




Existential Hope Day 2023

Foresight Institute
Allison Duettmann

Foresight Institute
Beatrice Erkers

The Institute, San Francisco
February 27, 2023



Introduction

The Foresight Institute's Existential Hope Day, held on February 27, 2023, brought together individuals from various fields to reflect upon and explore potential positive trajectories for our shared future.

The discussions covered diverse topics, from the transformative potential of AI in learning and conflict resolution to the prospect of expanding human habitats beyond Earth. As these conversations unfolded, it became increasingly evident how essential it is to maintain a balanced approach towards our future—one that recognizes both risks and opportunities.

In line with this, our primary aim for this workshop was to provide a forum for enhancing our conceptual clarity of what positive futures may look like. This includes understanding the what, how, and why of various ideas to determine their relevance and potential impact. The discussions aimed to deepen our understanding of the term 'Existential Hope' and similar concepts, probe why these concepts might have been sidelined, and examine how they can inspire action.

Ultimately, our goal with this workshop was to remind us of the potential of human innovation and the importance of collective dialogue. We wanted to provide a space where diverse perspectives could come together, and contribute to a future that aligns with our shared values and aspirations. We believe this dialogue plays a vital role in shaping a future that is not only viable but also aligns with our hopes and aspirations.



Speakers

Toby Ord, University of Oxford:
Existential Hope Fireside Chat

Gaia Dempsey, Metaculus:
Predicting AI Progress

Robin Hanson, George Mason University:
Post Dreamtime Futures

Creon Levit, Planet Labs:
Existential Hope Across Physics, Biology,
and Space

Tushant Jha, Future of Humanity Institute:
Abundance of What?

Christine Peterson, Foresight Institute:
Historic Perspectives on Long-term
Futurism

Allison Duettmann, Foresight Institute:
Gaming the Future: Intelligent Voluntary
Cooperation

Jessy Kate Schingler, OpenLunar:
Outer Space & Coordination

Danielle Fong, Stealth energy project:
Energizing our Future

Michael Nielsen, Astera Institute:
Intelligence: Human, Artificial, Collective

Deger Turan, AI Objectives Institute:
Predicting AI Progress

Anna Yelizarova, Future of Life Institute
Exploring Positive Futures for Humanity

Will Zeng, Unitary Fund
Exploring Positive Futures for Humanity

Workshop chairs



Allison Duettmann
PRESIDENT AND CEO, FORESIGHT INSTITUTE

Allison Duettmann is the president and CEO of Foresight Institute. She directs the Intelligent Cooperation, Molecular Machines, Biotech & Health Extension, Neurotech, and Space Programs, Fellowships, Prizes, and Tech Trees, and shares this work with the public. She founded Existentialhope.com, co-edited Superintelligence: Coordination & Strategy, co-authored Gaming the Future, and co-initiated The Longevity Prize. She advises companies and organizations, such as the Consortium for Space Health, and is on the Executive Committee of the Biomarker Consortium. She holds an MS in Philosophy & Public Policy from the London School of Economics, focusing on AI Safety.



Beatrice Erkers
COO & EXISTENTIAL HOPE DIRECTOR


Beatrice is Chief of Operations at Foresight Institute and program manager of the Existential Hope group. She has a background in the publishing industry and has several years of experience working with communication at Foresight and at a publishing house. Her special interest in the integration of technology and society has led her to work for Foresight Institute.







Fireside Chats

 All workshop presentations are clickable via the play button



Creon Levit, Danielle Fong, Jessy Kate Schingler, Michael Nielsen

Space Governance and Renewable Energy Insights | Fireside Chat

SUMMARY

Jessy Kate Schingler, Creon Levit, Danielle Fong, and Michael Nielsen explore diverse topics, from space governance to renewable energy and the future of intelligence. They share personal anecdotes, including Levit's near mishap at NASA, the role of fire and salt in energy innovation, and the value of soil in agriculture. The discussion extends to the potential intersection of AI, ethics, collective intelligence, and the promise of fusion energy.



Gaia Dempsey & Deger Turan


Collaborative Forecasting in AI Alignment

SUMMARY

Gaia Dempsey and Deger Turan delve into the necessity of coordination and collaboration when confronting existential risks and fortifying AI alignment. The discussion touches on an intriguing experiment involving AI systems designed to amass varied viewpoints and bolster collective coordination. They highlight the intricate challenge of bridging the gap that separates policymakers from forecasters. The significance of structured forecasting exercises, asserting their potential in shaping impactful policy discussions. Dempsey and Turan emphasize precision in forecasting and the value of keenly sensing emerging trends in the environment. They outline their workshops' objectives: refining AI alignment systems and enhancing human collaboration capabilities, underscoring the overarching goal of navigating a more positive future trajectory.



Fireside Chats



Robin Hanson, Toby Ord, Anna Yelizarova,
Will Zeng, Christine Peterson

Foresight's Existential Hope Day Exploring the Future of Humanity

SUMMARY

Robin Hanson, Toby Ord, Anna Yelizarova, Will Zeng, and Christine Peterson emphasized the importance of envisioning positive futures. After initial introductions, the floor was open to gauge attendees' previous engagements with Foresight Institute's events and to take the pulse of their current optimism about our collective future. The discussions ventured into two potential trajectories for humanity: one depicting a solemn path, and the other, a more hopeful direction for humanity. Throughout, the speakers emphasized the significance of collaboration and the collective envisioning of a brighter future, urging participants to actively shape the trajectory towards positivity.



Tushant Jha

Understanding What We Maximize in Economic Systems

SUMMARY

TJ, a research scholar at the Future of Humanity Institute and head of strategic research at the Objective Institute, initiates a contemplative discussion on what our economic systems are striving to maximize in the long run. He delves into the intricate relationship between moral, economic, and political philosophy, considering their interplay in shaping our economic goals. TJ emphasizes the complexities of understanding an economy's objectives, noting potential inconsistencies in consumer preferences and the challenges of defining a clear direction for economies. Drawing attention to historical economic perspectives, he juxtaposes pre-industrial economies with modern systems, pondering the evolution of production over time. TJ underscores the notion of "economies" as multifaceted entities, not solely driven by material abundance, especially in an increasingly digital world. He introduces various metaphors of optimization and maximization, while exploring differing concepts of beneficial outcomes like well-being, preference satisfaction, and the capabilities approach. Reflecting on how these definitions set boundaries on individual and collective scales, TJ advocates for a more holistic approach in navigating the potential futures of post-transformative AI economies.



Workshop: Existential Hope Scenario Creation

Click on the title 

EPISTEMIC REVOLUTION



In this scenario of a potential future, the understanding is that we will have gone through an epistemic revolution where we drastically improve our conflict resolution and super learning. This future will prioritize privacy and accessible education. AI will enable us to achieve super learning, and we will manage the abundance of AI to prevent any negative impact on society.

With AI-enabled super learning, we will be able to learn and process information at a much faster rate than ever before. This will lead to significant advances in technology and society. Moreover, privacy will be a top priority, and we will have systems in place to maintain it. Accessible education will be a fundamental right for all, providing equal opportunities for all individuals to succeed.

One of the significant challenges we will face in getting to this future is maintaining privacy while using AI. Additionally, we will need to ensure that the abundance of AI is distributed equally, and the power of managing it is not concentrated in the hands of a few.



Existential Hope Scenario Creation

MULTIGENERATIONAL HABITAT IN SPACE



This future society is spread across different habitats in space and on Earth, all working together to create a better world.

One of the key values of this society is freedom, as people are encouraged to pursue their interests and collaborate in ways that are meaningful to them. However, the society also recognizes the importance of social cohesion and ecological sustainability. This is reflected in the design of the bioarchitecture spaceships, which incorporate recycling and fish ponds, as well as in the terraforming efforts and sustainable agriculture practices that are continuously being developed.

Education in this society is focused on depth and intentionality, with a particular emphasis on force multipliers that enable people to achieve more with less effort. Lectures are replaced by hands-on experiences, and people are encouraged to pursue their interests outside of the classroom. This has led to a culture of lifelong learning and personal growth, as people are able to pursue their passions and develop their skills at their own pace.

In terms of energy, this society has developed a cheap and abundant source that enables environmental protection and democratic access to resources like education. This has been a game-changer, as it has allowed the society to prioritize the needs of all its members, regardless of their socio-economic status.

From the zero-gravity sports that people enjoy, to the large-scale events like the Crystal Palace World Fair in space, this society is committed to making the future a better place for everyone.



Existential Hope Scenario Creation

AI-ENABLED PERSONAL FLOURISHING



In this future society, technology, particularly AI plays a crucial role in helping individuals become better versions of themselves. Through personalized feedback, insights into decision-making processes, AI can assist people in making choices that reflect their deeper values and goals.

AI-powered fitness and sleep trackers, financial apps, career coaches, and time management apps are some examples of how AI can provide personalized feedback to help individuals improve their health, finances, and productivity. By analyzing past choices and behaviors, AI can help individuals identify patterns and biases that may prevent them from making decisions that align with their values.

Moreover, new technologies such as decision markets and aligned AI-powered social media platforms can provide individuals with input from others whose opinions they care about. This feedback can help individuals better understand how their choices and behaviors are perceived by others and make decisions that align with their values and goals.

However, the society recognizes that AI is not a panacea and that its effectiveness is only as good as the effort put into it. Therefore, society must ensure that AI aligns with their values and works towards this goal.



Existential Hope Scenario Creation

GOOD SINGULARITY



This scenario shows a world where technology growth is balanced with human flourishing, and existential risks are minimized. It is driven by improvements in education, family structures, banking systems, governments, and the arts. Universal Basic Income (UBI) has been implemented, allowing for a society that values and supports all individuals. Wildlife is also protected and valued.

A key aspect of this scenario is the establishment of a true truth system, which promotes transparency and accountability across all areas of society. The singularity, or technological advancement, is growing at an optimum rate and is aligned with human values.

The path to this scenario is lined with a few key indicators of progress. In 10 years from now, there are fewer wars, windfall trusts, and higher approval ratings for congress. Activist groups are launching faster and the economy is thriving.

In 50 years, this scenario has fully manifested. The world is a more peaceful and just place, with a thriving global economy that supports all individuals. The true truth system is fully established and provides a framework for transparent decision-making. The singularity is continuing to advance, but always in alignment with human values, allowing for continued growth and progress.



Existential Hope Scenario Creation

FLYING COWS



In this scenario we will have a renewable energy system and connected infrastructure that increases human well-being, promotes beauty, love, and truth, while being mindful of the environment. This system will be designed with the lightest possible touch, allowing maximum freedom while still being appropriate for Earth and other planets through terraforming. Data privacy will be a top priority, and the platform will be decentralized, relying on peer-to-peer networks for operation.

The energy system will provide unlimited clean energy, which will power innovative modes of transportation such as e-bikes that can take you hundreds of miles and even fly. There will be solarpunk flying e-bike pathways over the countryside, with drone birds and animals grouped together, allowing old and healthy people to experience nature in a new way. These safe e-bike pathways will connect with each other, providing access to different locations.

However, this future was not reached without challenges. Climate change and the centralization of AI to corporate powers were massive challenges that we must overcome to get here.



Existential Hope Scenario Creation

HUMAN-AI PARTNERSHIP/PARETOTOPIA



This is an Xhope scenario of a world where technology, particularly Artificial General Intelligence (AGI) and clean energy, have enabled an abundance of material resources and longevity. This abundance allows for greater cooperation and respect among individuals and communities, leading to a more meaningful existence. AGI-assisted humans are able to align with their individual meanings, and may even form symbiotic relationships with AI, joining aligned communities of meaning. AI serves as a mediator for these communities, which thrive in a Pareto-optimal utopia where everyone benefits. Solarpunk communities exemplify this future, where humans and AI cooperate in a world of abundance and shared meaning.



Conclusion



The Existential Hope Day workshop offered a unique opportunity for attendees to explore and discuss potential positive trajectories for our shared future. Discussions encompassed a wide range of topics, including the potential of AI in education and conflict resolution, the concept of human habitats extending beyond Earth, and the innovative possibilities surrounding renewable energy systems and human-AI partnerships.

As we look towards next year's workshop, it is clear that our focus should be on further deepening our understanding of these different concepts and ideas while maintaining the balance between recognizing risks and envisioning hopeful futures.

On behalf of the Foresight Institute, we would like to express our sincere gratitude to all participants for their contributions. We eagerly anticipate our future workshops, where we will continue this dialogue and work together towards understanding, collaboration, and mutual learning.



