2021 NSF CYBER-PHYSICAL SYSTEMS PRINCIPAL INVESTIGATORS' MEETING

CPS: Medium: Collaborative Research: Security vs. Privacy in Cyber-Physical Systems

AWARD #s 1929410 & 1837517

Alvaro A. Cardenas
Jairo Giraldo
Luis Burbano
Gabriel Torres
UNIVERSITY OF CALIFORNIA SANTA CRUZ

Murat Kantarcioglu
Mestan Celiktug
UT DALLAS

Jonathan Katz
Madison Krell
UNIVERSITY OF MARYLAND

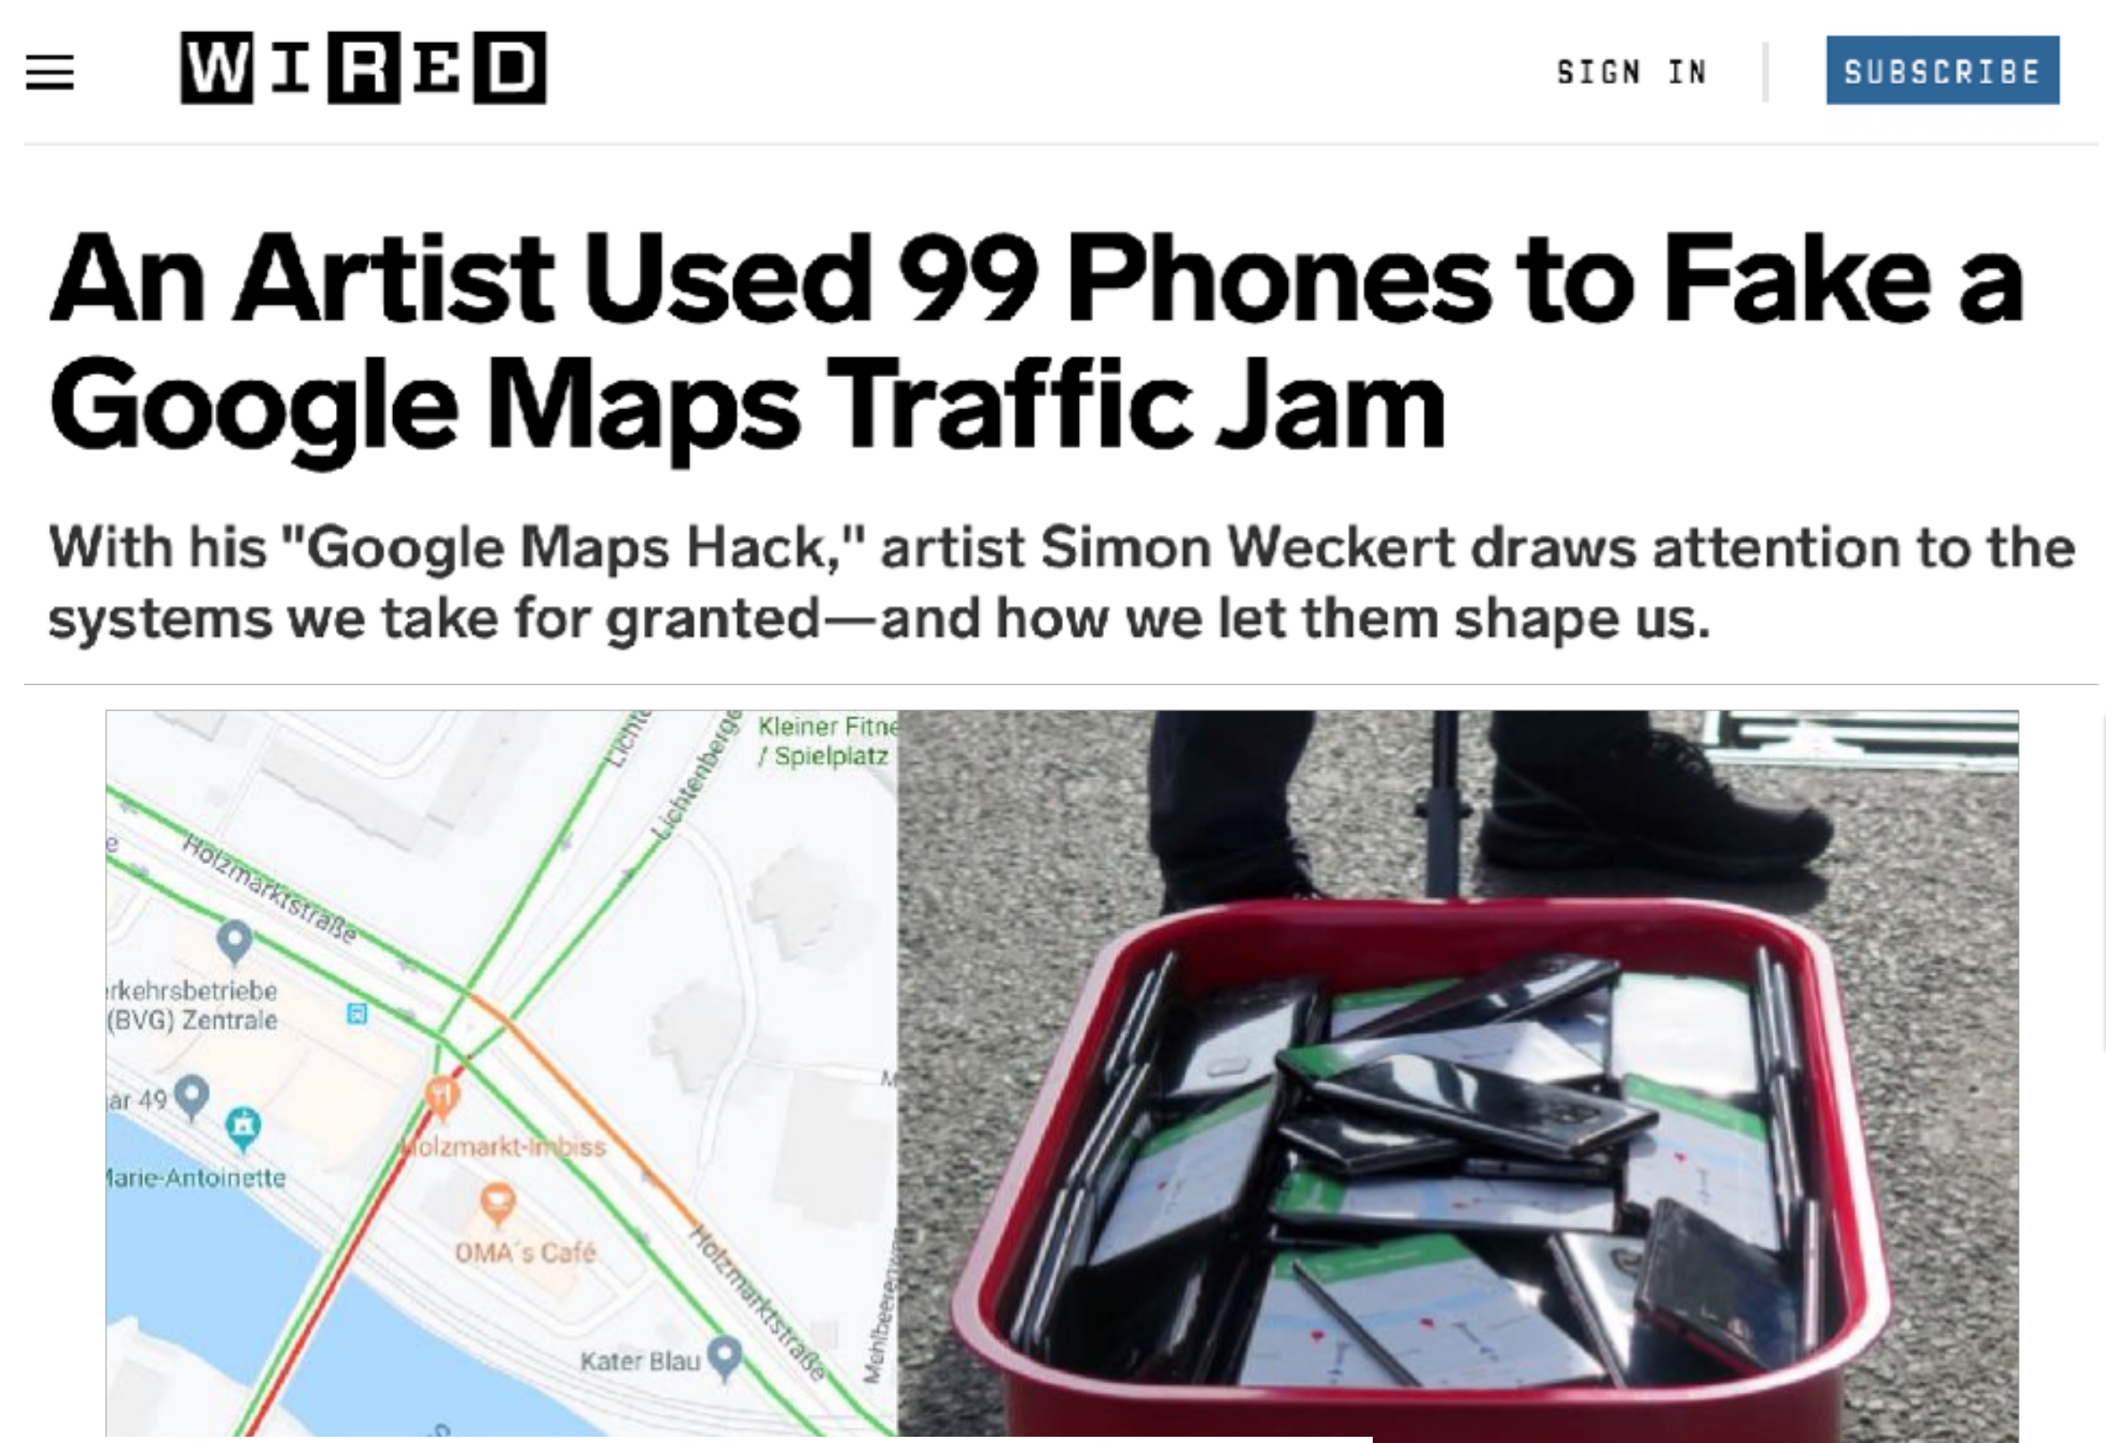

An Artist Used 99 Phones to Fake a Google Maps Traffic Jam

With his "Google Maps Hack," artist Simon Weckert draws attention to the systems we take for granted—and how we let them shape us.

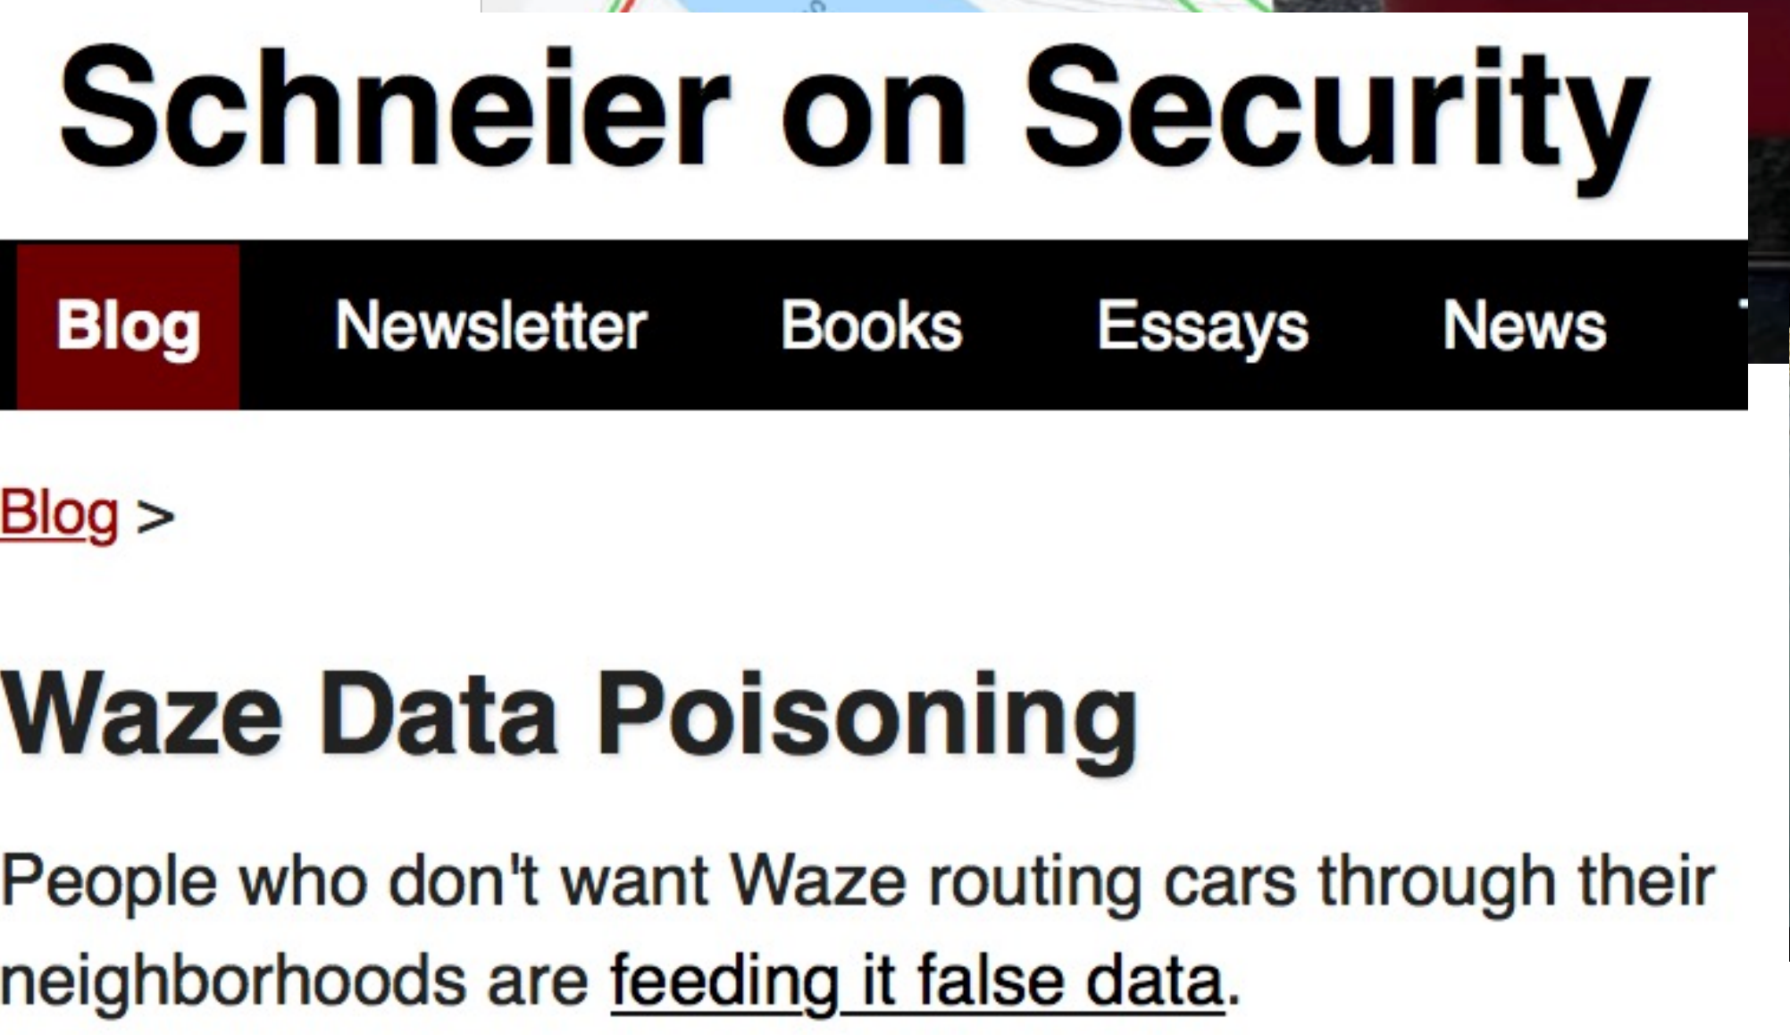

Blog >

Waze Data Poisoning

People who don't want Waze routing cars through their neighborhoods are feeding it false data.

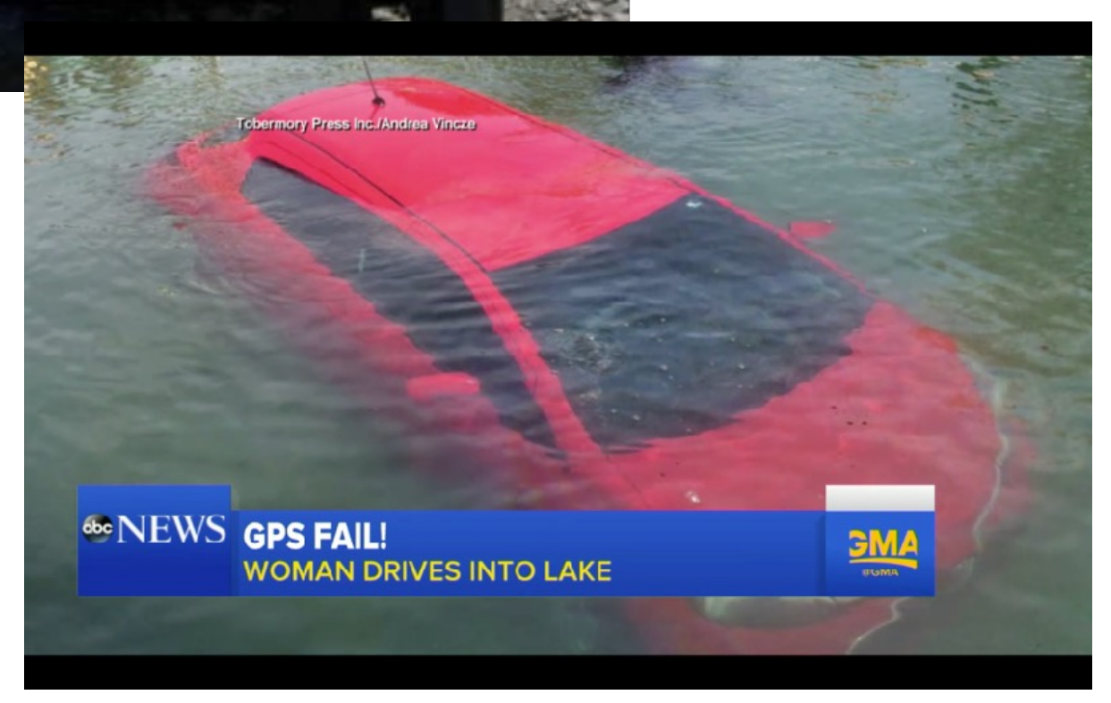

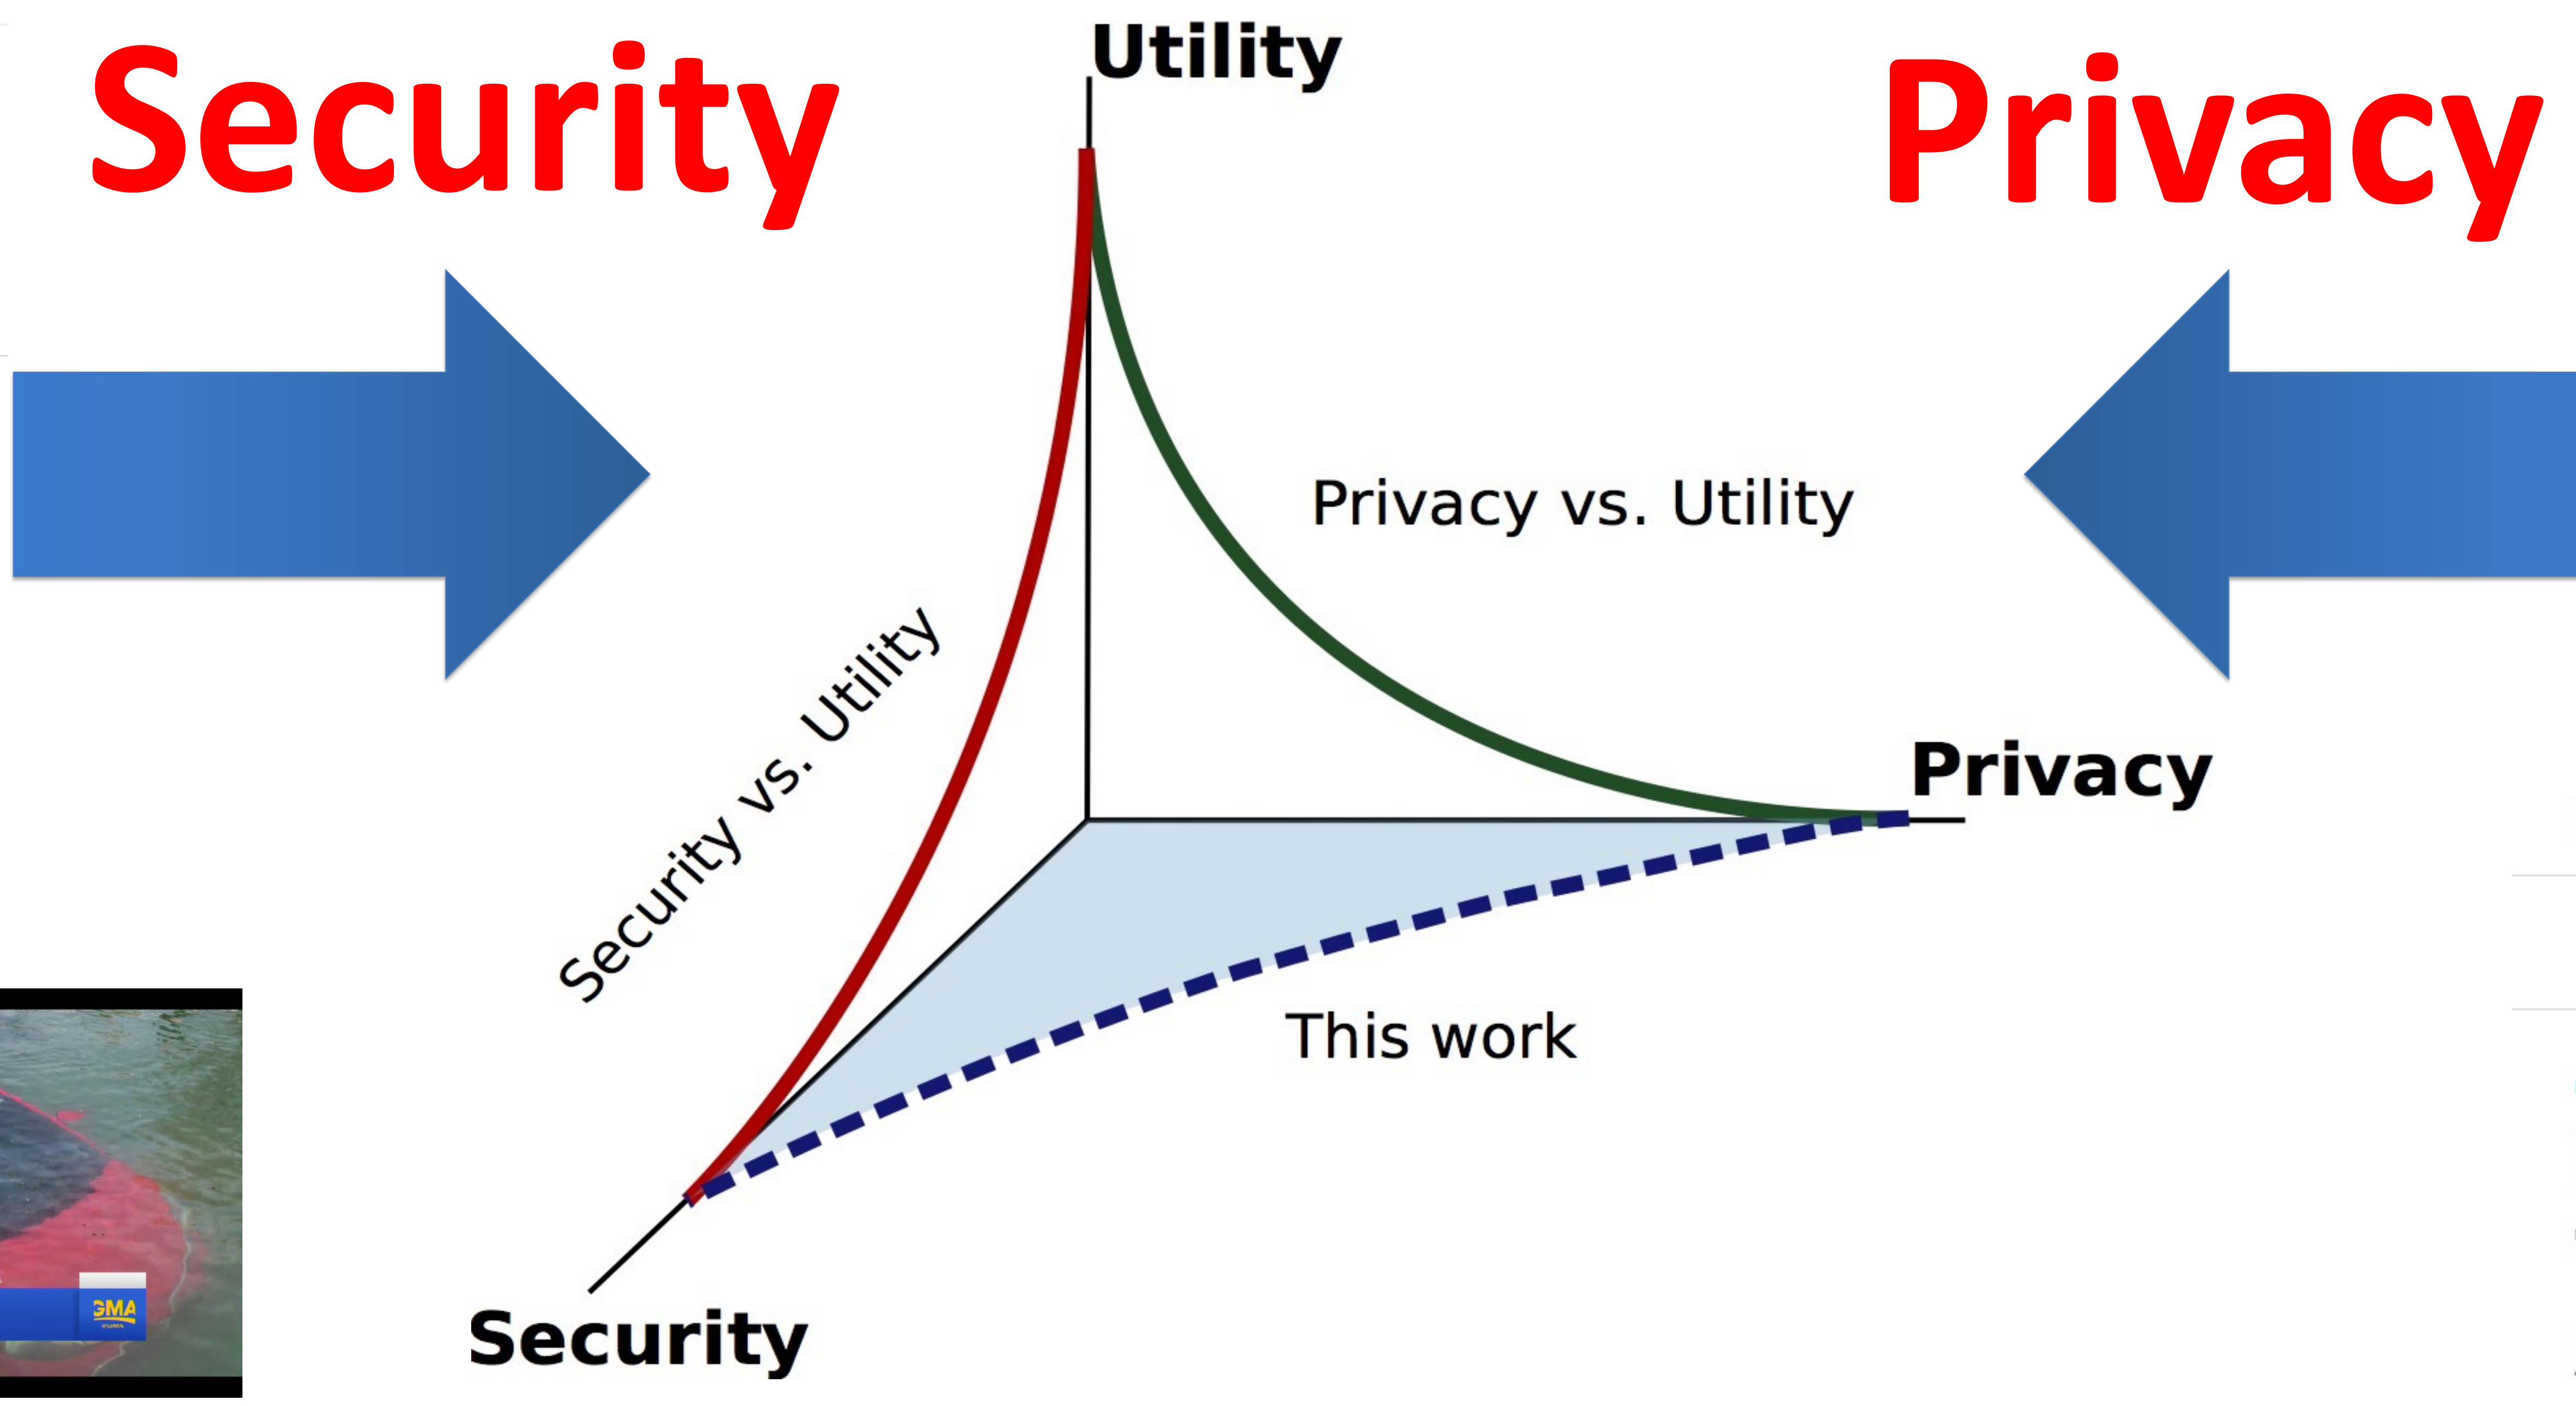

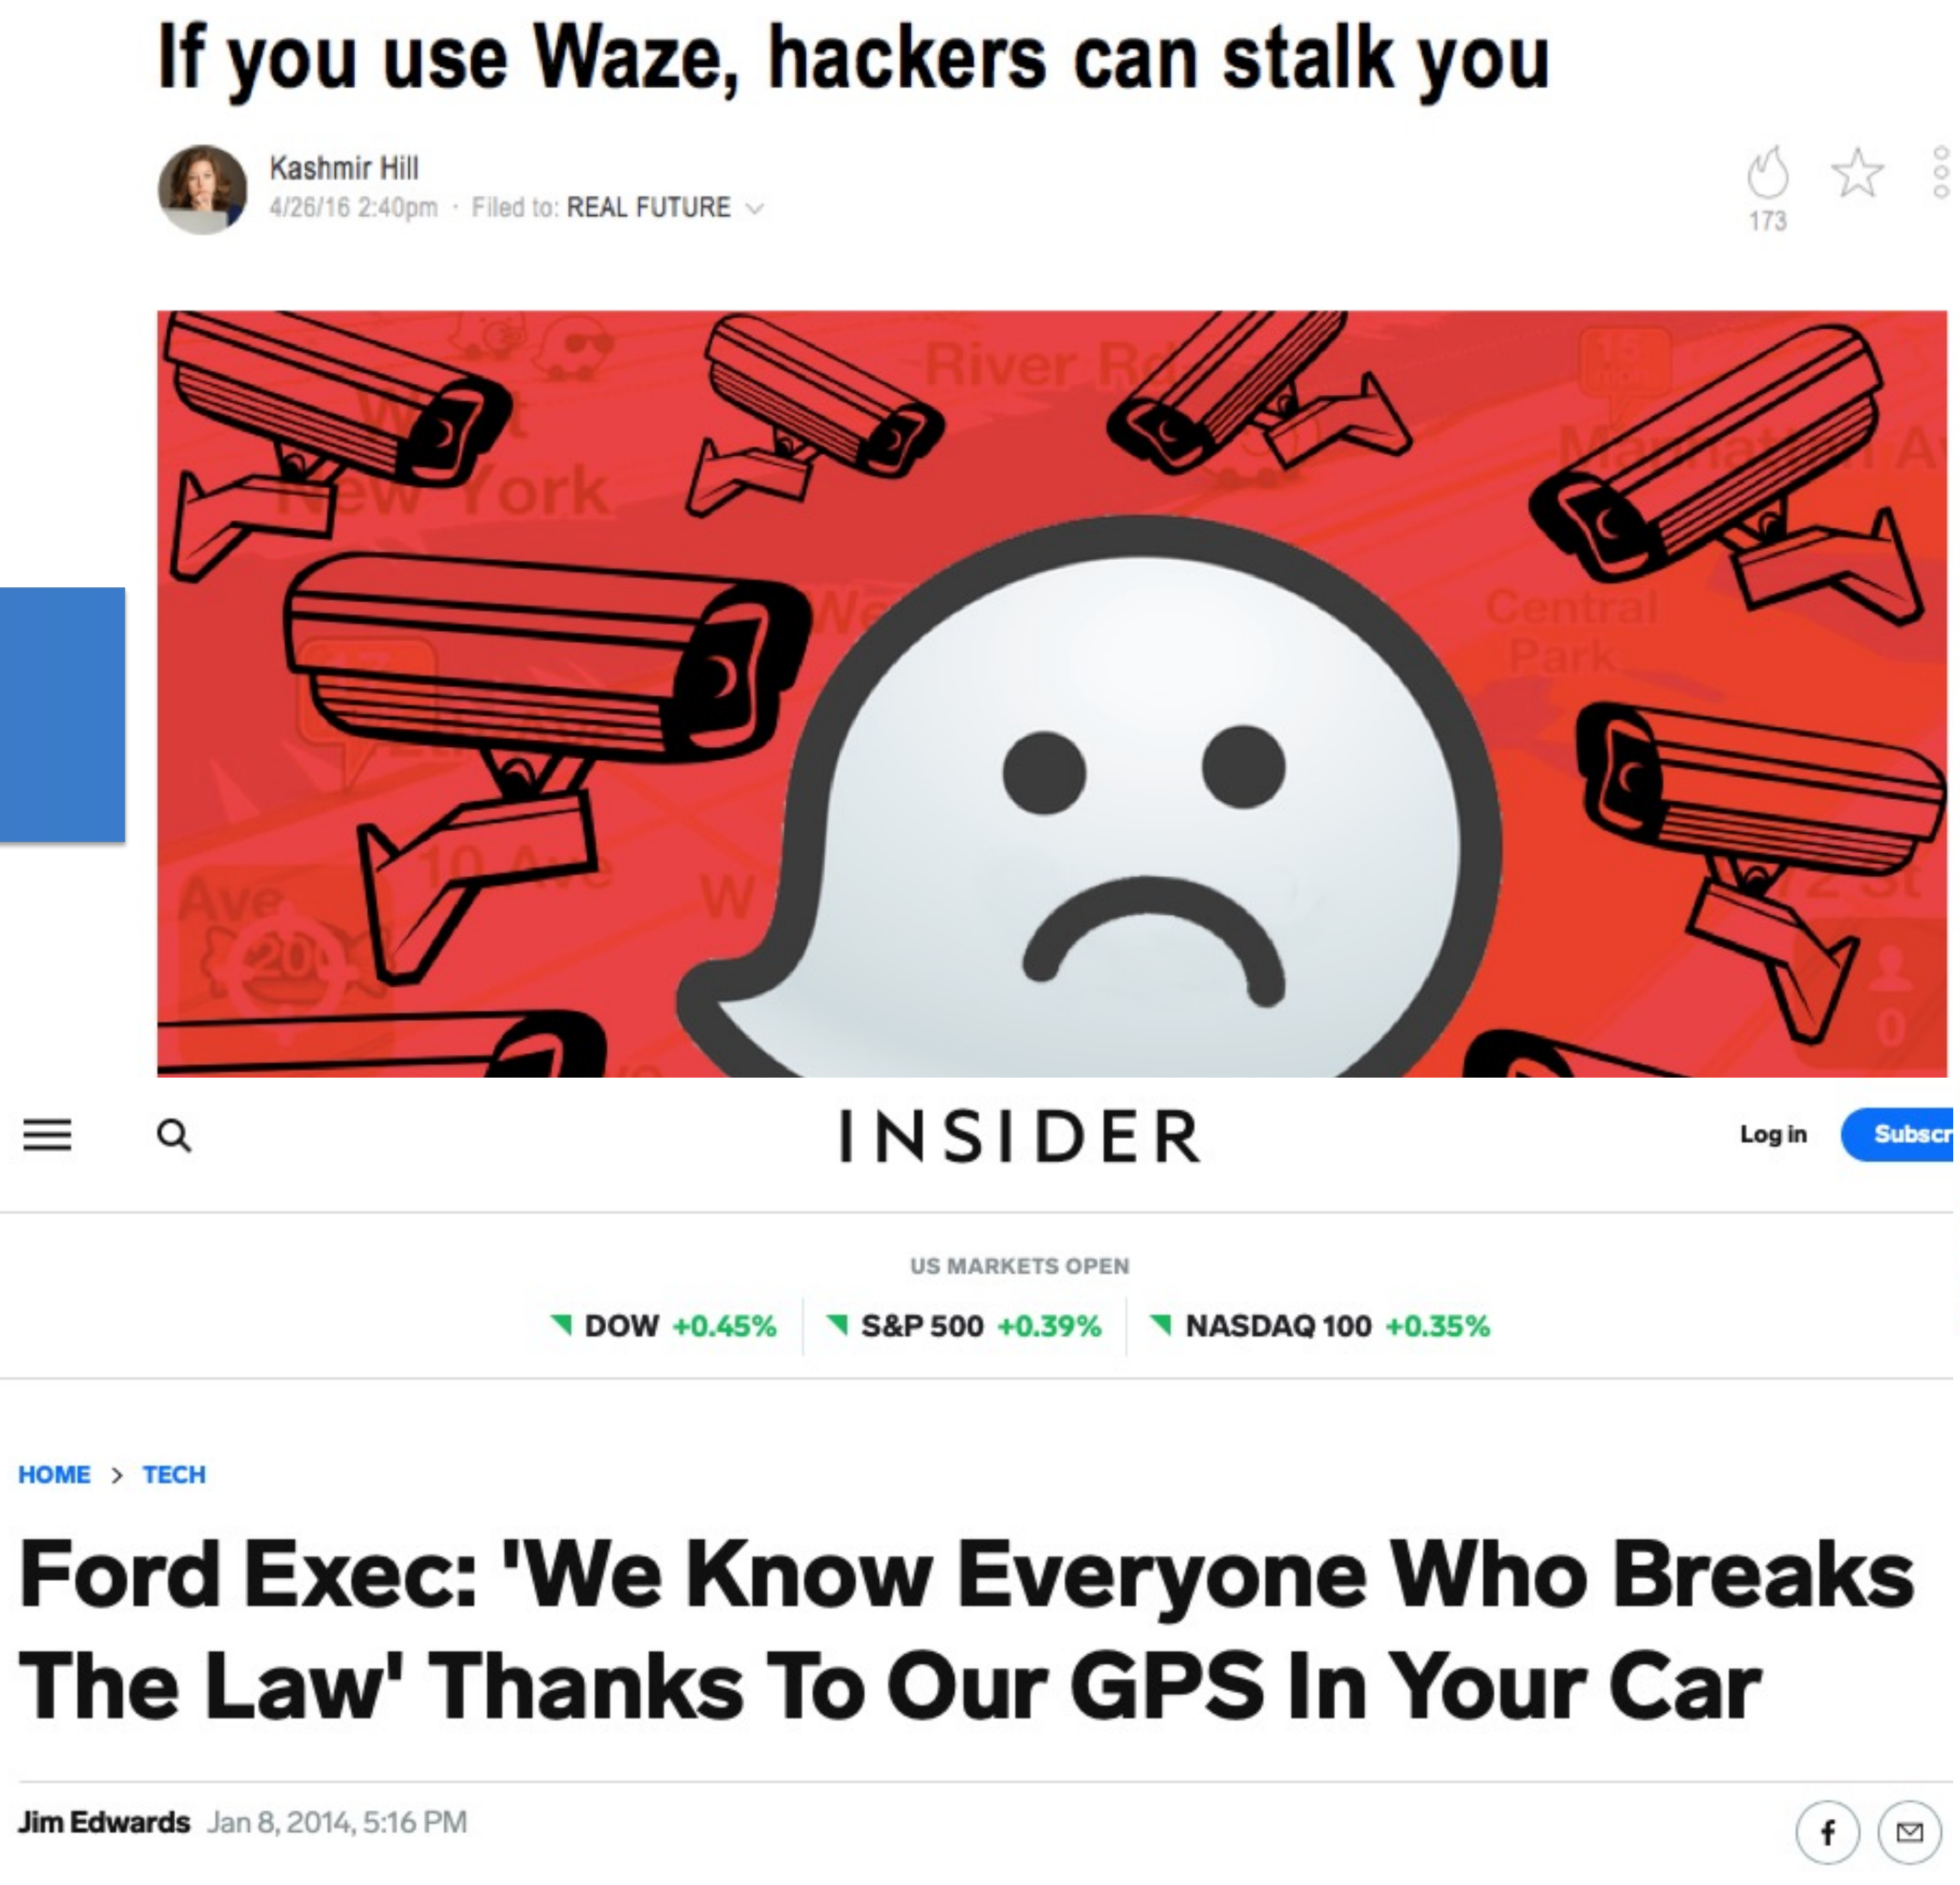

If you use Waze, hackers can stalk you

Ford Exec: 'We Know Everyone Who Breaks The Law' Thanks To Our GPS In Your Car

New Adversary Model:

- Consumer Data Protected by Differential Privacy
- Classical DP adversary is curious
  - Our adversary hides poisoning attacks in DP

Classical DP

$$\bar{\boldsymbol{Y}} \leftarrow \mathcal{M}(D)$$

$$\bar{\boldsymbol{Y}} \sim f_0$$

Attack

$$\boldsymbol{Y}^{\boldsymbol{a}} \text{ instead of } \bar{\boldsymbol{Y}}$$

Attack Goals:
Multi-criteria Optimization

$$\max_{f_a} E[\boldsymbol{Y}^a]$$

$$s.t.$$

$$D_{KL}(f_a \| f_0) \leq \gamma$$

$$f_a \in \mathcal{F}$$

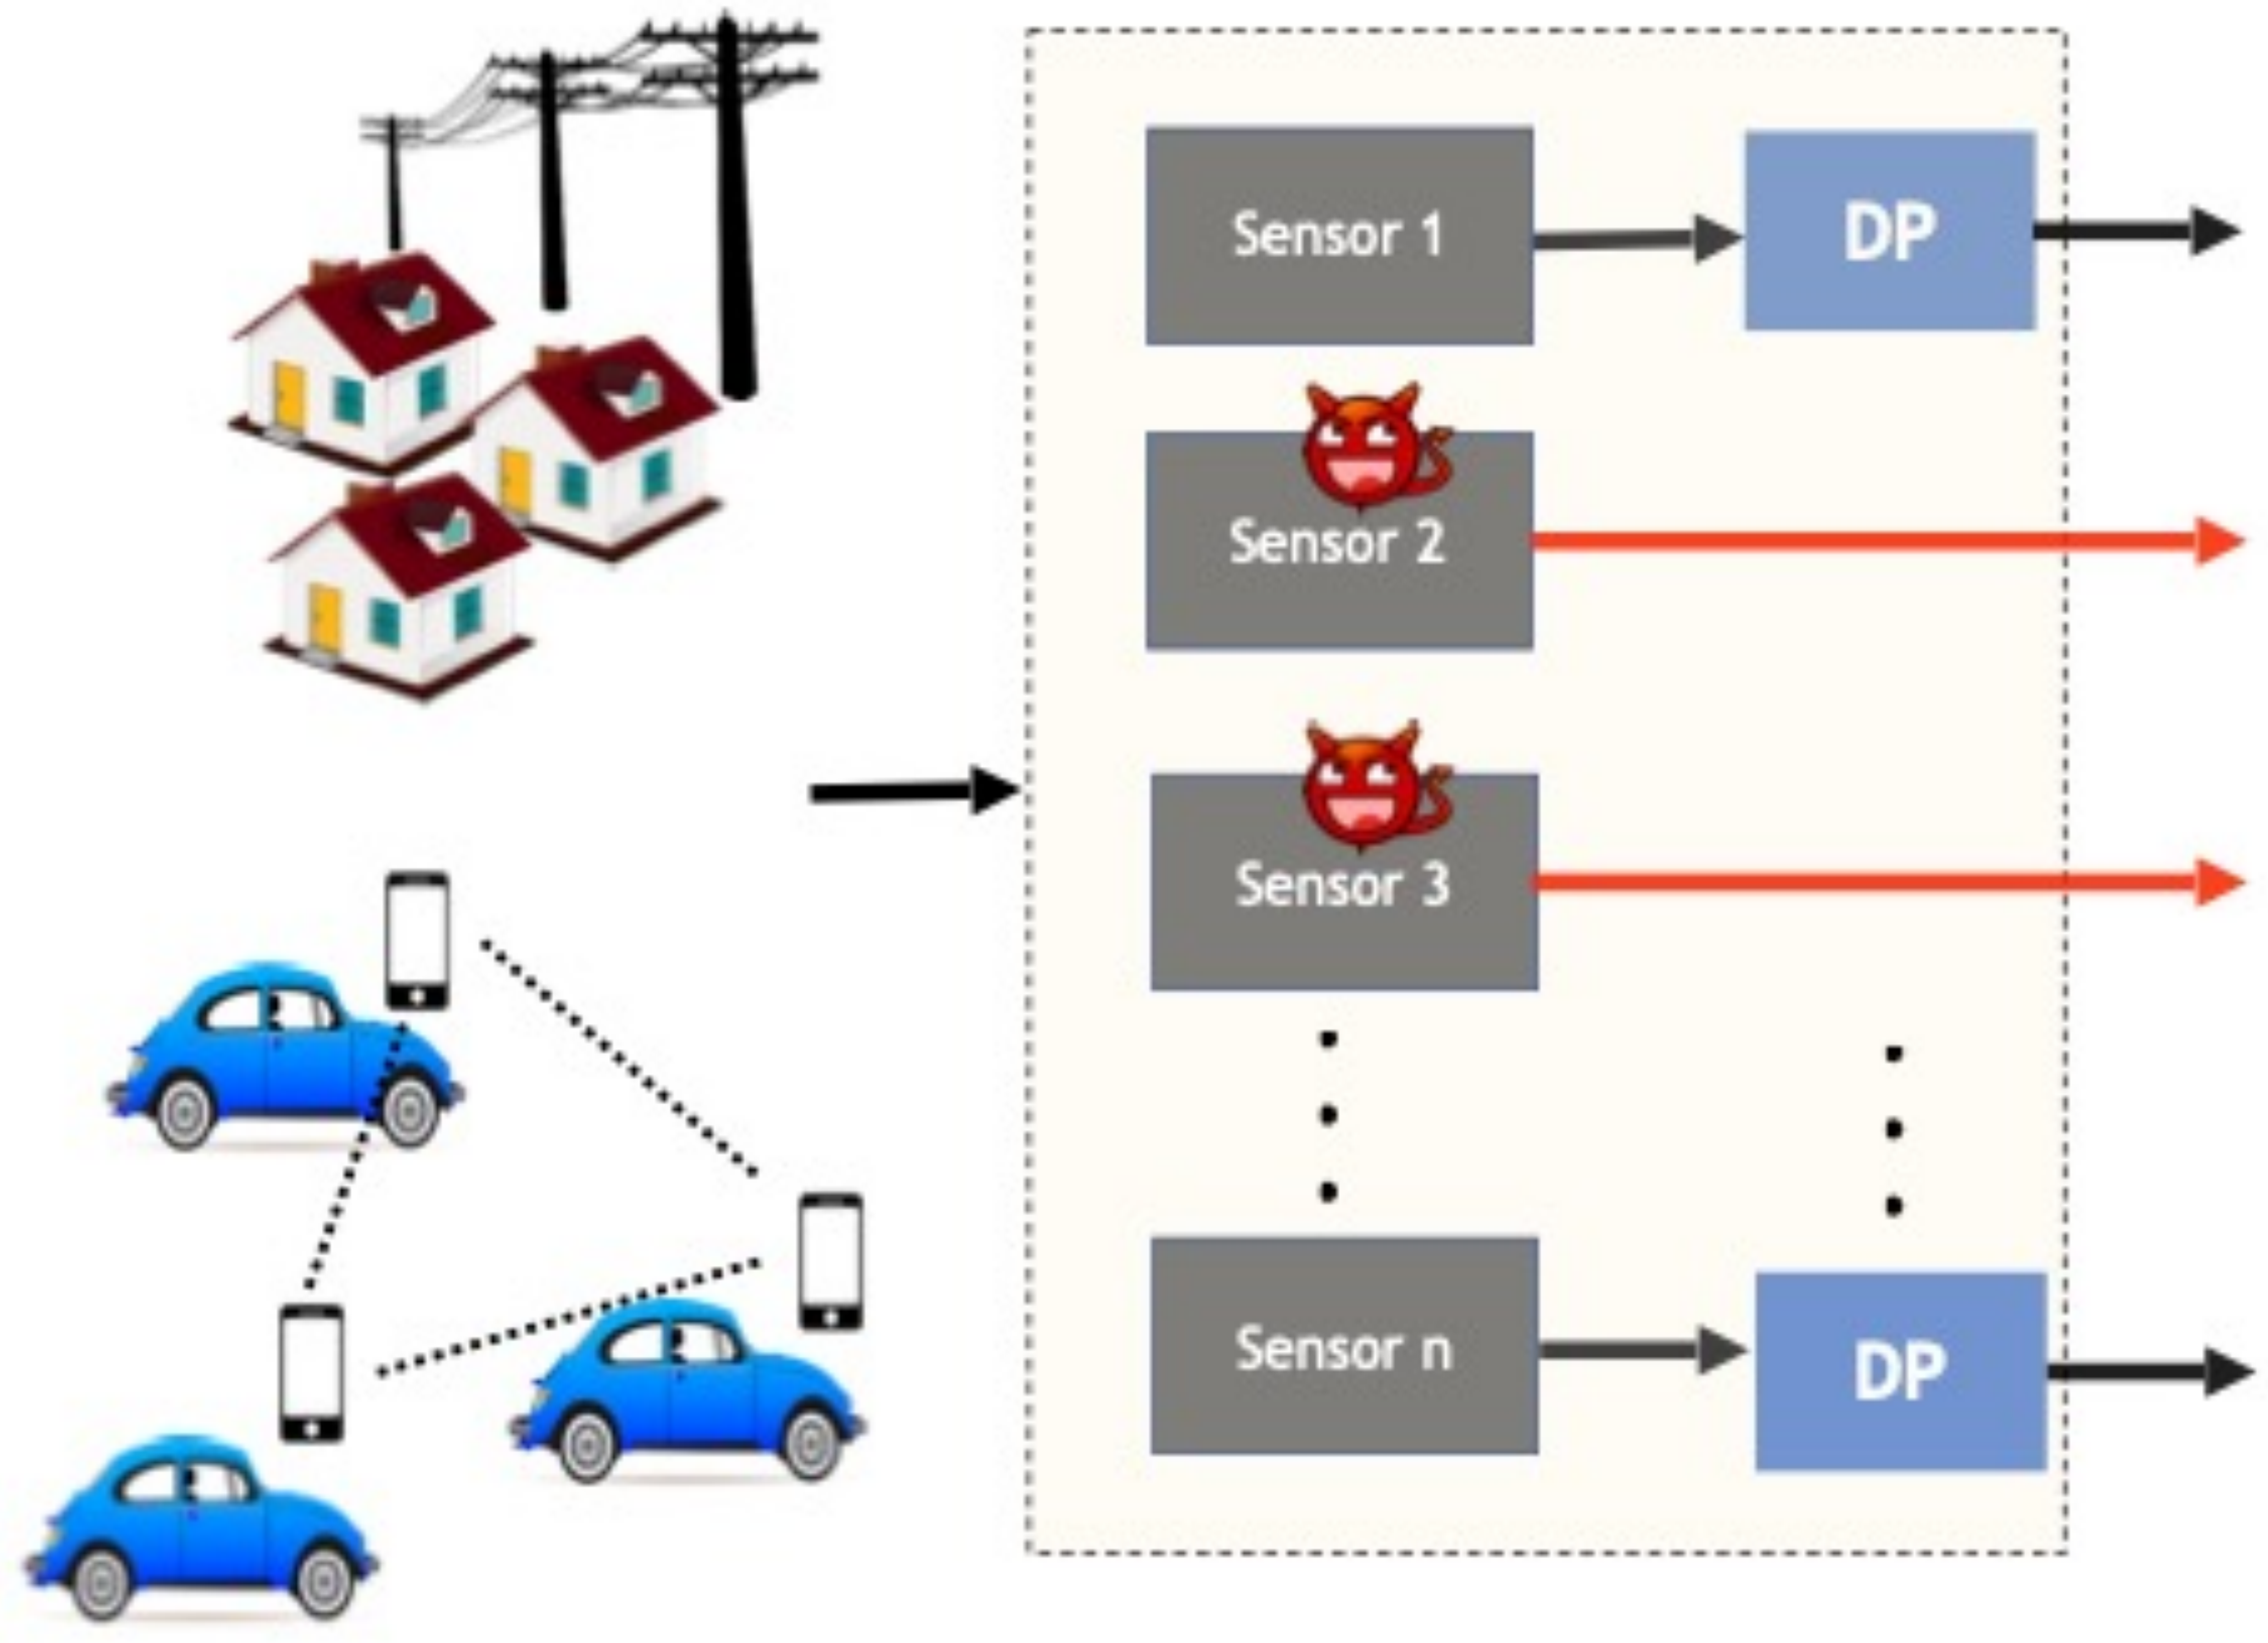

Optimal Attacks and Defenses:

- Variational methods are a useful tool to find the shape of functions

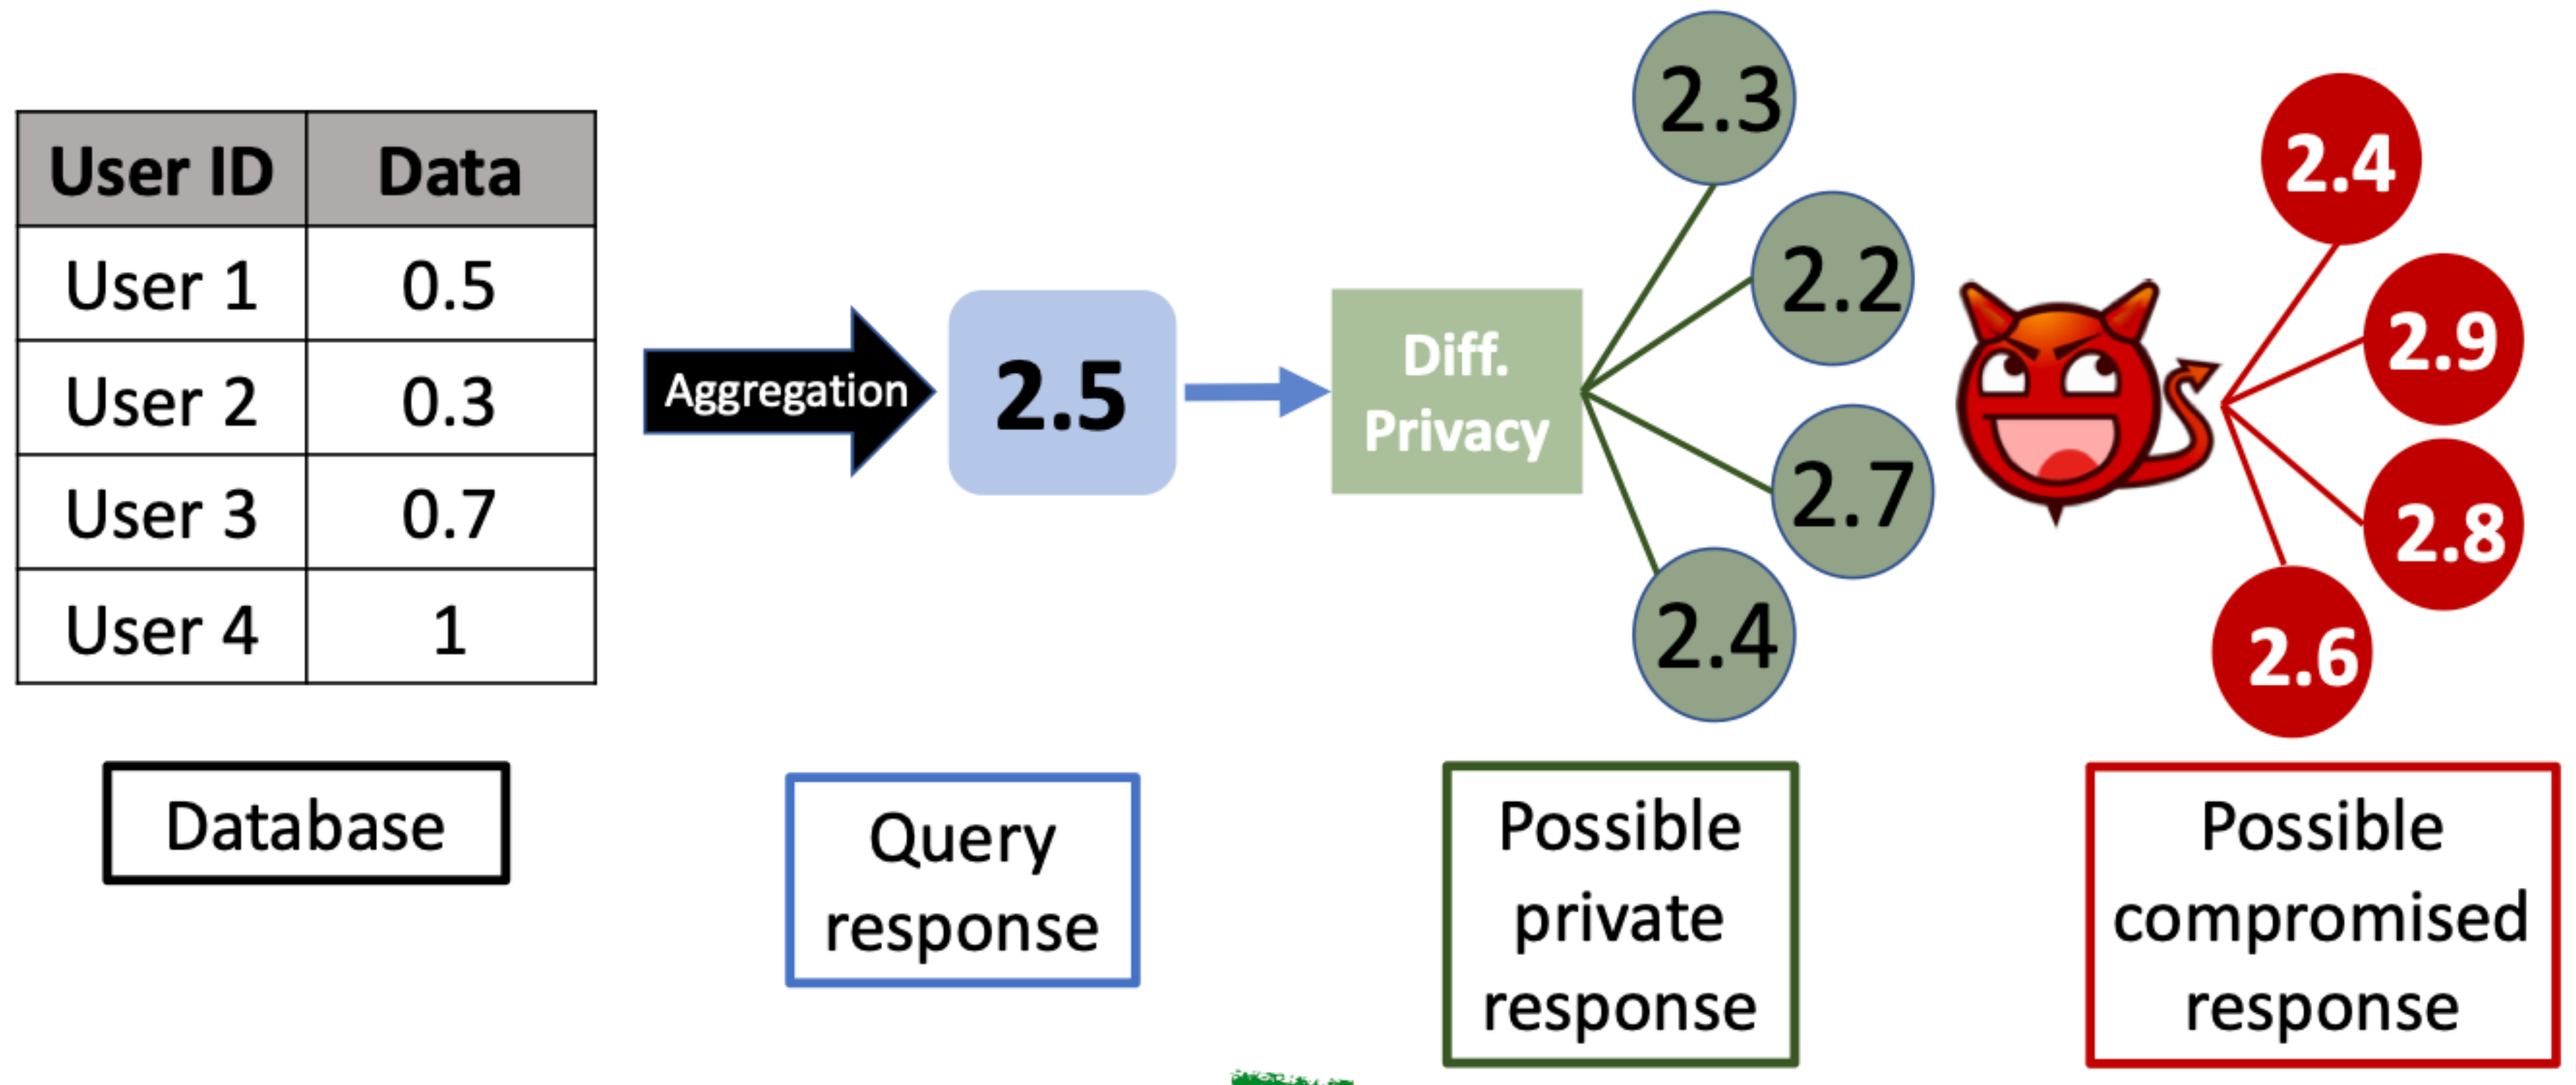

User ID	Data
User 1	0.5
User 2	0.3
User 3	0.7
User 4	1

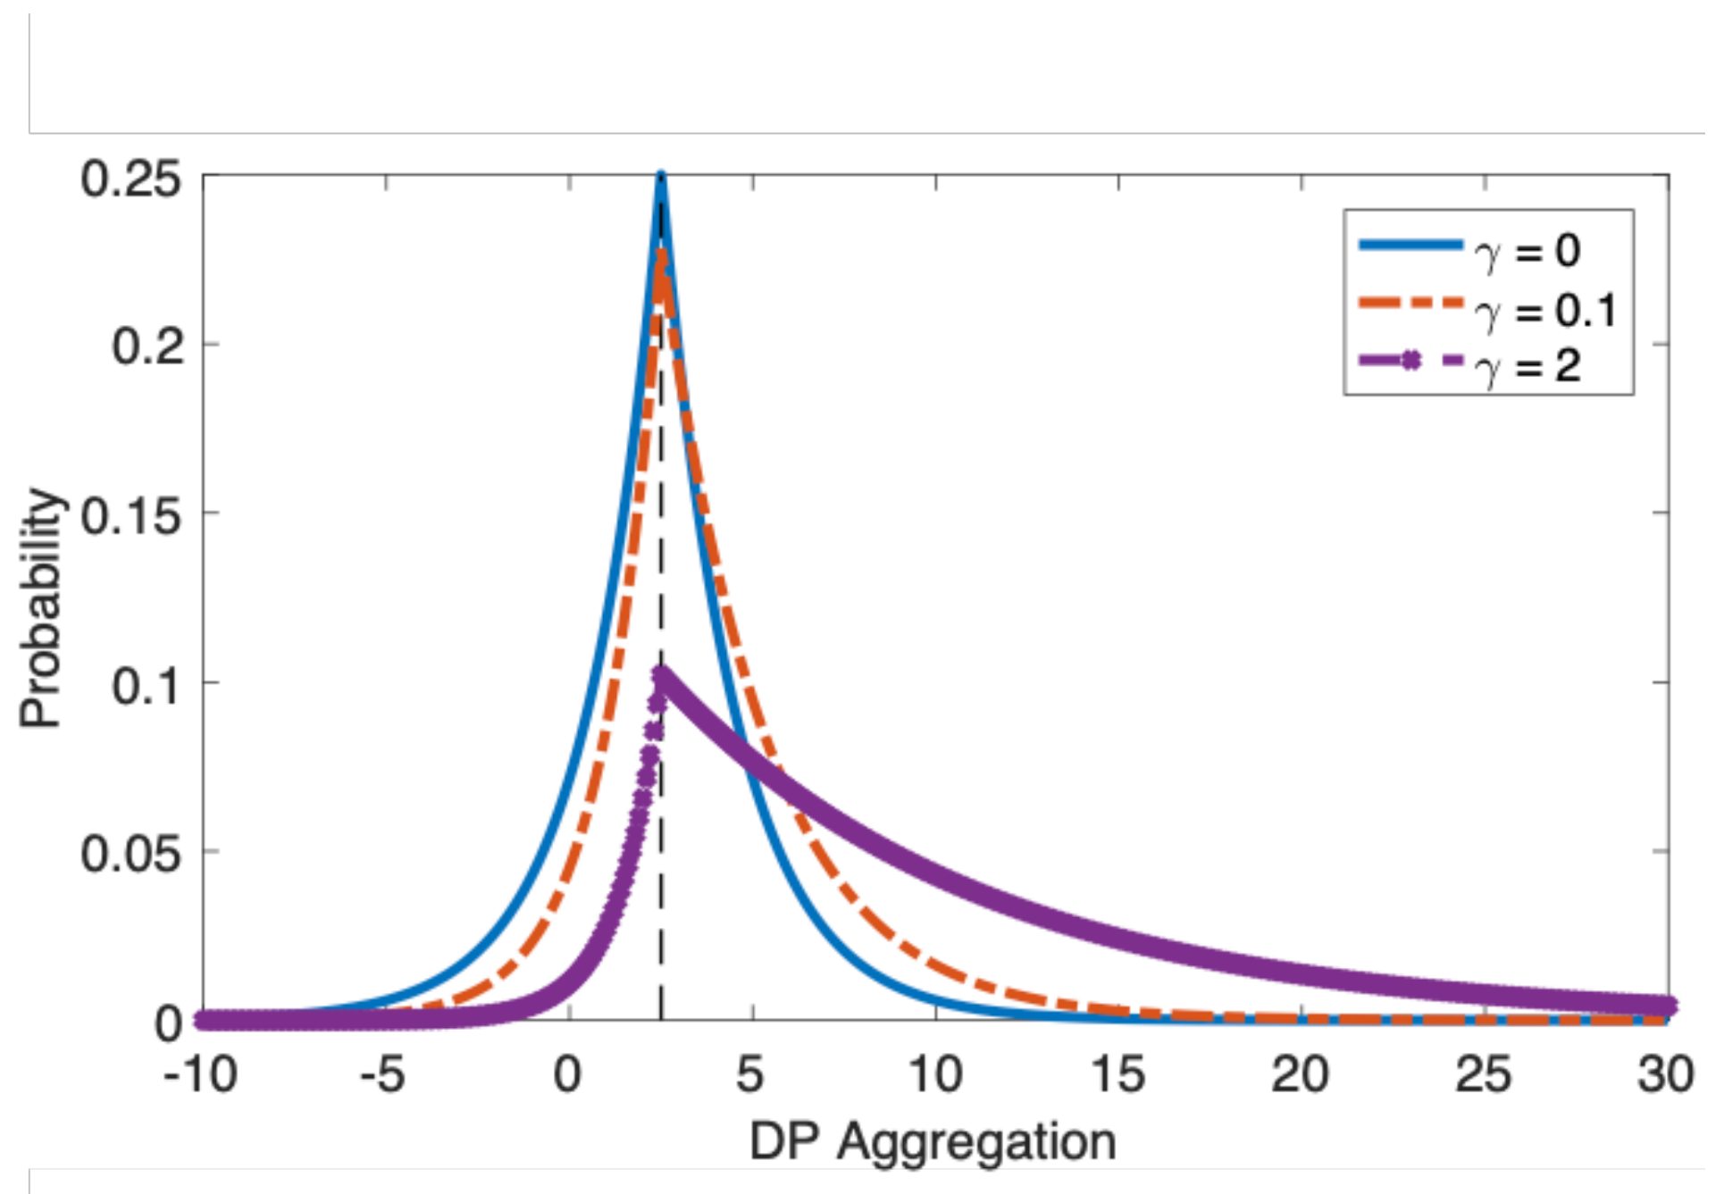

$$f_0(y) = \frac{1}{2b} e^{-|y-\theta|/b}$$

$$f_a^*(y) = \frac{\kappa_1^2 - b^2}{2b\kappa_1^2} e^{-\frac{|y-\theta|}{b} + \frac{(y-\theta)}{\kappa_1}}$$

κ_1 is the solution to $\frac{2b^2}{\kappa_1^2 - b^2} + \ln(1 - \frac{b^2}{\kappa_1^2}) = \gamma$

Traffic Estimation Example

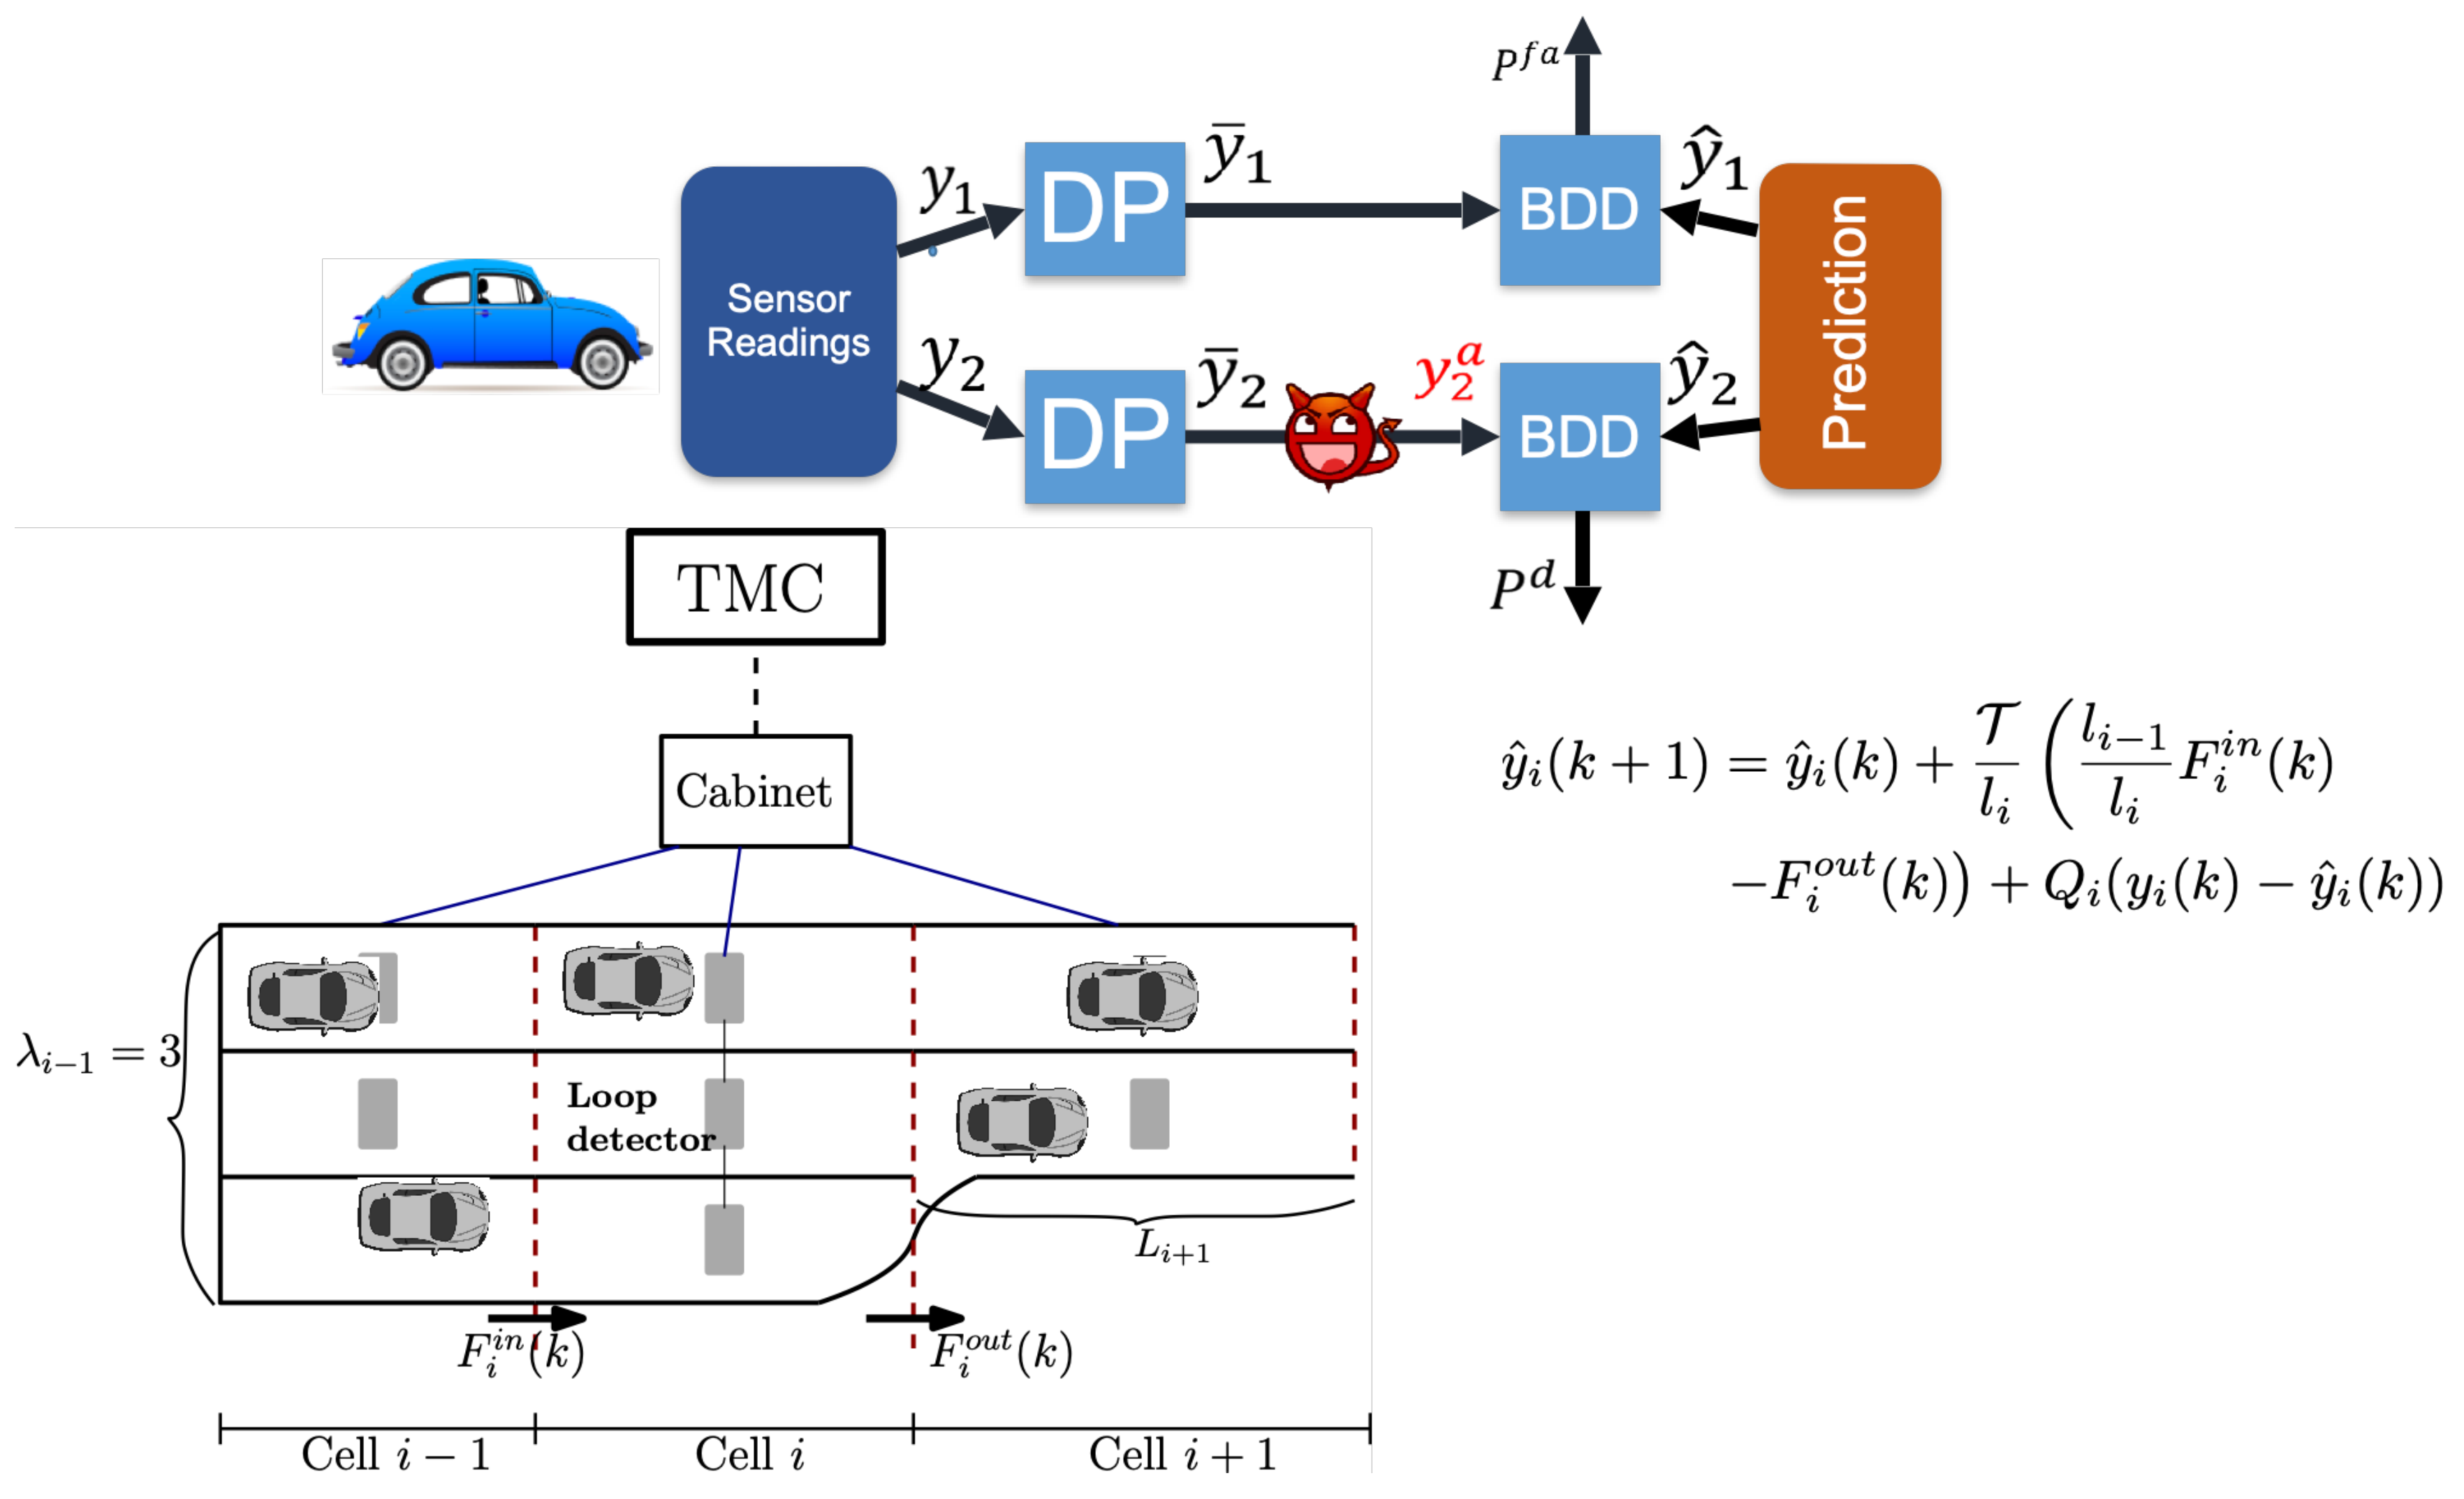

$$\hat{y}_i(k+1) = \hat{y}_i(k) + \frac{\mathcal{T}}{l_i}\left(\frac{l_{i-1}}{l_i}F_i^{in}(k) - F_i^{out}(k)\right) + Q_i(y_i(k) - \hat{y}_i(k))$$

Without DP the attack is limited

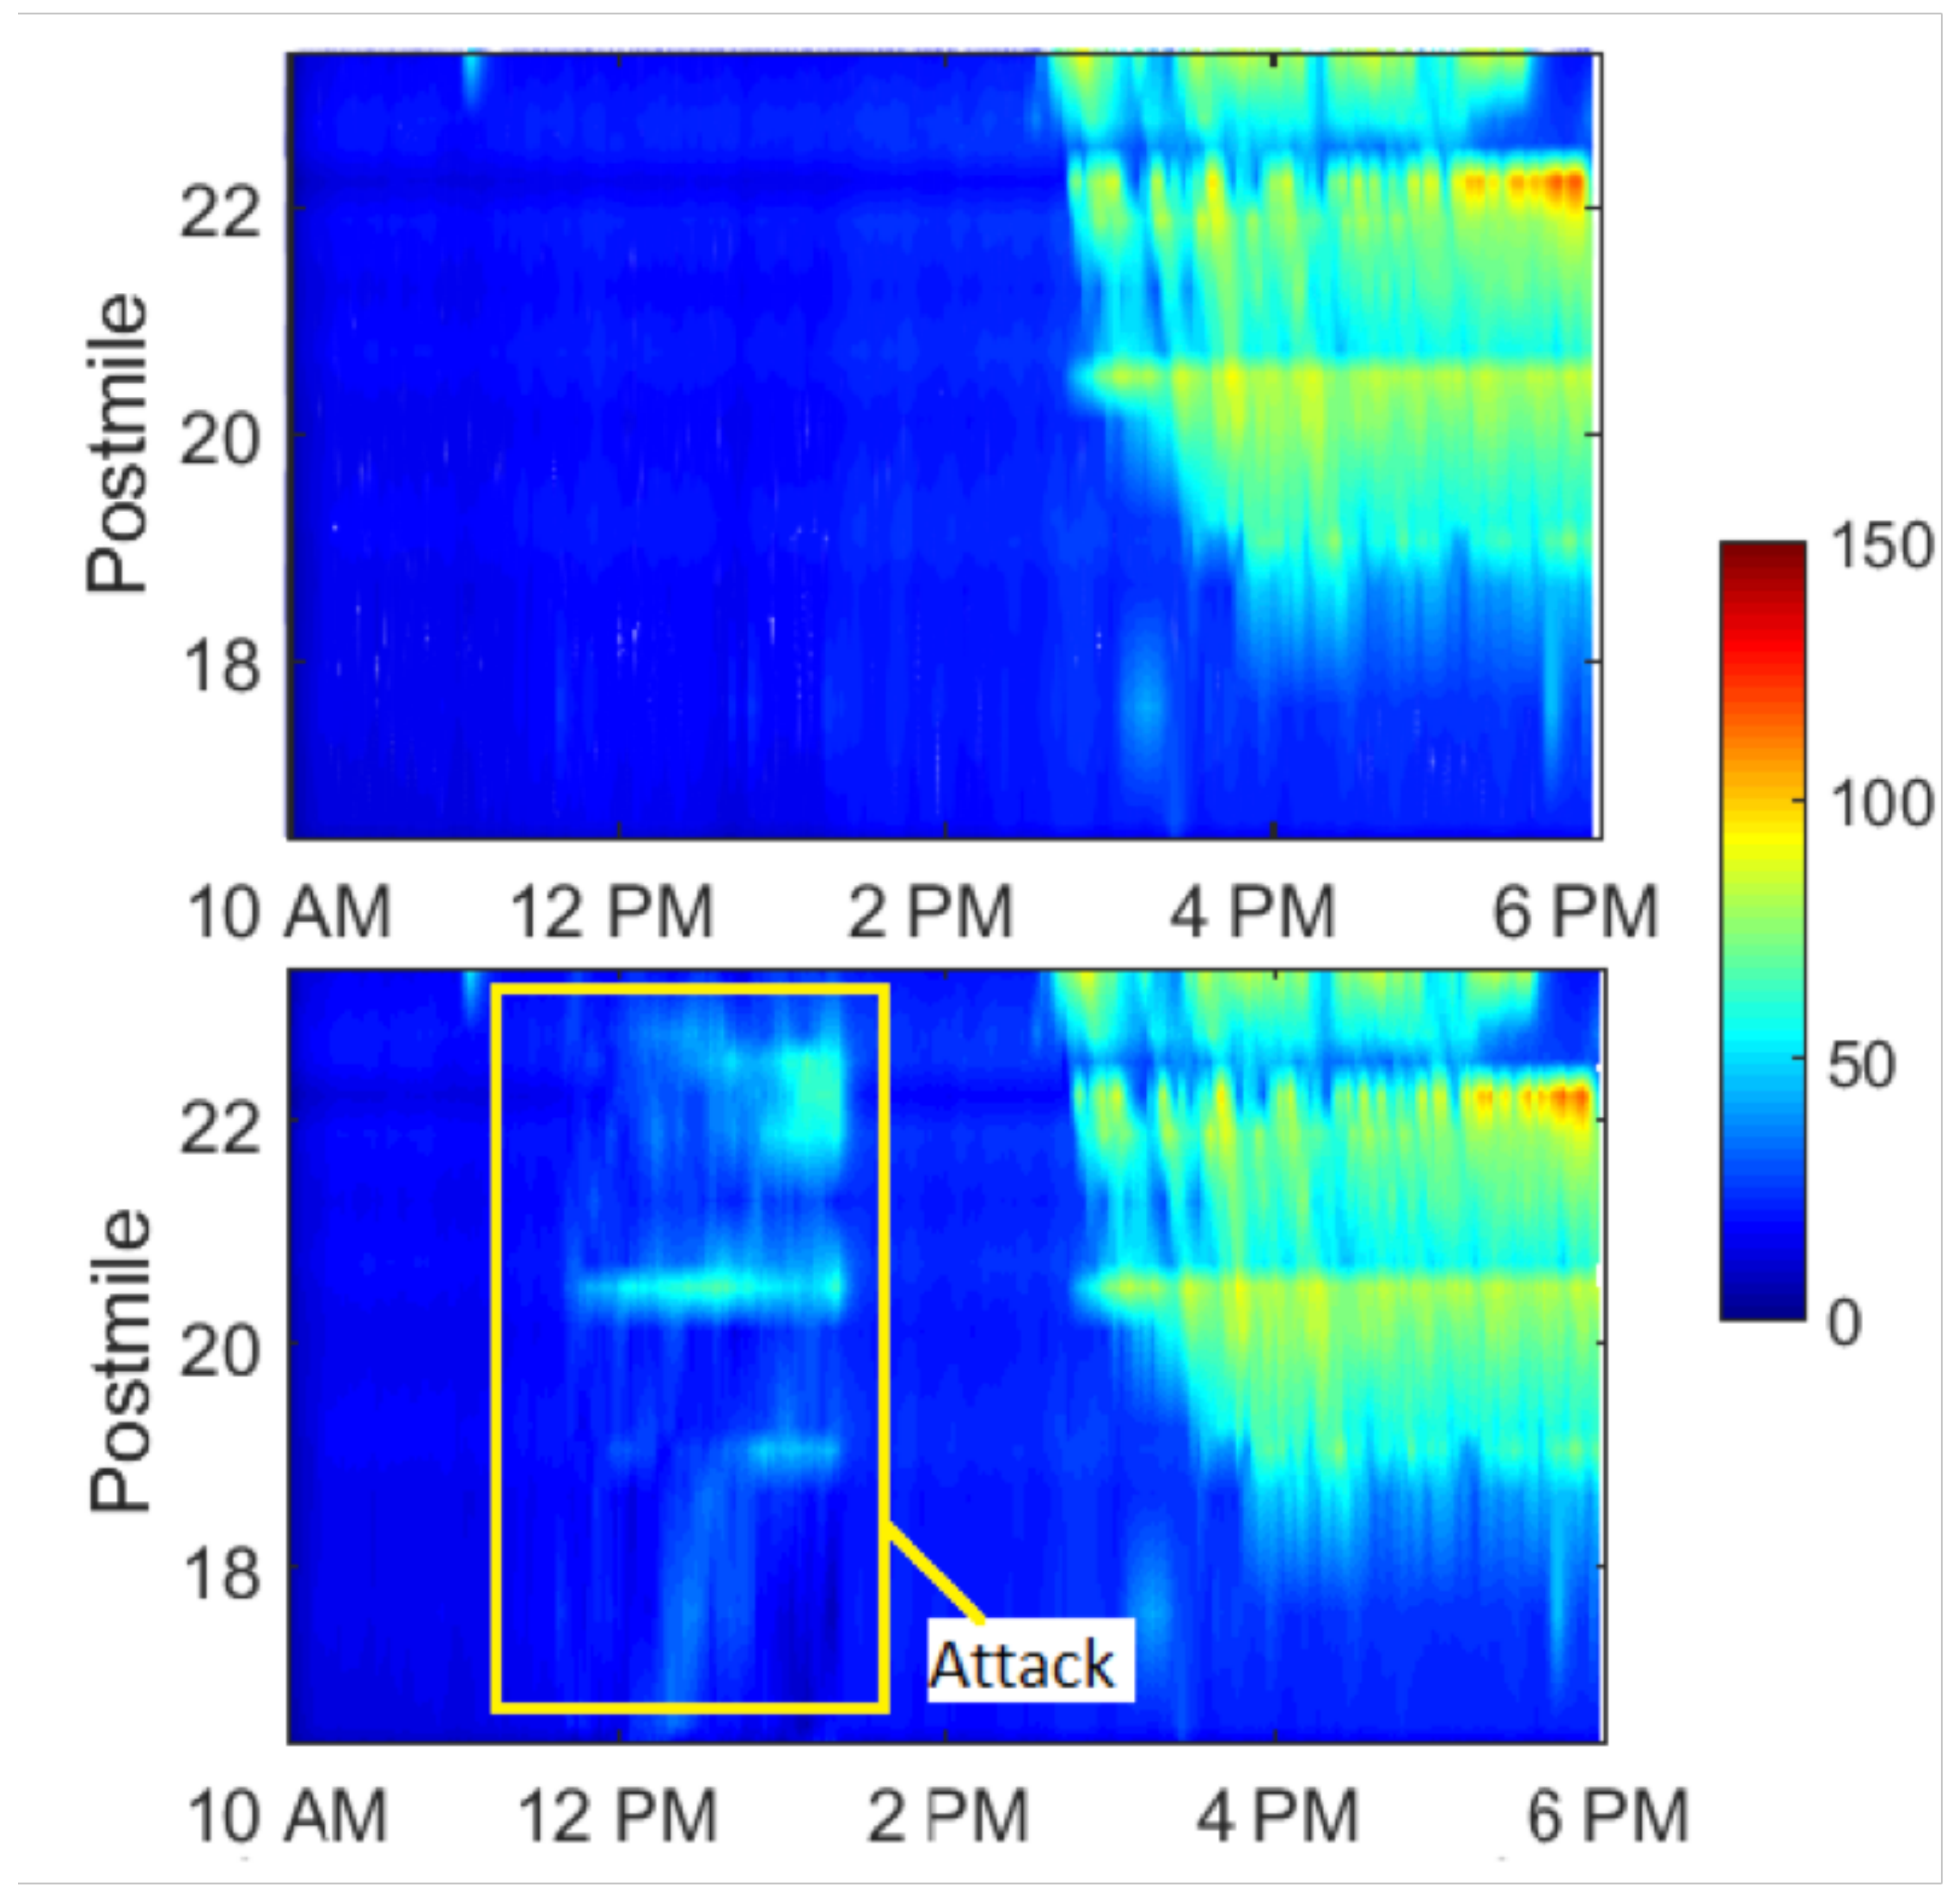

With DP, the attacker can lie more without detection

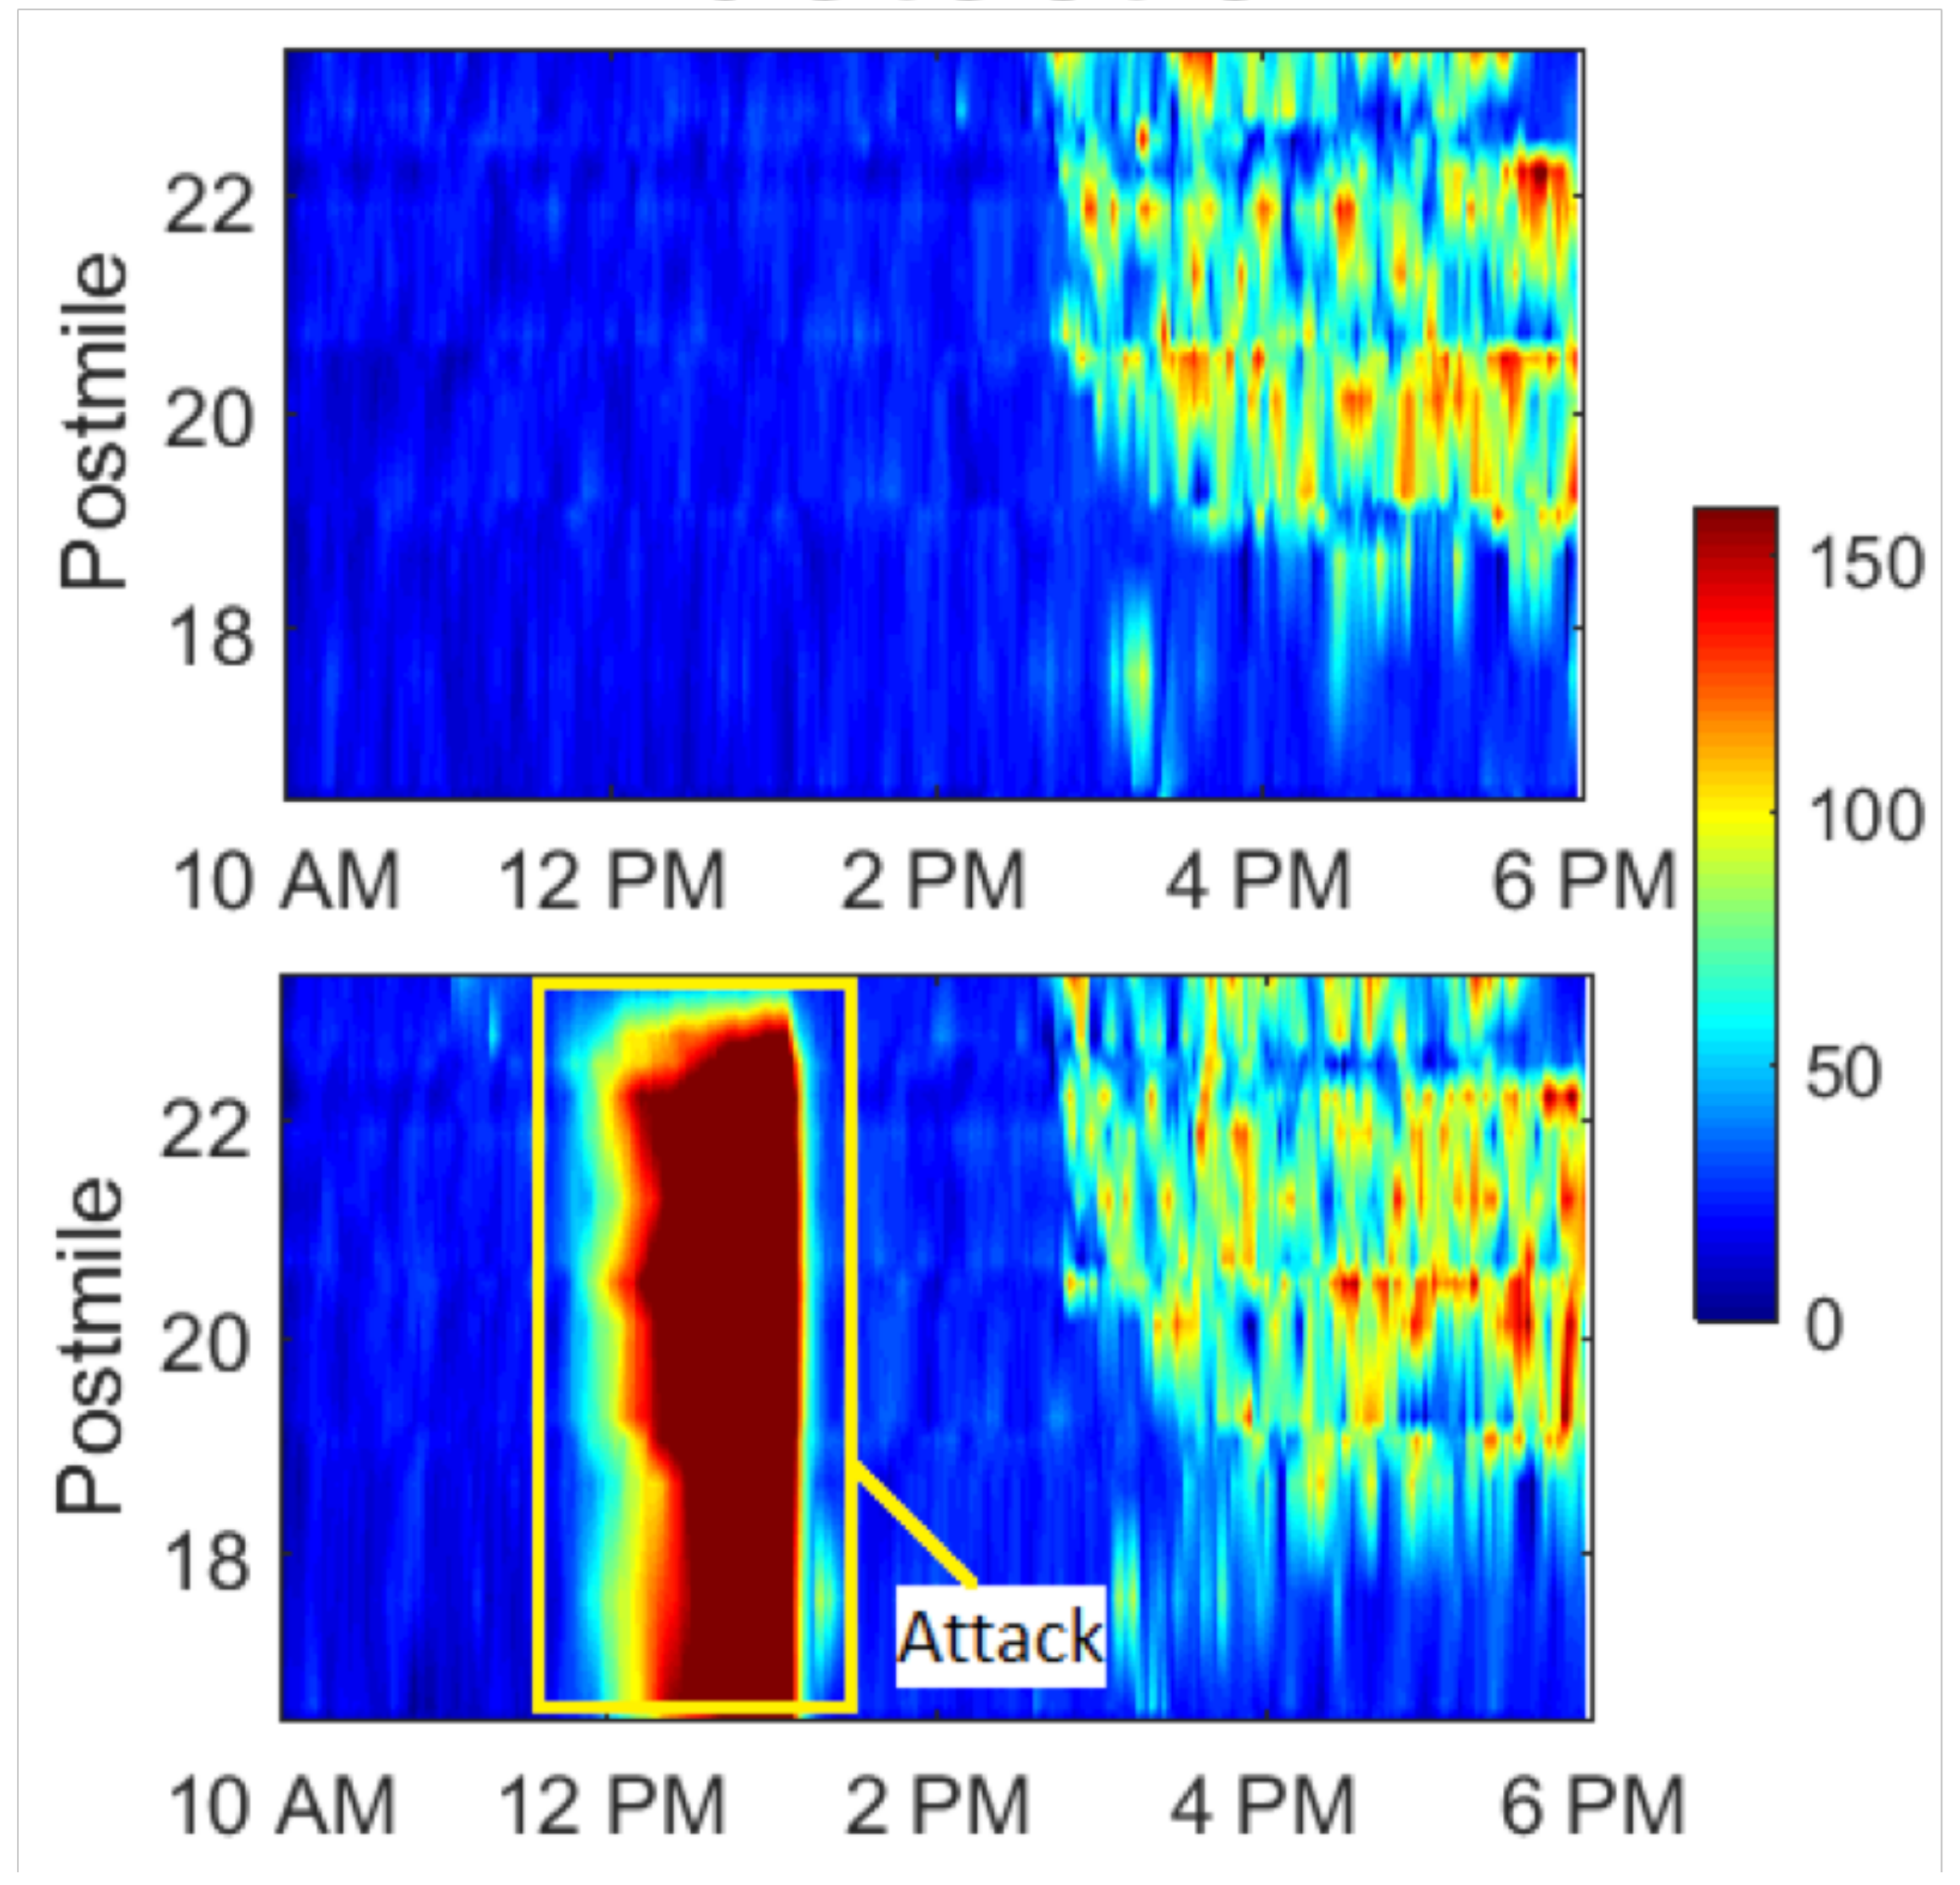

Optimal Defense:

- With classical defense

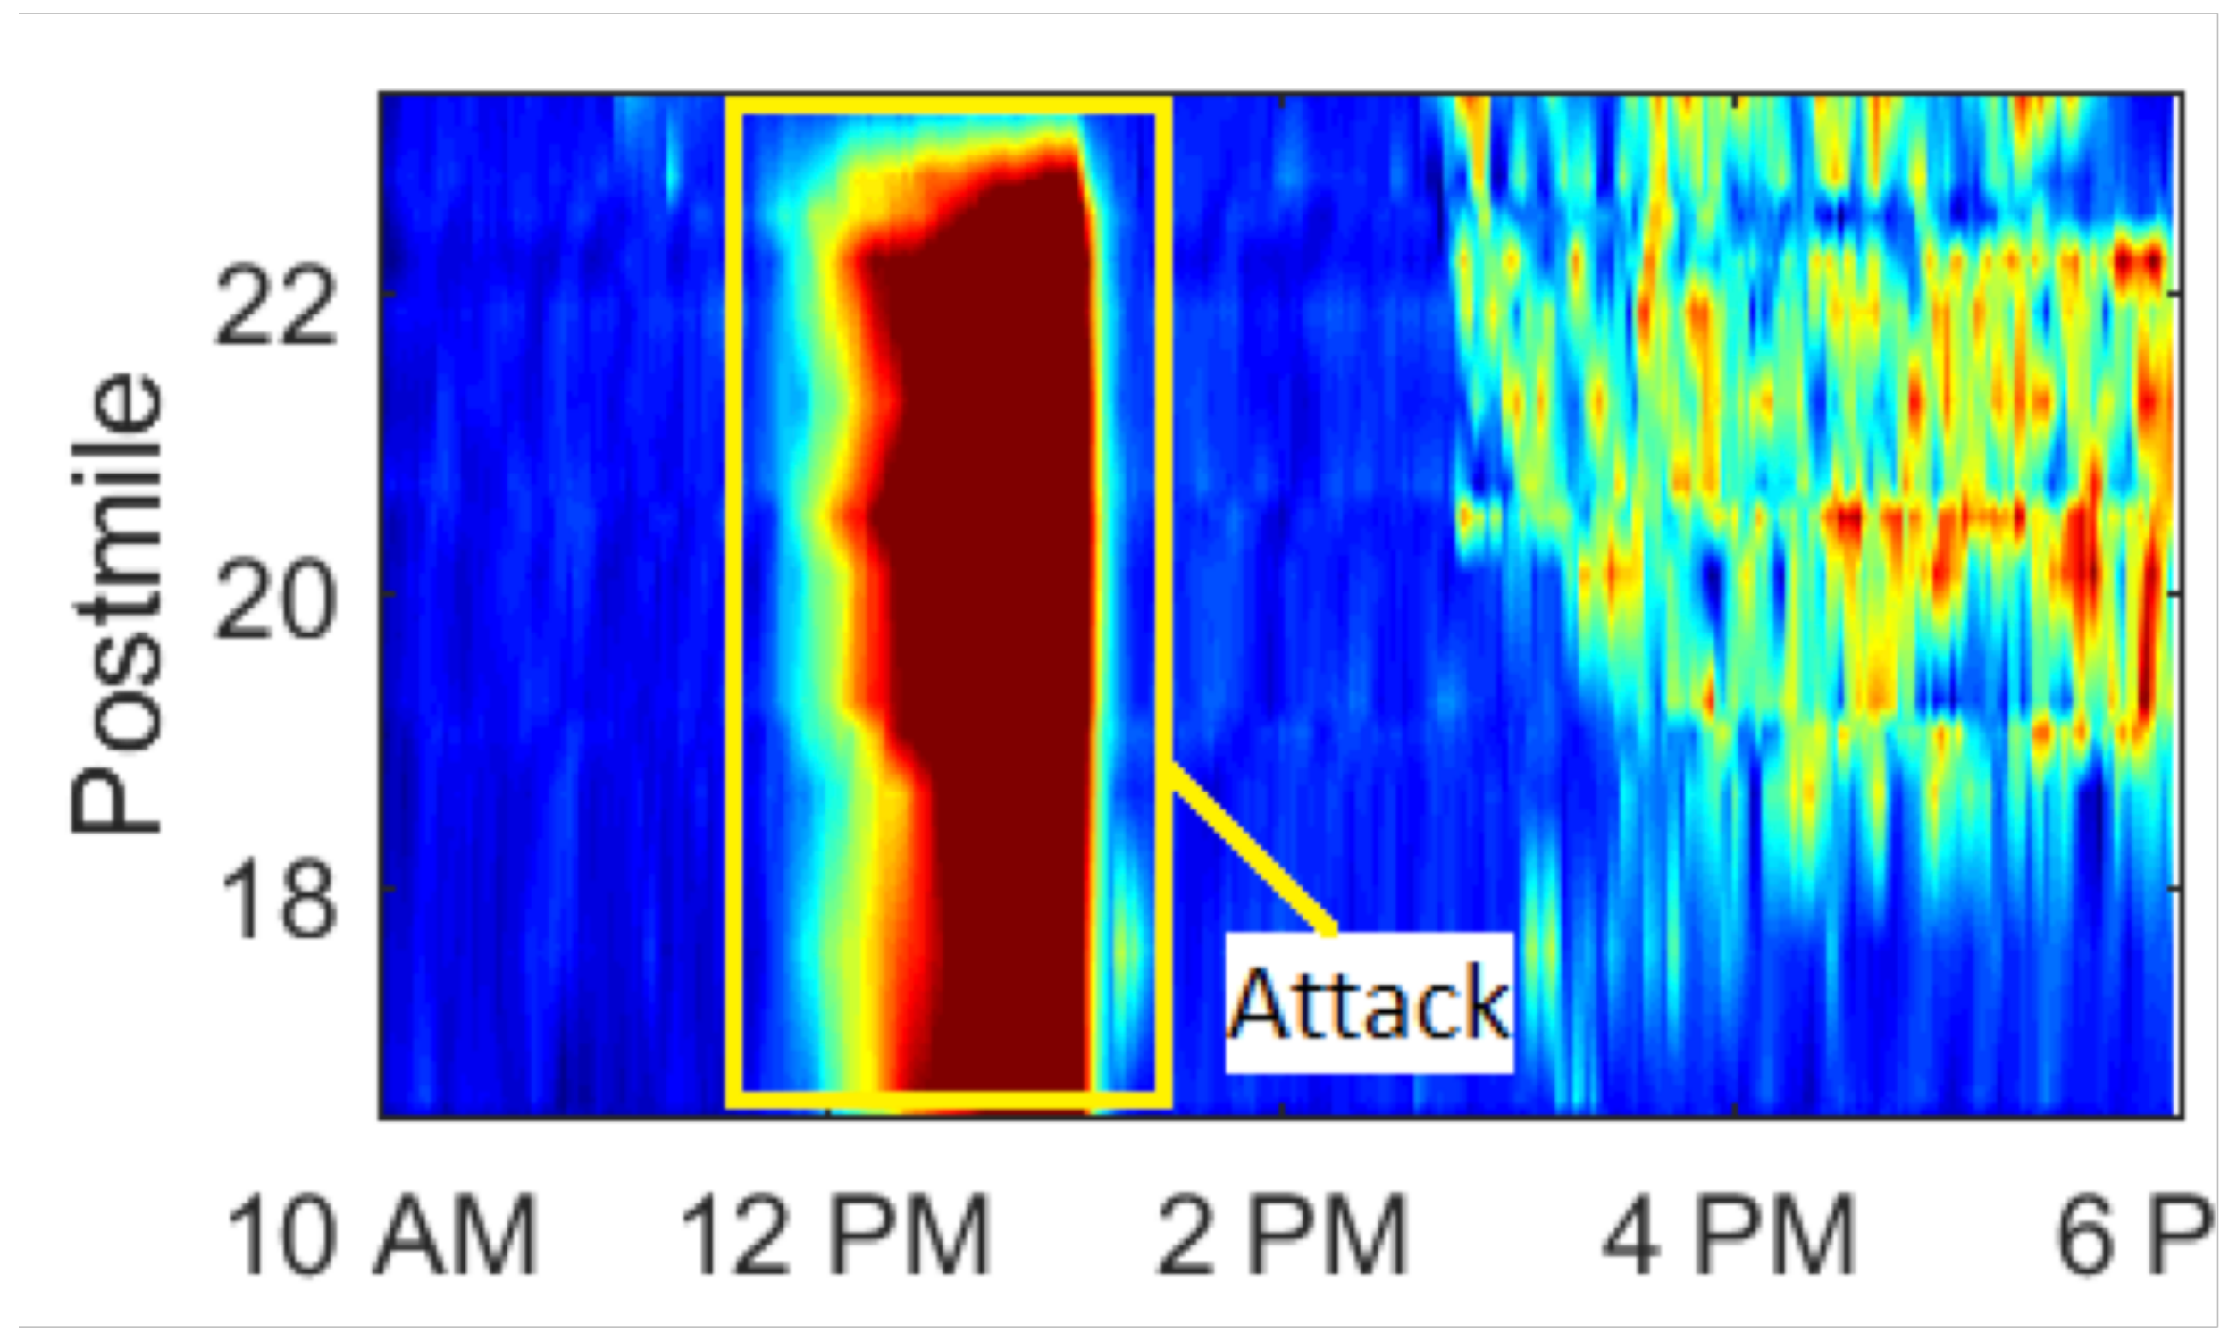

- With our defense

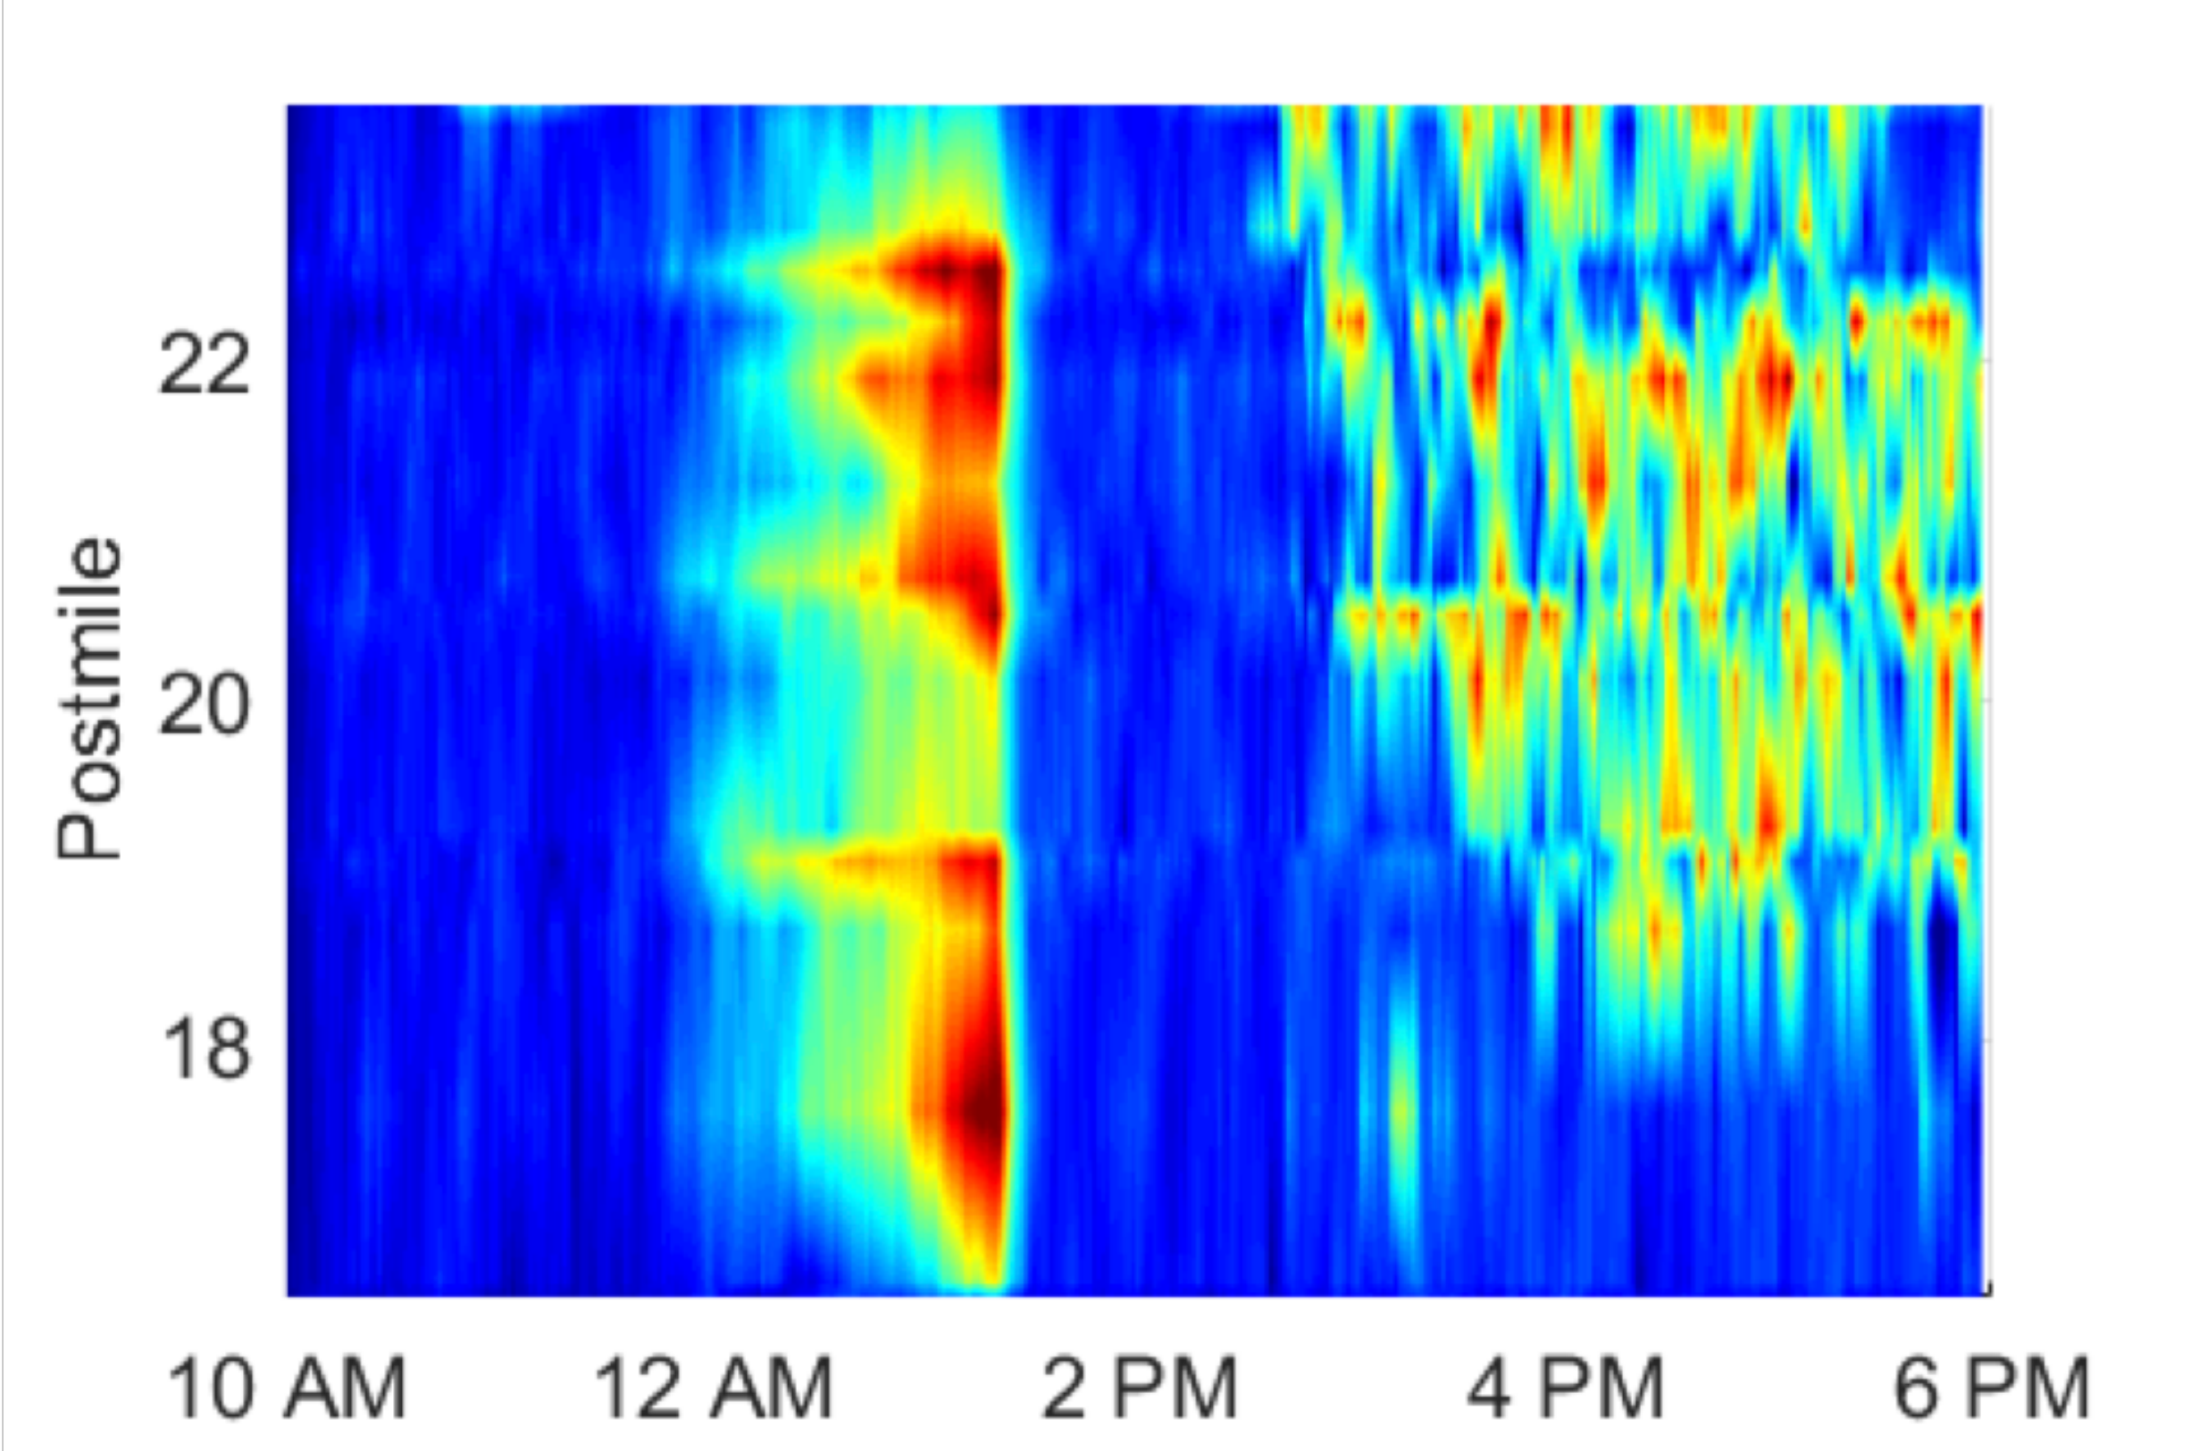

Recent Publications:

- Giraldo, Cardenas, Kantarcioglu, Katz. Adversarial Classification Under Differential Privacy. **NDSS 2020**
- Ozdayi, Kantarcioglu, Gel. Defending Against Backdoors in Federated Learning with Robust Learning Rate. **AAAI 2021**

Ongoing Work: Secure computation for attack-detection in control systems