Obfuscating Biological Sequences to Ethically Reveal Vulnerabilities

Pilot Event Synopsis

OBSERV is a multi-week virtual challenge designed to identify vulnerabilities in gene synthesis company sequence screening algorithms by leveraging the expertise of small teams of students and postdocs. This proposed pilot event is open to EBRC-affiliated students and postdocs and will test the feasibility, utility, format, and potential security risks and benefits of an event of this nature. Up to 20 teams of three students and postdocs each will be allowed to participate. Pending the results of this event, a future iteration of OBSERV could be expanded beyond the EBRC community to provide more regular stress testing of sequence screening systems.

Table of Contents

Overview

Purpose

OBSERV aims to build and promote a security-conscious community within engineering biology, enabling early-career researchers to collaborate with commercial gene synthesis providers and biosecurity experts to strengthen global security.

Motivation

OBSERV seeks to identify and mitigate sequence screening vulnerabilities by leveraging the expertise of students and postdocs on the cutting edge of research. This event proposes red-teaming, adversarial stress testing by qualified, benevolent actors, as a tool to strengthen commercial sequence screening methods. We believe that emerging biotechnology leaders bring creativity and fresh perspectives which uniquely enable them to identify unknown gaps in current sequence screening methodologies. Initial results of the red-teaming exercises will be shared in strict confidence with representatives from participating synthesis companies so that they can immediately patch any vulnerabilities that are identified. The event organizers may also provide the results to the broader security community under guidelines developed in consultation with security experts and government partners to ensure that knowledge of vulnerabilities is shared appropriately. The broader biotechnology community can utilize these findings to improve current screening practices and inform future development and governance of gene synthesis technologies.

Format

OBSERV challenges small teams of trainees to develop strategies to identify and obfuscate genetic sequences that, when assembled and integrated into biological systems, may pose security concerns. Select gene synthesis companies opt in to screen these sequences to identify vulnerabilities and use this information to strengthen their screening systems. A panel of expert judges evaluate the teams' sequences and obfuscation strategies and select winning teams based on the potential impact of the vulnerabilities that are identified. The broad findings from this event may be published publicly such that all gene synthesis companies can benefit from the lessons learned.

Event Organizers and Affiliates

This event is led by the Engineering Biology Research Consortium (EBRC) Student and Postdoc Association and Security Working Group. Graduate students Kathryn Brink (Rice University) and Bridget Luckie (University of California, Berkeley) conceived and created OBSERV. They can be contacted at observ@ebrc.org.

Participating sequence screening companies and organizations include Battelle, Ginkgo Bioworks, Joint Genome Institute, Raytheon BBN Technologies, Thermo Fisher Scientific, and Twist Bioscience. Prizes are generously sponsored by Twist Bioscience.

Participant Guidelines and Eligibility Requirements

Link to Registration Form

Participants must:

- Be students or postdocs at universities, non-profit research institutions, or National/US Government Labs located within the United States
- Be members of the EBRC Student and Postdoc Association *OR* be members of an EBRC-affiliated lab
- Attend an EBRC Malice Analysis workshop prior to the beginning of the competition (see details below)
- Attend a Kickoff Event that summarizes event rules and expectations (see details below)

Participants will be divided into teams of three members each. Members of pre-formed teams can indicate their preferred teammates on their registration forms. Participants who do not belong to a pre-formed team will be assigned to teams by the event organizers. Up to 20 teams will be allowed to participate on a first come, first served basis.

Competition Logistics

Malice Analysis Workshop

All participants will be required to attend a Malice Analysis (MA) workshop to participate in the competition. Malice Analysis provides participants with a framework for assessing security concerns and enables participants to practice adopting a security mindset. Workshop participants will also receive a certificate of completion that can be noted on their CV. Participation in any one of the upcoming workshops listed below satisfies the Malice Analysis workshop requirement. The upcoming events are:

- **April 28**, 1-4:30pm ET / 10am-1:30pm PT (Rice University)
- **April 29**, 12-3:30pm ET / 9am-12:30pm PT (UC Berkeley)
- **May 5**, 1-4:30pm ET / 10am-1:30pm PT (MIT)

While these workshops are affiliated with particular universities, each workshop is open to all OBSERV participants. Please sign up for the workshop that works best for your schedule. Event participants who have previously attended a Malice Analysis workshop are not required to attend a second workshop. Participants who are unable to attend any of the MA workshops but who would still like to participate in the competition should contact the event organizers (observ@ebrc.org), who may be able to make accommodations.

Kickoff Event

The virtual Kickoff Event will be a required event on **May 3**, **7-8:30pm ET / 4-5:30pm PT**, and will serve as the formal start of the competition. Event organizers will communicate an overview of the competition format and rules. The Kickoff Event will also include information about sequence screening algorithms and about policies surrounding information hazards. Participants who are unable to attend the Kickoff Event but who would still like to participate in the competition should contact the event organizers (observ@ebrc.org), who may be able to make accommodations.

Competition Tracks

The organizers have identified two general classes of vulnerabilities in gene synthesis screening approaches, which form the basis of two tracks for the competition. The first track seeks to identify obfuscation strategies that could enable a bad actor to obtain sequences that are currently regulated by the US government. The second track seeks to identify sequences that pose security concerns but are not currently subject to regulation by the US government.

Teams may choose to participate in one or both tracks. In Track 1, teams will have three opportunities to design obfuscation strategies and submit sequences based on those strategies to gene synthesis companies, who will return results indicating whether the sequences passed screening or raised flags. These multiple rounds will enable teams to probe the limits of screening systems and iterate on their designs. Track 2 sequences are designed to fall outside the boundaries of current sequence screening and regulation, and as a result, are not expected to raise flags. Thus Track 2 sequences will only be submitted once to the event organizers during the final submission round and, with some possible exceptions, will not be subjected to screening.

At the end of the competition, teams will submit each obfuscation strategy (for Track 1) or each set of related sequences of concern (for Track 2) that they develop as a separate write-up. These write-ups will be reviewed by a panel of judges from the EBRC Security Working Group. Judges will evaluate Track 1 and Track 2 submissions separately and choose one winning team and one runner-up team for each track.

Track 1: Vulnerabilities involving sequences currently regulated by the US government

In Track 1, participants will attempt to identify vulnerabilities in sequence screening that could enable an individual to obtain a sequence currently subject to regulation by the US government. Participants will attempt to obfuscate regulated sequences (e.g., by exploiting DNA assembly strategies or by any other means) so that they are not recognized by screening algorithms.

Gene synthesis companies currently screen for sequences on the following lists:

- Select Agents and Toxins List (FSAP)
- Commerce Control List (CCL), including the following:
    - 1C351 (p. 67)
    - 1C353 (p. 70)
    - 1C354 (p. 71)
    - 1C991 (p. 81)

Because the FSAP and CCL lists can change over time, participants will be provided with a copy of the latest lists at the start of the competition that they should use as the basis for determining whether a sequence is controlled. Sequence submissions will be evaluated based on the provided lists and evaluations will ignore any updates to these lists that may occur while the competition is in progress.

In Track 1, participants will submit a list of sequences that 1) are explicitly or implicitly controlled by these lists or 2) could be assembled or manipulated *in vitro* or *in vivo* so as to produce a sequence that is explicitly or implicitly controlled by these lists. Specifically, as noted in the CCL, sequences are implicitly controlled if they are a) unique to a controlled organism and/or b) able to endow or enhance pathogenicity.

Track 1 does not include metabolites. Enzymes or pathways that produce metabolites regulated by FSAP and CCL should be submitted through Track 2 instead.

Track 2: Vulnerabilities involving sequences not currently subject to regulation

The event organizers recognize that sequences of concern may exist that are not covered by current regulations. Track 2 aims to identify these sequences. In Track 2, participants will submit currently unregulated sequences that 1) are sequences of concern or 2) could be assembled or manipulated *in vitro* or *in vivo* so as to produce a DNA sequence, RNA sequence, protein, or biochemical compound that poses security concerns.

Sequences submitted via Track 2 should not be currently regulated by FSAP or CCL (see above). However, sequences of enzymes or pathways that are capable of producing metabolites that are controlled by FSAP or CCL are eligible to be submitted in Track 2, as current screening systems are not typically designed to flag these sequences.

Sequence Submissions

Participants will only submit sequences to the event organizers. Participants should **NOT** attempt to order sequences developed in the context of the OBSERV competition with participating (or non-participating) synthesis companies.

Track 1

Teams competing in Track 1 will have three opportunities to submit sequences for screening, each one week apart. This format will permit teams to learn and iterate on their sequence obfuscation strategies. For each sequence submission, teams will provide a FASTA file containing a list of up to 100 sequences to the event organizers prior to the indicated submission deadline. Teams may, but are not required to, use computational approaches to develop their sets of sequences. There is no minimum number of sequences required per submission, but teams are encouraged to maximize the amount of information that they gain about the screening process by submitting multiple sequences in each round. Teams are also encouraged to test multiple sequence obfuscation strategies in each round of submissions to increase their chances of identifying a successful obfuscation strategy. Where feasible, we recommend that participants attempt the same obfuscation strategies with several sequences of concern to provide an estimate of the magnitude of the vulnerability (e.g., percent of sequences of concern that do not raise flags for a given obfuscation strategy).

The event organizers will share sequence submissions with participating companies and return results to participants within 72 hours. All participants will be notified of their results simultaneously via email. The results will consist of a CSV file that contains four columns, namely:

1. Sequence name
2. Sequence
3. Flag: True/false field where “False” indicates that the sequence raised no flags in the screening process and “True” indicates that the sequence was flagged for further review.
4. Notes: Companies can use this field to provide additional context about their true/false finding as reported in “Flag”.

Submissions are subject to the following restrictions:

1. Sequences should be DNA sequences that contain only the four canonical DNA bases (adenine (A), cytosine (C), guanine (G), and thymine (T)). Degenerate sites (e.g., N) are not allowed.
2. All submitted sequences must be gene fragments greater than or equal to 200 base pairs and no longer than 5000 base pairs in length. For the purposes of this event, teams should assume that fragments shorter than 200 base pairs would automatically pass sequence screening. Sequences that are shorter than 200 base pairs should still be included in a team’s write-up (see below) if they would be required for assembly of a sequence of concern.
3. Teams should assume that all sequences are being screened by the same vendor.

Participants should ***only* submit sequences to the event organizers and should NOT attempt to order their sequences on the websites of participating (and non-participating) synthesis companies.**

Track 2

Because Track 2 sequences are designed to fall outside of normal screening processes, Track 2 sequences will only be submitted once, at the conclusion of the competition. During the evaluation process, judges can request that sequences be screened to validate that they are not currently regulated by FSAP or CCL. Sequences submitted to Track 2 that *do* fall under FSAP or CCL guidelines may be disqualified from consideration. However, in the event that a Track 2 submission is flagged as a sequence of concern but is *unregulated* based on the FSAP and CCL lists provided, the submission will not be disqualified.

Competition Scoring

After all sequence submission periods have ended and teams have received the results for each submission, teams will submit write-ups of their obfuscation strategies (Track 1) and sets of unregulated sequences (Track 2). Teams should submit a separate write-up for each strategy or set of related unregulated sequences that they attempt or identify during the competition. Write-ups should contain sufficient detail such that a skilled individual could replicate the team's strategy. Teams may include up to 1,000 words per write-up but are encouraged to be concise. Each write-up should be accompanied by a file containing all relevant sequences. Teams are encouraged to include relevant literature citations where possible to provide additional context for their sequence selections and strategies.

These write-ups will include:

1. Rationale for choosing the submitted sequence(s).
2. Workflow for how the submitted sequences could be used to generate a DNA sequence, RNA sequence, protein, or biochemical compound of concern *in vitro* or *in vivo*.
    a. For example: a cloning strategy to generate a controlled DNA sequence from a set of submitted sequences.
3. Justification for why the submitted sequence or sequences pose security risks and their projected societal impact.
4. *Bonus:* Suggested changes to the sequence screening process to mitigate the vulnerability. (up to 500 additional words)

Teams will be evaluated by a panel of judges comprising volunteers from the EBRC Security Working Group and other biosecurity experts. These judges will evaluate the write-ups using the following criteria:

1. *(Track 1 only)* Success at evading detection by gene synthesis companies: Submitted sequences do not raise flags for further review.
2. Feasibility of approach: Ease with which a semi-skilled actor could transform the submitted sequences into a security threat.
3. Creativity of approach: Originality in the team's design.
4. Societal impact: Impact on society if vulnerability were to be exploited.

5. Bonus points will be awarded to teams that suggest changes to the sequence screening process that could mitigate the vulnerability that they attempted to exploit/exploited.

In the event that no Track 1 team successfully identifies a vulnerability (i.e., submits a sequence of concern that does not raise flags), the other criteria will be used to identify winning teams. One winning team and a second runner-up will be chosen for each track.

Public Release of Winning Teams

Winners and runners-up will be announced to all participants approximately two weeks after the write-ups have been submitted. At that point, participants on the winning teams can opt in to having their names published on the event webpage. The event organizers will coordinate with the winning teams to provide them with prizes.

Prizes

Twist Bioscience has generously provided $5000 in prizes to be distributed as follows:

Track 1

- Winning team: $1500
- Runner up: $1000

Track 2

- Winning team: $1500
- Runner up: $1000

Academic Publication of Event Results

The event organizers may coordinate the publication of a summary of findings from the event. Details about the vulnerabilities identified during the event may be withheld due to security-related concerns. The event organizers will work with journal editors, appropriate government representatives, and security experts to ensure that published materials do not pose security risks.

Managing Information Hazards

Vulnerabilities in gene synthesis company screening approaches constitute information hazards, as they could be exploited by bad actors.

We will take several measures to prevent accidental or intentional release of information hazards, including:

1. Company identity will be anonymized to competition participants. Participants will communicate their sequence submissions to the organizers who will submit them to companies on the participants' behalf. This rule mitigates the risk that participants could accidentally or intentionally release information about company-specific vulnerabilities.

2. Participant data, including information about their university or research institution, will be collected and reviewed by the organizers and EBRC to ensure that participants meet the guidelines for participation (e.g., is a student or postdoc, is studying/employed at a US institution, is EBRC-affiliated). Only the event organizers will have access to participant data, and participant data will only be released beyond the event organizers (a) with the consent of the participant, (b) if the participant violates competition guidelines, or (c) as is otherwise required by law.

3. Participant information will be anonymized to gene synthesis companies when organizers submit sequence submissions. This rule prevents accidental or intentional release of identifiable information about participants.

4. Participants will receive a short training about information hazards, including what constitutes an information disclosure and consequences for the intentional disclosure of information hazards.

Disclosure Guidelines

1. All information about sequence submission results (including whether any submissions raise flags during screening, which companies experience breaches, etc.) must remain confidential. Participants may only share information about their sequence submission results with other members of their team or with the event organizers.

2. If participants believe that a vulnerability they identify poses a severe or immediate security threat, they must contact the event organizers (observ@ebrc.org). Any regulated sequence that has not been obfuscated that fails to raise flags during screening constitutes a significant vulnerability and should be reported immediately. The event organizers will communicate this threat at their discretion to appropriate entities including participating companies and the International Gene Synthesis Consortium (IGSC), such that affected companies can take immediate action to patch the vulnerability, and/or to legal authorities. If any participant is uncomfortable contacting the organizers, they may alternatively contact Clem Fortman (EBRC Director of Security Engagements, clem@ebrc.org) directly.

3. In the event of an intentional disclosure of confidential information (as described above), the participant and their team will be disqualified from the event. In coordination with the event organizers, the EBRC may also bar the participant from participating in any future

events of a similar nature. In severe cases, additional actions may be taken.

4. If a participant believes that confidential information was unintentionally disclosed (e.g., through malware), the participant should immediately contact the event organizers. The event organizers may investigate the incident and inform participating companies and/or legal authorities about the disclosure.

5. Participants who have questions about what constitutes an information disclosure should contact the event organizers.

Schedule

Mon	Tues	Wed	Thurs	Fri	Sat	Sun
May 3 Kickoff Event						
May 10 Round 1 sequence submissions due (*Track 1*)			**May 13** Round 1 sequence results returned to participants (*Track 1*)			
May 17 Round 2 sequence submissions due (*Track 1*)			**May 20** Round 2 sequence results returned to participants (*Track 1*)			
May 24 Round 3 sequence submissions due (*Track 1*)			**May 27** Round 3 sequence results returned to participants (*Track 1*)			
May 31 Write-ups and accompanying sequences due to event organizers (*Track 1 & 2*)	**June 1** Event organizers send write-ups and sequences to judges					

	June 15 Judges return scoring results			**June 18** Winners announced		