A machine learning approach for probing the unconscious:
evidence for unconscious lie detection being more accurate

Martin Spüler

Wilhelm-Schickard-Institute for Computer Science, University of Tübingen, Tübingen,
Germany

Abstract

This paper demonstrates a machine learning approach to the analysis of reaction times and shows how this approach can be used to probe different cognitive processes on a single-trial level. The probing method allows to measure how a stimulus is perceived, and when combined with priming, it allows to probe unconscious processes. The method is demonstrated on a publicly available dataset of a lie detection experiment, which provides evidence for unconscious lie detection being more accurate than conscious lie detection. It is further shown how this method can be used to predict human behavior, as a participant can be unconsciously probed to predict how the participant will respond consciously to a stimulus in the future.

Keywords: machine learning, probabilistic, deception detection, consciousness

A machine learning approach for probing the unconscious: evidence for unconscious lie detection being more accurate

Introduction

The idea of unconscious processes influencing human behavior dates back more than 300 years to philosophers like Leibnitz (Merikle and Reingold, 1992) and became popular by Freud's theory of the unconscious mind. Centuries later, studying an unconscious process still remains a major challenge (Timmermanns and Cleeremans, 2015): How to objectively measure what another person perceives unconsciously, when by definition that person is not aware of it? One way to demonstrate unconscious awareness is by priming, in which the processing of a target stimulus or a behavioral measure is influenced by the prior presentation of a prime stimulus, which is not consciously perceived by the participant (Merikle and Reingold, 1992). While the existence of the unconscious seems to be generally accepted (Schmidt, 2007), the question of how to demonstrate those processes and what they are capable of is still highly disputed and so far it is not possible to measure an unconscious process on a single-trial level.

One study investigating the unconscious process of lie detection was performed by ten Brinke et al. (2014) who showed videos of a lying or truth-telling suspect to participants and analyzed the reaction times in a word categorization task, where the participants were primed with an image of the person from the video. Ten Brinke et al. (2014) found a significant difference in the response time for congruent and incongruent trials, which led to them to conclude that unconscious lie detection is more accurate than conscious detection. This work sparked some comments that criticized the study (Levine and Bond, 2014) and did not agree with the conclusion.

In particular, Franz and von Luxburg (2015) argued that the unconscious prediction accuracy needs to be assessed and be greater than the conscious prediction accuracy to be able to claim that the unconscious process is more accurate. To underpin their argument they performed a classification of the congruency based on the reaction times, which resulted in classification accuracies around chance level and they thereby concluded that there is no evidence for accurate unconscious lie detection. While their general argument is important, their conclusion is only valid for linear classifiers.

In the course of this paper, a probabilistic approach to data analysis and a non-linear probabilistic classifier will be used to demonstrate that the data from (ten Brinke et al., 2014) actually shows evidence for unconscious lie detection being more accurate than conscious lie detection.

Probabilistic data analysis: beyond mean and standard deviation

When analyzing behavioral data, usually the mean of one condition is compared to the mean of the other condition. This was done by ten Brinke et al. (2014) who used a t-test to show that the mean of congruent trials is different to the mean of incongruent trials ($p=.027$, $t(65) = 2.26$). Similarly, Franz and von Luxburg (2015) used a linear classifier to test if the reaction time is above or below the mean. However, if the data does not follow a certain distribution, only looking at the mean can ignore effects in the data. In the following, a probabilistic approach to data analysis is introduced that goes beyond mean and standard deviation and thereby allows to uncover effects in the data that were previously overlooked.

Figure 1. Simplified depiction of experiment 2 performed by (ten Brinke et al., 2014) with regard to one suspect video: **A:** 6 Blocks were performed and in each block the participant watched two videos. **B:** After the videos, the participant performed 64 test trials which consisted of a word categorization task preceded by a masked prime consisting of the suspects face from one of the videos. Each face was used in 32 trials. **C:** After the test trials, the suspects faces were presented to the participant to judge the veracity of the suspects

Data

For this paper, the data of experiment 2 from (ten Brinke et al., 2014) were reanalyzed. Ten Brinke and colleagues performed an experiment where a participant watched a pair of videos with a suspect being interrogated and the suspect either lied or told the truth (see figure 1). After the videos, the participants completed 32 test trials for each video in which a target word was displayed and the participants had to rate if this word belongs to a truth-category (e.g. "honest") or to the lie-category (i.e. "deceitful"). Before the word presentation, a masked prime was presented which consisted of the face of the suspect from the video. The word was either congruent to the prime (both truth or lie) or incongruent (one truth, one lie). After the test trials, the suspects face was presented to the participant and they had to decide if the

suspect lied or told the truth in the video. This process was repeated 6 times for a total of 12 videos. The average accuracy of the veracity judgments of the participants was 49.62 % and thereby at chance level.

While ten Brinke and colleagues made the data publicly available, Franz and von Luxburg (2015) reanalyzed that data and made their analysis scripts publicly available including the dataset from (ten Brinke et al., 2014) in a different data format. So for this paper, the data from ten Brinke et al. was used in the data format provided by Franz and von Luxburg¹.

Due to inconsistencies in the participant identifiers only data from 65 participants could be used. Therefore the results in this paper can be different (e.g. response accuracy of 49.7 % instead of 49.6 %) than the ones reported by ten Brinke et al. (2014), who used 66 participants. However, the differences are only marginal and do not influence the interpretation of the results.

Obtaining probability distributions

To obtain probability distributions of the data, data binning was used on the reaction times. With a bin size of 1 ms, the number of trials was counted, for which a specific reaction time was observed and this number was divided by the total number of trials to obtain the probability $P(RT=x)$ of observing a trial with a reaction time of x . For plotting the probability distribution, the probability distribution was smoothed by a 30-point moving average. By performing this procedure separately on data from one condition (e.g. $C=1$) and the other condition ($C=0$), conditional probability distributions $P(RT=x|C=1)$ of observing a reaction time of x under the condition $C=1$ can be obtained.

¹ The publicly available dataset used in this paper was obtained from <https://osf.io/38kav/>

To obtain posterior probability distributions, the binning procedure was performed separately for all data, and for data under one condition ($C=1$). By dividing the number of occurrences for a reaction time of x in one condition by the number of occurrences in the complete dataset, the posterior probability $P(C=1|RT=x)$ of having one condition when observing a reaction time of x can be calculated. For bins that do not contain any values, a posterior probability of 50 % is assumed. To obtain a more robust estimation of the posterior probability distribution, the resulting distribution was smoothed by a 200-point moving average.

To test for a significant difference between two conditional probability distributions or if the posterior probability is significantly different than chance, a permutation test was performed. The vector containing information about the condition was randomly permuted 2000 times, the probability distributions were calculated and based on those distributions a 99.5 % confidence interval was obtained. If the difference between the two conditional distributions exceeds this interval, there is a significant difference in the probability distributions with $p < .005$. If the posterior probability exceeds that interval, the posterior probability is significantly different than chance ($p < .005$).

Results

Both analyses of (ten Brinke et al., 2014) and (Franz and von Luxburg, 2015) only investigated a congruency effect. But besides the congruency (C), there are three other variables that are of interest: the veracity (V), the participant's response (R) and the word category (W).

Instead of looking at the mean and standard deviation, it is more informative to look at how the data is distributed and for which reaction times there is an actual difference in the data distribution between two conditions. Looking at the probability distribution regarding the variable veracity (see Figure 2.B left) shows that there are significant differences ($p < .005$) in the distribution if the reaction time is between 200 ms and 500 ms, and if the reaction time is between 550 ms and 800 ms. Despite the significant differences, the probability distributions for the two conditions are very similar. To better visualize the differences in the data distribution, one can look at the posterior probability distribution (see Figure 2.B right). The posterior probability shows that if a participant reacts within 200 ms to 500 ms, it is more likely that they watched a video of a lying suspect. When the participants reacted within 550 ms to 800 ms, the probability was higher that the suspect told the truth. Similar information can be inferred from the posterior probability distributions regarding word category and the participant's response. The posterior probability distribution for congruency does not exceed the 99.5 % confidence interval.

Figure 2. Analysis of the reaction time (RT) using probability distribution (left) and posterior probability distribution (right) regarding different variables. **A:** word category (W), **B:** veracity (V), **C:** response (R), **D:** congruency (C). Significant differences ($p < .005$) are marked with gray background.

A non-linear probabilistic approach for classification

Based on the posterior probability distribution, a probabilistic classifier can be trained to predict the individual variables based on the reaction time.

Methods

To ensure that the classifier is not biased towards the majority class (the condition with more data points), data from the majority class is randomly removed to balance the data and ensure that the probability of one class in the training data is 50 %. Based on the balanced training data, the posterior probability distribution $P(V=I|RT=x)$ is obtained by the method described in the previous section. Again, the use of smoothing the distribution by a moving average should be highlighted, because it acts as a form of regularization that improves the generalization performance of the classifier and avoids overfitting. If $P(V=I|RT=x) > 0.5$, the trial is classified as $V=1$, or $V=0$ if $P(V=I|RT=x) < 0.5$. For a value of 0.5 it is randomly chosen if the outcome is 1 or 0 to avoid any class bias.

A participant-specific classifier can be obtained by only using data from one participant.

When information of multiple trials is integrated, the mean over the posterior probabilities of each trial is calculated and tested to be > 0.5 . When a participant-specific classifier using multiple trials was tested in this paper, data from a maximum of 16 trials was used. For some participants, there are conditions for which only 32 trials are available. If all 32 trials would have been used for testing, there would be no trials left for training the classifier, which is why the maximum number of trials to integrate information was set to 16.

For the evaluation of the classifier, it is important to separate training and testing data, which is why a leave-one-out cross-validation was performed. The classifier was trained on $n-1$ trials and the remaining trial was used for testing. This process was repeated, so that each trial was used for testing once. The resulting accuracy was calculated as the percentage of correctly

predicted trials reported. If a classifier integrated information over multiple trials, all those trials were excluded from the training set, so that $n-16$ trials were used for training and the classifier was tested on those 16 trials.

When testing a classifier regarding its accuracy, it should be avoided to define an arbitrary threshold above which the classifier is accurate and below which it is not. Instead it should be tested if the classifier is significantly more accurate than chance level, which can be tested by modeling classification accuracy as a binomial cumulative distribution (Combrisson and Jerbi, 2015) and checking how likely it is to observe a certain accuracy for a given sample size.

Results

The posterior probability distribution can further be used to train a non-linear probabilistic classifier. Based on training data we can model $P(W=1|RT=x)$, the posterior probability that $W=1$ if a reaction time (RT) of x has been observed. This allows us, by observing the reaction time to a single stimulus, to predict if it is more likely that the word belonged to the truth-category ($W=1$) or to the lie-category ($W=0$). Applying this to the data from ten Brinke et al. shows that, based on the reaction time to a single stimulus, the word category can be predicted from the reaction time with an accuracy of 53.2 %, which is significantly above chance level ($p < 10^{-9}$).

As there are significant differences in the posterior probability distributions for the veracity (V) of the suspect and the participant's direct response (R) regarding the veracity, we can also use those variables for prediction. By training a probabilistic classifier based on $P(V=1|RT=x)$, the veracity can be predicted with an accuracy of 52.0 %, which is significantly above chance level ($p < 10^{-9}$). The accuracy of the classifier can be increased by modeling $P(V=1|RT=x)$ independently for each participant to achieve an average accuracy of 53.0 %. As for each video 32 test trials were performed, information from multiple trials can be integrated to be able to predict the veracity with an accuracy of 55.6 %.

When training a probabilistic classifier based on $P(R=I|RT=x)$, the response (R) regarding the lie detection can be predicted with an accuracy of 50.7 % which is significantly above chance level ($p<.05$). By using a participant-specific classifier and integrating information over multiple trials, the classification accuracy can be improved to predict how the participant will respond with an accuracy of 54.1 %.

When training a classifier on $P(C=I|RT=x)$ the congruency of the prime can not be predicted as the accuracy of 49.7 % is not significantly different from chance level ($p>.05$).

A summary of the classification accuracies is also shown in table 1.

Table 1. Accuracy (Acc.) for predicting different variables using a non-linear probabilistic classifier. * indicates an accuracy significantly different from chance level with $p<.05$, ** with $p<10^{-3}$ and *** with $p<10^{-9}$. Results that are not significant ($p>.05$) are marked with †. Cohen's d is calculated using the standard deviation across participants.

	word category (W)		Response (R)		Veracity (V)		Congruency (C)	
	Acc.	d.	Acc.	d.	Acc.	d.	Acc.	d.
population-based (1 trial)	53.2 %	*** 0.318	50.7 %	* 0.101	52.0 %	*** 0.142	49.7 %	† -0.037
participant-specific (1 trial)	53.3 %	*** 0.613	51.8 %	** 0.359	53.0 %	*** 0.583	50.3 %	† 0.067
participant-specific (8 trials)			54.1 %	** 0.390	55.3 %	** 0.426		
participant-specific (16 trials)			53.4 %	* 0.249	55.6 %	** 0.369		

Probing a cognitive process

If we can train a classifier to predict a certain variable, we can use it to probe the cognitive process behind that variable. To demonstrate the rationale of this approach, we can use the experiment from (ten Brinke et al., 2014) to investigate the cognitive process regarding the word categorization task as an example.

To get familiar with the terminology, we first look at the traditional approach, which infers information about a cognitive process by asking the participant: We have the hypothesis that there exists a cognitive process P_W which is able to decide to which category (W) a word belongs to. To test this hypothesis, we need to probe the cognitive process, which we can do by using the probing method $\theta_R(\cdot)$ that asks the participant and infers information about P_W based on their response $R_W = \theta_R(P_W)$. As the response R_W reflects W with an accuracy of 97.2 % and the accuracy is significantly above chance level we can infer: 1. there exists a cognitive process P_W which can categorize words, and 2. P_W is able to categorize words with an accuracy of at least 97.2 %. This is only a lower bound, because we don't know the origin of the remaining inaccuracy of 2.8 %, which could be caused by P_W being inaccurate, the probing method $\theta_R(\cdot)$ being inaccurate or a combination of both.

Instead of the traditional approach to ask the participant and use $\theta_R(\cdot)$ as probing method, we can also use the probabilistic classifier as probing method $\theta_{RT}(\cdot)$ that uses the reaction time (RT) instead of the response (R). With $\hat{W} = \theta_{RT}(P_W)$ being the prediction of the classifier, $\hat{W}$ reflects W with an accuracy of 53.2 %. As this is significantly better than chance we can infer that P_W exists and that P_W can discriminate words with an accuracy of at least 53.2 %.

Although using the probabilistic classifier as probing method is more noisy than asking the participant, it has the benefit that we don't need to ask the participant and can probe a cognitive process without the participant being aware of it.

Probing unconscious lie detection

As the stimuli in the word categorization task were preceded by a masked prime that showed the face of a suspect, we can use the reaction time to that task to probe the process of lie detection. As previously shown, a probabilistic classifier can be used to predict V with an accuracy of 55.6 %. This means that there is a cognitive process P_V which is able to discriminate between lying and truth-telling suspects with an accuracy of at least 55.6 % and that this process can be probed without the participant being aware of it.

When looking at the participants' response, the average accuracy of the lie detection was 49.7 %. If the conscious lie detection works with an accuracy around chance level and there exists a cognitive process P_V that when unconsciously probed shows an accuracy of 55.6 % and is thereby significantly better than chance ($p < 10^{-4}$), this provides evidence for an unconscious lie detection process that is more accurate than conscious lie detection.

However, a common criticism against priming research (Levine and Bond, 2014; Newell and Shanks, 2014) is that the direct measure (the response R) can be inaccurate. If the participant watches the video, the results of the conscious lie detection process is stored in memory and associated with the suspect's face thereby affecting the reaction time to primed stimuli. But as the direct measure (R) is probed after the reaction to the primed stimuli, the results of a conscious lie detection process could have been present in memory at the time of the priming, but retrieving that information could have failed when asked later.

Therefore it is necessary to show that the inaccurate conscious lie detection performance is not caused by a flawed direct measure. This can be done by training a probabilistic classifier

to predict the response (R) regarding the lie detection, which shows that the conscious response can be predicted with an accuracy of 54.1 % which is significantly above chance level ($p < 10^{-3}$) and that there exists a cognitive process P_R related to the conscious veracity judgment.

As the classifiers used to predict V and R use the same reaction time, the cognitive processes P_R and P_V are probed at exactly the same moment. Further, there is no significant correlation between the two processes ($r = -0.0134$, $p = .60$). This provides some evidence for the existence of two uncorrelated cognitive processes P_V and P_R , with P_R being related to the conscious response R and P_V being related to the veracity of the suspect, which can be interpreted as P_R reflecting a conscious lie detection process working at chance level, P_V reflecting an unconscious lie detection process working with an accuracy of at least 55.6 % and both processes being independent. However, to strengthen that conclusion, it would be necessary to validate the results by probing both processes also after the direct measure, which is not possible with the present dataset but should be taken into account in future experiments.

Discussion

Recently, there are efforts in psychology to increase focus on prediction rather than on explanation (Yarkoni and Westfall, 2017). In line with these efforts, a probabilistic approach for the analysis of reaction times was presented in this paper and it was shown how the method can be used to probe conscious and unconscious cognitive processes. Although the probing method is very noisy, due to the high variance in reaction time, it allows to predict the result of a (conscious or unconscious) cognitive process on a single-trial level.

Based on publicly available data of the experiment from (ten Brinke et al., 2014), the method was demonstrated to provide evidence for unconscious lie detection being more accurate than

conscious lie detection and for both processes being independent. Further, it could be demonstrated that the approach can be used to predict how a participant will respond to a stimulus in the future without the participant being aware of the probing process.

Evidence for unconscious lie detection

Franz and von Luxburg (2015) stated that the accuracy of unconscious lie detection needs to be assessed and it needs to be shown that the accuracy of unconscious lie detection is higher than the accuracy of conscious lie detection. With the presented approach, a classification accuracy of 55.6 % could be reached, while the conscious lie detection had an accuracy of 49.7 %, which provides evidence for unconscious lie detection being more accurate.

However, the results presented here are different than that of Franz and von Luxburg (2015) who state that classification is around chance level on the dataset of ten Brinke et al. (2014). There are two key differences between the presented study and the one of (Franz and von Luxburg, 2015): 1. this study tried to predict the veracity directly, not the congruency and 2. a non-linear classifier was used instead of a linear one.

While Franz and von Luxburg (2015) tested different classifiers, they only tested linear classifiers that are based on a threshold. They state to have calculated an upper bound that is the maximum accuracy that can be achieved on that dataset, which is true for linear classifiers, but neglects the existence of non-linear classification methods.

Linear classifiers work well if the data for each class is normally-distributed, but might not work on data that does not follow a normal distribution. As was shown in this paper, a non-linear classifier can be used in such cases to classify data independent of the underlying distribution.

However, Franz and von Luxburg (2015) make an important point by stating that it needs to be shown that the unconscious lie detection accuracy is greater than the conscious lie

detection accuracy to claim that unconscious lie detection is more accurate. With the probabilistic approach it was possible to show that the unconscious lie detection accuracy is greater than the conscious lie detection accuracy, which is evidence for accurate unconscious lie detection and thereby supports (ten Brinke et al., 2014) who suggest that unconscious lie detection might be more accurate than conscious lie detection. Also, with an accuracy of 55.6 % the accuracy is greater than the 54 % accuracy that Bond and DePaulo (2006) found in a meta-analysis for conscious lie detection.

Due to the author's lack of background in psychology, the implications of unconscious lie detection are not further discussed here, but the reader is instead referred to (ten Brinke et al., 2014), who discussed this topic in their publication.

Nevertheless, independent research groups should repeat the experiment by ten Brinke et al. (2014) and analyze the data with the here presented methods to see if results can be reproduced and to see if the claim of accurate unconscious lie detection holds. As already mentioned above, the primed reaction times should be collected before and after the conscious response to ensure the reliability of the direct measure.

Looking beyond congruency

As much of priming research is concerned with congruency effects, this variable should also be discussed. When training a classifier on $P(C=I|RT=x)$ the congruency of the prime can be predicted with an accuracy of 49.7 %, which is not significantly different from chance level ($p>.05$). The congruency (C) is defined by the interaction of V and W. Thereby, even if a congruency effect is found in the data, one should look at the individual variables to identify if the congruency caused the effect or if the effect is caused by the individual variables. As demonstrated in this paper, the congruency effect found by ten Brinke et al. (2014) is actually not caused by the congruency, but by the individual variables.

Looking beyond congruency has shown that the reaction time is influenced by three different processes (conscious and unconscious lie detection, and word categorization). With the available dataset, the temporal relationship between those processes can also be investigated. As the conscious and unconscious lie detection happened when the participants watched the video, and the behavior (reaction time) to the primed stimuli is influenced by the result of those two processes, this is a nice example that the outcome of conscious and unconscious processes in the past can influence our behavior without us being aware of that influence. Due to this temporal relationship, we can also use this method to predict how a participant will respond consciously to a stimulus in the future.

Conclusion and further applications

This paper has demonstrated how machine learning methods can be applied to psychological data and foster the understanding of that data. Although the classification accuracies of the presented approach are significantly above chance level, with a maximum accuracy of 55.6 %, it is difficult to justify this approach to be used in some real applications.

Nevertheless, such potential applications should be discussed.

The first obvious application would be to improve human lie detection performance. As probing the unconscious lie detection is more accurate than asking the participant about their veracity judgment, the approach could be used to improve detection performance.

Another aspect of the probing method is the fact that it allows to probe how a stimulus is perceived by a participant and to predict how a participant will respond to a stimulus. If the results presented here can be reproduced with other aspects besides lie detection, e.g. if a stimulus is perceived positively or negatively, the presented method could be used to evaluate the participant's opinion without the participant being aware of this process which could potentially be applied for marketing applications.

The presented probabilistic approach, which is heavily inspired by machine learning, is a new tool for the analysis of reaction times and can be applied to other behavioral data. Due to the relative simplicity of the method, data from previous experiments can be reanalyzed to investigate different cognitive processes regarding their conscious and unconscious properties. Applying this probabilistic approach to data from other experiments hopefully gives more insight into unconscious processes, what they are capable of and how they influence human behavior.

Acknowledgments

This work was supported by the *Deutsche Forschungsgemeinschaft* (DFG SP 1533/2-1). The author wants to thank Leanne ten Brinke and colleagues for their open data policy. Without sharing the data, this work would not have been possible.

References

- Merikle P.M. and Reingold E.M. (1992) Measuring unconscious perceptual processes. Measuring unconscious perceptual processes In R. F. Bornstein & T. S. Pittman (Eds.), *Perception without awareness: Cognitive, clinical, and social perspectives* (pp. 55-80), New York: Guilford Press
- Timmermans B and Cleeremans A. (2015) How can we measure awareness? An overview of current methods. In M. Overgaard (Ed.), *Behavioural Methods in Consciousness Research*. (pp. 21–46), Oxford: Oxford University Press
- Schmidt T. (2007) Measuring unconscious cognition: Beyond the zero-awareness criterion. *Advances in Cognitive Psychology*, 3(1-2), 275.

- ten Brinke L, Stimson D and Carney DR (2014) Some evidence for unconscious lie detection. *Psychological science*, 25(5), 1098–1105.
- Levine TR and Bond CF. (2014) Direct and indirect measures of lie detection tell the same story: A reply to ten brinke, stimson, and carney (2014). *Psychological science* 2014; 25(10), 1960–1961.
- Franz VH and von Luxburg U. (2015) No evidence for unconscious lie detection: A significant difference does not imply accurate classification. *Psychological science*, 26(10), 1646–1648.
- Combrisson E and Jerbi K. (2015) Exceeding chance level by chance: The caveat of theoretical chance levels in brain signal classification and statistical assessment of decoding accuracy. *Journal of neuroscience methods* 2015, 250, 126–136.
- Newell BR and Shanks DR (2014). Unconscious influences on decision making: A critical review. *Behavioral and Brain Sciences*, 37(1), 1–19.
- Bond Jr CF and DePaulo BM. (2006) Accuracy of deception judgments. *Personality and social psychology Review*, 10(3), 214–234.
- Yarkoni, T., & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, 12(6), 1100-1122.