Episode Book

EPISODE #6687

How CRISPR Finds One Letter in Three Billion

PUBLISHED	19 November 2026 at 21:29 UTC
DURATION	31:37
LISTEN	myweirdprompts.com/episode/crispr-gene-editing-mechanism-15463/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	z-ai/glm-5.3
VOICES	omnivoice
PIPELINE	V5.36
SOURCES	18 cited by the research (listed at the end)
BOOK GENERATED	19 November 2026 at 22:18 UTC

SYNOPSIS

Half of sequenced bacterial genomes carry CRISPR, an ancient immune system that files viral DNA like wanted posters. This walk-through traces how researchers turned that bacterial defense into a programmable gene-editing tool, from Philippe Horvath's yogurt experiments at Danisco to Jennifer Doudna and Emmanuelle Charpentier fusing two RNAs into a single guide. It also breaks down the step-by-step search mechanism—how Cas9 scans for a three-letter PAM motif to find one target among three billion letters—and why base and prime editing move beyond simple DNA scissors.

TRANSCRIPT



Corn

The most important editing tool in biology was stolen. From bacteria. They'd been using it for millions of years, and we only noticed in the last couple of decades.



Herman

And once we noticed, we did the very human thing — took it apart, simplified it, and put it to work on ourselves.



Corn

Daniel's prompt this week is a mechanism question, and it's aimed squarely at my brother. He wants CRISPR explained in relatable terms — how a bacterial immune system becomes a tool that can find one exact spot in a genome of three billion letters, and either cut it or rewrite it. He wants the walk-through, step by step, and then the part people skip: how base editing and prime editing improve on the original scissors.



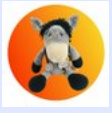
Herman

Which is the part I've been looking forward to, because the scissors framing is exactly what everybody gets wrong.



Corn

Before we go anywhere — give me the scale of the needle in the haystack, because I think people hear "three billion letters" and it doesn't land.



Herman

Here's an image for it. If you printed the human genome out at a normal font, a few thousand letters to a page, you'd have a stack of paper roughly the height of a two-story building. The edit you want is one letter on one page. And the tool has to find it, in the nucleus of a cell, in a fraction of that cell's lifetime, without reading every page.



Corn

And it does.



Herman

It does, with a trick we'll get to. But it's worth holding that image, because everything elegant about CRISPR is in service of that one search problem.



Corn

So let's start with the bacteria, because that's where the whole thing comes from. Herman, you're a retired pediatrician. Before you were editing anything, you were prescribing things. Was CRISPR on your radar in practice at all?



Herman

Not remotely. When I was in clinics, CRISPR was a microbiology curiosity. The clinical arrival is recent — the first CRISPR-based drug, Casgevy, was approved by the FDA at the end of twenty twenty-three, for sickle cell disease and beta thalassemia. Everything before that was bench work. And the bench work started with a odd observation.



Corn

Which was?



Herman

In the late eighties, researchers sequencing bacterial genomes kept finding these strange repeated stretches of DNA. The same short sequence, then a spacer of different DNA, then the repeat again, then another spacer — over and over, like a filing cabinet. Nobody knew what it was for. It looked almost decorative.



Corn

Bacteria don't decorate.



Herman

They don't. It took the better part of two decades to work out what those repeats were, and the answer is wonderful: it's a memory of past viral infections. The name itself is a description of the structure — clustered regularly interspaced short palindromic repeats. The repeats are the file dividers. Each spacer between them is a snippet of DNA from a bacteriophage, a virus that attacked that bacterium or one of its ancestors.



Corn

Hold on. The bacterium keeps a souvenir of its attacker.



Herman

Better than a souvenir — a wanted poster. When the same phage comes back, the bacterium transcribes that spacer into a small RNA, and that RNA guides a cutting protein to any DNA that matches it. Match found, DNA destroyed. It's adaptive immunity. Acquired immunity. A bacterium that survives an infection is immune to that phage, and so are its descendants, because the memory is written into the genome itself and inherited.



Corn

That's the part that should stop people. We teach that immune memory needs a whole adaptive immune system, white blood cells, antibodies — vertebrate machinery. Half the bacteria on earth have a simpler version of the same idea. Roughly half of sequenced bacterial genomes carry CRISPR, and in archaea it's nearly ninety percent. This is not exotic.



Herman

And think about the world they live in. There are phages everywhere — seawater, soil, your gut — and by some estimates they vastly outnumber bacteria. A bacterium without a memory system is living in a war zone with no intelligence files. The ones that kept files survived to be our ancestors' ancestors.



Corn

So this has been running continuously since before there were multicellular organisms.



Herman

Effectively, yes. It's one of the oldest immune strategies on the planet, and we catalogued it before we understood it. That happens more in biology than people think — the structure sat in sequence databases for years, filed under "repeats of unknown function," before anyone asked the right question about it.



Corn

And the proof came from a yogurt company, which I love.



Herman

Two thousand seven, Philippe Horvath, working at Danisco — a dairy company, in France. Their yogurt cultures were getting killed by phages, which is a real commercial problem — your entire product is a bacterial culture, and a phage is a predator of exactly that culture — and Horvath's team showed experimentally that bacteria exposed to a phage would file new spacers from that phage into their CRISPR array, and those bacteria then resisted the next wave of the same phage. The ones that didn't file a spacer, didn't.



Corn

The immune system was confirmed by someone protecting yogurt.



Herman

Industrial microbiology has quietly powered a lot of this field. Cheese and yogurt makers have been managing microbial populations, without knowing the genetics, for centuries. And the mechanism has two halves, and it's worth keeping them separate because the engineering only borrowed one of them.



Corn

Say them.



Herman

Adaptation and interference. Adaptation is the filing — grab a piece of invading phage DNA, slot it into the array as a new spacer. Interference is the retaliation — the whole array gets transcribed, chopped into individual CRISPR RNAs, and each one loads onto a Cas protein, which then patrols for matching DNA and cuts it. The Cas protein is the weapon. The RNA is the targeting system.



Corn

And when the engineers came in, which half did they take?



Herman

Only the interference half. The filing system — the piece that lets a bacterium update its own memory — nobody needed. In the lab, you write the memory yourself. You synthesize the guide RNA directly. The bacteria spent millions of years evolving a way to acquire new targets automatically, and we looked at it and said, we'll just type the address in.



Corn

Which is either flattering or rude to the bacteria, depending on your perspective.



Herman

It's the history of biotechnology in one gesture.



Corn

And the weapon we care about is Cas9.



Herman

Cas9, from Streptococcus pyogenes, the same bacterium that gives you strep throat. It works with two small RNAs — the CRISPR RNA and another one called the tracrRNA — so the natural complex is four components: the Cas9 protein, two RNAs, and the target DNA.



Corn

Four parts. In a lab that's three things that can go wrong before you've done anything.



Herman

Which is exactly what the twenty twelve work by Jennifer Doudna and Emmanuelle Charpentier fixed, and I think it's the most underrated moment in the whole story. Everyone points at the Nobel in twenty twenty, and that was for this — but the detail that changed everything is almost technical trivia. They fused the two RNAs into one molecule, a single-guide RNA.



Corn

Why does fusing two RNAs into one change the world?



Herman

Because it turned a four-part bacterial machine into a one-part programmable reagent. You want Cas9 to cut somewhere new? You don't redesign the protein. You synthesize a different short RNA — and synthesizing short RNA was already a routine, order-it-by-email procedure in every molecular biology lab on earth. The protein never changes. The address changes.



Corn

Walk me through what that meant day to day for a working lab.



Herman

Before, if you wanted to test an edit in ten different places, you had ten engineering problems. After, you had one protein sitting in a freezer and ten oligos in an envelope from a supplier. The barrier between "having an idea about a gene" and "testing that idea at the bench" collapsed from months to days. That's not an incremental improvement — that's the difference between a handful of specialized labs doing editing and every molecular biology lab on the planet doing it.



Corn

That's the whole revolution compressed into one sentence. The targeting information moved from protein to nucleic acid.



Herman

And that's why it steamrolled the older tools. Zinc finger nucleases and TALENs — the tools before CRISPR — could also target cuts to specific places, but each one required designing and building a custom protein for each target sequence. Months of protein engineering, per edit. CRISPR requires you to order a piece of RNA. It's the difference between commissioning a bespoke suit and printing a t-shirt.



Corn

You said once that CRISPR turned gene editing from a bespoke tailoring operation into a print-on-demand service. I remember because it was better than your usual output.



Herman

I stand by it. And it scales the way printing scales — you can target several sites at once just by delivering several guide RNAs with the same Cas9. Try doing that with zinc fingers, where every target is a fresh protein.



Corn

Fine. Walk me through the mechanism as it happens in a cell, in order. Daniel asked for the step-by-step and I want it clean.



Herman

Step one, delivery. You get the Cas9 protein and its guide RNA into a cell — that's an entire field of its own and we'll come back to it. Step two, scanning. Cas9 doesn't read the genome letter by letter looking for the guide's match. It first looks for something much shorter — a little DNA motif called the PAM, the protospacer adjacent motif. For the standard pyogenes Cas9 it's just three letters. Cas9 bumps along the DNA and only stops where it finds that motif.



Corn

How much does that shortcut actually save? You said the genome is a two-story building of paper.



Herman

It's the difference between reading every page and only opening pages with a colored tab on them. A three-letter motif will occur, statistically, every few hundred letters or so. So instead of interrogating three billion positions, Cas9 only fully checks a tiny fraction — the places where the keyword appears. It's a two-stage filter: cheap and indiscriminate first, expensive and precise second. Every fast search system works this way — that's roughly how you find a word in a document too. You don't read the paragraph; you scan for the shape of the word.



Corn

Why would an immune protein need a keyword before it checks the match? It has the guide. The guide is specific.



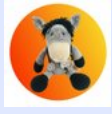
Herman

Because the guide has a terrible flaw — the bacterium's own CRISPR array contains the exact sequences the guide is drawn from. Without a keyword, the system would find its own memory, match it perfectly, and shred the archive. The PAM sits next to targets in phage DNA but not in the array itself. It's the self versus non-self check. Bacteria evolved it as a safety feature so their immune system wouldn't eat their own records.



Corn

And in human hands that safety feature became a fence around the yard.



Herman

You can only edit within shouting distance of a PAM, because Cas9 won't even park without one. People have since engineered Cas9 versions that recognize different PAM sequences, partly for exactly that reason. Step three: once Cas9 finds a PAM, it unwinds the DNA double helix right there and checks whether the guide RNA matches the strand opposite the PAM. Letter by letter. If the match doesn't hold, it lets go and moves on. If it does hold, step four — the cut. Cas9 has two nuclease domains, and each one cuts one strand, at the same height, facing each other. You get a clean double-strand break.



Corn

Clean meaning what, exactly? Both strands severed, nothing dangling?



Herman

Both strands cut at the same position, leaving what are called blunt ends — no overhang, no ragged edge. Which matters, because the cell's repair machinery treats a clean break and a messy break somewhat differently, and the cleanliness is part of what made CRISPR so immediately useful. It's predictable damage. And now the metaphor everyone learned dies. The scissors have done their work. Show me the edited gene.



Corn

You can't. There's nothing to show.



Herman

Right — I walked into my own trap there. That's the misconception Daniel's prompt is circling. Cas9 does not write anything. It cuts. Everything that happens next is done by the cell's own DNA repair machinery, and the cell has two very different moods about a double-strand break.



Corn

A double-strand break is bad news for a cell, though.



Herman

It's the worst news that isn't already a death sentence. A broken chromosome. The cell treats it as an emergency, and the fast emergency repair pathway is called non-homologous end joining — take the two broken ends and glue them back together, fast, consequences later. It glues well enough to keep the cell alive, but it's sloppy at the seam. Little bits get lost, little bits get added — what people call indels. If your goal is to knock a gene out, this is accidentally perfect: cut the gene, let the gluing scramble it, done. Knockouts are easy. That's the great trick of classic CRISPR — controlled damage plus opportunistic repair.



Corn

It's worth pausing on how strange that is as an engineering philosophy. You don't build the broken thing you want. You break a working thing and trust entropy to finish the job.



Herman

And the reason it works is that genes are mostly read in three-letter words. You don't have to destroy a gene to knock it out — you just have to shift the reading frame by one or two letters, and everything downstream becomes gibberish. A sloppy seam that deletes one letter is a complete knockout. The sloppiness of the repair and the fragility of the genetic code happen to cooperate.



Corn

And if you want a precise edit?



Herman

Then you need the other mood — homology-directed repair, HDR. You supply a template, a DNA sequence that looks like the region around the cut but carries the change you want, and the cell, if it's in the right phase and the mood strikes it, uses your template to rebuild the cut precisely. The problem is that HDR is slow, inefficient, and only really works in dividing cells. Most cells in your body aren't dividing.



Corn

How inefficient are we talking, roughly?



Herman

In many cell types, the precise repair pathway gets used a small fraction of the time — the emergency glue dominates. So even when you do everything right, deliver the perfect template, you come back and find the cell glued the break sloppily anyway. You're not choosing the outcome; you're weighting a lottery.



Corn

So the scoreboard reads: smash the gene, near-certain success. Fix one letter, and you're begging the cell to do the hard thing while it's panicking.



Herman

That's the gap in one line. It's why a therapy that knocks something out — like Casgevy disrupting a repressor gene in sickle cell — was first through the door, and precise one-letter corrections are still hard. Most disease-causing mutations are single letter changes. The tool we had was a hammer, and the job is watch repair.



Corn

The hammer works beautifully for a specific kind of hammering.



Herman

It does! And I want to be fair to it. The sickle cell example is a clever use of the hammer — instead of fixing the sickle mutation itself, which is a one-letter correction, you knock out a different gene entirely, one whose removal switches the cell back to making fetal hemoglobin. You sidestep the precision problem by choosing a target where blunt damage is the cure. That's real ingenuity. But the field looked at that repair bottleneck and said — the problem isn't the cut, the problem is that we're outsourcing the writing to a panicking cell. What if we didn't cut both strands at all?



Corn

Before you get there — one more thing about the blunt end of this. People also assume Cas9 cuts where you point it, every time, and only there. Off-target cutting.



Herman

It happens, and nobody's got a clean number for how often in clinical use — specificity is a design consideration in every protocol, guide sequences get screened against the whole genome to minimize near-matches. The honest framing is that off-target risk is managed, not eliminated. A guide with a near-match somewhere else in those three billion letters can sometimes cut there too, and the screening is about choosing guides whose near-matches are few and far from anything that matters.



Corn

So it's a statistical discipline, not a guarantee.



Herman

But the bigger design story of the last decade isn't making the cut safer. It's not cutting.



Corn

That's the pivot to base editing.



Herman

That's the pivot. Take Cas9 and break it on purpose. There's a version called dCas9 — dead Cas9 — where both cutting domains are disabled. It still finds the address, still binds the PAM, still unwinds and matches. It just can't cut. Now fuse a second enzyme onto it: a deaminase, a protein whose whole job is chemically converting one DNA letter into another.



Corn

A pencil welded to a locator.



Herman

That's the fair version of the metaphor. The cytosine base editor deaminates cytosine into uracil, which the cell reads and copies as thymine — so a C-G pair becomes a T-A pair. The adenine base editor does the mirror trick, converting adenine to something read as guanine. No double-strand break, no donor template, no HDR, no emergency. You've walked up to one letter and changed it in place while the DNA sits there intact.



Corn

And the catch?



Herman

It's narrow. Those are the only four directions it can go — C to T, A to G, and the reverses on the other strand. Those are called transitions. The other eight possible letter changes — transversions, like C to A — are outside its vocabulary, and so are insertions and deletions. A pencil, not a word processor. It's precise within a small box.



Corn

For anyone whose eyes glazed at transitions and transversions — give me the plain-language version.



Herman

DNA has four letters — A, T, G, C — and they pair up: A with T, G with C. Chemically, A and G are close cousins, and so are C and T. A transition is swapping a letter for its cousin. A transversion is swapping it for a member of the other family — a bigger chemical jump. And the deaminase enzymes, the pencils, can only manage the small jump within the family. The big jump is out of reach.



Corn

But for diseases caused by a transition mutation, that box is the whole job.



Herman

A striking share of known pathogenic mutations are exactly the transitions base editors can fix — G to A, C to T type changes are everywhere in the mutation catalogs. There are base editing therapies in clinical trials now for things like familial hypercholesterolemia, where a one-letter fix in the liver is the whole game.



Hilbert

Hilbert's Anteater Trivia. A giant anteater can eat up to about thirty thousand insects in a day, mostly ants and termites. I have never eaten thirty thousand of anything. I keep a list.



Corn

I'd like to see the list.



Herman

Getting back to the box problem — what if the mutation is a transversion? Or you need to delete six letters, or insert three? The base editor can't help you. You'd be back to the scissors and the begging. That's the gap prime editing was built to close, and it's the most mechanically beautiful thing in this whole story. I say that as a warning, because you're going to hear how I sound in a second.



Corn

I've heard the sound. Continue.



Herman

Prime editing — first published by David Liu's group in twenty nineteen, and the title of the paper is the pitch: search-and-replace genome editing without double-strand breaks or donor DNA. The machine has two parts. Part one is a Cas9 nickase — Cas9 with one cutting domain disabled, so it cuts one strand, not both. Part two is fused to that: a reverse transcriptase, an enzyme that copies RNA into DNA, taken from a mouse leukemia virus and engineered — five specific amino acid substitutions — to behave in a genome.



Corn

You said five substitutions with a straight face.



Herman

It's the kind of detail I live for, Corn, I won't apologize. The guide is where it gets clever. A normal guide RNA is an address and nothing else. The prime editing guide RNA — the pegRNA — carries the address, and then, hanging off the end of it, extra letters: two additions. One is a primer binding site. The other is a reverse transcriptase template — a stretch of RNA that literally spells out the edit you want written, letter by letter.



Corn

The address and the replacement text in the same envelope.



Herman

That's it. The mechanism then runs like this. The nickase finds the PAM, matches, and nicks one strand only — a single-strand nick, three letters upstream of the PAM. That nick leaves a free end of DNA. And here's the beautiful part: that free end grabs onto the primer binding site on the pegRNA — the DNA holds onto its own instruction sheet. Then the reverse transcriptase runs down the template on the pegRNA and synthesizes new DNA off it, writing the edited sequence directly onto the genome.



Corn

Let me make sure I have the geometry. The machine brings its own replacement strand, physically attached to the guide that found the spot. It doesn't hand the cell a blueprint and hope — it hands it a finished half.



Herman

That's the conceptual jump, yes. With HDR, you deliver a loose template and ask the cell to consult it under duress. With prime editing, the new text arrives already covalently stitched to the targeting molecule. The cell doesn't have to choose to use your template — the edit is physically attached to the address.



Corn

The cell is now holding two versions of that stretch. Old strand, new strand.



Herman

A heteroduplex — one old strand, one new strand with your edit in it. Then the cell's own housekeeping does the cleanup, and this is the part I find moving, because the mechanism outsources the last steps to three completely ordinary cellular enzymes. FEN1, a flap endonuclease, trims the old unedited flap. Ligase seals the nick. Mismatch repair resolves the disagreement between the strands. None of those enzymes know CRISPR exists. They've been doing their jobs for a billion years, and prime editing quietly borrowed them.



Corn

It's a Rube Goldberg machine built out of parts that were already in the room.



Herman

And it works. Scope: all twelve possible base-to-base conversions, plus targeted small insertions and deletions, with no double-strand break. PE3, the next generation, adds a second guide that nicks the opposite strand, and that alone boosts efficiency two to four times. The generations have kept coming — PE2 through PE6, plus something called PEmax — each one chasing the same things: efficiency, purity of the edited product, and how many places in the genome you can reach.



Corn

You've been waiting to be asked this. Efficiency numbers. Where does it actually stand, in a real cell, today?



Herman

The honest answer is that it's context-dependent, and anyone who quotes you one efficiency number is selling something. It depends on the pegRNA design, on the cell type, on which DNA repair pathways dominate in that cell. The same edit can work well in one cell type and poorly in its neighbor, because the neighbor's repair economy runs on different rules. The reviews all say the same thing in different words: prime editing is still in its infancy. It needs better efficiency and better delivery before it's a routine therapy.



Corn

You said earlier the delivery is a field of its own. Say more about that, because I think it's the unglamorous wall this whole program hits.



Herman

The editing chemistry is a solved-looking problem compared to getting the editor into the right cells, in the right amount, and only those cells. Cas9 plus guide is a big protein plus an RNA — that doesn't cross a cell membrane on its own. In therapies today you often take a patient's cells out, edit them in a dish, and put them back, because ex vivo delivery is easy — you control everything in the dish, and you can check the edit before the cells go back in. Casgevy works exactly that way: it's not a drug you swallow, it's a procedure. Editing cells inside a living body — in the brain, in muscle — that's where people use lipid



Herman (continued)

nanoparticles and viral vectors, and both have real limitations of size and targeting. And a prime editor is a bigger cargo than plain Cas9, which makes the size problem worse, not better.



Corn

So the fancier the pencil, the harder it is to mail.



Herman

There's also a wider family now that people should know exists, beyond even prime editing. The toolkit has diversified — Cas12 and Cas13 variants that target RNA instead of DNA, CRISPR-associated transposases that insert big chunks, and then the newest arrivals, discovered in twenty twenty-four: seekRNA and bridgeRNA systems, which pair the RNA-guided targeting idea with recombinases — enzymes that cut and paste DNA without the double-strand break at all.



Corn

Recombination instead of repair. The field's trajectory in one line: first cut and pray, then nick and pencil, and now don't even break the strand — swap it.



Herman

The ambition is moving from cut-and-repair toward write-what-you-want. There was a line in a review that stuck with me — our ability to edit genomes still lags behind our capacity to read them, but the gap is narrowing.



Corn

Sequencing got cheap twenty years before editing got precise. We've been living in the gap.



Herman

We still are, mostly. But the gap has an end shape now, which it didn't used to. In twenty ten you could sequence a patient's genome over a weekend and then do essentially nothing about the mutation you found. Today, for a growing set of mutations, there's at least a candidate tool with the right shape — whether it's a pencil, a search-and-replace, or a swap. The question left is delivery and efficiency, not concept.



Hilbert

Corn, you said the address and the replacement in the same envelope. How many times does the envelope have to arrive before anyone gets a delivery?



Corn

Once would be nice.



Hilbert

That's my whole feeling about it. My cousin Norbert worked thirty years in a molecular diagnostics lab, and I'll relay his opinion on all this, with the warning that Norbert is not to be trusted on anything involving yogurt. He says the mechanism papers are beautiful and the bottleneck is the courier, always was. His lab ruined a month of work once because a dry ice shipment sat on a dock over a weekend. He wrote the date and the word thawed on the tube and kept it on the shelf next to the fume hood as a reminder. He says the same thing about your PAM — the scissors were fine, the address didn't exist. Also I'm cutting that cough out of your last line, Herman, it clipped the mic.



Herman

The cough is gone. But your cousin's framing — the address didn't exist — that's the constraint people underestimate. Cas9 can be perfect and still useless if there's no PAM within reach of your target. The sequence you need to fix is where it is; the tool can only park where the motif lets it. That's why the engineered PAM variants matter as much as any efficiency gain. And the thawed tube is the same lesson one level up — the chemistry on the page and the chemistry that survived the dock are different products.



Hilbert

Norbert would agree, and he'd also say something about the reagent suppliers that I won't repeat. He's a good man with a bad temper about cold chains.



Corn

So the souvenir tube. Labeled thawed.



Hilbert

Norbert's. Not mine. I don't keep things. He kept it, and he'll tell you about it whether you ask or not. Anyway. The people doing that work with the dry ice and the couriers, they're the ones who actually get the letter there. Whatever the envelope says.



Corn

On the dry ice question — that's the real lesson hiding in Norbert's thawed tube, isn't it. The elegance that doesn't survive the dock is worth nothing.



Herman

The delivery people have a saying that amounts to the same thing. The best editor in the world is a paperweight in the wrong cell.



Corn

And the closing thought I'd leave Daniel with — the field is now racing to close the gap it stopped pretending about. Not whether we can read a genome, or even find an address in it, but whether the envelope arrives cold. That's where the next five years go.



Herman

On the mechanism, the summary is that the pencil and the search-and-replace both do the thing the scissors couldn't: they take the panicking repair machinery out of the loop. That's the conceptual shift, and it happened in under a decade.



Corn

Credit where it's due — our producer, Hilbert Flumingtop, keeps the desk running and the coughs out of the mix. If this was your kind of episode, go back for episode eleven, How Does Fine Tuning Work Anyway; and episode twelve, The AI Breakthrough. This has been My Weird Prompts. If you enjoyed this one, leave us a review, and the whole archive lives at my weird prompts dot com.



Herman

See you soon, everyone.



Corn

See you soon.

SOURCES

What the research for this episode read before the script was written. Primary sources first.

1. **Anzalone AV et al. Search-and-replace genome editing without double-strand breaks or donor DNA. *Nature*, 2019, Fig. 1 [primary]**
   https://www.nature.com/articles/s41586-019-1711-4/figures/1

2. **Wikipedia, CRISPR gene editing**
   retrieved 2026-11-19, last edited 2026-11-15.
   https://en.wikipedia.org/wiki/CRISPR_gene_editing

3. **Wikipedia, CRISPR**
   retrieved 2026-11-19, last edited 2026-11-14.
   https://en.wikipedia.org/wiki/CRISPR

4. **Wikipedia, Cas9**
   retrieved 2026-11-19, last edited 2026-11-15.
   https://en.wikipedia.org/wiki/Cas9

5. **Wikipedia, Genome editing**
   retrieved 2026-11-19, last edited 2026-11-16.
   https://en.wikipedia.org/wiki/Genome_editing

6. **Wikipedia, He Jiankui affair**
   retrieved 2026-11-19, last edited 2026-11-03.
   https://en.wikipedia.org/wiki/He_Jiankui_affair

7. **Chen PJ, Liu DR. Prime editing for precise and highly versatile genome manipulation. *Nature Reviews Genetics*, 2023. PMID 36344749**
   https://pubmed.ncbi.nlm.nih.gov/36344749/

8. **Villiger L, Joung J, Koblan L et al. CRISPR technologies for genome, epigenome and transcriptome editing. *Nature Reviews Molecular Cell Biology*, 2024. PMID 38308006**

https://pubmed.ncbi.nlm.nih.gov/38308006/

9. **Liu D, Cao D, Han R. Recent advances in therapeutic gene-editing technologies. *Molecular Therapy*, 2025. PMID 40119516**

https://pubmed.ncbi.nlm.nih.gov/40119516/

10. **Zhao Z, Shang P, Mohanraju P et al. Prime editing: advances and therapeutic applications. *Trends in Biotechnology*, 2023. PMID 37002157**

https://pubmed.ncbi.nlm.nih.gov/37002157/

11. **Mentani A, Maresca M, Shiriaeva A. Prime Editing: Mechanistic Insights and DNA Repair Modulation. *Cells*, 2025. PMID 39996750**

https://doi.org/10.3390/cells14040277

12. **Aliciaslan M, Erbasan E, Erendor F et al. Prime Editing: The Next Frontier in Precision Gene Therapy. *Journal of Gene Medicine*, 2025. PMID 41038616**

https://pubmed.ncbi.nlm.nih.gov/41038616/

13. **Aliciaslan M, Erbasan E, Erendor F et al. Advances in CRISPR Base Editing: From Molecular Evolution to Therapeutic Applications in Genomic Medicine. *Journal of Cellular and Molecular Medicine*, 2026. PMID 42033211**

https://pmc.ncbi.nlm.nih.gov/articles/PMC13109818/

14. **Ma X, Li L, Sun Y. CRISPR Technology: From Bench to Bedside Applications. *MedComm*, 2026. PMID 43011237**

https://pmc.ncbi.nlm.nih.gov/articles/PMC13748939/

15. **Li X, Wu S, Tian P. Base Editors for Engineering Industrial Microorganisms: Types, Applications, and Future Perspectives. *Biotechnology Journal*, 2026. PMID 42307052**

https://pubmed.ncbi.nlm.nih.gov/42307052/

16. **Ali Agha ASA, Hattab D, Bakhtiar A et al. Programmable RNA-Guided DNA Recombination: Mechanisms, Engineering, and Applications. *Biomedicines*, 2026. PMID 42792750**

https://pmc.ncbi.nlm.nih.gov/articles/PMC13604230/

17. **Prime editing: therapeutic advances and mechanistic insights. *Gene Therapy*, 2024**

https://www.nature.com/articles/s41434-024-00499-1

18. **Singh S, Tiwari H, Singh M et al. Expanding the Microbial Genomic Landscape and Biotechnological Applications of CRISPR-Cas Systems. *Biology*, 2026. PMID 42187710**

https://pmc.ncbi.nlm.nih.gov/articles/PMC13203232/