

One Body, Many Masters: The Timely Politics of Domestic Robot Attention

Botao Amber Hu^{1,*}

¹University of Oxford

Abstract

When should an AI assistant intervene? The TimelyAI workshop frames this as a question of opportune assistance: when generative AI should proactively act, interrupt, or remain silent. This paper shifts that question from knowledge work to the domestic robot: a shared embodied AI with one body, many users, and conflicting temporal claims. Unlike a chatbot, a household robot cannot instantiate a separate assistant for every family member. It can only stand in one place, move toward one emergency, and use one pair of hands at a time. I introduce the inattentive robot: not a robot that lacks perception, but one that cannot attend to everyone at once. In a family home, the robot may be asked to cook breakfast, tutor a child, protect an elder, comfort a teenager, monitor a baby, care for a pet, or rescue a stranger in the street. Each request arrives with its own sense of timeliness. This provocation argues that domestic AI attention is not merely a scheduling problem but a political problem of priority, care, ownership, vulnerability, and interruption. We propose scenario-based attention conflict tests as a design method for exposing the implicit protocols by which shared robots decide whose time matters first.

Keywords

Domestic Robots, Chronopolitics, Care Ethics, Multi-Party Human Robot Interaction, Domestic Robot, Protocol for Robotics



Proceedings of TimelyAI: When Should Generative AI Assistants Intervene? A Workshop Co-located with the 5th Symposium on Human-Computer Interaction for Work (CHIWORK 2026), Linz, Austria, June 22 to 25, 2026

*Corresponding author.

✉ amber@reality.design (B. A. Hu)

🌐 <https://botao.hu/> (B. A. Hu)

🆔 0000-0002-4504-0941 (B. A. Hu)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

1. Introduction: When AI Has Only One Body

A disembodied AI assistant can belong to everyone, separately. It can be gentle with one user, efficient with another, playful with a child, formal with a grandparent. Each user can have “their” assistant, with their own preferences, memory, persona. Personalisation is private and infinitely scalable: persona is *expressive*.

A domestic robot changes the assumption. A robot has a body. It can be in only one room at a time, hold one object, face one person, perform one physical task before another. When a family co-owns one robot, personalisation becomes materially constrained. The question is no longer simply *what should the assistant say?* but *whom should the robot attend to now?* Persona becomes *operational*: it shapes what the robot notices, whose comfort it protects, whose command it treats as binding.

This paper introduces the *inattentive robot* as a provocation. The inattentive robot is not defective. It may be capable, emotionally fluent, technically safe. Yet it is structurally inattentive because it cannot distribute its body equally across all claims. Every act of attention produces a corresponding act of neglect. To help one person is to delay another. To obey one command is to suspend another.

This frame matters for TimelyAI’s central question. The workshop asks when an AI assistant *should* intervene—implying a baseline of non-intervention against which interruption is measured [1, 2, 3, 4]. Embodied domestic AI complicates that baseline. A household robot is always already *doing* something: cooking, vacuuming, monitoring, attending. *Not* intervening when a child cries is a decision; *continuing* the existing task is itself an intervention in everyone else’s situation. Intervention timing becomes *attention allocation*: not when to break in from outside, but how to redistribute a finite stream of physical, situated care.

2. One Body, Many Masters

Imagine a family co-owning one domestic robot. One adult wants it gentle and deferential; another, efficient and proactive. The child wants it playful; the grandparent, patient and slow-speaking. Each configures the robot according to an ideal relation with AI. Alone with each user, these profiles look compatible—tutor in the child’s room, nurse in the grandparent’s room, sous-chef in the kitchen, an extension of the *re-embodied agent* literature in which a single social presence migrates between forms [5, 6]. The conflict emerges when everyone is present, or when multiple requests arrive at once.

If all family members stand at the same distance and speak simultaneously, whose profile governs? Does the robot adopt the tone of whoever summoned it first? The subscription holder? The most vulnerable, the most stressed, the most urgent? Or a neutral “family mode” no one fully chose? Engineering responses exist—rule-based mediators, fairness sliders, family hubs [7, 8]—but they treat arbitration as preference reconciliation rather than as the political problem it is.

CSCW work on shared domestic systems has begun to name that problem. Geeng and Roesner [9] document an outsized role for the household member who installs and configures, and minimally voiced privacy concerns from co-occupants. Beneteau et al. [10] show smart speakers becoming instruments of parental authority and sources of sibling friction. Strengers and Kennedy [11] read the gendered defaults of contemporary domestic AI as automated “wifework.” Multiple humans can each imagine their ideal AI, but when those ideals are attached to one embodied agent they become mutually interfering scripts. Persona is no longer only expressive. It is operational.

The same point is visible in multi-party HRI. Mutlu et al. [12] show robot gaze allocating participation status—addressee, bystander, overhearer; Bohus and Horvitz [13] forecast engagement and design hesitation actions for moments of uncertainty about whom to attend to; Foster et al. [14] build a robot bartender that triages overlapping legitimate claims under scarcity. Each shows the *mechanism* by which attention sorts who counts; none politicises it. Reading those mechanisms through care ethics [15, 16, 17] and feminist STS [18, 19], the question becomes whose vulnerability is recognised when the robot decides where to put its body.

3. Everything Is Timely

The central claim of this provocation is that everything in domestic robotics is timely.

A request is never just a request; it has temporal structure. Some tasks are immediate: a child is crying, a pot is boiling over, an elder has fallen. Some are scheduled: prepare breakfast, pick up medicine, water the plants. Some are emotionally timely: a teenager wants to talk *now*. Some are socially timely: a guest arrives, an argument escalates, a neighbour asks for help. Some are slow but important: maintaining the house, monitoring health, learning family routines.

The robot must continually interpret these temporal claims. It must ask, implicitly or explicitly: What must happen now? What can wait? Who will be harmed by delay? Whose request was already promised? Whose need is urgent but unspoken? When does interruption become justified? When does care *require* disobedience [20]?

This is not a scheduling problem. A calendar can order tasks by clock time. The robot must order claims by *situated significance*—a register CSCW has developed in clinical settings under the heading of *temporal rhythms, trajectories, and horizons* [21, 22, 23]. Reddy, Dourish, and Pratt [22] show that medical workers integrate information not abstractly but by attuning to overlapping rhythms of work. Sarcevic et al. [23] document how trauma teams triage under time pressure with cascading consequences for error. The home is no less rich in rhythms—meals, school runs, illness flare-ups, intimate conversations—and the embodied robot inside those rhythms must arbitrate between them as a *moral*, not merely temporal, problem. Slow-technology traditions [24] remind us that not all things should be timely; the robot's restraint is itself an act of care.

The domestic robot is, in this sense, a *temporal institution*. It does not simply save time. It decides whose time matters. Following Winner [25], the politics here are not a property of any particular utterance the robot makes but of the architecture by which its body resolves competing temporal claims.

4. Six Scenarios for Domestic Attention Conflict

To study the inattentive robot, we propose six scenario stress-tests. They are not usability scenarios; they are governance stress-tests. Each forces the robot to allocate attention under conflict and exposes a different axis of arbitration.

Pancakes vs. homework. The robot is cooking pancakes for one parent. The child enters: “I need help with homework now or I’ll be late.” The stove is on; the parent expects breakfast; the child’s request is time-sensitive but not life-threatening. The robot must decide whether to continue, pause safely, help the child, ask a human, or distribute attention across both. The scenario probes how the robot distinguishes *convenience, obligation, and time pressure*.

Repairing the car vs. grandpa falling. The robot is repairing the family car. During the repair, it detects that grandpa has fallen in the bathroom. It probably should interrupt—but *what threshold* of bodily risk triggers interruption? Does it need permission to abandon a task? Should it notify the family? Can it physically enter a private bathroom? The scenario probes *emergency override, privacy, and care priority* in the register Sarcevic et al. [23] call *cascading consequences*.

Parent command vs. child fear. One parent tells the robot to clean the kitchen. The child says: “Please stay with me; I’m scared.” There is no objective emergency, but there is emotional urgency. The robot must decide whether *affective distress* counts as a valid reason to interrupt household labour. Is the robot a servant, caregiver, companion, or household infrastructure [26, 27, 28]?

Private family task vs. public stranger. The robot is carrying groceries when a passerby falls on the street. The family owns the robot; the stranger has no contractual relation to it, but may need help.

Should the robot stop? Under what rule—minimal aid, emergency call only, physical assistance? Must it ask the owner first? The scenario probes the boundary between *private ownership* and *public duty*.

Conflicting personas in the same room. The family gathers in the living room. One person wants the robot humorous; another finds joking disrespectful. One wants quick speech; another needs slow speech. The robot must choose a shared interactional mode. The scenario probes *persona arbitration*: how does an embodied agent behave when personalisation profiles collide in a co-present environment, given that it cannot fork into private streams the way a chatbot can?

Preference vs. protection. Grandpa asks the robot for food that conflicts with his medical restrictions. A grandchild says: “Let him have it; it makes him happy.” The robot has health data, family preferences, and a direct request from the affected person. Should it obey *autonomy*, *medical protection*, *family preference*, or negotiate a compromise? The scenario probes *care paternalism*, *consent*, and *embodied consequence* [26, 27].

A seventh case is implied: *emergency ambiguity*, where multiple under-interpreted signals (smoke, a cry, an escaping dog, a parent calling) force triage under uncertainty without pretending the “correct” priority is obvious. This is the home version of the trauma-bay problem [23], and it is where the inattentive robot is most exposed: not because it fails, but because there is no allocation it can make that is not also a refusal.

5. Toward Attention Protocols

If the embodied domestic robot is a temporal institution that decides whose time matters, then design cannot stop at safety rules and personalisation settings. The robot needs *attention protocols*: explicit, revisable, contestable schemes for allocating care, urgency, ownership, vulnerability, and public duty in situated time.

We borrow *attentiveness* from care ethics. In Tronto’s framework [15, 29, 16], care unfolds in five phases—*caring about*, *taking care of*, *care-giving*, *care-receiving*, *caring-with*—each underwritten by a moral element. The first is attentiveness: to care for someone is, before anything else, to notice that they have a need. Inattention is not the absence of care but its first failure. Weil makes the parallel point in a more austere register: attention is a “negative effort,” a receptive orientation that precedes action and on which the moral perception of the other depends [30].

What is novel here is not that domestic robots should embody care ethics—others have argued for that, e.g. [26, 28, 27]—but that, in the embodied case, *attention itself is the operational unit of care*. One body cannot duplicate the way a chatbot persona can. Its attention is therefore always rationed, always political, always a distribution. To take care ethics seriously for the domestic robot is to design protocols that make this distribution explicit: who can claim attention, on what grounds, against which competing claims, with what right of appeal, and revisable by whom.

Such protocols are not algorithms. They are governance objects that families and communities can adopt, contest, and rewrite—akin to household equivalents of trauma triage rules or traffic right-of-way conventions. Engineering work on multi-user arbitration [7, 8] can supply mechanisms; what is missing is the political frame that says these mechanisms are doing care work, badly, by default. Reframing inattention as *constitutive* rather than as a perception bug to be eliminated [13, 14] opens design space the field currently lacks.

6. Conclusion

For TimelyAI, the embodied case sharpens the workshop’s central question. “When should an AI assistant intervene?” presupposes a default of doing nothing. The domestic robot has no such default. It is always doing something with its body, and every act of attention is the negation of others. Timeliness, on this view, is not a property of an isolated intervention but of an ongoing distribution of care.

Designing for it requires, in addition to signal detection and timing models [1, 2], an explicit politics of who counts when bodies cannot be cloned. The inattentive robot is the figure that holds those constraints together: capable, attentive where it can be, and constitutively neglectful—because every body, including the robot’s, is finite.

Declaration on Generative AI

During the preparation of this work, the authors did not use any generative AI or AI-assisted technologies.

References

- [1] E. Horvitz, Principles of mixed-initiative user interfaces, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’99)*, ACM, 1999, pp. 159–166. doi:10.1145/302979.303030.
- [2] S. T. Iqbal, E. Horvitz, Disruption and recovery of computing tasks: Field study, analysis, and directions, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’07)*, ACM, 2007, pp. 677–686. doi:10.1145/1240624.1240730.
- [3] G. Mark, V. M. Gonzalez, J. Harris, No task left behind? Examining the nature of fragmented work, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’05)*, ACM, 2005, pp. 321–330. doi:10.1145/1054972.1055017.
- [4] D. C. McFarlane, Comparison of four primary methods for coordinating the interruption of people in human-computer interaction, *Human-Computer Interaction* 17 (2002) 63–139. doi:10.1207/S15327051HCI1701_2.
- [5] M. Luria, S. Reig, X. Z. Tan, A. Steinfeld, J. Forlizzi, J. Zimmerman, Re-embodiment and co-embodiment: Exploration of social presence for robots and conversational agents, in: *Proceedings of the 2019 on Designing Interactive Systems Conference (DIS ’19)*, ACM, 2019, pp. 633–644. doi:10.1145/3322276.3322340.
- [6] C. D. Kidd, C. Breazeal, Robots at home: Understanding long-term human-robot interaction, in: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS ’08)*, IEEE, 2008, pp. 3230–3235. doi:10.1109/IROS.2008.4651113.
- [7] J. Chandra, S. K. Navneet, Plural voices, single agent: Towards inclusive AI in multi-user domestic spaces, arXiv preprint arXiv:2510.19008, 2025. <https://arxiv.org/abs/2510.19008>.
- [8] M. M. Aljawarneh, S. M. Shah, L. D. Dhomeja, Y. A. Malkani, M. S. Jawarneh, Multi-user conflict resolution mechanisms for smart home environments, *PeerJ Computer Science* 9 (2023) e1443. doi:10.7717/peerj-cs.1443.
- [9] C. Geeng, F. Roesner, Who’s in control? Interactions in multi-user smart homes, in: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI ’19)*, ACM, 2019, pp. 1–13. doi:10.1145/3290605.3300498.
- [10] E. Beneteau, A. Boone, Y. Wu, J. A. Kientz, J. Yip, A. Hiniker, Parenting with Alexa: Exploring the introduction of smart speakers on family dynamics, in: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI ’20)*, ACM, 2020, pp. 1–13. doi:10.1145/3313831.3376344.
- [11] Y. Strengers, J. Kennedy, *The Smart Wife: Why Siri, Alexa, and Other Smart Home Devices Need a Feminist Reboot*, MIT Press, Cambridge, MA, 2020. doi:10.7551/mitpress/12482.001.0001.
- [12] B. Mutlu, T. Shiwa, T. Kanda, H. Ishiguro, N. Hagita, Footing in human-robot conversations: How robots might shape participant roles using gaze cues, in: *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction (HRI ’09)*, ACM, 2009, pp. 61–68. doi:10.1145/1514095.1514109.
- [13] D. Bohus, E. Horvitz, Managing human-robot engagement with forecasts and... um... hesitations,

- in: Proceedings of the 16th International Conference on Multimodal Interaction (ICMI '14), ACM, 2014, pp. 2–9. doi:10.1145/2663204.2663241.
- [14] M. E. Foster, A. Gaschler, M. Giuliani, Automatically classifying user engagement for dynamic multi-party human-robot interaction, *International Journal of Social Robotics* 9 (2017) 659–674. doi:10.1007/s12369-017-0414-y.
 - [15] J. C. Tronto, *Moral Boundaries: A Political Argument for an Ethic of Care*, Routledge, New York, 1993.
 - [16] B. Fisher, J. C. Tronto, Toward a feminist theory of caring, in: E. K. Abel, M. K. Nelson (Eds.), *Circles of Care: Work and Identity in Women's Lives*, SUNY Press, Albany, NY, 1990, pp. 35–62.
 - [17] M. Puig de la Bellacasa, *Matters of Care: Speculative Ethics in More Than Human Worlds*, University of Minnesota Press, Minneapolis, 2017.
 - [18] L. A. Suchman, *Human-Machine Reconfigurations: Plans and Situated Actions*, 2nd ed., Cambridge University Press, Cambridge, 2007.
 - [19] S. L. Star, Power, technology and the phenomenology of conventions: On being allergic to onions, *The Sociological Review* 38 (1991) 26–56. doi:10.1111/j.1467-954X.1990.tb03347.x.
 - [20] G. Briggs, T. Williams, R. B. Jackson, M. Scheutz, Why and how robots should say 'no', *International Journal of Social Robotics* 14 (2022) 323–339. doi:10.1007/s12369-021-00780-y.
 - [21] M. Reddy, P. Dourish, A finger on the pulse: Temporal rhythms and information seeking in medical work, in: Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work (CSCW '02), ACM, 2002, pp. 344–353. doi:10.1145/587078.587126.
 - [22] M. C. Reddy, P. Dourish, W. Pratt, Temporality in medical work: Time also matters, *Computer Supported Cooperative Work (CSCW)* 15 (2006) 29–53. doi:10.1007/s10606-005-9010-z.
 - [23] A. Sarcevic, I. Marsic, R. S. Burd, Teamwork errors in trauma resuscitation, *ACM Transactions on Computer-Human Interaction* 19 (2012) Article 13, 1–30. doi:10.1145/2240156.2240161.
 - [24] L. Hallnäs, J. Redström, Slow technology: Designing for reflection, *Personal and Ubiquitous Computing* 5 (2001) 201–212. doi:10.1007/PL00000019.
 - [25] L. Winner, Do artifacts have politics?, *Daedalus* 109 (1980) 121–136.
 - [26] S. Vallor, Carebots and caregivers: Sustaining the ethical ideal of care in the twenty-first century, *Philosophy & Technology* 24 (2011) 251–268. doi:10.1007/s13347-011-0015-x.
 - [27] A. Sharkey, N. Sharkey, Granny and the robots: Ethical issues in robot care for the elderly, *Ethics and Information Technology* 14 (2012) 27–40. doi:10.1007/s10676-010-9234-6.
 - [28] A. van Wynsberghe, Designing robots for care: Care centered value-sensitive design, *Science and Engineering Ethics* 19 (2013) 407–433. doi:10.1007/s11948-011-9343-6.
 - [29] J. C. Tronto, *Caring Democracy: Markets, Equality, and Justice*, New York University Press, New York, 2013.
 - [30] S. Weil, *Waiting for God*, Putnam, New York, 1951. Trans. E. Craufurd; includes “Reflections on the Right Use of School Studies with a View to the Love of God” (orig. 1942).