

---

---

# City Forensics

By Sean M. Arietta, Alexei A. Efros, Ravi Ramamoorthi,  
and Maneesh Agrawala

---

---

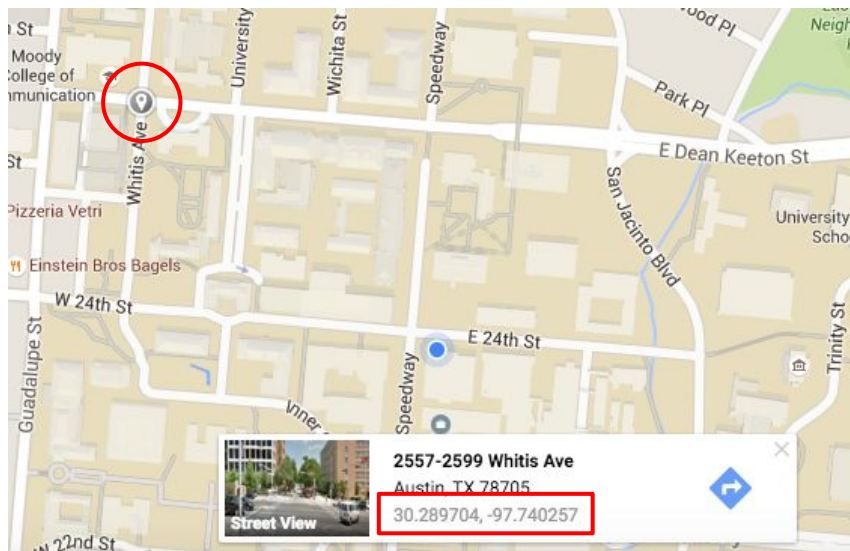
CS 381V Experiment Presentation

# Experiments Overview



- What **visual elements** distinguish a **university campus**?
  - a. Mining discriminative image patches.
  - b. Classifying campuses:
    - Which visual elements work best?
    - Examples where the algorithm succeeds and fails.
- **Campus forensics**: can we predict non-visual attributes?
  - a. Predicting student enrollment.
  - b. What visual elements work?

# Getting Data




- Google Maps API for Street View images:
  - <https://maps.googleapis.com/maps/api/streetview?size=640x640&location=30.289563,-97.7391009&heading=0>
- Linearly interpolate coordinates using Google Maps.



# Getting Data




~1200 images per campus.





# Extracting Patches




→ 48 patches → HOG + Lab descriptors

# Experiment 1: Patch Mining

- For each campus:
  - a. Randomly select 10k high-contrast patches.
  - b. Remove patches with high overlap with top 50 nearest neighbors.
  - c. Remove patches with more than 5 of the top 20 nearest neighbors.

# Experiment 1: Set 1



vs.



vs.



# Experiment 1: UT Austin




# Experiment 1: UT Austin



Randomly sampled patches.




# Experiment 1: UT Austin



Removed patches that have too much overlap with nearest neighbors.

# Experiment 1: UT Austin



Removed patches with too many negative nearest neighbors.

# Experiment 1: Texas A&M




# Experiment 1: Texas A&M






# Experiment 1: Stanford





# Experiment 1: Stanford



# Experiment 1: Set 2



vs.








vs.



# Experiment 1: CMU



# Experiment 1: CMU






# Experiment 1: NYU





# Experiment 1: NYU




# Experiment 1: Harvard



# Experiment 1: Harvard




# Experiment 1: Comparing Patches





# Experiment 1: UT Austin





# Experiment 1: Texas A&M



# Experiment 1: Stanford






# Experiment 1: CMU



# Experiment 1: CMU




Google

© Classroom Services



# Experiment 1: NYU





# Experiment 1: Harvard





# Experiment 1: Side-by-Side



# Experiment 1: Considerations


- Weather at time of Street View photography.
  - Color differences.
- Campus size and street coverage.
  - Interior images.
  - Panorama artifacts.
  - Not enough information.





# Experiment 2: Which Campus?

- Classify **top 1000 patches** for each campus.
- 800/200 training/validation split.



vs.




vs.









# Experiment 2: Which Campus?

SVM Results:







# Experiment 2: Which Campus?

CNN Results:


			
	<b>0.58</b>	0.10	0.32
	0.15	<b>0.57</b>	0.28
	0.24	0.23	<b>0.53</b>

# Experiment 2: Misclassifications






# Experiment 2: Which Campus?




6 errors

# Experiment 2: Which Campus?



18 errors


# Experiment 2: Which Campus?



29 errors



# Experiment 2: Which Campus?















46 errors



# Experiment 2: Which Campus?

CNN Results:













						
	<b>0.41</b>	0.06	0.11	0.06	0.18	0.18
	0.04	<b>0.58</b>	0.09	0.16	0.11	0.03
	0.12	0.10	<b>0.49</b>	0.10	0.14	0.05
	0.08	0.11	0.14	<b>0.59</b>	0.06	0.02
	0.05	0.10	0.11	0.10	<b>0.63</b>	0.01
	0.19	0.10	0.14	0.12	0.15	<b>0.30</b>













# Experiment 2: Which Campus?

- Classify **top 1000 patches + their top 5 nearest neighbors** for each campus.
- 4800/1200 training/validation split.

# Experiment 2: Which Campus?

CNN Results:

						
	<b>0.41</b>	0.06	0.11	0.06	0.18	0.18
	0.04	<b>0.58</b>	0.09	0.16	0.11	0.03
	0.12	0.10	<b>0.49</b>	0.10	0.14	0.05
	0.08	0.11	0.14	<b>0.59</b>	0.06	0.02
	0.05	0.10	0.11	0.10	<b>0.63</b>	0.01
	0.19	0.10	0.14	0.12	0.15	<b>0.30</b>

						
	<b>0.63</b>	0.02	0.05	0.02	0.09	0.19
	0.05	<b>0.51</b>	0.19	0.05	0.13	0.07
	0.14	0.05	<b>0.49</b>	0.08	0.12	0.12
	0.08	0.05	0.03	<b>0.66</b>	0.12	0.06
	0.10	0.04	0.05	0.04	<b>0.69</b>	0.08
	0.18	0.04	0.06	0.06	0.14	<b>0.52</b>

# Experiment 3: Campus Forensics

- Student enrollment: high or low?
  - Train on top 1000 patches + 5 nearest neighbors.
  - Test on new data set.

UT Austin	50,950	High
Texas A&M	58,577	High
Stanford	16,136	Low
CMU	13,285	Low
NYU	57,245	High
Harvard	21,000	Low



# Experiment 3: Campus Forensics

- Tested on new data (10k high-contrast patches).
  - Validation score (old data): 71.6%.
  - Test score (new data): **62.6%**.

Arizona State	83,308	High
Princeton	8,125	Low

# Experiment 3: Campus Forensics

Arizona State (high)

Correct




Wrong




# Experiment 3: Campus Forensics

Princeton (low)

Correct



Wrong





# Discussion / Future Experiments

- Patches vs. whole images?
- Using SVMs: could it work?
- Experimental limitations:
  - \*Small sample size (campuses are relatively small).
  - CPU time limitations.

# References

- [1] Arietta, Sean M., et al. "City forensics: Using visual elements to predict non-visual city attributes." *Visualization and Computer Graphics, IEEE Transactions on* 20.12 (2014): 2624-2633.
  
- [2] Doersch, Carl, et al. "What makes paris look like paris?." *ACM Transactions on Graphics* 31.4 (2012).