

Toward Mapping the Paths to AGI

Ross Gruetzemacher and David Paradice

¹ Auburn University, Auburn, Alabama, USA
rossg@auburn.edu, dparadice@auburn.edu

Abstract. There is substantial interest in the research community for a map of the paths to artificial general intelligence (AGI), however, no effort toward these ends has been entirely successful. This paper identifies an alternative technique called scenario network mapping that is well suited for the difficulties posed in mapping the paths to AGI. The method is discussed, and a modified version of scenario network mapping is proposed which is intended specifically for the purpose of mapping the paths to AGI. Finally, a scenario network mapping workshopping process is proposed to utilize this method and develop a map of the paths to AGI. This will hopefully lead to discussion and action in the research community for using it in a new effort to map the paths to AGI.

Keywords: AGI, scenario analysis, scenario mapping, technology roadmap.

1 Introduction

Technology roadmaps are a technology management technique that have been used with a large degree of success in a number of different technology research areas [24]. Primarily employed for informing resource allocation, they can also be used to structure and streamline the innovation process, to set targets and expectations, and to identify possible risks or potential roadblocks [27]. Technology roadmaps are perhaps even more valuable for developing artificial general intelligence (AGI) [25]. They can be used to: elucidate biases leading to research for near-term gains, illuminate dead ends in ongoing research, identify hidden problems or prizes in research plans, compare alternate paths, introduce young researchers to the field, align the community, etc. In short, technology roadmaps offer a powerful technology management tool for optimizing the development of AGI.

While many benefits could come from a roadmap to AGI, there are likely many paths to it rather than just one [11]. This poses a major challenge [1], and previous attempts have indicated that traditional technology roadmaps are insufficient for mapping the paths to AGI [10]. Consequently, the technique presented here does not generate merely a single path, but rather a lattice-like structure of interconnected possible paths to AGI. Specifically, this method produces a directed graph that includes two layers of nodes: one for AGI's technological components and another for its milestones. Unlike traditional roadmaps, this approach can enable comparison of the many possible paths to AGI. This would allow researchers to compare the required resources, risks, technological challenges and other crucial factors for developing AGI.

2 Background

A decade has passed since the first attempt to map the path to AGI was conceived [9]. This idea, intended to align the community, would lead to a small workshop in 2009 which built on work from earlier workshops [15, 16] to produce the first roadmap to AGI [1]. Organizers were disappointed that the resulting roadmap was not a straight-forward, road-like path, but rather like climbing the peak of a mountain range, with many possible paths, the easiest of which is difficult to tell from the bottom [11]. Although the results were not what organizers had hoped for [10], much progress has been made toward the milestones that were proposed as a result of the roadmap. In fact, some of the most impressive advances in the past ten years have been in general video-game learning [6, 21, 26], reading and grade school level tasks [23] – domains that represent over 50% of that roadmap. However, the amount of true progress that has been made in these domains is debatable, and progress made on the roadmap is uncertain. What is clear is that while the 2009 roadmap has proved to be a much better guide for AGI progress than forecasts [4], further improvements are still desirable.

Another roadmap toward AGI (or machine intelligence) was proposed in 2016 [20]. This roadmap did not use a structured group process like the 2009 workshop, but rather proposed a full training environment as well as the only end-to-end description of a process for training an AGI agent. However, it lacked concrete proposals for the more challenging tasks that were described, some of which would be critical to the agent. Other intelligence frameworks that have been proposed in the AGI research community, such as NARS, OpenCog or MicroPsi 2 could also be seen as roadmaps to AGI as envisioned by their developers [3, 14, 28]. In fact, one of the challenges that organizers of the 2009 workshop found was the difficulty to get participants to agree on a common direction because they each advocated their own roadmap since it was well suited for their own AGI framework [11]. Although neither a roadmap nor a framework, a 2017 study on creating human-like machines constitutes a significant contribution to the roadmap-oriented literature [17]. Rather than mapping the milestones or specifying a path, this study surveyed the requisite components for a brain-inspired AGI agent. The AI Roadmap Institute¹ has also created a roadmap, however, it is less technical and focuses more on the exploration of an AI Race. While this map was simply a flow chart of possible future scenarios during the development process of AGI, it may be the closest example to the output of this proposed workshoping technique.

All of the relevant previous studies on mapping the paths to AGI have one thing in common: none resulted in a map in the sense discussed here². Technology roadmapping [7] is an established and widely used technique from technology management literature that is useful for supporting strategic planning [22]. It has been used suc-

¹ The AI Roadmap Institute has also thoroughly identified the benefits and uses for roadmaps to AGI [25]. (www.roadmapinstitute.com).

² The notion of a map here more closely resembles a lattice than a flowchart or a technology roadmap. The following section discusses this further, and a generic map of this sort is depicted in Figure 1.

cessfully by numerous organizations and consortiums, including Philips Medical Systems [27] and the Semiconductor Industry Association [24], to foster innovation and to align industry innovation goals. However, the technology roadmapping process is not rigorous, is heavily reliant on visual aids and was considered unsuccessful in the previous attempt to use it for mapping the paths to AGI [10, 11]. Recent work has proposed a new class of scenario analysis techniques, called scenario mapping techniques, due to the common mapping properties they share [12]. These techniques are more suitable for mapping the paths to AGI. Generally, scenario analysis techniques are considered to be a powerful family of techniques that are commonly used by organizations to illuminate blind spots in strategic plans [5]. Their use may be able to identify blind spots in existing AGI frameworks that are difficult for the developers to see. Scenario network mapping (SNM) is a comprehensive, flexible approach for anticipating plausible futures in environments with high levels of uncertainty [2]. Recent work has suggested this technique to be better suited for mapping the paths to AGI than the technology roadmapping procedure or other scenario mapping processes due to its unique workshop style and its ability to model numerous entangled possible paths [13].

Given the progress in AI research over the past ten years and the promise of a new mapping process, we argue that a workshop should be held with AGI experts³ to conduct an updated mapping of the paths to AGI. To these ends, this paper proceeds by first introducing the SNM technique. Then, the outlines a modified SNM process that is specifically tailored for the mapping of the paths to AGI. The paper concludes by urging members of the AGI research community to participate in a workshop for developing a new map of the paths to AGI.

3 Scenario Network Mapping

SNM was first proposed in 2005 to improve upon standard scenario analysis techniques by enabling the use of a large number of possible scenarios, each representing a component of one possible pathway to a particular outcome [19]. Scenario network mapping is intended for scenario planning purposes, however, the technique can also be extended to concepts or ideas for new technologies. The map resulting from SNM is easily modified as the future unfolds by updating it with new events and repositioning the existing components and connections to accommodate the new events⁴. SNM is conducted via four half-day workshops, each ideally with 15-20 participants.

The result of SNM is a directed graph wherein the nodes are components of the pathways and the edges are the causal links between these components. SNM utilizes event trees and the holonic principle (explained below) to enable the generation of a

³ A development workshop has been conducted with early career AGI researchers which was used in the development of the method proposed here. More details can be found at www.rossgritz.com/snm-development-workshop. Further development workshops are recommended for refinement of the technique proposed here.

⁴ In the adapted technique that is the focus of this paper we are concerned with mapping future technologies rather than events.

large number of interconnected scenarios. Event trees are comprised of a hierarchy of antecedents (the roots), the central event (the trunk) and a hierarchy of outcomes (the branches). SNM maps are laid out horizontally so that the depiction of time may flow from left to right, improving readability when stacking these event trees with complex interactions. The holonic principle is another essential feature of SNM that means each node in the resulting graph is simultaneously both a component of the larger system and itself comprised of smaller systems. This principle implies that, if necessary, each component can be broken down further into its constituents for analyzing the relationships with other components in the graph. This is well suited for complex technologies that are poorly understood, and which may be best anticipated through their subcomponents.

The workshoping process is well-documented and includes a user manual for facilitators [18]. A slightly altered process has been developed and widely used for mapping complex networks of components involving interactions between micro and macro level system innovation for sustainability [8]. For the specific purpose proposed here, an altered process has also been developed. Fig. 1 below is adapted from a figure of a generic SNM structure in [8]. We have recreated this figure with modifications consistent with the adaptations for the purposes of this study.

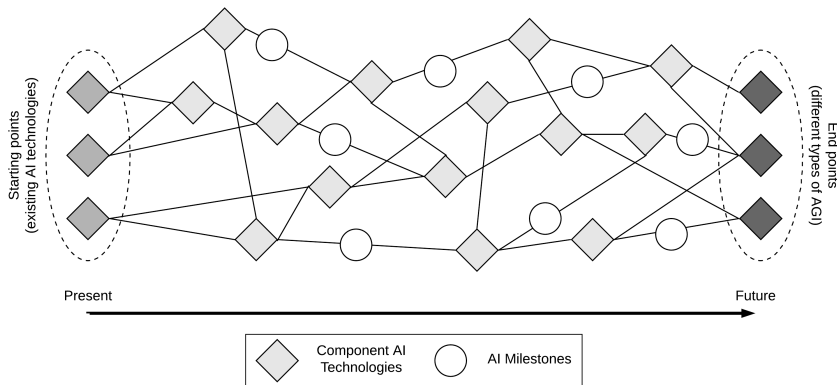


Fig. 1. This figure depicts an example of the lattice-like output from the scenario network mapping workshop that is proposed.

The structure depicted in Fig. 1 can be seen to demonstrate the lattice-like structure that has been described earlier. It can be seen that there is both a technology layer as well as a milestone layer. Fig. 1 depicts a generic map for demonstration purposes only; an actual map would be expected to have many more nodes for technologies as well as milestones. The following section outlines a similar altered SNM process for the purpose of mapping the paths to AGI.

4 AGI Scenario Network Mapping

The original SNM workshopping process involves four half-day workshops that are intended to be conducted over the span of up to four weeks, allowing one week in-between each workshop [19]. This is only reasonable for organizations, and consequently, the proposed AGI-SNM workshopping process proposed here is designed to be suitable for four consecutive half-day workshops over two days. However, the half-day workshops could also be spread out over as many as four weeks for organizations. A single AGI-SNM development workshop has been conducted with early career AGI researchers. The experience from this workshop has helped to develop the workshopping process described here.

The SNM workshopping process requires some specific resources in order to be conducted effectively. The most important resource is the experts. At least 10 are needed for diversity but larger groups can take longer and become more chaotic. Thus, it is recommended to stay between 15 and 20 experts⁵ [19]. Another important resource is the venue; a single room large enough for breakout groups is necessary to maintain an efficient process during breakout sessions. The process also requires a large amount of wall space and freedom from interruptions. Other equipment and materials include size A3 paper, multicolored fine tip markers, multicolored sticky notes, ribbon, masking tape, colored circular stickers (dots for voting), recording equipment (if desired) and a projector for the facilitator [18].

The original workshopping technique uses the first half-day exploring historical antecedents to the current state of events [18]. In general, a substantial amount of content in the original workshop manual had to be adapted for the unique purposes of mapping the paths to AGI⁶. Such variations are often necessary dependent on the use case. The original SNM component workshops from the SNM manual are below⁷.

- Workshop 1: Influences from past and present
 - Introduction
 - Unfinished business
 - Prouds and sorries
 - Scenarios of the recent past
 - Stakeholder map
 - Leaf of goals
- Workshop 2: Generating possibilities
 - Futures wheel
 - Defining paths
 - Backcasting
 - Midcasting
- Workshop 3: Mapping paths to the future

⁵ For organizations, it is suggested that well-informed outsiders are also included to give a diversity of perspectives.

⁶ For more details regarding the original technique, interested readers are encouraged to read the scenario network mapping manual found in [18].

⁷ This outline lacks implementation details because it is intended to serve for comparison.

- Introduction and review
- Grouping the event trees
- Linking the event trees
- Reviewing and digitizing the scenario map
- Workshop 4: Revealing the underlying layers
 - From event trees to scenarios
 - Finding the influences
 - Grouping the stakeholders
 - Finding the visions
 - Finding the worldviews
 - Review

The proposed modified SNM workshopping process for mapping the paths to AGI is designed to take place over as little as two days through four separate half-day workshop sessions. It roughly follows the process laid out in the SNM manual, however, the individual workshops have been modified significantly for the specific task of mapping the paths to AGI. The process is very tactile and utilizes post-its, colored ribbon, various sizes of paper, colored stickers for voting and other items that were described earlier⁸. The outcome is lattice-like map with a technology layer as well as a milestone layer. The outline below depicts the four independent workshops in this process. It includes more detail than the outline for the standard process so that it may be used for implementation⁹.

- Workshop 1: Identifying the present and future (approximately 3 hours)
 - Introduction to the workshopping process and supporting techniques
 - Mapping the core technologies that have led to the current state of AI
 - Identify the core technologies driving AI research
 - Split into groups for each of these research areas
 - List recent milestones in AI research for each technology group¹⁰
 - Vote on milestones using stickers¹¹
 - Create event trees for the most important of the milestones¹²
 - Link and combine the most important event trees
 - Results are pasted to wall
 - The facilitator guides the group in connecting the event trees¹³
- Workshop 2: Identifying paths to the future (approximately 5 hours)

⁸ See www.rossgritz.com/snm-development-workshop for examples.

⁹ A complete manual for use of the modified method requires further research and is beyond the scope of this introduction to the technique.

¹⁰ Milestones are written on a large sticky note.

¹¹ A fixed number of stickers is given to each participant to vote. Participants may use one or more sticker for each item they vote on.

¹² To create event trees, each large sticky note is placed at the center of a blank A3 sized sheet of paper. Smaller sticky notes are placed on the left and right for the antecedents and the outcomes, respectively. Different yet consistent colors are used for the left and right sticky notes.

¹³ The event trees are connected with ribbon.

- Identify the different visions for arriving at AGI
 - Split into groups for identifying different visions
 - Assign different visions to groups to explore further
- Forward-flow and backward-flow analysis
 - Split into two groups (one for forward-flow and another for backward-flow)
 - Forward-flow group identifies technologies that will likely be part of the development of AGI starting from the current state¹⁴
 - Begins from the results of Workshop 1
 - Backward-flow group identifies technologies that will likely lead to AGI
 - working backward from the different visions for AGI identified earlier
 - These results are pasted to the wall and duplicates or overlaps are condensed
 - Each technology is assigned by facilitator to a group for event tree creation
 - Each group creates event trees¹⁵ for these technologies
 - Groups can split further if needed (groups of 3-5 are ideal)
- Results are pasted to wall (from left to right) building on those from Workshop 1
- Workshop 3: Connecting the present and future (approximately 5 hours)
 - Introduction and facilitator notes from first two workshops
 - Reassess previous workshops' work
 - Technology groups split away to reassess their work
 - Modifications and updates are made if necessary
 - Split into forward-flow and backward-flow groups
 - Each group reassesses their previous work
 - Modifications and updates are made if necessary
 - The event trees are connected
 - The facilitator guides the group in connecting the event trees¹⁶
 - The most important elements are determined and the map is condensed
 - Each participant votes on the most important elements using stickers
 - Voting is done for the event trees as a whole and the subcomponents¹⁷
 - As a group, the facilitator goes through the event trees to determine what to combine and what to remove
 - Gaps and items for expansion are identified
 - As a group the facilitator helps to identify gaps between paths and the items in the current map that need breaking down further (using the holonic principle)
 - Groups split into breakout groups (size of 3-5 is ideal)
 - Gaps and items are assigned to each breakout group

¹⁴ Technologies are written on a large sticky note.

¹⁵ Event trees are created in the same way as for Workshop 1.

¹⁶ Different colored ribbon can be used for more complex mappings.

¹⁷ Different colored stickers are used for low and high priority items. A limited number of stickers is given to each participant. Stickers are to be placed directly on either the large or small sticky notes for each of the event trees. Stickers may be placed to overlap due to constraints on the size of sticky notes as long as the total number of votes is still clear. Sticky notes can also be rewritten and replaced in order to make room for stickers for voting if necessary.

- Each breakout group develops event trees for the items and gaps assigned
- Results are pasted to the wall (in-between the event trees they are intended to connect or adjacent to the items they breakdown)¹⁸
- Workshop 4: Mapping the paths and milestones (approximately 3 hours)
 - The most important links are determined and the map is finalized
 - Each participant votes on the most relevant event trees that were added
 - Participants also vote on whether certain elements need further attention
 - Entire group discusses the votes the facilitator guides discussion to condense the map
 - If necessary entire group or subgroups can address any elements that need further attention
 - When group is content then finalize map
 - Add milestones layer
 - Identify core future technology groups and split into subgroups for each
 - Each subgroup identifies potential milestones for their technology domain¹⁹
 - Milestones associated with technology paths are added
 - Participants vote on most relevant and plausible milestones
 - The facilitator guides the group in removing the unnecessary or unpopular milestones
 - Conduct concluding discussion about the process and outcome

Participants are encouraged to keep notes of their personal experience throughout the process in order to help to improve future efforts. In general, participants are encouraged to be creative and to not be conservative in suggesting technologies or milestones, or in creating the event trees. Irrelevant or unnecessary items will always be removed in the process of voting and condensing. It may be helpful for the facilitator to be familiar with brainstorming and creativity techniques in order to assist the group or breakout groups and to improve the overall outcome. It can also be beneficial to have a co-facilitator for the entire process due to the size of the ideal group. Particularly, a co-facilitator is highly recommended for Workshop 2, where groups need to split into forward-flow and backward-flow groups.

Following the workshop, it is necessary to digitize the results. The easiest way to do this is using a spreadsheet application [18]. More advanced techniques can include the use of visualization software packages. In order to create a map of the sort that is depicted in Fig. 1, this would be necessary. Such a visualization may have to be created manually due to the lack of automated software for digitizing scenario maps.

¹⁸ Connections are self-evident and the facilitator connects the event trees without the group's input.

¹⁹ Milestones and technologies are both written on large sticky notes – these sticky notes should each be of a distinctive and consistent color for the entire process.

5 Conclusion

This study presented an adapted scenario network mapping (SNM) workshopping process for mapping the paths to AGI (AGI-SNM). SNM is a comprehensive and flexible approach that comes from the family of scenario analysis techniques commonly used in technology forecasting and management. It is more rigorous and methodical than technology roadmapping which was used in an earlier coordinated effort to map the paths to AGI. Furthermore, it sufficiently addresses some of the challenges mentioned by organizers of the earlier attempt a decade ago. Specifically, it is intended to accommodate many intersecting paths and large numbers of scenarios.

Many may think the pursuit of a roadmap to AGI to be useless due to the results of previous efforts. Perhaps this is correct, but SNM does not produce a roadmap like previous efforts, rather, it produces a lattice-like map of intersecting possible paths. It does this by utilizing a powerful combination of group facilitation techniques for identifying things that may be difficult for independent researchers or researchers in standard group meetings to foresee on their own. Thus, SNM has the potential to aid all active members of the AGI research community by illuminating blind spots, hidden problems and hidden prizes that couldn't be found otherwise. It can also help in ways such as aligning the research community, providing a useful overview of the field to young researchers and refocusing research efforts on longer-term goals rather than goals for near-term gains. Simply participating in the AGI-SNM workshopping process can be a valuable experience to researchers as well²⁰. Future work should continue to refine and apply the process. We intend for this paper to foster discussion within the community about an effort to use it to conduct an updated mapping of the paths to AGI with leading experts in the field.

References

1. Adams, S., Arel, I., Bach, J., et al. "Mapping the landscape of human-level artificial general intelligence." *AI Mag.* Vol. 33, no. 1, 25-42 (2012). doi:10.1609/aimag.v33i1.2322
2. Amer, M., Daim, T., Jetter, A. "A review of scenario planning." *Futures.* **46**, 23-40 (2013). doi:10.1016/j.futures.2012.10.003
3. Bach, J. "MicroPsi 2: the next generation of the MicroPsi framework." In: *Proceedings of 5th Conf on AGI*, pp. 11-20. Springer, Berlin (2012). doi:10.1007/978-3-642-35506-6_2
4. Baum, S., Goertzel, B., Goertzel, T. "How long until human-level AI? Results from an expert assessment." *Tech Forecasting and Social Change.* Vol. 78, no. 1, 185-195 (2011). doi:10.1016/j.techfore.2010.09.006
5. Bradfield, R., Wright, G., Burt, G., et al. "The origins and evolution of scenario techniques in long range business planning." *Futures.* Vol. 37, no. 8, 795-812 (2005). doi:10.1016/j.futures.2012.10.003
6. Ecoffet, A., Huizinga, J., Lehman, J., et al. "Go-Explore: A New Approach for Hard-Exploration Problems." *arXiv preprint arXiv:1901.10995* (2019).

²⁰ A large majority of participants in the development AGI-SNM workshop felt the experience to be illuminating, enlightening and very valuable.

7. Garcia, M., Bray, O. Fundamentals of technology roadmapping. No. SAND-97-0665. Sandia National Labs, Albuquerque, NM, United States (1997). doi:10.2172/471364
8. Gaziulusoy, A., Boyle, C., McDowall, R. "System innovation for sustainability: a systemic double-flow scenario method for companies." *J. of Cleaner Production* 45 (2013): 104-116. doi:10.1016/j.jclepro.2012.05.013
9. Goertzel, B., Arel, I., Scheutz, M. "Toward a roadmap for human-level artificial general intelligence." Artificial General Intelligence Roadmap Initiative (2009).
10. Goertzel, B. *Ten Years to the Singularity If We Really, Really Try*. Humanity Press (2014).
11. Goertzel, B. *The AGI Revolution*. Humanity Press (2016).
12. Gruetzemacher, R., Paradise, D. "Alternative Techniques to Mapping Paths to HLAI." 2019. arXiv preprint arXiv:1905.00614 (2019).
13. Gruetzemacher, R. "A Holistic Framework for Forecasting Transformative AI." Forthcoming manuscript (2019).
14. Hart, D., Goertzel, B. "Opencog: A software framework for integrative artificial general intelligence." In: *AGI*, 468-472 (2008).
15. Laird, J., Wray, R., Marinier, R., Langley, P. "Claims and challenges in evaluating human-level intelligent systems." In: *Proceedings of 2nd Conf on AGI*, Atlantis Press (2009).
16. Laird, J., Wray, R. "Cognitive architecture requirements for achieving AGI." In: *Proceedings of 3rd Conf on AGI*, Atlantis Press (2010).
17. Lake, B., Ullman, T., Tenenbaum, J., Gershman, S. "Building machines that learn and think like people." *Behavioral and Brain Sciences*. 40 (2017).
18. List, D. *Scenario network mapping: the development of a methodology for social inquiry*. University of South Australia, Adelaide, SA, Australia (2005).
19. List, D. "Scenario network mapping." *Journal of Futures Studies*. Vol. 11, no. 4, 77-96 (2007). doi:10.1.1.390.6457&rep=rep1&type=pdf
20. Mikolov, T., Joulin, A., Baroni, M. "A roadmap towards machine intelligence." In: *Intl Conf on Intelligent Text Processing and Comp Linguistics*, pp. 29-61. Springer, Cham (2016). doi:10.1007/978-3-319-75477-2_2
21. Mnih, V., Kavukcuoglu, K., Silver, D., et al. "Human-level control through deep reinforcement learning." *Nature*. Vol. 518, no. 7540, 529- (2015). doi:10.1038/nature14236
22. Phaal, R., Farrukh, C., Probert, D. "Technology roadmapping—a planning framework for evolution and revolution." *Tech forecasting and social change*. Vol. 71, no. 1-2, 5-26 (2004). doi:10.1016/S0040-1625(03)00072-6
23. Radford, A., Wu, J., Child, R., et al. "Language Models are Unsupervised Multitask Learners." OpenAI Blog, (2019).
24. Roper, A., Cunningham, S., Porter, A. et al. *Forecasting and management of technology*. John Wiley & Sons (2011).
25. Rosa, M., Feyereisl, J., Collective, T.G. "A framework for searching for general artificial intelligence." arXiv preprint arXiv:1611.00685 (2016).
26. Vinyals, O., Babuschkin, I., Chung, J., et al. "AlphaStar: Mastering the Real-Time Strategy Game StarCraft II." DeepMind Blog, (2019).
27. Van der Duin, P.A.. *Qualitative futures research for innovation*. Eburon Uitgeverij BV (2006).
28. Wang, P., *Rigid flexibility: The logic of intelligence*. Springer Science & Business Media (2006).