Association for Information Systems
AIS Electronic Library (AISeL)

Wirtschaftsinformatik Proceedings 2015   Wirtschaftsinformatik

3-4-2015

E-Mail Tracking in Online Marketing - Methods, Detection, and Usage

Benjamin Fabian

Benedict Bender

Lars Weimann

Follow this and additional works at: http://aisel.aisnet.org/wi2015

E-Mail Tracking in Online Marketing – Methods, Detection, and Usage

Benjamin Fabian, Benedict Bender, and Lars Weimann

Institute of Information Systems
Humboldt-Universität zu Berlin
Spandauer Str.1, 10178 Berlin, Germany
`{bfabian,benderbe,weimannl}@wiwi.hu-berlin.de`

Abstract. E-Mail tracking uses personalized links and pictures for gathering information on user behavior, for example, where, when, on what kind of device, and how often an e-mail has been read. This information can be very useful for marketing purposes. On the other hand, privacy and security requirements of customers could be violated by tracking. This paper examines how e-mail tracking works, how it can be detected automatically, and to what extent it is used in German e-commerce. We develop a detection model and software tool in order to collect and analyze more than 600 newsletter e-mails from companies of several different industries. The results show that the usage of e-mail tracking in Germany is prevalent but also varies depending on the industry.

Keywords: E-Mail Tracking, Online Marketing, Privacy

1 Introduction

In modern e-commerce, customer data has become critical for business success [12]. Business Intelligence based on personalization is used to optimize market positions, to engage in price discrimination [22], and to recommend users products they might buy in the future. Tracking user behavior is an important aspect of online marketing [22]. This does not only affect users browsing a commercial web site and actually looking for specific products. Tracking is also used in emails and has become a powerful instrument for marketing and personalization [2]. E-mail communication is an important marketing channel since it is still widely and frequently used. Moreover, it is a cheap and time-efficient medium because an e-mail only has to be designed once and can afterwards spread fast and to many users in parallel with low costs [6, p. 19]. Thus, e-mails are an important way to inform users and try to influence buying decisions. But how can companies ensure that users actually read e-mails and receive marketing information?

E-Mail tracking enables them to remotely observe, for instance, if an e-mail is opened, the time when a user reads an e-mail, the program in which the user opens it, and could also identify links on which the user clicks [9, p. 316]. This information is very useful for a company in order to understand customer behavior in more depth. Tracking data can also be sold to data aggregators or other companies that are interested in enriching their own data on consumers and their behavior. In particular, data

on actively used e-mail addresses is so valuable that some companies even specialize in this business segment and are selling such addresses.

All of this could create massive privacy concerns on the consumer side, in particular with privacy-sensitive users [3]. Our paper investigates how important tracking techniques work, how they can be detected, and to what extent they are used by online businesses in Germany.

2 E-Mail Tracking: Literature and Technology

In the following we discuss related work. In technical literature, some work has been conducted on email tracking in the sense of assessing whether an e-mail has been delivered successfully [19] [21]. This includes tracking in the sense of technically tracing the forwarding of mail through different hosts, including extensions of the mail protocols in *Requests of Comments* (RFCs) such as RFC 3798, RFC 3461, and RFCs 3885-3887 [19, p. 17]. With *Message Disposition Notification* (MDN) [17], the sender can request an acknowledgement of mail reception or reading from the receiver. From a perspective of privacy it is important to note that the receiver can choose if such an acknowledgement is sent or not [20, p. 131].

Other literature discusses tracking in the sense of snooping on mails for extracting content from a mail conversation or transmission without being an intended or authorized recipient [9, p. 307], or for grouping and combining related content [4] [10]. Sometimes tracking also refers to determining the real-world position of a person or device [8] [20, p. 131], or the linking of e-mail addresses to real-world identities, and is also related to security and privacy issues of using e-mail addresses for identification [13] [18]. The term e-mail tracking is sometimes also used in the sense of identifying the location of an e-mail as intangible object in a complex messaging network, for example, in the context of designing new decentralized mail architectures, e.g., involving peer-to-peer concepts [7] [14].

In our paper, e-mail tracking is understood as the remote logging of e-mail opening or reading, typically without user notice or consent. For this specific use of the term the following sources could be identified in the literature. The term e-mail tracking is explained by [11], including an example how a specific large company uses an e-mail tracking service. In [19, p. 17] an external service is described which generates a link for the e-mail content and sends that link to the receiver who needs to click on it in order to retrieve the mail content; this access is logged. Such an approach is very similar to the tracking link method described below but does not seem practical for everyday e-mail usage. In [20] the reaction of common mail software and providers to tracking pixels is investigated.

In contrast, our paper focuses on empirical investigation of the senders of tracking mails. Our goal is to assess to what extent e-mail tracking is used by companies in different industries. In our paper, two principle methods of tracking in e-mails are considered: *tracking links* and *tracking pixels*. Both approaches are based on personalization. The sender places a piece of content, addressed by a personalized link, or a picture with a receiver-specific name on a web server and includes the corresponding

link in the mail. When the receiver opens the mail, he or she accesses the content, either by downloading the picture from the web server or by opening a personalized link in the mail. This access is logged by the web server and can be associated with the receiver's customer profile. From now on, identifiers such as IP addresses and cookies, or more advanced web browser "fingerprinting" can be applied for further personalized tracking [1] [24], even across different web sites [16].
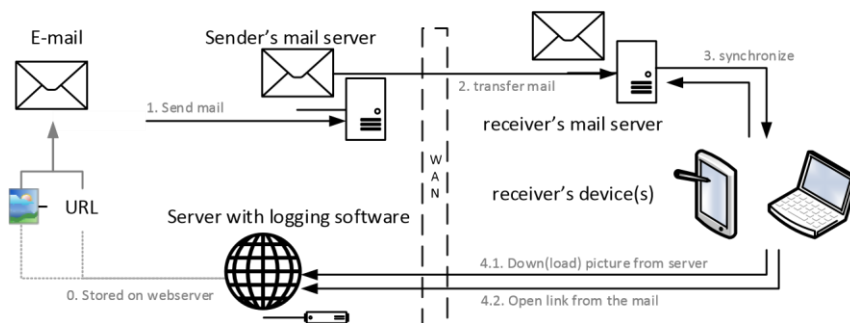

Fig. 1. How e-mail tracking works

Figure 1 illustrates these two main forms of e-mail tracking. When applying the *tracking link* method, the sender of an e-mail includes one or several links that refer to additional relevant content. Such a link has at least two components: an identifier for the content and an identifier for the receiver. These two components could also be combined into a single identifier. It is important that the user-identifying component is unique or that the combination of content identifier and user identifier is unique. The following example link contains an identifier for the content, the *cid* (content ID), and an identifier for the recipient of the mail, the *uid* (user ID): `http://www.example.com/content/view.php?cid=218471248&uid=38jd38 29d1`. The (combined) uniqueness requirement is important. If the same content will be referred to in another mail for another tracking target, the content ID can be reused and only a new user ID needs to be assigned.

After sending, the mail gets transferred to the outgoing mail server of the sender (1). The sender's mail server then attempts to transfer the mail to the receiver's mail server (2). Once this is successfully completed, the client can synchronize or download (depending on the protocol) the mail from the incoming mail server (3). After opening the mail, the user might click on links that are included in the mail (4.2). When the receiver opens the link and downloads the referred content from the server, this access is logged by the web server and can be associated with the receiver through the identifiers described above.

The *tracking pixel* method is also referred to as "web bug" method in privacy literature on web-user tracking [15]. It archetypically represents a stealthy tracking approach that can be transferred to new technology such as RFID [5]. Here, the sender of an e-mail includes a reference to a very small, usually invisible picture in the mail. Typically this picture does not have any border, has a width and height of 1 pixel, and

is transparent or at least has the same color as the background. It is also called a "pixel" because usually those images are in 1x1 format. To apply this technique, the *Uniform Resource Locator* (URL) of the picture needs to be unique. One way of realizing this is to assign a unique file name to the picture: `<img src=http://www.example.com/media/images/FI3RTG0-DOJFD3280.gif>`.

Once the sender has included a reference to the tracking picture, which is stored on a web server with a logging software (0), the mail is sent. The intermediate steps are the same as described above. When the receiver opens the mail on at least one device, the mail client displays the message and usually automatically downloads the referenced picture from the web server (4.1). This download access will be logged and, due to the unique identifier (typically the name), can be associated with the e-mail.

In practice, some problems may arise because many mail clients and providers offer options to block the download of external images [20]. However, users who would like to see other images referenced by the mail could also, inadvertently and without informed consent, download the tracking pixel once they choose to override this protection manually. This version of e-mail tracking only works if the mail was sent in the common HTML format.

Both techniques rely on the same principle of logging an access to a file on a web server. A client needs to use an application for accessing this file. The tracking link method triggers the client's default browser, whereas the tracking pixel approach uses an integrated HTML viewer. Both clients send some general information to the web server, for example, the application name and the device platform (from the *User-Agent-String*) [20, p. 131].

Furthermore, the client's IP address is transmitted, which allows to draw interesting conclusions for example about the location of the mail receiver and his or her Internet service provider. Another logged information is the access timestamp. By combining multiple log entries concerning the same file, the number of times a file was accessed and, therefore, how often an e-mail was viewed by the specific recipient can be derived. In general, file access itself gives the information that the e-mail was successfully delivered and opened by the client, and can be used for in-depth analyses of consumer behaviour [20, p. 130].

When applying the tracking link method, logging takes place at the moment a customer actively clicks on a link included in the mail. As a result, opening a link is a confirmation that a mail was read. With the tracking pixel method, the logging process is triggered when the mail client downloads the tiny picture to display the message which takes place when a customer opens the mail with his client. The combination of both methods could therefore improve data quality about a user's preference.

3 Experimental Design

In order to develop a detection method for tracking mails, and to compare several businesses with respect to their use of tracking, e-mails from different industry sectors had to be gathered. We chose not to use the content of our own personal e-mail inboxes because they may be biased by our personal interests and behaviour. Instead,

two new e-mail accounts were created at the same mail provider. This enabled comparison of links and ensured that there is a double check if both accounts receive similar e-mails. Both e-mail accounts shared the same properties except for a variation of the first name of the user. Other factors such as e-mail provider, age, and country were kept identical.

Both new e-mail addresses were used for registering to the same newsletters. It had to be ensured that enough e-mails would be received and that for every industry several companies would be selected. In total, 16 different industries were used, each including four different companies. Due to the fact that *Trade* industry is a very broad definition, it was sub-divided into the four sub-categories *All* (selling any kind of product), *Clothes, Electronic* and *Furniture*. The second e-mail address was used for comparison in order to investigate variations and to analyze the accuracy of the detection model. Table 1 gives an overview on the industries that have been investigated. After registration, for a five-month period from the end of 2013, companies were able to send newsletter e-mails, a selection of which we first analyzed manually by inspecting the HTML code.

Table 1. Overview of different industries and selected companies

No.	Industry	Representative companies			
1	E-Mail Germany	Telekom	GMX	Web.de	Freenet
2	Tourism	Expedia	Ab-in-den-Urlaub	Holidaycheck	Travel 24
3	Trading: All	Amazon	Ebay	Meinpaket	Otto
4	Trading: Clothes	Esprit	S. Oliver	Tommy Hilfiger	Kik24
5	Trading: Electronic	Alternate	Hardwareversand	Media Markt	Saturn
6	Trading: Furniture	IKEA	XXXL Lutz	Höffner	Möbel Boss
7	Healthcare	Bayer	Siemens Healthcare	GE Healthcare	Pfizer
8	Telecommunication	Vodafone	Mobilcom Debitel	Simyo	T-Mobile
9	Food	Vapiano	Block House	McDonald's	Burger King
10	IT	SAP	Sage	SalesForce	Lexware
11	Estate	Liegenschaftsfond	Immobilien Newsticker	Immobilienscout 24	Immowelt
12	Credit	AXA	Allianz	Sparkasse	Deutsche Bank
13	Automotive	Audi	BMW	Mercedes Benz	Toyota
14	Energy	EnBw	RWE	Vattenfall	E-on
15	Education	Berlitz	Lehrer online	Scoyo	SGD
16	Culture	Cinemaxx	Deutsche Oper	That's musical	Eventim

3.1 Detection Criteria

In our initial manual HTML code review, we analyzed every link and picture of 51 randomly selected mails. This enabled the possibility for comprehensive comparison of tracking techniques across industries and companies. Our code review in combination with input from the literature research constituted criteria for detecting e-mail tracking. For a tracking pixel, its border, width, and height are important properties. With tracking links, keywords such as the personal account name or the service provider are indicators. In addition, the difference between manual and automatic link

creation is an important criterion. This is induced by observing patterns in a pixel description or link, e.g., a change between upper and lower case or from characters to numbers and vice versa.

Table 2. Explanation of criteria

Criterion	Explanation
Switch between upper and lower case	Switching between upper and lower cases could indicate that this reference may not have been created manually and could have tracking potential. For example: `http://www.example.com/gp/r.html?R=`*`3D2P41Uo6HBRUzB`*`&C=`*`3D1VSYR`* *`mgw3HNG0`*`&=H=3D2OPCL0A4HDIAADJAOO5YIXRVNX4A`
Switch from characters to numbers and vice versa	Switching between characters and numbers could also indicate that this reference may not have been created manually and could have tracking potential. `http://www.example.com/gp/r.html?R=`*`3D2P41Uo6HBRUzB`*`&C=`*`3D1VSYR`* *`mgw3HNG0`*`&=H=`*`3D2OPCL0A4HDIAADJAOO5YIXRVNX4A`*

Border=0	Border of 0 displays the picture without any border, which makes it hard to detect for the user.
	
	`<img src=3D`
	`"http://www.example.com/g.html?uid=3DA.H.jwu.IU1w.CLk.3XmsTz`
	`remYG9OSEEMy4VT=Q" alt=3D""` *`border=3D"0"`* `height=3D"1"`
	`width=3D"1" />`

Width=1	Width of 1 ensures that the image is small enough to prevent noticing.
	
	`<img src=3D`
	`"http://www.example.com/g.html?uid=3DA.H.jwu.IU1w.CLk.3XmsTz`
	`remYG9OSEEMy4VT=Q" alt=3D"" border=3D"0" height=3D"1"`
	`width=3D"1"` `/>`

Height=1	Height of 1 ensures that the image is small enough to prevent noticing.
	
	`<img src=3D`
	`"http://www.example.com/g.html?uid=3DA.H.jwu.IU1w.CLk.3XmsTz`
	`remYG9OSEEMy4VT=Q" alt=3D"" border=3D"0"` *`height=3D"1"`*
	`width=3D"1" />`

Keywords such as "track", "code", service provider, or e-mail account	In addition, keywords such as "track" or the user's mail account name are used.
	
	`<img src=3D"https://`*`tracking`*`.example.com/op/0/XSPYOM6-`
	`XAZKIEC-145S7FE.gif" width=3D"1" height=3D"1" />`

A limitation of using image width or height is that if both values are modified, for instance, to 2, the threshold of 1 would not be matched and a tracking attempt could be missed. Correspondingly, an additional picture criterion "area" was introduced. Area calculates the product of height and width, and if the result is less than 10, the criterion is matched. As a further threshold, we require that a switch between upper and lower case or a switch between characters and numbers is used more often than twice. This allows for strings such as "Word2014" without triggering a tracking flag. With keywords, the identification model was configured to require three or more keyword occurrences in order to identify tracking. The following table summarizes the final criteria.

Table 3. Final criteria and thresholds for tracking mails

Criterion	Tracking Link	Tracking Pixel
Switch between upper and lower case: 3+ times	X	X
Switch from characters and numbers and vice versa: 3+ times	X	X
Border=0		X
Width=1		X
Height=1		X
Keywords, e.g., "track", "code", mail account: 3+		X

3.2   Detection Model

In the next step, the discovered criteria have been weighted (Tables 4 and 5). This reflects that some criteria are stronger indicators for tracking than others. This insight is based on the manual code review which showed that some values are used more often than others. As a consequence, the two switches between upper and lower case and characters and numbers are always included.

Table 4. Criteria weights for tracking link

Criterion	Weight
Switch between upper and lower case: 3+	30
Switch from characters and numbers and vice versa: 3+	40
Keywords: 3+	30
Sum	100

Table 5. Criteria weights for tracking pixel

Criterion	Weight
Switch between upper and lower case: 3+	30
Switch from characters and numbers and vice versa: 3+	40
Border=0	10
Width=1	40
Height=1	40
Area <10	40
Sum	200

The criteria and their weights were used to evaluate whether a tracking link or pixel is used or not. In general, if more than 60% of the identified criteria are matched, a link or pixel is considered a tracking attempt.

The following example illustrates the model and its usage. Assume that the HTML code of a pixel in one of the received e-mail contains the following: `<img src=3D` `"http://www.example.com/g.html?uid=3DA.H.jwu.IU1wCLk.3XmsTzremY` `G9osEeMy4VT=Q" alt=3D"" height=3D"4" width=3D"1"/>`. There are switches between upper and lower cases, even more than three times; thus, the first criterion is fulfilled. The next criterion, switches between characters and numbers, is also satisfied. The criterion border cannot be determined and is therefore not included in the model. However, the height is not 1 and thus this criterion does not match the threshold (`height=3D"4"`). But the criterion width is fulfilled (`width=3D"1"`). Moreover, criterion area is also satisfied because height is 4 and width is 1, which means area is 4, i.e., less than 10. The identification model is now used to evaluate whether this example is considered a tracking pixel or not. From the possible six criteria, five were found. This means a base score of 190. Out of these five criteria, four matched the criteria of the identification model. Summing up, the weight of these four is 150. In a next step, 150 is divided by 190, resulting in 0.79, which is higher than 0.6, and thus this picture is identified as a tracking pixel.

In summary, the manual code review is using the following approach: read every single e-mail manually, copy its links and pixels into a different file and compare for matching criteria of the identification model. This process renders the manual identification of tracking e-mails too time intensive for any larger-scale analysis. Therefore, we developed a software prototype to support the detection of email tracking.

4   Detection Prototype

The fundamental functionality of our detection prototype is the analysis of emails and its content, especially the mail source code. Our detection prototype is written in Java. Java is portable to any hardware, therefore the prototype could in theory run on any device which is able to execute a Java machine. The detection prototype is currently limited to only analyze e-mails from GMail, a web-hosted mail service offered by Google.

The following list describes the functions executed by our program:

1. The program asks the user for a valid combination of a user name and password. This combination is needed by GMail to access the user's e-mail account.
2. After the successful login to GMail, for which the Internet Message Access Protocol (IMAP) Application Programming Interface (API) is used, the program is accessing e-mails of the user account and starts to analyze them according to the identification model described earlier.
3. The program identifies e-mails in which either a tracking link or a tracking pixel was found and flags them with a star in the GMail inbox (Fig. 2).
4. It also displays a short analysis screen where the number of tracking links and tracking pixels is shown.

5. In addition, the program creates a log file in which every e-mail with a tracking approach is analyzed with respect to sender's location. The algorithm compares the sender's server location with a tracking link or a tracking pixel in the same e-mail. By this, the program investigates if the sender uses a different location for sending and tracking e-mails.

Fig. 2. (Anonymized) prototype screenshot: Flagging of tracking mails

This basic functionality was sufficient to conduct our study. The prototype can be customized to analyze local mailboxes and could also be extended to integration with further web-based mail services.

4.1 Evaluation of the Detection Prototype

In order to evaluate the prototype, the research team decided to first manually review and classify 51 randomly picked e-mails by an intense code review, and then let the software prototype run through the classified set of e-mails used as a ground truth or test set. We display the results for both tracking approaches separately. Table 6 represents the Confusion matrix [23] of detecting tracking pixels, while Table 7 refers to tracking links.

Table 6. Confusion matrix for detecting tracking pixels

		Predicted class	
		True	False
Actual class	True	35 (TP)	10 (FN)
	False	0 (FP)	6 (TN)

In order to evaluate the accuracy of the software prototype by a single metric, the following established metric for accuracy is applied: *Accuracy = (TP+FP) / (TP+TN+FN+FP)*. For the case of tracking pixels: Accuracy = (35+6)/51 = 0.8039.

This result illustrates that the software prototype was, compared to the manual code review, able to classify 80.39% of test e-mails correctly with respect to tracking pixels.

Table 7. Confusion matrix for detecting tracking links

		Predicted class	
		True	False
Actual class	True	33 (TP)	3 (FN)
	False	6 (FP)	9 (TN)

For tracking links, Accuracy = (33 + 9)/51 = 0.8235. Therefore, 82.35% of the test set e-mails have been classified correctly by the software with respect to tracking links.

5 Results

Overall, 604 different e-mails have been analyzed. Fig. 3 displays the number of e-mails received for each industry class. The *Trading* industry sent most of the emails in our sample. The top-3 e-mail senders are: *Trading: Electronic* with 150 received e-mails, *Trading: All* with 125 received e-mails and *Trading: Clothes* with 71 received e-mails. This result may be explained by the fact that the trading industry often changes their assortments and as a consequence is highly interested in marketing. Analyzing the same e-mails with respect to tracking, 591 of 604 e-mails are using at least one tracking approach, i.e., 97.85%. This means only 13 e-mails (2.15%) are not using any tracking method.
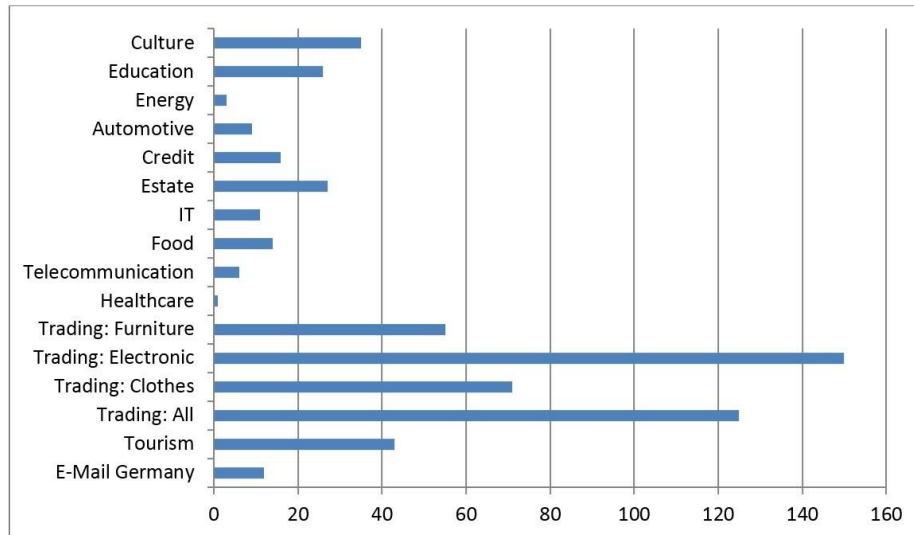
Fig. 3. E-Mails received, by industry

Then, the amount of tracking was analyzed in depth. In the 591 e-mails which are using tracking approaches, 421 tracking pixel and 428 tracking links were found; in total, 849 tracking approaches were identified. This result indicates that many e-mails are using more than one tracking approach, for instance a combination of tracking links and pixels. On average, 1.41 tracking approaches are included in a single e-mail. Fig. 4 displays the results.
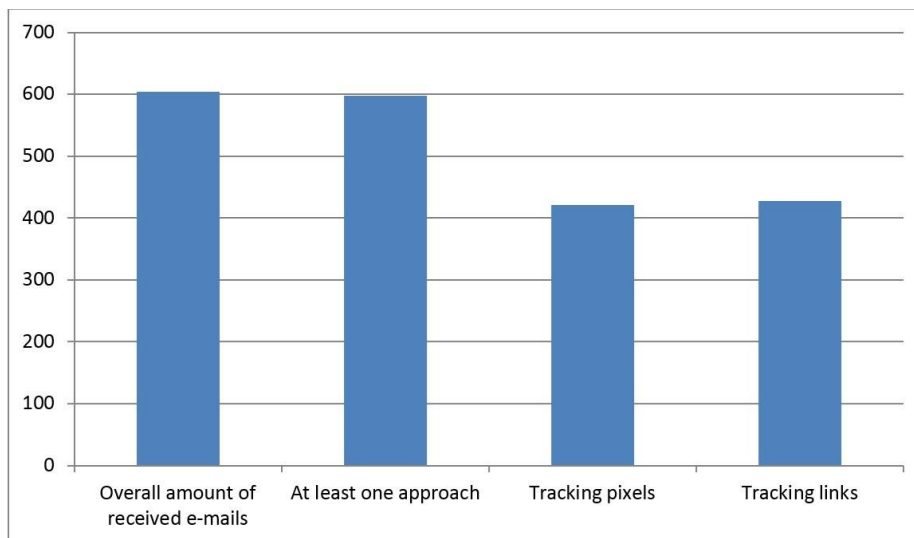

Fig. 4. Overview of tracking pixels and tracking links

The next step is the analysis of the different industries and how intensely they are using tracking. The following figure illustrates selected industries by their relative amount of tracking approaches per e-mail. Some industries are not illustrated because the amount of received e-mails was too low and thus the sample was not large enough to give a representative view. In general, more than ten e-mails had to be received in order to make an assumption about the industry sector. Thus, the industries *Healthcare, Telecommunication, Automotive*, and *Energy* are not represented. In addition, the relative amount of tracking approaches of all e-mails within one industry is selected. Otherwise, a higher absolute amount of e-mails in one industry would distort the results. Fig. 5 illustrates the relative amount of e-mails with tracking links and tracking pixels per industry. Both *Trading: Clothes* and *E-Mail Germany* sectors display a 100% ratio for both tracking approaches in our sample.

Finally, we investigated the locations of the servers that are hosting tracking pixels. In total, 268 locations could be determined. Table 8 lists the tracking server locations. Of course, the high percentage of tracking servers hosted in Germany is predictable. Furthermore, not surprising might be tracking servers in Austria due to the fact that several companies combining the German-speaking area to a business group and pool services within this area. An interesting result is the high percentage of servers in Ireland and the US.
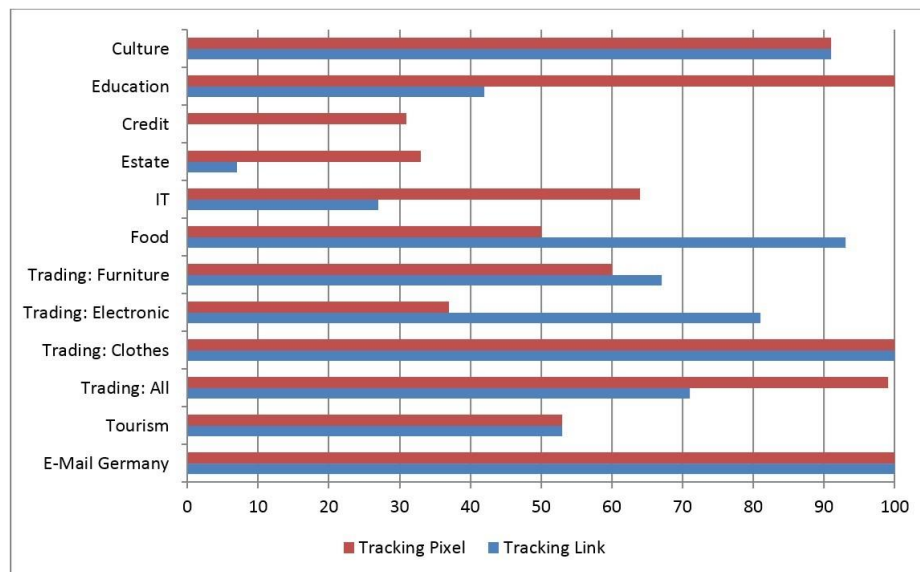

Fig. 5. Relative amount of tracking mails per industry

Table 8. Distribution of tracking server locations (for tracking pixels)

Country	Number of tracking pixels hosted	Percentage
Germany	185	67,27%
US	51	18,55%
Ireland	26	9,45%
Austria	5	1,85%
Belgium	1	0,36%

Fig. 6 shows the identified locations of the tracking servers we encountered. Our study confirms that several outsourcing service providers for e-mail tracking services do exist. Motivations for outsourcing could differ widely. One could be the necessity of choosing an offsite location for hosting tracking servers when tracking people's behaviour might constitute a violation of local law or privacy.
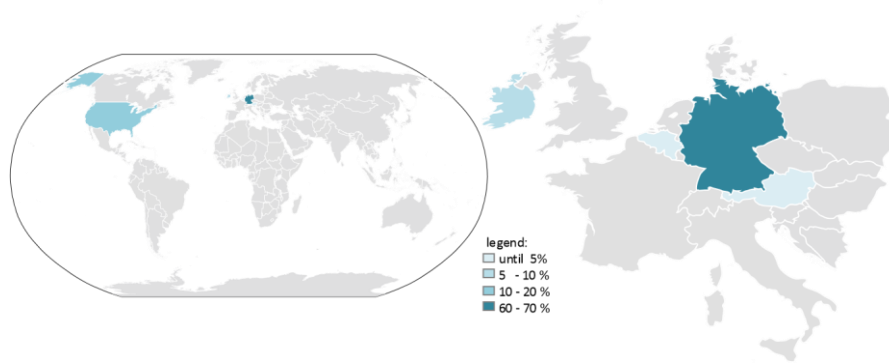

Fig. 6. Maps of tracking server locations

6 Limitations and Outlook

In our study, more than 600 e-mails have been analyzed with respect to the usage of tracking methods. In future work, an even broader and larger field of companies could be considered for analysis. Another possibility would be to focus on a specific industry and determine the use of e-mail tracking in depth. One could also analyze e-mails from different time periods and investigate how techniques and use of e-mail tracking changes over time.

Concerning our detection model, not all theoretical options are included and the model, therefore, will not detect some tracking mechanisms or will claim that tracking methods are used when in reality this is not the case. The implemented model has an accuracy of over 80% for detecting tracking pixels and an accuracy of over 82% for detecting tracking links with respect to our test set. Further improvements could be reached with optimization of the detection models in order to achieve a higher accuracy. In particular, determining whether a URL is personalized is of significant importance. Furthermore, the model could be extended to recognize further tracking

techniques. For example, a new tracking technique which has been revealed through our manual code review is to implement tracking pixels as the background of table cells.

For our analysis, a specialized software prototype was developed. This software prototype is integrated into GMail as an email provider. In order to enable end-users to flexibly detect tracking techniques in their e-mails, a vendor-independent software should be developed. One possibility is to implement different provider-specific APIs for a general detection software. Another approach is to develop a plugin for a specific mail client that would be provider-independent but mail client-dependent. An example for this scenario is a plugin for a mail client such as Thunderbird which offers powerful SDKs for extension development.

Another promising direction for future work is to develop methodologies for identifying similarities between tracking attempts. This would allow to deduce information about the software or service used by a company to manage e-mail marketing. Clustering of tracking pixels involves the study of schematic similarities and also reflects if they are provided by the company itself or by an external provider that can be identified, e.g., by a WHOIS query.

For an end-user, the question of how to protect his or her privacy is of great relevance. Up to now, the most effective protection against tracking links is to not open any personalized link in e-mails, and against tracking pixels – to disable the automatic download of external images in the e-mail client. However, with in-depth studies of tracking servers, corresponding blacklists could support automatic link filtering for increased privacy protection.

7   Conclusion

Our study shows that both tracking links and tracking pixels are widely used in commercial practice: 97.85% of all e-mails received in this study contained at least one e-mail tracking method. Concerning different industries, there are business sectors in our study where only a few mails but also others where the entirety of messages are containing tracking attempts.

References

1. Acar, G., Juarez, M., Nikiforakis, N., Diaz, C., Gürses, S., Piessens, F., Preneel, B.: FPDetective: Dusting the Web for Fingerprinters. In: Proceedings of the 20th ACM Conference on Computer and Communications Security (CCS 2013) (2013)
2. Ansari, A., Mela, C.F.: E-Customization. Journal of Marketing Research 40(2), 131–145 (2003)
3. Awad, N.F., Krishnan, M.S.: The Personalization Privacy Paradox: An Empirical Evaluation of Information Transparency and the Willingness to Be Profiled Online for Personalization. MIS Quarterly 30(1), 13–28 (03 2006)
4. Bacchelli, A., D'Ambros, M., Lanza, M.: Extracting Source Code from E-Mails. In: Proceedings of the 18th IEEE International Conference on Program Comprehension (ICPC). pp. 24–33 (2010)

5. Bauer, M., Fabian, B., Fischmann, M., Gürses, S.: Emerging Markets for RFID Traces. arXiv preprint cs/0606018 (2006)
6. Becker, L.: Professionelles E-Mail-Management: Von der individuellen Nutzung zur unternehmensweiten Anwendung. Gabler / GWV, Wiesbaden (2009)
7. Bercovici, S., Keidar, I., Tal, A.: Decentralized Electronic Mail. In: Workshops 26th IEEE International Conference on Distributed Computing Systems (ICDCS) (2006)
8. Chipperfield, T.R., Cooper, J.S., Foulger, M.G., Storms, A.C.: System and Method Related to Generating and Tracking an Email Campaign. https://www.google.com/patents/WO2001082112A3?cl=en (2003)
9. Cole, E., Nordfelt, M., Ring, S., Fair, T.: Cyber Spying Tracking Your Family's (Sometimes) Secret Online Lives. Elsevier Science (2005)
10. Cselle, G., Albrecht, K., Wattenhofer, R.: BuzzTrack: Topic detection and tracking in email. In: Proceedings of the 12th International Conference on Intelligent User Interfaces. pp. 190–197. ACM (2007)
11. Evers, J.: How HP Bugged E-Mail. http://news.cnet.com/How-HP-bugged-e-mail/2100-1029_3-6121048.html (2006)
12. Goldfarb, A., Tucker, C.: Privacy and Innovation. Innovation Policy and the Economy 12(1), 65–90 (January 2012)
13. Jin, L., Takabi, H., Joshi, J.B.: Security and Privacy Risks of Using E-Mail Address as an Identity. In: Proceedings of the Second International Conference on Social Computing (SocialCom). pp. 906–913. IEEE (2010)
14. Kangasharju, J., Ross, K.W., Turner, D.A.: Secure and Resilient Peer-to-Peer E-Mail Design and Implementation. In: Proceedings of the Third International Conference on Peer-to-Peer Computing (P2P 2003). pp. 184–191. IEEE (2003)
15. Martin, D., Wu, H., Alsaid, A.: Hidden Surveillance by Web Sites: Web Bugs in Contemporary Use. Communications of the ACM 46(12), 258–264 (Dec 2003)
16. Mayer, J., Mitchell, J.: Third-Party Web Tracking: Policy and Technology. In: Proceedings of the IEEE Symposium on Security and Privacy (SP 2012). pp. 413– 427 (May 2012)
17. Network Working Group: RFC 3798: Message Disposition Notification. https://tools.ietf.org/html/rfc3798 (2004)
18. Nguyen, D.H., Hayes, G.R.: Information Privacy in Institutional and End-User Tracking and Recording Technologies. Personal and Ubiquitous Computing 14(1), 53–72 (2009)
19. Oppliger, R.: Providing Certified Mail Services on the Internet. IEEE Security & Privacy 5(1), 16–22 (2007)
20. Schmidt, J.: E-Mail im Visier: Tracking im Alltag aufspüren und abstellen. c't Magazin für Computertechnik 22/2013, 130–135 (2013)
21. Surmacz, T.: Reliability of E-Mail Delivery in the Era of Spam. In: Proceedings of the 2nd International Conference on Dependability of Computer Systems (DepCoSREL-COMEX'07). pp. 198–204. IEEE (2007)
22. Taylor, C.R.: Consumer Privacy and the Market for Customer Information. RAND Journal of Economics pp. 631–650 (2004)
23. Visa, S., Ramsay, B., Ralescu, A.L., van der Knaap, Esther: Confusion Matrix-based Feature Selection. In: Proceedings of the 22nd Midwest Artificial Intelligence and Cognitive Science Conference (MAICS). pp. 120–127 (2011)
24. Yen, T.F., Xie, Y., Yu, F., Yu, R.P., Abadi, M.: Host Fingerprinting and Tracking on the Web: Privacy and Security Implications. In: NDSS Symposium 2012, San Diego, California, USA (2012)