Humans Learn Using Manifolds, Reluctantly

Bryan Gibson, Xiaojin Zhu,
Timothy Rogers, Charles Kalish, and Joseph Harrison

University of Wisconsin-Madison

A Familiar Task

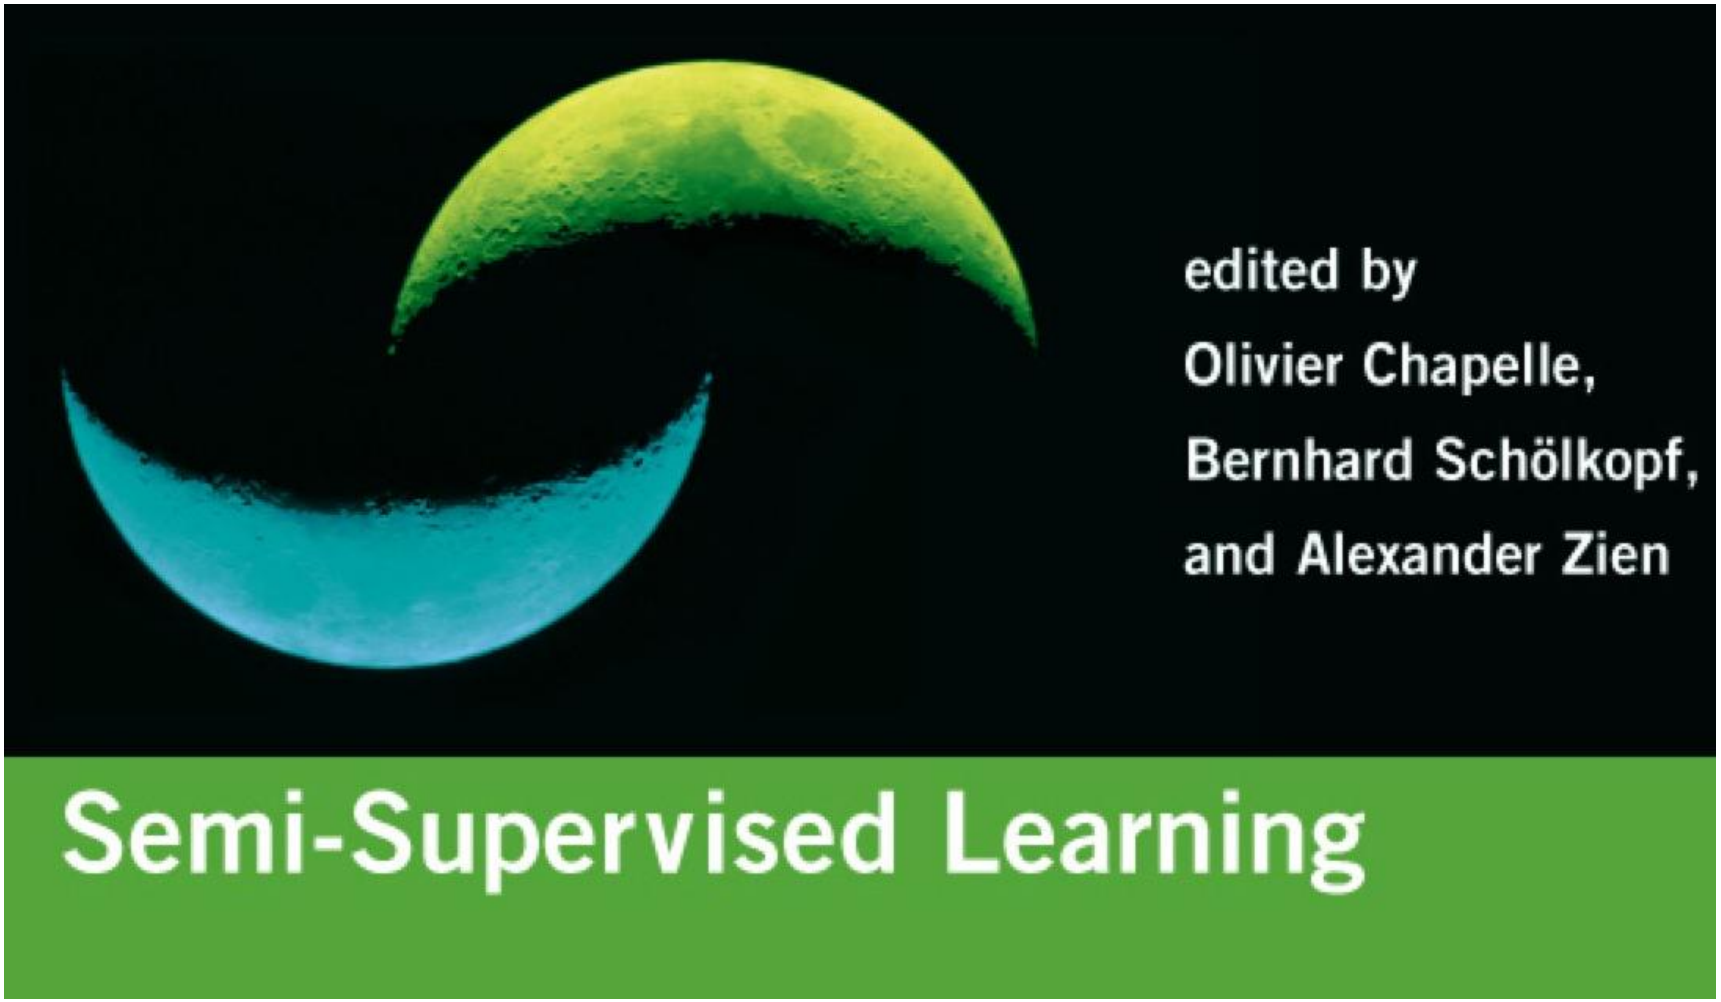

A Familiar Task

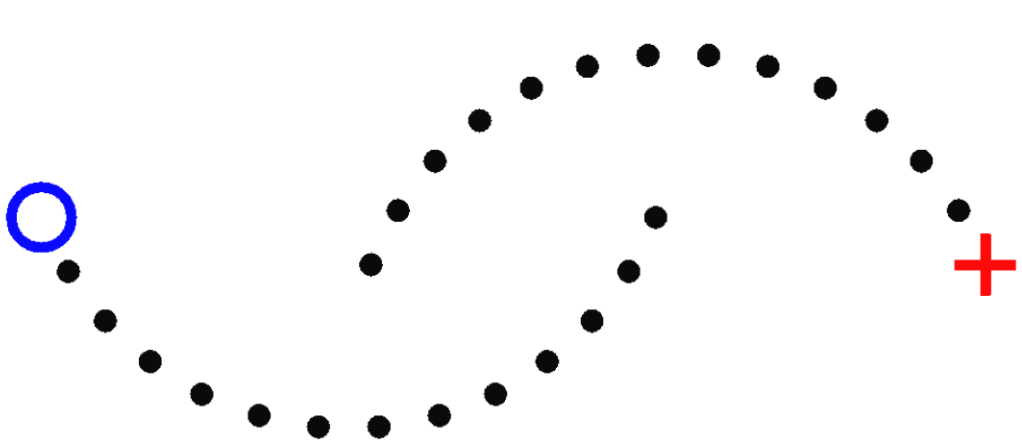

Question: which one would humans do?

Contribution: an empirical study of human manifold learning behaviors

if we choose supervised learning

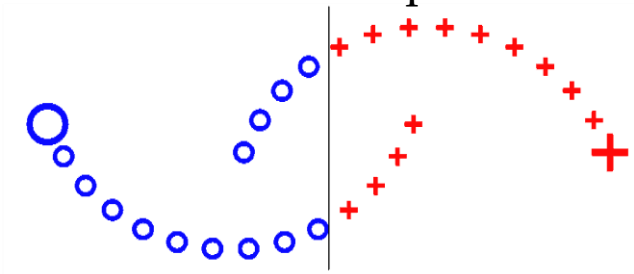

semi-supervised learning

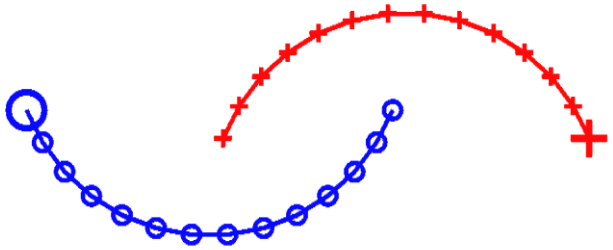

Manifolds are Common in Life

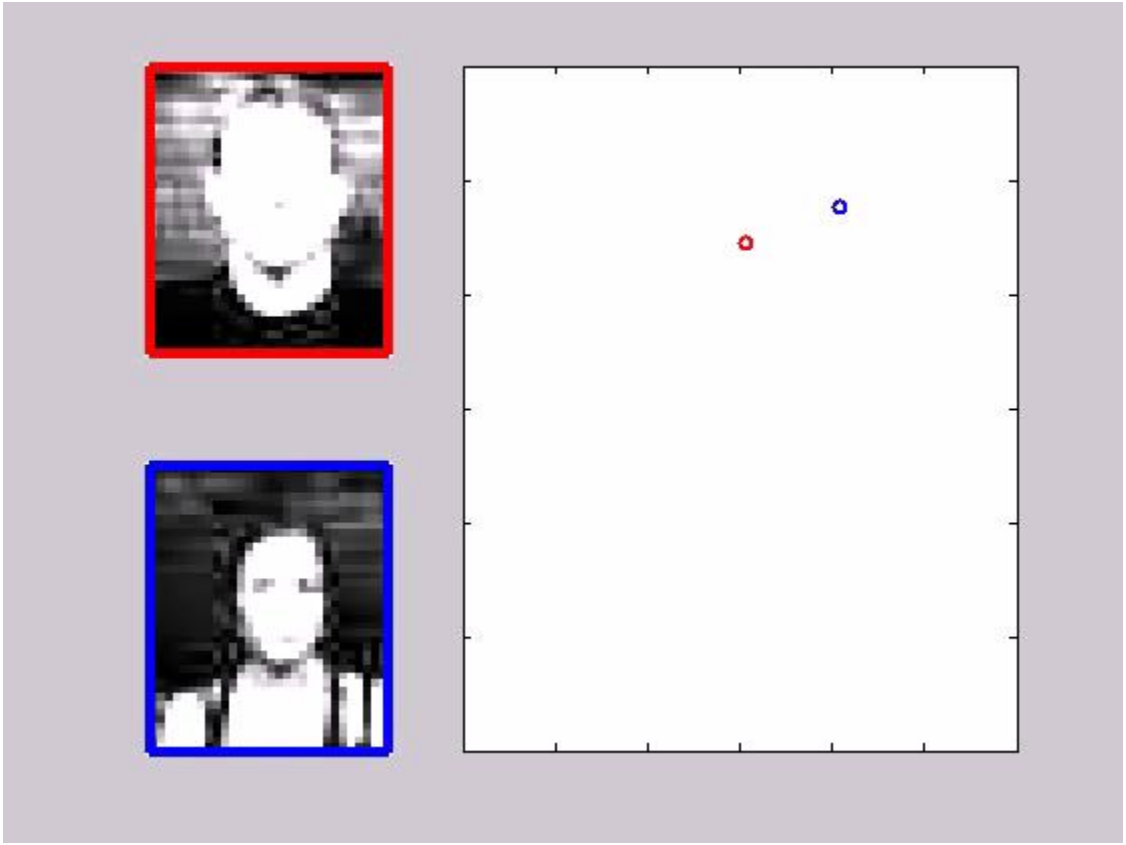

Manifolds are Common in Life

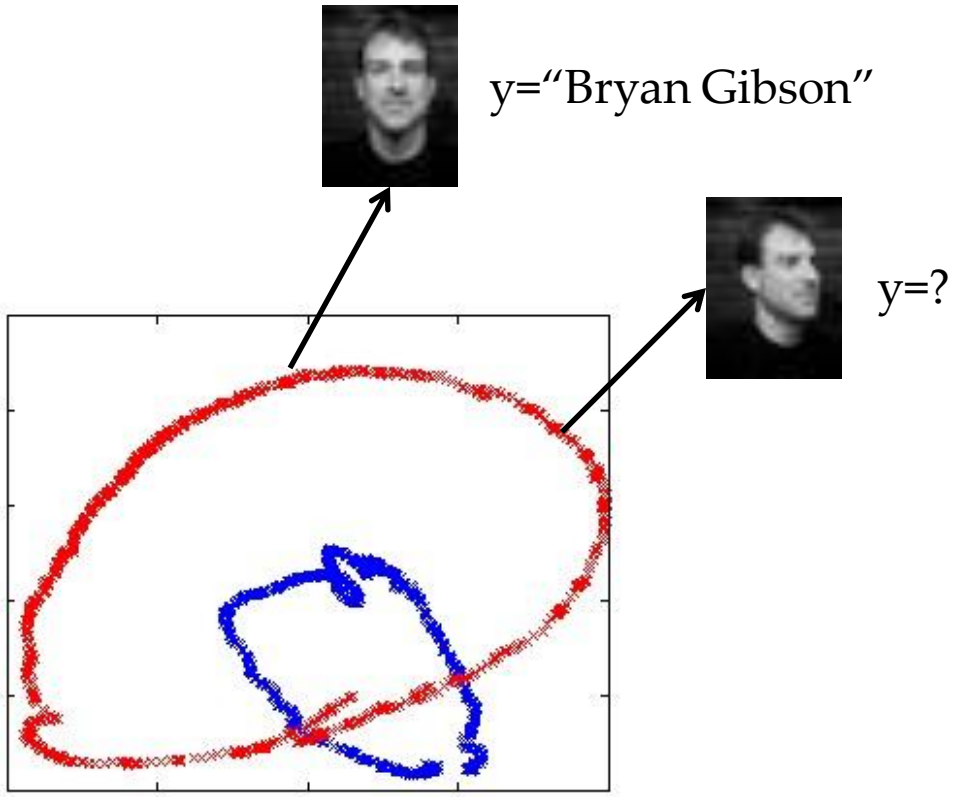

The Stimuli of Our Behavioral Experiments

- Not faces: avoid existing knowledge

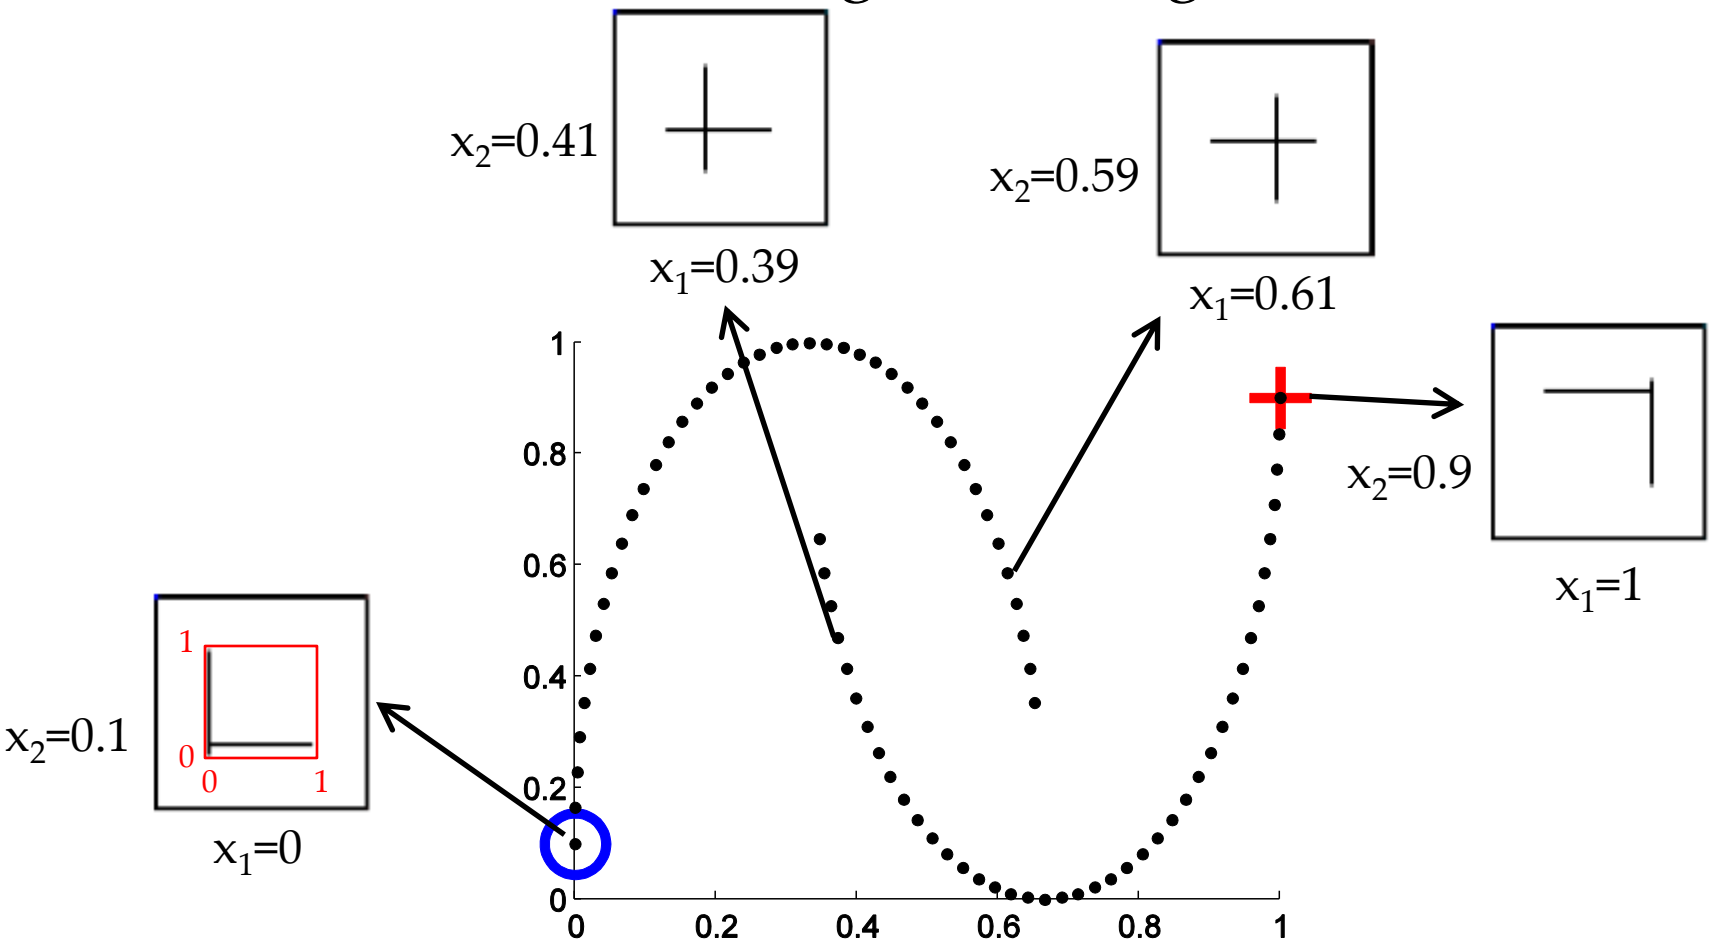

Human Behavioral Experiments

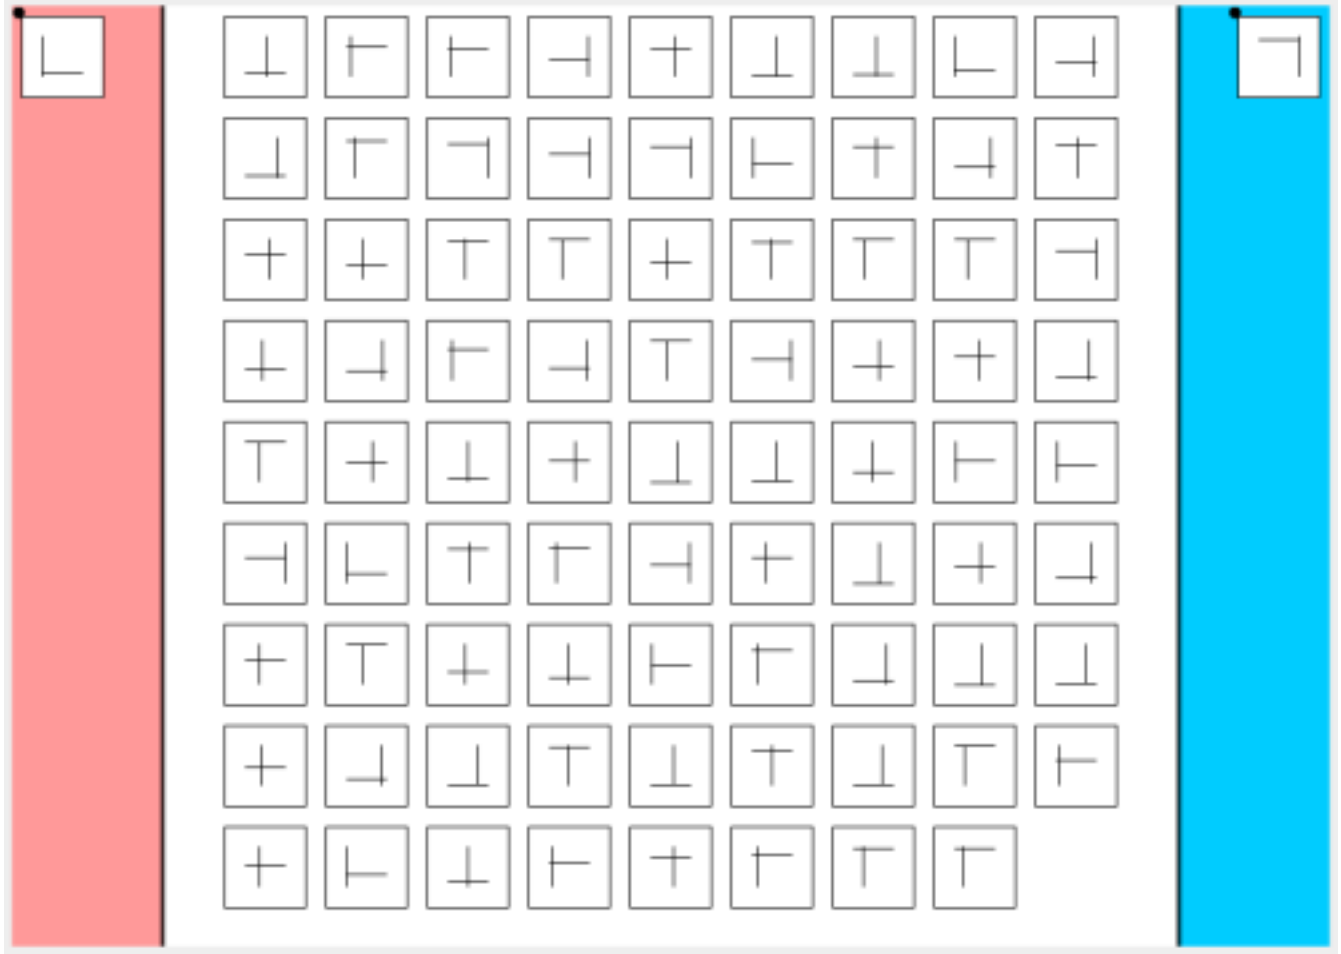

Human Behavioral Experiments

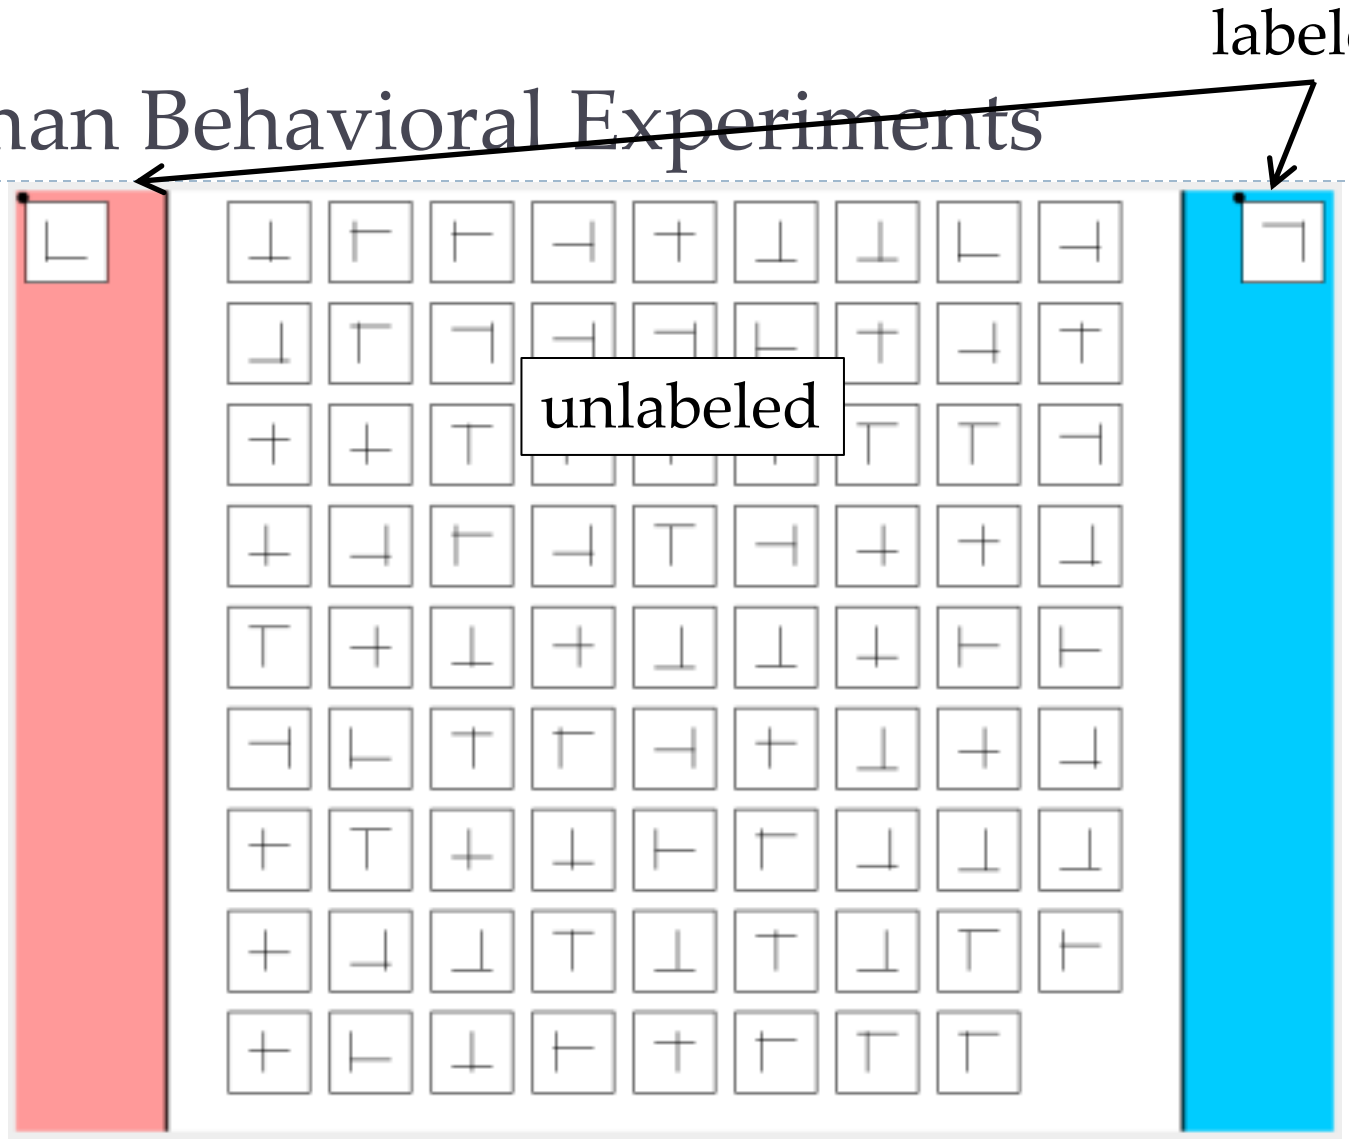

- Batch learning (instead of sequential)

Aggregated Results: Majority Vote

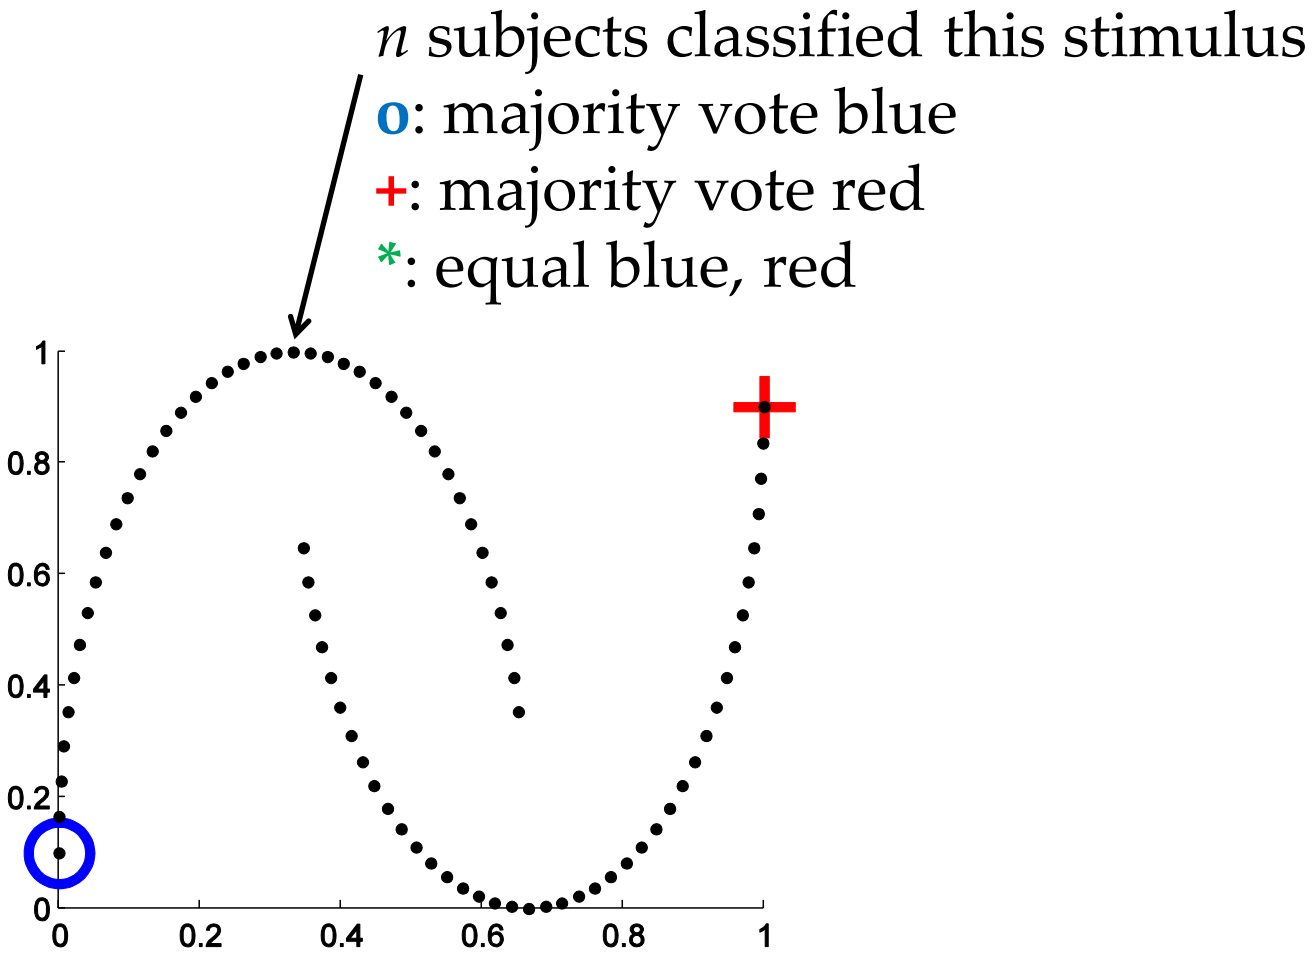

Aggregated Results: Majority Vote

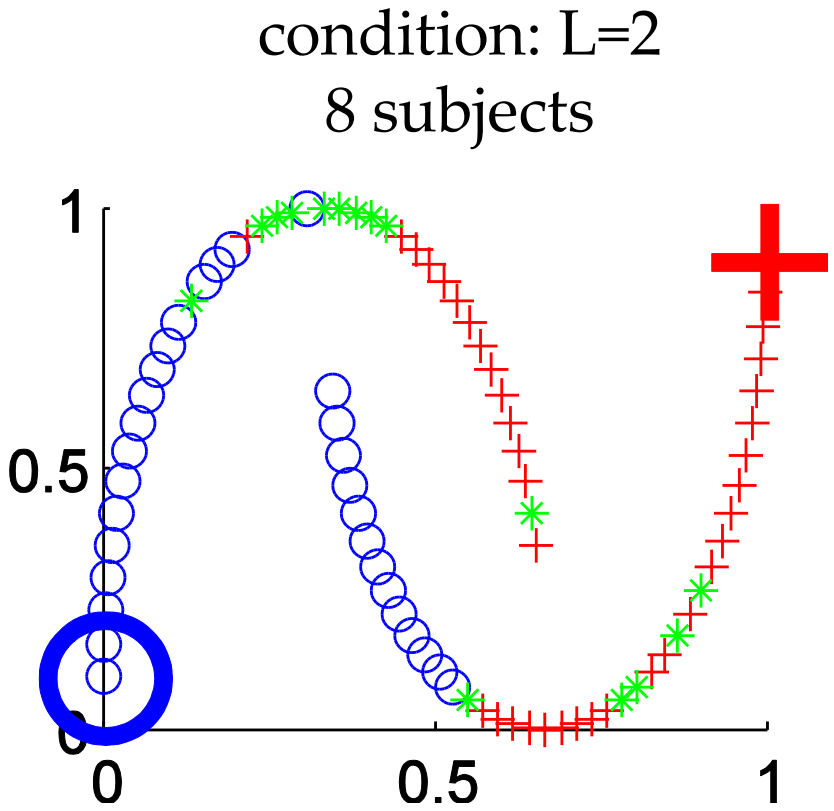

Individual Subject Fit

graph [Belkin et al.'06] | 1NN ℓ_2 [Nosofsky'86] | 1NN ℓ_1 | vertical | horizontal

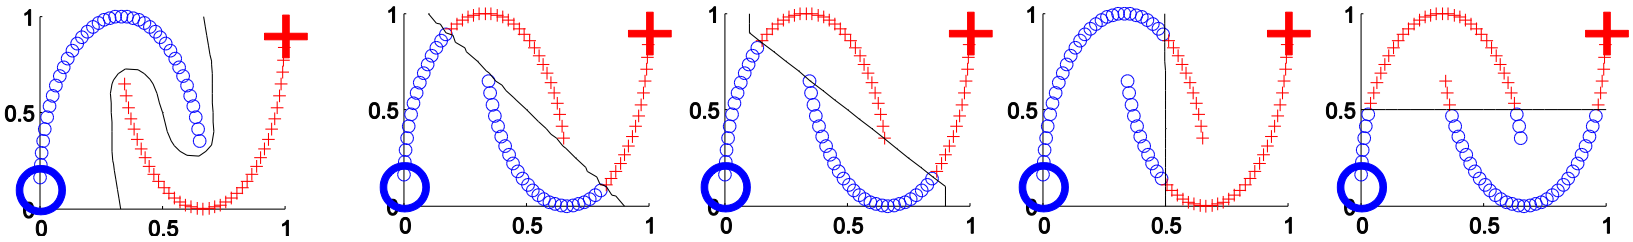

- If best fitting model *m* has accuracy >75%
  - then subject potentially uses *m*,
  - Otherwise subject uses "other"
- Percentage of subjects potentially using each model:

	graph	1NN ℓ_2	1NN ℓ_1	v	h	other
L=2	0	0.12	0	**0.5**	0.25	0.12

Individual Subject Fit

graph [Belkin et al.'06] | 1NN ℓ_2 [Nosofsky'86] | 1NN ℓ_1 | vertical | horizontal

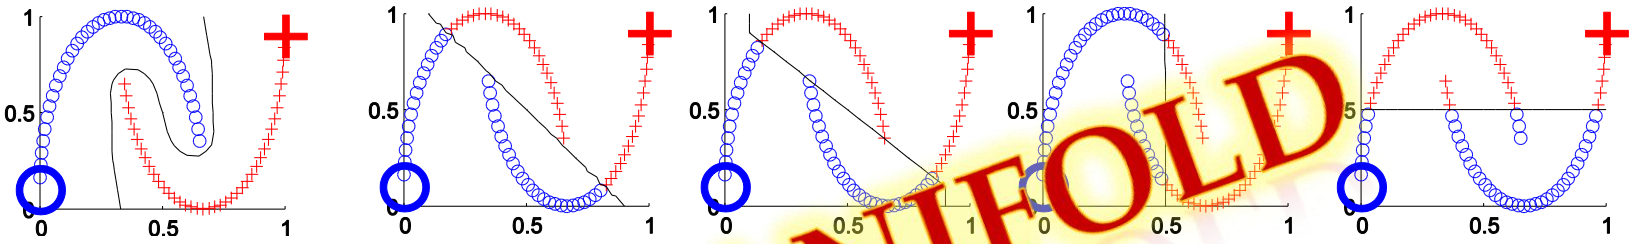

NOT MANIFOLD

- If best fitting model m has accuracy >75%
  - then subject potentially uses m,
  - Otherwise subject uses "other"
- Percentage of subjects potentially using each model:

	graph	1NN ℓ_2	1NN ℓ_1	v	h	other
L=2	0	0.12	0	**0.5**	0.25	0.12

New Experiment: 4 Labeled Points

- Intention: remove "vertical" or "horizontal"
- L=4:

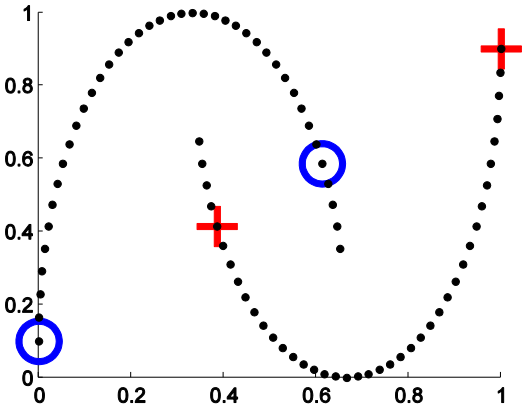

- These 4 labeled points chosen to maximize differences in model predictions:

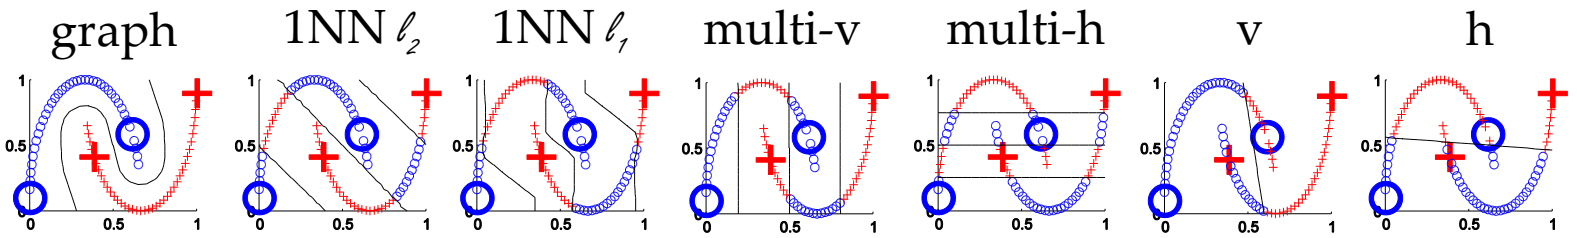

The Interface for L=4

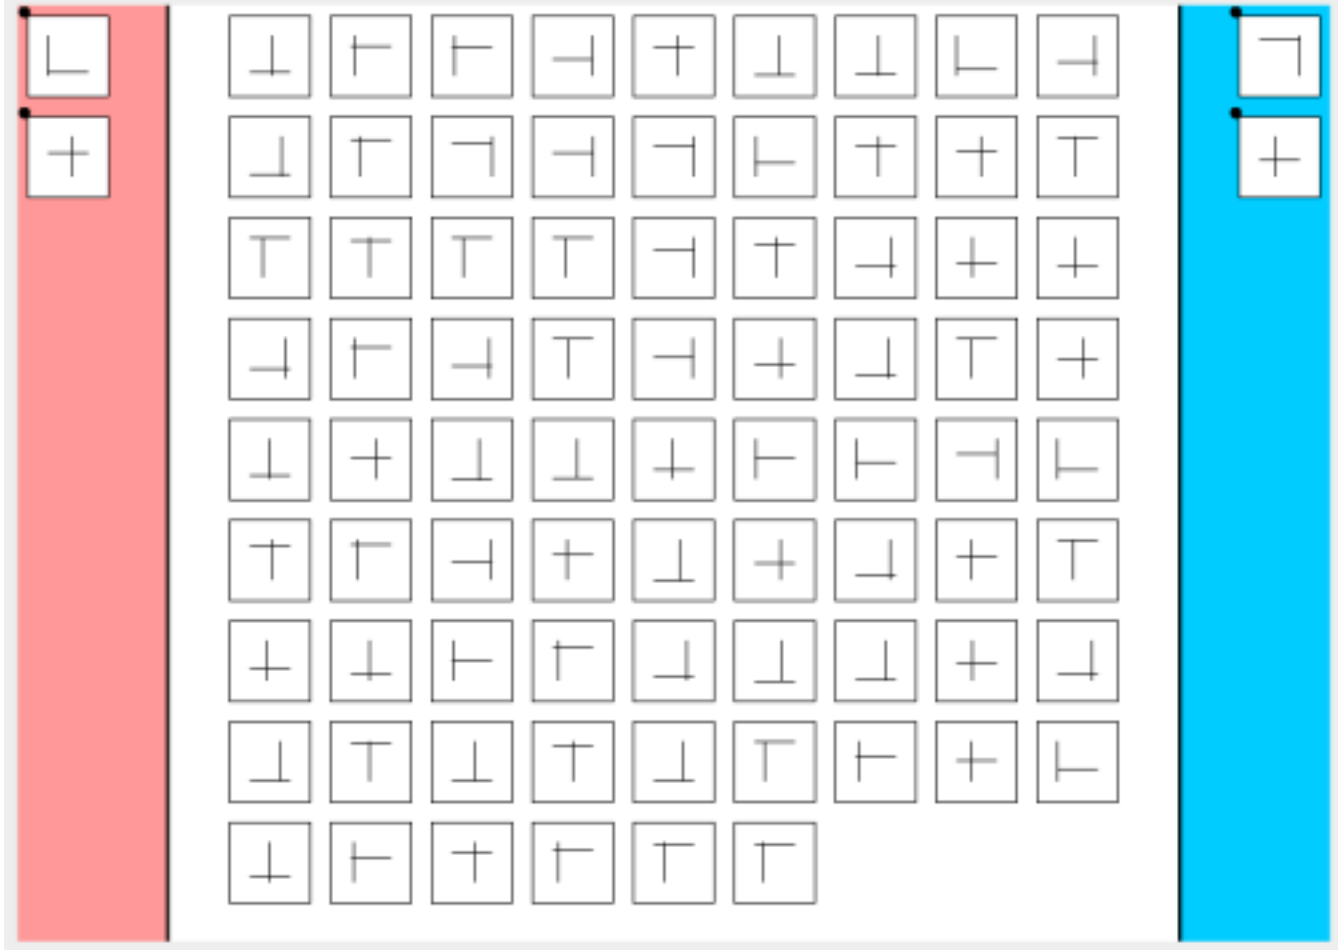

Results for L=4

- Majority vote:

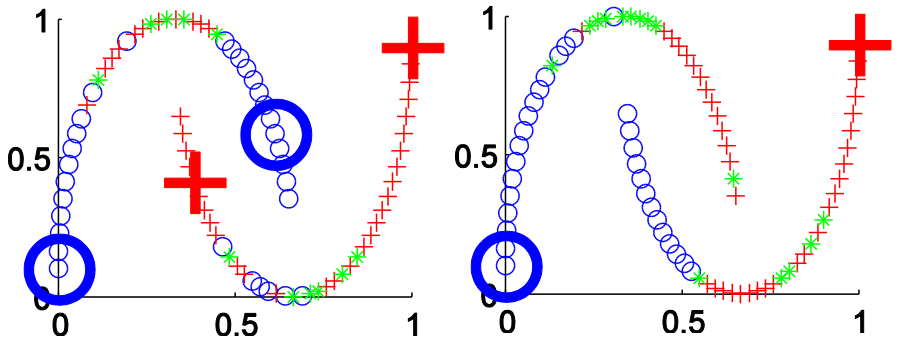

L=4, 24 subjects L=2, 8 subjects

- Individual fit:

	graph	1NN ℓ_2	1NN ℓ_1	multi-v	multi-h	v	h	other
L=4	0.25	0.25	0.12	0.12	0	0.04	0.08	**0.38**
L=2	0	0.12	0	-	-	**0.5**	0.25	0.12

Results for L=4

- Majority vote:

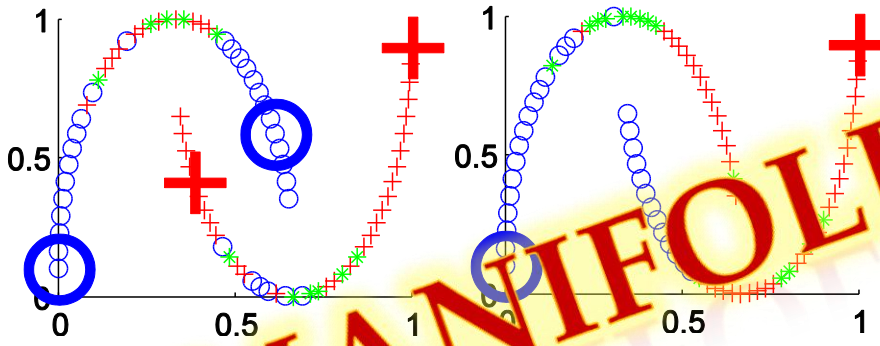

L=4, 2 subjects L=2, 8 subjects

- Individual fit:

	graph	1NN ℓ_2	1NN ℓ_1	multi-v	multi-h	v	h	other
L=4	0.25	0.25	0.12	0.12	0	0.04	0.08	**0.38**
L=2	0	0.12	0	-	-	**0.5**	0.25	0.12

New Experiment: Give People the Graph!

- εNN graph

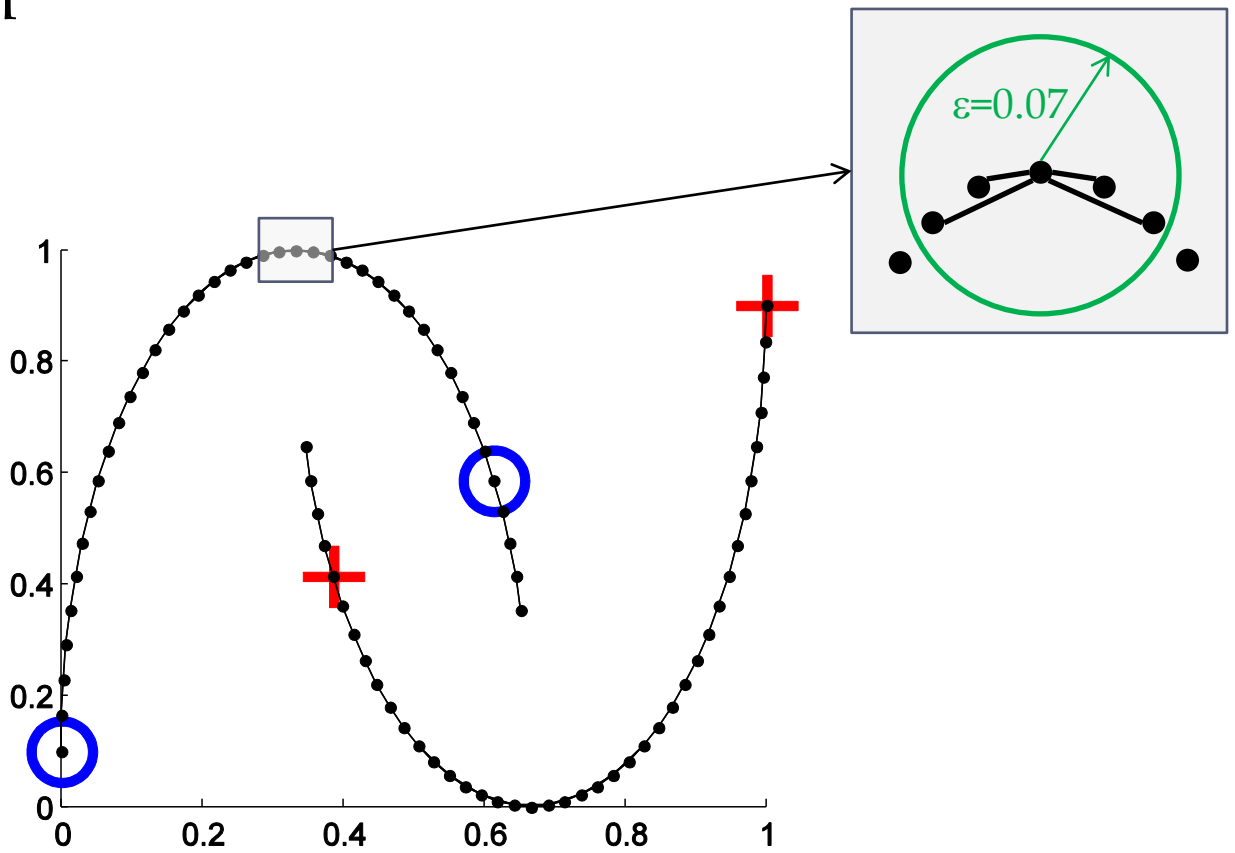

- 2 connected components with pure labels

Giving People the Graph via Highlighting

- Clicking on any item highlights all its neighbors

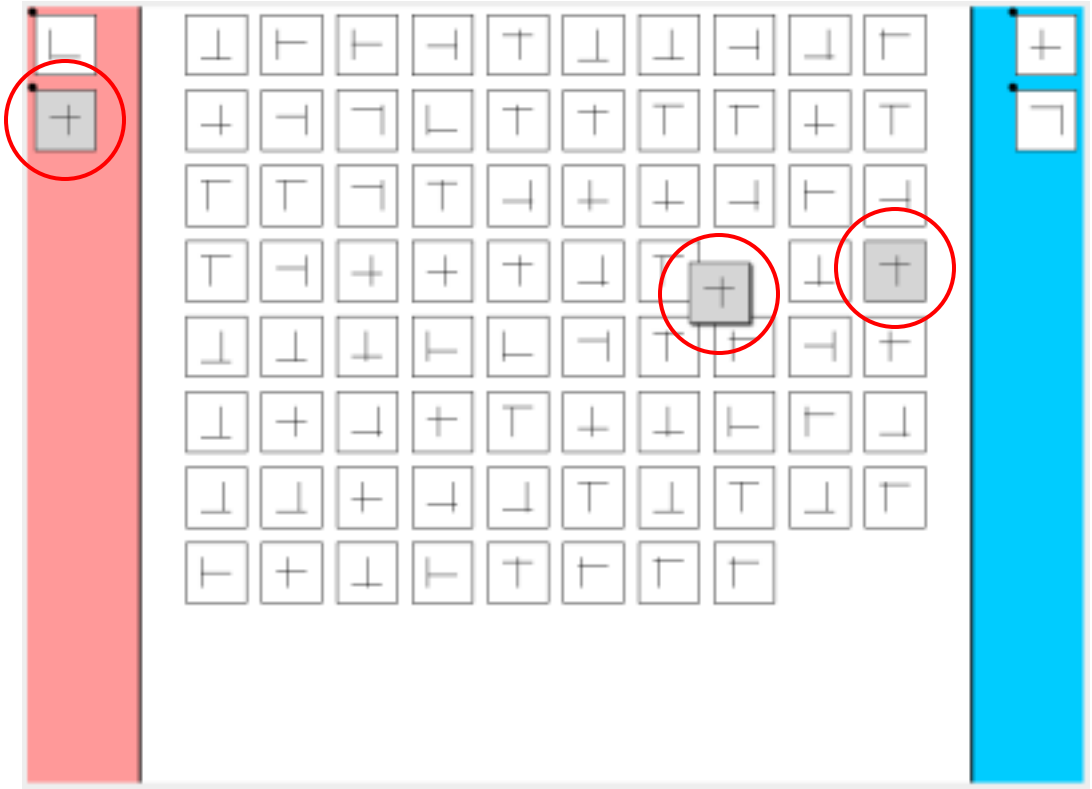

- Instructions to subjects: highlighting not necessarily mean "same class"

Results for (L=4, Hi)

- Majority vote:

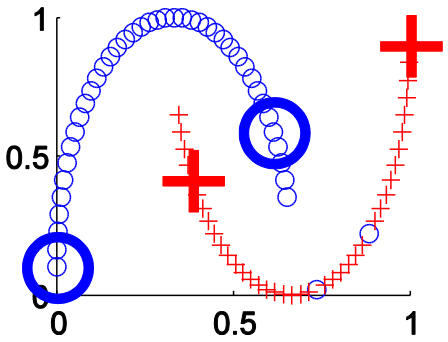

(L=4, Hi) 23 sub

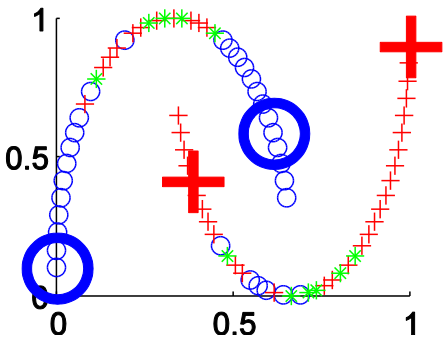

L=4, 24 subjects

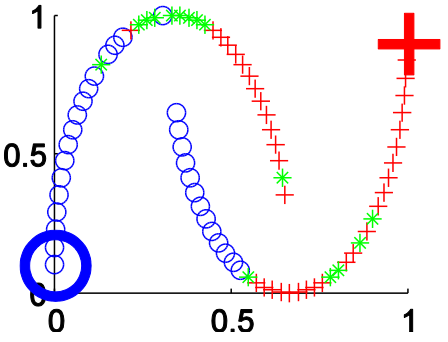

L=2, 8 subjects

- Individual fit:

	graph	1NN ℓ_2	1NN ℓ_1	multi-v	multi-h	v	h	other
L=4 Hi	**0.39**	0.09	0.09	0.04	0.04	0	0.13	0.22
L=4	0.25	0.25	0.12	0.12	0	0.04	0.08	**0.38**
L=2	0	0.12	0	**-**	**-**	**0.5**	0.25	0.12

Results for (L=4, Hi)

- Majority vote:

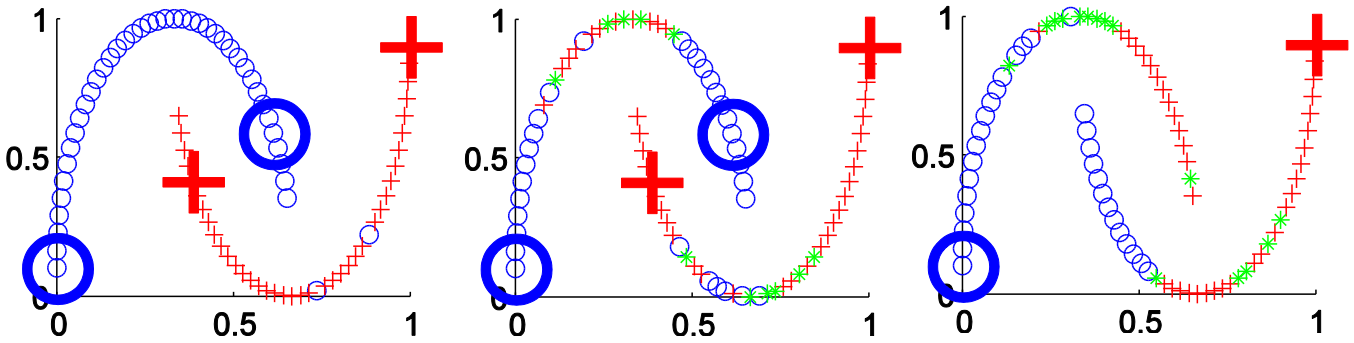

(L=4, Hi) 22 subjects L=4 24 subjects L=2 8 subjects

- Individual

MANIFOLD!

	graph	1NN ℓ_2	1NN ℓ_1	multi-v	multi-h	v	h	other
L=4 Hi	**0.39**	0.09	0.09	0.04	0.04	0	0.13	0.22
L=4	0.25	0.25	0.12	0.12	0	0.04	0.08	**0.38**
L=2	0	0.12	0	-	-	**0.5**	0.25	0.12

This is Awfully Suspicious

- Humans simply follow highlighting?!

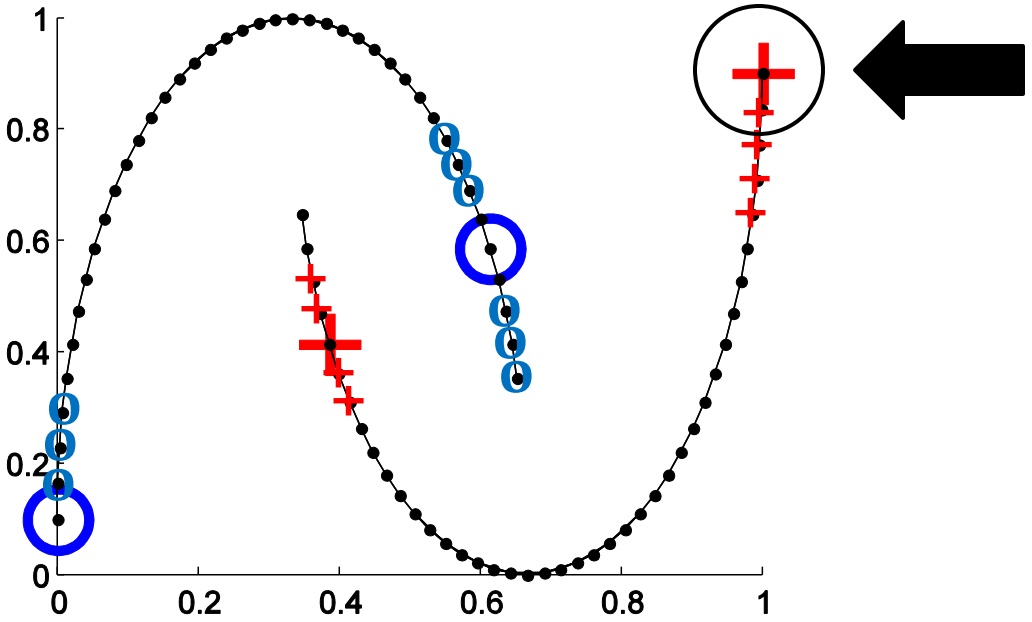

- No. Three pieces of evidence

Evidence 1: Leap-of-Faith Against Highlight

- A leap-of-faith (lof) move goes against labeled neighbors, or has no labeled neighbors

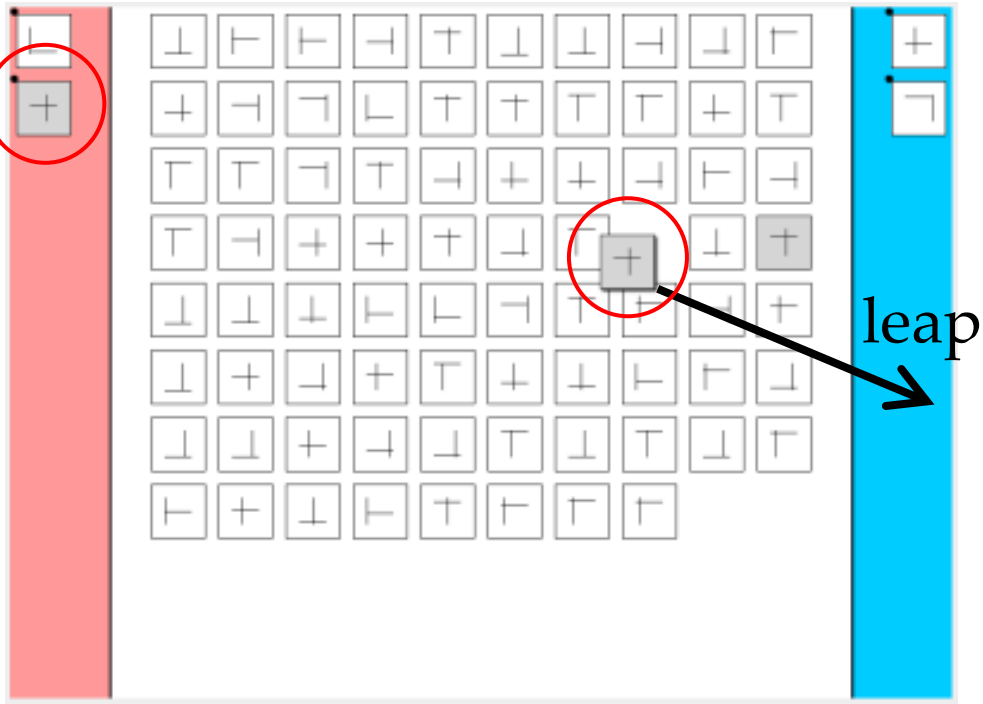

- If people simply follow highlighting, #lof=0
- In (L=4, Hi), average #lof=17 (20% of their moves)

Evidence 2: New Experiment (L=2, Hi)

- Same highlighting, but 2 labeled points
- Majority vote:

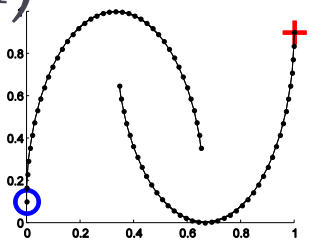

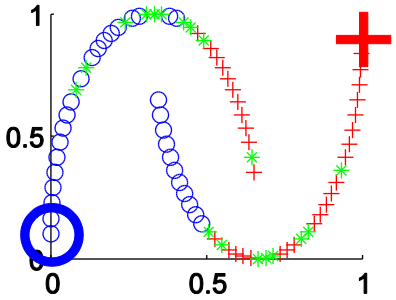

(L=2, Hi) 8 sub

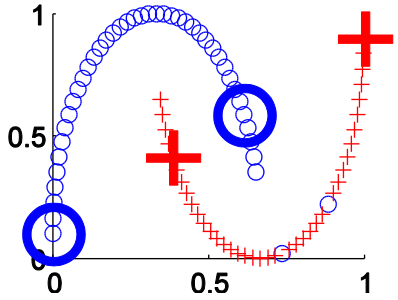

(L=4, Hi) 23 sub

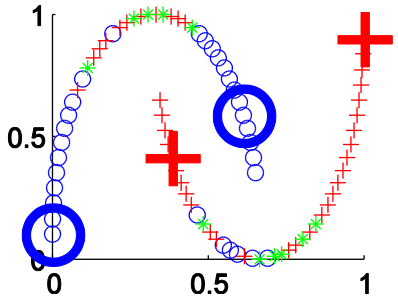

L=4, 24 subjects

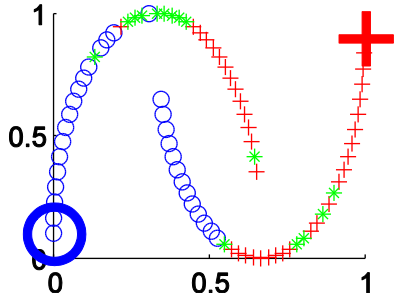

L=2, 8 subjects

- Individual fit:

	graph	1NN ℓ_2	1NN ℓ_1	multi-v	multi-h	v	h	other
L=2 Hi	0.12	0	0	-	-	**0.38**	0.25	0.25
L=4 Hi	**0.39**	0.09	0.09	0.04	0.04	0	0.13	0.22
L=4	0.25	0.25	0.12	0.12	0	0.04	0.08	**0.38**
L=2	0	0.12	0	**-**	**-**	**0.5**	0.25	0.12

Evidence 2: New Experiment (L=2, Hi)

- Same highlighting, but 2 labeled points
- Majority vote:

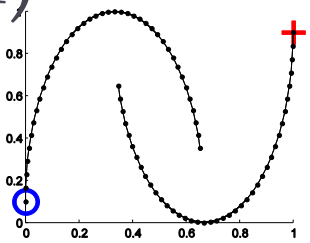

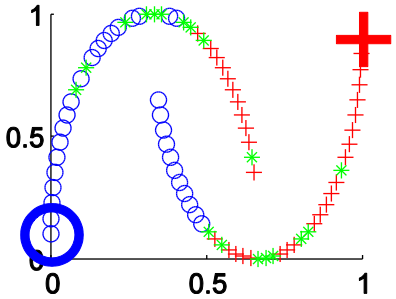

(L=2, Hi) 8 sub

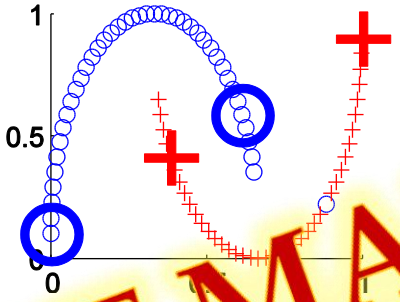

(L=4, Hi) 23 sub

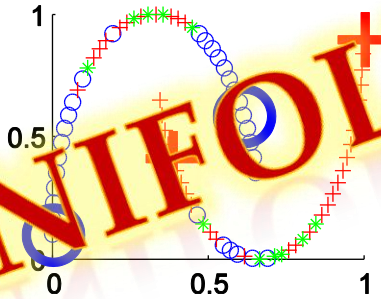

L=4, 24 subjects

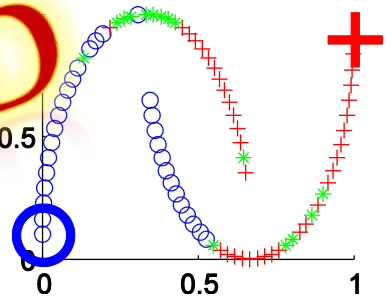

L=2, 8 subjects

NOT MANIFOLD

- Individual fit:

	graph	1NN ℓ_2	1NN ℓ_1	multi-v	multi-h	v	h	other
L=2 Hi	0.12	0	0	-	-	**0.38**	0.25	0.25
L=4 Hi	**0.39**	0.09	0.09	0.04	0.04	0	0.13	0.22
L=4	0.25	0.25	0.12	0.12	0	0.04	0.08	**0.38**
L=2	0	0.12	0	**-**	**-**	**0.5**	0.25	0.12

Evidence 3: New Experiment (Isomorphic Graph)

- Randomly permute the unlabeled nodes, "bring the edges with them"

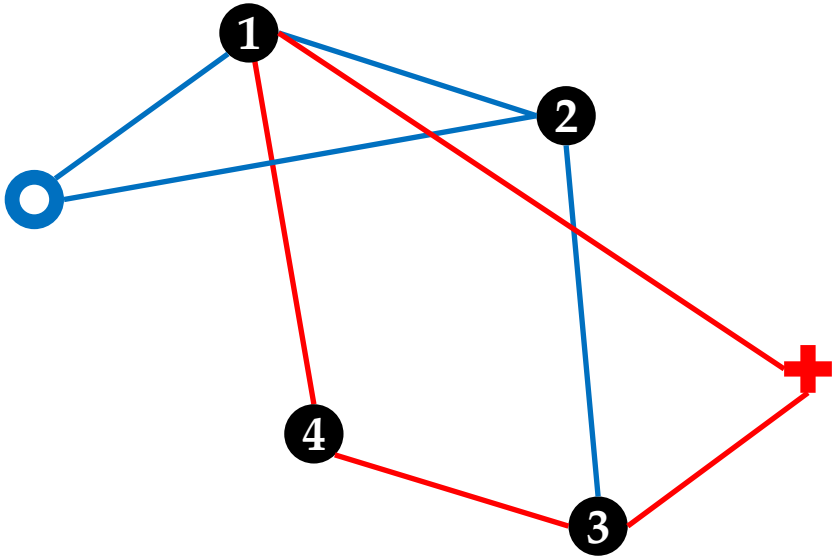

- If people simply follow highlighting, then everything should be the same: 1 is easy blue, 2 is harder but still blue, …

Evidence 3: New Experiment (Isomorphic Graph)

- (L=4, Iso) isomorphic to (L=4, Hi):

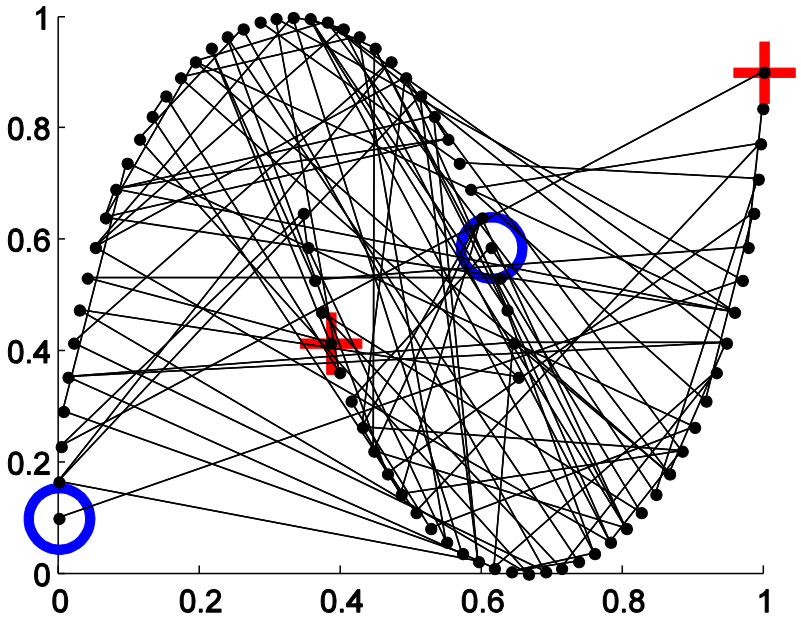

- There are really two "pure" connected components!

Evidence 3: New Experiment (Isomorphic Graph)

- The interface for (L=4, Iso)

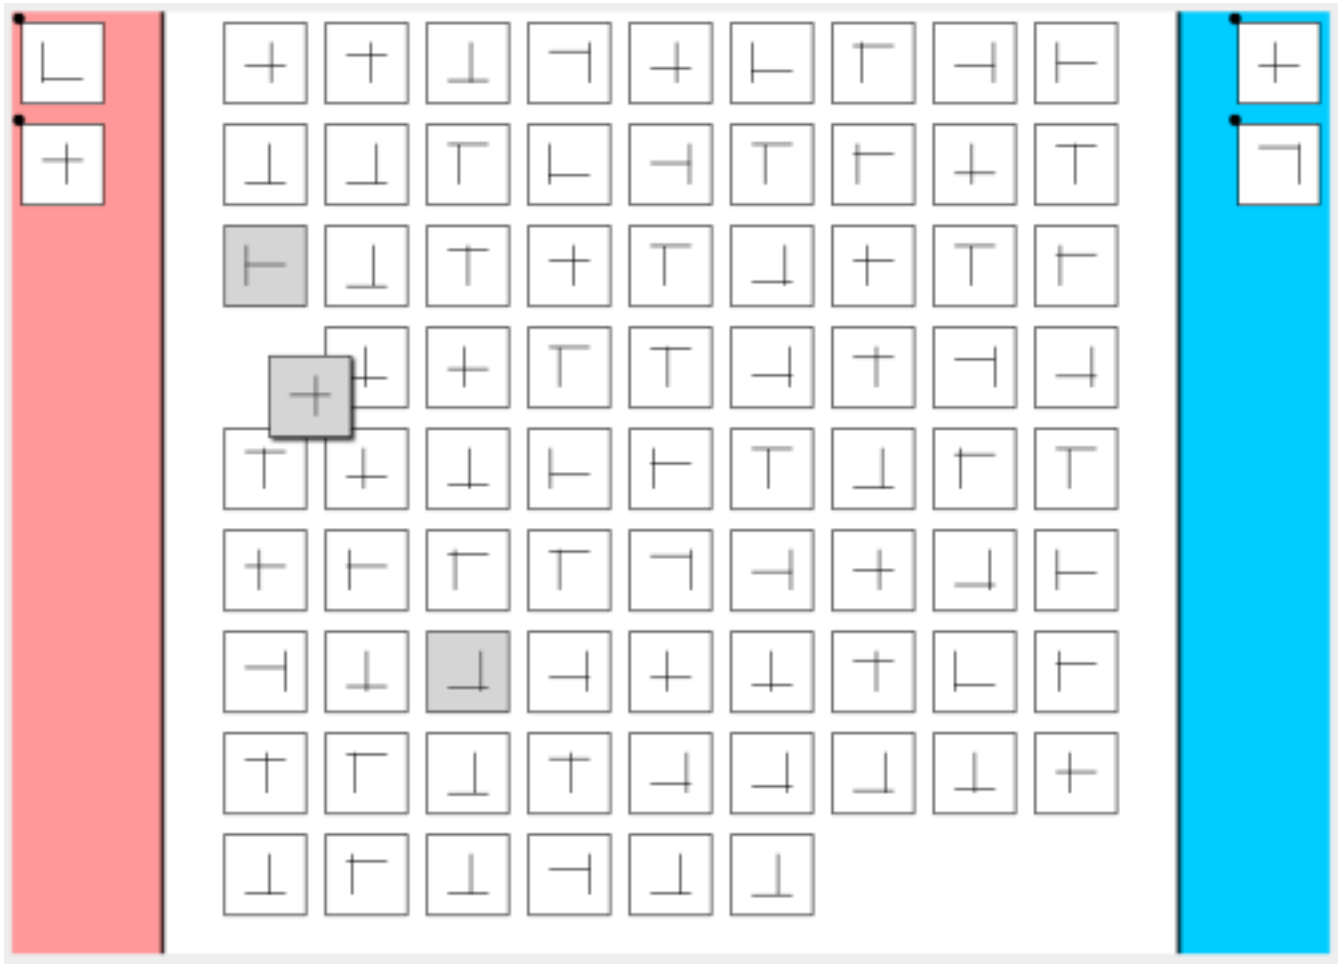

Evidence 3: New Experiment (Isomorphic Graph)

L=4, Hi (23 subjects)

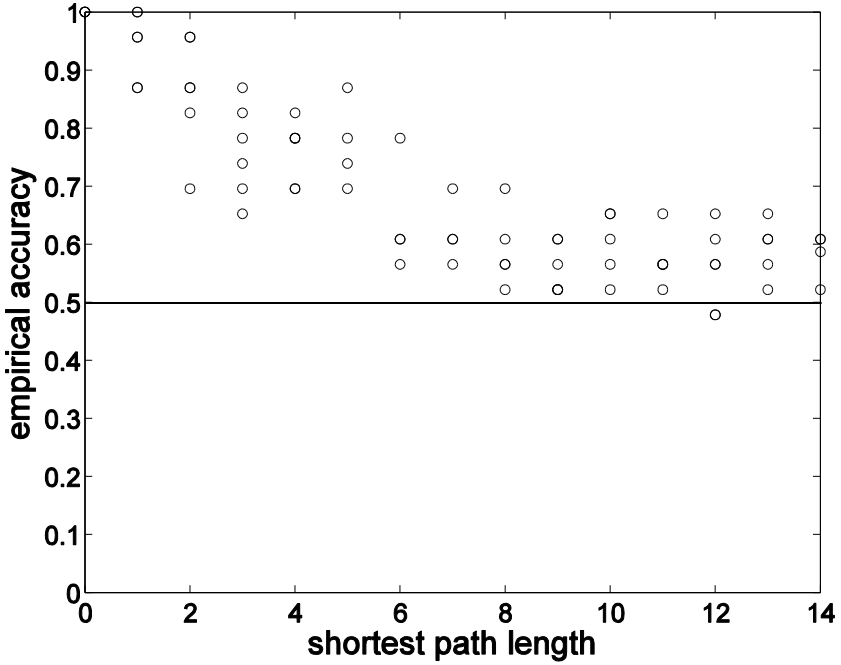

L=4, Iso (30 subjects)

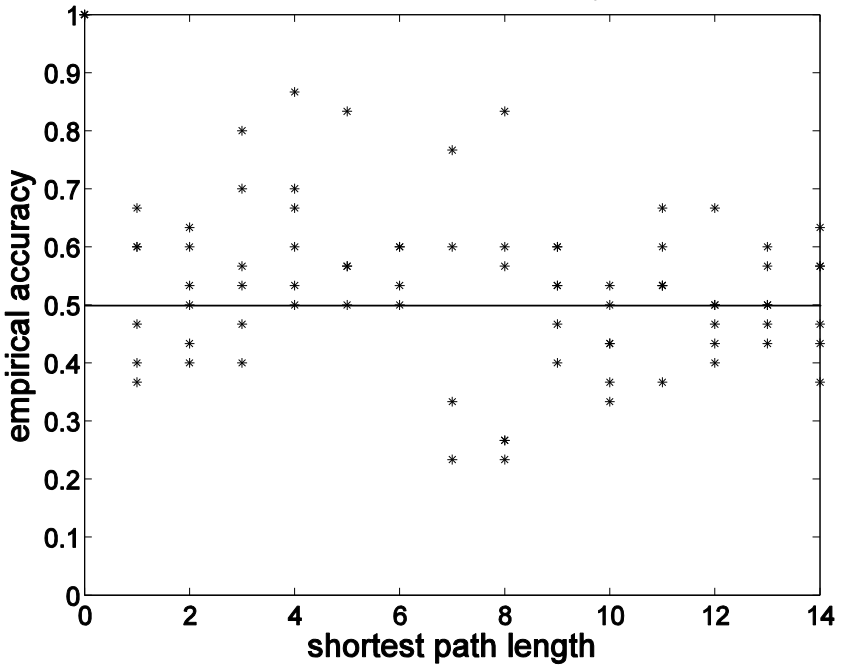

- Highlighting ignored if it contradicts with similarity

Why not Always Manifold Learning?

highlighting:	none	manifold	isomorphic random
L=2	✘	✘	
L=4	✘	✓	✘ (neither graphs)

- Possible explanation: Bayesian model selection

$$m^* = \underset{m}{\operatorname{argmax}}\, p(m \mid X, Y_L, H) \propto p(m \mid X, H) p(Y_L \mid m, X)$$

  - data-dependent prior:

$$p(m = v \mid \epsilon NN(X) \neq H) \gg p(m = g \mid \epsilon NN(X) \neq H)$$
$$p(m = v \mid \epsilon NN(X) = H) \gtrsim p(m = g \mid \epsilon NN(X) = H)$$

  - likelihood:

$$p(Y_{L=2} \mid m = v, X) = p(Y_{L=2} \mid m = g, X)$$
$$p(Y_{L=4} \mid m = v, X) < p(Y_{L=4} \mid m = g, X)$$

Summary

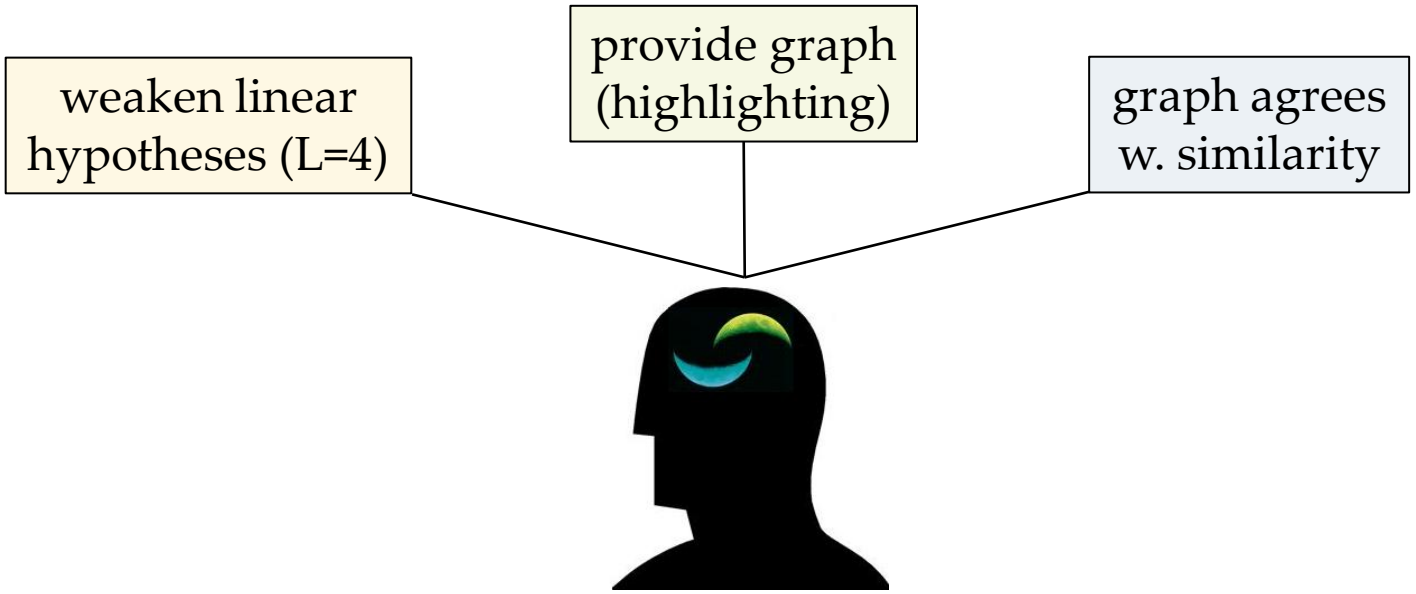
weaken linear hypotheses (L=4)
provide graph (highlighting)
graph agrees w. similarity

Thank you

(L=4, Iso)

- Majority Vote (30 subjects):

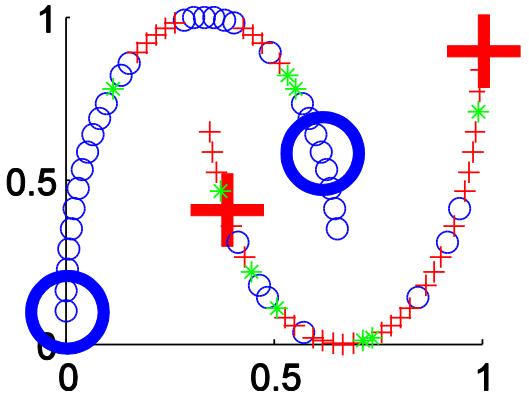

- Individual fit:

	graph	1NN ℓ_2	1NN ℓ_1	multi-v	multi-h	v	h	other
L=4 Iso	0.13	0.03	0.07	0	0	0.07	0.03	**0.67**
L=2 Hi	0.12	0	0	-	-	**0.38**	0.25	0.25
L=4 Hi	**0.39**	0.09	0.09	0.04	0.04	0	0.13	0.22
L=4	0.25	0.25	0.12	0.12	0	0.04	0.08	**0.38**
L=2	0	0.12	0	**-**	**-**	**0.5**	0.25	0.12