Episode Book

EPISODE #5925

Plus Codes, Geohash, and the Open Addressing Gap

PUBLISHED	13 October 2026 at 02:08 UTC
DURATION	27:27
LISTEN	myweirdprompts.com/episode/plus-codes-open-addressing-gap-38486/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.4
SOURCES	14 cited by the research (listed at the end)
BOOK GENERATED	19 October 2026 at 17:41 UTC

SYNOPSIS

Google's Open Location Code was designed to work completely offline, with no geocoding service and no proprietary database in the middle. Then the consumer interface shipped a shortened form that needs a locality name — and a locality name needs Google. We trace the mechanics of code shortening, the base-20 alphabet built by testing eight billion combinations against words in thirty languages, and why a hundred-metre ambiguity in Zurich is enough to break the whole promise. Then we go looking for the purer version: eleven language ports, a zero-dependency TypeScript library, an offline Garmin watch datafield — and the standalone consumer app that we could not confirm exists.

TRANSCRIPT



Corn

Here's a number that shouldn't still be true. The trick that makes a Plus Code short is the same trick that makes it fragile. Drop four digits off the front, and you've saved six characters at the cost of a dependency you cannot control.



Herman

And that trade is entirely invisible to most people using them.



Corn

Daniel's question picks up where we left the last one. He's been on a hunt, in his words, for the ideal open geographic addressing system. Something that spits out a short, human-readable, globally unique identifier for a precise spot, generated off a phone's GPS, decodable offline by whoever receives it, with no locality name, no proprietary database, no geocoding service standing in the middle.



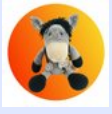
Herman

And the target resolution he names is fourteen metres by fourteen metres, though he says the exact precision is negotiable.



Corn

He also splits it into two questions, and the first one is sharper than it sounds. Has Google's consumer implementation obscured the real potential of Google's own open standard? Because his complaint isn't about the standard. It's about the interface. The shortened code they put in front of everyone.



Herman

And the second question is whether anyone else has built the technically purer version of the idea. Whether the ideal already exists, just badly presented, or whether there's a real gap.



Corn

So let's start with what Google actually shows people.



Herman

The full-length code is the whole point of the design. It's self-contained. Google's own README says a full code can be converted back to a location completely offline, without any data tables to look up and without any online services required.



Corn

Which is what you'd want in a flood.



Herman

Right. And then the same README describes the shortened form. Those can be used regionally, or in combination with a reference location that all users of the short code need to be aware of. And then the sentence that does all the damage. If the reference location is given in the form of a location name, use of a geocoding service might be necessary to recover the original location.



Corn

So the short code isn't wrong. It's just not finished.



Herman

It's incomplete by design. And the README draws the line itself. Geocoding services are not a part of the Open Location Code technology. They're a part of the product built on top of it. The locality dependency is introduced by the interface.



Corn

Which means the standard's defenders are correct and the users are still stuck.



Herman

So the arc for today is, why shortened codes exist at all, what the independent ecosystem actually offers, how the competing systems compare, and what all of it means for the one case where the answer really has to be right.



Corn

The emergency case. Walk me through the shortening, then. Mechanically.



Herman

You drop four or more leading digits from the front of a full code. What's left works on its own within roughly fifty kilometres of the place it belongs to. If you supply a city or a locality, it resolves globally, because the missing prefix can be reconstructed from the reference point. And the short code always keeps the plus sign, so the receiver can see how much is missing and compute the rest.



Corn

The plus is the seam.



Herman

And the precision numbers are worth having. A ten-digit code is one eight-thousandth of a degree, which at the equator is about thirteen point nine metres by thirteen point nine. That's Daniel's fourteen by fourteen. Add a digit and you get roughly three point four metres by two point seven. Add another and you're under a square metre.



Corn

So a person could hand over a code smaller than a doorway.



Herman

A person could. And the reason Google shortened them at all is usability. Nine seven M F plus two seven, Mumbai, India, is a thing you can say out loud and a stranger can write down. The full ten-digit string is not. The locality name is doing the work of memory.



Corn

It's doing the work of memory and borrowing against trust.



Herman

That's the trade. And the standard's own design paper is explicit that this trade was never supposed to be made. The requirements list reads like a manifesto. Complete, it says. The codes must have enough information on their own. Extra information such as street number, locality or country name could be helpful, but should not be required.



Corn

Should not be required. Somebody wrote that sentence and then somebody shipped Mumbai, India.



Herman

Same paper, three more. Functions offline. Easy to implement, and it should not depend on a single provider. Free, and it should not require a licence fee.



Corn

So the standard promised self-contained and the flagship interface leads with dependent.



Herman

And Google's own documentation concedes the failure mode. There's a worked example in the README about Zurich. Although one service may place Zurich close to the Google office, another may move it by a hundred metres or more, and this could be enough to prevent the original code being recovered.



Corn

A hundred metres. That's a different street.



Herman

It's a different building, in a case where the point of the code was to say exactly which one. And the README's advice is telling. Use a smaller neighbourhood feature rather than a large city name.



Corn

Which is good advice and not the one the interface follows.



Herman

The interface optimizes for a tourist typing a place they've heard of.



Corn

The character set is the part I find elegant. Tell me about that, because it's the bit people miss.



Herman

The alphabet is two three four five six seven eight nine C F G H J M P Q R V W X. Base twenty. No vowels, so you can't spell anything. And the confusable letters are gone. No I, no L, no O, no S, no Z, no zero and no one to fight with them. They picked it by generating every possible twenty-character combination, over eight billion of them, and scoring each one against ten thousand words from more than thirty languages. Anything that spelled a real word got thrown out.



Corn

They ran the world's languages through it and kept the set that says nothing.



Herman

Deliberately. Transcribing something out loud over a bad radio, you cannot accidentally read a word that isn't there. There's no word to read.



Corn

And then the imperfections, because there always are.



Herman

The character set is Latin-only, so you get the Cyrillic N that looks exactly like a Latin H, and that ambiguity is real the moment a code crosses a keyboard. There are discontinuities at the poles and at longitude one eighty. Area distortion grows at high latitude. And the scheme can't exactly represent latitude ninety.



Corn

The top of the planet is undefined.



Herman

The top of the planet is undefined, which matters to approximately nobody, and I still want it fixed.



Corn

Daniel also cites a discussion on the Google Maps subreddit from twenty twenty-three about how hard it is to get a full-length code out of Google's own interface.



Herman

And I want to be careful there. We could not verify that thread. Daniel cites it, and I don't doubt he read it, but I haven't independently confirmed the content, so treat that as his claim and not as something I checked.



Corn

Fair. And notice what it claims, though. Not that the full code is unavailable. That it's undiscoverable.



Herman

Which is a different kind of problem. The capability ships. The path to it doesn't.



Corn

So: has Google obscured its own standard? First pass, yes, but not by hiding anything. By leading with the version that needs them.



Herman

That's a fair first pass. Though I'd push on the word obscured.



Corn

Push.



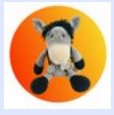
Herman

Shortened codes exist because the full code is hard to communicate by voice. That's not a marketing decision, that's a real constraint. Somebody in a call centre reading nine seven M F plus two seven plus a city gets it right. Somebody reading the full string gets it wrong twice.



Corn

So the shortened form is a concession to humans, not a lock-in.



Herman

It's a concession to humans that happens to require the vendor.



Corn

Both can be true. It's good for the caller and bad for the flood.



Herman

Yes. And here's the thing I keep circling. The dependency isn't the shortening, it's the name. You could shorten a code against another code. Against a known landmark set offline. The standard supports a reference location that isn't a name at all. Google chose the name.



Corn

Because the name is what people have. So now the second question. Who else built the pure version.



Herman

The independent ecosystem is real, and it's broader than people assume. The reference repo is multi-language. Java is about twenty-two percent of it, JavaScript fourteen and a half, Go nine, then PL/pgSQL, VBA, C++, Dart, C, Python, Rust, Ruby. There's even a Visual Basic for Applications implementation, which means somebody is generating Plus Codes inside Excel.



Corn

Of course there is.



Herman

And the third-party ports are the interesting layer. There's a JS and TypeScript library by janne, two to fifteen digit precision, zero dependencies, and it advertises that it matches Google's reference output. There's a Python package on PyPI, version twenty twenty-two point one point three, Apache licensed. There's a pure Kotlin Multiplatform port. And then there's the one I'd actually point a listener at, which is a Garmin Connect IQ datafield. It displays the Plus Code for your current location, and it doesn't use any network.



Corn

A watch. Offline. On a wrist.



Herman

A watch, offline, on a wrist, doing the thing Daniel asked about, and it isn't Google Maps.



Corn

So the libraries are there. Now ask the hard version of his question. An ordinary person with an ordinary Android phone, no Play Services, no Google Maps, no Google APIs. Can they do it?



Herman

This is where I have to be honest with you, and with Daniel, because he asked specifically for actual applications and not demonstrations.



Corn

Go on.



Herman

I could not confirm a standalone, mature, F-Droid-distributed Plus Codes app.



Corn

Say that again plainly.



Herman

There is no Plus Codes app in the F-Droid catalogue that I can point to and say, install this. The libraries exist in eleven languages. The Garmin datafield exists and is offline and shipping. But the thing Daniel actually asked for, a mature consumer application that generates and shares a full-length code with no Google components, I could not confirm exists. That's a genuine gap and I'm not going to paper over it.



Corn

So the algorithm is solved and the app is missing.



Herman

The algorithm is solved and the app is missing. Which is a stranger result than either of us expected going in.



Corn

What about the repo's own health? Because that tells you whether anyone's still home.



Herman

Four thousand three hundred and thirty-five stars, five hundred and twenty-nine forks, seventy contributors, fifty-six open issues. The latest tagged release is one point zero point four, from July twenty nineteen. But the main branch was pushed in March of this year.



Corn

So the standard is frozen and the plumbing is maintained.



Herman

The standard is stable, which is what you want in a standard. The release cadence is slow, which is what you'd expect from something Google treats as finished.



Corn

Now the field. Because if OLC isn't the answer, something else might be, and Daniel wants to know about the systems that tried.



Herman

Start with Geohash, because everybody's used it and nobody's checked it. Thirty-two character set, vowels included, so it can spell things. Not just profanity, though that's the famous one. The one that should worry you is that it can produce almost-words. There's a documented example where a Geohash looks like Z U R one C H. That's Zurich with a one in it.



Corn

Which a tired person reads as Zurich.



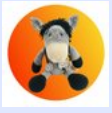
Herman

It's variable length. And it has five discontinuities, at the equator, longitude zero, longitude one eighty, and both poles. Those are coordinates where adjacent squares in the real world sit at opposite ends of the string.



Corn

Adjacent geography, distant codes.



Herman

That's the thing to remember. Two neighbours whose codes look nothing alike.



Corn

Geohash thirty-six.



Herman

Drops the vowels but gains mixed case, so now you have a capital O and a zero and a lower-case L in the same string. It can be one character shorter than a full Plus Code at similar accuracy, and it carries an optional altitude field and a checksum.



Corn

The checksum is the one good idea in the room.



Herman

The checksum is the one good idea in the room, and Plus Codes don't have one. There's an optional check digit in the specification but it isn't in the codes people actually use.



Corn

Keep going. MapCode.



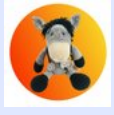
Herman

MapCode produces global codes about the same length as Plus Codes, which is respectable. The territory codes are shorter, but they need a maintained data file, and they're keyed to ISO three one six six country codes. Which works beautifully until a country is disputed, and then the address system is taking a side.



Corn

The map decides the border.



Herman

And one location can carry several codes at once. Schiphol Airport has a global MapCode, and then at least three territory forms. Four valid answers to where is the airport.



Corn

Open Post Code.



Herman

Global codes and country codes, and the two sets are completely dissimilar. Which means a global code typed into a country field will be valid, and will decode somewhere else entirely. It won't fail. It'll succeed, incorrectly.



Corn

That's worse than no answer.



Herman

That's worse than no answer, and it's a design choice rather than a bug.



Corn

Natural Area Code.



Herman

Proprietary, licence required, whitespace significant, and no truncation. Three of those are fixable and the licence isn't.



Corn

Maidenhead, then. Because that's the one that's been quietly right for decades.



Herman

That's the one that's been quietly right for decades. It's area-based, you can truncate it, it lands at similar accuracy and similar length to a Plus Code. Radio operators have used it forever. Its limit is that it's only formally defined to eight characters, so the precision stops where the standard stops.



Corn

And What3Words, which is what most people have actually heard of.



Herman

Three-by-three metre squares, proprietary algorithm, and the part that matters for this conversation. The offline SDK was announced late twenty fourteen, but it doesn't support the OneWords people actually bought. So the thing you licensed isn't the thing that works on the plane.



Corn

Three words is easier to say than a code. That's the whole pitch.



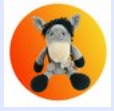
Herman

Three words is easier to say, and it's dramatically easier to say wrongly. Banana rabbit monkey is in Argentina. Banana monkey rabbit is in Russia. Same three words, order reversed, wrong hemisphere.



Corn

And this is where the published research comes in, because Daniel specifically wants the engineering literature and not vibes.



Herman

Rudy Arthur at Exeter, in PLOS ONE, October twenty twenty-three. A critical analysis of the What3Words geocoding algorithm. Two findings, and both are simulation-based. First, most simulated addresses have one or more word triples they could be confused with.



Herman (continued)

Only about a third of triples have no possible confusion at all, under typographic and homophone error modes.



Corn

So two thirds have a twin.



Herman

Two thirds have a twin somewhere. And around a quarter of addresses have more than three confusions, which matters because What3Words' own autosuggest offers three suggestions. Four confusions and the correct answer isn't on the list.



Corn

That's the number that should end the argument for emergency use.



Herman

That's the number that should slow it down, certainly. The second finding is about proximity. Confusable pairs near each other show up far more often than the one-in-two-and-a-half-million figure What3Words quotes. The cause is the linear-congruence assignment, and in a simulated area about fourteen and a half kilometres by four and a half, the peak of confusable pairs sat about ten kilometres apart. Two places you could walk between, whose codes are one letter apart.



Corn

So the errors aren't random. They cluster in exactly the geography where a mistake costs you the response.



Herman

That's the finding, stated plainly.



Corn

And then Arthur praises Plus Codes in the same paper.



Herman

He does, and it's the reason the citation is honest rather than an attack. He writes that it's obvious the Plus Codes P F Q nine plus six P seven, plus six P eight, plus six P nine are close together, while it's far from obvious that washed dot expand dot logs, poems dot joke dot urban, punks dot nails dot villa are neighbours.



Corn

Adjacent codes look adjacent.



Herman

That's the property the whole system is built on, and Arthur confirms it while demolishing the alternative.



Corn

But he also notices Google's own habit.



Herman

He does. He points out that Plus Codes are integrated into Google Maps and made more user-friendly by substituting the top-level information with standard place names and keeping the lower level for fine detail. His example is the Exeter computer science department, which comes out as UK, Exeter, P F Q nine plus six P seven.



Corn

So the researcher who defends the standard also documents the substitution.



Herman

Both things are in the paper. Nothing is being hidden.



Corn

And his recommendation for the lost walker?



Herman

The OS Grid Reference. Open source, widely understood, and critically, printed on most maps.



Corn

Printed on paper.



Herman

And that's the dimension I underweighted going in. Every phone-based identifier has the same failure: the battery dies, the phone goes in the river, and the code goes with it. A grid reference printed on a paper map in a rucksack survives all of that. The identifier and the medium are two separate design problems and we've been treating them as one.



Corn

Arthur also says something about the token set that we should quote, because it's the strongest argument against words.



Herman

Alphanumeric codes, he writes, could be much easier for non-native speakers and are less likely to be affected by accent, since the sender and receiver are only using tokens from a very restricted set.



Corn

Restricted set. That's the whole philosophy of the Plus Code alphabet in one phrase.



Herman

It's the same argument the designers made when they scored eight billion combinations against thirty languages. Arthur arrives at it from error modelling and lands in the same place.



Corn

And his closing caution.



Herman

That until such an evaluation has been carried out, he would caution against the widespread adoption of What3Words, especially in applications with a noisy communication channel between sender and receiver.



Corn

Noisy channel. Which is a radio in a storm, or a phone with one bar.



Herman

Which is every emergency.



Corn

Daniel also raises Uber's H3.



Herman

H3 is a spatial indexing system, and that's the right way to frame it. It's for databases and analysis, hexagons at multiple resolutions, and it's excellent at that job. But it's not designed to be said out loud by a person who is frightened. Nobody's reading a hexagon index over a radio.



Corn

Man walks into a shop with a hexagon index.



Herman

Man walks into a shop and the shopkeeper asks him to repeat it.



Corn

Back to Plus Codes, then. Because they do have real deployments, and we should say so.



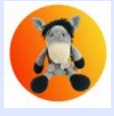
Herman

CEP Digital runs across fifty-five Brazilian favelas. Favela Brasil Xpress uses it to address a hundred and twenty thousand residents. Shelter Associates addressed one thousand seven hundred families in Rajendra Nagar. Those aren't pilots. Those are people receiving post.



Corn

In places where there was no address to receive post at.



Herman

Where there was no address at all. Which is the strongest argument for the system that exists. It's not a thought experiment about emergencies. It's the mechanism by which somebody gets a delivery.



Corn

And there's a comment from the Hacker News thread on the original announcement that's worth reading, because it's the commercial version of Daniel's objection.



Herman

Doctor underscore Fegg, in twenty fifteen. Neat idea, shame the data and algorithm is closed. In the Rio example, a postal service or courier would have to pay What3Words to license the words-to-latitude-longitude lookup before they could deliver anything.



Corn

Read that back. Before they could deliver anything.



Herman

That's the sentence. The address system sits between the courier and the door, and it charges rent.



Corn

That's the one I'd want a listener to remember from the whole episode. An address is infrastructure. Infrastructure that can deny service is a toll road wearing a map.



Herman

And the other comment, from aw3c2 in twenty fourteen. I find them all really questionable as they try to push non-free, centralised setups.



Corn

Which is the same complaint, ten years earlier, before any of this was settled.



Herman

Before any of this was settled, and it held up.



Corn

So let's do the ledger. Six properties, and you can't have all of them.



Herman

Global uniqueness, compactness, memorability, transcription reliability, precision, and offline independence. Shorten for memorability and you buy a dependency. Use words for memorability and you buy a dictionary, a language, and a confusion set. Go finer on precision and you lengthen the string, and every added character is another chance to drop one.



Corn

And the one people forget. A globally unique identifier doesn't mean the location is right.



Herman

That's the distinction Daniel flags and it's the one that gets lost. Uniqueness is a property of the encoding. Accuracy is a property of the fix. You can have a flawless, globally unique, offline-decodable code for a position that is forty metres wrong, because the phone was under a bridge.



Corn

Perfect address, wrong door.



Herman

Perfect address, wrong door, and the code did nothing wrong.



Corn

So where does that leave the two questions?



Herman

First question, yes, the consumer implementation has obscured the standard, and the obscuring is presentational rather than malicious. The standard's own requirements say the codes should be complete on their own. The flagship interface leads with an incomplete one. Google's own README concedes the recovery can fail. Arthur's paper praises the design while documenting Google's substitution of place names. Every source agrees on the shape of the thing.



Corn

And the second question?



Herman

Plus Codes already meet almost every requirement Daniel listed. Complete, offline, open, no licence, no single provider. The system he describes mostly exists, and the piece that doesn't exist is the app.



Corn

Which is the strangest answer available and therefore probably the true one.



Hilbert

It's not thirteen point nine metres.



Hilbert

At the latitude you'd actually be standing at, it's narrower. The grid's fixed in degrees, so the ground distance shrinks as you go north. Thirteen point nine is the equator number and everyone quotes it like it's the number.



Herman

That's right. It's a degree-fraction grid, so the east-west spacing tightens with latitude. At Jerusalem's latitude the cell is noticeably narrower on the ground than the equator figure.



Corn

You've been sitting on that one.



Hilbert

I've been sitting on a lot of things. My brother-in-law sells surveying equipment. Has for years. And he will not use any of these.



Herman

Why not?



Hilbert

He says a code is a promise that someone else's software still exists. A mark on the ground is a promise that the ground is still there.



Corn

That's better than anything we said.



Hilbert

He's also not to be trusted. He once told a customer a boundary line ran through a shed, charged him for the survey, and was right. The customer had built the shed four metres onto the neighbour's side. The neighbour had known for eleven years and said nothing, waiting.



Herman

Waiting for what?



Hilbert

A reason. When the survey landed, the neighbour produced a folder. Photographs, dated, going back eleven years. He'd been documenting a shed that didn't exist yet when he started.



Corn

That's not possible.



Hilbert

The foundation was poured first. It sat there bare for two years while the neighbour photographed it. That's what the dates show.



Herman

So the neighbour documented a hole for two years.



Hilbert

He documented a hole. Then a slab. Then a shed. Then he waited eight more years for someone to ask. My brother-in-law was the first person to formally agree with him.



Corn

What happened to the shed?



Hilbert

It's still there. The neighbour decided he preferred having it. He'd spent eleven years documenting it and didn't want to start over on something else. He asked for a written apology instead, which he got, and which he then framed.



Herman

He framed the apology.



Hilbert

He framed the apology and hung it inside the shed. My brother-in-law says it's the only boundary dispute he's ever closed where both parties ended up with a document they liked.



Corn

What's the code for the shed?



Hilbert

I've never looked it up. I know the code for the pub, though. Nine characters, I've had it memorized for years. Couldn't tell you my own phone number.



Herman

Why the pub?



Hilbert

It's where the folder got handed over. Seemed worth keeping.



Corn

So the ground was there before the code and outlasted the argument.



Hilbert

That's about the size of it. Though the survey still cost four hundred.



Corn

Hilbert, thank you, and thank you to Hilbert Flumingtop, our producer.



Herman

Before we go. The thing on the cutting-room floor, because it doesn't fit anywhere and it's the detail I keep thinking about. The Plus Code alphabet was chosen by generating every possible twenty-character combination and scoring each one against ten thousand words from more than thirty languages, specifically to make sure the codes never spell anything.



Corn

They ran the world's languages through it so it would say nothing.



Herman

That the only thing a Plus Code can ever mean is a place.



Corn

Here's where I land. The technology for a universal, open, human-readable geographic address already exists, and the reason it isn't a public utility isn't technical. It's that nobody's built the boring app around it.



Herman

The boring app is the whole thing. A library that matches the reference implementation in eleven languages is not a utility. A watch datafield that works offline is closer, but it's still not the thing a person reaches for.



Corn

The invention happened. The adoption didn't.



Herman

Which means Daniel's quest has an answer, and the answer is a to-do list rather than a research programme.



Corn

If you enjoyed this one, rate and review the show. It helps other listeners find us.



Herman

We'll be back soon.



Corn

If you want more of this, try episode fifty-nine twenty-three, Plus Codes; episode forty-three seventy-three, Why Jaffa Street Has Three Names; and episode twenty-four thirty-seven, The False Precision of GPS Coordinates. This has been My Weird Prompts.

SOURCES

What the research for this episode read before the script was written. Primary sources first.

1. **OLC repo README, Apache-2.0, last push 2026-03-30. [primary]**
https://github.com/google/open-location-code

2. **OLC definition paper (Rinckes & Bunge). [primary]**
https://github.com/google/open-location-code/blob/main/Documentation/Specification/olc_definition.adoc

3. **An Evaluation of Location Encoding Systems. [primary]**
https://github.com/google/open-location-code/blob/main/Documentation/Reference/comparison.adoc

4. **Google's consumer Plus Codes site (shortened-code presentation). [primary]**
https://plus.codes/

5. **Arthur, PLOS ONE 18(10):e0292491, 2023-10-25.**
https://pmc.ncbi.nlm.nih.gov/articles/PMC10599581/

6. **same paper, publisher version.**
https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0292491

7. **preprint of the What3Words analysis.**
https://arxiv.org/pdf/2308.16025

8. **JS/TS OLC implementation.**
https://github.com/janne/pluscodes

9. **Python OLC package v2022.1.3.**
https://pypi.org/project/pluscodes/

10. **Kotlin Multiplatform OLC port.**
https://github.com/aughtone/aughtone-openlocationcode

11. **HN thread on Google's OLC announcement (2015).**
https://news.ycombinator.com/item?id=9470158

12. **HN comments on What3Words closed data.**
https://news.ycombinator.com/item?id=9464465

13. **comparison of address/coordinate systems.**
https://www.colemanm.org/post/comparing-address-and-coordinate-systems/

14. **comparison table of geocoding systems.**
https://gridpointcode.com/compare