

460 A Property Analysis

461 A.1 Proof of Lemma 3.3

462 *Proof.* We can rewrite \mathcal{J} to

$$\mathcal{J} = \max_{\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{A}(\mathbf{x})} \left\{ \frac{1}{m} \sum_{i=1}^m P(\mathbf{x}, \mathbf{a}_i) G(\mathbf{x}, \mathbf{a}_i) + \frac{1}{m} \sum_{i=1}^m \min_{\lambda_i \geq 0} \lambda_i H(\mathbf{x}, \mathbf{a}_i) + \gamma R(\{S_{\mathcal{M}}(\mathbf{x}, \mathbf{a}_i)\}_{i=1}^m) \right\}. \quad (11)$$

463 We derive the fact that, for any i ,

$$\min_{\lambda_i \geq 0} \lambda_i R(\mathbf{x}, \mathbf{a}_i) = \begin{cases} -\infty & H(\mathbf{x}, \mathbf{a}_i) < 0 \\ 0 & \text{otherwise} \end{cases}$$

464 where $-\infty$ comes from setting $\lambda_i = \infty$ and 0 is obtained by setting $\lambda_{p_i} = 0$. By the linearity of
465 summation, we can further derive

$$\frac{1}{m} \sum_{i=1}^m \min_{\lambda_i \geq 0} \lambda_i R(\mathbf{x}, \mathbf{a}_i) = \begin{cases} -\infty & \exists i, H(\mathbf{x}, \mathbf{a}_i) < 0 \\ 0 & \text{otherwise} \end{cases}.$$

466 That is, if any constraint for the robustness is unsatisfied, the dual player will minimize the objective
467 towards $-\infty$; however, the primal player cannot optimize towards ∞ given that the limit of the gain
468 function and the diversity are finite. In other words, if the constraints are satisfied, the primal player
469 can freely optimize the objective. Once $H(\mathbf{x}, \mathbf{a}_i) \geq 0, \forall \mathbf{a}_i$ are satisfied, the objective becomes

$$\begin{aligned} \tilde{\mathcal{J}} &:= \max_{\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{A}^+(\mathbf{x})} \frac{1}{m} \sum_{i=1}^m P(\mathbf{x}, \mathbf{a}_i) G(\mathbf{x}, \mathbf{a}_i) + \gamma R(\{S_{\mathcal{M}}(\mathbf{x}, \mathbf{a}_i)\}_{i=1}^m) \\ &\geq \min_{\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{A}^+(\mathbf{x})} \frac{1}{m} \sum_{i=1}^m P(\mathbf{x}, \mathbf{a}_i) G(\mathbf{x}, \mathbf{a}_i) + \gamma R(\{S_{\mathcal{M}}(\mathbf{x}, \mathbf{a}_i)\}_{i=1}^m) \\ &> 0 \end{aligned} \quad (12)$$

470 as $P(\mathbf{x}, \mathbf{a}_i) > 0$ and $G(\mathbf{x}, \mathbf{a}_i) \geq 0$ for any $\mathbf{a}_i \in \mathbb{A}^+(\mathbf{x})$; also, $R \geq 0$ holds. We conclude the proof
471 here. \square

472 A.2 A Probabilistic Relaxation of Robustness

473 Absolute robustness is difficult to guarantee, and common practice is to relax this via a probabilistic
474 approach [15].

475 Assume there is a distribution over the sample space $\mathbb{B}_s(\mathbf{x}, \mathbf{a})$ denoted by $\Pr(\mathbb{B}_s(\mathbf{x}, \mathbf{a}))$. We write
476 $\mathbf{x}' \sim \Pr(\mathbb{B}_s(\mathbf{x}, \mathbf{a}))$ to indicate that \mathbf{x}' is sampled from the set $\mathbb{B}_s(\mathbf{x}, \mathbf{a})$ under the distribution P . Let
477 $\mathbb{E}[h(\mathbf{x}')|\mathbf{x}, \mathbf{a}]$ denote the expectation of \mathbf{x}' in this configuration. Hence, we modify Equation (6) to

$$\mathbb{E}[h(\mathbf{x}')|\mathbf{x}, \mathbf{a}] > \tilde{\psi}, \quad (13)$$

478 where $\tilde{\psi}$ is a function that adjusts the base score threshold ψ . It is crucial to have this threshold
479 function in order to consider the variance of scores in the neighbor set. Particularly, we would like
480 most neighbors to remain in a similarly "good" state, with low variance between them.

481 Moreover, we explicitly impose $h(S_{\mathcal{M}}(\mathbf{x}, \mathbf{a})) - \psi > 0$. It places a hard constraint to avoid the
482 case in which the neighbors of the semi-factual are robust, but the "semi-factual" itself has crossed
483 the decision boundary to become a counterfactual. Whilst somewhat unlikely, this situation is
484 theoretically possible, and requires consideration. In this case, H is re-written as H_p , which represents
485 a combination of (i) the probabilistic robustness, and (ii) the absolute robustness for the semi-factual
486 H_a such that:

$$H_p(\mathbf{x}, \mathbf{a}) = \mathbb{E}[h(\mathbf{x}')|\mathbf{x}, \mathbf{a}] - \tilde{\psi} \quad H_a(\mathbf{x}, \mathbf{a}) = h(S_{\mathcal{M}}(\mathbf{x}, \mathbf{a})) - \psi. \quad (14a)$$

487 In practice, H_p is still non-trivial to solve. Monte Carlo (MC) sampling is a common strategy to
488 apply here such that, by sampling a fixed sized batch $\mathbf{B} = \{\mathbf{x}' : \mathbf{x}' \sim \Pr(\mathbb{B}_s(\mathbf{x}, \mathbf{a}))\}$,

$$H_p(\mathbf{x}, \mathbf{a}) = \mathbb{E}[h(\mathbf{x}')|\mathbf{x}, \mathbf{a}] - \tilde{\psi} \approx (1/|\mathbf{B}|) \sum_{\mathbf{x}' \in \mathbf{B}} h(\mathbf{x}') - \tilde{\psi}. \quad (15)$$

489 This implies that we substitute an unbiased estimator for the population mean.

490 **B Actionability Constraints**

491 **B.1 Non-Causal**

492 Here we define the actionability constraints used in the various domains. It may be assumed that the
493 direction features are allowed to change corresponds with *positive gain*. We use various sized “action
494 sets” to fully test all algorithms in various setups. The German Credit data used 15 actionable features
495 to be closely in line with Mothilal et al. [28] whom allowed all features to be mutable. However, we
496 also used 7 on Lending Club, and 4 on Adult Census/Breast Cancer to test the algorithms in situations
497 with smaller action spaces also for completeness.

498 We ordered categorical features in a sensible fashion to “direct” semi-factual “even if” thinking, and
499 when we say a categorical feature could decrease/increase, we are referring to this pre-defined order.
500 If you are interested in the exact ordering, please refer to our code which contains all the lists, but
501 here we summarize. In reality however, a user must specify their exact actionability constraints, what
502 we have specified here is designed to be representative what is possible for the “average” individual.

503 **B.1.1 German Credit Dataset**

504 The continuous features used were ‘duration’, ‘amount’, ‘age’, the categorical ones were
505 ‘status’, ‘credit_history’, ‘purpose’, ‘savings’, ‘employment_duration’, ‘installment_rate’, ‘per-
506 sonal_status_sex’, ‘other_debtors’, ‘present_residence’, ‘property’, ‘other_installment_plans’, ‘housing’,
507 ‘number_credits’, ‘job’, ‘people_liable’, ‘telephone’, ‘foreign_worker’. As actionable features
508 for semi-factual recourse, we considered the following:

- 509 • *duration*: We allowed people to increase the duration of their loan.
- 510 • *amount*: We allowed people to increase the amount of their loan.
- 511 • *status*: We allowed people to move towards having lower status.
- 512 • *credit_history*: We allowed people to move towards e.g. having a late payment if their credit
513 history was otherwise good.
- 514 • *savings*: This feature was allowed to decrease.
- 515 • *employment_duration*: This feature was allowed to decrease in case people wanted to e.g.
516 start a new job.
- 517 • *installment_rate*: This feature was allowed to move towards lower payments.
- 518 • *other_debtors*: this feature was allowed to add another co-applicant.
- 519 • *present_residence*: This feature was allowed to move towards e.g. renting in case the user
520 desired to do so whilst searching for a new house with their loan.
- 521 • *property*: this feature was allowed to move towards having no property in case the user
522 desired to sell their house/car etc to help pay for e.g. a downpayment.
- 523 • *other_installment_plans*: This feature was allowed to add other installment plans.
- 524 • *housing*: this feature was allowed to move towards renting away from e.g. owning.
- 525 • *number_credits*: This feature was allowed to increase if the user desired to acquire more
526 credit cards.
- 527 • *job*: this feature was allowed to decrease in case the individual desired to get a different,
528 less demanding job within their institution, or indeed quite their job to e.g. start a business.
- 529 • *people_liable*: This feature was allowed to move towards more people being liable.

530 **B.1.2 Lending Club**

531 The continuous features used were ‘loan_amnt’, ‘pub_rec_bankruptcies’, ‘annual_inc’, ‘dti’, the
532 categorical ones were ‘emp_length’, ‘term’, ‘grade’, ‘home_ownership’, ‘purpose’. As actionable
533 features for semi-factual recourse, we considered the following:

- 534 • *home_ownership*: This feature was allowed to decrease towards e.g. renting.
- 535 • *annual_inc*: this feature was allowed to decrease if the person desired to e.g. work less
536 hours.

- 537 • *emp_length*: This feature was allowed to decrease in case the individual desired to change
538 careers.
- 539 • *dti*: dept to income ratio, this feature was allowed to increase.
- 540 • *pub_rec_bankruptcies*: This feature was allowed to increase in case the user decided they
541 wanted to declare bankruptcy to e.g. try and keep some assets.
- 542 • *loan_amnt*: this feature was allowed to increase.
- 543 • *term*: This feature was allowed to decrease.

544 B.1.3 Breast Cancer

545 The continuous features used were none, the categorical ones were ‘agegrp’, ‘density’, ‘race’,
546 ‘Hispanic’, ‘bmi’, ‘agefirst’, ‘nrelbc’, ‘brstproc’, ‘lastmamm’, ‘surgmeno’, ‘hrt’. As actionable
547 features for semi-factual recourse, we considered the following:

- 548 • *bmi*: This feature was allowed to move towards less healthy BMI levels in case the patient
549 e.g. has hypothyroidism.
- 550 • *brstproc*: this feature was allowed to move towards having had a previous breast procedure
551 in case the patient would like to do so or was advised.
- 552 • *hrt*: This feature was allowed to move towards starting HRT, in case a person may wish to
553 alleviate symptoms of the menopause.
- 554 • *agegrp*: this feature was allowed to get older in case the individual would like to take no
555 action confident that it would not lead to cancer in the next few years/decades.

556 B.2 Causal

557 In the causal setting, we allowed a user’s age to increase a maximum of 5 years to mimic the
558 motivating examples in the paper about a user having a bank loan accepted. In such a situation, the
559 user may want to e.g. work less hours over the next 5 years whilst they repay the loan, and still have
560 it accepted.

561 Next, we detail the direction features are allowed to change, and what direction corresponds to
562 *positive gain*.

563 B.2.1 Adult Income Census

564 We use the features “sex”, “age”, “native-country”, “marital-status”, “education-num”, “hours-per-
565 week”, which are the variables in the causal graph of Nabi & Shpitser [30]. We consider “age”
566 and “hours-per-week” as actionable. We allow “age” to increase a maximum of five years, and
567 “hours-per-week” to decrease.

568 For positive gain, we considered: Age, marital status, and education-num *increasing* corresponding to
569 positive gain, and hours-per-week *decreasing* corresponding to positive gain. A persons sex was seen
570 as neutral gain.

571 B.2.2 COMPAS

572 We use the features “age”, “race”, “sex” and “priors count”, which are the variables in the causal
573 graph of Nabi & Shpitser [30]. We consider “age” and “priors count” as actionable. As actionability
574 constraints, we assume that both features are non-negative and can only be increase. Age specifically
575 is only allowed to increase by 5 years for each individual.

576 For positive gain, we considered: Age and priors count increasing corresponding to positive gain. A
577 persons sex and race was seen as neutral gain.

578 C Hyperparameter Choices

579 C.1 Non-Causal

580 Here we note the values for the hyperparameters used in our demonstrations. All were obtained
581 though pilot grid-searches across each dataset. The hyperparameter choices are summarized in Table 1

Table 1: Hyperparameter Specifications

Data	λ_p	λ_s	γ_d	γ_p
German credit	30	10	1	$1e^{-1}$
Lending Club	30	10	1	$1e^{-1}$
Breast Cancer	10	10	10	$1e^{-1}$

582

583 For S-GEN itself, we used the same hyperparameters everywhere outside of the above table. The
584 number of generations spent searching for a solution was 20. The population size was fixed at
585 [12,24,48,72,96,120], for diversity sizes of [1,2,4,6,8,10], respectively. The mutation rate was 0.05.
586 The number of “elite” solutions passed on for each generation was 4. The probability of a crossover
587 happening was 0.5. The number of Monte Carlo trials for each instance was 100. The continuous
588 features were perturbed (in mutation or population initialization) by the output from sampling a
589 standard normal distribution with standard deviation equal to the max actionable feature value, minus
590 the min actionable feature value, multiplied by 0.05.

591 C.2 Causal

592 In our causal tests we chose λ as 1.0, and this was gradually decreased by a momentum of $\eta=0.9$
593 each iteration to put more emphasises on the maximization of gain.

594 D Algorithm Pseudocode

Algorithm 1 S-GEN: Genetic Algorithm to Generate Semi-Factual Recourse with Robustness and Diversity in a Non-Causal Model Agnostic Setting

Require: \mathbf{x} the user feature

Require: $h(\cdot)$ the predictive model

Require: m the expected number of suggestions

Require: n the number of candidates, $n > m$

Ensure: \mathbf{R}_{SF} the set of semi-factual(s)

- 1: Sample n candidates $\mathbf{D} \leftarrow \{\mathbf{x}_i \sim \mathbb{X}^*\}_{i=1}^n$
 - 2: **while** the stopping criterion is not satisfied **do**
 - 3: Obtain the fitness scores \mathbf{f} with respect to \mathbf{D}
 - 4: Save the fittest $\mathbf{x}^* \in \mathbf{D}$ according to \mathbf{f}
 - 5: Let \mathbf{D} evolve by *natural selection* according to \mathbf{f} , *crossover*, *mutation*, and *elitism* with \mathbf{x}^*
 - 6: **end while**
 - 7: Collect the best m unique candidates from $\{\mathbf{x}' \in \mathbf{D} : h(\mathbf{x}') = h(\mathbf{x}) = 1\}$ to \mathbf{R}_{SF} , according to the corresponding fitness scores in \mathbf{f}
 - 8: **if** $|\mathbf{R}_{SF}| < m$ **then**
 - 9: Complement \mathbf{R}_{SF} to m elements with \mathbf{x}' randomly drawn from \mathbf{R}_{SF}
 - 10: **end if**
-

595 E Code

596 For our full code used please see:

597 https://anonymous.4open.science/r/NeurIPS_2023-9F62/README.md

598 The ability to reproduce the results is given.

Algorithm 2 S-GEN: Algorithm to Generate Robust & Diverse Causal Semi-Factual Explanations for Differentiable Classifiers

Require: \mathbf{x} the user feature vector

Require: $h(\cdot)$ the predictive model

Require: \mathcal{M} the **differentiable** SCM

Require: ϵ the epsilon robustness

Require: η the momentum parameter

Require: τ the learning rate

Require: $\text{Proj}(\cdot)$ a projection that ensures the action is actionable

Ensure: \mathbf{R}_{SF} the set of semi-factual(s)

```
1:  $\mathbf{R}_{SF} \leftarrow \emptyset$ 
2:  $i \leftarrow 0$ 
3: for  $\mathbf{a} \in \mathbb{A}(\mathbf{x})$  do
4:   {Check if the initial action  $\mathbf{a}$  itself is a valid semi-factual}
5:   Sample a batch of neighbours from  $\mathbb{B}_s(\mathbf{x}, \mathbf{a})$ , denoted by  $\mathcal{B}$ 
6:   if  $h(S_{\mathcal{M}}(\mathbf{x}, \mathbf{a})) = 0$  or  $h(\mathbf{x}') = 0, \exists \mathbf{x}' \in \mathcal{B}_i$  then
7:     break
8:   end if
9:    $\mathbf{a}_i \leftarrow \mathbf{a}$ 
10:  while not converged do
11:    Sample a batch of neighbours from  $\mathbb{B}_s(\mathbf{x}, \mathbf{a}'_i)$ , denoted by  $\mathcal{B}_i$ 
12:    if  $h(S_{\mathcal{M}}(\mathbf{x}, \mathbf{a}'_i)) = 0$  or  $h(\mathbf{x}') = 0, \exists \mathbf{x}' \in \mathcal{B}_i$  then
13:      break
14:    end if
15:     $\mathbf{a}_i \leftarrow \mathbf{a}'_i$ 
16:     $\mathcal{J}_i \leftarrow -\lambda_i \mathcal{L}(h(\mathbf{x}'_i), h(\mathbf{x})) - \sum_{\mathbf{x}'_i \in \mathcal{B}_i} \frac{\lambda_i}{|\mathcal{B}_i|} \mathcal{L}(h(S_{\mathcal{M}}(\mathbf{x}, \mathbf{a}_i)), h(\mathbf{x})) + \hat{P}(\mathbf{x}, \mathbf{a}_i) \hat{G}(\mathbf{x}, \mathbf{a}_i)$ 
17:     $\mathbf{a}'_i \leftarrow \text{Proj}(\mathbf{a}_i + \tau \nabla_{\mathbf{a}_i} \mathcal{J}_i)$ 
18:     $\lambda_i \leftarrow \eta \lambda_i$ 
19:  end while
20:   $\mathbf{R}_{SF} \leftarrow \mathbf{R}_{SF} \cup \{S_{\mathcal{M}}(\mathbf{x}, \mathbf{a}_i)\}$ 
21:   $i \leftarrow i + 1$ 
22:  if  $i \geq m$  then
23:    break
24:  end if
25: end for
```

599 F Individual Dataset Results

600 The results are presented in Figure [4](#).

601 G Baselines

602 G.1 Non-Casual

603 **DiCE** Our modification to DiCE, starts by generating a counterfactual(s) for a query. Next, we
604 use the algorithm again, but on the generated counterfactual(s), to make them generate a second
605 counterfactual, which goes back over the decision boundary. In effect, this generates a semi-factual(s)
606 for a query.

607 **PIECE** Second, we use the PIECE framework by Kenny and Keane [\[21\]](#), but apply it to tabular
608 data. Following the authors, we divide the training data into two sets, the first corresponding to
609 those predicted as the original class c , and the second to those predicted as the counterfactual class
610 c' , these are again split into the respective features. Hence, if there are 2 classes, with 4 features,
611 there are $2 \times 4 = 8$ sets of data. These sets were then modeled using the best fit found for a Beta
612 distribution on continuous features, and a simple Categorical distribution for categorical features. To
613 generate a semi-factual predicted as c , we take the probability of each feature value in the query using
614 the models of the counterfactual class c' , and modify each to be its expected statistical value in c'

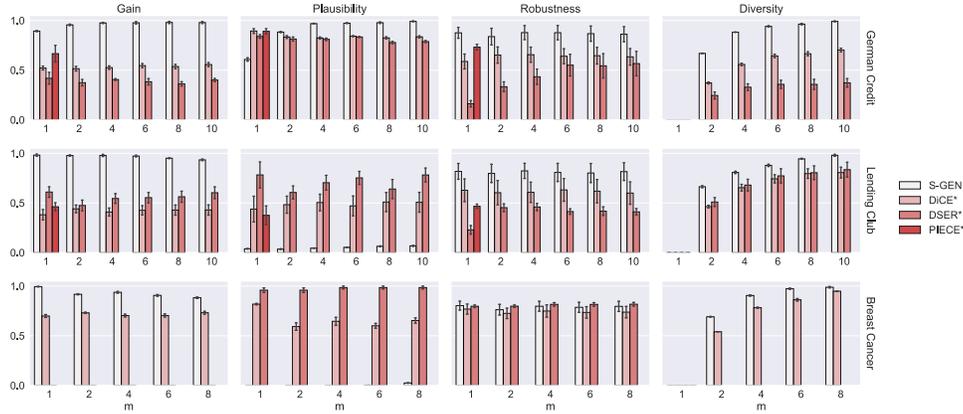


Figure 4: Results: The ability of S-GEN to create semi-factuals is compared to DiCE* and PIECE*. Overall, S-GEN does the best, achieving significantly better results to both baselines on 11/16 tests. Moreover, S-GEN was only significantly worse than either baseline on a single test (i.e., plausibility on German Credit), with the remaining four tests being competitive between methods. Standard error bars are shown.

615 one-by-one (from the lowest probability to the highest), until the next would take it over the decision
 616 boundary. In the case of continuous features, as done by (author?) [21], we take the probability as
 617 being the minimum of the two integrals either side of the feature value in the distribution. In the
 618 case the expected feature values lie outside the actionability range, we clip them to the closest value
 619 allowed.

620 **DSER** For Diverse Explanation of Reject [1] (DSER) we had to modify the the technique in two
 621 main ways. Most notably, the techniques doesn't deal with categorical features, so to overcome this,
 622 we optimized treating all one hot encoded features as real-valued, and then projected each categorical
 623 feature onto its nearest value. Next, the method addresses diversity by iterating all different sets
 624 of possible features, in our domains this is computationally intractable. Hence, we optimize one
 625 semi-factual at a time, each time pushing each solution as far as possible from those already found.

626 G.2 Causal

627 **Karimi et al. (2021)** The method by Karimi et al. [19] is a recourse method designed to minimize
 628 cost whilst traversing the decision boundary. To modify the technique, we simply stop the optimization
 629 when the next step would take it over the decision boundary.

630 **Dominguez et al. (2022)** The method by Dominguez et al. [15] is identical to Karimi et al. [19],
 631 but they add in a robustness component. Namely, they take an individual x , and solve an inner loss
 632 which means that an individual of distance $\epsilon = 0.1$ (in our tests) close to x , with the same recourse
 633 given, will also achieve recourse. We simply keep the same optimization process, but aim to solve
 634 a different objective. The objective we solve is to move towards the decision boundary, but when
 635 the recourse option causes either x or the individual close to it to cross the decision boundary, we
 636 terminate the optimization one step prior to this.

637 H Computational Costs

638 All tests were run on a MacBook Pro, Apple M1 Pro, 16 GB. Re-running tests will take less than 1
 639 day.

640 I User Study

641 Here we show our entire user study for complete transparency. We used the German Credit dataset,
642 but converted the currency into U.S. dollars since it was given to U.S. citizens to complete.

Intro Brief

Thank you for clicking on this study!

Do no take this study on a mobile phone, the tables and images wont display correctly.

Please don't take this study if you did a similar one recently.

You are free to leave at any time.

The study will take around 8mins.

You will be paid \$12 per hour for your efforts.

Thank you for your participation!

Enter ID

Please Enter Your Prolific ID Here

Introduction

Introduction



You are going to be shown **six situations** in which a person either has a loan application **approved**, or **rejected**.

You will then be shown **two** different pieces of information for each situation that a bank clerk *could tell* the person.

You are then asked to rate how **useful** each of these are. That is, could the information possibly be useful in any way? Or is it not that useful?

Each situation has 4 "features".

Feature Explanation

Features Used in Decision

The four "features" used to decide if each person has their loan **approved** or **rejected** are:

1: Duration: Over how long the applicant wishes to pay back the loan.

2: Amount: How much are they asking to loan from the bank.

3: Savings: How much money does the applicant have saved.

4: Credit Cards: How many credit cards does the applicant have.

These are the only features the bank clerk uses to make decisions.

Click Next

Now, Please Practice On The Next Question

Sample Question

Example Question

Lucas had his bank loan **accepted**, his features are:

Duration	12 Months
Amount	\$2,000
Savings	\$500
Credit Cards	2

The two possible things the bank clerk could tell him are:

Option 1: Even if you want to increase your **Duration** to 14 months, and **Amount** to \$3,000, we will still **accept** your loan application.

Option 2: If your **Savings** were \$100, and your **Credit Cards** 5, we would have **rejected** your loan application.

How **useful** is each option?

	Not Useful				Very Useful
Option 1	<input type="radio"/>				
Option 2	<input type="radio"/>				

Click Next 2

Please only participate in this study if you understand the instructions well

Block 15

Remember, the key question is how USEFUL is each option

Click Next 3

Click Next To Begin The Study

Question 1

Kate had her bank loan **accepted**, her features are:

Duration	6 Months
Amount	\$932
Savings	Over \$1000
Credit Cards	2-3

The two possible things the bank clerk could tell her are:

Option 1: Even if you want to increase your **Amount** to \$2,841, and increase your number of **Credit Cards** to 4-5, we will still **accept** your loan application.

Option 2: If your **Duration** was 44 months, and you had had **Savings** less than \$100, we would have **rejected** your loan application.

How **useful** is each option?

	Not Useful				Very Useful
Option 1	<input type="radio"/>				
Option 2	<input type="radio"/>				

Question 2

Paul had his bank loan **accepted**, his features are:

Duration	18 Months
Amount	\$1,239
Savings	Over \$1,000
Credit Cards	1

The two possible things the bank clerk could tell him are:

Option 1: Even if you want to increase your **Duration** to 21 Months, and lower your **Savings** to \$500-\$1,000, we will still **accept** your loan application.

Option 2: If you asked for an **Amount** of \$15,499 (or more), and had had 6 **Credit Cards**, we would have **rejected** your loan application.

How **useful** is each option?

	Not Useful				Very Useful
Option 1	<input type="radio"/>				
Option 2	<input type="radio"/>				

Question 3

Xue had her bank loan **accepted**, her features are:

Duration	9 Months
Amount	\$1549
Savings	Over \$1,000
Credit Cards	1

The two possible things the bank clerk could tell her are:

Option 1: Even if you want to increase your **Duration** to 25 Months, and increase your **Amount** to \$4,620, we will still accept your loan application.

Option 2: If you had had 3 **Credit Cards**, and no **Savings**, we would have rejected your loan application.

How useful is each option?

	Not Useful				Very Useful
Option 1	<input type="radio"/>				
Option 2	<input type="radio"/>				

Question 4

Siddarth had his bank loan rejected, his features are:

Duration	48 Months
Amount	\$6,143
Savings	None
Credit Cards	2-3

The two possible things the bank clerk could tell him are:

Option 1: Even if you increase your **Savings** to \$100, and lower your number of **Credit Cards** to 1, we will still reject your loan application.

Option 2: If you lower your **Duration** to 15 Months, and lower your **Amount** to \$4,627, we will accept your loan application.

How useful is each option?

	Not Useful				Very Useful
Option 1	<input type="radio"/>				
Option 2	<input type="radio"/>				

Question 5

Camila had her bank loan **rejected**, her features are:

Duration	60 Months
Amount	\$15,653
Savings	None
Credit Cards	2-3

The two possible things the bank clerk could tell her are:

Option 1: Even if you increase your **Duration** to 70 Months, and reduce your number of **Credit Cards** to 1, we will still **reject** your loan application.

Option 2: If you reduce your **Amount** to \$7,296 (or less), and you get \$1,000+ **Savings**, we will **accept** your loan application.

How **useful** is each option?

	Not Useful				Very Useful
Option 1	<input type="radio"/>				
Option 2	<input type="radio"/>				

Question 6

Angelo had his bank loan **rejected**, his features are:

Duration	60 Months
Amount	\$7,408
Savings	Less than \$100
Credit Cards	2

The two possible things the bank clerk could tell him are:

Option 1: Even if you decrease your **Amount** to \$6,505, and increase your **Savings** to over \$1000, we will still **reject** your loan application.

Option 2: If you lower your **Duration** to 5 Months, and you reduce your number of **Credit Cards** to 1, we will **accept** your loan application.

How **useful** is each option?

	Not Useful					Very Useful
Option 1	<input type="radio"/>					
Option 2	<input type="radio"/>					

Debrief

Debrief: You Have Reached The End

Thank you for your participation, this study was designed to evaluate what kind of explanation people prefer from an artificial intelligence system.