Spotting Faked Identities via Mouse Dynamics Using Complex Questions

Merylin Monaro
University of Padova
35100 Padova, Italy
merylin.monaro@gmail.com

Luciano Gamberini
University of Padova
35100 Padova, Italy
luciano.gamberini@unipd.it

Giuseppe Sartori
University of Padova
35100 Padova, Italy
giuseppe.sartori@unipd.it

The increment of criminals, including terrorists, crossing international borders using faked identities is a crucial issue. This paper validates a computerized technique to spot people who declare false identity information. Forty participants were asked to answer complex questions about their identity, clicking with the mouse on the correct alternative response on the computer screen. Half of the participants answered truthfully, while the others were instructed to lie. As long as the subject responded to questions, mouse dynamics were recorded. Because lying is cognitively demanding, liars had fewer cognitive resources available to analyse complex questions and to compute the response. As result, they showed a bad performance in the task compared with truth-tellers, revealing a greater number of errors, slower reaction times and larger mouse trajectories. Different machine learning classifiers were trained by a cross validation procedure, achieving a classification accuracy up to 90% in detecting liars.

Identity, Deception, Lie, Mouse.

1. INTRODUCTION

In the most recent literature about deception, a certain number of studies have focused on the detection of faked identity (Monaro et al., 2017c; Monaro et al., 2017d; Monaro et al., 2018a; Verschuere & Kleinberg, 2016). This trend follows the growing need to detect people who cross international borders using fraudulent documents (Hickey, 2015). Between them, there is a high number of criminals, including terrorists (2014). They generally enter US or Europe under a false name and despite the attempts by governments to increase security measures (2015), the issue remains open. The security measures adopted by the border patrol include mainly biometric passports and the cross-checks of information in databases (e.g., wanted lists, fingerprints). However, most of the people working for terrorist organizations are unknown. For this reason, researchers are now focusing on finding new techniques to detect faked identities, which do not require any background information about the suspect.

The first step in this direction was recently taken by Monaro et al. (Monaro et al., 2017c; Monaro et al., 2017a; Monaro et al., 2017b) who proposed a technique based on unexpected questions and the recording of mouse dynamics. Compared to the other existing cognitive-based lie detection methods (e.g., the autobiographical Implicit Association Test (aIAT; Agosta & Sartori, 2013) or the Concealed Information Test (CIT; Ben-Shakhar, 2012)) this was the first attempt to create a tool that works without any ground truth. In fact, both CIT and aIAT can identify which one between two alternative identities is true and which is false, but if and only if one of them corresponds to the actual identity of the subject. In other words, the true identity of the subject must be available to the examiner. However, this is a condition far removed from reality.

In a recent study (Monaro et al., 2017c), the authors asked participants to learn a new identity from a faked Italian identity (ID) card. Then, they were instructed to maintain the faked identity for the rest of the experiment. The experimental task consisted in a computerized questionnaire in which the subject had to respond questions about identity. Whereas liars responded according to the faked identity previously learned, another group of truth-tellers participants were asked to respond truthfully, that is according to their actual identity. Each question required a double-choice response (e.g., yes or no) that was given by subject clicking with the mouse on the box containing the right response. The questionnaire included three kinds of questions: control, expected and unexpected questions. Control questions were about the physical characteristics of the subject (e.g., eye colour), so all participants, including liars, had to respond truthfully in order to be credible and avoid to be unmasked. Expected questions concerned the information that was learned by the liars from the faked ID card (e.g.,

name, surname, date of birth). Finally, unexpected questions were about information not explicitly learned by the subjects from the faked ID card, but easily inferred (e.g., age, zodiac). During the response, mouse dynamics were recorded. Results showed that liars' mouse trajectories differed from those of truth-tellers for spatial and temporal features, such as the time to compute the response and the width of the curve traced by the mouse. Moreover, liars made a greater number of errors compared to truth-tellers, especially in unexpected questions. Training machine learning algorithms (e.g., Support Vector Machines, random forest, logistic regression) on mouse dynamics features, the authors obtained very high classification accuracies, detecting correctly over the 92% of liars.

Although the results are astonishing, the high accuracy is largely due to the effect of the unexpected questions. Unexpected questions are already used in the forensic setting during police questioning (Hartwig et al., 2007). The technique consists in asking questions to which the suspect cannot be prepared in advance and, as consequence, a liar will need more time to answer (Warmelink et al., 2013). Underneath the effectiveness of the unexpected questions, there is the cognitive load theory according to which the deception production requires increased cognitive resources compared to truth-telling (Blandón-Gitlin al., 2014). In fact, in responding unexpected questions, the suspect has to inhibit the actual response, to produce a new faked information and to monitor its consistency with other information previously given, or with objective facts (Vrij et al., 2009). The cognitive load increment, due to the computation of such mental processes, usually results in an increase of the response time and in committing a greater number of errors (Lancaster et al., 2013).

Although the faked identity fits well to be unmasked by unexpected questions, the use of these questions suffer from some limits, as they are difficult to apply in all deception detection situations. For example, in crimes that consist on having or not put in place an action (e.g., I dealt / I didn't deal drug in last few months), it is extremely difficult to fabricate unexpected questions. Such type of lies are known as 'lie of omission' and substantially consist in denying an action (Swol & Braun, 2014). Furthermore, in some cases, the crime details are unknown and the investigators have no elements to build unexpected questions.

The aim of this paper is to validate the mouse dynamics as a tool to detect deception using an alternative technique to unexpected questions to induce cognitive load in liars. Particularly, here we propose the use of complex sentences (Monaro et al., 2018b; Monaro et al., 2018c), which exceed the mentioned above limit of the unexpected questions.

To guarantee comparable results to those of Monaro et al. (2017c), we have followed the same experimental procedure focusing on the detection of faked identities and mouse dynamics recording.

1.1 The cognitive theory under complex sentences

In a previous study, Monaro et al. (Monaro et al., 2018b) have named "complex questions" the sentences that contain more than one information in the same phrase. For example, to investigate the identity one could ask a question about the name (e.g., Is Alice your name?) and a question about the place of birth (e.g., Were you born in Montréal?). A complex question encompasses both this information in the same sentence (e.g., Are you Alice born in April?).

Complex questions require greater cognitive resources compared to simple questions because subjects need to analyse each information one by one, labelling it as true or false. In other words, the subject has to monitor the plausibility of more than one information and retain it in working memory (Baddeley et al., 2014) to, finally, decide if the entire sentence is true or false. While truth-tellers can speedily carry out this sequence of mental operations, liars need more time to match the plausibility of each information with the lie they told (Williams et al., 2013). As result, liars have a bad performance, compared with truth-tellers, when they are involved in a decision task, making a greater number of errors and showing slower reaction times (Monaro et al., 2018b).

2. MATERIALS AND METHODS

In the following sections, the characteristics of the sample, the experimental procedure and the task performed by the subjects are described. We also report the dynamic features that were collected during the subjects' motor response.

2.1 Participants

Forty Italian volunteer participants took part in the experiment. They were recruited among the students of Padova University. The sample consisted of 22 female and 18 male, with an average age of 22.7 (SD=2.1) and average education level of 16.9 (SD=1.3). All subjects were Italian mother-tongue and right-handed.

Participants were randomly assigned to two different experimental conditions. Twenty participants were asked to perform the experimental task responding truthfully, while the other 20 were asked to lie about their identity.

Before the experiment, all participants agreed to the informed consent. The experimental procedure was

in accordance with the Declaration of Helsinki and was approved by the ethics committee for psychological research of Padova University Psychology Department.

2.2 Experimental Procedure

The experimental procedure follows that reported in Monaro et al. (2017c). Participants assigned to liars' group were asked to learn a faked identity from a false ID document. The ID card contained a real photo of the subject, aside from the basic faked information about identity (name, surname, date of birth, place of birth, residence address, occupation and marital status). There were no time restrictions to learn the new ID information, as the subjects were invited to take all the time they needed. When participants thought to be ready, they were asked to recall the faked identity twice. Between the first and the second recall, they performed a distracting task (mathematical operations). The examiner verified the correctness of the learned information and rectified any errors. Finally, participants were told that a second examiner was waiting for them in another room, unaware of their real identity. So, they were instructed to present themselves to the new examiner with the faked identity and respond to any questions according to it. On the contrary, participants that were assigned to the truth-teller condition were told to respond truthfully to all questions. Before the experiment, they were asked to provide their identity information compiling an ID document on which their photo was posted. After they performed the mathematical task, truth-tellers were moved to the other room with the second examiner to complete the computerized task.

The experimental task was implemented using MouseTracker software (Freeman & Ambady, 2010) and it was administered using a 13-inch laptop. It consisted of 60 questions in form of sentences to which the subjects had to respond using the mouse. To cause each question appear, participants were instructed to click on the "start" box that was located in the lower-central part of the computer screen. Then, the sentence showed up in the upper-central part of the screen together with two boxes, respectively on right and left, containing the possible responses (yes, no). Figure 1 shows the computer screen as it appeared to participants during the task. After each response, the mouse was automatically relocated on the "start" box.

2.3 Stimuli

Sixty stimuli were randomly presented to participants. Twenty stimuli were simple sentences consisting of only one personal information. Ten of them were true according to the identity declared by the subject, so they required to respond "yes" (e.g. "My name is Mary") and the other 10 were false, so

they required to respond "no" (e.g. "My name is Carol").

Figure 1: An example of the computer screen as appeared to the subjects during the experimental task. Clicking on "start" box, which corresponds to the X0,Y0 coordinate, the question showed up. The box containing the response labels (yes, no) remained fixed on the screen

Other 20 stimuli were complex sentences composed of two or three personal information (e.g. "I am Mary 29 years old, from Venice"). More in detail, 10 sentences contained two information (e.g., I am Mary and I am currently married") and the other 10 contained three information (e.g., I am Mary living in Milan, and I am married"). Complex sentences required a "yes" response (n=10) when all the information that composed the sentence were true (according to the identity declared by the subject), whereas they required a "no" response (n=10) when at least one of the information compounding the sentence was false. In other words, participants were asked to respond "yes" when the entire sentence was true and to respond "no" when in the sentences there was one or more false information. Finally, we introduced 20 control questions about the test situation (e.g. "I am involved in a computer task"). Ten was certainly true (e.g., "I am in front of a computer", in this case, the required response was "yes") and 10 certainly false (e.g., "I am climbing a mountain", in this case, the required response was "no"). Both liars and truth-tellers had to respond truthfully to control sentences.

An example of stimuli is reported in Table1. All the stimuli, including complex sentences, were built according to those reported by Monaro et al. (Monaro et al., 2018b).

2.4 Collected measures

During the response to each question, the *MouseTracker* software recorded the position of the mouse along the time (Freeman & Ambady, 2010). Because the length of the trajectory changes from trial to trial, the software normalizes each trajectory in 101 time frames. In other words, each time frame corresponds to an X,Y coordinate (e.g., X1,Y1 indicates the position of the mouse along x and y-axis in the first time frame).

Table 1: Example of control, simple and complex questions. It should be noticed that the required response (yes or no) is the same for liars and truth-tellers in all stimuli. In fact, subjects were asked to respond congruently with the identity they provided to the second examiner, regardless of whether it was true or false and previously learned. For example, if according to the faked identity the subject lives in Milan, he has to respond "yes" to the sentence "I live in Milan" and "no" to the sentence "I live in Toronto"

Type	Sentence	Response
Control	I am responding with the mouse	YES
Control	I am eating at the restaurant right now	NO
Simple	I live in Milan	YES
Simple	I live in Toronto	NO
Complex	I was born in Venice in 1987 and I live in Milan	YES
Complex	I was born in Venice in 1992 and I live in Toronto	NO

Moreover, the software recorded the following features:

- Errors: the total number of wrong responses given by the subject.
- Initiation time (IT): the time in milliseconds that occurs between the appearance of the question and the first mouse movement by the subject.
- Reaction time (RT): the time in milliseconds that occurs between the first mouse movement and the click in the response box.
- Maximum deviation (MD): the perpendicular distance between the actual trajectory (signed by the subject during the response) and the ideal trajectory.
- Time to maximum deviation (MD-time): the time in millisecond taken by the subject to reach the point of maximum deviation.
- Area under the curve (AUC): the geometrical area between the actual and the ideal trajectory.
- X-flip: the number of changes of direction of the trajectory along the x-axis.
- Y-flip: the number of changes of direction of the trajectory along the y-axis.
- Velocity: the minimum, maximum and average velocity along the x and y-axis during the response.
- Acceleration: the minimum, maximum and average acceleration along the x and y-axis during the response.

Finally, for each feature, we calculated the average value of the 10 stimuli, separately for control yes, control no, simple yes, simple no, complex yes and complex no sentences.

3. ANALYSIS AND RESULTS

First, we computed the averaged trajectories of the two experimental groups, graphically comparing liars with truth-tellers in control, simple and complex questions. A preliminary statistical analysis was run to investigate in which features liars and truth-tellers statistically differed. Then, a feature selection was performed to select variables to be entered in machine learning models. Finally, four machine learning classifiers were trained by a 10-fold cross-validation to distinguish liars from truth-tellers.

3.1 Analysis of Trajectories

The first comparison between liars and truth-tellers' motor response has been made observing their averaged mouse trajectories. Figure 2 shows the average trajectories of liars and truth-tellers, separately for control, simple and complex questions. It can be noticed that the two experimental groups have mostly overlapping trajectories for control and simple questions, whereas they differ in complex questions. In such stimuli, truth-tellers show straight trajectories from the origin to the response box. On the contrary, liars show wider trajectories, characterized by a greater AUC and MD. This visual pattern is in line with those found by Monaro et al. (Monaro et al., 2017c; Monaro et al., 2017a; Monaro et al., 2017b) observing motor trajectories on unexpected questions. Focusing on complex stimuli, we split the trajectories of questions requiring a "yes" response from those requiring a "no" response (see Figure 3). The plot reveals that liars have more erratic trajectories in responding questions requiring a "no" response compared to questions requiring a "yes" response. In the first response stage, they spend more time moving on the y-axis, with a very erratic route. Then, they deviate toward the chosen response box with a wider curve respect to truth-tellers (AUC liars M=0.83, SD=0.71 and truth-tellers M=0.38, SD=0.73; MD liars M=0.38, SD=0.28 and truth-tellers M=0.18, SD=0.29).

Finally, taking into account only complex questions that required a "no" response, we analysed the position of the mouse along the x and y-axis during the time, in search of time-points of the maximum difference between trajectories of truth-tellers and liars. As shown by Figure 4, the two groups had a maximum difference in the first half of the trajectory along the y-axis and in the last part of the trajectory along the x-axis. We identified as representative points of maximum separation the time frames Y10, Y21 and X75, X80.
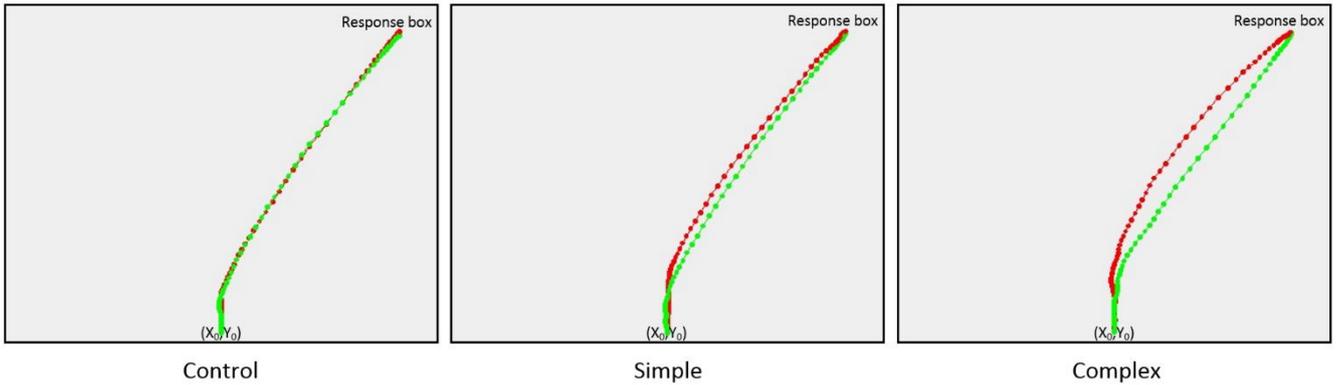
Figure 2: *Comparison of average trajectories of liars (in red) and truth-tellers (in green), respectively for control, simple and complex questions. Questions where subjects made errors were excluded from the plot*
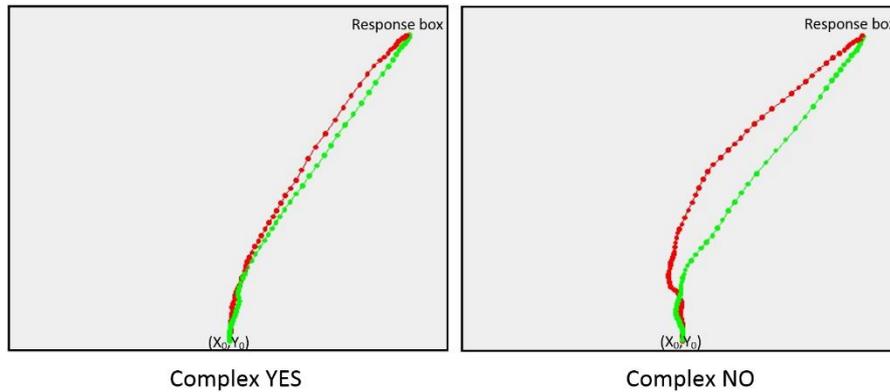

Figure 3: *Comparison of average trajectories of liars (in red) and truth-tellers (in green) for complex questions requiring a "yes" response and a "no" response. Questions where subjects made errors were excluded from the plot*
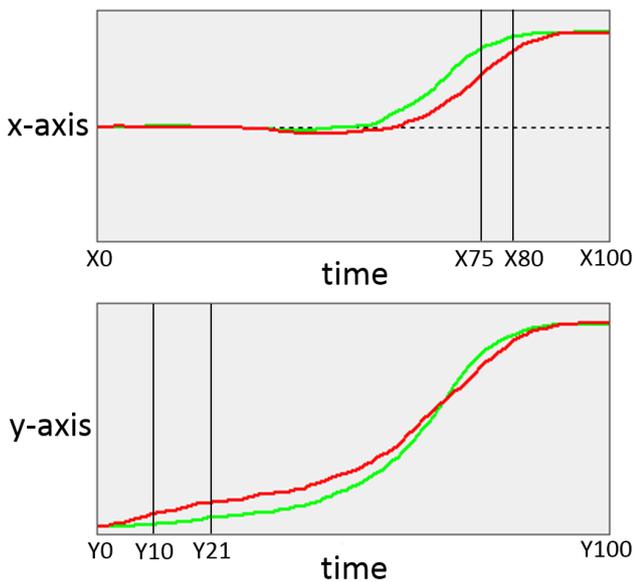

Figure 4: *Position of the mouse along x and y-axis for liars (in red) and truth-tellers (in green) during the time. The graphs refer only to complex questions that required a "no" response. The time frames of maximum separation between liars and truth-tellers are indicated ($X75$, $X80$, $Y10$, $Y21$)*

3.2 Statistics

An independent *t*-test was run for each one of the collected measures, in order to confirm whether the difference between the two experimental groups in complex questions that required a "no" response is statistically significant. A Welch's *t*-test was run using R software 'lsr' package: it adjusts the number of degrees of freedom when the variances are thought not to be equal to each other (Navarro, 2015). To avoid the multiple testing problem the correction of Bonferroni has been applied and the *p*-value has been set to 0.002. Results are reported in Table 2. Cohen's *d* has been calculated for an estimation of the effect-size (Sawilowsky, 2009).

Concerning control, simple questions, and complex questions that require a "yes" response, none of the collected measures have shown statistically significant results, except for the RT in complex "yes" ($p<.002$). For this reason, we entered in the feature selection only the variables related to complex "no" questions.

3.3 Feature Selection

In order to eliminate redundant features and to maximize the models' efficiency, a feature selection was run using WEKA 3.9 (Hall et al., 2009). In more detail, we used a correlation based feature selector (CFS) to sort the features that are more correlated

5

with the class to predict (liar vs truth-teller) and less correlated one to each other (Hall, 1999). Greedy Stepwise has been used as search method. From the list of the 24 total features, the CFS has selected the following: errors (r_{pb} = 0.43), X75 (r_{pb} = 0.69), minimum velocity on x-axis (r_{pb} = 0.35), minimum acceleration on y-axis (r_{pb} = 0.42). The correlation value (r_{pb}) of each feature with the dependent variable (liars vs truth-tellers) is reported.

Table 2: Results of the independent t-test comparing liars and truth-tellers on all the 24 collected variables. Data refer to complex questions that required a "no" response. The table reports t-value, p-value (significance level is set to 0.002 using Bonferroni's correction) and Cohen's d effect-size value. A Cohen's d magnitude of d < .2 indicate a small effect-size, d < .5 medium effect-size, d < .8 large effect-size

Feature	t-test (*t*-value, *p*-value)	Effect-size (Cohen's *d*)
Errors	t(22)=2.98, p=.007	d = 0.94
IT	t(34)=1.13, p=.266	d = 0.36
RT	t(37)=3.46, p<.002	d = 1.09
MD	t(38)=2.19, p=.035	d = 0.69
AUC	t(38)=1.97, p=.056	d = 0.62
MD time	t(32)=2.77, p=.009	d = 0.87
x-flip	t(37)=0.27, p=.784	d = 0.09
y-flip	t(36)=-0.02, p=.981	d = -0.01
X75	t(38)=5.86, p<.002	d = 1.85
X80	t(26)=3.98, p<.002	d = 1.26
Y10	t(21)=1.44, p=.165	d = 0.45
Y21	t(28)=1.76, p=.090	d = 0.55
Vel X min	t(35)=2.37, p=.023	d = -0.75
Vel X max	t(37)=2.37, p=.023	d = 0.75
Vel X mean	t(25)=3.02, p=.006	d = 0.95
Vel Y min	t(38)=1.07, p=.293	d = 0.34
Vel Y max	t(36)=1.92, p=.062	d = 0.61
Vel Y mean	t(36)=1.03, p=.308	d = 0.33
Acc X min	t(38)=2.92, p=.006	d = 0.92
Acc X max	t(38)=3.91, p<.002	d = 1.24
Acc X mean	t(31)=0.34, p=.736	d = 0.11
Acc Y min	t(31)=2.59, p=.015	d = 0.82
Acc Y max	t(28)=2.99, p=.006	d = 0.94
Acc Y mean	t(19)=0.90, p=.376	d = 0.29

3.4 Machine Learning Models

The four features selected above (errors, X75, minimum velocity on x-axis and minimum acceleration on y-axis) were entered as predictors in different machine learning classifiers. To compare our classification accuracies to those obtained by Monaro et al. (Monaro et al., 2017c), we have chosen the same four classifiers: random forest,

logistic, support vector machine (SVM), and logistic model tree (LMT) (Breiman, 2001; le Cessie & van Houwelingen, 1992; Keerthi et al., 2001; Landwehr et al., 2005). More information about the classifiers parameters are reported in supplementary material.

In a similar way to Monaro et al. (2017c), we have run a 10-fold cross-validation on the 40 participants. The classification accuracies that we obtained are the following: random forest = 90%, logistic =77.5%, SVM = 80%, LMT = 90% (see Table 3).

Table 3: Classification accuracies obtained from four different machine learning classifiers (random forest, logistic, SVM, LMT) performing a 10-fold cross-validation. The table reports the classification accuracy which corresponds to the true positive rate (TP Rate), precision, recall and F-measure (F)

Classifier	TP rate	Precision	Recall	F
Random Forest	0.900	0.904	0.900	0.900
Logistic	0.775	0.776	0.775	0.775
SVM	0.800	0.813	0.800	0.798
LMT	0.900	0.904	0.900	0.900

3.4.1. Alternative Models

One of the most discussed issues in lie detection concerns the resistance to countermeasures (Bowman et al., 2014). In fact, if the subject is aware of the indices that are measured by the lie detector, she could apply some strategies to beat it (Peth et al., 2016; Agosta et al., 2010). Detecting deception via mouse dynamics seems to be a promising technique, potentially resistant to countermeasures (Monaro et al., 2017c). Numerous indices are simultaneously recorded, and it is almost impossible for a human being to keep all them under control. Moreover, the broad range of indices allows building alternative classification models. Even if the subject knows in advance which indices will be recorded during the test, she cannot know which of them will be used to predict the outcome.

To argue this point, we have developed alternative machine learning models entering subsets of predictors different from that used above.

A first new subset of predictors has been selected taking out the four features that are more correlated with the dependent variable. These are X75 (r_{pb} = 0.69), X80 (r_{pb} = 0.54), maximum acceleration on x-axis (r_{pb} = 0.53) and RT (r_{pb} = 0.49). The 10-fold cross-validation, using this new four predictors, gives the following results: random forest = 90%, logistic = 82.5%, SVM = 82.5%, LMT = 85%.

A second subset of predictors has been chosen considering the features related to the amplitude of the trajectories: MD (r_{pb} = 0.33), AUC (r_{pb} = 0.30), MD-time (r_{pb} = 0.41). The 10-fold cross-validation

gives the following accuracies: random forest = 72.5%, logistic = 77.5%, SVM = 72.5%, LMT = 80%.

In the third set, we entered only features related to X and Y time frames: X75 (r_{pb} = 0.69), X80 (r_{pb} = 0.54), Y10 (r_{pb} = 0.23), Y21 (r_{pb} = 0.27). Accuracy in 10-fold cross-validation is the following: random forest = 90%, logistic = 82.5%, SVM = 77.5%, LMT = 80%.

4. DISCUSSION

In this paper, we have validated a computerized technique to spot people who declare false identity information via mouse dynamics. Particularly, we have verified whether complex questions are just as effective as unexpected questions in increasing cognitive load with a view to detecting liars when they respond to questions about identity using the mouse.

The experiment conducted by Monaro et al. (Monaro et al., 2017c) has been replicated, but complex questions (instead of unexpected questions) were asked to participants. Results indicated that complex questions are efficient in discriminating liars and truth-tellers, with a slightly lower accuracy comparing to unexpected questions. In fact, using an equal number of participants and the same classification algorithms, we have obtained accuracies ranging from 90% to 77.5% in the 10-fold cross-validation, whereas the accuracies reported in (Monaro et al., 2017c) range from 90% to 95%. Although an accuracy around 90% is not suitable for applications in the field of justice, it may be enough for security applications (e.g., screening for the verification of migrants' identity).

An interesting result concerns the evidence that liars and truth-tellers differ in mouse dynamics parameters only for complex questions that required a "no" response, or rather for complex questions that contain at least one information that is incoherent with the lie they told. In other words, liars need greater cognitive resources to identify one or more discrepancies with the lie they told, whereas they are skilled like truth-tellers in confirming their lie. Probably, it may be because the verification process (Nahari et al., 2014) (the careful monitoring of the congruence between the various information provided during the production of the lie) is cognitively heavier for negative responses. This result is consistent with that one obtained by Monaro et al. (Monaro et al., 2018b), who have used the complex sentences to detect liars for the first time. The authors found that liars had slower RT than truth-tellers in responding to complex questions, especially when the sentence required a "no" response.

In the present study, liars have shown wider mouse trajectories, a greater number of errors and they took more time to compute the response in the complex sentences that required a "no" response. The most predictive variables of deception are the reaction time, the position of the mouse along the x-axis during the late part of the motor response (X75 and X80), and the maximum acceleration on the x-axis. Other features, such as the number of errors, the minimum velocity on the x-axis and the minimum acceleration on y-axis have been also shown to be good predictors to detect liars.

To conclude, the detection of faked information via mouse dynamics is a promising technique, in particular as regards to the resistance to countermeasures. In fact, the high number of indices that are recorded from the mouse movements makes it harder to effectively control all the response parameters. We also demonstrated that different efficient classification models could be built. Anyway, countermeasures are a key point that has to be addressed with further studies, instructing explicitly participants to beat the lie detector. Additional studies are also needed considering a larger sample with a different technological background. In fact, in this study we mostly tested students who are daily immersed in computer and internet use. By contrast, a person who has never or rarely used a computer would have difficulty using the mouse, giving altogether a different trajectory. In particular, a more applied study should be done on people who cross borders.

5. REFERENCES

Agosta, S., Ghirardi, V., Zogmaister, C., Castiello, U., & Sartori, G. (2010). Detecting fakers of the autobiographical IAT. *Applied Cognitive Psychology 25*, 2, 299–306.

Agosta, S. & Sartori, G. (2013). The autobiographical IAT: a review. *Frontiers in psychology 4*, 519.

Baddeley, A., Eysenck, M.W., & Anderson, M.C. (2014). *Memory*. Psychology Press.

Ben-Shakhar, G. (2012). Current Research and Potential Applications of the Concealed Information Test: An Overview. *Frontiers in Psychology 3*, 342.

Blandón-Gitlin, I., Fenn, E., Masip, J., & Yoo, A.H. (2014). Cognitive-load approaches to detect deception: Searching for cognitive mechanisms. *Trends in Cognitive Sciences 18*, 9, 441–444.

Bowman, H., Filetti, M., Alsufyani, A., Janssen, D., & Su, L. (2014). Countering Countermeasures: Detecting Identity Lies by Detecting Conscious Breakthrough. *PLoS ONE 9*, 3, e90595.

Breiman, L. (2001). Random forest. *Machine

Learning 45, 1, 5–32.

le Cessie, S. & van Houwelingen, J.C. (1992). Ridge estimators in logistic regression. *Applied Statistics 41*, 1, 191–201.

Crossing borders: how terrorists use fake passports, visas, and other identity documents. (2014). *Frontline*. http://www.pbs.org/wgbh/pages/frontline/shows/trail/etc/fake.html.

Freeman, J.B. & Ambady, N. (2010). MouseTracker: software for studying real-time mouse-tracking method. *Behavior Research Methods 42*, 1, 226–241.

H.R.158 - Visa Waiver Program improvement and Terrorist Travel Prevention Act of 2015. (2015). *Congress.Gov*. https://www.congress.gov/bill/114th-congress/house-bill/158/text.

Hall, M.A. (1999). Correlation-based Feature Selection for Machine Learning. .

Hall, M.A., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., & Witten, I.H. (2009). The WEKA data mining software: an update. *ACM SIGKDD Explorations Newsletter 11*, 1, 10–18.

Hartwig, M., Granhag, P.A., & Strçmwall, L. (2007). Guilty and innocent suspects' strategies during interrogations. *Psychology, Crime & Law& Law 13*, 213–227.

Hickey, J.G. (2015). Report: 25 percent jump in illegals stopped at border for fake IDs. *Newsmax*. http://www.newsmax.com/Newsfront/illegals-border-fake-id/2015/01/08/id/617256/.

Keerthi, S.S., Shevade, S.K., Bhattacharyya, C., & Murthy, K.R.K. (2001). Improvements to platt's SMO algorithm for SVM classifier design. *Neural Computation 13*, 3, 637–649.

Lancaster, G.L.J., Vrij, A., Hope, L., & Waller, B. (2013). Sorting the Liars from the Truth Tellers: The Benefits of Asking Unanticipated Questions on Lie Detection. *Applied Cognitive Psychology 27*, 107–114.

Landwehr, N., Hall, M., & Frank, E. (2005). Logistic model trees. *Machine Learning 95*, 1–2, 161–205.

Monaro, M., Fugazza, F.I., Gamberini, L., & Sartori, G. (2017a). How Human-Mouse Interaction can Accurately Detect Faked Responses About Identity. In: L. Gamberini, A. Spagnolli, G. Jacucci, B. Blankertz and J. Freeman, eds., *Symbiotic Interaction. Symbiotic 2016. Lecture Notes in Computer Science, vol 9961*. Springer, Cham, 115–124.

Monaro, M., Galante, C., Spolaor, R., et al. (2018a). Covert lie detection using keyboard dynamics. *Scientific Reports 8, 1976*.

Monaro, M., Gamberini, L., & Sartori, G. (2017b). Identity Verification Using a Kinematic Memory Detection Technique. In: K. Hale and K. Stanney, eds., *Advances in Neuroergonomics and Cognitive Engineering. Advances in Intelligent Systems and Computing, vol 488*. Springer, Cham, 123–132.

Monaro, M., Gamberini, L., & Sartori, G. (2017c). The detection of faked identity using unexpected questions and mouse dynamics. *PLOS ONE 12*, 5, e0177851.

Monaro, M., Gamberini, L., Zecchinato, F., & Sartori, G. (2018b). False Identity Detection Using Complex Sentences. *Frontiers in Psychology 9*, 283.

Monaro, M., Spolaor, R., QianQian, L., Conti, M., Gamberini, L., & Sartori, G. (2017d). Type me the truth!: Detecting deceitful users via keystroke dynamics. *Proceedings of the 12th International Conference on Availability, Reliability and Security, ARES '17*.

Monaro, M., Toncini, A., Ferracuti, S., et al. (2018c). The detection of malingering: a new tool to identify made up depression. *Frontiers in Psychiatry*.

Nahari, G., Vrij, A., & Fisher, R.P. (2014). Exploiting liars' verbal strategies by examining the verifiability of details. *Legal and Criminological Psychology 19*, 2, 227–239.

Navarro, D. (2015). *Learning statistics with R: a tutorial for psychology students and other beginners*. University of Adelaide, Adelaide.

Peth, J., Suchotzki, K., & Matthias, G. (2016). Influence of countermeasures on the validity of the Concealed Information Test. *Psychophysiology 53*, 9, 1429–1440.

Sawilowsky, S. (2009). New effect size rules of thumb. *Journal of Modern Applied Statistical Methods 8*, 2, 467–474.

Swol, L.M. & Braun, M.T. (2014). Communicating deception: differences in language use, justifications, and questions for lies, omissions, and truths. *Group Decision and Negotiation 23*, 6, 1343–1367.

Verschuere, B. & Kleinberg, B. (2016). ID-Check: Online Concealed Information Test Reveals True Identity. *Journal of Forensic Sciences 61 Suppl 1*, S237-40.

Vrij, A., Leal, S., Granhag, P.A., et al. (2009). Outsmarting the liars: The benefit of asking unanticipated questions. *Law and Human*

Behavior 33, 2, 159–166.

Warmelink, L., Vrij, A., Mann, S., Leal, S., & Poletiek, F.H. (2013). The effects of unexpected questions on detecting familiar and unfamiliar lies. *Psychiatry, Psychology and Law 20*, 1, 29–35.

Williams, E.J., Bott, L.A., Patrick, J., & Lewis, M.B. (2013). Telling Lies: The Irrepressible Truth? *PLoS ONE 8*, 4, e60713.