


Day 3: Detecting Early Infections (Detect)

Wednesday 28 October · discussion 3:00 to 4:50 pm New York

THE DAY MAP

How do you find out a new outbreak has started, as early as possible?



Your deep-dive: page 3.3. Map one country's biosurveillance and find its holes. The United States, zoomed in on New York City. Its readings come first below; the other options follow.

15 readings here, about 139 minutes of reading. Discussion Wednesday 28 October, 3:00 to 4:50 pm New York, with Conor Spence.

[3.1-1](#) State of Pandemic Early Warning
[original paper →](#)

2 + 15 min

<u>3.1-2</u>	Toward a Global Pathogen Early Warning System original paper →	3 + 0 min
<u>3.1-3</u>	Scaling US pathogen detection original paper →	3 + 0 min
<u>3.1-4</u>	Social Media- and Internet-Based Disease Surveillance for Public Health original paper →	3 + 0 min
<u>3.1-5</u>	Nowcasting epidemics of novel pathogens: lessons from COVID-19 original paper →	3 + 0 min
<u>3.3-1</u>	Toward a Global Pathogen Early Warning System original paper →	2 + 20 min
<u>3.3-2</u>	Emerging COVID-19 success story: Vietnam's commitment to containment original paper →	3 + 13 min
<u>3.4-1</u>	Next Generation Sequencing original paper →	3 + 6 min
<u>3.4-2</u>	Nanopore Sequencing original paper →	3 + 6 min
<u>3.4-3</u>	Towards ubiquitous metagenomic sequencing: a technology roadmap original paper →	3 + 16 min
<u>3.4-4</u>	Would This Have Flagged COVID? original paper →	3 + 6 min

<u>3.5-1</u>	The Triple Rapid Framework for Pandemic Diagnostics original paper →	3 + 14 min
<u>3.5-2</u>	Massively multiplexed nucleic acid detection with Cas13 original paper →	3 + 0 min
<u>3.5-3</u>	Evaluation of the Visby medical COVID-19 point of care nucleic acid amplification test original paper →	3 + 0 min
<u>3.5-4</u>	Recent advances in loop-mediated isothermal amplification (LAMP) for rapid and efficient detection of pathogens original paper →	3 + 0 min

WHAT THE COURSE ASKS YOU TO DO

3.2 Detection deep-dive

Pick 1 of the following 3 topics to dive deeper into. We've given you some resources on the following pages as a starting point.

- **Option 1: Country-specific biosurveillance**
 - Map the current biosurveillance systems in a specific country and investigate how they could be improved
- **Option 2: Metagenomic sequencing technology roadmap**

- Study the current state of metagenomic sequencing and what would be required to make it as ubiquitous as PCR
- **Option 3: Rapid individual testing**
 - Explore emerging technologies for detecting individual cases

You will bring the results of your research to share in your group discussion.

THE READINGS

3.1-1 • CORE • COURSE TIME 30 MIN

State of Pandemic Early Warning

Jeff Kaufman • 2026

[Read the original paper →](#) • the course assigns: whole piece

Early warning should be built against a thinking attacker, which means a system that can detect any pathogen rather than a list of known ones, and the target to aim for is flagging an attack before roughly 1% of people are infected.

Average summary

This is a memo Jeff Kaufman presented at the Summer 2026 Biosecurity Summit outside Washington, lightly edited. He says up front it is **his own view, not a SecureBio house view.**

His starting point is that the threat to design against is an adversary, human or AI, who wants very large-scale harm. An attacker can look at the world's detection systems and aim at the weakest spot, so the landscape is "minimax": a system is only as good as its least sensitive link. That argues for some sensitivity to a wide range of attacks over very high sensitivity to a few. The best outcome is deterrence by denial, where an attacker drops the plan because they can see it would not work, and that requires both building the system and making it known you built it.

He sorts biosurveillance into three jobs. First, **initial detection of a stealth pandemic**, where a pathogen spreads unnoticed through most of the population before it causes serious harm. Second, **triggering the initial response**, where the problem is usually not that nobody saw anything but that nobody acted: in 1976 the US overreacted and vaccinated 45M people against swine flu, in 2014 Ebola in West Africa spread freely for three months before the extent was recognised, and in 2026 Ebola saw another three-month delay, partly because field PCR tests could not see it. Third, **ongoing suppression**, where monitoring tells you where to put resources and whether measures are working. He gives this third one the least attention, because it is the closest to what already exists.

Success, he argues, is not a threshold you cross but a cost-benefit ratio you worsen for the attacker. The sweet spot is a pathogen-agnostic system that flags an attack in time to protect enough essential workers to prevent civilizational collapse and buy time for vaccines and drugs. How early that has to be depends on how fast response moves: if response takes a month you need to flag before about 0.05% of people are infected; at two weeks, 2%; at seven days, even 10% would do.

Today he knows of only four systems doing untargeted metagenomic sequencing for biosurveillance: CASPER, Zephyr, ANTI-DOTE and the UK's mSCAPE. His list of what is missing is long, and he puts independent modeling of the right system scale and red teaming of existing systems at the top.

Why it's assigned

This is the day's one shared core reading, and it sets up the whole detect layer: prevention has failed, someone is already infected, and the question is how we would ever know. Kaufman gives the unit's three elements their shape (notice that something is circulating, work out what it is, find individual cases) and explains why a novel or pre-symptomatic pathogen is the hard case. It also names the technology the course keeps returning to, metagenomic sequencing, and is honest that the field is early.

In depth

The goal, and why the attacker changes the design

Kaufman wants defences robust to adversaries who want very large-scale harm with biology: actors, human or AI, who want to kill all humans, cause short-term incapacitation or long-term civilizational collapse, or who have strategies for sparing some while harming others. An actor need not be directly omnicidal to want this, for instance reducing response capacity during an AI takeover.

The existence of an attacker is itself a design constraint, because the attacker can assess the world's detection systems and plan accordingly. Pushed to the extreme this is a **minimax** landscape: the system is only as good as its weakest link, the place where it is least sensitive. He thinks that framing is right, and that it correctly prioritises broad coverage over deep coverage of a few threats, with two qualifications: where gaps depend on non-public choices you can keep strategic ambiguity so attackers cannot aim at them, and fewer gaps means fewer attacker options either way.

He separates two kinds of deterrence. **Deterrence by denial** is the strong form: an adversary who knows the attack would not accomplish their goal generally does not try it, which depends on detection genuinely thwarting the goal and on the system being known to

exist. **Deterrence by punishment** rests on attribution and retaliation, and works only on an attacker with something to lose: a state seeking strategic advantage might be deterred, while someone willing to die with everyone else has little left to threaten.

Three jobs for biosurveillance

Initial detection of stealth pandemics. A stealth pandemic spreads through most of the population with no or unremarkable symptoms before causing very serious effects. Whether such a pathogen can be built is an open question and he has heard different answers from different experts. Weighing the advances in biological design tools that have come with AI progress, and the slow and unpredictable path from unusual symptoms to attention and action, he thinks there is a significant chance that within the next five years many actors could cause stealth pandemics. He also warns that suspicious sequencing reads may not be enough to tell whether a stealth pandemic is underway: a pathogen might not advertise its own design, so you also need genome completion, estimates of what it would do in the body, and judgement about whether the genome suggests an intentional attack.

Triggering initial response. A pathogen does not have to be stealthy to be disastrous; it could simply be very hard to contain, the "wildfire" case, and could be released deliberately to cut capacity during a coup

or an AI takeover. Either way the time from discovery to serious response is critical, and early signals are ambiguous. His four examples: 1976 swine flu, where the US overreacted and vaccinated 45M people because monitoring could not show the virus was not spreading widely; 2014 Ebola in West Africa, three months of free spread because the extent was not recognised; 2026 Ebola, another three-month delay, partly because field PCR tests could not see it; and 2009 H1N1, which showed a working, though flu-specific, system enabling timely response. In a COVID-weary climate where public health fears losing credibility on a false alarm, he argues good monitoring helps avoid the default of waiting for more information.

Enabling ongoing suppression. Once a response is moving you need monitoring to place mitigations and see if they work. The value depends on how symptoms relate to infectiousness: if symptoms are absent or much delayed, response is essentially impossible without monitoring; if they are immediate and visible, monitoring is only moderately valuable, since you know you have a problem but not its extent. He deliberately spends little of the memo here, because today's capability is closest to what is needed.

What success looks like

There is no single winning threshold. Biosurveillance worsens the attacker's cost-benefit tradeoff; more

capable systems make an attack less likely to look worth the effort and reduce the harm if it happens. Within that he describes four regimes for stealth detection, at rising levels of investment: too slow to beat status-quo detection, which may still help trigger response but gives no initial-detection benefit; a bit faster but still too late for most people, which helps only where harm can be mitigated late; capable enough to protect enough workers that civilization does not collapse and countermeasures can be developed, which he calls significant benefit even though the attack is still massively disruptive; and an extremely capable system with very broad sampling that flags outbreaks while they are still small enough to contain.

None of these are hard boundaries. Luck smears them, for instance who happens to contribute to which samples, as does shared uncertainty, for instance how much a given pathogen would actually shed in a large population.

The figure

The illustrative chart is captioned "Illustrative chart of initial detection benefit by system capability". Its horizontal axis is deliberately vague, standing for resources invested rather than cumulative incidence at detection, because incidence is smeared and capability depends on more than sensitivity; the vertical direction is the detection benefit bought.

What to notice is that the four regimes appear as blurred bands rather than clean steps, and that the curve flattens: past the point where enough workers can be protected to avert collapse, extra sensitivity is still valuable but substantially less cost-effective.

So success is a pathogen-agnostic system that flags attacks in time to protect enough workers to prevent collapse and buy time for pathogen-specific mitigation. How early depends on response speed: a month from detection to action means flagging before about 0.05% of people are infected; two weeks means 2%; seven days means even 10% would do. He also notes that beating status-quo detection is a far higher bar for wildfire than for stealth, so on planning-relevant timescales he does not expect it to be economically feasible to build a system where biosurveillance is your first sign of a wildfire pathogen.

For triggering response, he admits we do not know much about what actually makes decision-makers move. Success is a very short time, ideally under a week, from flagging a catastrophic pathogen to the danger being recognised, PPE distributed, biohardening deployed, lockdowns instituted and countermeasure development started. Those are expensive actions, in money and in political capital and institutional trust, and monitoring's contribution is information: extent of spread, likely trajectory, and

initial wet lab work such as completing the genome. Response also requires trust, built over time, which is why non-catastrophic detections matter: they show the system tracks trends that match other evidence, surfaces real public health concerns, and turns up genuine engineered "benign positives". Most of the work in cutting response time, he says, is not biosurveillance at all; it is better agency plans, exercises that put the actual decision-makers through making these calls with incomplete information, and smoother inter-agency communication.

For suppression, he notes wastewater PCR was used by Australia, New Zealand and Singapore, the last down to building level, to guide COVID-19 response. Suppression monitoring is the same component as response triggering, only larger and finer-grained: knowing something is spreading in ten major US cities might spur action, but holding suppression needs a much more detailed picture.

What gets us there

Sampling. Municipal wastewater is very helpful because in most cities sewage is processed at a small number of sites, so one easily collected sample covers hundreds of thousands of people. But many pathogens do not shed heavily into wastewater and it is a very noisy sample type, so comprehensive detection at reasonable cost likely needs more sample types. SecureBio runs a nasal swab program;

beyond swabs he rates air, blood, aircraft wastewater and leftover material from clinical tests as the strongest candidates, while saying much could still be learned.

Lab technology. The technology has to be pathogen agnostic, because if you monitor only specific pathogens the adversary picks one you do not monitor, and with current technology that means sequencing. For viruses it is practical to concentrate particles by size and then do untargeted and enriched metagenomic sequencing, pushing the rest of the detection into software, which is what SecureBio does today. For bacteria he is pessimistic about copying the approach, at least in wastewater, because bacterial genomes are much larger and sewer bacteria cannot be physically separated from threatening bacteria the way viruses can. For mirror life, early spread might be very hard to recognise, and the important step would be learning it was spreading at all; he thinks a sensitive, relatively cheap assay could probably be developed, but nobody has started.

Computational technology. Sequencing moves the problem into the computer, which means more reads than humans could ever read, by his estimate 8 or more orders of magnitude. Some are easy to call concerning, for instance sequence unique to the smallpox genome turning up in wastewater; others need very sophisticated processing, such as

understanding the background well enough to flag a genome with no similarity to anything that exists. He gives four signal types: **Dangerous** (matching a known pathogen you would not expect in that sample), **Modified** (partial matches, perhaps engineered), **New** (never seen before, perhaps made from scratch), and **Growing** (becoming more common, perhaps spreading through people). This is hard because of data scale, uncertainty about what an attack looks like, and the need for fast analysis, but he expects AI to be very productive here, unlike the parts of the system that need slow real-world work.

What exists today

He knows of four systems doing untargeted metagenomic sequencing for biosurveillance.

CASPER (SecureBio with Marc Johnson's lab at the University of Missouri and other academic partners) covers wastewater, mostly municipal, from 49 facilities representing 24 US cities; virus-focused, very deep short-read sequencing. **Zephyr** (SecureBio with Helena Solo-Gabriele's lab at the University of Miami) is pooled nasal swabs from Boston and Miami collected mostly in public places, about 1,000 people weekly, virus-focused, long-read. **ANTI-DOTE** (DoW with PHC and SecureBio) is wastewater from five US military facilities, processed through the same pipeline as CASPER. **mSCAPE** (UK government) is bronchoalveolar lavage from UK hospitals, currently

relatively few samples, which limits sensitivity; alone among the four it sequences viruses and bacteria together, at relatively low depth with long reads, serving clinical practice, public health and biodefense. He adds that several hybrid-capture projects might catch an attack if the agent were similar enough to an existing pathogen, and that an outbreak can still become visible the old way, through symptoms in enough people, suspicious clusters, then sequencing that shows an edited-looking genome; that path is not well developed today and he would like to see investment in it.

What is missing

He prefaces the list with a caution that it is his current thinking, not as thoroughly considered as he would like, and asks readers not to overweight it. He also says the ideal structure is an ecosystem of organisations working in parallel, and that SecureBio can often share samples, sequence prepared nucleic acids or share data.

Roughly in descending order of how valuable he thinks non-SecureBio work would be: **modeling** of the necessary system scale and network design, which he most wants done outside SecureBio as a check on their work; **red teaming** existing systems against adversarially designed attacks; **parallel organisations** in other regions, especially Europe and Asia; **stealth bacterial detection**, which he

personally expects to be very hard; **parallel methods development**, where he would especially like deep learning tried, which is part of why much of their sequencing data is public (PRJNA1247874, PRJNA1379685); **more sample streams**, with blood, air and clinical lab discards most promising; and **response work**, figuring out what would get decision-makers to act fast, which he calls delicate, warning that someone coming in noisily and insensitively could set the field back.

Then, lower down: **genome completion**, since a short-read system like CASPER starts with only a suspicious fragment; **analysis tooling** to get from "this genome looks worrying" to whether it is engineered and what it would do in humans; easier **swab collection**, since Zephyr needs field samplers on street corners bringing in about 40 samples an hour, where workplaces or schools building sampling into the daily routine would scale far better; **clinical sequencing**, which he thinks will arrive on clinical merit eventually but far too slowly by default, and calls a commercial opportunity; **sensitivity increases**, where he expects large undiscovered gains; **deterrence tracking**, an automatically updated tracker of what frontier models say about whether an attack would succeed, on the logic that if models tell defenders the probability is low they are probably telling attackers the same; **partially targeted methods** such as hybrid capture, which might

sensibly form one component of a mature system; **cheaper protocols** and **shorter lab time**, where current work takes about a day on the bench and a day on the sequencer, with new machines coming in the roughly 8 hour range; **fungi, parasites and prions**, which he notes are generally much harder for an attacker; and **mirror life**, which current sequencing would not detect at all and which he ranks last only because he expects it to take long enough to develop that there will be time to build an assay later.

How this changes for a wildfire pandemic

During COVID-19 many groups built PCR-based targeted wastewater monitoring: in the US, NWSS at CDC, the philanthropic WastewaterSCAN and the private Biobot, while EU member states, Canada, Australia and others treat wastewater as standard public health. He estimates maybe two dozen countries have some wastewater PCR that could be retargeted to a new pathogen once its genome is known. None of it would move fast enough today for a wildfire pandemic: some delays are technical, such as stocking consumables, but most are organisational, such as policies that let production work be set aside, overtime budgeting, on-calls and rush-order deals with synthesis providers. Even in a serious emergency he calls ten days a good best case today, with three days possible given good preparation, and says

getting existing networks ready for rapid turnaround is really valuable work SecureBio has begun discussing.

The same metagenomic system built for stealth detection could also cut days off a wildfire response, because untargeted sequencing skips the physical-world changes that retargeting PCR requires, at the cost of much lower sensitivity. Once you know the target, PCR is significantly more sensitive, by about a factor of 100 for SARS-CoV-2, so a system able to flag a stealth pandemic before 1% infection could confirm a wildfire pandemic at about 0.01%. He is not sure even that is early enough; it might need 0.001% or lower, which argues either for a substantially larger investment in untargeted sequencing or for giving up on sequencing for this job. Deciding firmly would need estimates of what fraction of the population would already be infected when a wildfire pandemic is noticed, and how many doubling periods would pass before decision-makers acted without the system. He is not sure a firm decision is needed, and leans towards groups exploring both in parallel.

How this changes for ongoing suppression

The same COVID-era PCR wastewater capacity could serve suppression monitoring, which he thinks is already widely understood to be valuable and still under-invested. In a real emergency he expects governments to organise existing capacity reasonably well, if not as fast as ideal; his bigger worry is whether

the capacity is large enough. Monitoring ideally reaches neighborhood or building level, which needs a very large number of in-manhole composite samplers, and those are low-volume products from a small number of manufacturers, so you could not build many quickly in a crisis. The main job is building deployable capacity in advance, and the best bet is systems installed for everyday public health monitoring, because that proves they work as part of a larger system, spreads know-how, and delivers ongoing benefit, which may be an easier sell than stockpiling.

Appendix: how the 1% target is derived

He sets out three parameters and is clear that knowledge is limited at several steps. **Doubling period:** about 3 days, which he calls the high end for naturally occurring pathogens, and he notes you could argue lower, because a pathogen could be designed to spread faster, or higher, because of physical limits on spreading fast without attracting attention.

Fraction of workers to protect: 50%, following Patel et al. in *Physical Approaches to Civilian Biodefense*, who chose 50 percent of vital workers as a convenient midrange target because operators likely have flex capacity for a small fraction being absent but need more than single-digit percentages protected to keep national critical functions running; the difference between 35% and 70% is only a factor of two in

required sensitivity. **Response lag:** about 15 days, which he says is not where the world is today but is achievable, and here he makes the key point that every doubling during response demands twice the sensitivity, so because required sensitivity grows exponentially with lag, cutting the lag becomes the best use of marginal dollars after a relatively small initial investment. Together that gives SecureBio's long-standing target: flag a pathogen before very roughly 1% of people are infected, since 1% grown through five doublings (15 days at 3 days each) is 32%, under the 50% of workers you need to protect.

Words to know

- **Stealth pandemic.** A pathogen that spreads through most of the population unnoticed, with no or unremarkable symptoms, before causing very serious effects. It matters because detection, not treatment, is the binding constraint: by the time anyone feels ill it is already everywhere.
- **Wildfire pandemic.** A pathogen that is simply very hard to contain. It matters because the time between discovering it and acting is where the damage is decided.
- **Metagenomic sequencing.** Reading all the genetic material in a sample at once rather than testing for one organism you already suspect. It matters because it is the only currently practical way to be pathogen agnostic.

- **Pathogen agnostic.** A system that does not need to know in advance what it is looking for. It matters because a monitored list is a map of the gaps for an attacker.
- **Minimax.** Here, the idea that a defence is only as strong as its least sensitive point, so an attacker aims there. It matters because it argues for broad shallow coverage over narrow deep coverage.
- **Benign positive.** A real engineered sequence the system flags that turns out not to be an attack. It matters because these are how a system earns the trust it needs before anyone will act on a real alert.
- **Hybrid capture.** A partly targeted sequencing method that fishes out sequences resembling ones you specify, tolerating mismatches. It matters because it is far more sensitive than untargeted sequencing but only within the range of threats you anticipated.
- **Mirror life.** Organisms built with the mirror-image versions of life's molecules. It matters here only as a detection problem: current sequencing would not see it at all, so it would need an assay of its own.

Numbers in this reading

- Five years: the window in which he thinks there is a significant chance many actors could cause stealth pandemics.
- 45M people vaccinated in the 1976 swine flu overreaction, because monitoring could not show it

was not spreading widely.

- Three months of free spread in 2014 Ebola in West Africa, and another three-month delay in 2026 Ebola.
- 0.05%, 2% and 10%: the cumulative infection levels at which you would need to flag a pathogen if response takes a month, two weeks or seven days.
- Under a week: his target for the time from flagging a catastrophic pathogen to PPE, biohardening, lockdowns and countermeasure development being underway.
- 8 or more orders of magnitude: how far sequencing read volumes exceed what humans could evaluate.
- 49 wastewater facilities representing 24 US cities: CASPER's coverage.
- About 1,000 people weekly: Zephyr's pooled nasal swab intake from Boston and Miami.
- Five US military facilities: ANTI-DOTE's wastewater sites.
- About 40 samples per hour: what a Zephyr field sampling team brings in, which he gives as the reason to move collection into workplaces or schools.
- Ten days best case today, three days with good preparation: his estimate for retargeting existing PCR wastewater networks in an emergency.

- Two dozen countries: roughly how many have wastewater PCR that could be retargeted once a genome is known.
- A factor of about 100: how much more sensitive PCR was than sequencing for SARS-CoV-2 once the target was known, which turns a 1% stealth threshold into about 0.01% for wildfire confirmation.
- 0.001%: the sensitivity he thinks might actually be needed for wildfire, and cannot rule in or out.
- About 3 days: the doubling period they work from, the high end for naturally occurring pathogens.
- 50% of vital workers, from Patel et al., with 35% and 70% noted as only a factor of two apart in required sensitivity.
- About 15 days: the response lag they assume, which he says is not today's reality but is achievable.
- 1% infected: the detection target, because 1% through five doublings is 32%, under the 50% worker protection threshold.
- About 8 hours: the run time of new sequencing machines, against roughly a day on the bench plus a day on the sequencer today.

What the author is unsure about

A lot, and he says so. Whether a pathogen stealthy enough to spread unnoticed can actually be built is an

open question, and he has heard different answers from different experts. Whether suspicious reads would even tell you a stealth pandemic was underway is not guaranteed. He does not know what would actually make decision-makers act promptly. He says the "what's missing" list is not as thoroughly considered as he wishes and asks readers not to overweight it in their own decisions. He expects stealth bacterial detection to be very hard but notes people's estimates vary widely. He cannot say whether 0.01% is early enough for wildfire detection, and he does not think a firm choice between sequencing and PCR for that job is needed yet. On his own thresholds he is explicit that none are hard boundaries and that luck and shared uncertainty smear them. The whole memo is labelled his view rather than SecureBio's.

For the discussion

- The strongest point in the memo is that the attacker reads our system too, so a detection programme that monitors a list of known pathogens is effectively publishing its own gaps, which is why pathogen-agnostic sequencing is the design choice rather than a technical preference.
- Worth raising as a tension: he calculates that cutting the response lag beats buying sensitivity, since every doubling during response demands twice the sensitivity, yet almost all his concrete recommendations are sampling and sequencing

work and he admits the response side is "mostly not biosurveillance" and nobody knows what makes officials move.

3.1-2 · FURTHER · COURSE TIME 50 MIN

Toward a Global Pathogen Early Warning System

Natasha E. Bajema, William Beaver, Christine Parthemore · 2021

[Read the original paper →](#) · the course assigns: pages 4 to 18 (executive summary, Part I introduction, Part II through the genomic surveillance section)

The world's outbreak detection still depends on a clinician noticing something strange, and the authors argue the technologies to replace that are already here, so the real work now is coverage, interoperability, political will and money.

Average summary

This is a report from the Council on Strategic Risks, written a year and a half into COVID-19. The authors say the vision was set long ago, in the **World Health Organization's 2005 update to the International Health Regulations**, which called for every nation to be able to detect and report emerging outbreaks. New data-sharing systems, methods and tools followed. COVID-19 showed they were still not enough: delays

and weaknesses in detecting and tracking SARS-CoV-2 contributed to more than 4 million deaths and 187 million infections, strained relations between nations and severe economic damage, and they note the toll was still mounting as they wrote. Other biological threats, they add, could be far worse.

Their account of how detection actually works is the part the course cares about. The global system leans on clinicians noticing a suspicious pattern, then notifying public health authorities, then lab confirmation and epidemiology. It depends on so-called "astute clinicians" spotting atypical cases, and it is unreliable and slow. Symptoms alone are a poor signal, since many viruses produce similar ones, which is why COVID-19 kept being described as flu-like or pneumonia-like.

Days matter, because every case creates two to three more unless something interrupts it. Their contrast: in 2004 two likely distinct lab accidents with SARS, a virus slightly less contagious than SARS-CoV-2 but close to ten times deadlier, were caught and stopped before they spread beyond about a dozen people, using contact tracing of hundreds of possible cases and quarantine of the lab's workers. COVID-19, detected late, demanded a response compared to mobilisation for the world wars, including vaccine and therapeutic work projected to cost \$157 billion through 2025.

They set out the types of surveillance (indicator-based, event-based, genomic, sentinel, syndromic and vector), then describe what a global system would need: coverage of high-risk nodes such as major hospitals, zoonotic hotspots, high-volume transport hubs and countries holding dangerous pathogen specimens in biosafety level 3 and 4 labs; technologies that fit very different settings, from dense cities to places with no reliable electricity or internet; and interoperability across all of it. The hard parts they name are not technical: paper records, missing reagents, too few trained people, no sustainable funding for labs once built, countries that will not share data for fear of blame, and agencies inside one country that will not share with each other.

Their technology bet is genomic and then metagenomic sequencing, which reads all the DNA and RNA in a sample and can look for thousands of pathogens at once instead of the one you suspected.

Why it's assigned

This is the day's background on the detect layer: before you can argue about metagenomic sequencing or wastewater, you need to see how outbreak detection works today and why it is slow. The course note says exactly this, that we rely on medical workers noticing a suspicious pattern in ill people arriving at their clinic. Pages 4 to 18 give you the current system and the authors' vision of a stronger one; the same

report comes back on page 3.3 for its country case studies.

Words to know

- **Biosurveillance.** An umbrella term for collecting and analysing disease-related information. It matters because the authors show it covers wildly different activities, so two people can both claim to have it and mean different things.
- **Syndromic surveillance.** Watching patterns of symptoms, for example a jump in people arriving at emergency rooms with breathing problems, instead of waiting for diagnoses. It matters because it is faster but vague.
- **Event-based surveillance.** Finding signals in messy sources, which CDC describes as reports, stories, rumors and other information about health events that could be a serious risk. It matters because it can bypass slow official channels.
- **Sentinel surveillance.** A network of doctors and labs voluntarily sharing data on pathogens or trends. It matters because random sampling inside it caught the first community transmission of COVID-19 in Seattle.
- **Metagenomic sequencing.** Reading all the genetic material in a sample rather than testing for one suspect. It matters because it is the only method here that can find something nobody has seen before.

- **Interoperability.** Whether separate systems can connect and have their data combined. It matters because the authors treat it, not sequencing, as the binding constraint.

Numbers in this reading

- More than 4 million deaths and 187 million infections: the COVID-19 toll they attribute in part to detection and tracking failures, still mounting as they wrote.
- Two to three: how many new cases each case creates, absent disruptive action.
- Close to ten times deadlier: SARS compared with SARS-CoV-2, though slightly less contagious.
- More than a dozen people: the limit of spread from the two 2004 SARS lab accidents, stopped with contact tracing of hundreds of possible cases and quarantine of lab workers.
- \$157 billion through 2025: the projected cost of COVID-19 vaccine and therapeutics development and manufacturing.
- 23 days to detection and 35 days to lab confirmation: the median from estimated outbreak start, between 1996 and 2014, improving on average 7.3% per year, with evidence the pace has since slowed.
- 39 of hundreds of known human-infecting viruses had FDA-approved diagnostics as of October 2018, by an analysis in Nature.

- 5%: the share of samples some experts think genomic surveillance would need to cover to revolutionise global early warning.
- Less than a day: how fast metagenomic sequencing can get from clinical presentation to diagnostic confirmation.
- Up to a thousand pathogens at once for existing metagenomic products; IDbyDNA identifies 900 microorganisms from respiratory samples, with Karius similar.
- \$1.7 billion: the US investment announced in early 2021 to add sequencing capacity, initially to track COVID-19 variants.
- Five years: their horizon for metagenomic sequencing becoming available and in frequent use in every nation.

What the author is unsure about

The authors are careful to call this a snapshot of a changing landscape, drawn from an initial six-month mapping effort, and say fully mapping the global landscape would take significant resources. They say the assessment is not meant to cast blame. They state plainly that getting genomic surveillance to useful coverage is necessary but not sufficient, and that current costs and complexity of metagenomic sequencing are prohibitive, so deployment will be limited at first. The 5% coverage figure is attributed to what "some experts believe", not presented as a

finding. They also predict syndromic surveillance may become less relevant over time, which they present as a judgement about where the technology is going.

For the discussion

- The sharpest line in these pages is that the world's early warning system is a clinician's hunch with paperwork attached, and that the authors' own numbers (23 days to detection, 35 days to lab confirmation, improving 7.3% a year and slowing) show how little that has changed.
- Worth raising: almost nothing on their list of fixes is a laboratory problem. It is sustainable funding for labs after they are built, political will to share data, and trust between agencies inside the same country.
- A question worth asking the group: the 2004 SARS lab accidents were contained at a dozen people with contact tracing and quarantine, so how much of that success was detection and how much was that SARS makes people obviously, quickly ill.

3.1-3 · FURTHER · COURSE TIME 20 MIN

Scaling US pathogen detection

Simon Grimm, Lennart Justen, Jeff Kaufman · 2025

[Read the original paper →](#) · the course assigns: whole piece

For the **\$52M the President's FY 2026 Budget proposes for CDC's Biothreat Radar, the authors model a real national detection system and estimate it would catch a brand-new pathogen before 12 in 100,000 Americans were infected, and a known one before 1 in 100,000.**

Average summary

This is a costed design, not an opinion piece. The authors work at what was the Nucleic Acid Observatory and is now SecureBio Detection, and they take the \$52M line item for Biothreat Radar and build out what it would buy.

Their design deliberately uses collection programmes CDC already runs, and spends the new money on processing and analysis. Samples come from two sources. **Traveler-based Genomic Surveillance (TGS)** collects nasal swabs and wastewater from international travellers, which they like because travellers are disproportionately likely to be infected early in a pandemic, because nasal swabs and aircraft lavatory waste carry high pathogen concentrations, and because it already runs at busy airports with room to grow. The **National Wastewater Surveillance System (NWSS)** samples municipal wastewater and covers more than 100M people, and serves as the complement that can prove community transmission.

Three sample types: airplane waste through triturator sampling at BOS, JFK and SFO plus pooled airplane sampling at IAD and LAX, with ten more triturator sites added, covering more than 100,000 people a day for \$7.8M a year; nasal swabs at six current TGS airports (SEA, SFO, JFK, EWR, IAD, MIA) plus seven more sites, 400 swabs per site per day for 5,200 a day, at \$15.9M a year; and municipal wastewater at five NWSS sites, each covering more than 500K people for 2.5M in total, at \$840K a year. Everything is sequenced daily at two private laboratories, each running one Illumina NovaSeqX+ with a 10 billion read flow cell, for \$24.3M, with \$2.3M for storage, processing and analysis labour. Total \$50.8M including 35% overhead.

The performance claim splits in two. **Monitoring** means scanning for fragments matching known threats such as H5N1, smallpox or MERS, which they model as catching an introduction before 1 in 100,000 Americans are infected. **Discovery** means finding something nobody has ever seen, which they model as before 12 in 100,000. Both numbers come from 10,000 simulated outbreaks using influenza A characteristics, with detection defined as any part of the genome appearing in at least 2 samples across 2 sampling systems for monitoring, and 2x average genome coverage for discovery.

Four things have to be right: deep sequencing, because wastewater needs around 1 billion reads per sample to see past sewer and gut bacteria; centralised sequencing, because only a big machine gets the cost per read low enough; turnaround of 2 to 3 days, achieved by running around the clock and siting labs near the airports; and fast data sharing, so outside groups can cross-check and sort real threats from false positives.

Why it's assigned

This is the detect layer with a price tag on it, and it is the reading page 3.3 points to when it asks what the critical challenges would be in setting up a national biosurveillance system like the one described for the US. If you pick the United States for the deep dive, this is the concrete proposal to argue with: named airports, named programmes, a named machine, a cost per component, and a claimed sensitivity you can check against how fast a response would actually happen. Two of the sampling sites, JFK and EWR, serve New York.

Words to know

- **Pathogen-agnostic.** A system that does not need to know in advance what it is looking for. It matters because it is the difference between monitoring and discovery in this post.

- **Metagenomic sequencing (MGS).** Reading any nucleic acid, RNA or DNA, in a sample rather than testing for one target. It matters because it is what lets the same sample serve both known and unknown threats.
- **Triturator.** An airport ground facility that receives aircraft lavatory waste from trucks and grinds the solids before discharging into the municipal sewer. It matters because it is the physical place one of their cheapest, widest samples is taken; each triturator sample captures waste from roughly 7,500 people.
- **Reads, and read depth.** One read is one fragment of genetic sequence the machine outputs; depth is how many you take from a sample. It matters because sensitivity is bought in reads, and their whole cost model turns on it.
- **Cumulative incidence at detection.** What share of the population has ever been infected at the moment the system notices. It matters because it is the single number they use to score the system.
- **Benign positive.** A real engineered sequence that is not an attack. It matters because they found one: an edited cytomegalovirus sequence in wastewater, which the relevant company confirmed came from vaccine work in one of their catchment areas.

Numbers in this reading

- \$52M: the FY 2026 Budget allocation proposed for Biothreat Radar, split equally between TGS and AMD.
- Their split: \$23.7M to TGS, \$26.3M to AMD, \$0.8M to NWSS, \$50.8M total including 35% overhead.
- 12 in 100,000 Americans infected: detection threshold for a novel pathogen. 1 in 100,000: for a known pathogen entering the US.
- \$7.8M a year for airplane wastewater at 15 airports, covering 100,000 people per day; \$15.9M for nasal swabs at 13 airports, 400 per site per day, 5,200 a day; \$840K for five wastewater sites covering 2.5M people.
- \$24.3M for sequencing at two private laboratories; \$2.3M for storage, processing and analysis.
- 10 billion reads per day per laboratory on one NovaSeqX+ 10B flow cell; about 1 billion reads per wastewater sample, roughly 10% of a run; 150 million reads, or 1.5% of a flow cell, for a pooled swab sample.
- 10,000 simulated outbreaks, modelled on influenza A, with 50% of the genome assumed too similar to other pathogens to be useful for monitoring.
- Doubling times of only 3 days for many pandemic pathogens, which is why they target 2 to 3 day end-to-end turnaround.
- Their own pilots: 1,362 billion read pairs of wastewater MGS generated, 5,500 swabs analysed

for viral content, wastewater collaboration active in Boston, Chicago, Florida and 16 other sites, up to 70 swab samples per hour collected by two field samplers, and dashboards tracking dozens of viruses including 50+ rhinovirus serotypes.

What the author is unsure about

They are modelling, and they say so throughout. The design rests on publicly known information about Biothreat Radar plus their own assumptions. They leave clinical samples out of the model, not because clinical sequencing is unimportant, but because detailed public information on the Advanced Molecular Detection programme's throughput and processing speeds is limited; they add that pooled clinical sequencing could be a highly valuable extra layer. The performance figures depend on the pathogen behaving like influenza A on doubling time, shedding duration, shedding amount and genome length. They expect false positives from seasonal variation in common viruses, novel elements of the sewer microbiome and sequencing artifacts, and they describe a validation protocol that would need to be developed and iterated with other agencies rather than one that exists. They note the \$0.8M wastewater component is not in the President's Budget at all, it is their own addition.

For the discussion

- The useful move here is to turn the sensitivity claim into a question about response: a system that flags at 12 in 100,000 is only worth 12 in 100,000 if the government acts quickly, and this post's own scenarios depend on stockpiles being released, hospitals adjusting and the public being told.
- Worth raising: the thing they actually detected in a pilot was not a pathogen but a benign engineered sequence from a vaccine company in their own catchment area, which is exactly how you would expect a real system to spend most of its alerts.
- A question worth asking: the whole design hinges on centralising sequencing in two laboratories for cost reasons, so what happens to a 2 to 3 day turnaround if one of them goes down.

3.1-4 • FURTHER

Social Media– and Internet-Based Disease Surveillance for Public Health

Allison E. Aiello, Audrey Renson, Paul N. Zivich • 2020

[Read the original paper →](#) • the course assigns: whole review, summarized from a structured section-by-section account of the open-access article rather than from its original text

Search queries, tweets and symptom-reporting apps can see illness that never reaches a clinic, but this review's verdict is that they measure public attention as much as

disease, so their place is alongside traditional surveillance rather than instead of it.

Average summary

A note on the source first: this is summarized from a structured account another reader wrote while reading the open-access review in a browser, keeping all of its sections, rather than from the original text, so nothing here is a direct quotation. Its 118-item bibliography was not captured, though the key studies are named inline. The article is open access, so exact wording can be pulled from the DOI if it is ever needed.

Aiello, Renson and Zivich review what digital disease surveillance has actually delivered. They define it as using internet-based data to nowcast (track the present) or forecast (predict the future) disease incidence, and they sort the sources into three kinds: aggregate search trends such as Google searches and Wikipedia page views, social media posts, and participatory systems such as **Flu Near You** and **Influenzanet**, where people answer repeated symptom surveys. They note a distinction that runs through the whole review: searching for health information is not the same as being ill, and a tweet describing symptoms is better evidence than a visit to a Wikipedia page.

Their anchoring example is **Google Flu Trends**. CDC's weekly influenza-like illness reports usually lag

one to two weeks, so Google built a linear regression on 45 search queries highly correlated with the influenza time series, pruned by hand (basketball terms, which track flu seasonally, were removed manually). After its 2008 release it predicted influenza-like illness one to two weeks ahead for two seasons. Then it missed the 2009 A/H1N1 pandemic, and after a model update both versions vastly overestimated the 2012-2013 peak. The public site was removed in 2015.

From that failure they build a list of biases. **Concept drift:** search companies change their algorithms without notice, so models decay; a flu model trained on Wikipedia traffic went stale after four months and had to be retrained, and more data over a longer period did not help. **Nonindependence:** trending topics feed themselves, and a CDC press conference about H7N9 in April 2013 produced a spike in flu tweets unrelated to incidence. **Confounding:** Google trends showed a US cholera epidemic in 2007 that turned out to be Oprah Winfrey choosing *Love in the Time of Cholera* for her book club.

Representativeness: roughly 22% of US adults use Twitter, skewed to higher socioeconomic status, ages 30 to 44 and urban residents, which underrepresents exactly the groups at highest risk of infectious disease death; geolocation is off by default and enabled by only about 1 to 2% of users. Participatory systems are not exempt, since Influenzanet

participants look healthier and more advantaged than the population they stand for.

Their ethics section runs through beneficence, nonmaleficence, autonomy, equity and efficiency, and the sharpest point in it is that few digital systems have described what intervention followed an earlier detection, or whether it worked.

Why it's assigned

This is the detect layer's other route: not laboratories and sequencers, but what people type, post and report themselves. It explains why event-based and digital systems can see an outbreak before official diagnoses arrive, and it is honest about the limits, including the industry's most famous failure. It also closes the loop with the day's other readings, which keep saying that a signal only matters if someone acts on it.

Words to know

- **Nowcasting and forecasting.** Tracking the present state of incidence in near real time, versus predicting the future. It matters because a model can be excellent at one and worthless at the other, which is exactly what this review reports.
- **Concept drift.** When the thing you are measuring changes underneath a fixed model, so accuracy decays over time. It matters because search

companies change their algorithms with no notice, and online language shifts.

- **Participatory surveillance.** People volunteering symptom reports directly, as in Flu Near You and Influenzanet. It matters because it is the only digital source that collects individual demographics and risk factors, and the only one where consent is clean.
- **Influenza-like illness (ILI).** A syndromic case definition counted by providers, used as the ground truth most digital models are scored against. It matters because the comparator has its own biases, and prediction reproduces them.
- **Poststratification weighting.** Reweighting a non-random sample to look like the population. It matters because it is their main proposed fix for sampling bias, and the precedent they cite is a set of daily voter-intention polls run on the Xbox platform that still called the 2012 US election.
- **Digital divide.** Inequality in internet access and use, between and within countries. It matters because it decides who a digital system can see at all.

Numbers in this reading

- 45 unique search queries in the original Google Flu Trends regression, manually pruned for relevance.

- One to two weeks: how far ahead Google Flu Trends predicted influenza-like illness for the 2007-2008 and 2008-2009 seasons, and also the usual lag in CDC weekly reports.
- 2009 A/H1N1: the pandemic it missed; 2012-2013: the season both the original and updated models vastly overestimated; 2015: the public site's removal.
- Four months: how long before an influenza model trained on Wikipedia traffic went stale and needed retraining.
- Two to four weeks: the forecasting horizon beyond which accuracy degrades rapidly, across multiple diseases.
- Roughly 22% of US adults use Twitter; about 1 to 2% of users enable geolocation, and those who do are more likely to be men over 40.
- 1%: the random sample of raw tweets Twitter's data policy made public, which the review says turned it into a "model organism" for digital research.
- Six of nine teams in FluSight's first season used Twitter, alone or with other digital sources.
- Approximately 2,200 outpatient providers report CDC influenza-like illness percentages weekly, with figures often revised weeks or months later because some report late.
- Influenzanet was established in 2009 across 10 European countries, built on a Dutch and Belgian

platform launched in 2003-2004.

- Two to four weeks: the lead time traditional systems give, which the review calls insufficient for planning hospital surge capacity or vaccine manufacturing.
- Doubled accuracy: the improvement one team got in predicting the weekly direction of change by filtering media chatter out of influenza tweets.

What the author is unsure about

They are direct about the gaps. No study has examined how underrepresenting the very young and very old affects population-level flu estimates from internet data. No investigation has examined the relationship between enabling geolocation and disease risk, though they suggest it probably correlates. How often a model needs recalibrating is open. Sampling bias and nonrepresentativeness are the least explored of the biases, and the demographic variables needed to weight for them are usually not available from Google or Twitter data. They warn that rising internet use does not necessarily remove sampling bias, because the remaining holdouts become an increasingly unusual group. On ethics they say it is unclear whether consent can be set aside for individual-level health data drawn from nonhealth sources, and they call this a tumultuous moment for data privacy. Their overall verdict is a judgement, not a finding: a system driven solely by digital technology

seems unlikely now, and the largest impact is likely to come from integration with laboratory data, case reports and electronic health records. The disclosure statement notes that one author has consulted for Kinsa Inc. on studies of its smart thermometer, and that two authors received funding from GOJO Industries for hand hygiene research.

For the discussion

- The line worth carrying into the discussion is that digital sources mostly measure public awareness rather than disease, which is why a model can nowcast beautifully and forecast nothing.
- Worth raising as a tension: the sources that reach furthest beyond the clinic are the least representative, and they are least representative in exactly the populations with the highest infectious disease mortality.
- A question worth asking: this review's own beneficence test is that nobody has shown what intervention followed an earlier digital detection, so what would have to be in place for a one-week-earlier signal to change anything on the ground.

3.1-5 • FURTHER • COURSE TIME 40 MIN

Nowcasting epidemics of novel pathogens: lessons from COVID-19

Joseph T. Wu, Kathy Leung, Tommy T. Y. Lam et al. • 2021

[Read the original paper](#) → · the course assigns: whole piece (the article body; the 160-item reference list not summarized)

This is the list of questions you have to answer in the first weeks of an outbreak, in the order they arrive, with an honest account of how thin the data is each time you answer one.

Average summary

Wu, Leung, Lam and colleagues call this a Perspective and say plainly that they are sharing their views.

Nowcasting means working out the current state of an outbreak: what the pathogen is, how it spreads, how bad it is, how many people have it. Forecasting predicts the future; nowcasting tells you where you are. The article is organised as the questions themselves.

What is the pathogen and where did it emerge? The first job is outbreak investigation producing a detailed line list, and the most fundamental step is the **case definition**, because everything downstream rests on it. During China's first wave the definition was broadened several times, and each time the proportion of confirmed infections rose three- to sevenfold, which is why case counts from different periods or places cannot be compared naively. They also warn not to confuse where cases were first reported with where the pathogen emerged: SARS-

CoV was first identified in Hong Kong but traced to Guangdong.

Does it spread between humans? Clusters are a clue, not proof, since a common source or several animal spillovers can look the same. For COVID-19 the conclusive evidence was a Shenzhen patient with no Wuhan travel history who was infected by family members returning from Wuhan, showing at least two generations of human-to-human spread. They contrast H5N1, which has never achieved sustained human spread, and MERS-CoV, which has produced repeated explosive hospital clusters of more than 100 infections and four or more generations without becoming a pandemic.

Is there treatment? Early trials of lopinavir-ritonavir and remdesivir alone showed no significant benefit in severe disease. The WHO **SOLIDARITY** trial found four repurposed drugs had little or no effect on mortality, ventilation or hospital stay. The UK's **RECOVERY** trial found dexamethasone lowered 28-day mortality in patients needing respiratory support, while hydroxychloroquine and azithromycin did not help. They call both trials defining exemplars that set a new benchmark.

How easily does it spread? R, the average number of secondary cases per case, is what you want, and early estimates are inevitably uncertain. Their most striking example: the first published serial interval estimate

for COVID-19 rested on six infector-infectee pairs from the Wuhan line list, giving an R credible interval of 1.4 to 3.9.

Can it be contained? After Wuhan locked down on 23 January 2020, Tencent mobility data showed dozens to hundreds of cases had already reached multiple Chinese cities and that national epidemics were inevitable within weeks, which prompted a national network of cordons with active case finding.

How severe, how many infected, do measures work, and are vaccines available? China's infection fatality risk was estimated at 0.66%, rising with age, and from that plus R came the projections of more than 510,000 UK and 2.2 million US deaths without mitigation. Only ten of the first 425 confirmed cases on the Wuhan line list carried the data needed to estimate the incubation period, whose 95th percentile was about 12 days, which is where the 14-day quarantine comes from.

Why it's assigned

The course note says this article runs through the questions we would want answers to early in a pandemic, and that is exactly the second element of the detect layer: not noticing that something is circulating, but characterising it once found. Read it as the checklist behind the day's framing, and notice how many of the answers in the real case came from

tiny samples, revised definitions and borrowed mobility data.

Words to know

- **Nowcasting.** Assessing the present state of an outbreak from incomplete, delayed data. It matters because almost every decision early in an outbreak depends on a quantity nobody can measure directly yet.
- **Line list.** The case-by-case table of demographic, exposure, laboratory and clinical details built during outbreak investigation. It matters because the estimates everyone argues about are squeezed out of it, and it is the first thing to fall apart when investigators are overwhelmed.
- **Case definition.** The agreed criteria for counting someone as a case. It matters because changing it changes the epidemic curve, as China's three- to sevenfold jumps show.
- **R and R_0 .** The average number of further cases one case causes; R_0 is that number in a fully susceptible population. Below 1 the outbreak dies out, above 1 it can grow exponentially.
- **Generation time and serial interval.** The gap between infections in an infector-infectee pair, and the gap between their symptom onsets. It matters because the serial interval stands in for the generation time, since onset is easier to observe

than infection, and uncertainty here was a main source of uncertainty in R.

- **Infection fatality risk (IFR).** The share of all infected people who die, as opposed to the share of confirmed cases. It matters because it is what turns an epidemic size into a projected death toll.
- **Seroepidemiology.** Measuring antibodies in a population to infer how many have been infected. It matters because they call it the most direct and reliable way to estimate cumulative incidence, and because assay throughput made it mostly infeasible in real time.
- **PHSMs.** Public health and social measures: travel restriction, distancing, contact tracing, closures, curfews, lockdowns. It matters because the article treats their effectiveness as something to be nowcast, not assumed.
- **Infodemic.** Harm from misinterpretation or misuse of fast-moving information. It matters because they argue nowcasts themselves can fuel it.

Numbers in this reading

- Three- to sevenfold: the increase in the proportion of confirmed infections each time China broadened the COVID-19 case definition during the first wave.
- Six infector-infectee pairs: the basis of the first published COVID-19 serial interval estimate, giving a 95% credible interval for R of 1.4 to 3.9.

- Ten of the first 425 confirmed cases on the Wuhan line list provided data sufficient to estimate the incubation period.
- About 12 days: the 95th percentile of the COVID-19 incubation period, behind the at-least-14-day quarantine most jurisdictions required.
- About 40%: the share of infections that symptomatic cases passed on before showing symptoms, which is why they argue contact tracing has to be very fast and efficient.
- 0.66%: the estimated infection fatality risk in China, increasing with age.
- More than 510,000 UK deaths and 2.2 million US deaths: the projected toll of an unmitigated epidemic.
- 50%: the upward revision of the Wuhan death toll weeks after the national lockdown lifted, because deaths were hard to count without widespread testing.
- Around 10 days: the delay between infection and reporting for COVID-19, which is why the effect of measures does not show up immediately in case-based R.
- More than 100 infections and four or more generations: the size of some MERS-CoV hospital clusters, which still did not become a pandemic.
- At least 1 year: how long serum neutralizing antibodies were found to persist after severe or

mild disease; at least 5 months for protection from natural infection.

- 23 January 2020: the Wuhan lockdown, after which Tencent mobility data showed dozens to hundreds of cases had already reached multiple major Chinese cities.
- US\$16 trillion: the estimated economic damage of COVID-19, which they use to argue for funding standing nowcasting systems.
- 2016: the year the National Academy of Medicine-led Global Health Risk Framework made the same recommendation, alongside dozens of others that would have left the world better prepared.

What the author is unsure about

This is a careful article about uncertainty, and it names its own. The origin of SARS-CoV-2 remains elusive, largely because early animal and environmental samples were never available for sequencing, while there is some evidence it may have circulated elsewhere unnoticed. The transition from purely animal spillover to efficient human spread can be prolonged, so a snapshot at any single moment can be inconclusive. The proportion and infectiousness of asymptomatic cases remains poorly understood. Early R estimates are inevitably fraught with substantial uncertainty, and genomic approaches are limited by the phylogenetic signal in the sequences. Real-time severity estimation is hard

because it needs reliable death counts and an accurate epidemic size at once. Duration of vaccine-induced immunity was unknown at writing, protection from disease may not mean protection from transmission, vaccines could not be compared head to head across trials run in different conditions, and there were no trial data yet for pregnant women, children or the immunocompromised. On behaviour they note most studies used online convenience samples and were cross-sectional, so it is unclear whether mental health outcomes followed the epidemic or preceded it. Their most pointed observation is political rather than technical: there is an almost inverse correlation between countries' pre-COVID-19 assessed preparedness and their actual performance.

For the discussion

- The fact worth bringing is how thin the foundations were: a serial interval from six pairs, an incubation period from ten of the first 425 cases, and a case definition that moved the confirmed-case count three- to sevenfold each time it changed.
- Worth raising as a tension: they want an integrative nowcasting system, something like a Bloomberg terminal for epidemic intelligence, and in the same breath say it would compete for data scientists and software engineers against industries that pay far more.

- A question worth asking: the article's own conclusion is that the disconnect between policymakers, practitioners and scientists is the most important lesson, so what in the detect layer actually closes that gap rather than producing better numbers for nobody.

3.3-1 • CORE • COURSE TIME 30 MIN

Toward a Global Pathogen Early Warning System

Natasha E. Bajema, William Beaver, Christine Parthemore ·
2021

[Read the original paper →](#) · the course assigns: the later pages, printed 19 to 95 (the rest of the technology section, Part III The Landscape Today, Part IV Next Steps, and the Annex country case studies)

Country by country, the report shows that early warning is not one system anywhere, it is dozens of partly connected ones with different rules, different data and different people, and the countries that did well in COVID-19 mostly did well on testing, tracing and political will rather than on detection technology.

Average summary

These are the pages to read for the country deep-dive. The authors take the general argument from the

front of the report and show it in real places.

The United States comes first, and the point is fragmentation. There is no uniform approach.

Massachusetts is their good example: lab-based notification for specific pathogens, syndromic surveillance across every emergency department in the state, sixty sentinel sites for influenza, an event-based platform for frontline workers to flag an outbreak, and a web platform called MAVEN that joins the layers, used by 92% of local health boards. Other states and localities are severely under-served. Many states run no statewide syndromic surveillance at all. Where it exists, it varies: 16 of California's 58 counties do it at some level; Colorado runs it county by county with one person serving as ESSENCE coordinator for every participating county; Florida gets uploads from 266 of 268 emergency departments every 2 hours and has two epidemiologists per county chasing flagged visits; Illinois has complete emergency department coverage but a 2016 evaluation found under 30% of users had access and nobody logged in daily, which made it of limited value for early warning. **New York** is in good shape on paper: 95% of counties have emergency department syndromic surveillance, and the CDC rates New York City's as one of the best in the country, useful for ruling Ebola in and out in emergency rooms between 2014 and 2016 and for catching a rise in drug overdoses.

Then the world, grouped by theme rather than ranked. South Korea scaled to six hundred screening centres, 150 diagnostic laboratories and 110,000 tests a day. Southeast Asia is the One Health case, where human, animal and agricultural surveillance have to join up. Central Asia and the Caucasus are building on Soviet-era laboratories, with Georgia running almost 2 million tests by the end of 2020 in a country of under 5 million people. Africa is receiving a large influx of investment onto a mostly paper-based base. The Middle East has little cross-border legal obligation to share data. India shows how fragile progress is inside a federal system.

The Annex then walks through five outbreaks: COVID-19 across many countries, SARS from 2002 to 2004 including the 2004 Beijing lab accidents, H5N1 avian influenza where Thailand succeeded and Cambodia did not, Ebola in West Africa and the DRC, and mosquito-borne disease.

Why it's assigned

This is the core reading for Option 1, the country-specific biosurveillance deep dive, and the course note points you at exactly these later pages. The guiding question is how a novel virus would actually be detected and characterised in the country you pick, and the only way to answer it is to see how concrete a real national system is: which agency, which data stream, how often it updates, who reads it. For the

detect layer, these pages are also the evidence for the day's main claim, that detection is less a technology problem than a coverage, data-sharing and political-will problem.

In depth

Before the case studies: the rest of the technology section (printed 19 to 24)

The authors finish their technology survey before turning to real places. **CRISPR diagnostics** are the closest competitor to metagenomic sequencing as a catch-all method, because they can be "massively multiplexed", testing for hundreds or thousands of pathogens on one sample by matching genetic material against digital databases; the US Department of Defense programme DIGET aims at detecting 1,000 pathogens simultaneously, with UC Berkeley and the Broad Institute leading the research. Their stated weakness matters for this unit: they are less effective at detecting novel pathogens, and would need work to catch diseases genetically dissimilar from known threats or weapons engineered to avoid detection.

On wearables, they report NASCAR, the NBA and the WNBA using OURA Rings, originally sold as sleep trackers, to look for tiny temperature changes that might precede infection, and say Fitbits and Apple Watches have similar potential. They are blunt that

separating signal from noise is difficult with a high risk of false positives, and suggest piloting them where exposure risk is high, including labs working on pandemic-potential viruses, where the nuisance of false positives may be worth catching an escape.

On environmental detection, the US built **Biowatch** in 2003 to test air for pathogens of interest, and its weaknesses, shared by today's wastewater testing, are difficulty covering a wide variety of pathogens and immature automated collection. On digital detection, ProMED was first to report SARS, MERS, Ebola and Zika, but not COVID-19: on December 30 the automated **HealthMap** system, which scans news and social media, put out an advisory about pneumonia-like illness in Wuhan, and ProMED's report, which better characterised the threat as high priority, went up four hours later.

They close on interoperability, endorsing a federated ecosystem and citing the Milken Institute's June 2021 recommendation of a hub-and-spoke model with a central coordinating entity to link existing local and regional activities, fill gaps and provide a governance framework. What such a system needs: fast data standards, an integration hub linked to regional sub-systems, serious training, financial sustainability, enough data to develop countermeasures, and room for technology to change. Who plays which role, they say, is unclear today.

The United States: diverse systems, not one system

They open with a caution about their own evidence: fully mapping the global landscape would take significant resources and constant updating, online information often differs from what happens on the ground, and even publications from the past five years are often out of date.

Figure 1 shows US counties reporting to the CDC's National Syndromic Surveillance Program from April 2020 to April 2021. The figure's point is how low the bar is: counties shaded dark blue only needed to send at least one valid patient visit record in the last twelve months, which is far from enough data for early warning against a highly transmissible virus. Several states use ESSENCE, the Electronic Surveillance System for the Early Notification of Community-based Epidemics, which gives some commonality but no uniformity; where syndromic surveillance starts at county level, as in Colorado and California, neighbouring counties may track completely different things, one watching influenza-like illness, the next watching overdose and suicide rates.

The most instructive state cases, from the authors' interviews (the rest of the figures are in the numbers section below). **Florida** is the best picture of a system with a human attached: 266 of 268 emergency departments and 80 urgent care centres upload every

2 hours, alongside poison control, vital statistics, the National Weather Service and the reportable disease and STD systems; an algorithm flags Visits of Interest from words in the chief complaint and discharge fields, and two epidemiologists in each county get a report of the past 24 hours at 10 am daily; during Zika a state study found an additional 18 cases from triage notes and 11 from foreign travel triage notes. **Illinois** is the cautionary case: statewide since 2013, complete coverage from 185 emergency departments reporting every 15 minutes with strong legal frameworks, yet a 2016 evaluation found under 30% had access and nobody accessed it daily, making it of limited value for early warning. **Colorado** shows how thin the staffing is, with one person serving as ESSENCE coordinator for every participating county, and officials telling the authors they consider the software underutilised and want hospitals better incentivised. **Massachusetts** is their best case, with the authors' own caveat that its ability to detect an outbreak early is questionable; MAVEN runs a statistical algorithm over multiple sources and epidemiologists classify signals by threat potential and stage of investigation, and the state uses electronic laboratory reports rather than paper and fax. **Minnesota** shows the limits of consent: no formal system, data collected only since summer 2020 on existing hospital-capacity infrastructure, no legal framework and strong public opposition, possibly because of the detail collected, including social

security numbers. **New York** has most counties integrating laboratory, syndromic and electronic-record surveillance, 95% of counties with emergency department syndromic systems, and New York City rated by CDC as one of the best in the country, having helped identify and exclude possible Ebola cases in emergency rooms between 2014 and 2016 and flag a rise in drug overdoses and culture-negative meningitis, with the state running its own Communicable Disease Electronic Surveillance System. **Texas** shows triage under load: it adopted standard case definitions and, knowing some jurisdictions could not investigate everything, defined five priority levels from all confirmed cases down to suspect cases; its syndromic system recorded 326,928 visits matching COVID-19-Like Illness in a year, correlating with case counts at 0.83 but representing only about 15% of cases, and its Austin laboratory found the state's first B.1.1.7 case in a man with no travel history, meaning it was already circulating. And **Washington** is the one that should unsettle anyone defending the status quo: its RHINO system and legally required lab reporting were both beaten to the first confirmed community transmission in Seattle by the **Seattle Flu Study**, an academic project founded in 2018 that mailed at-home tests to people reporting flu-like symptoms, and which found SARS-CoV-2 in 25 of 2,353 specimens, 1.1%, between January 1 and March 9, 2020. California, Connecticut, Georgia, Louisiana, Maryland and Oklahoma each get

the same treatment, and the pattern holds: different software, different data streams, different legal basis, different coverage.

Their account of how US biosurveillance changed during COVID-19 is equally concrete. Federal tracking could not pivot to SARS-CoV-2, so academics, journalists and volunteers filled the gap: the Covid Tracking Project, run by volunteers and overseen by The Atlantic, hand-curated data as states kept changing what they reported, and disbanded in March 2021. CDC's own systems were disparate and in some cases outdated. HHS built the HHS Protect data hub, powered by eight commercial companies, which could add new data or users within days, and because HHS could threaten to withhold Medicare and Medicaid reimbursement, 96% of hospitals were reporting daily by the end of 2020. Their analogy for the limits of all this: PCR prevalence testing is a forecast telling you how much precipitation to expect, sequencing tells you whether it will be rain, snow or hail, and neither tells you whether a tornado is coming, because neither says what else is circulating.

Around the world

Northeast Asia. South Korea's system took on many features of the US one through cooperation with US defense, health and diplomatic agencies, but its MERS experience left it better prepared; working with the private sector it scaled to six hundred screening

and testing centres, 150 diagnostic laboratories and 110,000 tests per day as of November 2020. Vietnam succeeded with less testing, more tightly targeted around clusters. Japan's investment is shaped by its history, including a domestic terrorist group that tried to build chemical and biological weapons and the 2011 Triple Disaster, and its National Institute of Infectious Diseases receives syndromic, sentinel, indicator-based and event-based reporting. Their honest caveat: today's diagnostics work well once an outbreak is characterised, and it is much less clear how these systems would have performed if SARS-CoV-2 had started in either country.

Southeast Asia. Overlapping programmes run by governments, donors, international organisations and NGOs, with no network integrating them. Vietnam has a wildlife trade valued at \$1 billion annually and one of the region's highest livestock densities, and built sentinel surveillance of at-risk groups such as farming households and wet market and wildlife restaurant workers; the three-year VIZIONS study, supported by the Wellcome Sanger Institute, sequenced samples from high-risk people to see which viruses had taken the first step into a human host. Thailand's PODD programme collects mobile-phone reports of sick livestock or wild animals, now extended to human and environmental health.

Central Asia and the Caucasus. Soviet-era scientific infrastructure, repurposed after the collapse of the USSR and its industrial-scale bioweapons programme. Before US-supported work began in the early 2000s, reporting was paper filtering up from facility to province to nation. The US programme now called the Biological Threat Reduction Program funded the building of EIDSS from 2003, which includes clinical symptoms and, importantly, integrates human and veterinary reporting. Georgia's National Center for Disease Control implemented PCR testing in late January 2020, added a WHO-approved antigen test for rural areas without laboratories, ran almost 2 million tests by the end of 2020 in a country of under 5 million, and set up physical and virtual clinics for almost a quarter of a million positive cases.

Africa. In Kenya, metagenomic sequencing is being used to find which pathogens most often sicken and kill children, with the Chan-Zuckerberg Biohub and Gates Foundation extending it to Gambia, South Africa, Malawi and Madagascar. In a recent survey the single most cited challenge was data collection and dissemination; systems are largely paper-based, and Sierra Leone was the first country to move to electronic collection and sharing of outbreak data. Some public health experts warn about absorption capacity and long-term sustainability where records are not electronic or electricity is inconsistent. Africa CDC and the Institute of Pathogen Genomics are the

entities many hope will let parts of the continent leapfrog.

The Middle East. Home to MERS, first detected in Saudi Arabia in 2012, and to former weapons programmes in Iraq. Most nations lack legal requirements to share infectious disease data. Saudi Arabia is the exception, with years of event-based surveillance and syndromic surveillance tailored to its concerns, including the Hajj. MECIDS, launched in 2003 among Israel, the Palestinian Territories and Jordan, started with food-borne illness because that is easier to share than human disease trends that could reveal national weakness. The authors suggest the region could leapfrog to methods that do not use individual-level data, such as the AI-enhanced wastewater analysis in Abu Dhabi now extending into airports.

Latin America. In May 2021 six countries announced an alliance to digitise regional health systems, starting with a digital vaccination card usable in at least three of them. PAHO's COVID-19 Genomic Surveillance Regional Network, created in 2020, had identified 38,517 virus sequences in member states as of June 2021, with in-country sequencing in ten states and two regional laboratories, FIOCRUZ in Brazil and ISPCH in Chile.

India. In 2018 India detected and responded to a Nipah virus outbreak, praised by the global health

community. But administration of its universal healthcare is left to states, most care is privately provided, states and regions differ hugely in wealth and technology, the public system is overloaded, and there is no robust integrated tracking of animal and wildlife disease; data sharing is in some cases weekly rather than real time. India was near the front of the first wave through students returning from Wuhan, detection of those initial cases was lacking, and the country fell back on blunter measures. A government contact tracing app passed 100 million installs in six weeks and eventually roughly double that, about a third of smartphone users but under 14% of the population; by May manual contact tracing was falling behind and testing was finding a small fraction of cases. After the first wave subsided control was handed back to states, and the second wave from April 2021 was far larger, including the B.1.617 variant the WHO designated a variant of concern. They cite an estimate that COVID-19 in India alone pushed an additional 85 million people back into extreme poverty, wiping out five years of global progress.

They also describe the long tail of **tailored systems** built for one disease or population, some dating to the early-to-mid 1900s, including trade-driven agricultural systems that are hard to change because of commercial interests, and GISAID, which began with influenza and expanded into coronaviruses.

The USS Theodore Roosevelt is their case for unique environments. After the carrier's March 2020 Da Nang visit, which went ahead because the US CDC rated Vietnam low risk and which included a four-hundred-person hotel reception and thirty visiting reporters, 1,271 crew members tested positive over two months, twenty percent of the crew, and the ship was out of commission for much of that time. As many as 50 percent of cases were asymptomatic, and the Navy could never identify the chain of infection that brought the virus aboard, even afterwards, despite quarantining off the ship every sailor with confirmed exposure during the port call. Their lesson is the need for rapid, cheap, versatile presymptomatic detection that could have tested the whole crew at once.

On **coordination**, they list the WHO's new data hub announced May 5, 2021, built on the belief that the data to spot the next outbreak already exists and needs to come out of its silos; GOARN for capacity building; GISRS for influenza mutations; the Global Health Security Agenda, launched February 2014 with seventy countries, about a month before the West Africa Ebola outbreak was discovered; and CORDS, which connects regional detection networks and which they rate as a strong partner for piloting new approaches. They also flag an unresolved problem with commercial systems such as Metabiota: their data and tools are often closed and sometimes classified, and nobody knows how such companies

would stay profitable inside a global early warning system.

Their closing argument in Part III is about threat coverage. Each COVID-19 case produced another two to three infections, which they call relatively low; for malaria it is another seventeen. SARS-CoV-2 was hard to stop not because it spread fast but because it spread pre-symptomatically and asymptotically. The extreme case is HIV, where symptoms can take up to ten years and where 30 to 60 years may have passed between initial human-to-human transmission and identification. A system that waits for symptoms is also blind to attack: in May 1965 US special operations personnel walked into Washington International Airport with suitcases modified to disperse a simulant for smallpox, and a real version of that attack would only have become apparent two weeks later, as victims began getting sick.

Part IV: next steps

Seven recommendations: conduct further targeted landscape assessments, including in-country conversations that also build trust; set interoperability parameters as soon as possible, so the billions now going into sequencing become the leading edge of early warning; invest in diverse tools for diverse settings; map how tools apply in specific settings, with a template for ideal technologies, information flows and response timelines compared against what

exists; hook into enduring missions such as trade or security interests; launch confidence-building measures, including efforts to reduce fear of punishment for being transparent; and expand cooperative biological engagement programmes. Most of all, they say, this work must begin now.

The Annex: five outbreaks

The annex adds detail on past outbreaks across different operating settings. The **COVID-19 vignette's** own themes: China was likely unable to identify the virus before it spread internationally, and significant uncertainty about early transmission chains remains; countries thought to have top systems, including the US, could not track infections once they crossed borders; pre-symptomatic and asymptomatic spread is especially problematic; current systems do not sufficiently emphasise novel threats, where the window to act near the origin may be small; and subnational inequality matters as much as national capability.

Country by country, the cases that carry the argument. **Japan's** first patient was examined on January 6 with a fever already three days old, hospitalised on the 10th, tested on the 14th as tests were still being developed and confirmed on the 15th; its expert group identified poorly ventilated indoor spaces and close contact as the key routes well ahead of many western countries, and the country chose rapid cluster

identification and control. **New Zealand** diagnosed its first case on February 21, peaked at 424 new cases on March 23, then ran one of the world's most aggressive elimination strategies, using its National Close Contact Service, an app for people to track their own movements, sequencing for community transmission and new strains, and traffic and cell phone data to check whether distancing was working; by late April it was finding potential cases about 2 days before symptoms appeared. **Israel's** advanced records and hospital sequencing were blunted because COVID-19 looks like other respiratory disease; it pooled tests to stretch capacity and used cellphone location surveillance the authors call largely ineffective. **Russia** confirmed its first case through Rospotrebnadzor, whose coronavirus surveillance centre sits at the Vector laboratory almost 2,000km east of Moscow, a former Soviet bioweapons site; Putin's claim that spread was under control preceded the shutdown ordered on March 25. **Kazakhstan** closed borders with China early and imported no initial cases from its longest border, which the authors use to make their point about secondary spread, and sequenced only 53 of more than 1,000 early confirmed patients for want of funding and reagents. The **UK**, whose systems were seen as among the best in the world, squandered existing laboratory capacity by building its own facilities instead, but its early decision to prioritise sequencing let it find the more transmissible variant in December

2020 and led it to propose a "pandemic radar" with genomic sequencing as a key pillar. **Germany** cut its March 2020 peak of 6,000 daily cases to under 1,000 by May, then lost control in October when state governments each set more lenient policies. **Brazil** never imposed a national lockdown and had more than 17 million confirmed cases and more than 470,000 deaths in a country of 210 million, while **Argentina**, which had a pre-existing structure with national software giving real-time tracking and shut down early, still converged towards Brazil's detected-case rate, though with lower fatality. China, India and Iraq are covered too, and the figures for all of them are in the numbers section below.

The **SARS vignette's** lesson is that five months passed in China between the initial cases and identification of the pathogen, with no functioning biosurveillance system at the time; SARS was first diagnosed only as pneumonia of unknown cause and was still being reported as bacterial pneumonia in February 2003 despite early reports that antibiotics had no effect; it reached twenty-nine countries and ended with about 8,000 cases and 774 deaths. Canada's first case arrived in Toronto in February 2003; her son spent at least twenty hours in an emergency department without containment, and it took a full two weeks for hospital doctors to connect the illness to what was spreading in China. Canada had 438 cases and 44 deaths, 40% of cases among

healthcare workers, and a second surge in May and June from mixed messaging between levels of government. Its GPHIN event-based system did collect the initial report from China but left it unanalysed because it was in Chinese. Then the 2004 lab accidents at Beijing's National Institute of Virology: two staff infected in separate exposures three weeks apart, more than 500 people quarantined, eight infected, and the mother of the first victim died; investigators traced it to transfer from BSL-3 space to lower containment, where the detergent protocol had not inactivated the virus and nobody had verified that it had.

The **H5N1 vignette** turns on compensation, and it is the most transferable lesson in the annex: whether farmers report sick birds depends on whether reporting ruins them. Thailand culled hard, ran a media campaign and crucially reimbursed farmers at market value; 90% of its production is industrial and easier to regulate, it was protecting a \$1 billion a year poultry industry, it required any patient with flu-like symptoms to be asked about poultry exposure, sampled poultry from every village, chose not to vaccinate for fear of generating variants, and has had no outbreaks since 2008. Vietnam, poorer and more decentralised with 70 to 80% of poultry in small household flocks, compensated unevenly, so reporting and culling were incomplete and H5N1 became endemic in poultry. Cambodia paid no compensation

at all, so farmers hid outbreaks or sold birds on; a 2015 effort found 50 to 80% of live market environmental samples positive and 4.5% of poultry workers carrying antibodies without symptoms, meaning human cases were spreading undetected. Egypt authorised twenty-four different vaccines with no guidance on which to use, and as late as 2019 30% of subsistence farmers were still disposing of carcasses in local streams.

The **Ebola vignette** shows delay in low-resource and conflict settings even for a disease that is easy to notice. Guinea, Liberia and Sierra Leone had among the fewest physicians per capita in the world, worsened by staff reductions between 2011 and 2013 at the WHO regional office for Africa, and few Ebola-specific diagnostics existed. In Nigeria the index case lied about his medical history and the disease spread for 3 days before staff identified it, yet a public information campaign held it to nineteen further infections and seven deaths. In the DRC the outbreak began 600 kilometers east of previous ones inside an active conflict; attacks on health workers and misinformation blocked the response, contact tracing was impossible in conflict zones and was suspended at one point along with vaccination, fewer than a third of the population believed local authorities represented their interests and more than a quarter believed the outbreak was fictitious.

The **mosquito-borne vignette** makes the subtlest point for early warning: routine surveillance is mostly case counting after the fact. In the US, arboviral disease is reported to CDC through ArboNET, built in 2000 after West Nile arrived in 1999, and the authors state plainly that human infection is not a sensitive indicator of viral activity and so is not useful for early warning, because infections can be mild or asymptomatic and symptoms can take two weeks. Dead bird surveillance has been one of the earliest and most cost-effective indicators of West Nile activity, yet there is no national dead bird reporting system. In low-resource settings, early warning for these diseases rests on rainfall, humidity and temperature from meteorological departments and online climate libraries, beginning with FEWS NET maps in 2002 and Columbia University's Malaria Early Warning System. On genomics they note that in the Zika outbreak, even in a high-resource setting, there was no sign that sequencing informed the response: it was done either too early, before much was known, or retrospectively, months after the outbreak ended. Pollett and colleagues' five challenges close it out: research and public health objectives are not always aligned, funding for reagents, salaries, collection, equipment and training is hard to secure, results need to be produced in a form useful for decisions, linking routine clinical and epidemiological data to genomic data takes expertise, and the training pipeline needs sustaining.

Words to know

- **Syndromic surveillance.** Watching patterns of symptoms rather than confirmed diagnoses, usually from emergency departments. It matters because it is the fastest layer most US states have, and the report shows how unevenly it is built and used.
- **ESSENCE.** The Electronic Surveillance System for the Early Notification of Community-based Epidemics, the software many US states and counties run their syndromic surveillance on. It matters because shared software is not a shared system: neighbouring counties configure it for different purposes.
- **National Syndromic Surveillance Program (NSSP).** The CDC's collection and analysis hub that state and county systems feed. It matters because Figure 1's coverage map is a map of this programme, and its reporting bar is very low.
- **Visit of Interest.** Florida's term for an emergency department visit an algorithm flags from words in the chief complaint or discharge diagnosis. It matters because it is the clearest example in the report of a signal with a named human being whose job is to follow it up.
- **One Health.** Treating human, animal and environmental health as one system. It matters because most emerging infections in the regions described here come from animals.

- **Indicator-based and event-based surveillance.** Structured official reporting versus signals pulled from messy sources such as media and professional networks. It matters because Canada had the second kind and still missed a report that was sitting in its own system in Chinese.
- **Ring method.** Vaccinating the contacts, and the contacts of contacts, around a confirmed case. It matters because it is how Ebola vaccines were used once they existed.
- **Reassortment.** When two flu viruses infecting the same host swap genetic material and produce a new combination. It matters because it is the specific mechanism that could turn H5N1 into a human-transmissible pandemic strain.

Numbers in this reading

- 92% of Massachusetts local health boards use MAVEN; its syndromic system covers 100% of emergency departments; the state has sixty sentinel sites for influenza.
- 16 of California's 58 counties conduct syndromic surveillance at some level.
- 83 Colorado facilities report, roughly 90% coverage in participating counties, with one ESSENCE coordinator for all of them.
- 266 of 268 Florida emergency departments plus 80 urgent care centres upload every 2 hours; two epidemiologists per county review flagged visits

daily at 10 am; a Zika study found 18 extra cases from triage notes and 11 from foreign travel triage notes.

- Georgia: 135 facilities, 18 districts, 159 counties, 5 facilities outside the system, a median reporting timespan of 32 hours, 36 syndromes, 85 to 90% of data shared nationally.
- Illinois: 185 emergency departments, data every 15 minutes, 12 syndromes and 165 subsyndromes, but under 30% user access in a 2016 evaluation and nobody logging in daily.
- Maryland: complete coverage of 24 counties, 100% of emergency departments and 100% of schools, data in under 24 hours.
- Minnesota: 173 facilities, no legal framework, strong public opposition.
- New York: 95% of counties with emergency department syndromic surveillance; New York City rated by CDC among the best, used to exclude possible Ebola cases in emergency rooms between 2014 and 2016.
- Texas: 326,928 COVID-19-Like Illness visits in one year, correlation of 0.83 with case counts, about 15% of cases; five investigation priority levels; first B.1.1.7 case found January 7, 2021; almost 3 million cases and more than 51,000 deaths as of May 24, 2021.
- Seattle Flu Study: 2,353 specimens tested between January 1 and March 9, 2020, SARS-CoV-2 found in

25, or 1.1%.

- 96% of US hospitals reporting daily by the end of 2020, once Medicare and Medicaid reimbursement was at stake.
- South Korea: six hundred screening and testing centres, 150 diagnostic laboratories, 110,000 tests per day as of November 2020.
- Georgia (the country): almost 2 million tests by the end of 2020 in a population under 5 million; 25 viral strains whole-genome sequenced; almost a quarter of a million positive cases cared for.
- Kazakhstan: 53 of more than 1,000 early patients sequenced, 87% of those lineages from Asia; 430,000 cases in 19 million people; \$42 per test; 8,000 people on a quarantine app; 14,000 surveillance cameras in the capital.
- PAHO: 38,517 virus sequences identified in member states as of June 2021.
- India: over 100 million app installs in six weeks, eventually about double, a third of smartphone users and under 14% of the population; 85 million people estimated pushed back into extreme poverty.
- USS Theodore Roosevelt: 1,271 positive crew members, twenty percent of the crew, up to 50 percent asymptomatic, a four-hundred-person reception and thirty visiting reporters during the port call.

- Transmissibility comparison: two to three further infections per COVID-19 case, seventeen per malaria case; HIV can take up to ten years to show symptoms and may have gone 30-60 years before identification.
- SARS: five months from initial cases to identification in China, spread to twenty-nine countries, about 8,000 cases and 774 deaths; Canada 438 cases, 44 deaths, 40% healthcare workers; the 2004 Beijing lab accidents infected eight and quarantined more than 500.
- H5N1: around 60% fatality in humans; 1918-19 pandemic killed 50-100 million; Thailand 90% industrial production and a \$1 billion a year industry, no outbreaks since 2008; Vietnam 70 to 80% household flocks; Cambodia 50 to 80% of market environmental samples positive and 4.5% of poultry workers with antibodies; Egypt 359 cases, 120 deaths, 24 authorised vaccines, 30% of subsistence farmers still dumping carcasses in streams in 2019, antibodies in about 2% of the population.
- Ebola: more than 11,000 deaths of more than 28,600 cases in West Africa 2013-2016; 3,470 cases and 2,280 deaths in the DRC 2018-2020; 50% average case fatality, up to 90%; the DRC outbreak began 600 kilometers east of previous ones; fewer than a third of the population trusted

local authorities and more than a quarter thought the outbreak was fake.

- Mosquito-borne: more than 3,000 species on 6 continents, three genera responsible for most human disease, more than 400,000 malaria deaths annually as of 2019.

What the author is unsure about

They are careful about their own evidence base. They say mapping the whole landscape would take significant resources and constant updating, that online information often differs from the ground truth, and that even recent publications are often outdated, so Part III is deliberately a high-level summary with selected examples. They call their state case studies not comprehensive, only representative. They admit Massachusetts' ability to detect an outbreak early is questionable even though it is their best example, and they speculate openly about why MAVEN coverage is 92% rather than closer to 100%. They say it is unclear how well Japan's or South Korea's systems would have performed had SARS-CoV-2 started there. They flag as unresolved how private companies with closed or classified data could be integrated into a global system and stay profitable. They note significant uncertainty remains about early COVID-19 transmission chains and disagreement about the origin of SARS-CoV-2 that is unlikely ever to be settled. They say it is unclear how much Ebola-

specific diagnostics would have helped in Canada's SARS case given that its laboratories were overwhelmed anyway, and that the sensitivity of syndromic data for malaria case surveillance has not been sufficiently evaluated.

For the discussion

- The most useful thing to bring from these pages is how specific a real answer has to be: in Florida a named algorithm flags a Visit of Interest, a report lands at 10 am, and two epidemiologists per county follow it up, which is the level of detail the question about your own country is asking for.
- Worth raising as a tension: Illinois had complete emergency department coverage, 15-minute data and strong legal frameworks, and a 2016 evaluation found nobody looked at it daily, so coverage is not capability.

3.3-2 · CORE · COURSE TIME 30 MIN

Emerging COVID-19 success story: Vietnam's commitment to containment

Todd Pollack, Guy Thwaites, Maia Rabaa et al., with
Exemplars in Global Health · 2021

[Read the original paper →](#) · the course assigns: whole piece

Vietnam held COVID-19 to 1,465 confirmed cases and 35 deaths through the end of 2020

by acting on exposure rather than symptoms: trace contacts out to several degrees, quarantine them in government centres whether or not they are sick, and lock down the smallest area that will do.

Average summary

This is a guest post on Our World in Data written by in-country experts, several of them from Vietnam's Ministry of Health and National Institute of Hygiene and Epidemiology and from the Oxford University Clinical Research Unit, as part of the Exemplars in Global Health series. Our World in Data has since added a notice that it was published earlier in the pandemic using the data available then.

The headline numbers: as of December 31, 2020, Vietnam had reported **1,465 laboratory confirmed cases and 35 deaths** in a country of nearly 100 million. The authors credit a well-developed public health system, a decisive central government, and a containment strategy built on comprehensive testing, tracing and quarantining.

The detect half is deliberately targeted rather than broad. Testing scaled up where community transmission appeared, from two testing sites nationwide in late January 2020 to 120 by May, and 137 laboratories with a maximum daily capacity of 51,000 tests by October 22. Publicly funded

institutions developed at least four domestic test kits, the earliest an RT-LAMP kit on February 7 costing US\$15 and taking 70 minutes.

The contain half is where Vietnam differs from most countries. Contact tracing runs from **F0**, the infected person, through F1 (close contact), F2 and on to F5. The authors explain the arithmetic behind the urgency: the incubation period averages five days and infectiousness begins two days before symptoms, which leaves only three days from exposure to find and quarantine a contact. F1s are tested; if positive they are hospitalised, free of charge regardless of symptoms; if negative they go to a government quarantine centre for 14 days. F2s self-isolate at home for 14 days. Crucially, people were quarantined on **epidemiological risk, not symptoms**, which matters because the authors report 43 percent of cases never develop symptoms.

The scale of that is the number to sit with: by December 31, 2020, a total of **10,242,896 people had been placed under quarantine**, 211,983 in health care facilities, 4,296,302 in centralised facilities and 5,734,611 at home.

Lockdowns were early, small and specific: a commune of 10,000 in February when the country had 16 cases, a single street of 190 people, a village of 5,000, Hanoi's Bach Mai Hospital after 45 linked positives. The only real failure was Da Nang in July and August

2020, where hospital transmission spilled into the community after 99 days with no community spread; Da Nang General Hospital alone had at least 251 cases among patients, caregivers and staff.

Communication was constant and unusual: a prevention message added to every phone call in the country, a handwashing pop song, and by December 21 YouGov found 91 percent of Vietnamese people thought the government was responding very or somewhat well.

Why it's assigned

This is the third element of the detect layer: finding individual cases and acting fast enough to contain them. It is also the course's deliberate counterweight. The course page says this approach required draconian enforcement and was much less effective against the more infectious omicron variant, and that conventional wisdom holds containment is unlikely to work at all in a "wildfire" scenario. So read it as proof that testing, tracing and quarantine can work against a moderately transmissible pathogen when a government moves in days, and as a case study in what that costs.

In depth

What Vietnam brought to 2020

Vietnam is a country of nearly 100 million people that went from a centrally planned economy to a socialist-oriented market economy after the Doi Moi reforms of the mid-1980s. Public health expenditure per capita grew an average 9 percent per year between 2000 and 2016, and infant and maternal mortality both more than halved between 1990 and 2018.

It also had practice. Vietnam was the second country after China to face SARS, and after 63 cases and five deaths it was the first country the WHO declared SARS-free. Many interventions it pioneered then were reused for COVID-19, and the authors suggest that experience also left people more willing to comply with a central public health response.

After SARS it built the machinery: a national public health emergency operations centre and a national surveillance system, with the national centre and four regional centres running exercises and training for government stakeholders, which had already handled measles, Ebola, MERS and Zika. The data layer moved to a nearly real-time web-based system in 2009, and since 2016 hospitals must report notifiable diseases to a central database within 24 hours. In 2018, with the US CDC, Vietnam added **event-based surveillance**, which empowers members of the public, including teachers, pharmacists, religious leaders and traditional medicine healers, to report

public health events, with the aim of spotting clusters of people with similar symptoms.

The timeline

The first case was reported on January 23, 2020. A week later a national steering committee formed, initially meeting every two days, to run a whole-of-government strategy. In Vinh Phuc province, about an hour from Hanoi, provincial leaders locked down Son Loi commune, isolated patients and close contacts in quarantine camps for at least 14 days, and started community-wide screening at the first sign of community spread.

A second wave, imported from Europe and the United States, surfaced on March 6. Within a day of the first case being detected, the government had tracked and isolated about 200 close contacts.

By May 1, a hundred days in, Vietnam had confirmed just 270 cases despite extensive testing, with no community transmission since April 15. Restrictions loosened in May, including domestic travel.

On July 25 Vietnam marked 99 days without community transmission, and then Da Nang happened. Da Nang is a city of more than 1 million that receives about 8 million tourists a year, and it was unusually crowded in July because the government was promoting domestic tourism to offset lost international tourism. What began as

transmission inside hospitals spilled into the community, and in the last week of July new cases grew by about 30 percent, the fastest rate of the epidemic so far. Between July 25 and September 8, 551 locally transmitted cases were reported across 15 cities and provinces, with Da Nang and nearby Quang Nam worst affected, and approximately 98 percent of cases were either linked to major Da Nang hospitals or had visited the city. The exposure figure is the arresting one: from July 1 to July 27, an estimated more than 1.5 million people returned from Da Nang to other provinces, of whom 41,000 had visited Da Nang Hospital. A temporary hospital was built for suspected and mild cases and two others were designated for COVID-19 treatment under a special Ministry of Health committee. By September 10, 61,968 people were being monitored, 998 quarantined in health care facilities, 15,619 in centralised facilities and 45,351 at home.

Detect: tests made at home, used in clusters

In late January 2020 the Ministry of Science and Technology met virologists to push diagnostic development. From early February, publicly funded institutions produced at least four domestic test kits, validated by the Ministry of Defense and the National Institute of Hygiene and Epidemiology, with private firms Viet A and Thai Duong manufacturing. Most confirmation laboratories used in-house versions of

the WHO protocol, which the authors say let tests be administered widely without long waits.

The four kits, with the detail as given: February 7, 2020, Hanoi University of Science and Technology, RT-LAMP, US\$15, 70 minutes. March 3, Vietnam Academy of Science and Technology, real-time RT-PCR, less than US\$21, 80 minutes from receiving a sample. March 5, Military Medical University, commercialised by Viet A, US\$19 to \$25, RT-PCR and real-time RT-PCR, over one hour but with four times the sample capacity of the CDC kit; this one carried up to 80% of Vietnam's testing during the global stock-out, was certified by the European Union and others, and drew orders from 20 countries and territories. April 28, Thai Duong's RT-LAMP and RT-PCR kits.

Capacity went from two testing sites nationwide in late January to 120 by May. As of October 22, 137 laboratories could run RT-PCR with a maximum daily capacity of 51,000 tests, 62 of them designated screening laboratories and 75 confirmatory, with plans to extend testing into the lung hospital system using GeneXpert machines.

The strategy was not mass screening. Given low case numbers, testing was used to find clusters and stop wider transmission, and a single community case triggered contact tracing, commune-level lockdown and widespread local testing. Da Nang was the

exception: more than 700,000 tests between July 23 and September 24, with people who had visited high-risk areas required to file a health declaration through a mobile app and notify local health authorities for sampling. From September 3 to September 10, a third of all households in Da Nang, 72,492 of them, were tested along with 13,776 students and staff sitting graduation exams. All were negative.

Contain: tracing to five degrees

Contact tracing and quarantine are, in the authors' words, the key parts of containment, and Vietnam's tracing is uniquely comprehensive, running from F0 through F1, F2 and up to F5.

The time budget explains the intensity. The incubation period between exposure and symptom onset averages five days, and infectiousness begins two days before symptoms, so there are only three days from contact to find and quarantine someone before they can infect others.

The process: once a case (F0) is identified, local public health officials, supported by health professionals, security officers, the military and other civil servants, work with the patient to reconstruct who they may have infected over the past 14 days. Close contacts (F1), defined as people who came within approximately 6 feet (2 meters) or had prolonged contact of 30 or more minutes with a confirmed case,

are identified and tested. A positive F1 is isolated in hospital, and all COVID-19 patients in Vietnam are hospitalised at no cost regardless of symptoms. A negative F1 goes to a government-run quarantine centre for 14 days. F2s self-isolate at home for 14 days.

The authors single out the logic: Vietnam quarantined suspected cases on epidemiological risk, contact with a confirmed case or travel to an affected country, not on symptoms. With 43 percent of cases never developing symptoms, they suggest this was a key contributor to limiting community transmission early. They draw the contrast with SARS explicitly: isolating symptomatic people worked for SARS because it was infectious only after symptoms began, and that same strategy would have been inadequate for SARS-CoV-2.

From January 23 to May 1, 2020, over 200,000 people spent time in a quarantine facility. After Da Nang the numbers jumped, and resources were moved to localities with weaker health systems to trace the more than 1.5 million people linked to the outbreak. Ho Chi Minh City traced everyone returning from Da Nang and stratified them: people with respiratory symptoms or exposure to the three epicentre hospitals went to centralised quarantine and were tested, the rest were isolated at home and monitored by commune health staff. The cumulative figure as of December 31, 2020 was 10,242,896

people placed under quarantine: 211,983 in health care facilities, 4,296,302 in centralised facilities and 5,734,611 at home.

Two apps supported this. **NCovi**, launched March 10 with telecom companies, creates what the authors call a neighbourhood watch system: a map of cases and clusters, self-declared health status, reporting of suspected cases, and real-time movement of people under quarantine. They note it has drawn criticism from some privacy advocates. **Bluezone**, launched in mid-April by the cybersecurity firm Bkav, is Bluetooth-based and tells users if they came within approximately 6 feet of a confirmed case in the past 14 days; the Ministry of Information and Communications aimed for 60 percent of the population and it passed 20 million downloads by August 20.

Keeping it out of hospitals

The authors treat health care transmission as a containment strategy in its own right, and the sentence that carries it is about SARS: during the 2003 to 2004 outbreak dozens of Vietnamese health care workers were infected, and apart from the index patient, everyone in Vietnam who died of SARS was a doctor or a nurse. Over the following ten years Vietnam invested in hospital infection control through organisational systems, physical facilities, equipment and supplies, and training.

On February 19, 2020 the Ministry of Health issued national Guidelines for Infection Prevention and Control for COVID-19 in Healthcare Establishments, covering screening, admission and isolation of confirmed and suspected cases, isolation areas, PPE, cleaning and disinfection, waste management, sample handling and transport, prevention of laboratory-acquired infection, handling of remains, and guidance for family members and visitors.

It still failed in one place. Da Nang General Hospital was the epicentre of the Da Nang outbreak with at least 251 cases among patients, caregivers and staff, and the hospital was locked down in August.

The lockdowns, in their own table

Table 1 in the article lists the targeted lockdowns, and its argument is that they were early, small and tied to specific named cases rather than national and late. Son Loi Commune in Vinh Phuc, February 13 to March 4, 10,000 people, when the whole country had 16 cases and 6 of them were in Son Loi. Truc Bac Street in Hanoi, March 6 to 20, 190 people and 66 households, where the first case of the second wave lived. Two streets in Phan Thiet, March 13 to April 3, 150 people, 29 households. Van Lam 3 Village in Ninh Thuan, March 17 to April 14, 5,000 people, after two infections, with all 16 entrances to the village closed and monitored. Thua Loi Village in Ben Tre, March 23 to April 20, 1,600 people and 480 households, after a

17-year-old tested positive. Bach Mai Hospital in Hanoi, March 28 to April 11, 4,000 to 5,000 people, after 45 linked positives, with over 15,000 people associated with the hospital tested and 40,000 who had been in contact with it tracked down. Ha Loi Village in Hanoi province, April 7 to May 6, 10,000 people. And Dong Van District in Ha Giang, April 22 to 23, 7,600 people, locked down before test results came back and released the next day when they were negative, which the authors offer as an example of how fast authorities moved.

Nationally, Vietnam entered lockdown on April 1 for an initial 15 days, extended to 21 days in 28 of 63 provinces. During the Da Nang outbreak in August, 10 provinces were locked down.

Borders and movement

Inbound passengers from Wuhan were screened before Vietnam's first case was even reported. Visas for Chinese tourists stopped on January 30, a week after the first confirmed case. On January 31, at the end of the Lunar New Year holiday and with only five confirmed cases in the country, all schools nationwide were ordered to stay closed. Masks were recommended in early February, before the WHO or most countries, and made mandatory nationwide in mid-March. Flights to and from China were suspended February 1 and trains on February 5, when cases were still in single digits. Flights from the Schengen Area

and the UK were suspended March 15, and all visa issuance stopped March 18. Government quarantine centres took Vietnamese arrivals from China from February 4, from South Korea from March 1, and all international arrivals from March 20 to 22, with international flights diverted away from airports still used for domestic travel.

Messaging, and its hard edge

The Ministry of Health first warned citizens on January 9, before the first case. Since then the government communicated constantly: a short prevention statement added to every phone call placed in the country, direct texts, and heavy use of social media, with 64 million active Facebook users in Vietnam and 80 percent of smartphone users having the local app Zalo installed. In late February the National Institute of Occupational Safety and Health turned a well-known pop song into "Ghen Co Vy", a handwashing public service announcement with new lyrics by Khac Hung and choreography by Quang Dang, which became a TikTok dance challenge; the #Onhavanvui campaign has had more than 8 billion views. The messages stayed consistent, first "fighting the epidemic is like fighting against the enemy", later the "5K" formula of mask, hygiene, safe distance, gathering and health declaration, and by late 2020 a WHO-supported Safe Coexistence with COVID-19 initiative.

The authors do not hide the enforcement side. A 2018 cybersecurity law already gave the government substantial authority, and on April 14 Vietnam passed a decree setting fines for people who use social media to share false, untruthful, distorted or slanderous information; police have pursued more than 200 cases, and the ordinance has drawn opposition from Amnesty International and others. Alongside that, YouGov data as of December 21 found 91 percent of Vietnamese people believed the government was responding very or somewhat well.

What it cost

The Prime Minister put the revenue cost at nearly 200,000 billion dong, US\$9.4 billion. Tourism was hit hardest, in a country that normally welcomes 21 million international arrivals, and the upscale shops and hotels of Hanoi and Ho Chi Minh City stayed empty. Fruit exports fell by half between April and June 2020, with freight operators demanding two to four times normal prices. But Vietnam was among the few countries with positive GDP growth, from a high pre-pandemic baseline, and manufacturing was growing again by June. Local producers exported PPE: France imported hundreds of millions of cloth masks in May and the US government bought 4.5 million made-in-Vietnam PPE suits, before global demand dropped in August and Vietnamese mask exports fell by a third. International flights later reopened for

foreign workers such as LG engineers from South Korea, and some international investors pushed to shorten compulsory isolation for them.

Their lessons

Invest in public health infrastructure such as emergency operations centres and surveillance systems, so you start a crisis ahead. Act early, from border closures and masks to testing and lockdowns. Trace contacts thoroughly, which is what makes targeted containment possible. Quarantine on possible exposure rather than symptoms alone, especially the testing and quarantining of international travellers. Communicate clearly, consistently and seriously throughout. And use a whole-of-society approach that brings many sectors into decisions. Their closing claim: with the exception of Da Nang, Vietnam prevented community transmission for a full year while keeping the economy open enough to grow, and no other country of its size has matched that.

Words to know

- **F0 to F5.** Vietnam's labels for degrees of contact, from the infected person outward. It matters because most countries traced one degree, and tracing several is what turned a handful of cases into tens of thousands of quarantines.

- **Centralised quarantine.** Government-run facilities where exposed people stay for 14 days, as opposed to isolating at home. It matters because 4,296,302 of Vietnam's quarantines were of this kind, which is the policy's real weight.
- **Nosocomial transmission.** Infection acquired in a hospital. It matters because that is how the Da Nang outbreak started, and because Vietnam's SARS deaths were almost entirely clinical staff.
- **RT-LAMP.** Reverse transcription loop-mediated isothermal amplification, a way of amplifying viral genetic material without the temperature cycling PCR needs. It matters because the first domestic kit used it, at US\$15 and 70 minutes.
- **RT-PCR.** Reverse transcription polymerase chain reaction, the standard laboratory test that copies a known target sequence until it is detectable. It matters because it was the backbone of Vietnam's confirmatory testing.
- **Event-based surveillance.** Letting ordinary people, including teachers, pharmacists and religious leaders, report unusual health events. It matters because Vietnam built this in 2018, before it needed it.
- **Health declaration.** A required self-report of symptoms and travel, filed through an app. It matters because it was how Da Nang's 1.5 million returning visitors were sorted into who needed a test.

Numbers in this reading

- 1,465 laboratory confirmed cases and 35 deaths as of December 31, 2020, in a country of nearly 100 million.
- SARS: 63 cases and five deaths, and the first country the WHO declared SARS-free.
- Public health expenditure per capita rising an average 9 percent a year between 2000 and 2016; infant and maternal mortality more than halved between 1990 and 2018.
- 24 hours: the deadline for hospitals to report notifiable diseases to the central database, required since 2016.
- First case January 23, 2020; steering committee meeting every two days; about 200 close contacts tracked and isolated within a day of the second wave's first case on March 6.
- 270 confirmed cases at 100 days (May 1), with no community transmission since April 15; 99 days without community transmission by July 25.
- Da Nang: a city of more than 1 million with about 8 million tourists a year; new cases up about 30 percent in the last week of July; 551 locally transmitted cases across 15 cities and provinces between July 25 and September 8; approximately 98 percent linked to Da Nang hospitals or visits; more than 1.5 million people returned from Da Nang between July 1 and 27, 41,000 of them

having visited Da Nang Hospital; at least 251 cases at Da Nang General Hospital.

- September 10: 61,968 monitored, 998 quarantined in health care facilities, 15,619 in centralised facilities, 45,351 at home.
- Test kits: US\$15 and 70 minutes (RT-LAMP, February 7); less than US\$21 and 80 minutes (real-time RT-PCR, March 3); US\$19 to \$25 and over an hour with four times the CDC kit's sample capacity (Viet A, March 5), which carried up to 80% of Vietnam's testing during the global stock-out and drew orders from 20 countries and territories.
- Testing sites: two in late January, 120 by May, 137 laboratories by October 22 with a maximum 51,000 tests a day, 62 screening and 75 confirmatory.
- More than 700,000 tests between July 23 and September 24; 72,492 Da Nang households, a third of the city's, tested between September 3 and 10, plus 13,776 students and staff, all negative.
- Five days average incubation, infectiousness starting two days before symptoms, leaving three days to find a contact.
- F1 definition: within approximately 6 feet (2 meters), or prolonged contact of 30 or more minutes.
- 43 percent of cases never develop symptoms.
- Over 200,000 people in quarantine facilities between January 23 and May 1; 10,242,896 total

quarantined by December 31, 2020, split 211,983 in health care facilities, 4,296,302 centralised and 5,734,611 at home.

- Bluezone: a 60 percent population target, more than 20 million downloads by August 20.
- Lockdowns: 10,000 people in Son Loi when the country had 16 cases; 190 people on one Hanoi street; 4,000 to 5,000 at Bach Mai Hospital after 45 positives, with over 15,000 tested and 40,000 contacts tracked; 7,600 in Dong Van locked down before test results and released the next day.
Nationwide lockdown from April 1, 15 days extended to 21 in 28 of 63 provinces; 10 provinces locked down during the Da Nang outbreak.
- Five confirmed cases in the country when all schools were ordered closed on January 31; single digits when flights to China were suspended on February 1.
- 64 million active Facebook users; 80 percent of smartphone users with Zalo installed; more than 8 billion views on the #Onhavanvui campaign.
- More than 200 police cases pursued under the April 14 misinformation decree; 91 percent public approval of the response in YouGov data as of December 21.
- Nearly 200,000 billion dong (US\$9.4 billion) in lost revenue; 21 million international arrivals in a typical year; fruit exports down by half April to June with freight at two to four times normal prices; 4.5

million made-in-Vietnam PPE suits bought by the US government; mask exports down by a third in August.

What the author is unsure about

The authors are advocates for this case, and the hedges are narrow but real. They say the success has been attributed to several factors rather than proving causation, and that the high share of never-symptomatic cases "suggests" quarantining on exposure may have been a key contributor. They say the NCOVI app "may have helped" slow transmission, and note it drew criticism from privacy advocates; they report the misinformation decree's 200-plus police cases and the opposition from Amnesty International without resolving it. They acknowledge Vietnam had a high baseline GDP growth rate before the pandemic, which qualifies the growth claim, and they concede some of the success derives from Vietnam's unique context even while arguing the lessons generalise. Our World in Data's own notice is the largest caveat: the article reflects data available at the time, and it ends at the close of 2020, before the variants that changed what containment could achieve.

For the discussion

- The transferable idea is quarantining on exposure instead of symptoms, and the number that shows

what it costs is 10,242,896 people quarantined to hold 1,465 cases.

- Worth raising as a tension: Vietnam's one real failure started inside hospitals, which is also where its SARS deaths were concentrated, so the system that protects a population and the system that protects clinical staff are not the same system.

3.4-1 • CORE • COURSE TIME 8 MIN

Next Generation Sequencing

ClevaLab • 2023

[Read the original paper →](#) • the course assigns: whole video, summarized from its auto-generated captions

Reading a human genome took 32 years the first time and now takes about a day, because next generation sequencing cuts DNA into millions of short pieces, reads billions of them at once, and puts them back together against a reference.

Average summary

This summary is written from the video's auto-generated captions, so a few technical words come through garbled and are flagged where that happens.

The video opens with the comparison that makes the point. The Human Genome Project started in 1990 and took until 2000 to complete 85 percent of the

first genome; the gaps were filled in 2022, so the complete sequence took **32 years** in total. With **next generation sequencing (NGS)**, sequencing a person's entire genome takes about **a day**.

The difference is parallelism. NGS sequences billions of DNA strands at the same time; **Sanger sequencing**, all the Human Genome Project had, reads one strand at a time. The video is careful to add that NGS only works because the Human Genome Project produced a human reference sequence in the first place: NGS cuts DNA into small pieces, reads them, and assembles the pieces against that reference.

The workflow has three stages. **Sample and library preparation**: collect the sample, purify the DNA or RNA, check it is pure and not degraded (the captions say "undergraded"; the standard term is undegraded), reverse-transcribe RNA into DNA if needed, then cut the DNA into short pieces of a chosen size using high frequency sound waves or enzymes. Short known sequences called **adapters** are attached to both ends of every fragment. The adapters carry what the machine needs to start sequencing, plus an **index**, a barcode saying which sample the fragment came from. Unbound adapters are washed off, and there may be a PCR step to increase the amount of library. A good library is the right fragment size and concentrated enough to sequence.

Sequencing: the main instruments are Illumina's, using **sequencing by synthesis**. It happens on the glass surface of a **flow cell** coated with short DNA pieces called oligonucleotides that match the adapters. The library is split into single strands, binds to the flow cell, and each fragment is copied over and over in place to make a **cluster**, because one molecule's fluorescent signal would be too faint to see. Then the machine adds the four nucleotides G, C, T and A, each with a different colour tag and a terminator so only one can be added per cycle, plus DNA polymerase. The matching base binds, a camera photographs the colour of every cluster, the terminator is stripped off, and the cycle repeats for as many reads as the run is set for. The indexes are read, and in **paired-end sequencing** the fragment is also read from the other end.

Analysis: bad reads are filtered out, including overlapping clusters, clusters that lead or lag, low-intensity ones and nanowells holding more than one fragment. Surviving reads are **demultiplexed**, sorted back to their samples using the indexes, then mapped onto the reference genome. Two metrics matter: **read depth**, how many reads cover a given nucleotide, where 30x average is good for a whole genome and 1500x suits finding rare mutation events in cancer; and **coverage**, meaning no missing areas across the target.

Why it's assigned

Option 2 asks how the two main next generation sequencing technologies work, and this is the first half of that answer, the sequencing-by-synthesis method Illumina uses. The whole detect layer rests on this machinery: the metagenomic systems in the day's other readings are Illumina flow cells reading everything in a wastewater or swab sample, and their cost and turnaround arguments are arguments about reads, depth and how many samples share one flow cell. Watch it for three ideas in particular: parallelism is where the speed comes from, indexes are how hundreds of samples share one run, and depth is what you buy when you spend more.

In depth

The speed comparison

The video's frame is historical. The Human Genome Project began in 1990 and by 2000 had completed 85 percent of the first genome; the remaining gaps were filled in 2022, making 32 years from start to a genuinely complete sequence. The captions carry an inline "[correction]" marker at this point, so the speaker is correcting a date on screen. Against that, NGS sequences an entire human genome in a day.

Two reasons are given. First, volume: billions of DNA strands are sequenced simultaneously with NGS, where Sanger sequencing handles one at a time.

Second, and more interesting for this unit, NGS is only possible because the Human Genome Project built a reference. The basic principle is that DNA is cut into small pieces, each piece is sequenced, and the pieces are assembled by matching them against the reference sequence. That dependence is worth holding onto: a method that assembles against what is already known has to work differently when what it finds is new.

Preparing the sample and building the library

Samples are collected and the DNA or RNA is purified, then checked to confirm it is pure and not broken up. RNA has to be reverse-transcribed into DNA before it can be sequenced, since the chemistry reads DNA.

Then comes the **library**, which here means a collection of short DNA fragments taken from a long stretch of DNA, not a database. Making one has three moves. Cut the DNA into short pieces of a specified size, using either high frequency sound waves to shear it physically or enzymes to cut it chemically. Attach adapters, short known sequences, to each end of every fragment; the adapters contain the information the sequencer needs and an index identifying the sample. Wash away any adapters that did not bind. Depending on the application, a PCR step can be added to increase the amount of library material. Success is judged on two things: the

fragments are the correct size, and the concentration is high enough to sequence.

Sequencing by synthesis, step by step

The instruments are Illumina's, and the method is sequencing by synthesis: build the complementary strand one base at a time and watch which base goes on.

It happens on the glass surface of a flow cell, which is coated with two kinds of oligonucleotides, short DNA pieces that match the library's adapter sequences (the captions render these as "aligos"; the standard term is oligos). The library is denatured into single strands and flowed in. A strand binds one of the two oligos; the video calls that the forward strand. The reverse strand is synthesised, the forward strand is washed away, and the library is now anchored to the flow cell.

At this point there is a signal problem: a single molecule's fluorescence would be too dim to detect. So each unique fragment is amplified in place into a cluster of identical copies. This clonal amplification is a PCR run at a single temperature, where annealing, extension and melting are driven by changing the solution flowing over the flow cell instead of changing the temperature. The strand bends over and binds the second oligo to form a bridge, the strands are copied, the double strands are separated, and the cycle repeats until there is a localised cluster. Finally the

reverse strands are cut and washed off, leaving forward strands ready to read.

Reading works like this. A sequencing primer binds the forward strands. Fluorescent versions of the four nucleotides, G, C, T and A, flow in with DNA polymerase. Each nucleotide carries a different colour tag and a terminator, so only one base can be added per cycle. The complementary base binds, a camera reads and records each cluster's colour, a new solution removes the terminators, nucleotides and polymerase flow again, and the next base is added. These cycles continue for the read length the sequencer is set to.

Then the barcodes. The read sequences are washed away, the first index is sequenced and washed away, and if a single read is all that is needed, the run ends there. For paired-end sequencing the machine also reads the second index and then the fragment's reverse strand. There is no primer for the second index read, so a bridge is formed and the second oligo acts as the primer. The two index reads are unique dual indices, and they are what allow up to 384 samples to share the same flow cell.

Cleaning up and mapping

After sequencing, bad reads are filtered out: clusters that overlap, clusters that lead or lag the sequencing cycle, and clusters of low intensity. The video notes

that clusters cannot overlap on a patterned flow cell (the captions say "patent"; the standard term is patterned, meaning the surface is etched with fixed nanowells), but more than one library fragment can land in one nanowell, and those polyclonal wells are filtered out too.

Reads that pass the filter are demultiplexed, which means using the attached indexes to sort every read back to the sample it came from. Then the reads are mapped to the reference genome, where they align and overlap one another.

Paired-end sequencing helps here, because two reads come from the same original fragment and the alignment algorithm (the captions say "algorithm") knows they belong together. That lets longer stretches of DNA or RNA be analysed with more confidence that the alignment is right.

The two metrics

Read depth is the number of reads covering a given nucleotide, and average read depth is the average across the region sequenced. For whole genome sequencing, 30x average read depth is good. For detecting rare mutation events in cancer, 1500x average read depth is suitable, because a mutation present in a tiny fraction of cells will only show up if you read the same position many times over.

Coverage is the other, and the aim is simply to have no missing areas across the target DNA.

Where it is used

The video closes with the breadth: diagnosing cancer and rare disease, guiding cancer treatment, and research from ecology to botany to medical science. Both DNA and RNA can be sequenced, as a whole genome or whole transcriptome, as just the coding regions (called **exomes**), or as targeted genes. All types of RNA can be sequenced, including non-coding RNAs such as microRNAs and long non-coding RNA, and the method also extends to cell-free DNA, single cells, methylation and protein binding sites.

Words to know

- **Next generation sequencing (NGS).** Reading huge numbers of short DNA fragments in parallel and assembling them. It matters because every metagenomic system in this unit is built on it.
- **Sanger sequencing.** The older method that reads one strand at a time. It matters as the comparison that explains the 32 years versus one day.
- **Library.** A set of short DNA fragments cut from a longer stretch, with adapters attached. It matters because what you can detect later depends on how the library was built.
- **Adapter and index.** A short known sequence added to each fragment end, carrying what the

machine needs plus a barcode for the sample. It matters because indexes are how up to 384 samples share one flow cell, which is where cost per sample comes from.

- **Flow cell.** The glass surface inside the sequencer where the reaction and imaging happen. It matters because flow cell capacity, quoted elsewhere in this unit as billions of reads, is the unit biosurveillance budgets are measured in.
- **Cluster, and clonal amplification.** Copying one fragment in place many times so its fluorescent signal is bright enough to photograph. It matters because it explains why the machine reads clusters rather than single molecules.
- **Sequencing by synthesis.** Building the complementary strand one tagged base at a time and recording the colour. It matters because it is the Illumina method the course asks you to compare with nanopore sequencing.
- **Paired-end sequencing.** Reading the same fragment from both ends. It matters because knowing two reads belong together makes alignment more reliable over longer stretches.
- **Demultiplexing.** Sorting the mixed reads back to their samples using the indexes. It matters because without it a pooled run is unusable.
- **Read depth and coverage.** How many times each position is read, and whether any part of the target

is missed. It matters because sensitivity in this field is bought in depth.

Numbers in this reading

- 1990 to 2000: the Human Genome Project period in which 85 percent of the first genome was completed; 2022: when the gaps were filled; 32 years in total.
- About one day: how long NGS takes to sequence a person's entire genome.
- Billions of DNA strands sequenced simultaneously with NGS, against one at a time with Sanger sequencing.
- Up to 384 samples in the same flow cell, made possible by unique dual indices.
- 30x average read depth: good for whole genome sequencing. 1500x average read depth: suitable for detecting rare mutation events in cancer.
- Four nucleotides, G, C, T and A, each with its own colour tag and a terminator so only one is added per cycle.

What the author is unsure about

The video presents all of this as settled method and raises no doubts; it is an explainer, not an argument. The only hedges it offers are practical conditions rather than uncertainties: a library has to be the right size and concentration to work, a PCR step is added

only depending on the application, and read depth targets differ by purpose. The real uncertainty in this summary is the source: it is written from auto-generated captions, which garble several technical words, and three are flagged in the text above.

For the discussion

- The sentence worth repeating is that NGS only works because the Human Genome Project built a reference, which is precisely why finding something genuinely new is a different and harder problem than resequencing something known.
- Worth raising: indexes are the hidden economics of this unit, since up to 384 samples sharing one flow cell is what makes pooled wastewater and pooled swab surveillance affordable at all.
- A question worth asking: cancer work uses 1500x depth to catch a mutation present in a few cells, so what depth would you need to catch a pathogen present in a few people's contribution to a city's sewage.

3.4-2 · CORE · COURSE TIME 11 MIN

Nanopore Sequencing

BestDx Academy · 2023

[Read the original paper →](#) · the course assigns: whole video, summarized from its auto-generated captions

Nanopore sequencing threads a single strand of DNA or RNA through a tiny protein hole and reads the sequence from the dips in electrical current as each base passes, which makes it long-read, real-time and portable, at the cost of being less accurate per base than Illumina.

Average summary

This summary is written from the video's auto-generated captions, which have no punctuation and mis-hear several technical names; the correct terms are flagged below.

This is the second part of a series on third generation sequencing. The first covered SMRT sequencing (captions say "smart"; the standard term is SMRT, single-molecule real-time, the PacBio method). This one covers **Oxford Nanopore Technologies (ONT)** sequencing, which the captions render variously as "o andt", "ENT" and "on".

The core idea: a **nanopore** is a tiny protein channel sitting in an electrically resistant membrane that separates two chambers of salt solution. Ions flow through the pore, creating a steady electrical current. When a DNA or RNA molecule is pulled through the pore, each base in turn blocks it slightly differently, and that pattern of disruption, which the video calls a **squiggle**, is converted into a sequence. Every pore

has its own electronic sensor, and a flow cell holds thousands of pores.

The step that matters most for this unit: most sequencing methods need the target DNA amplified by PCR first, and nanopore does not. It can sequence a single native molecule of DNA or RNA.

The workflow: extract and purify the nucleic acid; repair the molecule ends and ligate on sequencing adapters (captions say "n repair" and "liation"; the standard terms are end repair and ligation). Those adapters carry a **motor protein**, an enzyme that unzips the molecule one base at a time and controls how fast it moves through the pore, plus a **tether** that acts as a flexible linker holding the molecule near the pore. The library is loaded onto the flow cell, the squiggles are recorded, and **base calling** software, ONT's Guppy, uses neural networks trained on large datasets to turn the electrical signal into A, C, G and T or U. Then standard bioinformatics: alignment, genome assembly, structural variant detection, transcriptome analysis or metagenomic analysis.

Its advantages, in the video's words: long reads, real-time data as it goes, and native molecules without amplification. Its stated disadvantage is blunt: a **higher error rate** than other high throughput technologies such as Illumina (captions say "alumina").

The video compares the two third generation methods. Both make long reads spanning tens of thousands of base pairs or more, both read single molecules without amplification, and both produce data in real time, which it says is especially useful for time-sensitive work such as pathogen detection and outbreak monitoring. The differences: SMRT watches fluorescently labelled bases being incorporated by a polymerase, nanopore measures current; SMRT has higher per-base accuracy; only nanopore can sequence RNA directly without first converting it to cDNA; and only nanopore is portable, with the pocket-sized **MinION** against PacBio's benchtop Sequel II.

Why it's assigned

This is the other half of the answer to Option 2's first guiding question, and the direct comparison the course's second question asks for: why both technologies are commercially viable with no clear winner. The trade-off is the whole point. Illumina gives you cheap, deep, accurate short reads, which is why it carries the big centralised wastewater and swab programmes in the day's other readings. Nanopore gives you long reads, results while the run is still going, direct RNA, and a sequencer you can carry, which is why one of the four real biosurveillance systems in the day's core reading uses long-read sequencing and why the UK's hospital programme does too.

In depth

The principle

Nanopore sequencing is described as a high throughput technology for real-time long-read sequencing of DNA or RNA, based on protein nanopores embedded in an electrically resistant membrane. The membrane separates two chambers filled with an electrolyte solution containing ions. The pore lets ions pass from one chamber to the other, which produces an ionic current across the membrane. When a nucleic acid molecule passes through the pore it changes that current, and the change can be measured and used to work out the sequence.

The video states the contrast with most other methods plainly: they require the target DNA to be amplified by PCR, and nanopore sequencing can read a single molecule of DNA or RNA without any amplification.

Sample preparation

DNA or RNA is extracted and purified from the sample. The captions then say that if the sample contains double-stranded DNA it must be converted into single-stranded DNA, and that for RNA samples the reverse transcription step into complementary DNA (cDNA) is not needed, which is the direct-RNA capability the comparison section later returns to.

Library preparation

Two steps: **end repair** and **adapter ligation**. The molecule's ends are repaired first to create a uniform structure, then specific sequencing adapters are ligated to them so the molecules can enter the pores and be sequenced.

Two components ride on the adapter and do the mechanical work. The **motor protein** is an enzyme attached to the adapter that controls the speed at which the strand passes through the nanopore; the video says it essentially unzips the double-stranded molecule one base at a time and feeds it through in a controlled way. The **tether sequence** is a flexible linker between the target molecule and the motor protein, which keeps the movement smooth and controlled.

That control is the real engineering trick of the whole method, and worth saying plainly: the signal is only readable if the molecule moves at a steady, known speed, so a protein is used as a brake.

Loading and the flow cell

The prepared library is loaded onto a flow cell containing thousands of nanopores. Each nanopore is a tiny protein channel in the membrane, and each one is wired to its own electronic sensor that measures the current change as molecules pass through. So throughput is the number of pores running in parallel,

in the same way that Illumina's throughput is the number of clusters on a flow cell.

Sequencing and base calling

As the molecule moves through a pore, each base temporarily occupies the narrowest part and disrupts the current in its own characteristic way, producing the **squiggle**: a raw electrical trace rather than a picture of coloured dots. The video grounds this in the physical difference between the bases' chemical properties and the structure of the pore.

Base calling is the step that turns that trace into letters. ONT supplies software called **Guppy** that uses neural networks (captions say "neuron networks") trained on large datasets to recognise the distinctive patterns of each nucleotide and of combinations of them. This is worth noticing: unlike sequencing by synthesis, where a camera reads one colour per cycle, nanopore's accuracy is partly a machine learning problem, because several bases sit in the pore at once and the signal has to be disentangled.

Because base calling happens as the molecule goes through, data is generated in real time rather than at the end of a run.

Analysis and uses

Once there is a sequence, it can be aligned, assembled and analysed with ordinary bioinformatics

tools: genome assembly, structural variant detection, transcriptome analysis or metagenomic analysis, depending on the study.

The clinical applications listed, which the video attributes to the combination of long reads, real-time data and portability: rapid pathogen detection and identification, antibiotic resistance profiling, cancer genomics and personalised medicine, HLA typing for transplantation, and prenatal and neonatal genetic testing.

Three platforms are named, in order of scale: **MinION**, a portable pocket-sized sequencer suited to fieldwork and small projects; **GridION**, a benchtop machine for larger projects and higher throughput than MinION; and **PromethION** (captions say "Promethean"), a high throughput instrument for large-scale projects and whole genome sequencing.

SMRT and nanopore compared

The video sets the two third generation technologies side by side.

What they share: both generate long reads spanning tens of thousands of base pairs or more; both sequence individual molecules without amplification; and both generate data in real time, which the video says is particularly useful in time-sensitive applications such as pathogen detection and outbreak monitoring.

Where they differ. **Method:** SMRT monitors the incorporation of fluorescently labelled nucleotides by DNA polymerase as a new strand is synthesised, while nanopore measures changes in ionic current and infers the sequence from them. **Accuracy:** SMRT generally has higher per-base accuracy than nanopore. **Direct RNA:** nanopore can sequence RNA directly without reverse transcription into cDNA, and SMRT currently cannot. **Portability:** ONT offers portable devices like the MinION, while SMRT platforms such as the PacBio Sequel II (captioned say "PacBio Sequel II") are benchtop instruments not designed to be carried.

The honest limitations

The closing section lists what third generation sequencing still gives up against short-read platforms, even granting its advantages of long reads, de novo genome assembly, direct RNA sequencing, detection of DNA modifications and real-time data.

Lower per-base accuracy than short-read platforms, which it says can be mitigated by generating higher coverage or by building a consensus sequence from multiple reads of the same molecule (captioned say "census"; the standard term is consensus). **Higher cost**, which has been decreasing but can still exceed short-read sequencing, particularly at large scale. **Computational requirements**, since long-read analysis is computationally intensive and needs

specialised bioinformatics tools, expertise and substantial compute for jobs like genome assembly and error correction. **Lower throughput** on some platforms than short-read machines, which limits large projects, though it notes newer instruments like PacBio Sequel II and PromethION are built for higher throughput.

Its conclusion is a straightforward claim for the field: third generation sequencing has transformed genomics through long reads and real-time data.

Words to know

- **Nanopore.** A tiny protein channel in an electrically resistant membrane, through which ions and a single nucleic acid strand pass. It matters because it is the whole sensor: no fluorescence, no camera, just current.
- **Ionic current.** The steady flow of charged particles through the pore, which each base disturbs differently. It matters because this is the raw measurement the sequence is inferred from.
- **Squiggle.** The raw electrical trace produced as a strand passes through the pore. It matters because it has to be decoded by software, not simply read off.
- **Base calling.** Converting that raw signal into A, C, G and T or U, which ONT does with neural networks in its Guppy software. It matters because accuracy

here is a model's accuracy, and it improves when the models improve rather than when the chemistry changes.

- **Motor protein.** An enzyme on the adapter that unzips the molecule and controls how fast it threads through the pore. It matters because a steady, known speed is what makes the signal interpretable.
- **Tether.** A flexible linker between the target molecule and the motor protein. It matters because it keeps the molecule in position for a controlled pass.
- **End repair and adapter ligation.** Tidying the molecule's ends and chemically joining the sequencing adapters on. It matters because this is the entire library preparation here, with no amplification step.
- **Long read.** A single read spanning tens of thousands of bases or more, rather than a few hundred. It matters because long reads can be assembled with no reference genome, which is the case for something nobody has seen before.
- **De novo assembly.** Reconstructing a genome from reads alone, without matching them to a known reference. It matters because the Illumina explainer said next generation sequencing works by assembling against a reference, and this is the alternative.

- **SMRT sequencing.** Single-molecule real-time sequencing, PacBio's method, which watches a polymerase add fluorescent bases to one molecule. It matters as the other third generation technology, and the one this video compares against throughout.
- **Consensus sequence.** One corrected sequence derived from several reads of the same molecule. It matters because it is the stated fix for nanopore's per-base error rate.

Numbers in this reading

- Thousands of nanopores in a single flow cell, each with its own electronic sensor.
- Tens of thousands of base pairs or more: the read length both third generation technologies reach.
- Three ONT platforms by scale: MinION (pocket-sized, fieldwork), GridION (benchtop, higher throughput), PromethION (high throughput, whole genome).
- One base at a time: the rate at which the motor protein unzips and feeds the molecule through the pore.

What the author is unsure about

The video is an explainer and claims no uncertainty about the method, but it is unusually direct about the technology's weaknesses: a higher error rate than Illumina, lower per-base accuracy than short-read

platforms generally, higher cost at scale, heavy computational requirements, and lower throughput on some platforms. It presents the fixes for those as mitigations rather than solutions, namely more coverage or consensus sequences from repeated reads of the same molecule. It also says cost has been decreasing without giving figures. The uncertainty in this summary is the source: the captions have no punctuation and mis-render SMRT, Illumina, PacBio Sequel II, PromethION, end repair, ligation, neural networks and consensus, each flagged where it appears.

For the discussion

- The useful frame is that there is no winner because the two methods fail in opposite directions: short reads are cheap, deep and accurate but assemble against a reference, while long reads are costlier and less accurate per base but can assemble something with no reference at all.
- Worth raising: nanopore's accuracy depends on neural networks decoding a squiggle, so this is one of the few places in the unit where the detection limit improves when a model improves rather than when a machine is replaced.
- A question worth asking: direct RNA sequencing without reverse transcription is unique to nanopore, so what does that buy you for an RNA

virus outbreak, where every other method has to copy the RNA into DNA first.

3.4-3 · CORE · COURSE TIME 40 MIN

Towards ubiquitous metagenomic sequencing: a technology roadmap

Nava Whiteford, Andrew Heron, Leonard McCline, Ales Flidr, Jacob Swett, Ajay Karpur · 2023

[Read the original paper →](#) · the course assigns: whole piece

If a sequencing device could match a PCR test on four numbers (sensitivity, under an hour, under \$10 a test, under \$10,000 an instrument) then every clinic could test for any pathogen at all, and the authors argue that is technically achievable but will not happen on its own, because the sequencing market is built to sell accuracy and read length to researchers instead.

Average summary

The authors start with the gap COVID-19 exposed. SARS-CoV-2 circulated in people as an unknown pneumonia, invisible to standard-of-care tests, for an estimated **4 to 8 weeks** before metagenomic sequencing identified a novel coronavirus in patient samples. Their thought experiment is the whole report: imagine every clinician and patient had a

simple device that could detect any virus, bacterium or other pathogen causing disease in a symptomatic patient.

Metagenomic sequencing (MGS) means reading all the genetic material in a sample instead of testing for one suspect. It already saves lives in hospital, but complex workflows and high costs keep it out of routine use. The benchmark they set it against is **qPCR**, the test that was the workhorse of COVID-19 testing: fast, sensitive, specific and cheap. PCR's fatal limitation is that it can only find what you designed it to find, so a new resistance gene, a new strain, or a completely new pathogen like SARS-CoV-2, HIV, Ebola or Marburg would go undetected.

Their target specification: sensitivity comparable to qPCR (ideally matching a **Ct of 35**, meaning very little virus present); a push-button workflow needing under 5 minutes of operator time; **time to answer under 1 hour**; cost of goods per test around **\$10**; an instrument ideally under **\$10,000**, compact and portable; and supply chains that are not overly complex.

No existing sequencer comes close to all of that at once. Their platform comparison shows the trade: an Illumina NovaSeq 6000 costs about \$50 per sample but takes 13 hours to run, while a MinION is ready after 10 minutes of library prep and runs in under an hour, at about \$250 per sample.

Their performance model is the useful part. Because MGS has to read through everything else in the sample to reach the pathogen, its sensitivity depends on the pathogen's **relative** abundance, not its absolute amount. In nasopharyngeal samples, SARS-CoV-2 RNA was typically 0.01% to 0.001% of the material, one fragment in 10,000 to one in 100,000. From that they conclude a device delivering **1 to 10 million reads per hour**, with reads as short as **25 base pairs** and single-base accuracy of only **95%**, would robustly detect most pathogens in respiratory samples. Detection needs far less accuracy and far shorter reads than research does.

Two technology families could get there, both single-molecule: **nanopore** sequencing, which already meets the speed and workflow targets but whose consumables cost too much because the pore arrays are expensive to manufacture, and **single-molecule optical** sequencing, currently sold in half-million-dollar instruments but fast, with a clearer path to a \$10 consumable if rebuilt for cheapness rather than accuracy.

Their last argument is economic: more than half of global sequencing capacity sits in about 1,500 Illumina NovaSeq machines serving researchers who want low cost per base on huge runs. Nobody is building the cheap, shallow, fast device, so they call for push mechanisms (direct funding, as with the

Human Genome Project) and pull mechanisms (advance market commitments, a promise to buy at volume).

Why it's assigned

This is the technical heart of Option 2, which asks what it would take to make metagenomic sequencing as ubiquitous as PCR. The course warns the option gets technical and suggests it for people who have studied biosciences at undergraduate level, so take the specifications as the thing to understand and the manufacturing detail as background. For the detect layer it answers a question the day's other readings leave open: those systems put sequencers in a handful of central laboratories reading pooled samples, and this report asks what would change if the sequencer were in the clinic, reading one patient at a time, on the day they walk in.

In depth

The problem, stated as a time delay

The opening is specific. SARS-CoV-2 circulated in human populations as an unknown pneumonia, undetected by standard-of-care diagnostics, for an estimated 4 to 8 weeks before unbiased metagenomic sequencing identified a novel coronavirus in infected samples. That delay cost time in a pandemic that went on to kill millions and cost trillions of dollars.

MGS has shown it can be a near-universal diagnostic and already informs clinical decisions across various symptoms and sample types, but complex workflows and high costs have kept it from wide adoption, so it had limited impact on managing COVID-19 at the point of care. Their framing sentence is worth keeping: regardless of their value for public health, sequencing-based tests have to become comparable with the standard of care before anyone will treat them as a real alternative.

They also say clearly what this report is not about. MGS can be used for early detection through wastewater sequencing or testing sentinel populations, and they agree those belong in a layered system, but this roadmap is only about MGS at the point of care on human respiratory samples. Four reasons: telling signal from noise is likely easier in a human sample than an environmental one; early detection needs large scale, and clinical usefulness lets scale grow inside the existing health economy without permanent subsidy; finding a pathogen in one person's sample immediately enables quarantine and contact tracing; and widely available MGS would make mass testing of a new pathogen possible from day zero, saving the months it takes to design and approve new PCR primers.

The model to copy: how PCR got everywhere

Their precedent is the point-of-care PCR machine, which started bulky, expensive, confined to specialised laboratories and hours long, and became an affordable push-button device needing no specialist. A WHO report credits real-time PCR with wider adoption of molecular diagnostics through rapidity, sensitivity, reproducibility and reduced carry-over contamination; the authors note that of those four, only sensitivity is not still a problem for MGS.

The concrete case is the **Cepheid GeneXpert**, originally developed to detect anthrax and then adapted to many diseases through collaboration with international organisations and philanthropies. Successive rounds of cost-optimisation made it practical in developing countries for tuberculosis and HIV, giving an installed base of some 22,000 devices before COVID-19, rising to 40,000 between 2019 and 2022, and because those machines were already in place, SARS-CoV-2 primers could be deployed without building new infrastructure. The user experience is the standard to beat: put the sample in a cartridge, insert the cartridge, wait about 45 minutes. PCR can also be multiplexed: the **BioFire FilmArray** detects 20 to 40 targets at once across respiratory, blood culture, cerebrospinal fluid and gastrointestinal samples.

And then the limit. PCR cannot in principle detect unknown or changing targets, because it needs primers designed against a known sequence. A novel

resistance gene or new viral strain needs new primers, and a completely novel pathogen would go completely undetected. Their examples of what that has meant: SARS-CoV-1 and 2, HIV, Ebola, Marburg.

The target specification

Six requirements for an MGS diagnostic: sensitivity comparable to qPCR, ideally matching a Ct of 35; a push-button workflow needing under 5 minutes of operator time and minimal skill; time to answer under 1 hour; cost of goods for one test comparable to qPCR at \$10; an affordable instrument, ideally under \$10,000, compact and ideally portable; and supply chains that are not overly complex.

Their platform comparison shows nothing meets the set, and the trade-offs run in every direction: Illumina's machines need 1.5 hours of library preparation and 1.5 to 13 hours of run time at \$30 to \$50 a sample with no direct RNA; Ion Torrent runs in 10 minutes but needs 2 hours of preparation; the ONT MinION needs only 10 minutes of preparation and runs in under an hour, with direct RNA, but costs about \$250 a sample; SeqLL, the former Helicos optical system, needs no library preparation at all but takes 7 days; PacBio's Sequel II runs in under an hour after 3 hours of preparation. The full figures are in the numbers section below.

They diagnose why. Sensitivity alone is achievable with enough depth but is incompatible with the cost and time targets. Workflows are complex sequential steps in sample and library preparation and usually need hours of trained expert work; the automated options, VolTRAX for nanopore and NeoPrep for Illumina, are expensive and for specialised laboratories. And cost per sample only falls below \$10 by running many samples together, which is impractical at the point of care because batching adds delay.

One structural point rules out most platforms immediately: anything that needs to grow colonies or clusters on a chip before reading loses at least 60 minutes to cluster generation alone, and needs a large set of reagents that would add cost and complexity to a sample-to-answer box. New entrants such as MGI and Singular Genomics will likely push prices down by competing on margins, but will not change that limit. So the natural place to look is **single-molecule** methods, which can have minimal library preparation and a real-time readout. As of 2023 the two classes are nanopore and single-molecule optical.

Today's cheapest single-sample instrument is ONT's **Flongle**, whose consumables sell for \$90 and probably cost around \$50 to make. Optimised clinical workflows with real-time MinION sequencing reach a

time to answer of 6 hours, driven almost entirely by sample and library preparation since the sequencing itself takes under an hour, and even workflows tuned for intensive care still need a clinical microbiology laboratory to run them.

The workflow problem, and the cartridge

The sharpest contrast with PCR is workflow complexity. Interviews with practitioners at the front of clinical MGS adoption told them that training new staff in MGS workflows takes months, and that results vary significantly with operator skill.

So they ask whether something like Cepheid's cartridge could be built around sequencing. Cepheid's own first and second generation instruments could not handle raw samples, and they solved it by designing a fluidic cartridge that processes a raw sample and holds the qPCR reaction tube, with reagents preloaded and no fluid connection to the instrument, which interfaces only through a reagent selection valve and a plunger. The cartridge's cost of goods including reagents has been estimated at \$10, and it is typically sold for more than \$20.

The figure

The report includes a cartridge sketch figure, and its argument is simplicity by subtraction: a general-purpose cartridge has to serve many assays, but an MGS cartridge could be fixed to one sample type

(nasopharyngeal, upper nasal or saliva), one input volume and RNA sequencing only, which strips the fluidics back to a short chain of steps. What the sketch lays out is that chain: break the cells open (**lysis**), pull out the nucleic acid (**extraction**), convert RNA into DNA (**reverse transcription**, unless direct RNA sequencing becomes more reliable), remove unwanted material (here, digesting DNA with **DNAase** so the RNA stands out), and in most cases add the sequencing adapters. For platforms needing a lot of input, such as nanopore, random amplification may be needed to reach a minimum amount without skewing what is in the sample. What to notice is their conclusion: stripped to essentials, the complexity need not be much greater than qPCR's, and building the cartridge and integrating the system into one box carries relatively little technical risk. The reason nobody has done it, they argue, is not the cartridge at all. It is that no platform can deliver enough sequencing depth cheaply enough without batching samples together.

How deep does it have to be?

This is the analytical core, and it rests on one distinction. PCR targets a short unique region, so it is insensitive to background: a sample full of human or bacterial material barely affects it, and its limit of detection depends on the **absolute** amount of target present. MGS has to read through background

fragments until it happens to read the target, so its sensitivity depends on the target's **relative** abundance among everything else in the sample.

In typical human clinical samples, host nucleic acids outnumber the pathogen's by orders of magnitude. SARS-CoV-2 RNA was typically between 0.01% of nasopharyngeal sample material, one fragment in 10,000, and 0.001%, one in 100,000. But viral load spans some 8 orders of magnitude, and loads as low as tens of copies per millilitre were found in more than 1% of cases.

They also flag that MGS and PCR give different kinds of answer. PCR gives a yes or no about presence, and indirectly an abundance. MGS gives a count of fragments matching a pathogen, alongside fragments of the patient and their harmless microbiome, so someone has to judge whether that count means infection. In practice cutoffs will be set per pathogen as data accumulates on what is normal in uninfected people, and the two published examples they give, in the numbers section below, use reads per million and deviation from the rest of a cohort.

The depth arithmetic is simple and worth following once: expected pathogen reads equal the pathogen's share of the sample multiplied by the number of reads in the run. If one fragment in 100,000 is the pathogen and you take 1 million reads, you expect about 10 reads. Their modelled table maps target

fraction to expected reads and to a predicted PCR Ct value, and the row that matters is one target in a million, which gives 10 reads at a 10 million read run and a predicted Ct of 35.8; the rest of the rows are in the numbers section. As a rule of thumb they note Ct 25 means a high viral load and values over 35 are of disputed clinical relevance.

Reading off that table: for a respiratory sample with 10 ng of nucleic acid, a device delivering 10 million reads in the available time yields more than 10 expected reads at Ct 35. A device managing only 1 million reads would be unreliable at that level, though it could still be clinically useful where a limit of detection ten times higher is acceptable. With a one hour budget, that lands on the throughput target of 1 to 10 million reads per hour for most respiratory applications. They are explicit that this does not generalise: sample types with more human RNA, such as blood, where low-abundance pathogens like HIV live, need one or two orders of magnitude more depth. They also note that at 40 million reads per sample, MGS has been shown to match or beat PCR's sensitivity, which is the point of comparison their cheap device is trading away.

Depletion is the lever that could change all of this: removing known non-target material from a library lowers the depth you need and raises sensitivity for the pathogen, and ribosomal RNA is the obvious

target because it can be 60 to 95% of an RNA sample. Practitioners pointed them to Qiagen's FastSelect, which removes more than 95% of host and bacterial rRNA in 14 minutes, with a reported 10-fold reduction in required depth in cerebrospinal fluid. The catch is cost, over \$50 per sample, though laboratories report similar results after diluting the reagent tenfold. The authors call cheaper rapid depletion kits a high development priority, and then deliberately assume no depletion in the rest of their analysis, which makes their depth targets conservative.

Read length and accuracy matter much less than the market thinks

Sequencers also differ in read length and single-base accuracy, both highly desirable in research, where a single mutation may be the object of study. Their simulations say that for classifying a known virus or spotting a high abundance of unknown material, the requirements are comparatively modest, and that less of one can be made up with more of the other.

The specific claim: a read length of only 25 base pairs with 95% single-base accuracy is sufficient for unique pathogen detection, given enough depth. And the field precedent is stronger still: highly sensitive detection of emerging pathogens has been demonstrated in the field with error rates as high as 20 to 30%, using early long-read nanopore

sequencing during Ebola surveillance. The divergence between what detection needs and what research buys is their central lever on development direction.

Nanopore: fast and simple already, too expensive to consume

In nanopore sequencing, bases are read as a strand passes through a protein pore in a membrane, so throughput is set by how fast bases pass through and how many pores run in parallel. Speed is already near its ceiling, because the strand has to be slowed by a motor protein for the electrical signal to be interpretable at all.

Long reads, nanopore's research advantage, with standard average read lengths over 1,000 bases and record reads over a million, are an active hindrance here: a long fragment occupies a pore for longer, so to a first approximation the number of fragments you read scales with the inverse of average read length. The design conclusion is to add a fragmentation step to sample preparation deliberately, cutting the material shorter to get more reads and a better chance of catching a rare pathogen.

The arithmetic: a MinION has 512 active channels, and at the standard translocation speed of about 400 bases per second with pores well saturated, you need an average fragment length of about 700 bases to get 1 million reads in an hour, and about 70 bases to get

10 million. The cheaper Flongle has approximately 100 active channels, implying some 2 million reads in 60 minutes; elsewhere the report gives the Flongle's pore count as 126, so treat it as roughly a hundred.

So nanopore already meets the targets for time to answer, throughput and simple library preparation, given a sufficiently fragmented sample. The barrier is the cost of consumables, driven by the cost of manufacturing the nanopore arrays, which scales linearly with chip area. That leaves two levers: use less area, or use cheaper materials and fabrication. From published patents they infer ONT's design needs MEMS fabrication facilities able to handle novel materials, so the cost is likely driven by limited global MEMS capacity rather than by materials as such. Alternatives worth trying in parallel: silicon substrate arrays for thousands or tens of thousands of pores, or, for low pore counts, silver electrodes fabricated directly onto printed circuit boards, which might scale to about 100 channels though it is uncertain beyond 16. Their expectation is that a ground-up reimplementaion aiming at cost would land much lower.

Single-molecule optical: surprising candidate, clearer path to \$10

Single-molecule optical sequencing was commercialised by Helicos, now defunct, and PacBio. Including it looks odd, since these have been costly

fridge-sized instruments with library preparation over three hours. The authors' argument is that this is market demand, not physics: those instruments were built for high single-base accuracy and long reads. PacBio's SMRTBell library preparation exists to push accuracy above 99.9% and read length to 15 kb, neither of which MGS needs. Drop those requirements and optical library preparation can be as fast as nanopore's, around 5 minutes, with accuracy upward of 80% compensated by read lengths of hundreds of bases, and Helicos-derived approaches have worked with minimal preparation at error rates below 5%.

On speed, optical wins: many more sensors can work in parallel, PacBio commonly images more than 1 million fragments at once, and at 10 nucleotides per second a run could finish in under a minute at MGS read lengths. Optical also needs a less complex chip and has a clearer path to a \$10 consumable.

The challenge is the instrument. PacBio machines currently exceed \$500,000, but getting under \$10,000 looks feasible if the instrument is reimaged from scratch: photonic chips need not be monolithically integrated, and consumer-grade cameras available for \$300 have enough resolution for the accuracy an MGS diagnostic needs.

Throughput would then be limited by what a cheap camera can see, but 1 to 10 million sensing regions, and therefore reads, looks achievable. They single out

electro-optical zero-mode waveguides as a promising emerging approach that may need much less input material, potentially removing the need for random amplification.

They are careful to say they expect viable solutions in the next 5 to 10 years from these two categories, while not ruling out a colony-based method, if cluster generation got an order of magnitude faster, or unexpectedly fast progress on something new such as solid-state nanopores.

Why the market will not do this by itself

The last section is about incentives, and the authors say it draws on their own experience of the sequencing landscape plus expert conversations.

Development follows customers, and the customers are research and high-end diagnostics such as cancer, where extremely high per-base accuracy matters because every base may be informative, and where read length and depth requirements are also much higher than detection needs. Time to answer of hours or days is not a constraint in research, where workflows are built around enormous runs and what matters is low cost per base, which favours large high-end instruments. The sequencing market is comparatively small, dominated by a few players with relatively enforceable intellectual property, so freedom to operate is limited. And infectious disease

diagnostics would mean lower margins and lower barriers to entry, with a potentially large volume but no clear demand signal to justify the investment needed, which they put at at least tens of millions of dollars for a working product.

Their capacity estimate makes the imbalance concrete. More than half of global sequencing capacity in bases per year is accounted for by Illumina NovaSeq alone, from some 1,500 instruments in large laboratories, each costing \$1 million. Even ONT, known for miniaturised devices like the MinION and Flongle, drew 55.7% of its 2021 revenue from its 67 PromethION instruments.

Two fixes. **Push:** direct public investment in research and development, as the Human Genome Project and then the \$1,000 Genome Project did, with a \$10 MGS device as a challenging but achievable goal to organise grants and support around. **Pull:** advance market commitments, which work here because the target product profile is easy to define and achievable in under 5 years. Building a fabrication plant from scratch to make low-cost nanopore devices currently costs hundreds of millions and cannot be justified without credible demand, but at volumes comparable to qPCR diagnostics, more than 1 million units a month, it could be. Advance market commitments plus credible signals of interest from national and international organisations could unlock that.

Their conclusion is modest and clear: current technology is not suited to widespread adoption, but single-molecule optical and nanopore sequencing have a good chance of reaching a combination of cost and performance that would put this in clinicians' hands worldwide.

Words to know

- **Metagenomic sequencing (MGS).** Reading all the genetic material in a sample at once, then working out what is in it, instead of testing for one suspected organism.
- **qPCR, and Ct value.** Quantitative polymerase chain reaction copies a chosen target sequence over and over until it can be detected. The Ct, or cycle threshold, is how many doubling cycles that took, so a high Ct means there was very little of the target to start with. Ct 25 is a high viral load, and above 35 the clinical meaning is disputed, which is why Ct 35 is the sensitivity target here.
- **Limit of detection (LoD).** The smallest amount of something a test can reliably find.
- **Relative versus absolute abundance.** How much of the sample is the pathogen, versus how much pathogen is there in total. PCR cares about the second, MGS about the first, which is the single most important asymmetry in this report.
- **Read, and sequencing depth.** One read is one fragment the machine sequences; depth is how

many reads you take. More depth costs more money and more time, and buys sensitivity.

- **Reads per million (rpm).** A pathogen's reads scaled to the size of the run, so results from different-sized runs can be compared. It is the unit the published clinical cutoffs are written in.
- **Depletion.** Deliberately removing known non-target material, especially ribosomal RNA, before sequencing, so more of the run is spent on what you care about.
- **Ribosomal RNA (rRNA).** The highly repetitive RNA from the cell's protein-making machinery, which can be 60 to 95% of an RNA sample and is therefore the obvious thing to deplete.
- **COGS.** Cost of goods sold: what it costs to make one test, as opposed to its price.
- **Library preparation.** The chemistry that turns extracted nucleic acid into something a sequencer can read. Here it is the main thing standing between a sample and an answer in under an hour.
- **Lysis, extraction, reverse transcription, DNAase.** Breaking cells open; pulling the nucleic acid out; copying RNA into DNA; and an enzyme that chews up DNA so RNA stands out. These are the cartridge's steps.
- **Colony-based versus single-molecule sequencing.** Methods that must first copy each fragment into a cluster bright enough to see, versus methods that read one molecule directly.

Cluster generation alone costs at least an hour, which is why the report bets on single-molecule.

- **Translocation speed.** How fast a strand passes through a nanopore, about 400 bases per second, deliberately slowed by a motor protein so the signal can be read.
- **MEMS fabrication.** Micro-electromechanical systems manufacturing, the specialised plants that can build tiny mechanical and fluidic structures. Global capacity is limited, which the authors think is what makes nanopore arrays expensive.
- **Zero-mode waveguide.** An optical structure that confines light to a space smaller than its own wavelength, so a single molecule can be watched. The electro-optical version may cut how much sample material is needed.
- **Advance market commitment.** A binding promise to buy a product at a set price if someone builds it, used to create a market where none exists yet.
- **Push and pull mechanisms.** Paying for the research directly, versus promising to buy the result.

Numbers in this reading

- 4 to 8 weeks: how long SARS-CoV-2 is estimated to have circulated as an unknown pneumonia before metagenomic sequencing identified it.
- The target specification: qPCR-comparable sensitivity at Ct 35, under 5 minutes of operator

time, under 1 hour to answer, about \$10 per test, under \$10,000 per instrument.

- The performance targets: 1 to 10 million reads per hour, read lengths greater than 25 base pairs, accuracy greater than 95%.
- GeneXpert: some 22,000 devices installed before COVID-19, rising to 40,000 between 2019 and 2022; about 45 minutes to a result; cartridge cost of goods estimated at \$10, sold for more than \$20.
- BioFire FilmArray: 20 to 40 targets detected simultaneously.
- Platform comparison: NovaSeq 6000, 1.5 hour prep, 13 hour run, \$50; MiSeq, 1.5 and 1.5 hours, \$30; iSeq, 1.5 hours and 2 hours, \$50; Ion Torrent GeneStudio S5, 2 hour prep and a 10 minute run, \$50; ONT MinION, 10 minute prep, under 1 hour, \$250; SeqLL, no prep but a 7 day run, \$50; PacBio Sequel II, 3 hour prep, under 1 hour, \$50.
- At least 60 minutes for cluster generation alone on colony-based platforms.
- Flongle: consumables at \$90, estimated \$50 to make, approximately 100 active channels (given as 126 elsewhere in the report), implying some 2 million reads in 60 minutes.
- 6 hours: the best time to answer for an optimised clinical MinION workflow, almost all of it sample and library preparation.

- MinION: 512 active channels, about 400 bases per second translocation, needing about 700 base average fragments for 1 million reads an hour and about 70 for 10 million.
- 40 million reads per sample: the depth at which MGS has matched or beaten PCR sensitivity.
- SARS-CoV-2 in nasopharyngeal samples: 0.01% to 0.001% of material, one fragment in 10,000 to one in 100,000; viral load spanning some 8 orders of magnitude, with loads as low as tens of copies per millilitre in more than 1% of cases.
- Modelled depth table: one target in 10 million gives 1 read at 10 million reads (predicted Ct 39.1); one in a million gives 10 reads (Ct 35.8); one in 100,000 gives 100 reads (Ct 32.4); one in 10,000 gives 1,000 reads (Ct 29.1); one in 1,000 gives 10,000 reads (Ct 25.7).
- Published clinical cutoffs: Z of 2 or more, or 10 reads per million, for bacteria, and 1 read per million for viruses and fungi; or non-overlapping reads from 3 or more distinct genomic regions.
- rRNA at 60 to 95% of an RNA sample; FastSelect removes more than 95% in 14 minutes at over \$50 a sample, with a reported 10-fold depth reduction from 10 to 20 million reads down to 1 to 2 million.
- 20 to 30% error rates: what early long-read nanopore sequencing managed during field Ebola surveillance while still detecting pathogens sensitively.

- Nanopore research read lengths: over 1,000 bases on average, with the longest recorded over a million.
- SMRTBell preparation targets: accuracy above 99.9% and 15 kb reads, neither needed for detection; without them, roughly 5 minute preparation, accuracy upward of 80%, read lengths of hundreds of bases; Helicos-derived methods at error rates below 5% with reads over 25 bases.
- PacBio optical: more than 1 million fragments imaged in parallel, 10 nucleotides per second, a run finishing in under a minute at MGS read lengths; instruments currently over \$500,000; consumer cameras at \$300 judged good enough.
- Silver-electrode printed circuit boards: demonstrated beyond doubt only to 16 channels, possibly scalable to about 100.
- Illumina NovaSeq: more than half of global sequencing capacity in bases per year, from some 1,500 instruments; \$1 million per instrument, up to 6 Tb per run, error rate on the order of 0.1%, more than \$5,000 per run.
- ONT in 2021: 55.7% of revenue from 67 PromethION instruments, each up to 12 Tb, costing \$225,000 to \$450,000.
- Modelled global totals: 30,412 instruments, 6,071,676 runs a year, 2,220,290.8 Tb a year.
- Investment scale: at least tens of millions of dollars for a working product, hundreds of millions to build

a fabrication plant, justified at volumes above 1 million units a month; 5 to 10 years for a viable platform, under 5 years for an advance market commitment target, a 10-year frame for ubiquity.

What the author is unsure about

They are careful about the boundaries of their own claims. The depth conclusions apply to viral pathogens in respiratory samples and explicitly do not generalise to blood, where low-abundance pathogens need one or two orders of magnitude more depth. The sensitivity numbers come from their own simulations and a literature review provided as supplementary material, and they assume no host depletion on purpose, which they say makes the targets conservative; they also note there are not yet sufficient studies of FastSelect's effectiveness across sample types, and that the tenfold dilution result is a laboratory report. They say clinical significance cutoffs will have to be established pathogen by pathogen as data on normal abundance accumulates, so the judgement step is unresolved. On technology they will not pick a winner: they expect solutions from nanopore or single-molecule optical in 5 to 10 years but cannot confidently rule out a colony-based method or something new such as solid-state nanopores, they say it is difficult to predict which cost-reduction approach will work, and they flag technical uncertainty about scaling printed-circuit

nanopore arrays past 16 channels. They infer ONT's manufacturing constraints from patents, which they label an inference. They also admit it is unclear whether any of this would look attractive to ONT itself, and that the market analysis is informed largely by the authors' own experience of the field and conversations with experts. The internal inconsistency in the Flongle's channel count, about 100 in one place and 126 in another, is in the source.

For the discussion

- The reframing worth bringing is that detection and research want opposite things: research pays for accuracy and read length, detection needs neither, and this report's claim is that 25-base reads at 95% accuracy would do the job if you had enough of them fast enough.
- Worth raising as a tension: MGS sensitivity depends on relative abundance, so the pathogen competes with the patient's own RNA for space in the run, which means the cheapest route to sensitivity may be a 14 minute depletion step rather than a better sequencer.
- A question worth asking: the report deliberately leaves out wastewater and sentinel populations to focus on the clinic, so does putting a cheap sequencer in every clinic replace the centralised

metagenomic systems in the day's other readings, or feed them.

3.4-4 • CORE • COURSE TIME 5 MIN

Would This Have Flagged COVID?

Jeff Kaufman • 2026

[Read the original paper](#) → • the course assigns: whole piece

Yes, both of SecureBio's detection algorithms would have flagged SARS-CoV-2, but one of them would have needed a hundred times more sequencing to do it, and the reason is the single most important thing about a detection algorithm: what fraction of a pathogen's reads it can learn anything from.

Average summary

Kaufman says people always ask this question when he explains the algorithms used on untargeted metagenomic sequencing data, and he likes it, both because a system should at least have caught the last pandemic and because answering it exposes a key determinant of sensitivity.

He splits an alert into two functions. The **physical** one: actually sequence a read that came from the pathogen. The **algorithmic** one: recognise that the read indicates a problem. He notes the physical side can fail in several ways, if nobody shedding is

contributing to your samples, if your methods are not sensitive to that pathogen, if your libraries lack complexity, or if you are not sequencing deeply enough, and then sets those aside to focus on the algorithms.

The crucial contrast is abundance. When researchers in China first identified the SARS-CoV-2 genome, they had untargeted metagenomic data from bronchoalveolar lavage fluid taken from a single very sick patient, and **22% of the reads were SARS-CoV-2**. In wastewater, by contrast, SecureBio found that **fewer than 0.00002% of reads matched SARS-CoV-2 when 1% of people were infected**. So where the original researchers could assemble a whole genome, a biosurveillance system has to make its first call from individual reads.

The algorithm that would have been most sensitive is **Clades of Concern (CoC)**, which compares reads against known pathogens. SARS-CoV-2 was not in any database yet, so the nearest available genome would have been bat SARS coronavirus ZC45, uploaded in early 2018. The two align well except for the first half of the spike protein gene, positions 21,697 to 23,074, which contains the receptor binding domain. That novel stretch is 1,378 base pairs of a 29,903 base pair genome, about **5%**. So CoC would have flagged reads from the other 95% as suspiciously close to some kind of SARS-like pathogen, and Kaufman says that

with so much of the genome informative, we can confidently say it would have flagged COVID.

Without CoC, **Chimera Detection (CD)** would probably have caught it. CD flags reads where only part matches a known pathogen, so it would have flagged reads crossing the two discontinuities. But the arithmetic is harsh: SecureBio's median read is 170 base pairs after adapter trimming, and you need about 30 base pairs either side of a junction to tell that something is unexpected, so only about **1%** of reads are informative.

That gap, 95% of reads informative to one algorithm and about 1% to the other, is his real subject. He applies it to two hypothetical ways of finding engineered viruses: looking for reads coding for proteins a model predicts would bind mammalian host cell receptors, or looking for reads with no similarity to anything normally seen in that sample type. The first is nearly blind by design, because most of a viral genome codes for internal machinery rather than surface proteins. The second can use most of the genome, so it detects far earlier, which he says is a strong reason to prioritise it.

Why it's assigned

Option 2's third guiding question asks whether a Chinese version of today's SecureBio Detection system would have flagged COVID before Chinese

clinicians did. This post answers the algorithmic half of that directly, and it is the piece that connects the metagenomic machinery to the detect layer's actual question. It also gives you the number that makes the whole unit's economics make sense: one sick patient's lung fluid was 22% virus, while city wastewater at 1% population infection is under 0.00002%, which is why sensitivity in biosurveillance is bought in reads.

In depth

Two things have to happen for an alert

Kaufman frames a system like **CASPER**, SecureBio's wastewater programme, as needing two functions to produce an alert. Physically, it has to sequence at least one read that came from the pathogen. Algorithmically, it has to recognise that read as a problem.

He is explicit that the physical side is a significant concern, and lists four ways it fails: nobody who is sick is shedding into your samples; your sequencing methods are not sensitive enough to this particular pathogen; your libraries do not have enough complexity; or you are not sequencing deeply enough. He says they have put substantial work into those questions, and then sets them aside, because this post is about the algorithms.

The abundance problem, in two numbers

The original identification of SARS-CoV-2 used untargeted metagenomic sequencing, which is why this question is answerable at all. But that data came from bronchoalveolar lavage fluid, meaning fluid washed out of the lungs, from a single very sick patient, and 22% of the reads in it were SARS-CoV-2.

SecureBio's own measurement from wastewater is the other end of the scale: fewer than 0.00002% of reads matched SARS-CoV-2 at a point when 1% of the population was infected.

The consequence is methodological, not just a matter of degree. With 22% of reads from the virus you can assemble the genome and look at the whole thing. At 0.00002% you are looking at individual reads, one fragment at a time, and your algorithm has to decide whether a single 170-base fragment is worrying.

Clades of Concern, and the 5% blind spot

Clades of Concern compares sequencing reads against known pathogens, and it is the algorithm that would have been most sensitive here. The complication is that SARS-CoV-2 itself was not in any database when it started spreading, so the closest available reference would have been bat SARS coronavirus ZC45 (accession MG772933), uploaded in early 2018.

The figure

The post carries two images, neither with a caption. The first, placed right after the alignment discussion, shows SARS-CoV-2 laid against ZC45 along the genome: the two track each other closely across almost the whole length, and the one conspicuous break is the first half of the spike protein gene, positions 21,697 to 23,074, which contains the receptor binding domain, the part of the spike that latches onto a human cell. What to notice is the proportion: that novel stretch is 1,378 base pairs out of a 29,903 base pair genome, about 5% of it, so the picture is 95% familiar and 5% new. The second image, following the Chimera Detection paragraph, is a comparison of what the two algorithms can use, and the text's own figures are the point: 95% of the genome informative to Clades of Concern, about 1% to Chimera Detection.

So Clades of Concern would have flagged reads from the remaining 95% of the genome as suspiciously close to some kind of SARS-like pathogen. It would not have been able to identify reads falling entirely inside the novel region, because there was nothing to compare them to. Kaufman's verdict: even with that blind spot, so much of the genome is informative that we can confidently say this would have flagged COVID.

Chimera Detection, and why it needs a hundred times more data

If Clades of Concern were unavailable, Chimera Detection would likely have caught it instead. CD flags reads where only a portion matches a known pathogen, which is a signature of something stitched together. Because the alignment against ZC45 breaks at positions 21,697 and 23,074, CD would likely have flagged reads crossing either of those points.

The sensitivity cost is arithmetic. SecureBio's median read is 170 base pairs after adapter trimming, and you probably need at least 30 base pairs on either side of a discontinuity to reliably tell that something is unexpected. So only reads straddling a junction with enough context on both sides are informative, roughly 1% of reads.

He shows the calculation in a footnote and labels it rough rather than a simulation, noting that real reads vary in length and that sequencing coverage is not uniform along a genome. With 170-base reads and a 29,903-base genome there are 29,734 possible read positions. Requiring 30 bases of context on each side leaves 111 positions per junction that contain it with enough context, and with two junctions that is 222 of 29,734 possible reads, or 0.7%.

He also adds an important footnote about what CD's hit would and would not mean: he makes no claim about SARS-CoV-2's origin, because although CD was designed to flag engineered pathogens, the changes it reacts to could come from natural recombination,

strongly divergent selection pressure, or engineering, and the algorithm cannot tell the difference.

The general lesson: what fraction of reads are useful

The 95% versus 1% contrast is what Kaufman actually wants to teach, and he calls it one of the most important ways detection algorithms vary.

His worked example is two hypothetical approaches to finding viruses engineered from scratch. Method A: identify reads coding for proteins that a model predicts would bind to mammalian host cell receptors. Method B: identify reads with no sequence similarity to anything normally observed in that sample type, then assess whether they look potentially viral.

He is candid that neither is easy to implement, let alone to run reliably inside a cost-effective system processing hundreds of billions of reads weekly with acceptably few false positives. But their theoretical sensitivity differs dramatically. Method A can only learn from a small fraction of a pathogen's reads, because most of a viral genome codes for internal machinery rather than the surface proteins that mediate entry into a cell. Method B can learn from most of the sequence, which allows detection far earlier, and he says that is a strong reason to prioritise B.

His closing answer to the title question: both algorithms would likely have flagged COVID, but CD would have needed two orders of magnitude more sequencing reads. That difference, he says, is a critical consideration both in modelling whether a real-world deployment is cost-effective and in deciding which algorithms to develop.

Words to know

- **Untargeted metagenomic sequencing.** Reading whatever genetic material is in a sample without deciding in advance what to look for. It matters because it is the only approach that can find something not yet in any database.
- **Read.** One fragment of sequence the machine produces, here a median of 170 base pairs after trimming. It matters because at low abundance a system's judgement has to be made one read at a time.
- **Clades of Concern (CoC).** SecureBio's algorithm comparing reads to known pathogens and their relatives. It matters because it can use almost any part of a genome that resembles something known.
- **Chimera Detection (CD).** SecureBio's algorithm for reads where only part matches a known pathogen. It matters because it finds joins rather than resemblance, which makes it powerful but data-hungry.

- **Receptor binding domain.** The part of a virus's spike protein that attaches to a human cell. It matters here because it sits inside the one stretch of SARS-CoV-2 that had no close relative in the databases.
- **Discontinuity, or junction.** The point in a read where it stops matching the reference. It matters because that boundary is the entire signal Chimera Detection works from.
- **Adapter trimming.** Cutting the synthetic sequencing adapters off a read before analysis. It matters only because it sets the usable read length, 170 bases here.
- **Library complexity.** How many genuinely different molecules are in the prepared sample, rather than copies of the same few. It matters because low complexity means deeper sequencing adds no new information.

Numbers in this reading

- 22%: the share of reads that were SARS-CoV-2 in the original Chinese metagenomic data, from bronchoalveolar lavage fluid from a single very sick patient.
- Fewer than 0.00002%: the share of wastewater reads matching SARS-CoV-2 when 1% of people were infected.
- 29,903 base pairs: the SARS-CoV-2 genome. 1,378 base pairs: the novel region at positions 21,697 to

23,074, about 5% of the genome.

- 95%: the share of the genome informative to Clades of Concern. About 1%: the share informative to Chimera Detection.
- 170 base pairs: the median read length after adapter trimming. 30 base pairs: the context needed either side of a junction.
- The footnote's arithmetic: 29,734 possible read positions, 111 per junction with enough context, 222 of 29,734 across two junctions, or 0.7%.
- Two orders of magnitude: how much more sequencing Chimera Detection would have needed than Clades of Concern.
- Bat SARS coronavirus ZC45 (MG772933), uploaded in early 2018: the closest reference genome available at the time.
- Hundreds of billions of reads weekly: the scale at which any of these algorithms has to run with acceptably few false positives.

What the author is unsure about

He hedges where the evidence is thin and says so. Clades of Concern "would have" flagged it, which he states confidently; Chimera Detection "would likely" have flagged it. The 1% informative-read figure comes from a calculation he explicitly labels rough rather than a simulation, noting that real reads vary in length and that coverage is not uniform along a genome. He sets the whole physical side of sensitivity aside rather

than resolving it, listing four ways it can fail. He says neither of the hypothetical engineered-virus methods is easy to implement or to run reliably at scale with few false positives. And he is careful to disclaim the inference people would be tempted to draw: a Chimera Detection hit says nothing about origin, because natural recombination, strong selection pressure and engineering all look the same to the algorithm.

For the discussion

- The two numbers to bring are 22% and 0.000002%: the same technology that identified SARS-CoV-2 from one patient's lung fluid is looking for the same virus a million times more thinly diluted in a city's sewage, which is the whole difference between diagnosis and surveillance.
- Worth raising: the algorithm that needs no prior knowledge of the pathogen is the one that needs a hundred times more sequencing, so generality and sensitivity trade against each other in detection just as they do in testing.
- A question worth asking: this answers whether the algorithm would have recognised COVID, not whether a sample containing it would ever have reached the sequencer, so what would the sampling side have had to look like in Wuhan in December 2019.

The Triple Rapid Framework for Pandemic Diagnostics

Janika Schmitt, Jake Swett · 2024

[Read the original paper →](#) · the course assigns: whole piece
(the course says skim it, summarized in full here)

Stop treating months-long test development as a fact of life: the authors propose a standard of 10,000 at-home tests deployed within 10 days of an outbreak being detected, each giving an answer in under a minute, and set out the five US policy changes that would make it possible.

Average summary

The course says to skim this and then pick a diagnostic technology from the optional resources. It is summarized in full here.

The three rapids, which is the thing to remember: - **Rapid reconfiguration.** The test platform can be adapted to a new pathogen or variant without rebuilding it, with updates possible in hours or a few days, and regulators able to approve the updated version quickly. - **Rapid deployment.** The tests are **point-of-person**, meaning a person uses one themselves at home or anywhere else, with simplified supply chains so they can be distributed across very

different settings. - **Rapid results.** An answer in under a minute.

Their target: **10,000 at-home tests within 10 days of outbreak detection, each giving results within 1 minute**, distributed first to high-risk areas so cases can be identified and isolated immediately. They call it ambitious and say it rests on an extensive literature review and interviews with more than 20 experts across government, the private sector, nonprofits and academia.

The case for urgency is built from COVID-19's testing failure. Scientists estimate that by March 2020, US testing had identified less than 1% of all SARS-CoV-2 infections. The first US patient hospitalised with COVID-19 of unknown origin did not initially meet CDC testing criteria, which cost a four-day delay; the CDC's own PCR kits had a design flaw and were likely contaminated; and the FDA did not relax its rules on tests from academic laboratories and private companies until February 29, 2020. By March 17, 2020, the US had tested around 400 people per million while South Korea had tested more than 5,000 per million. Modelling they cite estimates a nationwide antigen testing strategy from June to December 2020 could have saved more than 100,000 lives and added at least \$395 billion to GDP.

They acknowledge progress since: the RADx programme and HHS investments put cheap at-home

tests into production, and between January and March 2022 the Postal Service delivered free tests to more than half of US households. But the first over-the-counter COVID-19 rapid test took eight months to get Emergency Use Authorization, which misses the 2022 National Biodefense Strategy's 90-day goal, and even BARDA's more ambitious target, more than 100 million pathogen-specific tests within 45 days of a public health emergency declaration, is too slow: 45 days after COVID-19 was declared a public health emergency of international concern, New York City closed its schools and states were circulating guidance on who would get a ventilator.

Their five recommendations: research agencies (NIH, BARDA, ARPA-H, DARPA) should fund a broad portfolio including instrument-free rapid PCR, at-home isothermal amplification, multiplexed CRISPR-based tests and better antigen tests; Congress should make RADx permanent inside NIH; ASPR and CMS should create durable demand through pre-purchase agreements and long-term reimbursement; ASPR should run an agile, reconfigurable Strategic National Stockpile; and FDA should create a regulatory path that judges a public health screening test on its public health benefit rather than holding it to the standards of an individual clinical diagnosis.

Why it's assigned

This is the core reading for Option 3, rapid individual testing, and it sets the terms for the optional technologies on the same page: when you look at Cas13 multiplexing, the Visby device or LAMP, this framework is what you measure them against. For the detect layer it covers the third element of the day, finding individual cases to contain spread, and it makes the argument that the alternative to testing is not nothing, it is lockdowns, travel bans and quarantines imposed on people whose infection status nobody knows.

In depth

Why a new framework

They begin by granting that the ideal diagnostic is not one-size-fits-all: which attributes matter most, development speed, affordability, accessibility, depends on the use case and the outbreak. Existing frameworks reflect that. The WHO's **(RE)ASSURED** criteria define the ideal test for resource-limited settings, and epidemiologists have set out principles for public health testing where the goal shifts from diagnosing an individual to improving population health.

Their claim is that none of these addresses **outbreak control** specifically. COVID-19 left most countries flying blind through the critical early days, so public health systems reached for broad interventions:

lockdowns and travel restrictions. The 2024 mpox surge and H5N1 bird flu show the same persistent inability to scale testing fast for emerging and novel pathogens. Their proposed shift in mindset: instead of accepting months-long development, aim to deploy tests within days of detection, which would move public health strategy from reactive and sweeping to proactive and precise.

The record they are arguing from

The time to stop a pathogen spreading is at the outset, when we know the least. Their specific indictment of the US response: by March 2020 testing had identified less than 1% of all SARS-CoV-2 infections. The CDC initially set strict PCR testing criteria because of shortages, and the first US patient hospitalised with COVID-19 of unknown origin, possibly indicating community transmission, failed to meet those criteria, producing a four-day diagnostic delay. The CDC's own PCR kits had a basic design flaw and were likely contaminated. Not until February 29, 2020 did the FDA relax its rules on using tests from academic laboratories and private companies. By March 17, 2020, the US had tested around 400 people per million; South Korea had tested more than 5,000 per million.

They set South Korea as the counterexample: it implemented widespread testing and contact tracing swiftly and flattened its curve at a manageable social

cost. And they cite modelling estimating a hypothetical nationwide antigen testing strategy running June to December 2020 could have saved more than 100,000 lives and increased GDP by at least \$395 billion. The same delays, they say, hindered responses to the 2022 and 2024 mpox outbreaks and the 2024 H5N1 outbreak in US cattle.

Their framing line is worth holding onto: in an era of precision cancer treatment and gene therapies, we should not have to quarantine people whose infection status is unknown.

They credit the progress that has happened. RADx, the Rapid Acceleration of Diagnostics programme, and the HHS Testing and Diagnostics Working Group drove strategic investment into at-home tests and domestic manufacturing, and between January and March 2022 the Postal Service delivered free COVID-19 tests to more than half of US households. But it took eight months for the first over-the-counter rapid test to receive Emergency Use Authorization, short of the 2022 National Biodefense Strategy's goal of rapid point-of-need tests within 90 days of identifying a biological incident. BARDA's own goal is more ambitious, more than 100 million pathogen-specific tests within 45 days of a public health emergency declaration, and the authors use a single fact to argue that even that is too slow: 45 days after COVID-19 was declared a Public Health Emergency of International

Concern, New York City shut down its public school system and state health departments were circulating guidelines on who would receive ventilator care if ventilators ran short.

The three rapids, in their own terms

Rapid reconfiguration. Pathogens change, so the test has to. Developers should build platforms where an update takes a few days without overhauling the whole system, and regulators should be able to approve the updated test swiftly. A footnote gives the mechanism behind this: nucleic-acid amplification tests can be reconfigured for a new pathogen from sequence data alone, whereas antigen tests take longer because designing the antibodies can take weeks to months.

Rapid deployment. In an outbreak, scalable distribution and decentralised use matter as much as accuracy. TRF tests are point-of-person with simplified supply chains, so they work in diverse and resource-limited settings. At-home testing lets people make risk-based decisions privately and reduces inequality in access; over-the-counter tests can link a diagnosis to timely treatment where treatment exists; and they enable repeated testing, which matters for healthcare staff with repeated exposures, for pathogens with incubation periods of several days, and for pathogens that transmit without symptoms.

Rapid results. Under one minute, so action can be taken immediately. Their argument here is behavioural rather than technical and is the most quotable part of the piece: for entry screening and contact tracing, a wait of 15 minutes versus one minute can decide whether screening happens at all, because 15 minutes needs planning in advance, somewhere for people to wait, and may discourage anyone from bothering. A one-minute test makes wide adoption feasible.

They add a peacetime use. TRF tests could also help with endemic disease such as influenza and RSV, where respiratory illness costs an estimated more than \$100 billion a year in lost work and medical costs in the US alone, with reconfiguration done seasonally to track variants and subtypes. That matters for the policy argument, because a test platform that earns its keep every winter does not need to be stockpiled on faith.

Why now

A recurring theme from their expert interviews was that the US is at risk of squandering what it learned about testing during COVID-19: as urgency faded, so did government support, and innovative companies that took taxpayer money during the pandemic now face a combination of demand and regulatory uncertainty they describe as deadly.

Their example is **Lucira**. In February 2023 the FDA approved the first over-the-counter test for both COVID-19 and flu, made by Lucira, on the very day Lucira filed for bankruptcy, citing the prolonged FDA approval process as a contributing factor.

They note the CDC's Traveler-based Genomic Surveillance programme and National Wastewater Surveillance System as promising, while saying gaps persist: H5N1 had infected 442 cattle herds in 14 states with 44 human cases reported, and by November no rapid at-home tests specifically detecting H5N1 were available.

They also flag the FDA's move to tighten regulation of **laboratory-developed tests**. They call the FDA's intention to ensure quality commendable, while reporting that hospital and public health laboratories warn the approach could slow test development in an emergency. The history matters: clinical laboratories have traditionally developed and run their own tests without FDA oversight, regulated instead by CMS through the Clinical Laboratory Improvement Amendments, but in early 2020 the FDA suggested Emergency Use Authorization was needed for SARS-CoV-2 laboratory-developed tests, creating a bottleneck exactly when expansion was crucial.

Where TRF tests sit in the wider picture

They are explicit that these tests are one layer, not the system. An ideal biosurveillance system has individual clinical testing for symptomatic patients to inform care, public health screening of asymptomatic people in settings like hospitals, schools and nursing homes, and surveillance testing to track spread. The tools span lab-based diagnostics such as centralised PCR, point-of-care testing where the patient is being seen, and point-of-person testing for individual use.

Lab-based tests complement TRF tests rather than competing, especially early. South Korea set up hundreds of high-capacity PCR screening facilities quickly, which let it test early and often. They cite the Testing Playbook for Biological Emergencies: the first FDA-authorised pathogen-specific test should reach the US Laboratory Response Network and Public Health Laboratories within 48 hours of determining a potential biological incident, with lab PCR acting as the interim measure until at-home tests arrive.

They also sort tests by how many things they look for.

Targeted tests find one specific pathogen.

Multiplexed tests find several at once, which is valuable for respiratory or gastrointestinal illness where different pathogens look alike. **Pathogen-agnostic** tests aim to detect a wide range including unknown ones, usually through metagenomic sequencing.

Their judgement, which connects this page to Option 2: targeted and multiplexed tests, not pathogen-agnostic ones, are the likely route to the triple-rapid combination. Pathogen-agnostic tests are not in routine clinical use because of high costs and complex workflows, and while government programmes are investing and experts have published a roadmap for a low-cost rapid pathogen-agnostic sequencing device, such a device is projected to be at least several years away. They add that during an emerging outbreak attention usually narrows to one specific pathogen anyway, which is what targeted and multiplexed tests are for.

The five recommendations

One: fund a broad research portfolio. It is premature to pick one detection method or sample type, so invest across antigen tests, nucleic-acid amplification tests (NAAT) and array-based tests, and across saliva, blood and breath samples. The 2022 National Biodefense Strategy and Implementation Plan proposes a diagnostics joint capabilities plan across agencies, and they name NIH, BARDA and its DRIVe division, ARPA-H, and the Department of Defense with its Chemical and Biological Defense Program and DARPA as well-positioned, noting BARDA's own strategic objective to invest in tests that can be rapidly adapted to emerging threats and move testing closer to the patient.

On the technologies: NAAT tests beat current antigen tests on limit of detection and on speed of reconfiguration, and recent advances include instrument-free rapid PCR and isothermal amplification. Array-based tests such as CRISPR-based ones can be massively multiplexed, detecting hundreds or even thousands of pathogens at once. Neither at-home NAAT nor CRISPR tests are in widespread use, and the key challenge besides further research is driving production costs down to compete with traditional antigen tests, which needs both government research support and long-term reimbursement.

Antigen tests will stay in the portfolio on their track record, but developing one currently takes weeks to months and often involves immunising animals such as rabbits with the target antigen to raise the antibodies. Building libraries of broad-spectrum antibodies against pathogens of concern in advance, or finding ways to identify and scale antibodies fast, could cut that. They also back continued work on lateral flow optimisation and on synthetic alternatives to antibodies such as aptamers and oligomeric peptides, which are more robust and quicker to manufacture and could go into lateral flow assays or sensor-coupled tests like breath tests.

Two: make RADx permanent within NIH. Until RADx there was no specialised HHS programme for

diagnostics. RADx-backed companies supplied more than 6.7 billion tests to the US market from late 2020 through the fourth quarter of 2022. RADx de-risks and validates tests and walks companies through optimising performance and winning FDA approval; its Independent Test Assessment Program, run with the FDA, accelerated regulatory review including for over-the-counter and multiplexed tests. Its other value was distribution: with accessible design principles, the RADx-UP programme for underserved populations, and the CDC's Increasing Community Access to Testing programme, tests actually reached communities quickly.

Three: build durable demand. The COVID-19 partnerships with test manufacturers and commercial laboratories came late and were short-term. Public health experts have proposed a National Diagnostics Action Plan with pre-event and routine contracts to reduce market uncertainty; ASPR could fund standing capability in commercial and academic laboratories through pre-purchase agreements, keeping expertise and infrastructure alive between pandemics. CMS should be able to expedite reimbursement for tests against emerging pathogens and raise payment rates, as it did for lab-based testing during COVID-19, and to encourage multiplexed diagnostics they suggest paying 1.5 times the standard rate for a pathogen-specific test. They note philanthropy's role too, citing the Gates Foundation's funding of FIND, the

Foundation for Innovative New Diagnostics, in developing and validating rapid tests for the developing world.

Four: make the stockpile agile and reconfigurable.

Supply chain bottlenecks limited test availability early in COVID-19, but as they note, you cannot stockpile pathogen-specific tests for a pathogen you have not met. An agile Strategic National Stockpile would keep a warm manufacturing base and readiness among industry partners. The stockpile currently manages most inventory itself and pays significantly to replace expiring supplies, so its Vendor-Managed Inventory system, where manufacturers hold stock on site and rotate it, should be expanded. What it holds should include basic components such as the nitrocellulose used in lateral flow tests and swabs, which were a critical bottleneck during COVID-19, overseen by ASPR's new Office of Industrial Base Management and Supply Chain. And because modular, automated manufacturing of NAAT tests and next-generation CRISPR diagnostics now allows kits to be adapted quickly, the stockpile should hold reconfigurable test kits, not only components. A footnote notes other agencies can be tasked with stockpiling, naming HERA in the European Union and Africa CDC.

Five: let the FDA judge public health tests as public health tools. Their evidence: in June 2020 the CDC estimated actual US COVID-19 cases were

ten times higher than confirmed ones, which means that despite PCR's high sensitivity it could detect at best 10% of infections, because of limited accessibility. They quote Rochelle Walensky, whom they describe as then-CDC director, saying in September 2020 that for screening to prevent SARS-CoV-2 outbreaks, saliva-based antigen testing is better than the PCR swab. Yet the FDA authorised the first at-home rapid test only on November 17, 2020, eight months into the pandemic, by which time Slovakia and the UK were doing mass testing with rapid tests. From October 2020 Slovakia tested large parts of its population of 5.5 million with rapid antigen tests and found more than 50,000 positive cases; after two rounds of testing with isolation and quarantine of household members of those testing positive, infection prevalence fell by about 80%.

The structural problem they identify is that the FDA has no specific regulatory framework for public health screening tests, so over-the-counter tests face the same clinical trial and analytic accuracy standards as diagnostics meant for individual clinical care, and during COVID-19 that one-size-fits-all approach made some requirements practically unattainable for over-the-counter tests, delaying authorisation and restricting availability.

Four fixes. Issue internal guidance on how to weigh public health benefit for infectious disease

diagnostics, balancing speed, accessibility and scalability against the risk of false positives and negatives, so such tests are judged differently from individual diagnostics; or have Congress direct the FDA to create a new authorisation pathway for public health surveillance and screening tools. Mirror the May 2024 draft guidance on a Platform Technology Designation Program for drugs by simplifying how a new pathogen is added to an existing over-the-counter platform, which matters because COVID-19 produced the first at-home NAAT tests and those adapt more easily than antigen tests. Rebalance the burden of proof so tests showing efficacy in pre-clinical data can be authorised first and then monitored in real time, with the FDA revoking the authorisation of any test that underperforms once better alternatives appear. And take on public education, because public perception fixates on accuracy alone, when what matters for public health testing is speed, frequency, accessibility and accuracy together; they frame at-home testing as a privacy-preserving option that sidesteps the political tension around vaccines and mask mandates.

Their closing: the economics of prevention favour building these tools before they are needed, and the question is no longer whether better tests are possible but whether there is the collective will to make them.

Words to know

- **Point-of-person test.** A test a person performs on themselves, wherever they are, as opposed to point-of-care (where a clinician sees you) or lab-based. It matters because deployment and willingness to test, not just accuracy, decide whether an outbreak is seen.
- **Rapid reconfiguration.** Adapting an existing test platform to a new pathogen or variant without rebuilding it. It matters because it is the difference between a test arriving in days and arriving in months.
- **Antigen test.** Detects a piece of the pathogen itself, usually a protein, typically on a lateral flow strip. Fast and cheap, less sensitive, and slow to develop because the antibodies have to be made, often by immunising animals.
- **NAAT (nucleic-acid amplification test).** Any test that copies the pathogen's genetic material until it is detectable; PCR is the familiar one. It matters because a NAAT can be redesigned from sequence data alone, which is what makes rapid reconfiguration possible.
- **Isothermal amplification.** Amplifying genetic material at one constant temperature, so no thermal cycler is needed. It matters because it is what lets a NAAT work at home.
- **Multiplexed, targeted, pathogen-agnostic.** Tests that look for several named things at once, one

named thing, or anything at all. It matters because the authors conclude the first two, not the third, are the realistic route to a one-minute at-home test.

- **Emergency Use Authorization (EUA).** The FDA mechanism for allowing an unapproved test or product during a declared emergency. It matters because the eight months it took for the first over-the-counter COVID-19 test is their central regulatory complaint.
- **Laboratory-developed test (LDT).** A test a clinical laboratory designs and runs itself, historically overseen by CMS under CLIA rather than by the FDA. It matters because the FDA's assertion of authority over SARS-CoV-2 LDTs in early 2020 created a bottleneck.
- **Lateral flow assay.** The strip format of a home test, where liquid wicks along nitrocellulose past a capture line. It matters here because nitrocellulose is one of the raw materials they want stockpiled.
- **Aptamers and oligomeric peptides.** Synthetic molecules that bind a target the way an antibody does, but are more robust and faster to manufacture. It matters because they could remove the animal-immunisation step that makes antigen tests slow to develop.
- **Strategic National Stockpile, and Vendor-Managed Inventory.** The US emergency stockpile, and the arrangement where manufacturers hold and rotate stock on their own premises. It matters

because you cannot stockpile a test for an unknown pathogen, so what you stockpile is components and capacity.

- **Warm manufacturing base.** Production lines and skills kept running between emergencies so they can scale fast. It matters because it is the alternative to stockpiling finished product.

Numbers in this reading

- The target: 10,000 at-home point-of-person tests within 10 days of outbreak detection, each giving results within 1 minute.
- More than 20 experts interviewed across government, the private sector, nonprofits and academia, alongside an extensive literature review.
- Less than 1%: the share of all US SARS-CoV-2 infections that testing had identified by March 2020.
- Four days: the diagnostic delay for the first US patient hospitalised with COVID-19 of unknown origin, who did not initially meet CDC testing criteria.
- February 29, 2020: when the FDA relaxed its rules on tests from academic laboratories and private companies.
- March 17, 2020: the US had tested around 400 people per million, South Korea more than 5,000 per million.

- More than 100,000 lives saved and at least \$395 billion added to GDP: the modelled effect of a nationwide antigen testing strategy from June to December 2020.
- More than half of US households received free COVID-19 tests by post between January and March 2022.
- Eight months: how long the first over-the-counter COVID-19 rapid test took to receive Emergency Use Authorization, authorised November 17, 2020.
- 90 days: the 2022 National Biodefense Strategy's goal for developing rapid point-of-need tests. 45 days: BARDA's goal for deploying more than 100 million pathogen-specific tests, and also how long after the international public health emergency declaration New York City shut its school system.
- 15 minutes versus one minute: the wait that decides whether entry screening gets implemented at all.
- More than \$100 billion a year: the estimated US cost of respiratory disease in lost work and medical costs.
- February 2023: the FDA approved Lucira's first over-the-counter COVID-19 and flu test on the day the company filed for bankruptcy.
- H5N1: 442 cattle herds infected across 14 states, 44 human cases reported, and no rapid at-home H5N1 tests available by November.

- 48 hours: the Testing Playbook's target for getting the first authorised pathogen-specific test to the Laboratory Response Network and Public Health Laboratories.
- More than 6.7 billion tests supplied to the US market by RADx-backed companies from late 2020 through the fourth quarter of 2022.
- 1.5 times the standard pathogen-specific rate: their suggested CMS payment for multiplexed tests.
- Ten times: how much higher the CDC estimated actual US cases were than confirmed cases in June 2020, meaning PCR detected at best 10% of infections.
- Slovakia: a population of 5.5 million, mass rapid antigen testing from October 2020, more than 50,000 positive cases found, and prevalence down about 80% after two rounds with isolation and household quarantine.
- Two to five years: their timeframe for achieving TRF tests. At least several years: how far off a low-cost rapid pathogen-agnostic sequencing device is projected to be.

What the author is unsure about

They call the goal ambitious themselves and rest it on a literature review and expert interviews rather than a demonstrated prototype, and they note explicitly that their conclusions do not necessarily imply endorsement or unanimous consensus from the

people they interviewed. Much of the piece is recommendation rather than finding: the two-to-five-year timeframe, the policy package and the claim that the goal is within reach given resources and attention are all judgements. They say it is premature to narrow to one detection method or sample type, which is an admission that nobody knows which technology wins. They acknowledge at-home NAAT and CRISPR tests are not in widespread use and that cost is the open problem. They present the FDA tightening of laboratory-developed test rules as a genuine tension, calling the intention commendable while reporting laboratories' warnings. The lives-saved and GDP figures are modelling estimates, and the share of infections detected is an estimate. One attribution to note: they describe Rochelle Walensky as then-CDC director for a September 2020 statement, and that is how the reading presents it.

For the discussion

- The strongest idea to carry in is that one minute versus fifteen is not a technical detail but the thing that decides whether screening is attempted at all, because 15 minutes requires a plan, a waiting space and someone willing to organise it.
- Worth raising as a tension: the authors want tests judged on public health benefit rather than individual diagnostic accuracy, which means deliberately accepting more false results in

exchange for speed and reach, and that is a trade the public has not been asked about.

3.5-2 • FURTHER

Massively multiplexed nucleic acid detection with Cas13

Cheri M. Ackerman et al. • 2020

[Read the original paper →](#) • the course assigns: abstract and main text (the methods sections, panel design appendices and reference list not summarized)

A platform called CARMEN lets one chip run more than 4,500 tests at once, and the authors use it to tell apart all 169 human viruses that had at least ten published genomes, at a reagent cost per test more than 300 times lower than the same CRISPR test run in plates.

Average summary

The problem the authors start from: most circulating pathogens go undetected. Of the viral species known to infect humans, 576 had been sequenced and 169 had at least ten published genomes by October 2018, and only **39 had FDA-approved diagnostics**.

Sequencing can find anything but is expensive and logistically heavy; fast cheap tests, whether CRISPR-based, antigen, PCR or LAMP, each look for only one or a few targets per reaction. They want something

that is both highly multiplexed and easy to scale across hundreds of samples.

Their answer is **CARMEN**, Combinatorial Arrayed Reactions for Multiplexed Evaluation of Nucleic acids, combined with **Cas13** detection. The trick is physical self-organisation. Each amplified sample, and each Cas13 detection mix containing a guide **crRNA** and a reporter, is tagged with its own fluorescent colour code and emulsified into 1-nanolitre droplets. All the droplets are pooled and loaded into a microwell array chip in one pipetting step. Each well takes two droplets at random, so the array spontaneously forms every pairwise combination of samples and guides, in replicate, with no robot and no further handling. A microscope reads the colour codes to find out what is in each well, an electric field merges the pairs to start all reactions at once, and the microscope then watches the fluorescence as Cas13, once it recognises its target, cuts the reporter.

The performance claims: attomolar sensitivity on Zika, matching SHERLOCK and PCR-based assays; **1,050 colour codes** built from ratios of four commercial fluorophores, with 99.5% of droplets correctly classified after filtering that keeps 94% of them; 1,000 tests on a standard chip and more than **4,500 statistically replicated tests** on their massive-capacity chip; and a reagent cost per test more than **300-fold** lower than standard multiwell-plate

SHERLOCK, with fewer pipetting steps and faster turnaround.

Then four demonstrations. The **human virus panel**: 169 crRNAs against synthetic sequences for every species, 30,912 tests across 8 chips, with 157 of 167 crRNAs (94%) selective for their targets in the second round and a median area under the curve of 0.997.

SARS-CoV-2: a new test was folded into the coronavirus panel during manuscript review, and one chip can test more than 400 samples against that panel. **Clinical samples**: 58 plasma, serum, throat and nasal swab samples treated as unknowns, with 11,236 of 11,268 interpretable tests (99.7%) concordant with next-generation sequencing.

Influenza A subtyping: all H1 to H16 and N1 to N9 subtypes identified, 100% concordance with CDC testing on 20 swabs from the 2018-2019 season. **HIV**: six drug-resistance mutations found in parallel in 22 patient samples, plus a panel for 21 integrase mutations.

Why it's assigned

This is the paper the course page links when it asks you to imagine a cheap rapid test that checks for the hundred most common diseases at once, and then sends anything that tests negative for everything to a central laboratory for metagenomic sequencing. It is the strongest published evidence that massive multiplexing is real rather than hypothetical, and the

Triple Rapid Framework reading names multiplexed CRISPR-based tests as one of the technologies worth funding. Read it as the upper bound of what targeted testing can do, and notice that it is still a laboratory platform with microscopes and emulsions, not something you use at home.

Words to know

- **Multiplexed.** Testing for many targets in one go. It matters because everything in this paper is about trading single-target simplicity for scale.
- **CRISPR-Cas13.** An enzyme that, guided to a matching RNA sequence, starts cutting RNA around it. It matters because that collateral cutting is the signal: the enzyme shreds a reporter molecule and the well lights up.
- **crRNA (CRISPR RNA).** The short guide sequence that tells Cas13 what to look for. It matters because adding a new pathogen means designing a new crRNA, not rebuilding the platform, which is what let them add SARS-CoV-2 mid-review.
- **SHERLOCK.** The earlier Cas13-based detection method this builds on, run conventionally in plates. It matters as the cost and sensitivity benchmark.
- **Microwell array, and emulsion droplets.** A chip of tiny wells, and samples suspended as nanolitre droplets in oil. It matters because random pairing of two droplets per well is what generates every combination without a liquid-handling robot.

- **Colour code.** A ratio of fluorescent dyes used as an optical barcode so the microscope can tell which droplet is which. It matters because the number of codes sets the number of samples and guides one chip can hold.
- **Amplification (PCR or RPA).** Copying the genetic material before detection; recombinase polymerase amplification is an alternative to PCR that works at lower temperature. It matters because CARMEN's inputs are already-amplified samples, which is a step at-home testing would still have to solve.
- **Attomolar.** A concentration of about one billionth of a billionth of a mole per litre, so vanishingly few molecules. It matters because it is the sensitivity claim that puts this level with PCR.
- **Drug-resistance mutation (DRM), and SNP.** A single-letter genetic change that makes a drug stop working, and a single-letter difference generally. It matters because it shows the method can tell apart sequences differing by one base, not just one pathogen from another.
- **Area under the curve (AUC).** A measure of how well a test separates positives from negatives, where 1.0 is perfect. It matters because their median of 0.997 across 167 guides is the headline accuracy figure.

Numbers in this reading

- 576 human-infecting viral species sequenced and 169 with at least 10 published genomes by October 2018, against only 39 with FDA-approved diagnostics.
- More than 4,500 crRNA-target pairs on a single array; 1,000 tests on a standard-capacity chip.
- 1-nanolitre droplets; 1,050 colour codes from 4 fluorophores; 99.5% of droplets correctly classified after filtering that retains 94%.
- More than 300-fold reduction in reagent cost per test compared with standard multiwell-plate SHERLOCK.
- Attomolar sensitivity on Zika sequences.
- The human virus panel: 169 crRNAs, 184 total PCR products including controls, 30,912 tests across 8 chips, two rounds of testing with designs improved for 11 species (6.5%), 97.2% concordance between rounds for unchanged designs, and in round two 157 of 167 crRNAs (94%) selective with signal above a threshold of 6 standard deviations over background, median AUC 0.997.
- More than 400 samples testable in parallel against the coronavirus panel on one chip.
- Clinical evaluation: 58 plasma, serum, throat and nasal swab samples, amplified with all 15 primer pools and pooled in sets of three into five metapools per patient; next-generation sequencing at more than 2 million reads per sample; 11,236 of 11,268 interpretable tests concordant, 99.7%.

- Influenza A: subtypes H1 to H16 and N1 to N9 identified, covering 144 possible combinations; 32 of 35 (91.4%) synthetic sequences spanning more than 90% of N1 to N9 diversity identified; 20 throat and nasal swabs from the 2018-2019 season subtyped with 100% concordance with CDC reverse transcription quantitative PCR and in-house sequencing.
- HIV: six reverse transcriptase drug-resistance mutations detected in parallel across 22 patient samples, at 90% concordance with Sanger sequencing and 86% with next-generation sequencing, rising to 93% and 93% when filtering to sequences with up to three mismatches against the design; plus a panel detecting 21 integrase mutations across 9 composite synthetic targets.

What the author is unsure about

They are careful about the limits of a targeted method. Sequence heterogeneity at the target site is a challenge every targeted nucleic acid test faces, and their answer is to add more guides rather than to claim it is solved. Their HIV results show the problem concretely: the assay was designed against HIV subtype B but tested on samples from patients with subtype G, and concordance improved only after filtering out sequences with more than three mismatches against the design. In the clinical samples, both their method and sequencing found

viruses nobody knew were in the samples, and the authors say plainly that results from this panel will have to be integrated with patient symptoms and medical expertise to be used well in clinical settings, and that adoption of such panels alongside clinical care will require careful contextualised interpretation by experts. The crRNA design tool they rely on, ADAPT, is cited as a manuscript in preparation. Their vision of region- and outbreak-specific panels testing thousands of samples from selected populations, including animal vectors and reservoirs, is explicitly a future they imagine rather than something demonstrated here.

For the discussion

- The number worth bringing is 169 viruses with sequenced genomes against 39 with approved diagnostics, because it says the bottleneck is not knowing what the pathogens are, it is having a test for them.
- Worth raising: this platform scales by making every reaction tiny, and the 300-fold cost drop comes from miniaturisation rather than from better chemistry, which is the same argument the sequencing roadmap makes about cheap shallow devices.
- A question worth asking: CARMEN is massively multiplexed but still needs amplified samples, droplet emulsions and a fluorescence microscope,

so what would have to change for anything like this to be the one-minute at-home test the Triple Rapid Framework asks for.

3.5-3 • FURTHER

Evaluation of the Visby medical COVID-19 point of care nucleic acid amplification test

Brooke M. Katzman et al. • 2021

[Read the original paper →](#) • the course assigns: whole paper

A Mayo Clinic team ran 100 leftover patient swabs through a disposable, single-use PCR device and found it agreed with central laboratory PCR on 98 of them, which makes it one of the few point-of-care tests with laboratory-grade sensitivity.

Average summary

This is a small, careful validation study, and it is useful here because it shows what a real rapid molecular test can and cannot do.

The device is the **Visby Medical COVID-19 point of care test**, a single-use, fully disposable RT-PCR assay that was granted emergency use authorization for point-of-care use on **February 8, 2021**. At the time of writing it was one of eight rapid point-of-care nucleic acid amplification tests the FDA had authorised; the

authors list all eight, including the Lucira all-in-one kit, Abbott ID NOW, Cepheid Xpert Xpress and the Cue test. What makes it unusual is that it is a genuine PCR test in a throwaway plastic device: the sample mixes with freeze-dried reagents containing primers for the virus's nucleocapsid region plus an internal control, the device thermocycles to copy any virus present, and the result appears as a **colour change** rather than a fluorescence reading, produced by an enzyme reaction along a flow channel. Results come in approximately 30 minutes.

The study design was deliberately modest. They took 100 residual nasopharyngeal swab specimens from Mayo Clinic patients, 30 previously positive and 70 previously negative, kept refrigerated for up to 3 days, and ran each on both the Visby device and a laboratory reference PCR method, either Mayo's own authorised laboratory-developed test or Roche's cobas platform.

The headline results: the Visby test correctly identified **29 of 30 positives and 69 of 70 negatives**, giving overall agreement of **98.0%**, positive percent agreement 96.7%, negative percent agreement 98.6%, and a Cohen's kappa of 0.95, which falls in the "almost perfect agreement" band. The one missed positive had a reference value of Cp 33, indicating a low viral load, though two other samples at Cp 35

were caught. The one false positive has no clear explanation.

Two further findings matter for deployment. **Invalid results:** only 2.0% of samples were invalid on first run, and both matched the reference on repeat; the authors contrast that with reported invalid rates of 8.6% for the Cue test and 8.6% for Abbott ID NOW.

Sensitivity: the manufacturer claims a limit of detection of 1112 genomic copies per millilitre, and in an abbreviated experiment with heat-inactivated virus the device detected everything down to **500 copies/mL**. For comparison, they cite a published limit of detection of 20,000 copies/mL for Abbott ID NOW, and note another lateral flow rapid test detected only 66.7% of samples in the Cp 30 to 35 range.

Their conclusion: high concordance with central laboratory PCR, a low invalid rate, better analytical sensitivity than some other authorised point-of-care devices, and portability and ease of use that suit settings staffed by people who are not laboratory scientists.

Why it's assigned

Page 3.5 asks you to pick one diagnostic technology from the optional resources and ask why it is better than PCR or a rapid antigen test, how it would need to be deployed early in a pandemic, and why it has not

replaced what we already use. This paper is the evidence base for that argument about one real device, and it is the only reading on the page that measures a product rather than proposing one. It also puts a number on the trade the Triple Rapid Framework is making: this test has close to laboratory sensitivity and takes about 30 minutes, against the framework's target of under a minute.

Words to know

- **Point of care (POC).** Testing done where the patient is being seen, rather than in a central laboratory. It matters because this is one step short of the at-home point-of-person testing the day's core reading argues for.
- **NAAT, and RT-qPCR.** A nucleic acid amplification test copies the pathogen's genetic material until it is detectable; reverse transcription quantitative PCR is the laboratory standard version for an RNA virus. It matters because this device is a NAAT, not an antigen test, which is why its sensitivity is comparable to a laboratory's.
- **Limit of detection (LoD).** The lowest concentration a test reliably finds, here in genomic copies per millilitre. It matters because it is the single number that separates this device from weaker rapid tests.
- **Cp or Ct value.** The crossing point, or cycle threshold: how many amplification cycles the

laboratory machine needed before it saw a signal, so a higher number means less virus. It matters because the one missed positive sat at Cp 33, near the limit of positivity.

- **Positive and negative percent agreement.** How often the new test agrees with the reference on samples the reference called positive, and on those it called negative. It matters because with no perfect truth standard, agreement is used instead of sensitivity and specificity.
- **Cohen's kappa.** A measure of agreement that corrects for agreement by chance, where 0.81 to 1.00 is called almost perfect on the scale they use. It matters because 0.95 is the strength of this study's main claim.
- **Invalid result.** A run the device itself rejects, requiring a repeat. It matters for deployment, because every invalid result is another swab, another half hour and another device.
- **Nasopharyngeal swab, and viral transport media.** A swab taken deep behind the nose, placed in a liquid that preserves the virus in transit. It matters because these were leftover laboratory specimens in transport fluid, not swabs taken fresh for this device.
- **Lyophilized reagents.** Reagents freeze-dried into the device so they need no cold chain and activate when the sample is added. It matters because it is what makes a disposable PCR device possible.

Numbers in this reading

- 100 residual nasopharyngeal specimens: 30 previously positive, 70 previously negative, refrigerated at 2 to 8 °C for up to 3 days, with all testing completed within 72 hours of collection.
- 29 of 30 positives and 69 of 70 negatives correctly identified.
- 98.0% overall agreement, 96.7% positive percent agreement, 98.6% negative percent agreement, Cohen's kappa 0.95 with a 95% confidence interval of 0.89 to 1.00.
- Cp 33: the reference value of the one sample Visby missed; Cp 35: the value of two samples it did detect; Cp 40 or below: how positive was defined on both reference methods.
- Approximately 30 minutes for a result.
- 2.0% invalid on first run, against reported invalid rates of 8.6% for the Cue COVID-19 test, 8.6% for Abbott ID NOW, and 1.1% for Visby in another evaluation.
- Limit of detection: 1112 genomic copies/mL claimed by the manufacturer at a 95% detection rate, and 75% at 500 copies/mL; their dilutions at 500, 1000 and 10,000 copies/mL were all positive, in triplicate at the two lower levels and duplicate at the highest.
- Comparisons cited: 100 copies/mL analytical sensitivity for Visby in another study, 20,000

copies/mL for Abbott ID NOW, and only 66.7% detection between Ct 30 and 35 for the Accula lateral flow test.

- Another evaluation of 78 specimens reported 95% sensitivity, 100% specificity and 96.1% overall concordance.
- Eight: the number of rapid point-of-care nucleic acid amplification tests authorised by the FDA at the time of writing.
- February 8, 2021: the date of the device's point-of-care emergency use authorization.

What the author is unsure about

They set out their limitations plainly. It is a small, single-centre study using a random sample of leftover, non-frozen clinical specimens, not a prospective trial with recruited patients. Only nasopharyngeal swabs were used, because nasal and mid-turbinate samples were not available at the time, even though the device is intended for those too. They had no information on whether patients had symptoms, or how long since symptom onset, because of the terms of their review board approval. The limit of detection experiment was significantly abbreviated, and they say a full formal experiment might have verified a lower value. They have no explanation for the single false positive, and note that discrepancies could come from differences in sample processing, buffers or limits of detection between the assays. Drawing on their own earlier

work, they say some rapid nucleic acid tests are more prone to false positives than central laboratory tests, so a 1 to 2% false positive rate was not unexpected and may need to be factored into point-of-care testing protocols. They call for further evaluation with symptomatic and asymptomatic patients tested prospectively on both methods. One internal wording difference to note: the protocol says results take approximately 30 minutes, while the conclusion describes the device as giving an assessment within minutes.

For the discussion

- The point worth making is that this device shows laboratory-grade sensitivity does fit inside a disposable plastic box, so the obstacle to rapid molecular testing is not physics; it is cost, authorisation and demand.
- Worth raising: the authors treat a 1 to 2% false positive rate as expected for rapid molecular tests and something to design testing protocols around, which is a more honest framing than the public conversation about test accuracy usually allows.
- A question worth asking: an 8.6% invalid rate on two other authorised devices means roughly one person in twelve has to swab again, so how much does that alone decide whether screening works in a real clinic or airport.

Recent advances in loop-mediated isothermal amplification (LAMP) for rapid and efficient detection of pathogens

Nupur Garg, Farhan Jalees Ahmad, Sudeshna Kar • 2022

[Read the original paper →](#) • the course assigns: whole review (the three results tables summarized in prose; the reference list not summarized)

LAMP does what PCR does, copying a pathogen's genetic material until you can see it, but at one constant temperature in under an hour, so it runs in a water bath and can be read by eye, and this review catalogues how far that has been pushed.

Average summary

PCR is the gold standard for finding a pathogen's genetic material, and the authors' complaint is practical: it needs an expensive thermocycler to ratchet the temperature up and down, trained staff, and long reaction times, which rules it out in resource-poor and non-laboratory settings.

Loop-mediated isothermal amplification (LAMP), first developed by Notomi and colleagues in 2000, avoids the temperature cycling entirely. It uses a DNA polymerase called **Bst** that can push its way along a strand and displace it, plus primers that recognise six

specific sequences on the target, and the whole reaction runs at a single temperature in simple equipment such as a water bath or heating block, in **less than an hour**. An extra pair of loop primers speeds it up by creating more starting points. Its other advantage over PCR is that it amplifies the target even when non-target sequences are present.

Its weakness is the mirror image of its design. Because LAMP uses many primers, they can pair with each other and amplify with no template at all, producing **false positives**. Much of the technical advance in this review is about controlling that.

The best part of the review for a non-specialist is the catalogue of ways to **read the result**, which is what makes LAMP field-usable. Gel electrophoresis or fluorescent dyes in a laboratory. **Turbidity** you can see with the naked eye, because amplification produces pyrophosphate ions that form an insoluble complex with magnesium and cloud the tube. **Colour change**: calcein turning yellow to green, metal indicators like hydroxy naphthol blue, or pH-sensitive dyes such as phenol red, since the reaction shifts the pH. **Lateral flow dipsticks**, which the authors prefer for point-of-care because they are user-friendly and affordable, where labelled products and gold nanoparticles produce a red stripe. And electrochemical chips, FRET probes, ELISA

combinations, and gold nanoparticles that shift red to blue.

The applications span bacteria, viruses, parasites and fungi. They highlight **neglected tropical diseases**, a group of around 20 diseases concentrated in the poorest regions, where LAMP has been used for early diagnosis, therapy monitoring and contact tracing, and where their table gives sensitivities down to single eggs or femtograms of DNA. **TB-LAMP** is a WHO-recommended kit, completed in under an hour and read under UV light. Their RT-LAMP table lists times and sensitivities for highly virulent viruses, from West Nile in 17 minutes to Ebola in 15 to 20 minutes and MERS-CoV at 4 RNA copies.

Three directions of advance: **multiplex LAMP**, detecting several targets at once, with analysis under 30 minutes, sensitivity of 1 to 10 copies, tolerance for unprocessed saliva or tissue and compatibility with smartphones; **real-time LAMP**, including a hand-held colorimetric device run from a smartphone app and validated on SARS-CoV-2 samples; and **LAMP-on-a-chip**, including microfluidic, paper-based and digital versions, one of which reaches single-molecule detection in nanolitre droplets.

Why it's assigned

Page 3.5 asks you to pick one diagnostic technology and work out why it beats PCR or a rapid antigen test,

how it would have to be deployed early in a pandemic, and why it has not replaced what we use now. LAMP is the best-evidenced answer to the first question, since it is already in WHO-recommended kits and was behind one of the test kits Vietnam built in a month. It is also the method behind the at-home isothermal amplification tests the Triple Rapid Framework names as a technology worth funding, which makes this review the ground truth for that recommendation.

Words to know

- **Isothermal amplification.** Copying genetic material at one constant temperature. It matters because the thermocycler is the expensive, power-hungry box that keeps PCR in laboratories.
- **LAMP.** Loop-mediated isothermal amplification: the most widely used isothermal method, recognising six sequences on the target with multiple primers.
- **RT-LAMP.** The version that starts by copying RNA into DNA with reverse transcriptase, all in one tube at one temperature. It matters because the pathogens in this unit, including influenza and coronaviruses, are RNA viruses.
- **Bst polymerase.** The copying enzyme LAMP uses, which can displace the strand ahead of it instead of needing heat to separate the strands. It matters because strand displacement is what removes the need for temperature cycling.

- **Primer, and primer dimer.** A short piece of DNA that tells the enzyme where to start; a primer dimer is two primers sticking to each other instead. It matters because LAMP's many primers make dimers, and dimers make false positives.
- **Amplicon.** The copied product of the reaction, the thing you are trying to see.
- **Turbidity.** Cloudiness in the tube, here from magnesium pyrophosphate forming as the reaction runs. It matters because it is a readout that needs no instrument at all.
- **Colorimetric readout.** Reading the result as a colour change, whether from a dye bound to the product, a metal ion indicator or a pH-sensitive dye. It matters because it is what makes a test readable by eye in a village clinic, and because the authors concede colour judgement is subjective in field conditions.
- **Lateral flow dipstick (LFD).** The strip format, where labelled product and gold nanoparticles produce a visible line. It matters because the authors consider it the most practical point-of-care readout.
- **Multiplex LAMP (mLAMP).** Using more than one primer set to detect several targets at once, which also improves specificity by requiring more than one gene to match. It matters for telling closely related species apart.

- **FRET, and self-quenching probes.** Fluorescence resonance energy transfer, where one dye hands its energy to a quencher, and probes that only light up once bound to their target. It matters because these are the main fixes for LAMP's false positive problem, giving sequence-specific confirmation rather than just "something amplified".
- **Digital LAMP.** Splitting a sample into many tiny chambers or droplets and reading each one, which allows counting down to single molecules.
- **ASSURED criteria.** The WHO's checklist for a diagnostic in a low-resource setting: affordable, sensitive, specific, user-friendly, rapid and robust, equipment-free and delivered to those who need it. It matters because satisfying it is the authors' closing claim for LAMP.

Numbers in this reading

- 2000: when Notomi and colleagues first developed LAMP. Six target sequences recognised; the reaction completes in less than an hour at a single temperature.
- Around 20: the number of neglected tropical diseases, where the review concentrates its case.
- Sensitivities reported for neglected tropical disease assays include 5 femtograms and 50 femtograms for *Trypanosoma cruzi* strains, 50 parasites per millilitre, 10 to 100 femtograms for *Echinococcus*, a single *Fasciola* egg, a single

whipworm egg, 1 picogram for *Onchocerca*, 0.1 picogram for *Schistosoma haematobium*, 0.02 nanograms for scabies, and reported sensitivities of 72% and 80% for two African sleeping sickness subspecies and 94.4% for *Leishmania*.

- The RT-LAMP table: Zika in 20 minutes at 10 to the minus 5 PFU; Ebola in 15 to 20 minutes at 256 copies; HIV in 60 minutes at 10,000 RNA copies per millilitre; influenza in 60 minutes at 0.1 RNA copies; hepatitis B in 15 minutes at 2.2 femtograms per microlitre; dengue in 15 minutes at 0.8 femtograms per microlitre; HPV in 60 minutes at 10 copies; H1N1 in under 40 minutes at 10 RNA copies; encephalitis in 70 minutes at 50 picograms; MERS-CoV in 60 minutes at 4 RNA copies; West Nile virus in 17 minutes at 0.1 PFU.
- Multiplex LAMP characteristics: analysis under 30 minutes, sensitivity of 1 to 10 copies, a large dynamic range, use of unprocessed saliva or tissue, and smartphone compatibility.
- A multiplex assay distinguishing tuberculosis from non-tuberculous mycobacteria reported 98.97% analytical sensitivity and 100% specificity.
- A multiplex influenza assay gave specimen-to-result in 40 minutes, with 12 minutes of amplification and 10 minutes of specimen preparation, against 90 to 120 minutes for PCR, at single genome copy sensitivity.

- Other multiplex figures: 10,000 copies per microlitre for most colistin-resistance genes and 100,000 for one, 100,000 copies for a SARS-CoV-2 assay, 98.0% sensitivity and 98.1% specificity for foot-and-mouth disease, 100 femtograms for Salmonella and Shigella in milk, 94.62% for influenza A and 97.50% for influenza B, and 100 colony-forming units per microlitre for Acinetobacter baumannii.
- A multiplex lateral flow biosensor for Leptospira reached a detection limit of 3.95×10 to the minus 1 genomic equivalents per millilitre, showing three red lines when positive and two when negative.
- Real-time results: a four-serotype dengue assay 100% reproducible and specific with 95.8% sensitivity; a self-dequenching probe assay for varicella-zoster virus at 96.8% clinical sensitivity and 100% specificity; a FRET assay detecting white spot syndrome virus at 100 copies.
- A microfluidic RT-LAMP chip detected porcine coronaviruses in 40 minutes; a droplet array chip reached a single-molecule detection limit.

What the author is unsure about

This is a review rather than an experiment, so its claims are other people's results, and the authors are open about the method's problems. They state plainly that LAMP suffers from non-specific binding through primer dimers, that using multiple primers raises the

risk of primer-primer hybridisation and template-free amplification, and that this leads to false positives, which they call one of the method's major disadvantages. They say colorimetric readings may be variable and subjective in field settings, which is why they favour dipsticks. They acknowledge limitations in the advanced versions too, naming complex fabrication, cross-contamination and variability of results, and say further work is in progress on those. Their search strategy is declared and narrow: MEDLINE, NCBI and PubMed with a set of LAMP search terms, mostly articles from the previous five years. The closing claim that LAMP satisfies the WHO's ASSURED criteria is their judgement, not a measured result. One inconsistency sits in the source itself: the introduction says LAMP uses three primer pairs, while the methods section describes four primers recognising six sequences plus an additional loop-primer pair.

For the discussion

- The useful point is that LAMP's appeal is not chemistry but infrastructure: taking out the temperature cycling takes out the thermocycler, the laboratory and the technician, which is why it reached neglected tropical diseases long before anything else molecular did.
- Worth raising as a tension: the same multiple primers that let LAMP work at one temperature are

what generate its false positives, so the fixes in this review, probes, quenched primers and multiplexing, all add back some of the cost and complexity that made it attractive.

- A question worth asking: a WHO-recommended LAMP kit for tuberculosis has existed for years and the method can hit single genome copies in 40 minutes, so why is the home test most people own still an antigen strip.