

The background of the cover is a composite image of Earth from space. The left side shows a bright, curved horizon of the planet, with swirling white and grey cloud patterns over the oceans. The right side shows a dark, starry night sky with a dense, glowing spiral of golden-yellow city lights, representing a global view of human civilization.

AN ANTHOLOGY OF GLOBAL RISK

EDITED BY
SJ BEARD AND TOM HOBSON



<https://www.openbookpublishers.com>

©2024 SJ Beard and Tom Hobson

Copyright of individual chapters is maintained by the chapter's authors



This work is licensed under an Attribution-NonCommercial 4.0 International (CC BY-NC 4.0). This license allows you to share, copy, distribute and transmit the text; to adapt the text for non-commercial purposes of the text providing attribution is made to the authors (but not in any way that suggests that they endorse you or your use of the work). Attribution should include the following information:

SJ Beard and Tom Hobson (eds), *An Anthology of Global Risk*. Cambridge, UK: Open Book Publishers, 2024, <https://doi.org/10.11647/OBP.0360>

Copyright and permissions for the reuse of many of the images included in this publication differ from the above. This information is provided in the captions and in the list of illustrations. Every effort has been made to identify and contact copyright holders and any omission or error will be corrected if notification is made to the publisher.

Further details about CC BY-NC licenses are available at <http://creativecommons.org/licenses/by-nc/4.0/>

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at <https://archive.org/web>

Digital material and resources associated with this volume are available at <https://doi.org/10.11647/OBP.0360#resources>

ISBN Paperback: 978-1-80511-114-6

ISBN Hardback: 978-1-80511-115-3

ISBN Digital (PDF): 978-1-80511-116-0

ISBN Digital eBook (EPUB): 978-1-80511-117-7

ISBN XML: 978-1-80511-119-1

ISBN HTML: 978-1-80511-120-7

DOI: 10.11647/OBP.0360

Cover image: Javier Miranda, Alien planet, June 18, 2022, <https://unsplash.com/photos/nc1zsYGkLFA>

Cover design: Jeevanjot Kaur Nagpal

8. Exploring Artificial Intelligence Futures

Shahar Avin

Highlights:

- This chapter provides a survey and initial categorisation of tools for exploring different futures for Artificial Intelligence, drawing mainly on work in the humanities.
- While no tools exist to reliably predict the future of AI, they can still help us expand our range of possible futures to reduce unexpected surprises and create common languages and models that enable constructive conversations about the kinds of futures we should try to occupy or avoid.
- Fictional narratives have long dominated thinking about AI futures but vary greatly in their degree of realism. They tend to suffer from a range of issues, including the need to entertain audiences, pressure to embody AI in physical forms like robots, a lack of diversity in authorship and representation, and limited accountability for their claims.
- Researchers from a variety of disciplines have produced high quality work drawing on insights from their fields. However, their predictions tend to fare poorly due to factors such as biases, partial perspectives, non-linear trends, and hidden feedback mechanisms. Furthermore, disagreement between experts can have a paralysing effect for audiences.

- Group-based futures exploration can address some of these challenges using techniques like expert surveys, polling, interdisciplinary futures exercises, and expert elicitation. There are also opportunities to extrapolate futures from past and current data trends. Perhaps the most promising tools at our disposal are participatory futuring tools, including workshops, scenarios, and role-plays; however, these need to be realistic, integrative, and data-driven.

This chapter provides an alternative perspective on futures and foresight techniques, drawing more on work in the humanities, and their applications to Artificial Intelligence. The chapter's identification of high-quality scenario role-plays as an important methodological tool led directly to the development of Intelligence Rising (<https://intelligencerising.org>). The development of this tool is described in *Exploring AI Futures Through Role Play*¹ and is the subject of ongoing research, with further papers forthcoming. Group-based, collaborative and collective forms of knowledge generation and futures exploration are also discussed in Chapters 11 and 16.

1. Introduction

"Artificial Intelligence" (AI) is one of the more hyped-up terms in our current world, across academia, industry, policy and society.² The interest in AI, which long predates the current fascination, has given rise to numerous tools and methods to explore the potential futures of the technology, and its impact on human lives in a great variety of domains. While such visions are often drawn to utopian or dystopian extremes, more nuanced perspectives are also plentiful and varied, drawing on the history of the field, measurable progress and domain-specific expertise to extrapolate into possible future trends.

This chapter presents a survey of the different methods available for the exploration of AI futures, from narrative fiction in novels and movies, through disciplinary expert study of e.g. economic or philosophical aspects of AI futures, to integrative, interdisciplinary and participatory methods of exploring AI futures.

I begin in this section with setting common terms and boundaries for the discussion: the boundaries of “Artificial Intelligence” for the purposes of this chapter, certain contemporary technologies and trends that help ground and define the space of exploration, and an outline of the utopian and dystopian extremes that bound the current imagination of AI futures. I then go through each method of futures exploration in turn, providing a few examples and discussing some of the advantages and shortcomings of each. I conclude with a summary of the different methods and suggestions of strategies that may help furnish us with better information and expectations as we progress into a future shaped by AI.

1.1 Defining Artificial Intelligence

Given the newfound interest in AI, it is important to remember the history of AI as a field of research originating from work during the Second World War on computation and encryption, and the visions of the field’s founders of machines that can learn and think like humans.³

While a precise definition of AI is elusive, I will satisfy myself with an analogy to artificial hearts and lungs: machines that can perform (some of) the functions of biological systems, in this case the human or animal brain/nervous system, while at the same time lacking other functions and often differing significantly in shape, material and other properties; this behavioural definition coheres well with the imitation game, or Turing test, that focuses on the machine “passing as” a human in the performance of a specific, delineated task within a specific, delineated domain. As the tasks become more vague, multifaceted and rich, and the domain becomes wider and less well defined, we move on the spectrum from narrow to general intelligence.⁴

The history of the field of AI research shows how wrong we tend to be, *a priori*, about which tasks are going to be easy, and which will be hard, for a machine to perform intelligently.⁵ Breakthroughs in the field are often indexed to new exemplars of classes of tasks being successfully automated — for example, game playing⁶ or image classification.⁷

1.2 Contemporary Artificial Intelligence

The current AI hype cycle is dominated by machine learning, and in particular by deep learning.⁸ Relying on artificial neural networks, which emerged as broadly neurologically inspired algorithms in the second half of the 20th century,⁹ these methods gained newfound success with the increasing availability of fast hardware and of large labelled datasets.¹⁰

In recent years we have seen increasing applications of deep learning in image classification, captioning, text comprehension, machine translation, and other domains. In essence, the statistically driven pattern recognition afforded by these technologies presented a sharp break from previous conceptions of AI as logic/rule-based, and a transition from the domain of explicit expert knowledge to domains of split-second recognition and response tasks (including, for example, driving-related tasks). However, the revolution also touched on expert domains that rely on pattern recognition, including medical image diagnosis¹¹ and Go game-play.¹²

Alongside these broadly positive developments, we have seen more ethically questionable applications, including in speech¹³ and video synthesis¹⁴ that mimics existing individuals, in learning to execute cyber attacks,¹⁵ and in profiling and tracking individuals and crowds based on visual, behavioural and social patterns.¹⁶ Existing and near future technologies enable a range of malicious use cases which require expanded or novel policy responses.¹⁷

1.3 Possible Artificial Intelligence futures

As we look further into the future, our imagination is guided by common tropes and narratives that predate the AI revolution.¹⁸

On the utopian end, super-intelligent thinking machines that have our interests as their guide, or with which we merge, could solve problems that have previously proven too hard to us mere humans, from challenges of environmental management and sustainability, to advanced energy sources and manufacturing techniques, to new forms of non-violent communication and new worlds of entertainment, to medical and biological advances that will make diseases a thing of

the past, including the most terrifying disease of all — ageing and death.¹⁹

On the dystopian end, robotic armies, efficient and entirely lacking in compassion, coupled with the ability to tailor propaganda to every individual in every context on a massive scale, suggest a future captured by the power-hungry, ruthless few, with no hope of freedom or revolution.²⁰

Worse still, if we ever create super-intelligent artificial systems, yet fail to align them with humanity's best interests, we may unleash a process of relentless optimisation, which will (gradually or rapidly) make our planet an uninhabitable environment for humans.²¹

The danger with extreme utopian and dystopian visions of technology futures is that they chart out what biologist Drew Endy called “the half pipe of doom”,²² a dynamic where all attention is focused on these extreme visions. More attention is warranted for mapping out the rich and complex space in between these extremes.

2. Exploring Artificial Intelligence Futures

We are not mere bystanders in this technological revolution. The futures we occupy will be futures of our own making, by action or inaction. To take meaningful action, we must come prepared with a range of alternatives, intervention points, a map of powerful actors and frameworks of critique. As the technical advances increasingly become widely accessible (at least on some level), it is our responsibility, as scholars, policy makers, and citizens, to engage with the technical literature and communities, to make sure our input is informed and realistic.

While it is the responsibility of the technical community to engage audiences affected by their creation (which, in the context of AI technologies, seems to be everyone), there is also a responsibility for those in the relevant positions to furnish decision-makers (again, broadly construed) with rich and diverse, yet fact-based and informed, futures narratives, maps and scenarios. Below I will survey a variety of tools available to us for exploring such futures, pointing out a few examples for each and considering advantages and limitations for each tool.

As a general note, this survey aims to be illustrative and comprehensive, but does not claim to be exhaustive. The examples chosen are by no means representative or exemplary — they are strongly biased by my regional, linguistic and disciplinary familiarity and preferences. Nonetheless, I hope the overall categorisation, and analysis of merits and limitations, will generalise across languages, regions and disciplines. I look forward to similar surveys from other perspectives and standpoints.

2.1 Fictional narratives

Probably the most widely recognised source of AI futures is fictional narratives, across different media such as print (novels, short stories, and graphic novels), music, films and television. These would often fall within the science-fiction genre, or one of its numerous sub-genres. A few examples, chosen somewhat carelessly from the vast trove of AI fictions, include Asimov's *Robot* series, Leckie's *Imperial Radch* trilogy, Banks' *Culture* novels, Wells' *Murderbot Diaries* series, *The Jetsons*, the *Terminator* franchise of movies and TV series, the movie *Metropolis*, and the musical concept series of the same name by Monée.

Works vary greatly in their degree of realism, from those rich in heavily researched details, to those that deploy fantastical technology as a tool to explore some other topic of interest, such as emotions, power relations, agency or consciousness. As such, fictional AI narratives can be both a source of broadened horizons and challenging ethical questions, but also a source of harm when it comes to exploring *our* AI futures — they can anchor us to extreme, implausible or misleading narratives, and, when they gain widespread popularity, can prevent more nuanced or different narratives from gaining attention.

The challenge for fictional AI narratives to provide useful guidance is further aggravated by four sources: the need to entertain, the pressure to embody, a lack of diversity, and a limited accountability.

2.1.1 *The need to entertain*

Authors and scriptwriters need to eat and pay rent, and the amount of remuneration they receive is linked to the popularity of their creations, either directly through sales or indirectly through the

likelihood of contracting. Especially with high-budget production costs, e.g. in Hollywood films,²³ scripts are likely to be more popular if they elicit a positive response from a broad audience, i.e. when they entertain. There is no *prima facie* reason to think that what makes for good entertainment also makes for a useful guide for the future, and many factors are likely to point to these two coming apart, such as the cognitive load of complexity and other cognitive biases,²⁴ or the appeal of extremes.²⁵

2.1.2 *The pressure to embody*

Especially in visual media, but also in written form, narratives are made more accessible if the AI technologies discussed are somehow concretised or embodied, e.g. in the form of robots, androids, cyborgs or other machine bodies.²⁶ Such embodiment serves as a useful tool for exploring a range of pertinent issues, but also runs the risk of distracting us from other forms of intelligence that are less easy to make tangible, such as algorithms, computer networks, swarm intelligence and adaptive complex systems. The pressure to embody relates to, and is made complicated by, the proliferation of embodied instances and fictions of Artificial Intelligence, either as commercial products²⁷ or as artistic creations of robots and thinking machines in visual and physical forms — for example, robot toys or the illustrations that accompany news articles and publications. In general, as per my definition in the beginning, our understanding of Artificial Intelligence should focus on action and behaviour rather than form, though there are good arguments suggesting the two are linked.²⁸

2.1.3 *Lack of diversity*

While narrative fictions may well provide us with the most rich and diverse exploration of possible AI futures, we should be mindful that not all identities and perspectives are represented in fictional narratives, and that the mere existence of a work does not readily translate into widespread adoption; narratives, like individuals, groups and world views, can be marginalised. While science fiction has been one of the outlets for heterodox and marginalised groups

to make their voices heard,²⁹ this is not universally welcome,³⁰ and the distribution of attention is still heavily skewed towards the most popular works.³¹

2.1.4 Limited accountability

Creators of fictional narratives receive feedback from two main sources, their audience (through purchases and engagement with their works) and their critics. While these sources of feedback may occasionally comment or reflect on a work's ability to guide individuals and publics as they prepare for the future, this is not seen as a main aim of the works not an essential part of it.³² In particular, there is little recognition of the possible harms that can follow misleading representations, though it is reasonable to argue that such harms are limited, especially in the absence of better guidance, and the fact that experts deliberately aiming to provide such guidance tend to fare quite poorly (Armstrong and Sotala, 2015).

2.2 Single-discipline futures explorations

As part of the phenomenon of AI hype, we are seeing an increase in the number of non-fiction books exploring the potential implications of Artificial Intelligence for the future, though of course such books have been published since before the field became established in academia, and previous "AI summers" have led to previous periods of increased publication. The authors who publish on the topic come from a wide range of disciplines, and deploy varying methods and arguments from diverse sources. These contribute to a richer understanding of what is, at heart, a multifaceted phenomenon.

For example, AI researchers spend just as much time on the history and sociology of the field, and on dispelling misconceptions, as they do on laying down observations and arguments with relevance for the future;³³ mathematicians and physicists focus on the world as seen through the lens of information, models and mathematics, and the AI futures that such a perspective underwrites;³⁴ technologists focus on underlying technology trends and quantitative predictions;³⁵ risk

analysts explore the various pathways by which AI technologies could lead to future catastrophes;³⁶ economists focus on the impacts of AI technologies on the economy, productivity and jobs;³⁷ self-published, self-proclaimed business thought-leaders share their advice for the future;³⁸ political commentators write manifestos arguing for a particular future;³⁹ and philosophers examine the very nature of intelligence, and what happens when we extrapolate our understanding of it, and related concepts, into future capabilities that exceed what evolution has been able to generate.⁴⁰

While the quality of research and arguments presented in such works tends to be high (as academic and public reputations are at stake), any predictions presented in such works tend to fare poorly, due to numerous factors including biases, partial perspectives, non-linear and discontinuous trends, hidden feedback mechanisms, and limited ability to calibrate predictions.⁴¹ Furthermore, disagreement between experts, while to be expected given the uncertainties involved, can have a paralysing effect for audiences, a fact that can be exploited.⁴²

If fictional narratives are best seen as a rich and fertile ground for futures imagination (as long as we do not get too distracted by the flashy and popular), expert explorations provide a rich toolset of arguments, trends and perspectives with which we can approach the future with an informed, critical stance, as long as we appreciate the deep uncertainty involved and avoid taking any trend or prediction at face value.

2.3 Group-based futures exploration

The nature of the problem being addressed — what are possible AI futures and which ones we should aim for or avoid (and how) — is inherently complex, multi-faceted and interdisciplinary. It is therefore natural to explore this problem through utilising diverse groups. There are various methods to do this, each with advantages and disadvantages (Rowe and Beard, 2018).

2.3.1 *Expert surveys*

What do different individuals think about the future of AI? One way to find out is to ask them. While survey design is not an easy task, we have the ability to improve upon past designs, and regularly update our questions, the target community, and the knowledge on which they draw (as more experience is gained over time).

Surveys amongst experts have been used in particular to explore questions of timing and broad assessment of impact — when will certain capabilities become available, and will they have a positive or negative impact?⁴³ As surveys only tell us *what* people think, rather than *why* they think it, they are best treated not as a calibrated prediction of the future (as all estimates could be flawed in the same way), but rather a useful data point about what beliefs are prevalent right now, which in itself is useful for exploring what beliefs might hold currency in the future, and how these might affect the future of AI.

2.3.2 *Public polling*

Public polling aims to examine both public understanding of the technology, the desirability of possible applications and concerns about possible uses and misuses of the technology.⁴⁴ While it may be tempting to interpret these polls as “hard data” on public preferences, it should be remembered that many factors affect responses.⁴⁵ In the Royal Society study cited above, conducted by Ipsos Mori, poll findings were compared with surveys of focus groups that had in-depth interactions with experts and structured discussions around the survey questions. Such practices bring polling closer to participatory futures workshops, discussed below.

2.3.3 *Interdisciplinary futures studies*

Often, we would want to go beyond an aggregate of single points-of-view, aiming for a more holistic understanding of some aspect of the future of AI through interactions between experts. Such interactions can be one-off or long standing, and they can be more or less structured (Rowe and Beard, 2018). An example of a broad-scoped, long-term academically led interdisciplinary study is the Stanford

100-year study of Artificial Intelligence.⁴⁶ An example of a more focused study is the workshop that led to the report on the potential for malicious use of Artificial Intelligence.⁴⁷ While such studies offer a depth advantage over surveys, and a diversity advantage over single-domain studies, they still face challenges of scope and inclusion: too narrow focus, on either topic or participants, can lead to a narrow or partial view, while too broad scoping and inclusion can make the process unmanageable.⁴⁸

2.3.4 Evidence synthesis and expert elicitation

With a growing evidence base relevant to AI futures, policy-making and policy-guiding bodies are beginning to conduct structured evidence synthesis studies.⁴⁹ The methodologies for conducting such studies have been improved over the years in other evidence-reliant policy domains, and many lessons can be ported over, such as making evidence synthesis more inclusive, rigorous, transparent and accessible.⁵⁰

We are also seeing efforts from governments to solicit expertise from a broad range of source, as early fact-finding steps that could lead to or inform policy in this space.⁵¹ While such efforts are welcome — both in their interdisciplinary and participatory nature — through their democratic mandate, and through the proximity of expertise and accountable decision making, it should be noted that results still very much depend on the experts in the room, that such exercises tend to avoid areas of high uncertainty or disagreement (which may be the areas demanding most attention), and that the issues are often global and open in nature, limiting the effectiveness of national strategy and regulation.

2.4 Extrapolating from past and current data trends

While historical trends may provide only a limited guide to the future when it comes to emerging technologies,⁵² it is still useful to have an up-to-date understanding of the state-of-the-art, especially when the field is progressing at a rapid pace leaving many outside the cutting edge with an out-dated view of what contemporary capabilities are

(and are not). This is a constructive and interdisciplinary effort, as the tools to measure performance of AI technologies are just as much in flux as the technology itself. Measurements of the technology focus either on performance⁵³ or the resource use of the technology in terms of data or compute,⁵⁴ though other dimensions could also be measured.⁵⁵ Other efforts go beyond the technology itself and also track the ecosystem in which the technology is developed, looking at hardware, conference attendance numbers, publications, enrolment, etc.⁵⁶

2.5 Interactive futures narratives and scenarios

For most of the futures exploration tools described above, the audience is passive, and is being communicated at via text or vision and sound. Even surveys of the public often involve only localised and limited contributions from each individual. However, there also exist tools that enable the audience to take a more active role, either in a pre-defined narrative or in the co-creation of narratives. The emphasis on greater public participation is a key tenant of responsible research and innovation⁵⁷ and it applies with force to the field of Artificial Intelligence.⁵⁸

2.5.1 Participatory futures workshops

On the more formal end, participatory future workshops,⁵⁹ or one of the numerous variations on the theme,⁶⁰ go through a structured engagement between different stakeholders. These reflect the (originally more corporate and less open) processes of scenario planning.⁶¹ Similar to scenario planning, where participants explore a range of possible futures as a team, wargaming⁶² and drama theory⁶³ use role-play to place participants in opposing roles, to explore what strategies may emerge or investigate novel opportunities for cooperation and resolution. While the author knows of no such exercises on long-term AI futures, nearer-term exercises — for example, on autonomous driving — are already taking place.⁶⁴ When such exercises have the support of government and buy-in from both experts and non-experts, they can prove to be

highly valuable tools in preparing for AI futures; indeed, they come close to certain visions of the ideal interaction between science and society.⁶⁵ However, they also require significant resources and expertise to carry out well.

2.5.2 Interactive fictions

At the less participatory end, but still allowing the audience to play a more active role, are interactive fictions, especially in the medium of video games. While Artificial Intelligence, as a long-standing science fiction trope, has been depicted in video games for decades, recent games incorporate more of the nuanced arguments presented about the potential futures and characteristics of AI.

For example, *The Red Strings Club* explores fundamental questions of machine ethics in an interactive dialogue with the player,⁶⁶ and *Universal Paperclips* allows the player to experience a thought experiment created to explore the “orthogonality thesis”,⁶⁷ the argument that arbitrarily high levels of intelligence are compatible with a wide range of ultimate goals, including ones that would seem to us foolish or nonsensical.⁶⁸

Other video games focus less on the narrative element, but rather present a rich simulator in which Artificial Intelligence is one of many technologies available to the player, allowing the exploration of a wide range of future AI scenarios and their interplay with other systems such as diplomacy or resource management. Examples include *Stellaris*,⁶⁹ in which Artificial Intelligence technologies are available to the player as they establish their galactic empire, or the *Superintelligence* mod⁷⁰ for *Sid Meier's Civilisation V*,⁷¹ which allows the player, in the shoes of a world leader, to gain a strategic advantage using AI and achieve a scientific victory by creating an Artificial Superintelligence, while risking the creation of an unsafe superintelligence which could lead to an existential catastrophe.

2.5.3 Role-play scenarios

While video games allow audiences to take a more active role in the exploration of possible AI futures within the game environment, they hardly satisfy the call for public participation in jointly imagining and constructing the future of emerging technologies. To explore AI futures in a collaborative and inclusive manner, experts and audiences must explore them together. One way to achieve this is through the joint exploration of narratives in role-play games.

Scenarios that have been developed with expert participation through any of the methods above, or through other means, can be circulated more broadly as templates for role-play games amongst interested parties. At the hobbyist level, game systems such as *Revolt of the Machines*⁷² and *Mutant: Mechatron*⁷³ allow players to collectively explore a possible AI future. While these are often very entertaining, they may fall into the same failures as narrative fictions. It seems there is currently an unmet need for realistic and engaging AI futures role-play game systems.

3. Summary and Conclusion

As AI hype drives utopian and dystopian visions, while rapid technological progress and adoption leaves many of us uncertain about the future impacts on our lives, the need for rich, informative, and grounded AI futures narratives is clear. It is also clear that there is a wide range of tools to develop such narratives, many of which are available to creators and experts outside the AI research community. It is less clear, however, how best to utilise each of the available tools, with what urgency and in which domains. The table below summarises the different tools surveyed above, with their respective advantages and limitations.

Table 1: Tools to develop Artificial Intelligence futures narratives.

Tool	Existing abundance	Skills and resources required	Advantages	Limitations
Fictional narratives	Overly abundant	Creative writing, production costs for film	Unbridled imagination, (relatively) open participation	Lack of realism, pull to extremes, lack of accountability, lack of diversity, skewed popularity distribution
Single-discipline futures exploration	Growing rapidly, though some disciplines are still missing	Domain expertise, familiarity with AI, forecasting skills	Deep dives into relevant facts and arguments	Predictive power is poor, disagreements can paralyse, not easy to integrate across disciplines
Surveys	Few key studies	Survey design, resources to carry out the survey	Aggregate evidence can counteract some biases, present a snapshot of current beliefs	Survey design is hard, topic in flux, misunderstanding is commonplace; poor predictive power
Interdisciplinary futures exploration	Few but growing rapidly	Interdisciplinary facilitation, network of stakeholders, time and geographic availability	Holistic view of complex topics, opportunity to directly engage with policy-makers and other key stakeholders	Risk of groupthink, conservatism; scoping is difficult: too narrow and miss opportunities and challenges, too broad and becomes intractable
Evidence synthesis	Few	Access to studies in a range of disciplines and expertise to assess them and communicate findings	Evidence-based holistic picture drawing on a wide range of works, prepared with policy in mind	Time and labour intensive, evidence may be partial and rapidly changing, best practices still evolving

Tool	Existing abundance	Skills and resources required	Advantages	Limitations
Extrapolating data trends	Few key hubs, abundant but disperse data	Familiarity with the field and the techniques of AI, measurement platforms, data harvesting and curation	Historical and contemporary measurements can be largely uncontested, informative	Difficult to extrapolate from past trends due to non-linearity, feedback, potential for discontinuity; need to constantly evolve and adapt measurements
Participatory futures workshops	None on long-term AI, few on short term issues such as self-driving cars	Buy-in from experts and non-expert participants, budget for workshops, facilitation skills, time of participants	Participatory, expert-informed exploration of future scenarios, legitimacy for policy guidance	Difficult to get buy-in and time commitment from experts and stakeholders, requires significant investment to tutor non-experts
Interactive fictions	Several, though few with realistic representations informed by recent advances	Game development skills and budget	Audience takes an active role, can explore alternatives, simulators offer a combinatorial explosion of options	Similar to fictional narratives, plus limitations of what can be represented effectively with limited skills and budget
Role-play scenarios	Few	Facilitation, game/scenario design	Stakeholders can come together to co-explore possible futures	Information gaps in the group can slow down or derail the conversation, strongly depends on the available expertise and facilitation skills

As can be expected, no tool is strictly better than all other tools. Some provide more evidence-based, deep analysis, but tend to be limited in the range of questions they can cover and place barriers on participation. Others allow for more diverse and integrative perspectives, but tend to preclude detailed and in-depth analysis or come at a very high cost in terms of time and facilitation. Instead of judging individual futures narratives in isolation, it may be more useful to look at the entire ecosystem of future AI narratives, asking whether certain narratives are dominating our imagination without sufficient warrant, or if there are tools and narratives that are underutilised or gaining insufficient attention. At present, it seems that not enough attention is being given to data-driven, realistic, integrative, and participatory scenario role-plays, which can build on and integrate a range of other tools and narratives and make them more accessible to a wider audience in a more nuanced way. A more balanced portfolio is called for.

As we act to critique and curate the ecosystem of AI futures, we should keep in mind the aims of these narratives: beyond entertainment and education, there are real ongoing processes of technological development and deployment that currently have, are likely to continue to have, significant social impacts. These processes are not isolated from the societies in which they take place, and the interactions between technology developers, policymakers, diverse stakeholders and numerous publics are mediated and shaped by the futures narratives each group has access to. Thus, AI futures narratives play a crucial role in making sure we arrive at futures of our own choosing, that reflect our values and preferences, that minimise frictions along the path, and that do not take us by surprise. Thus, critique and curation of AI futures is an integral part of the process of responsible development of Artificial Intelligence, a part in which humanities scholars have a significant role to play.

Notes and References

- 1 Avin, S., R. Gruetzmacher and J. Fox. 'Exploring AI futures through role play', *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* (February 2020), pp. 8–14.
- 2 Shoham, Y., R. Perrault, E. Brynjolfsson and J. Clark. *Artificial Intelligence Index — 2017 Annual Report* (2017). <https://aiindex.org/2017-report.pdf>

- 3 Turing, A.M. 'Computing machinery and intelligence', *Mind*, 49 (1950): 433–60.
- 4 Legg, S. and M. Hutter. 'Universal intelligence: A definition of machine intelligence', *Minds and Machines*, 17(4) (2007): 391–444. <https://doi.org/10.1007/s11023-007-9079-x>
- 5 Minsky, M. *Society of Mind*. Simon and Schuster (1988); Moravec, H. *Mind Children: The Future of Robot and Human Intelligence*. Harvard University Press (1988).
- 6 Campbell, M., A. J. Hoane Jr and F. H. Hsu. 'Deep blue', *Artificial Intelligence*, 134(1–2) (2002): 57–83; Silver, D., J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez ... and Y. Chen. 'Mastering the game of Go without human knowledge', *Nature*, 550(7676) (2017): 354. <https://doi.org/10.15368/theses.2018.47>
- 7 Krizhevsky, A., I. Sutskever and G. E. Hinton. 'Imagenet classification with deep convolutional neural networks', *Advances in Neural Information Processing Systems* (2012), pp. 1097–1105. <https://doi.org/10.1145/3065386>
- 8 LeCun, Y., Y. Bengio and G. Hinton. 'Deep learning', *Nature*, 521(7553) (2015): 436. <https://doi.org/10.1038/nature14539>
- 9 Lippmann, R. (1987). An introduction to computing with neural nets. *IEEE Assp magazine*, 4(2), 4–22.
- 10 Amodei, D. & Hernandez, D. (2018) AI and Compute. Open AI blog. Retrieved from <https://blog.openai.com/ai-and-compute/>.
- 11 Esteva, A., B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau and S. Thrun. 'Dermatologist-level classification of skin cancer with deep neural networks', *Nature*, 542(7639) (2017): 115. <https://doi.org/10.1038/nature21056>
- 12 Silver et al. (2017).
- 13 Lyrebird. *We Create the Most Realistic Artificial Voices in the World* (2018). <https://lyrebird.ai/>
- 14 Suwajanakorn, S., S. M. Seitz and I. Kemelmacher-Shlizerman. 'Synthesizing Obama: Learning lip sync from audio', *ACM Transactions on Graphics (TOG)*, 36(4) (2017): 95. <https://doi.org/10.1145/3072959.3073640>
- 15 Frazee, D. *Cyber Grand Challenge (CGC)* (2018). <https://www.darpa.mil/program/cyber-grand-challenge>
- 16 Zhang, C., H. Li, X. Wang and X. Yang. 'Cross-scene crowd counting via deep convolutional neural networks', *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2015), pp. 833–41. <https://doi.org/10.1109/cvpr.2015.7298684>
- 17 Brundage, M., S. Avin, J. Clark, H. Toner, P. Eckersley, B. Garfinkel, ... and H. Anderson. 'The malicious use of artificial intelligence: Forecasting, prevention, and mitigation', *arXiv preprint arXiv:1802.07228* (2018).
- 18 Cave, S. & Dihal, K. (2018). Ancient dreams of intelligent machines: 3,000 years of robots. *Nature*, 559(7715), 473. <https://doi.org/10.1038/d41586-018-05773-y>
- 19 Kurzweil, R. *The Singularity Is Near*. Gerald Duckworth & Co (2010).
- 20 Mozur, P. 'Inside China's dystopian dreams: AI, shame and lots of cameras', *New York Times* (8th July 2018). Turchin, A. and D. Denkenberger. 'Classification of global catastrophic risks connected with artificial intelligence', *AI & SOCIETY* (2018): 1–17. <https://doi.org/10.1007/s00146-018-0845-5>
- 21 Bostrom, N. *Superintelligence: Paths, Dangers, Strategies* (2014).

- 22 Endy, D. *Synthetic Biology — What Should We Be Vibrating About?* TEDxStanford (2014). https://www.youtube.com/watch?v=rf5fTe_i7aA
- 23 De Vany, A. *Hollywood Economics: How Extreme Uncertainty Shapes the Film Industry*. Routledge (2004).
- 24 Yudkowsky, E. 'Artificial intelligence as a positive and negative factor in global risk', *Global Catastrophic Risks*, 1(303) (2008): 184. <https://doi.org/10.1093/oso/9780198570509.003.0021>
- 25 Kareiva, P. and V. Carranza. 'Existential risk due to ecosystem collapse: Nature strikes back', *Futures* (2018); Needham, D. and J. Weitzdörfer (eds.). *Extremes* (Vol. 31). Cambridge University Press (2019). <https://doi.org/10.1016/j.futures.2018.01.001>
- 26 Kakoudaki, D. *Anatomy of a Robot: Literature, Cinema, and the Cultural Work of Artificial People*. Rutgers University Press (2014).
- 27 Harris, J. '16 AI bots with human names', *Chatbots Life* (2017). <https://chatbotslife.com/10-ai-bots-with-human-names-7efd7047be34>
- 28 Shanahan, M. *Embodiment and the Inner Life: Cognition and Consciousness in the Space of Possible Minds*. Oxford University Press (2010).
- 29 Rose, H. 'Science fiction's memory of the future', *Contested Futures: A Sociology of Prospective Techno-Science*. Ashgate (2000): 157–74.
- 30 Oleszczuk, A. 'Sad and rabid puppies: Politicization of the Hugo Award nomination procedure', *New Horizons in English Studies*, (2) (2017): 127. <https://doi.org/10.17951/nh.2017.2.127>
- 31 De Vany (2004).
- 32 Kirby, D. A. *Lab Coats in Hollywood: Science, Scientists, and Cinema*. MIT Press (2011).
- 33 Boden, M. A. *AI: Its Nature and Future*. Oxford University Press (2016); Domingos, P. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. Basic Books (2015); Shanahan, M. *The Technological Singularity*. MIT Press (2015).
- 34 Fry, H. *Hello World: How to Be Human in the Age of the Machine*. Penguin (2018); Tegmark, M. *Life 3.0: Being Human in the Age of Artificial Intelligence*. Knopf (2017).
- 35 Kurzweil (2010).
- 36 Barrett, A. M. and S. D. Baum. 'A model of pathways to artificial superintelligence catastrophe for risk and decision analysis', *Journal of Experimental & Theoretical Artificial Intelligence*, 29(2) (2017): 397–414; Turchin and Denkenberger (2018). <https://doi.org/10.1080/0952813x.2016.1186228>
- 37 Brynjolfsson, E. and A. McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton & Company (2014); Hanson, R. *The Age of Em: Work, Love, and Life When Robots Rule the Earth*. Oxford University Press (2016).
- 38 Hyacinth, B. T. *The Future of Leadership: Rise of Automation, Robotics and Artificial Intelligence*. MBA Caribbean Organisation (2017); Rouhianien, L. *Artificial Intelligence: 101 Things You Must Know Today About Our Future*. Createspace Independent Publishing Platform (2018).
- 39 Srnicek, N. and A. Williams. *Inventing the Future: Postcapitalism and a World Without Work*. Verso Books (2015); Bastani, A. *Fully Automated Luxury Communism: A Manifesto*. Verso (2018).

- 40 Bostrom (2014).
- 41 Armstrong, S. and K. Sotala. 'How we're predicting AI — or failing to', *Beyond Artificial Intelligence*. Springer (2015), pp. 11–29. https://doi.org/10.1007/978-3-319-09668-1_2; Beard, S., T. Rowe and J. Fox. 'An analysis and evaluation of methods currently used to quantify the likelihood of existential hazards', *Futures*, 115 (2020): 102469. <https://doi.org/10.1016/j.futures.2019.102469>; Yudkowsky, E. 'There's no fire alarm for artificial general intelligence', *Machine Intelligence Research Institute* (2017). <https://intelligence.org/2017/10/13/fire-alarm/>
- 42 Baum, S. 'Superintelligence skepticism as a political tool', *Information*, 9(9) (2018): 209.
- 43 Grace, K., J. Salvatier, A. Dafoe, B. Zhang and O. Evans. 'When will AI exceed human performance? Evidence from AI experts', *arXiv preprint arXiv:1705.08807* (2017); Müller, V. C. and N. Bostrom. 'Future progress in artificial intelligence: A survey of expert opinion', *Fundamental Issues of Artificial Intelligence*. Springer (2016), pp. 555–72.
- 44 The Royal Society. *Public Views of Machine Learning* (2017). <https://royalsociety.org/~media/policy/projects/machine-learning/publications/public-views-of-machine-learning-ipsos-mori.pdf>
- 45 Achen, C. H. and L. M. Bartels. *Democracy for Realists: Why Elections Do Not Produce Responsive Government* (Vol. 4). Princeton University Press (2017).
- 46 Grosz, B. J. and P. Stone. 'A century long commitment to assessing Artificial Intelligence and its impact on society', *arXiv preprint arXiv:1808.07899* (2018).
- 47 Brundage, Avin et al. (2018).
- 48 Collins, H. M. and R. Evans. 'The third wave of science studies: Studies of expertise and experience', *Social Studies of Science*, 32(2) (2002): 235–96. <https://doi.org/10.1177/0306312702032002003>; Owens, S. 'Three thoughts on the Third Wave', *Critical policy studies*, 5(3) (2011): 329–33. <https://doi.org/10.1080/19460171.2011.606307>
- 49 British Academy and The Royal Society. *The Impact of Artificial Intelligence on Work* (2018). <https://royalsociety.org/~media/policy/projects/ai-and-work/evidence-synthesis-the-impact-of-AI-on-work.PDF>
- 50 Donnelly, C. A., I. Boyd, P. Campbell, C. Craig, P. Vallance, M. Walport ... and C. Wormald. 'Four principles to make evidence synthesis more useful for policy', *Nature*, 558(7710) (2018): 361. <https://doi.org/10.1038/d41586-018-05414-4>; Sutherland, W. J. and C. F. Wordley. 'A fresh approach to evidence synthesis', *Nature*, 558 (2018): 364–66. <https://doi.org/10.1038/d41586-018-05472-8>
- 51 Felten, E. and T. Lyons. *The Administration's Report on the Future of Artificial Intelligence* (2016). <https://obamawhitehouse.archives.gov/blog/2016/10/12/administrations-report-future-artificial-intelligence>; House of Lords. 'AI in the UK: Ready, willing and able?', *House of Lords Select Committee on Artificial Intelligence Report of Session 2017–19* (2018).
- 52 Farmer, J. D. and F. Lafond. 'How predictable is technological progress?', *Research Policy*, 45(3) (2016): 647–65. <https://doi.org/10.2139/ssrn.2566810>
- 53 Eckersley, P. and Y. Nasser et al. *EFF AI Progress Measurement Project* (2017). <https://eff.org/ai/metrics>
- 54 Amodei and Hernandez (2018).

- 55 Martínez-Plumed, F., S. Avin, M. Brundage, A. Dafoe, S. ÓhÉigeartaigh and J. Hernández-Orallo. 'Accounting for the Neglected Dimensions of AI Progress', *arXiv preprint arXiv:1806.00610* (2018).
- 56 Benaich, N. and I. Hogarth. *The State of Artificial Intelligence in 2018: A Good Old-Fashioned Report* (2018). <https://www.stateof.ai/>; Shoham et al. (2017).
- 57 Owen, R., P. Macnaghten and J. Stilgoe. 'Responsible research and innovation: From science in society to science for society, with society', *Science and Public Policy*, 39(6) (2012): 751–60. <https://doi.org/10.4324/9781003074960-11>
- 58 Stilgoe, J. 'Machine learning, social learning and the governance of self-driving cars', *Social Studies of Science*, 48(1) (2018): 25–56. <https://doi.org/10.2139/ssrn.2937316>
- 59 Jungk, R. and N. Müllert. *Future Workshops: How to Create Desirable Futures*. Institute for Social Inventions (1987).
- 60 Nikolova, B. 'The rise and promise of participatory foresight', *European Journal of Futures Research*, 2(1) (2014): 33. <https://doi.org/10.1007/s40309-013-0033-2>; Oliverio, V. *Participatory Foresight*. Centre for Strategic Futures (2017). <https://www.csf.gov.sg/our-work/Publications/Publication/Index/participatory-foresight>
- 61 Amer, M., T. U. Daim and A. Jetter. 'A review of scenario planning', *Futures*, 46 (2013): 23–40. <https://doi.org/10.1016/j.futures.2012.10.003>
- 62 Perla, P. *The Art of Wargaming: A Guide for Professionals and Hobbyists* (Vol. 89028818). Naval Institute Press (1990).
- 63 Bryant, J. *The Six Dilemmas of Collaboration: Inter-Organisational Relationships as Drama*. Wiley (2002).
- 64 Cohen, T., J. Stilgoe and C. Cavoli. 'Reframing the governance of automotive automation: Insights from UK stakeholder workshops', *Journal of Responsible Innovation* (2018): 1–23. <https://doi.org/10.1080/23299460.2018.1495030>
- 65 Kitcher, P. *Science in a Democratic Society*. Prometheus Books (2011).
- 66 Deconstructeam. *The Red Strings Club* [PC game]. Devolver Digital (2018).
- 67 Lantz, F. (2017) Universal Paperclips [online video game]. Retrieved from <http://www.decisionproblem.com/paperclips/>
- 68 Bostrom (2014).
- 69 Paradox Interactive. *Stellaris* [video game] (2016).
- 70 Shapira, S. and S. Avin. *Superintelligence* [video game mod] (2017). <https://steamcommunity.com/sharedfiles/filedetails/?id=1215263272>
- 71 Firaxis Games. *Sid Meier's Civilization V* [PC game] (2010).
- 72 Fantasy Flight Games. *End of the World: Revolt of the Machines* [roleplaying game] (2016).
- 73 Ligan, F. *Mutant: Mechatron* [roleplaying game] (2017).

