

Artificial Intelligence Exploration (AIE) Opportunity
DARPA-PA-19-03-05
Gamebreaker

I. Opportunity Description

The Defense Advanced Research Projects Agency (DARPA) is issuing a new Artificial Intelligence Exploration (AIE) Opportunity entitled, *Gamebreaker*, which invites submissions of innovative basic or applied research concepts in the technical domain of automating game balance to explore new capabilities/tactics/rule modifications that are most destabilizing to the game or simulation.

This *Gamebreaker* AIE Opportunity is being issued under the DARPA Program Announcement for AIE, DARPA-PA-19-03, which details the AIE Program's overall intent and provides evaluation and submission instructions in addition to those outlined in this solicitation. To view the original DARPA Program Announcement for AIE, visit beta.SAM.gov (formerly FedBizOpps) under solicitation number DARPA-PA-19-03.

All proposals in response to the technical area(s) described herein will be submitted to the *Gamebreaker* AIE Opportunity, solicitation number DARPA-PA-19-03-05, and if selected, will result in an award of an Other Transaction (OT) for Prototype Projects, not to exceed \$1,000,000. This total award value includes Government funding and performer cost share, if required, or if proposed.

A. Introduction

The *Gamebreaker* AIE Opportunity seeks to develop and apply Artificial Intelligence (AI) to existing open-world video games to quantitatively assess game balance, identify parameters that significantly contribute to balance, and explore new capabilities/tactics/rule modifications that are most destabilizing to the game.

For *Gamebreaker*, game balance is defined as an inherent property of the game that reflects the win/loss ratio of players of equal skill level based on strategies and tactics employed within the game. For example, in a balanced game, if the skill level of both players is equal, each player will win approximately 50 percent of the total games played. Similarly, if the skill level of both players is equal but one player wins disproportionately due to inherent advantages arising from a condition of the game construct, the game is unbalanced.

The commercial gaming industry has a long-standing interest in maintaining game balance as balanced games are typically more entertaining, and market pressures help drive their development. Moreover, the contemporary method for assessing and balancing games is a trial-and-error approach, thus representing an opportunity for the application of AI.

Normally, game developers release and observe an initial configuration of the game in large-scale play. Then, developers gather high-level win/loss statistics while players provide feedback about elements of the game that are overpowered or imbalanced. Finally, updates to the game are made in which elements are buffed (performance increases) or nerfed (performance decreases) to achieve game balance. To date, little quantitative modeling of game balance exists, and research on the application of AI algorithms to automating game balance assessment (formally referred to

as quantitative balance analysis) is extremely limited.

Department of Defense (DoD) applications of automating quantitative balance assessment are plentiful and range from identifying and mitigating adversary capabilities to generating methods for diagnostically assessing the impact of new defense technologies or changes in force design/posture. In future conflicts, DoD investment is designed to maximize imbalance to create an advantage or to seek equilibrium when an adversary is seeking an advantage. New AI algorithms inspired by *Gamebreaker* could help develop winning warfighting strategies when the adversary's objectives – i.e. the “rules of the game” – are not clearly known. By exploiting game balance, *Gamebreaker* addresses an existing gap in AI and data analytics research as applied to current wargaming and simulation.

Complex, multi-domain modeling and simulation (M&S) environments currently under development by several DARPA programs aim to create a useful “[Mosaic](#)” model within which to experiment on new warfighting constructs using distributed, adaptive, all-domain force composition, tactics, and strategies – yet these models do not currently exist. It is reasonable to assume, however, that once these simulation environments reach maturity, an “AlphaMosaic” equivalent will be capable of searching for optimal strategies and tactics in the same way [AlphaGo](#) and [AlphaStar](#) agents have already proven effective in exploiting their respective game environments. Yet this does not take into account modifications to the construct of the game itself. While AI techniques have demonstrated the ability to master games and models of increasing complexity, there has been almost no research in the application of AI to game modification.

Future Mosaic models and wargames will likely include many of the same attributes of current, complex, real-time strategy (RTS) video games. As a result, contemporary RTS video games serve as surrogate engagement models, since they simulate force-on-force engagements with heterogeneous Red/Blue platform and weapon systems, while accounting for geography and imperfect information.

B. Objective/Scope

Proposers should select an existing game representative of anticipated Mosaic models while also remaining tractable. *Gamebreaker* proposers should use AI algorithms to automate the assessment of game balance, identify parameters that significantly contribute to game balance, and evaluate the effect of perturbations such as the introduction of a new weapon, capability, or rule change on the balance of the game. The resulting methodology would allow for the derivation of an effective game balance state equation. In order to accomplish this, *Gamebreaker* intends to rely on the prior successes from other fields such as finance, aircraft design, or epidemiology, where complex relationships (often with parameter spaces that are impossible to search by brute force) have been derived based on the observation of results.¹

Once a proposer determines the effective state equation for game balance, they can recommend the minimum modifications to the game that will result in an imbalanced condition. Proposers may choose to articulate the relative influence of different game elements in any number of different ways, and tornado plots or parameter coefficients could indicate relative sensitivities. Once complete, proposers will generate predictions to Government-provided independent

¹ M. Raissi, P. Perdikaris, and G.E. Karniadakis, Physics Informed Deep Learning (Parts I and II): Data-Driven Discovery of Nonlinear Partial Differential Equations, arXiv: 1711.10566v1, 28 Nov 2017

perturbations to avoid trivial or absurd solutions and impose natural constraints to control scope.

To determine extensibility and limitations of an approach, each proposer will implement their methodology in at least two separate games of different types (for example, one might begin with a fantasy game and then apply the method to a sports game).

Phase 1 will focus on the primary game, will last nine (9) months, and will conclude with the analysis of automated game balance assessment results including measurement of key contributors/sources of balance. Phase 2 will last another nine (9) months and will conclude with results of the automated introduction and measurement of perturbing effects and extensibility to the secondary game.

C. Structure

Proposals submitted to DARPA-PA-19-03-05 in response to the technical area of this *Gamebreaker* AIE Opportunity must be UNCLASSIFIED and must address two independent and sequential project phases: Phase 1 Feasibility Study (Base Period) and Phase 2 Proof of Concept (Option Period). The Periods of Performance for these two Phases are nine months for the Phase 1 Base Period and nine months for the Phase 2 Option Period. Combined Phase 1 Base and Phase 2 Option efforts for the *Gamebreaker* AIE Opportunity should not exceed 18 months. The Phase 1 (Base) award value is limited to \$500K. The Phase 2 (Option) award value is limited to \$500K. The total award value for the combined Phase 1 and Phase 2 is limited to \$1,000,000. This total award value includes Government funding and performer cost share, if required, or if proposed.

Proposals will be evaluated to the extent that the proposed approach addresses the Technical Area Descriptions identified below. Phase 1 activities will be evaluated to determine feasibility of the approach and whether to select the Option for Phase 2. Additionally, all proposals submitted in response to *Gamebreaker* AIE Opportunity will be evaluated and selected in accordance with Section 6 of the DARPA AIE Program Announcement, DARPA-PA-19-03, available on beta.SAM.gov here:

<https://beta.sam.gov/opp/5a5b919dcc337814fe57eb70c147bd72/view>

D. Technical Area Description

Game developers typically use a combination of anecdotal evidence, experience, and computerized methodologies such as Monte Carlo plays to coarsely assess and incrementally adjust the bounds of the game to maintain its balance. Detection of the existence of game imbalance can be difficult, and pinpointing the source can be challenging. To restore or maintain game balance, game developers iteratively decrease or increase the performance of various game elements (often referred to as “nerf” or “buff,” respectively) to counter an overpowered (OP) or underpowered (UP) environment. Unfortunately, this is often an intensively manual process.

Gamebreaker proposals are expected to use AI algorithms to automate the assessment of game balance, identify those parameters that most significantly contribute to balance, and evaluate perturbations such as the introduction of a new weapon, capability, or rule change and predict their effect on the balance of the game. Proposals must include the following:

- Identification and description of primary and secondary games to be used as a surrogate engagement model (at least one of these should be an RTS open world game) including justification for how it will be utilized to accomplish the objectives. Examples might include the planned use of existing open source toolkits like OpenAI Gym, access to

source code, an existing application programming interface, or the ability to modify parameters to automate the assessment of game balance. Proposals should also include a description and feasibility of any planned modification(s) required.

- A proposed integrated game balancing approach (consisting, for example, of context assessment, environment establishment, a combination of manual and automated balancing techniques, and data analysis). For more details see Preuss et al.² or Beyer et al.³ Proposers should include a plan to deal with complexity introduced via heterogeneity. For example, a symmetrical game is one in which both players have access to the identical load out of capabilities. This can be thought of as an Order of Battle (OoB). A game, or scenario, can be extremely complex but easily assessed to be balanced if it is symmetrical. As the relative OoBs become more heterogeneous, an additional dimension of complexity is introduced. It is expected that proposers may begin with simple deviations of heterogeneity between “forces”, testing the fitness of their approach as game conditions are altered and the degree of heterogeneity is increased. In this approach, selecting a game environment that can be easily modified is likely requisite.
- Description of training data acquisition plan (e.g. self-play, public game play, archived game play, etc.). The atom of analysis is a game replay as it captures all of the data from the game, including outcomes and often includes interim measures of success. Proposers will likely need to develop suitable agents to play the game. While this is not the primary focus of the effort, it will likely be required to generate the large sets of game replays within a reasonable timeframe, and research suggests that general game players may exist to help limit program scope.⁴ Proposers should address agent development in cases for which human play baselines may or may not exist, and should determine if/how the potential lack of human play may impact agent development. The quality of the agent(s) will also matter. For example, the more creative the agent(s) the more thorough exploration of total strategy space. Additionally, use of agent play should allow for control of player skill level.
- Outline of the selected algorithmic approach(es) for automating game balance assessment with the integrated game balancing process, including high-level application, limitations, and tradeoffs associated with implementation.
- Metrics beyond the win-loss ratio to assess the relative contributions of game characteristics and furthermore the effect of different game modifications or perturbations (such as the introduction of a new weapon, capability, or rule change). If possible, this should include a description of how to generate an empirically-based model of game balance to reduce the retraining burden. Examples of typical analysis tools can be found in Deepmind’s AlphaStar blog.⁵
- Detailed plan for creating and/or acquiring the baseline dataset that will be used to

² M. Preuss, T. Pfeiffer, V. Volz and N. Pflanzl, "Integrated Balancing of an RTS Game: Case Study and Toolbox Refinement," 2018 IEEE Conference on Computational Intelligence and Games (CIG), Maastricht, 2018, pp. 1-8. doi: 10.1109/CIG.2018.8490426

³ M. Beyer et al., "An integrated process for game balancing," 2016 IEEE Conference on Computational Intelligence and Games (CIG), Santorini, 2016, pp. 1-8. doi: 10.1109/CIG.2016.7860425

⁴ M. Genesereth, N. Love, and B. Pell, "General Game Playing," *AI Magazine* 26(2), 2005, pp 62-72.

⁵ <https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/>

introduce, predict, and evaluate perturbations (such as the introduction of a new weapon, capability, or rule change). This can be thought of as an approach for gamebreaking. This should include expected computational requirements, specific instances/scenarios, and their method of selection that will enable an impact assessment of Government-provided independent perturbations (to demonstrate that the solutions are flexible and not an artifact of a handcrafted solution).

- An assessment of the portability of the methodology to include anticipated limitations on how it might be applied to other open world games. Proposed plan for application of the methodology to the secondary RTS game.

As a risk reduction for the inevitable application to wargames, consideration should be given to an RTS game representative of a campaign-level military engagement. *Gamebreaker* performers are expected to use existing environments to the greatest extent possible, lightly tailored if necessary to maximize synergy with the intent of the program. A representative game environment should strive to satisfy as many of the following criteria as possible:

- RTS open-world environment;
- Effects-chain-based game: win/loss based on aggregate closing of effects chains where the accumulated magnitude of effects over some time/space determines winner;
- Allow for the prescription of heterogeneous asset capabilities/attributes (such as speed, range, field of regard, cost, damage, reliability, survivability, etc.);
- Capable of both individual and team behavior exploration, multi-polar games (i.e. games in which there are more than two players/teams);
- Provide some analogous representation of different operating domains (such as air, maritime, ground, space, cyber, etc.) with different fundamental limitations and time constants;
- Optional distribution of effects chain functions (i.e. sense, decide, act decisions do not have to be self-contained on a single asset); and,
- Faster than real-time capability.

Proposals should include methods to quantitatively measure balance within the game/scenario environment in at least two dimensions:

- How various capabilities (of platforms, weapons, etc.) contribute to game balance (i.e., if Blue uses an A, Blue win percentage increases by X%), and
- How various strategies and tactics contribute to game balance (i.e., if Blue does A, B, and C, Blue win percentage increases by X%).

E. Schedule/Milestones

Proposers must address the following Research Project Objectives, metrics, and deliverables, along with fixed payable milestones in their proposals. The task structure must be consistent across the proposed schedule, Task Description Document (TDD), and the Vol. 2 - Price Volume. If selected for award negotiation, the fixed payable milestones provided in Section 9.C of the Vol. 2 – Price Volume will be directly incorporated into Attachment 2 of the OT agreement (“Schedule of Milestones and Payments”). Please see the sample OT for Prototype provided as an attachment to AIE Program Announcement, DARPA-PA-19-03, available on beta.SAM.gov.

For planning and budgetary purposes, proposers should assume a program start date of **May 1, 2020**. Schedules will be synchronized across performers, as required, and monitored/revised as necessary throughout the program. Proposals must include delivery schedules for Phase 1 and Phase 2 that include timelines for preliminary (to facilitate inspection by the Program Manager) and final (to facilitate evaluation) release of deliverables.

Phase 1 Milestones include:

- Month 1 After Contract Award (ACA): Project kick-off: technical review of proposed schedule and approaches. Provide a comprehensive description of the proposed game balancing techniques that will be explored, with in-depth support of published literature and identification of literature gaps. Include an overview and comparison of current state of the art pertaining to the proposed approach especially as it compares to game balancing techniques. Define and justify primary game environment selection and identify any existing or to be developed infrastructure required to support research. Provide methodology for implementation of proposed game balancing techniques and highlight relative merits and risks associated with each. Identify and address major challenges (e.g. baseline for comparison, acquisition of training data, dealing with state space complexity, metrics for assessing balance) and provide plan to address each. Provide proposed schedule to include roles and responsibilities of project team, key contributors, and any subcontractors. Update the research plan to include details of the exploration that will be performed by Milestone 2.
- Month 4 ACA: Interim progress report: Provide an interim report describing the maturing game balance techniques and estimated performance. An update on the status of the game environment and any customization should be provided. This milestone should include identification of metrics to be used for the duration of the effort based on the proposed approach(es). Updates should be provided for any significant deviations from the original proposed plan based on insights gained during the first several months of research. Progress should include acquisition of initial training data sets, initial implementation of strategy for dealing with state space complexity, as well as key theoretical foundations identified. Provide strategy for "gamebreaking" based on insights developed from development, initial implementation of game balancing techniques. Update the research plan to include details of the exploration that will be performed by Milestone 3.
- Month 7 ACA: Interim review progress report: Provide an interim report describing the maturing game balance techniques and estimated performance. An update on the status of the game environment and any further customization should be provided. This milestone should include refinement of metrics to be used for the duration of the effort based on the proposed approach(es). Updates should be provided for any significant deviations from the original proposed plan based on insights gained during the first several months of research. Progress should include analysis results for game balance measurement for static game configuration further substantiated to hold up under game modification. Provide results or insights for initial implementation of "gamebreaking" strategy, including metrics and analysis tools for sense making. Update the research plan to include details of the exploration that will be performed by Milestone 4.
- Month 9 ACA: Feasibility study final review: This milestone should present a compelling case that game balance measurement and "gamebreaking" offers significant performance gains over the state-of-the-art and that sufficient conclusions can be derived from analysis. Evidence should be provided that the techniques developed are extensible to

other games, engagement models. Submit Phase 1 Final Report in preparation for the Phase 2 option. Phase 1 final report should summarize the approach for phase 2 tasks.

Phase 2 Milestones include:

- Month 10: Phase 2 kick-off: technical review of proposed schedule and approaches including plan to incorporate government-provided independent perturbations. Report on automated game balancing lessons learned, plan for continued technique development, and continued simulations/analysis. Provide a comprehensive description of the proposed gamebreaking techniques that will be implemented along with insights from implementation of automated game balancing. Provide plan to incorporate government-provided independent game modifications. Provide plan for application of methodology to secondary game environment, identify any existing or to be developed infrastructure required to support research. Update the research plan to include details of the exploration that will be performed by Milestone 6.
- Month 13: Interim review progress report: Provide an interim report documenting analysis and results of maturing automated game balance (secondary game) gamebreaking (primary game) techniques and estimated performance. An update on the status of the game environments and any further customization should be provided. This milestone should include refinement of both game balancing and gamebreaking metrics. Progress should include analysis results for gamebreaking measurement for both performer and government-provided game perturbations. Update the research plan to include details of the exploration that will be performed by Milestone 7.
- Month 16: Interim review progress report: Interim progress report on all aspects of the effort. Provide an interim report documenting analysis and results of maturing automated game balance (primary and secondary games) gamebreaking (primary and secondary games) techniques and estimated performance. This milestone should include refinement of both game balancing and gamebreaking metrics. Progress should include analysis results for gamebreaking measurement for both performer and government-provided game perturbations. Update the research plan to include details of the exploration that will be performed by Milestone 8.
- Month 18: Proof-of concept final review: This milestone should report on analysis, insights, and results of automated game balance and "gamebreaking" methodologies as well as summarize key findings achieved throughout the effort. Performers should present evidence based recommendations on the extensibility of the approach, future application and research areas including a cost and schedule implementation plan. Documentation should be provided sufficient for those knowledgeable in the field but did not participate in the study to continue the research at some later date. A technical data package will be provided that includes software, firmware, and simulation code. A demonstration of techniques developed and implemented throughout the effort is required. Submit Phase 2 final report.

All proposals must include the following meetings and travel in the proposed schedule and costs:

- To foster collaboration between teams and disseminate program developments, a two-day Principal Investigator (PI) meeting will be held approximately every six months, with locations split between the East and West Coasts of the United States.
- Regular teleconference meetings will be scheduled with the Government team for

progress reporting, as well as problem identification and mitigation. Proposers should also anticipate at least one site visit per Phase by the DARPA Program Manager during which they will have the opportunity to demonstrate progress towards agreed-upon milestones.

F. Deliverables

At the final reviews for Phase 1 and Phase 2, performers are expected to provide at a minimum the following deliverables:

- A technical description (and discussion of evolution) of the integrated game balancing process including algorithms and approach proposed to automate balance assessment.
- The algorithmic approach used to build agents (reinforcement learning, genetic fuzzy tree, etc.) including considerations such as:
 - The training/learning process;
 - Training convergence rates, computational requirements, and memory requirements for initial training;
 - Computational and memory requirements to run the trained AI system;
 - Types and amounts of data required for training;
 - Ability to update or extend the AI decision system without complete re-training;
 - Ability to adapt “in real time” to unexpected adversary behavior or capabilities;
 - Explainability and understandability of the decision process;
 - Ability to replicate and implement known effective tactics;
 - Ability to discover novel effective tactics; and
 - Ability to prove behavioral properties and performance of the trained system.
- A use case description of the implementation of the integrated game balancing process and documentation of results.
- Analysis tools and results for assessing game balance; these may include registered reports, experimental protocols, publications, intermediate and final versions of software libraries, code, and APIs, including documentation and user manuals, and/or a comprehensive assemblage of design documents, models, modeling data and results, and model validation data.
- Extensibility assessment applying methodology to the proposed secondary game as well as a plan for how the method might be applied to other gaming or simulation environments with an emphasis on the potential portability to those often used in military applications.

Proposers are expected to provide additional deliverables (plans, drafts, tests, etc.) to be delivered during intermediate reviews (listed above in milestones) to facilitate final review deliveries.

II. Award Information

Selected proposals that are successfully negotiated will result in award of an OT for Prototype Project. Refer to Section 3 of the AIE Program Announcement (DARPA-PA-19-03) for information on awards that may result from proposals submitted in response to this notice.

Proposers must review the model OT for Prototype agreement provided as an attachment to the AIE Program Announcement (DARPA-PA-19-03) prior to submitting a proposal. DARPA has provided the model OT in order to expedite the negotiation and award process and ensure

DARPA achieves the goal of AIE, which is to enable DARPA to initiate a new investment in less than 90 days from idea inception. The model OT is representative of the terms and conditions that DARPA intends to award for all AIE Opportunities. The task description document, schedule of milestones and payments, and data rights assertions requested under Volumes 1, 2, and 3 of the AIE Opportunity will be included as attachments to the OT agreement upon negotiation and award.

Proposers may suggest edits to the model OT for consideration by DARPA and provide a copy of the model OT with track changes as part of their proposal package. Please note that suggested edits, however, may not be accepted by DARPA. The Government reserves the right to remove a proposal from award consideration should the parties fail to reach agreement on OT award terms and conditions. If edits to the model OT are not provided as part of the proposal package, DARPA assumes that the proposer has reviewed and accepted the award terms and conditions to which they may have to adhere and the sample OT agreement provided as an attachment, indicating agreement (in principle) with the listed terms and conditions applicable to the specific award instrument.

In order to ensure that DARPA achieves the AIE goal of award within 90 days from the posting date (**February 3, 2020**) of this announcement, DARPA reserves the right to cease negotiations when an award is not executed by both parties (DARPA and the selected organization) on or before **May 1, 2020**.

III. Eligibility

Refer to Section 4 of the AIE Program Announcement (DARPA-PA-19-03) for eligibility information.

IV. AIE Opportunity Responses

Responses to the *Gamebreaker* AIE Opportunity must be submitted as full proposals to DARPA-PA-19-03 as described therein. All proposals must be UNCLASSIFIED.

A. Proposal Content and Format

All proposals submitted in response to this notice must comply with the content and format instructions in Section 5 of the AIE Program Announcement (DARPA-PA-19-03). All proposals must use the templates provided as Attachments to the PA and follow the instructions therein.

Information not explicitly requested in the AIE Program Announcement (DARPA-PA-19-03), its attachments, or this notice may not be evaluated.

B. Proposal Submission Instructions

Refer to Section 5 of the AIE Program Announcement (DARPA-PA-19-03) for proposal submission instructions.

C. Proposal Date and Time

Proposals in response to this notice are due no later than 4:00 PM, ET on March 3, 2020. Full proposal packages as described in Section 5 of the AIE Program Announcement (DARPA-PA-

19-03) must be submitted per the instructions outlined therein *and received by DARPA* no later than the above time and date.

Proposals received after this time and date may not be reviewed.

Proposers are warned that the proposal deadline outlined herein is in Eastern Time and will be strictly enforced. When planning a response to this notice, proposers should take into account that some parts of the submission process may take from one business day to one month to complete.

V. Proposal Evaluation and Selection

Proposals will be evaluated and selected in accordance with Section 6 of the AIE Program Announcement (DARPA-PA-19-03). Proposers will be notified of the results of this process as described in Section 7.1 of that same announcement.

VI. Administrative and National Policy Requirements

Section 7.2 of the AIE Program Announcement (DARPA-PA-19-03) provides information on Administrative and National Policy Requirements that may be applicable for proposal submission as well as performance under an award.

VII. Point of Contact Information

LtCol Daniel “Animal” Javorsek, Program Manager, DARPA/STO, Gamebreaker@darpa.mil.

VIII. Frequently Asked Questions (FAQs)

All technical, contractual, and administrative questions regarding this notice must be emailed to Gamebreaker@darpa.mil. Emails sent directly to the Program Manager or any other address may result in delayed or no response.

All questions must be in English and must include name, email address, and the telephone number of a point of contact. DARPA will attempt to answer questions publically in a timely manner; however, questions submitted within seven (7) days of the proposal due date listed herein may not be answered.

DARPA will post a Frequently Asked Questions (FAQ) list and informational video under the AIE Opportunity on the DARPA/STO Opportunities webpage at <https://www.darpa.mil/work-with-us/opportunities>. The list will be updated on an ongoing basis until one week prior to the proposal due date. In addition to the FAQ specific to this notice, proposers should also review the Program Announcement for AIE General FAQ list on the DARPA/STO Opportunities page under the Program Announcement for AIE (DARPA-PA-19-03).