Coastal Virginia

Explore Privacy-Preserving in Deep Image Retrieval Systems

PI Cong Wang, Old Dominion University
PI Qun Li, College of William and Mary
PI Janet-Walters Williams, Hampton University

Coastal Virginia

- The goal of the project is to enhance the privacy and security of large-scale image retrieval systems.

- The project is structured as 1): protecting the privacy of users from malicious image retrieval and 2): defending against malicious image queries.

Coastal Virginia

- Initial focus area and assets
  Background: AI Applications are everywhere – online shopping (use image to search image, e.g., finding a similar online product without knowing its name.)
- Online database maintains a large-scale collection of images to return as the search results (might include private images). E.g., Google Image search maintains 7-day storage of user's image queries.
- Third parties can retrieve those "private" images with similarity search.
- Research problem: How to prevent the hackers/third parties from retrieving private images from the database?

Coastal Virginia

- We are working on the research Goal 1: prevent the private images in the database from being matched (returned) to the attacker using adversarial examples (by injecting random perturbations into user's image to spoof the AI algorithm).
- We have re-designed the optimization goal of the existing method to drive private images into a subspace away from all the categories.
- We have tested the performance under the white-box scenario (all the parameters about the image search engine is known).
- Now we are focusing on the second part: Black-Box Transferability.
- We do not know the model and its parameters.