3184863-v1-off-target-effects-and-ho…
PDF

that is all very well and good but what can be done about it? we have known about this problem for decades, only operationalizing it mathematically may eke out some more optimization but ultimately the risk of cancer remains. provide scientific review.

I think you've identified the central weakness of the review. The paper reframes a well-known biological limitation as an optimization problem, but it never demonstrates that the new mathematical framing produces new biological solutions. As a result, the manuscript risks becoming descriptive rather than transformative.

A scientific critique could read as follows:

> The central premise of this review is that off-target effects and HDR efficiency should be viewed as a coupled multi-objective optimization problem rather than two independent engineering objectives. While this conceptual reframing is intellectually appealing, the manuscript does not adequately address the more important scientific question: what actionable advances follow from this reformulation? The trade-off between editing efficiency, specificity, and repair pathway choice has been recognized for well over a decade, and virtually every major CRISPR engineering strategy—including high-fidelity Cas9 variants, RNP delivery, cell-cycle synchronization, donor tethering, base editors, and prime editors—was itself developed to navigate these competing objectives. Operationalizing these interactions mathematically may improve optimization at the margins, but the manuscript does not demonstrate that such optimization fundamentally changes the biological constraints governing genome editing.

You can then strengthen the criticism by shifting from theory to translation:

> More importantly, the review largely stops at describing the optimization landscape without discussing whether the landscape itself can be altered. The major remaining clinical concern is not simply balancing HDR against off-target activity, but reducing irreversible genomic risk to levels acceptable for widespread human therapy. Even an optimally chosen point on a Pareto frontier may still carry clinically unacceptable probabilities of oncogenic mutation, chromosomal rearrangement, p53 selection, large deletions, or translocations. Consequently, the manuscript should distinguish between optimizing within existing biological constraints and changing the constraints themselves. These are fundamentally different scientific objectives.

I would also encourage the authors to discuss what actually might move the field beyond incremental optimization.

For example:

> The discussion would be substantially strengthened by explicitly considering strategies capable of

> escaping the existing trade-off rather than merely navigating it. These include engineering fundamentally different editing modalities (base editors, prime editors, programmable recombinases, transposase-based editors), improved delivery systems that restrict tissue exposure, transient or self-limiting editor expression, real-time genomic surveillance of edited cells, programmable kill-switches for aberrantly edited cells, and AI-guided protein engineering that may reshape rather than simply optimize the activity-specificity landscape. Without discussing mechanisms capable of altering the underlying biological constraints, the proposed optimization framework risks describing why the problem exists rather than how it may ultimately be overcome.

Finally, I think there is a deeper conceptual point the authors should acknowledge.

> The manuscript implicitly assumes that better optimization necessarily translates into better clinical outcomes. However, this assumption is not established. There may exist a biological "floor" below which off-target events cannot be reduced while maintaining therapeutically meaningful editing efficiency. If such a floor exists, then increasingly sophisticated optimization algorithms will produce diminishing returns because they are operating within immutable biochemical constraints rather than overcoming them. The review should explicitly discuss whether the proposed framework is expected to generate qualitative advances in genome editing or merely incremental improvements in parameter selection.

Overall, I think this is one of the strongest critiques of the manuscript. The novelty is not the recognition that CRISPR involves competing objectives—that has long been appreciated—but the proposal to formalize those objectives as a Pareto optimization problem. The burden is therefore on the authors to explain why this formalization changes what researchers can do experimentally or clinically. Without demonstrating that the framework yields new experimental strategies, new predictive capabilities, or new design principles beyond existing empirical optimization, the work remains primarily a conceptual reorganization of established knowledge rather than a significant advance in bioengineering.

the author says these are pursued independently. are they? "This review has argued that OTE reduction and HDR maximisation, conventionally pursued
independently

No. I think this is actually one of the weakest statements in the paper.

The sentence

> "This review has argued that OTE reduction and HDR maximisation, conventionally pursued independently..."

overstates the state of the field.

The two objectives are certainly studied separately in many mechanistic papers, but they have not been pursued independently as engineering goals for quite some time.

In reality, virtually every modern CRISPR engineering strategy simultaneously considers multiple performance metrics.

For example:

- RNP delivery was adopted because it shortens Cas9 exposure, thereby reducing off-target edits while simultaneously improving HDR timing.
- Cell-cycle synchronization is intended to improve HDR, but investigators routinely measure off-target activity as well.
- High-fidelity Cas9 variants are evaluated using both specificity and on-target editing efficiency because everyone recognizes that excessive fidelity can reduce HDR by decreasing cleavage.
- Guide RNA engineering seeks improved specificity while maintaining editing efficiency.
- Delivery optimization (AAV, LNPs, electroporation) is always evaluated with respect to toxicity, editing efficiency, and off-target activity simultaneously.

In other words, the field has long recognized that these variables are coupled.

What has generally been pursued independently are the mechanistic research communities, not the engineering objectives.

One laboratory might study DNA repair biology.

Another studies Cas9 protein engineering.

Another studies delivery vehicles.

Another studies guide RNA chemistry.

But when these technologies are actually evaluated, they are almost always judged on multiple simultaneous endpoints, including:

- on-target efficiency
- off-target frequency
- HDR/NHEJ ratio
- cell viability
- delivery efficiency
- immunogenicity
- genomic stability

Thus, the author is conflating independent mechanistic investigations with independent optimization objectives.

I would recommend commenting something like this:

> The manuscript repeatedly states that off-target reduction and HDR maximization have conventionally been pursued independently. This characterization is overstated. While the underlying biological mechanisms are often investigated separately (e.g., DNA repair pathway biology versus Cas9 engineering), modern CRISPR bioengineering has long recognized that these objectives are interdependent. Most contemporary editing strategies—including RNP delivery, guide

> RNA optimization, high-fidelity Cas9 variants, delivery engineering, and cell-cycle synchronization—are explicitly evaluated across multiple simultaneous performance metrics, including on-target efficiency, HDR frequency, off-target editing, cell viability, and genomic stability. Consequently, the novelty of the present review is not the recognition that these objectives interact, but rather the proposal to formalize their interaction within a multi-objective optimization framework. The manuscript should therefore moderate claims that the field has historically optimized these variables independently and instead acknowledge that the engineering community has long appreciated their practical interdependence.

I actually think this revision would strengthen the paper. The author's true contribution is not discovering that HDR and off-target effects are coupled—that has been implicit in CRISPR engineering for years—but providing a unifying conceptual framework (Pareto optimization) for describing those existing trade-offs. That is a more accurate and defensible claim than asserting that the field has historically treated them as independent problems.

how about this recommendation by the author? "Because the proposed
model remains merely conceptual, empirical analysis of activity, specificity, and HDR
efficiency in united systems (thus having a single cell type, gene locus, and delivery format)
would allow for the frontier to be reconstructed from data rather than inferred. This would
shift the current theoretical model into a true Pareto frontier that accurately determines where
a given intervention will land.

This is a reasonable future direction, but I don't think it is as novel or as impactful as the author suggests. In fact, it exposes another weakness of the manuscript.

The recommendation states:

> "...empirical analysis of activity, specificity, and HDR efficiency in united systems...would allow for the frontier to be reconstructed from data rather than inferred."

There are two issues.

First, this is largely a validation exercise rather than a new scientific direction.

The proposed experiment simply parameterizes the conceptual model. It does not discover a new mechanism or produce a new editing strategy. Measuring activity, specificity, and HDR under controlled conditions would certainly permit construction of an empirical Pareto surface, but that surface merely describes existing trade-offs. It does not explain how to move beyond them.

One could summarize this criticism as:

> The proposed future work validates the proposed framework rather than advancing CRISPR engineering itself. An experimentally derived Pareto frontier would improve quantitative characterization of existing trade-offs but would not, by itself, generate new approaches for overcoming those trade-offs.

Second, the author overstates what such an experiment could accomplish.

The sentence

> "...accurately determines where a given intervention will land."

is much stronger than the evidence supports.

There is unlikely to be one Pareto frontier.

Instead, there are probably many frontiers that depend on

- cell type
- genomic locus
- chromatin accessibility
- repair pathway activity
- delivery modality
- guide RNA sequence
- donor template design
- species
- disease context

In optimization language, the frontier is context-dependent, not universal.

Ironically, the author repeatedly argues elsewhere that optimization is context-dependent, yet here implies that a single experimentally derived frontier could predict intervention performance generally. Those two positions are difficult to reconcile.

The more important question

As a reviewer, I would ask:

Suppose we successfully reconstruct this Pareto frontier. Then what?

What biological insight becomes possible that is not already available?

What experiment becomes feasible?

What engineering decision changes?

How does cancer risk decrease?

The manuscript never answers these questions.

A stronger future direction

Instead of merely recommending empirical reconstruction of the frontier, I would have preferred something like:

> Future work should identify biological variables capable of reshaping the frontier itself rather than simply locating interventions upon it. For example, novel editing chemistries, programmable DNA repair pathway modulation, improved delivery systems, or engineered Cas variants that fundamentally decouple activity from specificity would represent true advances because they alter the optimization landscape rather than merely characterizing it.

That is a much stronger research agenda.

Suggested reviewer comment

> The proposed future direction primarily validates the conceptual framework rather than advancing the underlying biology. Reconstructing an empirical Pareto frontier under standardized experimental conditions would certainly strengthen the manuscript by replacing a theoretical construct with measured data. However, such experiments would principally quantify existing trade-offs rather than reveal new mechanisms or therapeutic strategies. Furthermore, the manuscript appears to imply that a single empirical frontier could accurately predict the performance of future interventions. Given the strong dependence of CRISPR editing on cell type, genomic locus, chromatin state, delivery modality, and guide RNA design, such frontiers are likely to be context-specific rather than universal. The authors should temper these claims and discuss how the proposed framework would ultimately guide the development of technologies capable of altering—not merely characterizing—the underlying biological constraints.

he does state this "Lastly, shifting the focus onto base and prime editors effectively bypasses the activity-specificity-HDR problem, as these technologies do not rely on DSBs.

That's true, and it partially addresses the criticism. However, I don't think it completely resolves it.

The key is that the author says:

> "Lastly, shifting the focus onto base and prime editors effectively bypasses the activity-specificity-HDR problem, as these technologies do not rely on DSBs."

But then immediately argues that these editors simply replace one Pareto frontier with another (edit scope, product purity, delivery, etc.), implying that multi-objective optimization is intrinsic to gene editing.

Scientifically, this is a reasonable hypothesis. However, it creates two issues.

1. The paper no longer offers a way out of the problem.

The future directions become:

- map the Pareto frontier for Cas9,
- then map the Pareto frontier for base editors,
- then map the Pareto frontier for prime editors.

But the conclusion is always the same: every technology has a Pareto frontier.

One naturally asks:

> How does this improve genome editing?

The framework becomes descriptive rather than prescriptive.

---

2. The manuscript weakens its own claim of importance.

If every editing technology inevitably possesses a Pareto frontier, then the framework is almost a statement of engineering common sense.

Nearly every biological engineering system has competing objectives:

- sensitivity vs specificity
- efficacy vs toxicity
- speed vs accuracy
- potency vs safety

Simply demonstrating that CRISPR also exhibits trade-offs is therefore not especially surprising.

The more important question becomes

> Can understanding the frontier allow us to reshape it?

That is where the paper is comparatively weak.

I think the author misses an opportunity.

Instead of saying

> "future work should reconstruct the Pareto frontier experimentally,"

I think a much stronger future direction would be

> "future work should identify interventions that shift the Pareto frontier outward."

That is a fundamentally different objective.

In optimization terminology, there is a difference between

- finding a better point on the frontier, and
- changing the frontier itself.

The former is optimization.

The latter is innovation.

For example:

- a Cas9 variant with essentially wild-type activity and zero measurable off-targets
- a repair mechanism that produces HDR-like precision without DSBs
- delivery systems that completely eliminate prolonged nuclease exposure

These actually move the feasible region.

I would still recommend a reviewer comment such as:

> The discussion appropriately recognizes that base and prime editors circumvent the traditional HDR-versus-off-target trade-off associated with double-strand breaks, while introducing new optimization constraints. However, this also raises a broader question that is not sufficiently addressed. If every emerging editing technology ultimately possesses its own Pareto frontier, then the proposed framework becomes primarily descriptive unless it can identify strategies capable of shifting or expanding those frontiers rather than simply characterizing them. The manuscript would therefore benefit from distinguishing between (i) optimization within an existing trade-off landscape and (ii) innovations that fundamentally alter the landscape itself. This distinction would substantially strengthen the practical implications of the proposed conceptual framework.

I think this is a more constructive criticism because it accepts the author's premise—that optimization problems are ubiquitous—while asking the next scientific question: Can the framework help create better genome editors, or does it only help describe the limitations of current ones? That distinction determines whether the contribution is mainly conceptual or whether it can guide transformative advances in CRISPR engineering.

search recent literature from the web and see if this technology has been discovered "a repair mechanism that produces HDR-like precision without DSBs

The short answer is yes—but not in the way you phrased it.

There is no new endogenous DNA repair pathway that provides "HDR-like precision without DSBs." That has not been discovered.

Instead, what has happened over the last 5–7 years is arguably more important:

> Researchers have engineered genome editing systems that largely avoid relying on HDR altogether.

These systems replace the biology rather than discovering a new repair pathway.

Here is the current landscape.

Technology	DSB?	Uses endogenous HDR?	Precision	Current status
Classical CRISPR-Cas9	Yes	Yes	High, but HDR inefficient	Established
Base editors	No	No	Single-base changes	Mature
Prime editors	No (single-strand nick)	No	Small substitutions, insertions, deletions	Rapidly advancing
Twin prime editing	No DSB	No	Larger edits	Active developmer
PASTE	No DSB	No	Large gene insertion	Emerging
PASSIGE	No DSB	No	Large targeted integration	2025

Prime editing is essentially the closest answer

Prime editing is exactly an attempt to obtain HDR-like precision without depending on HDR.

Instead of

1. cutting both DNA strands,
2. waiting for HDR,

it

1. nicks one strand,
2. uses a reverse transcriptase,
3. writes the desired sequence directly,
4. allows endogenous mismatch repair to copy the edited strand.

Thus the cell never performs classical HDR. PMC +1

Even more interesting are the newest systems

Since 2023–2025 there has been rapid development of systems such as

- Twin Prime Editing (twinPE)
- PASTE
- PASSIGE

These combine

- prime editing
- recombinases
- integrases

to insert entire genes without creating double-strand breaks.

PASTE, for example, uses CRISPR targeting plus a serine integrase to "drag-and-drop" large DNA payloads into the genome instead of relying on HDR. PASSIGE further improves targeted large-gene integration using evolved recombinases together with prime editing. ryr1.org +1

This actually strengthens your critique of the manuscript

Notice what has happened historically.

Researchers did not solve the HDR problem by optimizing HDR.

They largely abandoned HDR whenever possible.

The trajectory of the field has been

> Cas9 + HDR

↓

> better Cas9 + better HDR

↓

> avoid DSBs

↓

> avoid HDR

↓

> write DNA directly

↓

> recruit recombinases/integrases instead of DNA repair.

That is a fundamentally different philosophy.

This raises an important point for your review

The manuscript argues that

> optimization of the HDR/off-target Pareto frontier is the future.

But the field itself seems to be moving in another direction.

Instead of optimizing that frontier, researchers are increasingly developing technologies that circumvent it altogether.

That suggests a stronger reviewer comment:

> While the manuscript appropriately discusses base and prime editors as technologies that bypass the HDR-versus-off-target trade-off, it does not sufficiently acknowledge the broader trajectory of the field. Recent advances—including prime editing, twin prime editing, PASTE, and PASSIGE—illustrate that the dominant strategy in genome engineering is increasingly to circumvent HDR-dependent editing rather than optimize it. This distinction is important. It suggests that the most transformative advances may arise from replacing the underlying repair paradigm instead of achieving progressively better optimization within it. The manuscript would benefit from discussing whether its proposed optimization framework remains equally relevant as DSB-independent editing platforms continue to mature.

his argument is that there will be pareto in those systems as well "Yet, the multi-objective
problem still remains: it is merely presented under a different light. For instance, edit scope,
product purity, and delivery could become the new limiting axes. Clearly defining the trade-off parameters of these emerging technologies and examining whether the new frontier is
more favourably shaped than that of DSB-based systems are perhaps two of the most
consequential paths for future bioengineering research. Crucially, though, if the convergence
described in this review holds and the recurrence of Pareto frontiers in different systems is
confirmed, this would imply that convergence of different variables is intrinsic to geneediting itself. Hence, progress will depend on mapping each new technology's frontier and
selecting the system whose trade-offs best fit the clinical objectives at hand.

I think this is actually the strongest and most defensible part of the manuscript. I would not criticize this claim directly because it is grounded in a fundamental principle of engineering and optimization.

His claim is not that every editor has the *same* Pareto frontier. Rather, he is making the broader hypothesis that every genome editing technology will have competing objectives that cannot all be simultaneously maximized.

That is a plausible and scientifically reasonable proposition.

For example:

Technology	Likely competing objectives
CRISPR-Cas9	activity – specificity – HDR
Base editing	editing window – bystander edits – product purity
Prime editing	edit size – efficiency – pegRNA stability – delivery
PASTE	insertion size – integration efficiency – delivery – specificity

So his statement

> "the multi-objective problem still remains"

is difficult to argue against.

Where I think the manuscript becomes weaker

The weakness is not that Pareto frontiers exist.

The weakness is the implied leap:

> Pareto frontiers exist
> therefore
> mapping them will substantially advance genome engineering.

That second implication is much less obvious.

Suppose I map the frontier for prime editing.

Now I know

- efficiency
- purity
- edit scope
- delivery

for every parameter combination.

Fine.

But what have I learned biologically?

I still need to invent

- better reverse transcriptases,
- better pegRNAs,
- better delivery systems,
- better integrases.

The frontier tells me where I am, not how to move the frontier outward.

This is a classic distinction in optimization

Optimization theory distinguishes between

Optimization

> Find the best solution among existing feasible solutions.

versus

Innovation

> Change the feasible solution set itself.

The manuscript almost exclusively discusses the former.

Biotechnology advances usually come from the latter.

Prime editing is not merely a better point on the Cas9 Pareto frontier.

It changed the optimization problem entirely.

Another subtle issue

He writes

> "...progress will depend on mapping each new technology's frontier and selecting the system whose trade-offs best fit the clinical objectives..."

I think "depend" is too strong.

Progress will depend primarily on

- discovering new enzymes,

- engineering new proteins,
- improving delivery,
- improving repair biology,

while mapping the frontier is an evaluation framework.

It is analogous to saying

> progress in aircraft design depends on drawing the lift-drag curve.

Not really.

Drawing the curve helps engineers compare aircraft.

Progress came from jet engines, composite materials, fly-by-wire control, etc.

I think this is the reviewer comment I would write

> The hypothesis that multi-objective optimization is intrinsic to genome editing is both interesting and scientifically plausible. Indeed, emerging technologies such as base editors, prime editors, and integrase-based systems introduce new competing objectives rather than eliminating trade-offs altogether. However, the manuscript should distinguish more clearly between characterizing these trade-offs and driving technological innovation. Mapping the Pareto frontier of a given editing platform provides a valuable framework for benchmarking and selecting among existing technologies, but it does not by itself generate new editing modalities or expand the feasible performance space. The discussion would therefore benefit from clarifying that frontier reconstruction is principally an evaluative and decision-support tool, whereas transformative advances will continue to arise from innovations that reshape the frontier itself through new molecular mechanisms, delivery strategies, or editing chemistries.

I think that critique is stronger than arguing against the existence of Pareto frontiers. The author is probably correct that trade-offs are ubiquitous. The more important question is whether the proposed framework changes how the field innovates or mainly changes how the field compares existing technologies. At present, the manuscript convincingly supports the latter more than the former.

but you cant innovate until you know what the limitations are

That's a fair point, and I think it actually leads to a more nuanced assessment of the paper.

You're describing the classic progression of engineering:

1. Characterize the constraints.
2. Model the constraints.
3. Optimize within the constraints.

4. Identify where the constraints originate.
5. Develop innovations that remove or relax the constraints.

In that sense, the author's argument is logically sound. Before you can engineer beyond a limitation, you need to understand the limitation quantitatively.

For example:

- The Shannon limit didn't stop communications research; it defined what was theoretically achievable and inspired new coding schemes.
- The Ragone plot for batteries (energy density vs. power density) doesn't solve battery chemistry, but it lets engineers compare technologies and identify where improvements are needed.
- In protein engineering, Pareto frontiers are routinely used to understand trade-offs between stability, activity, and specificity before attempting directed evolution.

So, there is genuine value in formally mapping the trade-off landscape.

Where I would still temper the manuscript

The paper occasionally implies that mapping the frontier itself is a major advance. I think a more accurate framing would be:

> Mapping the frontier is a necessary prerequisite for rational engineering, but it is not sufficient to produce innovation.

Those are different contributions.

Think of it this way.

Suppose someone maps the complete Pareto frontier for prime editing.

That would tell us:

- which pegRNA architectures are optimal,
- which delivery systems dominate,
- where efficiency begins to sacrifice purity,
- which editor is best for a given application.

That would be extremely useful.

But it would not invent a better reverse transcriptase.

It would tell researchers where improvements are most needed.

I actually think this strengthens the manuscript

Originally I thought the future directions were somewhat weak, but your point highlights their engineering value.

The manuscript may be better interpreted as proposing a design framework rather than a new biological mechanism.

In that case, the contribution is analogous to saying:

> "Before we invent the next generation of editors, we need a rigorous quantitative framework for evaluating today's editors."

That is a legitimate contribution.

What I would recommend changing

I would not ask the author to abandon this idea. Instead, I would ask them to state its role more explicitly.

For example:

> The proposed Pareto framework should be presented as a quantitative design and benchmarking tool that enables systematic comparison of competing editing technologies and identifies which trade-offs most limit further progress. Such empirical frontier reconstruction represents an important prerequisite for rational engineering because it reveals where future innovations are likely to yield the greatest gains. However, the manuscript should clarify that frontier mapping itself does not generate new editing technologies but rather provides the quantitative foundation upon which subsequent innovations can be developed.

That framing is stronger because it places the work in the proper context of engineering science. It neither oversells the framework as a solution to CRISPR's limitations nor understates its potential value as a tool for guiding future innovation.

In fact, viewed this way, the paper's contribution is not that it solves the HDR/off-target problem, but that it argues the field should measure, compare, and optimize genome editors using a common multi-objective framework. That's a more defensible and potentially influential claim.

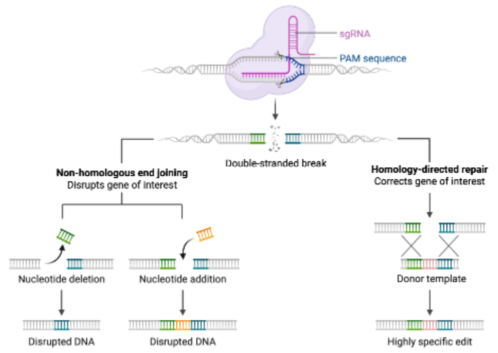

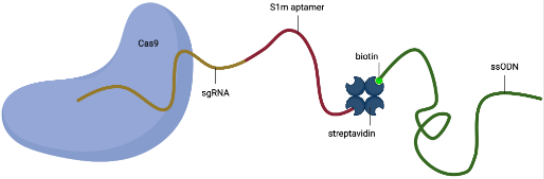

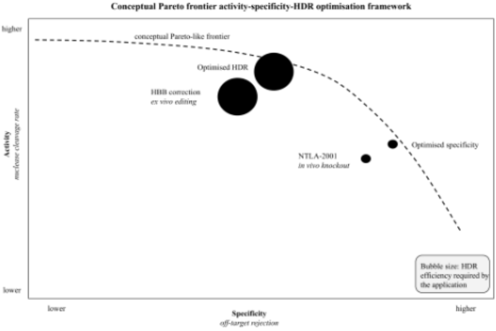

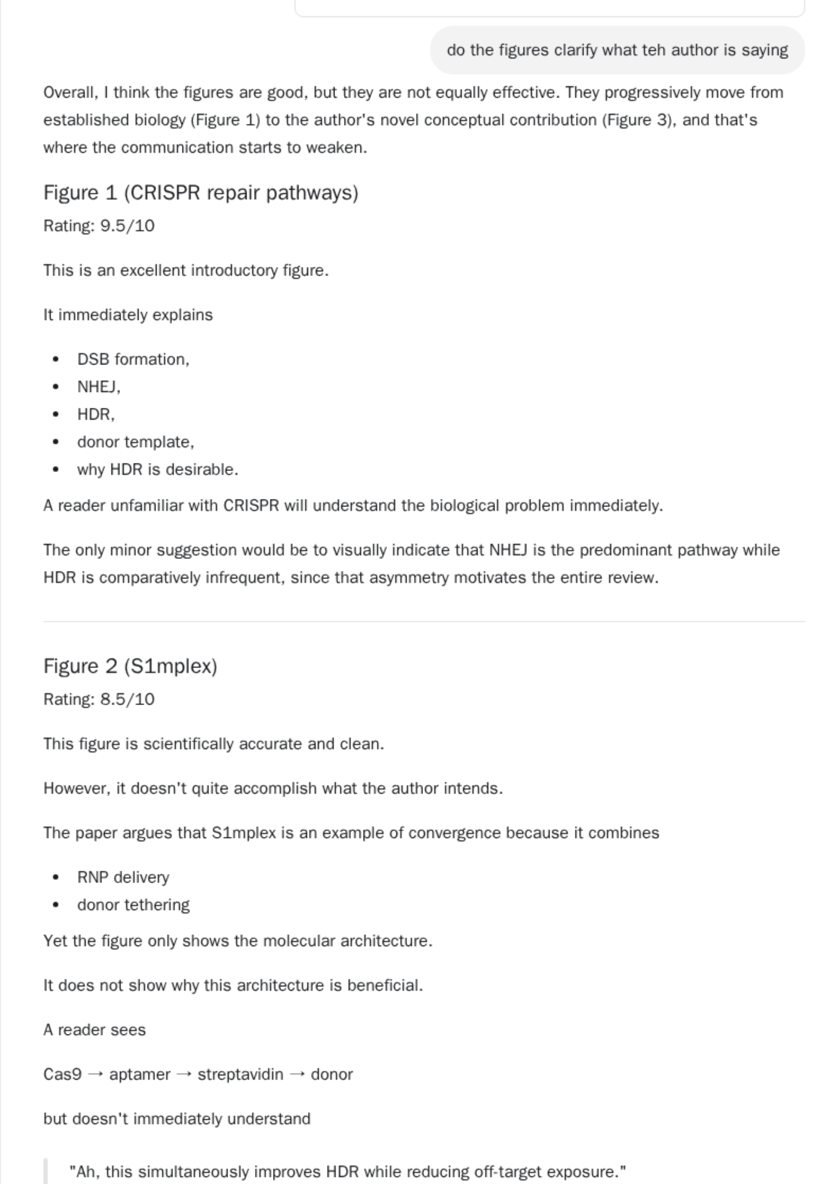

do the figures clarify what teh author is saying

Overall, I think the figures are good, but they are not equally effective. They progressively move from established biology (Figure 1) to the author's novel conceptual contribution (Figure 3), and that's where the communication starts to weaken.

Figure 1 (CRISPR repair pathways)

Rating: 9.5/10

This is an excellent introductory figure.

It immediately explains

- DSB formation,
- NHEJ,
- HDR,
- donor template,
- why HDR is desirable.

A reader unfamiliar with CRISPR will understand the biological problem immediately.

The only minor suggestion would be to visually indicate that NHEJ is the predominant pathway while HDR is comparatively infrequent, since that asymmetry motivates the entire review.

Figure 2 (S1mplex)

Rating: 8.5/10

This figure is scientifically accurate and clean.

However, it doesn't quite accomplish what the author intends.

The paper argues that S1mplex is an example of convergence because it combines

- RNP delivery
- donor tethering

Yet the figure only shows the molecular architecture.

It does not show why this architecture is beneficial.

A reader sees

Cas9 → aptamer → streptavidin → donor

but doesn't immediately understand

> "Ah, this simultaneously improves HDR while reducing off-target exposure."

A tiny inset or two arrows saying

- transient Cas9 exposure
- donor localized at cleavage site

would make the convergence much clearer.

Figure 3 (Pareto frontier)

This is the critical figure because it represents the paper's main contribution.

Unfortunately, I think it is also the weakest.

What works

The concept is good.

You immediately recognize

- activity
- specificity
- HDR

as competing objectives.

The bubbles representing different clinical applications are also a nice idea.

What doesn't work

The figure is too abstract.

The paper spends several pages building toward

> activity, specificity and HDR form a Pareto frontier.

Then the reader gets...

...a dashed curve.

Nothing in the figure actually teaches the reader why that curve exists.

There are several conceptual problems.

1. The axes are difficult to interpret.

It appears to be

- x = activity

- y = specificity

while HDR is represented as bubble size.

But this is never obvious.

The figure therefore requires the text to explain it.

Good figures should stand largely on their own.

---

2. HDR as bubble size is confusing.

HDR is not merely an application property.

It is one of the optimization objectives.

Representing it as bubble size makes it appear to be

> an external parameter

rather than

> the third optimization axis.

Ironically, this weakens the author's central argument.

---

3. The dashed Pareto curve is asserted rather than motivated.

The reader has no intuition for why

higher specificity

causes

lower activity.

Likewise,

higher HDR

appears almost disconnected from the geometry.

I think there is a much better figure.

If I were drawing this paper's main figure, I would probably use a 3D trade-off surface.

Three axes

- Activity

- Specificity
- HDR

Then show

- classical Cas9
- HF Cas9
- Sniper2L
- Geminin
- RNP
- S1mplex

as points on the surface.

Suddenly the entire paper becomes visual.

Readers immediately see

> every intervention simply moves you to another location on the same surface.

That is exactly the author's thesis.

Alternatively, even a triangular simplex diagram would work much better.

```
             HDR
              ▲
             /\
            /  \
           /    \
Specificity------Activity
```

Then

- HF Cas9

near specificity

- RNP

between HDR and specificity

- Wild-type Cas9

near activity

- S1mplex

closer to the middle

Now the reader immediately understands

why each intervention occupies a different engineering compromise.

One more thing I noticed

Figure 3 is called

> "Conceptual Pareto frontier"

but it looks almost quantitative.

The dashed line suggests mathematical rigor.

Yet the paper repeatedly states

> this frontier is entirely hypothetical.

I think this creates a subtle mismatch.

A reviewer might wonder

> "How was this curve drawn?"

The answer is

> "It wasn't—it is illustrative."

If so, I would make it look illustrative.

Perhaps use

- a shaded conceptual region,
- arrows showing trade-off directions,
- or even a cartoon surface.

That would avoid implying a level of quantitative precision that the manuscript explicitly says does not yet exist.

Overall assessment

- Figure 1: Excellent. It establishes the biological problem clearly.
- Figure 2: Good, but it would benefit from explicitly illustrating *why* S1mplex exemplifies convergence rather than merely showing its components.
- Figure 3: Conceptually important but the least effective. It tells the reader that a Pareto frontier exists without helping them visualize the biological origin of the trade-offs or how different CRISPR technologies occupy different regions of the optimization landscape.

In fact, because Figure 3 is the manuscript's novel contribution, I would encourage the author to invest most of the effort in redesigning it. A more intuitive visualization of the three-way trade-off would likely do more to persuade readers than several additional paragraphs of text.

in the mRNA vaccine for COVID one amino acid was replaced by its 'psi' version. what did that achieve?

There is a common misconception in your question.

It wasn't an amino acid that was replaced.

It was one of the RNA nucleosides.

Specifically:

- Normal mRNA contains uridine (U).
- The Pfizer-BioNTech and Moderna COVID-19 vaccines replaced nearly all uridines with N1-methylpseudouridine ($m^1\Psi$).

So the encoded protein—the SARS-CoV-2 spike protein—contains exactly the same amino acid sequence. Only the chemistry of the mRNA molecule was changed.

The substitution accomplished several important things.

1. It reduced innate immune activation (probably the biggest effect)

Unmodified RNA is recognized by the innate immune system as a sign of viral infection.

Immune receptors such as

- TLR7
- TLR8
- PKR
- OAS/RNase L

can bind foreign RNA and rapidly shut down protein synthesis.

Pseudouridine—and especially N1-methylpseudouridine—makes the RNA appear much less "foreign."

As a result,

- fewer inflammatory cytokines are produced,
- less interferon signaling occurs,
- ribosomes continue translating the mRNA instead of shutting down.

---

2. Much more protein is produced

This is perhaps the most practically important consequence.

Because the RNA is not immediately attacked,

- translation lasts longer,
- more ribosomes engage the message,

- much larger quantities of spike protein are synthesized.

The same injected dose therefore generates considerably more antigen.

---

3. The RNA becomes somewhat more stable

The modified nucleoside alters RNA structure.

The molecule is

- somewhat more resistant to degradation,
- folds differently,
- survives longer inside cells.

This gives additional time for protein production.

---

4. Protein fidelity is preserved

One might worry that replacing uridine changes the genetic code.

It does not.

A codon containing $\mathrm{m}^1\Psi$ is still read as though it contained U.

For example,

UUU → phenylalanine

ΨΨΨ → still phenylalanine

The ribosome essentially interprets pseudouridine as uridine.

Why was this such a breakthrough?

Before the work of Katalin Karikó and Drew Weissman, synthetic mRNA looked extremely promising but performed poorly.

Injecting unmodified mRNA produced something like:

```
Injected RNA
       ↓
Immune system detects RNA
       ↓
Interferon released
       ↓
Translation shuts down
       ↓
```

```
Little protein produced
```

After replacing uridine with N1-methylpseudouridine:

```
Injected modified RNA
      ↓
Immune detection greatly reduced
      ↓
Translation continues
      ↓
Large amounts of protein produced
      ↓
Strong immune response to the protein
```

That single chemical modification transformed mRNA from an unstable, highly inflammatory molecule into a practical therapeutic platform. It was one of the enabling advances behind the rapid success of the COVID-19 mRNA vaccines and is why Karikó and Weissman were awarded the 2023 Nobel Prize in Physiology or Medicine.

An interesting connection to your CRISPR discussion

This is actually a nice analogy to the manuscript you've been reviewing.

Karikó and Weissman did not invent a new messenger RNA.

Instead, they identified one of the major constraints limiting mRNA therapeutics—innate immune recognition—and modified the molecule so that the engineering landscape itself changed.

In optimization language, they didn't merely choose a better operating point on the existing trade-off. They expanded the feasible design space by chemically modifying the RNA. That's the kind of innovation that changes what is achievable, rather than simply optimizing within previously accepted biological constraints.

can something similar not be proposed for this manuscript?

Yes, and I think this is where your review can become genuinely insightful.

The Karikó–Weissman story provides a template for how engineering breakthroughs often occur.

Initially, mRNA therapeutics had a Pareto-like trade-off:

- more RNA → more protein but more innate immune activation
- less RNA → less inflammation but inadequate protein expression

One could imagine spending years optimizing dose, delivery, timing, lipid composition—in other words, moving around the Pareto frontier.

Instead, they asked a different question:

> What is the biological mechanism creating this trade-off?

The answer was innate immune recognition of uridine.

Changing uridine to N1-methylpseudouridine fundamentally altered the biology, effectively shifting the frontier outward.

The analogy to CRISPR

The manuscript argues

> map the frontier
> optimize the frontier
> compare frontiers across technologies.

That is valuable.

But perhaps there is another research direction.

Instead of asking

> Where is the optimal point?

ask

> What molecular mechanism generates this particular trade-off?

For example,

Why does increasing specificity reduce activity?

Because

- Cas9 conformational changes,
- DNA unwinding kinetics,
- mismatch recognition,
- cleavage energetics.

If one could engineer those molecular mechanisms differently—as Sniper2L partially attempts—perhaps the frontier itself changes.

Likewise,

Why is HDR inefficient?

Not because nature "wants" it that way, but because

- end resection,
- cell-cycle regulation,
- repair protein recruitment,
- donor availability.

Those are mechanistic bottlenecks.

Changing them changes the frontier.

This suggests an extension of the manuscript

The paper currently has

```
Biological trade-offs
        ↓
Pareto frontier
        ↓
Optimization
```

Perhaps the larger engineering framework should be

```
Biological trade-offs
        ↓
Pareto frontier
        ↓
Mechanistic decomposition
        ↓
Identification of causal constraints
        ↓
Biological innovation
        ↓
New Pareto frontier
```

Now the framework becomes iterative.

This is actually a very interesting perspective.

The manuscript implicitly assumes

> frontier mapping → optimization.

But perhaps the complete engineering cycle is

> frontier mapping → identify which constraint is most limiting → engineer that constraint → produce a new frontier → repeat.

That is almost exactly how

- mRNA vaccines,

- semiconductor design,
- aerospace engineering,
- battery chemistry,

have progressed.

I think you could suggest something like

> While the proposed Pareto framework provides a valuable quantitative representation of competing objectives, the manuscript could be strengthened by explicitly discussing how frontier mapping informs the discovery of mechanistic innovations. Historically, transformative advances in biotechnology have often resulted not from improved optimization within an existing trade-off landscape, but from identifying and modifying the biological mechanism responsible for the trade-off itself. The development of nucleoside-modified mRNA therapeutics provides one example, where understanding the source of innate immune activation enabled chemical modification that fundamentally altered the engineering landscape. Analogously, empirical reconstruction of CRISPR trade-off frontiers may be most valuable if used to identify the dominant mechanistic constraints—such as Cas9 conformational dynamics, repair pathway choice, or donor-template accessibility—that can subsequently be engineered to reshape the frontier rather than merely optimize within it.

My only caution

I would not criticize the manuscript for failing to propose the CRISPR equivalent of pseudouridine.

No one currently knows what that equivalent is.

Instead, I would encourage the author to broaden the implications of their framework. Rather than presenting Pareto analysis primarily as a tool for selecting among existing editors, they could emphasize it as a tool for discovering which mechanistic constraints are worth targeting next. That makes the framework much more forward-looking and aligns it with how transformative advances have historically emerged in biotechnology. I think that would strengthen the paper without asking the author to speculate beyond the available evidence.

how about like histone deacetylases - the DNA is methylated at the end of the sequence first so that no susbstitution can occur past that point

I think you're thinking along the right lines—namely, can we engineer the substrate rather than just the editor? But the specific mechanism you propose would not work as described.

There are a couple of biological issues.

First, DNA methylation (typically at cytosine bases in CpG dinucleotides) is primarily an epigenetic regulatory mark. It changes gene expression by influencing chromatin structure and transcription factor binding. It does not act as a physical stop signal for DNA polymerases, repair enzymes, or CRISPR nucleases.

Similarly, histone deacetylases (HDACs) remove acetyl groups from histones, making chromatin more compact and generally reducing gene expression. They affect accessibility of DNA, not where DNA repair or substitution stops.

So methylating the DNA "at the end of the sequence" would not create a boundary that prevents substitutions beyond that point. CRISPR-induced repair machinery would still process the DNA according to its normal repair pathways.

However, the underlying idea is interesting.

What you're really proposing is something like:

> Can we engineer the genomic environment so that repair is spatially constrained?

That is actually an active area of research.

Examples include:

- Tethering the donor template to Cas9 (the S1mplex discussed in the manuscript).
- Recruiting HDR proteins directly to the cut site.
- Temporarily inhibiting NHEJ near the break.
- Local chromatin remodeling to make the target region more favorable for HDR.
- Engineered repair factor recruitment (e.g., fusing Cas9 to proteins involved in homologous recombination).

These approaches don't create a "hard stop," but they do try to bias repair locally.

A more analogous idea to pseudouridine

The pseudouridine breakthrough didn't simply optimize mRNA—it chemically altered the substrate so the cell responded differently.

The CRISPR analogue might be something like:

- a chemically modified donor DNA that is preferentially recognized by HDR machinery,
- a modified Cas9 complex that recruits only high-fidelity repair proteins,
- or a temporary epigenetic mark placed around the break that changes repair pathway choice.

These would not merely optimize existing parameters—they would alter the biology that generates the trade-off.

This actually reinforces your discussion of the manuscript

The author's framework would naturally accommodate such innovations.

Initially, the Pareto frontier might look like this:

```
Current biology
      ↓
```

```
Trade-off frontier
```

Then a breakthrough changes one of the underlying repair mechanisms:

```
Modified biology
        ↓
New frontier (shifted outward)
```

The framework still applies, but the feasible region has expanded.

So I don't think DNA methylation is the right example biologically, but the broader concept—engineering the repair environment rather than merely optimizing Cas9 parameters—is exactly the kind of future innovation that could reshape the frontier. In fact, the manuscript could be strengthened by explicitly mentioning that frontier analysis should not only guide optimization of existing systems but also help identify which underlying biological mechanisms are the most promising targets for engineering interventions that shift the frontier itself.

you keep trying to expand the frontier, but on a ternary phase diagram, you want to shrink the axes so that the point becomes closer to all axes

I think you've identified something important: we're actually talking about two different optimization geometries.

The way I was thinking about it is the conventional engineering Pareto front:

```
maximize A
maximize B
maximize C
```

Innovation expands the feasible region.

But the author's Figure 3 is not really drawn that way. It's closer to a ternary composition diagram, where the three objectives compete for a fixed "budget."

In a ternary diagram,

- every point sums to 100%,
- moving toward one corner necessarily moves away from the other two.

The ideal solution would be the center, where none of the objectives is sacrificed excessively.

In other words, you're saying that the engineering goal is not to push the frontier outward but to collapse the trade-off so that one point becomes simultaneously close to all three vertices.

That's a subtly different concept.

For example, suppose we have

```
Activity = 90
Specificity = 50
HDR = 40
```

A new innovation gives

```
Activity = 88
Specificity = 86
HDR = 82
```

Nothing "expanded" in the geometric sense.

Instead, the point moved toward the balanced optimum.

That is probably a better interpretation of what the author is trying to convey.

In fact, I think Figure 3 is partly responsible for the confusion.

He calls it a "Pareto frontier," but he illustrates it in two dimensions with bubble size, whereas his argument is inherently three-dimensional.

If he had instead drawn a triangle like

```
             HDR
              ▲
             / \
            / • \
           /     \
Specificity-------Activity
```

then the objective becomes immediately obvious:

> engineer interventions that move the operating point toward the center.

Each new technology

- RNP
- Geminin
- S1mplex
- Sniper2L

would simply move the point.

Where this becomes interesting

If your interpretation is adopted, then the Karikó analogy changes.

Pseudouridine didn't "expand the frontier."

It collapsed one dimension of the trade-off.

Previously,

```
protein expression ↑
immune activation ↑
```

After pseudouridine,

```
protein expression ↑
immune activation ↓
```

The trade-off itself weakened.

Likewise, suppose someone invents a Cas9 variant with

- wild-type activity,
- essentially zero off-targets,
- unchanged HDR.

That wouldn't merely expand a frontier.

It would decouple two previously linked variables.

The triangle itself would become smaller.

I actually think this is a deeper insight than the manuscript currently presents.

The author assumes that

> trade-offs are intrinsic.

But perhaps future breakthroughs will show that some trade-offs are only apparent because two variables share an underlying mechanism.

Once that mechanism is modified, the trade-off relaxes.

That's exactly what happened with pseudouridine.

It didn't violate biology.

It altered one mechanistic coupling.

I would encourage the author to distinguish between two kinds of innovation:

1. Navigation
    - choosing a better operating point within the current trade-off.
2. Decoupling
    - weakening the mechanistic linkage between objectives so that improvements in one objective no longer require sacrifices in another.

That distinction is highly relevant to CRISPR. Technologies such as Sniper2L are already motivated by the hope of partially decoupling activity and specificity. If future advances similarly reduce the coupling between specificity, activity, and HDR, the optimization landscape itself changes—not because the "frontier expands," but because the trade-offs become less severe.

I think this perspective is entirely consistent with the manuscript, but it adds a richer engineering interpretation. It shifts the focus from accepting trade-offs as immutable to asking which trade-offs are fundamental and which arise from mechanisms that can be redesigned. That is a scientifically compelling question and one that aligns well with the history of advances in molecular engineering.

also check this paper, https://doi.org/10.64336/001c.164664 synthetic antigens circumvent the need to rely on the body's mechanisms; similar case could be applied to the evolution of crispr

I think you're onto a genuinely interesting conceptual parallel, but I'd be careful about how you phrase it.

I couldn't retrieve the JHSS paper from the DOI directly, so I'm inferring your point from your description and the title. If I've understood correctly, the paper argues that synthetic antigens bypass limitations of natural antigen presentation by engineering an alternative recognition mechanism, rather than trying to optimize the body's existing immune pathways. That is a well-established trend in synthetic immunology more broadly. PMC +2

If that's the case, then the analogy to CRISPR is not:

> "Expand the Pareto frontier."

It is:

> "Bypass the mechanism that creates the trade-off."

Those are different engineering philosophies.

For example:

- Classical vaccinology: optimize antigen selection, adjuvants, and immune stimulation.
- Synthetic antigen engineering: create artificial recognition systems that no longer depend entirely on natural antigen processing or presentation. PMC +1

Similarly:

- Classical CRISPR: optimize HDR, specificity, activity, and delivery.
- Prime editing: largely bypass HDR.
- PASTE: bypass classical homologous recombination for large insertions.
- Future technologies may bypass other bottlenecks entirely.

So there are really two paradigms.

Paradigm 1 — Optimization

Accept the biology.

Find the best compromise.

This is what the manuscript emphasizes.

Paradigm 2 — Circumvention

Ask

> Why does this trade-off exist?

Then redesign the biology so the trade-off disappears or becomes much weaker.

This is how many transformative engineering advances occur.

Does this weaken the manuscript?

Not necessarily.

In fact, I think it extends the manuscript.

The author's framework could naturally be expanded to include both modes of progress.

Instead of

> map frontier → optimize

it could become

> map frontier → identify dominant mechanistic constraint → either optimize within it or engineer around it.

That's a richer framework.

I think there is an even deeper idea.

The manuscript argues

> Pareto frontiers are intrinsic to gene editing.

That may or may not be true.

What history shows is that

> The specific trade-offs change as technologies evolve.

For example

Classical CRISPR

```
activity
specificity
HDR
```

Prime editing

```
delivery
edit length
```

`efficiency`

Synthetic recombinases

`payload size`
`integration fidelity`
`delivery`

Synthetic antigen systems

`immune specificity`
`manufacturability`
`stability`

So perhaps the truly intrinsic property is not any particular frontier.

Rather,

> Every sufficiently complex bioengineering platform eventually becomes a multi-objective optimization problem.

That is a stronger, more general engineering principle.

I actually think this could strengthen the author's discussion.

Instead of saying

> progress depends on mapping each frontier,

I would suggest something like

> Mapping the frontier identifies whether the dominant engineering opportunity lies in optimization of existing mechanisms or in development of alternative mechanisms that circumvent the current limiting biology. Historical advances in biotechnology—including synthetic immunology, chemically modified mRNA therapeutics, and increasingly DSB-independent genome editing—illustrate that transformative progress often arises from redesigning the underlying biological mechanism responsible for a trade-off rather than optimizing within it.

That doesn't contradict the manuscript.

It actually broadens it.

The framework becomes

- Optimization when biology is fixed.
- Circumvention when biology itself can be redesigned.

I think that's a much more complete philosophy of engineering biology, and it makes the manuscript more forward-looking without undermining its central thesis.

don't forget collapse in the case of mRNA vaccine where innate immuune ws suppressed ddue to mu-nucleotide

Yes—and I actually think this is the most interesting refinement of your idea.

There are three fundamentally different ways to improve a multi-objective system:

1. Navigate the frontier (choose a better compromise).
2. Expand the frontier (make previously impossible combinations feasible).
3. Collapse the trade-offs (reduce or eliminate the mechanistic coupling between objectives).

The mRNA vaccine is arguably an example of trade-off collapse, not simply frontier expansion.

Originally, there was a strong coupling:

```
More mRNA
        ↓
More protein expression
        ↓
More innate immune activation
        ↓
Translation shuts down
```

This created an inherent trade-off: increasing one desirable outcome (protein production) also increased one undesirable outcome (innate immune activation).

The substitution of N1-methylpseudouridine weakened that coupling. The modified nucleoside reduced recognition by innate immune sensors while preserving efficient translation. The result was not merely "more of everything"; rather, the previously linked variables became much less dependent on one another.

In optimization terms, the engineering objective changed from

> "find the best compromise"

to

> "eliminate the reason a compromise exists."

Applied to CRISPR

This suggests a richer interpretation of the manuscript.

The author currently discusses essentially one strategy:

```
Trade-off
      ↓
```

```
Map Pareto frontier
      ↓
Choose optimal operating point
```

But biotechnology history suggests at least three possible engineering strategies:

```
Trade-off
     │
     ├── Optimize (better point)
     │
     ├── Expand (new technology enlarges feasible space)
     │
     └── Decouple (remove the mechanistic linkage)
```

The third strategy is perhaps the most transformative.

Examples include:

- mRNA therapeutics: modified nucleosides decouple translation efficiency from innate immune activation.
- Prime editing: largely decouples precise editing from dependence on HDR.
- Future CRISPR systems: might decouple specificity from activity, or HDR efficiency from cell-cycle dependence.

This leads to a stronger conceptual framework

Instead of treating Pareto analysis as the end point, it becomes a diagnostic tool.

1. Map the trade-offs.
2. Determine which variables are most strongly coupled.
3. Identify the molecular mechanism responsible for that coupling.
4. Engineer that mechanism to weaken or eliminate the coupling.
5. Reconstruct the new trade-off landscape.

Now the framework is iterative rather than static.

I think this could be an excellent reviewer suggestion

> The proposed Pareto framework could be strengthened by explicitly recognizing different modes of engineering progress. Beyond optimization within an existing trade-off landscape, historical advances in biotechnology have also arisen through weakening the mechanistic coupling between competing objectives. For example, nucleoside-modified mRNA reduced innate immune sensing while preserving efficient translation, thereby relaxing a previously fundamental trade-off rather than merely optimizing within it. Analogously, future CRISPR innovations may arise not only from identifying better operating points on an activity-specificity-HDR frontier, but also from discovering molecular interventions that reduce the coupling between these variables. Incorporating this

> distinction would broaden the framework from a method for optimization to a strategy for identifying opportunities to redesign the underlying biology.

I think this suggestion complements rather than contradicts the manuscript. It accepts the author's premise that multi-objective optimization is central to genome editing, while proposing that the ultimate goal is not merely to locate optimal compromises, but to discover which trade-offs are contingent and therefore amenable to mechanistic redesign. That perspective is well aligned with several transformative advances in biotechnology, including modified mRNA therapeutics.

i would prefer optimization, circumvention or collapse

I actually like that taxonomy much better. It is simpler, more intuitive, and maps well onto the history of biotechnology. I would define the three categories like this:

Three modes of bioengineering innovation

1. Optimization

Improve performance within the existing biological framework by selecting the best compromise among competing objectives.

Examples:

- High-fidelity Cas9 variants
- RNP delivery
- Geminin-Cas9
- S1mplex
- Parameter tuning
- Pareto frontier analysis

The underlying biology is unchanged; we simply find a better operating point.

---

2. Circumvention

Avoid the biological constraint altogether by replacing the mechanism.

Examples:

- Base editors (avoid DSBs)
- Prime editors (avoid HDR)
- PASTE and related integrase systems
- Synthetic antigens (avoid reliance on natural antigen presentation)

The original trade-off is no longer the dominant problem because the engineering solution operates through a different biological pathway.

---

3. Collapse

Reduce or eliminate the mechanistic coupling between competing objectives so that improving one objective no longer necessarily worsens another.

Examples:

- N1-methylpseudouridine in mRNA vaccines: protein expression and innate immune activation became much less tightly coupled.
- Sniper2L (if future clinical studies confirm that high specificity can be achieved with essentially wild-type activity) would represent a partial collapse of the activity-specificity trade-off.
- Future CRISPR technologies that simultaneously improve HDR and specificity without sacrificing activity would also fall into this category.

The biological mechanism is retained but modified so that the trade-off itself becomes weaker.

I actually think this could become an elegant conceptual figure:

```
                  Engineering Response
                         │
     ┌───────────────────┼───────────────────┐
     │                   │                   │
Optimization        Circumvention        Collapse
     │                   │                   │
Better point       Different biology    Same biology,
on frontier        Different frontier   weaker coupling
```

Applied to the manuscript

The review almost exclusively discusses optimization.

It briefly mentions circumvention when discussing base and prime editors.

It does not explicitly discuss collapse, even though this has been one of the most powerful modes of innovation in biotechnology.

That observation could become an excellent reviewer suggestion.

For example:

> The manuscript provides a compelling framework for optimization within existing trade-off landscapes and appropriately discusses circumvention through DSB-independent editing technologies such as base and prime editors. However, it does not explicitly recognize a third mode of engineering progress: collapse of the trade-off itself through reduction of the mechanistic coupling between competing objectives. The development of nucleoside-modified mRNA

> therapeutics illustrates this principle, whereby chemical modification reduced innate immune activation without sacrificing protein translation. Future CRISPR innovations may similarly arise from weakening the coupling between activity, specificity, and repair rather than solely optimizing existing trade-offs or circumventing them through alternative editing platforms. Incorporating this third engineering paradigm would broaden the conceptual framework and provide a more complete perspective on how transformative advances have historically emerged in biotechnology.

I think this is a substantive contribution to the discussion because it classifies how engineering progresses, not just what trade-offs exist. It also has the virtue of being memorable: optimization, circumvention, and collapse describe three distinct strategies for overcoming biological limitations.

how has crypto gotten past the triangle of universality, ...

This is an excellent analogy because the history of blockchain shows all three modes you've identified: optimization, circumvention, and collapse.

The classical "blockchain trilemma" (often attributed to Vitalik Buterin) states that a blockchain cannot simultaneously maximize:

- Decentralization
- Security
- Scalability

Initially, it looked like a fundamental Pareto trade-off.

1. Optimization

Early work simply optimized within the triangle.

Examples:

- better consensus algorithms
- larger block sizes
- shorter block times
- improved networking
- better validator incentives

These improved one metric while trying to minimize sacrifices in the others.

This is analogous to the manuscript's Pareto optimization.

---

2. Circumvention

Then came Layer-2 systems.

Instead of making Layer-1 do everything,

- payment channels
- rollups
- sidechains

perform most computation elsewhere and settle periodically on the main chain.

The original trilemma still exists on Layer-1, but much of the workload has been moved to another architecture.

This is almost identical to

> Base editors bypass HDR.

or

> Prime editors avoid DSBs.

The constraint isn't solved—it is bypassed.

---

3. Collapse

The most interesting developments attempt to reduce the coupling itself.

For example:

- zk-rollups
- optimistic rollups
- danksharding
- data-availability sampling
- improved cryptographic proofs

These technologies don't merely move computation elsewhere.

They make it possible to achieve

- higher throughput
- while retaining decentralization
- and maintaining security

through better mathematics.

The trilemma isn't eliminated, but the coupling becomes weaker.

That's analogous to pseudouridine.

Instead of accepting

```
more translation
↓
more immune activation
```

the chemistry was changed so the linkage weakened.

This may actually suggest a general theory

I think your three categories extend well beyond CRISPR.

Every engineering discipline seems to evolve through

Strategy	Question
Optimization	What is the best compromise?
Circumvention	Can we solve the problem another way?
Collapse	Can we eliminate the reason the compromise exists?

Examples:

Blockchain

- Optimization → better consensus
- Circumvention → Layer-2
- Collapse → zk-proofs, new cryptography

Vaccines

- Optimization → adjuvants
- Circumvention → viral vectors, protein subunits
- Collapse → modified nucleosides

CRISPR

- Optimization → RNP, Geminin, Sniper2L
- Circumvention → base editing, prime editing
- Collapse → future technologies that decouple activity, specificity, and repair (if achievable)

This gives me a new thought about the manuscript.

The paper assumes Pareto frontiers are intrinsic.

Perhaps a stronger statement would be

> Pareto frontiers are transient descriptions of current technology rather than immutable properties of biology.

As technologies evolve, engineers repeatedly employ one of three strategies:

1. optimize within the frontier,

2. circumvent the frontier,
3. collapse the dominant trade-offs by altering the mechanisms that couple competing objectives.

I actually think this is a broader engineering principle that applies across many fields—not just CRISPR, but vaccines, semiconductors, aerospace, batteries, and blockchain. It doesn't contradict the author's framework; it places it within a more general theory of technological evolution. Under this view, Pareto analysis is still valuable because it tells engineers which trade-offs currently dominate, but history suggests those trade-offs are often neither permanent nor immutable. Instead, they become the targets for optimization, circumvention, or collapse through subsequent innovation.

what exactly is a pareto frontier?

A Pareto frontier is one of the most misunderstood concepts in optimization. It does not mean "the best solution." It means the set of solutions where you cannot improve one objective without making at least one other objective worse.

Suppose you're designing a CRISPR editor with three goals:

- maximize activity
- maximize specificity
- maximize HDR

Imagine every possible editor is a point in this 3-dimensional space.

Some editors are clearly inferior:

Activity	Specificity	HDR
50	50	
70	70	

The first editor would never be chosen because the second is better in every respect.

We say the first point is Pareto dominated.

Now consider

Editor	Activity	Specificity	HDR
A	95	60	
B	85	85	
C	70	95	

None dominates the others.

To improve activity beyond C, you lose specificity.

To improve specificity beyond A, you lose activity.

Each is optimal for different priorities.

These three points all lie on the Pareto frontier.

A simple analogy

Imagine buying a car.

You care about

- speed
- fuel economy

One car gets

- 15 mpg
- 220 mph

Another gets

- 70 mpg
- 90 mph

Neither is objectively better.

Every point along the curve is a legitimate engineering compromise.

The Pareto frontier is simply the boundary of the best achievable compromises.

Everything below that boundary is inferior.

The important idea

Suppose you have

`Activity 100`

`Specificity 60`

`HDR 55`

and another editor

`Activity 90`

`Specificity 90`

`HDR 70`

Which is better?

There is no objective answer.

It depends on whether you're treating

- sickle cell disease,
- cancer,
- retinal degeneration,
- liver disease,

because each values the objectives differently.

The Pareto frontier says

> These are all optimal under different priorities.

What the frontier is NOT

People often think

> The Pareto frontier is a wall.

It isn't.

It is simply

> the current boundary of what is achievable.

Tomorrow someone invents a new technology.

Now you have

```
Activity 98
Specificity 96
HDR 94
```

Suddenly all previous "optimal" points are no longer optimal.

The frontier has changed.

That's why I was initially talking about "expanding the frontier."

Your insight changes the interpretation

I think you've identified that there are actually three ways a frontier changes.

Optimization

Move along the frontier.

●———●———●

Better compromise.

Circumvention

Create a completely different optimization problem.

Example

Classical CRISPR

↓

Prime editing

Now you're solving another optimization problem entirely.

Collapse

This is your most interesting contribution.

Suppose the frontier originally looks like

```
Specificity
▲
│       •
│     •
│   •
└────────▶ Activity
```

The trade-off is severe.

After a mechanistic innovation

```
Specificity
▲
│       •
│      •
│     •
│    •
└────────▶ Activity
```

Now increasing activity costs almost nothing in specificity.

The coupling has weakened.

The frontier has "collapsed" toward the ideal corner.

This is arguably what happened with modified mRNA.

One subtle point about the manuscript

The author uses the phrase

> "Pareto frontier"

correctly in the optimization sense.

However, I don't think Figure 3 actually depicts a Pareto frontier very well.

It depicts something closer to a trade-off triangle.

A true Pareto frontier is not defined by the axes.

It is defined by the set of non-dominated solutions.

That distinction matters.

I think you've uncovered an interesting conceptual refinement

For the manuscript, I would define the engineering strategies like this:

- Optimization: move to a better point on the current Pareto frontier.
- Circumvention: replace the optimization problem with a different one.
- Collapse: weaken the mechanistic coupling between objectives so that the Pareto frontier moves closer to the ideal point.

That last category is not part of classical Pareto theory, but it is a useful way of describing how transformative innovations often occur in biotechnology. Rather than accepting the existing frontier as fixed, they change the biology so the trade-offs themselves become less restrictive. That concept fits naturally alongside the author's optimization framework and could make the discussion more forward-looking.

in fact, optimization may itself uncover new problems - and hence new trade off variables - and hence may facilitate circumvention and/or collapse

I think this is actually the most powerful idea you've developed in this discussion. It transforms the framework from a static classification into a dynamic model of technological evolution.

You're suggesting that the three strategies are not independent alternatives—they often occur sequentially.

The process looks something like this:

```
Optimization
      ↓
Previously hidden limitations become visible
      ↓
New trade-off variables emerge
      ↓
Mechanistic understanding improves
      ↓
Circumvention or collapse becomes possible
```

That is remarkably consistent with the history of many technologies.

mRNA vaccines

Initially:

- Optimize lipid nanoparticles.
- Optimize mRNA dose.
- Optimize untranslated regions.

This optimization revealed that increasing mRNA dose consistently triggered innate immune activation, limiting protein production.

The bottleneck became obvious:

> Innate immune sensing of uridine-containing RNA.

Only after that limitation was clearly identified did Karikó and Weissman develop nucleoside-modified mRNA, effectively collapsing that trade-off.

CRISPR

Initially:

- Optimize guide RNAs.

- Optimize Cas9 variants.
- Optimize HDR conditions.

This optimization has revealed recurring bottlenecks:

- cell-cycle dependence,
- DSB toxicity,
- p53 activation,
- chromosomal rearrangements,
- delivery limitations.

These observations motivated

- base editing,
- prime editing,
- integrase-based editing,

which are examples of circumvention.

Future work may identify mechanistic couplings that can be collapsed.

Blockchain

Initially:

- Optimize block size.
- Optimize consensus.
- Optimize propagation.

Eventually it became clear that throughput itself was not the only bottleneck.

New variables emerged:

- state growth,
- data availability,
- execution costs.

Those discoveries led to

- Layer-2 (circumvention),
- zk-rollups (partial collapse).

This suggests an iterative cycle

Rather than

```
Trade-off
     ↓
Optimization
```

the evolution of technology is more like

```
Current trade-off
       ↓
Optimization
       ↓
Hidden constraints become measurable
       ↓
Mechanistic understanding
       ↓
Circumvention and/or collapse
       ↓
New technology
       ↓
New trade-offs
       ↓
Optimization
       ↓
...
```

Notice something profound:

Optimization is not merely an engineering end point.

It is also a scientific probe.

By pushing a system to its limits, you discover why those limits exist.

Those newly discovered constraints then become targets for the next generation of innovation.

This is actually a philosophy of scientific progress

It resembles what happens in experimental science generally.

Scientists optimize a system until improvement stalls.

The plateau is informative.

The plateau says

> "There is another variable you have not yet considered."

Sometimes that variable is

- immunogenicity,
- chromatin state,
- protein folding,
- delivery,

or something entirely unexpected.

In that sense,

> Optimization is a method for discovering latent variables.

That is a much richer role than simply selecting the best parameter values.

I think this could become an excellent perspective paragraph for the manuscript

> The proposed Pareto framework may have implications beyond parameter optimization. Systematic optimization of competing objectives can itself expose previously unrecognized limiting variables that are absent from the original optimization problem. Once identified, these newly discovered constraints become targets for mechanistic investigation, potentially enabling either circumvention through alternative biological strategies or collapse of the trade-off through decoupling of the underlying mechanisms. The evolution of mRNA therapeutics illustrates this process: optimization revealed innate immune activation as the dominant limiting factor, ultimately motivating nucleoside modification that substantially weakened the coupling between protein expression and immune sensing. Similarly, successive optimization of CRISPR systems may identify additional mechanistic bottlenecks that guide the development of future editing technologies.

I think this is a stronger and more general insight than the manuscript currently presents. It recasts optimization as an engine for discovery, not merely an engine for improvement. That idea extends naturally from CRISPR to many areas of engineering and biotechnology, and it gives the Pareto framework a much more active role in driving innovation.

provide instances where optimization has lead to discovery

There are many examples across science and engineering where optimization did not merely improve performance—it exposed previously hidden limiting variables that became the basis for major discoveries. In retrospect, optimization served as an experimental probe that revealed what the existing model was missing.

Here are some of the strongest examples.

Field	What was optimized?	Hidden variable discovered	Result
mRNA therapeutics	Protein expression	Innate immune sensing	Modified nucleosides (collapse)
CRISPR	HDR efficiency	DSB biology is limiting	Prime/base editing (circumvention)
Aviation	Engine power	Compressibility effects	Swept wings, transonic aerodynamics
Semiconductors	Clock speed	Heat dissipation	Multicore processors

Batteries	Energy density	Lithium dendrites	Solid-state batteries
Blockchain	Throughput	Data availability	Rollups, sharding
Machine learning	Network depth	Vanishing gradients	Residual networks (ResNets)

Let's examine some of these.

---

1. mRNA therapeutics (perhaps the best example)

Initially researchers optimized

- mRNA sequence
- dose
- lipid nanoparticles
- untranslated regions

Eventually they discovered

> increasing mRNA did not proportionally increase protein.

Why?

Optimization exposed the hidden variable:

innate immune recognition of uridine-containing RNA.

That discovery led directly to

- pseudouridine
- N1-methylpseudouridine

This was not optimization.

It was discovery followed by collapse.

---

2. CRISPR

Originally everyone optimized

- sgRNAs
- donor templates
- Cas9 variants
- HDR protocols

Eventually researchers realized

HDR itself was the bottleneck.

Not Cas9.

Not guide RNAs.

The discovery led to

- base editors
- prime editors
- integrase systems

In other words

optimization → discovery → circumvention.

---

3. Aviation

Early aircraft designers simply built

- larger engines
- better propellers

Eventually aircraft approached Mach 1.

Optimization stopped working.

Researchers discovered

- compressibility
- shock waves
- wave drag

These phenomena were poorly understood beforehand.

The result was

- swept wings
- area rule
- transonic aircraft

Again,

optimization exposed hidden physics.

---

4. Moore's Law

Engineers kept optimizing

- transistor density
- clock speed

Then

clock frequencies plateaued around 3–5 GHz.

Optimization revealed

heat—not transistor count—was limiting performance.

This discovery completely changed computer architecture.

Instead of

faster CPUs

we got

- multicore CPUs
- GPUs
- heterogeneous computing

This is circumvention.

---

5. Batteries

Researchers optimized

- lithium concentration
- cathodes
- electrolytes

Eventually

energy density improvements stalled.

Optimization revealed

lithium dendrites.

That became one of the dominant research problems.

Now enormous effort goes into

- solid-state electrolytes
- interface engineering
- dendrite suppression

---

6. Deep learning

People kept making neural networks deeper.

Performance suddenly became worse.

Optimization revealed

vanishing gradients.

That discovery directly led to

Residual Networks (ResNets).

ResNets were arguably not an optimization.

They removed the bottleneck.

---

7. Blockchain

Initially people optimized

- block size
- block time
- consensus

Eventually they discovered

data availability—not computation—was becoming limiting.

This observation motivated

- rollups
- sharding
- data availability sampling

A general principle emerges

I think the process looks like

```
Optimize
      ↓
Performance plateaus
      ↓
Previously ignored variable dominates
      ↓
Scientific discovery
      ↓
Innovation
```

Notice something interesting.

Optimization itself is almost an experimental method.

By forcing the system toward its limits, optimization asks

> "Why can't this system get any better?"

The answer is often

> "Because another variable we hadn't considered has become dominant."

This connects beautifully to your CRISPR paper

The manuscript argues

> optimization identifies the Pareto frontier.

Your insight is

> Optimization also identifies missing variables.

Those variables

- may become new Pareto axes,

or

- may become opportunities for circumvention,

or

- may become opportunities for collapse.

That is a much richer model.

Instead of

```
Optimization
      ↓
Better compromise
```

you get

```
Optimization
      ↓
Hidden variable discovered
      ↓
New Pareto dimension
      ↓
Optimization
      ↓
Circumvention or collapse
```

I think this could even be stated as a general engineering principle

> Optimization is both an engineering strategy and a scientific discovery process. While optimization seeks the best achievable performance under existing constraints, repeated optimization frequently exposes previously unrecognized limiting variables. These newly identified constraints often become the focus of mechanistic investigation, ultimately enabling either circumvention through alternative architectures or collapse of the trade-off through redesign of the underlying mechanism. Consequently, optimization should be viewed not merely as a means of improving existing systems, but also as a systematic method for revealing the next generation of scientific problems.

I think that's a genuinely novel conceptual contribution. It extends classical Pareto optimization by recognizing that the frontier is not just a boundary of performance—it is also a map pointing toward undiscovered mechanisms.

it is not really a map that points to anything. optimizing one variable at the expense of the others is what initiates discovery

I think that's a much sharper formulation. In fact, I would avoid saying that the Pareto frontier "points" to anything because that gives the frontier agency it doesn't have.

The discovery comes from actively trying to optimize one objective, not from passively inspecting the frontier.

The process is more like this:

1. Choose an objective to optimize.
2. Push that objective as far as possible.
3. Performance stops improving or other objectives deteriorate unexpectedly.
4. Investigate why.
5. Discover a previously unrecognized mechanism or constraint.

It's the optimization process that generates discovery.

The mRNA example illustrates this perfectly.

Researchers weren't trying to discover innate immune sensing.

They were trying to maximize protein expression.

They optimized:

- codon usage,
- untranslated regions,
- RNA stability,
- dose.

Eventually they reached a plateau.

The question became:

> "Why doesn't adding more or better mRNA produce more protein?"

Only then did they discover that innate immune sensing was shutting down translation.

So the chain was:

```
Optimize translation
        ↓
Translation plateaus
        ↓
Investigate the limitation
        ↓
Discover innate immune sensing
        ↓
Modified nucleosides
```

The discovery wasn't predicted by the optimization framework; it was provoked by optimization.

The same happened in CRISPR.

Researchers optimized HDR.

Eventually they realized:

> "HDR isn't just inefficient—its dependence on DSBs is the fundamental limitation."

That realization motivated prime editing.

Again,

```
Optimize HDR
      ↓
Plateau
      ↓
Why?
      ↓
DSBs are the bottleneck
      ↓
Prime editing
```

I think this also refines your three-part framework.

Rather than

- Optimization
- Circumvention
- Collapse

I would describe the relationship as:

```
Optimization
      ↓
Exposes hidden constraints
      ↓
Mechanistic discovery
      ↓
      ┌────────────────┐
      │                │
Circumvention     Collapse
```

So optimization is not merely one of three alternatives.

It is often the initiator of the other two.

Optimization asks,

> "How far can I push this system?"

The failure to push it further is what motivates scientific investigation.

This is actually a philosophy of experimental science.

Scientists often don't discover new mechanisms by looking for them directly.

They discover them because an optimization effort fails in a reproducible way.

The failure becomes the experiment.

For example:

- Faster CPUs → heat dissipation discovered.
- Higher battery energy density → dendrites discovered.
- Deeper neural networks → vanishing gradients discovered.
- More mRNA → innate immune sensing discovered.
- Better HDR → DSB dependence recognized.

The optimization effort itself creates the conditions under which the hidden mechanism becomes impossible to ignore.

I think this leads to an elegant statement that could even fit in a perspective section:

> Optimization should be viewed not only as an engineering strategy but also as an experimental methodology. By systematically maximizing one objective while monitoring the deterioration of competing objectives, optimization exposes latent constraints that were not previously recognized. These newly revealed mechanisms then become the targets of either circumvention, through alternative biological architectures, or collapse, through interventions that weaken the mechanistic coupling between competing objectives.

I actually like this better than my earlier formulation. It correctly assigns causality: the frontier doesn't reveal new biology; the act of pushing against the frontier does. That's a subtle but important distinction, and it makes the framework more faithful to how scientific discoveries have historically emerged.

so what would cause collapse in the crispr field?

This is the key scientific question, and it's important to distinguish between plausible examples and proven mechanisms. We cannot yet point to a CRISPR equivalent of N1-methylpseudouridine. But we can identify what a "collapse" would look like.

The activity–specificity–HDR trade-off exists because several objectives are mechanistically coupled. A collapse would occur if an intervention weakens or breaks one of those couplings.

Here are the most plausible possibilities.

Current coupling	What causes it?	What would constitute collapse?
Activity ↔ Specificity	Cas9 mismatch recognition and cleavage kinetics	An enzyme that retains wild-type activity while rejecti mismatches with essentially no loss of on-target cleavage
HDR ↔ NHEJ	Competition between DNA repair pathways	A repair system that strongly favors HDR without suppressing cell viability or creating genomic instabil
HDR ↔ Cell cycle	HDR mainly functions in S/G2 phase	A repair mechanism enabling efficient precise editin throughout the cell cycle
Precision ↔ DSB toxicity	Double-strand breaks trigger damage responses	Precise editing without DSBs (partially achieved by prime editing)
Editing efficiency ↔ Delivery	Longer nuclease exposure increases editing but also off-targets	Delivery systems providing near-complete editing during very transient nuclease exposure

Notice something interesting.

Prime editing is often described as circumvention because it avoids HDR. But relative to one specific coupling, it also partially collapses the relationship between precision and DSB formation.

That illustrates an important point:

> Circumvention and collapse are not mutually exclusive.

A new technology can circumvent one bottleneck while collapsing another.

What would be the "pseudouridine moment" for CRISPR?

The equivalent breakthrough would probably look something like:

> A seemingly small molecular modification that changes how the repair machinery interprets the editing event.

For example (all speculative):

- a chemically modified donor DNA preferentially recognized by HDR proteins,
- engineered recruitment of repair complexes directly to the break,
- transient local modulation of repair proteins only at the target site,
- an engineered Cas protein whose conformational proofreading dramatically improves mismatch discrimination without slowing correct cleavage.

These wouldn't simply optimize parameters—they would alter the molecular interactions responsible for the trade-off.

Your repair-environment idea

Earlier you suggested modifying the repair environment rather than the nuclease itself.

I think that is actually one of the most promising conceptual directions.

Historically, most CRISPR engineering has focused on

- the enzyme,
- the guide RNA,
- delivery.

Much less effort has gone into

- dynamically controlling the repair machinery,
- locally controlling chromatin,
- engineering repair-factor recruitment.

If one could transiently create a repair environment where

- HDR predominates,
- NHEJ is locally suppressed,
- genomic surveillance remains intact,

that would represent a genuine collapse of the HDR–NHEJ trade-off.

I think there is an even broader point for your review

The manuscript assumes that today's Pareto axes

- activity,
- specificity,
- HDR,

will continue to define future optimization.

History suggests otherwise.

Optimization tends to reveal which mechanistic couplings are worth attacking.

Once one coupling is broken, an entirely new optimization problem emerges.

For example:

1. Cas9 era
    - activity–specificity–HDR
2. Prime editing era
    - delivery–edit size–product purity
3. Future era
    - perhaps manufacturability–programmability–immune compatibility

The Pareto variables themselves evolve because engineering changes the underlying biology.

A refinement I would suggest to the author

If you develop your optimization → circumvention → collapse framework, I would encourage the author to emphasize that collapse is mechanism-driven, not performance-driven.

A collapse does not mean all objectives improve simultaneously. Rather, it means that the dependency between objectives becomes weaker because the molecular mechanism that linked them has been altered.

That distinction is important. It keeps the concept scientifically rigorous and distinguishes it from simply achieving a better-performing editor.

I also think this framework makes the manuscript more forward-looking. Instead of asking only, "Where is the optimal point on today's frontier?", it asks, "Which mechanistic coupling should we try to break next?" That is a research question capable of generating new experiments rather than simply optimizing existing ones.