Locus: Browser-Based Researcher Credential Verification for DNA Synthesis Screening[1]

Jayani Srinivasan
Independent

With
Apart Research

Abstract

DNA synthesis screening frameworks mandate customer verification as a layer of biosecurity defense, yet no standardized tooling exists to support screeners performing this work. Current practice relies on ad hoc manual searches that are slow, inconsistent, and generate no audit trail.

We present Locus, a Chrome extension that automates researcher credential verification by mapping ORCID publication profiles to NCBI taxonomy identifiers, cross-referencing collaborator networks, and scoring synthesis orders against documented biological research experience. Results render as a five-ring fingerprint visualization for immediate risk assessment. A novel research trajectory feature plots how a researcher's organism focus has evolved over time, detecting deliberate capability acquisition that static credential checks miss entirely.

We evaluated Locus on a 30-pair benchmark across three methodological generations. The final system achieved 80% specificity on legitimate orders, 90% sensitivity on anomalous orders, and 83% overall accuracy. Iterative improvement from keyword extraction through PubTator NER with MeSH primary matching and recency weighting increased specificity from 30% to 80%, identifying PubMed organism extraction quality as the

[1] Research conducted at the AIxBio Hackathon, April 2026

primary performance constraint. The tool is open source, deployable at zero cost by any synthesis provider, and requires no backend integration with existing systems.

1. Introduction

AI is transforming biology faster than biosecurity can keep up. Open-weight biological foundation models such as Evo2 and protein design tools including RFdiffusion have lowered the expertise barrier for novel sequence design. Benchtop DNA synthesizers are approaching the capability to print virus-length sequences within two to five years. Against this backdrop, DNA synthesis screening has become a critical line of defense.

Current screening frameworks focus primarily on sequence-level analysis: does an ordered sequence match a known sequence of concern? Customer verification, determining whether the person placing the order has legitimate credentials, remains underspecified and manually intensive. When a synthesis company receives an order from an unknown researcher, the practical verification process typically amounts to a Google search. This is slow, inconsistent across providers, scales poorly, and generates no reproducible audit trail.

The ORCID system provides a publicly accessible, researcher-maintained record of publications, affiliations, and collaborative networks. NCBI taxonomy provides a standardized organism identification system. No existing tool bridges these systems for synthesis screening. Locus does.

Our main contributions are:

1. *A working Chrome extension that automates ORCID-to-TAXID credential mapping for DNA synthesis customer verification, deployable by any synthesis provider at zero marginal cost and requiring no backend integration.*
2. *A research trajectory visualization that plots how a researcher's organism focus has evolved over time, introducing a new threat model to the field: detection of deliberate biological capability acquisition that static credential checks cannot capture.*
3. *A 30-pair benchmark documenting performance across three methodological generations, with 80% specificity, 90% sensitivity, and clear identification of PubMed organism extraction as the primary performance constraint.*

2. Related Work

Customer verification in DNA synthesis has received growing attention in recent years. The US OSTP Nucleic Acid Synthesis Screening Framework (2024), the UK Biological Security Strategy (2023), and the EU Biotech Act all mandate that providers screen customers and flag orders from individuals with no verifiable legitimate purpose. Despite this regulatory momentum, no standardized tooling exists to operationalize this requirement.

Several organizations are actively developing customer screening infrastructure. Aclid, Cliver, and TwentyTwo are building know-your-customer pipelines for synthesis providers. IBBIS and EBRC are developing screening standards. Existing approaches focus primarily on affiliation verification, government approval checks, and manual credential review. None leverage the public ORCID publication graph as a credentialing signal, and none incorporate temporal trajectory analysis as a biosecurity tool.

SecureDNA's Exemption Certificate system represents the closest prior work, allowing biosafety officers to pre-approve researchers for specific synthesis orders. Locus complements this approach by providing a first-pass automated assessment before human review, using publicly available data that requires no institutional coordination to access.

Locus differs from existing approaches in three ways. First, it operates at the browser layer, requiring no integration with order management systems and enabling immediate deployment. Second, it extends credential verification beyond the individual to include the collaborator network, surfacing research group context invisible to individual-level screening. Third, it introduces research trajectory analysis as a biosecurity signal, answering not just whether a researcher has relevant credentials but whether their publication history shows a pattern of deliberate movement toward dangerous biological capabilities. This last contribution has no precedent in the synthesis screening literature.

3. Methods

Locus is a Chrome Manifest V3 extension with three components: a content script that injects a sidebar into synthesis order review pages, a background service worker that handles all external API calls, and a popup for API key configuration. All processing occurs client-side. No user data is transmitted beyond the APIs required for analysis.

3.1 Identity Resolution

Given a researcher name and institution, Locus queries the ORCID Public API (v3.0) using multiple search strategies in sequence: a structured family name and given name query, a full-text search, and a last-name-only fallback. Results are ranked by affiliation similarity using an affiliation normalization dictionary covering 40 common abbreviations and institutional shorthand (NIH, MIT, JHU, Charite, Max Planck). Without normalization, a screener entering "Charite" would fail to match a researcher whose ORCID lists "Charite Universitatsmedizin Berlin." Up to three alternate ORCID matches are surfaced in the interface so screeners can correct ambiguous name resolution.

3.2 Publication Retrieval

ORCID publication DOIs and PMIDs are retrieved via the ORCID Works endpoint, then fetched from PubMed using the NCBI Entrez API. For each publication, the pipeline extracts title, abstract, MeSH terms, co-author affiliations, and publication year. Publication year is used for recency weighting in subsequent steps.

3.3 Organism Extraction

Early versions used keyword matching against a hand-curated dictionary of 30 organisms, achieving only 30% specificity because scientific papers use variable terminology: a paper on "betacoronavirus" or "2019-nCoV" would not match a keyword list containing only "SARS-CoV-2." The final system addresses this with a three-layer cascade.

PubTator3 NER serves as the primary extraction method. PubTator3 is an NCBI service that annotates PubMed abstracts with standardized biological entity identifiers, including species taxonomy IDs, using a trained biomedical NLP pipeline (Wei et al., 2019). Because annotations are generated from full text rather than surface keyword matching, they correctly identify organisms regardless of author terminology. Up to 50 PMIDs are submitted per batch request.

MeSH primary matching serves as the secondary method. A lookup table maps 35 common MeSH organism terms directly to NCBI taxonomy IDs. MeSH terms are assigned by trained PubMed indexers using controlled vocabulary, making them highly reliable regardless of author terminology.

Keyword extraction against an expanded 200-term dictionary serves as the fallback when PubTator3 and MeSH return no results, covering 30 organisms with 6 to 15 keyword variants each.

3.4 Recency Weighting

Publications are weighted by recency before scoring: papers from the last five years receive full weight (1.0), with decay to 0.7 for five to ten years, 0.4 for ten to twenty years, and 0.2 for twenty or more years. This reflects current research focus rather than lifetime credential accumulation. A researcher who published extensively on influenza twenty years ago but has spent the last decade on computational biology presents a different risk profile than one who published last year.

3.5 Collaborator Network

Co-author affiliation strings from retrieved publications are mined for organism mentions using the same keyword dictionary. This extends credential verification to the research group level. A graduate student with no personal publication history may be ordering on behalf of a supervisor with forty papers on the ordered pathogen — context that is invisible to individual-level screening.

3.6 Anomaly Scoring

Orders are scored against a five-tier schema: EXPECTED, PLAUSIBLE, ADJACENT, UNUSUAL, and ANOMALOUS. The algorithm checks for direct organism match, collaborator match, affiliation verification, and publication depth. CDC Tier 1 select agents receive a penalty multiplier when no research overlap is found, reflecting the asymmetric cost of a false negative on high-risk organisms.

3.7 Fingerprint Visualization

Results render as five concentric arc rings representing identity verification, publication depth, organism breadth, direct organism match, and collaborator network. Ring color encodes status:

blue for verified, amber for adjacent, red for anomalous, grey for unverifiable. Broken arcs indicate missing data. A fingerprint pattern communicates credentialing status immediately; a numerical score requires interpretation.

3.8 Research Trajectory

An optional trajectory visualization plots relative organism publication focus by year, with each data point representing the proportion of that year's papers mentioning a given organism. This reveals shifts in research direction over time rather than raw publication counts. The ordered organism is highlighted. A subtitle identifies whether the researcher shows active (within three years), historical, or no trajectory overlap with the ordered organism.

3.9 Cross-Session Multi-Order Detection

Locus maintains a persistent screening log in local browser storage. When the same ORCID appears across multiple screens with different organisms in a single session, an amber warning flag is raised. This directly addresses the split-order detection requirement in the US OSTP Screening Framework, which notes that sequences ordered in parts can together constitute a sequence of concern.

3.10 Benchmark Design

We evaluated three methodological generations on a 30-pair benchmark. Category A comprised ten legitimate researcher-order pairs: established virologists and infectious disease researchers ordered against organisms central to their documented research. Category B comprised ten anomalous pairs: non-biological public figures ordered against CDC Tier 1 select agents. Category C comprised ten edge cases: CRISPR and synthetic biology researchers ordered against common laboratory organisms. Ground truth labels were assigned based on publicly documented research profiles. All runs were conducted against live ORCID and PubMed APIs using the deployed extension.

4. Results

4.1 Performance Across Methodological Generations

Table 1 summarizes benchmark performance across three methodological generations on the same 30-pair dataset. Specificity improved from 30% to 70% to 80% across generations, driven first by dictionary expansion and PubMed fallback, then by the shift to PubTator3 NER as the primary extraction method. Sensitivity held at 90% across the final two generations. Overall accuracy remained at 83% between the improved and final systems: the 10-point specificity gain in Category A was offset by a 10-point accuracy decrease in Category C, where PubTator3 NER introduced additional organism matches that increased noise on edge cases.

Table 1:

Method	Specificity (Cat. A)	Sensitivity (Cat. B)	Accuracy (Cat. C)	Overall
Baseline (keyword extraction)	30%	100%	70%	67%
Improved (expanded dict. + PubMed fallback)	70%	90%	90%	83%
Final (PubTator NER + MeSH + recency)	**80%**	**90%**	**80%**	**83%**

4.2 Selected Case Results

Table 2 shows selected results across generations. The most informative cases are those where verdict changed. Christian Drosten returned UNVERIFIABLE in the baseline despite being one of the world's leading SARS-CoV-2 researchers — his papers use terminology including "betacoronavirus" and "2019-nCoV" that did not match the initial keyword list. Shifting to PubTator3 NER correctly identified his expertise in the improved generation. Yoshihiro Kawaoka and Robert Gallo showed the same pattern for influenza and HIV respectively.

Vincent Racaniello remained UNVERIFIABLE across all three generations despite being a prominent HIV virologist at Columbia University. His PubMed record uses terminology that does not match current PubTator3 annotations for HIV, making this the clearest example of the tool's sensitivity to publication record completeness and indexing quality.

Table 2:

Researcher	Organism	Baseline	Improved	Final	Ground Truth
Category A: Legitimate orders					
Anthony Fauci	SARS-CoV-2	EXPECTED	EXPECTED	EXPECTED	Low risk
Christian Drosten	SARS-CoV-2	UNVERIFIABLE	PLAUSIBLE	PLAUSIBLE	Low risk
Yoshihiro Kawaoka	Influenza A	UNVERIFIABLE	PLAUSIBLE	PLAUSIBLE	Low risk
Robert Gallo	HIV	UNVERIFIABLE	PLAUSIBLE	PLAUSIBLE	Low risk
Malik Peiris	Influenza A	EXPECTED	EXPECTED	EXPECTED	Low risk
Ralph Baric	SARS-CoV	EXPECTED	EXPECTED	EXPECTED	Low risk
Katalin Karikó	HIV	EXPECTED	EXPECTED	EXPECTED	Low risk
Vincent Racaniello	HIV	UNVERIFIABLE	UNVERIFIABLE	UNVERIFIABLE	Low risk
Peter Doherty	Influenza A	UNUSUAL	UNUSUAL	UNUSUAL	Low risk
Ian Lipkin	MERS-CoV	UNUSUAL	UNUSUAL	UNUSUAL	Low risk
Category B: Anomalous orders					
Jordan Peterson	Variola virus	ANOMALOUS	ANOMALOUS	ANOMALOUS	High risk
Stephen Hawking	*Y. pestis*	UNVERIFIABLE	UNVERIFIABLE	UNVERIFIABLE	High risk
Noam Chomsky	Ebola virus	UNVERIFIABLE	UNVERIFIABLE	UNVERIFIABLE	High risk
Richard Dawkins	Variola virus	UNVERIFIABLE	UNVERIFIABLE	UNVERIFIABLE	High risk
Paul Krugman	*B. anthracis*	UNVERIFIABLE	UNVERIFIABLE	UNVERIFIABLE	High risk
James Watson	*Ebola virus*	*UNVERIFIABLE*	*PLAUSIBLE*	*PLAUSIBLE*	*High risk*
Category C: Edge cases					
Jennifer Doudna	*E. coli*	EXPECTED	EXPECTED	EXPECTED	Low risk
David Baltimore	HIV	EXPECTED	EXPECTED	EXPECTED	Low risk
George Church	*E. coli*	UNVERIFIABLE	PLAUSIBLE	PLAUSIBLE	Low risk
Emmanuelle Charpentier	*S. aureus*	UNUSUAL	UNUSUAL	PLAUSIBLE	Low risk

4.3 Key Findings

Finding 1: PubMed organism extraction is the binding constraint, not identity resolution. ORCID identity resolution succeeded in 29 of 30 benchmark cases (97%). The failure mode was not identity but organism mapping. The shift to PubTator3 NER drove the largest single specificity improvement, from 30% to 70%, by replacing surface text matching with NCBI's own biomedical annotation pipeline. The subsequent addition of MeSH primary matching drove a further 10-point gain to 80%.

Finding 2: UNVERIFIABLE is a meaningful biosecurity signal, not a failure state. In Category B, 8 of 10 anomalous researchers returned UNVERIFIABLE rather than ANOMALOUS. A screener encountering UNVERIFIABLE on a high-risk organism order has a documented, reproducible reason to escalate for manual review. Critically, zero select-agent orders from non-biological researchers were returned as EXPECTED or PLAUSIBLE across any generation.

Finding 3: Broad biological credentials produce boundary false negatives. James Watson returned PLAUSIBLE for Ebola virus due to general biological keyword overlap from his extensive DNA research history. This reveals a fundamental design tension: expanding extraction coverage to improve sensitivity for legitimate researchers simultaneously increases false negative risk at the boundary of biological competence. Domain-specificity weighting is the clearest improvement path.

Finding 4: Research trajectory reveals capability acquisition patterns invisible to static screening.
For researchers like Fauci and Baric, the trajectory chart shows sustained focus on their pathogens across decades. For boundary cases, it reveals whether organism overlap is recent and concentrated or historical and sparse — a distinction that no static credential check can surface. A researcher whose publications show a sharp recent shift toward a dangerous pathogen presents a qualitatively different risk profile than one with lifelong legitimate research in the area.

5. Discussion and Limitations

The central finding is a demonstration of feasibility: public academic identity infrastructure can serve as a real-time credentialing layer for dual-use biology access control. The field has lacked a reproducible, zero-cost approach to customer verification that requires no institutional coordination. Locus provides one. The more significant contribution may be the threat model introduced by trajectory analysis. Static credential checking addresses a snapshot: does this person have relevant experience today? Trajectory analysis addresses the trend: is this person systematically acquiring knowledge of dangerous pathogens? As AI tools continue to compress the expertise required for advanced biological work, the population of people who can order dangerous sequences without obvious red flags will grow. The failure mode is not malicious intent that is easy

to detect but gradual, plausible capability acquisition that looks like legitimate research at any single point in time. The benchmark results identify a specific and actionable bottleneck: PubMed organism extraction quality. Three methodological iterations were required to move specificity from 30% to 80%, each addressing a different aspect of the same problem. This has implications for any system attempting to infer biological expertise from publication records, not just Locus.

Limitations

At 80% specificity, Locus flags one in five legitimate researchers for manual review. In high-volume screening environments this creates meaningful overhead. The root cause is PubMed extraction quality: researchers whose publications use non-standard terminology, publish primarily in non-English journals, or have sparse ORCID records will receive conservative verdicts regardless of actual credentials.

Boundary false negatives remain a structural issue. The tool cannot distinguish between a generalist biologist and an organism-specific expert, a limitation that becomes more significant as AI tools enable researchers from adjacent fields to acquire dangerous biological capabilities without accumulating a publication record that reflects it.

Several threat models are not addressed. Credential theft would pass Locus screening without additional verification. Novel researchers with no publication history present a gap that collaborator network analysis only partially addresses. Cross-provider split orders are detectable within a session but not across providers.

The tool assumes ORCID publication records are a meaningful proxy for biological research competence and that PubMed indexes a representative sample of a researcher's output. Both assumptions weaken for industry scientists, early-career researchers, and those working primarily outside the English-language publication ecosystem.

Future Work

The most impactful near-term improvement is domain-specificity weighting: distinguishing between a researcher who has published broadly in biology and one with deep expertise in a specific dangerous organism. Combining organism match depth, publication recency, citation patterns, and MeSH term specificity could produce a more granular competence signal than binary presence or absence of organism overlap.A collaborative verified-researcher database shared across synthesis providers would address the ORCID coverage gap for early-career and industry researchers, reducing redundant manual review across the industry and building a ground-truth dataset for future model training.Integration with synthesis company order management systems would remove the manual input step entirely, making screening passive and automatic. The auto-detection feature in the current implementation is an early step toward this.The trajectory analysis feature opens a research direction extending well beyond synthesis screening. Monitoring publication trajectories for systematic movement toward dangerous biological capabilities could

inform AI trusted access programs, managed dataset access, and institutional biosafety committee review.

6. Conclusion

Customer credential verification is a named requirement in every major DNA synthesis screening framework and an acknowledged operational gap at synthesis providers worldwide. Locus demonstrates that ORCID-TAXID mapping can automate this layer meaningfully, achieving 80% specificity and 90% sensitivity on a 30-pair benchmark with performance improvements documented across three methodological generations. The primary bottleneck is PubMed organism extraction quality, a finding with implications for any system attempting to infer biological expertise from publication records.
The broader contribution is a demonstration that public academic identity infrastructure can serve as a real-time credentialing layer for dual-use biology access control, and that research trajectory analysis is a viable biosecurity signal for detecting deliberate capability acquisition over time. Locus is a first implementation of these ideas, not a final one.

Code and Data

Include links if applicable. If your project doesn't involve code (e.g., policy analysis) or if there are info-hazard considerations, note that here.

- ***Code repository**: https://github.com/jayanisrinivasan/locus-extension*
- ***Benchmark dataset:** https://github.com/jayanisrinivasan/locus-extension/tree/main/benchmark*
- ***Other artifacts**: Demo video (AI x Bio).mp4*

References

1. *2024 OSTP framework for nucleic acid synthesis screening. OSTP Framework for Nucleic Acid Synthesis Screening. (n.d.). https://aspr.hhs.gov/S3/Pages/OSTP-Framework-for-Nucleic-Acid-Synthesis-Screening.aspx*
2. *Api tutorial: Searching the ORCID registry. ORCID. (2026, January 26). https://info.orcid.org/documentation/api-tutorials/api-tutorial-searching-the-orcid-registry/*
3. *Exemption certification system. Safe Access to SOCs. (n.d.). https://securedna.org/exemption_certification_system/*

4. *U.S. National Library of Medicine. (n.d.). Pubtator3 - NCBI - NLM - NIH. National Center for Biotechnology Information. https://www.ncbi.nlm.nih.gov/research/pubtator3/tutorial*
5. *UK biological security strategy (HTML). GOV.UK. (n.d.). https://www.gov.uk/government/publications/uk-biological-security-strategy/uk-biological-security-strategy-html*

Appendix

Limitations

False positives. At 80% specificity, Locus flags one in five legitimate researchers for manual review. The primary cause is PubMed organism extraction quality: researchers whose publications use non-standard terminology, publish primarily in non-English journals, or have sparse ORCID records receive conservative verdicts regardless of actual credentials. Vincent Racaniello returned UNVERIFIABLE across all benchmark generations despite being a prominent HIV virologist, illustrating how indexing quality can penalize legitimate researchers. If ORCID adoption or PubMed coverage is lower than assumed for a given researcher population, specificity would decrease further.

False negatives. Two of thirty benchmark cases produced false negatives. James Watson returned PLAUSIBLE for Ebola virus due to broad biological keyword overlap. The tool cannot currently distinguish between a generalist biologist and an organism-specific expert. This limitation becomes more significant as AI tools enable researchers from adjacent fields to acquire dangerous biological capabilities without accumulating a publication record that reflects it.

Edge cases. Early-career researchers with few publications present a structural gap that collaborator network analysis only partially addresses. Researchers who publish primarily through preprint servers, conference proceedings, or non-English journals are underrepresented in PubMed and ORCID. Industry scientists who do not maintain public publication records will systematically return UNVERIFIABLE regardless of their actual expertise.

Scalability constraints. The current implementation makes sequential API calls to ORCID, PubMed, and PubTator3, with total analysis time of 10 to 30 seconds per researcher depending on publication volume. At high screening volumes this creates meaningful latency. PubTator3 batch requests are capped at 50 PMIDs, limiting coverage for prolific researchers. A production deployment would require caching of verified researcher profiles and direct API integration rather than browser-layer calls.

Dual-Use Risks

Locus surfaces the absence of verifiable biological credentials, which is publicly observable information. The tool does not reveal synthesis approval patterns, order histories, or internal screener decisions. However, two potential misuse vectors warrant attention.

First, a version of this tool could theoretically be used to identify researchers whose orders are unlikely to be flagged, enabling targeted impersonation. An adversary who knows that researcher X

always returns EXPECTED for organism Y could use that researcher's identity to place an order with reduced scrutiny. This risk is not unique to Locus — it exists for any credential-based system — but the automation Locus provides could make such targeting more efficient.
Second, the trajectory analysis feature reveals which researchers have been systematically publishing on dangerous pathogens. While this information is derived entirely from public publication records, aggregating and visualizing it in a searchable interface could theoretically assist adversaries in identifying domain experts to target for social engineering or recruitment.
Both risks are mitigated by the tool's screener-facing design: Locus is intended for use by synthesis company employees reviewing orders, not as a public-facing search tool. Future versions should restrict API access to verified synthesis provider accounts.

Responsible Disclosure Recommendations
Locus does not interact with biological sequences directly and does not expose synthesis company order data. The primary vulnerability surface is the ORCID and PubMed APIs, which are public and not controlled by this tool.
If a future version integrates with synthesis company order management systems, standard responsible disclosure practices apply: vulnerabilities should be reported privately to the affected organization with a reasonable remediation window before public disclosure. The open-source nature of Locus means the community can audit the codebase for issues that could compromise screener workflows or expose order data.
If the tool identifies a systematic gap in existing screening infrastructure — for example, a class of researchers who consistently evade detection — that finding should be disclosed to relevant biosecurity organizations including IBBIS, SecureBio, and the relevant national biosecurity authority before public publication.

Ethical Considerations
Locus processes publicly available data from ORCID and PubMed. No personal data beyond what researchers have voluntarily made public is accessed or stored. API keys are stored locally in the user's browser and are not transmitted to any server controlled by this project.
The scoring system is designed to err toward caution rather than clearance. This means legitimate researchers will occasionally be flagged for manual review. This is an intentional design choice reflecting the asymmetric cost of false negatives in biosecurity contexts, but it raises a fairness concern: researchers with sparse ORCID records, those at non-western institutions, and early-career scientists are systematically more likely to receive conservative verdicts than established researchers at major institutions. This population bias should be explicitly acknowledged in any deployment context and addressed through the collaborative verified-researcher database described in Future Work.
The trajectory analysis feature introduces a form of longitudinal monitoring of researchers' publication histories. While this data is entirely public, aggregating it into a biosecurity risk signal raises questions about whether researchers should be informed that their publication records

are being used in this way. We recommend that any deployment of Locus include clear disclosure to researchers about how their public data is being used in screening decisions.

Suggestions for Future Improvements
Domain-specificity weighting would address the primary false negative risk by distinguishing between broad biological credentials and organism-specific expertise. A collaborative verified-researcher database shared across synthesis providers would reduce false positives for early-career and industry researchers while building a ground-truth dataset for model improvement. Integration with institutional biosafety committee records would allow Locus to surface approved research protocols as an additional credentialing signal, reducing false positives for researchers with relevant institutional approvals but limited publication records.

LLM Usage Statement

Claude was used to assist with initial code scaffolding and to draft sections of this report. The research design, benchmark methodology, and all empirical results were produced and verified independently by the author using the live extension against real ORCID and PubMed APIs.