1. ## Many Roads, Not One Super-highway

   The paper's core insight is that relying solely on ever-bigger chatbots to "wake up" into general intelligence is risky. AGI may arise from several different technical recipes working together. That matters for PSA because it widens the toolbox: instead of waiting for the next giant model, educators and clinicians can combine smaller, cheaper, purpose-built systems that already exist.

2. ## Three Practical Yardsticks for AGI

   RAND uses three everyday measures—human-level skill across many tasks, the ability to replace costly labor, and the emergence of new skills on its own. Keeping these yardsticks in mind helps a lay reader judge hype: if a model only chats well but can't safely run a clinic billing workflow, it isn't AGI yet.

3. ## Emergent "Magic" May Be a Mirage

   When researchers switched to smoother scoring methods, the dramatic "jumps" in ability reported for large language models mostly disappeared. For schools and hospitals that need predictability, this means pilot projects should emphasize measured outcomes, not headline benchmarks.

4. ## The Over-Confidence Trap

   As models grow, they not only get smarter—they also grow surer of wrong answers. In high-stakes settings (dosage calculators, student transcripts) that false confidence can be deadly or costly. Early warning systems that flag uncertainty will be essential companions to any AGI-flavored application.

5. ## The "Data Wall" Is Closer Than It Looks

   Human-written text that models train on could run out within the decade. After that, models would be learning from their own recycled echoes, degrading quality. Community-sourced education datasets or privacy-preserving health records become strategic assets, suggesting local initiatives to curate fresh, high-quality data.

6. ## The Energy Wall

   Training a single frontier model already rivals the annual electricity use of a small town. Clinics, schools, or nonprofits that adopt compact or edge-based AI will avoid skyrocketing utility bills, reduce carbon footprint, and keep services running during outages.

7. ## Physics-Informed & Causal Hybrids

   By baking in the laws of motion or cause-and-effect, these models reason about the real world instead of word statistics. A rehabilitation robot that "knows" knee-joint physics can coach exercises safely; a causal tutor can show how diet choices affect glucose.

8. ## Cognitive AI—Learning Like People

   Cognitive architectures copy the brain's trick of forming concepts, storing memories, and learning continuously. Such systems could serve as lifelong study partners that adapt as a student's knowledge grows, or as virtual health coaches that remember a patient's history across many appointments.

9. ## Information-Lattice Learning (ILL)

   ILL can discover clear, textbook-style rules from a handful of examples—no giant dataset needed. A small rural clinic could train a local ILL model on a few dozen anonymized cases to spot region-specific disease patterns, while teachers might generate new practice problems on the fly.

10. **Reinforcement Learning—Trial, Error, Mastery**

    RL agents improve by experimenting in virtual worlds before acting in reality. In education this powers personalized game-based curricula that adapt to each learner; in healthcare it enables robots that practice hundreds of needle insertions in simulation before assisting nurses.

11. **Neurosymbolic Reasoning**

    Blending pattern-hungry neural nets with rule-based logic yields systems that both "sense" and "explain." Imagine a reading-comprehension tutor that not only answers "who-did-what" but shows the logical chain so the student sees why. In diagnostics, the same transparency can satisfy regulators and reassure patients.

12. **Embodied AI—Learning by Doing**

    Robots and sensor-rich avatars gather tactile, visual, and auditory experience, grounding concepts in the physical world. Classroom science can move from 2-D screens to hands-on labs with affordable robot arms; physical-therapy bots can adjust pressure based on real-time feedback, personalizing care.

13. **Neuromorphic and Bio-Computing**

    Brain-inspired chips and cultured neurons promise huge energy savings and new kinds of adaptability. Battery-powered health monitors or off-grid learning devices could run complex models locally for days, widening access in underserved areas without costly data centers.

14. **Hybrid Stacks Beat Silver Bullets**

    The strongest candidates for AGI mix several of the above approaches—e.g., an LLM front-end that calls a physics engine, guided by reinforcement learning and running on neuromorphic hardware. For nonprofits this modularity means starting small and swapping in better pieces over time rather than waiting for one "perfect" model.

15. **Policy Hedge: Plan for Uncertainty**

    Because no one can guarantee which path will win, RAND urges governments (and by extension organizations) to keep options open—fund diverse research, avoid lock-in to any single vendor, and build flexible governance. PSA can mirror that strategy: pilot multiple AI tools, maintain open standards, and cultivate local expertise to pivot as the field evolves.