AUTONOMOUS ORTHOPEDICS

Beyond Geometric Precision

The Next Control Layer in Orthopedic Robotics

A scientific and strategic analysis of RACER-Knee, arthroplasty adoption, and the case for clinically consequential mechanical intelligence

Kambiz Behzadi, MD

Orthopedic Surgeon | Autonomous Orthopedics Program

August 2026

Controlled non-confidential discussion draft

Executive thesis

The conclusion. Orthopedic robotics should not be judged by whether it can reproduce a plan more precisely. It should be judged by whether the additional system changes a clinically consequential decision, improves a patient outcome, or materially reduces risk. RACER-Knee shows that geometric precision can improve while patient benefit remains unchanged. The next frontier is therefore not more geometry layered onto the same procedure; it is real-time mechanical intelligence at the bone-tool and bone-implant interface.

The 2026 RACER-Knee trial is a pivotal result for surgical robotics. In the largest participant-masked and assessor-masked randomized comparison of a contemporary robotic-arm total knee replacement system with conventional instruments, robotic assistance produced more accurate execution of planned alignment. Yet it did not produce a clinically meaningful improvement in the Forgotten Joint Score or the trial's major secondary patient outcomes at 12 months. It increased operative time and cost, and it was unlikely to be cost-effective.[1]

That is not evidence that precision is useless, that all robots are equivalent, or that longer-term benefit is impossible. It is evidence that a technically successful intermediate variable can be a clinically weak product proposition. A robot can do exactly what it was designed to do and still fail to change what matters to the patient.

This distinction matters even more in total hip arthroplasty. Robotic hip adoption remains lower than robotic knee adoption in comparable U.S. claims data.[3] Recent hip evidence again shows the same pattern: greater accuracy in component position and reconstruction, but no advantage in the proportion of patients achieving a minimal clinically important improvement in a 132-patient multicenter randomized trial.[4] Meanwhile, the high-force tasks that determine press-fit preparation, seating, fixation, and fracture margin remain largely manual and qualitative.

Figure 1. Two different control problems

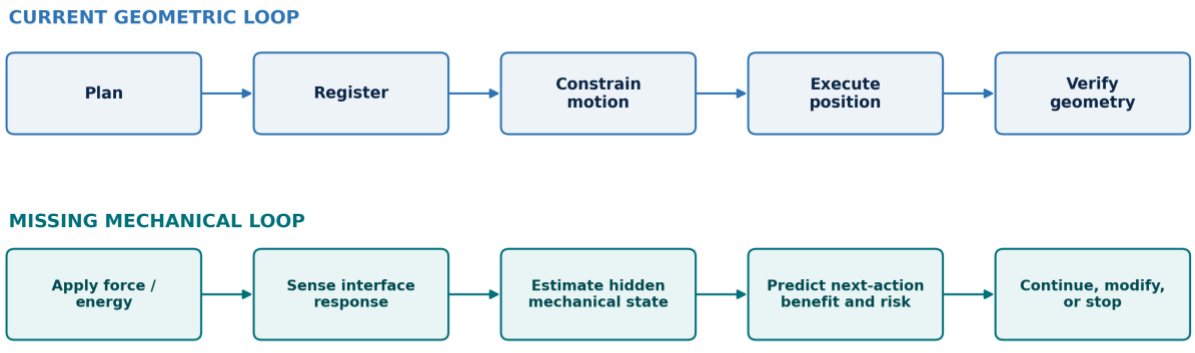

Geometry establishes where the tool or implant should be. Mechanical intelligence determines what the action is doing to the patient.

The current geometric loop and the proposed mechanical loop are complementary. AO does not discard planning, navigation, or geometric control; it adds the missing estimation and decision layer for force-mediated interaction.

1. The RACER result: a precision-to-outcome translation failure

RACER-Knee randomized 339 patients across ten hospitals and 33 surgeons to Mako robotic-arm-assisted total knee replacement (rTKR) or conventional total knee replacement (cTKR). Participants and assessors were masked; sham pin incisions and identical preoperative CT pathways were used to preserve masking. The trial was pragmatic: surgeons could use their usual alignment strategy, and participating surgeons had performed at least ten Mako cases. The primary endpoint was the 12-month Forgotten Joint Score (FJS), with a prespecified target between-group difference of 12 points.[1]

Table 1. RACER-Knee: improved geometric precision without demonstrated first-year patient benefit

Endpoint	Robotic TKR	Conventional TKR	Adjusted result
Forgotten Joint Score, 12 months	49.2	50.2	-1.5 (95% CI -7.5 to 4.5); p=0.62
Total NHS/PSS cost, 0-12 months	£8,979.69	£8,140.66	+£839.03 (95% CI £370.95 to £1,307.11)
Skin incision to final dressing	92 min	83 min	+10.5 min (95% CI 6.8 to 14.3)
Absolute planned HKA-angle error	2.0°	2.8°	-0.8° (95% CI -1.4 to -0.3)
Participants with a serious adverse event	16 (10%)	16 (9%)	No detected safety advantage

Source: Parsons et al., RACER-Knee, The Lancet (2026). HKA = hip-knee-ankle; NHS = National Health Service; PSS = Personal Social Services. Negative FJS values favor conventional TKR. The trial evaluated Mako-assisted TKR as delivered in routine practice.

The primary result was not merely statistically nonsignificant. The adjusted difference was -1.5 points, favoring conventional TKR, and the 95% confidence interval excluded the prespecified 12-point target. Secondary outcomes—including early pain, opioid use, discharge timing, Oxford Knee Score, quality of life, participation, and pain intensity—showed no clear clinical advantage. Serious adverse events occurred in 16 participants in each group.[1]

At the same time, the robot improved the accuracy of achieving the planned hip-knee-ankle angle: absolute error averaged 2.0 degrees with the robot and 2.8 degrees conventionally. This juxtaposition is the central finding for technology strategy. The machine improved the metric it directly controlled, but that metric did not drive a measurable first-year patient benefit in the tested population and workflow.

What RACER does—and does not—prove

It does prove: For Mako-assisted TKR as used in this pragmatic trial, added geometric execution precision did not yield clinically meaningful benefit through 12 months, increased the initial admission and total 12-month cost, and added operative time.

It does not prove: That every robotic platform, alignment philosophy, patient subgroup, or future implementation lacks value. RACER did not standardize or test one personalized alignment strategy, evaluated one system, and has not yet reported long-term revision outcomes.

The proper inference: Robotics requires a stronger causal chain from controlled variable to patient outcome. Geometric precision is an engineering capability; its clinical value depends on whether the chosen geometric target is itself outcome-determining.

2. Adoption has outrun proof of benefit

Robotic TKA is no longer a fringe technology. RACER cites 2024 registry estimates of robotic assistance in 6% of TKR procedures in the United Kingdom, 16% in the United States, and 41% in Australia.[1] Those rates represent substantial community uptake—especially in the United States and Australia—even though high-quality randomized evidence has not established a first-year patient-reported or cost-effectiveness advantage. RACER's updated meta-analysis of the two available randomized trials, RACER and ROAM, likewise found no difference in patient-reported outcomes at three or 12 months.[1]

Table 2. Adoption is real, but exact rates depend on registry and coding method

Source and year	Robotic knee	Robotic hip	Interpretation
National registry rates cited by RACER, 2024	UK 6%; U.S. 16%; Australia 41%	Not reported in RACER	Robotic TKA has achieved substantial uptake in some markets.
U.S. PearlDiver all-payer claims, 2022	3.45% of TKA	2.36% of THA	Within one dataset, robotic knee use was about 46% higher than robotic hip use.
U.S. National Inpatient Sample, 2022	Not analyzed in the THA study	6.7% of THA	Robotic THA remained a minority practice; this estimate should not be compared directly with claims rates.

Sources: Parsons et al. [1], Gordon et al. [3], and Maman et al. [5]. Rates differ because databases, coding rules, settings, and denominators differ. The within-dataset knee-versus-hip comparison is more defensible than cross-study percentage comparisons.

The adoption gap should not be overstated numerically because coding methods produce materially different absolute estimates. But the direction is consistent: robotic knee replacement has gained more traction than robotic hip replacement. In a single all-payer U.S. claims dataset, robotic assistance represented 3.45% of TKA and 2.36% of THA in 2022.[3] A separate inpatient study estimated that robotic THA reached 6.7% in 2022, still a minority of procedures.[5]

The key point is not that knee surgeons made a poor choice or that hip surgeons resisted innovation. It is that adoption can be driven by marketing, institutional competition, capital strategy, surgeon preference, perceived workflow value, and the intuitive appeal of precision before patient benefit is demonstrated. The stronger the adoption, the more important it becomes to ask whether complexity is attached to the correct clinical uncertainty.

3. Why geometry can improve while outcomes remain unchanged

3.1 Precision is not the same as target validity

A robot can reduce variance around a target. It cannot by itself establish that the target is biologically or mechanically optimal for a particular patient. In knee arthroplasty, the ideal alignment strategy—and whether it should differ among patients—remains contested. RACER explicitly cautions that it was not a trial of any one alignment philosophy.[1] Improving the accuracy of an uncertain target may narrow technical dispersion without changing the distribution of patient outcomes.

3.2 Mature manual procedures create a high incremental-benefit threshold

Conventional arthroplasty already produces large improvements for most patients. Incremental technology must therefore overcome a high baseline. When the conventional method is reliable in experienced hands, a new system must either affect a failure mode that manual technique does not reliably control, deliver benefit to a definable high-risk subgroup, reduce resource use, or create a new capability. Marginal geometric improvement alone may be absorbed by biological variation, rehabilitation, pain mechanisms, implant design, soft-tissue behavior, and ceiling effects in clinical outcomes.

3.3 The decisive mechanical steps remain outside the robotic loop

Many current systems plan, register, guide, or constrain spatial execution while the surgeon still performs the highest-force and most mechanically consequential steps: reaming, broaching, impaction, press-fit seating, and endpoint determination. A revealing first-generation premise was that if a robot created a dimensionally accurate cavity, the surgeon could complete insertion using a 'reasonable' amount of force. The phrase exposes the unresolved variable. Reasonable for which bone, which interface, which impact, and which moment in the procedure?

Geometry describes pose, dimensions, and trajectory. It does not directly reveal effective stiffness, local contact, interference fit, friction, seating progression, fixation gain, micromotion risk, microdamage, or proximity to fracture. The same nominal preparation and implant can create different interface mechanics in different patients. When these variables remain hidden, the robot automates the measurable perimeter while the surgeon still manages the mechanistic core through experience, sound, resistance, advancement, recoil, and feel.

4. Total hip replacement exposes the missing control variable

The hip is where the geometry-mechanics distinction becomes most obvious. A cup can be within its planned inclination and anteversion while primary fixation remains inadequate. A broach can follow the correct trajectory while hoop stress approaches fracture. An implant can advance with repeated impacts

yet gain little additional stability. Conversely, stopping too early can leave inadequate press fit and micromotion. Component position is necessary, but it is not a sufficient state description of the interface.

A 2024 analysis of seven national arthroplasty registries covering 197,624 revised THAs reported aseptic loosening as 35.1% of revision indications and periprosthetic fracture as 11.4%—46.5% combined.[6] These categories are multifactorial and cannot be attributed wholesale to intraoperative force management. Nevertheless, they bracket a clinically important mechanical spectrum: inadequate effective fixation at one end and excessive local loading or structural failure at the other. The mechanically controllable fraction is the target; its magnitude must be established, not assumed.

Recent hip evidence reinforces the translation problem without establishing that robotics is ineffective. In a 132-patient multicenter randomized trial, robotic THA improved cup-position accuracy, limb-length control, offset restoration, and component-size prediction. Harris Hip Scores were higher at three and 12 months, but the proportion achieving a minimal clinically important improvement was not different.[4] A national observational analysis found shorter inpatient stays and fewer coded in-hospital complications, but higher charges; its nonrandomized design cannot establish long-term superiority.[5] The evidence is mixed, and the clearest benefit remains technical accuracy rather than durable patient value.

5. The missing layer: mechanical intelligence

The design shift. The product should not be 'a more complex robot.' The product should be a reduction in clinically consequential uncertainty: Is the interface underprepared, stably seated, still gaining fixation, or approaching irreversible damage—and what should happen next?

Autonomous Orthopedics (AO) proposes a complementary control layer for force-mediated surgery. The system senses mechanical interaction, estimates a patient-specific hidden state, predicts the consequence of candidate actions, constrains those actions within a safety envelope, and supports a decision to continue, modify, redirect, vibrate, pause, or stop. The architecture begins with measurement and decision support; autonomy is earned only after state estimation and causal benefit are validated.

5.1 From theoryless skill to theoryfull intervention

Current high-force orthopedic technique is often theoryless in a specific sense: it works through accumulated experience, but the governing mechanics are not fully identified, measured, or controlled in the individual patient. Expert performance is real; the limitation is that its state variables and termination rules are difficult to verbalize, transfer, calibrate, or audit. A theoryfull intervention identifies the relevant variables, tests causal mechanisms against ground truth, and reduces the procedure to measurable, predictive, controllable steps.

5.2 The AO state-estimation problem

At each interaction event, the applied action—force, torque, energy, impulse, displacement, velocity, vibration, direction, and time—produces multimodal measurements. Force-displacement trajectories, advancement, recoil, acoustic response, vibration, kinematics, and motor signals are not isolated numbers; together they are observations of a changing contact system. Candidate latent states include effective stiffness or compliance, contact condition, interface conformity, frictional state, cortical engagement, seating progression, stability gain, and damage proximity.

The scientific question is observability: do the available signals contain enough information to estimate the clinically relevant state with calibrated uncertainty? Only after that question is answered should the system predict the marginal benefit and risk of the next action. A simpler threshold model should defeat no baseline before a complex estimator is introduced; a complex estimator should advance only if it produces a material, reproducible improvement over simple methods.

5.3 Why this could create a larger effect than another geometry layer

Mechanical intelligence targets decisions that are currently unmeasured, patient-specific, and causally close to fixation and structural damage. It therefore has the potential—not the guarantee—to produce a larger effect size than incremental gains in already-precise positioning. It could also distribute a form of expert sensorimotor calibration across surgeons and institutions: manual surgery becomes measurable; measurable surgery becomes predictable; predictable surgery becomes controllable; controlled surgery creates a responsible pathway toward bounded autonomy.

AO therefore operates at two learning scales: within a procedure, new measurements update the patient-specific mechanical-state estimate; across procedures, de-identified signals, actions, independent ground truth, and longitudinal outcomes are retained and curated to train and challenge candidate models across patients, implants, surgeons, and sites. As with autonomous vehicles, the advantage compounds as diverse real-world cases expose rare conditions and failure boundaries—but in medicine, updated models must be prospectively validated, version-locked, authorized, and monitored before deployment rather than learning unconstrained during a case.

Table 3. What is established, what is inferred, and what must still be proven

Proposition	Status	Scientific meaning
RACER's robotic TKA improved alignment precision.	**Established**	The system improved the process variable it was engineered to control.
RACER's robotic TKA did not improve 12-month patient-reported outcomes and was more costly.	**Established**	Technical precision did not translate into meaningful first-year clinical value in that implementation.
Robotic TKA is more widely adopted than robotic THA.	**Supported**	Multiple datasets show greater knee uptake, although exact rates vary by method.
Lower hip uptake partly reflects failure to address mechanical decision-making.	**Inference**	Plausible and consistent with workflow, but adoption data alone cannot establish causation.
Mechanical-state intelligence will improve fixation and reduce damage.	**Testable hypothesis**	AO must prove observability, predictive validity, causal benefit, safety, and value before added complexity is justified.

6. A complexity threshold for medical robotics

A central engineering principle follows from the evidence: medicine should accept algorithmic and regulatory complexity only when the clinical advantage is substantial. That principle should be treated as a design constraint, not a rhetorical caution. Every layer of sensing, modeling, actuation, connectivity, and autonomy creates acquisition cost, operating-room burden, integration work, maintenance, cybersecurity exposure, human-factors risk, verification obligations, and regulatory debt.

The appropriate objective is clinical value density: validated benefit per unit of total system burden. A mechanically intelligent system should advance only if it clears five gates.

Gate 1 — Consequence

The targeted decision must affect a clinically important failure mode, functional outcome, safety event, or resource burden—not merely a more elegant process metric.

Gate 2 — Observability

The relevant mechanical state must be estimable from practical signals with quantified uncertainty and independent ground truth.

Gate 3 — Causality

Changing the recommended action or endpoint must change fixation, damage, or another mechanistically linked outcome—not simply correlate with it.

Gate 4 — Incremental effect

The system must outperform manual technique, fixed-energy protocols, simple force limits, and other low-complexity baselines by a prespecified clinically meaningful margin.

Gate 5 — Net value

The benefit must justify added time, cost, workflow disruption, failure modes, training, and regulatory complexity. If it does not, the system should simplify or stop.

7. A falsifiable validation program

AO should not begin by asking a surgeon or regulator to trust an autonomous system. It should begin by testing whether the hidden mechanical states are measurable and whether acting on those estimates changes independent ground truth. A disciplined program would proceed through the following stages.

Stage 0 — Measurement without recommendation

Instrument conventional preparation and implantation. Synchronize force, displacement, vibration, acoustic, kinematic, and motor data. Establish signal quality, calibration, repeatability, timing, and failure detection without changing surgical decisions.

Stage 1 — Observability against independent ground truth

Use fixtures, synthetic bone, and cadaveric models to test estimation of seating progression, fixation strength, micromotion, microdamage, fracture proximity, and endpoint. Ground truth should include direct displacement, pull-out or torsional stability, imaging, strain, damage inspection, and controlled failure testing as appropriate.

Stage 2 — Comparative prediction

Compare simple thresholds, classical state estimators, physics-informed models, and data-driven approaches prospectively. Lock models before testing. Require calibration, transfer across specimens and implant designs, and explicit out-of-domain detection.

Stage 3 — Surgeon decision support

Test whether blinded mechanical-state estimates improve endpoint recognition, fixation consistency, or damage avoidance relative to manual judgment. Measure false alarms, missed hazards, cognitive load, operating time, overrides, and trust calibration—not only algorithm accuracy.

Stage 4 — Safety-constrained assistance

Permit bounded adaptation of force, energy, pulse shape, vibration, direction, or termination only after the earlier stages establish observability and benefit. Use hard limits, sensor-agreement gates, uncertainty thresholds, surgeon authorization, fallback modes, audit logs, locked versions, and rollback.

Prespecified simplify-or-stop criteria

The program should stop or simplify if mechanical states are not observable across clinically relevant heterogeneity; if estimates fail to generalize across implants, specimens, operators, or sites; if simple thresholds perform as well as advanced models; if the intervention improves surrogates without changing fixation or damage; or if workflow and regulatory burden exceed the measurable benefit. These criteria are not concessions. They are what distinguish scientific development from complexity for its own sake.

8. Strategic and commercial implications

8.1 For MedTech

The strategic risk is continued investment in capital-intensive geometry platforms whose differentiation converges while outcome evidence remains modest. Mechanical intelligence offers a potentially different competitive axis: implant-specific knowledge of preparation, fixation, damage margin, and endpoint. It can begin as sensorized instruments or procedure modules that strengthen existing implant and robotic franchises rather than require replacement of the entire operating-room stack.

The commercial advantage would come from integration, not an isolated algorithm: protected instrumentation and actuation, synchronized mechanical data, validated state estimators, implant-specific control maps, regulatory evidence, surgeon workflow, and longitudinal linkage between intraoperative signals and outcomes. That combination is harder to reproduce than software alone and can create recurring value through instruments, disposables, enabling technology, and data-supported performance improvement.

8.2 For surgeons and health systems

Surgeons do not need technology that merely makes an already-visible action more elaborate. They need help with uncertainty they cannot directly observe: whether bone preparation is sufficient, whether fixation is increasing, whether another impact will help, and when additional force becomes harmful. Health systems need evidence that the additional workflow reduces complications, revision burden, variability, or resource use enough to justify acquisition and operating cost.

8.3 For Big Tech: a different physical AI specimen

Big Tech's physical AI agenda is searching for systems that do more than generate predictions on a screen: systems that perceive, reason, and act in the physical world under uncertainty. The dominant proving grounds are humanoid labor substitution, weaponized drones, autonomous vehicles, and factory robots. Each is important, but each carries a limitation as a general specimen of physical intelligence: social resistance, destructive purpose, open-world risk, or operation inside an engineered and repeatable environment. AO offers a different proving ground—a bounded, surgeon-supervised, safety-critical agent

acting on heterogeneous biological material, where success means reducing uncertainty and preventable harm rather than replacing labor, selecting a target, or repeating a factory task.

Scientifically, AO compresses many of physical AI's hardest problems into a constrained domain: multimodal perception; contact-rich dynamics; latent-state estimation; prediction under distribution shift; calibrated uncertainty; real-time edge inference; safety-constrained control; human authorization; and post-action verification. Bone-tool and bone-implant interaction is partially observable, patient-specific, nonlinear, and history-dependent. Actions change the state being estimated, and an error can be irreversible. Yet the action vocabulary, anatomy, instruments, workflow, and candidate ground truths can be bounded far more tightly than in general-purpose humanoid or open-world autonomy. That combination—high physical complexity with a narrow, auditable action space—makes AO potentially valuable as both a medical program and a rigorous physical AI benchmark.

The Big Tech proposition. AO is a compact real-world test of whether physical AI can convert multimodal signals into a trustworthy estimate of hidden state, choose within hard safety constraints, defer when uncertainty is excessive, and verify the physical consequence of action.

The strategic opportunity is broader than supplying a model. Big Tech can contribute foundation-model methods for multimodal time series, physics-informed learning, simulation and digital twins, edge computing, accelerators, data infrastructure, runtime assurance, and model monitoring. In return, AO could provide something the field lacks: longitudinal, action-linked data from consequential contact; explicit failure boundaries; independent mechanical ground truth; and a regulated pathway for demonstrating that physical AI can be useful, measurable, and governable. The clinical channel may remain with MedTech and health systems, while the enabling intelligence, compute, safety architecture, and cross-domain learning infrastructure create a distinct role for Big Tech.

This opportunity should not be overstated. Orthopedic data are expensive, heterogeneous, and difficult to label; clinical deployment is slow; and a general foundation model may add no value over a simpler estimator. The Big Tech thesis is therefore conditional: AO matters only if a bounded program first proves observability, prospective generalization, causal benefit, and safe abstention. If those gates are cleared, AO becomes more than another surgical application. It becomes evidence that physical AI can operate responsibly in the real world where matter is variable, consequences are high, and uncertainty cannot be hidden behind a simulation score.

8.4 For investors

The investment thesis is not 'robotics is growing.' RACER shows why category growth can coexist with weak clinical differentiation. The stronger thesis is that orthopedic robotics has digitized geometry but not the force domain. If AO can establish observability and causal control of the mechanical interface, it can become a foundational layer across implants, instruments, robotics, and high-force procedures. If it cannot, the complexity should not be funded beyond the evidence.

9. Principal risks and failure modes

Attribution risk. Aseptic loosening and periprosthetic fracture are multifactorial; registry revision categories do not quantify the preventable intraoperative mechanical fraction.

Observability risk. Signals may be noisy, nonunique, implant-specific, approach-specific, or insufficient to distinguish stable seating from occult damage.

Ground-truth risk. The states that matter clinically may be difficult to measure during development, and convenient surrogates may not predict durable fixation or fracture.

Generalization risk. Performance may degrade across bone quality, anatomy, implants, instruments, operators, and institutions. Out-of-domain detection is therefore central, not optional.

Human-factors risk. An inaccurate or poorly explained recommendation can create automation bias, alarm fatigue, or delayed action. Surgeon authority, clear confidence, and rapid fallback must remain part of the control architecture.

Regulatory and commercial risk. The more adaptive and autonomous the system becomes, the greater the validation, monitoring, liability, integration, and reimbursement burden. Progressive autonomy must be evidence-gated.

10. Conclusion: complexity must purchase consequence

RACER-Knee should not be read as the end of orthopedic robotics. It should be read as a design correction. The trial demonstrates that a contemporary robot can improve geometric precision, gain meaningful adoption, add cost and operative time, and still fail to deliver a clinically meaningful first-year patient advantage. That is the signature of a translation gap, not necessarily a failed machine.

The response should not be to add more planning layers, more screens, more degrees of freedom, or more artificial intelligence without changing the causal target. The response should be to identify the uncertainty closest to preventable harm and durable function. In high-force arthroplasty, that uncertainty is often mechanical: what the instrument is doing to the patient-specific interface, what state now exists, and whether the next action will improve fixation or cause damage.

> **The governing principle.** Robotics is not the clinical product. Reduced uncertainty at a consequential decision is the product. Complexity is justified only when it purchases consequence.

Orthopedic robots increasingly know where the instrument and implant are. The next generation must determine what force is doing to the patient—and whether more force should be applied at all.

References

1. Parsons H, Metcalfe A, Griffin J, et al. Robotic-arm-assisted versus conventional total knee replacement (RACER-Knee): a pragmatic, multicentre, participant-masked and assessor-masked, superiority, randomised controlled trial. The Lancet. Published online August 20, 2026. https://doi.org/10.1016/S0140-6736(26)00986-4

2. American Joint Replacement Registry (AJRR). 2025 Annual Report. American Academy of Orthopaedic Surgeons; 2025. https://www.aaos.org/registries/publications/ajrr-annual-report/

3. Gordon AM, Nian P, Baidya J, et al. Trends of robotic total joint arthroplasty utilization in the United States from 2010 to 2022: a nationwide assessment. Journal of Robotic Surgery. 2025;19:155. https://doi.org/10.1007/s11701-025-02313-5

4. Zhang Y, Li S, Liu X, et al. Robot-Assisted Total Hip Arthroplasty Versus Conventional Surgery in Terms of Surgical Accuracy, Function, Trauma, and Complications: A Prospective, Multicenter, Parallel-Group, Open-Label Randomized Controlled Trial. Journal of Arthroplasty. 2026;41(3):841-848.e1. https://doi.org/10.1016/j.arth.2025.07.029

5. Maman D, Steinfeld Y, Berkovich Y. Robotic-assisted total hip arthroplasty in the United States: a nationwide propensity-matched analysis of adoption, outcomes, and complications. Langenbeck's Archives of Surgery. 2026;411:65. https://doi.org/10.1007/s00423-025-03963-7

6. Hauer G, Rasic L, Klim S, et al. Septic complications are on the rise and aseptic loosening has decreased in total joint arthroplasty: an updated complication-based analysis using worldwide arthroplasty registers. Archives of Orthopaedic and Trauma Surgery. 2024;144:5199-5204. https://doi.org/10.1007/s00402-024-05379-2

Source material reviewed

The analysis also incorporates the supplied RACER-Knee article, Autonomous Orthopedics technical materials, and historical background concerning orthopedic robotics. The AO architecture and proposed validation program remain hypotheses requiring independent scientific testing.