

From Scanning Guidelines to Action: A Robotic Ultrasound Agent with LLM-Based Reasoning

Yuan Bi^{1,2*}, Yiping Zhou^{1*}, Pei Liu¹, Feng Li^{1,2}, Zhongliang Jiang³, and Nassir Navab¹

¹ Chair for Computer-Aided Medical Procedures and Augmented Reality, Technical University of Munich, Munich, Germany

yuan.bi@tum.com

² Munich Center for Machine Learning, Munich, Germany

³ Medical Intelligence and Robotic Cognition Lab, The University of Hong Kong, Hong Kong, China

Abstract. Robotic ultrasound offers advantages over free-hand scanning, including improved reproducibility and reduced operator dependency. In clinical practice, US acquisition relies heavily on the sonographer’s experience and situational judgment. When transferring this process to robotic systems, such expertise is often encoded explicitly through fixed procedures and task-specific models, yielding pipelines that can be difficult to adapt to new scanning tasks. In this work, we propose a unified framework for autonomous robotic US scanning that leverages a large language model (LLM)-based agent to interpret US scanning guidelines and execute scans by dynamically invoking a set of provided software tools. Instead of encoding fixed scanning procedures, the LLM agent retrieves and reasons over guideline steps from scanning handbooks and adapts its planning decisions based on observations and the current scanning state. This enables the system to handle variable and decision-dependent workflows, such as adjusting scanning strategies, repeating steps, or selecting the appropriate next tool call in response to image quality or anatomical findings. Because the reasoning underlying tool selection is also critical for transparent and trustworthy planning, we further fine tune the LLM agent using a RL based strategy to improve both its reasoning quality and the correctness of tool selection and parameterization, while maintaining robust generalization to unseen guidelines and related tasks. We first validate the approach via verbal execution on 10 ultrasound scanning guidelines, assessing reasoning as well as tool selection and parameterization, and showing the benefit of RL fine tuning. We then demonstrate real world feasibility on robotic scanning of the gallbladder, spine, and kidney. Overall, the framework follows diverse guidelines and enables reliable autonomous scanning across multiple anatomical targets within a unified system. Code: <https://github.com/yuan-12138/RUSSAgent>; Video: <https://youtu.be/pfM0c4e2IGA>.

Keywords: robotic ultrasound · embodied AI · LLM.

* These authors contributed equally to this work.

1 Introduction

Ultrasound (US) is widely used in clinical practice due to its real-time feedback, portability, and lack of ionizing radiation. However, free-hand scanning is highly operator dependent, with image quality varying with the sonographer’s expertise, which limits standardization across operators and sites and can lead to fatigue and work-related musculoskeletal injuries. Robotic US offers a promising alternative by providing precise and repeatable motion for more standardized and reproducible scanning, while integrated force sensing and control can regulate probe–tissue contact to improve safety and support consistent image quality [1,15].

Most prior approaches toward autonomous US scanning follow a rule-based design. These systems typically use an external RGB-D camera to reconstruct the patient surface and a robotic arm to manipulate the US probe. Scanning is then guided by anatomy-specific, hand-crafted protocols, often coupled with dedicated US segmentation models. Most prior approaches to autonomous US scanning are rule based. They typically use an external RGBD camera to reconstruct the patient surface [21] and a robotic arm to manipulate the probe, with scanning guided by anatomy-specific hand-crafted protocols, often combined with dedicated segmentation models. Representative systems include robotic spine US [32], autonomous carotid scanning with tailored protocols [13], and breathing compensation for aortic US [30], as well as additional components such as force-based vertebral localization [29] and pulsation-informed artery segmentation [12]. While effective, these rule-based designs are often task-specific and transfer poorly to new anatomies and scanning objectives.

Beyond hand-crafted rule-based designs, robot learning methods have been explored to acquire scanning behaviors from interaction [3] or expert demonstrations. Early work often trained reinforcement learning (RL) agents in simulation, such as RL navigation toward standard spine planes [19] and VesNetRL for carotid standard views [2]. More recent studies increasingly learn from expert demonstrations to mitigate sim-to-real gaps, including inverse RL to learn the “language of sonography” [14] and Diffusion Policy for robotic carotid scanning [7]. While effective in targeted scenarios, achieving flexible behavior and robust generalization across diverse scanning tasks remains challenging.

Large language models (LLMs) have demonstrated strong general-purpose capabilities and can benefit robotics by enabling flexible task understanding and high-level decision making [23]. In navigation and manipulation, LLMs are often integrated in a zero-shot manner as high-level planners for tool selection and long-horizon reasoning, while low-level controllers and perception modules execute the actions [25,26,22,27]. However, this transfer is less direct in the medical domain, where general-purpose LLMs may lack domain-specific knowledge. As a result, early attempts to integrate LLMs into robotic US have incorporated clinical scanning guidelines to appropriately constrain the planning process and improve reliability [31]. Chen *et al.* [6] proposed USPilot, an LLM-based framework for robotic US scanning across multiple organs. However, the reasoning capability of the LLM agent is often not explicitly modeled or evaluated in prior

work. As a result, existing systems typically emphasize guideline execution and tool use, while offering less insight into the agent’s intermediate reasoning that underlies its planning and tool selection decisions.

In this work, we propose a unified agentic framework for autonomous robotic US, in which an LLM interprets scanning guidelines and translates them into actions by dynamically selecting from a set of perception and control tools. By retrieving protocol steps from handbooks and reasoning over the current scanning state and observations, the agent can execute variable, decision-dependent workflows (e.g., repeating steps, adapting strategies, or adjusting force based on image quality) rather than following a fixed sequence. To strengthen reasoning over planning decisions, we further introduce a RL strategy to fine-tune the LLM agent for more reliable reasoning and tool-calling behavior. We evaluate the proposed framework through verbal interaction experiments on unseen US scanning protocols, and we further demonstrate end-to-end feasibility in real-world robotic experiments on human volunteers across three representative examinations: gallbladder, spine, and kidney.

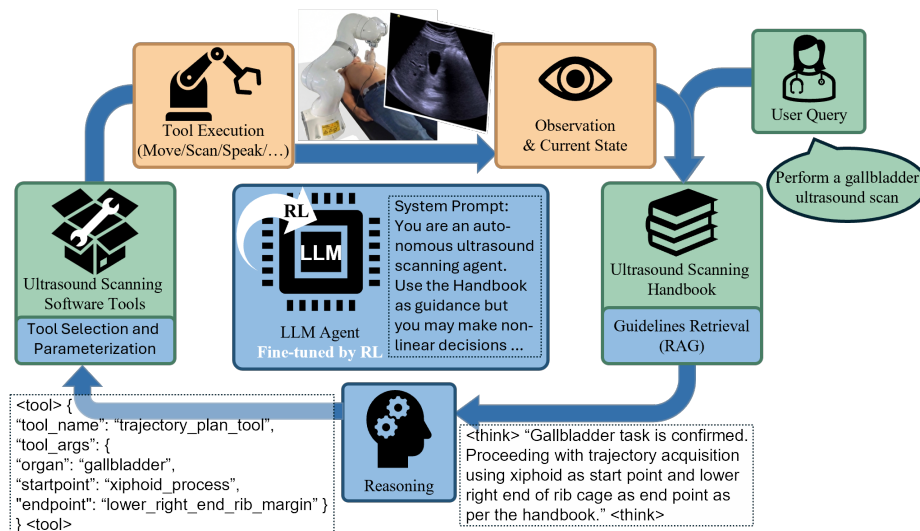


Fig. 1. Overview of the proposed guideline-driven agentic robotic US system. An RL fine tuned LLM retrieves guideline steps from the scanning handbook, reasons over the current state and observations, and generates structured tool calls with parameters. Executed tools update the state, enabling closed loop, non linear scanning workflows.

2 Methods

LLMs have strong reasoning and decision-making capabilities. We leverage an LLM as a high-level agent for autonomous US scanning by equipping it with soft-

ware tools for image processing, perception, and robot control, and by providing access to a database of scanning guidelines to supply domain knowledge. Conditioned on the user query and current observations, the agent retrieves guideline steps, selects and parameterizes tools, and executes the requested scanning task. We further apply RL-based fine tuning to adapt a small LLM to robotic US, improving planning-oriented reasoning and tool calling while still generalizing robustly to unseen guidelines and producing plausible reasoning traces to support transparency and clinician inspection.

2.1 Agentic LLM for Robotic US

Our proposed framework follows an agentic paradigm similar to [24,9,8], but targets autonomous robotic US scanning through guideline-driven planning and tool use rather than verbal or image analysis tasks. It consists of three components: (1) robotic and perception tools for path generation, patient-robot interaction, etc.; (2) a guideline database containing US workflow steps; and (3) an LLM that serves as the high-level planner. At initialization, a prefix prompt defines the LLM role for sequential problem solving and tool-based response generation. Given a user query, the system retrieves relevant guideline instructions via retrieval-augmented generation [18] and the agent selects and parameterizes tools accordingly. At each step, the agent reasons over the current observations and outputs a `<think>` reasoning trace followed by a `<tool>` token with a structured JSON tool call. Tool outputs are appended to the history to guide subsequent decisions, enabling variable, decision-dependent workflows such as repeating steps, adapting scanning strategies, and adjusting trajectory or force based on US observations until the guideline objectives are met.

2.2 Tools Design

In order to give the LLM the ability to perform automatic US scanning and analyses the scanning status, specific tools are designed for imaging processing and robot execution.




Fig. 2. Trajectory planning pipeline for initial scan generation.

Trajectory Planning Tool is designed to generate an initial US scanning trajectory based on the input query. As shown in Fig. 2, first, the RGBD camera mounted on the robot end effector is used to capture the patient body surface. Due to the limited field of view of the RGBD camera, multiple views are acquired and fused to reconstruct a whole body surface representation of the patient. A skeletal atlas is then registered to the patient using an open-source model [20,17], which enables the detection of anatomical landmarks for trajectory planning. Based on LLM-provided inputs specifying the target organ and the start and end points of the scan defined with respect to anatomical landmarks, the tool determines a scanning path and projects it into 3D space using the patient point cloud. In summary, the inputs to the tool include the target organ and landmark-referenced start and end points of the scanning trajectory, and the output is a 3D US scanning trajectory.

Robot Execution Tool receives either a scanning trajectory or a target pose as input, together with the execution speed. It then executes the corresponding motion command to move the robot to the specified pose or along the specified trajectory for US scanning or RGBD image acquisition.

Contact Adjustment Tool evaluates US image quality and probe contact quality based on the US confidence map [16], and then adjusts the robot pose to establish proper contact with the patient surface [5,4]. The input to this tool is the US image, and the output is adjusted robot pose or trajectory.

Voice Guidance Tool generates the verbal instructions needed to guide the patient during scanning. These instructions include breath-hold cues and body pose adjustment guidance. US scanning is a highly interactive procedure, and effective communication with the patient is essential for adjusting body posture and obtaining an appropriate acoustic window. For example, such guidance is particularly important for gallbladder scanning, because without proper instructions, such as a deep-inhalation breath hold, it is often difficult to visualize the gallbladder beneath the rib cage and avoid acoustic shadowing caused by bone. The input to the tool is the message that needs to be conveyed to the patient.

Trajectory Refinement Tool evaluates the results of the initial scan and refines the trajectory when necessary. Since the initial scan is planned only based on the registered skeletal atlas and anatomical priors, the target organ may not appear in the center of the US image. Therefore, after the initial scan, a target-organ-specific segmentation model is applied to localize the organ in 3D space, and the scanning trajectory is then replanned to enable a refined scan. If the organ is already sufficiently centered and well visualized in the initial scan, no refinement is performed. The input to the tool is the tracked US data from the initial scan, and the output is a replanned trajectory when it is necessary.

2.3 RL-based Fine-tuning

Given the response time constraints in robotic applications, using a very large LLM is often impractical. Although the proposed framework is not a vision language action model that requires a response frequency above 20 Hz, an agentic

system for robotic US still benefits from sufficiently fast response time on accessible GPU hardware. Therefore, fine tuning a smaller model is necessary to obtain reliable reasoning and tool calling behavior, since smaller models typically have more limited capacity and generalization than larger models. At the same time, the fine tuned model should still robustly generalize to other guidelines and related tasks.

To this end, we construct a guideline driven tool use dataset by collecting stepwise agent traces from a set of training US scanning guidelines. For each guideline step, a larger LLM is used to generate a reference response consisting of `<think>` reasoning and a `<tool>` JSON command with the correct tool and parameters. We first perform supervised fine tuning (SFT) on these reference traces to teach the small model the structured output format and basic tool use behavior. We then apply RL, which is increasingly used to improve reasoning behavior and response correctness in LLMs [10]. Specifically, we adopt Proximal Policy Optimization to further optimize the agent. We define a dense reward for each predicted tool call by comparing the structured output $\hat{y} = \{\hat{t}, \hat{\mathbf{a}}\}$ to a reference tool call $y = \{t, \mathbf{a}\}$, where t and $\hat{t}$ are tool names and $\mathbf{a}, \hat{\mathbf{a}}$ are JSON argument dictionaries. Let $\mathcal{A}$ and $\hat{\mathcal{A}}$ denote the sets of argument keys in $\mathbf{a}$ and $\hat{\mathbf{a}}$, and let $\mathcal{K} \subseteq \mathcal{A}$ be the set of scored keys with tool-specific weights w_k .

$$r(\hat{y}, y) = \begin{cases} -0.5, & \hat{t} \neq t \\ \text{clip}_{[0,1]}(0.1 + 0.9 s_{\text{args}}), & \hat{t} = t \end{cases} \quad (1)$$

This equation assigns negative reward if the predicted tool name is incorrect, and otherwise combines a small base reward with a dense argument score.

$$s_{\text{args}} = \text{clip}_{[0,1]}(0.5 s_{\text{pres}} + 0.5 s_{\text{corr}} - 0.1 s_{\text{extra}}) \quad (2)$$

This equation defines the dense argument score as a mixture of key presence and value correctness, with a penalty for spurious extra arguments.

$$s_{\text{pres}} = \frac{\sum_{k \in \mathcal{K}} w_k \mathbf{1}[k \in \hat{\mathcal{A}}]}{\sum_{k \in \mathcal{K}} w_k}, \quad s_{\text{corr}} = \frac{\sum_{k \in \mathcal{K}} w_k \mathbf{1}[\hat{a}_k \approx a_k]}{\sum_{k \in \mathcal{K}} w_k} \quad (3)$$

Here, s_{pres} rewards including the required argument keys, while s_{corr} rewards correct argument values, using exact matching for discrete fields and tolerance-based matching for selected numeric fields in the relation $\hat{a}_k \approx a_k$.

$$s_{\text{extra}} = \frac{|\hat{\mathcal{A}} \setminus \mathcal{A}|}{\max\{1, |\hat{\mathcal{A}}|\}}. \quad (4)$$

This term penalizes extra argument keys that are not expected by the reference tool call, encouraging concise and executable tool commands. The reward is defined based on tool calling correctness with respect to the reference tool calling JSON, assigning higher reward to outputs that select the correct tool and provide valid arguments, and penalizing incorrect tool selection, malformed JSON, or

incorrect argument values. The model is fine tuned with LoRA [11] using a rank of 16 for SFT and a rank of 8 for RL.

3 Experiments and Results

Verbal Interaction Experiments We create 20 US scanning guidelines and split them into 10 training guidelines and 10 held-out testing guidelines, with no overlap in guideline content. The tool set remains identical across training and testing. To obtain reference outputs, we use Qwen3-30B-A3B-Instruct-2507 [28] to generate guideline-conditioned agent traces consisting of `<think>` reasoning and `<tool>` JSON commands. We fine-tune Qwen3-4B-Instruct-2507 [28] on the training guidelines and evaluate on the held-out guidelines to assess robustness to unseen guideline instructions, measuring tool selection and parameterization accuracy together with structured output validity.

We then compare the base model and the RL fine tuned model under identical guideline and tool settings. Specifically, we report the overall scanning task success rate, defined as successful completion of the guideline specified scan. In addition, we evaluate step-level tool use by measuring, at each guideline step, whether the agent selects the correct tool and provides correct tool arguments. The results are shown in Tab. 1. Overall, RL fine tuning leads to a clear improvement across both task success and step-level tool calling accuracy.

Table 1. Performance comparison before and after fine-tuning.

	Base Model	After Fine Tuning
Step-wise Accuracy	0.6512	0.8973
Overall Success Rate	0.5384	0.9230

We also include qualitative examples comparing tool calls generated by the base model and the RL fine tuned model (Fig. 3), illustrating differences in tool parameterization and reasoning at representative guideline step. Qualitatively, the base model often produces minimal or uninformative reasoning traces, whereas the RL fine tuned model generates more coherent and task-relevant rationales alongside its tool calls.

Real-world Robotic Experiments We further validate the proposed system in real-world robotic US experiments on three human volunteers, demonstrating autonomous scanning of the gallbladder, spine, and kidney, with each organ scanned twice per volunteer. Ethical approval was obtained from the institutional review board of the Technical University of Munich, and all participants provided written informed consent. The scanning setup is shown in Fig. 2. We use a KUKA LBR iiwa robot equipped with a Siemens ACUSON Juniper US system and a 5C1 probe mounted via a 3D printed holder, with US images acquired through an Epiphan frame grabber. An RGBD camera, Azure Kinect by Microsoft, is

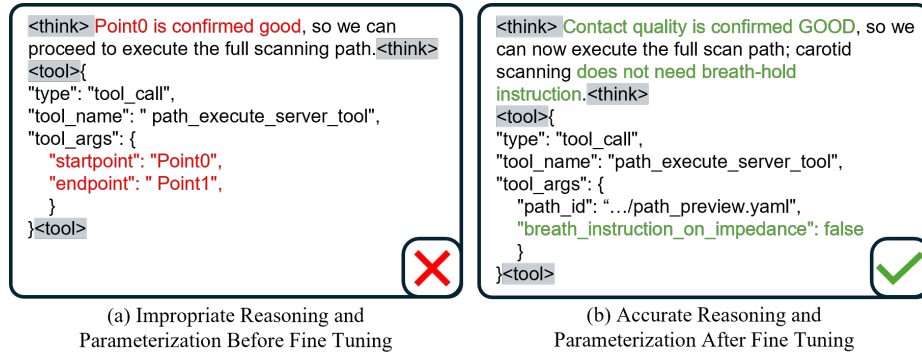


Fig. 3. Verbal Experiments Examples.

mounted on the robot end effector to capture the 3D surface of the scanning target. We evaluate task success based on whether the system can correctly visualize the target organ in the US image. One exemplary scanning scenario is shown in Fig. 4. All spine and kidney experiments were successful, resulting in 100% success rate. For gallbladder scanning, one case failed, resulting in a success rate of 83.3%. This is mainly due to the higher procedural complexity, as breathing-induced organ motion can shift the gallbladder out of the acoustic window, making the scan highly dependent on accurate breath-hold instructions and close patient-robot coordination.




Fig. 4. Real Scenario Experiments Examples.

4 Conclusion

In this work, we presented a guideline-driven agentic framework for autonomous robotic ultrasound that uses an RL fine tuned LLM to retrieve scanning guidelines, reason over the current scanning state, and generate structured tool calls for execution. Experiments on held-out guidelines and real-world volunteer studies across gallbladder, spine, and kidney scanning demonstrate that the proposed system can robustly follow diverse guidelines and achieve reliable autonomous scanning within a unified pipeline.

Acknowledgement This work was supported in part by the Multi-Scale Medical Robotics Center, AIR@InnoHK, Hong Kong.

Disclosure of Interests The authors have no competing interests to declare that are relevant to the content of this article.

Informed consent Informed consent was obtained from all individual participants included in the study.

References

1. Bi, Y., Jiang, Z., Duelmer, F., Huang, D., Navab, N.: Machine learning in robotic ultrasound imaging: Challenges and perspectives. *Annual Review of Control, Robotics, and Autonomous Systems* **7** (2024)
2. Bi, Y., Jiang, Z., Gao, Y., Wendler, T., Karlas, A., Navab, N.: Vesnet-rl: Simulation-based reinforcement learning for real-world us probe navigation. *IEEE Robotics and Automation Letters* **7**(3), 6638–6645 (2022)
3. Bi, Y., Qian, C., Zhang, Z., Navab, N., Jiang, Z.: Autonomous path planning for intercostal robotic ultrasound imaging using reinforcement learning. *Scientific Reports* **16**(1), 6356 (2026)
4. Bi, Y., Su, Y., Navab, N., Jiang, Z.: Gaze-guided robotic vascular ultrasound leveraging human intention estimation. *IEEE Robotics and Automation Letters* (2025)
5. Chatelain, P., Krupa, A., Navab, N.: Confidence-driven control of an ultrasound probe. *IEEE Transactions on Robotics* **33**(6), 1410–1424 (2017)
6. Chen, M., Fan, S., Cao, G., Liu, Y.h., Liu, H.: Uspilot: An embodied robotic assistant ultrasound system with large language model enhanced graph planner. *arXiv preprint arXiv:2502.12498* (2025)
7. Chen, R., Yan, X., Lv, K., Huang, G., Li, Z., Li, X.: Ultradp: Generalizable carotid ultrasound scanning with force-aware diffusion policy. In: *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 20074–20080. IEEE (2025)
8. Daghyani, M., Wang, L., Hashemi, N., Medhat, B., Abdelsamad, B., Velez, E.R., Li, X., Tsang, M.Y., Luong, C., Tsang, T.S., et al.: Echoagent: Guideline-centric reasoning agent for echocardiography measurement and interpretation. *arXiv preprint arXiv:2511.13948* (2025)

9. Fallahpour, A., Ma, J., Munim, A., Lyu, H., Wang, B.: Medrax: Medical reasoning agent for chest x-ray. *arXiv preprint arXiv:2502.02673* (2025)
10. Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, Q., Xu, R., Zhang, R., Ma, S., Bi, X., et al.: Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948* (2025)
11. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. *Iclr* **1**(2), 3 (2022)
12. Huang, D., Bi, Y., Navab, N., Jiang, Z.: Motion magnification in robotic sonography: Enabling pulsation-aware artery segmentation. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 6565–6570. IEEE (2023)
13. Huang, Q., Gao, B., Wang, M.: Robot-assisted autonomous ultrasound imaging for carotid artery. *IEEE Transactions on Instrumentation and Measurement* **73**, 1–9 (2024)
14. Jiang, Z., Bi, Y., Zhou, M., Hu, Y., Burke, M., Navab, N.: Intelligent robotic sonographer: Mutual information-based disentangled reward learning from few demonstrations. *The International Journal of Robotics Research* **43**(7), 981–1002 (2024)
15. Jiang, Z., Salcudean, S.E., Navab, N.: Robotic ultrasound imaging: State-of-the-art and future perspectives. *Medical image analysis* **89**, 102878 (2023)
16. Karamalis, A., Wein, W., Klein, T., Navab, N.: Ultrasound confidence maps using random walks. *Medical image analysis* **16**(6), 1101–1112 (2012)
17. Keller, M., Werling, K., Shin, S., Delp, S., Pujades, S., Liu, C.K., Black, M.J.: From skin to skeleton: Towards biomechanically accurate 3d digital humans. *ACM Transactions on Graphics (TOG)* **42**(6), 1–12 (2023)
18. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.t., Rocktäschel, T., et al.: Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems* **33**, 9459–9474 (2020)
19. Li, K., Wang, J., Xu, Y., Qin, H., Liu, D., Liu, L., Meng, M.Q.H.: Autonomous navigation of an ultrasound probe towards standard scan planes with deep reinforcement learning. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 8302–8308. IEEE (2021)
20. Li, Z., Liu, J., Zhang, Z., Xu, S., Yan, Y.: Cliff: Carrying location information in full frames into human pose and shape estimation. In: *European Conference on Computer Vision (ECCV)*. pp. 590–606. Springer, Cham (2022)
21. Nardi, D., Lamon, E., Fontanelli, D., Saveriano, M., Palopoli, L.: An anatomy-aware shared control approach for assisted teleoperation of lung ultrasound examinations. *IEEE Robotics and Automation Letters* **11**(3), 2570–2577 (2026)
22. Rana, K., Haviland, J., Garg, S., Abou-Chakra, J., Reid, I., Suenderhauf, N.: Say-plan: Grounding large language models using 3d scene graphs for scalable robot task planning. *arXiv preprint arXiv:2307.06135* (2023)
23. Raptis, E.K., Kapoutsis, A.C., Kosmatopoulos, E.B.: Agentic llm-based robotic systems for real-world applications: a review on their agenticness and ethics. *Frontiers in Robotics and AI* **12**, 1605405 (2025)
24. Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Hambro, E., Zettlemoyer, L., Cancedda, N., Scialom, T.: Toolformer: Language models can teach themselves to use tools. *Advances in neural information processing systems* **36**, 68539–68551 (2023)
25. Shah, D., Osiński, B., Levine, S., et al.: Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action. In: *Conference on robot learning*. pp. 492–504. PMLR (2023)

26. Singh, H., Das, R.J., Han, M., Nakov, P., Laptev, I.: Malmm: Multi-agent large language models for zero-shot robotic manipulation. In: 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 20386–20393. IEEE (2025)
27. Stone, A., Xiao, T., Lu, Y., Gopalakrishnan, K., Lee, K.H., Vuong, Q., Wohllhart, P., Kirmani, S., Zitkovich, B., Xia, F., et al.: Open-world object manipulation using pre-trained vision-language models. arXiv preprint arXiv:2303.00905 (2023)
28. Team, Q.: Qwen3 technical report (2025), <https://arxiv.org/abs/2505.09388>
29. Tirindelli, M., Victorova, M., Esteban, J., Kim, S.T., Navarro-Alarcon, D., Zheng, Y.P., Navab, N.: Force-ultrasound fusion: Bringing spine robotic-us to the next “level”. IEEE Robotics and Automation Letters **5**(4), 5661–5668 (2020)
30. Velikova, Y., Azampour, M.F., Simson, W., Esposito, M., Navab, N.: Implicit neural representations for breathing-compensated volume reconstruction in robotic ultrasound. In: 2024 IEEE International Conference on Robotics and Automation (ICRA). pp. 1316–1322. IEEE (2024)
31. Xu, H., Wu, J., Cao, G., Chen, Z., Lei, Z., Liu, H.: Transforming surgical interventions with embodied intelligence for ultrasound robotics. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 703–713. Springer (2024)
32. Yang, C., Jiang, M., Chen, M., Fu, M., Li, J., Huang, Q.: Automatic 3-d imaging and measurement of human spines with a robotic ultrasound system. IEEE Transactions on Instrumentation and Measurement **70**, 1–13 (2021)