Summary: What can we learn from a billion agents?

We stand at a crossroads of two possible futures: one where advanced technology enables unprecedented human coordination and flourishing, as imagined in Star Trek, and another where it leads to fragmentation and inequality, as depicted in Elysium. Within a decade, we're projected to have over 60 billion AI agents augmenting human society - a transformation that will fundamentally reshape how we coordinate and collaborate at scale.

To realize this future, we must move beyond current AI paradigms. While Large Language Models have dramatically enhanced individual capabilities, they remain constrained by human cognitive limits - the same barriers that have always restricted our ability to coordinate at scale. We need a fundamentally different approach: protocol-centric intelligence that progresses through five distinct levels. Starting from enhancing individual awareness, through coordinating immediate circles and managing bounded networks, to enabling global orchestration and finally shaping collective physical reality - each level expands our ability to coordinate while preserving individual choice.

Imagine this progression in 2030: your AI assistant suggests adding calcium-rich foods based on your dietary patterns (individual awareness), coordinates your departure time with nearby shoppers (local interaction), connects with hundreds of stores to optimize delivery schedules (social coordination), processes millions of shopping patterns to prevent congestion (global orchestration), and ultimately shapes physical spaces and systems to enhance everyone's experience (physical integration). This isn't just about convenience - it's about enabling coordination at previously impossible scales.

These advances bring significant risks: unprecedented surveillance, new forms of inequality through "protocol privilege," and the subtle erosion of individual choice. However, through privacy-preserving computation, decentralized control, and universal access by design, we can ensure these systems enhance rather than restrict human agency. The difference between Star Trek's abundance and Elysium's scarcity lies not in technological capability, but in how we enable human coordination at scale. This future isn't predetermined—it's a protocol we get to write.

What can we learn from a billion agents?

1 Thesis: Scale, Complexity and Collective Behavior

Imagine a typical morning grocery run in 2030. As you check your phone, your personal AI assistant suggests adding calcium-rich foods to your list, having noted your recent dietary patterns. Your phone gently nudges leaving at 8:45 AM instead of your usual 8:30 - a suggestion being subtly coordinated with hundreds of other shoppers in your neighborhood. As you navigate the store, the layout on your augmented reality display subtly adjusts, not just to optimize your path but to orchestrate a delicate dance of all shoppers, preventing congestion before it occurs. During a surge in local flu cases, your path from door-to-door is automatically adjusted to maintain safe distances from other shoppers, while the store's ventilation systems adapt in real-time to minimize transmission risks.

This isn't science fiction. As billions of AI agents begin to augment human society - managing our calendars, guiding food choices, coordinating healthcare - we're witnessing an unprecedented transformation. This future isn't about replacing human interaction with automation, but enabling unprecedented coordination when we choose to come together. When combined with current technology adoption trends, we are heading to a future where each person could have 10 or more specialized AI agents, adding the equivalent of 60 billion "working agents" to our global systems within a decade. This explosion in artificial agents presents us with a fundamental choice: will these technological capabilities enhance or diminish human autonomy?

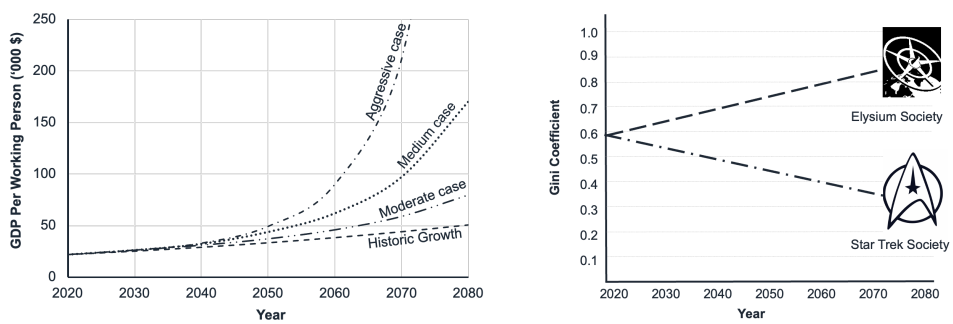

Figure 1: (a) Projected impact of AI agents on global workforce and economic growth, showing acceleration paths to post-scarcity across different adoption scenarios (b) Potential societal outcomes measured by Gini coefficient, contrasting possible futures between an equitable Star Trek-like society and a stratified Elysium-like world. In the age of AI, will we build an Elysium or Star Trek society?

The stakes of this choice manifest vividly in competing visions of our technological future. The Star Trek [19] universe depicts an optimistic 23rd-century where advanced coordination protocols have eliminated material wants, allowing humanity to focus on exploration and advancement. In stark contrast, the film Elysium [1] presents a dystopian 2154 where Earth is overpopulated and impoverished, with wealthy elite enjoying advanced medical technology in an orbital habitat. The tragedy isn't technological limitation - their medical pods could heal any illness - but rather the failure to develop protocols for equitable access and distribution. The difference - Star Trek's abundance

Preprint. Under review.

versus Elysium's scarcity - lies not in technological capability, but in their ability to coordinate human activity and resources at unprecedented scales.

Yet, throughout history, humans have been limited to meaningfully maintaining a few hundred stable relationships - a constraint known as Dunbar's number [12, 13]. Even in our hyper-connected digital age, while we can theoretically access millions through social media, our cognitive architecture remains fundamentally limited. As our technological reach expands exponentially through billions of AI agents, our human capacity to understand and guide these interactions remains fixed. The solution, however, lies not in making individual agents smarter or in replicating human cognition, but in discovering protocols that enable beneficial coordination at scales far beyond human cognitive limits - protocols that enhance rather than replace human agency.

From our morning grocery run to global pandemic response, the future of human flourishing depends on our ability to coordinate effectively at unprecedented scales while preserving individual autonomy. This requires a fundamental shift in how we think about artificial intelligence - moving from enhancing individual capabilities to enabling better collective outcomes through what we call protocol-centric intelligence.

2 From Natural Protocols to Human Coordination

Nature offers profound insights into solving massive-scale coordination challenges [14]. Consider how army ants construct living bridges: each ant follows a remarkably simple protocol - if there's an ant in front, cross the bridge; if not, become part of the bridge. This minimal set of rules, when executed by thousands of ants simultaneously, creates remarkably resilient structures that no individual ant could comprehend or design. A back-of-the-envelope calculation reveals that the strength-to-weight ratio of an ant bridge is 1,000-24,000 times higher than that of a human-built concrete bridge [23].

Similar patterns emerge across diverse natural systems: multicellular organisms coordinate millions of cells through shared genetic protocols, fish schools balance individual survival with group movement through simple interaction rules, and bacterial colonies achieve sophisticated collective behavior despite limited individual capability and conflicting incentives [18, 25, 11]. Our greatest societal challenges mirror this need for coordination beyond individual capability. The COVID-19 pandemic demonstrated how individual outcomes depend not on any single decision but on millions of interlinked choices about testing, isolation, and vaccination. Climate change presents similar dynamics, where individual benefits from carbon emissions create collective harm. These aren't failures of individual decision-making, but our inability to coordinate at scale.

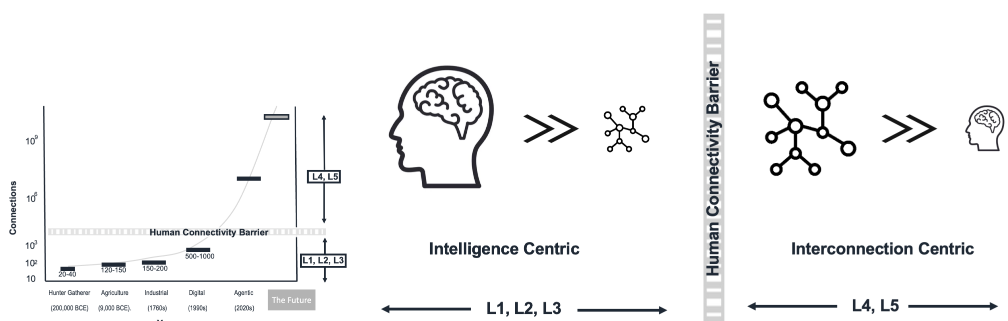

Figure 2: (a) The Human Connectivity Barrier: cognitive limits of human recognition and relationship maintenance (b) Evolution of agentic systems from operating within human cognitive bounds to enabling coordination beyond these natural limits through protocol-centric approaches

Throughout history, humans have developed increasingly sophisticated protocols to transcend individual limitations. Early writing systems enabled asynchronous coordination across time and space. Railway timetables synchronized society across vast distances. The Internet's TCP/IP [4, 10] protocols enabled global information exchange without central control. Yet each advance, while expanding our collective reach, hit the same fundamental barrier: human cognitive limits. Medieval guild masters could train only a handful of apprentices. Railway operators could manage only a limited number of routes. Even in today's digital age, while social media theoretically connects billions, meaningful engagement remains bounded by Dunbar's number - our cognitive limit of around 150 stable relationships and 1500 recognizable faces [1].

This constraint isn't just about individual relationships; it fundamentally limits the complexity of protocols themselves. Traditional protocols must remain comprehensible to the humans who design and oversee them. Even sophisticated systems like financial markets or supply chains must ultimately operate within human cognitive bounds. This creates a crucial tension: as our technological reach expands exponentially through billions of AI agents, our human capacity to understand and guide these interactions remains fixed.

The imminent emergence of billions of AI agents presents an opportunity to transcend these cognitive limits entirely. Unlike historical protocols that had to remain "human-readable", agentic protocols can operate at complexities and scales far beyond human comprehension while still producing beneficial outcomes. However, current AI approaches focused on enhancing individual agent intelligence miss this opportunity. Even sophisticated language models and multi-agent systems remain constrained to human-scale interactions. We need a fundamental shift from making individual agents smarter to enabling smarter interactions - a shift toward protocol-centric intelligence

3 The Path to Protocol-centric Intelligence

Current AI dominated by Large Language Models (LLMs) [3] has dramatically enhanced individual capabilities - from writing assistance to coding help to personal scheduling. However, these systems are designed to operate within human cognitive bounds, limited to enhancing direct person-to-person interactions. Even recent works in multi-agent AI [20, 24, 2, 15] demonstrate sophisticated individual behaviors but they remain constrained to small populations (10-1000 of agents). While these can be deployed within human interaction patterns to - schedule your day, analyze your diet, or recommend a shopping list - they fundamentally operate as amplified individual intelligences. This prevents utility in real-world societal challenges involving millions of individuals. Here, LLM-based agents face the same challenges that have always plagued human coordination: increasing complexity, conflicting goals, and cascading unintended consequences [5].

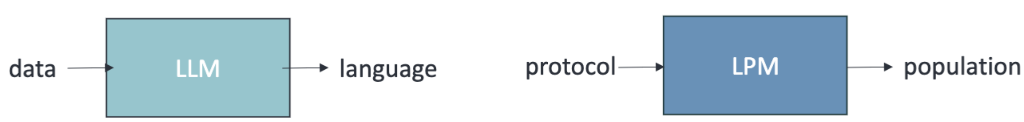

Figure 3: Contrasting approaches of Large Language Models (LLMs) and Large Population Models (LPMs): from data-centric capability enhancement to protocol-centric behavior orchestration

This limitation has sparked an emerging research direction: Large Population Models (LPMs) [9]. Rather than simply enhancing individual agent capabilities, LPMs aim to discover protocols that can inform individual decisions with unprecedented global context. The distinction is crucial: while LLMs process data to enhance individual capabilities, LPMs process protocols to shape collective behavior. This enables LPM-agents to coordinate at unprecedented scales and rest on three key innovations:

- **Differentiable Population Simulation**: Unlike traditional agent-based simulations that merely predict outcomes, LPMs can trace how small changes ripple through entire populations. When millions of agents interact, these systems maintain end-to-end gradients that reveal unintuitive but effective coordination strategies [7, 22]. For instance, they might discover that slightly adjusting just 5% of shoppers' timing can prevent all congestion in a neighborhood, or that prioritizing testing speed over accuracy during a pandemic leads to better collective outcomes.
- **Adaptive Protocol Fields**: Traditional systems rely on fixed rules that quickly become outdated [17]. LPMs instead create "protocol fields" - living systems of rules that continuously

[1]https://www.theatlantic.com/family/archive/2021/05/robin-dunbar-explains-circles-friendship-dunbars-number/618931/

evolve based on real-world feedback [8]. These protocols adapt to changing conditions while preserving individual privacy through decentralized computation. Imagine traffic systems that don't just predict congestion but actively prevent it with subtle, coordinated adjustments across thousands of vehicles.

- **Digital-Physical Bridge**: LPMs bridge the gap between simulation and reality. They first discover effective protocols through massive-scale simulation, then implement them through decentralized networks of edge devices [6]. This creates a continuous feedback loop: protocols shape physical behavior, real-world outcomes refine simulations, and simulations improve protocols. When your phone suggests a slight adjustment to your schedule, it's participating in a living protocol network that's continuously learning and adapting

This technological foundation enables a progression through five distinct levels of social orchestration. Each level expands our ability to participate in increasingly sophisticated collective behavior while preserving individual agency. Most importantly, these levels represent a path to transcend the coordination limits that have constrained human societies throughout history - not by replacing human decision-making, but by enabling unprecedented levels of beneficial coordination when we choose to participate.

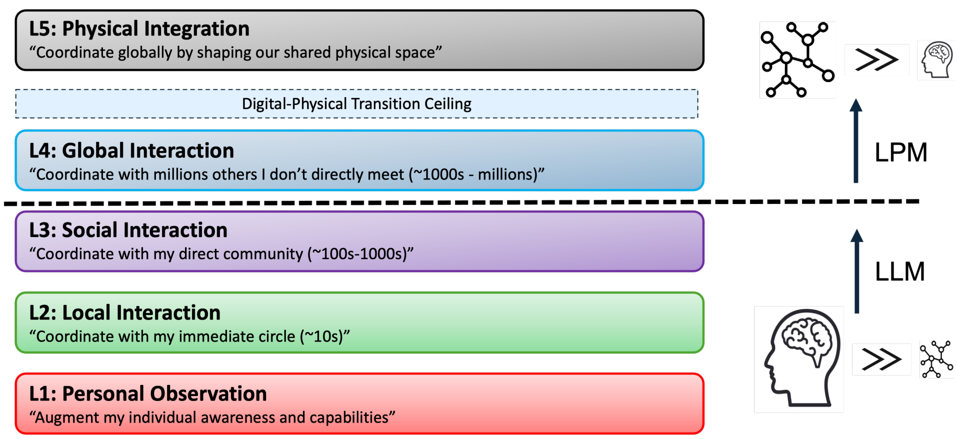

Figure 4: The Individual-Interconnection Transition: From Local to Global Context in Individual Decision-Making. Lower levels (L1-L3) are LLM-based and focus on enhancing individual and small-group capabilities. Higher levels (L4-L5) are LPM-based to enable coordination beyond human cognitive limits by orchestrating millions of simultaneous interactions - in digital (L4) and physical worlds (L5).

4 The Five Levels of Agentic Systems

AI agents will evolve to mirror how humans naturally process information and make decisions - progressing from personal observation to increasingly wider circles of influence. This evolution represents a crucial research direction: how can we develop AI systems that expand individual capabilities while preserving human agency? We propose a five-level framework for this progression, transitioning from current LLM-based personal assistants to future LPM-enabled contextual awareness.

To explore how this framework might manifest in practice, let's project ourselves to 2030 and follow two individuals navigating complex challenges through increasingly sophisticated AI agents: Sarah Chen, a restaurant owner in Kansas City managing through a pandemic, and Michael Roberts, a software engineer planning his first Coachella experience. Their hypothetical journeys illustrate both the potential and challenges of each level.

Level 1: Personal Observation - Enhancing Individual Awareness Current LLM-based systems with retrieval-augmented generation [16] already augment our sensing capabilities, like lane departure warnings in cars. Sarah's COVID analyst agent continuously monitors local transmission patterns, translating complex epidemiological data into clear, actionable insights for her restaurant. Meanwhile,

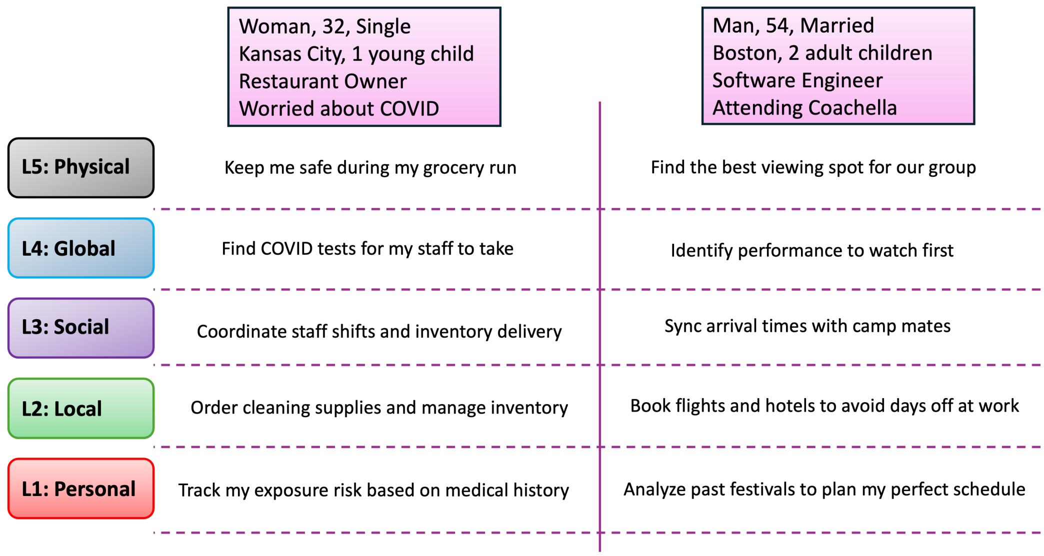

Figure 5: From Information to Orchestration: How agentic systems evolve in real-world scenarios. Following two individuals - a restaurant owner navigating COVID-19 protocols and a music-lover planning his trip to Coachella - we see how interactions progress from simple information gathering (L1) through automated tasks (L2), small group coordination (L3), population-scale interaction (L4), and finally to active shaping of physical behavior (L5). Each level reduces cognitive burden while increasing the system's capacity to coordinate beneficial outcomes at scale.

Michael's festival expert analyzes years of Coachella data to create an optimal experience strategy based on his preferences and comfort needs.

Level 2: Local Interaction - Coordinating Immediate Circles These systems handle specific tasks with contextual intelligence similar to cruise control maintaining vehicle speed. Sarah's specialized agents automatically adjust supply orders based on safety protocols and manage staff health monitoring. Michael's travel agents coordinate accommodations and transport timing for his immediate group, automatically securing better options when available, enabled by emerging tool-use capability [21]

Level 3: Social Coordination - Managing Bounded Networks A key research milestone is extending current multi-agent LLM [26, 24] capabilities, to manage bounded networks effectively, like traffic-aware cruise control responding to nearby vehicles. This will enable Sarah's restaurant to join a network of hundreds of nearby businesses, coordinating delivery schedules to minimize cross-exposure and sharing real-time safety alerts. Michael's agents seamlessly synchronize movement patterns with his extended festival group, smoothly reorganizing meetups while respecting everyone's preferences. This level defines the limit of human cognition.

Level 4: Global Orchestration - Discovering Optimal Strategies These agents will help each individual to coordinate synchronously with millions to optimize decisions, similar to how autonomous vehicles route paths in well-mapped areas. This is being enabled by LPM frameworks like AgentTorch [9] which enable million-scale agent simulation to discover unintuitive patterns. When Sarah asks "Which COVID tests should my staff take today?", her planning agent would process global data to optimize testing strategies. At Coachella, Michael's agent would analyze real-time movement data from 250,000 attendees to identify optimal viewing spots and timing.

Level 5: Physical Integration - Shaping Collective Reality Our most ambitious research goal envisions LPMs bridging digital recommendations and physical reality. Just as fully autonomous vehicles might transform traffic patterns, L5 systems will create living feedback loops between individual actions and physical reality. Sarah's phone doesn't just predict crowd patterns – it actively coordinates with other shoppers' devices to create dynamic "immunity networks," with testing stations adjusting operations based on neighborhood transmission patterns. At Coachella, Michael's agent participates in peer-to-peer protocols that subtly coordinate crowd movements, preventing bottlenecks while ensuring everyone stays comfortable.

This progression represents a fundamental shift: from enhancing individual capabilities (L1-L3) to enabling beneficial collective behavior at unprecedented scales (L4-L5). The key distinction lies in the transition from L3 to L4-L5: while earlier levels operate within human cognitive bounds, higher levels discover and implement protocols that transcend these limitations while preserving individual choice.

5 L4-L5 Transition: From Understanding to Shaping Behavior

While the progression through Levels 1-4 represents expanding circles of influence, each level fundamentally operates in the digital realm - discovering protocols to coordinate among AI agents. Powered by LPMs' ability to enable agents to coordinate millions of digital interactions, Level 5 marks a transformative leap: extending this orchestration into physical reality. The fundamental distinction between Level 4 and Level 5 lies not in the number of coordinated interactions, but in their nature - from coordinating digital information flows to actively shaping physical protocols themselves. This isn't just about processing more data or expanding coordination scope - it's about transforming how artificial intelligence enables individual agents to harmoniously interact with and shape physical reality.

This distinction becomes clear through our examples. At Level 4, Sarah's restaurant management agent can process real-time pandemic data and interactions of millions to recommend optimal testing strategies. But it's limited to prediction and recommendation - when it suggests using rapid tests despite lower accuracy, it's because that's the best coordination possible given how others might behave. At Level 5, the agent doesn't just predict behavior - it actively shapes physical interactions. Testing stations dynamically adjust their operations, ventilation systems create coordinated airflow patterns, and customer devices form living "immunity networks" that prevent transmission hotspots before they form.

This evolution mirrors the evolution of navigation technology with self-driving vehicles. Traditional navigation apps (with current Level 4 vehicles) can analyze traffic data from millions of vehicles to suggest optimal routes to individual drivers. But being limited to digital coordination, they often create "ghost traffic jams" when too many vehicles follow similar recommendations. Yet at Level 5 (as full self-driving emerges), each driver's choices will become part of a living protocol network - their individual routing decisions subtly influencing traffic signals, lane configurations, and other drivers' suggestions in real-time, creating naturally flowing traffic patterns without central control.

Dimension	State
Level 1: Personal	Stable
Level 2: Local	Current
Level 3: Social	Current
Level 4: Global	Emerging
Level 5: Physical	Future

(a) Technology Readiness for agentic systems

Dimension	Level 4: Orchestrate	Level 5: Harmonize
Primary Domain	Digital environments	Physical reality
Individual's Role	System observer	System shaper
Decision Support	Shows possible outcomes	Guides collective evolution
Infrastructure	Works within existing systems	Creates new protocols
Real-world Impact	Through recommendations	Through direct integration
Feedback Loop	One-way learning	Continuous adaptation

(b) The capability transition from L4 to L5 systems

Figure 6: Technology readiness and capability transition from L4 to L5 systems: from predictive recommendations to active physical-world orchestration.

Similarly, Michael's festival experience demonstrates this transition from digital to physical orchestration. At L4, his app processes crowd movement data to suggest optimal viewing locations. At L5, his individual choices about when and where to move become part of a dynamic physical choreography - his device participating in subtle peer-to-peer protocols that naturally prevent congestion while preserving his freedom to explore the festival as he wishes.

The key breakthrough enabling this transition is the protocol-centric design of LPMs. Unlike traditional AI that focuses on enhancing individual decision-making in isolation, LPMs discover protocols that allow individual choices to naturally harmonize through local interactions. They maintain privacy and agency while enabling unprecedented coordination across both digital and physical domains.

6 Risks and Net Benefits of Protocol-Centric Intelligence

The development of protocol-centric intelligence presents three fundamental challenges that must be carefully addressed:

- **Privacy and Control**: While these systems enable decentralized coordination, they could potentially enable unprecedented surveillance. A system designed to coordinate shopping patterns could track movements; one designed for pandemic response could enable unwanted social monitoring. The solution lies in privacy-preserving computation and decentralized protocols that make such abuse technically impossible, not just prohibited.
- **Digital Divide 2.0**: Without careful design, these coordination capabilities could create "protocol privilege" - where those with access to advanced protocols optimize their lives at the expense of others. Similar to how internet access defines opportunities today, protocol access could become a new axis of inequality.
- **The Agency Paradox**: In pursuing collective efficiency, poorly designed systems could subtly restrict individual choice. The challenge is maintaining true optionality while enabling coordination benefits.

However, the potential benefits of protocol-centric intelligence are transformative:

- **Public Health**: Rather than choosing between individual freedom and collective safety, adaptive protocols could enable fine-grained pandemic responses that maintain both. Privacy-preserving contact tracing networks could provide early warnings while protecting individual privacy.
- **Climate Action**: Instead of relying on regulations or individual sacrifice, these systems could discover protocols that make sustainable choices naturally advantageous. From optimizing shared transportation to coordinating energy use, they could help address climate change while improving quality of life.
- **Urban Living**: Cities could become living systems that adapt to inhabitants' needs, preventing congestion before it forms and dynamically adjusting public spaces to enhance density without sacrificing livability.

The net positive impact emerges from three key principles: i) Technical decentralization that prevents central control, ii) Enhancement of human agency rather than replacement, iii) Universal access by design to reduce inequalities.

7 Conclusion: Choosing Our Protocol-Centric Future

Return to our opening contrast: Star Trek's abundance through coordination versus Elysium's scarcity through fragmentation. The difference lies not in raw technological capability - both societies had advanced AI - but in their ability to coordinate human activity at scale while preserving individual agency. As billions of AI agents enter our society, we face the same choice. Will these agents enhance our autonomy by enabling unprecedented coordination, or diminish it through fragmentation and control? The answer lies not in making individual agents smarter, but in discovering protocols that enable beneficial collective behavior while preserving individual choice.

Our morning grocery run in 2030 could reflect either future. In one, AI agents subtly coordinate millions of individual choices to prevent congestion, reduce waste, and enhance everyone's experience. In another, these same technologies create invisible barriers, privileging some while excluding others. The difference lies in the protocols we discover and implement today.

The future of artificial intelligence depends not on replicating human cognition, but on enabling unprecedented human coordination. By focusing on protocol-centric intelligence - built on principles of decentralization, enhancement, and universal access - we can ensure that technological advancement serves to expand rather than constrain human potential. The choice between abundance through coordination and scarcity through fragmentation isn't predetermined - it's a protocol we get to write.

References

[1] Elysium. Wikipedia, 2013.

[2] Altera AL, Andrew Ahn, Nic Becker, Stephanie Carroll, Nico Christie, Manuel Cortes, Arda Demirci, Melissa Du, Frankie Li, Shuying Luo, et al. Project sid: Many-agent simulations toward ai civilization. *arXiv preprint arXiv:2411.00114*, 2024.

[3] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.

[4] Vinton G Cerf and Robert E Kahn. A protocol for packet network intercommunication. *IEEE Transactions on Communications*, 22(5):637–648, 1974.

[5] Ayush Chopra, Shashank Kumar, Nurullah Giray-Kuru, Ramesh Raskar, and Arnau Quera-Bofarull. On the limits of agency in agent-based models. In *Proceedings of the 24th International Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS '25, 2025.

[6] Ayush Chopra, Arnau Quera-Bofarull, Nurullah Giray-Kuru, Michael Wooldridge, and Ramesh Raskar. Private agent-based modeling. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS '24, 2024.

[7] Ayush Chopra, Alexander Rodríguez, Jayakumar Subramanian, Arnau Quera-Bofarull, Balaji Krishnamurthy, B Aditya Prakash, and Ramesh Raskar. Differentiable agent-based epidemiology. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS '23, 2023.

[8] Ayush Chopra, Jayakumar Subramanian, Balaji Krishnamurthy, and Ramesh Raskar. flame: A framework for learning in agent-based models. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, pages 391–399, 2024.

[9] Ayush Chopra and team. AgentTorch: Large population models, 2024.

[10] David D Clark. The design philosophy of the darpa internet protocols. *ACM SIGCOMM Computer Communication Review*, 18(4):106–114, 1988.

[11] Iain D Couzin, Jens Krause, Nigel R Franks, and Simon A Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.

[12] Robin I M Dunbar. Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, 22(6):469–493, 1992.

[13] Russell A Hill and Robin I M Dunbar. Social network size in humans. *Human Nature*, 14(1):53–72, 2003.

[14] Bert Hölldobler and Edward O Wilson. *The Ants*. Harvard University Press, 1990.

[15] Yizhe Huang, Xingbo Wang, Hao Liu, Fanqi Kong, Aoyang Qin, Min Tang, Xiaoxi Wang, Song-Chun Zhu, Mingjie Bi, Siyuan Qi, et al. Adasociety: An adaptive environment with social structures for multi-agent decision-making. *arXiv preprint arXiv:2411.03865*, 2024.

[16] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.

[17] Eric S Maskin. Mechanism design: How to implement social goals. *American Economic Review*, 98(3):567–576, 2008.

[18] Richard E Michod. Evolution of individuality during the transition from unicellular to multicellular life. *Proceedings of the National Academy of Sciences*, 104(suppl 1):8613–8618, 2007.

[19] Michael Okuda and Denise Okuda. *The Star Trek Encyclopedia*. Simon and Schuster, 2011.

[20] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22, 2023.

[21] Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*, 2023.

[22] Arnau Quera-Bofarull, Ayush Chopra, Joseph Aylett-Bullock, Carolina Cuesta-Lazaro, Anisoara Calinescu, Ramesh Raskar, and Michael Wooldridge. Don't simulate twice: one-shot sensitivity analyses via automatic differentiation. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS '23, 2023.

[23] Chris R Reid, Matthew J Lutz, Scott Powell, Albert B Kao, Iain D Couzin, and Simon Garnier. Army ants dynamically adjust living bridges in response to a cost-benefit trade-off. *Proceedings of the National Academy of Science*, November 2015.

[24] Alexander Sasha Vezhnevets, John P Agapiou, Avia Aharon, Ron Ziv, Jayd Matyas, Edgar A Duéñez-Guzmán, William A Cunningham, Simon Osindero, Danny Karmon, and Joel Z Leibo. Generative agent-based modeling with actions grounded in physical, social, or digital space using concordia. *arXiv preprint arXiv:2312.03664*, 2023.

[25] Stuart A West, Ashleigh S Griffin, Andy Gardner, and Stephen P Diggle. Social evolution theory for microorganisms. *Nature Reviews Microbiology*, 4(8):597–607, 2006.

[26] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. Autogen: Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*, 2023.