

Bloom Preview: A Low-Cost LLM-Powered Social Robot

Xiangfei Kong
Rex Gonzalez
Nicolas Echeverria
Sofia Cobo Navas
Thuc Anh Nguyen
RARE Lab
University of South Florida
Tampa, FL, USA

Divyamshu Shrestha
Sebastian Ramirez-Vallejo
Deekshita Senthil Kumar
Andrew Teixeira
Thao Nguyen
RARE Lab
University of South Florida
Tampa, FL, USA

Deep Akbari
Hong Wang
Hongdao Meng
Zhao Han
zhaohan@usf.edu
RARE Lab
University of South Florida
Tampa, FL, USA

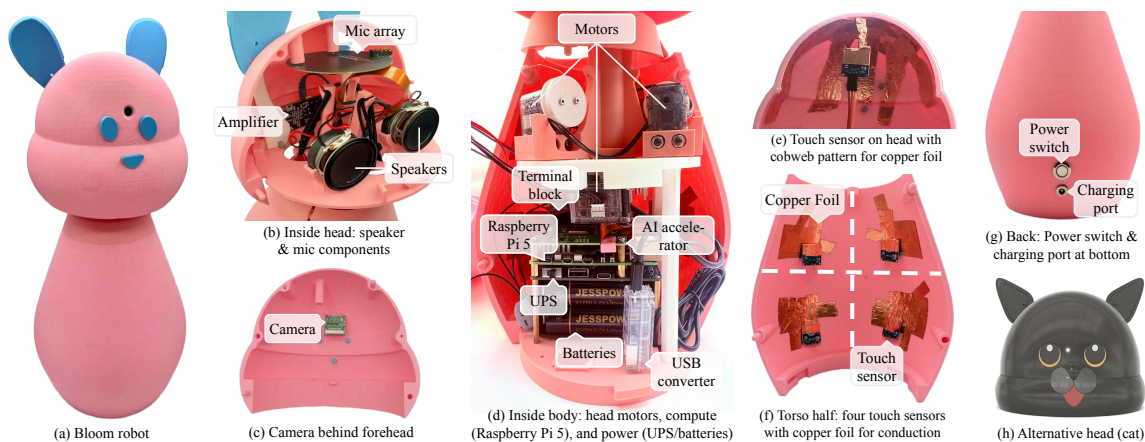


Figure 1: Bloom Robot Overview: Shell design (a), internal components (b-f), and a customized cat head example (g)

ABSTRACT

HRI research shows that social robots support human companionship and emotional well-being. However, the cost of high-end social robots limits their accessibility, and, yet, existing low-cost platforms cannot fully support users emotionally due to limited interaction capabilities. In this demo paper with accompanying video, we showcase Bloom, a low-cost social robot that combines touch sensing with LLM-driven conversation to facilitate more expressive and emotionally engaging interactions with humans. Bloom also integrates a customizable 3D-printed shell with a flexible software pipeline that enables touch-responsive movement, LLM-driven dialogue, and face tracking during conversations, offering more expressive interaction capabilities at a low cost, which remains easy to fabricate and adaptable across a wide range of societal applications, including mental health, companion, and healthcare.

CCS CONCEPTS

• **Computer systems organization** → *Robotics*; • **Computing methodologies** → *Artificial intelligence*.

Please use nonacm option or ACM Engage class to enable CC licenses
This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.
HRI' 26 Companion, March 16–19, 2026, Edinburgh, Scotland, UK
© 2026 Copyright held by the owner/author(s).
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>



KEYWORDS

social robot, human-robot interaction (HRI)

ACM Reference Format:

Xiangfei Kong, Rex Gonzalez, Nicolas Echeverria, Sofia Cobo Navas, Thuc Anh Nguyen, Divyamshu Shrestha, Sebastian Ramirez-Vallejo, Deekshita Senthil Kumar, Andrew Teixeira, Thao Nguyen, Deep Akbari, Hong Wang, Hongdao Meng, and Zhao Han. 2026. Bloom Preview: A Low-Cost LLM-Powered Social Robot. In *Companion of the 2026 ACM/IEEE International Conference on Human-Robot Interaction (HRI' 26 Companion)*, March 16–19, 2026, Edinburgh, Scotland, UK. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Social robots are greatly beneficial, e.g., supporting companionship [5] and improving emotional well-being [4] in a wide range of application domains such as education [2], healthcare [3], and entertainment [9]. However, the high cost remains a financial barrier that limits their accessibility to the much-needed target population. Even compact desktop social robots are not affordable: 64cm QTrobot (USD \$12k), 41cm Furhat (\$28k), and 35cm Misty (\$8k).

To democratize social robots, recent efforts are made to low-cost platforms like Blossom (\$250) [8] and Reachy Mini (\$449/\$299 with/without compute). However, they remain limited in engaging capabilities like affective interaction through touch and usefulness with built-in adaptive conversational intelligence. Touch offers an immediate and intuitive channel for emotional connection [11],

Table 1: Feature Comparison of Low-Cost Social Robots

Feature	Blossom	Reachy Mini Lite	Bloom Lite (ours)
Compute (WiFi)	–	–	51
Speaker	–	51	51
Mic	–	51	51
Camera	–	51	51
Head Movement	51	51	51
Total	\$250	\$299	\$315 (\$260 no compute)

Table 2: Cost for Three Versions of Bloom

Feature	Component	Cost	Lite	Regular	AI
Compute	Raspberry Pi 5 + SD	\$55	51	51	51
Head movement	4 × Servo Motor	\$124	51	51	51
Audio	U2D2 (USB converter)	\$37	51	51	51
	5W Stereo Speakers	\$7	51	51	51
	Amplifier & DAC	\$10+\$6	51	51	51
	2-/4-Mic Array	\$12/\$90	51	51 (2-)	51 (4-)
Vision	Camera	\$35	51	51	51
Touch	Touch Sensors + Copper	\$9	51	51	51
General	PLA + Wires + Hardware	\$20	51	51	51
Power	2-/4-Cell UPS	\$27/\$45	-	51 (2-)	51 (4-)
	2/4 Batteries	\$12/\$24	-	51 (2)	51 (4)
	AC Adapter + Hardware	\$19	-	51	51
AI	AI HAT+ & Active Cooler	\$70+\$5	-	-	51
Total			\$315	\$373	\$556

providing comfort and reducing negative feelings [12]. Meanwhile, modern generative AI, beyond scripted dialogue [6], offers context-aware, adaptive, emotionally attuned responses.

In this work, we thus developed Bloom (Fig. 1). Leveraging Blossom [8]’s head actuation and internal structures, Bloom’s novelty lies in contributing a touch-sensitive and holdable hard-shell body with critical multimodal features including audio, vision, affective touch, and modern LLM-driven AI conversation capabilities. Maintaining a low cost starting from \$310 and competitive with Blossom and Reachy mini (Table 1), the new features still make it an accessible platform that democratizes robot companions for a broader impact. To further demonstrate the usefulness, we developed three applications to show the generalizability of Bloom in three much-needed domains in our society: college student mental health, elderly companion, and healthy recipe recommendation.

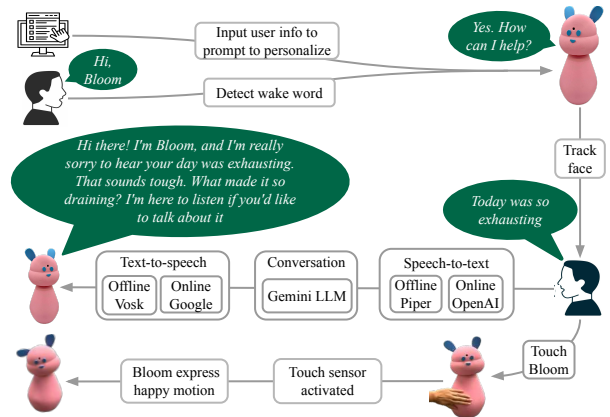
2 BLOOM ROBOT

Bloom is a social robot with a rigid, customizable 3D-printed shell optimized for holdability, touch. With a 191mm diameter, it stands 410mm tall with ears or 370mm without. There are three versions: *Bloom Lite* is the lowest-cost configuration with touch sensing, expressive motion, and LLM-driven conversation. *Bloom Regular* adds power management for wireless use. *Bloom AI* further integrates an AI accelerator for local and real-time computer vision tasks like facial tracking for engagement. Table 2 shows the bill of materials.

2.1 Enabling Hardware

Compute: Bloom has a Raspberry Pi 5 (Pi) single-board computer for its compact form factor, cost (\$50 vs. NVIDIA Jetson Orin Nano’s \$249), and compute power (2.4GHz quad-core CPU).

Motors: Inherent from Blossom, Bloom has three compact DY-NAMIXEL XL-320 servo motors for head orientation and one for

**Figure 2: Overview of Bloom’s software and behavior pipeline with one of three applications: elderly companion.**

head turning. They feature daisy-chaining to facilitate wiring and use a U2D2 converter to connect to Pi’s USB port.

Audio: Bloom uses a ReSpeaker microphone array for voice detection, direction of arrival estimation, beamforming, and noise reduction, allowing the robot to locate a speaker and turn its head toward the voice to convey attentive listening. For clear audio output, we designed a stereo system with a UDA1334 digital-to-analog converter (DAC), a PAM8406 amplifier and two 5W speakers.

Vision: Bloom uses a 12MP Camera Module 3 Wide with high-dynamic range, autofocus, and a 120° lens to track users’ faces.

Touch: Bloom has nine TTP223B capacitive touch sensors via an ESP32 microcontroller and copper foil for increased conduction through the robot shell. Eight sensors are on two torso halves (see Fig. 1f) and one is on the head with cobweb pattern for coverage.

Power: Bloom uses a Waveshare UPS battery management system for wireless use. It features continuous charging and we also added a power switch and charging port on Bloom’s back.

AI Accelerator: To improve AI power, we added an AI HAT+ with a 13-TOPS Hailo-8 AI accelerator. To avoid overheating, we used an Active Cooler with an aluminum heatsink and a fan.

2.2 Software and Robot Behavior

Bloom’s software stack integrates sensing, wake word, speech, face tracking, touch behavior, LLM-powered conversation, and a web interface to support interactions, as shown in Fig. 2.

Speech-to-Text (STT) and Wake Word: Bloom uses a dual speech-to-text pipeline that supports both lightweight offline Vosk run locally and more accurate online Google Speech recognition. This design enables seamless backend switching and robust voice interaction regardless of noise level or network availability. The transcribed text is then passed to the LLM for Bloom’s conversational response. To activate Bloom, users say “Hi, Bloom”. We trained a model using the EfficientWord-Net wakeword engine with multiple voices and varied pronunciations of “Bloom” improves activation accuracy. Unlike picovoice on our first try, EfficientWord-Net does not require online accounts.

Text-to-Speech (TTS): Similar to our STT design, Bloom uses natural OpenAI Text to Speech (TTS) for online synthesis and Piper

for fast offline speech synthesis and so it can adapt to different network conditions at deployment. To maintain responsiveness, we used audio streaming for Bloom to speak while receiving speech responses. Bloom also speaks “thinking” phrases like *um* for higher network latency to show users that it hears them.

Conversation Capability Powered by LLM: Bloom generates responses by calling the Gemini 2.5 Flash Lite, and its conversational behavior is guided by prompt designs following three prompt patterns [10]: *Persona*, *Meta Language Creation*, and *Context Manager*. *Persona* defines a warm companion that can express emotions through movement. *Meta Language Creation* ensures Bloom speaks in clear, everyday language to maintain accessibility. *Context Manager* helps Bloom track user input, maintain topic continuity, and respond appropriately while avoiding sensitive content.

Face Tracking: To create an interaction that conveys a sense of natural eye contact, we implemented face detection running on the AI accelerator, enabling real-time face tracking for Bloom to continuously follow and engage with interactants.

Touch Behavior: To enable touch behavior, Bloom integrates its touch sensing hardware with an event-driven software pipeline that translates touch into expressive and natural robot motions. Upon touch, the system triggers a callback that selects one of eleven happy head movements from Blossom’s movement library [8], each conveying positive affect and encouraging warm, companion-like interaction, e.g., slightly bouncing head.

Web Interface: To support personalized and contextually meaningful conversations, we implemented a web interface using Flask to quickly input user information and import it directly into Bloom’s LLM prompts, e.g., names, demographics, preferences, and interests.

3 APPLICATION DOMAIN DEMONSTRATIONS

To show the usefulness of Bloom, we developed three societal applications: college student storytelling for mental health, reducing loneliness in older adults, and recommending healthy recipes.

Storytelling for Mental Health: College students face increasing mental health challenges, yet many hesitate to seek professional support due to financial barriers or limited access [7]. Bloom’s friendly appearance and gentle expressive design, combined with empathetic LLM-driven conversation, offer accessible support that feels safe and personal. Through the storytelling prompt for college students, Bloom can help them reflect on their emotions, feel understood, and develop positive emotional habits.

Elderly Companion: Beyond college students, Bloom also suits older adults who often face social isolation: Many older adults experience loneliness as a result of reduced mobility, decreasing social networks, or living alone [1]. Bloom’s affective touch sensing and LLM-powered personalized conversation can offer a comforting companion in their daily life, encouraging older adults to experience social connection and emotional comfort.

Healthy Recipe Recommender: Bloom could also serve as an interactive household social robot to address the challenge parents face in balancing nutritional goals with convenience and children’s preferences. Specifically, the wireless feature allows Bloom to be freely taken to the kitchen or dining area to recommend healthy recipes during cooking or meal decision process. By combining this portable access with LLM-driven conversational intelligence, Bloom



Figure 3: Ongoing work to improve head actuation.

can provide personalized guidance to assist parents in balancing nutritional goals with convenience and children’s preferences.

4 NEXT STEPS AND CONCLUSION

To keep reducing latency, we are working on audio streaming for both offline and online STT, and having expressive thinking motions while the LLM generates the response. We are also increasing the accuracy of head movement, which currently uses Blossom’s tensile string connection to motors. After considering linear-actuator and rigid-link approaches, we selected the rigid-link approach due to feasibility and compatibility with existing servo motors. We have prototyped and tested this mechanism using 3D-printed links and updated ball-and-socket joints, show in Fig 3. We will continue internal structure changes to accommodate the improvement.

Looking forward, we are collaborating with industry partners to deploy Bloom in senior living communities to collect subjective feedback about its touch, dialogue, and overall experience. Bloom’s integration of affective touch, expressive movement, and LLM-driven conversation at low cost offers a novel contribution with meaningful real-world impact on social robot companionship.

REFERENCES

- [1] R. A. Ananto and J. E. Young. 2021. We can do better! an initial survey highlighting an opportunity for more HRI Work on loneliness. In *HRI Companion*.
- [2] T. Belpaeme, J. Kennedy, A. Ramachandran, B. Scassellati, and F. Tanaka. 2018. Social robots for education: A review. *Science Robotics* 3, 21 (2018).
- [3] Carlos A Cifuentes, Maria J Pinto, Nathalia Céspedes, and Marcela Múnera. 2020. Social robots in therapy and care. *Current Robotics Reports* (2020).
- [4] Jill A Dosso, Jaya N Kailley, Gabriella K Guerra, and Julie M Robillard. 2023. Older adult perspectives on emotion and stigma in social robots. *Frontiers in Psychiatry* (2023).
- [5] H. Lalwani, M. Saleh, and H. Salam. 2025. A study companion for productivity: exploring the role of a social robot for college students with ADHD. In *HRI*.
- [6] Emmanuel Pot, Jérôme Monceaux, Rodolphe Gelin, and Bruno Maisonnier. 2009. Choregraphe: a graphical tool for humanoid robot programming. In *RO-MAN*.
- [7] Nahal Salimi, Bryan Gere, William Talley, and Bridget Iriooogbe. 2023. College students mental health challenges: Concerns and considerations in the COVID-19 pandemic. *Journal of College Student Psychotherapy* (2023).
- [8] Michael Suguitan and Guy Hoffman. 2019. Blossom: A handcrafted open-source robot. *ACM THRI* (2019).
- [9] John Vilk and Naomi T Fitter. 2020. Comedians in cafes getting data: evaluating timing and adaptivity in real-world robot comedy performance. In *HRI 2020*.
- [10] J White. 2023. A prompt pattern catalog to enhance prompt engineering with ChatGPT. *arXiv preprint arXiv:2302.11382* (2023).
- [11] Christian JAM Willemsse and Jan BF Van Erp. 2019. Social touch in human-robot interaction: Robot-initiated touches can induce positive responses without extensive prior bonding. *International journal of social robotics* (2019).
- [12] R. Wullenkord, M. R. Fraune, F. Eyssele, and S. Šabanović. 2016. Getting in touch: How imagined, actual, and physical contact affect evaluations of robots. In *RO-MAN*.