

AI-based Re-identification of Behavioral Clickstream Data

Stefan Vamosi

WU Vienna University of Economics and Business

Michael Platzer

MOSTLY AI

Thomas Reutterer

WU Vienna University of Economics and Business

Cite as:

Vamosi Stefan, Platzer Michael, Reutterer Thomas (2022), AI-based Re-identification of Behavioral Clickstream Data. *Proceedings of the European Marketing Academy*, 51st, (106830)

Paper from the 51st Annual EMAC Conference, Budapest, May 24-27, 2022



AI-based Re-identification of Behavioral Clickstream Data

Abstract

AI-based face recognition, i.e., the re-identification of individuals within images, is an already well established technology for video surveillance, for user authentication, for tagging photos of friends, etc. This paper demonstrates that similar techniques can be applied to successfully re-identify individuals purely based on their behavioral patterns. In contrast to de-anonymization attacks based on record linkage, these methods do not require any overlap in data points between a released dataset and an identified auxiliary dataset. The mere resemblance of behavioral patterns between records is sufficient to correctly attribute behavioral data to identified individuals. Further, we can demonstrate that data perturbation does not provide protection, unless a significant share of data utility is being destroyed. These findings call for sincere cautions when sharing actual behavioral data with third parties, as modern-day privacy regulations, like the GDPR, define their scope based on the ability to re-identify. This has also strong implications for the Marketing domain, when dealing with potentially re-identify-able data sources like shopping behavior, clickstream data or cookies. We also demonstrate how synthetic data can offer a viable alternative, that is shown to be resilient against our introduced AI-based re-identification attacks.

Keywords: Privacy, Re-Identification, Clickstream Behavior

Track: Methods, Modelling & Marketing Analytics

1 Introduction

The steady rise of digital native business formats witnesses the tremendous opportunities offered by an exploding amount and variety of individual-level, behavioral micro-data accruing in a broad range of industries. For example, firms like Amazon, Netflix or Meta track the behavior of their customers to derive personalized recommendations and targeted marketing actions. Other companies realize that sharing customer information with other parties (e.g., linked with “Internet of Things” elements, such as mobile tracking meters, medical or fitness devices, etc.) can create synergies for both sides. Likewise, the non-profit sector and research institutions increasingly rely on the availability or “share-ability” of publicly available or open behavioral data (Beaulieu-Jones et al., 2019).

However, all these benefits are in strong contrast to the legitimate desire of individuals to protect their privacy and to refrain from sharing their personal data (Wieringa et al., 2021). In the vein of the Facebook-Cambridge Analytica scandal, also firms have increasingly become sensitive to protect their customer data against re-identification attacks and their brands against a loss in customer trust (Schneider et al., 2017, 2018). All these concerns lead to modern privacy regulations (in particular EU’s GDPR and California’s CCPA) which impose very strict standards for data anonymization. Both the GDPR and the CCPA do not specify any specific process for anonymization, but they demand the outcome to be irreversible to prevent re-identification of individuals by all the means reasonably likely to be used (Finck & Pallas, 2020).

Against this background, the key challenge for many firms is to keep benefiting from data-driven marketing while maintaining the privacy of their customers’ data. Previous research already showed that conventional perturbation techniques (such as adding random noise, masking, or obfuscation) fail to do so in the presence of high dimensional, highly correlated data, which typically arise when observing individuals over an extended period of time, i.e. for sequential personal data. For example, Narayanan & Shmatikov (2006) document successful re-identification attacks in the context of Netflix user’s movie ratings, De Mon-

tjoye et al. (2013) for human mobility traces and De Montjoye et al. (2015) for credit card retail transaction data.

In this research, we extend this perspective and show that a powerful general-purpose, AI-based model recently proposed by Vamosi et al. (2021) is capable to re-identify behavioral data in a highly effective way. As we will demonstrate in detail, this makes standard “anonymization” techniques inapt to protect individual-level sequential data against a new breed of AI-based pattern attacks. We also show that data synthetization can help to manage the trade-off between preserving the useful information in the original data, while reducing the risk of violating privacy.

2 Behavioral Pattern Attacks

In 2006 Netflix released an “anonymized” dataset to participants of a Machine Learning competition. The dataset consisted of 470,000 Netflix users, and their complete history of over 100 million movie ratings across 18,000 movies. Despite the fact, that only a subset of the actual customer base has been released, despite the omission of all customer-level attributes, and despite the injection of random noise to the data, researchers were quickly able to perform and publish a successful re-identification attack (Narayanan & Shmatikov, 2006). They achieved so by fuzzy matching the released de-identified records with publicly available, yet identified records, that were obtained from a popular movie rating website, IMDB. For every successful linkage they would then have been able to expose the complete Netflix rating history of users, who only shared a fraction of their ratings publicly via IMDB.

The above attack scenario requires that there is a temporal overlap (coinciding events) between the released and the auxiliary dataset. We take this further, and show that a successful linkage doesn’t necessarily rely on shared records, but can also be performed by extracting and comparing behavioral patterns among disjoint set of records. Any existence of overlapping data points, particular at subject-level

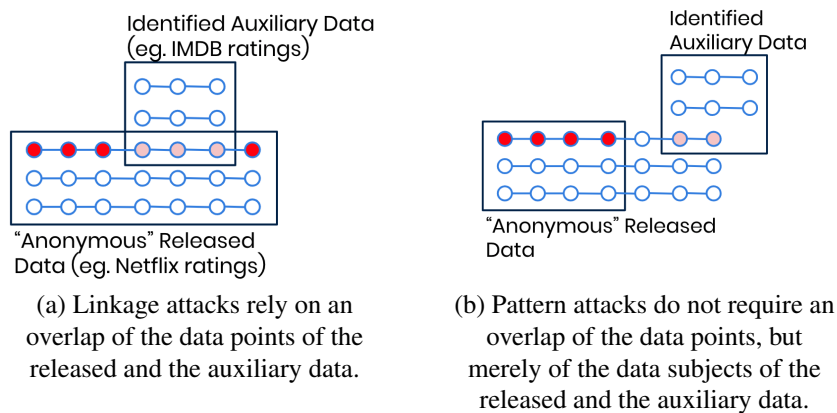


Figure 1: Angle of Attack

like basic sociodemographic attributes, would further increase the likelihood of success. The attack scenario (see Figure 1) is outlined as follows:

1. Organization releases an "anonymous" behavioral dataset D for period P1
2. Attacker learns characteristic patterns of individuals from D
3. Attacker obtains auxiliary data A of a known user X for period P2
4. Attacker extracts and matches the characteristic patterns of D and A
5. If successful, attacker can then reveal activities of user X within D

Vamosi et al. (2021) introduced an AI model, based on a Triplet-Loss Recurrent Neural Networks (TL-RNN), that can learn and abstract characteristic patterns of individuals from sequential data. That model optimizes for an embedding space, so that two different sequences of the same user end up closer together than two sequence of two different users. A similar approach has been introduced and shown to be effective for face recognition in Schroff et al. (2015). And the same idea can now be applied to re-identify individuals based on behavioral categorical data, i.e., based on sequences of structured data records. Once the model is being

trained, any sequence of data can then be mapped into a corresponding embedding space, yielding a numeric latent vector that represents the most characteristic traits of that sequence. The re-identification can then be simply performed via a Nearest-Neighbor Search in that embedding space.

3 Empirical Re-Identification Study

For our empirical demonstration, we use 2018 data from the Comscore Web Browser Panel, which provides continuous tracking of the visited websites among its panelists. We assume that only the visited sequence of domains (i.e., no timestamps, no visit duration, etc.) for January to June data (P1) is being released for 4'000 active, yet “anonymous” panelists, and we attempt to re-identify 1'000 individual panelists based on their observed July data (P2). Note, that there is no overlap in period P1 and P2, and each individual exhibits a different set of records during these periods. The question is then, whether the patterns within the data are characteristic and specific enough, so that we are still able to link individuals across periods. In total we have between 500 and 2'500 visits per panelist recorded, with an average of 1'290 visits, all across 115'000 distinct domains.

First of, the TL-RNN model is fitted to P1 data, with the records of each panelist being split by calendar week into several sequences. The Triplet Loss then optimizes the weights of the neural network, so that the week-wise records of the same panelist are embedded in its latent space closer together than the records of different panelists. The fitted TL-RNN model is then used to map the last week of P1, and the full July information of P2 into its 128 dimensional embedding space, to then search for nearest neighbors therein.

A pure random guess would result in only a 1 in a 4,000 chance of a successful match, i.e., a 0.025% probability of identifying the correct P1 panelist for a given P2 panelist. However, after training the TL-RNN model on the P1 data of those 4'000 customers, and use the corresponding embeddings to identify nearest neighbors of P1 sequences and P2 sequences, we are able to correctly link

Permutation	N1N	N3N	N5N	N10N
0%	49.9%	60.6%	65.6%	71.5%
10%	42.8%	52.0%	56.4%	63.8%
20%	34.0%	46.5%	52.9%	60.4%
30%	26.6%	36.4%	40.5%	48.7%
60%	1.1%	1.4%	1.5%	1.8%

Table 1: Ten, five, three and nearest neighbor(s) that include the correctly re-identified individual. In percent of totally 1,000 auxiliary users.

499 customers - thus have a success of re-identification of 49.9%. Relaxing the success criteria by considering the five nearest neighbors within P1, then we find matches for 65.6% of the panelists. This demonstrates that the re-identification on behavioral traits alone, without any overlap in data points, is indeed possible and becoming practical. The addition of any overlapping information, like subject-level attributes, would only further boost the rate for a successful match.

The next question is then, whether a data publisher can prevent re-identification by injecting random noise to the released dataset D. While this has already been shown to be of limited effect to protect against linkage attacks Narayanan & Shmatikov (2006), a pattern attack is expected to be even more resilient against noise. We thus constructed a perturbed dataset D' by replacing any data point of an individual with 30% probability with a data point from any other subject. Even though this is already a highly destructive mechanism, 26.6% of individuals could still be re-identified. Even when going up to 60% of data points being randomly permuted, 1.1% of individuals were successfully re-identified. Thus, the presented type of re-identification attack is robust against noise, where all results of the re-identification task are presented in Table 1.

A completely different approach towards privacy-safe sharing of granular-level data is the emerging domain of AI-based data synthetization. A generative

model is being fitted to the original dataset, that then allows an arbitrary number of new, yet statistically representative data records to be created. Synthetic data is particularly known for its applications for image (Karras et al., 2017) and text (Brown et al., 2020) generation, but can be applied to structured behavioral data (Lin et al., 2020) just as well. As there is no 1:1 relationship between actual and synthetic subjects, the re-identification is by definition not possible. Yet, the learning algorithm might still leak information on individuals into the released dataset via memorization, and thus empirical validation methods to assess the privacy of synthetic data are being introduced (see Platzer & Reutterer (2021) and Alaa et al. (2021)). For this study, we pick up on the privacy concept, that the synthetic subjects shall not be any closer to the training subjects than to holdout subjects. And the embedding space of TL-RNN provides us with the relevant, task-specific distance measure, yielding a sensitive metric towards subject-level memorization.

Thus, we split the June data of the 4,000 panelists into equally sized training and holdout set, and generate 2,000 synthetic subjects based on the former with a publicly available synthetization software¹. We then measure the distances of each synthetic subject to its nearest training subject, as well as to its nearest holdout subject within the previously constructed TL-RNN embedding space. The resulting average distance to training is 0.731, which is near identical to the average distance to holdout of 0.737. Another metric to investigate, is the share of synthetic records that ends up closer to a training subject than to a holdout subject. As we've split the original dataset 50/50, but are not using the holdout for synthetization, the expected lower bound for that share is at 50%. In our empirical study, 51.8% of synthetic subjects end up closer to a training subject, and resp. 48.2% closer to a holdout. Thus, the synthetic records are shown to be nearly just as likely close to a training subject than close to a holdout subject, that was never seen before.

This last metric is particularly remarkable, as even for a perturbed dataset, for which 90% of data points were being randomly replaced, we still have 57.5%

¹<https://mostly.ai/>

Flip	10%	20%	30%	40%	50%	60%	70%	80%	90%	synt.
DCR share	98.5%	89.9%	70.6%	62.6%	58.7%	58.1%	57.6%	57.5%	55.4%	51.8%

Table 2: Privacy Test based on Embedding

of records being closer to the original dataset, than to another records. According to Figure 2, its cumulative distribution function confirms the conclusion, that synthetic data set is the closest to the holdout compared to perturbed records.

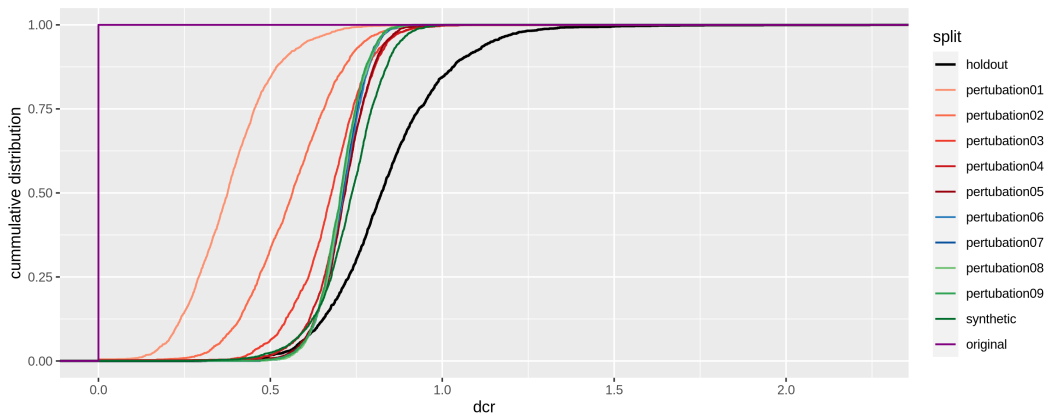


Figure 2: Cumulative distribution function of the distance to closest ratio for different permutations, synthetic data and holdout.

4 Conclusions

We have demonstrated how novel AI models can be leveraged to successfully re-identify behavioral data. In addition, we have shown that basic anonymization techniques, such as random perturbation, offer little safety against such pattern attacks. With these methods becoming increasingly stronger, and more commonly

available, this has wider implication for what data can still be considered truly anonymous in the context of modern-day privacy regulations. Finally, we have also shown that data synthetization can offer a viable alternative to releasing data, as it defies even strong empirical privacy evaluations. The conference presentation will cover content, that exceeds this written demonstration.

References

- Ahmed M Alaa, Boris van Breugel, Evgeny Saveliev, and Mihaela van der Schaar. How faithful is your synthetic data? sample-level metrics for evaluating and auditing generative models. *arXiv preprint arXiv:2102.08921*, 2021.
- Brett K Beaulieu-Jones, Zhiwei Steven Wu, Chris Williams, Ran Lee, Sanjeev P Bhavnani, James Brian Byrd, and Casey S Greene. Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, 12(7):e005122, 2019.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Yves-Alexandre De Montjoye, César A Hidalgo, Michel Verleysen, and Vincent D Blondel. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports*, 3(1):1–5, 2013.
- Yves-Alexandre De Montjoye, Laura Radaelli, Vivek Kumar Singh, and Alex “Sandy” Pentland. Unique in the shopping mall: On the reidentifiability of credit card metadata. *Science*, 347(6221):536–539, 2015.
- Michèle Finck and Frank Pallas. They who must not be identified—distinguishing personal from non-personal data under the gdpr. *International Data Privacy Law*, 2020.

- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- Zinan Lin, Alankar Jain, Chen Wang, Giulia Fanti, and Vyas Sekar. Using gans for sharing networked time series data: Challenges, initial promise, and open questions. In *Proceedings of the ACM Internet Measurement Conference*, pp. 464–483, 2020.
- Arvind Narayanan and Vitaly Shmatikov. How to break anonymity of the netflix prize dataset. corr, abs. *arXiv preprint cs/0610105*, 2006.
- Michael Platzer and Thomas Reutterer. Holdout-based empirical assessment of mixed-type synthetic data. *Frontiers in big Data*, 4, 2021.
- Matthew J Schneider, Sharan Jagpal, Sachin Gupta, Shaobo Li, and Yan Yu. Protecting customer privacy when marketing with second-party data. *International Journal of Research in Marketing*, 34(3):593–603, 2017.
- Matthew J Schneider, Sharan Jagpal, Sachin Gupta, Shaobo Li, and Yan Yu. A flexible method for protecting marketing data: An application to point-of-sale data. *Marketing Science*, 37(1):153–171, 2018.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 815–823, 2015.
- Stefan Vamosi, Thomas Reutterer, and Michael Platzer. A deep recurrent triplet loss model for learning sequence similarities for user-identification. 2021.
- Jaap Wieringa, PK Kannan, Xiao Ma, Thomas Reutterer, Hans Risselada, and Bernd Skiera. Data analytics in a privacy-concerned world. *Journal of Business Research*, 122:915–925, 2021.