AEDA

Sven Haadem
sven@aeda.no
+47 99 10 75 23

Can insurance companies cooperate and still provide value?

Data is shared among each contributor in a fair way

While adhering to GDPR
& Supervisory Requirements

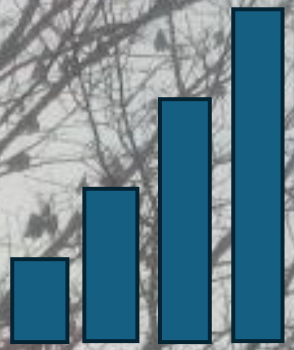

More Players More Value

AEDA

Being a Central Part always us to provide add on value

Model Builder

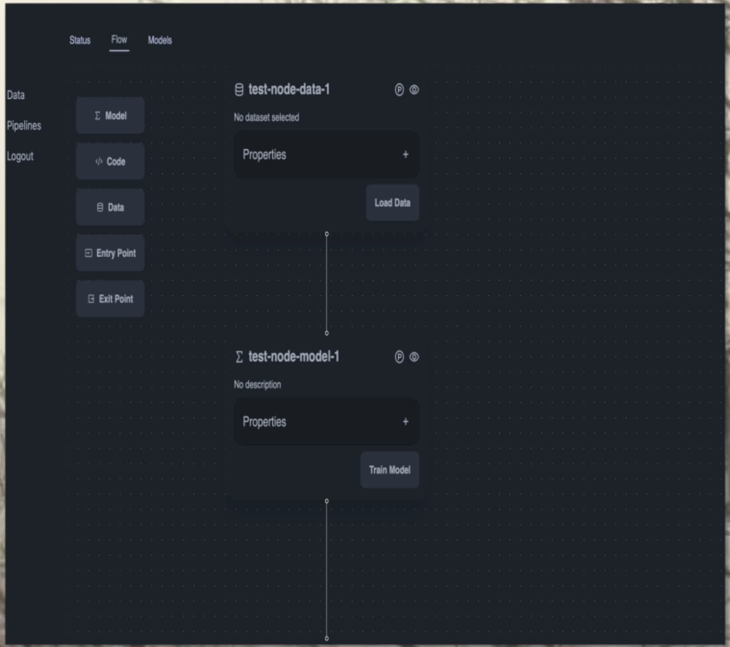

Dashboard

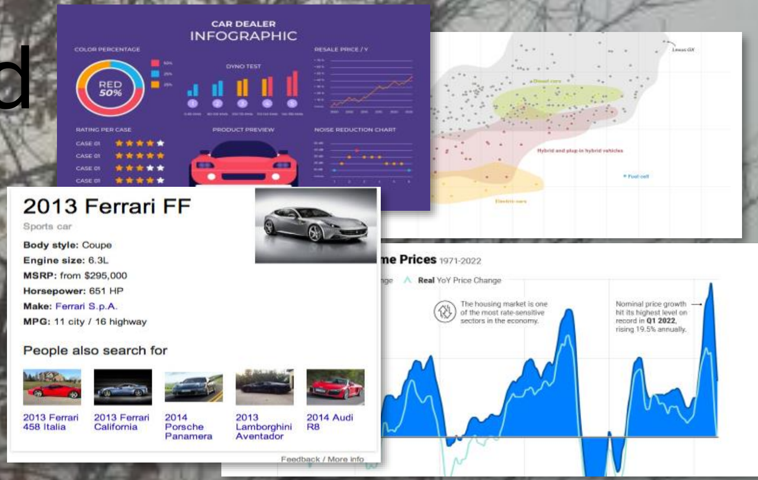

AEDA

Anonymization

Deep Learning with Differential Privacy

October 25, 2016

Martín Abadi* H. Brendan McMahan* Andy Chu* Ilya Mironov* Li Zhang* Ian Goodfellow† Kunal Talwar*

ABSTRACT

Machine learning techniques based on neural networks are achieving remarkable results in a wide variety of domains. Often, the training of models requires large, representative datasets, which may be crowdsourced and contain sensitive information. The models should not expose private information in these datasets. Addressing this goal, we develop new algorithmic techniques for learning and a refined analysis of

1. We demonstrate that, by tracking detailed information (higher moments) of the privacy loss, we can obtain much tighter estimates on the overall privacy loss, both asymptotically and empirically.
2. We improve the computational efficiency of differentially private training by introducing new techniques. These techniques include efficient algorithms for computing gradients for individual training examples, sub-

Clustering

Efficient k-Anonymization Using Clustering Techniques*

Ji-Won Byun[1], Ashish Kamra[2], Elisa Bertino[1], and Ninghui Li[1]

[1] CERIAS and Computer Science, Purdue University
{byunj, bertino, ninghui}@cs.purdue.edu
[2] CERIAS and Electrical and Computer Engineering, Purdue University
akamra@purdue.edu

Abstract. k-anonymization techniques have been the focus of intense research in the last few years. An important requirement for such techniques is to ensure anonymization of data while at the same time min-

Dynamic Tech Stack

Solid Theoretical Foundation

Differential Privacy for Deep and Federated Learning: A Survey

AHMED EL OUADRHIRI AND AHMED ABDELHADI, (Senior Member, IEEE)
Department of Engineering Technology, University of Houston, Houston, TX 77204, USA
Corresponding author: Ahmed El Ouadrhiri (aeouadrh@central.uh.edu)

ABSTRACT Users' privacy is vulnerable at all stages of the deep learning process. Sensitive information of users may be disclosed during data collection, during training, or even after releasing the trained learning model. Differential privacy (DP) is one of the main approaches proven to ensure strong privacy protection in data analysis. DP protects the users' privacy by adding noise to the original dataset or the learning parameters.

AEDA

Sven Haadem
sven@aeda.no
+47 99 10 75 23