Simulation as a Service for Decision Support & Training

Dr. D M Culley, Mr. P Cavanagh, Mr. R Warner
Improbable, London E1 6EG
UNITED KINGDOM

daveculley@improbable.io

ABSTRACT

The insights gathered in synthetic worlds will reveal the most effective way to achieve your mission objectives. From strategic campaigns to tactical actions, governments and international bodies will have moved to a state of near perpetual virtual rehearsal and experimentation. Whole force preparedness will rapidly evolve and improve inside highly realistic synthetic environments. This future capability relies on building a synthetic environment where governments, industry and academia can collaborate in co-creating the models and simulations needed to build realistic scenarios for an increasingly complex and demanding battlespace. The platform on which this work will be built is the key enabler that will unlock the potential of future decision makers and warfighters.

This paper discusses the complexities around constructing this base 'Simulation as a Service', describing both current and future approaches to a suite of products that fulfil the requirement to build integrated multi-domain, highly realistic training and planning scenarios across multiple levels of war. Many technical, environmental and cultural challenges exist in the development, adoption and use of such a platform, but the potential benefits this brings to governments' shared understanding, consensus forming, decision making and training make facing these challenges a critical priority for the future.

1.0 INTRODUCTION

Information and communication technologies allow wide and instantaneous interaction, facilitating new opportunities for greater interconnectedness and collaboration across society but increasing its sensitivity to destabilising shocks. Increasingly accessible computing and the big data revolution offer opportunities to transform governments' approach to national security and resilience both at home and abroad, but our ability to fully utilise these resources to promote stability is lagging behind our needs. We find ourselves in an analytical arms race; with an imperative to understand better and faster and to act when needed in a targeted, decisive and effective way.

The major challenges are threefold: firstly, greater interconnection has led to ever deeper interdependence between economies, institutions and all manner of critical socio-technical systems. Traditional disciplinary thinking is often too limited to support adequate understanding of this densely woven fabric and to support effective decision making. Secondly, in order to react quickly and effectively, joined up processes are required that span tools and job functions. The insight produced by an analyst using a Machine Learning application, for example, needs to be compiled with other sources, reviewed, planned against and acted upon. This requires a connected and interdependent suite of technological capabilities such as cloud compute, storage, modelling and simulation. Finally, the barrier to entry to fully utilise the apparatus of data driven decision making is often an advanced degree in a highly technical field. As quickly as tools become easier to use, data volumes grow, analytical methods evolve and complexity increases. To the greatest extent possible tools need to be accessible, straight forward to use and facilitate easy collaboration.

This paper presents platform enabled Synthetic Environments, which are an approach to Simulation as a Service that seeks to address these challenges. Synthetic Environments bring together the broadest range of cutting edge synthetics and analytical technologies in a way that is accessible to users, dissolves traditionally stovepiped disciplines and unifies the decision making process across organisations. The following section introduces Synthetic Environments, while later sections consider how they help bridge cross-disciplinary and cross-functional boundaries and promote accessibility to a broad range of users. A companion paper at this conference, "Single Synthetic Environment for Operational Decision Support" [STO-MP-MSG-177-9] discusses the practical development of a Synthetic Environment for decision support.

2.0 SYNTHETIC ENVIRONMENTS

Synthetic Environments amplify humans' capabilities in policy setting, decision-making, and training, by using virtual worlds which represent the scale, complexity, and integrated nature of the modern world. They are ecosystems of interconnected data, artificial intelligence technologies, machine learning systems and constructive models such as digital twins and simulators. These ecosystems are hosted on a platform and delivered as a service to users via user interfaces that are carefully tailored to support their needs and empower them in their work.

The platform simplifies management and maintenance of the technologies across the ecosystem, abstracting the complications of system integration and deployment. For the user, it simplifies access, delivering a broad range of capabilities to users in a single place, as a service and ensures a coherent and consistent user experience across previously siloed products. The increased interconnection between previously distinct products helps join up the people who use them, encouraging greater collaboration across and between organisations.

Synthetic Environments are engineered to help users explore real-world scenarios by letting them visualise, interact with and experiment on rich simulations. Data and analytics let users see the whole picture and provide a deeper understanding of the situation as it stands. Closed simulation allows us to play this forward and enables potentially thousands of permutations of a scenario to be explored quickly to evaluate risks and design optimal response strategies. Meanwhile with human-in-the-loop simulation coupled with detailed 3D rendering, users can even interact directly with the environment, enabling immersive training and rehearsal of the proposed response.

Since hundreds or even thousands of users can work in and experience this shared world at the same time, Synthetic Environments also promote shared understanding, improved consensus-forming, optimised decision-making and interoperation within and between organisations. The goal is to use technology to support, amplify and apply uniquely human attributes: imagination, intuition and experience.

2.1 Platform for Synthetic Environments

The technology which makes these wide ranging benefits possible is Improbable's Synthetic Environment Platform. It incorporates and extends SpatialOS, a framework for building massive multiplayer online games, that has been developed to enable Simulation as a Service for decision support and training. It sets open standards for integration of technologies and provides various deployment options and supporting functionality such as data storage. The core benefits of the platform is that it takes care of the communication and synchronisation between models and other technologies, as well as facilitating the resulting simulation scale and complexity by efficiently handling the distributed computing and computational load balancing.

The platform enables simulation to be delivered as a service by addressing many of the historical limitations of large-scale simulation approaches. These were based on the federation of individual simulations, often running on standard PCs, through interoperability standards such as the Distributed Interactive Simulation (DIS) protocols and / or the High Level Architecture (HLA). These 'traditional' simulations come with a logistical cost in the form of coordinating the availability of resources (operators and machines) and the connectivity of disparate networks to allow interoperability. The Synthetic Environment Platform removes these problems, once integrated, technologies can be made straightforwardly available as a service at the point of need.

Improbable's Synthetic Environment Platform is built using open standards and can integrate a range of models, simulations, and AI/ML technologies from different sources across government, industry and academia. The platform approach enables and promotes modularity in technological assets, modularity in turn encourages division of labour between more, specialised companies providing clearly bounded, focussed capabilities. Each capability can be developed at lower risk and leverage the supporting technologies provided by the platform rather than having to build their own. Integration with the platform also means the capability becomes part of the ecosystem, can interact with other capabilities and is part of a whole that is greater than itself. This open-platform approach to unifying and managing these assets inoculates against vendor lock-in and also ensures that the Environments can draw on the most reliable, relevant and up-to-date content to grow and evolve continuously.

Development of this ecosystem of assets is, by its very nature, a multi-disciplinary and multi-stakeholder enterprise. It means bringing together the best capabilities and expertise from across industry, academia and governments. Over time, this ecosystem will become richer, and the ability to substitute and upgrade parts of it would ensure it remains a cutting-edge capability.

Open standards are therefore crucial to encouraging a thriving ecosystem, but do not, however, suggest 'open access'. The platform and deployment model are designed with robust security at their core, and configured and deployed to meet the security and access requirements pertinent to the work being undertaken, whether that involves a secure cloud, on-premises or bespoke solutions.

3.0 FACILITATING CROSS FUNCTIONAL PROCESSES

Synthetic Environments allow consolidation of the whole 'value chain' from data collection to decision implementation into a single place to allow a stronger relationship between suppliers and consumers of information, analysis and plans. This section aims to demonstrate how this works and the benefits it might bring.

Machine Learning (ML) has, in recent years, become increasingly popular, both benefiting from and fuelling an explosion in data-collection capabilities across a wide range of sectors. In this family of analytical techniques, highly flexible models are fitted to vast volumes of data, leveraging advances in computation and statistical algorithms. They are excellent at finding patterns in data, flagging trends and anomalies for further investigation, and helping us better understand the world as it really is.

Powerful though they may be, however, ML technologies can only work on the data they are given. In isolation and without broader situational context, it can sometimes be hard to trust ML and to identify blind spots in the analysis it provides. Interconnecting ML technologies into a broader Synthetic Environment helps mitigate this weakness by maximising the user's contextual understanding, therefore increasing the scope and quality of the insights.

This was demonstrated in a programme in which Improbable developed a Synthetic Environment application to help military operational-level planning. In this project, a natural language processing (NLP) technology (a type of ML) developed by a specialist vendor was integrated onto the platform. It continuously combs through vast amounts of text data such as news articles and intelligence reports and could extract salient details such as 'what', 'where' and 'who'. Since the tool is integrated with the Synthetic Environment, this information is then presented to the user along with the broader context of many other data sources. In this way, something as abstract as a population's political sentiment might be overlaid as a layer on the geospatial dashboard alongside other data layers of interest. At a glance, an analyst or decision maker can see as broad and comprehensive picture as required.

The ecosystem for this application was not, however, limited to analytical technologies but extended to synthetics. Constructive simulation modules included population sentiment, logistics, civilian infrastructure and computer generated forces (CGF) sourced from a specialist provider. Based on the information developed by analysts through use of the analytical tools, other user groups, for example planners, can evaluate and improve potential interventions by 'playing them out' using constructive simulation. The candidate plan can then be passed to the relevant operatives who can rehearse and perfect its execution in a 3D rendered, fully immersive, interactive simulation of the scenario.

The vision is that future projects will demonstrate integration of live force tracking and Command and Control (C2) systems so that the whole lifecycle of a plan is facilitated by a Synthetic Environment. This would ensure that all analysts, planners, operatives and commanders involved in the plan from inception to execution have access to all the relevant information and the shared ecosystem of the best possible tools.

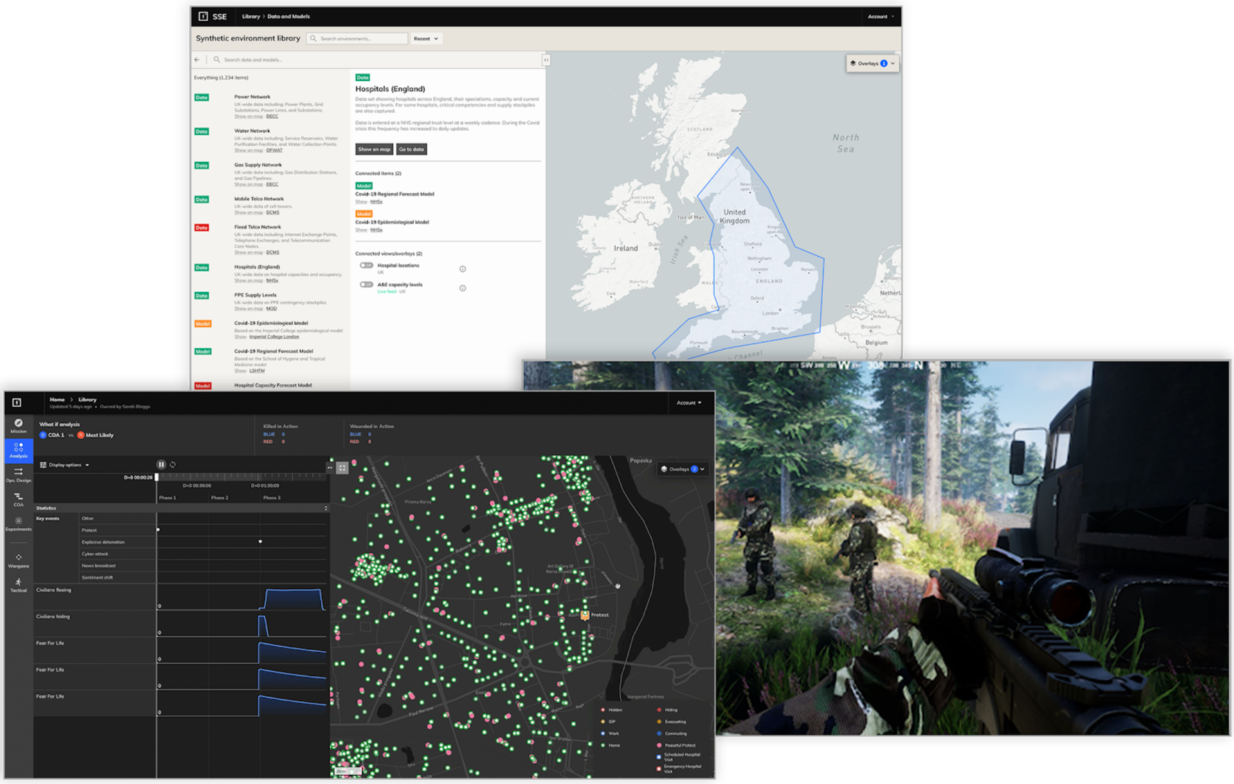

Figure 3-1: Synthetic Environments are presented to users according to their needs. An analyst may see data and dashboards, but on-the-ground personnel may experience their environment as a richly detailed virtual world in which they can train and rehearse safely, frequently, collectively and cost-effectively.

4.0 BREAKING DOWN DISCIPLINARY BARRIERS

To cope with a world comprising increasingly interconnected systems, our synthetics and analytical technologies must be equally joined up. In performing a strategic assessment it is natural to decompose the task into political, military, economic, social, information, infrastructure and potentially a host of other disciplines. This allows the work to be divided among highly focussed specialists but creates the challenge of recomposing the work so as to understand the full picture. By its nature, governmental decision making - especially in Defence and National Security - requires this holistic understanding.

In the training space there are also great benefits to tackling siloed synthetics and facilitating straightforward and performant interoperation between models and systems and across domains. Military plans are multi-faceted, multi-domain, complex artefacts and to train with this level of complexity in the real world is prohibitively complicated and expensive.

In Synthetic Environments, the Common Data Model (CDM), acts as an agreed schema for data components used by multiple models, facilitating interoperability in a manner similar to the High Level Architecture (HLA) Federation Object Models (FOM) used in traditional modelling and simulation applications. Technological assets such as models will interact via this schema but may have a separate internal schema for properties that do not need to be shared with other models or components. The CDM includes aspects such as identity, position and health status representation, ensuring a common understanding of what these attributes mean across disparate parts of the technology ecosystem. The CDM makes use of an existing enumeration standard from the Simulation Interoperability Standards Organisation (SISO) [SISO-REF-010-2019], to identify entities.

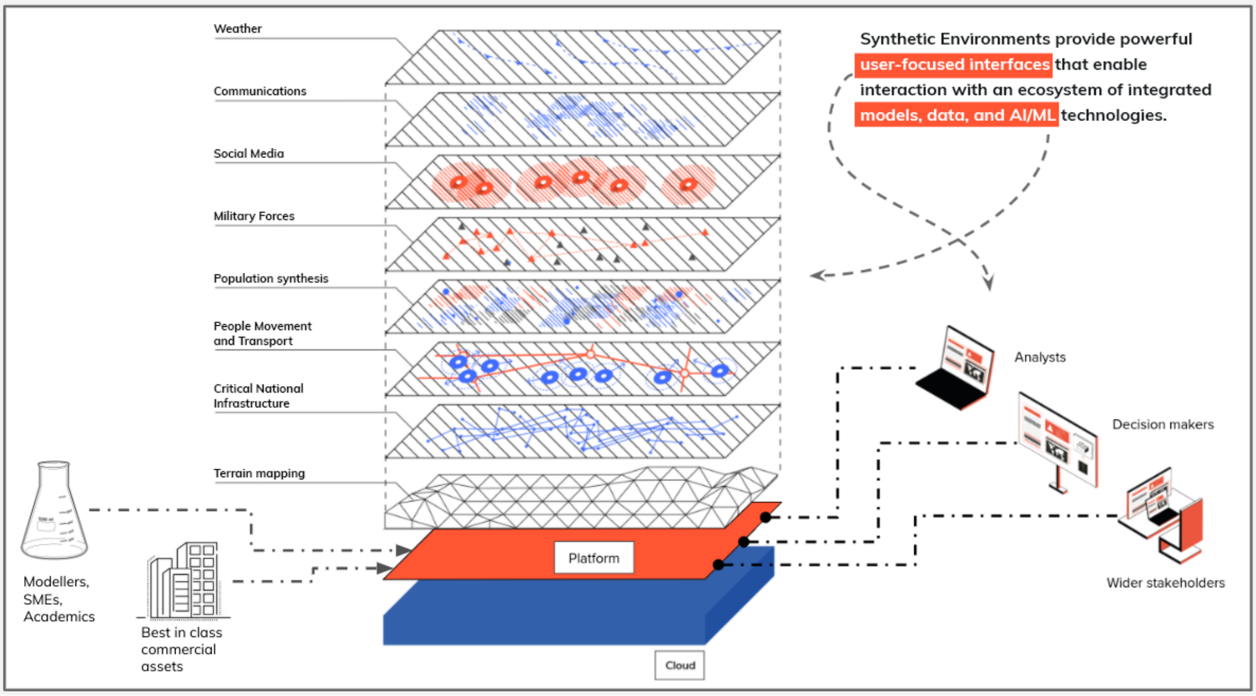

Figure 4-1: Synthetic Environments comprise an ecosystem of technological assets that span disciplines and let users explore decision options safely and cost-effectively in a virtual world before taking action in the real one.

4.1 Cascading consequences

As an example of the practical importance of capturing cross-disciplinary interactions, consider a country's power network and the myriad downstream systems that require a power supply. The cascading consequences of an interruption in supply to these downstream systems could be catastrophic but requires a cross-disciplinary approach to effectively reason about it. An interruption in national gas supply may cause a power station to go offline, causing power cuts in certain regions, this will affect traffic as traffic lights go off line, health care as hospitals are affected, telecommunications as masts and exchanges lose power. Over longer time periods there may be civil unrest, population movement, shifts in political sentiment or epidemics.

Many of the downstream factors being explored in this example are based on social systems that are hard to model, and the models that do exist are often perceived to be unreliable. As a result, it can be difficult to justify their utility and applicability in informing decision makers. When Synthetic Environments have capabilities and models in this area they must be presented to users carefully and with measures in place to allow for validation and verification. However, there is no doubt that understanding potential impacts on and responses of social systems is essential to the rich holistic picture that Synthetic Environments seek to reveal. There can be very high value in capturing these systems and enabling experimentation and exploration with their dynamics via simulation, even with nascent models and innovative applications. Ultimately the decision maker is in the loop but with access to Synthetic Environments, they are also empowered with additional data and analysis that goes beyond traditional, contained disciplines.

4.2 Rich training experiences

While discussion has so far been focussed on decision making applications, SpatialOS originated in the gaming industry where human-in-the-loop (or 'interactive') simulation execution creates detailed and responsive virtual worlds that players can experience. Synthetic Environments can support both decision making and training as well as link these activities together – using real data, situations and plans to generate realistic training scenarios. This interconnection enables a 'single' Synthetic Environment supporting a range of applications operating across land, air, sea, space, cyber and information domains. Social systems models described in the previous section can support synthetic populations of civilians who have realistic patterns of life, communicate sentiment on social media, consume civilian infrastructure services.

The modern warfighter should feel better prepared to face real situations and trainers should feel increased confidence in the skills their students have acquired within these dynamic Synthetic Environments. The lines will become blurred with integrated live, virtual and constructive (LVC) training - actions that occur in one world have consequences in another. This capability is critical for achieving enhanced training outcomes at reduced cost. Virtual training will be immersive and have the highest standards for realism, drawing technology from the commercial gaming sector where the amount of investment and market demand drive innovation at an ever increasing pace.

5.0 EMPOWERING THE USER

Previous sections have shown how Synthetic Environments enable a diverse range of capabilities, as a result, the pool of potential users may span entire organisations. Each job function and role will have different requirements and expectations of the system and want to engage with it differently. A trainee just needs to be able to connect to the synthetic exercise easily, while a planner needs a much more full application to explore the common operating picture, author plans, run and evaluate simulations. The modularity of Synthetic Environments extends to the user interfaces (UIs), each of which can be thought of as a 'window' into the underlying virtual world. These user interfaces must empower the user to understand and navigate the complexity of the technology so that, rather than distracting or misleading it empowers and augments their work. Designing such tailored interfaces requires a deep mutual understanding between technology providers and prospective users, and new ways of working on technology projects in which the customer is an integrated member of the delivery team and giving continuous input and feedback.

5.1 Decision support not decision making

Critical to the successful use of analytics technologies in decision support is communication to the user of the reliability of the analytical techniques being used and the assumptions that underpin them. When these become opaque to the user the value of the analysis is greatly diminished; either analysts become over dependent on the results, switching off critical faculties, or lose trust and disregard the analysis altogether.

Tailored user interfaces help to mitigate this by placing modelling and computational analysis in the context of the end to end decision process. The user's journey through the application to perform a planning task can be designed and structured so as to encourage critical thought. For example, each time the user is presented with analysis, simulation outcomes or data, the UI can ensure they are encouraged to critically evaluate it, interpret it and to record that analysis. This means that the user is not exposed to 'raw science' without context or caveat. Furthermore, existing planning processes are not replaced by use of the Synthetic Environment, they are simply supported by it. These processes, already established and robust, are reinforced and ensure the resulting plans and decisions will be based on a balance of objective and subjective analysis by the user, supported by the technology. As a concrete example, simulation output statistics of a proposed tactical plan are provided to the user as evidence based upon which they make their human-expert evaluation of the proposed mission's measures of effectiveness.

This pragmatic approach balances what Davis & Blumenthal [1991] describe as the 'art' and 'science' aspects of governmental and military decision making. Scientific models and automation technologies - such as data processing - are pushed as far as they can to provide useful evidence to underpin decisions. These tools are accessed from the application interface that has been specifically designed for the intended user; to provide targeted contextualisation and appropriate guidance and warnings, and to solicit their more creative and nuanced expert opinion as a lens through which the computational analysis is interpreted.

Related to this is the need for explainability. Military stakeholders are alert to and distrustful of so-called 'black box' analysis tools - those that output an answer to a problem but cannot provide the user with a user-centric and understandable explanation of how it was reached. Depending on the needs of the user, interfaces can provide detailed documentation for models, can replay simulations to see how events unfolded and can stage experiments that investigate natural (or model-caused) variation in a scenario (via Monte Carlo runs) or sensitivity of particular scenario parameters (via parameter sweeps). The analyst (or team of analysts) contributing to a plan can explore and understand how outcomes are reached and can therefore determine how simulation evidence feeds into their plan and judgment. In some cases, the simulation may simply stimulate consideration of a factor that the planner hadn't considered: This could be as simple as accounting for traffic on the road when planning route timings, which is potentially easy to forget with a paper-based map, but jumps out of a dynamic simulation.

The ultimate plan or decision remains, therefore, a human led effort that is supported by data and simulation that is explained and explorable to the analyst. They can use this computational aid to the degree to which they feel comfortable with it and the UI allows them to record and present their thinking in familiar, doctrinally established formats. Disparities between the human-designed plan and the supporting data and simulations may be something that senior decision makers wish to explore and this should be actively encouraged: it stimulates critical analysis of the mechanics of a scenario and the plan to tackle it and should foster a deeper situational understanding.

5.2 Solution design process

In order to achieve the tailored user experience, technology providers, customers and users must work closely together to develop mutual understanding of what is needed and what can be achieved. The solution design process (SDP) aims to facilitate multidisciplinary collaboration between teams and people with different skill sets in order to design better technological solutions. It helps us to iterate and refine our solution quickly with confidence, while better fostering shared understanding between technologists and users. The SDP follows the 'double diamond' design model which encourages users in "divergent" thinking to generate many potential ideas, before "convergent" refinement to narrow down options to the best solution.

The process focuses on a particular high-level requirement or capability area and works through discovery, definition, refinement and specification phases to achieve a detailed description of the feature and design to be implemented. In the discovery phase, market, user, and requirements research is undertaken to understand the problem and lay the foundation for solution development and implementation. In the define phase, this research is consolidated into a more focussed understanding of the opportunities to deliver user value. Continual user feedback and expert input ensures that user concerns and market context are properly considered as a solution begins to form. The solution is then iteratively refined based on usability and feasibility considerations as the design becomes more detailed and better understood. By the end of the process, sufficient detail is achieved to understand the scope of work, begin implementation of the feature and estimate when stakeholders can trial and evaluate releases.

The ability to be able to iterate and refine solutions quickly, collaboratively and with confidence is invaluable. It is more efficient and less risky to deliver small pieces of capability and value continuously and incrementally rather than delivering a solution only when it is 'perfect' and 'ready'. For customer organisations this close collaboration may be a new way of working and can require a proactive culture shift to ensure that all relevant stakeholders are involved, discussions are open and decisions and actions are aligned.

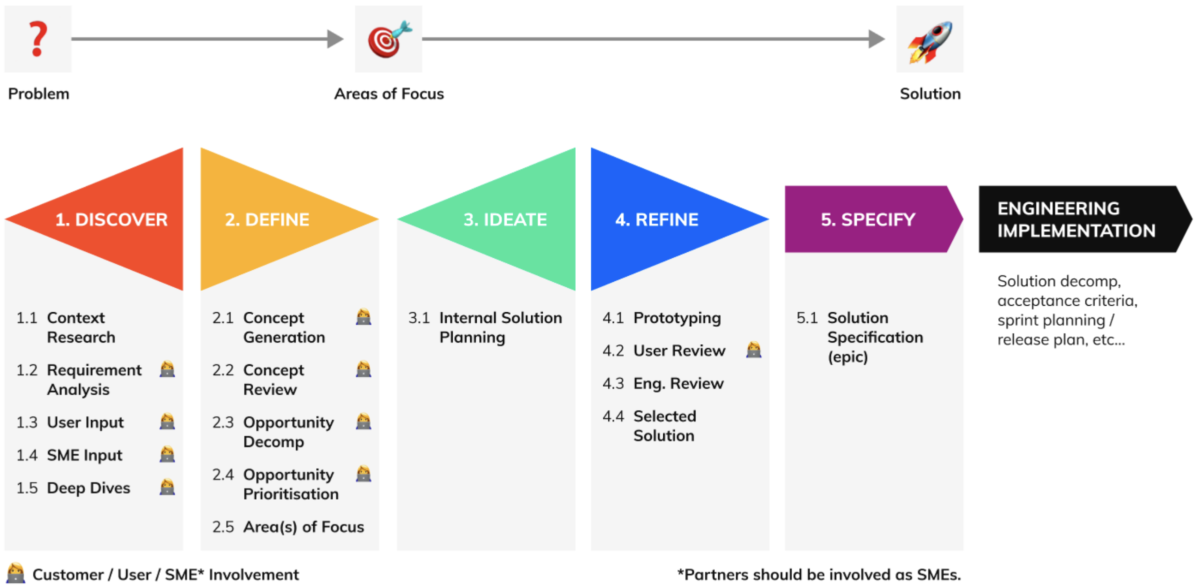

Figure 5-1: The solution design process alternates divergent and convergent thinking to develop mutual understanding between technologists and users and to specify high value solutions from loosely defined requirements.

5.3 Many users collaborating in a shared Synthetic Environment

Step changes in collaboration and communication are some of the greatest benefits of cloud-based technology. The Synthetic Environment Platform approach realises these benefits by providing a single, highly flexible, platform for all synthetics use-cases, decreasing both development & maintenance cost and time to operating capability. Depending on security constraints and the deployment environment, a cloud deployment means that participants can connect from anywhere, potentially using anything; from AR or VR in planning use cases, to human-in-the-loop simulators in interactive training / plan-rehearsal simulations.

A singular cloud-deployed Synthetic Environment can act as a universal operations room for a host of remotely located users. The potential benefits of this have already been discussed in previous sections, but for forward deployed forces on the move, these cannot be understated. In evolving crises new information is continuously coming in, and disseminating that to where it needs to be is a singular challenge that prevents a common understanding of the situation amongst those who are trying to address it.

Since Synthetic Environments are a new capability, work is ongoing with users to understand better how they want to share and collaborate using them. There are also, clearly, significant security and compliance factors that need to be considered. Other industries and their processes may provide helpful examples of how use of the system could be developed in future. For example in engineering design groups - such as Formula One racing teams - large teams of specialists and highly technical decision makers at different levels of detail and from different domains collaborate to design high performance, highly complicated artefacts. Military plans can be thought of as similarly complicated artefacts requiring similarly collaborative cross-domain and varying detail design.

In engineering, ‘concurrent design’ ensures the high-level design feeds down and across domains, and low-level design proceeds concurrently, feeding up when necessary. The alternative to this approach is called ‘over the fence’ and involves teams working in isolation with narrow concerns and ‘throwing’ their work to the next team. Because planning assets are digital and shared across users, a Synthetic Environment offers the opportunity for military and governmental planning to become more concurrent, by joining strategic, operational and tactical planning processes.

5.4 Digitised planning and training records

The Synthetic Environment can help organisations identify and respond to situations in a joined up way. It can offer indicators and warnings of potential problems through a single surface that visually integrates multiple live data feeds and serves as a common view of the world as it evolves in real time. Information and notes can be drafted directly on top of the data layer, facilitating collaborative and faster strategic assessment efforts. Response options can then be quickly explored and understood via in-depth analysis of the environment using exploratory modelling and simulation. Based on this increased understanding, detailed plans can formulated and refined collaboratively, interactively and iteratively; aided by simulation and analytics to quickly identify key outcomes and risks.

One of the major benefits of computer-supported planning and training is that outputs and artefacts such as plans, policy options and records of decisions are stored digitally. Compared to analogue archives, digital ones are easier to search and can be accessible from anywhere, at any time. Furthermore, these artefacts can be dynamic, updating automatically as variables change and new data emerges.

For example, several contingency plans may have been developed, agreed and stored on the system that pertain to potential disaster relief scenarios in a particular location. As time passes, if the system is connected to the relevant data streams, the information in the plans could be continuously updated. Even relatively simple information like the organisational structure of the local government and contact details for key figures could be kept up to date automatically, which would reduce friction in a crisis scenario.

Perhaps more powerfully, if the relevant data streams are connected to the Synthetic Environment Platform, then ML technologies could be set to raise a flag when measurable underlying assumptions of a stored, digitised plan no longer hold true. The disaster relief plans from the above example may be contingent upon logistical assumptions that certain resources are available in certain places in certain quantities. A separate crisis – for example Covid-19 – may have seen this material moved or depleted, in which case alert systems would notify the responsible owners of the contingency plans that a vulnerability has opened up against this set of pre-planned disaster-relief scenarios.

Synthetic training – as an augmentation of live training – generates a rich, digitised record that facilitates similar benefits, enabling a much more complete understanding of the training level, skill sets and proficiency of operational personnel. As the Synthetic Environment is used more and more, this data will accumulate. Thanks to ML, it's possible to extract the maximum value from it: training data can be fed back into the planning process, leading to better-trained personnel who can respond faster based on improved plans with improved outcomes.

The common environment for plans and training can also be leveraged proactively. High-impact and/or high-likelihood scenarios (whether civil or military) can be identified and earmarked, and synthetic training exercises developed and rehearsed with the relevant services. This would provide a step-change towards perpetual readiness.

6.0 CONCLUSION - TOWARDS PERPETUAL READINESS

Governments and defence departments are increasingly challenged by fast-moving, evolving threats, and require transformational changes to keep pace. Stovepiped synthetics and paper-based analysis must be replaced with collaborative tools that use authoritative content, and support continuous improvement.

Scalable, adaptable and ever-evolving Synthetic Environments are designed to bring together planning, training and rehearsal across and between organisations, be they military forces or government departments. This will enable faster problem identification and evaluation, improve consensus forming around the need for action and enable rigorous testing of proposed actions with minimal cost. It wall also facilitate effective communication of warnings, assessments and response options to decision makers. The goal is for an improved shared understanding of the key objectives and risks of a briefing, to accelerate option assessment, speed up decision making and shorten the loop from problem identification to aligned action.

The benefits of this pan-organisational ethos is that different roles and functions can work collaboratively and concurrently on a project across government. This enables a more agile and integrated response to both seize opportunities, and respond to challenges. Use of such technology ensures that knowledge and provenance of decisions is recorded, and that institutional understanding is both retained and made more easily accessible. This, coupled with the opportunities for cheaper and more accessible synthetic training, facilitates a state of perpetual readiness to achieve strategic objectives.

Finally, the platform-enabled ecosystem approach reduces the barrier to developing and trialling new technologies, facilitating the development of a truly evolving capability that draws on the very best governmental, academic and industrial institutions. The ability for the ecosystem to evolve and grow ensures it can be sustained as a cutting-edge capability which supports and maintains security and prosperity in the information age.

REFERENCES

Davis, P.K. and Blumenthal, D., 1991. The base of sand problem: A white paper on the state of military combat modeling (No. RAND/N-3148-OSD/DARPA). DEFENSE ADVANCED RESEARCH PROJECTS AGENCY ARLINGTON VA.

SISO, G., 2020. Reference for Enumerations for Simulation Interoperability. SISO-REF-010-2020, Simulation Interoperability Standards Organization.

Single Synthetic Environment for Operational Decision Support, Welton, D., Judd, G. and Cavanagh P.J. [STO-MP-MSG-177-9]