

Detecting Deception in Text Messaging:
The Role of Context, Relationships, the Lie and the Liar

Lindsay Reynolds¹

Jeff Hancock¹

Megan French¹

Madeline E. Smith²

Jeremy Birnholtz²

¹Cornell University

²Northwestern University

Abstract

Text messaging has become wildly popular in the United States, and it has some unique characteristics that may impact interpersonal communication such as deception. Here we report how deception detection operates in text messaging, specifically exploring the roles of context, relationship closeness, message content, and one's liar type on deception detection. We addressed these questions using a novel data collection method which relied on the archives of undergraduate students' sent text messages. We find that people are overall accurate at identifying text messages, primarily because most texts are truthful. For those text messages that are deceptive, our participants performed poorly in identifying them, correctly identifying less than 20% of lies. Beyond this central contribution, this study also demonstrates that the other factors described above are important for assessing text messages, which adds an important dimension to the deception detection research.

Detecting Deception in Text Messaging:

The Role of Context, Relationships, the Lie and the Liar

Is it possible to detect deception in text messaging, perhaps the most quickly adopted communication medium in human history (Smith, 2011)? The extensive research on deception detection, spanning over fifty years, is not encouraging: people tend to detect lies at a rate roughly 54% of the time, when at chance they have 50-50 odds of being correct (Bond & DePaulo, 2006). Given that text messages lack the usual nonverbal cues and are, by design, quite short (160 characters), it seems unlikely that deception detection accuracy in text messaging should be any different than what prior work has found, if not worse.

There are several important factors for text messaging, however, that differentiate it from the typical kind of messages used in standard deception detection paradigms. In the standard paradigm participants are asked to judge a series of stimuli, such as a videotaped recoding of a person denying stealing a wallet (e.g., Burgoon, Blair, Qin, & Nunamaker, 2003), on whether the speaker is being truthful or deceptive. Usually the participant has little to no context other than the stimulus, and must make their decision based only on the cues and information available in the stimulus. Our interest, however, is in text messages that have been exchanged between communication partners who know each other and are part of an ongoing interaction situated within their social life. Recent work suggests that this kind of context can play an important role in deception detection (Levine & McCornack, 1992; Vrij, 1994). For example, in one study the availability of context drastically increased a participant's ability to detect deception (Levine & McCornack, 1992). Can the contextually embedded nature of text messages between communication partners play a similar role in improving deception detection for this novel

medium? If this is the case, we should expect higher deception detection accuracy for text messages that are received as part of an ongoing interaction than for text messages that are from a partner but not part of an ongoing interaction (e.g., a message sent from the partner to another person) Furthermore, even lower rates of deception detection should be expected for text messages produced by a random stranger.

Other factors emerging in deception research that may affect text messaging detection accuracy include the nature of the *relationship to the partner* and the actual content or nature of the lie. For example, there is an ongoing debate in the literature about whether it is more difficult to detect deception from an intimate than from a stranger (e.g., Miller, Mongeau, & Sleight, 1986). There is also recognition that lies can be about very different *content*, such as one's feelings versus one's stated achievements, and that lies can be enacted with quite different *methods*, from an outright lie (e.g., "I got an A on the exam" when in fact the speaker received a lower grade) to much more subtle forms (e.g., "I don't know how I did on the exam" when the speaker knew they received a low grade) (DePaulo et al, 1996; Hancock, 2004). Are some kinds of content more difficult to detect than others, and are subtle lies more difficult to detect than more blatant ones?

Finally, a recent phenomenon highlighted in deception research is the *prolific liar*, which refers to people that lie at substantially higher rates (e.g., two standard deviations) in their everyday communication than most people (Serota, Levine, & Boster, 2010). While these individuals make up only a small fraction of the population, does their proficiency in lying affect their ability to detect the lies of others? Their frequent use of lying may improve their detection accuracy on the one hand, or their enhanced experience might bias their perceptions of how often others lie and interfere with their ability to perceive when others are being truthful.

In the present study we seek to address these questions using a novel method of analyzing real texts exchanged between communication partners. We asked pairs of participants to produce recent messages they had sent one another via text messaging, and to retrospectively identify which of these texts was deceptive. We then asked their partner to identify any messages they suspected to be deceptive. This procedure allowed us to calculate detection accuracy for messages exchanged between the partners as part of their everyday communication (high context messages), which we compared to their ability to detect deception in text messages sent by their partner to other individuals (low context messages) or to text messages sent by random others (no context messages). The procedure also allowed us to examine the role of the partners' relationship on their ability to detect one another's lies, and to examine how the content and type of the lie affect detection accuracy. Lastly, we identified a small subset of participants as prolific liars that told two standard deviations more lies on average than other participants and assessed their deception detection accuracy relative to more typical liars. Given the small number of prolific liars, this last analysis was limited to a descriptive assessment.

Context

Research on deception detection has found that accuracy rates of assessing the truthfulness or deceptiveness of a message averages only slightly above chance (Bond & DePaulo, 2006; Kraut, 1980). Research has looked at accuracy rates under different circumstances, yet even under circumstances that would seem most favorable to those making the assessments, such as when judges receive detection training prior to making their assessments, accuracy rates still do not exceed 75% (Fielder & Walka, 1993).

Miller et. al (1986) and Levine and McCornack (1992) discuss the problem with studying deception detection using impersonal, laboratory studies. Traditionally, deception detection has

been studied by having a person state a number of truths and lies (typically, the truthful and deceptive statements involved either correctly reporting or falsifying the answers to a survey they took) while being video recorded. Participant judges then watch the tape to assess whether the statements are true or false. In these situations, the participant judges had no relationship with the message communicator and no additional information about the context of a message. This means that, essentially, participants were making assessments about random messages without any of the additional information that individuals are afforded when making assessments about messages in natural settings.

One type of information available in natural settings that may affect the accuracy of the assessment of messages is the amount of context available to the message judge. Park, Levine, McCornack, Morrison, and Ferrara (2002) argue that when detecting deception in natural settings, judges rely on external or contextual information, such as the consistency of the message with existing knowledge or other evidence. Furthermore, they argue that natural deception is rarely detected immediately, but rather slowly uncovered over time. It is important to note that their conclusion about the importance of context was based on interviews about how individuals uncovered past deceptions, not on examining how different levels of context impact judges' accuracy on specific messages.

In the present study we extend this line of research by comparing deception detection accuracy for text messages exchanged between communication partners, which obviously have high levels of contextual information, with deception detection accuracy for messages in which they have lower levels of context. Specifically, we also examined detection accuracy for text messages sent by one of the partners to another person. Here, the context is low since the participant never received this particular message, but the participant knows the sender of the

message, and the content may pertain to issues that the participant is privy to, such as information common to a social group that includes the participant and the partner. For instance, if Adam and Ben are the pair of participants, there are several ways that Ben may still be able to detect deceptive messages that Adam sent to Chris; Ben and Adam could have talked about the deception previously, Chris may have said something to Ben that he knew to be false, etc. We also compared high-context messages with no-context messages by asking participants to identify deceptive text messages sent by random other people to whom the participants had no relation and no additional information. We expected performance to be lowest in this no-context condition.

Relationship

Another type of information available in natural settings which may affect the accuracy of the assessment of messages is the relationship between the judge and the communicator of a message. There are several reasons why researchers believe that a closer relationship between communicator and judge would impact the judge's ability to make an accurate assessment about the truthfulness of that message. However, there are conflicting views about whether being in a closer relationship with a person increases or decreases detection accuracy.

Some researchers have suggested that being in a closer relationship should help improve detection deception. First, personal and impersonal relationship types have different understandings about the person who communicates a message (Miller et. al 1986). People in impersonal relationships have very little information at their disposal to utilize when making assessments about the truthfulness of a message. Therefore, they often rely on what are often thought of as general cues to deception, such as whether the message communicator maintains eye contact, but these cues have been repeatedly shown to not be particularly helpful in

deception detection (Kraut, 1980, Fielder and Walka, 1993). People in personal relationships, however, do not have to rely only on these generalized and faulty cues to deception, as they have much more information to draw on when making assessments about the truthfulness of messages. For instance, through their interactions over time, they have expectations about the behavior and idiosyncratic tendencies of the communicator of a message, which allows them to compare suspicious behavior with that person's normal tendencies (Miller & Steinberg, 1975). Furthermore, in long-term relationships (either platonic or romantic), judges likely have historical evidence about how known truthful and deceptive statements are presented to others (Miller et. al 1986).

On the other hand, despite the intuitive belief that relationship closeness improves detection accuracy, some studies have found that being in a more personal relationship decreases judges' ability to detect deception among others. McCornack and Parks (1986) developed a model of deception detection within personal relationships that argued that as relationship closeness increases, people becomes increasingly confident in their ability to detect deception from the message communicator. As a result of this increased confidence, they are more likely to succumb to the "truth bias," and believe that the message communicator will never lie to them, which leads to a decrease in the ability to detect deceptive messages. (Levine and McCornack, 1992)

Additionally, the setting and the type of lies told may impact the conclusions that can be drawn about the effect of relationship closeness on deception detection. Miller et al. (1986) argue that in laboratory studies of deception detection among unrelated individuals, the message communicator has little concern for the implication of their lies, and thus has little incentive to effectively conceal a lie. Additionally, they argue that in this setting, those attempting to detect

message deception also have little incentive to effectively assess those messages, as deception in this situation has very little impact on their life. We believe these arguments can be extended to laboratory studies of people who do have a personal relationship, because both the circumstances surrounding the lie (the message communicator was instructed to deliver deceptive messages) as well as the nature of the lies (false answers to a survey about the message communicator's Machiavellian beliefs) create a situation in which the relational consequences of delivering and detecting deception are artificially low.

Kinds of Lies

Although deception research is dominated by studies focusing on detection, the available findings on the nature of lies reveals that lies can be about a variety of topics, and told in myriad ways. In one of the more comprehensive studies on the content and style of lies in everyday communication, DePaulo et al. (1996) observed that the content of lies could be reliably categorized into roughly five categories: 1) lies about feelings and opinions, which refer to lies about a person's private beliefs, opinions or emotions (e.g., "I told him I missed him when I really don't"); 2) lies about actions, plans and whereabouts, which refer to descriptions of a person's activities, both planned and past (e.g., "Said I had sent the check this morning"); 3) lies about achievements, which often refer to lies that cover up a person's shortcomings or underplay strong performances (e.g., "I told him I had done poorly on my calculus homework when I had aced it"); 4) lies about explanations or reasons for actions, which refer to explanations for a person's usually negative behavior (e.g., "I told her I had to quit because my parents wanted me to"), and 5) lies about facts and possessions, which refer to deceptive but straightforward statements of fact (e.g., "I said that I did not have change for a dollar.") (all examples from DePaulo et al., 1996).

How might the content of messages of text message lies affect how frequently they are detected? One approach is to consider that people often rely on outside information to uncover deception (Park et al., 2002). This process is likely much more possible when the message in question contains content that is verifiable, such as facts or actions, and less possible when the message contains purely subjective content, such as feelings. In this study, we examine how detection accuracy rates vary as a function of text messages' content as identified by DePaulo et al (1996).

Prolific Liars

Lastly, recent evidence suggests that individuals vary substantially in how frequently they lie to others, with a small subset of individuals being prolific liars (Serota et al., 2010.) Prolific liars tell a disproportionately large number of lies compared to others. Past research has found that one-half of all reported lies are told by only 5.3% of American adults (Serota et al., 2010). These studies have been large national surveys that ask participants to report how many lies they tell daily. Due to the difference in design between past prolific liar research and the current study, we defined prolific liars as those whose rate of deception is more than two standard deviations above the mean (Reynolds et al., 2013). These individuals, who produce far more deceptive messages than typical users, may also have a unique approach to their assessment of messages as either truthful or deceptive. A final aim of this study is to determine whether prolific liars judge text messages differently than more typical liars.

It is unclear whether prolific liars would be better or worse than typical liars at accurately identifying deceptive and truthful text messages, and the nascent research to date has not examined this question. On one hand, they could be adept at deciphering whether or not people are being honest. As prolific liars, they are aware of the circumstances and strategies they use to

deceive others. Their personal knowledge about deception could lead them to more easily identify when someone else is employing the same tactics. A heightened skill at evaluating veracity would lead prolific liars to be more accurate overall, including being more accurate at detecting both truthful and deceptive messages.

On the other hand, prolific liars could be prejudiced by the false-consensus effect. The false-consensus effect states that people tend to overestimate the extent to which their behavior, attitude, and beliefs are shared by the majority (Ross, Green, & House, 1997). Prolific liars' higher rate of deception may lead them to believe that other people are just as deceptive, resulting in increased suspicion towards' the truthfulness of others' text messages. If this is the case, prolific liars will have a reduced truth bias and identify more text messages as deceptive compared to typical liars. Given that most messages are truthful; prolific liars would be less accurate overall, but especially less accurate at identifying truthful text messages. Identifying more messages as deceptive should improve accuracy for lies given that it increases the likelihood that deceptive text messages are labeled as deceptive. However, prolific liars could easily mistakenly identify both truthful messages as deceptive and deceptive text messages as truthful.

The Current Study

As noted above, we used a novel design in the present study, in which we asked pairs of participants who regularly text-messaged each other to share the text messages that they exchanged with each other, as well as the text messages that each of them shared with one other individual. Participants indicated which of the text messages that they sent were deceptive, and subsequently tried to detect which of the text messages that their partner had sent was deceptive. They also assessed whether text messages sent by unknown individuals, drawn from a different

study, were truthful or deceptive. This procedure allowed us to create three context conditions: *high context*, where the participant knew both the sender of the text message and had received the text message, *low context*, where the participant knew the sender of the text message (their partner) but had not previously received the text message, and *no context*, where the participant did not know the sender nor had previously received the text message (i.e., had no additional information beyond the actual text of the text message).

With this approach we focus on four factors that may impact the accuracy of assessing whether text messages are truthful or deceptive in a natural setting. We were first interested in the effect of the degree of context a message judge has about the message communicator and the circumstances under which a text message was sent on judges' ability to accurately detect deception. Given that external and contextual information are important cues that should aid in the detection of deception (Park et al, 2002), we hypothesized that:

H1: Detection accuracy will be higher when the message judge has more context about a text message than low or no context.

Next, we wanted to explore more deeply into the conflicting arguments regarding the effect of relationship closeness on detection accuracy. Miller et al. (1986) argue that increased relationship closeness should lead to improved detection accuracy because judges have evidence about a communicator's past truths and lies, but Levine and McCornack (1992) argue that being in a personal relationship actually leads to a decrease in detection accuracy because judges are more likely to succumb to problems associated with truth bias. Therefore, we posited that:

H2A: Detection accuracy will be higher when the message judge has a closer relationship with the message communicator

H2B: Detection accuracy will be lower when the message judge has a closer relationship with the message communicator

Furthermore, we were curious about whether the content of the text messages, using categories described in DePaulo et. al (1996), affected judge's accuracy regarding the deceptiveness of text messages. For example, the content of some text messages may be more easily verifiable than others, leading to differences in lie detection rates.

RQ1: Does the content of text messages affect detection accuracy?

Lastly, we were interested in exploring how the message receiver's liar type impacts how they perceive deceptiveness or truthfulness in others. We were specifically interested in the assessments of prolific liars, the small subset of individuals who lie very frequently to others (Serota et al., 2010), and whether they varied from typical liars. For example, prolific liars may be more adept at successfully detecting deception in others than typical liars, or they may be overly suspicious about deception, lowering their accuracy for truthful text messages. Although these are both reasonable hypotheses, because the number of prolific liars was expected, by definition, to be low, we were not able to conduct parametric statistics to test hypotheses. We therefore state a research question and address it with a descriptive analysis of the data:

RQ2: Do prolific liars assess received text messages differently than typical liars?

Method

Participants

Eighty-two pairs participated in this study, which totaled 164 students who all were students at a large U.S. university. Of the 126 participants who provided demographic data, 91 were female, ranged from 18 to 34 years old, and had used some form of text-based messaging

for an average of 5.86 years. Students were recruited via an on-campus web-based recruitment system, and signed up for the study in pairs. Participants were instructed that pairs should be comprised of two people with a friendly—but specifically non-romantic—relationship, and who exchanged text-based messages on a regular basis. All received either course credit or \$10 for their participation.

Procedure

Once both members of a pair arrived for the study, they were verbally briefed on the procedure of the survey, specifically that their partner would be reading the text messages that they enter into the survey. Partners were then separated and seated at individual computers for the duration of the study. Participants first filled out an online consent form, then read a short tutorial that provided definitions and examples of deceptive content. Next, they completed a short questionnaire about demographics (e.g., age, gender) and their text messaging behavior (e.g., how long they have used text messaging, the people with whom they most often exchange text messages). They then completed a two-phase text messaging survey, as described below.

Phase 1: Participants entered the last 15 text messages that they sent to their survey partner, as well as the last 15 text messages that they sent to an individual of their choice who was not their study partner, for a total of 30 messages (some participants reported fewer than 30 messages, however, if they had sent less than 15 messages to either recipient). They were asked about their relationships with the two people with whom they exchanged messages, both in terms of their categorical relationships with those people (e.g., “family member,” “acquaintance,” “close friend”) as well as the closeness of their relationships (measured using a 5-point scale, anchored by (1) “not close at all” and (5) “very close.”). For each message that participants entered, they were asked whether the message was deceptive (measured using a 6 point scale,

anchored by (0) “not deceptive” and (5) “extremely deceptive.”) If a message was marked as deceptive, participants were also asked to explain why the message was deceptive.

Phase 2: In this phase, the messages that survey partners entered in phase 1 were exchanged, such that each participant viewed the messages that their partner had entered into the survey during phase 1. Participants then read text messages that fell into one of three context conditions: 1) *high context messages*, which were the 15 messages that their survey partner had sent to them previously; 2) *low context messages*, which were the 15 messages that their survey partner had sent to someone else; and 3) *no context messages*, which were the 15 messages sent from an anonymous third participant selected at random from a previous separate study (Reynolds, et. al 2011). In this last condition participants had neither information about the sender nor previously seen the message and had to rely solely on information in the message when making their decision. For each message read, participants were asked whether they believed the message was deceptive (using the same scale noted above), and if they believed a message was deceptive, they were asked to explain why they thought so.

Message Coding

Overall, participants provided 4920 messages in phase 1, and read and rated 7124 messages in phase 2 of the study. Messages were coded for jocularity (e.g., Hancock et. al, 2009), and the categories presented by DePaulo et al (1996). Two independent coders performed the coding task. They rated a sample of messages until they reached 80% agreement, and then they each independently rated the remaining messages. Inter-rater reliability is reported in each section.

Jocularity: Participants reported sending 533 deceptive messages, and received 443 messages that they believed to be deceptive. Messages were coded as jocular if they were clearly

not intended to create a false belief in the recipient (e.g., the message “lol” was not technically true because the participant was not actually laughing out loud, but was not likely intended to mislead the recipient). Inter-rater reliability for jocularity was acceptable (Kappa= .84 for outgoing messages, .79 for incoming messages). There were 65 outgoing messages and 39 incoming messages (messages that participants read and rated that were sent by their partner) that were coded as jocular and not used in subsequent analyses.

Message Content: Messages which were rated by their senders as deceptive were coded for the content category presented in DePaulo et. al (1996), such as whether the lie regarded feelings, achievements, actions, explanations or facts. Messages were coded for content based on the explanations participants provided about why a message was deceptive; messages with no or an insufficient explanation were not coded (Kappa = .65.)

Results

Traditionally, studies in deception detection have evaluated only a judge’s overall accuracy; that is, their overall ability to distinguish accurately both truthful and deceptive messages. However, Levine, Park and McCormack (1999) argue that combining both truthful detection and deception detection leads to an incomplete picture of accuracy. This is because people tend to have a truth bias in interpreting the statements of others, and thus people are often much more accurate at detecting truthful messages than they are at detecting deceptive messages. Therefore, for all results below, we present the overall accuracy rate along with the truthful and deception accuracy rates separately.

To test H1, three linear mixed models were created, one for each dependent variable of overall accuracy, truth accuracy, and deception accuracy. In each of the models, text messages

were nested within participants, participants were a random variable nested within pairs, and pairs were set as a random variable. Context type was included as a fixed variable.

H1 predicted that more context for a text message would improve detection accuracy. In terms of overall accuracy, this hypothesis was supported as judges with increased context had higher overall accuracy, $F[2, 327.4] = 33.03, p < .01$. The descriptive data are presented in Table 1. Post-hoc tests revealed that the 7.2% difference in accuracy between high context or no context was significant ($p < .001$), and the 6.1% difference in accuracy between when judges had low context or no context was also significant ($p < .001$). There was, however, no significant difference when judges had high or low context, suggesting that knowing the sender of a text message provided sufficient context to improve overall deception detection accuracy, and surprisingly that receiving the text message did not improve accuracy over not having previously received it.

In terms of truth accuracy, the hypothesis was also supported, as judges with more context had significantly higher accuracy for truthful text messages, $F[2, 325.1] = 20.17, p < .001$. Post-hoc tests revealed that the 5.7% difference between when judges had high context or no context was significant ($p < .001$), and that the 5.1% difference between when judges had low or no context was also significant ($p < .001$), but that there was no significant difference between when judges had high or low context ($p > .10$). Similar to the results with overall accuracy, these findings indicate that knowing the sender of a message greatly improves one's ability to correctly identify a truthful text message, but that knowing more information about the circumstances under which a text message was sent does not further improve accuracy for these types of messages.

For deception accuracy, the hypothesis was not supported, as there were no significant differences in accuracy for deceptive text messages among the different context types, $F[2, 220.713] = 2.086, p < .10$. One possible interpretation of these findings is that without having any context, participants may have been more suspicious about text messages, and therefore believed that more random text messages were deceptive

Our next area of interest was the effect of relationship on deception detection for text messages. To test H2A and H2B, three linear mixed models were created using overall accuracy, truth accuracy, and deception accuracy as dependent variables separately. In each of the models, text messages were nested within participants, participants were a random variable nested within pairs, and pairs were set as a random variable. Relationship closeness was included as a fixed variable. The random text messages drawn from the separate study were excluded from these models, as the judges had no relationship with the senders of those messages.

H2A and H2B were contrasting hypotheses about the impact of relationship closeness and detection accuracy. For overall accuracy, H2A was partially supported as the data revealed a marginally significant association between relationship closeness and overall accuracy, $F[1, 321.3] = 2.90, p < .10$. For each unit increase in relational closeness there was a 1.06% increase in overall accuracy. In terms of truth accuracy, neither hypothesis was supported, with no significant differences in the accuracy of truthful messages by relationship closeness, $F[1, 316.6] = 1.90, p > .10$. For deception accuracy, neither hypothesis was supported, with no significant differences by relationship closeness, $F[1, 194.0] = .104, p > .10$. Taken together, these results suggest a small overall advantage for detecting text messages from partners with whom we share a closer relationship, and that this advantage is not due to particular improvements in either truth or lie accuracy.

Our first research question concerned how the content of deceptive text messages may impact judges' deception detection accuracy. In order to test this, a generalized estimating equation model was created for the content of a lie described by DePaulo et. al (1996.) By definition only lie accuracy was analyzed as only deceptive text messages could be coded for their deceptive content.

The content of a deceptive text message was a significant predictor of deception detection accuracy, $\chi^2(4) = 15.12, p < .005$. While detection accuracy for lies was low in general, as described above, post hoc t-tests revealed that the accuracy for achievements ($M=6.0\%$, $SE=3.2\%$) was significantly lower than all other content categories (feelings ($M=19.0\%$, $SE=2.7\%$), $p < .005$; actions ($M=21.0\%$, $SE=2.7$), $p < .001$; explanations ($M=29.0\%$, $SE=5.3\%$), $p < .001$; facts ($M=21.0\%$, $SE=5.4\%$), $p < .05$). Furthermore, accuracy for lies about feelings was marginally worse than for explanations, $p < .1$. There were no significant pair-wise differences in accuracy among any of the other content categories.

Taken together these results suggest that it is particularly difficult to know when someone is lying about their achievements, which on its face is a surprising result given that achievements are typically verifiable. A closer examination of the achievement lies revealed, however, that the achievements in this student sample often included performance on school-related assignments (e.g., "I missed the last essay q's." when in truth many more than the last essay questions were missed), which is information that is typically not verifiable by others. Indeed, grades are carefully protected personal information with strict limits on how that information can be shared by professors or institutions. To test the possibility that achievement lies tend to be unverifiable, we coded all of the deceptions that had been considered achievement lies to for whether the receivers of the messages could in fact verify the claim in the message ($Kappa = .81$). 71% of the

achievement lies in this sample were unverifiable, which may explain this otherwise surprising finding.

Prolific Liars

To investigate RQ2, prolific liars were identified as any participant with a rate of lying two standard deviations above the mean ($M = 11\%$, $SD = 11\%$). Out of our sample of 164, there were five prolific liars. Due to the large discrepancy in sample size between prolific and typical liars, only descriptive data is presented (see Table 2). Claims of differences between the prolific and typical liars are therefore observational only and are not supported by inferential statistics.

In terms of overall accuracy, prolific liars were less accurate ($M = 66\%$, $SD = 18\%$) than typical liars ($M = 86\%$, $SD = 11\%$), suggesting that their detection rate overall was lower than typical liars. In addition, prolific liars ($M = 75\%$, $SD = 20\%$) were less accurate at judging truthful text messages than typical liars ($M = 93\%$, $SD = 7\%$). The decreased rate is due to the fact that prolific liars (21.5%) were three times more likely than typical liars (6%) to identify a truthful text message as a deceptive, suggesting that prolific liars are more suspicious than typical liars of others' text messages, resulting in a reduced truth bias. Given that most statements are truthful, a reduced truth bias results in prolific liars to be less accurate overall. In the case of the much less frequent deceptive texts, there was little difference in the rate of detection of deceptive text messages between prolific liars ($M = 14\%$, $SD = 11\%$) and typical liars ($M = 20\%$, $SD = 8\%$), suggesting that their suspicion that others are lying did not facilitate their detection of deceptive text messages

Discussion

This study is the first, to the best of our knowledge, explored deception detection in text messages. We examined the overall rates of detection the impact of contextual information,

relationship closeness, message content, and liar type on a judge's ability to rate text messages as truthful or deceptive. Using a novel method in which pairs of participants who regularly exchanged text messages judged the veracity of text messages their partner sent to themselves, text messages their partner sent to another person, and randomly selected text messages sent by a person unknown to the participant, we established three context conditions that varied in terms of how much additional information the judge would have regarding the circumstances under which a text message was sent.

One key finding from this study is that context affected judges' overall detection accuracy and truthful detection accuracy, but did not impact the accuracy of the detection of deceptive messages. When judges knew the sender of a message, their ability to correctly judge truthful text messages increased, although it did not matter whether these messages were sent to themselves or others, as the additional context of receiving the text message was not enough to increase accuracy. This result is interesting because it suggests several possibilities. First is that because the partners were often friends, their shared background and social milieu may be sufficient context for classifying text messages as truthful or false. That is, even though one does not receive a message, one may be familiar enough with the friend and their other communication partners that they have sufficient contextual information to make veracity judgments.

An alternative possibility is that judges were not attending to the context, but rather to the sender. It may be the case that the results of our context manipulation are driven by an unintended relationship closeness manipulation, since judges were closer to their partners than they were to randomly selected strangers. If so, our findings could be explained by McCornack and Parks' Model of Relational Deception, which predicts that increased relational closeness

leads to an increased truth bias. According to the model, when participants knew the sender, they may have felt more confident in their detection ability, and thus were more likely to succumb to the truth bias (McCornack & Parks, 1986). Given that most of the text messages that were judged were truthful, this would explain both the increased accuracy in detecting truthful text messages and increased accuracy overall. An increased truth bias would also explain the effect not extending to the accuracy of detecting deceptive text messages.

From our direct test of relationship closeness, we found that increased relationship closeness was marginally associated with improved overall accuracy, but that there was no significant association between relationship closeness and accuracy for either truthful or deceptive text messages specifically. One likely reason for these findings is that there was not a great deal of reported relationship variance; most participants rated their partner as fairly close, which limits our ability to explore the effect of relationship closeness on message detection.

To further examine the circumstances surrounding detection rates, we explored how deception accuracy rates were impacted by various properties of lies. Interestingly, we found that deception detection was markedly worse when a lie was about achievements than any other type of content. This was a somewhat surprising finding, as we initially thought that lies which included verifiable information, such as one's achievements, would be easier to detect than lies which contained less verifiable information. Upon subsequent analysis, we discovered that most of the achievement lies in this sample were in fact unverifiable, as they were largely about academic achievements. As a result, the message judges (which were often friends or classmates) would be unlikely to have access to the types of information that could verify these text messages as lies.

Finally, we looked at the difference in detection rates between prolific and typical liars. Although we did not perform any statistical analyses, the difference in the accuracy rates suggests that prolific liars are less accurate overall and less accurate at judging truthful text messages. The difference largely lies in the fact that prolific liars are more likely to mistakenly identify a truthful text message as deceptive. Their reduced truth bias may be a result of the false consensus effect, which would lead them to believe that others lie as much as they do. As a result, they are overly suspicious of others' veracity. It is unclear what impact their suspicion is having on their relationships. Do they have lower levels of trust overall? Are their friends aware of their prolific lying, as well as their perception of them as liars? In addition, does the misconception that their friends lie at a similar rate extend to their beliefs regarding the general population? Does the false-consensus effect ease their guilt regarding their deception? Future research could investigate the nuances of prolific liars' approach to detection and its impacts on their relationships.

Limitations and Future Work

The current study used a convenience sample of undergraduates. It may be the case that undergraduates differ in their patterns of deception using text messaging (and by extension, process of lie detection) from other populations. In addition, few participants reported low levels of relationship closeness with either their partner or their "other" person. Both of these factors limit our ability to generalize our findings to a broad range of interpersonal relationships. Future studies should try to include a more diverse sample with a wider range of ages and relationship types.

Lastly, there were several issues regarding our investigation into the effects of context on deception rates. Our results suggest that there was no difference in accuracy rates between our

high and low context conditions. As described above, this may be due to the reliance on the sender, rather than on the text message, to inform detection judgments. However, it may be because there was too much noise between those conditions to capture the effect. For example, we did not ask participants about whether they could identify who their partner's "other" person was. This may have muddled our low context condition somewhat, as some participants may have known the other person (and thus much more about the situation), while others did not. We thought that not being an active party in the conversation would be enough to create different layers of context, but this may not have been the case. Future research could look more carefully at the degrees of context made available to participants to more fully explore the effects of limited context on message accuracy. This would help establish the influence of context, above and beyond the influence of relationship closeness. Despite this limitation, our data clearly show that context matters; knowing the sender of the message improved detection accuracy significantly over messages for which one knew nothing about the sender.

Conclusion

Text messaging has become an incredibly widely adopted communication medium for Americans (Smith, 2011). Here we report on the first study to examine how deception detection operates in text messaging. We find that people are overall accurate at identifying text messages, primarily because most texts are truthful. For those text messages that are deceptive our participants performed poorly in identifying them, correctly identifying less than 20% of lies. Beyond this central contribution, this study also demonstrates that context, relationships, and the content of the message are important factors in classifying text messages. While this may seem straightforward, research on deception detection has largely ignored the importance of context or content on deception detection, and research such as this adds an important dimension to

deception detection research in general. Lastly, the emergence of work on the prolific liar highlights that not all liars are the same, and the present research provides some preliminary indication that the kind of liar you are affects how good you are at detecting lies.

References

- Bond, C. F., & DePaulo, B. M. (2006). Accuracy of deception judgments. *Personality and Social Psychology Review, 10*(3), 214-234.
- Burgoon, J.K., Blair, J.P., Qin, T., Nunamaker, J. (2003). Detecting Deception through Linguistic Analysis. *Lecture Notes in Computer Science, 2665*, 91-1001.
- DePaulo, B.M., Kashy, D.A., Kirkendol, S.E., Wyer, M.M., & Epstein, J.A., (1996). Lying in everyday life. *Journal of Personality and Social Psychology, 70*, 979-995.
- Fielder, K. & Walka, I. (1993). Training lie detectors to use nonverbal cues instead of global heuristics. *Human Communication Research, 20*, 199-223.
- Hancock, J. Birnholtz, J. Bazarova, N., Guillory, J., Perlin, J., & Amos, B. (2009). Butler lies: Awareness, deception and design. *Proceedings of CHI, 517-526*.
- Hancock, J., Thom-Santelli, J., & Ritchie, T. (2004). Deception and design: The impact of communication technology on lying behavior. *Proceedings of CHI, 129-134*.
- Kraut, R. (1980). Humans as lie detectors. *Journal of Communication, 30*, 209-218.
- Levine, T.R. & McCornack, S.A. (1992). Linking love and lies: A formal test of the McCornack and Parks model of deception detection. *Journal of Social and Personal Relationships, 9*, 143-154.
- McCornack, S. A., & Parks, M. R. (1986). Deception detection and relationship development: The other side of trust. In M. L. McLaughlin (Ed.), *Communication yearbook 9* (pp. 377-389). Beverly Hills, CA: Sage.

- Miller, G.R., Mongeau, P.A., & Sleight, C. (1986). Fudging with friends and lying to lovers: Deceptive communication in personal relationships. *Journal of Social and Personal Relationships*, 3, 495-512.
- Miller, G.R. & Steinberg, M. (1975). *Between People: A New Analysis of Interpersonal Communication*. Science Research Associates: Chicago, IL.
- Park, H. S., Levine, T. R., McCornack, S. A., Morrison, K., & Ferrara, M. (2002). How people really detect lies. *Communication Monographs*, 69, 144-157.
- Ross, L., Green, D., & House, P. (1977). The 'False-Consensus' effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13, 279-301.
- Reynolds, L., Gillette, S., Marder, J., Miles, Z., Vodenski, P., Weintraub, A., Birnholtz, J., & Hancock, J. (2011). Contact stratification and deception: Blackberry messenger versus SMS use among students. *Proceedings of CSCW*, 221-224.
- Reynolds, L., Hancock, J., Birnholtz, J., and Smith, M.E. (2013). Everyday Deception or Prolific Liars? A New Approach to the Prevalence of Deception. Paper presented at ICA '13.
- Serota, K.B., Levine, T.R., & Boster, F.J. (2010). The prevalence of lying in America: Three studies of self-reported lies. *Human Communication Research*, 36, 2-25.
- Smith (2011). Americans and Text Messaging. Pew Internet & American Life Project, Sept 19, 2011, <http://pewinternet.org/Reports/2011/Cell-Phone-Texting-2011.aspx>
- Vrij, A. (1994). The impact of information and setting on detection of deception by police detectives. *Journal of Nonverbal Behavior*, 18, 117-136.

Tables

Table 1.

Message receivers' mean (standard error) overall, truth, and lie accuracy by context type.

	Overall Accuracy		Truth Accuracy		Lie Accuracy	
	M	(SE)	M	(SE)	M	(SE)
High Context	86.6%	(0.9%)	93.0%	(1.0%)	19.0%	(3.0%)
Low Context	85.5%	(0.9%)	92.4%	(1.0%)	13.4%	(3.0%)
No Context	75.1%	(0.9%)	87.2%	(1.0%)	16.2%	(3.0%)

Table 2.

Crosstabulation of Deceptiveness Judgments by Liar Type and Reality

Judgments Made by Receiver	Reality	
	Truth	Lie
Truth		
	Prolific	64%
	Typical	84%
Lie		
	Prolific	12%
	Typical	6%

Note. Table 2 shows the percentage of deceptiveness judgments for each type of liar. Judgment refers to the receiver's perception of the veracity, while reality refers to the actual veracity identified by the sender.