
Amazon Research Awards

Prime Video call for proposals — Fall 2022

Pushing the boundaries of science and technology

 Share

About this CFP:

With the mission to be the world's most loved entertainment service, Prime Video continually strives to delight its customers by offering them the most engaging video-watching experiences. Prime Video works to solve a broad range of cutting-edge technical problems. We welcome proposals related to the following broad research areas in order to accelerate progress in the state of the art of video-watching.

Research Area 1 : Anomaly Detection and Insights

Undisrupted entertainment via large-scale anomaly detection

Overview: Customers should be able to reliably stream content at all times using any device where the application is available. This scope results in a combinatorial explosion of metrics that individually describe the quality of service across different marketplaces, regions, and devices. Monitoring such metrics in real-time is crucial for ensuring that any faults are immediately identified and resolved with minimal user disruption, and we are investing heavily in building robust and scalable anomaly detection tools to meet this requirement.

Research sub-areas: We are looking to fund research in the following sub-areas:

a. Anomaly detection on intermittent time series: Balancing the trade-off between maintaining high precision while also minimizing time-to-detection is a key challenge for time series anomaly detection. This is especially pertinent to the case of intermittent time series, whereby seasonality patterns in the time series may only become observable when data is aggregated over coarser periods of time.

Examples of research questions include:

- How can representations of time series at different levels of granularity be leveraged for detecting anomalies?
- For metrics with different levels of sparsity, how can we ensure that optimal time-to-detection is achieved for both smooth and intermittent metrics?
- How can we ensure that derived anomaly scores remain consistent and interpretable across all time series?

b. Anomaly detection for multivariate time series: Collections of monitorable time series are often closely correlated or even generated from the same underlying process, and would benefit from being jointly modeled. However, this may also introduce additional complexity when individual time series have dimensions that are either missing at random or altogether. Adapting to such data challenges is crucial to enabling robust anomaly detection across groupings of related time series.

Examples of research questions include:

- Are generative models effective for multivariate time series anomaly detection?
- How can we apply multivariate forecasting-based anomaly detection when individual multi-dimensional time series have missing or additional dimensions?
- How can we convert global anomaly scores emitted from a multivariate anomaly detection model to individual scores across specific dimensions?

c. Anomaly detection in time series with cold start

Monitoring capabilities to be enabled without extended warm-up period are important. The robustness of newly-deployed anomaly detection models often relies on having sufficient examples of normal and anomalous data before launch, but how do we handle cases where only few anomaly samples are available (if any at all)? The notion of transfer learning or meta-learning is commonplace in other domains, but has received little attention for time series anomaly detection.

Examples of research questions include:

- How can we jointly monitor new time series with ramp-up behavior (e.g. sessions on newly-launched devices) alongside related metrics having longer and smoother historical data?
- How can we determine the decision boundary for anomaly detection in a time series with no previous anomalies?
- As more labels become available over time, how can we efficiently augment already-existing models without having to retrain them from scratch?

Proposal Requirements: Proposals should be prepared according to the [proposal template](#). In addition, to submit a proposal for this CFP, please also include the following information:

1. Describe current applications of your work (e.g, libraries, codebases and industry code).
2. What are potential applications of your work?
3. What assumptions are made by your work (e.g., that affect soundness, precision, and/or scalability)?
4. If your work involves the development and maintenance of a tool:
 - a. What license is your tool released under?
 - b. What on-boarding/tutorial material is available?
 - c. Is your tool actively maintained (commits within last 3 months)? How many active contributors does your project have?
5. What data are you planning to work with?

Research Area 2 : Personalization and Discovery

Personalized recommendations and discovery

Overview: The mission of this call for proposal is to improve engagement by providing relevant, personalized and timely recommendations. We guide users to discover content, and to stream stories relevant to their interests -current and emerging- while providing users a personalized experience that is transparent and builds trust in our brand.

Research sub-areas: We are looking to fund research in the following sub-areas:

a. Representation learning for title/user understanding: This topic includes but is not limited to the following research areas: Developing new representation techniques for different entities in RS; Understanding and evaluating existing representations, e.g. probing representations for generalization, compositionality & robustness, adversarial evaluation, analysis of representations; Efficient learning of representations and inference with respect to training and inference time, model size, amount of training data, etc.

b. Reinforcement learning for title recommendation/page composition:

Reinforcement Learning (RL) is a sequential decision making technique which maximizes the notion of long-term rewards. Framing title/carousel recommendations as building RL agents that maximize user satisfaction will enable us to explore and extend recent RL developments. Under this call for proposals, we are specifically seeking proposals on 1. RL Applications in recommender systems 2. Page composition 2. Real-world challenges and best practices for RL e.g. effective real-world exploration strategies, the role of offline and online metrics for diagnostics and modeling, Real-time inference and scalable ML workflows, hyper-parameter tuning for RL, interpretability, scalability and exploratory data analysis. 3. RL algorithms and evaluation, e.g. data driven, offline, and batch reinforcement learning; off-policy learning and counterfactual evaluation; deep RL and multi-arm Bandits; bandits for non-stationary environments.

c. New deep learning architecture for recommendation: Deep learning based recommendation system architectures make use of multiple simpler approaches in order to remediate the shortcomings of any single approach to extracting, transforming and vectorizing a large corpus of data into a useful recommendation for an end user. The topic will focus on new learning paradigms & architecture for recommendation systems.

d. Trustworthy AI for recommendation: Recommendation systems may lead to undesired counter-effects on users, items, producers, platforms, or even the society at large, such as compromised user trust due to non-transparency, unfair treatment of different consumers, or producers, privacy concerns due to extensive use of user's private data for personalization et al. All of these create an urgent need for Trustworthy Recommender Systems (TRS) so as to mitigate or avoid such adverse impacts and risks. The topic will focus on fairness and bias, explainability, diversity, causal inference, accountability, and privacy across recommendation systems.

Proposal Requirements: Proposals should be prepared according to the [proposal template](#).

Research Area 3 : Video Quality Analysis

Eliminating media quality defects

Overview: The ability to accurately detect and characterize media quality defects across both new and existing formats is critical.

Research sub-areas: We are looking to fund research in the following sub-areas:

a. Extended parameter space: While new media formats such as UHD, HDR and Dolby Atmos deliver unparalleled immersive experiences, there is a limited understanding of their interaction within the wider extended parameter space (e.g. resolution, bitrate, content etc) and the impact on media quality. This will be a key driver in the development of novel technologies to identify audio and/or visual quality defects across this space.

b. Context: Signal processing techniques are typically applied on audio, video and/or closed captions to detect media quality defects. Context (e.g. genre, creative intent, metadata), and interactions between these signal modalities, are often ignored. Methodologies which can exploit and disentangle this additional information will likely be more precise, robust and explainable.

c. Defect synthesis: There exists a long tail of media quality defects, and as a consequence obtaining ground truth examples can be prohibitive and expensive. Synthesizing these defects poses a number of challenges including defect modeling, data imbalance and a potential domain gap between synthetic and real data. Additionally it can be difficult to both characterize and have confidence in detecting low-prevalence and/or unseen defects. Addressing these problems is crucial in ensuring that consistent experience is delivered to all users.

Proposal Requirements: Proposals should be prepared according to the [proposal template](#), and can be video, audio and/or audiovisual based.

Timeline

Submission period: September 16 to October 19, 2022

Decision letters will be sent out March 2023

Award details

Selected Principal Investigators (PIs) may receive the following:

- Unrestricted funds, no more than \$50,000 USD on average
- AWS Promotional Credits, no more than \$40,000 USD on average
- Training resources, including AWS tutorials and hands-on sessions with Amazon scientists and engineers

Awards are structured as one-year unrestricted gifts. The budget should include a list of expected costs specified in USD, and should not include administrative overhead

costs. The final award amount will be determined by the awards panel.

Eligibility requirements

Please refer to the ARA Program rules on the [FAQ page](#).

Proposal requirements

Proposals should be prepared according to the [proposal template](#). In addition, to submit a proposal for this CFP, please also include the following information:

1. Please list the open-source tools you plan to contribute to.
2. Please list the AWS ML tools you will use.

Selection criteria

ARA will make the funding decisions based on the potential impact to the research community and quality of the scientific content.

Expectations from recipients

To the extent deemed reasonable, Award recipients should acknowledge the support from ARA. Award recipients will inform ARA of publications, presentations, code and data releases, blogs/social media posts, and other speaking engagements referencing the results of the supported research or the Award. Award recipients are expected to provide updates and feedback to ARA via surveys or reports on the status of their research. Award recipients will have an opportunity to work with ARA on an informational statement about the awarded project that may be used to generate visibility for their institutions and ARA.

How to apply

When you're ready to submit your proposal, click the button below and follow the instructions on the site.

[News and features](#)

[Publications](#)

[Conferences](#)

[Collaborations](#)

[FAQs](#)

[Research Awards](#)

[Amazon Developer](#)

[Amazon Web Services](#)

[Newsletter](#)

[RSS](#)



Have feedback for Amazon Science?

We want to hear from you.

[Take survey](#)