

Optimizing diversity and quality

Goal: Help decision-maker develop selection criteria

- E.g., for college admissions, image search, job search
- Ideally, should yield **high-quality** and **diverse** results

Challenge: True quality of each criterion may be unknown

- Decision-maker must rely on heuristics and intuition
- Makes it difficult to simultaneously optimize for quality and diversity

Example: Consider choosing college admissions criteria

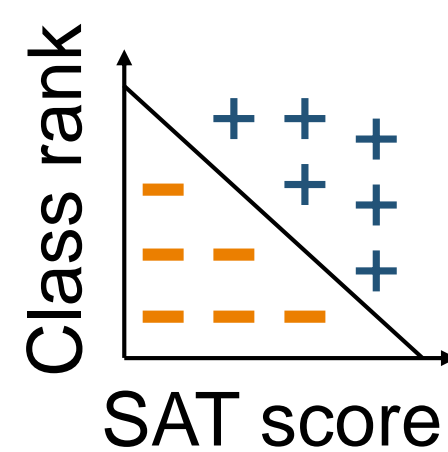
- Extremely limited information about any candidate's "true quality"
- Intuitively, searching for SAT > 1400 might yield high-quality results *but might return few candidates from minority groups*
- May be similar queries yielding high-quality, diverse candidates
E.g., SAT > 1000 and class rank in top 10%

Our model

Expert chooses a criterion, e.g.:

- ACT score above 56
- Programmer job applicants with 5+ years experience

Specifically, a criterion is a function $t : X \rightarrow R$



Every example $x \in X$ also associated with sensitive attributes (e.g., race)

Algorithmic framework

Suggests similar criteria with better (estimated) diversity. Relies on:

1. Application-specific similarity function
Measures how substitutable any two criteria are
2. Function measuring diversity of any criterion's results

Optimize similarity function while meeting diversity constraint

Application: Image search

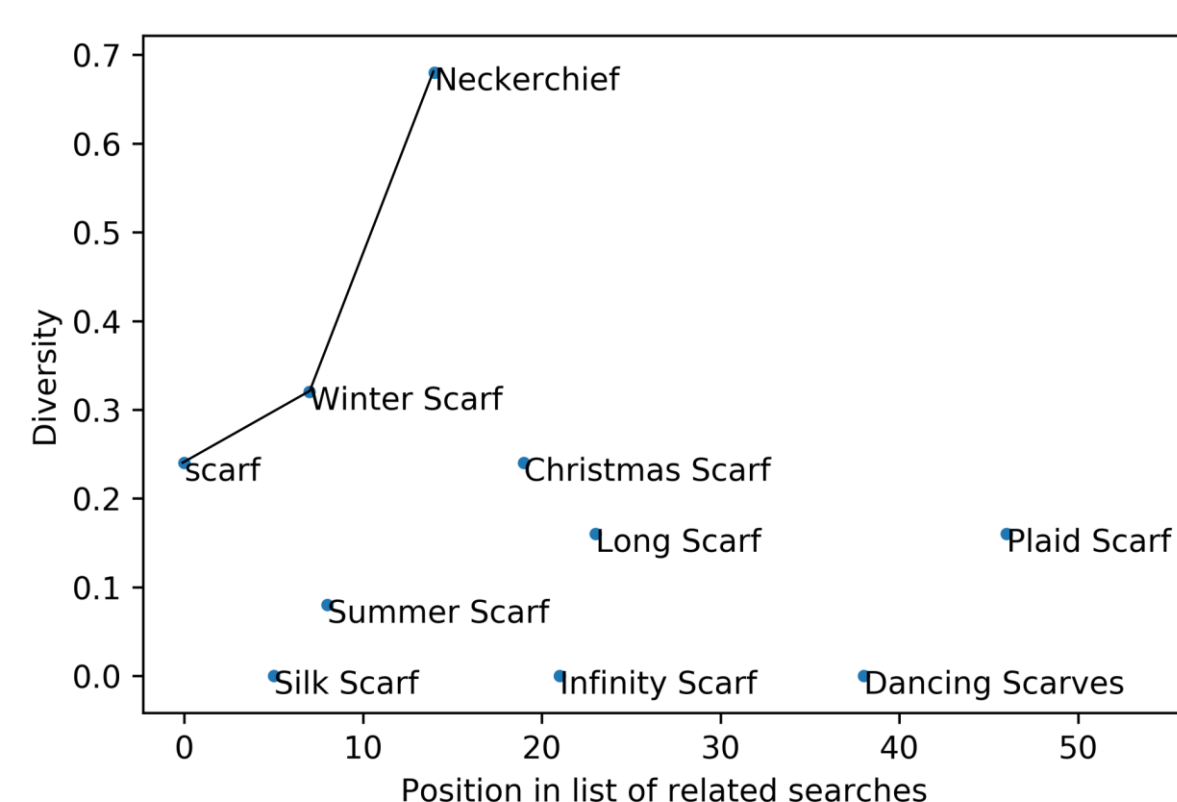
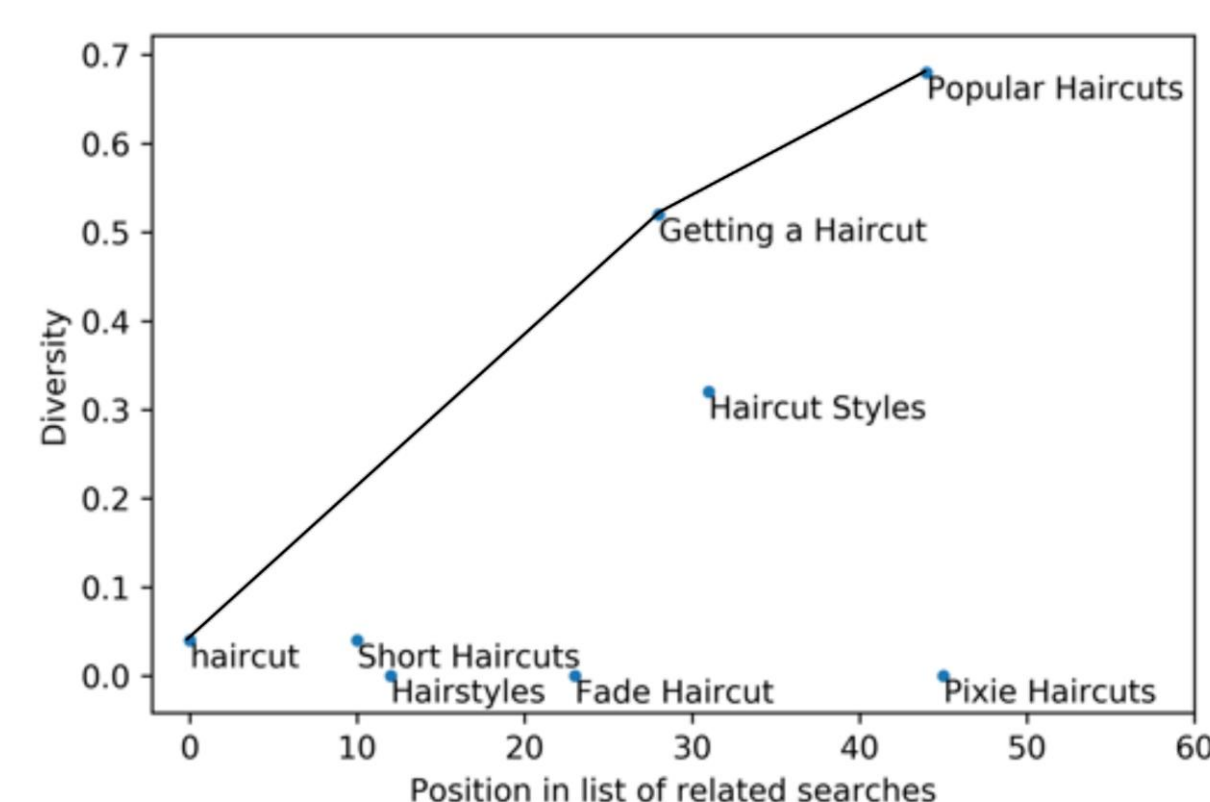
Similarity based on search suggestions:

Criteria t and t' are similar if t' high in list of searches related to t

Diversity of a criterion: How close to 50-50 is gender ratio of top 50 images?



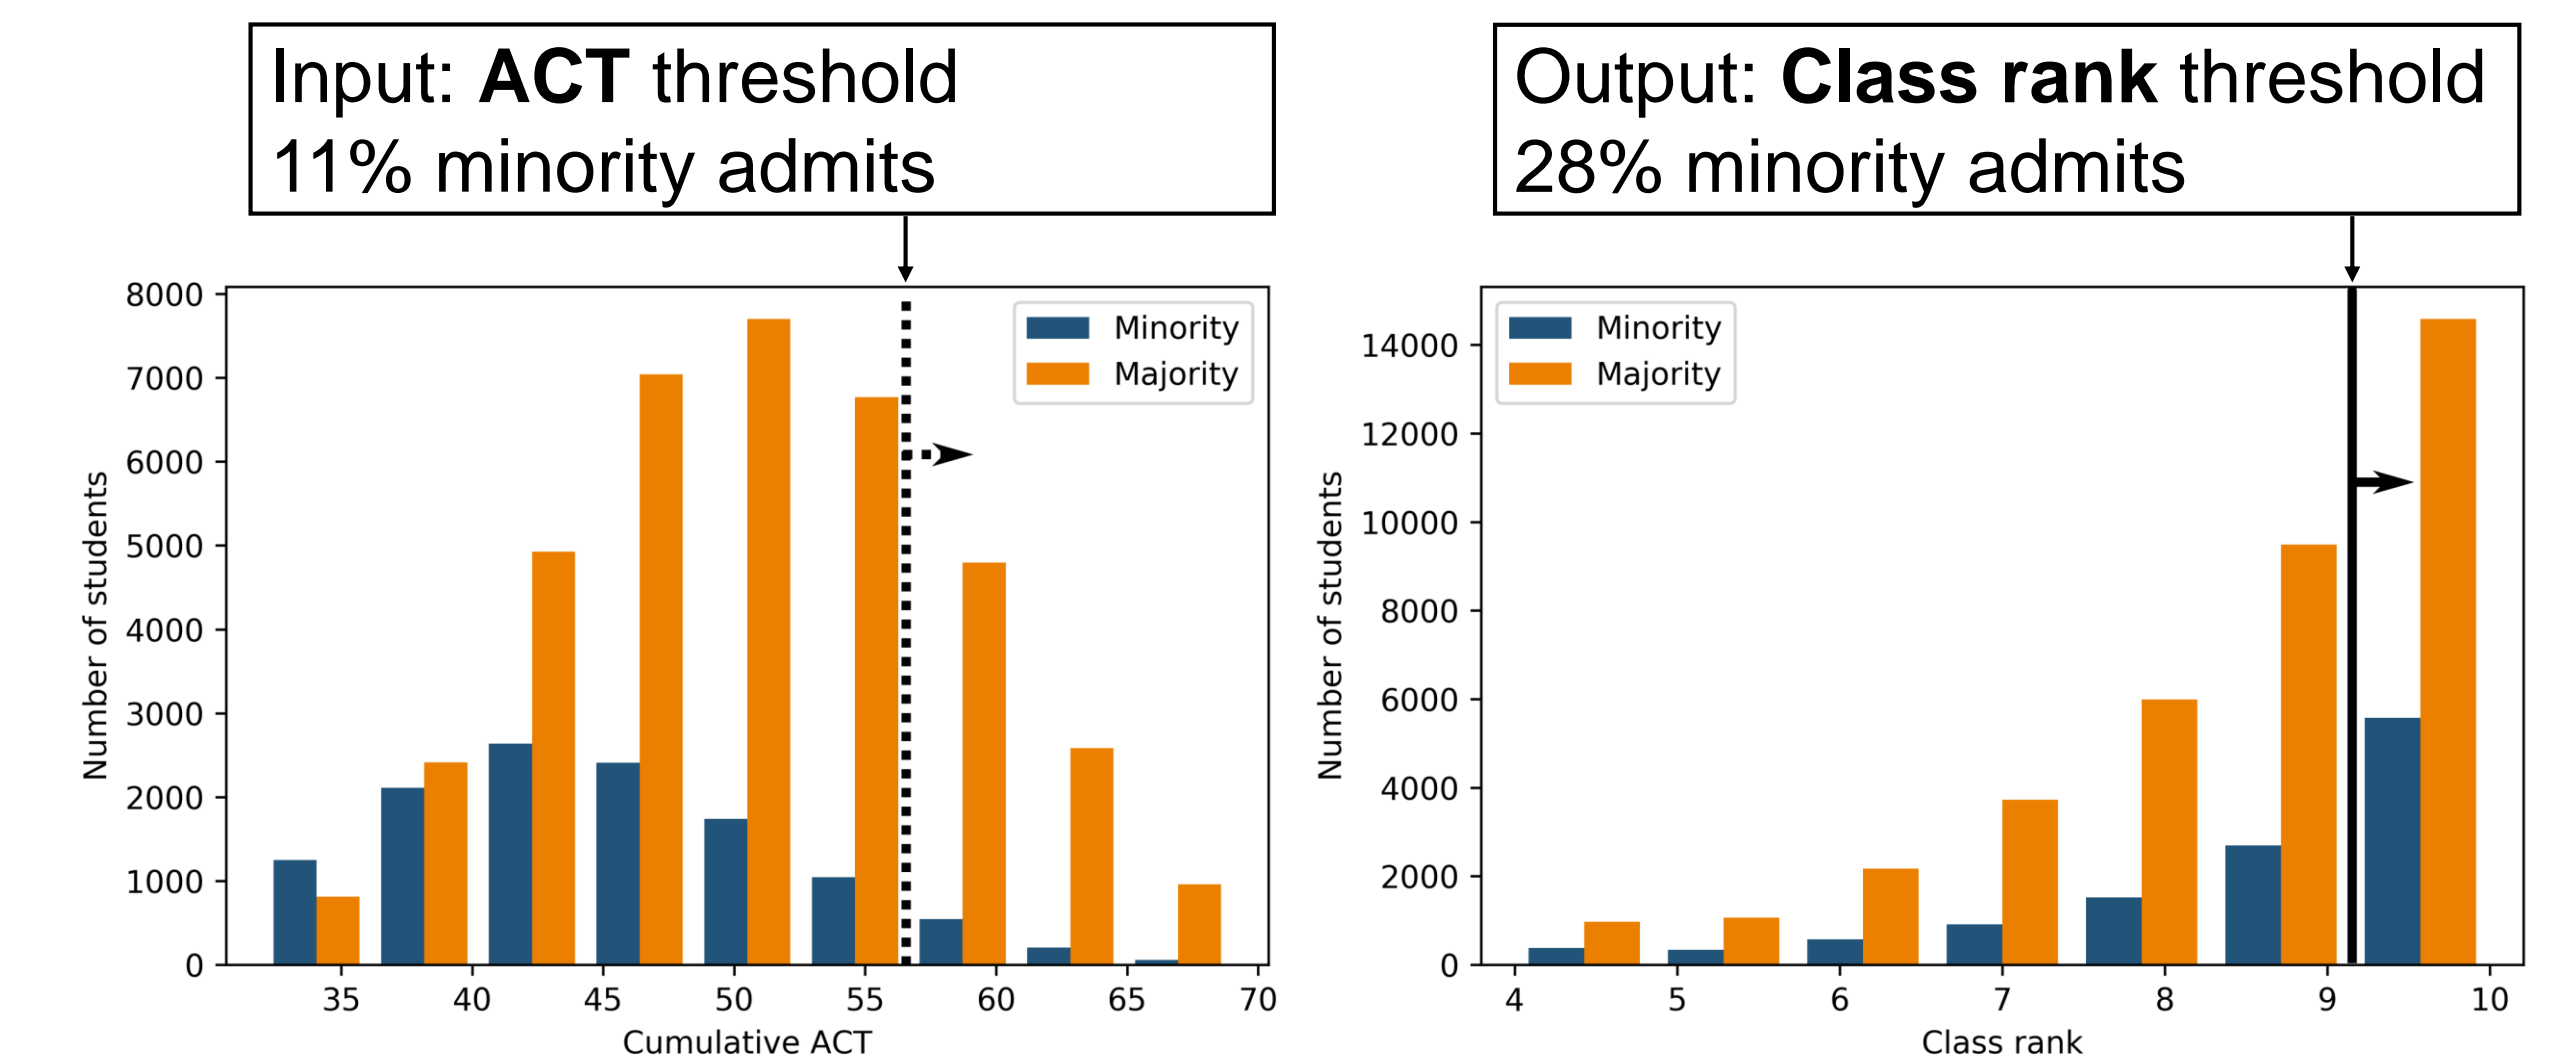
Movie director more similar to **director** than **executive** to **director**



Application: College admissions

- Run experiments on a dataset of UT Austin applicants
- **Observable** features: SAT score, ACT score, class rank
- **Protected** feature: Minority (black/Hispanic) or majority (Caucasian) group
- Each criterion t accepts set E_t of students
- **Similarity** between criteria t and t' :
Function of symmetric difference between admits: $|E_t \cup E_{t'}| - |E_t \cap E_{t'}|$
- **Diversity** of criterion: Fraction of admits from minority group

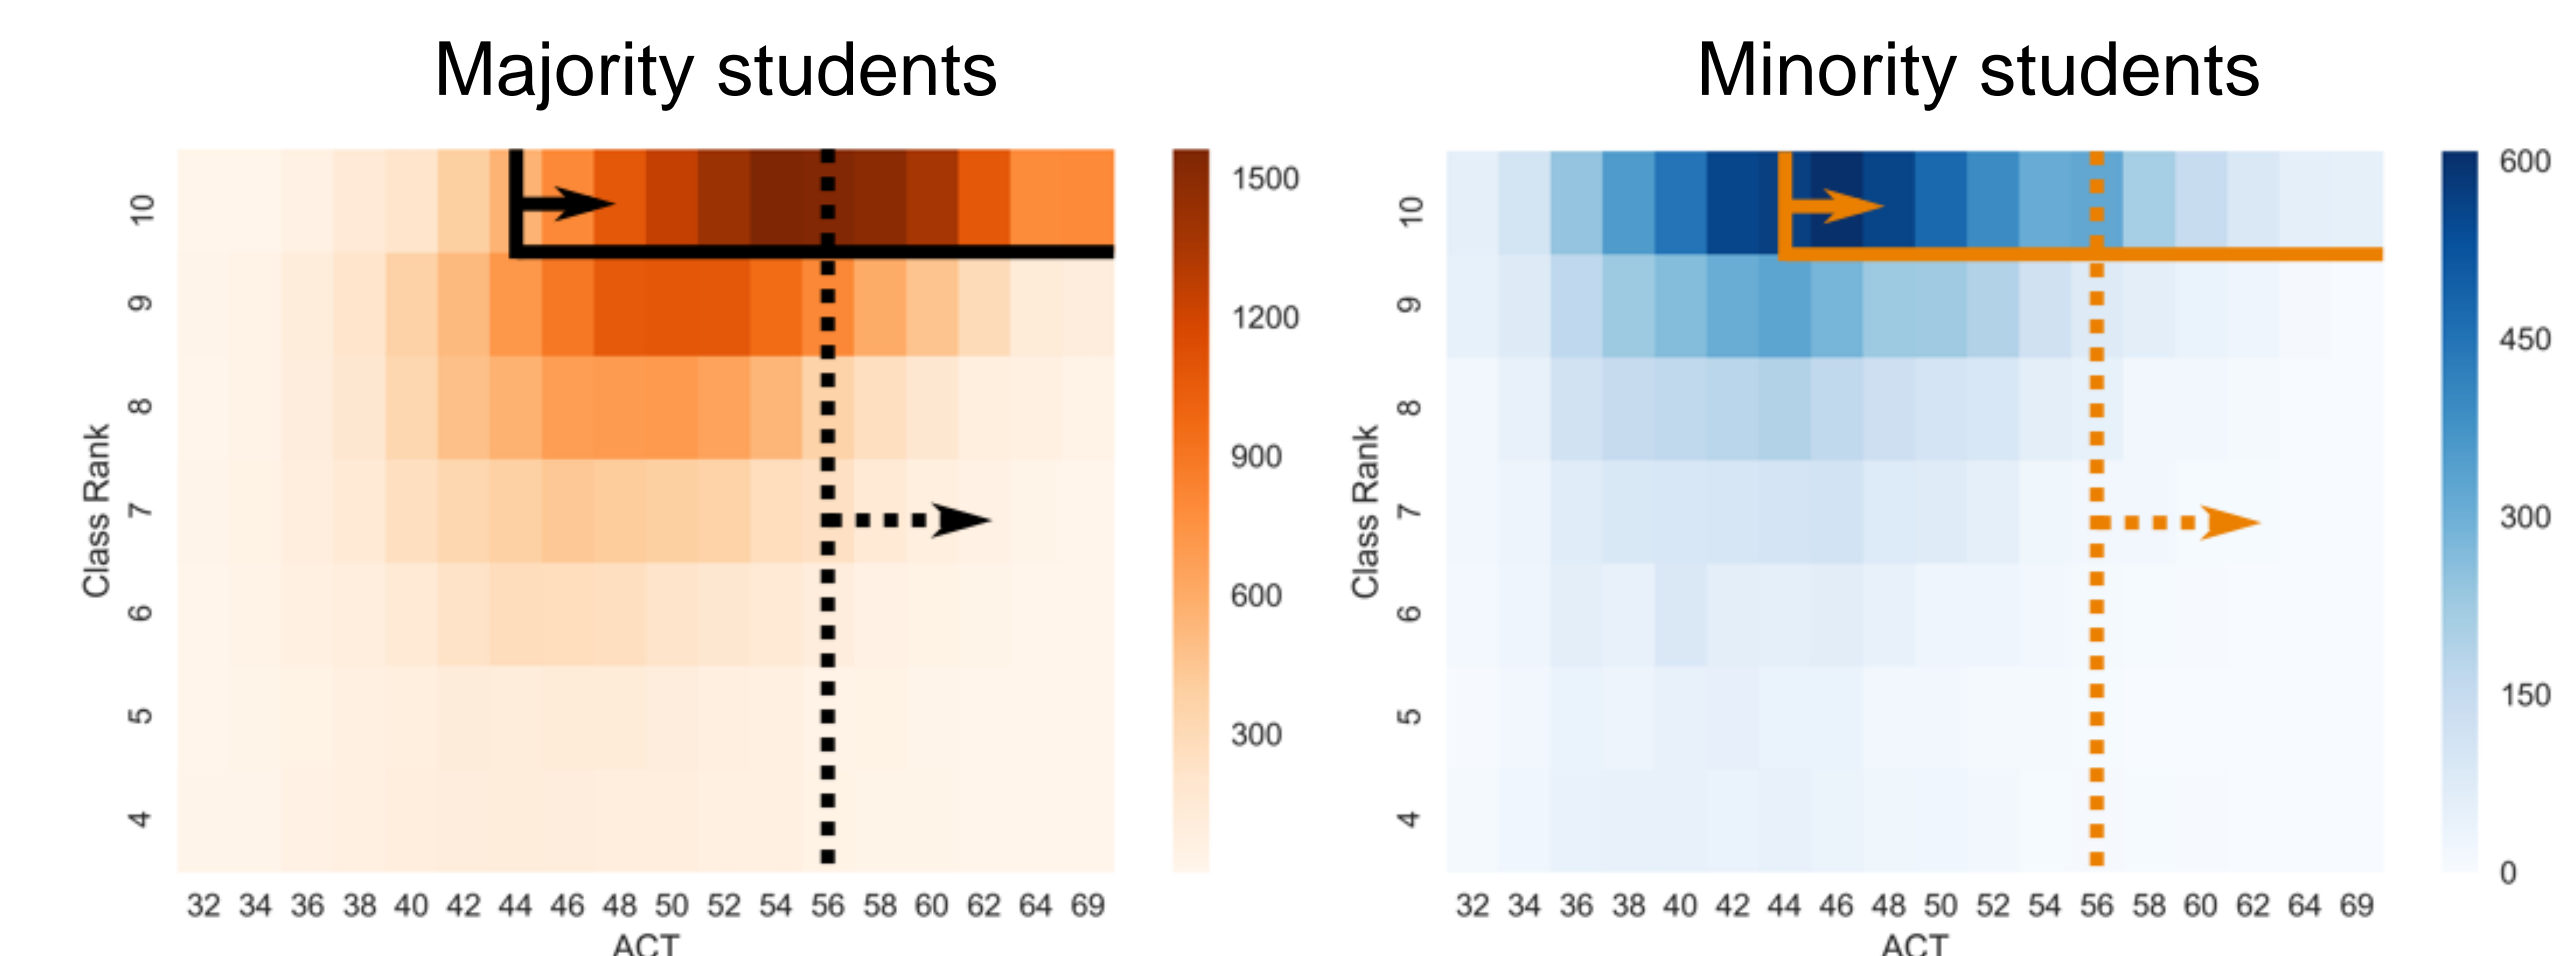
Input criterion: Axis-aligned hyperplanes in observable feature space



We rediscover TX's "top 10% rule":

TX students in top 10% guaranteed admission into any public university


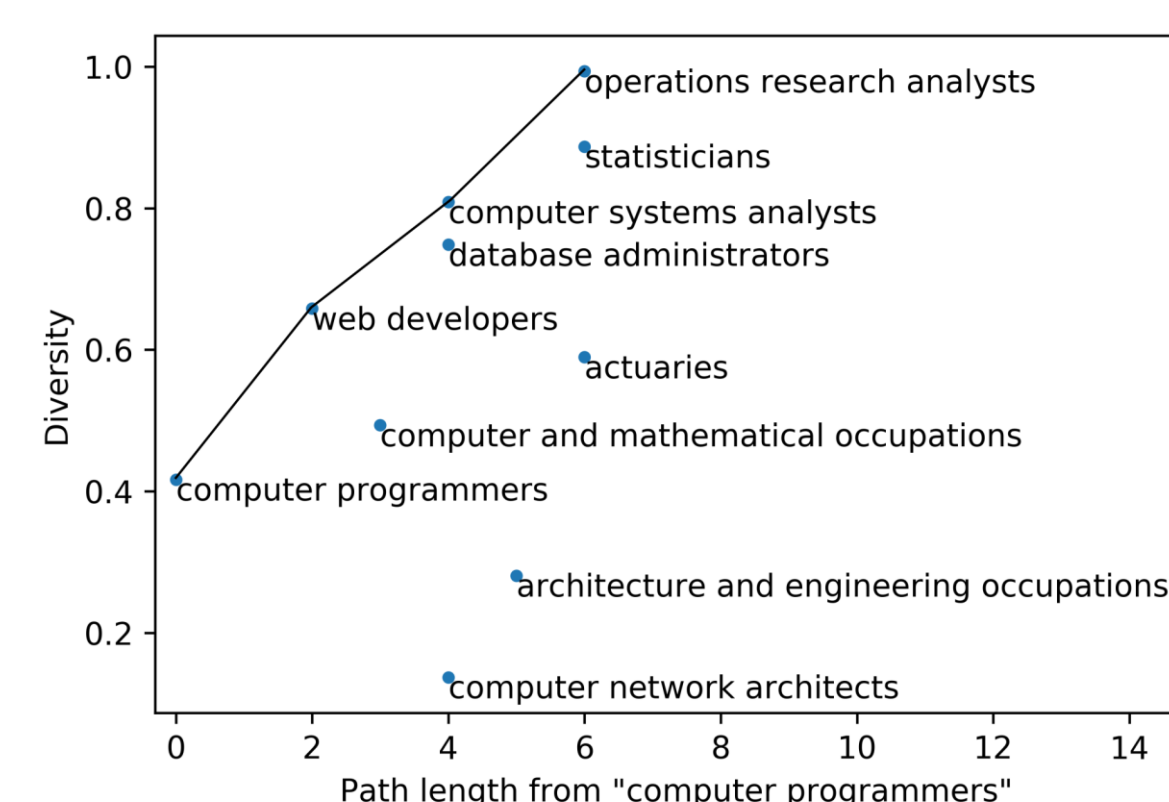
Input criterion: Intersection of axis-aligned hyperplanes



- Input: ACT threshold 11% minority admits
- Output: Relaxed ACT but stricter CR cutoff 15% minority admits

Application: Job applicant search

- Each criterion is a job title
- **Similarity** based on DOL's Standard Occupational Classification System
- **Diversity:** how close to 50-50 is gender ratio (based on DOL statistics)?



Diversity estimation

Provable guarantees for estimating criteria diversity given historical data

- Algorithm can use diversity estimates instead of true diversity scores
- May be preferable (or mandatory) to use estimates, not protected attributes

- Bound dataset size sufficient for estimation
- Relate sample complexity to "intrinsic complexity" (VC dimension) of set of criteria

