

AI-driven antibody design from an antigen sequence

Pavel Sinitcyn^{1,2}, Albert J. R. Heck^{*2}

<https://doi.org/10.15302/vita.2025.12.0004>

Large protein language models are used to generate fully human, paired-chain antibodies directly from an antigen sequence, with measurable binding *in vitro*. The monoclonal antibody generator tool and validation study by Wasdin *et al.*¹ published in *Cell* provides a strong experimental case for this direction and underscores the need for a critical assessment of structure prediction-like, community benchmark to compare *de novo* antibody design strategies.

The human immune system protects us against pathogens but also contributes to surveillance against cancer and other diseases. A central component of this immune response is provided by antibodies, produced by plasma B cells and secreted into circulation and mucosal tissues. Antibodies recognize antigens mainly through complementarity-determining regions — short, hypervariable stretches in the variable domains that interact with the antigens. Through affinity maturation, B cells iteratively refine these hypervariable regions via somatic hypermutation and selection, yielding high-affinity IgG-producing plasma cells. In principle, the combinatorics of V(D)J recombination, mutation, and chain-pairing allow an enormous diversity of antibody sequences, far beyond the number of cells in the body, even if only a small fraction may be ever realized in an individual². Current estimates suggest that an individual's antibody repertoire contains at least about 10⁹ distinct antibodies³.

Recombinant antibodies have become one of the most successful therapeutic modalities, with use in oncology, autoimmunity, and infectious diseases. Yet, finding “the right antibody” is rarely just a question of antigen binding. Therapeutic function also depends on factors such as neutralization mechanism, epitope choice, Fc-mediated effector functions, stability, manufacturability, and safety liabilities⁴. From both basic and translational perspectives, this leads to a key challenge — how can we efficiently identify antibody sequences that meet a specific therapeutic goal?

So far, both experimental and computational antibody discovery strategies have remained time-consuming and resource-intensive. Many computational methods start from known antibodies, require antigen structural information, or focus on optimizing already identified binders rather than generating complete *de novo* heavy–light chain pairs⁵. Approaches that claim “from-scratch” generation often face different limitations — unclear generalization, limited experimental validation, or reliance on strong structural assumptions about the interacting partners. As a result, it has remained uncertain whether fully sequence-based models, given only an antigen amino acid sequence, can reliably design paired, fully functional antibodies.

Wasdin *et al.* report in *Cell* a sequence-based model, monoclonal antibody generator (MAGE), that generates paired antibody variable regions using only the antigen sequence as the input¹. MAGE is obtained by fine-tuning ProGen2⁶, an auto-regressive protein language model (PLM) pretrained on diverse protein sequences (Fig. 1). The core training objective remains the same as in standard decoder-only language modeling — next-token (next-amino acid) prediction. Fine-tuning exposes the model to paired antibody–antigen sequence examples, encouraging it to learn the statistical relationships between antigen sequence patterns and antibody variable-region sequence features associated with binding. At inference time, MAGE produces both heavy and light chain variable sequences together, without requiring (i) a starting antibody template, (ii) antigen structure, or (iii) a predefined epitope. This is conceptually appealing, as it mirrors how discovery teams often wish they could operate at the early stages in a program — generate diverse human antibody candidates quickly, then let experiments select what works.

A major strength of the current work is its emphasis on experimental validation across multiple antigens. The authors evaluated MAGE on three viral targets: SARS-CoV-2 (spike RBD), RSV (prefusion F), and influenza (hemagglutinin). For each antigen, they expressed and tested a panel of designed antibodies, reporting that a subset showed measurable binding *in vitro*, with some predicted candidates also showing functional activities (e.g., neutralization of viral infection). Notably, the study also reports binding to an influenza variant that was not present in the training set, suggesting at least limited ability to generalize beyond highly represented targets. In addition to binding assays, structural characterization on RSV supports that the designed antibodies can adopt plausible binding modes, with contributions from both chains. Taken together, these results move the field beyond “promising *in silico* scores.” MAGE provides evidence that sequence-only generation can yield experimentally confirmed, human paired-chain binders across multiple target classes.

Wasdin *et al.* also outline the limitations that matter for interpretation and future development. The fine-tuning data are heavily enriched for coronavirus-related antibodies, including substantial representation of receptor-binding-domain (RBD)-focused examples. This raises a legitimate concern, whether strong performance on RBD could partly reflect proximity to the training distribution. The influenza result on an unseen variant is encouraging, but broader claims of generality will require systematic testing across targets with minimal historical antibody data. Because the model is not explicitly guided by structure or epitope definition, it offers limited control over where on the antigen it binds or how it binds. In addition, the

1. AI Technology for Life, Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands. 2. Biomolecular Mass Spectrometry and Proteomics, Bijvoet Center for Biomolecular Research and Utrecht Institute for Pharmaceutical Sciences, Utrecht University, Utrecht, The Netherlands. *Correspondence: Albert J. R. Heck (a.j.r.heck@uu.nl).

Published: February 2, 2026

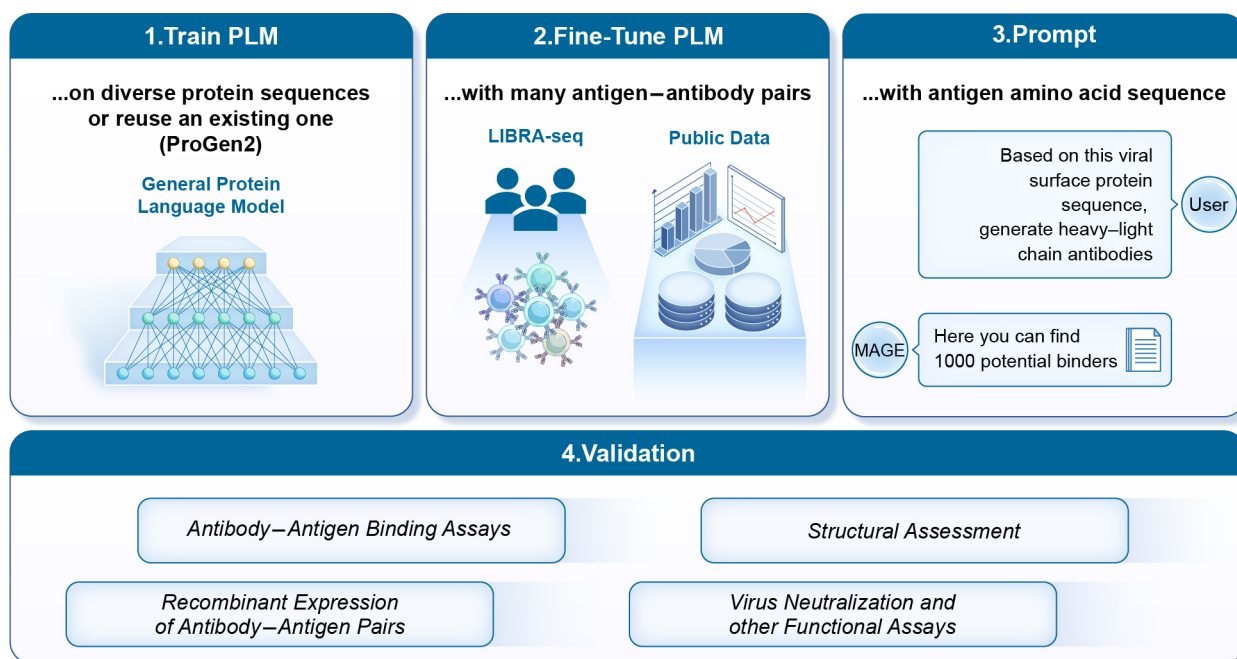


Fig. 1. MAGE — workflow for sequence-based *de novo* antibody generation and validation. (1) A general protein language model (PLM) is trained on diverse protein sequences, or an existing pretrained PLM (ProGen2) is reused. (2) The PLM is fine-tuned on large collections of paired antigen–antibody sequence data, combining high-throughput datasets (LIBRA-seq) with public antibody–antigen resources (CoVAbDab, SAbDab). (3) At inference, the user provides an antigen amino acid sequence as a prompt; MAGE generates an unranked set of paired heavy–light chain antibody candidates (typically hundreds to thousands) predicted to bind the target. (4) Candidates are recombinantly expressed and experimentally validated using binding assays, functional readouts such as neutralization assays, and structural assessment where feasible.

reported candidate list is not ranked by affinity, since affinity prediction was not a training objective. This can be acceptable for early “binder discovery,” but it does not guarantee the therapeutic developability, as binding is indirectly necessary but not sufficient. Properties such as affinity, avidity, stability, polyreactivity, aggregation propensity, expression yield, and immunogenicity risk still require downstream assessment and often iterative engineering⁷.

MAGE arrives at a moment when many groups — academic labs and biotech companies — are pursuing *de novo* antibody design using different strategies, like a sequence-only generation, structure-conditioned diffusion, docking-guided optimization, hybrid pipelines, and more⁵. In this environment, progress is hard to compare because evaluations are typically not blinded and often differ in target selection, screening budgets, and success metrics.

We would like to highlight that protein structure prediction faced a similar problem decades ago, and the community benefited enormously from a critical assessment of structure prediction (CASP)⁸, which created a shared, neutral framework for blinded assessment. A comparable effort for *de novo* antibody design could provide blinded antigen targets (including novel variants and “hard” nonviral antigens), standardized reporting, common quality metrics, and optional experimental validation where feasible. The primary goal would not be to rank winners for marketing, but to establish an unbiased scoreboard of what works today, what fails consistently, and which innovations truly move the needle.

If template-free antibody design becomes reliable, it could compress the earliest stages of discovery from months and years to weeks and help bring antibody therapeutics to emerging pathogens and long-neglected targets. The most plausible near-term path is a hybrid workflow — generative models propose diverse, fully human sequence candidates;

high-throughput experiments identify true (functional) binders; and established engineering refines these hits into developable leads. Just as importantly, MAGE is unlikely to be the final word. Structure- and epitope-guided approaches such as RFdiffusion-based antibody design offer complementary strengths — control over binding geometry and direct structural grounding⁹. The next advances may come not from choosing one paradigm, but from integrating them, turning sequence-based creativity and structure-based precision into a shared community-proven design engine for faster, more reliable therapeutics.

COMPETING INTERESTS

The authors declare no competing interests.

REFERENCES

1. Wasdin, P.T. et al. *Cell* **188**, 7206–7221.e16 (2025).
2. Bondt, A. et al. *Cell Systems* **12**, 1131–1143.e5 (2021).
3. Briney, B. et al. *Nature* **566**, 393–397 (2019).
4. Castelli, M.S. et al. *Pharmacol. Res. Perspect.* **7**, e00535 (2019).
5. Vecchietti, L.F. et al. *mAbs* **17**, 2528902 (2025).
6. Nijkamp, E. et al. *Cell Systems* **14**, 968–978.e3 (2023).
7. Raybould, M.I.J. et al. *Proc. Natl. Acad. Sci. USA* **116**, 4025–4030 (2019).
8. Kryshchak, A. et al. *Proteins* **94**, 5–14 (2025).
9. Bennett, N.R. et al. *Nature* **649**, 183–193 (2025).

ADDITIONAL INFORMATION

Correspondence and requests for materials should be addressed to Albert J. R. Heck. **Reprints and permission information** is available at <https://www.vita-journal.com/>.

© The Author(s) 2026. Published by Higher Education Press. This is an Open Access article distributed under the terms of the CC BY license (<https://creativecommons.org/licenses/by/4.0/>).