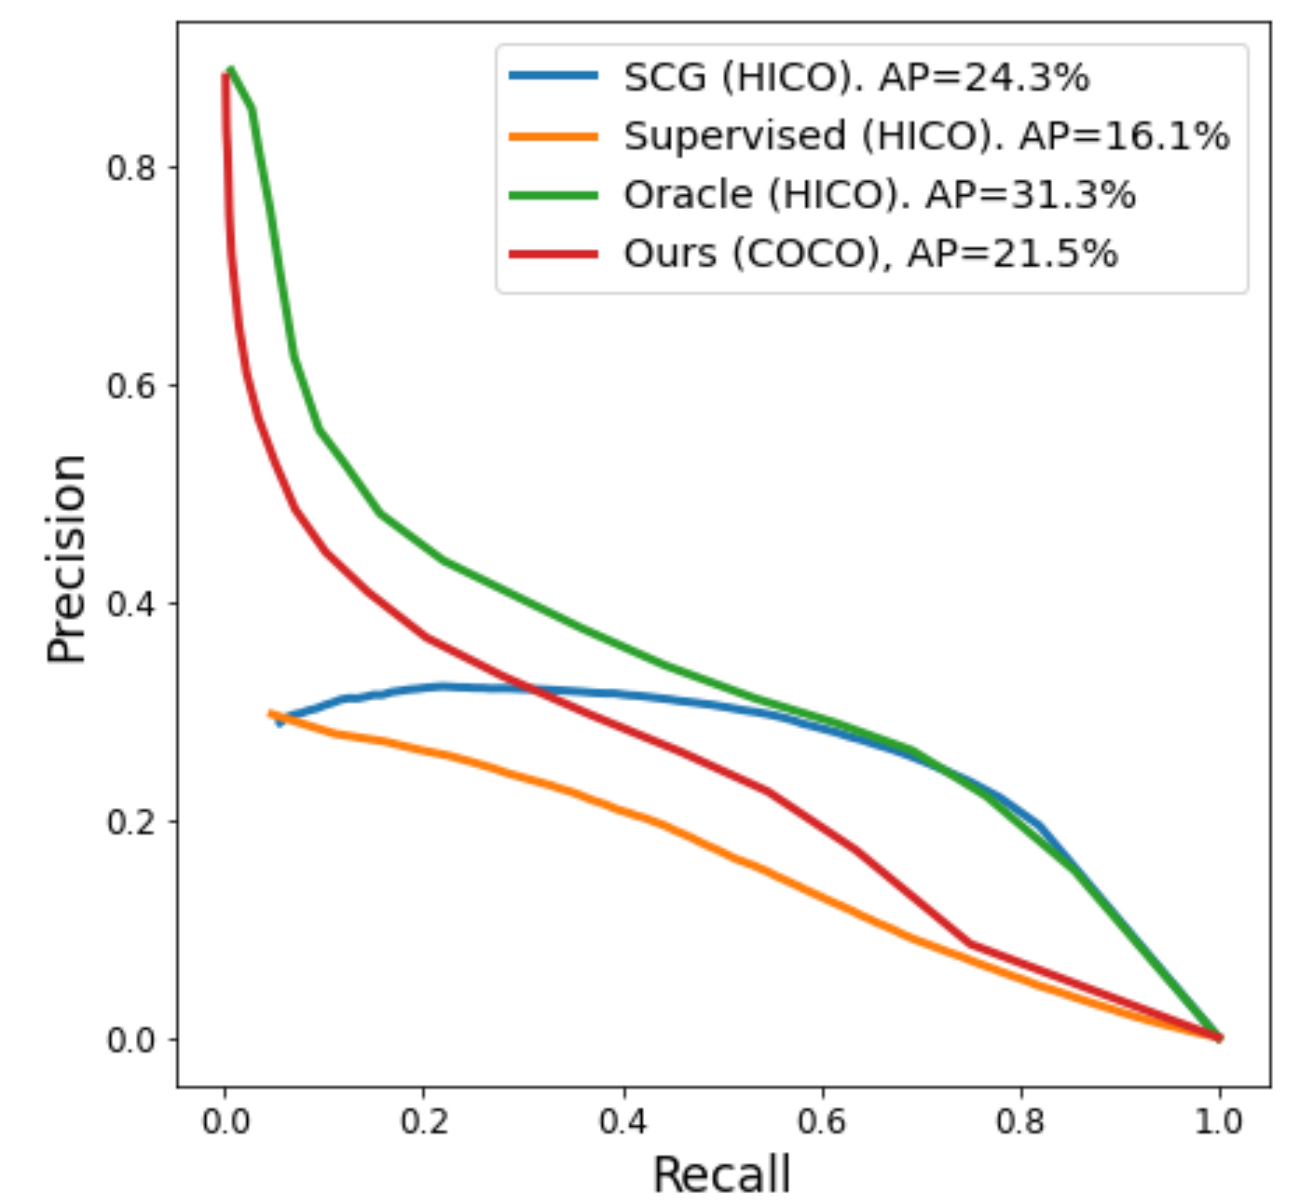




Contextual Visual Feature Learning for Zero-Shot Recognition of Human-Object Interactions

Tsung-Wei Ke Dong-Jin Kim Stella X. Yu Liang Gou Liu Ren

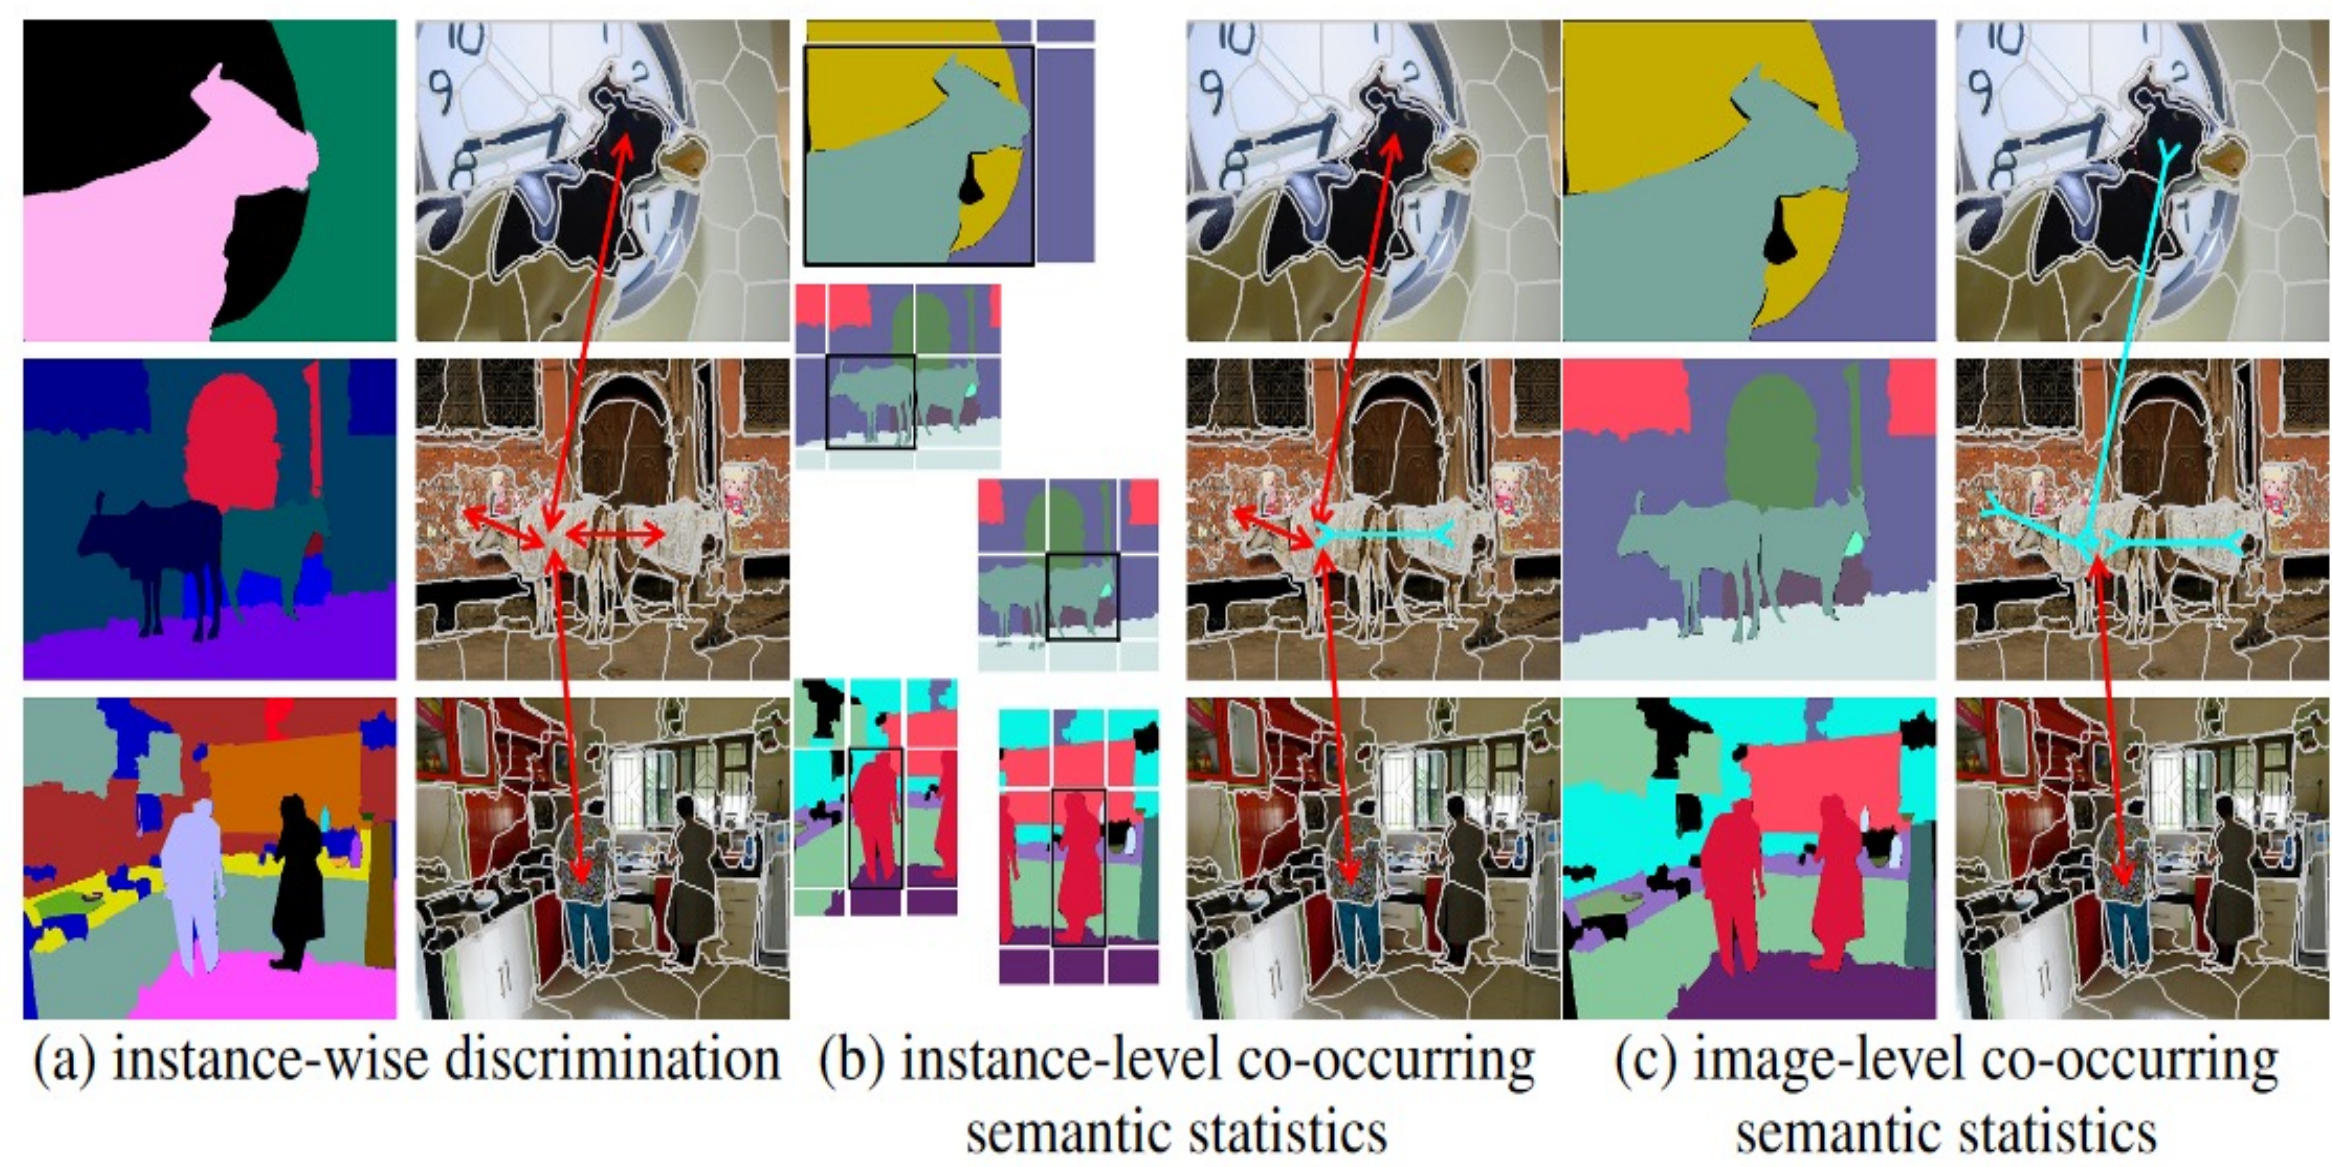
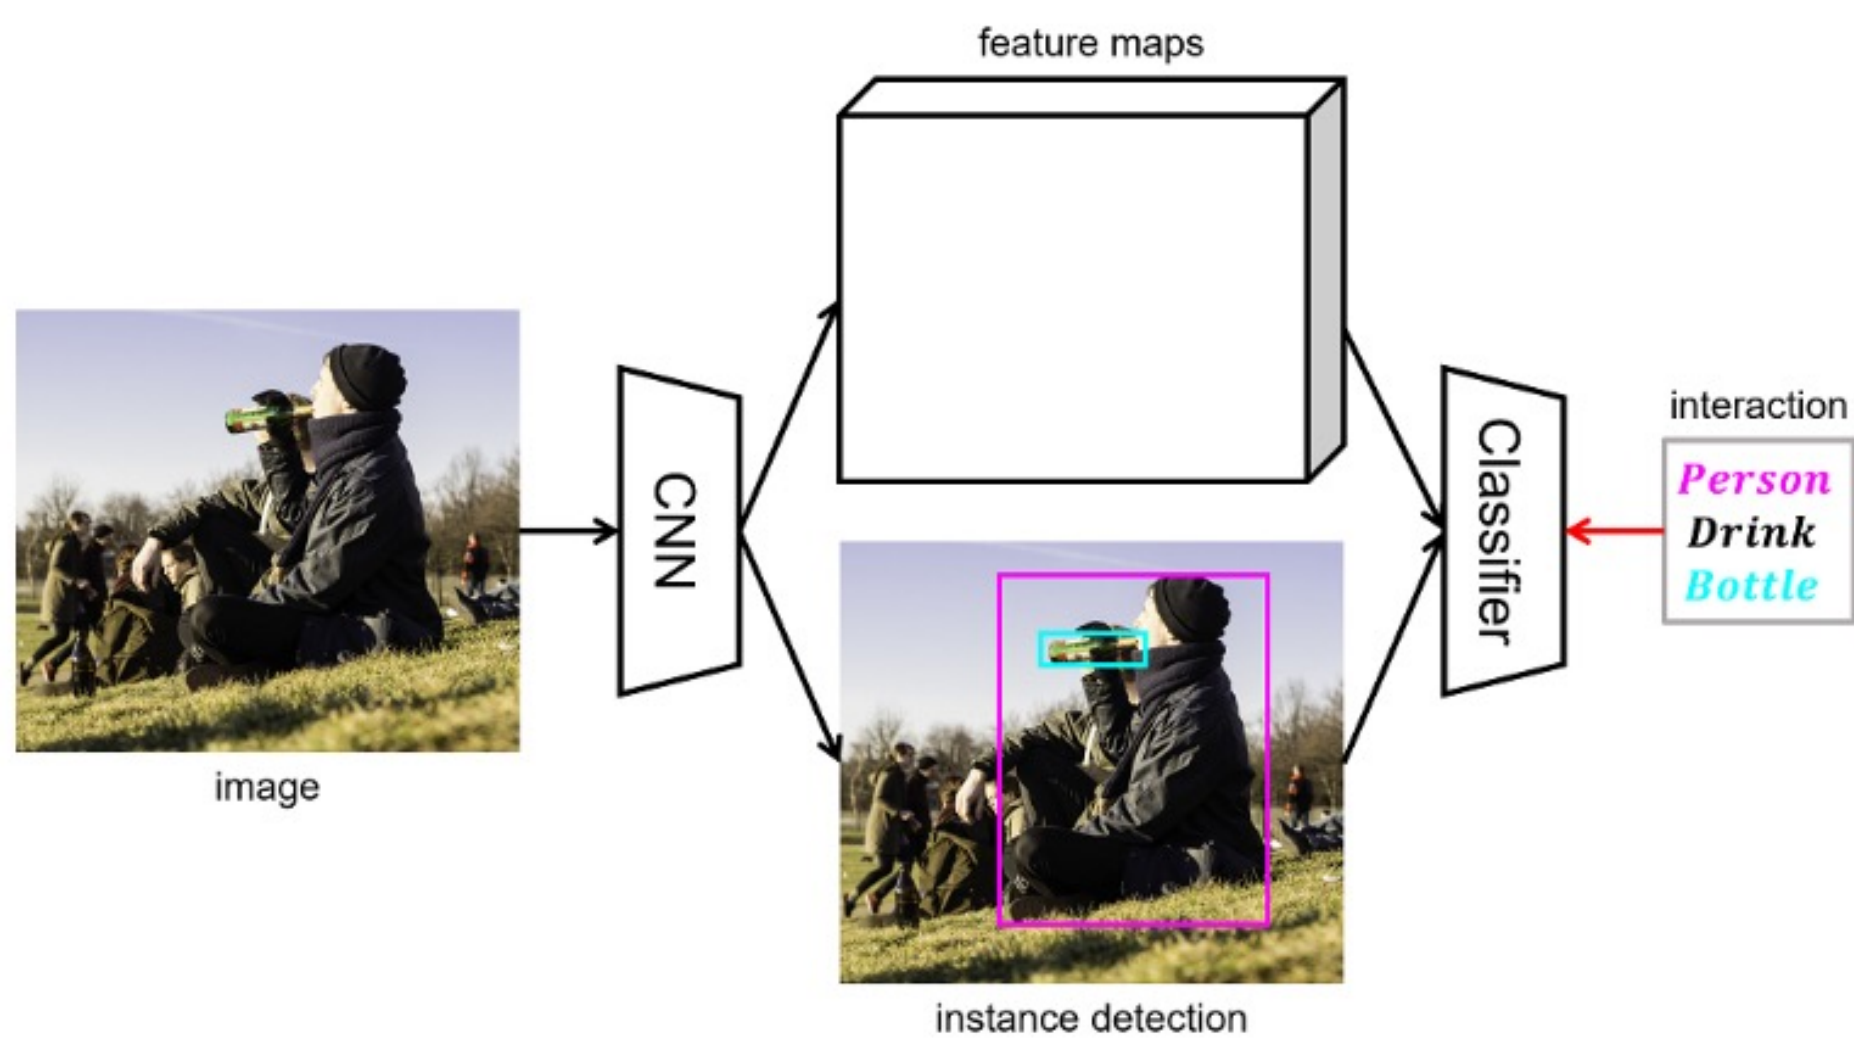


High-level Contextual Semantics Emerge from Our Learned Features



Baselines: Classify Human-Object Pairs.
Training Needs Interaction Labels.

We Contrast Pixels on 3 Types of Relationships



Our Work: Maps Pixels to Latent Features.
Training Doesn't Need Interaction Labels.

We Model Contexts with Spatially Co-occurred Semantics

