

*Stellingen behorende bij het proefschrift: Explainable Computer Vision for Cities*

1. Geographical leakage is the largest confounder in urban vision modelling and evaluations should always be performed on data that is geographically disjoint from the training data.
2. The arbitrary nature of the data science pipeline for preprocessing street view images and linking it to socio-economic data is detrimental to both the reproducibility of research as well as its practical use.
3. The low signal-to-noise-ratio of street view images has a strong relation to camera proximity, which in itself is influenced by the width of the street.
4. Evaluation of AI methods with human participants is affected by a positivity bias, as participants are hesitant to provide negative feedback out of fear of ruining what they believe to be the researcher's hypothesis.
5. The over/under sampling boundary of the Google Street View coverage distribution in Johannesburg follows the apartheid line.
6. Every researcher using street view images has encountered a point at which they deem privacy violations to be collateral damage.
7. Uncovering relevant visual elements from street view imagery is like trying to win a toy from a claw machine by sticking your hand through the prize chute.
8. The ambition perpetuated by big tech CEO's of replacing x% of the workforce with AI is pathetic and embarrassingly lays bare their perceptions regarding what it means to be an expert, what a *real* job entails, what the capabilities of their own models are, or what the value of work is within a society.
9. Feedback can be broadly divided into two categories: People who know *what* is wrong, and people who can tell you exactly what to change so it is *right*. Gravitate to the latter and be wary of people of the first category masquerading as the second.
10. Discipline is not to do something every day, but to miss a day and keep going.
11. While it takes guts to quit, there is merit to finishing things.
12. Unfunny comedy is rarely true.