Social Engineering in a Post-Phishing Era: Ambient Tactical Deception Attacks

Filipo Sharevski, Paige Treebridge, Peter Jachim, Audrey Li, Adam Babin, Jessica Westbrook

Divergent Design Lab

DePaul University

Abstract

It is an ordinary day working from home, and you are part of a team that regularly interacts over email. Since this is your main line of communication, the company trained you to spot phishing emails. You've learned to skip over emails that exhibit obvious phishing red flags: suspicious links, attachments, grammar errors, etc. You just received an email from your boss about a major project on which you play a critical role. The email is more demanding than usual, even impolite. Your boss has generally seemed more upset with you lately, so you approach them to express your concerns and clear the air. Your boss is not receptive to your feedback. This causes a rift that impacts your working relationship, compromising the effectiveness and productivity of the entire team. You have been a victim of an *Ambient Tactical Deception (ATD)* attack. We developed and tested a proof-of-concept social engineering attack targeting web-based email users. The attack is executed through a malicious browser extension that acts as a man-in-the-middle and reformats the textual content to alter the emotional tone of the email. The objective of ATD is not stealing credentials or data. ATD seeks to coerce a user to the attacker's desired behavior via the subtle manipulation of trusted interpersonal relationships. This goes beyond simple phishing and this paper reports the findings from study that investigated an ATD attack on the politeness strategy used in work emails.

Keywords: social engineering, Ambient Tactical Deception (ATD), politeness

theory, persuasion, computer-mediated communication

---

1. Introduction

Modern workplaces focus on winning the employees' loyalty by allowing flexible hours and working from home, among other things [1]. Many companies have scattered teams working remotely. The lack of shared office space reduces the face-to-face communication, heightening the use of email or other types of online communication [2]. Reciprocity and other forms of conditional cooperation are still at stake in formal settings, however, now employees have the option to use asynchronous mode (email) in addition to synchronous mode of communication (face-to-race, telephone, or instant messaging). Asynchronous messages, like email, are of particular interest for social engineering because they avoid most of the nonverbal cues that reveal malicious intentions [3]. No wonder that more than 50% of the overall email traffic is spam and phishing [4].

Phishers try to cooperate with their potential victims and gain their compliance to yield their credentials, install malicious software, or send money [5]. People and employees dependent on email communication learned to spot phishing emails. Emails are now routinely checked for sender's email address, a digital signature, grammar errors, suspicious links and attachments, any kind of unwarranted urgency for taking action [6]. What if all of this is intact? Is there still a room for social engineering through emails? We believe there is, if the attacker phishes not for what the email receiver *has* (e.g. credentials, system permissions, or money), but instead what they *perceive, think, or feel*. We call this new type of social engineering Ambient Tactical Deception (ATD).

In the physical world, *tactical deception* refers to the "misrepresentation of the state of the world to another individual and it allows adversaries to exploit conditional cooperation by tactically misrepresenting their intentions" [7]. Attackers can bring the tactical deception in the cyber world as the phishers did through email communication 24 years ago [8]. To remain undetected, attackers have to reside in the computing ambience, e.g. the trusted commu-

nication interfaces for routing information exchange online [9]. This is quite different than traditional phishing - the ATD attacker is not the sender but the man-in-the-middle. The social engineering objective is different too - instead of impersonating the sender, the ATD attacker uses an already established conditional cooperation the sender has with the receiver and silently manipulates the textual content exchanged. The ATD attackers are not phishers *per se*; their intention is not to steal but rather to make people (un)happy with a person, a project, or an event. To the same objectives, although through different means, we have already witnessed adversaries during the 2016 US presidential elections and the Brexit campaign [10], [11].

In the next section, we are introducing the concept and the technical implementation of the ATD attack together with the threat model and the most likely victims. In section 3 we discuss the results from a study where we investigated the plausibility of the ATD attack. Section 4 provides an analytical comparison between the ATD and the conventional social engineering exploits to highlight the transition into a post-phishing era of human manipulation online. In Section 5 we conceptualize a defense-in-depth prototype to counter ATD attacks. This is a promising way to protect against this sort of cyberattack that also has innovative properties beneficial beyond the realm of cybersecurity alone. We conclude the paper in Section 6 discussing the evolution of ATD.

2. Ambient Tactical Deception

2.1. Concept

An attacker can use malicious software to act as a man-in-the-middle in online communication, particularly in exchanging information through a web browser. Instead of merely "listening," to the data that flows between two people, they can induce misperception. We call this new form of exploit Ambient Tactical Deception (ATD). Ambient, because the malware-based extension allows the user to routinely complete their web tasks without considering any changes in the interface. Tactical deception, because the attacker silently alters

3

a "honest" content (e.g. email, social media post, or a website) to misrepresent the state of the world to another individual. ATD, in cybersecurity terms, targets the integrity of the communication content.

An ATD attack works in several steps as shown in Figure 1. In the first step, the attacker employs legitimacy-by-design (legitimate both in visual design and in meeting what the user expects to see from a legitimate application) to persuade a victim to install a benign web browser extension for a standard utility that requires text manipulation permissions from the user (Sticky Notes, for example). The attacker bets on the fact that roughly 17% of users pay attention to permissions during installation and only 3% understand how permissions correspond to security risks [12]. The ATD attacker phishes for victim's system permissions, but not to exfiltrate data, install ransomware, or cause any particular damage to their system. ATD works because developing extensions for browsers like Chrome or Mozilla is free and a benign extension can pass all the security checks before publishing. ATD also exploits the browsers are already trusted applications and most antivirus products give it a free pass.
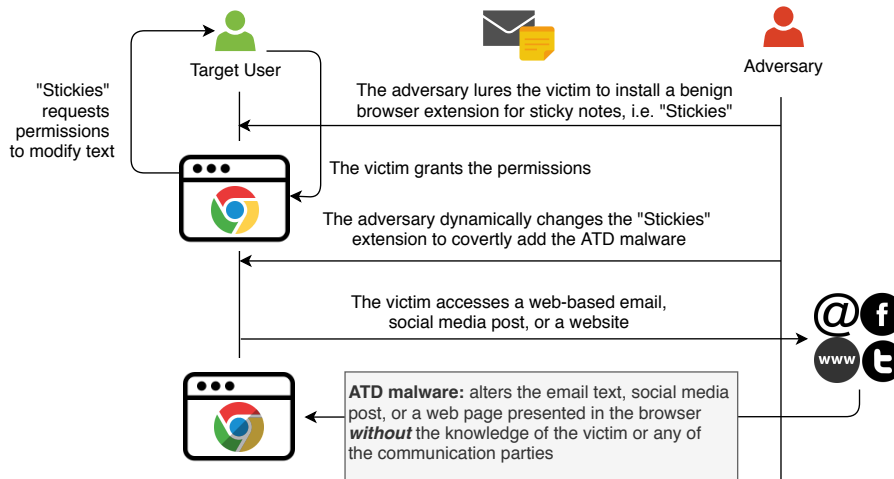

Figure 1: The ATD Attack Flowchart.

In the second step, the attacker changes the behavior of the extension dy-

namically and uses the previously issued permissions to manipulate any text as part of the ATD attack. It is important to note that text is only changed as means to manipulate the tone but not the facts. ATD works insofar as the original text remains clear so as not to raise red flags, and the shift in tone still opens the opportunity for subtle coercion.

2.2. Implementation

The ATD extension, written in JavaScript, changes the tone of textual content by inserting, rearranging and/or swapping words with synonyms detected on a web page. The extension parses out the HMTL for predefined words or word patterns and renders a version of the HTML with the targeted swap or word rearrangement. An example application of ATD is shown in Figure 2a (ATD extension "off") and Figure 2b (ATD extension "on") swapping "disagree" with "strongly oppose him" and "love" with "hold dear" in a public Twitter post [13]. The victim has no reason to question the legitimacy of the tweet because it comes form a trusted source.

(a) ATD extension "off" (b) ATD extension "on"

Figure 2: Dynamic manipulation of text in a social media post with the ATD extension

Borrowing from Orwell's *Politics and the English Language*, the simple idea for the ATD alternation in this example is to make the original message sound less direct and lessen its effect - countering the principle of political writing to "never use a long word or metaphor where a short one will do. [14]" An ATD

attacker interested in meddling with political Twitter messaging can use, for example, the *The New American Lexicon* playbook or jargon characteristic of pecific political party [15]. Altering words does indeed affect an individual's perception of political messaging - a study measuring trends in the partisanship of congressional speech found that it is fairly easy for an observer to infer a congressperson's party from a single utterance [16].

We didn't take this route for our study and instead focused on email requests in formal settings. One of the motivations was the peculiar incident with John Podesta, Hillary Clinton's campaign chef, where a Russian hacking group was able to retrieve a decade of his emails [17]. He received a phishing email claiming that hackers had tried to infiltrate his Gmail account, and sender provided a link to reset his password. Suspicious of potential phishing, he rightfully forwarded the email to the IT staff for further investigation. But then their reply contained a typo: it said the email was "legitimate" (instead of "illegitimate") so Podesta should proceed to change the password (and with that, to reveal his new password to the hacking group). Carrying out an email request in formal settings, like a political campaign headquarters, can have devastating results. We believe that after this incident there is a lot of protection against phishing attacks. But what if an attacker intentionally plants typos not to steal information but to change attitudes and perceptions?

2.3. Attacker's Profile

An ATD attack could be specifically targeted to an individual and most of our examples take that perspective. However, ATD is a variable tactic and would likely be more successful if employed by an individual or group who wished to in some way disrupt another group, a company, or a team of people. Our background research showed that it is much easier to make communication more negative than it is to make communication more positive [citation redacted]. This works to the advantage of an attacker who wants to make people unhappy with a person, a project, or a company. The attacker(s) would want to create discord within a group, and might employ ATD as part of a larger effort. It

could be used to slow down a competitor, poach disgruntled workers, or turn some parts of a group against a specific leader.

2.4. Threat Model and Victims

The ATD attacks originally stem from efforts in which social networks were used to distribute deceptive material through ads and target people of interest. The most infamous example is with Cambridge Analytica, a political data firm, gained access to private information on more than 87 million Facebook users [10]. The firm then offered tools to interested parties that could identify the personalities of voters and influence their political opinion [18]. Another similar case for ad-based ATD attack was when the UK Labour Party campaign chiefs believed that digital ads requested by party leader Jeremy Corbyn were too expensive. Instead, they ran the ads so that only Corbyn and his team would see them, using "individually-targeted, hyper-specific ads through Facebook" [11]. Outside of the political arena, a company called Spinner helps customers manipulate their loved ones with a variety of ads aiming to "boost their intimate life, quit their jobs, or buy their kids a dog" [19].

A malware-based ATD variant is less costly, can target technologies beyond social network platforms, and can have a bigger impact (more granular microtargeting, for example, where the ATD is invoked for a person of interest). We focused on developing a malware-based browser extension because they are low-tech and require minimal investment. In general, the ATD attacks do not need to be approached as a browser extension. A malicious actor may achieve a man-in-the-middle advantage over a victim through a malware for desktop, mobile-email clients, or as an insider threat. For example, the ATD attack can be executed using the LightNeuron malware which allows the attacker to read and modify any email passing through a compromised mail server [20]. Another attack vector for ATD might be a variant of a keylogger/spelling software.

The ATD malware, if deployed successfully, is independent of the email sender, social media account, or webpage source, the protections in the communication, and the specificity of the textual content. The most likely victims of

7

ambient tactical deception are "users that have abandoned traditional interme-
diaries," such as newspapers and other sources that include editorial judgement
of the information provided [21]. Any relationship in which people rely on web
browsers for email correspondence and online communication would be good
ATD targets. It is estimated that one in five Americans communicates excul-
sively online [22].

2.5. System Security Exploit

ATD compromises the system security or on both system and application
level. On a system level ATD compromises the access control - it uses permis-
sions to manipulate text that was granted by the targeted user for a seemingly
bogus application. On an application level, ATD acts as a man-in-the-middle
exploit that targets the integrity of a message (social media post, email, web
page). Based on the on impersonation technique, ATD can be classified as a
spoofing-based man-in-the-middle attack [23]. This is an attack in which the at-
tacker intercepts a legitimate communication between two hosts by the means of
a spoofing attack, and controls transferred data, while hosts are not aware of a
middle man existence. The "middle man" ATD is the malware-based extension:
it intercepts a legitimate HTML document between a legitimate server and a
target user right before it is rendered in the browser window, and dynamically
controls how the target user views the text when looking at the social media
post, web page, or email.

3. Ambient Tactical Deception in Formal Email Requests: A Study

3.1. Politness in Formal Emails

Email communication is asynchronous, limits emotional inference (some-
times requires emojis), and allows senders to plan and revise messages for
grammar, mechanics, clarity, and politeness [24]. This makes email a highly
preferable vector for a malware-based ATD attack. Our initial focus is on po-
liteness in email because it is a critical component of human communication

8

and personal discourse, especially in formal settings [25]. The theory of politeness suggests that people use various politeness strategies to mitigate any face threatening acts when they initiate requests (any acts, including written text, which in some way threaten the 'face' or self-esteem, autonomy, or freedom of another person) [26]. There are four strategies in doing so, ordered from least to most polite (the budget request email examples are added for illustration):

- **Bald-on-record** - a way of speaking that is clear, direct, and concise, e.g. "We need a budget, now!"

- **Positive politeness** - a redress directed to the recipient's desire to be liked, appreciated, approved, e.g. "Jake, we need a budget. Let's finalize it for the proposal today?"

- **Negative politeness** - a redress directed to the recipient's desire to not to be imposed upon, intruded, or otherwise put upon e.g. "Jake, I know you are busy, but would you be willing to meet with me for just an hour? We need a budget for the proposal - the deadline is today."

- **Off-record** - the receiver is given full autonomy to decide how to act upon the request. For example, a sender writing: "Proposals that include budgets are more likely to receive funding" tries to implicitly note to the receiver that they need a budget to submit a complete proposal.

According to the theory of politeness, the requestor or email sender considers three factors when choosing a politeness strategy to craft a request:

- *Degree of imposition* - ranking of impositions by the degree to which they are considered to interfere with one's self-determination or approval

- *Power of the receiver over the requestor* - the degree to which the receiver can impose his/her own plans at the expense of the sender's plans

- *Social distance between the receiver and the requestor* - usually the frequency and type of interaction between them (or how close they are)

We chose to work with email discourse in formal settings for several reasons. A case study of politeness in formal settings found that the aforementioned strategies have been largely employed in the Enron email corpus [27]. Further, a receiver in a formal setting is willing to carry out an email request. This is important to eliminate cases where a receiver discards a request as irrelevant, which is the crucial difference between ATD and traditional phishing emails. An analysis of the email responsiveness of the Enron email corpus suggests that receivers are willing to carry requests in formal settings and with a response generated within a short period of time [28]. Additionally, a receiver in formal settings can easily verify the sender's email address. The social engineering research shows that the sender's email address verification is one of the main cues in deciding the legitimacy of an email [29], [30], [31]. This is important because the emails in the ATD form of social engineering are in fact coming from legitimate senders (colleagues in the workplace), which allows the ATD extension to work on the direct route of persuasion when manipulating the politeness strategy used in a formal email request (see Section 4).

3.2. Study Design

We conducted a preliminary phenomenological study where the malware-based ATD browser extension was used to alter an email request in formal settings. Our objective was to investigate the plausibility of ATD as a new type of social engineering exploit. A convenience sample of 36 participants agreed to participate in the study [32]. The inclusion criteria required participants to be 18 years old or above, to have at least one year of experience working in formal settings and communicating over email, and be a native English speaker (the ATD was developed to alter text written in English language and the theory of politeness pertains to western, english speaking cultures [27]). The study was advertised as a "study in email effectiveness in a workplace" to prevent any influence the full knowledge about the ATD attack might have on the participants' response. The research involved minimal risk, was approved by the IRB, and the participants were debriefed on the overall study immediately after they

provided their response.

Each participant was asked to imagine an email discourse between colleagues in formal settings. Each participant was presented a screen of the Chrome browser in a Windows operating system, and a Web Outlook client already opened in the browser. First, the participant was presented an email request with a bald-on politeness strategy as shown in Figure 3. After reading this email, each participant was asked to verbally answer the following questions:

- What, in your opinion, is the degree of imposition in the email?

- What, in your opinion, is the power distance between the email sender and receiver?

- What, in your opinion, is the social distance between the email sender and receiver?
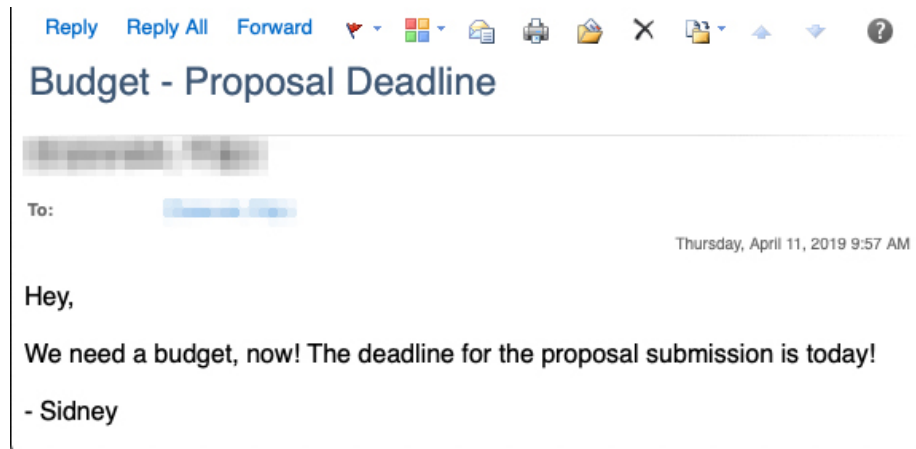

Figure 3: The email with a *bald-on* politeness strategy.

Next, each participant was shown a second email request, which was altered by the ATD browser extension to employ a negative politeness strategy as shown in Figure 4. Finally, the participants were again asked to answer the three questions above relative to the second email request, The verbal answers were recorded, transcribed, and coded for later analysis. We choose to work

with interview-style data collection to enable the participants to elaborate their choice. The participation took less than 30 minutes.
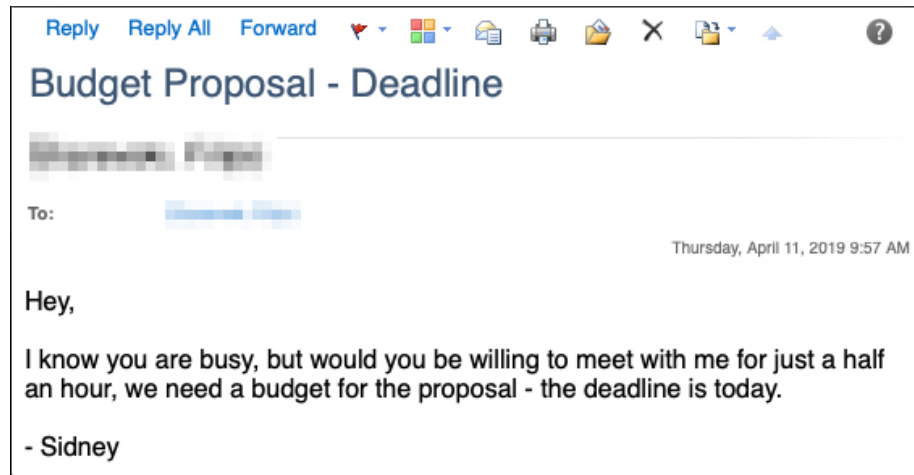

Figure 4: The email with a *negative* politeness strategy.

The email request we used was from a sidneyt@company.com email address, contained only grammatically correct textual content, without links, attachments images, or emojis, so the participants were able to legitimacy of the email. We crafted the email request to contain a neutral phrase, for example "We need a budget". This was important so the participants are able to recognize the literal meaning of the request in the first place [33].

We choose the bold-on-record and the negative politeness strategy because they are the least and the most polite in a direct way (the off-the-record strategy is actually the most polite but it can introduce ambiguity given that it is left to the receiver to interpret the content; we wanted to avoid this). For full realization of the ATD attack, the malicious extension extracted the key request phrases "we need budget", "proposal," "deadline," and "now," from the first email and rearrange the remaining content of the email to read as if it was written with a negative politeness strategy.

3.3. Results

The results are summarized in Table 1. The answers on the first question are coded as "The degree of imposition in the first email is small/large", on the second as "The sender has less/more/equal power than the receiver", and on the third as "The sender and the receiver are very close/close/distant" [33].

Table 1: ATD Study Results

Politeness Factor	Email 1	Email 2	Responses
degree of imposition	small	large	32
	large	large	4
power distance	more	less	22
	more	equal	7
	more	more	4
	less	less	3
social distance	very close	very close	2
	very close	close	16
	very close	distant	6
	close	very close	4
	close	distant	3
	distant	close	4
	distant	distant	1

89% of the participants reported a change in the degree of imposition and social distance between the sender and the receiver between both emails. For the degree of imposition, the overwhelming perception was that the sender in the first email sounded "very direct", while in the second email the sender changed the tone "by mincing words" so "to be more considerate of the effort needed by the receiver." Interestingly, four participants perceived no change in the degree of imposition, reporting that "the receiver has to send a budget by the end of the day in any case, regardless of how the email sounds."

A smaller percentage, 80.55%, notice a change in the power distance - 22 reporting that the sender has more power in the first email while the receiver is the one that has more power in the second email (seven reported that the sender and the receiver are equal in power in the second email). Four participants reported that in both emails, the sender has more power of the receiver, stating that "the email is probably sent by a boss". Three participants reported that the sender has less power in both emails, stating that "the sender needs something from the receiver".

The social distance between the sender and the receiver increased in 69.45% of the responses between the first and the second email (either from very close to close, very close to distant, or close to distant). The social distance between the first and the second email decreased for 22.23% of the responses (close to very close and distant to close). Only 8.32% of the responses reported no change in the social distance between the sender and the receiver in both emails (3 participants). The perception for very close sender and receiver was that they are "casual, probably know each other very well to communicate like that". The perception for close sender and receiver seem that they are "somewhat friendly at work and they at least know each other". The perception for largest social distance was that the sender and the receiver are "distant and more formal."

3.4. Plausibility of Ambient Tactical Deception in Formal Emails

Our main objective in this study was to investigate whether an ATD attack manipulating the politeness strategy used to craft a formal email is a plausible new form of social engineering. Specifically, we wanted to know if the perception on the three factors used to access politeness - the degree of imposition, power distance, and social distance - will change if the ATD extension alters the politeness strategy in a formal email request. Overall, 89% of the participants reporting change in the degree of imposition, 68.75% change in the power distance, and 69.45% change in social distance.

The general impression is that ATD is in fact a plausible new type of social engineering exploit with a huge potential for practical realization. The high

percentage of reported change in perception of the degree of imposition indicates that an ATD attacker can create an implicated context of formal domination over a victim. Not everybody in a formal setting becomes victim to a social engineering attack, however, and this holds true also for the ATD attack - four participants reported no change on how they are implied in the context of the task ("a budget has to be delivered by the end of the day in any case, regardless on how the email sounds").

Similarly, but to a larger account, 14 participants reported no change in the perception of power distance between the sender and the receiver. Seven reported that they are equal, four that the sender has more power, and three that the sender has less power than the receiver between the two emails. One reason might be the choice of the formal setting for the study. Another reason might be the generational differences and the reported years of experience working in formal settings (M = 8.027, SD = 6.669). These remarks hold true in general for the perception on the social distance between the sender and the receiver. In this case, the dominant impression was that the social distance increased (to a variable degree) between the two emails in 25 of the responses, decreased in 8 of the responses, and it remained unchanged in only 3 of the responses.

It is important to keep in mind that these findings may, or may not, illustrate that the the silent alternation of the politeness strategy is the sole factor contributing to the overall change of perception. It way very well be that the ATD acted as a behavioral intervention, or a "nudge", in the cybersecurity behavior [34], not to promote a best behavior practice, but rather to "socially engineer" a behavior to the objective of the ATD attacker. Other factors, including the study design, the content of the email, or the choice of politeness strategies, may have contributed in the what participants reported about the politeness factors in both emails. We take these cautions when analyzing these results and they should be seen as initial support of our idea to investigate the plausibility of ATD attacks rather than an authoritative test.

4. The Post-Phishing Era of Social Engineering

4.1. ATD and Social Engineering Theory

"Social engineering seeks to persuade people or gain a victim's compliance" [5]. The ATD attack works the persuasion through a silent, man-in-the-middle manipulation of text exchanged between two parties. The crucial difference in the ATD attacks is that the attacker is not "phishing" for a one-shot compliance but for a continuous compliance; the gain is not what the victim has, rather, the gain is what the victim perceives. Unlike the traditional phishing, the ATD actually wants the victim to be engaged instead of mindlessly responding to a request. The twist with ATD is that this engagement must not cross the *deception judgment threshold*, otherwise the victim will "temporarily abandon the truth-default approach to seemingly legitimate emails, scrutinize the content, and cognitively retrieve and/or seek evidence to assess honesty-deceit" [11]. The truth-default theory suggest that people presume others to be honest because they either don't think of deception as a possibility during communicating or because there is insufficient evidence lending them unable to prove they are being deceived [11].

According to the Elaboration Likelihood Model (ELM), social engineering utilizes the "peripheral route" of persuasion to successfully engage the victim [35] The ELM distinguishes "central" from "peripheral" routes of persuasion, where a central route encourages an elaborative analysis of a message's content, and a peripheral one is a form of persuasion that does not encourage elaboration (i.e. extensive cognitive analysis) of the message content. Rather, it solicits acceptance of a message based on some adjunct element, such as perceived credibility, likeability, or attractiveness of the message sender, or "a catchy" phrase or slogan. Quite the opposite then the traditional social engineering, the ATD utilizes the central route to encourage the victim to elaborately analyze the message's content. In our preliminary study, the "victim" was encouraged to analyze three politeness factors in the email message: the degree of imposition, the power of the sender over the sender, and the social distance between them.

The social engineering theory resulted from testing six hypotheses to identify the traditional social engineering behavior [36]. The theory posits that people with higher normative commitment, continuance commitment, affective commitment, trust, obedience, and reactance succumb more frequently to social engineering attacks. Normative commitment comes from a reciprocal exchange with a target, where someone will expend effort and perform actions because it is customary or is obligatory. In ATD attacks, the normative commitment doesn't take the form of typical "give-and-take" in phishing email victims because the victims are not obliged for reciprocal exchange. Rather, the objective of ATD is to induce a context of discourse to the objective of the attacker. For example, in our study of altering politeness strategies, the victims have been asked to provide a budget (which was their part of the job) by the end of the day. All subjects stated they will provide the budget if in the position, but when the email was more polite most of the subjects felt they have more power over the requestor (instead of less power, which was reported for an email with a less polite strategy).

ATD attacks are more related to the victims' continuance and affective commitment. With continuance commitment, people become psychologically vested in a decision they have made and maintain consistency in behaviors related to it [35]. So, in case of the workplace ATD tested in our study, victims are willing to comply with email requests coming from verifiable work addresses. The objective of the ATD attack is to engineer the victims to psychologically vest into behavior of the choice of the attacker by manipulating the politeness strategy in an email, without the knowledge of the sender. For example, if a victim feels the sender of an email has a bigger power than them, they might continuously prioritize any email request sent by them. In the same manner, ATD can engineer the victims to psychologically attach to others by manipulating whom they like and identify with. Using our study as an example, the ATD attack was successful in changing the perception on the social distance between the sender and the receiver, which leads to affective commitment [37].

ATD attacks rely heavily on trust. Trust brings cognitive comfort that

17

limits variety of thought and action and attentiveness to detail" [38]. The ATD attack, unlike phishing, provides cognitive comfort to a degree that is necessary to "nudge" the victim to think about the context of a discourse or a request instead of the content. For example, the ATD victims can verify the sender (email address, no attachments, mutual workflow, etc.) but might wonder whether the became distant with the sender if they address them with more polite emails all of a sudden. Of course, victims can be under-trusting and any such a change might trigger the deception judgment threshold, in which case the ATD attacker needs to race to change a strategy of politeness before it is fully detected.

However, some aspects from the obedience to authority theory can counter the potential of under-trusting victims uncovering the ATD attack. Obedience creates actions in deference to those who have perceived coercive power [39]. The ATD itself doesn't create a perception of coercive power, but simply power of compliance. Victims, especially in workplaces, obey commands or requests to simply avoid a negative consequence such as disciplinary action (no one wants to "drop the ball"). As such, the ATD attacker can also play on the reactance of the victim, acting on a scarcity item such as time. In our study the ATD alternated the emails to demand budget from the receiver either by the "end of the day" (more polite) or "now because the deadline the budget is needed is today" (less polite). The less polite variant induced perception of higher power of the sender prompting the subjects to feel that the request "is big" and "the receiver better get that budget done".

4.2. ATD in Social Engineering Taxonomy

The advanced social engineering taxonomy places the attacks in three categories: channel, operator, and type [2]. From the channels, ATD certainly can be conveyed through websites and emails, but also social networks (when accessed through the web). Web extensions targeting particular functional alternation of social networks like Facebook have been developed in the past [40], [41]. Simply by modifying the alternation to focus on the content rather than

18

certain features (number of likes, dates of posts, etc), ATD can be employed though the social network channel.

Operators in the taxonomy can be either humans or software. ATD clearly belongs to the later given it is delivered through a malware-based web browser extension. Software operators in this taxonomy are brought because of their advantage in automating attacks and reaching within a short period of time to a considerably higher number of victims than with purely human attacks. The ATD attacks indeed are automated but not for the purpose of reaching to many victims (thought possible), rather, the ATD attacks are automated to micro-target specific victims and work in the context of personal discourse or tailored web content narrative.

The proposed taxonomy recognizes four types of social engineering attacks: physical, technical, social, and socio-technical. ATD is a new type of a social engineering exploit that is socio-technical in nature, but not as described in the taxonomy. Phishing is identified as one of the most common combination of social and technical approaches, emphasizing spear-phishing campaigns an advanced, more sophisticated version. In spear-phishing, the attacker creates highly targeted messages carried out after initial data-mining about a victim. They point an example where social networking sites were mined on students and then a message was sent that looked like it had been sent by one of their friends. While ATD certainly can use data mining to fine tune the manipulation of politeness strategies, it does so to learn not to make "look-a-like" changes but to make sure the victim doesn't cross the deception judgment threshold. The social component in ATD is the trust and the footprint of a relationship established in a workplace or an assumed credibility of a technology.

5. ATD Detection and Prevention

5.1. Detection

As described in this paper, ATD is highly likely to elude a conventional automatic security detection for several reasons. First, ATD operates on the

HTML textual content after all of the email protections (e.g. digital signatures, end-to-end encryption, or provider-enabled encryption) have been checked and verified on the victim's computer and in their web browser. Second, any operational security monitoring like intrusion or anomaly detection is focused on data exfiltration or patterns of unusual traffic. ATD is not trying to exfiltrate any data but to infiltrate in a subtle way and change the usual data exchanged in a workplace. Third, the ATD attacker can always revert back to the bogus functionality of the browser extension if they suspect that a vulnerability scanning or a forensic investigation is taking place.

5.2. Defense - Conceptual Prototype

This and our previous paper focused on ATD intended to point out what we believe is a possible attack vector against individuals and groups that has not yet emerged in the wild, but that has analogs in other deception and information warfare contexts [42]. We believe that the threat of ambient tactical deception is an inherent risk of computer-mediated communication, particularly as artificial intelligence and machine learning enable software to parse and edit text toward particular emotional tone. We have focused on linguistic politeness as one vector of changing communication without detection. As part of this effort, and the high likelihood that ATD eludes conventional automatic security detection, we considered how such attacks might be defended against. Our prototype for ATD defense is based on a layered protection approach. The prototype is developed in such a manner that each defensive technique interlocks with and supports all the others [43].

5.3. Education and Training

As with any information warfare tactic, awareness of the potential of attack is an advantage to the defender. We assume most people are aware that the words they type move across computer networks, but we question whether much thought is put into how easily those words might be changed along the way. We have pointed out one low cost method of gaining a man-in-the-middle advantage

and changing how an email reader may perceive the linguistic politeness of an email sender. We do not believe the threat of such an attack is imminent, but that those who consider computer and network security should be aware that artificial intelligence and machine learning can be applied to social engineering in ways that would previously have seemed like science fiction.

A practical training session for detecting ATD revolves around crossing the deception judgment threshold and scrutinizing the email communication. The traditional phishing training is focused on quick visual assessments for the most reliable indicators like URLs, grammar, padlocks for https, links, and attachments. As we suggested in the abstract, this is already in place for ATD. The focus for the ATD training is thus on the analysis of the email content in the context of interpersonal communication. One size won't fit all because the degree of imposition, power, and social distance might change over time.

However, the deception judgment can be calibrated based on how polite someone has been in previous emails to set expectations for the ongoing and future email discourse. Compared to the traditional social engineering victims, the ATD ones have the advantage to actually approach (call, text, meet) the email sender and confront about potential change in the tone of the email discourse (or discuss the senders' behaviors with other colleagues). This, in our opinion, is an empowering strategy, and we suggest that any ATD training includes out-of-band email verification. Certainly, this might make the formal interaction cumbersome and redundant, but that is a very small price to quickly cross the deception judgment threshold.

5.4. Linguistic Politeness Check

Since linguistic politeness follows patterns [26], [25], it should be possible to develop software that checks the linguistic politeness of incoming and outgoing email. Whether it is possible to do so effectively is beyond the scope of this paper, but we can imagine a rudimentary prototype that searches for and points out linguistic politeness with a basic scale of the least polite (bold-on-record) to the most polite (off-the-record) strategy. Employed as a browser or

email software extension, it would serve as a data visualization supporting or contradicting a user's perception of the tone of the email. At minimum, this would assist users in triggering the deception judgment. However, there are some obvious problems with implementation as a browser extension, notably a race condition with the ATD extension. If two or more extensions use the textual content the same page, only one wins, and that can result in many cases where the ATD will work but the linguistic politeness check will not (or vice versa).

5.5. Individual Sensitivity Footprint (ISF)

An idiolect is the "totality of the possible utterances of one speaker at one time in using a language to interact with one other speaker" [44]. Forensic authorship attribution, a sub-field of forensic linguistics, "is the process in which linguists set out to identify the author(s) of disputed texts using identifiable features of linguistic style, ranging from word frequencies to preferred syntactic structures." [45]. Researchers have already attempted to identify email authors using idiolect properties of their sentence structures by studying the Enron email corpus [46].

Identifying a change in linguistic politeness requires a measurement of existing linguistic politeness: how polite someone has been in email previously sets expectations for how polite others expect them to be. We have developed a hypothetical measurement and named it an Individual Sensitivity Footprint (ISF). As part of a conceptual prototype we explored how this measurement might be used to screen for an ATD attack. Based on the work in [46] and [45], we have propose a hypothetical measurement and named it an Individual Sensitivity Footprint (ISF). As part of a conceptual prototype we posit that this measurement might be used to screen for an ATD attack, especially in the formal settings where we established the plausibility for such social engineering.

An ISF is a measurement of a focused idiolect, an individual's unique way of communicating as it relates to their linguistic politeness habits. A prototype software could employ corpus approaches to an individual's past email, as in

[46], but focused on linguistic politeness structures. This software, analogous to an anti-virus or anti-malware protection would scan incoming email for linguistic politeness, and then compare the use of each individual's linguistic politeness against their previously established ISF. While a person may vary their linguistic politeness depending on the situation, repeated derivation from an established ISF would indicate a change, and possibly an ATD attack. An ISF-based defense would alert the person who received the email that the sender's ISF profile indicates that the current email is either more or less polite than expected. This software would have advantages beyond ATD defense, and could also be used by an email sender to ensure they are not being rude, or that they do not seem manipulative. In a team situation, an ISF-based system could allow managers to monitor the tone of group communication, and the atmosphere it creates.

6. Conclusion and Discussion

In this paper we presented the findings on a new type of social engineering attack that we baptized as Ambient Tactical Deception or ATD. ATD works though a malicious intermediary - a web browser extension, email exchange malware, or malicious application - to manipulate a textual content exchanged between two parties online. The objective is not to merely listen (and compromise the confidentiality, like in traditional phishing) but to change the tone of the messages, for example emails (compromise integrity). Not stealing credentials by manipulating behavior is an objective attackers so far accomplished by microtargeted ads or internet trolling. In the aftermath of the election meddling and Brexit, we believe that attackers will likely try to use another vector for the same objective. Email is a great candidate because conveys sound and direct communication and as such allows deception that withstands the potential of scrutiny of the content.

We tested and confirmed the plausibility of the ATD attack in formal email communication where the attacker, acting in a man-in-the-middle fashion (not

as the sender) alters the linguistic politeness in a single email to manipulate the perception of the receiver on three factors: the degree of imposition, the power distance, and the social distance between the sender and the receiver. To counter ATD attacks in formal settings, we propose a conceptual defense-in-depth prototype that includes three protection layers. Education and training is a layer already used in protection from traditional phishing. For ATD, we recommend email receivers to scrutinize the email content for sharp changes in the tone of the email and seek out-of-band email verification thorough phone call, face-to-face interaction, or asking others about the sender's recent behavior.

Another layer that is used in traditional phishing is automated detection [47]. The automated detection is specific to the ATD strategy used - for the one we tested it requires linguistic politeness checks of incoming and outgoing email. This might prove ineffective in some cases, so we propose a third layer of defense. Identifying a change in linguistic politeness requires a measurement of existing linguistic politeness, something we call Individual Sensitivity Footprint or ISF. ISF works by identifying unique features of linguistic style for email senders from their past emails. Analogous to an anti-virus or anti-malware protection a software employing ISF checks would scan incoming email for linguistic politeness and then compare the use of each individual's linguistic politeness against their established ISF to detect any change, which might be a result of an active ATD attack. This is a conceptual prototype and we are fully aware that it is far from practical realization, but we want to illuminate any future research that wants to explore ATD countermeasures in addition to new ATD vectors.

A fully functioning ATD attack would require far more finesse to remain in the ambience. Artificial intelligence, particularly as it has been developed to allow "ambience" could be employed to enable the ATD malware to keep track of conversations, allowing for a bidirectional, or even multidirectional ATD attack (cc-ed emails, for example). ATD could function across multiple accounts and devices, limited only by the ability of the adversaries to gain and hold man-in-the-middle positions for each device or communication vector. However, given the processing speed of contemporary computers and improvements in AI that

24

can edit still images, video, and sound, as well as the increase in the amount of time people spend experiencing reality via some form of computer mediation, we believe ATD in future can be employed beyond simple manipulation of text. We can imagine a future in which an ATD attack changed the reality a victim perceived through "smart glasses" or "smart contacts." The ATD concept holds across any computer-mediated reality and the ATD threat grows to the degree people trust what the see, hear, feel, and perceive from any source that has passed through a computer of some sort.

References

References

[1] S. Lewis, Flexible working arrangements: Implementation, outcomes, and management, International Review of Industrial and Organizational Psychology 2003 18 (2003) 1–28. `doi:10.1002/0470013346.ch1`.

[2] K. Krombholz, H. Hobel, M. Huber, E. Weippl, Advanced Social Engineering Attacks, Journal of Information Security and Applications 18 (C) (2015) 113–122. `doi:10.1016/j.jisa.2014.09.005`.

[3] J. T. Hancock, A. Gonzales, Deception in computer-mediated communication, Pragmatics of computer-mediated communication 9 (2013) 363. `doi:10.1080/01638530701739181`.

[4] M. Vergelis, T. Shcherbakova, T. Sidorina, Spam and phishing in 2018 (2019).
URL `https://securelist.com/spam-and-phishing-in-2018/89701/`

[5] M. Workman, Gaining Access with Social Engineering: An Empirical Study of the Threat, Information System Security 16 (6) (2007) 315–331. `doi:10.1080/10658980701788165`.

[6] J. S. Downs, M. B. Holbrook, L. F. Cranor, Decision strategies and susceptibility to phishing, in: Proceedings of the second symposium on usable privacy and security, ACM, 2006, pp. 79–90. `doi:10.1145/1143120.1143131`.

[7] L. McNally, A. L. Jackson, Cooperation creates selection for tactical deception, Proceedings. Biological sciences 280 (1762). `doi:10.1098/rspb.2013.0699`.

[8] Symantec, A Brief History of Phishing: Part I (2007).
URL `https://www.symantec.com/connect/blogs/brief-history-phishing-part-i`

[9] D. J. Cook, J. C. Augusto, V. R. Jakkula, Ambient intelligence: Technologies, applications, and opportunities, Pervasive and Mobile Computing 5 (4) (2009) 277–298. `doi:https://doi.org/10.1016/j.pmcj.2009.04.001`.

[10] H. Kozlowska, The Cambridge Analytica scandal affected nearly 40 million more people than we thought (2018).
URL `https://qz.com/1245049/the-cambridge-analytica-scandal-affected-87-million-people-facebook-says/`

[11] T. Baldwin, Ctrl Alt Delete: How Politics and the Media Crashed our Democracy, Oxford University Press, Oxford, UK, 2018.

[12] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, D. Wagner, Android Permissions: User Attention, Comprehension, and Behavior, in: Proceedings of the Eighth Symposium on Usable Privacy and Security, SOUPS '12, ACM, New York, NY, USA, 2012, pp. 3:1—-3:14. `doi:10.1145/2335356.2335360`.

[13] A. Ocasio-Cortez, Donald trump has decided he does not want to be president of the united states. he does not want to be a president to those who disagree. and he'd rather see most americans leave than handle our nation's enshrined tradition of dissent. but we don't leave the things we love., `https://twitter.com/AOC/status/1151110467147980801` (July 2019).

26

[14] G. Orwell, Politics and the English language, 1946.

[15] M. Abadi, Democrats and republicans speak different languages — and it helps explain why we're so divided, `https://www.businessinsider.com/political-language-rhetoric-framing-messaging-lakoff-luntz-2017-8` (August 2017).

[16] M. Gentzkow, J. M. Shapiro, M. Taddy, Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech, NBER Working Papers 22423, National Bureau of Economic Research, Inc (Jul. 2016).
URL `https://ideas.repec.org/p/nbr/nberwo/22423.html`

[17] E. Lipton, D. Sagner, S. Shane, The perfect weapon: How russian cyberpower invaded the u.s., `https://www.nytimes.com/2016/12/13/us/politics/russia-hack-election-dnc.html?referer=` (December 2016).

[18] K. Granville, Facebook and Cambridge Analytica: What You Need to Know as Fallout Widens (2018).
URL `https://www.nytimes.com/2018/03/19/technology/facebook-cambridge-analytica-explained.html`

[19] S. Chandler, Facebook is helping husbands 'brainwash' their wives with targeted ads, `https://www.dailydot.com/debug/husband-brainwash-wife-spinner-ads-facebook/` (2019).

[20] M. Faou, CTurla LightNeuron: An email too far (2019).
URL `https://www.welivesecurity.com/2019/05/07/turla-lightneuron-email-too-far/`

[21] H. Lin, J. Kerr, On Cyber-Enabled Information/Influence warfare and manipulation, in: Oxford Handbook of Cybersecurity, Oxford University Press, Oxford, UK, 2018.

[22] E. Shearer, Social media outpaces print newspapers in the u.s. as a news source, `https://www.pewresearch.org/fact-tank/2018/12/10/social-`

`media-outpaces-print-newspapers-in-the-u-s-as-a-news-source/`
(December 2018).

[23] M. Conti, N. Dragoni, V. Lesyk, A survey of man in the middle attacks, IEEE Communications Surveys & Tutorials 18 (3) (2016) 2027–2051. `doi: 10.1109/COMST.2016.2548426`.

[24] S. Biesenbach-Lucas, Students writing emails to faculty: An examination of e-politeness among native and non-native speakers of english, Language Learning & Technology 11 (2) (2015) 59–81. `doi:10125/44104`.

[25] J.-R. Park, Linguistic politeness and face-work in computer mediated communication, part 2: An application of the theoretical framework, Journal of the American Society for Information Science and Technology 59 (14) (2008) 2199–2209. `doi:10.1002/asi.20926`.

[26] P. Brown, S. C. Levinson, Politeness: Some universals in language usage, Vol. 4, Cambridge University Press, Cambridge, UK, 1987.

[27] K. Peterson, M. Hohensee, F. Xia, Email formality in the workplace: a case study on the enron corpus, in: Proceedings of the Workshop on Languages in Social Media, Association for Computational Linguistics, 2011, pp. 86–95.

[28] Y. M. Kalman, G. Ravid, D. R. Raban, S. Rafaeli, Pauses and response latencies: A chronemic analysis of asynchronous CMC, Journal of Computer-Mediated Commmunication 12 (1) (2006) 1–23. `doi:10.1111/j.1083-6101.2006.00312.x`.

[29] B. B. Gupta, N. A. G. Arachchilage, K. E. Psannis, Defending against phishing attacks: taxonomy of methods, current issues and future directions, Telecommunication Systems 67 (2) (2017) 247–267. `doi:10.1007/s11235-017-0334-z`.

[30] S. Sheng, M. Holbrook, P. Kumaraguru, L. F. Cranor, J. Downs, Who falls for phish?: a demographic analysis of phishing susceptibility and effective-

ness of interventions, in: Proceedings of the 28th international conference on Human factors in computing systems - CHI '10, ACM Press, New York, New York, USA, 2010, p. 373. `doi:10.1145/1753326.1753383`.

[31] K. W. Duthler, The politeness of requests made via email and voice-mail: Support for the hyperpersonal model, Journal of Computer-Mediated Communication 11 (2) (2006) 500–521. `doi:10.1111/j.1083-6101.2006.00024.x`.

[32] J. W. Creswell, Research Design: Qualitative, Quantitative, and Mixed Methods Approaches, SAGE Publications, Thousand Oaks, California, 2014.

[33] T. Holtgraves, J. N. Yang, Interpersonal underpinnings of request strategies: general principles and differences due to culture and gender, J. Pers. Soc. Psychol. 62 (2) (1992) 246–256. `doi:10.1037/0022-3514.62.2.246`.

[34] L. Coventry, P. Briggs, D. Jeske, A. van Moorsel, SCENE: A Structured Means for Creating and Evaluating Behavioral Nudges in a Cyber Security Environment, in: A. Marcus (Ed.), Design, User Experience, and Usability. Theories, Methods, and Tools for Designing the User Experience, Springer International Publishing, 2014, pp. 229–239. `doi:10.1007/978-3-319-07668-3_23`.

[35] J. T. Cacioppo, R. E. Petty, C. F. Kao, R. Rodriguez, Central and peripheral routes to persuasion: An individual difference perspective. (1986). `doi:10.1037/0022-3514.51.5.1032`.

[36] M. Workman, Wisecrackers: A theory-grounded investigation of phishing and pretext social engineering threats to information security., Journal of the American Society for Information Science & Technology 59 (4) (2008) 662–674. `doi:10.1002/asi.20779`.

[37] N. J. Allen, J. P. Meyer, The measurement and antecedents of affective, continuance and normative commitment to the organization, Jour-

nal of Occupational Psychology 63 (1) (1990) 1–18. `doi:10.1111/j.2044-8325.1990.tb00506.x`.

[38] R. Krishnan, X. Martin, N. G. Noorderhaven, When Does Trust Matter to Alliance Performance?, Academy of Management Journal 49 (5) (2006) 894–917. `doi:10.5465/amj.2006.22798171`.

[39] J. N. Weatherly, K. Miller, T. W. McDonald, Social Influence as Stimulus Control, Behavior and Social Issues 9 (1) (1999) 25–45. `doi:10.5210/bsi.v9i1.135`.

[40] B. Grosser, Facebook Demetricator — benjamin grosser (2018).
URL `https://bengrosser.com/projects/facebook-demetricator/`

[41] D. Sucher, Jailbreak the Patriarchy (2015).
URL `https://github.com/DanielleSucher/Jailbreak-the-Patriarchy`

[42] Joint Chiefs of Staff, Military Deception - Joint Publication 3-13.4, Tech. rep., Joint Chiefs of Staff, Washington DC (2012).

[43] M. R. Stytz, Considering defense in depth for software applications, IEEE Security Privacy 2 (1) (2004) 72–75. `doi:10.1109/MSECP.2004.1264860`.

[44] B. Bloch, A set of postulates for phonemic analysis, Language 24 (1) (1948) 3–46. `doi:10.2307/410284`.

[45] A. Johnson, D. Wright, Identifying idiolect in forensic authorship attribution : an n-gram textbite approach, Language and Law 1 (1) (2014) 37–69.

[46] D. Wright, Using word n-grams to identify authors and idiolects: A corpus approach to a forensic linguistic problem, International Journal of Corpus Linguistics 22 (2) (2017) 212–241. `doi:10.1075/ijcl.22.2.03wri`.

[47] Y. Zhang, J. I. Hong, L. F. Cranor, Cantina: A Content-based Approach to Detecting Phishing Web Sites, in: Proceedings of the 16th International

Conference on World Wide Web, WWW '07, ACM, New York, NY, USA, 2007, pp. 639–648. `doi:10.1145/1242572.1242659`.