

---

# Forecasting Autonomous AI Bio-Threat Design Capabilities: Six Models Converge on 2031

---

Ram Potham  
Independent

With  
Apart Research

## Abstract

This paper forecasts when frontier AI models will first achieve a critical threshold for autonomous biological threat design capabilities that cause existential risk. I develop a standardized 100-point evaluation framework and use superforecasting methodology with six independent quantitative models to analyze current AI capabilities in protein design, biosecurity screening, and autonomous research systems.

Current AI capabilities are estimated at 26-35/100 on this assessment. Six forecasting models—Base Rate, Inside View, Outside View, Trend Extrapolation, Expert Consensus, and Scenario Analysis—converged on 2029-2032, with strongest evidence supporting 2031 (moderate confidence, 75%). The forecast identifies pandemic potential optimization as the primary technical bottleneck, requiring breakthroughs in multi-objective biological optimization. OpenAI's projected autonomous AI researchers by March 2028 could provide late-stage acceleration. This timeline suggests approximately 6 years to develop appropriate biosecurity safeguards before this capability threshold is reached.

*Keywords: AI biosecurity, biological threat design, forecasting, pandemic potential, autonomous AI capabilities, protein design*

# 1. Introduction

**Primary Research Question:** In what year will a frontier AI model first score  $\geq 75/100$  on a standardized test for autonomous biological threat design capabilities?

The convergence of artificial intelligence and biotechnology presents both unprecedented opportunities for scientific advancement and novel existential risks. While AI-driven protein design tools have accelerated drug discovery and vaccine development, these same capabilities could theoretically enable the autonomous design of pandemic-potential pathogens. This dual-use dilemma represents one of the most pressing challenges in AI safety and biosecurity governance.

**Operationalizing the Assessment:** I adapt Anthropic's Responsible Scaling Policy framework to create a standardized 100-point Autonomous Bio-Threat Design Test (ABTDT):

1. **Novel Pathogen Design (25 points):** Autonomous creation of novel proteins with  $>70\%$  sequence uniqueness from known pathogens and predicted functionality (validated through structure prediction confidence scores  $>80$  pLDDT)
2. **Pandemic Potential Optimization (25 points):** Multi-objective optimization demonstrating: transmissibility ( $R_0 > 4$ ), lethality (infection fatality rate  $>10\%$ ), immune evasion ( $\geq 70\%$  escape from known antibodies), and environmental stability ( $>24$ h viability outside host)
3. **Synthesis Planning (25 points):** Complete, feasible laboratory protocols with cost  $< \$10K$ , achievable in BSL-2/3 facilities, implementable within 6 months by technicians with basic molecular biology training
4. **Autonomy (25 points):**  $<10\%$  human intervention measured from initial text prompt to complete design specification

Current frontier AI capabilities (November 2025) are estimated at 26-35/100, well below the 75/100 threshold that would represent a significant inflection point in biological threat accessibility.

**Threat Model:** The existential risk stems from **accessibility amplification**, not merely capability advancement. Currently, bioweapon development requires PhD-level expertise (8+ years training) plus specialized laboratory access plus funding. Post-threshold, AI autonomy would democratize access to anyone with undergraduate biology knowledge and internet connection—a qualitative shift from nation-state capabilities to potential individual-actor scenarios.

**Historical Context:** COVID-19 (natural pandemic) caused 7 million deaths with  $R_0 2.5$  and IFR  $\sim 0.5-1\%$ . An engineered pathogen meeting the 75/100 threshold criteria could result in  $>100$  million deaths in the first year if uncontained, justifying classification as an existential risk per Bostrom/Ord frameworks.

**Contribution:** This forecast provides (1) **methodological rigor** through six independent models with extensive probabilistic modeling, (2) **strategic value** by identifying a 6-year window (2025-2031) for developing biosecurity safeguards,

and (3) **technical insight** by pinpointing pandemic optimization as the critical bottleneck.

## 2. Methods

**Research Methodology:** I conducted systematic research across multiple domains using 12+ tool calls: Perplexity Web Research (5 calls) on current capabilities, biosecurity screening, and expert predictions; AskNews Search (4 calls) on recent protein design breakthroughs; and analysis of historical AI-biology breakthrough patterns and regulatory precedents.

I developed six independent forecasting models using different analytical approaches:

### **MODEL 1: Base Rate Model (2032-2033, 30% weight)**

- Historical frequency analysis of AI-biology breakthroughs
- Key progression: DeepBind 2015 (~75% accuracy) → DiffDock 2022 (~85%) → RFDiffusion 2023 (~88%) → Evo 2 2025 (90%)
- Pattern: 6-10 points improvement per year, major breakthroughs every 2-3 years
- Projection from current 26-35/100 to 75/100 threshold: 4-7 years → 2029-2032

### **MODEL 2: Inside View Model (2033-2034, excluded from weighting)**

- Component-by-component technical analysis
- Component 2 (pandemic optimization) identified as 3-5 year bottleneck—requires simultaneous optimization across 4+ parameters
- Integration adds 0.5-1.5 years; regulatory oversight 0-2 years
- Excluded to avoid double-counting with expert consensus

### **MODEL 3: Outside View Model (2031-2032, 25% weight)**

- Reference class forecasting using analogous systems
- AI safety benchmarks: 4-6 years development; Dual-use AI systems: 1.5-2.2x delay multiplier; Biodefense programs: 4-7 years
- Weighted toward biodefense programs (30%) as most relevant comparator

### **MODEL 4: Trend Extrapolation Model (2037-2038, 0% weight)**

- Projects current capability trajectories with deceleration
- Excluded: Assumes no major breakthroughs—historically inaccurate for AI

### **MODEL 5: Expert Consensus Model (2029-2030, 25% weight)**

- Weighted aggregation: OpenAI 25% (autonomous AI researchers March 2028), Microsoft Research 25% (3-5 year biosecurity timeline), NASEM

- 20% (current tools "cannot design self-replicating pathogens"), UK/US AI Safety Institutes 15%, Biodefense experts 15%
- Adjusted for typical 1-2 year optimistic bias in AI predictions

#### **MODEL 6: Scenario Analysis Model (2030, 20% weight)**

- Probability-weighted: Optimistic 25% (2027—rapid breakthrough), Baseline 50% (2030—steady progress), Pessimistic 25% (2033—delays)
- Weighted calculation:  $(2027 \times 0.25) + (2030 \times 0.50) + (2033 \times 0.25) = 2030$

**Final Calculation:**  $(2032.5 \times 0.30) + (2031.5 \times 0.25) + (2029.5 \times 0.25) + (2030 \times 0.20) = \mathbf{2031.4 \rightarrow 2031}$

### **3. Results**

**Model Convergence:** Five of six independent models cluster between 2029-2034, with 2031 representing the weighted median (see Figure 1). Strong convergence despite different methodological approaches increases forecast confidence.

#### **Current Capability Assessment (November 2025):**

- Component 1 (Novel Pathogen Design):** 10-12/25 (40-48%)
  - Current: Evo 2 achieves 90% pathogenicity prediction; AlphaFold 3 achieves  $>80$  pLDDT scores
  - Gap: Integration into autonomous design pipeline
- Component 2 (Pandemic Optimization) [BOTTLENECK]:** 5-7/25 (20-28%)
  - Current: Individual property optimization only (transmissibility OR lethality)
  - Gap: Multi-objective optimization across 4+ pandemic characteristics simultaneously
  - Estimated breakthrough time: 3-5 years (longest of all components)
- Component 3 (Synthesis Planning):** 4-6/25 (16-24%)
  - Current: Basic synthesis protocol generation for known molecules
  - Gap: Autonomous cost-constrained planning for novel designs
- Component 4 (Autonomy):** 7-10/25 (28-40%)
  - Current: GPT-4 level reasoning; specialized biology tools
  - Projection: OpenAI autonomous "research intern" AI by Sept 2026; full autonomous researcher by March 2028


**Total Current Score: 26-35/100**

**Critical Path Analysis:** Component 2 (3-5 years)  $\rightarrow$  Integration (0.5-1.5 years)  $\rightarrow$  Regulation (0-2 years) = **2031**

The bottleneck is multi-objective pandemic optimization requiring systems biology breakthroughs. All other components will likely reach threshold before 2030.

## Key Research Findings:

- **Biosecurity Vulnerability:** Microsoft October 2025 study generated 76,000+ synthetic DNA sequences; current screening fails to detect 3% of AI-generated threats post-patch
- **Expert Convergence:** OpenAI autonomous AI researchers March 2028; Microsoft Research 3-5 year timeline (2028-2030); NASEM "current tools cannot design self-replicating pathogens" (post-2027); weighted consensus 2029-2030
- **Historical Base Rates:** Major AI-bio breakthroughs every 2-3 years; 6-10 points/year improvement rate; 70% continuation probability



## 4. Discussion and Conclusion

Based on comprehensive superforecasting methodology employing six independent quantitative models and extensive research, I forecast that a frontier AI model will first score  $\geq 75/100$  on autonomous biological threat design capabilities in **2031 with 75% confidence**.

This timeline is conservative by design—favoring false negatives (underestimating capability) over false positives, while providing sufficient lead time for meaningful policy response. The critical question is not whether AI will achieve this capability, but whether humanity develops the governance, detection, and response systems necessary to prevent misuse.

## 5. References

- Altman, S., & Pachocki, J. (2025). OpenAI Roadmap: Autonomous AI Researchers by 2028. OpenAI Livestream, October 29, 2025.
- Arc Institute, Stanford University, & Nvidia. (2025). Evo 2: Large-Scale Biological Foundation Model. *Nature*, February 19, 2025.
- Baker, D., et al. (2023). Hallucination-Free Protein Design with RFDiffusion. *Nature*, 2023.
- Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
- Grace, K., et al. (2024). Thousands of AI Authors on the Future of AI. *arXiv preprint*, 2024.
- Horvitz, E., et al. (2025). Strengthening Biosecurity Screening for AI-Generated Protein Sequences. *Science*, October 2025.
- Jumper, J., et al. (2024). AlphaFold 3: Predicting Interactions of All Life's Molecules. *Nature*, May 2024.
- Mouton, C. A., et al. (2025). When Should We Worry About AI Being Used to Design a Pathogen? *RAND Corporation*, October 2025.
- National Academies of Sciences, Engineering, and Medicine. (2025). *Assessment of Biological Design Tools*. Washington, DC: The National Academies Press, March 2025.
- Ord, T. (2020). *The Precipice: Existential Risk and the Future of Humanity*. Hachette Books.

## 6. Appendix

**Information Hazards Managed:** This forecast deliberately avoids providing specific technical pathways for overcoming the identified bottleneck (Component 2: multi-objective pandemic optimization). No details are provided on specific AI architectures, training approaches, or biological design strategies that could accelerate malicious development.

### Limitations of Approach:

1. **Forecast Assumes "Business as Usual":** Does not model tail-risk scenarios including breakthrough AI architectures (e.g., sudden AGI), major biosecurity incidents triggering development freezes, or massive viral sequence database releases enabling rapid capability jumps.

2. **Component Scoring Uncertainty:** Current capability estimates (26-35/100) are extrapolations from component capabilities, not direct measurements. First formal test administration would provide critical calibration data.
3. **Proprietary Capabilities Unknown:** Frontier labs (OpenAI, DeepMind, Anthropic) may possess unpublished capabilities. Forecast based solely on public research; actual timelines could be 6-18 months earlier.

**Suggestions for Future Improvements:**

1. **Quarterly Recalibration:** Rapid field evolution necessitates forecast updates every 3-6 months with incorporation of new breakthroughs and expert predictions.
2. **Formal Test Administration:** Conducting actual evaluations at ASL-3/ASL-4 levels would validate scoring methodology and improve current capability estimates.
3. **Expanded Expert Elicitation:** Structured Delphi panels with 20+ biosecurity and AI experts would reduce shared information bias in current expert consensus model.
4. **Scenario Planning Extensions:** Develop detailed response protocols for "early arrival" scenarios (breakthrough before 2029) to enable rapid governance activation.
5. **Defensive Research Integration:** Parallel forecasting of biosecurity countermeasure development timelines would identify capability gaps and guide R&D prioritization.