


Learning Perceptual Inference by Contrasting

<http://wellyzhang.github.io/project/copinet.html>

Chi Zhang*, Baoxiong Jia*, Feng Gao, Yixin Zhu, Hongjing Lu, Song-Chun Zhu
(* indicates equal contribution)

Poster: 10:45 AM - 12:45 PM @ East Exhibition Hall B + C #193






Motivation

- Raven's Progressive Matrices is a representative task on “thinking in pictures”, or spatial-temporal reasoning.
- Study on contrast effects on cognitive science, biology, and computer science.
- Interplay between perception and inference detailed in Carpenter *et al.* for humans to solve RPM.
- Missing permutation-invariance in computational models.

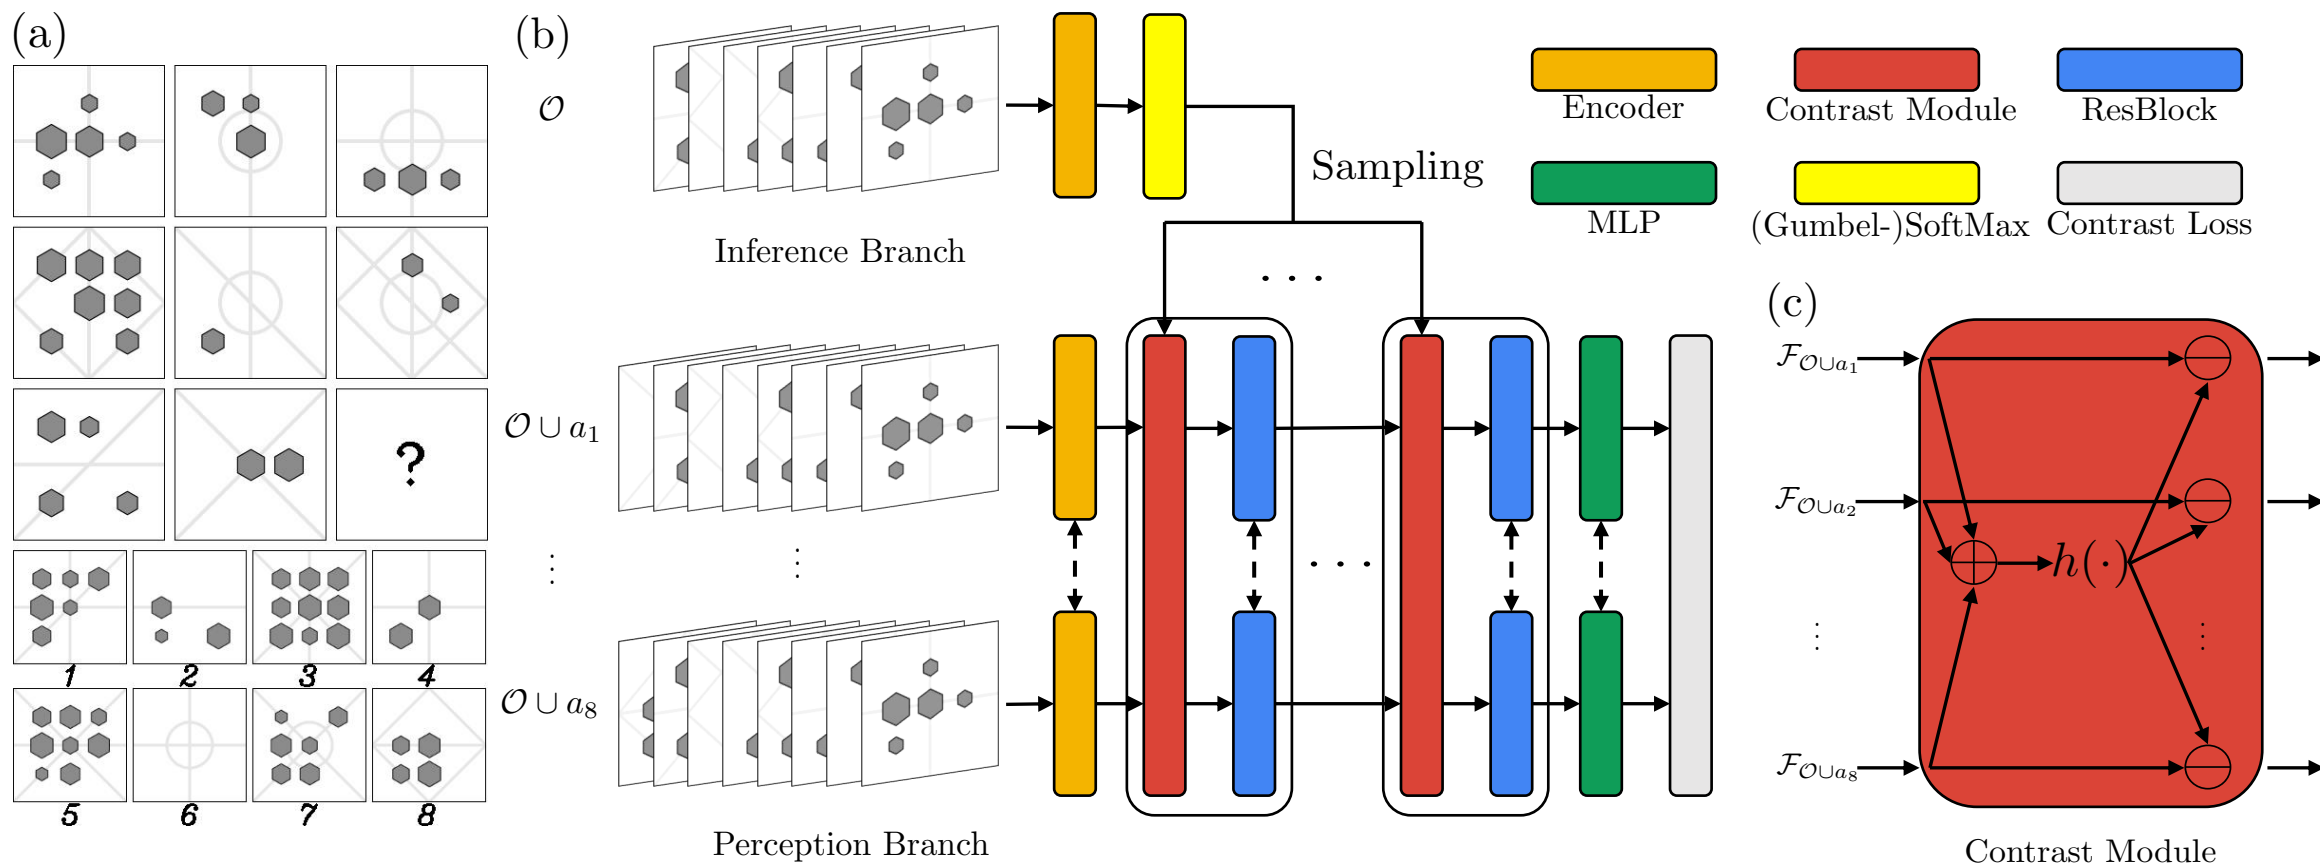
CoPINet: Formulation




A ranking perspective:

$$p(a_{\star}|\mathcal{O}) \geq p(a'|\mathcal{O}), \quad \forall a' \in \mathcal{A}, a' \neq a_{\star}$$

CoPINet: Overview



CoPINet: Contrast on Model



$$\text{Contrast}(\mathcal{F}_{O \cup a}) = \mathcal{F}_{O \cup a} - h \left(\sum_{a' \in \mathcal{A}} \mathcal{F}_{O \cup a'} \right)$$



CoPINet: Contrast on Objective

- If we model $p(a|\mathcal{O}) = \frac{1}{Z} \exp(f(\mathcal{O} \cup a))$
- Take log on both side of $p(a_\star|\mathcal{O}) \geq p(a'|\mathcal{O}), \quad \forall a' \in \mathcal{A}, a' \neq a_\star$
 $\log p(a_\star|\mathcal{O}) - \log p(a'|\mathcal{O}) = f(\mathcal{O} \cup a_\star) - f(\mathcal{O} \cup a') \geq 0, \quad \forall a' \in \mathcal{A}, a' \neq a_\star$
- Take to ***infinity***
 $f(\mathcal{O} \cup a_\star) - f(\mathcal{O} \cup a') \rightarrow \infty \iff \sigma(f(\mathcal{O} \cup a_\star) - f(\mathcal{O} \cup a')) \rightarrow 1$
- Add baseline and turn it into its sufficient conditions
 $f(\mathcal{O} \cup a_\star) - b(\mathcal{O} \cup a_\star) \rightarrow \infty \iff \sigma(f(\mathcal{O} \cup a_\star) - b(\mathcal{O} \cup a_\star)) \rightarrow 1$
 $f(\mathcal{O} \cup a') - b(\mathcal{O} \cup a') \rightarrow -\infty \iff \sigma(f(\mathcal{O} \cup a') - b(\mathcal{O} \cup a')) \rightarrow 0$
- Final objective
 $\ell = \log(\sigma(f(\mathcal{O} \cup a_\star) - b(\mathcal{O} \cup a_\star))) + \sum_{a' \in \mathcal{A}, a' \neq a_\star} \log(1 - \sigma(f(\mathcal{O} \cup a') - b(\mathcal{O} \cup a')))$



CoPINet: Perceptual Inference

- Assume each of N attributes is governed by one of M rules

$$p(\mathcal{T}|\mathcal{O}) = \prod_{i=1}^N p(t_i|\mathcal{O})$$

- Modify the form of distribution to take rules as hidden variables

$$\log p(a|\mathcal{O}) = \log \sum_{\mathcal{T}} p(a|\mathcal{T}, \mathcal{O})p(\mathcal{T}|\mathcal{O}) = \log \mathbb{E}_{\mathcal{T} \sim p(\mathcal{T}|\mathcal{O})} [p(a|\mathcal{T}, \mathcal{O})]$$

- Sample rules using Gumbel-SoftMax and optimize

$$\ell = \log(\sigma(f(\mathcal{O} \cup a_*, \hat{\mathcal{T}}) - b(\mathcal{O} \cup a_*))) + \sum_{a' \in \mathcal{A}, a' \neq a_*} \log(1 - \sigma(f(\mathcal{O} \cup a', \hat{\mathcal{T}}) - b(\mathcal{O} \cup a')))$$



Performance: PGM

General performance

Method	CNN	LSTM	ResNet	Wild-ResNet	WReN-NoTag-Aux	CoPINet
Acc	33.00%	35.80%	42.00%	48.00%	49.10%	56.37%

Ablation study

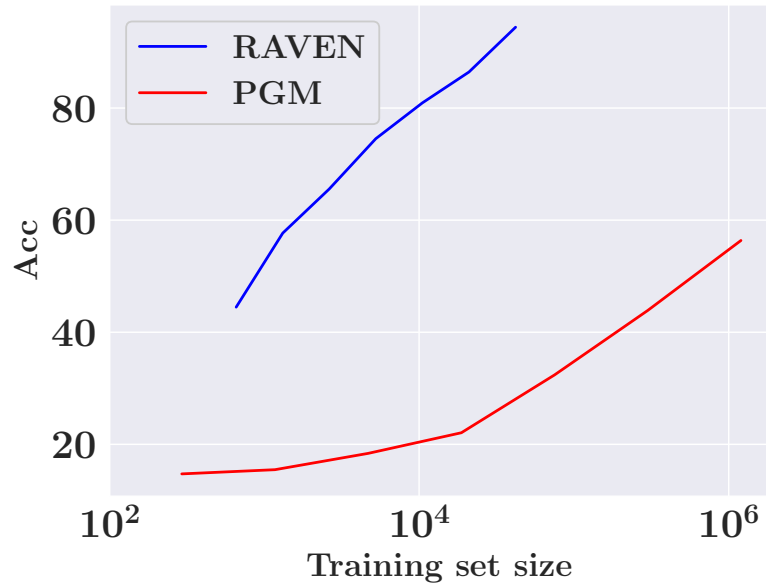
Method	WReN-NoTag-NoAux	WReN-NoTag-Aux	WReN-Tag-NoAux	WReN-Tag-Aux
Acc	39.25%	49.10%	62.45%	77.94%

Method	CoPINet-Backbone-XE	CoPINet-Contrast-XE	CoPINet-Contrast-CL	CoPINet
Acc	42.10%	51.04%	54.19%	56.37%



Performance: Effect of Dataset Size

Acc vs dataset size



Dataset size on RAVEN

Training set size	Acc
658	44.48%
1,316	57.69%
2,625	65.55%
5,250	74.53%
10,500	80.92%
21,000	86.43%

Dataset size on PGM

Training set size	Acc
293	14.73%
1,172	15.48%
4,688	18.39%
18,750	22.07%
75,000	32.39%
300,000	43.89%



Thanks!

See you later at East Exhibition Hall B + C #193

