



AFRL-RI-RS-TR-2024-007

AI-HUMAN COLLABORATIVE CYBER-PHYSICAL ENGINEERING AND LEARNING (ACCEL) ENVIRONMENT

SRI INTERNATIONAL

JANUARY 2024

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the AFRL Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2024-007 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE CHIEF ENGINEER:

/ S /

STEVEN L. DRAGER
Work Unit Manager

/ S /

NICK P. KOWALCHUK
RIT Chief Engineer
Computing & Communications Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE JANUARY 2024	2. REPORT TYPE FINAL TECHNICAL REPORT	3. DATES COVERED START DATE SEPTEMBER 2020 END DATE MARCH 2023	
4. TITLE AND SUBTITLE AI-HUMAN COLLABORATIVE CYBER-PHYSICAL ENGINEERING AND LEARNING (ACCEL) ENVIRONMENT			
5a. CONTRACT NUMBER FA8750-20-C-0155	5b. GRANT NUMBER N/A	5c. PROGRAM ELEMENT NUMBER 62383E	
5d. PROJECT NUMBER	5e. TASK NUMBER	5f. WORK UNIT NUMBER R31C	
6. AUTHOR(S) Ben Southall, Bogdan Matei			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SRI International 333 Ravenswood Ave. Menlo Park CA 94025-3452			8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/RITA 525 Brooks Road Rome NY 13441-4505		10. SPONSOR/MONITOR'S ACRONYM(S) DARPA 675 North Randolph St Arlington VA 22203 AFRL/RI & DARPA	11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-RI-RS-TR-2024-007
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited. PA#: AFRL-2023-6385, AFRL-2023-1186 Date Cleared: 16 JAN 2024			
13. SUPPLEMENTARY NOTES			
14. ABSTRACT Current state-of-the-art design methodologies for cyber-physical systems (CPS) have several limitations, especially for handling complex systems consisting of many sub-modules from different domains. Designer teams for CPS typically consist of subject matter experts (SMEs) with narrow expertise in each other domain (e.g., computer-aided design (CAD), electrical systems, and control systems). A fundamental limitation of past approaches is that the knowledge is siloed and is hard to transfer to another design team such that they benefit from past collective expertise. Central to the approach taken on this effort is the focus on humans and enabling them to think about high-level design ideas, thus freeing them from highly intensive tasks requiring specialized expertise that can be outsourced to an artificial intelligence (AI) Co-Designer.			
15. SUBJECT TERMS Human-AI Collaborative Cyber-Physical Engineering (ACCEL), Symbiotic Design of Cyber-Physical Systems (SDCPS), corpus browsing, neural network, cyber-physical systems			
16. SECURITY CLASSIFICATION OF:		17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	SAR 65
19a. NAME OF RESPONSIBLE PERSON STEVEN L. DRAGER			19b. PHONE NUMBER (Include area code) N/A

Table of Contents

List of Figures	ii
1 Summary.....	1
2 Introduction.....	2
3 Methods, Assumptions, and Procedures	4
3.1 Requirements for the Symbiotic Design User Interface	5
3.2 Design problems and sandboxes.....	6
3.2.1 Southwest Research Institute Unmanned Air Vehicle.....	6
3.2.2 Systems & Technology Research LLC UUV.....	6
3.2.3 Proxy 2D Car problem.....	7
3.3 Support for single users and design teams.....	8
3.4 System architecture	8
3.5 Workflows and interfaces	9
3.5.1 Design initiation	12
3.5.2 Design analysis and exploration interfaces	13
3.5.3 Corpus browsing and filtering.....	13
3.5.4 Design Exploration	16
3.5.5 Infinite shelf	20
3.6 Interfaces for specific design problems	22
3.6.1 UAV/M – flight dynamics	22
3.6.2 UUV design composition.....	30
3.6.3 Unmanned Underwater Vehicle mission review.....	42
4 Results and Discussion.....	47
4.1 Proxy car design problem: Human-AI experiments with reinforcement learning	47
4.1.1 Human-AI interaction system	47
4.1.2 User study design.....	48
4.1.3 Results and Discussion	49

4.2 Experiments comparing part assembly using traditional 2D interfaces vs the 3D model assembler.	50
5 Conclusion	52
Appendix A. Physics-driven Numerical Simulation using Neural Networks for Fast Design Space Exploration	53
Appendix B. List of publications (in preparation or submitted)	57
List of Symbols, Abbreviations, and Acronyms.....	58

List of Figures

Figure 1 ACCEL vision addressing TA2 three main UIs for SDCPS.....	3
Figure 2 ACCEL, as proposed, detailed architecture.....	3
Figure 3 Three stages of the ACCEL symbiotic interface: Design Space Creation, Design Composition, and Design Space Exploration.	5
Figure 4 Symbiotic UI technical requirements.....	6
Figure 5 Proxy 2D car problem enabled quick experimentation with different concepts, such as parameter sweeping, as well as ideation and development ideas for the 2D sketch interface.	8
Figure 6 ACCEL high-level system architecture. Blue elements comprise the data storage backend. Orange elements are various UIs and co-ordination servers. Green elements indicate Representational State Transfer (REST) Application Programming Interface (API) connections to external services from TA1 and TA3.	9
Figure 7 High-level breakdown of design workflows.....	9
Figure 8 Expanded workflow for ACCEL, detailing some of the stages of the basic blocks.	11
Figure 9 Setting up the workflow in VR. The main user, represented by gloves in the center, is picking up tools to be employed during design.....	11
Figure 10 Design composition in VR. Multiple users can participate. The inset in the lower right shows a 3 rd person view of the design process from the room.	12
Figure 11 Design composition in a web-browser using a traditional input modality with keyboard and mouse.	13
Figure 12 VR browsing for candidate design.	14
Figure 13 VR search/browse for suitable parts.	14
Figure 14 Organizing corpus for easy browsing and search and viewing in VR.....	15
Figure 15 Web based design comparison allows the designer to assess each shape and parameters of several designs.....	15
Figure 16 Design browsing in a web interface, comparison of multiple designs in terms of the flight path and accomplishing mission.	16

Figure 17 Design filtering in the browser UI. The histograms show good/bad designs. With selection on the histogram, a large number of bad designs are able to be removed.	17
Figure 18 Design comparison in ACCEL browser UI.	17
Figure 19 The filtering UI.	18
Figure 20 Histogram selectors allow focus on particular designs.	19
Figure 21 Tabular views allow comparison.	19
Figure 22 Side-by-side design comparison.	20
Figure 23 The main panel of the infinite shelf, with data loading function button on left, sliding design browser in the middle, and parameter visualizer on the right.	21
Figure 24 VR system browsing the UAV dataset with the shelf display.	21
Figure 25 VR system browsing the UAV dataset with the sphere display.	22
Figure 26 The experiment selector.	23
Figure 27 The Flight Dynamics Model overview tool.	25
Figure 28 Zoomed-in view of the graphical overview and its controls.	26
Figure 29 Flight Dynamic Model detail views. Top: flight path review. Middle: 3D view. Bottom: Flight Dynamic Model plots.	28
Figure 30 The design evolution UI, showing two generations of designs.	29
Figure 31 UUV design tool overview.	31
Figure 32 The history section of the UI.	32
Figure 33 Fairing design work area.	34
Figure 34 Payload work area.	35
Figure 35 Propulsion work area.	36
Figure 36 Energy and battery work area.	37
Figure 37 Pressure vessel work area.	38
Figure 38 Pitch control work area.	39
Figure 39 Foam work area.	40
Figure 40 The packing and fairing modification work area.	41
Figure 41 Performance review user interface view.	43
Figure 42 Single run review user interface view.	44
Figure 43 Zoomed view of the unmanned underwater vehicle running aground.	45
Figure 44 Control editor user interface view.	46
Figure 45 Comparison of an unsuccessful control strategy (left) versus a successful control strategy (right).	47
Figure 46 Proxy 2D car climbs across obstacles on the terrain.	48
Figure 47 User interface for the proxy 2D car design.	48
Figure 48 The total number of changes on each design variables.	49
Figure 49 Interactions for users without and users with access to AI Designers.	50
Figure 50 Assembly process using traditional 2D interface with the help of Gizmos.	51
Figure 51 Assembly process using 3D VR interface assisted with various manipulation tools. ..	51

1 Summary

Current state-of-the-art design methodologies for cyber-physical systems (CPS) have several limitations, especially for handling complex systems consisting of many sub-modules from different domains. Designer teams for CPS typically consist of subject matter experts (SMEs) with narrow expertise in each other domain (e.g., Computer-Aided Design (CAD), electrical systems, and control systems). A fundamental limitation of past approaches is that the knowledge is siloed and is hard to transfer to another design team such that they benefit from past collective expertise. There is a large body of designs corpora available for many domains, however any designer or SME will have only a limited capability to access them due to their number and the complexity in searching/browsing, especially for components and domains that don't have representations in a parameterized form that is easily comprehended by humans. For example, it is relatively simple to search for designs of electric engines that have a certain power and size. However, it is much harder to browse and select complex 3-dimensional (3D) shapes of objects with many components, and further incorporate them in the context for which they were designed.

Central to our Human-AI Collaborative Cyber-Physical Engineering and Learning (ACCEL) approach is a focus on humans, enabling them to think about high-level design ideas and freeing them from highly intensive tasks that require specialized expertise that can be outsourced to an artificial intelligence (AI) Co-Designer. Under the Defense Advanced Research Projects Agency (DARPA) Symbiotic Design of Cyber-Physical Systems (SDCPS) program, AI Co-Designers to date have fallen into three broad groups: (i) fully AI design pipelines that generate fully formed designs with little opportunity today for human influence outside of filtering and selection of generated designs, (ii) sub-system or component designers that allow users to specify requirements for key elements of a whole design (e.g. pressure vessels, rotors, or fairing for an unmanned underwater vehicle (UUV)), and (iii) fast surrogate models that replace computationally intensive conventional design tools, allowing users to interactively modify designs with AI providing fast feedback on the results of the users' modification. For class (i) we have developed interfaces that allow users to rapidly sort and filter from 1000s of designs to a handful of interest on the basis of various clustering and ranking methods that are directed by user interest. For class (ii) we have developed design workbench interfaces that link together AI component Co-Designers to produce complete designs – these interfaces track design progress and alert users to incomplete sub-systems or inconsistencies between different parts of the design (e.g., when the required battery for a UUV exceeds the volume available for storing that battery). This same design workbench can integrate AI Co-Designers of class (iii), allowing for interactive design of a UUV fairing as part of our UUV design workbench.

In experimental hackathons during the program our tools were found to be useful to design teams in managing the volume of AI-produced designs, and, in the case of the UUV design, for managing the component AI Designers. We also provide experimental data from a user study on Human-AI collaboration on a separate problem of designing a 2-dimensional (2D) car to traverse uneven terrain.

2 Introduction

Current state-of-the-art design methodologies for cyber-physical systems have several limitations, especially for handling complex systems consisting of many sub-modules from different domains. Designer teams for CPS typically consist of subject matter experts with narrow expertise in each other domain (e.g., Computer-Aided Design, electrical systems, and control systems). A fundamental limitation of past approaches is that the knowledge is siloed and is hard to transfer to another design team such that they benefit from past collective expertise.

There is a large body of designs corpora available for many domains, however any designer, SME, will have only a limited capability to access them due to their number and the complexity in searching/browsing, especially for components and domains that don't have representations in a parameterized form that is easily comprehended by humans. For example, it is relatively simple to search for designs of electric engines that have a certain power and size. However, it is much harder to browse and select complex 3D shapes of objects with many components, and further incorporate them in the context for which they were designed.

This results in many systems designed using partial knowledge about some domains. Moreover, the real value of design corpora and software (e.g., CAD designs, simulation tools) cannot be fully realized unless they are consumable by humans for exploring large design spaces in collaborative settings. New design spaces are not explored because of substantial risks and costs involved. Design tools are built from the standpoint of an isolated domain expert and fail to present the design data to key stakeholders in a holistic manner.

Central to our approach is to focus on the humans and enable them to think about high-level design ideas and free them from highly intensive tasks that require specialized expertise that can be outsourced to an artificial intelligence Co-Designer.

The AI Co-Designer, Technical Area 1 (TA1), and human designers, Technical Area 2 (TA2), require very different spaces in which they operate, process, and interact with designs. This is because the TA1 spaces must be optimized for machine learning, optimization, and search algorithms. The AI Co-Designer space will be very high-dimensional (possibly, with thousands or tens of thousands of dimensions) that are not usable by humans. On the other hand, TA2 will require design spaces that are optimized for human interaction and facilitate communication to/from the TA1 teams. We expect the TA1 representations to mix AI Co-Designer with traditional computational systems. To address key challenges in the many ways and different computational systems that will generate the TA1 representations, we propose an "in-between" space, for translation and mapping, enabling symbiosis between (AI + traditional co-designer) and (humans), called Human-centric Design Symbiotic design space (or Human Design Space), which is shown in Figure 1. ACCEL will employ generative deep learning techniques based on Variational Autoencoders (VAEs) to abstract features associated with the co-designer spaces and map them into human-centric exploration spaces that are meaningful to the human designers with respect to the final design goal and are mappable back to the co-designer space such that human and co-designer can communicate between different spaces, optimal for each with respect to functionality. The architecture to support this vision is illustrated in Figure 2.




Figure 1 ACCEL vision addressing TA2 three main UIs for SDCPS.




Figure 2 ACCEL, as proposed, detailed architecture.

ACCEL employs multi-modal, multi-user interfaces using tablets, Augmented Reality/Virtual Reality (AR/VR) glasses, Kinect sensors that exploit computer vision and machine learning algorithms for intuitive navigation and interaction using sketches, gestures, voice, indexing and organization of large heterogeneous corpora for selecting design seeds. Our proposed user interfaces (UIs) will foster creativity, aesthetics, and lower cognitive load, and counter design bias with several novel features, e.g.:

- Non-traditional inputs such as Sketch based inputs and multi-domain design space that links the input specs and the overall design space (addressing TA 1.1 needs for design mining and design space construction).
- Cross-domain optimization of design using Graph Constraint networks and automated evaluation of performance specs using machine learning (addressing TA 1.2 needs for design composition).
- Low dimensional manifolds that help explore refinements to improve performance and viability (addressing TA 1.3 needs for design exploration).

3 Methods, Assumptions, and Procedures

Symbiotic design interfaces support the three phases in the ACCEL workflow, as shown in Figure 3. Correspondingly, our TA2 vision is to build three symbiosis interfaces supporting design exchange in various stages between human designers and AI co-designers. Our aim is to augment the humans and increase their intuitive understanding of the design through the mediated AI - and - ideally the rationale the AI exposes to the designer in visual ways. We envision a future of human augmentation and immersive technologies that is cross-disciplinary supporting design hierarchies “by blending digital and physical environments to extend and enable new spatial, remote, and cognitive capabilities of humans, richness of Human-AI interactions, and workflows.” We call the overall framework as Symbiotic DesignVRse and workflows as DesignFlow.




Figure 3 Three stages of the ACCEL symbiotic interface: Design Space Creation, Design Composition, and Design Space Exploration.

3.1 Requirements for the Symbiotic Design User Interface

For the human side there are several requirements for the symbiotic UI:

- **Human intuition:** Support new concept ideation and concept transition
 - An example of this is a Sketch-based interface that enables a designer to ideate a design using a tablet. Next, the initial design can be used to browse for similar parts in a corpus database and extrude it to 3D using AI techniques.
- **Collaboration and Creativity:** Support multi-device multi-team collaboration
 - We implemented a cloud-based system that ensures that all team members are synchronized. This means that somebody wearing Virtual Reality (VR) glasses and performing manipulation will share the same design with somebody doing design.
- **Take advantage of different capabilities of human-machine (H-M) interaction:** Support design space exploration through different visualization modalities
 - We have implemented and demonstrated UI employing desktop (mouse + kb), tablets, and VR goggles, assessing the effectiveness of each modality during the workflow in Phase II of SDCPS.
- **Respect Legacy:** Reusing design knowledge through multi-modality retrieval
 - function(ontology)/part/sketch-based retrieval.
- **Symbiosis Sand Box:** Support easy design composing in a 3D environment.
 - VR-based design sandbox

A visual of these requirements is shown in Figure 4.




Figure 4 Symbiotic UI technical requirements.

3.2 Design problems and sandboxes

3.2.1 Southwest Research Institute Unmanned Air Vehicle

The Southwest Research Institute (SwRI) TA3 team provided design challenges for unmanned air vehicles (UAV) and urban air mobility (UAM); UAV challenges were provided for the year one and Fall 2022 Hackathons, while the UAM challenge was set as the end of phase 1 demonstration. For both vehicle types, a corpus was provided that was comprised of a catalog of aircraft parts (many parameterized), a graph-based schema for composing parts into designs, and a software pipeline that can take design graphs as input, generate 3D CAD models, and run sets of simulations of the design through a flight dynamics model (FDM) oracle. These simulations provide both detailed flight characteristics for designs together with simulated performance on TA3-provided test flight paths, and scalar scores summarizing the aircraft's ability to traverse the flight path. For the first hackathon and end of phase 1 demonstration, the design challenges were based upon vehicle performance over a set of flight paths, each of the paths designed to test particular flight characteristics. In the Fall 2022 hackathon, the test challenge was a more realistic delivery scenario. A UAV would fly between two points carrying a cargo, land for unloading, and then take the return path with an empty cargo container. The test flight path included both takeoff and landing, curved flight, and long straight sections.

3.2.2 Systems & Technology Research LLC UUV

The Systems & Technology Research LLC (STR) Technical Area 3 (TA3) team provided a total of three design challenges during phase one of the program; a choice of two problems for the mid-phase hackathon, and a single challenge for the end of phase demonstration, a variation upon which was used for the Fall 2022 hackathon.

3.2.2.1 Hackathon challenges – year one

The mid-phase hackathon gave teams a choice of two variations on marine survey. The first problem centered upon a sub-arctic water sampling mission whereby a host vessel travels to the pole carrying a survey vessel. On reaching the pole, the host vessel attaches itself to the ice, bores an antenna through the surface to air to establish communication with a remote monitoring station, and releases the survey vessel that embarks on a survey route sampling water properties using a conductivity, temperature, and depth (CTD) sensor, periodically sending the collected data back to the host vessel, which relays them to the science team. To address the challenge, teams had to develop designs for both the host and survey UUVs.

The second hackathon challenge involved design of a fleet of vessels to conduct a seafloor mapping mission where each vessel conducts multiple survey sorties, leaving a shore station and transiting to the survey area, conducting sea floor scans with sonar arrays, then returning to the shore station for battery recharge before setting out for the next sortie. To meet this challenge, teams had to develop designs for a fleet of UUVs that could work together to conduct sorties with constraints on total fleet mass, turnover time for recharging. Design parameters include sonar array characteristics, operating depth for the vessel, transit, and survey speeds.

For the hackathon, we partnered with TA1 teams from SRI International (SRI) and the University of California, Berkeley (UCB), demonstrating how the UIs presented below could be used to browse large numbers of AI-generated designs, sort and filter them to designs of interest and compose fleets for testing, in addition to browsing fleets generated by the AI co-designers.

3.2.2.2 Phase 1 demonstration challenge

The end of phase demonstration challenge centered upon delivering a large (~1m diameter) seafloor seismology sensor to the seabed near the North Pole and then returning to base. Teams were tasked with designing a vessel capable of completing this mission within given time and location accuracy constraints under conditions with variable ocean currents and noisy, low frequency localization updates. For this challenge, we partnered with TA1 teams from SRI and Siemens Corporation (Siemens) to demonstrate our UIs for mission simulator understanding and UUV controller development, as well as browsing and filtering a corpus of AI-generated designs.

3.2.2.3 Year 2 hackathon

The hackathon for year 2 was a perturbation on that of the Phase 1 demonstration challenge – a similar mission, to deposit an undersea sensor at a given location, but with a more complex route over a more realistic subsea terrain. A further wrinkle was to account for the effects of adding a (towed) sonar array to the UUV. We worked extensively with the SRI TA1 team to develop an integrated Human-drive, AI-assisted UUV design tool.

3.2.3 Proxy 2D Car problem

To aid with experimentation and development of interaction interfaces and supporting AI, we also worked with a proxy problem, namely designing a 2D car that can traverse a path without getting stuck on the terrain, shown in Figure 5. Early experiments supported sketch-based design, which were later augmented with a parameterized design interface supported by a reinforcement learning-based AI co-designer.




Figure 5 Proxy 2D car problem enabled quick experimentation with different concepts, such as parameter sweeping, as well as ideation and development ideas for the 2D sketch interface.

3.3 Support for single users and design teams

Our approach supports both single-user and team design. We have both browser-based and Virtual Reality-based interfaces. The browser-based interface offers single user support when used in the desktop but can also be viewed as panels in VR for team-based design sessions. The VR environment allows for a shared experience when working in distributed teams that can also be more convenient for local teams, given the immersive views of designs available to VR users. The virtual space offers at one end of the spectrum an ‘infinite’ desktop across which to spread tools and views, and at the other, allows multiple users to share a viewpoint in space so that all team members can have exactly the same view of a particular object or analysis, which can help eliminate ambiguity that can arise even when team members are physically present in the same room.

3.4 System architecture

Our TA2 design interface is a distributed application comprised of a number of back-end and UI elements. The overall architecture is sketched in Figure 6. Design data is stored in a MongoDB document database (which stores design descriptions including high-level summary data for a design, its available test results, and pointers to identification (IDs) of any parent design from which a design may have been derived) and Minio (where we store design graphs, binary objects such as 3D models, graphical experimental outputs, vector experimental data). Our UI is distributed as a collection of Web and VR applications. The main ACCEL design browser acts as a focal point to co-ordinate the design storage with our TA1 and TA3 partners and hosts the design summary browser where users can review the overall structure of the design space available and prompt deeper exploration and design refinement using the editors and advanced browsers available in the VR environment and their browser-based equivalents for the non-VR user.




Figure 6 ACCEL high-level system architecture. Blue elements comprise the data storage backend. Orange elements are various UIs and co-ordination servers. Green elements indicate Representational State Transfer (REST) Application Programming Interface (API) connections to external services from TA1 and TA3.

3.5 Workflows and interfaces




Figure 7 High-level breakdown of design workflows.

As shown in Figure 7, we have partitioned the design process into three stages:

- Design initiation: a human designer or AI Designer creates a new design.
- Corpus browsing and filtering given a catalog of existing designs and test results, human designers whittle down to a few designs of interest for closer inspection and possibly further development.
- Design exploration: the human and AI Designers interact to explore variations on designs selected by the human designer.

These stages do not necessarily flow as a linear chain; after exploring some designs and producing variants, users may wish to return to corpus browsing to see how their new designs compare to pre-existing ones. At any stage of the process, the human designer may have an insight that inspires a completely new design and returns to design initiation; or new AI Designers may come online that can initiate new types of design or offer better interactive assistance to help human designers flesh out new design ideas.

Each of the stages outlined above may be approached using multiple specific workflows; we shall provide some example workflows for each stage as motivation for the specific UI elements that we have developed to enable human and AI Designers to collaborate throughout the design process. As described in the previous section, we have UI elements that can be used on the desktop, in VR or both. For example, 2D data plots and design data tabulation is implemented as web apps that can be used inside the browser by a desktop user, or brought up on virtual screens inside the VR environment for VR users. Some sample UIs, and their use in the workflow, are shown in Figure 8. For some other tasks we have separate implementations for VR and desktop – for example the component assembly interface has a browser-based implementation for use with mouse and keyboard and a 2D projection of 3D data, while the VR implementation allows for much more direct, fully 3D, interactions. As illustrated in Figure 9, our VR environment offers an interface that allows users to arrange the specific tools they need into targeted workflows to address particular design problems (and then reconfigure later on when the emphasis changes).




Figure 8 Expanded workflow for ACCEL, detailing some of the stages of the basic blocks.



Figure 9 Setting up the workflow in VR. The main user, represented by gloves in the center, is picking up tools to be employed during design.

3.5.1 Design initiation

Our system can import an initial set of designs from AI Designers, or from a dataset of those composed by humans using TA3 tools. In addition, we offer a set of design composition tools within the browser and VR environments.

During design composition, the designer(s) are using different modalities to refine a design or create a new one from component parts. We have explored two ways to perform design composition. The first is AI-assisted design completion. Users place components in an approximate position using VR or a web-browser. Then the partial design is sent to TA1 (SRI TA1 in this case) that completes the design and then returns the completed design to our TA2 interface. The second approach can take a complete design and have the users changing and swapping components. The advantage of the UI we have created is that we can employ both VR and traditional approaches to create designs and have everybody collaborate in a shared space.

We illustrate the design composition as shown in Figure 10 for VR and Figure 11 for a web-browsing interface.



Figure 10 Design composition in VR. Multiple users can participate. The inset in the lower right shows a 3rd person view of the design process from the room.




Figure 11 Design composition in a web-browser using a traditional input modality with keyboard and mouse.

3.5.2 Design analysis and exploration interfaces

Once imported into our tools, UAV designs can be further explored and annotated. The following exploration tools are available:

- Component swap – select a component in the design (e.g., a battery) and swap it for another.
- Parameter sweeps – for those parameterized parts in a design, initiate sweeps of parameter values to explore the possibilities of a given design topology.

In addition, the design completion capability (see ‘Design initiation’, above) can be utilized when adding new components to a design – users place the component of interest in their desired spatial location, and the AI Designer searches for complete designs that include the component, including addition of extra connecting parts if necessary. Finally, the interfaces support component removal from designs.

3.5.3 Corpus browsing and filtering.

During the corpus browsing and filtering, the users analyze existing corpora of designs (e.g., for UAM), selecting a few designs that are promising which can be edited or replaced with a desired part by browsing through similar parts. Figure 12, 13 and Figure 14 show different views of a browse/search for parts suitable for the candidate design. The similarity between different parts/designs can be done using different criteria such as based on shape, function, etc. SRI has also developed a web-based UI for browsing which was demonstrated during the final PI meeting for Phase I, and is shown in Figure 15 and Figure 16.



Figure 12 VR browsing for candidate design.




Figure 13 VR search/browse for suitable parts.

Integrate with legacy tools and knowledge contained elsewhere


- Organizing large corpus (COTs) and prior designs



- Supporting cross modality browsing:



VR interface



Web interface

- Supporting multi-modality browsing/searching



- Efficient searching/sorting/filtering




Figure 14 Organizing corpus for easy browsing and search and viewing in VR.




Figure 15 Web based design comparison allows the designer to assess each shape and parameters of several designs.




Figure 16 Design browsing in a web interface, comparison of multiple designs in terms of the flight path and accomplishing mission.

3.5.4 Design Exploration

The role of design exploration is to take a number of complete designs and vary different parameters (continuous or discrete) in order to find a design that is optimal with respect to the requirements. During the first stage of exploration, we employ either a surrogate model or the TA3 pipeline, provided by SwRI/STR, to evaluate various designs by modifying the parameters of the design (e.g., changing the size of the UAM, etc.). Then the TA2 user employs the symbiotic interface to assess areas in the parameter space where the designs meet or fail the requirements. By performing different exploration (using histograms or projections in a lower dimensional space, the designer can instruct TA1 on what regions in the parameter space are more likely to contain a good design. Finally, the TA2 users select a few designs for Comparison and further exploration.

We have used several modalities for design exploration, such as:

- Pair plots
- Histogram selection
- Projection in a low-dimensional manifold (e.g., t-distributed stochastic neighbor embedding (tSNE))
- Sorting on table
- History
- Comparison

Two examples are shown in Figure 17 (histogram intersection) and Figure 18 (comparison).




Figure 17 Design filtering in the browser UI. The histograms show good/bad designs. With selection on the histogram, a large number of bad designs are able to be removed.




Figure 18 Design comparison in ACCEL browser UI.

3.5.4.1 Workflow example – UUV corpus filtering and design comparison with Siemens TA1 outputs

We will illustrate one possible filtering workflow with a study of finding a known-good design from a set of design candidates provided by the Siemens TA1 team. In this example we will use histogram selections to first whittle the data from a large number of designs (205) to a handful, utilizing tabular sorting to find 3 designs for close comparison. The initial data reduction performed here with histograms could equally be conducted by selections from pairwise data plots or projection plots.

In four steps we go from 205 designs to a design comparison UI that shows 3 designs side-by-side, including the known viable design and two others that are near-misses that could become viable with adjustment to part placements.

1. Design overview table and histogram plots:

Our starting point is the filtering UI, shown in Figure 19, which consists of a table containing key performance characteristics of all designs alongside histogram plots of those variables. The table can be sorted by any column, and users can select subsets of all designs by making selections via mouse-dragging actions over the histogram data. Finally, clicking any row of the table will add the corresponding design to our design comparison window.



Figure 19 The filtering UI.

2. Focus on fewer designs using histogram selections:

In this view, Figure 20 illustrates how one might zoom in on selections made using the histogram plots. We have made selections on the two histograms highlighted in red, the net buoyancy over

mass histogram (centered around 0) and the histogram that shows the vertical difference between the center of buoyancy and center of mass (we select positive values to capture those vehicles that will remain upright) – the precise bounds of the selection depend upon the histogram boundaries, so this initial filtering may capture some designs that have center of buoyancy below center of mass. The orange overlays on the remaining histograms show the histograms of data from the designs that have been captured by our histogram selections. In this case, there are only 5 designs selected, so the orange overlay histograms are correspondingly sparse.




Figure 20 Histogram selectors allow focus on particular designs.

3. Sort remaining designs in tabular view and select for comparison:

The data table, shown in Figure 21, updates live as histogram selections are made to show only those designs captured by our selections. The top table shows the 5 designs that we have filtered down to from the original 205. We sort the data along the column relating center of buoyancy to center of mass to clarify which of our selected designs have center of buoyancy above center of mass and click the table row corresponding to each of the suitable designs.

#	Compare	Design ID	Concept	hotel_power	drag_force	thrust	thruster_effice	max_speed	energy_capaci	cog_cob_dz	collision_volum	net_buoyancy	battery_pv_vol	success	mission_end_e
133		1166	propeller	1.6	16.84618736	16.84618736	0.835703	2.5	2952000000	0.0383887479	0.0977622213	-0.0256196627	2.3336646609	0	2413613484.5
148		2065	propeller	1.6	12.97610512	12.97610512	0.841151	2.5	1440000000	0.0020467753	0.0552392065	0.0489233396	1.7900772096	1	1027809346.4
167		128	propeller	1.6	16.65418528	16.65418528	0.836004	2.5	2170080000	0.0135571456	0	-0.0025287522	0.8639285962	1	1635925607.6
187		3085	propeller	1.6	8.2028878399	8.2028878399	0.843509	2.5	1926000000	-0.0072459092	0.0734054816	-0.0404491648	1.4840140933	0	1666653969.4
199		4224	propeller	1.6	27.44175072	27.44175072	0.817675	2.5	3592800000	-0.0003336094	0.1534283651	0.1277182340	7.7522255973	1	2691631158.7

Sort

#	Compare	Design ID	Concept	hotel_power	drag_force	thrust	thruster_effice	max_speed	energy_capaci	cog_cob_dz	collision_volum	net_buoyancy	battery_pv_vol	success	mission_end_e
187		3085	propeller	1.6	8.2028878399	8.2028878399	0.843509	2.5	1926000000	-0.0072459092	0.0734054816	-0.0404491648	1.4840140933	0	1666653969.4
199		4224	propeller	1.6	27.44175072	27.44175072	0.817675	2.5	3592800000	-0.0003336094	0.1534283651	0.1277182340	7.7522255973	1	2691631158.7
148		2065	propeller	1.6	12.97610512	12.97610512	0.841151	2.5	1440000000	0.0020467753	0.0552392065	0.0489233396	1.7900772096	1	1027809346.4
167		128	propeller	1.6	16.65418528	16.65418528	0.836004	2.5	2170080000	0.0135571456	0	-0.0025287522	0.8639285962	1	1635925607.6
133		1166	propeller	1.6	16.84618736	16.84618736	0.835703	2.5	2952000000	0.0383887479	0.0977622213	-0.0256196627	2.3336646609	0	2413613484.5

Figure 21 Tabular views allow comparison.

4. Compare designs of interest side by side:

The three selected designs are shown side-by-side on our comparison UI, as illustrated in Figure 22. The bar charts allow for rapid visual comparison of particular attributes, and the 3D model renderings (which can be zoomed, panned, rotated, and made transparent for inspection of internal components). We can see that design ID 128 (the center design) is our feasible design; the other two clearly require more work to ensure all parts are packed within the fairing.




Figure 22 Side-by-side design comparison.

3.5.5 Infinite shelf

The Infinite shelf is the core element of the Virtual Reality-based design exploration system. It offers users an immersive experience of exploring the design spaces using the VR headset and the controller. In contrast to traditional 2D browsing and visualization methods the infinite shelf provides a natural way of visualizing the metadata and manipulating the 3D object simultaneously. Users can browse a dataset containing thousands of designs, search and retrieve the design parameters, visualize the CAD models, and re-order the dataset based on the sorting metrics. With the inherent advantages of visualization and manipulation afforded by VR, the infinite shelf can be an efficient media for browsing a large design dataset generated by AI Designers.

Design of the infinite shelf metaphor

Interface design

Figure 23 illustrates the interface of the infinite shelf. The user interacts with the system using the control panel shown in the middle of Figure 24. This panel consists of three fundamental functions, data loading, sliding browser, and parameter visualizer. The data loading function allows users to load an existing dataset from a local or web repository, for example, loading the UAV data created by the SwRI team. The data are processed during the loading to match the format of the VR system. The sliding browser gives the user an overview of the dataset. They can sort and retrieve the designs using different metrics. As shown in Figure 23, the parameter visualizer is an additional panel that attaches to the right-hand controller. The parameter visualizer provides information on the design parameters that are loaded from the initial dataset. Considering each design can have tens of parameters, a scroll-down menu is used in the parametric visualizer to browse those parameters.



Figure 23 The main panel of the infinite shelf, with data loading function button on left, sliding design browser in the middle, and parameter visualizer on the right.




Figure 24 VR system browsing the UAV dataset with the shelf display.

Visualization design

To use the spatial advantage of the VR space, we re-designed the visualization interface for the design browsing using an infinite shelf to display the design data in a 3D way. There are two display options, the shelf display, and the sphere display. Both options inherit the 2D control panel allowing users to switch between datasets and customize the sorting and retrieval metrics. Figure 24 shows a user browsing the UAV dataset using the shelf display. The shelf consists of hundreds of bubbles that highlight the UAV CAD models. Figure 25 shows the user browsing the UAV dataset using the sphere display. In the sphere display, users can use a miniature sphere on their left hand to rotate the entire sphere. So, they can access interesting designs without moving around in the virtual environment.




Figure 25 VR system browsing the UAV dataset with the sphere display.

3.6 Interfaces for specific design problems

Many of the interfaces described in the previous section have been applied to multiple design problems – available variables, contents of design graphs, and CAD models change between problems, but the basic principles of display and interaction are equally applicable across problem domains. There is also a role for more specialized interfaces that service the needs of a specific design problem or sub-problem; we have developed a number of these, and they are described below.

3.6.1 UAV/M – flight dynamics

To support design analysis and design selection for the UAV/UAM problem, we developed a tool that allows users to browse results of the flight dynamics model provided by SwRI as an oracle for both scoring the performance of air vehicles on challenge flight paths and for producing detailed analysis of a vehicle's flight performance across a range of 'trim-states', where the FDM code searches for settings of vehicle parameters (propeller speeds, wing angles) that allow for stable flight (if possible) across a range of forward speeds (0 to 50 meters per second). While the flight path score delivers a single value that can be used to compare designs, it is calculated on a

single path and is not necessarily a good predictor of the general airworthiness of the vehicle. The path score may also hide relative efficiency of designs, for example, if two air vehicles have similar score, but one of them must run its motors at near full power for the duration of the flight, while the other has more modest power usage, then the second design may be preferable since it is likely to be more durable over an extended period of use. The trim state data is more informative than the path score, but voluminous – for each of the 51 trim states there are 10s of measurements. In the Fall 2022 hackathon, we developed a tool to help subject matter experts and the design teams (in particular, our tool was used by the SRI TA1 team during the hackathon) navigate the detailed FDM data from designs numbering in the 1000s. For the TA3 teams this function is useful when comparing submitted designs from the TA1 teams. The SRI TA1 team, whose workflow was AI-driven design generation (1000s of designs) with human down-selection to 10s of designs for final AI-based optimization for the challenge flight path, the tool was used by the human designers for the down-selection process, i.e., to identify a small number of most promising designs from the 1000s generated by the AI Designer.

The tool that we developed included three main functions – experiment selection, FDM overview, and design evolution. Experiment selection, shown in Figure 26, tabulates summaries of available experiments alongside plots of particular performance figures requested by the TA3 subject-matter experts, and allows the SMEs to select sub-sets of available experimental data for further analysis, on the basis of familiar performance plots. The FDM overview function provides a sortable, filterable, graphical overview of the FDM results for all designs and allows users to perform a deep dive into the detailed results of single designs selected by double-clicking on the graphical overview. The design evolution UI was developed for the SRI TA1 team whose approach involves producing multiple generations of (up to many 1000s of) designs, the goal being that design quality improves over generations. The UI allowed the SRI TA1 team to rapidly confirm this generational design trajectory. We'll illustrate each of the three functions in the following sections.




Figure 26 The experiment selector.

3.6.1.1 Experiment selector

Figure 26 shows the experiment selector, which provides an overview of all FDM experiments conducted. The upper part of the UI (highlighted with an orange surround) offers several plots favored by the TA3 SMEs, and the lower part of the UI (blue surround) tabulates summaries of the experiments. Users can select particular experiments either by clicking points (or using an elastic-band tool) on the plots, or by clicking rows on the table.

The SwRI oracle flight dynamics model is executed in a batch manner on a single design topology – this topology being some set of fixed components (e.g., specific motors) linked together by parameterized components (e.g., struts). The design may also include parameterized components such as wings of variable length, or rotors of variable radius. For purposes of the FDM oracle, the design also includes the controller parameters for executing a flight path. Inputs to the FDM are the design topology, together with a text file of study parameters that specifies possibly several combinations of values for the parameterized parts of the control variables. The FDM is executed for each parameter setting, resulting in a set of results that we'll call an experiment. For example, when evaluating the challenge problem of transporting cargo from A-to-B and then returning from B-to-A with an empty cargo container, this file has one run with cargo mass set to 500 grams, and a second run with 0 cargo mass, representing the A-to-B and B-to-A flights. The SwRI SMEs may conduct design sweeps by, providing an input file that has several lines where parameters of interest are varied over lines of the file (e.g., wingspan, or length of some key connecting strut, or both), and the FDM is run for each. The experiment selector of our FDM review tool tabulates each experiment (the truncated table highlighted with a blue surround in Figure 26), presenting summary values and min/max values of certain flight variables over the set of runs captured by each study. This table allows users to quickly review different experiments, for example, to get a quick overview of the performance of different TA1-generated designs; the table can be sorted by various columns and there are filters based on experiment name if a particular set of results is known to be of interest. The upper part of the experiment selector shows summary plots frequently used by the SwRI SMEs to review designs. There are plots for reviewing how forces vary over all experiments, how control parameters are set over experiments, and what specific speeds, accelerations and landing speeds were selected in each experiment. Users can choose experiments of interest by either selecting them from the table (using checkboxes in the leftmost column), or by selecting them by clicking points (or surrounding several points with an elastic-band tool) in the graphical plots. Selected experiments are highlighted in the table and displayed with a different glyph from unselected experiments in the plots. As we shall see in the next section, selections made on this page can be propagated to the FDM overview UI.

3.6.1.2 Flight Dynamics Model overview

The FDM review tool, illustrated in Figure 27, is comprised of an overview area on the left, with a detailed view section on the right. In this example, the user has selected three different designs for closer study; there is one collapsible card for each selected design. For illustrative purposes,

we've chosen to show a different analysis for each design, flight path review at the top, 3D views in the middle, and detailed graphical data at the bottom. The analysis views are tabs on each card, allowing for side-by-side comparison of, e.g., detailed graphs for several designs.




Figure 27 The Flight Dynamics Model overview tool.

Users can sort and filter designs (including a filter that will include only those designs selected in the Experiment Selection UI) and select designs of interest for more detailed inspection in the detailed view. For each selected design, a detailed view card is displayed that has tabs to show 3D model views, flight path overview and detailed parameter plots for a design.

The overview has a graphical display of the FDM trim states, making for very rapid assessment, e.g., which designs have more valid trim states than others, and therefore are more likely to be controllable over a wider range of flight conditions. The graphical overview can be colorized according to several different variables: trim state validity, motor current ratio, motor power ratio, wing load ratio, and peak battery current. Trim state validity simply colorizes valid trim

states green and invalid states gray. The other options colorize from green (low) to red (high) according to the value of the selected parameter. The overview is also sortable and can be filtered.

In the zoomed-in example shown in Figure 28, battery peak current is used to colorize the graphical overview (in blue surround at the bottom of the screen grab) – those designs with trim states that have higher peak currents immediately leap out in red. Each row of this display corresponds to a single run of the FDM. The columns of the display correspond to each of the trim-state speeds assessed by the FDM, from 0 to 50 meters per second; gray cells indicate that the FDM could not control the vehicle at the given speed. Again, it is immediately clear which designs have more valid trims than others. The overview can display 1000s of designs, and users can zoom the display in and out, and scroll (when zoomed in) to get an overview of large datasets at different scales. Hovering over a row of the overview will (optionally, via a control on the plot) show a tooltip with detailed numerical data on the corresponding design.



Figure 28 Zoomed-in view of the graphical overview and its controls.

Above the graphical overview, surrounded by orange in Figure 28, we have a set of controls for helping users navigate the data. The 'Display' drop-down at the top chooses how the overview should be colorized (trim state validity, motor current ratio, motor power ratio, wing load ratio and peak battery current are the options), there is a sort drop-down, where designs are sorted

by number of propellers, with other options being: maximum distance flown at any trim speed, maximum hover time, lift-to-weight ratio, path score, battery remaining at mission end, number of wings, number of valid trim states, and number of interferences. To reduce the number of FDM results included in the graphical overview, the filters section allows users to control which different results are included in the overview. In our example we have filters (disabled) on maximum flight distance flown at any trim speed, and on maximum hover time; if enabled, these filters would show only those designs whose performance exceeds the specified limits. Other filtering options include path score, battery remaining and lift-to-weight ratio. This section of the UI also includes a checkbox that will display only those FDM results from experiments selected in the Experiment Selector UI (this filter is applied before those noted above – allowing users to start from selected experiments and then down-select further by hover time etc. as desired).

Finally, by double-clicking on a row of the graphical FDM overview, we can open a detailed view for more in-depth analysis of the selected result.

The FDM detail viewers offer 3 views that provide more information about individual runs of the FDM. Figure 29 shows examples of each view for three different designs from the Fall 2022 Hackathon – note that the cards containing these plots have tabs for these views, so users can select multiple designs and, for example, compare the flight path data for each one.

The flight path viewer, Figure 29 (top), shows a 3D plot of the trajectory computed by the FDM. The line plot can be colorized by various FDM parameters – in this example, the data are colorized according to which trim state the vehicle is operating in at every point along its path. Other parameters to colorize the data include instantaneous values for vehicle pose, velocity and lift, drag and thrust forces. To the right of the 3D plot is a summary of flight performance together with the controller parameters used (if the vehicle is using the SwRI-provided linear quadratic regulator (LQR) controller).

The second detail view shown in Figure 29 (middle) is the 3D view. This viewer renders the Stereo Lithography (STL) model of the air vehicle in an interactive display that allows users to rotate the object, zoom in and out and so on.






Figure 29 Flight Dynamic Model detail views. Top: flight path review. Middle: 3D view. Bottom: Flight Dynamic Model plots.

The third detail view shown, Figure 29 (bottom), is the FDM plots view. This allows users to see detailed plots of values computed by the FDM for each trim state. In the example above, the plots shown are Wing Loads, Drag, Thrust, Lift, Pitch and Battery Current. The design being inspected has two wings, hence two plots for wing load. For drag, lift, and thrust, there are both body and world values, so size plots. In this example, there is a single battery; if a design has multiple batteries, then there would be a line plot for each on the Battery Current axes. Users can either select one plot at a time from the 'Plots' widget, or they can use the drop-down menu button to select a preset, from Forces (drag, thrust, lift), Accelerations (linear and angular in separate axes), Flight time, distance (as the name suggests – max flight time and distance attainable at each trim state, on two axes) and Electrical (Battery current and motor current on two axes). Presets can be augmented with plots chosen one-at-a-time, and have individual plots removed. Once a user has set up plots for a particular result, the next result opened by double-clicking on the FDM graphical overview will inherit that plot layout, making side-by-side comparison of FDM data straightforward.

3.6.1.3 Design Evolution UI

The design evolution UI, shown in Figure 30, is a straightforward UI that helped the SRI TA1 team verify progress from their AI design generator and design filters. The example shows two generations of designs, generated, and filtered by the SRI TA1's AI. At this stage of design generation, designs are being filtered to select for hover capability.




Figure 30 The design evolution UI, showing two generations of designs.

Generation 0 includes 5,517 designs, and our screengrab shows a histogram of hover times. The pale brown bar shows the sum proportion of all vehicles with non-zero hover times. In generation 0, fewer than 2% of generated designs can hover. For generation 1, we plot the same histogram, and see that now 20% of designs can hover – a tenfold improvement over the previous

generation, confirming that the quality of AI designs is improving as expected. Further generations can be ingested and inspected in this way. This example shows only hover time, but histograms of other FDM variables are also available for similar types of generation summarization. In addition, the generation cards have a tab showing the same force plots found in the experiment selector UI, used by the SMEs to quickly assess designs, as an alternative view of population evolution.

3.6.1.4 The UI in use during the Fall 2022 Hackathon

We shared the UI with the TA3 SMEs and SRI's TA1 team. While time pressure did not allow for formal evaluation, TA3 SMEs provided feedback that was instrumental in developing the experiment selector and found the FDM overview UI to be promising. Again, without formal evaluation, we received positive feedback from the SRI TA1 team, who reported that the FDM overview tool made the human-in-the-loop portions of their design pipeline much more efficient than their in-house tools.

3.6.2 UUV design composition

In the Phase 2 hackathon, we collaborated with the SRI TA1 team to develop a human-driven, AI-assisted UUV design tool. The tool allows a human user to specify mission constraints and assemble a UUV design from component parts, with AI assistance for:

- Vessel fairing design,
- Propeller design for propulsion system,
- Pressure vessel design,
- AI Surrogate model for fairing computational fluid dynamics (CFD) analysis (enabling interactive fine-tuning of fairing by human designer as required).

Mission

name
obs_placement

mission_depth_m
1300

mean_speed_m_per_s
1.1

mission_duration_days
55.871212121211

mission_distance_km
5310

Dependent variable
Speed Duration Distance

net_buoyancy_margin
0.02

pitch_requirement_deg
16

Environment

water_density_kg_per_m3
1027

UUV design

UUV name
Mission1_50pc_effy_edits

Save Rename New

Summary

drag_at_mission_speed_N
61.110

propulsion_efficiency
0.801

hotel_load_W
23.000

total_energy_J
891027272.727

num_pvs
2.000

total_mass_kg
1350.307

net_buoyancy_margin
-0.002

Fairing
Payload
Propulsion
Batteries
PV volume
Mass shifter
Foam
Packing

Figure 31 UUV design tool overview.

The tool guides the design process from initial energy estimation through vessel packing. The overview area of the UI, as seen in Figure 31, shows mission requirements, key vessel properties and a set of indicators showing the status (good/warning/error – in this example all are green indicating a good design) of UUV subsystems throughout, making clear when changes to one subsystem propagate to others. The UI also maintains a design history, as shown in Figure 32, storing to a database the design at each subsystem edit. A visual display of this history, as a series of subsystem status indicators, is shown in a panel under the mission specification area of the UI. This same history panel also includes a graphing area where users can plot vessel summary values as a function of design history. When changes are made in the tabbed work areas, the UI stores modified designs to a history database, together with the name of the subsystem that was modified. This view shows the indicator status at each stage of the design’s history. Users can also choose to plot how key design parameters change across the design history.

index ▲	F ▲	Pload ▲	Prop ▲	Bttry ▲	PV ▲	Shftr ▲	Foam ▲	Pack ▲
0								
1								
2								
3								
4								
5								
6								
7								
8								
9								
10								
11								
12								
13								

Figure 32 The history section of the UI.

Returning to the vessel design workflow, energy estimation is initialized by users entering mission requirements (duration, average speed, required sensor and comms subsystems) together with ‘reasonable’ initial estimates for vessel drag and propulsion efficiency. AI co-designers are invoked to produce concrete fairing and propeller designs that will meet or exceed the ‘reasonable’ initial estimates, and the human co-designer selects a specific fairing and propulsion system; at this point we have a more realistic energy requirement, and the UI provides the required battery volume/mass to satisfy that requirement (with a user-specified safety margin), and the co-designer can split the total required battery mass into separate battery packs if desired. With battery packs and dry hotel components known, the AI co-designer is invoked to provide an array of pressure vessel designs that satisfy the volume requirements for battery and hotel components and are capable of withstanding pressures at the specified mission depth (plus some safety margin). Aided by calculators built into the UI, the human co-designer moves on to add mass shifters to allow for pitch/roll control, and to specify syntactic foam pieces such that the vessel is near neutrally buoyant. The final phase of the process is to solve the human-driven packing problem – arranging the specified components such that the vessel has proper balance, permitting it to both stay stably upright and to attain required pitch range. The packing UI allows the human co-designer to modify the shape of the fairing to better accommodate the arrangement of components, for example, reducing the overall volume of the fairing to eliminate wasted space, or reshaping a specific part of the vessel to allow more room for a component arrangement that satisfies the balance requirements. When making these changes, the UI keeps track of changes to fairing mass and volume (that will affect overall balance of the vessel) and it also invokes an AI CFD surrogate model that allows for rapid re-computation of drag for the new fairing shape, allowing the human co-designer to see what effects the modified fairing shape may have on vessel energy requirements. If energy needs increase, we may need to iterate over the

selection of battery, pressure vessel design and so on, in order to meet mission goals. If energy needs decrease, the human co-designer can choose to stay with current component selections or enter a new iteration where they reduce battery needs, leading to smaller pressure vessels, further lowering the mass of the vehicle, and likely reducing construction and transportation costs.

3.6.2.1 UI work areas

To the right of the mission and UUV summary, as seen in Figure 31, are a set of tabbed work areas for each sub-problem of the design. The following sections describe each of the available work areas through example. The example visualizes an essentially complete design with screen grabs captured as figures to illustrate a subset of available actions, and their effects on the UI indicators.

Fairing work area

Figure 33 illustrates the Fairing design work area. Tabs across the top of the UI are ordered left-to-right (matching the top-to-bottom ordering of the subsystem indicators) to reflect the design workflow. Energy needs due to hotel and locomotive power are set first, in order to properly specify battery and pressure vessel volume, with buoyancy control and packing as the natural conclusion to the process. In Figure 33, the fairing panel shows controls that, together with estimated drag used for initial energy estimation, set targets for the AI fairing designer, which can be invoked from the UI. Results from the designer are loaded into the fairing catalog (they can be renamed) and a 3D rendering of the fairing used in this design is shown.




Figure 33 Fairing design work area.

Payload work area

The payload work area, shown in Figure 34, allows users to select required mission payload modules from the column of drop-down buttons on the left-hand side. In this example, the user is in the process of choosing a doppler velocity log (DVL) module to add to the design. Specific components for each module type are loaded from a database of components derived from information provided by STR, filtered by mission parameters (e.g., maximum mission depth) for suitability. Both the background of the 'Add DVL' button and the payload indicator on the UI are red, indicating that the design does not yet include a DVL. On the right-hand side of the work area is an array of module 'cards' for each module currently in the design. Cards for dry components (e.g., INS_0, GPS_0, and Radio_0) include a 'Pressure Vessel (PV)' drop-down that indicates the pressure vessel to which each component has been assigned. Cards for electrical components include a duty cycle entry, used in hotel power calculations. Underneath the drop-down buttons for modules is a small worksheet for specification of a sonar array, whose drag is added to that of the fairing when computing required locomotion power. A checkbox indicates whether the array is included in the vessel design.

Fairing

Payload

Propulsion

Batteries

PV volume

Mass shifter

Foam

Packing

Fairing

Payload modules

Propulsion

Energy and Battery

Pressure Vessels

Pitch control

Foam

Packing

Add INS

Add DVL

Teledyne_DVL_Tasman_300kHz

Teledyne_DVL_Tasman_600kHz

Nortek_DVL1000-4000m

Add Altimeter

Add AcousticComms

Add Radio

Sonar array

☒ Include in vessel

Standoff distance (m)

7500

Wavelength (m)

0.375

Hydrophone diameter (m)

0.0075

Array length (m): 27.9375

Num. hydrophones: 150.0

Array drag at mission speed (N): 4.694498733097613

☐ Array inside fairing

INS_0

Mfr.: Ixblue

Part: C3

Mass (kg): 1.6

Power (W): 12.0

Duty cycle (%)

100

PV

PV0

Remove from design

GPS_0

Mfr.: Garmin

Part: GPS 15H

Mass (kg): 0.015

Power (W): 0.396

Duty cycle (%)

1

PV

PV0

Remove from design

Radio_0

Mfr.: Iridium

Part: Core 9523

Mass (kg): 0.032

Power (W): 2.3

Duty cycle (%)

1

PV

PV0

Remove from design

CTD_0

Mfr.: SeaBird

Part: SBE 61 Deep Argo CTD

Mass (kg): 2.8

Displacement (m3): 0.0042712

Net buoyancy (kg): 1.5865

Pressure rating: 7000.0

Power (W): 0.24

Duty cycle (%)

100

Remove from design

Altimeter_0

Mfr.: KONGSBERG

Part: 975-71010000

Mass (kg): 2.4

Displacement (m3): 0.0011982

Net buoyancy (kg): -1.1695

Pressure rating: 3000.0

Power (W): 6.25

Duty cycle (%)

100

Remove from design

AcousticComms_0

Mfr.: Benthos

Part: ATM-966

Mass (kg): 7.3

Displacement (m3): 0.0047299

Net buoyancy (kg): -2.4424

Pressure rating: 6000.0

Power (W): 0.55

Duty cycle (%)

100

Remove from design

OBS

Mass (kg): 345.4

Displacement (m3): 0.12292

Net buoyancy (kg): -219.16

Figure 34 Payload work area.

Propulsion work area

The propulsion work area, shown in Figure 35, offers a simple interface to the AI propeller designer. Drag force and required power at mission speed (computed from current vessel properties) are shown, with a button to request options from the AI propeller designer, and the resulting design catalog is shown in a table – expanding a row of the table shows a rendering of the propeller.




Figure 35 Propulsion work area.

Energy and battery work area

The energy and battery work area, shown in Figure 36, summarizes the computed energy needs (from propulsion power and hotel load) and allows the user to choose a battery type and number of packs into which batteries should be divided. Both hotel and locomotion energy can be manually overridden, which has been done in this case to add a 20% safety margin to the computed energy needs). Each battery pack has a corresponding card (shown on the right-hand side) that includes basic information and a drop-down menu to assign the battery to a pressure vessel (this drop-down will simply show 'not assigned' if pressure vessels have yet to be added to the design). In this example, the 'Batteries' indicator is yellow, indicating that there are sufficient batteries, but not all have been assigned to pressure vessels. Removing a battery from a pressure vessel also upsets the balance and buoyancy of the design, turning the 'Packing' indicator red.

In this design	
Battery pack 0	Battery pack 1
Battery mass kg: 209	Battery mass kg: 209
Battery volume m3: 0.143	Battery volume m3: 0.143
Battery type: Lithium_Thionyl_Chloride_D_Cells	Battery type: Lithium_Thionyl_Chloride_D_Cells
PV PV0	PV Not assigned

Figure 36 Energy and battery work area.

Pressure vessel work area

In the pressure vessel work area, Figure 37), the layout is like that of the propulsion work area. In the top left is an interface to request designs from the AI PV designer, to the right an array of cards, one for each PV included in the design (the header above these cards shows the volume enclosed by all PVs vs the volume required to enclose all dry payload components and batteries), and a table showing the catalog of available designs created by the AI Designer, together with controls to filter the catalog by volume enclosed and operating depth. The 'PV designer' card allows users to size a PV manually, or as shown here, to specify the PV volume by choosing from a table of battery packs and dry components we wish to package in the PV. This total required volume and the maximum mission depth can be sent to the AI PV designer, which then populates the catalog with suitable options.

Fairing

Payload

Propulsion

Batteries

PV volume

Mass shifter

Foam

Packing

Fairing

Payload modules

Propulsion

Energy and Battery

Pressure Vessels

Pitch control

Foam

Packing

PV designer

☒ Use volume of selected opts

name	volume
Battery0	0.143068
Battery1	0.143068
INS_0	0.0004
GPS_0	0.000014
Radio_0	0.000037

Volume margin

1

PV Volume (m3)

0.144

Request AI designs

In this design (0.2959 of 0.2866 m3 required)

PV0

Mass (kg): 39.90

Displacement (m3): 0.21

Interior volume (m3): 0.148

Depth rating (m): 1300

Remove

PV1

Mass (kg): 39.90

Displacement (m3): 0.21

Interior volume (m3): 0.148

Depth rating (m): 1300

Remove

PV design catalog

stiffener_separation	displacement	cyl_length	min_pcr	num_stiffeners	depth_rating	mass	stiffener_type	name	inner_volume	outer_diam	num_in_design		
0.086406	0.225047	1.0	17,691,903.075508	10	1,300.0	58.587334	External	PV1872	0.163846	0.422831	0	+	—
0.081861	0.232237	1.0	17,757,619.475648	10	1,300.0	56.021166	External	PV1873	0.163846	0.422831	0	+	—
0.091962	0.219156	1.0	17,631,812.009304	10	1,300.0	60.718142	External	PV1895	0.163846	0.422831	0	+	—
0.092936	0.287965	1.0	18,596,894.577435	10	1,300.0	71.442414	External	PV2040	0.197845	0.462686	0	+	—
0.095651	0.22	1.0	18,075,704.185756	10	1,300.0	71.099107	Internal	PV2041	0.151522	0.462686	0	+	—
0.090095	0.22	1.0	18,160,359.582259	10	1,300.0	67.647986	Internal	PV2042	0.145931	0.462686	0	+	—
0.091064	0.22	1.0	17,958,924.307824	10	1,300.0	68.294138	Internal	PV2043	0.145931	0.462686	0	+	—
0.092936	0.22	1.0	19,396,243.917676	10	1,300.0	71.003261	Internal	PV2044	0.139561	0.462686	0	+	—
0.090095	0.283861	1.0	18,160,359.582259	10	1,300.0	68.420672	External	PV2045	0.200122	0.462686	0	+	—
0.091064	0.2825	1.0	17,958,924.307824	10	1,300.0	68.978116	External	PV2046	0.199361	0.462686	0	+	—
0.092936	0.289341	1.0	19,396,243.917676	10	1,300.0	71.615155	External	PV2047	0.197845	0.462686	0	+	—
0.095651	0.22	1.0	18,063,700.070895	10	1,300.0	71.438886	Internal	PV2048	0.150581	0.462686	0	+	—
0.090095	0.22	1.0	18,150,336.512711	10	1,300.0	67.880075	Internal	PV2049	0.145011	0.462686	0	+	—
0.086406	0.226236	1.0	17,679,187.812199	10	1,300.0	58.916352	External	PV1896	0.163846	0.422831	0	+	—
0.091064	0.22	1.0	18,783,239.921692	10	1,300.0	68.501139	Internal	PV2050	0.145011	0.462686	0	+	—
0.090095	0.285225	1.0	18,150,336.512711	10	1,300.0	68.671411	External	PV2051	0.200122	0.462686	0	+	—
0.091064	0.283861	1.0	18,783,239.921692	10	1,300.0	69.201554	External	PV2052	0.199361	0.462686	0	+	—
0.098178	0.26	1.0	17,714,043.397441	10	1,300.0	80.302451	Internal	PV2053	0.174577	0.49847	0	+	—
0.08916	0.26	1.0	18,647,960.999905	10	1,300.0	80.964542	Internal	PV2054	0.174577	0.49847	0	+	—
0.100121	0.26	1.0	17,620,509.308549	10	1,300.0	81.729733	Internal	PV2055	0.174577	0.49847	0	+	—

Operating Depth (m):

1200.0 m

1950.0 m

Inner Volume: 0.12876 m3

0.28659 m3

First

Prev

1

2

3

4

5

Next

Last

Figure 37 Pressure vessel work area.

Approved for Public Release; Distribution Unlimited.

38

Pitch control work area

The pitch control work area, as seen in Figure 38, includes a simple calculator on the left-hand side and an array of cards describing the shifters included in the current design. The calculator allows the user to design a simple pitch shifter composed of a tube to enclose a mass. The user can specify material type for both the tube and the mass, and to specify the mass value for the shift mass (the dimensions of the shift mass are computed automatically from the inner radius of the tube, the shift mass density, and its mass).

The interface shows the 'Pitch control' tab selected. The 'Mass shifter design' form on the left includes the following fields:

- Tube material: Aluminum 6061
- Density (kg/m3): 2700.0
- Shifter tube length (m): 2
- Shifter tube outer radius (m): 0.076
- Shifter tube wall thickness (m): 0.005
- Shift mass material: Lead
- Density (kg/m3): 11000.0
- Shift mass (kg): 59.2285

At the bottom of the form is a button labeled 'Add shifter to UVV'. To the right, the 'In this design' section lists two shifters:

- MassShifter0**
 - Total mass (kg): 66.20
 - Total displacement (m3): 0.00797
 - Movable mass (kg): 59.2
 - Length (m): 1.2
 - Radius (m): 0.076
 - Remove button
- MassShifter1**
 - Total mass (kg): 66.20
 - Total displacement (m3): 0.00797
 - Movable mass (kg): 59.2
 - Length (m): 1.2
 - Radius (m): 0.076
 - Remove button

Figure 38 Pitch control work area.

Foam work area

The foam work area, illustrated in Figure 39, allows the human designer to create syntactic foam pieces sufficient to balance the buoyancy of the vessel, and as components to be used to balance the UUV during the packing process. As with other work areas, the left-hand side has a simple interface to design foam components – either cuboid or cylindrical – with the foam material chosen according to mission depth. Two indicators help the user understand how much foam is required. In the calculator ‘Foam volume for neutral buoyancy’ indicates how much more foam is required for the displacement of all wet components to balance the mass of all components at the mission-specified water density. In the ‘Summary’ section of the UI (shown in the inset) ‘net_buoyancy_margin’ shows, as a percentage of vehicle mass, the current discrepancy between mass and buoyancy. If this fraction exceeds a value assigned as part of the mission specification, then the foam indicator will turn red.

The screenshot displays the 'Foam designer' interface within a software application. The left sidebar contains a vertical menu with options: Fairing, Payload, Propulsion, Batteries, PV volume, Mass shifter, Foam, and Packing. The 'Foam' option is highlighted with a green bar. The main area is titled 'Foam designer' and includes a sub-header 'Cuboid or cylindrical foam'. Below this, it states 'Foam material chosen according to mission depth' and 'Foam type: BZ-24'. The 'Foam density kg/m3' is set to 385.0. Input fields for 'Length (m)' (0.5), 'Width/diameter (m)' (0.5), and 'Height (m) (if 0, foam will be a cylinder)' (0) are present. Calculated values show 'Foam displacement: 0.09817', 'Net buoyancy (kg): 63.028', and 'Foam volume for neutral buoyancy (m3): 0.004806644185954556'. An 'Add foam' button is at the bottom of this section.

The 'In this design' section lists four foam components (foam0, foam1, foam3, foam4) with their respective properties:

Component	Density (kg/m3)	Displacement (m3)	Net buoyancy (kg)	Diameter (m)	Length (m)	Depth rating (m)
foam0	385.0	0.17279	110.9296365982556	1.0	0.22	2000.0
foam1	385.0	0.17279	110.9296365982556	1.0	0.22	2000.0
foam3	385.0	0.098175	63.028202612645224	0.5	0.5	2000.0
foam4	385.0	0.031416	20.169024836046475	0.4	0.25	2000.0

Each component has a 'Remove' button below its details.

The 'Summary' inset on the right provides a high-level overview of system metrics:

Metric	Value
drag at mission speed N	61.110
propulsion_efficiency	0.801
hotel_load_W	23.000
total_energy_J	891027272.727
num_pvs	2.000
total_mass_kg	1350.307
net_buoyancy_margin	-0.002
Displacement (m3)	0.031416
Net buoyancy (kg)	20.169024836046475
Diameter (m)	0.4
Length (m)	0.25
Depth rating (m)	2000.0

Figure 39 Foam work area.

Packing work area

In the packing and fairing modification work area, illustrated by Figure 40, parts to pack are tabulated on the left, where their positions in the vessel can be updated. Longitudinal (top and side) views, as well as a cross-sectional view of the packed contents and vessel fairing serve as a visual reference. Underneath these views are an indicator showing the pitch range available with the current packing, together with plots showing relationship between the vessel's center of buoyancy and center of gravity in the up-down and left-right directions. Should the co-designer wish to modify the shape of the fairing to either increased space for packing, or to eliminate empty volume in the fairing, this is enabled by controls surrounding the three views of the vessel, that control axis-specific values of the Myring-parameterized fairing shape. After modifying the shape of the fairing, the user can summon the CFD surrogate model for rapid evaluation of drag properties of the new shape and can then commit the new fairing to the design (and add it to the design catalog) if desired.




Figure 40 The packing and fairing modification work area.

3.6.2.2 Using the UI during the Fall 2022 hackathon

During the Fall 2022 hackathon, the SRI TA1 team used the UI (with our assistance) to develop their submitted UUV designs. Although these findings are informal observations, we found that:

- The UI allowed for designs to be completed rapidly (with some delays for the non-real-time AI components to generate design catalogs).

- The organization of the UI to start from energy estimation allowed for many quick ‘what-if’ experiments regarding potential mission speeds and fairing designs.
- Having manual overrides for the energy estimates was useful when we needed to add a safety margin for battery capacity (although a dedicated energy margin entry would have made this process more transparent).
- The subsystem status indicators were informative, allowing us to see how changes to one subsystem would propagate through to others.
- The UI served as a central focus during group design discussions; separating the different subsystems into their own separate work areas helped keep group focus, while the status indicators made clear to all participants what else needs attention once the current subsystem activities are finished.

3.6.3 Unmanned Underwater Vehicle mission review.

The STR UUV design problem includes a stochastic mission simulator (ocean currents and vehicle localization system behavior is subject to random perturbances) that generates two sets of outputs; the first is a summary of a particular vessel’s performance on the mission, and the second is a detailed set of tracking data including a time-history of vehicle performance over a given trial. The track data includes vessel location, estimated location, energy usage, speed, and other time-varying measurements, measured at tens of thousands of time intervals along the vehicle’s course.

The summary output UI, shown in Figure 41, includes data such as pass/fail metrics for mission sub-goals (in this case, successful placement of a payload and return to base within given time frames and spatial boundaries), together with more detailed statistics measuring pertinent variables with respect to the mission goals. The overall success rate of a given design is shown above the summary data table for each run of the simulator. Binary pass (green) and fail (red) measures are indicated by colored squares alongside goal-relevant numerical data, where metrics responsible for failures are color coded red. Data can be sorted in any column. Checkboxes below the table allow users to hide success or failure runs from the table for more focused browsing of good and bad runs.

Success rate: 12.60%

#	Run ID	Success	Sufficient energy	Payload delivered	Returns to start	On time	Drop distance (m)	Return distance (m)	Elapsed time (days)	Energy remaining (m)	Final depth (m)
0	0	■	■	■	■	■	28355.11	44871.91	60d, 0h 0.22m	230m	201.93
1	1	■	■	■	■	■	33408.99	63189.07	60d, 0h 0.57m	97m	201.96
2	10	■	■	■	■	■	26048.92	63944.76	60d, 0h 0.63m	242m	201.96
3	100	■	■	■	■	■	31661.18	40918.18	60d, 0h 0.99m	245m	201.94
4	101	■	■	■	■	■	32995.30	49541.10	60d, 0h 0.03m	229m	201.94
5	102	■	■	■	■	■	25854.58	60134.68	60d, 0h 0.08m	97m	201.96
6	103	■	■	■	■	■	29263.71	94804.86	60d, 0h 0.59m	223m	201.93
7	104	■	■	■	■	■	41789.08	44567.70	60d, 0h 0.51m	257m	201.95
8	105	■	■	■	■	■	37649.45	60925.94	60d, 0h 0.47m	243m	201.94
9	106	■	■	■	■	■	24132.35	46032.62	60d, 0h 0.63m	245m	201.95
10	107	■	■	■	■	■	23355.28	81722.05	60d, 0h 0.95m	225m	201.94
11	108	■	■	■	■	■	6350.79	44050.37	52d, 8h 15.17m	690m	73.98
12	109	■	■	■	■	■	26123.62	79767.79	60d, 0h 0.99m	225m	201.95
13	11	■	■	■	■	■	23793.14	50315.48	60d, 0h 0.39m	254m	201.95
14	110	■	■	■	■	■	30022.31	56969.44	60d, 0h 0.28m	262m	201.94

Show in table: ☒ Successes ☒ Failures

Figure 41 Performance review user interface view.

In addition to developing a physical design for a vessel, TA1 teams were tasked with developing control strategies for these vessels that would be tested in the simulation; a single design could be run through the simulator many times to generate performance statistics over the range of random variation provided in the simulator. To allow teams to review the overall performance of a vessel (physical vessel and control algorithm) in the simulator, and to debug and refine control strategies, we developed a UI consisting of three distinct views:

1. Performance review page allows the designer to quickly understand the performance of the design over many runs.
2. Single run review allows for detailed inspection of a selected simulation run which helps to better understand features of that run, e.g., to diagnose faults that lead to mission failure, or to find out why a particular run may produce particularly good results.
3. Controller editor: a UI that allows designers to edit the control algorithm code (either to adjust parameters, or even re-write the algorithm itself), re-run the simulator and inspect the simulation outputs for the new control method.

The following paragraphs describe usage of these three views to provide an example of error diagnosis and controller refinement, illustrating the workflow of the UI.

The performance review page, Figure 41, is designed to provide designers with insight into the overall performance of their design in the simulator; color coding is used for quick visual understanding of success and failure cases, with separate columns for each binary pass/fail condition, and the tabulated data can be sorted along any column of the table, with color-coding in red indicating when a goal-relevant value is in violation of mission requirements. There are UI controls to allow designers to focus on only failure cases or only successes, and at the top of the page a success rate that summarizes the fraction of simulation runs that the current design completed with all goals satisfied. In this example, the table includes 1,000 runs generated by an early version of a control algorithm designed by our TA1 colleagues at SRI (their final control

algorithm had a success rate more than 99.9% over 10,000 runs). To take a closer look at any run in the table, the designer simply clicks the corresponding row of the table, and a single run review UI is generated for that run (with sufficient screen space, or within the less limited confines of a VR environment, many such runs can be opened at once for side-by-side comparison of several runs).

As shown in Figure 42, the single run review UI gives users an in-depth look at the performance of the vessel and controller in each run of the simulation. The left panel of the UI contains a scalable, rotating 3D plot of the estimated UUV trajectory (red line), the true vehicle trajectory (multi-colored line with circular markers), the sea floor (brown surface), and average water currents (blue cones). The true trajectory can be colored by any of the available track measurements, such as remaining energy, vehicle speed, localization error and so on. The size of the markers along the true trajectory are used to indicate the value of a second variable (which can be chosen in the same way as the trajectory color scheme). In this case, the color of the trajectory indicates UUV depth, and the radius of the markers is reflective of vehicle speed. The right-hand side of the UI includes a set of card widgets that detail individual pass/fail metrics for a run followed by plot controls that allow users to: select which variables influence trajectory color and marker size; turn on and off the sea floor and water current displays; restrict the time range of data displayed, included shortcuts to show only the outbound journey or the return journey. Below the plot controls is an overhead (zoomable, panning) 2D plot of the same data, colorized to match that of the main 3D display, and that matches the time window of the main 3D plot. A button at the bottom of the view allows users to open an editor UI to modify the control algorithm.




Figure 42 Single run review user interface view.

In the example shown in Figure 42, we have used the zoom and pan controls to show the final weeks of a failed mission where the UUV ran aground in the final stages of returning to its starting point. The true trajectory is shown in purple (this color indicates vehicle depth), and marker radius indicates relative speed. The red line shows the estimated vehicle trajectory. We see that the vehicle gets a position fix (indicated by the estimated trajectory snapping to the true trajectory) and the vehicle increases speed (indicated by the marker radius increasing) dramatically, and heads quickly into shallow waters, eventually running aground short of its goal, illustrated by the zoomed-in view of Figure 43.




Figure 43 Zoomed view of the unmanned underwater vehicle running aground.

This particular mission is set in a location where the seabed depth increases from the start position along latitude. TA-1 devised a new strategy for the final homing behavior of the vessel; close distance on the home position first in longitude (maintaining a safer depth) and then closing the final distance in latitude at a low speed. Such a strategy can be implemented and tested in the control editor UI, shown in Figure 44, summoned by clicking the button beneath the plot area of the single run review UI.

The control editor UI, illustrated in Figure 44, allows the designer to modify the controller algorithm code and start a new run of the simulator from within the UI. Two editor panes show the control algorithm code and a separate file containing the algorithm's tunable parameters. When users are satisfied with their edits, they can run the simulation with the new controller from the UI, where text outputs from the simulator are shown in a console window. Once the simulation is finished, the "View trajectory" button on the UI becomes active, allowing users to open a new single run review UI for the freshly generated simulator outputs. In our example scenario, we have modified the controller code to implement the new homing strategy, and the simulator outcome is successful. We can bring up the successful run trajectory and compare side-by-side with the previously failed attempt shown earlier.




Figure 44 Control editor user interface view.

The plots shown in Figure 45 provide detailed views of the overhead 2D trajectory plot for a failed run (left) and a run that succeeds after modification of the control strategy (right). In both cases, the plots are color-coded for speed – yellow is fastest, orange slower and purple slowest. The original control strategy heads South-West at high speed, eventually running aground short of the mission goal. The new strategy, which attempts to close distance at speed in regions of constant depth (lines of latitude) and with more cautious excursion into shallower water (in the North-to-South direction) is seen clearly in the overhead view on the right of Figure 45, and in this case leads to successful complement of the mission.




Figure 45 Comparison of an unsuccessful control strategy (left) versus a successful control strategy (right).

4 Results and Discussion

4.1 Proxy car design problem: Human-AI experiments with reinforcement learning

The Human-AI symbiotic design system involves humans preparing data, evaluating the model performance, and refining the model output. However, there is a lack of systematic study to understand the human experience, needs, and expectations when humans and AI work symbiotically on a design task.

This Human-AI experiment aims to investigate the human experience and performance when working with and without AI. By asking the user to explore a design space of a 2D model car, with and without a reinforcement learning-based AI model, we discovered that human designers show common interests in specific design variables when exploring the design space. The frequency of their attempts to explore some design variables is higher than other parameters. We also showed that the users benefit from working with AI Designers by making more exploration attempts.

4.1.1 Human-AI interaction system

We designed a proxy of the Human-AI symbiotic design system to study the Human-AI interaction. The proxy system targets a design space exploration problem. The human designer uses the system to find a set of design variables that makes the 2D car travel the longest distance on rocky terrain, as shown in Figure 46. The system consists of a web-based interface for interaction, a 2D physics simulator, and a reinforcement learning model. Figure 47 shows the interface of the system.




Figure 46 Proxy 2D car climbs across obstacles on the terrain.



Figure 47 User interface for the proxy 2D car design.

4.1.2 User study design

The users were divided into two groups. One group had access to the AI Designers, and the other group did not. The users that had access to the AI Designer were able to use the sliders to change the design variables of the AI Designer, as shown in Figure 47 (1). The other users can only use the sliders to control the design variables, as shown in Figure 47 (2). Both groups of users can see the run-time simulation result to evaluate their decision. Figure 47 (3), (4), (5) and (6) show various plots available to users, plotting pairs of variables against each other, or bar charts showing values of selected design variables for each design explored. Figure 47 (3) illustrates how clicking on a plot element will bring up a graphical review pane that shows a video of the corresponding design's performance on the course.

Six design variables are included in this problem, including *wheel size*, *wheel speed*, *wheel friction*, *chassis height*, *chassis width*, and *car height*. These design variables are initially set up with upper and lower limits, the same as real design tasks. The objective was to find the global optimal design variables for the travel distance. In solving, the users had three challenges to deal with. First, the design space was non-linear with multiple local minimums. Second, constraints between design variables had to be considered. Lastly, the problem had multiple objectives, including maximizing the travel distance and using fewer materials to achieve the goal.

4.1.3 Results and Discussion

The number of interactions over design variables, including wheel size, motor speed, friction, chassis height, width, and car height, was collected during studies. Users show different interests in the design parameters, as shown in the pie chart of Figure 48. Users conducted the most parameter changes on the wheel size, 28%, and the least parameter changes on the chassis height, 7%. The experiments show that humans fixate on a few parameters which they think are important. The AI Designer might help humans open to more combinations of parameters. On the other hand, humans provide constrained knowledge that can lead to faster human-compatible exploration.




Figure 48 The total number of changes on each design variables.

Interactions of each user across groups with AI co-designer and without AI co-designer are different, as shown in Figure 49. The user attempts more parametric cases with AI than without AI providing optimized design. It is potentially because, with AI support, human designers spend less time analyzing the simulation result and could focus more on exploring the design space.




Figure 49 Interactions for users without and users with access to AI Designers.

4.2 Experiments comparing part assembly using traditional 2D interfaces vs the 3D model assembler.

The goal of these experiments was to investigate interactions with virtual objects in a seated position through different mediums. By allowing users to assemble a cube made of puzzle pieces using traditional 2D interfaces, as well as 3D interfaces such as VR, we investigated the benefits of completing CAD assemblies through the help of mixed reality.

The first experiment involved assembling a cube made of puzzle pieces using a mouse and keyboard. Blocks were manipulated using 3D gizmos that allow for precise adjustments through translation and rotation. Figure 50 depicts three stages of the assembly process. Since the ability of the user to solve this puzzle was not being tested, but merely the interface through which they solved it, they were given a set of step-by-step instructions on how to solve the puzzle. The number and type of keypresses the user used to achieve this task were recorded, as well as the time taken to complete the task.




Figure 50 Assembly process using traditional 2D interface with the help of Gizmos.

The second experiment, illustrated in Figure 51, involved assembling the same puzzle, but in Virtual Reality using an Oculus Quest 2 headset. The 3D interface consisted of 3 main parts; a large depiction of the blocks in front of the user, a smaller design table upon which the user can move and manipulate the blocks using their hands, and a multi-manipulation tool on the side. The users were provided easy access to the instructions. The focus of this experiment was to investigate the impact of advantageous perspective in 3D space. The multi-manipulation tool allowed for the rotation and change of scale of the design-table. The user was able to interact with multiple blocks at the same time in the design stage. The left side of Figure 51 depicts a user interacting with multiple blocks on the design table, while the right side of Figure 51 depicts a user rotating their perspective with the help of the multi-manipulation tool to assemble the puzzle easier. With the experiments conducted thus far, users have greatly preferred performing tasks using the 3D interface, with comments being made about how a combination of both would serve to combine the best of perspective and fine control.




Figure 51 Assembly process using 3D VR interface assisted with various manipulation tools.

5 Conclusion

The premature end of the program curtailed much of our experimental work, but that conducted to date, and our experience of using our UIs in the hackathons allow for some tentative conclusions. Our 2D car experiments (detailed in section 4) indicate that there can be some benefit to a collaborative AI co-designer: when users can generate designs by constraining an AI design agent, they tend to perform more experiments and explore deeper into the design space than when making fully manual designs. Extending design space exploration could be a pathway to more novel designs in the future.

From our hackathon experience, we observed that when multiple independent AI co-designers are involved (e.g. for the UUV design problem, where there were specialized AIs for various UUV sub-systems), having a UI that focuses on filtering and selection of AI-generated component designs one subsystem at a time, while having always-visible status indicators for all subsystems, helped the team focus on subproblems while alerting them when design choices for the current sub-system of interest invalidate choices made for other subsystems. In addition, having manual overrides for certain key design parameters allowed for quick posing of ‘what-if?’ questions, which were valuable in quickly establishing new design needs when the towed sonar array was added to the mission.

Further formal experimentation is required to make more concrete conclusions, including measurement of not just design performance, but also human factors.

One final observation is that the current state of the Human-AI interaction is not that sophisticated. The sandboxing of the problem into TA1/TA2, with hackathon metrics largely based upon design performance did not encourage deep integration of AI and UI components; the AIs were able to generate large numbers of designs but take very limited (if any) feedback from the human designer, who effectively became a shopper, with TA2 UIs becoming searchable design catalogs.

Appendix A. Physics-driven Numerical Simulation using Neural Networks for Fast Design Space Exploration

List of Figures

Figure A-1 Hull shape optimization design space.	55
Figure A-2 Structural design space.	55
Figure A-3 Physics-driven neural network system structure.....	56
Figure A-4 RitzNet for a 2D structure and comparison with FEA results.....	57
Figure A-5 DuNN for a L-shape structure with stress singularity.	57
Figure A-6 LSNN for SHCL with a non-convex flux $f(u)=u^3/3$	57
Figure A-7 Parametric design space exploration using RitzNet.....	57

The goal of this work is to address the problem of stress simulation and hydrodynamic simulation using Physics-driven neural network methods. Traditionally, these problems are solved using the finite element, the finite difference, or the finite volume methods.

To support a large design space exploration, a numerical method should be able to conduct a high-fidelity continuous simulation with varying design parameters. E.g., In a hull shape optimization, once an initial concept is picked, a corresponding parametric design space is created, shown in Figure A-1. Similarly, for a structural problem, once a design space is created for an initial structure, we want to perform high-fidelity structural analysis in the high dimensional design space, as shown in Figure A-2, to support design optimization. Traditional numerical schemes rely on a fixed mesh or grid which is hard to adapt to a varying domain prescribed by changing design variables, due to the conformity requirement of a mesh.




Figure A-1 Hull shape optimization design space.




Figure A-2 Structural design space.

We designed several physics-driven numerical methods for supporting structure analysis and hydrodynamic problems. For the latter, we studied only the scalar hyperbolic conservation law problem as a preliminary investigation. The system structure is shown in Figure A-3. For the structure problem, we devised two different loss functions using different differential operators to solve for structural displacement (RitzNet) or structural stress dual neural network (DuNN) for the independent variables. For the scalar version of hydrodynamic problems (scalar hyperbolic conservation law), we used a least square formulation combined with a discrete divergence operator for the flux. The three network losses functions are listed below.




Figure A-3 Physics-driven neural network system structure.

(I) RitzNet Loss:

$$J(v) = \frac{1}{2} \left\{ \int_{\Omega} \sigma(v) : \varepsilon(v) dx + \mu \|v - g_D\|_{1/2, \Gamma_D}^2 \right\} - (f, v) - \int_{\Gamma_N} g_N \cdot v ds$$

Where g_D and g_N are boundary conditions for displacement and traction forces, and f is the body force.

(II) DuNN loss:

$$J(v) = \frac{1}{2} \|\mathcal{C}^{1/2} \varepsilon(v)\|_{0, \Omega}^2 - (f, v) - \int_{\Gamma_N} g_N \cdot v ds.$$

(III) Least-squares neural network (LSNN) loss for non-linear scalar hyperbolic conservation law (SHCL):

$$\mathcal{L}(v; g) = \|\text{div } f(v)\|_{0, \Omega}^2 + \|v - g\|_{0, \Gamma_-}^2.$$

Where $f(v)$ is the spatial and temporal flux, and g is the inflow boundary condition.

Some of the numerical results using the above three physics-driven methods are listed here. Figure A-4 shows a 2D membrane under tension and its structural response simulated using RitzNet. Due to the stress concentration around the hole, RitzNet results are comparable to high-order adaptive p-element finite element analysis (FEA) methods. Figure A-5 shows the results for an L-shape plate with re-entrant corner singularities. DuNN can handle singular problems and incompressible materials well. Figure A-6 lists preliminary results of using LSNN and the discrete divergence operator to solve non-linear SHCLs and proves its efficacy in simulating shocks, rarefaction wave and compound wave. Our future plan is to extend the method to hydrodynamic systems for help solving the design space exploration problem posed on Figure A-1. Our initial design space exploration problem was conducted using RitzNet with only one design variable. The result shows that the physics-driven neural network (NN) method supports a parametric simulation problem more efficiently, compared with the traditional mesh-based methods.




Figure A-4 RitzNet for a 2D structure and comparison with FEA results.




Figure A-5 DuNN for a L-shape structure with stress singularity.




Figure A-6 LSNN for SHCL with a non-convex flux $f(u)=u^3/3$.




Figure A-7 Parametric design space exploration using RitzNet.

Appendix B. List of publications (in preparation or submitted)

1. S. Chidambaram, P. Raja, K. Ramani, Interactive Assembly in VR through asymmetric interactions along with a WIMP Interface Desktop, In Preparation, to be submitted to the Computer Aided Design Journal.
2. M. Liu, Z. Cai, and K. Ramani, *Deep Ritz Neural Network for Linear Elasticity with Adaptive Quadrature*, submitted to the Journal of Computer Methods in Applied Mechanics and Engineering.
3. M. Liu, Z. Cai, and K. Ramani, *Dual Neural Networks (DuNN) Method for Elliptic PDEs*, in preparation for the Journal of Computer Methods in Applied Mechanics and Engineering.
4. R. Duan, J. Shi, and K. Ramani, *Symbiotic Explorer: Human-guided Reinforcement Learning for Engineering Design Space Exploration*, In Preparation, to be submitted to *The ACM CHI Conference on Human Factors in Computing Systems*.

List of Symbols, Abbreviations, and Acronyms

2D	Two-Dimensional
3D	Three-Dimensional
ACCEL	Human-AI Collaborative Cyber-Physical Engineering and Learning
AI	Artificial Intelligence
API	Application Programming Interface
AR/VR	Augmented Reality/Virtual Reality
CAD	Computer-Aided Design
CFD	Computational Fluid Dynamics
CPS	Cyber-Physical Systems
CTD	Conductivity, Temperature and Depth
DARPA	Defense Advanced Research Projects Agency
DuNN	Dual Neural Networks
DVL	Doppler Velocity Log
FDM	Flight Dynamics Model
FEA	Finite Element Analysis
H-M	Human-Machine
ID	Identification
LQR	Linear Quadratic Regulator
LSNN	Least-Squares Neural Network
NN	Neural Network
PV	Pressure Vessel
REST	Representational State Transfer
SDCPS	Symbiotic Design of Cyber-Physical Systems
SHCL	Scalar Hyperbolic Conservation Law
Siemens	Siemens Corporation
SME	Subject Matter Expert
SRI	SRI International
STL	Stereo Lithography (3D File Format)
STR	Systems & Technology Research LLC

SWRI	Southwest Research Institute
TA	Technical Area
TA1	Technical Area 1 – AI Designer
TA2	Technical Area 2 – Human-AI Interfaces
TA3	Technical Area 3 – Design Corpus, Challenges and Assessment
tSNE	t-Distributed Stochastic Neighbor Embedding
UAM	Unmanned Air Mobility
UAV	Unmanned Aerial Vehicle
UCB	University of California, Berkeley
UI	User Interface
UUV	Unmanned Underwater Vehicle
VAE	Variational Autoencoder
VR	Virtual Reality