MITRE

Integrated Cognitive-neuroscience Architectures for Understanding Sensemaking (ICArUS):

Phase 2 Challenge Problem Walkthrough

Kevin Burns
Craig Bonaceto

November, 2014

Table of Contents

1 Introduction

This document provides an introduction to the Integrated Cognitive-neuroscience Architectures for Understanding Sensemaking (ICArUS) Phase 2 challenge problem. The pages include screen shots from the tutorial that participants viewed before and during experiments. This tutorial is embedded in the Graphical User Interface (GUI) used in experiments.

The present document is not meant to stand alone. For a detailed description of the ICArUS Phase 2 challenge problem, please refer to the following document:

> Burns, K. (2014). *ICArUS Phase 2 Challenge Problem Design and Test Specification*. MITRE Technical Report, MTR140412.

For an overview of all materials developed in ICArUS Phase 2 Test & Evaluation (T&E), including the software and data from human experiments, please refer to the following document:

> Burns, K., Fine, M., Bonaceto, C., & Oertel, C. (2014). *Integrated Cognitive-neuroscience Architectures for Understanding Sensemaking (ICArUS): Overview of Test and Evaluation Materials*. MITRE Technical Report, MTR140409.

These documents are available at: http://www.mitre.org/publications.

There are two major sections of the tutorial appearing in the following pages. The first section, *Exam Tutorial*, provides an overview of the "missions" (i.e., game tasks) including scoring and tactics as well intelligence information that is provided as part of the experiment. Participants viewed this portion of the tutorial at the start of the experiment. The second section, *Mission Instructions*, provides detailed instructions for five missions that were played in succession by each participant in the experiment. The instructions for each mission were viewed immediately before the start of that mission. Both the *Exam Tutorial* and the *Mission Instructions* were also made available to participants (via the GUI) for references purposes throughout the experiment.

All materials used in human experiments were approved by MITRE's Institutional Review Board prior to use.

2 Exam Tutorial

Exam Tutorial

Overview of the Task

Scoring and Tactics

Intelligence Reports

Your Role

In this experiment, you will be playing the role of "Blue" intelligence and operations against a "Red" adversary.

The Red-Blue conflict is played out in a series of trials, where each trial represents one day's cycle of Blue intelligence and operations.

On each trial, Red may or may not attack a Blue point inside the Blue border.

This point is a location of planned Blue activity, which represents a potential target for Red attack.

Your Goal

Red is an insurgent force that plays offense, by either **attacking** or **not attacking** at the Blue location.

Blue is a counter-insurgent force that plays defense, by either diverting or not diverting from the Blue location.

On each trial, each side must choose its own action without knowing the other side's choice of action.

A trial is scored with points awarded to Red, Blue, or neither – as discussed in the next section of this Exam Tutorial: **Scoring and Tactics.**

Points accumulate over trials of a mission. The experiment will include several missions in different areas of interest, each with different Blue borders and different Red behaviors.

<u>Your goal is to maximize the Blue score (minimize the Red score).</u>

Trials

On each trial, you will be given a series of intelligence reports for use in estimating the probability of Red attack. These reports are listed on the left panel of the screen and described in a later section of this Exam Tutorial: **Intelligence Reports.**
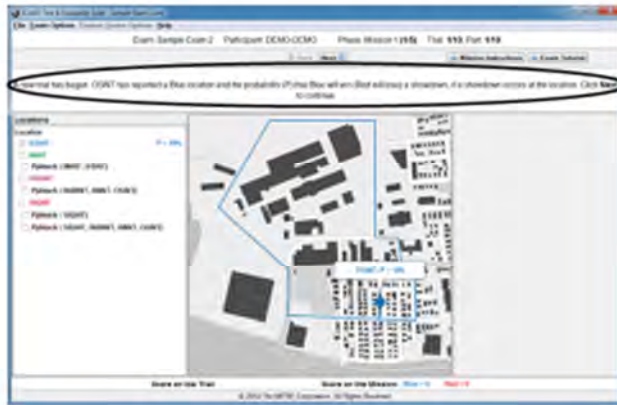
Missions

The experiment involves several missions, each placing different demands on Blue intelligence and operations. The mission and trial number are displayed in a status bar (circled below).

Message

At each stage of a trial, a message (circled below) will direct you on how to proceed. Directions for each mission are given, before the mission, in **Mission Instructions**.

Response

Sometimes a message will prompt you for a response. You will enter your response in a panel on the right side of the screen (circled below) per **Mission Instructions**.
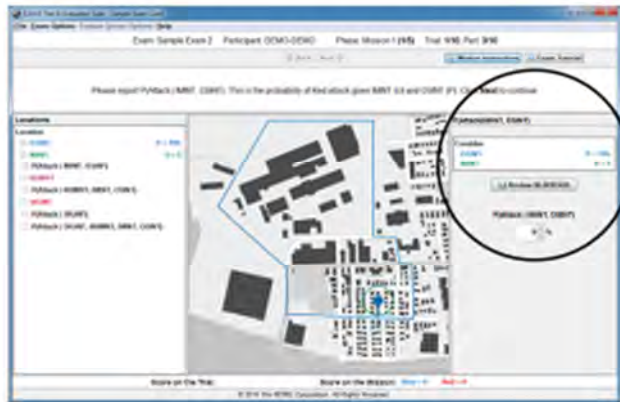
Review

You may review this **Exam Tutorial** or **Mission Instructions** at any point during the experiment. To do so, just click one of the help buttons (circled below).

Exam Tutorial

Overview of the Task
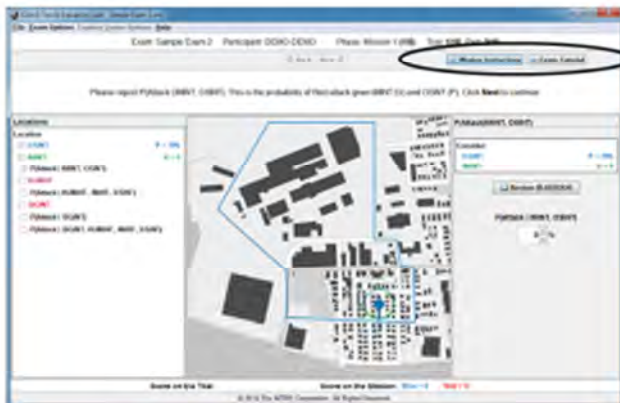
Scoring and Tactics

Intelligence Reports

Showdown: P and U

With two possible Red actions and two possible Blue actions, there are four possible combinations on each trial:

	Attack	Not Attack
Not Divert	Showdown	No attack
Divert	Blue diverts	No attack but Blue diverts

A **showdown** occurs when Red attacks and Blue does not divert.

The outcome of a showdown depends on two factors:

One factor is the likelihood that Blue can successfully defend the location, quantified as a **probability P**.

Another factor is the importance of the location, quantified as a **utility U**.

Numerical values for P and U are known by both Red and Blue, before choosing their actions. These values are important because they will affect Red and Blue decisions.

Scoring: Blue and Red

When a showdown occurs, the computer uses P for the trial to choose a winner probabilistically.

For example, if P = 20% then the computer will (on average) choose "Blue" as the winner on 2 of 10 trials and "Red" as the winner on 8 of 10 trials.

The winner of the showdown scores +U points; the loser scores -U points.

For the other three cases, where no showdown occurs, the score for Blue and Red is a number (-1, 0, or +1) that does not depend on P or U.

	Attack	Not Attack
Not Divert	+U (-U) or -U (+U)	0 (0)
Divert	0 (0)	-1 (+1)

The matrix shows Blue (Red) scores for all combinations of actions.

Note that any points won by Red are points lost by Blue, and vice versa.

Tactics: The BLUEBOOK

Your choice to divert or not, on each trial, depends on the probability of Red attack.

Red's action depends on P and U, per a table of attack probabilities that are programmed into the computer.

For example, if the probability of Red attack is 80%, then the computer will (on average) choose Red's action as "Attack" on 8 of 10 trials and "Not Attack" on 2 of 10 trials.

These probabilities are referred to as Red's **tactics**.

Probability of Red Attack

P > 25%	20%	40%
P ≤ 25%	60%	80%
	U = 2 or 3	U = 4 or 5

You will be given a table of Red's tactics in a Blue handbook. This Blue handbook, which reflects Blue's knowledge of Red's tactics, is called the BLUEBOOK.

Exam Tutorial

Overview of the Task

Scoring and Tactics

Intelligence Reports

Four Sources of Intelligence

Each trial of each mission will present one or two Blue locations that Red may possibly attack. The location(s) will vary from trial to trial.

For each location, intelligence reports will be provided for your use in estimating the probability of Red attack.

Besides the BLUEBOOK, discussed earlier under **Tactics**, there are four sources of intelligence described in the following slides:

OSINT   Open-Source Intelligence
IMINT   Image Intelligence
HUMINT Human Intelligence
SIGINT   Signal Intelligence

OSINT: Open-Source Intelligence

On a given trial, one or two locations of planned Blue activity are reported by Open-Source Intelligence (OSINT). These locations are known to both Red and Blue.

The locations are chosen by the computer, so they are not under your control.

However, you do control whether Blue will divert or not divert from each location.

At a given location, OSINT reports the probability P that Blue can defend the location against Red attack.

This probability P, known to both Red and Blue, governs the outcome of a showdown. P is the probability that Blue will win (Red will lose).

P increases with the distance (shortest straight line) that Red must penetrate into the Blue region in order to attack the location.

OSINT - Example

This location is close to the Blue border, OSINT reports P = 11%.

OSINT: P = 11%

IMINT: Image Intelligence

Each location of possible attack carries an associated importance, measured by utility U.

This utility U determines the number of points won or lost on a showdown.

U depends on the density of buildings at the Blue location.

The value of U is reported to you by Image Intelligence (IMINT).

Like P from OSINT, the U from IMINT is known by both Red and Blue.

However, all remaining intelligence differs between Red and Blue.

IMINT - Example

The IMINT circle (green) represents the area over which density is calculated. Here the location has a high density of buildings, IMINT reports U = 5.

BLUEBOOK

As noted earlier under **Tactics**, the probability of Red attack depends on P and U per the BLUEBOOK.

Therefore, after receiving OSINT (P) and IMINT (U), you can use the BLUEBOOK to determine the probability of Red attack.

Note: Throughout the experiment, P will always be less than 50%, and U will always be equal to 2, 3, 4, or 5.

Probability of Red Attack

	U = 2 or 3	U = 4 or 5
P > 25%	20%	40%
P ≤ 25%	60%	80%

BLUEBOOK - Example

This location has P = 11% and U = 5. Per the BLUEBOOK, the probability of Red attack is 80%, and this cell of matrix is highlighted by the computer display.

Probability of Red Attack

	U = 2 or 3	U = 4 or 5
P > 25%	20%	40%
P ≤ 25%	60%	80%

12

HUMINT: Human Intelligence

As described above, the BLUEBOOK gives the probability of Red attack, assuming Red has the capability to attack. **But Red may not be capable.**

Red's capability to attack is quantified by another probability, which increases with time from the last attack. This is reported to you by **Human Intelligence (HUMINT)**.

The computer uses this probability to determine whether Red has the capability to attack – "yes" or "no".

For example, if HUMINT reports a capability of 70%, then (on average) the computer will choose "yes" on 7 of 10 trials and "no" on 3 of 10 trials.

If the answer is "no", Red will not attack because Red does not have the capability to attack.

If the answer is "yes", the computer will use the BLUEBOOK value of Red attack probability to choose probabilistically whether Red actually attacks or not.

HUMINT - Example

This trial follows three consecutive trials without a Red attack. HUMINT reports a probability of 100% that Red is capable of attacking on this trial.

SIGINT: Signal Intelligence

Red's action (Attack or Not Attack) will be chosen by the computer, based on HUMINT and the BLUEBOOK as described above.

After Red's action has been chosen, you will be given a final intelligence report before making your choice (Divert or Not Divert).

This final intelligence report comes from **Signal Intelligence (SIGINT)** about Red's activity in the vicinity of the Blue location.

The SIGINT report is based on sensors that detect communications, such as cell phone calls, associated with the Red action to Attack or Not Attack.

More specifically, SIGINT will report either "Chatter" or "Silence", as an indication of whether Red has chosen to Attack or Not Attack.

However, these signals are known to be of limited reliability, per a table of likelihoods (see the next chart).

SIGINT - Example

The SIGINT circle (red) indicates the area over which SIGINT is collected. Here SIGINT reports "Chatter". Based on the table of SIGINT reliabilities, this signal is likely to be received (60% probability) if Red has chosen to Attack, and unlikely to be received (20% probability) if Red has chosen to Not Attack.
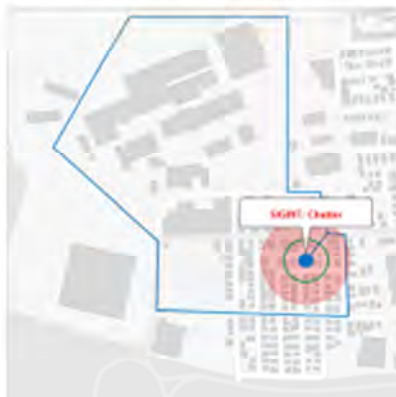
Likelihood of Signal

	Attack	Not Attack
Signal = Chatter	60%	20%
Signal = Silence	40%	80%

Notation for P(Attack|INTS)

Throughout the experiment, you will be asked to estimate the probability of Red attack given some intelligence, denoted **P(Attack|INTS)**.

In this notation, the symbol "**|**" means "given".

The word "**INTS**" refers to the intelligence you have been given, which you should consider in estimating the probability of Red attack.

Here are three examples:

P(Attack|SIGINT) is the probability of Red Attack given only SIGINT.

P(Attack|IMINT, OSINT) is the probability of Red Attack given IMINT and OSINT.

P(Attack|HUMINT, IMINT, OSINT) is the probability of Red Attack given HUMINT, IMINT, and OSINT.

3   Mission Instructions

Mission 1: Diagnosing the Chances of Attack

Mission 1 is an introductory mission in which you will play only the role of Blue intelligence.

A computer program will play the role of Blue operations and choose a Blue action (Divert or Not Divert) on each trial.

Like you, the computer program chooses a Blue action (Divert or Not Divert) without knowing what action Red has chosen (Attack or Not Attack).

Your job in Blue intelligence is to **estimate the probability of Red attack** on each trial.

P(Attack|IMINT, OSINT)

To begin a trial, OSINT reports P and IMINT reports U. Based on these values, and using the BLUEBOOK, you are asked to report **P(Attack|IMINT, OSINT)**. This is the probability of Red attack given IMINT (U) and OSINT (P).
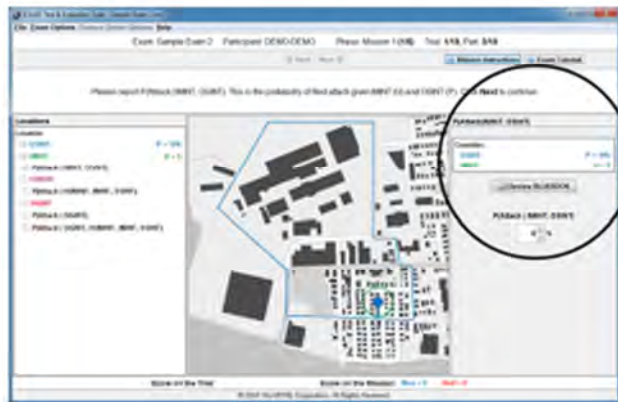
P(Attack | HUMINT, IMINT, OSINT)

Next you receive a HUMINT report of Red's attack capability. You are asked to report **P(Attack | HUMINT, IMINT, OSINT)**. This is the probability of Red attack given Red's capability (from HUMINT) and the BLUEBOOK probability (from IMINT and OSINT).

P(Attack | SIGINT)

Next you receive a SIGINT report about Red's recent activity, and you are asked to report **P(Attack | SIGINT)**. This is the probability of Red attack given **ONLY** the SIGINT report of Chatter or Silence.

P(Attack|SIGINT, HUMINT, IMINT, OSINT)

Finally, you are asked to report **P(Attack|SIGINT, HUMINT, IMINT, OSINT).** This is the probability of Red attack given all four intelligence reports.

Outcome

The computer will choose actions for Red and Blue. Both actions will be shown and points will be awarded based on the combination. In a showdown, the computer will choose a winner probabilistically using the value of P from OSINT.

Trials

Mission 1 will last for 10 trials.

Mission 2: Deciding to Divert or Not

Mission 2 requires that you play the role of Blue intelligence and operations.

On each trial, you will estimate the probability of Red attack, and then **choose a Blue action (Divert or Not Divert).**

The area of interest and behavior of Red are different for Mission 2 than for Mission 1.

The sequence of intelligence reports you will receive is similar to Mission 1. But there are some key differences as described below.

Two Red Styles

In Mission 2, the BLUEBOOK provides information about two styles of Red tactics: "passive" and "aggressive".

As controlled by the computer, Red is always either passive or aggressive in Mission 2.

The Red style will remain the same throughout the mission. But you will not know for sure which style it is.

On the first trial of Mission 2, you should assume that each style is equally likely. As the mission proceeds and you observe Red actions, you will be asked to report which style you think is most likely Red's style.

Passive:
Probability of Red Attack

	U = 2 or 3	U = 4 or 5
P > 25%	20%	30%
P ≤ 25%	40%	50%

Aggressive:
Probability of Red Attack

	U = 2 or 3	U = 4 or 5
P > 25%	50%	60%
P ≤ 25%	70%	80%

Red's Style

To begin a trial, OSINT reports P and IMINT reports U. On each trial after the first trial, you are asked to report which style you think is **most likely Red's style.**

P(Attack|IMINT, OSINT)

Next you are asked to use the BLUEBOOK and report **P(Attack|IMINT, OSINT)**, based on U (from IMINT) and P (from OSINT). This is the probability of Red attack, given IMINT and OSINT.

P(Attack|HUMINT, IMINT, OSINT)

Next you receive a HUMINT report of Red's attack capability. You are asked to report **P(Attack|HUMINT, IMINT, OSINT)**. This is the probability of Red attack given HUMINT, IMINT, and OSINT.

P(Attack|SIGINT, HUMINT, IMINT, OSINT)

Next you receive a SIGINT report about Red's recent activity. Unlike Mission 1, you are not asked to report P(Attack|SIGINT). Instead you are asked to consider all intelligence sources and report **P(Attack|SIGINT, HUMINT, IMINT, OSINT)**.

Divert or Not Divert

Finally, after estimating the Red attack probability based on all intelligence sources, you are asked to choose a Blue action – **either Divert or Not Divert**.
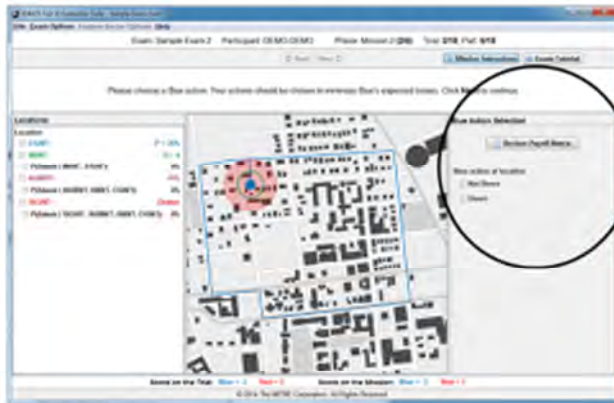
Outcome

The computer will display your Blue action, along with the Red action, and award points based on the combination. In a showdown, the computer will choose a winner probabilistically using the value of P from OSINT.

Trials

Mission 2 will last for 10 trials.

Mission 3: Defending Dual Locations

In Mission 3, there are two Blue locations to be considered on each trial. Red may attack either location, or neither location, but Red cannot attack both locations on a single trial.

You will need to **choose a Blue action (Divert or Not Divert) at each location.**

The area of interest and behavior of Red are different for Mission 3 than for Mission 2.

The sequence of intelligence reports you will receive is similar to Mission 2. But there are some key differences as described below.

Two Locations

In Mission 3, like Mission 1, Red's style is known to be "neutral" (between passive and aggressive).

However, each trial of Mission 3 presents two locations of possible Red attack.

Therefore, the BLUEBOOK values for probability of Red attack at each location are half the values of Mission 1, ranging from 10% to 40% rather than 20% to 80%.

Note: The sum of attack probabilities for the two locations is less than 100%, because Red can attack only one (or neither) location.

Probability of Red Attack at Location 1

	U = 2 or 3	U = 4 or 5
P > 25%	10%	20%
P ≤ 25%	30%	40%

Probability of Red Attack at Location 2

|	U = 2 or 3	U = 4 or 5
P > 25%	10%	20%
P ≤ 25%	30%	40%

P(Attack|IMINT, OSINT)

To begin a trial, OSINT reports P and IMINT reports U, at each of two locations. Using the BLUEBOOK, you are asked to report **P(Attack|IMINT, OSINT)** at each location. This is the probability of Red attack given IMINT (U) and OSINT (P).

P(Attack|HUMINT, IMINT, OSINT)

Next you receive a HUMINT report of Red's capability to attack. This report applies at each location, because Red can only attack one (or neither) location. You are asked to report P(Attack|HUMINT, IMINT, OSINT) at each location.
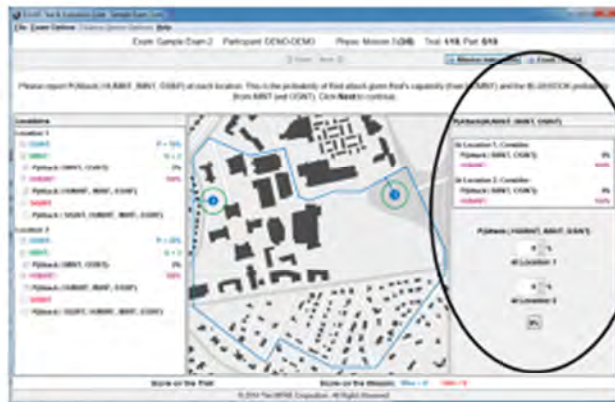
Obtaining SIGINT

After reporting P(Attack | HUMINT, IMINT, OSINT) at each location, you are asked to make a choice about SIGINT. You can only obtain SIGINT at one location, so you must choose one of the two locations.

P(Attack | SIGINT, HUMINT, IMINT, OSINT)

After receiving SIGINT, you are asked to consider all intelligence sources and report **P(Attack | SIGINT, HUMINT, IMINT, OSINT)** at each location. This is the probability of Red attack given all four intelligence sources.

Divert or Not Divert

Finally, after estimating the Red attack probability from all intelligence sources, you are asked to choose a Blue action at each location – **either Divert or Not Divert**. You can choose either action at each location.

Outcome

The computer will display your Blue actions, along with the Red actions, and award points based on the combinations. In a showdown at a location, the computer will choose a winner probabilistically using the value of P from OSINT at that location.

Trials

Mission 3 will last for 10 trials.

Mission 4: Detecting a Change in Style

In Mission 4, there is only one location of possible Red attack on each trial.

Red's style is either "passive" or "aggressive". But the **Red style will change** at some point in the mission.

There will only be one change: from passive to aggressive; or from aggressive to passive.

To help you detect Red's style and the change in style, you are allowed to create and inspect **"batch plots"** (described below) showing outcomes of previous trials.
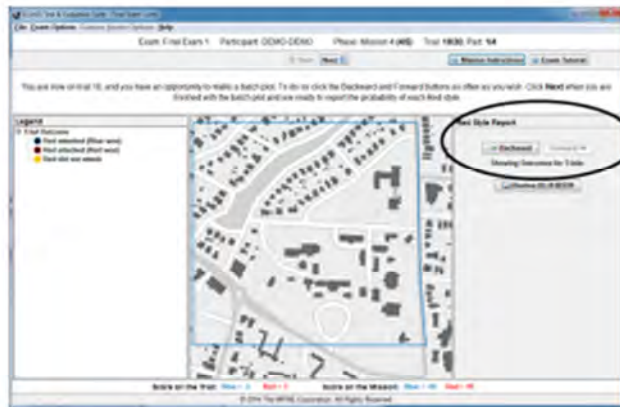
Passive:
Probability of Red Attack

	$U = 2$ or 3	$U = 4$ or 5
$P > 25\%$	20%	30%
$P \leq 25\%$	40%	50%

Aggressive:
Probability of Red Attack

	$U = 2$ or 3	$U = 4$ or 5
$P > 25\%$	50%	60%
$P \leq 25\%$	70%	80%

Batch Plot

<u>IMPORTANT:</u> Mission 4 lasts for 30 trials, and you are only allowed to make batch plots at the start of trials 10, 20, and 30.


Creating a Batch Plot

At the start of trials 10, 20, and 30, the display will change to show a color-coded legend of dot types (circled below on the left of the screen) along with "Backward" and "Forward" buttons (circled below on the right of the screen).

Display Outcomes (Backward)

To display outcomes, click the "Backward" button. Each click "Backward" displays an additional outcome, moving backward one trial per click.
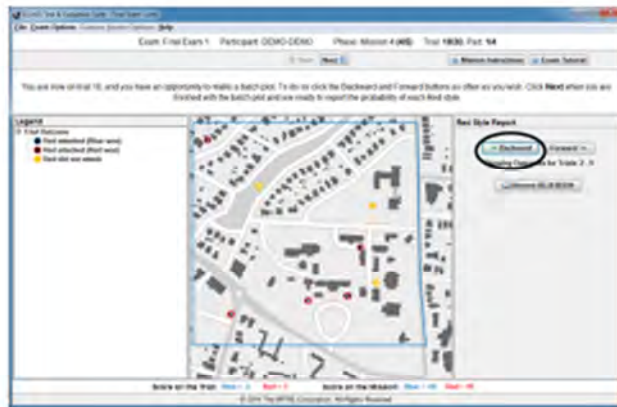

Display Outcomes (Forward)

Clicking "Forward" clears all but the last outcome displayed (going backward). From that outcome, each click "Forward" displays an additional outcome, moving forward one trial per click.

Display Outcomes (Backward and Forward)

Clicking "Backward" clears all but the last outcome displayed (going forward). From that outcome, each click "Backward" displays an additional outcome, moving backward one trial per click. You can go "Backward" or "Forward" as often as you wish.
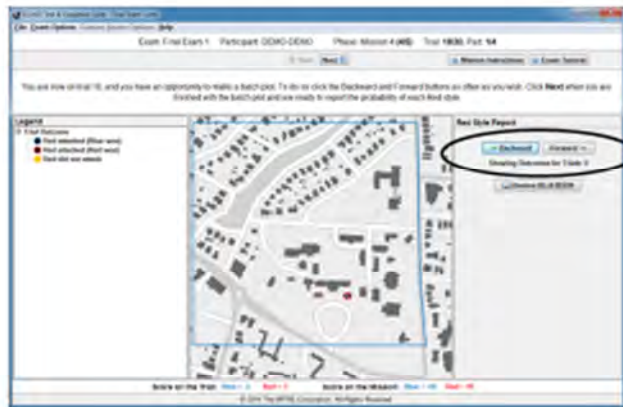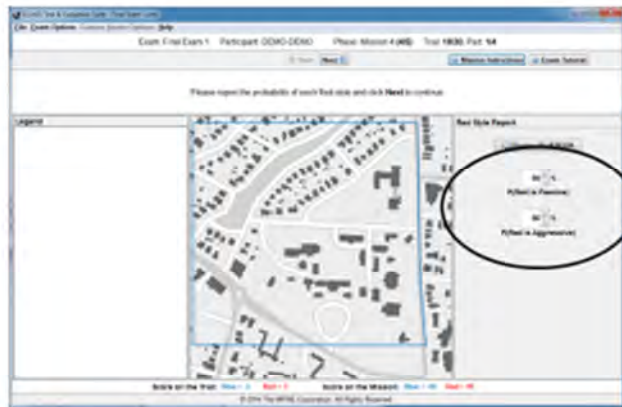

Inspecting a Batch Plot

Whenever dots are displayed, you can click on a dot to obtain further details. These details include the trial number as well as all the intelligence reports that were provided during the trial. Details can only be displayed for one dot at a time.

Red's Style

After creating and inspecting a batch plot (or not), you are asked to report the **probability of each Red style** for the current trial. Note that the "default" response is whatever you indicated on the previous trial, which the computer will remember.
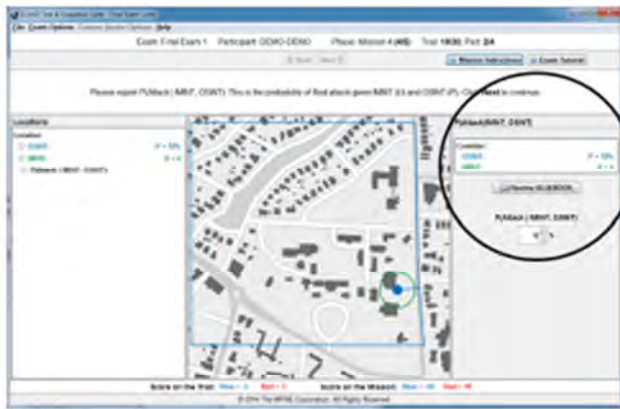

Current Trial

After you report the probability of each Red style, a batch plot will be replaced by the standard display for the current trial.

P(Attack|IMINT, OSINT)

Each new trial provides intelligence reports from two sources: OSINT and IMINT. You are asked to consider both intelligence sources and report: **P(Attack|IMINT, OSINT)**. This is the probability of Red attack given IMINT and OSINT.

Divert or Not Divert

Then you are asked to choose a Blue action – **either Divert or Not Divert.**

Outcome

The computer will display your Blue action (Divert or Not Divert), along with the Red action (Attack or Not Attack), and award points based on the combination.

Trials

IMPORTANT: Mission 4 lasts for 30 trials, and you are only allowed to make batch plots at the start of trials 10, 20, and 30.

Mission 5: Detecting a Different Change

Mission 5 is similar to Mission 4, except the Red styles are different.

One style is called "P-sensitive", because the Red attack probability depends only on P. The other style is called "U-sensitive", because the Red attack probability depends only on U.

There will be one change during Mission 5: from P-sensitive to U-sensitive; or from U-sensitive to P-sensitive.

To help you detect Red's style and the change in style, you are allowed to create and inspect **"batch plots"** showing outcomes of previous trials.

P-sensitive:
Probability of Red Attack

	U = 2 or 3	U = 4 or 5
P > 25%	40%	40%
P ≤ 25%	60%	60%

U-sensitive:
Probability of Red Attack

	U = 2 or 3	U = 4 or 5
P > 25%	20%	80%
P ≤ 25%	20%	80%

Batch Plot

<u>IMPORTANT:</u> **Mission 5 lasts for 40 trials, and you are only allowed to make batch plots at the start of trials 10, 20, 30, and 40.** The procedure for making batch plots is the same as in Mission 4.

Red's Style

After creating a batch plot (or not on most trials), you are asked to report the **probability of each Red style** for the current trial. Note that the "default" response is whatever you indicated on the previous trial, which the computer will remember.

Current Trial

After you report the probability of each Red style, a batch plot will be replaced by the standard display for the current trial.

P(Attack | IMINT, OSINT)

Each new trial provides intelligence reports from two sources: OSINT and IMINT. You are asked to consider both intelligence sources and report: **P(Attack | IMINT, OSINT)**. This is the probability of Red attack given IMINT and OSINT.

Divert or Not Divert

Then you are asked to choose a Blue action – **either Divert or Not Divert.**

Outcome

The computer will display your Blue action (Divert or Not Divert), along with the Red action (Attack or Not Attack), and award points based on the combination.

Trials

IMPORTANT: Mission 5 lasts for 40 trials, and you are only allowed to make batch plots at the start of trials 10, 20, 30, and 40.