

# Re[DP]: Realistic Data Mining Under Differential Privacy

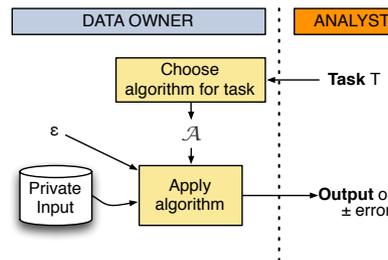
## Challenge:

- Differential privacy: a gold standard for analysis of sensitive data
- Practical adoption by industry/government is rare
- How to close gap between theory and practice?

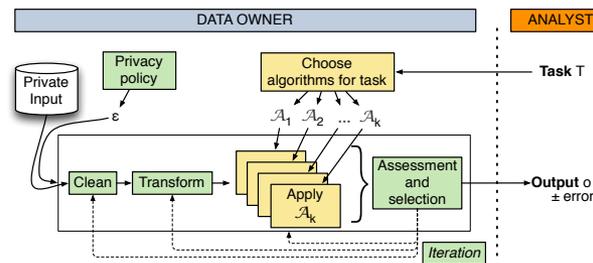
## Solution:

- Focus on under-studied aspects of data mining workflows
- Develop guidelines for privacy policies and standards for utility
- Design robust benchmarks for principled evaluation

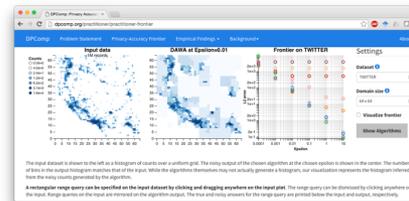
### World according to DP research



### Real World



### DPComp



## Scientific Impact:

- End-to-end privacy for realistic data mining workflows
- Sound and reproducible methodology for evaluation
- Produce error-optimal algorithms for data mining tasks including
  - classification/regression,
  - graph analysis, and
  - mining mobility traces

## Broader Impact:

- Expand practical usefulness of privacy algorithms and ultimately increase data sharing
- DPComp.org: online resource for practitioners and researchers to assess state-of-the-art of differentially private algorithms
- Real-world deployment: collaboration with US Census
- Education and outreach
  - Tutorial at VLDB 2016: Differential Privacy in the Wild
  - Grad and undergraduate courses
  - REU