University of Massachusetts Amherst

ICDCS 2021

Preserving Privacy in Personalized Models for Distributed Mobile Services

Akanksha Atrey, Prashant Shenoy, David Jensen

University of Massachusetts Amherst

International Conference on Distributed Computing Systems 2021

Cloud-Based Mobile Services

- Distributed between cloud and end-device
- Use contexts (e.g. location services)
- Many use prediction of future contexts, not just current context, to tailor service

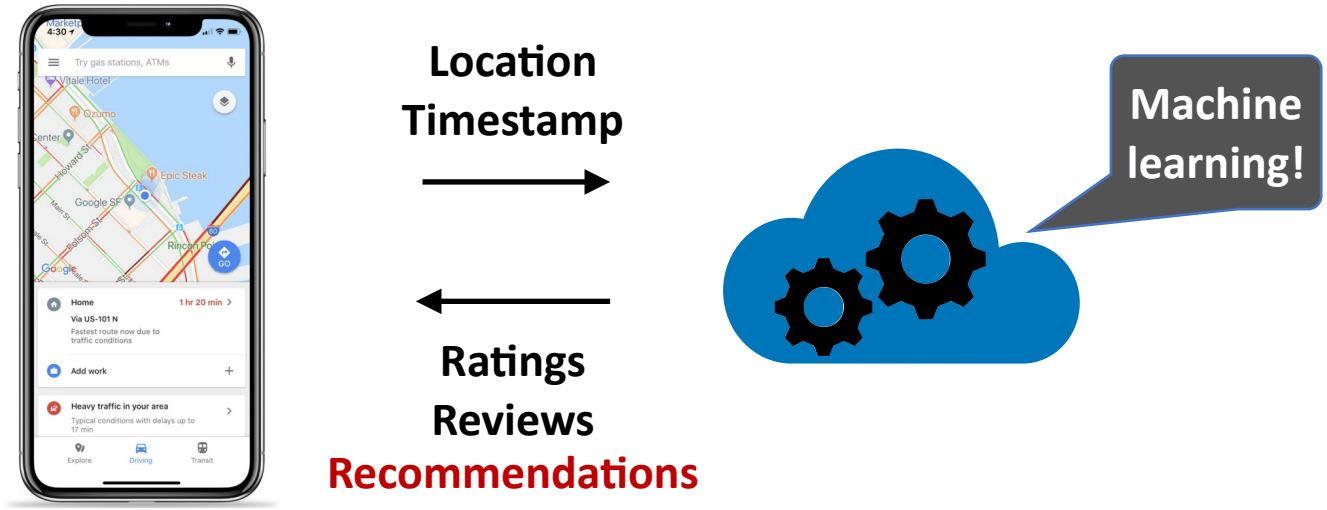

Modeling User Behavior via Personalized Models

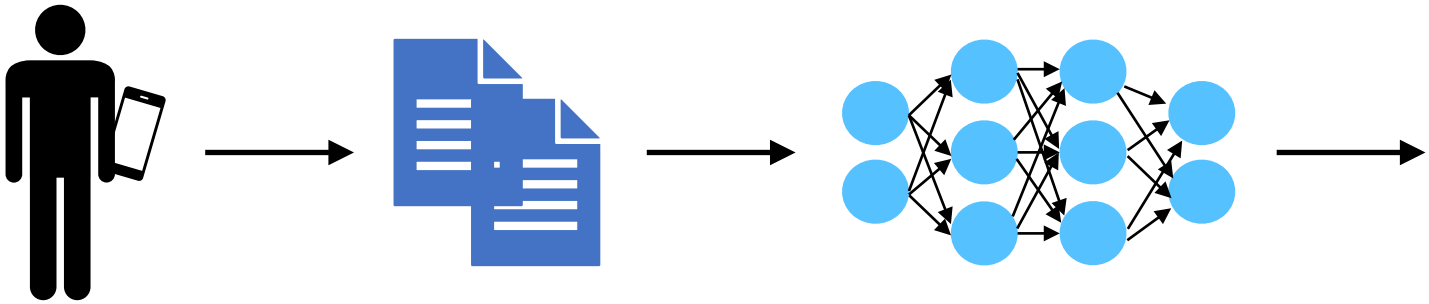

- ✓ Encode user-specific behavior
- ✓ Offer better efficacy for users with dissimilar behavior
- ✓ Less computationally expensive
- ? Privacy

Modeling User Behavior via Personalized Models

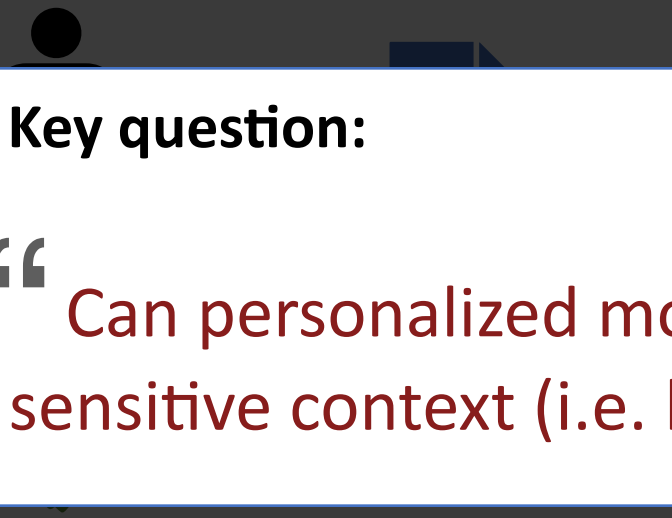

Key question:

"Can personalized mobile services be exploited to leak sensitive context (i.e. location) information about a user?"

- ✓ Offer better efficacy for users with dissimilar behavior
- ✓ Less computationally expensive
- ? Privacy

Outline

- Motivation
- **Privacy attack**
- Privacy-Preserving Personalization via *Pelican*
- Conclusion

Model Inversion Attacks

Key idea: learn sensitive features in input data using output of a trained model

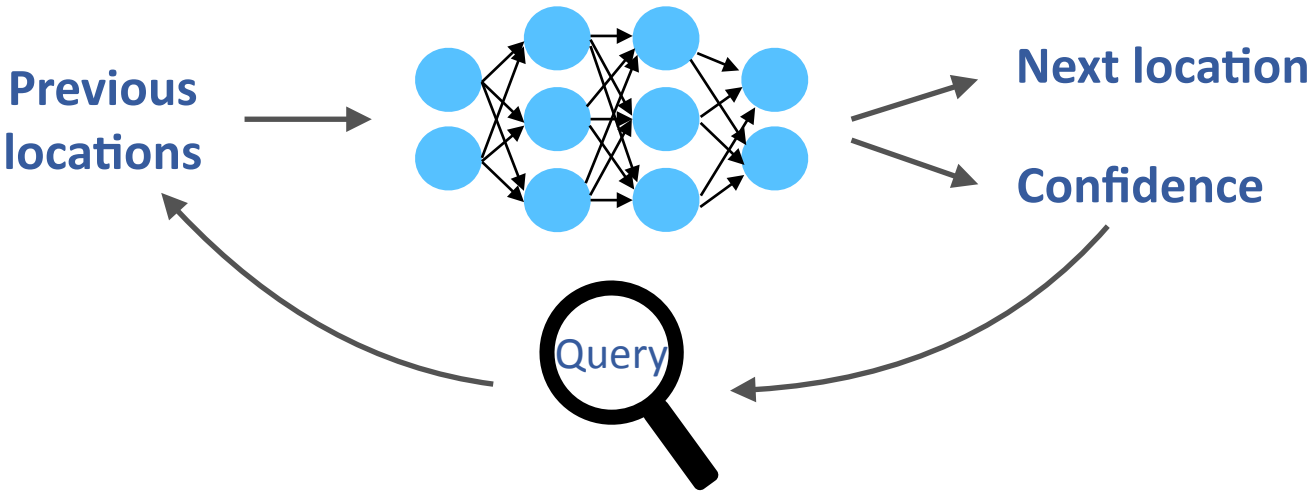

Approach: Maximize model's confidence on correct output to determine best input value

Ours: Time-Series Inversion Attack

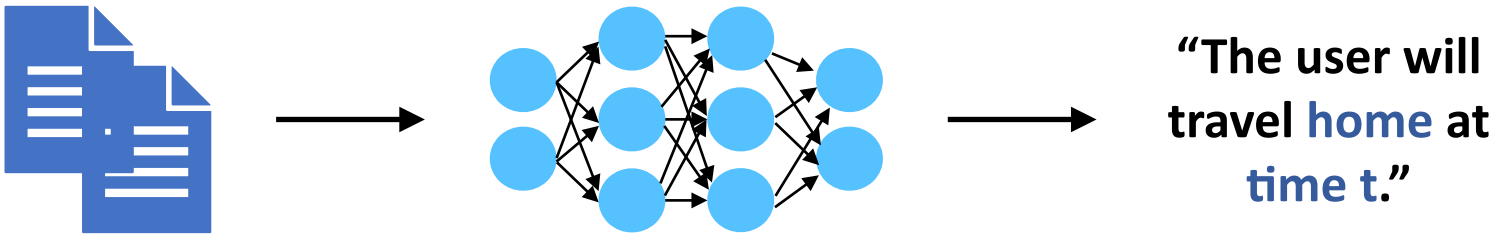

$x_t = \{f_1, f_2, \ldots, f_n\}$

$M : x_{t-2}, x_{t-1} \rightarrow l_t$

- Full access or no access to all data features at a time step
- Use cross correlation between sequences to infer features

Results — Attack Efficacy

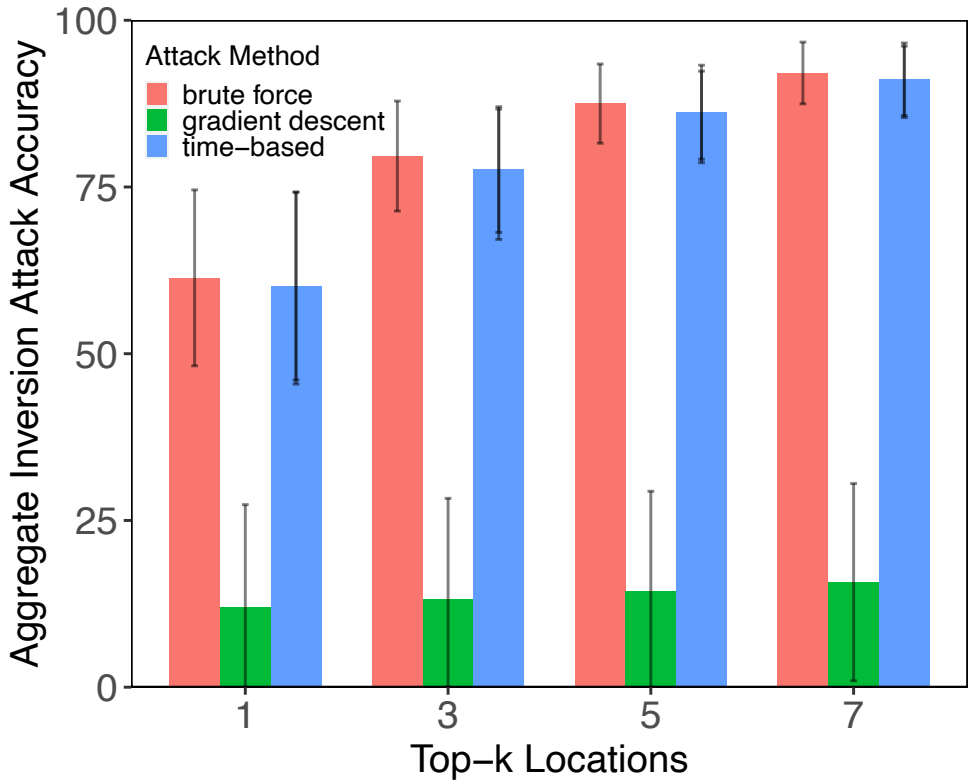

Method	Runtime (hours)
Brute Force	82.18
Gradient Descent	6.27
Time-Based	0.68

Time-based privacy attack is computationally efficient and effective with 77.61% accuracy for top-3 estimates.

Key Insight:

"Context-aware personalized models can be easily exploited with limited information for users with highly correlated mobility patterns."

Outline

- Motivation
- Privacy attack
- **Privacy-Preserving Personalization via *Pelican***
- Conclusion

Proposed Privacy-Preserving Personalization Framework: **Pelican**

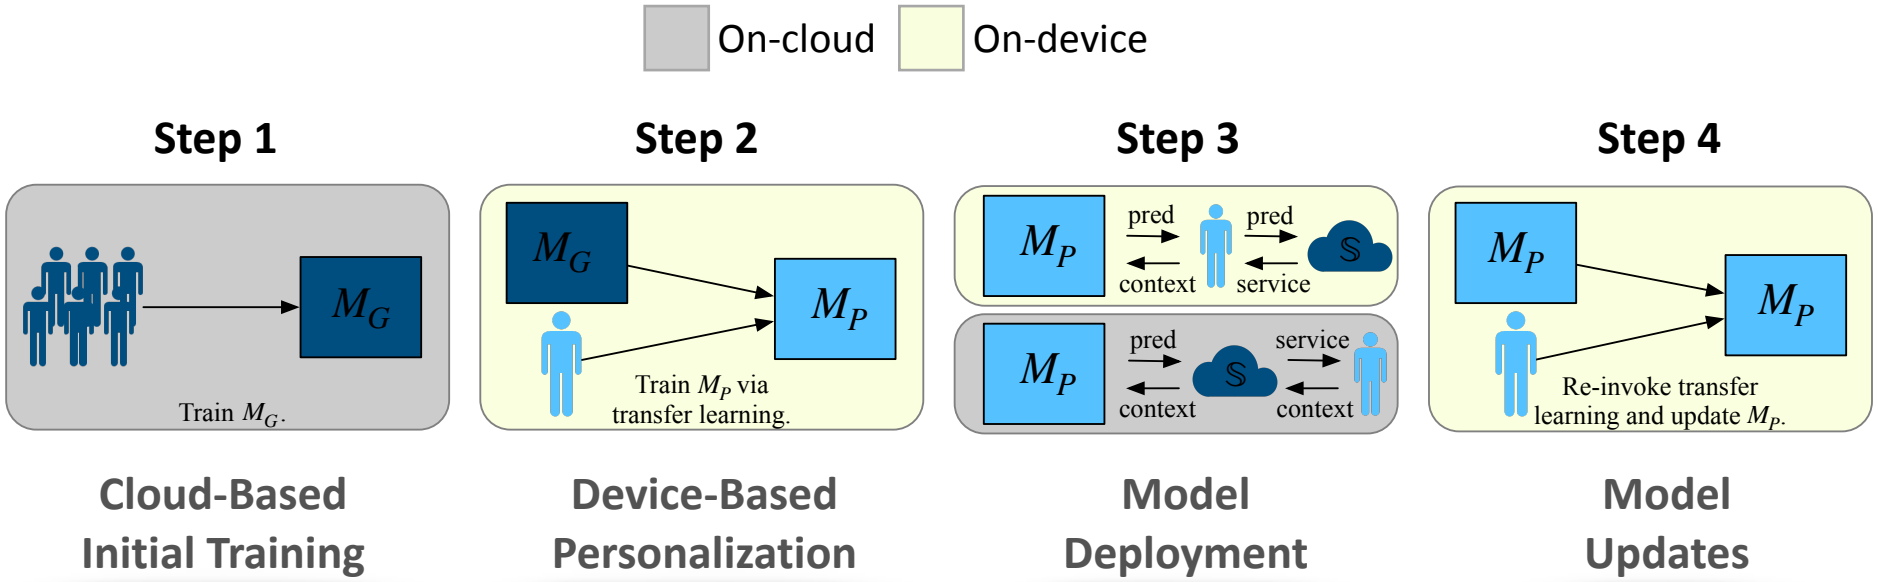

Step 1: Train a general model using training data from many users

Step 2: Personalize the general model using transfer learning methods

Step 3: On-device or cloud deployment with privacy enhancements

Step 4: Re-invoke transfer learning process to update the model with new data

Privacy Enhancement

Privacy tuner $\mathcal{T}$ **that scales the output before Softmax layer.**

$$p_i = \frac{\exp\left(z_i/\mathcal{T}\right)}{\sum_i \exp\left(z_i/\mathcal{T}\right)}$$

$\mathcal{T}$ changes sensitivity to the different outputs at inference time only.

Results — Privacy Enhancement

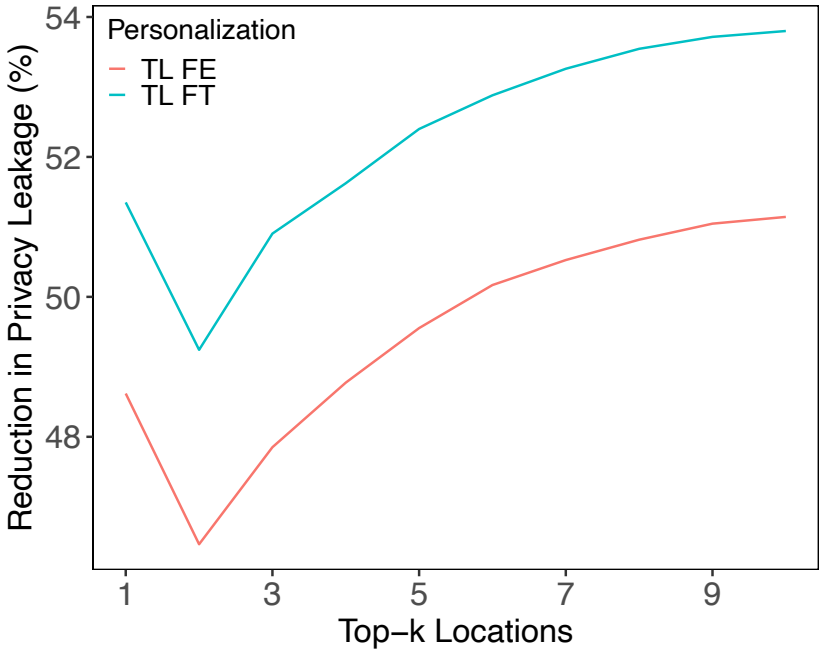

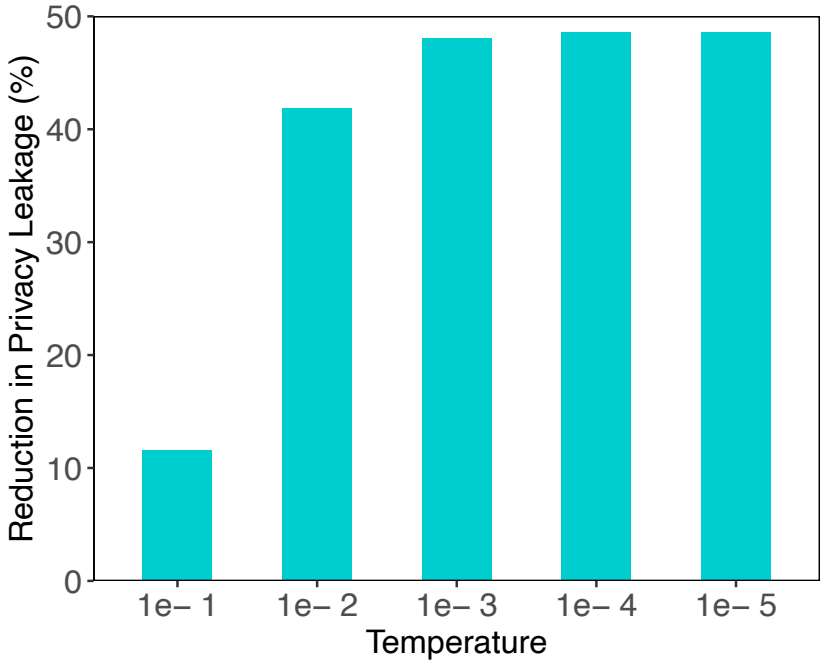

Pelican **is able to thwart privacy attacks in personalized models with up to 75.41% reduction in leakage while achieving state-of-the-art performance.**

Conclusion

- Personalized models are capable of leaking private information despite being trained in a privacy-preserving manner.
- The nature of context-aware applications results in more leakage of private information than expected by user.
- The proposed distributed framework learns and deploys transfer learning-based personalized ML models in a privacy preserving manner on resource-constrained mobile devices.

Preserving Privacy in Personalized Models for Distributed Mobile Services

Akanksha Atrey
University of Massachusetts Amherst
aatrey@cs.umass.edu

Prashant Shenoy
University of Massachusetts Amherst
shenoy@cs.umass.edu

David Jensen
University of Massachusetts Amherst
jensen@cs.umass.edu

***Abstract*—The ubiquity of mobile devices has led to the proliferation of mobile services that provide personalized and context-aware content to their users. Modern mobile services are distributed between end-devices, such as smartphones, and remote servers that reside in the cloud. Such services thrive on their ability to predict future contexts to pre-fetch content or make context-specific recommendations. An increasingly common method to predict future contexts, such as location, is via machine learning (ML) models. Recent work in context prediction has focused on ML model personalization where a personalized model is learned for each individual user in order to tailor predictions or recommendations to a user's mobile behavior. While the use of personalized models increases efficacy of the mobile service, we argue that it increases privacy risk since a personalized model encodes contextual behavior unique to each user. To demonstrate these privacy risks, we present several attribute inference-based privacy attacks and show that such attacks can leak privacy with up to 78% efficacy for top-3 predictions. We present *Pelican*, a privacy-preserving personalization system for context-aware mobile services that leverages both device and cloud resources to personalize ML models while minimizing the risk of privacy leakage for users. We evaluate *Pelican* using real world traces for location-aware mobile services and show that *Pelican* can substantially reduce privacy leakage by up to 75%.**

***Index Terms*—cloud-based mobile services, personalized ML models, privacy, deep learning, context-awareness**

I. Introduction

The ubiquitous nature of smartphones and smart devices, such as wearables, have led to a plethora of online mobile services in various domains including fitness, entertainment, news and smart homes. Such mobile services tend to be distributed between the end-device and the cloud with front-end components running on the devices as mobile applications and back-end components running on cloud servers. Modern mobile services are often context-aware to provide tailored content or service to users based on their current context. For example, it is common for a restaurant recommendation service to use location as its context when recommending nearby eateries. While the use of *current* context in mobile services is common, mobile services have begun to use machine learning (ML) models to predict *future* contexts (e.g., a user's next or future location(s)) and provide tailored recommendation based on these prediction (e.g., suggest directions or store closing time of predicted future location).

Machine learning has been used in mobile services for tasks such as next location prediction [1], medical disease detection [2] and language modeling [3]. The popularity of deep learning has established the use of aggregated data from a large number of users to train and deploy a general ML model that makes predictions for context-aware services for a broad range of users. A more recent trend in the field is to use personalized models on a per-user basis rather than a general model to further improve the efficacy of the service. In this scenario, rather than using a single ML model for all users, a model is personalized for each user using training data specific to the user. For instance, a user's frequently visited locations in a mobile service or a user's viewing history in a streaming service can be used to develop personalized ML models.

While model personalization is a growing trend in mobile and Internet of Things services, in this paper, we examine the implications of such an approach on the privacy of individuals. We argue that personalized ML models encode sensitive information in the single-user context traces used as training data and mobile services that use such personalized models can leak privacy information through a class of privacy attacks known as model inversion. Model inversion attacks exploit a trained ML model to infer sensitive attributes [4]. While ML researchers have studied inversion attacks in other contexts, they have not been studied or demonstrated for time-series models that are commonplace in mobile applications. Our work formalizes and demonstrates such attacks for personalized mobile services by showing how they can leak sensitive context (i.e. location) information about a user. To the best of our knowledge, privacy implications of personalized models in distributed mobile services have not been previously studied.

Motivated by the need to ensure the privacy of personalized ML models, we present *Pelican*, an end-to-end system for training and deploying personalized ML models for context-aware mobile services. Our system enhances user privacy by performing sensitive personalized training on a user's device and adding privacy enhancements to personalized models to further reduce and prevent inversion attacks from leaking sensitive user information. Our system is also designed to allow low overhead model updates to improve model accuracy while safeguarding privacy. Finally, our system leverages the device and cloud architecture of mobile services when personalizing ML models to enhance user privacy. In design and implementation of *Pelican*, we make the following contributions:

C1 We adapt low-resource transfer learning methods to train and execute personalized ML models on resource-