Facebook Has a Child Predation Problem
The platform can be quicker at recommending groups built around child predation than it is to remove them.

https://www.wired.com/story/facebook-has-a-child-predation-problem/

By Lara Putnam, Wired
March 13, 2022

WHILE TRYING TO map the extent and impact of place-based Facebook groups where QAnon and allied disinformation spread, I went looking for Facebook groups with names including 10, 11, or 12. This was part of my work with the Pitt Disinformation Lab, and I was thinking of the 10th, 11th, or 12th wards of the city of Pittsburgh. What appeared instead was a group named "Buscando novi@ de 9,10,11,12,13 años." Looking for a 9-year-old girlfriend? What?

The page's aesthetic was cartoon cute: oversized eyes with long lashes, hearts, and pastels. The posts that made explicit references to photographed genitalia were gamified and spangled with emoticons: "See your age in this list? Type it into the replies and I'll show 'it' to you."
Most often posts were just doorways to connection, the real danger offstage. "Looking for a perverted girlfriend of 11," read one post, with purple background and heart emojis. Replies asked for friend requests to continue via Messenger, or offered entry to private groups or WhatsApp chats—away from the eyes of even a digital passerby.

This was not some outlaw 8Chan message board. It was cheerfully findable on Facebook. And, I began discovering in alarm, it was not the only one. Indeed, as late as January 2022—three months into my efforts to get action taken against them—if I searched 11, 12, 13 on the platform, 23 of the first 30 results were groups targeting children of those ages, with group names that included the words boyfriend/girlfriend, novio/a, or niños/niñas, sometimes along with 'pervertidos,' 'hot,' etc. They totaled over 81,000 members.

You may have assumed that 18 years in, Facebook (now part of Meta) would have basic checks in place so that creating a group whose name announces the goal of seeking children for intimate contact triggers scrutiny. Especially since, according to Facebook's own policies, no one under 13 is supposed to be on the platform at all. Everyone interacting in such a group is by definition a child violating Facebook policies by being on Facebook, an adult violating Facebook policies by impersonating a child, or an adult openly acting as an adult as they violate Facebook policy (and multiple state and international laws) by seeking sexualized contact with children.

Surely due diligence would dictate proactive steps to prevent the creation of such groups, backed up by quick action to remove any that get through once they are flagged and reported. I

would have thought so. Until I stumbled into these groups and began, with rising disbelief, to find it impossible to get them taken down.

Children are sharing personal images and contact information in a sexualized digital space, and being induced to join private groups or chats where further images and actions will be solicited and exchanged.

Even as debate over Congress' Earn It Act calls attention to the use of digital channels to distribute sexually explicit materials, we are failing to grapple with a seismic shift in the ways child sexual abuse materials are generated. Forty-five percent of US children aged 9 to 12 report using Facebook every day. (That fact alone makes mockery of Facebook's claim that they work actively to keep children under 13 off the platform.) According to recent research, over a quarter of 9- to 12-year-olds report having experienced sexual solicitation online. One in eight report having been asked to send a nude photo or video; one in 10 report having been asked to join a sexually explicit livestream. Smartphones, internet access, and Facebook together now reach into children's hands and homes and create new spaces for active predation. At scale.

OF COURSE I reported the group I had accidentally uncovered. I used Facebook's on-platform system, tagging it as containing "nudity or sexual activity" which (next menu) "involves a child." An automated response came back days later. The group had been reviewed and did not violate any "specific community standards." If I continued to encounter content "offensive or distasteful to you"—was my taste the problem here?—I should report that specific content, not the group as a whole.

"Buscando novi@ de 9,10,11,12,13 años" had 7,900 members when I reported it. By the time Facebook replied that it did not violate community standards, it had 9,000.

So I tweeted at Facebook and the Facebook newsroom. I DMed people I didn't know but thought might have access to people inside Facebook. I tagged journalists. And I reported through the platform's protocol a dozen more groups, some with thousands of users: groups I found not through sexually explicit search terms but just by typing "11 12 13" into the Groups search bar.

What became ever clearer as I struggled to get action is that technology's limits were not the problem. The full power of AI-driven algorithms was on display, but it was working to expand, not reduce, child endangerment. Because even as reply after reply hit my inbox denying grounds for action, new child sexualization groups began getting recommended to me as "Groups You May Like."

Each new group recommended to me had the same mix of cartoon-filled come-ons, emotional grooming, and gamified invites to share sexual materials as the groups I had reported. Some were in Spanish, some in English, others in Tagalog. When I searched for a translation of "hanap jowa," the name of a series of groups, it led me to an article from the Philippines reporting on efforts by Reddit users to get child-endangering Facebook groups removed there.

If your local mall had a whole section of storefronts advertising "Boys and girls 10, 11, 12 years old, come find your sexy romance here"—with open doors leading back into a warren of hidden photo booths—and the mall owners set up a free on-demand shuttle service to pick up any child at any time—would we shrug and say, oh well, nothing to be done? Blame the parents, look away?

The problem is that the social media platforms that are shaping our expanded connectivity (and sometimes, [subsidizing it](#), as Facebook has by providing limited free internet service in some developing markets) create exactly the kind of semipublic, semiprivate spaces where we know child endangerment happens. Some 10 percent of children who are victims of sexual abuse are abused by strangers, another 30 percent by family members. The majority, though, are abused by acquaintances: people who have occasion for repeated contact that builds trust and emotional leverage, who can create opportunities to move out of the public eye to behind closed doors. Facebook groups—and the ecosystem of private chats and channels they feed into—allow strangers to become acquaintances, at scale, with private rooms only a click away.

The recommendations showed that AI fed by internal data recognized exactly the group characteristics I had recognized—capturing patterns of predation across language and region for frictionless boost. Meanwhile, as Facebook's recommendation engines function like a seamless Uber for abusers, the safety side functions like the DMV circa 1990: manual data entry, inaction as default.

Backchannel outreach to a person with connections inside Facebook was the one step that seemed to get action. That person took my concern seriously. A week later the largest groups began disappearing. But within months, new ones just as large had replaced them. Most recently, the largest of this latest wave of groups have disappeared, either taken down or taken to "secret group" status—it's impossible to know which. We're back to a smaller number of groups—content identical to the old ones, and again growing steadily despite my repeated reports of them through Facebook's "safety" tool.

My efforts may have had nothing to do with even the limited takedowns that have occurred. Who knows? There is zero transparency, which is part of the problem. Screenshots I took may be the only external evidence of dozens of groups with thousands of members and hundreds of engagements daily that flourished on Facebook for months unaddressed. ("We do not tolerate child exploitation, including child sexual abuse material or inappropriate interactions between adults and minors, on our platforms," a Meta spokesperson wrote in a statement to WIRED. "We encourage anyone who sees content they think breaks our rules to report it using our in-app reporting tools.")

As an outside researcher, it's not only that I don't have access to Facebook's vast data flows or algorithmic specs. The company doesn't even share the most basic information on the pace of creation, scale, or takedown of public, private, or secret groups. Moreover, as academics we are governed by ethics and rules that limit what materials our research assistants can be exposed to: or what identifiable images we can store, all the more so when dealing with

children, who are considered categorically incapable of providing informed consent to researchers.

To say that Facebook's interactions with children are not governed by such niceties is an understatement.

Even here, in what shouldn't be an edge case in the slightest—groups built around sexual grooming of children too young to be on the platform at all—Facebook is neither proactively set to prevent harm nor consistent in acting when flagged. That tells us more than any press release about how the balance between engineering for protection and engineering for expansion is working in practice, and it should make us very afraid.

Facebook is desperate to attract more young users. It cannot afford to lose the rising generation to TikTok. Mark Zuckerberg's vision for Meta as virtual reality emporium leans into the lure of multiplayer games. The ease through which gamification pulls children in is on sickening display in the groups I have seen. How will outside eyes even know how dangerous the metaverse becomes?

Recent proposals like the Platform Accountability and Transparency Act, drawing on frameworks developed by Brookings and others, would mandate some basic information access that would be the first step toward accountability. But given the interlocking complexity of mutable algorithms and stacked internal policy choices that determine how platforms actually work, effective external regulation seems far less attainable than revolt from within. (This seems to be the intuition behind new initiatives like the Integrity Institute.)

If public shaming is the best route available, we better figure out how to ramp it up fast. I'm a US academic with institutional backing, time to spend, and some public platform, and I found it impossible to get sustained action against these groups. How are parents in Tamaulipas or South Texas going to get traction against the predators reaching into their children's lives—or against the company without which these opportunities for harm would not exist?

I've found that if you talk about child sexual predation by strangers on the internet loudly enough, concerned friends will start telling you you are sounding like a QAnon believer. It's worth pausing to think about. Collective panics about children in danger recur repeatedly in history: I've written about the surge of fears of witchcraft and child blood sacrifice in the early 20th-century Caribbean. The specific details of these panics are not meaningless. Rather, they reflect genuine fears, often projecting onto a single group of supposed evildoers what is in reality a much more diffuse pattern of vulnerability.

Believing in enemies you can act against can feel empowering at least. Sitting with the knowledge that no one knows how to stop the mall from hell next door is just terrifying.