

Claude designs 17 lab-confirmed protein binders against scorpion neurotoxins

Source copy kept by wheresthe.ai for the story "Claude designs 17 lab-confirmed protein binders against scorpion neurotoxins" (<https://wheresthe.ai/d/us-anthropic-claude-designs-protein-binders-against-scorpion-neurotoxins-2026-10/>).

Source: <https://www.anthropic.com/research/claude-scorpion-toxin-binder-design>

Text saved 25 Oct 2026 (the page as published).

SHA-256 of the saved text: 8da2a7054b4867aeee5471390be7a60d2f206544fea1322a5393be5a02ef5afb

Documents the source offers: <https://www-cdn.anthropic.com/31ac578806cffd8b3d92d8affe6edc4c04d502e1.pdf>

Summary: As part of our work experimenting with how Claude can help accelerate drug discovery, we asked an internal research model to design minibinders (small proteins that attach tightly to a target) against scorpion neurotoxins. Claude worked from our prompt and four follow-ups, which we developed in collaboration with the financial investment research firm Citrini Research. Teams at Adaptic Bio and Twist Bioscience experimentally tested 211 designs for binding. Of these, 17 met our criteria. We're releasing all 211 designs and their data as a resource for the field.

In August, we showed how Claude designed protein minibinders after being given a single prompt. A month before we released those results, a team from Citrini Research had approached us with an interesting question: could Claude (and AI drug discovery more generally) enable a group of motivated nonexperts to tackle a genuine problem in human health? Members of our life science research team set out to answer this, first by brainstorming problems with Claude. We landed on asking for Claude's help designing novel protein binders against scorpion neurotoxins, with the hope that they could one day be used to develop a scorpion antivenom. Scorpions sting 1.2 million people or more every year, resulting in more than 3,250 deaths, mostly among children in parts of the world with limited access to medical care. New treatments are expensive to make, and are also unlikely to earn back the cost of developing them. With so little investment, scorpion envenomation has been described as a neglected tropical disease. We started by giving Claude the same prompt we'd used in our earlier protein-binder study, after modifying it with Citrini. 1 Over the next 64 hours, an internal research model, operating within Claude Science, worked mainly autonomously, although we directed its work with four follow-ups. Under our supervision, Claude analyzed the research data, generated the figures, and drafted the final report. Claude's design campaign produced 17 protein binders that each bind to one of four scorpion neurotoxins that are known to be dangerous to humans. Throughout the process, Claude generated 318,484 protein designs and selected 180 of them for testing. Our team added 31 more, with Claude's help. Our partners at Adaptic Bio and Twist Bioscience, two contract research organizations (CROs), lab-tested all 211 designs to see if they would bind to the toxins they were designed against, and 17 designs met our criteria for a binder. We have not confirmed whether these binders neutralize scorpion venom in a cellular environment, so we do not claim to have produced an antivenom. We sent Claude's work to two outside experts, Susana Vázquez Torres and Timothy P. Jenkins, who said they felt this work was a promising first step. If AI can now design credible binders against several hard-to-target toxins in days and they can be confirmed in the lab within weeks, this could, in theory, speed up discovery for envenomation and other neglected diseases. The challenge, as always, is turning those binders into safe, affordable, and widely accessible drugs. We share more detail in our technical report , and are making all 211 designs and their data available . Designing a binder against almost any target no longer seems unrealistic. We can now obtain very tight binders remarkably quickly, and binding itself is no longer the main bottleneck. For venomous animals such as snakes and scorpions, whose bites and stings kill or seriously injure large numbers of people, particularly in underserved regions, the real impact will come from tackling the parts of the therapeutic pipeline that go beyond binding. The challenge is to use these extraordinary tools to develop molecules that can become real drugs, ones that are safe, affordable, and widely accessible.

Susana Vázquez Torres , Postdoctoral Researcher, Structural Biology Programme, Spanish National Cancer Research Centre (CNIO), Madrid, Spain. Designing de novo binders against five toxins from four species of scorpions Scorpion venom contains multiple protein neurotoxins. The neurotoxins that are most likely to harm or kill humans attach to sodium channels , large proteins embedded in the membranes of nerve cells that allow them to fire. When sodium channels are disrupted, nerves fire at the wrong time—this releases large amounts of stress chemicals, like adrenaline. In severe cases, this can cause heart failure and excess fluid buildup in the lungs, occasionally resulting in death. There are some antivenoms available for scorpion stings. They're generally made

using the antibodies of immunized animals, a process that can take months. Venom is milked from thousands of scorpions and used to immunize horses; scientists then purify the resulting antibodies from the horses' blood. Traditional antivenoms can work when given early, but they're usually effective for only a small number of neurotoxins. They're also expensive to produce, and in many regions are costly or in short supply. De novo design offers a potential solution. A binder is a molecule that attaches tightly to a specific target; antibodies are one natural example. In principle, a binder that targets the part of a neurotoxin that latches onto sodium channels could stop the toxin from reaching the channel. In de novo design, protein binders are generated from scratch against a specific target. These binders are designed computationally and can be optimized for druglike properties such as potency and stability, which could make them easier to supply in regions where stings most often occur. Our approach was inspired by an earlier effort that designed de novo proteins that neutralized snake neurotoxins and protected mice. To our knowledge, no de novo -designed binders against scorpion neurotoxins have been described. For this campaign, we chose to target five scorpion-venom neurotoxins across four regions: Mexico (Cn2), Brazil (Ts1), North Africa (AaH I and AaH II), and the Middle East (LqhIII). We also chose to aim the designs at each toxin's channel-binding site. Claude picked the toxin structures to use for design by searching the Protein Data Bank (PDB) as well as the AlphaFold database (AFDB). For each toxin, it selected five hotspots, specific amino acid positions on the toxin that its binder designs were meant to latch onto.

We were primarily wondering if Claude could design binders against individual toxins. But we were also curious if Claude could design binders that targeted multiple toxins. There's a practical reason for trying to design multitoxin binders, namely that it would reduce the number of molecules needed for a universal antivenom cocktail—which would be the most cost-effective and regionally neutral treatment. We selected these five scorpion toxins (Cn2, Ts1, AaH I, AaH II, and LqhIII) as a challenging test of multitoxin design: they share one fold for the overall shape of the protein, but any two are only 26% to 62% identical in their amino acid sequences. To see how far Claude could go, we asked it to design four kinds of binders: binders against each of the five toxins individually; binders that targeted the β -family toxins Ts1 and Cn2; binders that targeted the α -family toxins AaH II, AaH I, and LqhIII; and binders that targeted all five toxins simultaneously. Designing credible binders against five toxins in parallel, in a matter of days, is a great showcase of the AI-driven paradigm shift we are seeing in therapeutic discovery. This is especially true for neglected diseases such as animal envenomings, which have been largely starved of innovation due to low commercial incentives. Being able to agentially tackle the uniquely broad target space across animal toxins provides a significant speed-up in discovery, and brings realistic therapeutic solutions much closer.

Timothy P. Jenkins, Associate Professor, Technical University of Denmark. Claude ran the design campaign with considerable autonomy. Our initial prompt defined the neurotoxin targets, the protein design and folding tools, and the rules for novelty, scoring, and filtering. Beyond choosing the structures and amino acid hotspots, Claude also set the quality thresholds a design had to clear, decided how much to lean on each available design tool, and developed the strategies for designing binders against multiple toxins. Ultimately, Claude used 12 design tools to generate 318,484 designs, filtering out any that resembled a toxin or a known protein. Claude co-folded and screened 553,052 design-toxin pairs, predicting each pair's structure as one complex. Sixty-four hours after we gave Claude our initial prompt, it produced its final list of 180 designs. The campaign used on the order of 10,000 GPU hours. 2 Claude worked autonomously for 24 hours following our first prompt. Its results encouraged us to prompt Claude to continue its campaign and look for better designs. Our follow-up prompts were mainly focused on directing Claude to design a multitoxin binder. Claude selected most of the 180 designs—including our strongest binder (Cn2-B1)—only after our third follow-up. At our request, Claude sent the protein sequences to Twist Bioscience (in emails indicating they were written by Claude), which produced the proteins in order to see how well they would bind to the toxins. We later worked with Claude to add another 31 designs, increasing the total to 211 designs: 60 designs against a single toxin (12 per toxin), 30 designs against the two β -family toxins, 30 designs against the three α -family toxins, and 91 designs against all five toxins. Of the 211 designs our teams assessed, 17 met our criteria for a binder, each against a single toxin. None of the designs made against multiple toxins met our criteria for more than one. (Beyond being limited by its current capabilities, there could be many explanations for why Claude struggled with this, including our prompting, the limitations of available design tools, and even the biological feasibility of multitoxin design.) Seven designs were binders against Cn2, four against AaH II, four against AaH I, and two against LqhIII. Among the 60 single-toxin designs, 11 were binders (18%). This rate is consistent with typical protein design campaigns, including our previous work.

Testing Claude's designs in the lab After completing the campaign, we sent the designs to Adaptic Bio and Twist Bioscience to test them for binding. Each group produced the designs in their labs, but in different formats. Adaptic Bio expressed them in a cell-free system, which uses protein-making machinery extracted from bacteria

rather than living cells. Twist Bioscience expressed the designs in mammalian cells (in this case, human kidney cells) as fusions to Fc, the stem of an antibody. These are complementary ways to measure the underlying binding process, and we combined them in order to improve confidence in our results. Both labs measured binding by surface plasmon resonance (SPR), in which a toxin flows over the design while the SPR instrument tracks in real time how quickly a toxin binds to and detaches from it. The resulting trace, called a sensorgram, produces a measurement known as the dissociation constant (KD), which is a measurement for binding strength (or affinity). The lower the KD, the tighter the binding. When it came to evaluating the resulting designs, we set out strict criteria for labeling a design a binder against a toxin. The design needed to be a hit at both labs —that is, it needed to show evidence of binding against that toxin in every replicate, and every KD needed to be 1 μ M or tighter. We also required that the design show no signal against CLEC12A, an unrelated human protein used to measure off-target binding or “stickiness.” Claude’s Cn2 binders likely bind the same site as a toxin-neutralizing antibody. Antibodies and nanobodies remain the standard for de novo designs to beat. Three such molecules have reported affinities (KD) below 1 nM (nanomolar) against Cn2 or AaH II, and they have been shown to protect mice from scorpion toxins. Our two strongest binders, which we call Cn2-B1 and Cn2-B2, target Cn2, a neurotoxin from the Mexican scorpion *Centruroides noxius*. They each bound Cn2 tightly at both Adaptiv Bio and Twist Bioscience, with affinities (KD) in the single-digit nanomolar range. One way to find out whether two molecules (e.g., an antibody and a miniprotein) bind the same site on a target is with a competition experiment, which tests whether both binders can attach to the target at the same time. If the presence of one molecule reduces the other molecule’s binding, then they likely compete for the same binding site. Our competition experiments suggest our Cn2 binders each attach on or near the same epitope (binding site) as an antibody called LR, which has been reported to neutralize Cn2 in mice. LR’s binding site covers the part of Cn2 that attaches to the sodium channel and causes nerves to misfire. We have not tested our strongest binders, Cn2-B1 and Cn2-B2, in a cellular environment or in live animals, but the fact that they compete with LR suggests that they will likely neutralize Cn2.

A binder is only the first step to developing an antivenom. Our results show binding only against recombinant toxins, which are grown in the lab from each toxin’s gene rather than purified from venom. The two biggest limitations of this study are that we didn’t test any design for neutralization in vitro (against cells grown in the lab) or for efficacy in vivo (in living animals, such as mice). If we were to advance one design toward a therapy, Cn2-B1 would be one of our strongest candidates. First, though, we would need to confirm the design also binds toxin purified from venom. We would then assess it for function in vitro with patch-clamp recordings (which can measure the electrical current passing through sodium channels) to see whether it stops the toxin from acting on the channel. Then, we would need to test the binder in mice to see whether it actually protects them. Beyond neutralizing the toxin, a treatment must also be safe, stable, affordable, and available where stings occur. Reaching patients will also take cooperation from clinicians, funders, and policymakers. We are releasing all 211 designs and their binding data so that other researchers can build on our work. We hope this demonstration of Claude’s capabilities, along with the data we are releasing, will help accelerate further research into antivenom design. Further reading The following resources provide further technical depth and more detailed information about the results described above: Technical report ; Data release: all 211 designs, their predicted structures and scores, and experimental data ; Our previous posts on protein design and uplifting biomolecular modeling .

Footnotes 01 Citrini has written an article on our collaboration, here . 02 We realize this is much more than most researchers have access to. Separately, we have made these models about 4x faster on average, and shown that a simplified prompt using a single agent and a single GPU over 24 hours reaches comparable in silico scores. 03 LqhIII was assessed only by Adaptiv Bio.