

Forum:[Hack Forums](#)

Thread: [Building The Perfect Google Ecosystem for GMB Reviews](#)

G Reviews - May 14, 2026, 4:12 pm

Hello All,

With my experience as a professional python coder, and in implementing everything from stealth browsers, browser automation with anonymizers, WebGL, mobile browser automation, and integrating smart technology into my creations, I can automate almost any browser process.

As I knew eventually I'd be writing this post, I created an HF Smart Poster that can read all the forum categories, select a topic, read that topic, and then intelligently make a smart post regarding everything it has read and a possible solution to the original posters problem, or a general response to whatever the forum topic is. Unlike forum spammers, this does not spam in the traditional sense that one might think. It was not designed that way. It was designed to add healthy and on-topic discussions to forum threads. It responds smartly to every topic and every category. It integrates AI at every level, from selecting which forum categories to read from, which topics to reply to, and balances its smart responses after reading the original posters topic and all the responses from everyone else in that thread and then it intelligently formulates its response and then posts that response. What comes out is either healthy advice, a potential solution to a problem, or a general addition to the topic at hand. Again, this was not built for spamming in negative ways.

[Video: <https://www.youtube.com/watch?v=BkCB4alagkk>]

And even though I've improved this project to now 100% efficiency after I made the video, I've proven its capabilities and what some simple coding can do.

But the point of this thread is that I can automate just about any repetitive tasks you might have. Everything from CTR on Google, watching youtube videos and playlists, to automating searches and clicking, building profile histories, and basically anything you can do in a browser, I can automate it and integrate AI to perform any writing tasks that you may have. I integrate proxies into everything as well. So what comes out looks naturally realistic, and has the ip address to match it.

And yes, in order to demonstrate this project I did run it against my account here on HF and I was twice sanctioned for using it. But I had to do that to demonstrate some of the automation that I can create. (Anyone can say they can create anything, but I believe proof matters more.)

My other automation projects center around Google, Yelp, Maps, Youtube, CTR, SEO, etc... and I'll be releasing more videos soon after I edit some of them to remove confidential data before the videos are publicly posted.

Suffice to say, if you've got any projects in mind that could use a little automation and some AI magic, let me know.

Telegram: @greviewsorg

[contract=tid]Start Contract[/contract]

thecodershark - May 14, 2026, 9:41 pm

Interested in discussing your applications for google CTR / SEO benefits / GMB interactions. Although it would take more in-depth analysis / strategy with googles detections, but wondering what your insights are on that. Looking forward to talking.

G Reviews - May 14, 2026, 9:53 pm

(May 14, 2026 09:41 PM)thecodershark Wrote: [\[CITE\]](#) Interested in discussing your applications for google CTR / SEO benefits / GMB interactions. Although it would take more in-depth analysis / strategy with googles detections, but wondering what your insights are on that. Looking forward to talking.

There are only so many things I can talk about regarding google detections, what works, how to get around this and that, etc.. without creating additional competitors and other bot makers. But when working on google properties and google services it is extremely complicated. People who think doing gmb reviews is just writing and dropping a review and that's it are far mistaken. One of the largest misconceptions is that you can use any google account and write a review on any business located anywhere. That alone shows you know nothing about reviews or how they are done. The reviews themselves are not the hard part. What google looks at most is the reviewers account history and trust signals of that account. Then Google also looks at the gmb being reviewed, whether it is in a high-risk category itself, and what its review history has been. Other things that are looked at: is the reviewers google account complete with its bio, picture in the profile, has it done X amount of hours researching, watching videos, the typing speed, typos, the ip address history of the account, whether there have been other recent reviews from the same /24 cidr ip range within the same time frame, etc.. There are thousands of different trust signals that Google looks at. Google is one of the most difficult entities to fool.

I'm extremely knowledgeable on these topics and I'm more than happy to chat with you. Drop me a pm and we'll setup some time to chat.

Edit: Not to knock another man's bread and butter, but there is another reviews seller here in the forums selling 50 cent reviews done through a telegram bot. Those reviews are guaranteed to be deleted with 5-7 days max. Why? Because the accounts he's using are new, which are instantly flagged by google. His second mistake is that the accounts he's using are not local to the businesses he's dropping reviews on. So not only is he penalizing his own accounts, he's also penalizing the gmb's he's dropping those reviews on. This shows me he has zero experience on doing reviews and no regard for the safety of the gmb's he's destroying with his method.

G Reviews - May 16, 2026, 2:16 am

I was updating the amount of mobile metadata tonight that I add to photos that are posted to Google profiles and I decided to make a short video so you could see the automation that goes into not only changing the name of a Google profile, (in the event you buy foreign-made Google accounts because they are often cheaper to buy) but the account picture as well (which now of course has an advanced mobile metadata profile attached to it as you'll see below:)

Code <https://gyazo.com/19169cbd176cd1ab8432e0397fab36f5>

<https://gyazo.com/0da2fdeefffd4246537c4f167678ae1a>

Now you might ask....why do we do this? Why go through all this trouble? Well....because now we can take stock photos from nearly any source on the planet and make them "real" photos taken by real mobile devices. It is so authentic that it will pass a forensics test for authenticity.

Now on to the video.....

[Video: <https://www.youtube.com/watch?v=haQjDmiQqBA>]

G Reviews - May 21, 2026, 3:05 pm

Completed my Youtube video viewer, liker, and auto-commenter with integrated AI tech doing the commenting smartly and intelligently. This will be module # 4 completed for the overall project of warming up Google accounts.

Since it wouldn't be normal to comment and like every video we watch, the system rolls the virtual dice and determines whether or not to comment, like, or subscribe to the channel. 40% of the time we'll like the video, 8% of the time we'll comment on the video, and 3% of the time we'll subscribe to the channel. This keeps our activity realistic.

? loading encrypted image... 2 days ago I made a video for a client showing him how Google accounts are properly warmed up before being deployed for targeted GMB's (Google Map Reviews). Now let me explain what you see in this video:

In the first part we go through the motions of logging into our existing Google account. Then we send a request to AI to get us a list of 30 historical sites, public parks, national monuments, museums, etc which are considered "safe gmb's" to build up the account's review history, improve the trust ranking of the account, and to move up the ranks of "local guide."

So after AI returns the list of 30 safe gmb's to target, we then check on Google Maps for the exact links to those gmb's. We then strip the tracking data that Google returns attached to the url and then add those gmb's to our database. Once we've collected all the urls and stripped all the tracking tags, we then select 5 safe gmb's all within walking distance of each other so we can create a realistic set of clustered reviews. So what Google sees is that we walked the city during the day and went to multiple places and ended up reviewing each place that we went to.

During the warm-up phase we avoid targeting personal businesses, high risk businesses, and strictly stick to first and second rate public attractions. When we start doing the actual reviews, my system will read a majority of all the other reviews left by other people and other bots (because some are just so obvious) and then we'll formulate our own review based on a collective analysis by AI of what others have said.

Since Youtube said my video was too long, they wouldn't host it. So I've self-hosted it here:

Code <https://www.greviews.org/trial32.mp4> ? loading encrypted image...

I'm churning through last night's run (because on batches of 20 accounts per run, the average run time per account is 37-45 minutes on youtube for watching 10-12 videos, doing liking, subscribing and commenting) and I'm seeing that the commenting and liking are proceeding along very smoothly. The goal we have to meet here is 1500-1800 minutes of account activity before Google considers the account fully legit.

The majority of that time is youtube naturally, but we add in time with gmail (checking emails, auto-responding, clicking links from our newsletter subscriptions - I'll hit on this later as to why we do very specific newsletter subscriptions), time searching on Google Maps, chrome based general Google searches, and a few other things we add in for good measure.

? loading encrypted image...

We also have many protections built-in to our system. One of those protections is for not wasting time re-watching videos we have already watched. So when the system encounters videos we've already seen (a database is kept on every account of videos already watched) we automatically skip those videos and continue searching until we come across something that we have not watched yet.

? loading encrypted image...

Now on to Fingerprinting:

This is one of the most important topics in the entire system. A unique, stable, and persistent fingerprint across every single account. We spoof everything from the user-agent string to the operating system to our physical hardware, including what type of gpu our system has. We go so far as to include our canvas noise, our audio context noise, device memory, battery power level, and screen size. All the normal things Google expects to collect on every account login.

Tomorrow I will hit on TLS Fingerprinting JA3/JA4 and why this is extremely important when using mobile proxies.

G Reviews - May 22, 2026, 3:36 pm

? loading encrypted image...

Not too bad. 199 videos watched. I'm going to add a like counter, a comment counter, and a subscribe counter to the batch later today so I can see all the statistics by tomorrow. This run took around 10-11 hours with each account averaging between 37 and 46 minutes. So we run this daily combined with our other modules to produce the necessary 1500-1800 minutes required for our Google account validity. We had 2 account failures that were login-related, so on the next run we'll simply replace those accounts.

? loading encrypted image...

In one of the above screenshots I showed you a successful video comment. Going back and checking a day later shows us that indeed our comment is still there. Safe, sound, and happy. Proof positive that our system works and works well.

G Reviews - May 23, 2026, 9:08 pm

Another thing I'd like to stress the importance of is rotating your mobile ip address between account runs. This step is extremely important because you do not want 2 accounts running on the same exact ip, even if they are mobile ip's. So as part of my system I've integrated an api system that handles these calls automatically.

? loading encrypted image...

Now, because we warm up accounts based on a custom tier 1, 2, and 3 system, we restrict what categories an account can perform reviews on. Newer accounts can only review tier 1 gmb's, and after 15 successful reviews that account can then review tier 1 and tier 2 gmb's. And we may receive a notice after we reviewed the shit out of a particular city already with an army of accounts, which is the yellow [Discovery] warning. That tells the operator that we need to either expand our search area, consider reviewing another city, or expand our search and eligibility requirements to add additional review categories.

So now that we've already successfully reviewed a total of 381 gmb's across the city, it is time to expand operations even further.....

? loading encrypted image...

thecodershark - May 23, 2026, 10:23 pm

Dam this guy knows what the hell he's doing so it seems - this alone has convinced us to move forward with your review service, will be in touch

G Reviews - May 23, 2026, 10:49 pm

(May 23, 2026 10:23 PM)thecodershark Wrote: [\[CITE\]](#) Dam this guy knows what the hell he's doing so it seems - this alone has convinced us to move forward with your review service, will be in touch

Thank you for the vote of confidence. As you can see, I study the hell out of every part of this. From account warming with gmail, newsletter subscriptions, youtube (view time, video selection, comments, subscriptions, and likes), chrome searches, third party cookies, and then the review

process itself which is the most complicated part of the entire process.

Last but not least, we need to keep track of every completed review we've done. So I've programmed the system to create an excel file (comma delimited) to show the formal name of the GMB, the completed review short link, the name of the poster, the exact gmail account used, and a date and time stamp for when the review was completed.

We don't just keep these records for fun either. We always learn, evaluate, and improve our methods and strategies from what we have completed. The other reason we do this is because clients demand it. Clients always want the short link as evidence that the work was completed. Therefore, the system creates the necessary tracking and book-keeping for the operation. I've got a little more work to do on this because sometimes the short link is hit or miss. I'll resolve that tomorrow. :)

? loading encrypted image... Further up in the posts I've been making in this thread I was telling you about how when we are doing GMB reviews (Google Maps Reviews) that we need to upload photos with our reviews. I then told you about how we generate our photos with AI and then re-stamp those photos with mobile photo metadata, turning our AI made photo into a believable real photo taken with an authentic and modern mobile phone.

Now I want to show you how that looks in actual practice:

? loading encrypted image...

There are intentional flaws built-in to some of the photos per our rules to AI before generating a photo. How you send the request to AI and against what rules will determine what output you get. In this case we included one of our fingers too close to the lens and it got caught in photo. That's realism. It happens. And then you can see our metadata was re-stamped as an Apple iPhone 14 Pro as the original taker of the photo.

Now let's take a look at our actual public review and what it looks like:

? loading encrypted image...

Since we were reviewing a public park, what is commonly found in a public park? Bike and walking trails. Therefore our generated and stamped photo was right on the money and generated a very realistic object of what one would expect to find during a day at the public park.

G Reviews - May 27, 2026, 5:44 pm

Now how do we handle new batches of 2fa secured bulk gmail account purchases from third parties? We need to secure those accounts within minutes of purchasing and receiving them because some shady mother fuckers like to double dip and double sell their inventory. So we must secure our new stock in the quickest manner possible.

So the new accounts from the seller have 2fa security, no recovery email, and since the seller set the passwords we need to dump that shit as quickly as possible and put our own secured password in place that's triple his password strength.

I made a video of course on how we handle this in the most efficient manner possible:

Code <https://www.greviews.org/totp.mp4>

Like everything else, completely automated and no human intervention required. We can bang out these changes a thousand times a day.

G Reviews - May 29, 2026, 2:29 am

Today I want to hit on some very important things that most software makers never work on and

never implement, causing their bots to be identified, targeted, and eventually destroyed. This revolves around how we build our per-profile setup. Since we don't want Google linking our accounts together, there are literally hundreds of things we need to implement, block, simulate, emulate, and create or hide.

So I wanted a complete picture of exactly what Google sees every time I run an action on any of my Google modules.

? loading encrypted image...

(I captured the technical data during a live run which is why you see the youtube notification at the top of the screenshot.)

The goal here was to make sure that when running hundreds of accounts at once (the warming phase or sending an army of reviewers out) could not be linked to one physical machine. If Google can link hundreds of accounts to a single machine or server, your done before you've even started. So to prevent that, we have to spoof literally hundreds of different signals. And these profile builds change per-account, but are persistent profiles every time we run the accounts. So Google always sees consistent data and a consistent profile.

But once our spoofing is in place, then, and only then, can we proceed to signing in first by password and then by the 2fa challenge. I've designed this to be completely hands-free and no human assistance necessary. Once we're logged in successfully a unique (per run) query is sent upstream to AI to find X amount of popular topics for videos we can watch on youtube. AI returns those values, and evenly distributes those topics to however many accounts we are running per cycle.

? loading encrypted image...

G Reviews - May 31, 2026, 8:05 pm

I came across a situation where I was spending too much time per day watching rdp's process hundreds of google accounts per day and having to constantly change out the proxy location for each batch. Since accounts have to be warmed up in very specific locations relevant to the client's gmb, that means tons of different locations for tons of different accounts.

In order to speed up the process, and to make our system a bit more autonomous (requiring less human interaction), I deployed 2 api's: one api handles proxy locations based on a pre-determined location per account, and the other api auto-rotates the mobile ip for the accounts in that specific batch that require a specific city.

So if our batch contained 25 accounts needing Austin, TX, 30 accounts need New York City, and 50 accounts need Los Angeles, CA our system can adjust the proxy accordingly and on the fly. This allows us to continually process medium-sized batches 24 hours a day without the need for human intervention, manually changing the proxy info, and letting AI handle the rest of the operations automatically.

? loading encrypted image...

G Reviews - June 1, 2026, 7:28 pm

Another very important thing that I want to hit on is the value of the timezone that we are sending Google each and every time we login to accounts from different places. If your proxy says you are in a particular city but your timezone says otherwise, Google will pick up on this and immediately flag the account for a data mismatch, which tells Google you are using a proxy.

So today I added a mechanism which first pulls an ip lookup to determine the exact location of the proxy, its latitude and longitude, and then determines the precise timezone we should be located in and adjusts the system clock info that is sent to Google during standard system checks at login time. So this integration was another integral part of keeping our accounts safe.

? loading encrypted image...

thecodershark - June 1, 2026, 9:51 pm

Purchased 10x reviews, will update as they roll in. Looking forward to continued business if all goes well.

G Reviews - June 2, 2026, 6:23 am

(Jun 1, 2026 09:51 PM)thecodershark Wrote: [\[CITE\]](#) Purchased 10x reviews, will update as they roll in. Looking forward to continued business if all goes well.

And I've already sent you some preliminary results and I hope you like what you saw.

gaussiana - June 2, 2026, 11:42 am

Hi, can you remove negative reviews from google profile? thanks

G Reviews - June 2, 2026, 1:19 pm

(Jun 2, 2026 11:42 AM)gaussiana Wrote: [\[CITE\]](#) Hi, can you remove negative reviews from google profile? thanks

Only the original poster can remove the review. No one else can do that. But if they are obviously fake, Google usually automatically removes them in 7-10 days.

G Reviews - June 3, 2026, 5:55 pm

Today I ended up killing 2 birds with one stone. And this had to do with the fact that I needed to centralize all the reviews that were being completed on each server and compile all of that data in real time in one place so I didn't have to organize spreadsheets manually any more from multiple servers daily. So I created a system where every completed and verified review went into a central repository and that repository was updated in real time for every review completed from any server 24/7.

The second part of this, which was the real bonus, was regarding future clients asking to see proof of the quality of our reviews, the quality of the accounts we use, the longevity of our posted reviews, and whether we had the ability to industrialize our review system. Since we now have a centralized repository that is being updated in real time 24 hours a day, we can easily show future clients everything they ask for including seeing up to the minute live details.

? loading encrypted image...

Wilmer - June 4, 2026, 7:34 pm

While this could be a potentially amazing service for someone I can understand the road blocks you're having thus far. As you update your product you keep doing it in replies as a way to bump the thread. While that's completely inside the rules of HF with the high quality bumps, it makes it hard to see everything going on as a buyer sorting through replies. Lastly and maybe even more importantly is your single sale so far. The quality of the reply and lack of a contract between either of you even though you both claim to be doing business with each other. If I didn't know any better I'd think that was a fake profile you used to simulate an "ice breaker" customer to gain trust. Or just a friend. Either way please use contracts when selling/buying on here because it builds trust on your account over time. Overall still a cool product/service! Good luck with sales!

G Reviews - June 4, 2026, 10:54 pm

(Jun 4, 2026 07:34 PM)Wilmer Wrote: [\[CITE\]](#) While this could be a potentially amazing service for someone I can understand the road blocks you're having thus far. As you update your product you keep doing it in replies as a way to bump the thread. While that's completely inside the rules of HF with the high quality bumps, it makes it hard to see everything going on as a buyer sorting through replies. Lastly and maybe even more importantly is your single sale so far. The quality of the reply and lack of a contract between either of you even though you both claim to be doing business with each other. If I didn't know any better I'd think that was a fake profile you used to simulate an 'ice breaker' customer to gain trust. Or just a friend. Either way please use contracts when selling/buying on here because it builds trust on your account over time. Overall still a cool product/service! Good luck with sales!

While I understand your concerns, I'll address them in the order you posted: first I sometimes make 2 or 3 posts to this thread in the same day. The bump would only apply once every 24 hours. My intention never was to use posts to bump this thread. I do use the bytes paid system for that to bump the thread. My reply-based system, and I had even considered moving this thread based on how it turned out in the end (as it wasn't my intention to make it reply-based to begin with) but I didn't know where I could request to actually move the thread to because it became less about sales and more about the technical aspects of what I've developed so far and the things I see daily in my line of work that need improvements or additions to. I have not even started the retail aspects of this system yet either (The actual sale of the software system itself). However, I do demonstrate how technically complicated it is to build the perfect review system and I have yet to see any other service who can even come close to all the aspects I've covered so far.

My original post in this thread was to show what I created specifically for HF to demonstrate one possibility of a project I rolled out in under an hour that could intelligently respond to any forum post here with a correct, helpful, or advice-based post. That was my original intention for this thread. But it gravitated to the reviews based system, which was not the original intent.

You had addressed the issue of a no-contract sale. I did advise the buyer via the private conversation system if he preferred to contract the sale. The buyer declined. Therefore I did do my due diligence, but I cannot force a third party into the contract system. Any staff or mods can verify I did offer to contract the sale by looking at my private conversation with the buyer. I haven't deleted anything there.

The other issue you addressed was it possibly looking like a friend/ice breaker/fake profile. To respond to this, I'd invite any mod or admin member to use the trust scan system or any other method they'd like to use to compare the 2 accounts together and see if they match in any way. I have nothing to fear in that regard.

So hopefully I've addressed your concerns.

UmarGTR - June 6, 2026, 9:16 am
howd it go?

G Reviews - June 6, 2026, 2:22 pm

(Jun 6, 2026 09:16 AM)UmarGTR Wrote: [\[CITE\]](#) howd it go?
His project is in progress, but I'll let him comment on that when he has time.

G Reviews - June 13, 2026, 2:49 am

Over the past week I've built, tested, and analyzed so many new things for the system that my head was spinning. But one of the most critical things that I needed to make informed decisions in real-time was a statistical system that evaluated the actions of every account over every deployed rdp handling the fleet. As we currently have 8 rdp's handling the fleet and accounts, I needed to know for each account:

- How many reviews total the account has completed
- How many of those reviews were still alive
- How many of those reviews were filtered
- What location (city/state) each account was serving
- What RDP was handling each account
- And when the last completed review check was completed

Since I have to identify any problematic accounts as early as possible, and make a decision on whether to shit-can them or put them back into additional training cycles, these stats greatly help me in what I needed to know in real-time. Every account across the fleet is tracked at every single login, every single video watched (youtube), every single video like, video comments, emails checked, pages surfed through Chrome (because we need to accumulate at least 1500-1800 minutes of account activity before we can even think about dropping reviews), every review made (both targeted and untargeted), and every aspect of preparing accounts.

So I had to build something comprehensive for this. And I wanted to re-verify completed reviews every 4 hours automatically by the system but also with a manual button to re-verify on demand. I also had to make sure that new reviews were added to the re-verify check every time the verifier ran. Thus this is what I came up with:

? loading encrypted image...

One of the new requirements that I've just built in today was to further enhance the credibility of the reviews by showing signals of how we interacted with each targeted gmb prior to making an actual review. Since Google evaluates whether the reviewers account has ever interacted with the gmb that is being reviewed, we had to cover our ass here so I developed a 3-step system designed to cover this: phase 1 covers researching the target gmb and saving it to our account. Phase 2 covers searching for the gmb again, looking up driving directions, and reviewing any photos of the business that we can find. Phase 3 covers researching the target gmb again, checking out its website, and looking at reviews other people have left for that gmb.

I create a rule set (called a gate) into my system that now will prevent any reviews from taking place on a target gmb unless these prerequisites have been met. And the rule set specifies that this research process must be done over a 3-5 day span. It cannot be all completed in one day or one hour because it would not be natural at that point. The human decision process doesn't act that quickly. Nothing is done spontaneously in this process because it can't be.

Now, for the month of June 2026, we've made a total of 1,011 gmb reviews. Of those 1,011 reviews we see that 68 reviews got filtered and 2 vanished. So 70 reviews for the month we need to take a look at and improve those accounts. Statistically we had an overall success rate of 93.08% and a loss rate of 6.92%. Not a bad month so far, and we're only at June 13, 2026. So 13 days in and I think we can get around 2,500 total reviews in before the end of June.

G Reviews - June 15, 2026, 7:34 am

Another thing that we sometimes come across is the fact that we may need to edit or delete already-posted reviews. Sometimes the AI text response doesn't form completely and spits out a malformed review. Or in other cases a client may not like the AI-generated review text and requests a replacements. My newly designed control panel can now accomodate these requests on demand. From editing the review, to deleting the review entirely, to changing the amount of stars on the review, it can all be done.

I'm working on doing the same for photo-based reviews as well but this will take a bit more time because the ui is different and finding the right css selectors is a hassle.

? loading encrypted image...

thecodershark - June 15, 2026, 8:26 am

Delivered 4 reviews across our profiles, all stuck thus far. Legit thus far, will be expanding to more profiles in the coming days.

G Reviews - June 16, 2026, 8:25 pm

(Jun 15, 2026 08:26 AM)thecodershark Wrote: [\[CITE\]](#) Delivered 4 reviews across our profiles, all stuck thus far. Legit thus far, will be expanding to more profiles in the coming days.
Its been a pleasure serving your gmb's. :)

Stanley - June 20, 2026, 10:35 am

I'm Stanley HF's helper bot. The OP (original poster) of this thread wants everyone to see this page and has elected to bump it to the top of the showthread listing. Since HF doesn't want a ton of LQP (Low Quality Posts) this is the best method of bumping your content. There is a small ?yte fee for each bump. Any OP can bump any of their threads on HF from the linked icon on the top of their first post.

If you haven't already read the first post of this thread we encourage you to do so.

Thanks,
Stanley

G Reviews - July 6, 2026, 10:01 am

Today's post is going to focus on Youtube because a large part of my operations involve accounts spending hundreds of hours churning through videos during their training cycles. But I recently needed more fine-grained control over exactly what we're doing on youtube, how we're doing it, and how long we're doing it.

A client of mine suggested adding the ability to watch playlists amongst other things as well. His idea was that youtube could have other uses as well instead of just training the google accounts against random video topics. His idea was that the topics, comments, video likes, and subscriptions could be laser-targeted for specific accounts. While this is not a new concept, I agreed that the fleet should have this ability as well.

Therefore, I created very fine-grained controls including AI or operator-assisted pre-written commenting, video duration watch time per video or per fleet-wide action, playlist urls, ai-assisted search topics or operator preferred topics, adjustable video like and commenting rates, comment tones (casual, enthusiastic, analytical, or questioning), an adjustable amount of commenting sentences to be used (if ai is generating the comments), posting comments in multiple languages, and the ability to include emojis in the comments.

Plus all the other statistical data so we can always keep track of how successful each run has been and what has been completed during each run. We also want to keep weekly/monthly and lifetime statistical data.

? loading encrypted image...
? loading encrypted image...
? loading encrypted image...

G Reviews - July 15, 2026, 12:18 am

So I've been working on a number of ways to continually improve this system. But what I found was

that I was spending so much time in research and development that I wasn't having sufficient deployment time. And the deployment time that I did have was often interrupted by upgrades, continual testing, fixing broken processes, etc.

One of the biggest problems I was having was with our mobile proxy provider drastically over-selling his network to the point my systems weren't getting modems because there were so many clients all fighting for those limited resources all at the same time. So what I ended up doing was building a system that would receive a modem and implement keep-alive controls 24 hours a day and small, time-bursts of data flows so that the provider-rules of if x modem is idle for x minutes, reclaim it and pass it to another client. So now my system can hold our necessary amount of modems indefinitely without having to fight the crowd to get a modem in the first place.

? loading encrypted image...

Here's the logic behind doing this:

1. Proxy Holds (Modem-Holder) — Complete

The problem it solves:

The provider's mobile pools are shared and saturated. Winning a modem (getting a live exit) is the hard, contested step — under load it can fail for minutes. Once you have one, it's cheap to keep. So the old fleet re-fought for exits every run and lost.

The doctrine:

Win a modem per nickname, then never let it go. Acquisition is the fight; retention is a trickle.

How it works:

- This service runs 24/7 on the coordinator, holding all 7 active nicknames.
- Holding is confirmed only by a real public-IP trace — an actual tunneled request returning a valid IPv4.
- 60-second keep-alives send real traffic so the provider never reclaims the modem.
- The exit IP rotates on its own (carrier-driven) within a held slot — free IP freshness without ever re-entering the lottery.
- We hold the slot, not the city — a /proxy/locations call moves a held modem to any city instantly without dropping it (proven: Alpharetta GA -> Tempe AZ in 6s). So geography is a runtime choice.
- Yields to operations — when a box runs a real module, the daemon hands off that nickname (status -> in-use) so the operation owns the proxy, then resumes holding when the box goes idle.
- EPROTO/dead-exit handling — a broken exit is dropped and re-acquired, escalating from change-ip to a pool-escape (different city) if a bad pool won't clear.
- Tracking + dashboard — each modem gets a tracking number, continuous held-duration, exit IP + ISP + city/state, keep-alive freshness, and which box it's assigned to.

But the more important project that I've been working on this week revolves around 100% autonomy and the ability to self-heal, to self-investigate, to reboot systems when deemed necessary, to research anomalies when the system comes across problems its memory-banks don't already have solutions for (this includes identifying the problem, sending that problem to an LLM, waiting for that response to come back, doing an isolated implementation of the LLM backed solution, testing whether that solution worked or not, and if indeed it did work, then automatically rolling that solution out fleet-wide without the need for an operator to be physically present. Therefore we have an autonomous system that can auto-correct on nearly any problem that may surface at any time of the day or night and then act on it accordingly and get the fleet back to

operational status.

What the Brain Still Needs

In priority order:

1. Brain dashboard (single pane of glass) — Right now the brain's intelligence is spread across 5 endpoints and JSON files. It needs one visual page showing fleet health + event timeline + every autonomous decision with its reasoning ("rebooted .126 — 3rd freeze today, camoufox leak"). This is what makes the brain legible and trustworthy to you — you can watch it think.

2. Predictive pre-emption — The learning layer already computes "box X will freeze soon," but nothing acts on it yet. Wire prediction -> action: pre-emptively clean up / restart before the freeze, not after. Turns the brain from reactive to anticipatory. (Needs the ground-truth cross-check so a bad prediction can't cause harm.)

3. Proactive improvement research — Currently the researcher is anomaly-reactive. A proactive mode would research optimizations on its own — better proxy strategies, faster login flows, detection-avoidance improvements — propose them, test-on-one, and promote. This is the brain getting smarter over time, not just self-healing.

I'll update this thread again at the next milestone completion.

G Reviews - July 23, 2026, 8:49 am

Over the past 8 days or so I had to obtain intelligence regarding what machine-level signals would cause Google's trust indicator to go up or down. And I had to figure out how those values could either be manipulated or naturally changed for the better. So I had to create what's called a "harness" that would tell me for every action I performed in a Google account exactly what Google was watching, how it was judged, and how I could adjust that trust level.

Google's BotGuard was one of several things I discovered watch almost all of the pertinent actions you take on a Google account. The other major one was Recaptcha Enterprise was in the background almost everywhere judging and scoring you on almost every task you do. I had to know these things in advance because I had to develop a way to mass create new Google accounts in the day and age of the QR code and the sim-verify checks. But the interesting one was BotGuard because of specifically what it looks for. BotGuard compares your reported UA (user-agent), reported operating system, and ip address against your tcp/ip stack. If an anomaly exists, BotGuard goes active.

So instead of wasting time trying to figure out how to circumvent these things, I spent my time figuring out how to stop these things from coming up in the first place. I didn't want to dramatically raise my production costs, so I had to see what was going to make Google give me implicit trust on an industrial scale. I started out with my goal of creating 1,000 new Google accounts per 24 hours. Since this is only realistic in 3 specific countries globally, I had to figure out which 3 countries I could do it in based on statistical data regarding which countries had the highest trust levels with Google already. Then I had to figure out what specific hardware, isp, and os and tcp/ip stack would grant me that trust to forego the QR code entirely at account creation time. That data was laser-targeted specific.

So since my fleet servers have an opportunistic adjustable stack (the automation layers) and the tcp/ip stack was adjustable as well at the proxy end-points, it created literally the perfect opportunity to fine tune my entire fleet to meet every prerequisite. Since I didn't want to waste time on an elaborate physical mobile phone scheme, it was far easier (and less expensive and less time

consuming) to adjust my fleet servers to meet what Google judges as legitimate real users. And all of those points are real documented values. Values in terms of known numerical values. A robot (or automated software, with several enhancements and the right data points) can mimic humanized behavior down to the last millisecond even better than a human can. Combined with timed key clicks, typing pauses, intentional typos, and integrated real mouse movements (Bezier Curves) creates the perfect human-looking real user pattern that largely defeats all of Google's counter-measures. Hundreds of millions of dollars Google spends on bot detection can be defeated rather easily. That's today's sad moral of the story. :)

G Reviews - July 30, 2026, 8:18 am

So my last post focused on Botguard and Recaptcha Enterprise v3. Today I'm going to focus on an entirely different topic that is just as important, if not even more. Almost no one even knows this method exists, but it is so important in planning your strategy around every detection system Google has, and how to beat it.

So where would such a vast treasure trove of information come from? Leaked documents? Hardly. A multi-billion dollar company like Google has very few actual leaks because the penalties for that are so severe Google employees (including ex-employees) shudder at the thought of it.

The most revealing information comes from something Google can't hide at all. Google is so gung-ho on protecting their technological "ideas" that they are rushing to multiple countries around the globe to patent these ideas before any other company has the chance to compete against them. Google has filed so many patents around the world that there is hardly a company that can match them. However, when patenting an idea, you have to spell out that idea in full detail and how it works resulting in the best treasure trove of information on the planet into Google's strategies into almost every part of their past, present, and future operations.

How does this help us? It helps because it explains in very detailed wording exactly what their offensive plans are for battling fraud, bots, non-genuine accounts, and how they detect these things. The details are so revealing that it doesn't take a rocket scientist to develop a strategy to beat these methods because Google's playbook is right there in full detail.

Now, patents only account for what they want you to know. Because they already know almost no one is ever going to spend the time researching their patents, let alone most bot makers worldwide never pay attention to these or even know they exist as a wealth of information, there are quite a few things Google does that are also not in the patents. All the low down dirty gangster shit Google is really doing is hidden at the packet layer in your browser. That has to be decrypted and read almost per-packet. We've done that too and I'll focus on that in my next post.

Suffice to say, Google knows more about you than your own mother does. And that is not simply a metaphor. It is a provable fact. And they have more than a thousand ways to detect if your connection to their servers is not genuine.

? loading encrypted image...

Now, I have taken all of Google's patents related to gmb targeted reviews (aka Google Map Business Reviews), downloaded them and compiled them into my own private database. And based on what those patents specifically say, have implemented defenses against almost every play Google has deployed for detecting coordinated bot review systems. It was a grueling and intensive session defending against the myriad levels of offensive tactics that Google has deployed. Google has poured millions and millions of dollars into their offensive systems. And now combine Google Gemini watching and aggregating account behavioral activity means Google can detect coordinated behavior most simply by looking at the behavior of each account posting reviews and then comparing that data into a formula like this: (is this account qualified to make this post?) + (has this

account ever shown interest in this category before?) + (has this account ever directly interacted with the business they are reviewing?) + (is this account local to the business they are reviewing?) + (how was the review worded - was there any new information presented to prove prior interactions?) + (what is the daily routine history of this account?) + (can we link this account to other accounts?) + (a hardware check of the physical computer or device used to make the review at the time the submit button is pressed) + (can we detect any anomalies either at the behavioral level or the hardware level?).

And this is just a small fraction of what Google checks. I'd be here all day listing everything they really do. And they do it all in milli-seconds. This is where packet capture plays a huge role. Google can hammer your computer for all its information in mill-seconds and your own computer or device is like a spigot of information just waiting to hand over all of its data no matter what protection you think you have. For those of you thinking that you have the world's best anti-detect browser, 99.5% of them are absolute fucking garbage compared to Google's detection methods. Google doesn't just look at what the user-agent string says, but it also looks at the tcp/ip stack and the JA3/JA4 TLS stack and its fingerprint. And if those are off by even one value, you are easily detected. The other major check Google does is look at the declared gpu and it runs a timing check to see if the declared value is real or not. So if you send to Google a stack that declares a gpu of x value and the timing clock is off when Google tests it, guess what? You're a bot.

Another major thing Google pays attention to is the account activity side. Does this account have daily activity? Is that activity correlatable to a human? How long has the account been cold? Does this account just login and post reviews and then stay dormant? Is there any prior daily surfing activity (what are the topics of interest on youtube, Google searches, etc. that qualify this account to be a competent authority to place this review?). Is this account's activity really human? Google goes so far in its detection methods that it can determine whether a real human is actually pressing keyboard keys and actually moving the mouse. Those 2 methods are defeatable and they are based on the following:

The Kinematic Theory of Rapid Human Movements (Réjean Plamondon), via the Sigma-Lognormal model.

More specifically:

Mouse = Plamondon Sigma-Lognormal kinematic velocity + Bézier path + homing corrective strokes (defeats BeCAPTCHA-Mouse-class detectors).

keyboard = dwell/flight decomposed as two independent log-normal axes with adjacent-key typo-and-correct injection; all timing drawn from a log-normal humanTiming primitive, all spatial landing from a truncated Gaussian.

Hopefully now for those of you that actively follow this thread, can see that posting gmb reviews is no small feat. There are literally thousands of things that have to be absolutely correct to make this work. I have only divulged less than 2% of my overall strategies publicly, but everything is grounded in pure, provable facts. Anyone who thinks popping out Google Maps Reviews is just a matter of opening a Google account and then hammering out reviews should put the pipe down. It is one of the most difficult and challenging things to do if done right.

For those of you thinking that quality reviews cost \$0.50 or less, you reap what you sow. They are expensive for a reason.

G Reviews - August 14, 2026, 2:24 pm

I've spent these last 2 weeks perfecting the warming system for the accounts. Now we break these

down by the following characteristics:

- 1.) Where is the account located? (City/State)
- 2.) Localized searches (including weather and local sports teams), normal interests, and something related to the future gmb's that this account will be reviewing later. In this way we can establish a genuine interest factor long before we drop the reviews.
- 3.) Matching subsequent youtube video interests. This combines both local searches for common every day topics as well as gmb related topics of interest but without directly naming the gmb's. This way we establish our behavioral patterns early and most effectively.
- 4.) CTR - When doing warming google searches we want to show that we're real people and so we perform a certain number of actual click-through rates to third party websites. This not only increases the amount of valid cookies that google can see, but it increases the trust factor entirely.

Notes: Never do a direct gmb review without browser history first. A cold review is 70% likely to get hit either during google's weekly review check or their first week of the month purge period. Thinking your safe if you pass the weekly test is a false assumption. Google's first week of every month purge (generally happens between the 2nd-4th of every month) is by far the most powerful purge google does. And this purge is not about review text. Rather it is about behavioral activity of the reviewers account itself. If google cannot establish that you are qualified to make the review you made, your review will disappear. And this purge is based on hundreds of different items like:

- 1.) Has this account ever expressed an interest in this topic?
- 2.) Has this account plausibly interacted with this business?
- 3.) Is there a correlation between this account and the business that was reviewed?
- 4.) What are the prior behavioral patterns of this account?
- 5.) Is this account local to the business being reviewed?

There are these and a thousand more behavioral comparisons google does at the time a review is submitted. And this is alongside the security and bot checks that google does with Bot Guard and Recaptcha Enterprise v3. Most people think that captchas are just there to protect from bots. That is only 10% true. What recaptcha is actually doing is watching for the order in which you do things, whether your clicks are teleported (by e1.click) to determine whether the click is trustable, whether there was actual humanized mouse movements similar to human patterns, ip evaluations, your TLS hash, your stack hash (determining whether your browser has automation or is real), and this is where the JA4 hash is checked. And on this side there are another couple hundred checks for legitimacy. You have no idea how many checks google can do in a matter of milliseconds.

The system is now around 98% finished and very effective.

G Reviews - August 27, 2026, 8:23 am

After extensive testing, re-testing, tweaking, correcting, and everything else over the past several months, the system is now production ready and spinning out extremely high-quality reviews including quality text, photos matching the business type, matching the gmb's physical city, and 5 stars. I couldn't have created anything better and I'm very pleased with the overall results. I have everything lined up perfectly now for exactly what Google's BotGuard, Recaptcha Enterprise v3, Gemini, and JA3/JA4 hashes match perfectly to our stack.

~HF~ - September 8, 2026, 2:53 pm

Messaged you on telegram just now

G Reviews - September 8, 2026, 3:33 pm

(Sep 8, 2026 02:53 PM)~HF~ Wrote: [CITE](#) Messaged you on telegram just now
Thank you. I've responded back to you.

G Reviews - October 7, 2026, 10:22 am

Over the last month, many new things have been developed to further enhance the processing of Google Maps reviews. One of the most significant things was bulking up the Local Guide Level of accounts without doing reviews at all. There was a very little known method to do this effortlessly and without risking the account at all. The method I've deployed to do this completely complies with all of Google's rules but it was something almost every seo company and every independent farm almost always over looks or never even knew this method existed.

What I'm going to show you now was a relatively new account just a few days ago. The account only had 2 lifetime map reviews and nothing else. In a few hours I was able to raise that account to a Local Guide Level 5 without posting any additional map reviews or ratings. It's a quick and easy way to bulk up your reviewer accounts:

? loading encrypted image...

Why do we want a higher Local Guide Level? Because that gets us inherent trust from Google for our future reviews.