~2.5-fold of the widely used native cysteine protease TEV.

Major points:

1. The authors tested the importance of the designed cysteine nucleophile, histidine base, and glutamine/asparagine oxyanion contact by mutating each residue individually to alanine. The authors mention that mutation of Cys or His to Ala abolished activity, however, the plots shown in Figure 2c (top panel) indicates that product formation is still observed for Gawi_09 (for Dokki_15 the scale is different, so it is much difficult to visually assess). The essential role of Cys/His is not confirmed in this Gawi_09 design.

It is not uncommon to see only a reduction in catalytic activity and not complete ablation. Even for native enzymes this occurs. We will change the language from “abolished” to “significantly reduced”, which still suggests that the catalytic site is involved in catalysis and on-mechanism.

In the revised manuscript we now also include data for an improved variant, Dokki-v2, in which the catalytic knockouts are substantially clearer. Mutation of the Cys, His, or Asp to alanine results in activity indistinguishable from background (Fig. 4c).

2. The X-ray structures confirmed the structure of the designs with RMSD below 1.2 Å. However, the structural panels shown in Figure 2 and 3 should be improved by adding more information regarding key catalytic distances.

We agree that clearer depiction of key catalytic geometries would strengthen Figures 2 and 3. However, interpretation of catalytic distances directly from the crystal structures is complicated by artifacts, including oxidation of the catalytic cysteine and coordination of Zn^{2+} from the crystallization buffer, both of which perturb the active-site geometry. To more accurately assess catalytically relevant parameters, we therefore analyzed structural models of the first tetrahedral intermediate and quantified key distances, angles, and dihedral angles. These geometrical features, which correlate with catalytic activity, are presented in Extended Data Fig. 10.

3. In some of the structures displacement of His is observed, which the authors assume it is caused by the presence of a Zn^{2+} ion present in the crystallization solution. However, in their discussion, when justifying the low catalytic constants of the designs, they hypothesized that this displacement of His is the major factor contributing to reduced catalytic efficiency, which they want to improve in future design efforts using alternative scaffolds. In our opinion, this is mandatory for publishing this paper in Nature.

We now provide multiple lines of evidence that incomplete active-site preorganization limits turnover in the first-generation designs. Although crystallography revealed non-productive histidine conformations, and while those observed off-target conformations were potentially induced by Zn^{2+} and/or cysteine oxidation, imprecise histidine preorganization more generally could not be ruled out. New experiments and computational analysis suggest that substitutions around the active-site may help preorganize the catalytic machinery, as PLACER ensemble analysis revealed that the optimized variant, Dokki-v2, exhibits a more tightly clustered ensemble of productive histidine and oxyanion sidechain rotamers. These improvements in catalytic geometry and preorganization correlate with the observed increase in turnover, indicating that incomplete catalytic preorganization is likely at least partly responsible for the modest activity of the initial designs.

The references are appropriate and the whole manuscript is well-written and easy to read.

We appreciate the reviewer's positive assessment that the references are appropriate, and the manuscript is well written and easy to read.

Responses to Reviewers' Comments on the Cysteine protease Manuscript

The reviewers' comments are shown in blue italics, and our responses are shown in regular black font.

Reviewer #1

Given the similarities between the papers, below we assess both together.

The authors have considerably revised the manuscripts, adding data on structures that exhibit high accuracy, mutations that improve catalytic rates, and molecular dynamics simulations that help explain the partial improvement in catalytic rates in mutants based on active-site preorganization. These are important additions.

We thank Reviewer #1 for these constructive suggestions, which prompted additions that have substantially strengthened the manuscript. We are grateful for the opportunity to incorporate these analyses and experiments, and believe they have improved both the depth and clarity of the work.

Yet, even after optimization by systematic mutagenesis, the designed enzymes remain 2-3 orders of magnitude less active than the natural counterpart from which the active site was derived. This is the most relevant comparison, particularly because several recent studies have already demonstrated that RFDiffusion can produce de novo enzymes with low k_{cat} . We agree that the catalytic parameters of natural enzymes vary, but it is important to remember that much of this variation is due to trade-offs: TEV trades off rate for very high specificity, and the specificity tests carried out by the authors are far from sufficient to suggest they have achieved similar levels. We do not doubt that hydrolysing amides presents a greater challenge in terms of activation barriers than previous studies on hydrolases.

We agree that comparison to native enzymes provides an important benchmark, and we have revised the manuscript to clarify this point. To contextualize our designs, we compared the turnover number of Dokki-v2 with k_{cat} values for representative sequence-specific cysteine proteases, including Ulp1, Mpro, TEV and HRV3C, together with their cognate substrates (Fig. 5a and Supplementary Table 5). This analysis places Dokki-v2 at the lower end of the range observed for native sequence-specific cysteine proteases, approximately 1-2 orders of magnitude below the most active enzymes.

Despite this gap, Dokki-v2 achieves turnover within the lower native range for the challenging hydrolysis of a peptide amide bond. We view this as an important advance, particularly as previous RFDiffusion-based enzyme-design efforts have targeted more activated substrates, whereas peptide-bond hydrolysis represents a long-standing challenge for enzyme design.

But the methods are very similar to those of the previous studies and exactly the same in the two companion papers, limiting novelty, and questioning why the authors did not integrate them into one paper.

Although the overall computational framework is conceptually related to recent de novo hydrolase-design studies, the specific RFDiffusion model, design pipeline, tool implementation and filtering criteria applied here are distinct. We adapted our workflow specifically for cysteine proteases by explicitly building and filtering active sites based on the Cys–His–Asp catalytic triad, substrate-binding geometry, and apo-state predictions. **This tailored approach achieved a high experimental success rate, with 13 active designs identified among 69 tested (19%),** exceeding the success rates reported in related studies. We also note that, although this study and the companion metalloprotease manuscript share broad conceptual similarities, they target distinct catalytic mechanisms. Combining the two studies would obscure the specific mechanistic and design pipeline details essential to each enzyme class.

Furthermore, the gap from the natural enzymes from which the active site was derived limits confidence in the authors' conclusions that low rates in the final designs are due to low preorganization. Given that all recent de novo hydrolases have shown low rates, there may be a ceiling to the rates that are feasible using the current generation of de novo fold design methods. Identifying such missing factors may explain what makes natural enzymes so proficient, beyond the accuracy of the active site, and suggest ways to break the rate ceiling.

In the revised manuscript, we further clarify the contribution of active-site preorganization to the observed gains in activity. While our previous version focused mainly on the Y41 and S44 mutations, which appear to improve the preorganization of the catalytic histidine, we now also discuss how the T99 and P67 mutations may enhance oxyanion-hole preorganization. Specifically, our analysis shows that the backbone carbonyl of the relevant glutamine residue

becomes progressively better stabilized by hydrogen-bonding interactions with Y20 across the Dokki series, correlating with increased activity (Fig. 4e). Nevertheless, the side-chain geometry of this residue still differs from the highly conserved configurations observed in natural proteases (Fig. 4h), consistent with the modest decrease in activity observed for the Gln knockout. These results suggest that the current oxyanion-hole interaction contributes only partially to catalysis (Fig. 4c). For Dokki-v2, redesigning the scaffold around the oxyanion hole to better fix the catalytic Gln geometry may be necessary to reduce the gap to native enzymes.

Furthermore, in the revised manuscript, we investigated the gap between Ulp1 and our design. In Ulp1, there is a large additional interface between the folded SUMO domain and the enzyme. To assess how this SUMO-mediated interface contributes to catalysis, we measured Ulp1 kinetics using eGFP-peptide-mScarlet substrates in which the cleavage-site peptide was presented either alone or within the context of an intervening SUMO domain. The peptide-only substrate contained six residues, comparable in length to the Dokki-v2 substrate, and retained the local Ulp1 recognition sequence (Fig. 5b). Ulp1 robustly cleaved the SUMO-containing substrate, with a k_{cat} of approximately 6.1 s^{-1} , whereas no detectable activity was observed toward the corresponding substrate lacking SUMO (Fig. 5c).

Together, these analyses suggest that efficient catalysis depends not only on active-site preorganization, but also on substrate-binding interfaces that constrain the cleavage-site peptide and enforce productive positioning of the scissile bond. We found that Dokki-v2 forms a smaller substrate-binding interface with fewer hydrophobic contacts than native proteases (Fig. 5e), providing another mechanistic explanation for the remaining activity gap relative to native enzymes.

Regarding generality, except for one designed enzyme, all were generated to target custom-designed substrates (two-sided designs), and the design process relied on the simplifying assumption that the substrate associates with the enzyme through β -sheet interactions. It is not clear from the writing whether the targeted human protein (one-sided design) was chosen because of similarity to some of the previously designed substrates, adding another limitation to generality.

Many proteases recognize substrates through extended β -strand-like backbone hydrogen-bonding interactions. As shown in Fig. 5b, TEV and HRV3C employ sandwich-like, strand-paired architectures to engage the substrate, while Ulp1 and Mpro stabilize the substrate

backbone via extensive hydrogen bonds. We adopted this architecture not as a simplifying assumption, but because it reflects a common strategy used by natural proteases to facilitate productive substrate binding.

Given the concerns on rates and generality, the authors' assertions that the work opens the way to therapeutic and industrial applications are not tenable.

Widely used laboratory cysteine proteases provide a useful point of reference. For example, TEV protease has a reported k_{cat} of 0.16 s^{-1} , only ~2.5-fold higher than that of Dokki-v2. Similarly, Dokki-v2 falls within the range of turnover numbers for HRV3C (Fig. 5a), indicating that its absolute rate lies within the lower range of native sequence-specific cysteine proteases. We agree, however, that such comparisons do not establish therapeutic or industrial utility, which would require further improvements in catalytic efficiency and validation under application-relevant conditions.

Reviewer #3

Additional experimental work has been performed and included. A mutagenesis study was performed for 20 residues around the active site, which led to the discovery of a variant “Dokki-v2” that has four additional mutations compared to the designed Dokki-15 enzyme. The Dokki-v2 variant showed a 1-2 orders of magnitude increased k_{cat} , attributed to better preorganization of active site residues. The MM parameters of this improved variant are only 1 order of magnitude below that of the natural TEV protease, which is an impressive result. These results and their discussion was included in the manuscript.


We thank Reviewer #3 for these positive comments. We greatly appreciate their careful review and their recognition of the additional experimental and mechanistic work included in the revision.

I have the feeling the effect of the His rotamers is overemphasized in the manuscript. If in the Dokki-15 only 6-10% of the His occupies the non-productive rotameric states, i.e. >90% is in the correct “active” rotameric state, then it is hard to understand the more than an order of magnitude improvement of activity observed for Dokki-v2, which has 100% occupancy of His in the active rotameric state. I am convinced that the controlling the His rotameric states is an important contributor to the increased activity of Dokki-v2, but in combination with other structural effects. The authors do note this briefly in the manuscript, by mentioning “Gln65 oxyanion hole contacts with shorter, less variable 233 H-bond distances” as possible additional contributor. But the impression given in the manuscript is still very much that His rotameric states are the main cause of the increased activity, whereas in their rebuttal to reviewer 1 they note that “it does seem like at least part of the gap in activity was a result of imprecise positioning of the catalytic triad in Dokki-v2.” I recommend to make the statements in the manuscript similarly more nuanced.

We agree that the role of the catalytic His rotameric state should be described more carefully and with greater nuance and we agree that it is unlikely to fully account for the more than order-of-magnitude improvement observed by only improved His preorganization.

In the previous version, our discussion emphasized the Y41 and S44 mutations, which appeared to enhance His preorganization. We now provide a more complete picture by also considering the contributions of the T99 and P67 mutations, which likely improve oxyanion-hole

preorganization. Our current analysis suggests that the backbone carbonyl of Q65, the oxyanion-hole residue, is progressively better stabilized through hydrogen-bonding interactions with Y20, contributing to the observed activity increase. The manuscript has been revised accordingly to clarify that His rotamer preorganization is one contributing factor, but not the sole determinant, of the enhanced activity of Dokki-v2, and to present a more balanced mechanistic interpretation.



We have also revised the manuscript to make clear that a substantial activity gap remains relative to the native precursor protease Ulp1. This gap likely reflects, at least in part, suboptimal geometry in the oxyanion hole of the Dokki designs compared to Ulp1 (Fig. 4h). Consistent with this view, the Gln knockout variant showed only a modest decrease in activity, indicating that the current oxyanion-hole interaction contributes relatively little to catalysis in its present form (Fig. 4c).

With regard to the response to my comments:

In my review, I noted the similarity of this manuscript to the companion manuscript on metalloprotease design: "The companion metalloprotease manuscript seemed a lot more fleshed out than this cysteine protease manuscript, and outperforms it with both scope and attention to detail, while the methods are the same or similar. For example, there is a less detailed investigation of substrate selectivity and no application to more biologically relevant reactions, although this prospect is raised in the discussion."

In their rebuttal, the authors do not really address this issue. They acknowledge the similarity in methods used, but emphasize the challenge of sub-angstrom design of the Cys-His-Asp catalytic triad that is addressed in this manuscript. While this is certainly true, the precise design of catalytic triads has been reported before by the Baker lab, e.g. for the design of Serine hydrolases (Science 2025, 10.1126/science.adu2454), using another approach. The current work, while impressive with regard to the results, does feel a bit as another, indeed “less fleshed out”, example of the application of the same methodology as used in the companion metalloprotease study – even though the quality and impact of the work itself, in isolation, is undisputed.

We apologize for not adequately addressing this point in our previous rebuttal and appreciate the reviewer's clarification. In the revised manuscript, we have expanded the substrate-selectivity analysis and clarified the specific contribution of this work. Although related computational strategies have been reported, the current study tackles a distinct and more demanding chemical challenge: the selective hydrolysis of a peptide amide bond, as opposed to activated ester substrates. We also clarify that cysteine proteases impose different geometric requirements than serine hydrolases, including the use of a different histidine nitrogen for nucleophile activation and a longer optimal Cys–His separation

We agree with the reviewer that this manuscript is narrower in scope than the companion metalloprotease study in terms of immediate biological applications. However, we view its primary contribution as demonstrating that de novo design can recapitulate key features of a canonical class of native proteases – namely cysteine proteases – to achieve selective hydrolysis of unactivated peptide bonds, thereby providing a foundation for future efforts. To further strengthen this point, we have added an analysis comparing substrate engagement in our designs with that of native sequence-specific proteases (Fig. 5). This comparison delineates general substrate-binding strategies and highlights design principles that can guide future work.

With regard to the response to my comment 3, about why not all transition states in the design process are included as reported before for the serine hydrolases (Science 2025, 10.1126/science.adu2454): I understand the reasoning provided, i.e. to focus on the first tetrahedral intermediate because that is “considered to approximate the rate-limiting transition state for peptide bond hydrolysis”, and the results show that this was a good choice. But it still

feels a little bit like a step back. In the Science 2025 paper the importance of including all transition states in the design process was emphasized as key to achieving the most active designed enzymes and here this is only mentioned as a possibility to further improve the design. In itself this is fine, but it feels a bit unsatisfying.

We agree that consideration of additional catalytic states is important for understanding the remaining activity gap and for guiding further improvements. In response, we have extended our analysis to include PLACER-generated ensembles spanning multiple states of the catalytic cycle, including the first transition-state (TS1), the acyl–enzyme intermediate (AEI), the second transition-state (TS2), and the and apo state (Extended Data Fig. 10a-c).

Among these states, TS1 most clearly distinguished Ulp1 from the Dokki designs and best matched the experimental activity trend, consistent with the initial nucleophilic attack being the primary rate-limiting step in peptide-bond hydrolysis. AEI and TS2 showed similar but less pronounced trends. In Dokki-v2, the dominant unproductive configuration involves an elongated oxyanion-hole distance, leading to a lower fraction of productive TS1-like ensembles compared to Ulp1 (36% versus 100%; Fig. 4f). Thus, our expanded analysis suggests that imperfect TS1-state oxyanion-hole preorganization remains a major limitation of the current designs.

A final note to the authors: a version of the manuscript with the changes highlighted, or at least a clear indication of where in the manuscript the changes were made would have been helpful.

We apologize for not previously providing a version with highlighted changes or clearly indicating where revisions had been made. In the revised manuscript, all changes are now clearly marked.

Reviewer #4

In this revised version of the manuscript, the authors provide computational evidence suggesting that incomplete active-site preorganization in the computational variants limits the activity of the first-generation designs. They further suggest that introducing mutations within the active-site pocket can improve the positioning of catalytic rotamers (see below) and thereby enhance the activity of the computational designs. However, as also noted for the related metalloprotease study, much of this additional improvement appears to have been achieved experimentally through a brute-force approach, involving the testing of 185 single-site variants at a selected set of positions within the active-site pocket (Tyr41, Ser44, Pro67, and Thr99).

The purpose of the mutational screen was not simply to search broadly for improved activity, but to test whether local perturbations within the designed active-site pocket could alleviate specific geometric limitations in the first-generation designs. The outcome was informative in both respects: it identified improved variants and, importantly, provided mechanistic support for the idea that suboptimal active-site preorganization contributes to the remaining activity gap.

A particularly concerning aspect is that none of these positions were sampled in the original design calculations because ProteinMPNN assigned them low probabilities. This raises questions about the reliability of the computational pipeline, since positions that ultimately proved important for improving active-site preorganization were not prioritized by ProteinMPNN.

ProteinMPNN is designed primarily to generate sequences compatible with a target fold, rather than to optimize catalytic function or active-site energetics. As a result, it does not explicitly account for the detailed interactions required for precise catalytic preorganization. We view this result not as evidence that the overall design pipeline is unreliable, particularly given that ProteinMPNN remains one of the most effective tools available for sequence generation on designed scaffolds.

Another important concern is that the mechanistic basis by which these mutations enhance active-site preorganization remains inconclusive. Much of the discussion focuses on the different side-chain rotamers of the catalytic histidine as predicted by PLACER, ranging from 100% in Ulp1 to 90% in Dokki, 94% in Dokki-15—where, notably, two conformations are also

predicted for the catalytic Asp—and 100% in Dokki-v2. However, analysis of the catalytic Cys-His distance shows that this distance progressively increases along the Ulp1 (3.85 Å) < Dokki < Dokki-15 < Dokki-v2 (4.2 Å) series, which does not support the conclusion of improved preorganization of the catalytic Cys-His dyad. Likewise, the discussion of how the Y41T and S44V mutations increase activity remains largely speculative, as does the proposed effect of P67V and T99F on oxyanion-hole preorganization.

Overall, although this revised version represents an improvement over the original submission, we do not believe that the manuscript is suitable for publication in Nature.

Our discussion of improved Cys–His dyad preorganization was intended to refer primarily to the rotameric state and positioning of the catalytic histidine, rather than to the absolute Cys–His distance alone. We appreciate the reviewer’s point that the observed increase in Cys–His distance warrants careful interpretation. However, these distances remain within the range observed in native papain structures, including approximately 3.8 Å in 1PPN, 3.9 Å in 1CVZ and 4.2 Å in 1BP4. Accordingly, we defined productive states using a Cys–His distance cutoff of <4.2 Å.

As noted above, our extended analysis indicates that, with increasing activity across the Dokki series, the backbone carbonyl of Q65 in the oxyanion hole becomes progressively better stabilized by hydrogen-bonding interactions. This suggests that activity improvements cannot be attributed solely to enhanced His preorganization, but also reflect improved organization of the oxyanion-hole architecture.

At the same time, we have revised the manuscript to make clear that a substantial activity gap remains relative to the native precursor protease Ulp1, likely due in part to suboptimal oxyanion-hole geometry in the Dokki designs. In particular, the side-chain configuration differs from that observed in Ulp1 (Fig. 4h). Consistent with this, the Gln knockout variant showed only a modest decrease in activity, suggesting that the designed oxyanion-hole interaction makes a relatively limited contribution to catalysis in its current form (Fig. 4c).

Also, as we discussed above, we investigated the activity gap between Ulp1 and our design by assessing the contribution of the SUMO-mediated binding interface. And we found that Ulp1 efficiently cleaved the SUMO-containing substrate, with a k_{cat} of approximately 6.1 s^{-1} , whereas no detectable activity was observed for the peptide-only substrate of comparable length to the

Dokki-v2 substrate (Fig. 5b,c). These results suggest that efficient catalysis depends not only on active-site preorganization, but also on substrate-binding interfaces that constrain the cleavage-site peptide and enforce productive scissile-bond positioning. Consistent with this, Dokki-v2 forms a smaller substrate-binding interface with fewer hydrophobic contacts than native proteases (Fig. 5e), providing an additional mechanistic explanation for the remaining activity gap.

Responses to Reviewers' Comments on the Cysteine protease Manuscript

The reviewers' comments are shown in blue italics, and our responses are shown in regular black font.

Reviewer #1

Given the similarities between the papers, below we assess both together.

The authors have considerably revised the manuscripts, adding data on structures that exhibit high accuracy, mutations that improve catalytic rates, and molecular dynamics simulations that help explain the partial improvement in catalytic rates in mutants based on active-site preorganization. These are important additions.

We thank Reviewer #1 for these constructive suggestions, which prompted additions that have substantially strengthened the manuscript. We are grateful for the opportunity to incorporate these analyses and experiments, and believe they have improved both the depth and clarity of the work.

Yet, even after optimization by systematic mutagenesis, the designed enzymes remain 2-3 orders of magnitude less active than the natural counterpart from which the active site was derived. This is the most relevant comparison, particularly because several recent studies have already demonstrated that RFDiffusion can produce de novo enzymes with low k_{cat} . We agree that the catalytic parameters of natural enzymes vary, but it is important to remember that much of this variation is due to trade-offs: TEV trades off rate for very high specificity, and the specificity tests carried out by the authors are far from sufficient to suggest they have achieved similar levels. We do not doubt that hydrolysing amides presents a greater challenge in terms of activation barriers than previous studies on hydrolases.

We agree that comparison to native enzymes provides an important benchmark, and we have revised the manuscript to clarify this point. To contextualize our designs, we compared the turnover number of Dokki-v2 with k_{cat} values for representative sequence-specific cysteine proteases, including Ulp1, Mpro, TEV and HRV3C, together with their cognate substrates (Fig. 4g and Supplementary Table 5). This analysis places Dokki-v2 at the lower end of the range

observed for native sequence-specific cysteine proteases, approximately 1-2 orders of magnitude below the best native enzymes.

Despite this gap, Dokki-v2 achieves turnover rates within the lower end of the native range for native enzymes that cleave unactivated peptide bonds. We view this as an important advance that addresses a long-standing challenge for enzyme design, particularly because previous RFDiffusion-based enzyme-design efforts targeted only more activated substrates.

But the methods are very similar to those of the previous studies and exactly the same in the two companion papers, limiting novelty, and questioning why the authors did not integrate them into one paper.

Although the overall computational framework is conceptually related to recent work on the design of de novo esterases, the specific RFDiffusion model, design pipeline, tool implementation and filtering criteria applied here are distinct. We adapted our workflow specifically for cysteine proteases by explicitly building and filtering active sites based on the Cys–His–Asp catalytic triad, substrate-binding geometry, and apo-state predictions. **This tailored approach achieved a high experimental success rate, with 13 active designs identified among 69 tested (19%),** exceeding the success rates reported in related studies. We also note that, although this study and the companion metalloprotease manuscript share broad conceptual similarities, they target distinct catalytic mechanisms. Combining the two studies would obscure the specific mechanistic and design pipeline details essential to each enzyme class.

Furthermore, the gap from the natural enzymes from which the active site was derived limits confidence in the authors' conclusions that low rates in the final designs are due to low preorganization. Given that all recent de novo hydrolases have shown low rates, there may be a ceiling to the rates that are feasible using the current generation of de novo fold design methods. Identifying such missing factors may explain what makes natural enzymes so proficient, beyond the accuracy of the active site, and suggest ways to break the rate ceiling.

In the revised manuscript, we further clarify the contribution of active-site preorganization to the observed gains in activity. While our previous version focused mainly on the Y41 and S44

mutations, which appear to improve the preorganization of the catalytic histidine, we now also discuss how the T99 and P67 mutations may enhance oxyanion-hole preorganization. Specifically, our analysis shows that the backbone carbonyl of the relevant glutamine residue becomes progressively better stabilized by hydrogen-bonding interactions with Y20 across the Dokki series, correlating with increased activity (Fig. 4e). Nevertheless, the side-chain geometry of this residue still differs from the highly conserved configurations observed in natural proteases (Extended Data Fig. 9e), consistent with the modest decrease in activity observed for the Gln knockout. These results suggest that the current oxyanion-hole interaction contributes only partially to catalysis (Fig. 4c). For Dokki-v2, redesigning the scaffold around the oxyanion hole to better fix the catalytic Gln geometry may be necessary to reduce the gap to native enzymes.

Furthermore, in the revised manuscript, we investigated the gap between Ulp1 and our design. In Ulp1, there is a large additional interface between the folded SUMO domain and the enzyme. To assess how this SUMO-mediated interface contributes to catalysis, we measured Ulp1 kinetics using eGFP–peptide–mScarlet substrates in which the cleavage-site peptide was presented either alone or within the context of an intervening SUMO domain. The peptide-only substrate contained six residues, comparable in length to the Dokki-v2 substrate, and retained the local Ulp1 recognition sequence (Fig. 4h). Ulp1 robustly cleaved the SUMO-containing substrate, with a k_{cat} of approximately 6.1 s^{-1} , whereas no detectable activity was observed toward the corresponding substrate lacking SUMO (Fig. 4i).

Together, these analyses suggest that efficient catalysis depends not only on active-site preorganization, but also on substrate-binding interfaces that constrain the cleavage-site peptide and enforce productive positioning of the scissile bond. We found that Dokki-v2 forms a smaller substrate-binding interface with fewer hydrophobic contacts than native proteases (Fig. 4i), providing another mechanistic explanation for the remaining activity gap relative to native enzymes.

Regarding generality, except for one designed enzyme, all were generated to target custom-designed substrates (two-sided designs), and the design process relied on the simplifying assumption that the substrate associates with the enzyme through β -sheet interactions. It is not clear from the writing whether the targeted human protein (one-sided design) was chosen

because of similarity to some of the previously designed substrates, adding another limitation to generality.

Many proteases recognize substrates through extended β -strand-like backbone hydrogen-bonding interactions. As shown in Extended Data Fig. 9h, TEV and HRV3C employ sandwich-like, strand-paired architectures to engage the substrate, while Ulp1 and Mpro stabilize the substrate backbone via extensive hydrogen bonds. We adopted this architecture not as a simplifying assumption, but because it reflects a common strategy used by natural proteases to facilitate productive substrate binding.

Given the concerns on rates and generality, the authors' assertions that the work opens the way to therapeutic and industrial applications are not tenable.

Widely used laboratory cysteine proteases provide a useful point of reference. For example, TEV protease has a reported k_{cat} of 0.16 s^{-1} , only ~ 2.5 -fold higher than that of Dokki-v2. Similarly, Dokki-v2 falls within the range of turnover numbers for HRV3C (Fig. 4g), indicating that its absolute rate lies within the lower range of native sequence-specific cysteine proteases.

We have also added cell-based reporter data on Fig. 5 showing that all Dokki-series proteases function in mammalian cells, where they cleave an H2B–sfGFP–substrate–mScarlet reporter and control mScarlet relocalization in a catalytic Cys-dependent manner. These results provide evidence that de novo designed proteases can function in living cells and support their potential use as programmable tools for cellular engineering and protein-trafficking applications.

Reviewer #3

Additional experimental work has been performed and included. A mutagenesis study was performed for 20 residues around the active site, which led to the discovery of a variant “Dokki-v2” that has four additional mutations compared to the designed Dokki-15 enzyme. The Dokki-v2 variant showed a 1-2 orders of magnitude increased k_{cat} , attributed to better preorganization of active site residues. The MM parameters of this improved variant are only 1 order of magnitude below that of the natural TEV protease, which is an impressive result. These results and their discussion was included in the manuscript.


We thank Reviewer #3 for these positive comments. We greatly appreciate their careful review and their recognition of the additional experimental and mechanistic work included in the revision.

I have the feeling the effect of the His rotamers is overemphasized in the manuscript. If in the Dokki-15 only 6-10% of the His occupies the non-productive rotameric states, i.e. >90% is in the correct “active” rotameric state, then it is hard to understand the more than an order of magnitude improvement of activity observed for Dokki-v2, which has 100% occupancy of His in the active rotameric state. I am convinced that the controlling the His rotameric states is an important contributor to the increased activity of Dokki-v2, but in combination with other structural effects. The authors do note this briefly in the manuscript, by mentioning “Gln65 oxyanion hole contacts with shorter, less variable 233 H-bond distances” as possible additional contributor. But the impression given in the manuscript is still very much that His rotameric states are the main cause of the increased activity, whereas in their rebuttal to reviewer 1 they note that “it does seem like at least part of the gap in activity was a result of imprecise positioning of the catalytic triad in Dokki-v2.” I recommend to make the statements in the manuscript similarly more nuanced.

We agree that the role of the catalytic His rotameric state should be described more carefully and with greater nuance and we agree that it is unlikely to fully account for the more than order-of-magnitude improvement observed by only improved His preorganization.

In the previous version, our discussion emphasized the Y41 and S44 mutations, which appeared to enhance His preorganization. We now provide a more complete picture by also

considering the contributions of the T99 and P67 mutations, which likely improve oxyanion-hole preorganization. Our current analysis suggests that the backbone carbonyl of Q65, the oxyanion-hole residue, is progressively better stabilized through hydrogen-bonding interactions with Y20, contributing to the observed activity increase. The manuscript has been revised accordingly to clarify that His rotamer preorganization is one contributing factor, but not the sole determinant, of the enhanced activity of Dokki-v2, and to present a more balanced mechanistic interpretation.



We have also revised the manuscript to make clear that a substantial activity gap remains relative to the native precursor protease Ulp1. This gap likely reflects, at least in part, suboptimal geometry in the oxyanion hole of the Dokki designs compared to Ulp1 (Extended Data Fig. 9e). Consistent with this view, the Gln knockout variant showed only a modest decrease in activity, indicating that the current oxyanion-hole interaction contributes relatively little to catalysis in its present form (Fig. 4c).

With regard to the response to my comments:

In my review, I noted the similarity of this manuscript to the companion manuscript on metalloprotease design: “The companion metalloprotease manuscript seemed a lot more fleshed out than this cysteine protease manuscript, and outperforms it with both scope and attention to detail, while the methods are the same or similar. For example, there is a less detailed investigation of substrate selectivity and no application to more biologically relevant reactions, although this prospect is raised in the discussion.”

In their rebuttal, the authors do not really address this issue. They acknowledge the similarity in methods used, but emphasize the challenge of sub-angstrom design of the Cys-His-Asp catalytic triad that is addressed in this manuscript. While this is certainly true, the precise design of catalytic triads has been reported before by the Baker lab, e.g. for the design of Serine hydrolases (Science 2025, 10.1126/science.adu2454), using another approach. The current work, while impressive with regard to the results, does feel a bit as another, indeed “less fleshed out”, example of the application of the same methodology as used in the companion metalloprotease study – even though the quality and impact of the work itself, in isolation, is undisputed.

We apologize for not adequately addressing this point in our previous rebuttal and appreciate the reviewer's clarification. In the revised manuscript, we have expanded the substrate-selectivity analysis and clarified the specific contribution of this work (Fig. 2h). Although related computational strategies have been reported, the current study tackles a distinct and more demanding chemical challenge: the selective hydrolysis of a peptide amide bond, as opposed to activated ester substrates. We also clarify that cysteine proteases impose different geometric requirements than serine hydrolases, including the use of a different histidine nitrogen for nucleophile activation and a longer optimal Cys–His separation

We agree with the reviewer that this manuscript is narrower in scope than the companion metalloprotease study in terms of immediate biological applications. However, we view its primary contribution as demonstrating that de novo design can recapitulate key features of a canonical class of native proteases – namely cysteine proteases – to achieve selective hydrolysis of unactivated peptide bonds, thereby providing a foundation for future efforts. To further strengthen this point, we have added an analysis comparing substrate engagement in our designs with that of native sequence-specific proteases (Fig. 4h,i). This comparison delineates general substrate-binding strategies and highlights design principles that can guide future work.

We have also added cell-based reporter data (Fig. 5) showing that all Dokki-series proteases function in mammalian cells, where they cleave an H2B–sfGFP–substrate–mScarlet reporter and control mScarlet relocalization in a catalytic Cys-dependent manner. Together, we believe these additions strengthen the manuscript by expanding the analysis of substrate recognition, and demonstrating cellular activity of the designed proteases.

With regard to the response to my comment 3, about why not all transition states in the design process are included as reported before for the serine hydrolases (Science 2025, 10.1126/science.adu2454): I understand the reasoning provided, i.e. to focus on the first tetrahedral intermediate because that is “considered to approximate the rate-limiting transition state for peptide bond hydrolysis”, and the results show that this was a good choice. But it still feels a little bit like a step back. In the Science 2025 paper the importance of including all transition states in the design process was emphasized as key to achieving the most active designed enzymes and here this is only mentioned as a possibility to further improve the design. In itself this is fine, but it feels a bit unsatisfying.

We agree that consideration of additional catalytic states is important for understanding the remaining activity gap and for guiding further improvements. In response, we have extended our analysis to include PLACER-generated ensembles spanning multiple states of the catalytic cycle, including the first transition-state (TS1), the acyl–enzyme intermediate (AEI), the second transition-state (TS2), and the and apo state (Extended Data Fig. 9a–c).

Among these states, TS1 most clearly distinguished Ulp1 from the Dokki designs and best matched the experimental activity trend, consistent with the initial nucleophilic attack being the primary rate-limiting step in peptide-bond hydrolysis. AEI and TS2 showed similar but less pronounced trends. In Dokki-v2, the dominant unproductive configuration involves an elongated oxyanion-hole distance, leading to a lower fraction of productive TS1-like ensembles compared to Ulp1 (36% versus 100%; Fig. 4f). Thus, our expanded analysis suggests that imperfect TS1-state oxyanion-hole preorganization remains a major limitation of the current designs.

A final note to the authors: a version of the manuscript with the changes highlighted, or at least a clear indication of where in the manuscript the changes were made would have been helpful.

We apologize for not previously providing a version with highlighted changes or clearly indicating where revisions had been made. In the revised manuscript, changes are now clearly marked.

Reviewer #4

In this revised version of the manuscript, the authors provide computational evidence suggesting that incomplete active-site preorganization in the computational variants limits the activity of the first-generation designs. They further suggest that introducing mutations within the active-site pocket can improve the positioning of catalytic rotamers (see below) and thereby enhance the activity of the computational designs. However, as also noted for the related metalloprotease study, much of this additional improvement appears to have been achieved experimentally through a brute-force approach, involving the testing of 185 single-site variants at a selected set of positions within the active-site pocket (Tyr41, Ser44, Pro67, and Thr99).

The purpose of the mutational screen was not simply to search broadly for improved activity, but to test whether local perturbations within the designed active-site pocket could alleviate specific geometric limitations in the first-generation designs. The outcome was informative in both respects: it identified improved variants and, importantly, provided mechanistic support for the idea that suboptimal active-site preorganization contributes to the remaining activity gap.

A particularly concerning aspect is that none of these positions were sampled in the original design calculations because ProteinMPNN assigned them low probabilities. This raises questions about the reliability of the computational pipeline, since positions that ultimately proved important for improving active-site preorganization were not prioritized by ProteinMPNN.

Most mutations introduced during the mutational screen significantly reduced enzyme activity (Fig. 4a), indicating that the original ProteinMPNN-designed sequence was generally well suited to maintaining the activity. However, because ProteinMPNN is primarily optimized to generate sequences compatible with a specified backbone fold, it may not fully capture subtle catalytic features such as active-site preorganization. We therefore do not interpret this result as evidence that the overall design pipeline is unreliable, especially given that ProteinMPNN remains one of the most effective tools available for sequence generation on designed scaffolds.

Another important concern is that the mechanistic basis by which these mutations enhance active-site preorganization remains inconclusive. Much of the discussion focuses on the different side-chain rotamers of the catalytic histidine as predicted by PLACER, ranging from 100% in Ulp1 to 90% in Dokki, 94% in Dokki-15—where, notably, two conformations are also predicted for the catalytic Asp—and 100% in Dokki-v2. However, analysis of the catalytic Cys-His distance shows that this distance progressively increases along the Ulp1 (3.85 Å) < Dokki < Dokki-15 < Dokki-v2 (4.2 Å) series, which does not support the conclusion of improved preorganization of the catalytic Cys-His dyad. Likewise, the discussion of how the Y41T and S44V mutations increase activity remains largely speculative, as does the proposed effect of P67V and T99F on oxyanion-hole preorganization.

Overall, although this revised version represents an improvement over the original submission, we do not believe that the manuscript is suitable for publication in Nature.

Our discussion of improved Cys–His dyad preorganization was intended to refer primarily to the rotameric state and positioning of the catalytic histidine, rather than to the absolute Cys–His distance alone. We appreciate the reviewer’s point that the observed increase in Cys–His distance warrants careful interpretation. However, these distances remain within the range observed in native papain structures, including approximately 3.8 Å in 1PPN, 3.9 Å in 1CVZ and 4.2 Å in 1BP4. Accordingly, we defined productive states using a Cys–His distance cutoff of <4.2 Å.

As noted above, our extended analysis indicates that, with increasing activity across the Dokki series, the backbone carbonyl of Q65 in the oxyanion hole becomes progressively better stabilized by hydrogen-bonding interactions. This suggests that activity improvements cannot be attributed solely to enhance His preorganization, but also reflect improved organization of the oxyanion-hole architecture.

At the same time, we have revised the manuscript to make clear that a substantial activity gap remains relative to the native precursor protease Ulp1, likely due in part to suboptimal oxyanion-hole geometry in the Dokki designs. In particular, the side-chain configuration differs from that observed in Ulp1 (Extended Data Fig. 9e). Consistent with this, the Gln knockout variant showed only a modest decrease in activity, suggesting that the designed oxyanion-hole interaction makes a relatively limited contribution to catalysis in its current form (Fig. 4c).

Also, as we discussed above, we investigated the activity gap between Ulp1 and our design by assessing the contribution of the SUMO-mediated binding interface. And we found that Ulp1 efficiently cleaved the SUMO-containing substrate, with a k_{cat} of approximately 6.1 s^{-1} , whereas no detectable activity was observed for the peptide-only substrate of comparable length to the Dokki-v2 substrate (Fig. 4h,i). These results suggest that efficient catalysis depends not only on active-site preorganization, but also on substrate-binding interfaces that constrain the cleavage-site peptide and enforce productive scissile-bond positioning. Consistent with this, Dokki-v2 forms a smaller substrate-binding interface with fewer hydrophobic contacts than native proteases (Fig. 4i), providing an additional mechanistic explanation for the remaining activity gap.

Responses to Reviewers' Comments on the Cysteine protease Manuscript

The reviewers' comments are shown in blue italics, and our responses are shown in regular black font.

Referees' comments:

Referee #1 (Remarks to the Author):

The current revision adds some computational analyses but not substantial new data to counter the criticisms of the previous reviews. The authors' assertion that the low K_{cat} is true also for TEV protease is, as noted in a previous review cycle, not justifiable, as TEV protease exhibits extremely high selectivity, which is why it is so useful in molecular biology. Such levels of specificity were not demonstrated for the current work.

The current revision includes substantial new experimental evidence, most notably the cell-based reporter assays presented in Fig. 5. These experiments demonstrate that all tested Dokki-series proteases are functional in mammalian cells and cleave their intended substrates in a cellular environment. These results directly address the previous concern regarding the biological relevance of the modest in vitro turnover rates and demonstrate that cellular activity is not limited to a single optimized design.

We agree that catalytic turnover and substrate specificity are distinct properties. Our intention in citing TEV proteases was solely to provide context for catalytic turnover, not to suggest that our designed proteases exhibit TEV-equivalent specificity. Although the experiments presented in Fig. 2h demonstrate a degree of sequence discrimination, the central advance of this study is the establishment of a general route for generating sequence-targeted cysteine protease activity in de novo designed proteins. Comprehensive characterization and optimization of substrate specificity represent important directions for future work but are beyond the scope of the present study.

The authors' computational analyses suggest that the performance gap relative to comparable natural enzymes stems from insufficient preorganization and incomplete substrate positioning. This is an almost self-evident conclusion that appears in many enzyme design studies. The critical test is whether the study can address these deficiencies and improve the enzymes, but this is not shown here.

The purpose of our computational analyses was to determine whether specific computationally measurable differences among closely related design variants could explain their experimentally observed differences in catalytic activity.

Because these proteins were generated computationally, establishing a design-by-design relationship between the modeled catalytic configurations and the experimental measurements is important for evaluating which features of the design process succeeded and which remain limiting. We therefore report analyses that produced clear, interpretable relationships between computational metrics and experimental activity. These results provide a mechanistic basis for understanding the performance differences among variants and identify specific structural features—particularly catalytic-site preorganization and productive substrate positioning—that can be incorporated more explicitly into future design and optimization cycles.

We believe that connecting experimentally measured activity with specific computational features across designed variants provides actionable information for improving subsequent generations of de novo proteases.

We previously noted that the methods are not sufficiently distinct from previous studies by the Baker lab. The authors explain that they implemented a different theozyme in this case, but this is not a substantial methodological difference. Given that in this and previous studies, the catalytic rates were low, and this study does not address this fundamental problem nor propose a substantially new approach, we do not see sufficient merit for publication in Nature.

The advance of this work lies in demonstrating that the computational framework can be extended to the chemically demanding problem of sequence-specific amide-bond hydrolysis. As described in our previous rebuttal, the integrated computational pipeline developed for this task yielded active proteases at an experimental success rate of 19%, exceeding those reported in previous de novo enzyme-design studies. Moreover, the most active designs achieved rate enhancements over the uncatalysed reaction ($k_{\text{cat}}/k_{\text{uncat}}$) of up to 3×10^7 , exceeding the values

reported in several recent de novo enzyme-design studies (Lauko et al., *Science*, 2025; Listov et al., *Nature*, 2025; Kim et al., *Nature*, 2026; Braun et al., *Nature*, 2026).

We further demonstrate that all tested Dokki-series proteases are functional in mammalian cells and cleave their intended substrates in a cellular environment. These results extend the relevance of the designs beyond purified in vitro systems and establish their potential for future synthetic-biology and biomedical applications.

We therefore believe that the combination of a high initial experimental success rate, substantial rate enhancement over the uncatalysed reaction, a demonstrated strategy for improving catalytic activity, and activity in mammalian cells represents a meaningful advance in de novo enzyme design. Collectively, these findings provide a foundation for developing increasingly efficient and functionally useful designed proteases.

Referee #7 (Remarks to the Author):

We thank the referee for the careful inspection of the crystallographic models and electron-density maps. We appreciate these detailed observations and agree that they provide a more nuanced interpretation of the local conformational differences. We have revised the structural models and corresponding discussion as described below.

• HC12: Overall the details are right and well explained in the manuscript. A minor detail: Gly91 adjacent to Cys94 (catalytic residue) has an unmodeled peptide flip; this may be related to the small difference density accumulating on Cys94, which may be indicative of the start of radiation damage. The related structure Kal C103A complex_D_1000301017_model-annotate_P1 was supplied but not mentioned in main text the manuscript. It is the same complex with the Cys mutated and indeed has the peptide flip as the main/only conformation.

We appreciate the referee's detailed insight of these structures and electron density. We agree that an alternative conformation of the main chain near Gly91 can be modeled. The related structure, *Kal C103A complex_D_1000301017_model-annotate_P1*, supports the presence of this alternative conformation/dynamic behavior. We also agree the difference density near Cys94 may be indicative of the onset of radiation damage possibly during the diffraction experiment. Although the 'Kal C103A complex' structure was supplied, we did not discuss it in

the main text because it does not materially affect the central findings or conclusions of the study and is therefore not essential to the main narrative.

• *Gawi_09: Overall I agree with the interpretation. Only one chain has alt conf for catalytic Cys, justified by slightly more pronounced density than the other chain. Overall the active site density features are very similar between the two chains, further supporting the interpretation and the Zn displacement hypothesis.*

We appreciate the referee's confirmation of our interpretation.

• *HC34/Dokki WT apo_D_1000301020_model-annotate_P1 looks more like Dokki_15 C100A complex_D_1000300811_model-annotate_P1. This despite the cysteine sulfonic acid on HC34. This suggests that the larger His shift of the Dokki_15_WT structure may be due to the absence of substrate, not just the damaged Cys. Furthermore, the absence of the substrate and the Cys sulfonic are accompanied by not just a side chain shift. Rather, the whole beta hairpin His45 sits on shifts. None of this materially changes any conclusions of the paper or even this section; I would just would reword as I think it is a bit more complicated than what they are explaining here.*

We agree that the structural differences cannot be attributed solely to cysteine oxidation. In apo Dokki-15, oxidation and the absence of substrate coincide with displacement of the His-bearing β -hairpin. We have revised the manuscript to describe these coupled structural effects and have removed wording that attributed the shift primarily to cysteine oxidation.

o For the Cys sulfonic, the authors may want to speculate whether this is caused by radiation damage or accumulates during the purification/crystallization. There may be ways to salvage the dataset and resolve this if it is radiation damage, which can address some of the issues they are seeing. For example, reprocessing the first half of the data, etc. James Fraser (UCSF)

We thank the referee for this helpful suggestion. To assess the possibility that cysteine oxidation potentially arose from radiation damage during data collection, we reprocessed the dataset using the first 50%, 30% and 20% of the collected diffraction images. We did not observe any significant changes in the electron density around the modified Cys residue. However, because the diffraction data were collected following multiple raster frames/exposures, we cannot

exclude the possibility that oxidation occurred during data collection. We appreciate the referee's suggestion, which prompted this additional analysis and helped us evaluate the possible origins of the observed cysteine modification. This clarification does not affect the overall structural conclusions but provides a more accurate interpretation of the observed cysteine modification.

We thank the referee again for these careful and constructive observations, which helped improve both the crystallographic models and the precision of our structural interpretation.