Episode Book

EPISODE #6399

Smart Home Emergency Buttons and Mental Labor

PUBLISHED	1 November 2026 at 14:08 UTC
DURATION	17:53
LISTEN	myweirdprompts.com/episode/smart-home-mental-labor-50947/
HOSTS	Corn, Herman, Desmond, Hilbert
SCRIPT MODEL	anthropic/claude-sonnet-5.5
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.26
SOURCES	24 cited by the research (listed at the end)
BOOK GENERATED	2 November 2026 at 06:29 UTC

SYNOPSIS

A two-tier Home Assistant setup divides true household emergencies from everyday coordination tasks. Separating rare panic buttons from chore relay panels prevents both alert fatigue and invisible push notification throttling. Meanwhile, delegating domestic chores like daycare pickups or groceries to language models raises deeper questions: whether smart home agents actually ease mental labor or merely automate the default division of invisible household tracking.

TRANSCRIPT



Corn

We did the pager episode a while back, and the case for a dedicated, rare channel held up. Now Daniel's gone and wired it into the house. He wants a two-tier Home Assistant build. Tier one is an emergency button with no AI in it: lights come on, a critical notification hits Hannah's phone, done. Tier two is a daily panel where one tap says "I'll do pickup today," an agent relays it to her and updates the calendar.



Corn

He calls that second part "dystopian and weird," and he's firm that the agent isn't replacing talking. It's there to clear the small stuff so the real conversations get room. He also wants to know what existing tools he could build on.



Herman

It's buildable, which is the dull part. What I care about is whether the second tier lowers the load or just makes the same work quicker.



Corn

Before that, why two tiers at all? One button, different urgency levels.



Herman

Because the emergency channel has to stay rare. It's his focus-work idea applied to a house. Important signals impossible to miss, everything else silent. But in a house two people disagree about what counts as important. Pickup is urgent to the person at the daycare gate. It isn't an emergency.



Corn

So the lights come on when somebody's running ten minutes late.



Herman

Yes, and then the lights mean nothing. Within a month Hannah would be sighing at the ceiling every time they flicked on. There's a technical reason too, and it's the best argument for his split.



Corn

Go on.



Herman

Android's default notifications won't ring a stationary phone until the screen turns on. Doze mode. A phone on a desk for an hour just sits there quietly. To override it you set the priority to high and the time to live to zero. Then there's the alarm stream channel, which makes a sound even when the phone is on vibrate or silent.



Corn

Fine. So you use those every time and it always rings.



Herman

No. Firebase throttles you if you hammer high-priority pushes, and nobody can tell you what trips it. A community thread says almost those exact words. So the more you cry wolf, the less the real alert gets through. Reliability depends on rarity.



Corn

So it's not just that she'd get annoyed. The phone itself starts ignoring you.



Herman

Right, two separate failures stacked. The human tunes it out, and the infrastructure quietly downgrades it. And you'd never see the second one happen. You'd find out on the one night it mattered.



Corn

A smoke alarm that goes off for toast eventually gets the battery pulled.



Herman

Same thing, yes. And the battery pull is invisible too. Nobody announces it. The alarm just sits on the ceiling looking like it works. iOS is stronger out of the box. Critical notifications, since iOS 12, sit above everything on the lock screen and sound even through Do Not Disturb and mute. Is Hannah on an iPhone? I can't remember.



Corn

I don't know either. I'd guess Android, just because Daniel's the one asking.



Herman

Doesn't matter. The notification is the backup anyway. The lights are the real signal. A phone in a bag can fail, but a whole flat lighting up in the middle of a work call is hard to ignore.



Corn

Unless she's on camera. Then the lights are the least of it.



Herman

Ha. Fair, though a call is exactly when you'd want it. There's a good precedent for the whole thing. A blogger at rastrillo.ca built a toddler panic button with an IKEA SOMRIG Zigbee button and an input boolean helper. Pressing it fires a critical notification with a Deactivate action.



Corn

And if nobody deactivates it?



Herman

That's the clever bit. A second automation waits a minute, and if panic mode is still on, it emails an emergency contact through the carrier's email-to-SMS gateway. A dead-man's switch. He trained his four-year-old with role-play and tests the button whenever he tests the smoke detectors.



Corn

Testing a panic button the way you test a smoke detector. I'd have forgotten by the second month.



Herman

Which is why it belongs on the same schedule. One reminder, one evening, every button and every alarm. Hardware is cheap, too. A Tuya four-button Zigbee switch, an IKEA shortcut button, Aqara. A forum thread has Daniel's pattern: long-press any button and every light in the house comes on, short-press cancels.



Corn

Cancel is the part that matters. Ezra's fifteen months old. At some point he'll find a button on a shelf.



Herman

Long press only, then. Anything he can slap should do nothing.



Corn

And if he does manage a long press, the cancel has to be easy for a half-asleep adult.



Herman

Yes. Panic buttons fail in both directions. Too hard to trigger, or too easy to trigger and too hard to stand down. The second one is how you train people to disable the whole system.



Corn

How does that play out in practice? Somebody just unplugs the hub?



Herman

Usually something quieter. They mute the notification channel, or pull the batteries from the button "just for tonight." Then it's a Thursday three weeks later and nobody remembers it's off. So the cancel path has to be so painless that nobody is ever tempted.



Corn

Right, before we get into the pager, let's have Desmond's Money Tip.

D

Desmond

Desmond's Money Tip. Reuse one sheet of aluminium foil for a full calendar year. A roll costs about two pounds and holds roughly thirty sheets, so a sheet a day would run you twenty four pounds a year. Reusing one sheet saves you very nearly all of it, call it twenty three pounds and change, which over a working lifetime is a bit over seven hundred pounds.

D

Desmond

I have a sheet in the kitchen drawer that has been in service since two thousand and nineteen. It has gone grey and it has gone crisp and it has developed a smell I will not describe. My wife has stopped asking.

D

Desmond

Money is mostly just things you did not do, and then one day it is a pension.



Corn

I'm going to try it.



Herman

Corn, no. That's appalling.



Corn

Back to the panel, then. This is the part I'm suspicious of.



Herman

Same. The panel isn't a messaging problem, it's a mental-load problem. Researchers define mental labor as the anticipating and monitoring: knowing it's a pickup day, deciding whose turn it is, noticing the daycare bag is empty. A one-tap "I'll do pickup" removes the cost of telling her. It doesn't touch the cost of noticing.



Corn

So if Daniel has to remember to press it, nothing's changed. He's got a button now.



Herman

A very nice button for the same invisible work. Where the agent could earn its place is the side effects. Update the calendar, check whose turn it is, flag that pickup collides with a meeting. A plain templated notification already relays the message, so the AI adds nothing there.



Corn

Give me the collision case. What does that look like on an actual Tuesday?



Herman

Daniel taps "I'll do pickup." The agent reads the calendar, sees he has a client call running to half past four, and the gate closes at four forty-five with a twenty-minute drive. It doesn't refuse. It writes the pickup in, tells him the timing is tight, and tells Hannah what's been claimed. That's a side effect a template can't produce, because it needs to read something and compare.



Corn

And Hannah sees that too? Not just "Daniel's got it"?



Herman

That's a design choice, and I'd say yes. "Daniel's got pickup, tight with his call" lets her decide whether to quietly build in a backup. Without it she's either trusting blindly or checking, and checking is the monitoring work we're trying to remove.



Corn

How does the plumbing work? Daniel will want to know.



Herman

Home Assistant documents the loop. A chat message event triggers an automation, the message goes to the conversation agent through the conversation process action, and the reply goes back. You can define custom intents with descriptions that guide the language model, and an intent script can call the calendar's get events service and return a structured answer. So "I'll do pickup" becomes a named intent that checks the calendar and writes to it.



Corn

And if the agent gets it wrong? Books pickup on a day she's already doing it?



Herman

Then you've got two people at the gate. Herman the retired pediatrician says that's a confidently wrong system, and somebody's accountable. Keep the agent's job narrow and make it confirm back what it wrote.



Corn

I'd want it to say what it did. A machine that does the work while nobody's watching is my ideal, and that's exactly why I'd want it checked.



Herman

You've spent a lot of time thinking about this.



Corn

I've thought about it for about as long as I've been awake today.



Herman

So, the literature. A review in Sex Roles in twenty twenty-three went through thirty-one articles on gendered mental labor. And a Croatian team surveyed three hundred sixty-nine employed mothers on whether AI tools ease that load or reinforce the inequality. I don't know their results. I only know the framing, and it's the right one.



Corn

Which is?



Herman

If the agent tracks whose turn it is, it might actually redistribute. If it just receives what the one who remembers tells it, it automates the same default. There's a counterargument from a paper in Merits this year. It says generative AI can recognize and redistribute routine support work that's always been feminized and unseen, a lineage back to Judy Brady's "I Want a Wife." Delegating support work is how professional success has always been enabled. Usually by a woman, invisibly.



Corn

So Daniel's "dystopian" worry is half right. Getting a wife-shaped assistant out of software is arguably the point.



Herman

And arguably the problem, if the household's wife was always the assistant. It depends who the agent works for.



Corn

Meaning whose instructions shape it.



Herman

Whose habits it learns, really. If Hannah has always been the one who remembers, and the agent is trained on what gets entered, it learns that she's the default and quietly offers her the next task. A fair setup would have it counting. Last six pickups, who did them, and it says so out loud.



Corn

Says it out loud how? A weekly scorecard sounds like a fight waiting to happen.



Herman

Not a scorecard, a mirror. Something like "Daniel did four of the last six pickups" appearing when the button's pressed. No verdict, just the count. People can argue with each other, but it's hard to argue with a tally they both can see.



Corn

Groceries is where it gets uncomfortable, though. He mentions growing into shared spaces and shopping.



Herman

Yes. A shared list is a shared artifact. A "we're out of milk" button is a template. The moment the agent notices you're low and decides what to buy, it's doing anticipation work, which is the actual mental load.



Corn

Is that bad? I'd like something that noticed the milk.



Herman

Then it's doing the work, and you'd better agree who gets to override it. Say it buys the oat milk Daniel likes and Hannah can't stand. Now there's a decision made by a system neither of them chose. My line is that the tool relays and records. It doesn't anticipate and decide, at least until you've both said you want it to.



Corn

Hannah would have to say so. Not Daniel on her behalf.



Herman

Exactly that. Now, what's out there? The self-hosted family organizer space is crowded. Yuvomi's the most adopted: over two thousand stars, MIT license, version two point eighty-three released this week. Twenty modules, tasks, calendar with two-way Google and CalDAV sync, shopping, pantry, a wall mode for a kitchen tablet, and an API tokens module. It runs on two hundred fifty-six megabytes of RAM and one SQLite file.



Corn

That's a lot of modules for a file.



Herman

Mantel is MIT too, pitched as a free private alternative to Skylight and Cozi. Then FamilyWell, OpenFamily, Kitchen Board and a handful more. Calendar, shared list, kitchen display, push. None of them advertises a one-tap "I'll do pickup" that an agent relays and writes to the calendar.



Corn

So the gap is the handshake, and he doesn't need a new app.



Herman

And that shrinks the build. Use Yuvomi or Mantel for the panel and the data, and add a thin Home Assistant layer for the button and the agent. The API tokens module is what makes that possible, because it lets the agent write to the calendar without screen-scraping anything.



Corn

Order of operations?



Herman

Emergency button first. Zigbee button, a helper, critical notification on the alarm stream, all lights on. Test it until you trust it, and test it again with the smoke detectors. Then three templated intents: I'll do pickup, I'll do drop-off, running late. Each fires a conversation call that updates the calendar and notifies her. Groceries only after the handshake is reliable.



Corn

Emergency button's a weekend. The intents are a week.



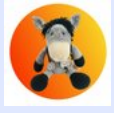
Herman

Probably, and that's not a startup.



Hilbert

"Solved" is the word I'd dispute. Nobody's solved the button. Your Zigbee thing is a shortcut button on a shelf.



Herman

Fair, I said solved about the plumbing.



Hilbert

Mm. I did one of these for my brother-in-law's mother, in the autumn. A button on the wall by her chair, and if nobody pressed the second button in ninety seconds it rang a neighbour, a man called Aldous Penhale, who once walked to Hebron Road and back in thirty-one minutes and still brings it up. I set sixty first. She needs ninety to get across the kitchen. Nobody writes that down in a manual.



Corn

Ninety seconds because of the kitchen floor.



Hilbert

The lino. Also she doesn't hurry for anyone. I keep a plastic clicker on my desk and press it when I've done something, and write the count on a pad. The phone keeps losing the number. It isn't connected to anything. A claw doesn't take to the tiny buttons anyway, so I bought the big one, four pounds fifty from the stationers on Jaffa Road.



Herman

Ninety seconds is a real design parameter. The whole dead-man's switch depends on a window that fits one person's kitchen.



Corn

And the rastrillo setup used a minute. Fine for a flat with a toddler in it, probably wrong for her.



Hilbert

Wrong for her, yes. Sixty gave us two false calls to Aldous in the first week. He was very gracious about it, which was worse.



Herman

That's the same failure as the lights, isn't it? Every false call spends a bit of Aldous's goodwill.



Hilbert

It does. A neighbour has only so many walks in him. The calendar was never the hard part. Pressing the button is. You can have the agent do the calendar if you like. Someone still has to remember to press. The people I worked with who were good at that kept a pad.



Corn

That's the noticing cost, and a pad beats a model at it.



Hilbert

The pad's got no opinion about what I did. It just has the count.



Herman

Which is the agent's job description, honestly. Record, don't decide.



Hilbert

Aldous still has the number of the button. It's his now. I haven't got the clicker's box either.



Corn

Something from the research that didn't fit. The Home Assistant assist satellite can start a conversation with a scoped system prompt, so one button press opens a pre-briefed chat instead of a blank one. That could be Daniel's pickup button.



Herman

The open question for me is social. The emergency channel only works if it's rare, so how do two people agree that it never gets used for logistics? That's not a software problem. Neither is who remembers to press the pickup button. Ninety seconds across a kitchen is a number I'd copy, though.



Corn

Thanks to Hilbert Flumingtop for producing. If this was your kind of episode, go back for episode four fifty-seven, Why a 90s Pager Beats Your Smartphone in an Emergency; and episode thirty-five, The Privacy Gap. This has been My Weird Prompts.



Herman

Find us at my weird prompts dot com, and a review helps.



Corn

See you tomorrow.

SOURCES

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

1. **Home Assistant Companion Docs, Critical notifications [primary]**
   https://companion.home-assistant.io/docs/notifications/critical-notifications/

2. **HA Companion Docs, Actionable Notifications [primary]**
   https://companion.home-assistant.io/docs/notifications/actionable-notifications/

3. **HA Companion Docs, Notification Commands [primary]**
   https://companion.home-assistant.io/docs/notifications/notification-commands/

4. **HA Notify integration docs [primary]**
   https://www.home-assistant.io/integrations/notify/

5. **HA docs, `conversation.process` incl. chat_message_received example [primary]**
   https://github.com/home-assistant/home-assistant.io/blob/current/source/_actions/conversation.process.markdown

6. **HA 2024.6 release notes, custom intents for LLM agents [primary]**
   https://github.com/home-assistant/home-assistant.io/blob/current/source/_posts/2024-06-05-release-20246.markdown

7. **HA AI agents for the smart home [primary]**
https://github.com/home-assistant/home-assistant.io/blob/current/source/_posts/2024-06-07-ai-agents-for-the-smart-home.markdown

8. **Yuvomi repo, v2.83.0 (2026-10-29) [primary]**
https://github.com/ulsklyc/yuvomi

9. **Rastrillo's Projects, toddler-friendly panic button (undated)**
https://rastrillo.ca/creating-a-toddler-friendly-panic-button-in-home-assistant/

10. **HA Community, Doze mode / Firebase throttling thread**
https://community.home-assistant.io/t/notifications-not-delivered-due-to-doze-mode/563525

11. **HA Community, emergency mode all-buttons automation**
https://community.home-assistant.io/t/automation-for-all-buttons/650728

12. **HA Community, Tuya TS0044 blueprint (updated 2025-12-01)**
https://community.home-assistant.io/t/zigbee2mqtt-tuya-4-button-scene-switch-ts0044-rdeangel/663302

13. **Yuvomi project site**
https://yuvomi.cloud/

14. **Mantel README**
https://github.com/heyitsmiike101/mantel/blob/main/README.md

15. **FamilyWell**
https://familywell.org/

16. **OpenFamily**
https://github.com/NexaFlowFrance/OpenFamily/

17. **FamilyHub**
https://github.com/joe-bor/FamilyHub/

18. **Betts Board**
https://github.com/Riley-D-Betts/betts-board

19. **Tribu**
https://itsdnns.github.io/tribu/

20. **Kitchen Board**
https://gitlab.com/g33kphr33k/kitchen-board

21. **Reich-Stiebert, Froehlich & Voltmer, *Sex Roles*, 2023**
https://doi.org/10.1007/s11199-023-01362-0

22. **Klasnić, Dremel & Adamović, *Sociologija i prostor*, 2026**
https://doi.org/10.5673/sip.64.1.5

23. **Klasnić, Dremel, Žažar et al., 2025, AI and mental labour**
https://doi.org/10.5937/snds25007k

24. **Kostihová & Makarevitch, *Merits*, 2026**
https://doi.org/10.3390/merits6020017