MEDICAL ROBOTS

The grand challenges of learning medical robot autonomy

Pierre E. Dupont[1]* and Alperen Degirmenci[2,3]

Most medical robots depend on human operators for sensing, decision-making, and action during procedures. Future progress depends on enabling robots to take on these capabilities. Although learning-based approaches provide remarkable promise toward achieving this goal, notable challenges must be addressed to unlock these robots' full potential in clinical settings.

INTRODUCTION

Medical robots have been in clinical use for more than a quarter of a century, but the number of robots deployed remains limited compared with their widespread integration in other industries. Although the largest medical robotics firm, Intuitive Surgical, has sold ~7500 robots over 30 years performing 17 million cumulative surgeries (*1*), this is dwarfed by the 15 million traditional laparoscopic procedures performed annually (*2*) and pales in comparison with the more than 1 million autonomous robots deployed by Amazon (*3*), where robots handle more than 1 billion packages per year (*4*). Why the disparity? Industrial robots leverage autonomy to augment productivity, whereas most surgical robots serve as teleoperated extensions of the surgeon.

So why not just automate medical robots? The challenge is that the classical tools used to implement autonomy only work for tasks that are simple and highly structured. Outside of medicine, successful robotic automation often means adapting the product and process to match these limitations—not just trying to make a robot do exactly what a human would do. This involves simplifying the task motions and minimizing the sensing and modeling requirements. For example, in warehouse automation, because it is challenging to pick up individual parts from a bin, warehouse robots often pick up shelves and follow grid lines on the warehouse floor to carry them to a human packing the boxes (*5*). As a second example, products to be assembled by robots are designed with snap-fit connectors instead of screws to simplify assembly motions.

In medicine, it is not possible to redesign the patient, and it is risky to radically change surgical techniques that have been developed over many decades. Consequently, existing examples of medical robot automation correspond to procedures that naturally fit the capabilities of classical autonomy. Noncontact procedures, comparable to a robot painting a car, are one such category. For example, the Accuray Cyberknife is an automated radiosurgery robot that moves around the patient delivering radiation to an internal tissue target on the basis of a precomputed plan (*6*). As a second example, in automated corneal flap creation and reshaping during LASIK (laser-assisted in situ keratomileusis), a robotic laser follows preplanned programs that compensate for eye motion.

Tasks requiring simple interactions with tissue can also be automated. An example is the automated selection and harvesting of follicular units for subsequent transplant during hair restoration (*7*). A third category of procedures amenable to automation is those that involve reshaping hard tissue in a manner comparable to computer numerical control machining. For example, in hip and knee replacement, robots with bone milling tools are used to create the cavities for the implants. Unlike the previous examples, these systems use shared autonomy, with the clinician holding the robot arm and guiding it through its motions (*8*). The robot prevents the clinician from moving outside the preplanned cutting zone, which enables milling bone cavities with higher accuracy than can be achieved by hand.

Up to the present, extending automation to more general surgical tasks has proven daunting because classical automation tools depend on deriving accurate and computationally efficient models of the system to be automated. Although such models are feasible for milling bone, it is much more challenging to accurately model the grasping and cutting of soft tissue in real time. The recent introduction of learning-based methods is transformative because these methods enable handling higher levels of complexity without necessitating analytical models of the underlying processes.

THE STRUCTURE OF AUTOMATION

To understand how learning-based methods can affect medical robot autonomy, Fig. 1 compares teleoperation, classical autonomy, and learned autonomy. Autonomy itself can be defined as a cycle of three modules: sense, think, and act (*9*), which are in some fields also called perception, planning, and control (*10*, *11*). In this cycle, the sense module performs state estimation—building a description of the current condition of the robot, patient, and procedure from sensor data. On the basis of this perceived state, the think module interprets the current state description in the context of the procedural plan and decides what the robot should do next. The act module converts proposed actions into the associated robot motor commands.

In teleoperation (Fig. 1A), the human operator is responsible for carrying out the sense-think-act cycle. They interpret real-time imaging to ascertain the spatial relationship between the robotic instruments and anatomy. On the basis of this understanding, the operator mentally plans subsequent actions and anticipates their outcomes. They act

---

[1]Department of Cardiovascular Surgery, Boston Children's Hospital, Harvard Medical School, Boston, MA 02115, USA. [2]NVIDIA Corporation, Santa Clara, CA 95051, USA. [3]Harvard John A. Paulson School of Engineering and Applied Sciences, Allston, MA 02134, USA.
*Corresponding author. Email: pierre.dupont@childrens.harvard.edu

by manipulating the input devices of the surgical console to command the robot's tool movements. The only automation consists of low-level physics-based models used by the sense and act modules for robot state estimation and motion control.

In classical automation (Fig. 1B), most of the thinking occurs in the preprocedural planning phase. Here, the clinician uses graphical software running optimization algorithms to develop a procedural plan. Intraoperatively, the role of the think module is limited to small adjustments in the plan, e.g., to account for patient motion. The human operator is no longer responsible for intraoperative sensing and action. Instead, these modules use conventional image and signal processing, analytical kinematic and dynamics models, and control theory. An exception is orthopedics, where the human operator moves the cutting tool collaboratively with the robot.

In learned autonomy (Fig. 1C), some or all of the hand-engineered algorithms of the sense-think-act cycle are replaced with data-driven, learned components. Alternatively, a single end-to-end learned model may be used that maps sensor inputs directly to control outputs [e.g., PilotNet for autonomous driving (*12*) or recent large-scale robotics initiatives, such as the π_0 vision-language-action model (*13*)]. A currently popular approach is to use separate, learned think and act modules that operate hierarchically at different time horizons [as proposed in (*14*)]: A fast (>100 Hz) reactive policy (system 1) serves as the act module, whereas a slow (<10 Hz) reasoning policy (system 2) serves as the think module [e.g., Figure Helix (*15*) and NVIDIA GR00T (*16*)]. Although classical autonomous system modules communicate using physical representations (e.g., positions, distances, and velocities), learning-based modules communicate via information-rich vectors embedded in a self-learned latent space. Using decoders and encoders, these representations can be converted to and from natural language and graphical representations to enable human interpretability and oversight. The clinician retains a critical supervisory role, reviewing and validating the diagnosis and proposed navigation plan, monitoring procedural progress, and issuing corrective instructions as needed through a natural language interface.

THE CHALLENGES OF LEARNED AUTOMATION

Replacing classical, model-based control with learned autonomy alleviates the burden of precise system modeling and identification, which can be infeasible for complex scenarios. However, it also introduces substantial challenges inherent to data-driven methods. Developing and deploying learned systems within the safety-critical context of medicine necessitate addressing several fundamental challenges around data, operational robustness, human-robot collaboration, and safety.

Data: Quality, diversity, and volume

In learned autonomy, engineering complexity shifts substantially from explicit system modeling toward collection, labeling, and curation of datasets for training and evaluation. Learned methods generate input-output mappings entirely on the basis of training data; therefore, data quantity must be sufficient to capture task complexity, data quality must ensure that the mapping accurately reflects surgical reality, and, perhaps most critically, data diversity must be adequate for robust generalization and safe handling of less common but clinically important scenarios (e.g., anatomical variations). Generating such data presents major challenges. Simulation can yield large data volumes (both for training and testing) but often struggles with the sim-to-real gap (*17*), where a mismatch between the real and synthetic data distributions degrades system performance. Expert demonstrations used for imitation learning can provide high-quality data, but acquiring and annotating a sufficient volume and diversity of expert behaviors are capital intensive.

A promising strategy to reduce the data demands of learned autonomy is to leverage large-scale pretrained foundation models [e.g., V-JEPA 2 (*18*), GR00T (*16*), and π_0 (*13*)]. Trained on internet-scale, multimodal datasets, these models can potentially be fine-tuned for specific clinical applications with substantially reduced data requirements. The effectiveness of this approach can be amplified through open-sourcing high-quality clinical datasets [such as SurgVU from Intuitive Surgical (*19*)], which not only enable broader participation by the research community but also provide the domain-specific data needed to adapt foundation models to the surgical setting.

Beyond initial model training, deployed systems must support mechanisms for continuing improvement. Data feedback loops—such as systematic logging of system disengagements (e.g., because of human operator intervention), retries, or failure events—can drive the creation of targeted datasets that improve model robustness, particularly in rare or complex clinical scenarios. This paradigm mirrors strategies used in autonomous driving (e.g., Tesla's fleet learning) but poses unique challenges in the clinical domain, particularly with respect to safe system disengagement, ensuring patient privacy, and facilitating secure data sharing across institutions.

Handling edge cases: Adaptability, reasoning, and fallback strategies

Given that collecting data for every conceivable eventuality (i.e., "infinite data") is impossible because of the long tail of such events, autonomous systems will end up in states that were not observed in the training data. This is especially critical in sequential decision-making tasks such as robotics, where initial errors compound over time, potentially driving the system progressively further outside its training manifold and into unfamiliar states, where system behavior is undefined—a problem known as covariate shift (*20*). Although increasing data diversity aims to reduce unfamiliar long-tail events, a direct approach to enhancing safety involves the learned system actively monitoring its inputs and state to recognize whether it is operating inside or outside the distribution of its training data. Current research explores methods like uncertainty estimation using Bayesian neural networks or model ensembles, measuring latent space distance or reconstruction error from autoencoders, or using other anomaly detection techniques to quantify state familiarity (*21*, *22*). Upon detecting a potential out-of-distribution state, the system needs strategies to safely recover if possible or to actively request human intervention or guidance. This is not just a challenge in medical robotics but in machine intelligence in general.

Complementing these technical approaches for detection and response, a pragmatic strategy for mitigating risks, particularly as learning methods mature, involves carefully selecting initial automation targets. Focusing early efforts on less complex procedures or well-defined subtasks—perhaps using robots and tools requiring simpler sensing and control—can limit the risks or the probability of the system facing novel long-tail events. This mirrors evolutionary trends in medical device development, where the implementation of minimally invasive approaches has led to the development of tools and devices requiring fewer and simpler motions to deploy, e.g., staplers replacing needle and suture, endovascular stents replacing vascular grafts, and transcatheter heart valves replacing surgical valve implantation. Such an incremental approach, starting with lower-complexity tasks,

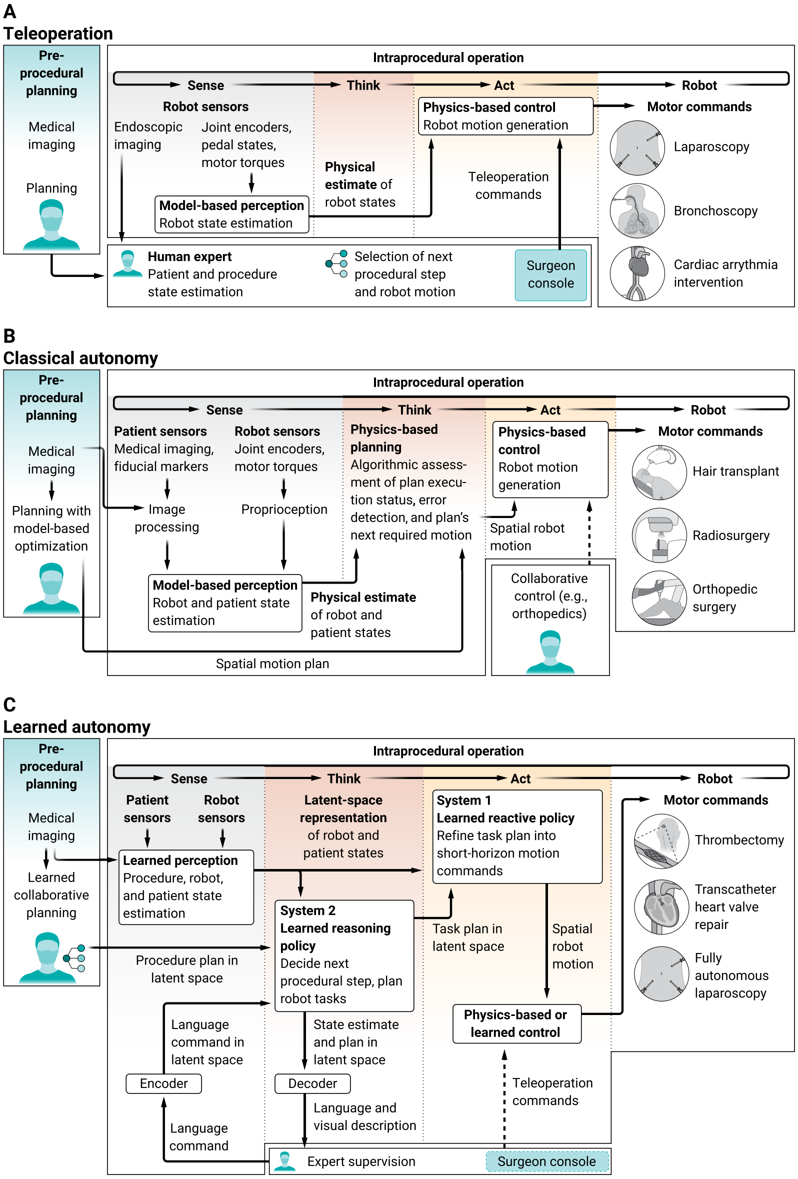

Fig. 1. Evolution of medical robot automation. Medical procedures comprise two distinct phases: preprocedural planning and intraprocedural operation. Autonomy emerges from the ability to repeatedly sense, think, and act (*9*) during intraprocedural operation. With increasing autonomy, the distribution of sense-think-act responsibilities shifts progressively from human to machine, expanding the repertoire of automated medical procedures. (**A**) Teleoperation. Most medical robots today are teleoperated, only capable of mimicking the motions of a human operator through low-level sensing and actuation. The human expert receives sensory data about the surgical environment ("sense"); on the basis of their perception and knowledge of the patient, they formulate a strategy for carrying out the medical procedure ("think") and issue commands to the robot by manipulating input devices ("act") that the robot executes through a physics-based controller. (**B**) Classical autonomy. Current implementations of autonomy target simple-to-model procedures and use non-AI algorithms. Robot and patient states are estimated using physics-based models and traditional image processing. "Think" adapts the preoperative plan to disturbances such as respiratory movement using model-based optimization, and the motion plan is converted into low-level motor commands using physics-based kinematic and dynamic robot models, with the clinician supervising or sharing control. Applications of classical autonomy are limited by the need for accurate analytic models of the robot, patient, and procedure. (**C**) Learned autonomy. Future learning-based automation targets procedures for which classical model-based methods fail. Learned perception encodes multimodal sensor inputs and medical images into latent representations; a learned reasoning policy generates long-horizon task-oriented plans in the latent space, which are then refined into short-horizon, high-frequency motion commands by a learned reactive policy. The motion commands are converted into low-level motor commands through either a classical or learned model of the robot. Human expert supervision is enabled by decoding the system's latent space into human-interpretable formats and allowing the expert to collaborate with the system, either through language instructions or direct teleoperation commands through a physical interface.

CREDIT: A. MASTIN/*SCIENCE ROBOTICS*

may facilitate safer adoption of autonomous systems and progressively expand their capabilities.

Collaborative autonomy, explainability, and trust

Medical robots will likely function as collaborative partners for the foreseeable future. Therefore, a core challenge is defining effective mechanisms for how autonomy can be safely shared with a human operator. This requires developing intuitive interfaces and blending control fluidly between the human operator and the learned system by predicting the operator's intent and arbitrating it with the operator's input (*23*). For example, LINGO-2 (*24*) continuously provides an explanation of its driving decisions in natural language and can receive linguistic instructions to alter its behavior. Conditions for clinicians to build trust in these systems must be analyzed, which may include the system communicating its perceived state of the world, intended actions, and the constraints it satisfied in a human-understandable representation. Understanding why the system proposes a particular action can help clinicians verify system behavior, predict potential issues, and dictate corrective actions.

Testing, verification, and regulatory approval

Validating the safety and efficacy of systems incorporating learned components before clinical adoption presents new challenges compared with traditional software. Defining adequate test coverage across high-dimensional input spaces (e.g., surgical video), determining the correct outputs and appropriate target metrics for test procedures, and making tests repeatable are notable hurdles. Diverse and well-curated datasets, standardized benchmarking suites, and metrics that correlate with end-to-end performance are necessary for objective evaluation and comparison of different approaches. Examples from autonomous driving are the Waymo Open Dataset (*25*), nuScenes (*26*), and the CARLA simulator (*27*). Applying verification methods such as input fuzzing to neural networks remains challenging because of the dimensionality of inputs (*28*).

Developing standardized testing and verification methods will provide regulatory bodies with the tools to effectively and efficiently evaluate learning-based systems (*29*). It will also lead to the creation of streamlined regulatory systems for incorporating learning-based updates over the life cycle of the system.

Regulatory frameworks typically require software to be "locked" at deployment, where updates to the neural network weights or architecture would typically require a new certification application. However, the US Food and Drug Administration has recently provided guidance (*30*) recommending that modifications to artificial intelligence (AI)–enabled device software functions be included in a predetermined change control plan, which should describe the planned AI modifications; the associated methodology to develop, validate, and implement those modifications; and an assessment of the impact of those modifications.

TRANSLATIONAL PATHWAY

Although learning provides the potential to broadly extend robotic autonomy in medical procedures, a critical question is which procedures should be targeted first. The most obvious use case is procedures for which the cost of a teleoperated robot can already be justified. For these systems, e.g., laparoscopic robots, autonomous functionality can be added over time in the same way driver assistance technologies have been added to cars.

The more challenging use case is to identify procedures for which a robot is only justified if it is autonomous. Such procedures would have to be high volume and provide high reimbursements. Furthermore, robotic automation would need to convey substantial benefits in terms of improved patient outcomes, reduced costs, or enhanced treatment availability. One such class of procedures is those that are needed on an emergent basis and for which specialized clinical staff are not available onsite but can provide guidance remotely. For example, although emergency mechanical thrombectomy after ischemic stroke is most successful when performed within an hour of presentation, many community hospitals are located more than an hour from stroke centers. A second class of procedures to be considered for automation is those with steep learning curves and that require a high case load for maintaining skills. For example, the operator performance of transcatheter valve repair procedures can continue to improve over hundreds of cases (*31*). For such procedures, autonomous robotic assistance could enhance geographic accessibility by enabling low-volume operators to perform at high skill levels.

As medical robot autonomy develops, it may be possible to use this technology in the same way it is used outside of medicine—to increase productivity by enabling a clinician to substantially increase the number of procedures they can perform per day. Although the clinical oversight of robots is critical at this point, advances in learning-based autonomy may enable a future tipping point leading to expanded access and leveling off in health care costs.

REFERENCES AND NOTES

1. Annual Report 2024 (Intuitive Surgical, 2024); https://isrg.intuitive.com/static-files/500ff989-ad91-4b32-a59e-f94a34d75997.
2. L. Boberg, J. Singh, A. Montgomery, P. Bentzer, Environmental impact of single-use, reusable, and mixed trocar systems used for laparoscopic cholecystectomies. *PLOS ONE* **7**, e0271601 (2022).
3. S. Herrera, "Amazon is on the cusp of using more robots than humans in its warehouses," *Wall Street Journal*, 30 June 2025; www.wsj.com/tech/amazon-warehouse-robots-automation-942b814f.
4. J. Quinlivan, "How Amazon deploys collaborative robots in its operations to benefit employees and customers" (Amazon, 26 June 2023); www.aboutamazon.com/news/operations/how-amazon-deploys-robots-in-its-operations-facilities.
5. R. D'Andrea, Guest editorial: A revolution in the warehouse: A retrospective on Kiva systems and the grand challenges ahead. *IEEE Trans. Autom. Sci. Eng.* **9**, 638–639 (2012).
6. W. Kilby, M. Naylor, J. R. Dooley, C. R. Maurer Jr, S. Sayeh, "A technical overview of the CyberKnife system" in *Handbook of Robotic and Image-Guided Surgery* (Elsevier, 2020), pp. 15–38.
7. P. T. Rose, B. Nusbaum, Robotic hair restoration. *Dermatol. Clin.* **32**, 97–107 (2014).
8. M. Roche, The MAKO robotic-arm knee arthroplasty system. *Arch. Orthop. Trauma Surg.* **141**, 2043–2047 (2021).
9. M. Siegel, "The sense-think-act paradigm revisited" in *Proceedings of the 1st International Workshop on Robotic Sensing* (IEEE, 2003), pp. 5.
10. S. D. Pendleton, H. Andersen, X. Du, X. Shen, M. Meghjani, Y. H. Eng, D. Rus, M. H. Ang, Perception, planning, control, and coordination for autonomous vehicles. *Machines* **5**, 6 (2017).
11. A. Tampuu, T. Matiisen, M. Semikin, D. Fishman, N. Muhammad, A survey of end-to-end driving: Architectures and training methods. *IEEE Trans. Neural Netw. Learn. Syst.* **33**, 1364–1384 (2020).
12. M. Bojarski, C. Chen, J. Daw, A. Degirmenci, J. Deri, B. Firner, B. Flepp, S. Gogri, J. Hong, L. Jackel, Z. Jia, B. Lee, B. Liu, F. Liu, U. Muller, S. Payne, N.K.N. Prasad, A. Provodin, J. Roach, T. Rvachov, N. Tadimeti, J. van Engelen, H. Wen, E. Yang, Z. Yang, The NVIDIA PilotNet experiments. arXiv:2010.08776 [cs.CV] (2020).
13. K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L.X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, U. Zhilinsky, π_0: A vision-language-action flow model for general robot control. arXiv:2410.24164 [cs.LG] (2024).
14. Y. LeCun, "A path towards autonomous machine intelligence," Open Review (2022); https://openreview.net/pdf?id=BZ5a1r-kVsf.
15. Figure, Helix: A vision-language-action model for generalist humanoid control, 20 February 2025; www.figure.ai/news/helix.
16. NVIDIA, J. Bjorck, F. Castañeda, N. Cherniadev, X. Da, R. Ding, L.J. Fan, Y. Fang, D. Fox, F. Hu, S. Huang, J. Jang, Z. Jiang, J. Kautz, K. Kundalia, L. Lao, Z. Li, Z. Lin, K. Lin, G. Liu, E. Llontop, L. Magne, A. Mandlekar, A. Narayan, S. Nasiriany, S. Reed, You L. Tan, G. Wang, Z. Wang, J. Wang, Q. Wang, J. Xiang, Y. Xie, Y. Xu, Z. Xu, S. Ye, Z. Yu, A. Zhang, H. Zhang, Y. Zhao, R. Zheng, Y. Zhu, GR00T N1: An open foundation model for generalist humanoid robots. arXiv:2503.14734 [cs.RO] (2025).
17. H. Choi, C. Crump, C. Duriez, A. Elmquist, G. Hager, D. Han, F. Hearl, J. Hodgins, A. Jain, F. Leve, C. Li, F. Meier, D. Negrut, L. Righetti, A. Rodriguez, J. Tan, J. Trinkle, On the use of simulation in robotics: Opportunities, challenges, and suggestions for moving forward. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e1907856118 (2021).
18. M. Assran, A. Bardes, D. Fan, Q. Garrido, R. Howes, M., K., M. Muckley, A. Rizvi, C. Roberts, K. Sinha, A. Zholus, S. Arnaud, A. Gejji, A. Martin, F. Robert Hogan, D. Dugas, P. Bojanowski, V. Khalidov, P. Labatut, F. Massa, M. Szafraniec, K. Krishnakumar, Y. Li, X. Ma, S. Chandar, F. Meier, Y. LeCun, M. Rabbat, N. Ballas, V-JEPA 2: Self-supervised video models enable understanding, prediction and planning. arXiv: 2506.09985 [cs.AI] (2025).
19. A. Zia, M. Berniker, R. Nespolo, C. Perreault, Z. Wang, B. Mueller, R. Schmidt, K. Bhattacharyya, X. Liu, A. Jarc, Surgical visual understanding (SurgVU) dataset. arXiv:2501.09209 [cs.CV] (2025).
20. J. Spencer, S. Choudhury, A. Venkatraman, B. Ziebart, J.A. Bagnell, Feedback in imitation learning: The three regimes of covariate shift. arXiv:2102.02872 [cs.RO] (2021).
21. J. Gawlikowski, C. R. N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher, M. Shahzad, W. Yang, R. Bamler, X. X. Zhu, A survey of uncertainty in deep neural networks. *Artif. Intell. Rev.* **56**, 1513–1589 (2023).
22. J. Yang, K. Zhou, Y. Li, Z. Liu, Generalized out-of-distribution detection: A survey. *Int. J. Comput. Vis.* **132**, 5635–5662 (2024).
23. A. D. Dragan, S. S. Srinivasa, A policy-blending formalism for shared control. *Int. J. Robot. Res.* **32**, 790–805 (2013).
24. LINGO-2: Driving with natural language (Wayve, 17 April 2024); https://wayve.ai/thinking/lingo-2-driving-with-language.
25. P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, D. Anguelov, "Scalability in perception for autonomous driving: Waymo open dataset" in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, 2020), pp. 2446–2454.
26. H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V.E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, O. Beijbom, "nuScenes: A multimodal dataset for autonomous driving" in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, 2020), pp. 11621–11631.
27. A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, V. Koltun, "CARLA: An open urban driving simulator" in *Proceedings of the 1st Annual Conference on Robot Learning* (MLResearchPress, 2017), pp. 1–16.

28. F. Tambon, G. Laberge, L. An, A. Nikanjam, P. S. N. Mindom, Y. Pequignot, F. Khomh, G. Antoniol, E. Merlo, F. Laviolette, How to certify machine learning based safety-critical systems? A systematic literature review. *Autom. Softw. Eng.* **29**, 38 (2022).
29. H. J. Marcus, P. T. Ramirez, D. Z. Khan, H. L. Horsfall, J. G. Hanrahan, S. C. Williams, D. J. Beard, R. Bhat, K. Catchpole, A. Cook, K. Hutchison, J. Martin, T. Melvin, D. Stoyanov, M. Rovers, N. Raison, P. Dasgupta, D. Noonan, D. Stocken, G. Sturt, A. Vanhoestenberghe, B. Vasey, P. McCulloch, The IDEAL framework for surgical robotics: Development, comparative evaluation and long-term monitoring. *Nat. Med.* **30**, 61–75 (2024).
30. Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions (FDA, 2024); www.fda.gov/media/166704/download.
31. A. K. Chhatriwalla, S. Vemulapalli, D. R. Holmes Jr., D. Dai, Z. Li, G. Ailawadi, D. Glower, S. Kar, M. J. Mack, J. Rymer, A. S. Kosinski, P. Sorajja, Institutional experience with transcatheter mitral valve repair and clinical outcomes: Insights from the TVT registry. *JACC Cardiovasc. Interv.* **12**, 1342–1352 (2019).

Acknowledgments

Funding: National Institutes of Health grants R01HL167925 and R01HL124020 (to P.E.D.). **Author contributions:** Both authors contributed equally to this work. **Competing interests:** P.E.D. is a co-inventor on US Patent 12274413B2, "Autonomous robotic catheter for minimally invasive interventions." A.D. holds stock in NVIDIA Corp., which is developing technology for medical robot autonomy.

10.1126/scirobotics.adz8279

Science Robotics

The grand challenges of learning medical robot autonomy

Pierre E. Dupont and Alperen Degirmenci

Sci. Robot. **10** (104), eadz8279. DOI: 10.1126/scirobotics.adz8279

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.