

- How should we rethink saliency benchmarks?
- What new tasks should we evaluate saliency models on?
- Can saliency help other visual recognition tasks, or are they going to help saliency instead?
- How can we use sequences and temporal information?
- Is saliency ready for production-grade applications? i.e., can I build saliency into my app today?
- What about training population-specific or user-specific saliency models?
- Can we use saliency for task decoding? Diagnosis?
- How much bigger do the datasets need to be? What are we missing?
- Will mobile eye trackers solve our problems?
- What old-school ideas have we forgotten about?