

# Supplemental Material for The Topology and Language of Relationships in the Visual Genome Dataset

David Abou Chacra  
University of Waterloo  
Ontario, Canada

d2abouchacra at uwaterloo.ca

John Zelek  
University of Waterloo  
Ontario, Canada

jzelek at uwaterloo.ca

## 1. Full Figures

In this supplemental material we provide the full figure for the inverse relationship analysis on all 50 relationship predicates, see Figure 1. We also provide the full topological analysis in Figure 2 and the directional analysis in Figure 3. These figures are best viewed on a screen where you can zoom into them.




Figure 1. A heatmap of the occurrence of inverse relationships for all 50 predicates in the VG200 dataset. The numbers in parentheses are the total occurrences for this predicate in the dataset. For every row the value in the heat map reflects the ratio of: (inverse relationship occurrences of the row predicate with the predicate in the column) to (total occurrences of the predicate in the row).

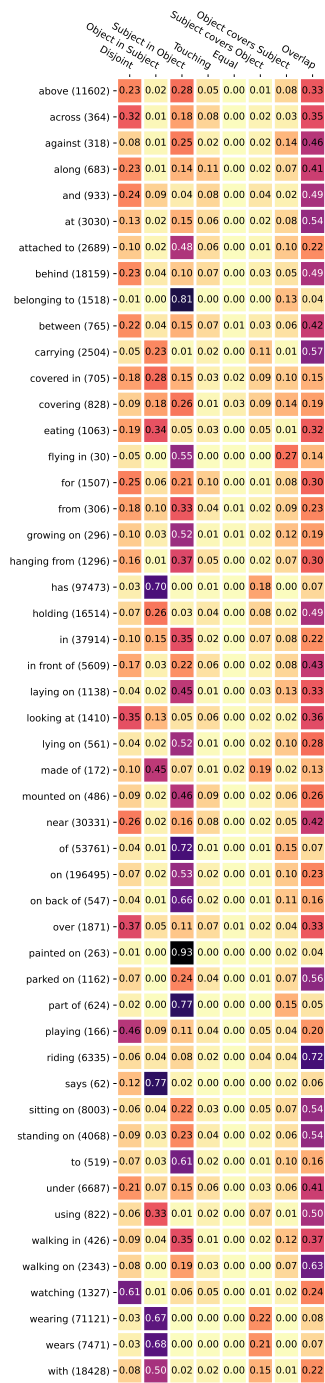


Figure 2. A heatmap of the occurrence of topological relationships between bounding boxes related by all 50 predicates in the VG200 dataset. The values shown in the heatmap are the portions of the total occurrences of the row predicate that exhibit the specific topological configuration in the column. The values in parenthesis next to the predicate names are the total occurrences of that predicate.

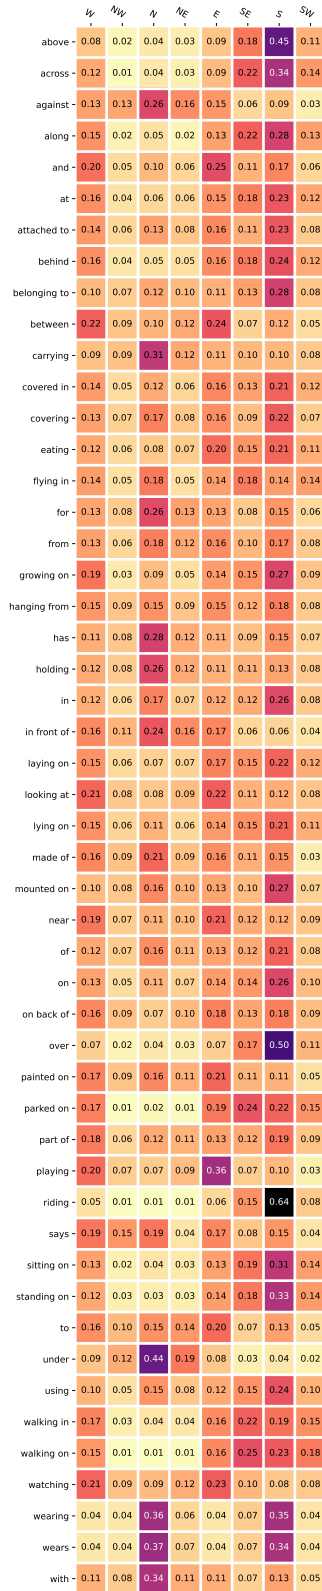


Figure 3. A heatmap of the angles between subject and object for all 50 relationship predicates in the VG200 dataset. The values shown in the heatmap are the portions of the total occurrences of the row predicate that exhibit the specific directional configuration in the column.