



Community Targeted Phishing

A Middle Ground Between Massive and Spear Phishing Through Natural Language Generation

Giaretta, Alberto; Dragoni, Nicola

Published in:

Proceedings of 6th International Conference in Software Engineering for Defence Applications - SEDA 2018

Link to article, DOI:

[10.1007/978-3-030-14687-0_8](https://doi.org/10.1007/978-3-030-14687-0_8)

Publication date:

2020

Document Version

Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):

Giaretta, A., & Dragoni, N. (2020). Community Targeted Phishing: A Middle Ground Between Massive and Spear Phishing Through Natural Language Generation. In P. Ciancarini, M. Mazzara, A. Messina, A. Sillitti, & G. Succi (Eds.), *Proceedings of 6th International Conference in Software Engineering for Defence Applications - SEDA 2018* (pp. 86-93). Springer. https://doi.org/10.1007/978-3-030-14687-0_8

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Community Targeted Phishing

A Middle Ground Between Massive and Spear Phishing through Natural Language Generation

Alberto Giaretta¹ and Nicola Dragoni^{1,2}

¹ Centre for Applied Autonomous Sensor Systems (AASS),
Örebro University, Sweden, alberto.giaretta@oru.se,

² DTU Compute, Technical University of Denmark,
Denmark ndra@dtu.dk

Abstract. Looking at today phishing panorama, we are able to identify two diametrically opposed approaches. On the one hand, massive phishing targets as many people as possible with generic and preformed texts. On the other hand, spear phishing targets high-value victims with hand-crafted emails. While nowadays these two worlds partially intersect, we envision a future where Natural Language Generation (NLG) techniques will enable attackers to target populous communities with machine-tailored emails. In this paper, we introduce what we call Community Targeted Phishing (CTP), alongside with some workflows that exhibit how NLG techniques can craft such emails. Furthermore, we show how Advanced NLG techniques could provide phishers new powerful tools to bring up to the surface new information from complex datasets, and use such information to threaten victims' private data.

Keywords: phishing, NLG, Natural Language Generation

1 Introduction

Even though many years have passed since the Internet has come into our lives, some of its atavistic problems still have to be addressed. Phishing is one of them, and we foresee that its severity could even worsen in the coming years.

On the one hand, we have a massive phishing approach where attackers fill simple templates with some basic information about the target, aiming to create simple and sound emails with the minimum effort possible.

On the other hand, we have *spear phishing* emails where attackers experienced in social-engineering techniques write articulate and specific emails. Even though these emails are carefully forged and hard to detect, the creation process requires such an effort that they are reserved only for high-level targets. Spear phishing techniques are more and more popular into hackers' toolboxes, as an example they are essential to strike *Advanced Persistence Threats (APTs)* [6].

Although they have the same goal, until today these two worlds have been somehow distinct from each other: massive production on one side, craftsmanship on the other one. We envision a future where a middle ground will be prominent,

thanks to sophisticated production processes that go under the name of *Natural Language Generation* techniques, which will enable to target specific categories with machine-tailored messages. We call this new approach to phishing *Community Targeted Phishing (CTP)*, and we show in Figure 1 an example of this evolution, as we foresee it.

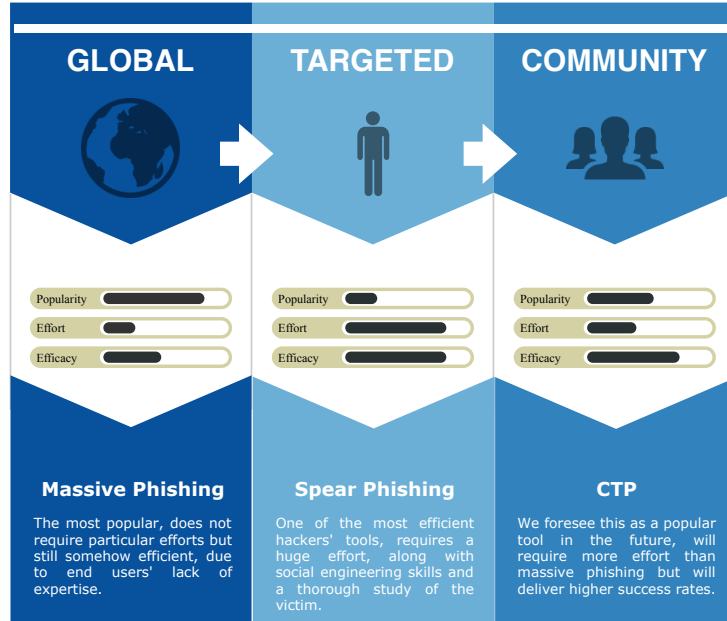


Fig. 1. Phishing naturally evolved from a massive approach to a manually hand-crafted tool, used by hackers to target high value victims. We foresee a midway approach in the future, which combines effectiveness and cheapness. Value bars present an estimate of each characteristic, not actual statistic data.

To the best of our knowledge, at the time of writing only one recent work about NLG applied to emails generation has been published [2]. In that work, authors show that NLG-based emails can be extremely effective to deceive people. Apart from such work, no other studies dig into how NLG fits into the big picture, nor how smart partitioning the targets could dangerously combine with NLG techniques. Strictly related, Natural Language Processing (NLP) goes in the opposite direction of NLG by extracting information from written texts, and it has been proposed as a solution to develop spam filters [8].

This paper is organized as follows: Section 2 gives a basic overview about NLG and its core concepts, useful to read through this paper. Section 3 digs into our predictions introducing our CTP concept with related workflows, in order to

clarify how Template-Filling NLG could enable attackers to target the victims in the future. Then, Section 4 goes even further, making some hypothesis about how Advanced NLG techniques could enable attackers to gain victims' trust and increase the success rate of their phishing attempts. Finally, Section 5 wraps up our conclusive thoughts and lists some directions that we would like to take in the future.

2 Natural Language Generation (NLG)

First of all, what do we talk about when we talk about Natural Language Generation? A simple definition is: "Natural language generation (NLG) systems generate texts in English and other human languages" [5], yet this task is easier said than done. As a matter of fact, NLG systems have to make quite a number of different choices, in order to achieve the most suitable text for the specific purpose. As an example, psycho-linguistic models of language comprehension show that reflective pronouns usually improve a text readability and thus should be normally used, but such pronouns should be usually avoided in safety-critical texts, like operation manuals for nuclear power plants [5].

From a high-level perspective, NLG is traditionally divided in three macro-categories, namely *Canned Text*, *Template-Driven (or Template Filling)*, and *Advanced (or Proper) NLG*. Canned Text is the most basic approach, where a text is pre-generated (usually by a human being) and is used when the right moment comes. Even though different texts can be merged with some *glue texts*, this technique is hardly flexible and unable to adapt even to little domain variations. As previously pointed out [7], "Canned text is the borderline case of a template without gaps".

Talking about Template-Driven, this approach consists into defining some templates that exhibit *gaps*, to be filled with the correct information gathered from an external source, such as a database. The third approach, the Advanced NLG, is sometimes called proper NLG to highlight the fact that a NLG system implies a large variety of choices, and the two aforementioned approaches oversimplify the matter by using preformed texts. In particular, some authors argue that Template-Driven approaches are not proper NLG and tend to dismiss them, while others affirm that Template-Driven systems have been so much refined through the years that the distinction between them and pure NLG is too blurred to hold, nowadays [7].

No matter what kind of approach is chosen, even though NLG systems are far from being widespread, they have been used in many applications. One of the classic examples are weather forecasts, which are automatically generated from meaningful database entries; other fields implemented NLG in their routines, such as soccer, financial, technical, and scientific automatic reports.

3 Aiming at Groups: Community Targeted Phishing (CTP)

What massive phishing lacks is a deeper comprehension about recipients' relationships, characteristic which is totally neglected nowadays. Yet, dividing users in meaningful subsets could help phishers to impersonate some group members, and attack the others, in a much more effective way. This is not done yet, mostly because traditional phishing is still effective, but spam filters are proving themselves more and more reliable [1] [3]. One popular approach to develop such filters, is to apply the Bayesian filtering. In order to discern a legitimate email from a malicious one, this process checks every single word against a table, which contains the most commonly used words in spam emails. Even though advanced adversarial machine learning techniques can be used to poison the whole process [4], this approach is effective.

Such filters require training and leverage the strong points of phishing, which are massiveness and repetitiveness, to get statistically meaningful data. If phishers could smartly divide this huge user blob in smart subsets, they could forge more complex emails and deprive Bayesian spam filters of the aforementioned vital statistical data. We call this approach Community Targeted Phishing (CTP), and in the following subsection we provide some examples on how it could be applied.

3.1 Case Study: Scientific Community

One community which we are familiar with is the scientific one. As a matter of fact, researchers working in the same research field mostly share many interests and know quite a number of colleagues. Furthermore, the email exchange for collaborations, sharing of ideas and call for papers are the routine. On the average, if computer science researchers are enough knowledgeable about security and spam, members of other scientific fields usually exhibit very basic computer skills, which makes them easy targets for malicious attackers.

Scientists already receive a lot of scam emails from shady conferences and journals, which name resembles a legitimate venue and invite the recipients to submit their work. Sometimes people willingly submit to pay-for-publish venues, some other times they are fooled into believing that they are actually submitting to an honest venue. But these phishing emails are really simple in their form, as attackers just fill the first line with something along the line of "Dear Prof. Smith" without any kind of variation in the preformed text and without trying to impersonate an actual researcher. Furthermore, in many cases the attackers even fail to parse correctly the title, as it often happens for PhD Students to be addressed as Dr., Prof., or similarly wrong titles. Succeeding into impersonate a real person through an automated process, on the other hand, is a far more serious (and hard to achieve) threat and this is what we address in our paper.

As a case study, we propose a mechanism that fills two proposed templates by utilizing Template-Driven NLG techniques, and the related workflows that

show how the whole email forging process could take place. As a matter of fact, we believe that even such basic templates could successfully fool a good number of people. The main problem that we see with this approach is that it still requires some human effort, to understand how a community usually interacts and what subsets should be excluded or included into the attack. As an example, in Subsection 3.1 we will see that targeting very close colleagues with that particular template should be avoided.

Template 1 The first proposal takes advantage of the fact that researchers know many names in their field, but the relationships tend to be very weak. Yet, networking is an important part in scientists' lives, therefore a proposal from a respected name in the field could be an easy lure for a young researcher. The template is: "Dear [Colleague Name], I've read your recent work entitled [Manuscript Title] and I found it quite interesting. I've come up with some ideas about that same topic and I would be enthusiastic to work with you on this. At [Fake Google Drive Link] you will find a brief recap about my insights. Hope to hear from you soon, best regards [Scholar Name]". As before, the link will redirect the user to a fake login page which resembles Google login page. Figure 2 shows the workflow to produce the proposed text.

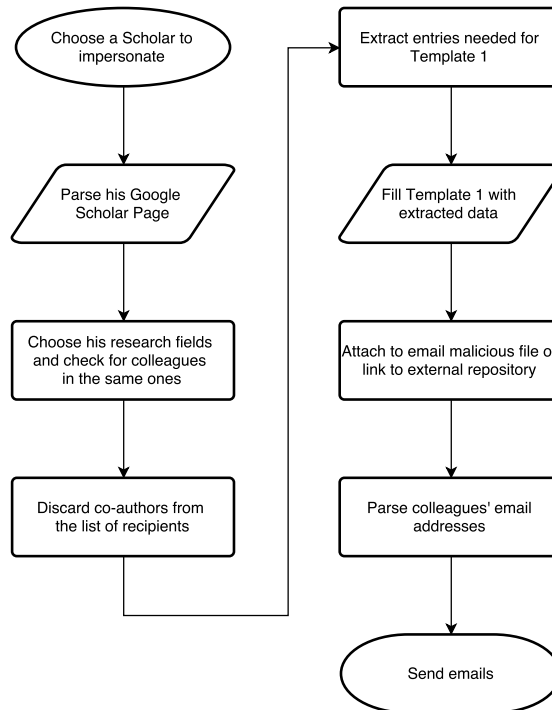


Fig. 2. Workflow of Template-driven NLG applied to Template 1.

Template 2 The second approach is riskier because the attacker has to impersonate someone close to the victim. A forged email could sound way too formal (e.g., nicknames are normally used between acquaintances), but it could pay since colleagues often read privately each others’ works, as a form of preliminary peer-review. The template is: “Hi [Colleague Name], as you know I’m working on [Topic Name] and I’ve written the attached paper. I would like to have some suggestions from you about it, since I’m a little uncertain about the solidity of the whole work. At [Fake Google Drive Link] you will find my draft. Thanks so much in advance, see you soon [Scholar Name]”. As before, the link will redirect the user to a fake login page which resembles Google login page. Figure 3 shows the workflow to produce the proposed text.

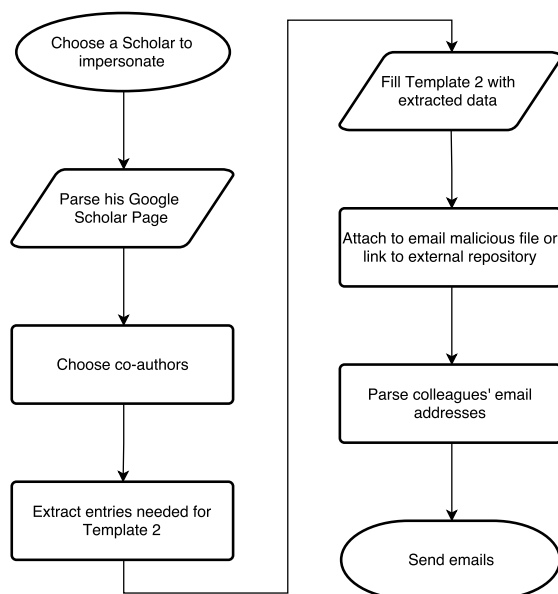


Fig. 3. Workflow of Template-driven NLG applied to Template 2.

4 Advanced NLG

In this section, we take one step more in our predictions and depict a scenario where Advanced NLG is used to create more convincing phishing emails.

As an example, an attacker could parse Google Scholar and look for researchers’ h-index trend during the past 5 years and their performance with respect to their colleagues in the same field. With such data, the attacker could use NLG techniques to summarize the trends in a natural language and incorporate this new information in a phishing email. The result would be something

along the line of “Dear [Target Name], I am Prof. [Reputable Name] from [Reputable University], we are currently expanding our lab and we are evaluating some possible candidates. I have personally checked your Google Scholar account and noticed that your h-index is high with respect of the average in your field, and considerably grew over the last 5-years span. Therefore, we would like to propose you a position. You can enrol at the following link: [Phishing Link]”. If the victim is successfully tricked, they might fill the recruitment form and give up valuable information to the attacker.

Besides the research community, similar techniques could be applied to social network users. For example, an attacker could use NLG techniques to identify and summarize a specific interest of a social network account and forward a private message along the line of: “Hello [Victim], I’ve seen that you’re really into [Music Genre]. I’m an event organizer and I would like to inform you that the next weekend, at the [Big Square] in [Victim Hometown], [Good Music Player] will give a short exhibition. All the details in the following link, feel free to attend! [Phishing Link]”. The phishing link would be a fake login page which mimics the social network one, created by the attacker. If the victim believes that a disconnection due to technical issues happened, they might try to login again, which would result in stolen credentials.

Despite the very simplistic nature of the aforementioned examples, the capability of translating heterogeneous information into natural language text can play a huge part in the next future of automatic phishing generation. Indeed, through Advanced NLG techniques attackers do much more than creating automatic text: they pinpoint submerged information which would take a huge effort to manually discover.

5 Conclusion

In this paper, firstly we have given a brief overview about the concept of Natural Language Generation (NLG) techniques and how these could be used to forge phishing emails. Moreover, we have expressed our concerns about how NLG could help hackers to easily target a community, such as the scientific one, proposing the concept of Community Targeted Phishing (CTP). We have provided two workflows that show how a Template-driven approach could be used to create phishing messages, and then we analysed how Advanced NLG techniques could enable phishers to achieve complex (thus, harder to spot) emails.

Future works will investigate different NLG techniques, in order to understand which ones provide the best results in terms of credibility and effectiveness. Moreover, we would like to investigate the variety that Advanced NLG techniques can ensure in text generation, as suggested in Section 4, and evaluate how traditional Bayesian spam filters perform against such phishing texts.

This work is just a first step, useful to introduce the very concept of CTP and raise awareness about a problem that might soon show up. We believe that using a NLG approach to target people with similar interests could be worthwhile,

since it would allow to create emails more effective than the general ones, and cheaper than the spear phishing ones.

References

1. Almomani, A., Gupta, B.B., Atawneh, S., Meulenberg, A., Almomani, E.: A survey of phishing email filtering techniques. *IEEE Communications Surveys Tutorials* 15(4), 2070–2090 (Fourth 2013)
2. Baki, S., Verma, R., Mukherjee, A., Gnawali, O.: Scaling and effectiveness of email masquerade attacks: Exploiting natural language generation. In: *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*. pp. 469–482. *ASIA CCS '17*, ACM, New York, NY, USA (2017), <http://doi.acm.org/10.1145/3052973.3053037>
3. Blanzieri, E., Bryl, A.: A survey of learning-based techniques of email spam filtering. *Artificial Intelligence Review* 29(1), 63–92 (Mar 2008), <https://doi.org/10.1007/s10462-009-9109-6>
4. Nelson, B., Barreno, M., Chi, F.J., Joseph, A.D., Rubinstein, B.I.P., Saini, U., Sutton, C., Tygar, J.D., Xia, K.: Exploiting machine learning to subvert your spam filter. In: *Proceedings of the 1st Usenix Workshop on Large-Scale Exploits and Emergent Threats*. pp. 7:1–7:9. *LEET'08*, USENIX Association, Berkeley, CA, USA (2008), <http://dl.acm.org/citation.cfm?id=1387709.1387716>
5. Reiter, E.: *Natural Language Generation*, pp. 574–598. Wiley-Blackwell (2010), <http://dx.doi.org/10.1002/9781444324044.ch20>
6. Sood, A.K., Enbody, R.J.: Targeted cyberattacks: A superset of advanced persistent threats. *IEEE Security Privacy* 11(1), 54–61 (Jan 2013)
7. Van Deemter, K., Krahmer, E., Theune, M.: Real versus template-based natural language generation: A false opposition? *Comput. Linguist.* 31(1), 15–24 (Mar 2005), <http://dx.doi.org/10.1162/0891201053630291>
8. Verma, R., Shashidhar, N., Hossain, N.: *Detecting Phishing Emails the Natural Language Way*, pp. 824–841. Springer Berlin Heidelberg, Berlin, Heidelberg (2012), https://doi.org/10.1007/978-3-642-33167-1_47