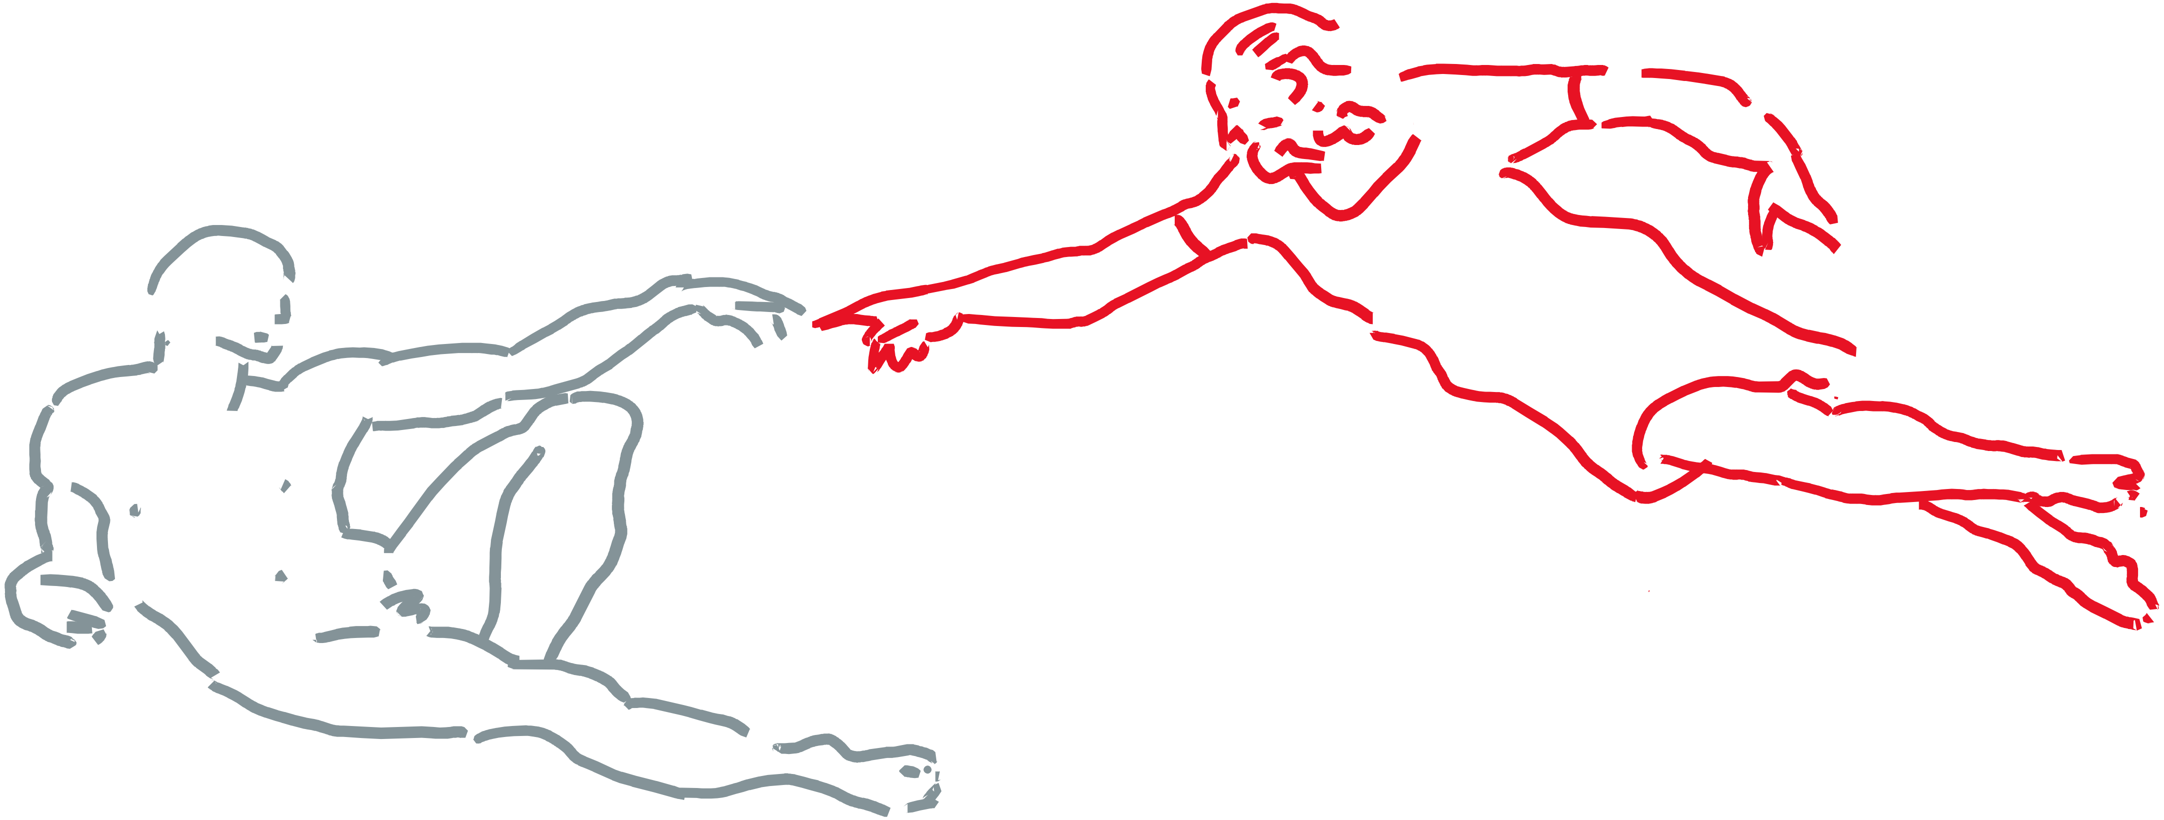

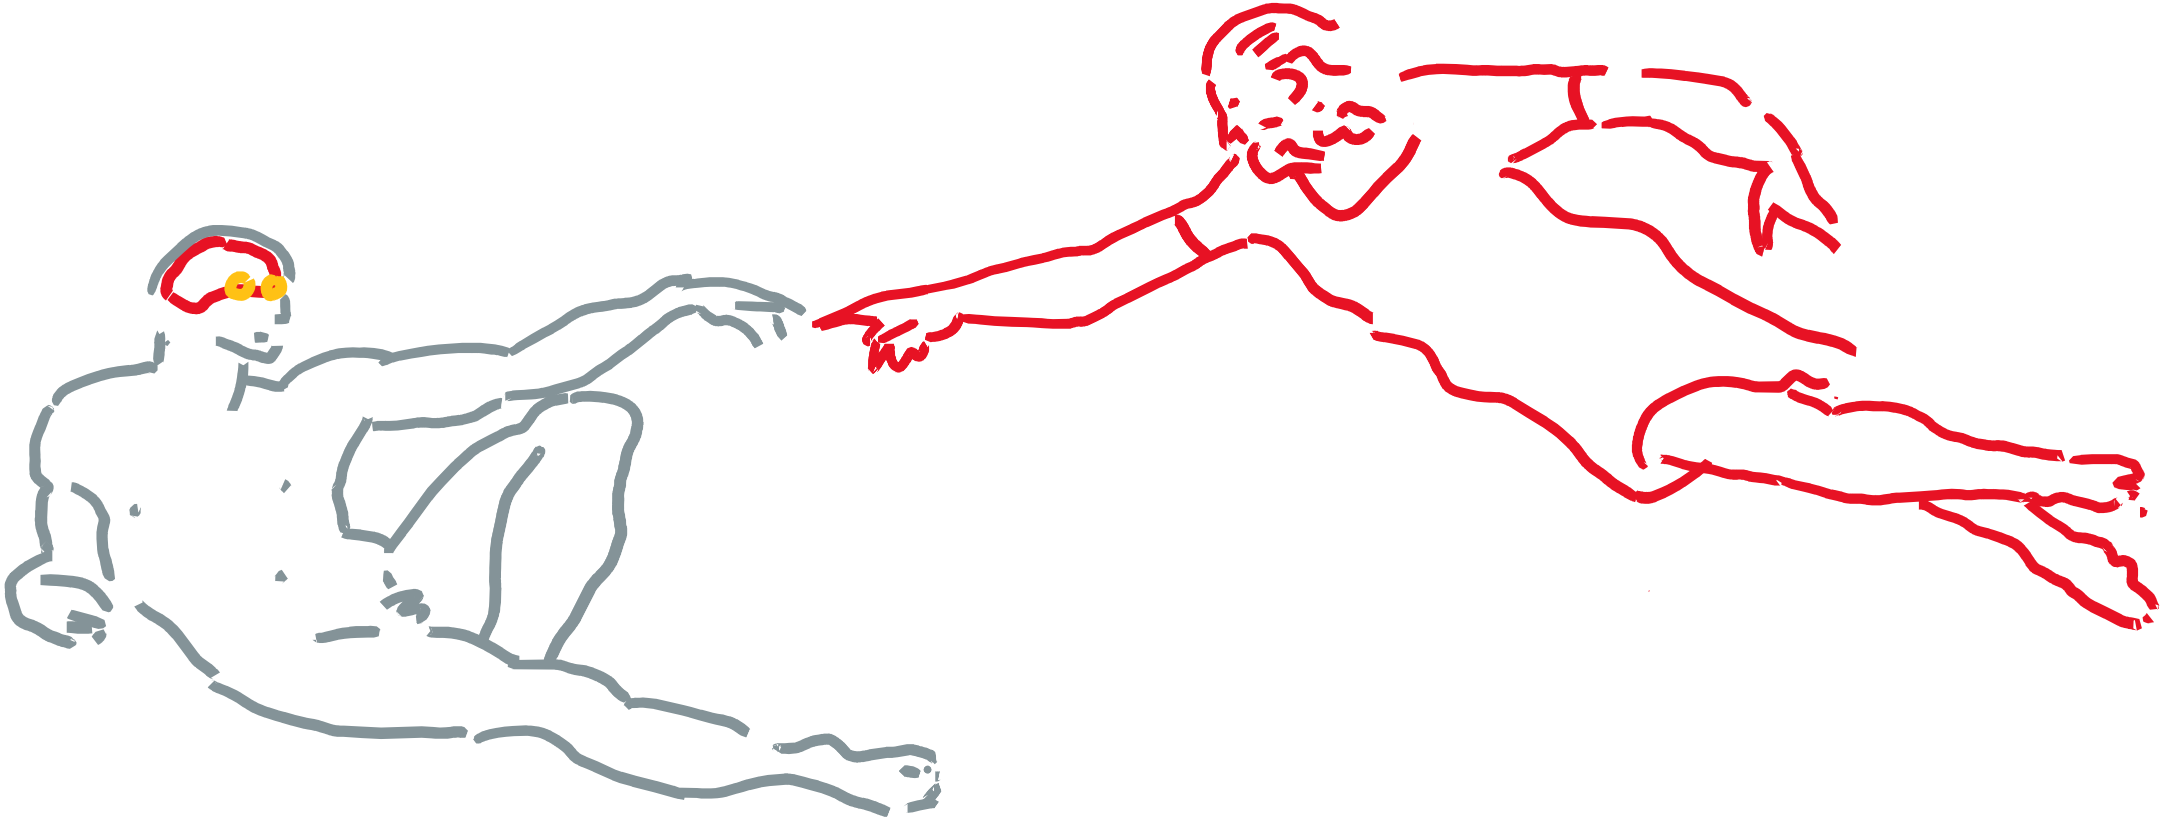

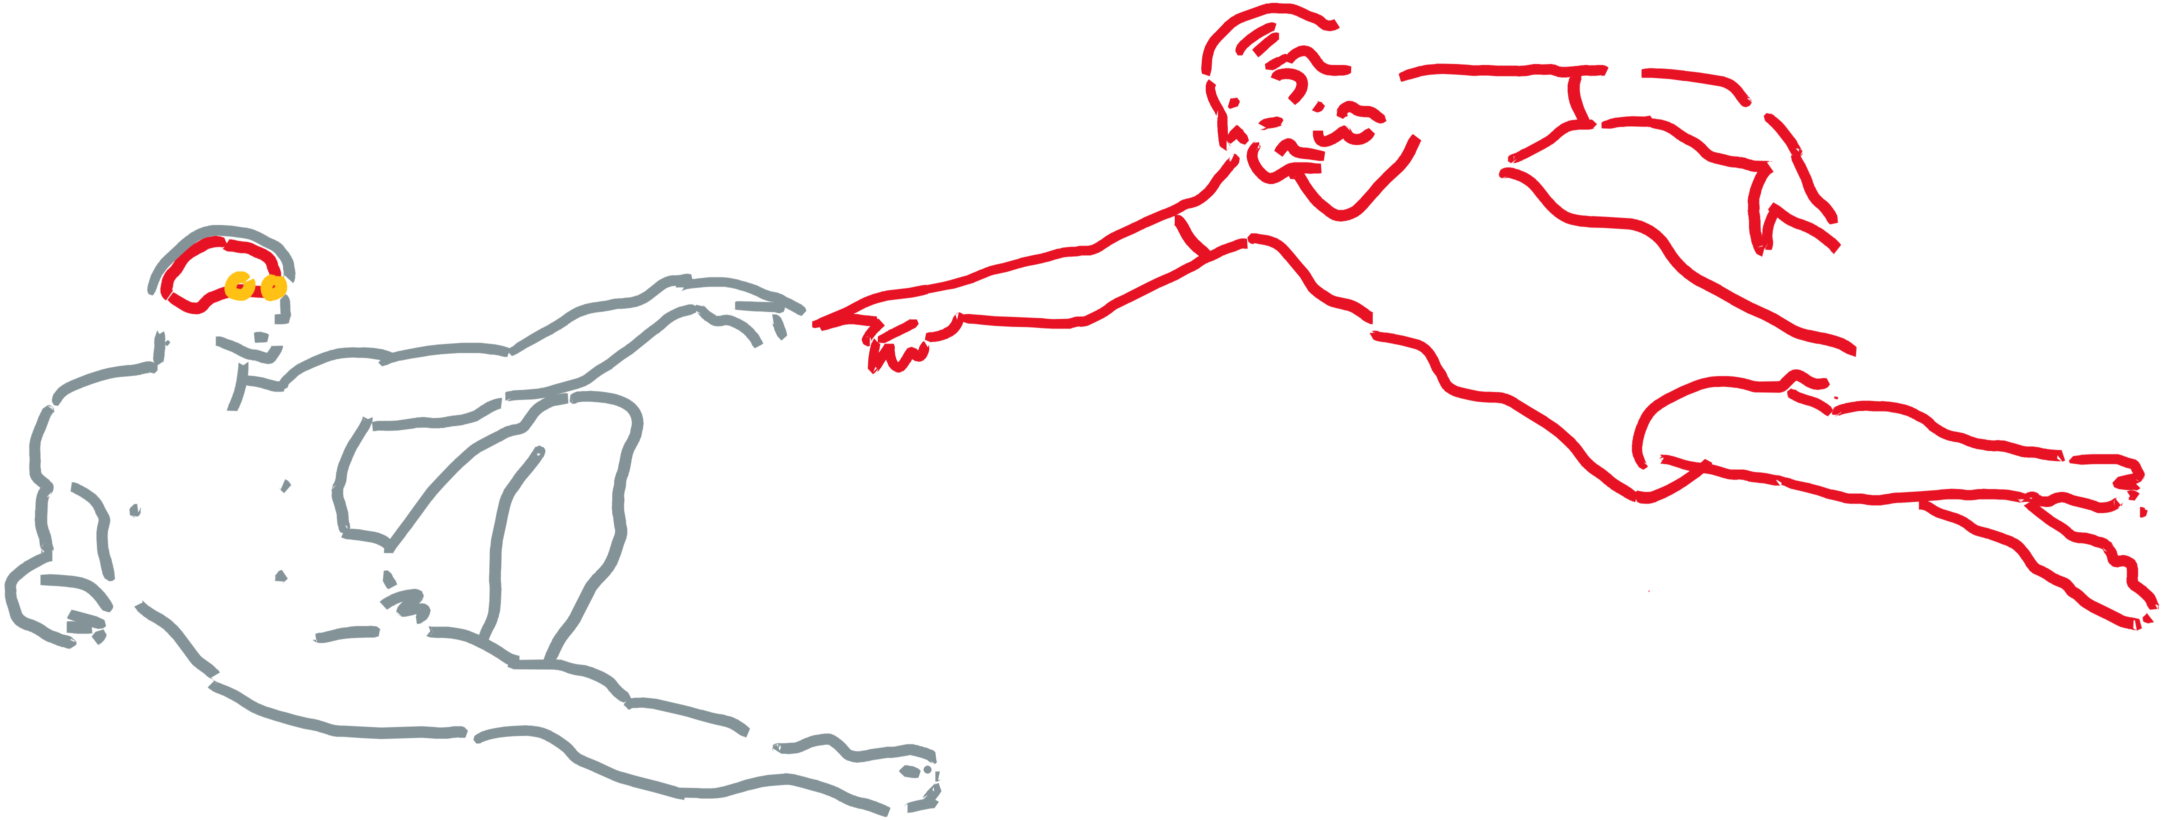

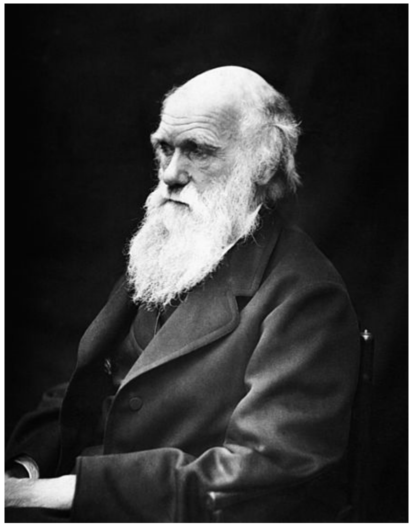

CHARLES DARWIN!

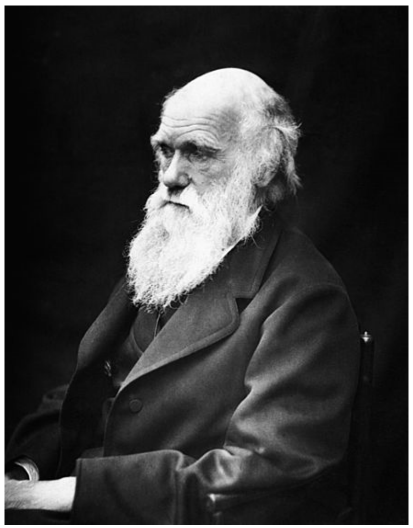

CHARLES DARWIN!

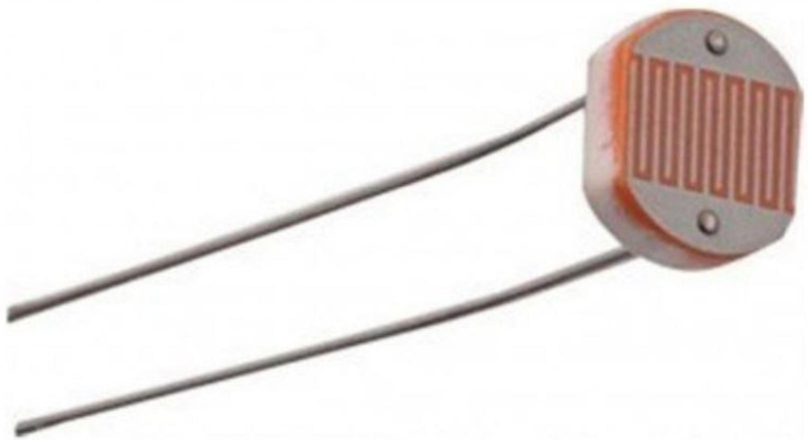

SENSOR FOTORESISTENCIA LDR GL-5528

Fotoresistencia que permite medir niveles de luz.

0,25 €

0,21 € (IVA no incluido)

Fracciona tu pago desde 50,00 €

SEQURA

Estado: Nuevo

Fabricante: tiendatec

Referencia: GL5528

EAN: 8472496014380

Disponible, recíbelo el lunes 4

WHAT WOULD YOU DO ?

WHAT WOULD YOU DO?

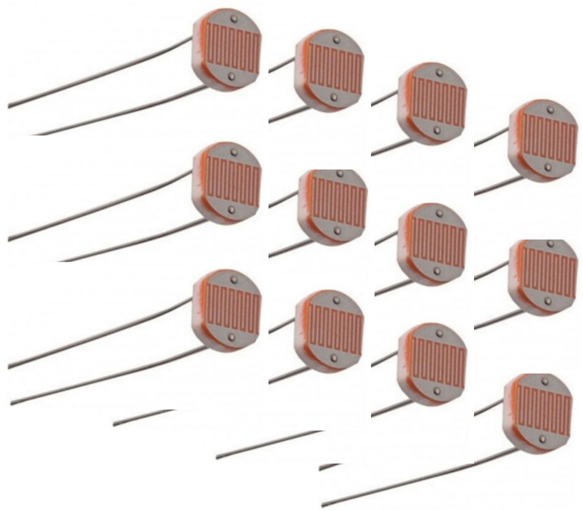

WHAT WOULD YOU DO ?

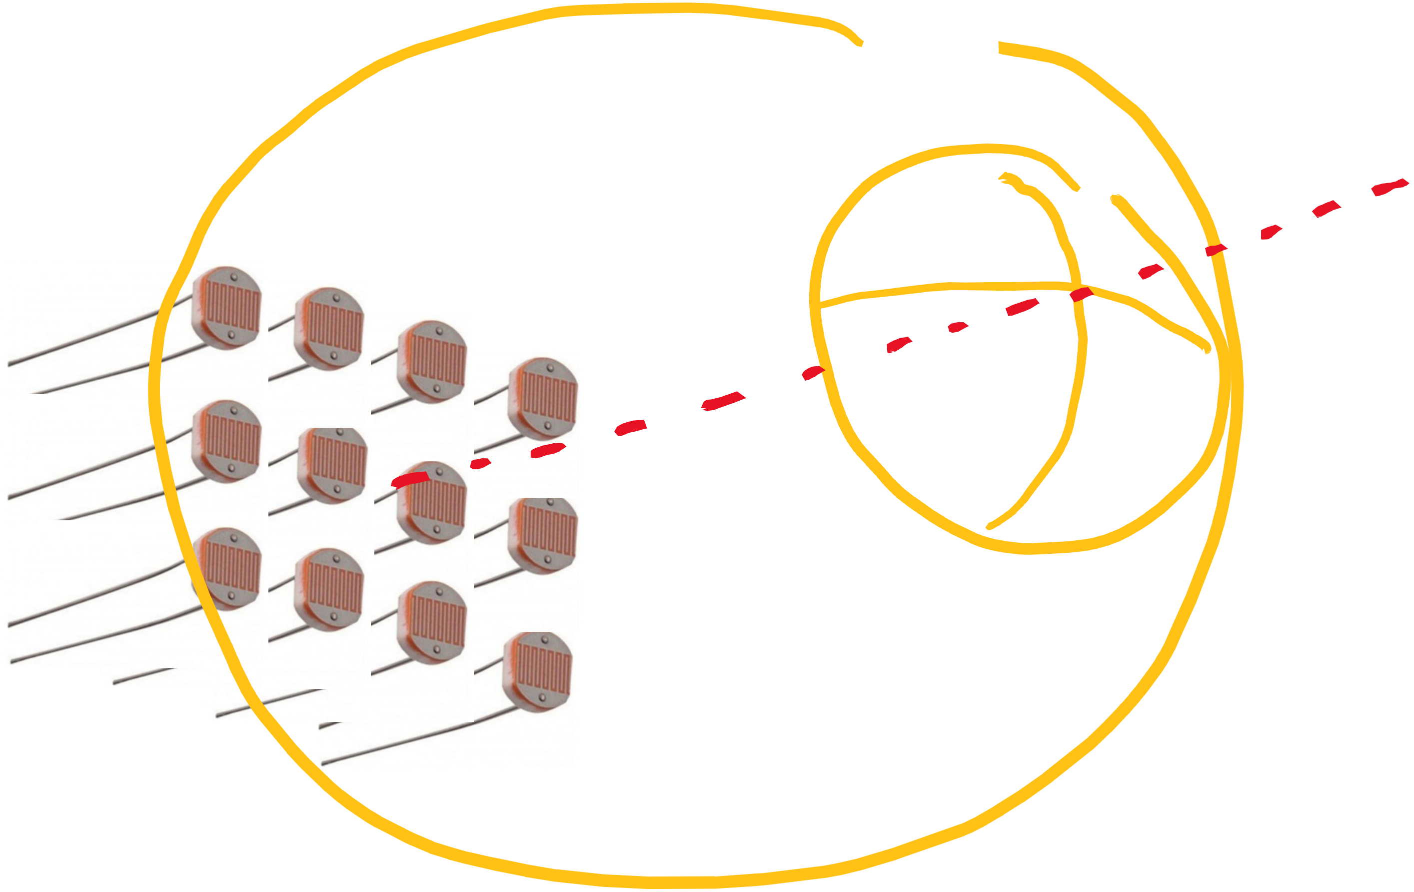

WHAT WOULD YOU DO ?

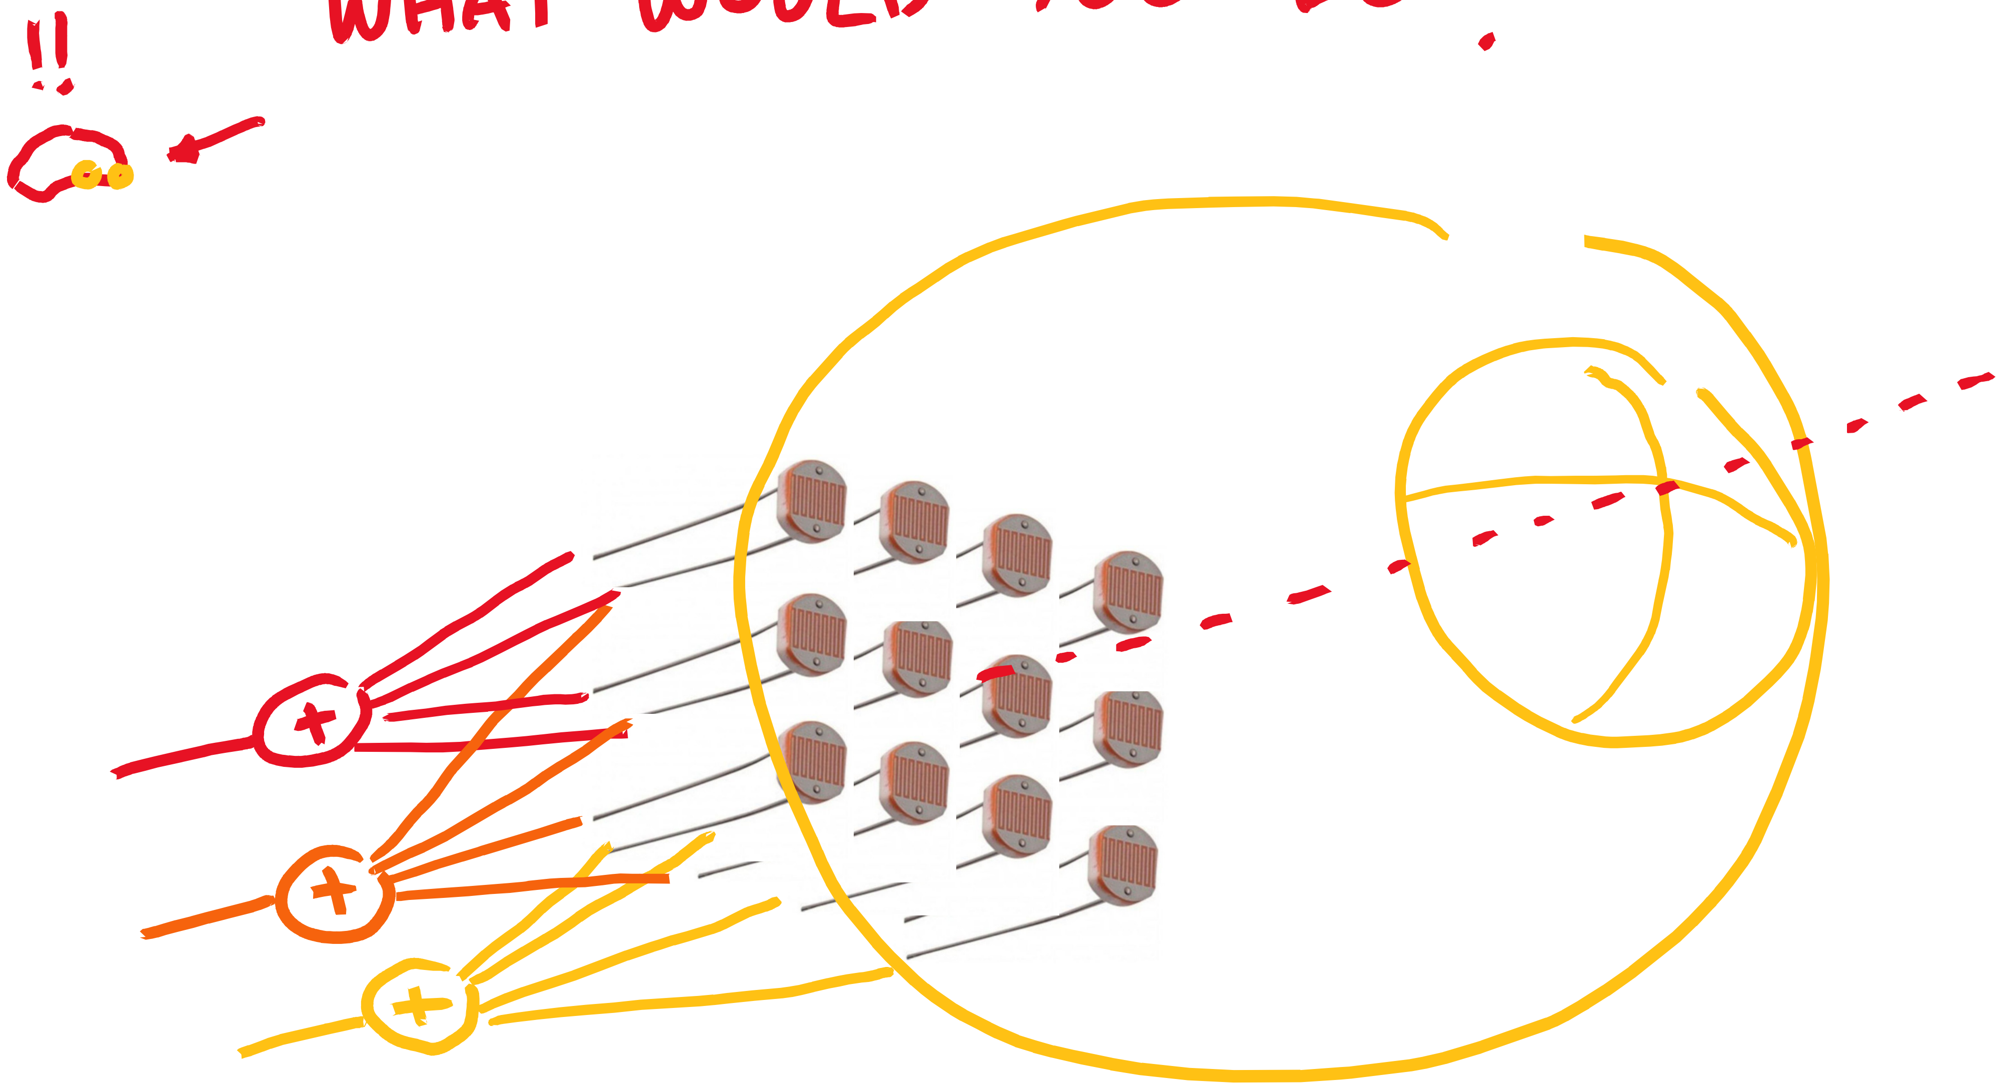

WHAT WOULD YOU DO?

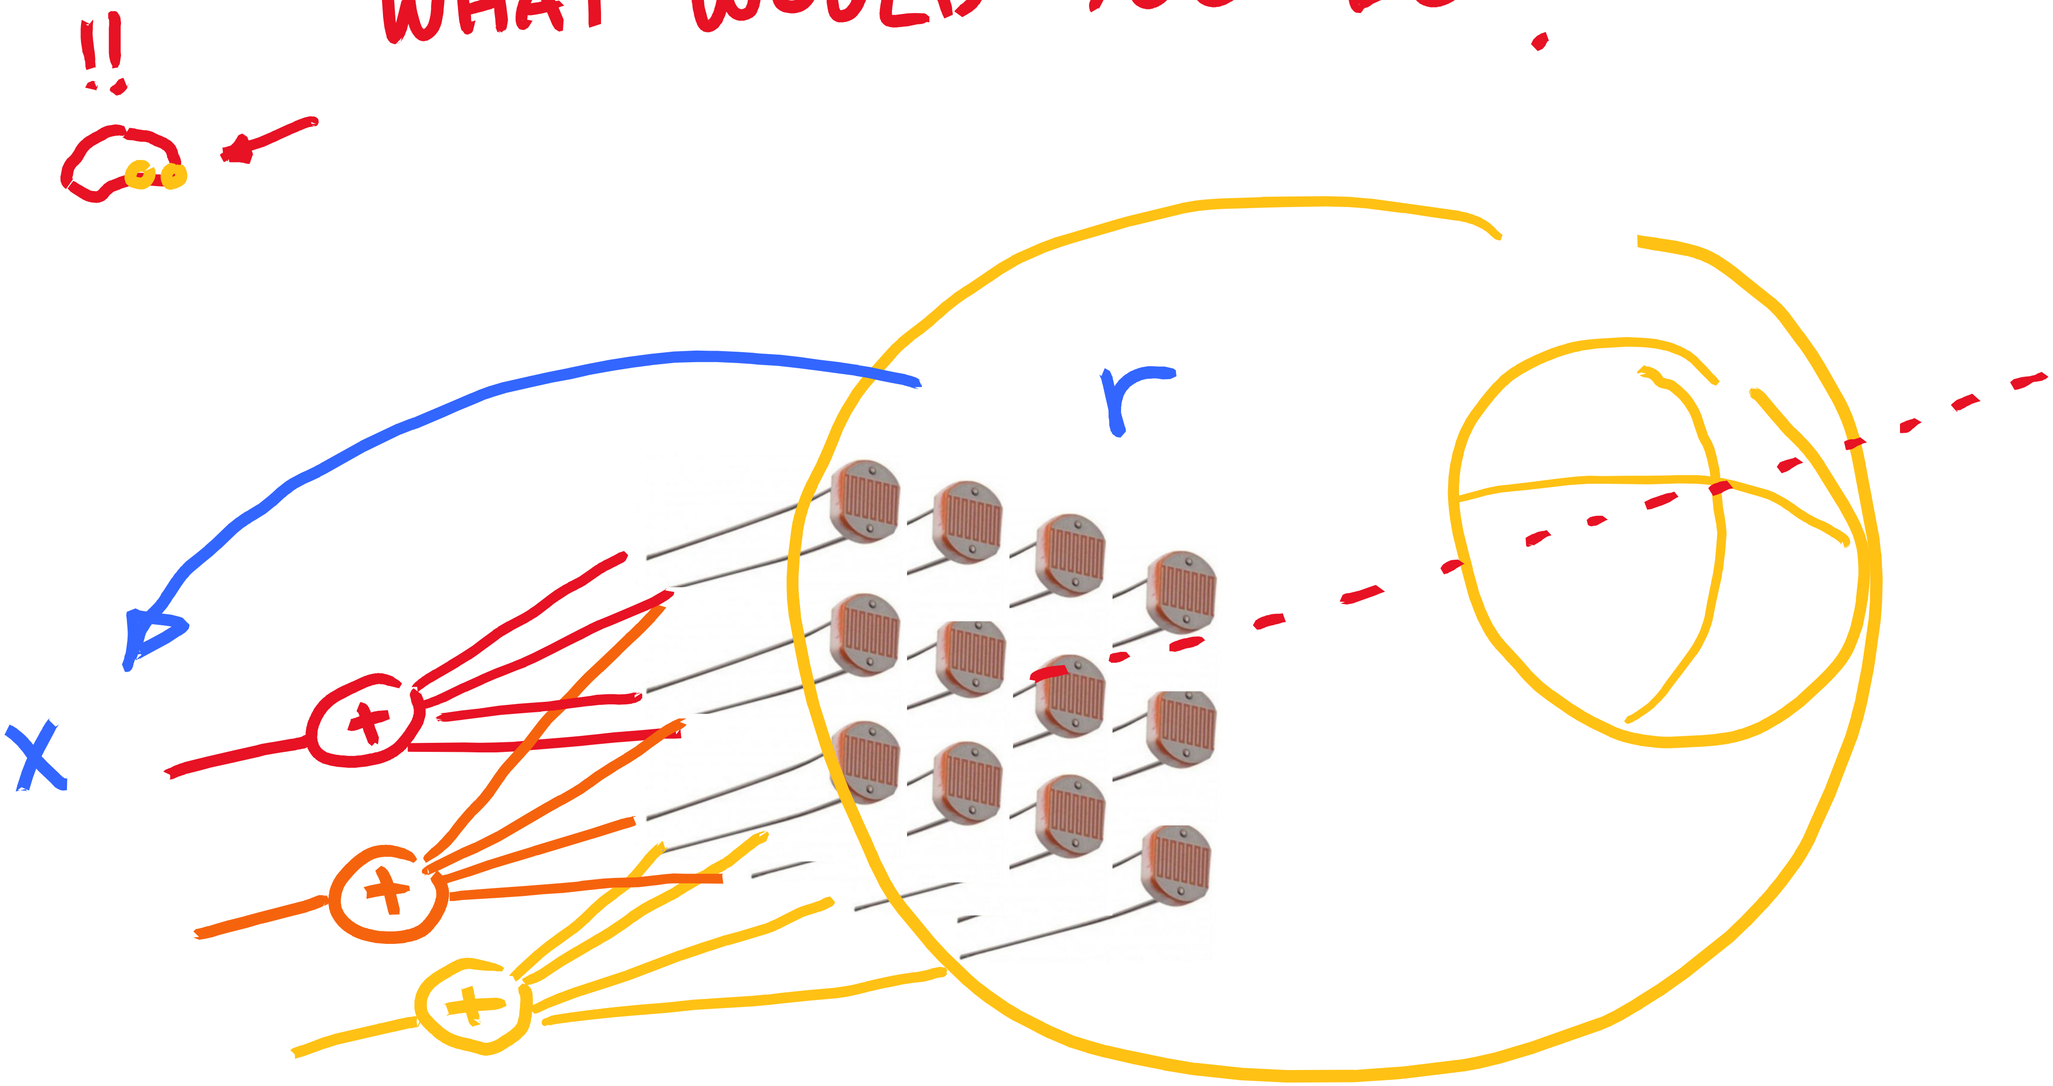

WHAT WOULD YOU DO ?

$$r \xrightarrow{S_\theta} x = S_\theta(r) + n$$

VISUAL ILLUSIONS IN HUMAN AND ARTIFICIAL NEURAL NETWORKS

JESÚS MALO

VNIVERSITAT DE VALÈNCIA

SPAIN

OUTLINE

1. THEORY: Sensible transforms in visual neuroscience
   - Uniformizat.
   - Gaussianizat.
2. Visual illusions from context
   - Examples
     - Texture
     - Motion
     - Color & Brightness
   - Empirical descriptions
3. RESULTS I: Human illusions from infomax and errormin.
4. Illusions in Convolutional Neural Networks
   - Physiology
   - Psychophysics
5. RESULTS II: Artificial illusions from "artificial" CSFs
6. DISCUSSION & CONCLUSIONS
   - Generality vs illusions
   - Vision & Stats in GANs
   - Goals in Neuroscience
7. Other stuff we do
   - Psychophysics - MAD
   - Physiology - fMRI

1 THEORY: Sensible transforms in visual neuroscience

1.1 Information flow with noisy sensors $\Rightarrow$ $\begin{cases} \text{- Maximize entropy} \\ \text{- Minimize redundancy} \end{cases}$

1.2 Sensible transforms $\begin{cases} \text{- Uniformization} \\ \text{- Gaussianization} \end{cases}$

1 THEORY: Sensible transforms in visual neuroscience: Information flow in noise

$$r = \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} \xrightarrow{S_\theta} x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

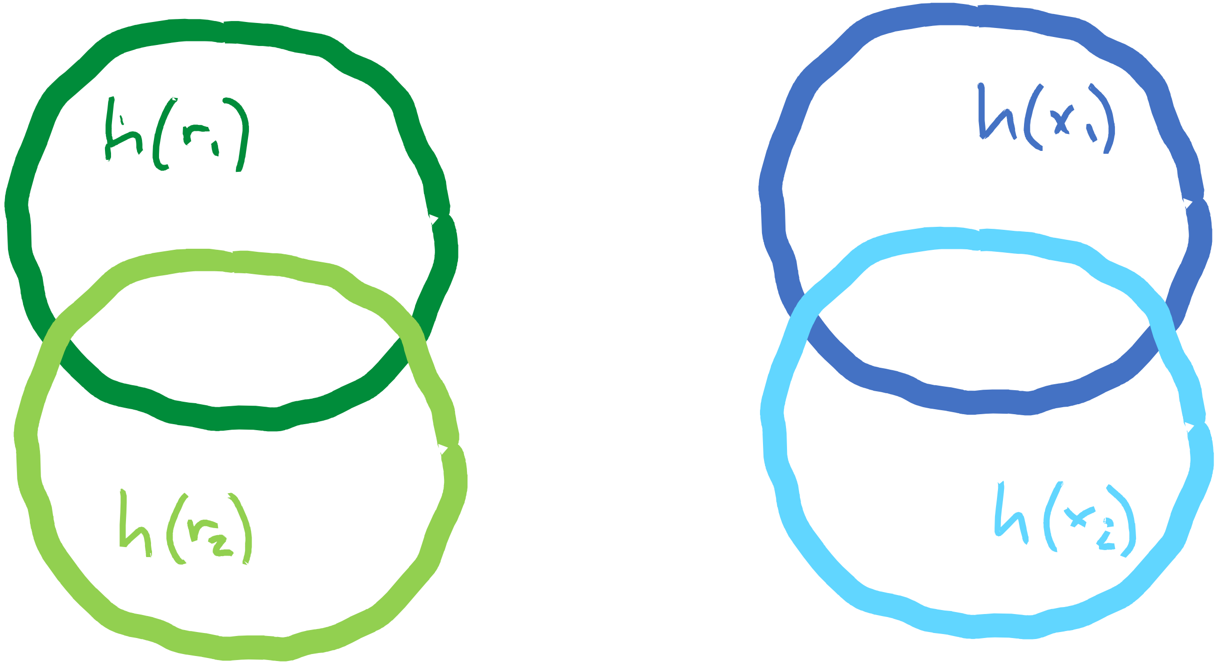

① THEORY: Sensible transforms in visual neuroscience: Information flow in noise

$$r = \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} \xrightarrow{S_\theta} x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

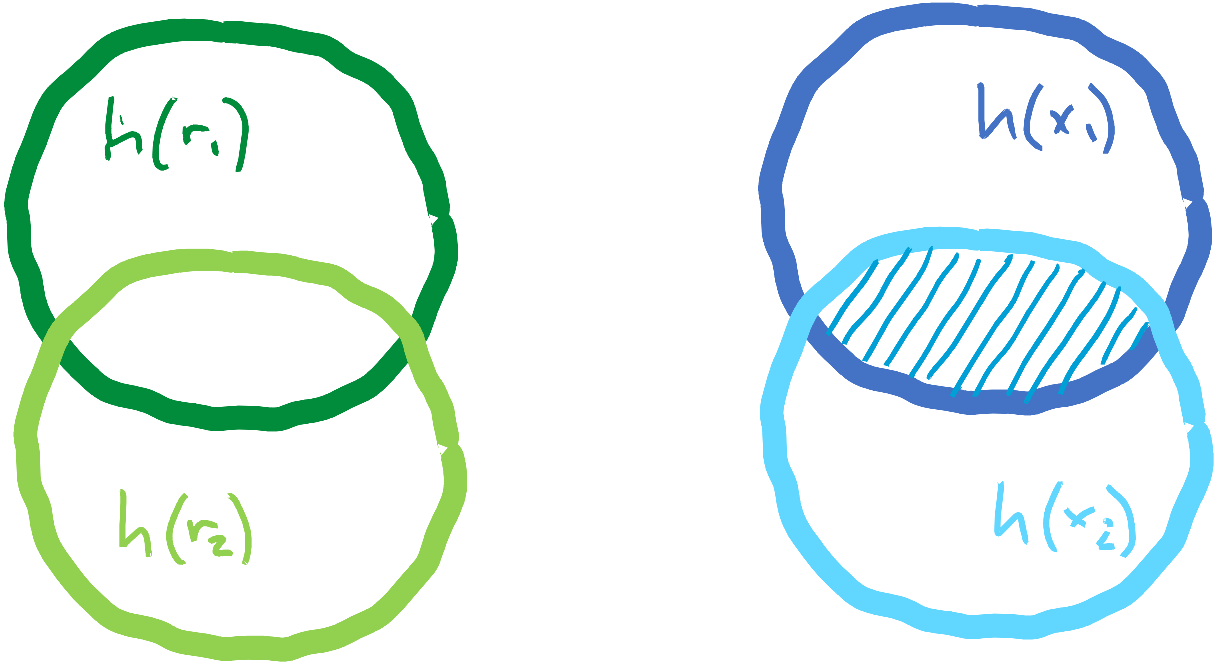

TOTAL CORRELATION ≡ Redundancy within a vector $T(x) = \sum_i h(x_i) - h(x)$

1 **THEORY:** Sensible transforms in visual neuroscience: Information flow in noise

$$r = \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} \xrightarrow{S_\theta} x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

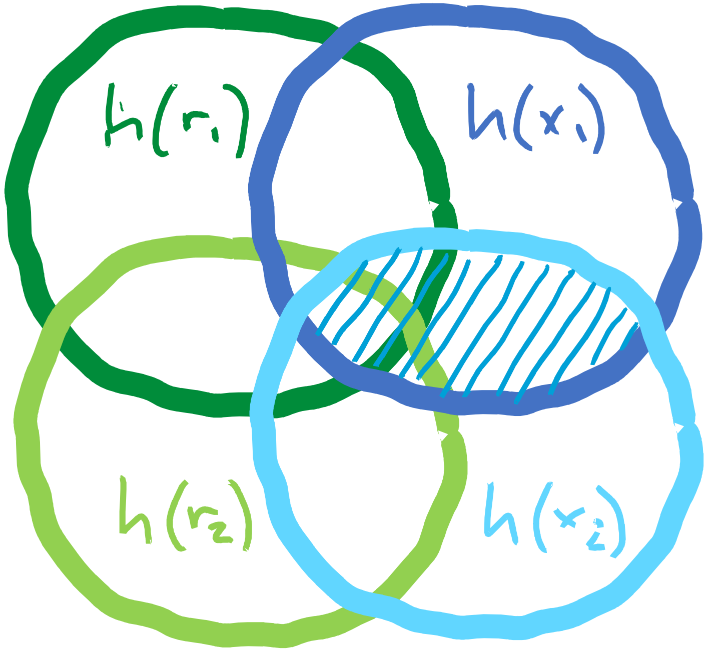

TOTAL CORRELATION ≡ Redundancy within a vector $T(x) = \sum_i h(x_i) - h(x)$

① THEORY: Sensible transforms in visual neuroscience: Information flow in noise

$$r = \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} \xrightarrow{S_\theta} x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

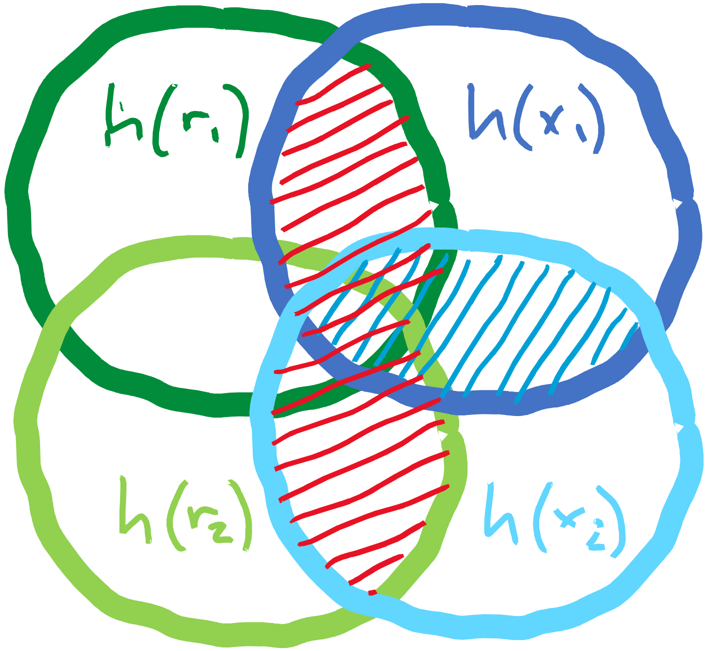

TOTAL CORRELATION ≡ Redundancy within a vector $T(x) = \sum_i h(x_i) - h(x)$

MUTUAL INFORMATION ≡ Info shared by two vectors $I(r, x) = h(r) + h(x) - h([r, x])$

(1) THEORY: Sensible transforms in visual neuroscience: Information flow in noise

The "design" problem:

$$r \xrightarrow{S} x = S(r) + n$$

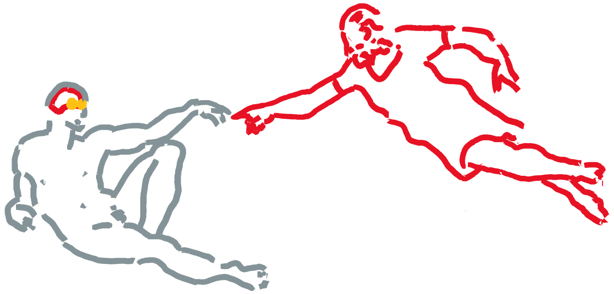

please maximize $I(r, x)$!

J. Malo (2019) Spatio-chromatic information available from different neural layers
https://arxiv.org/abs/1910.01559

① THEORY: Sensible transforms in visual neuroscience: Information flow in noise

The "design" problem:

$$r \xrightarrow{S} x = S(r) + n$$

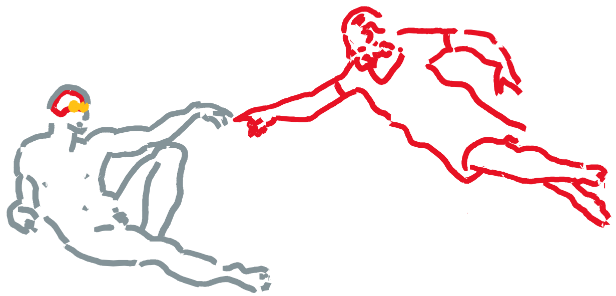

please maximize $I(r,x)$!

$$I(r,x) = h(r) + E_r\left[\log_2 \|\nabla S\|\right] - \left(h(n) - E_n\left[D_{KL}\left(p(s(r)) \,|\, p(s(r)+n)\right)\right]\right)$$

J. Malo (2019) Spatio-chromatic information available from different neural layers
https://arxiv.org/abs/1910.01559

1 THEORY: Sensible transforms in visual neuroscience: Information flow in noise

The "design" problem:

$$r \xrightarrow{S} x = S(r) + n$$

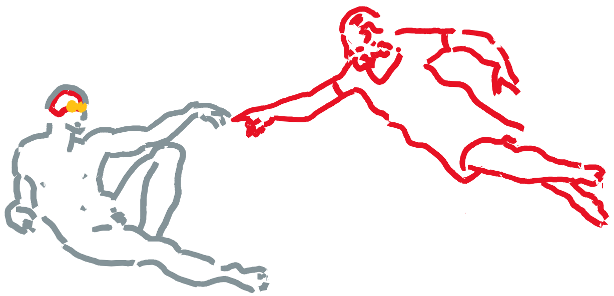

please maximize $I(r,x)$!

$$I(r,x) = h(r) + E_r\left[\log_2 \|\nabla S\|\right] - \left(h(n) - E_n\left[D_{KL}\left(p(S(r)) \,\middle|\, p(S(r)+n)\right)\right]\right)$$

$$I(r,x) = \underbrace{\sum_i h(x_i)}_{(1)} - \underbrace{T(x)}_{(2)} - h(n) \Rightarrow \begin{cases} (1) \text{ MAXIMIZE ENTROPY} \\ (2) \text{ MINIMIZE REDUNDANCY} \end{cases}$$

J. Malo (2019) Spatio-chromatic information available from different neural layers
https://arxiv.org/abs/1910.01559

1 THEORY: Sensible transforms in visual neuroscience:

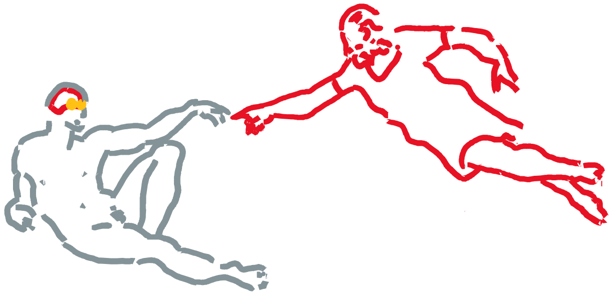

please maximize $I(r,x)$!

$\Rightarrow$ (1) MAXIMIZE ENTROPY

(2) MINIMIZE REDUNDANCY

$\Rightarrow$ UNIFORMIZATION

GAUSSIANIZATION

J. Malo (2019) Spatio-chromatic information available from different neural layers
https://arxiv.org/abs/1910.01559

1 THEORY: Sensible transforms in visual neuroscience: UNIFORMIZATION

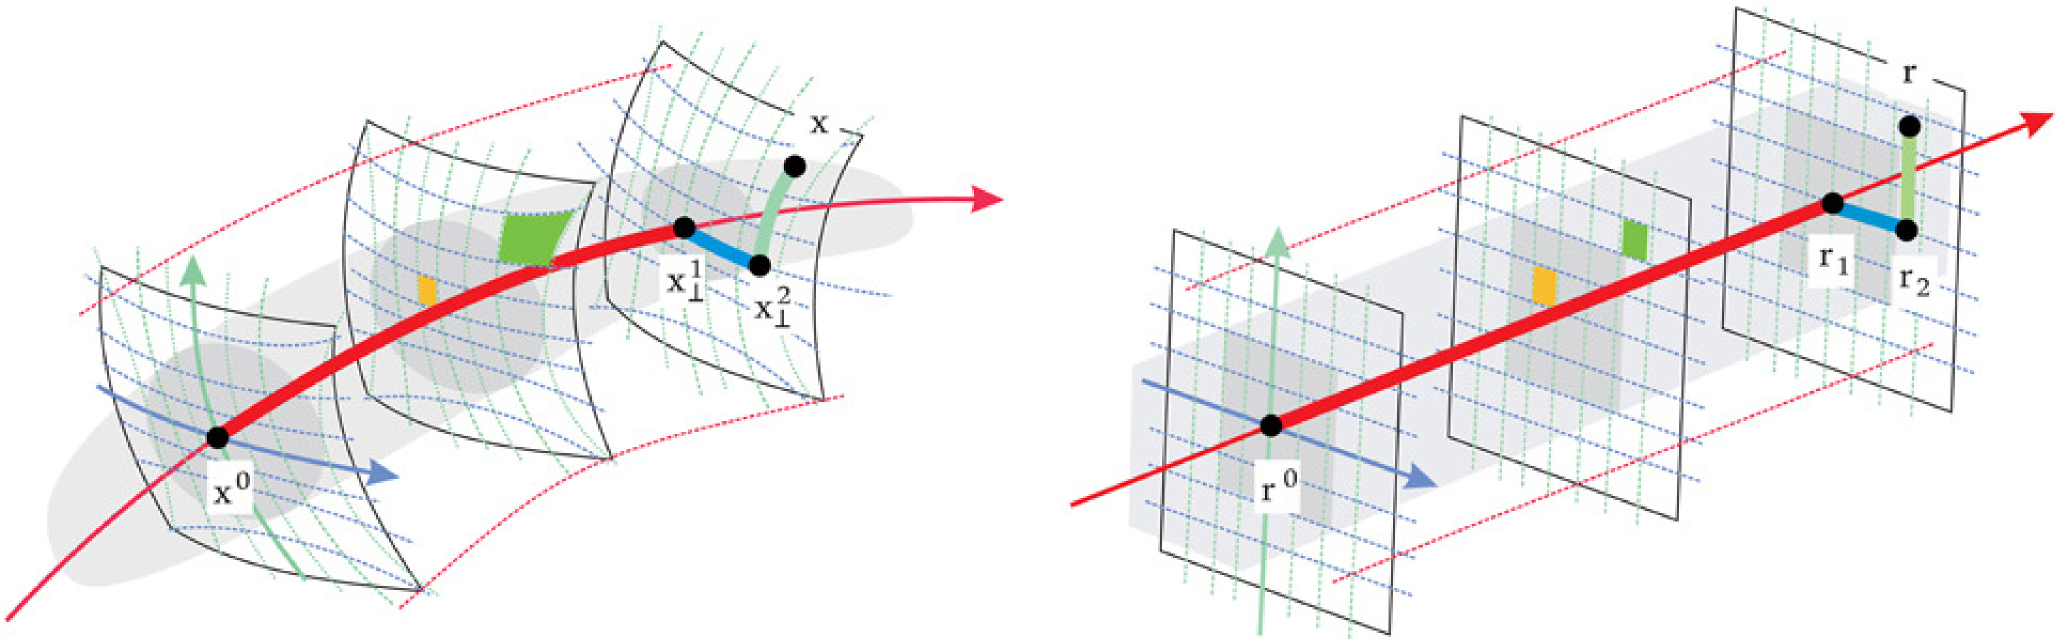

J. Malo & J. Gutiérrez (2006) V1-nonlinearities emerge from Local-to-Global ICA
Network: Comp. Neur. Syst. Vol. 17, 85–102
V. Laparra, J. Malo et al. (2012) Nonlinearities and adaptation in color vision from Sequential Principal Curves Analysis
Neural Comput. Vol. 24, 2751–2788. doi: 10.1162/NECO_a_00342
V. Laparra & J. Malo (2015) Visual aftereffects and sensory nonlinearities from a single statistical framework
Front. Hum. Neurosci., https://doi.org/10.3389/fnhum.2015.00557

1 THEORY: Sensible transforms in visual neuroscience: GAUSSIANIZATION

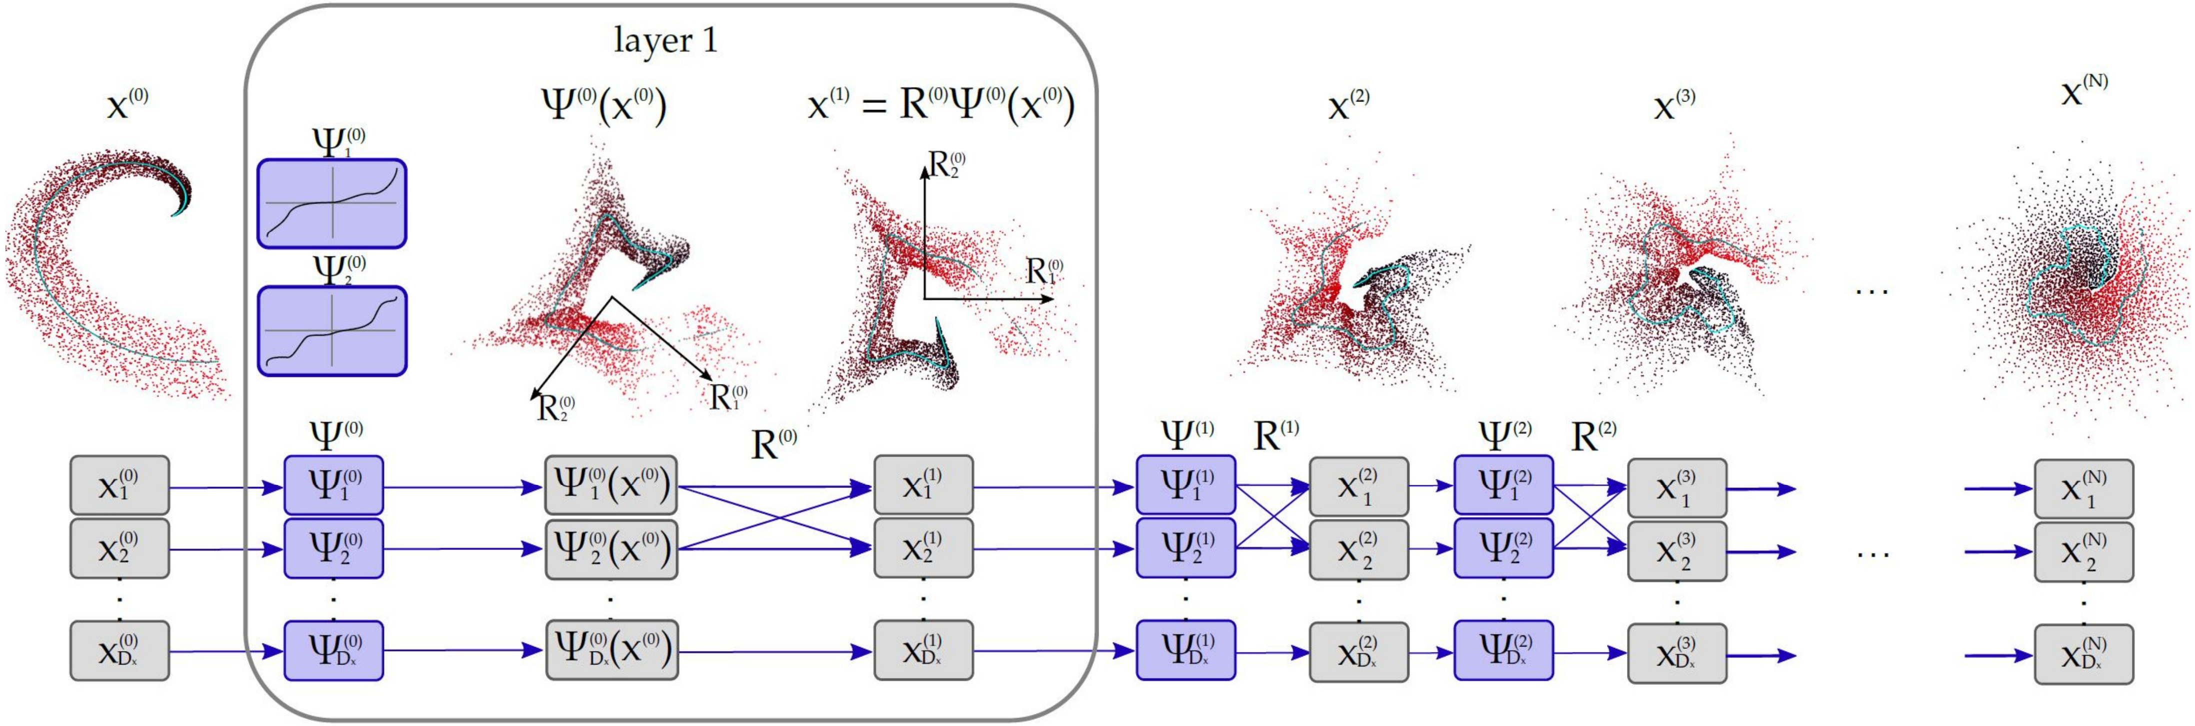

V. Laparra, G. Camps & J. Malo (2011) Rotation-based Iterative Gaussianization: from ICA to Random Rotations
IEEE Trans. Neural Nets. 22(4):537-549, 2011

E. Johnson, V. Laparra, J. Malo et al. (2019) Information theory in density destructors.
7th Int. Conf. Mach. Learn., **ICML 2019**, Workshop on Invertible Normalization Flows

① THEORY: Sensible transforms in visual neuroscience:

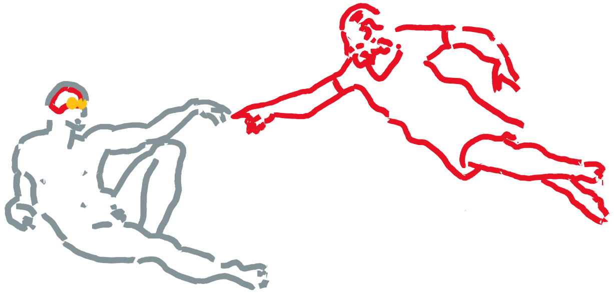

maximize $I(r,x)$! $\Longrightarrow$ $\left\{ \begin{array}{l} \text{(1) MAXIMIZE ENTROPY} \\ \\ \text{(2) MINIMIZE REDUNDANCY} \end{array} \right\}$ $\Longrightarrow$

UNIFORMIZATION

GAUSSIANIZATION

2 Visual illusions from context

2.1 Examples

- Texture
- Motion
- Color & Brightness

2.2 Empirical descriptions

2 Visual illusions from context 2.1 EXAMPLES

2.1 Examples

- Texture — Campbell & Blakemore J. Physiol. 69
- Motion — Mather et al. Trends. Cogn. Sci. 08
- Color & Brightness — Loomis Vis. Res. 74
  Ware & Cowan Vis. Res. 82

http://isp.uv.es/after_effects/

2 Visual illusions from context 2.2 EMPIRICAL DESCRIPTIONS: Adaptive Nonlinearities

$$r \xrightarrow{L} C \xrightarrow{N} x$$

LINEAR $\quad C = L \cdot r$

NON-LINEAR $\quad x_i = \dfrac{c_i^{\gamma}}{b_i + \sum_j H_{ij}\, c_j^{\gamma}}$

Heeger & Carandini 94

Carandini & Heeger 12

Simoncelli & Heeger 98

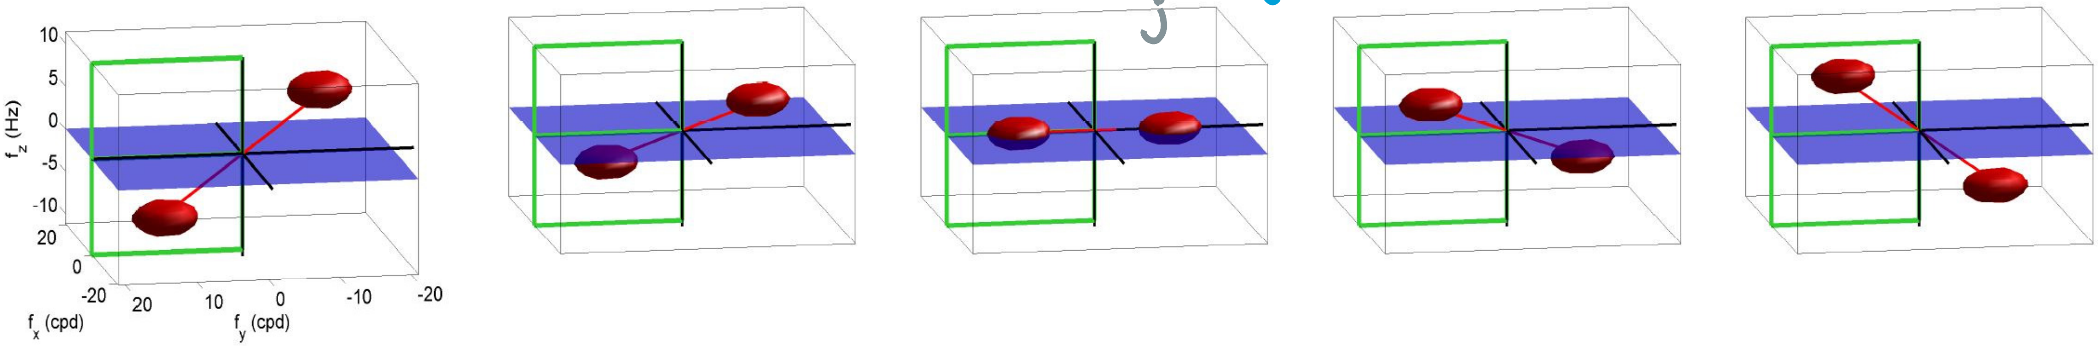

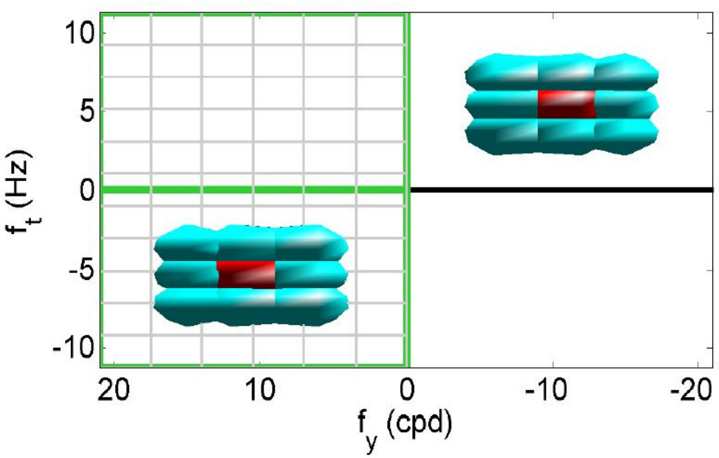

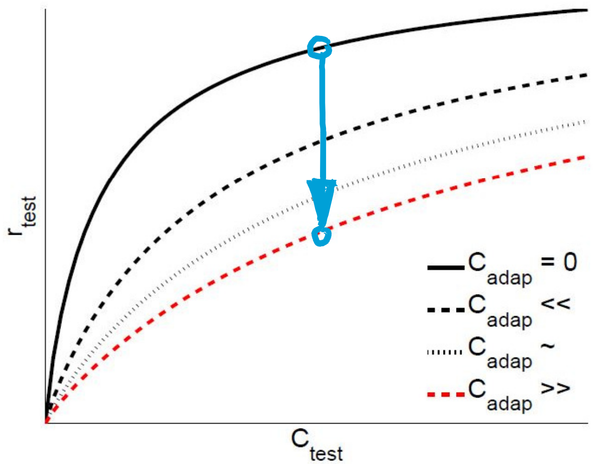

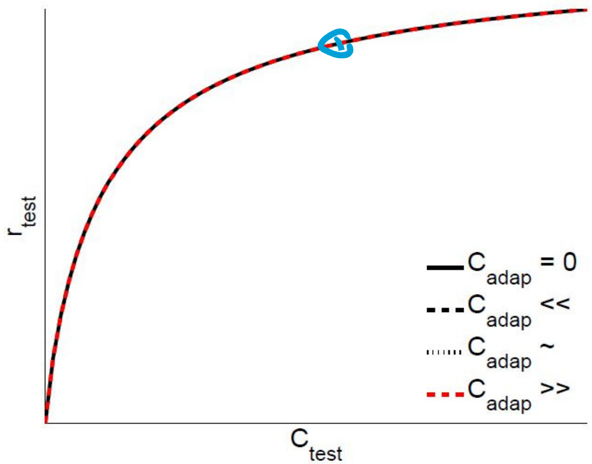

2 Visual illusions from context 2.2 EMPIRICAL DESCRIPTIONS : Adaptive Nonlinearities

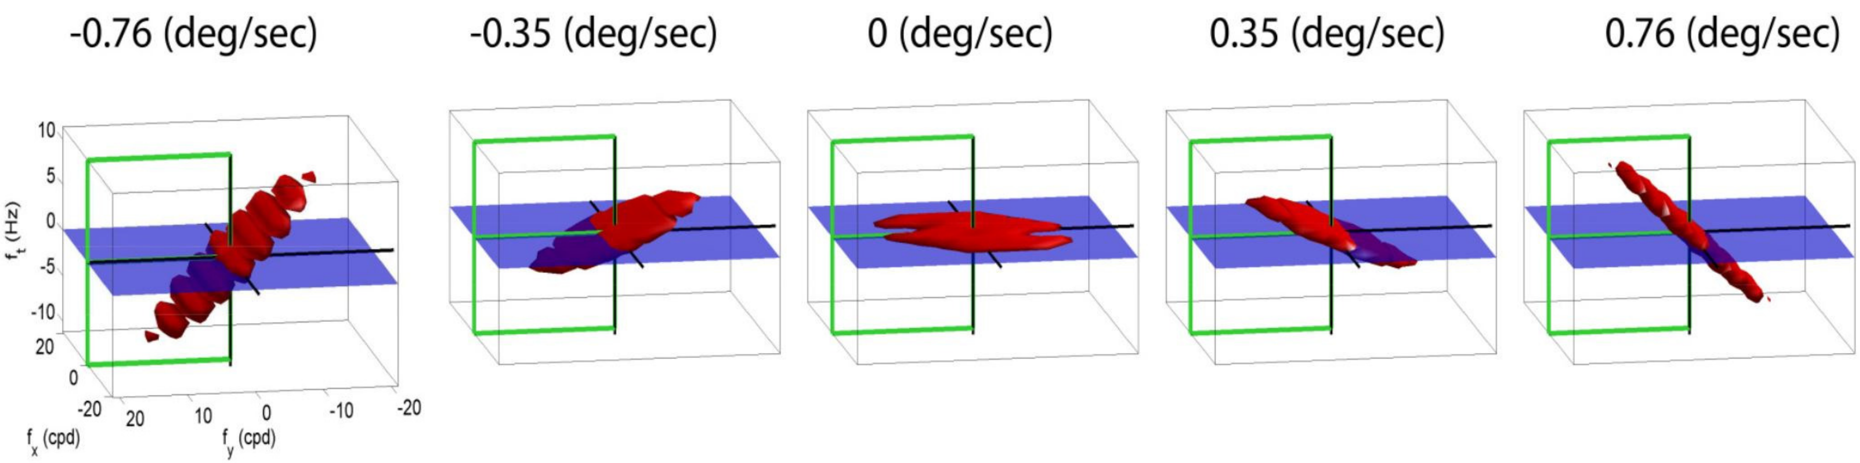

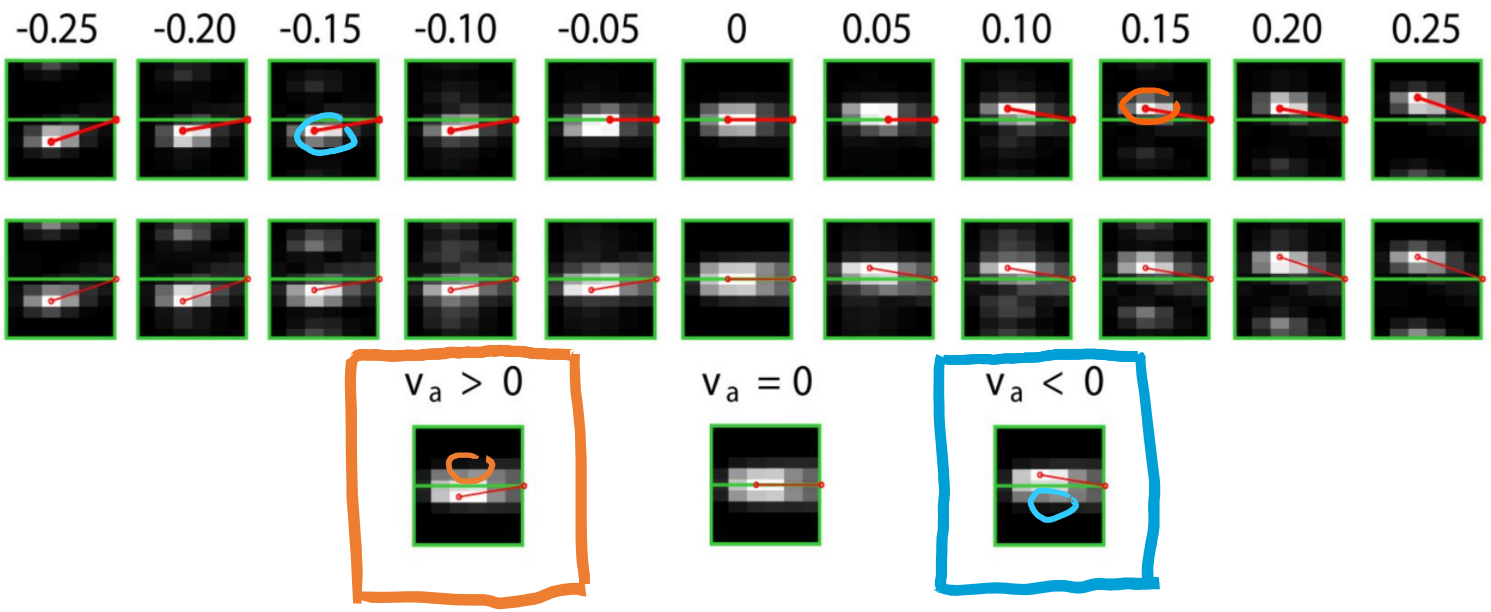

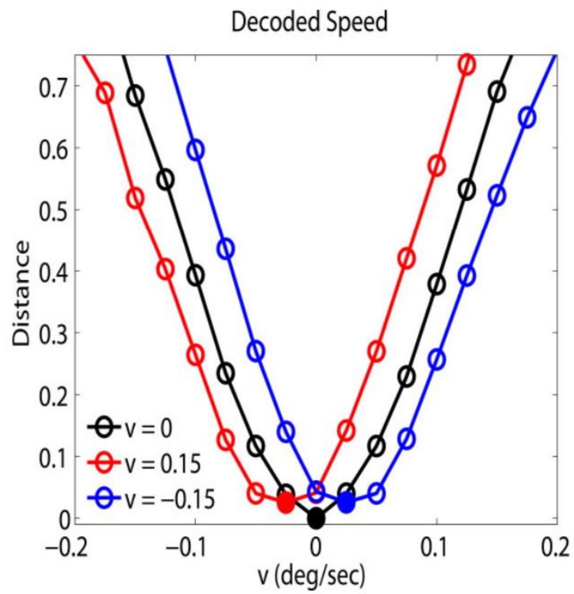

(2) Visual illusions from context 2.2 EMPIRICAL DESCRIPTIONS: Adaptive Nonlinearities

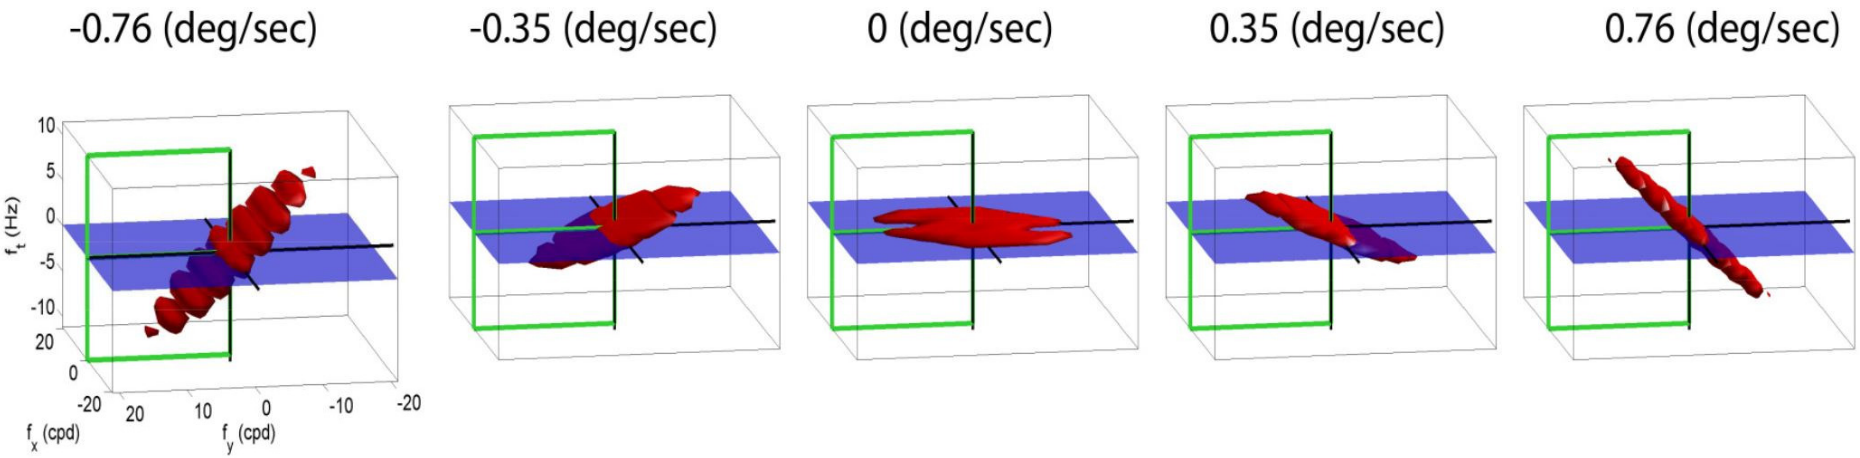

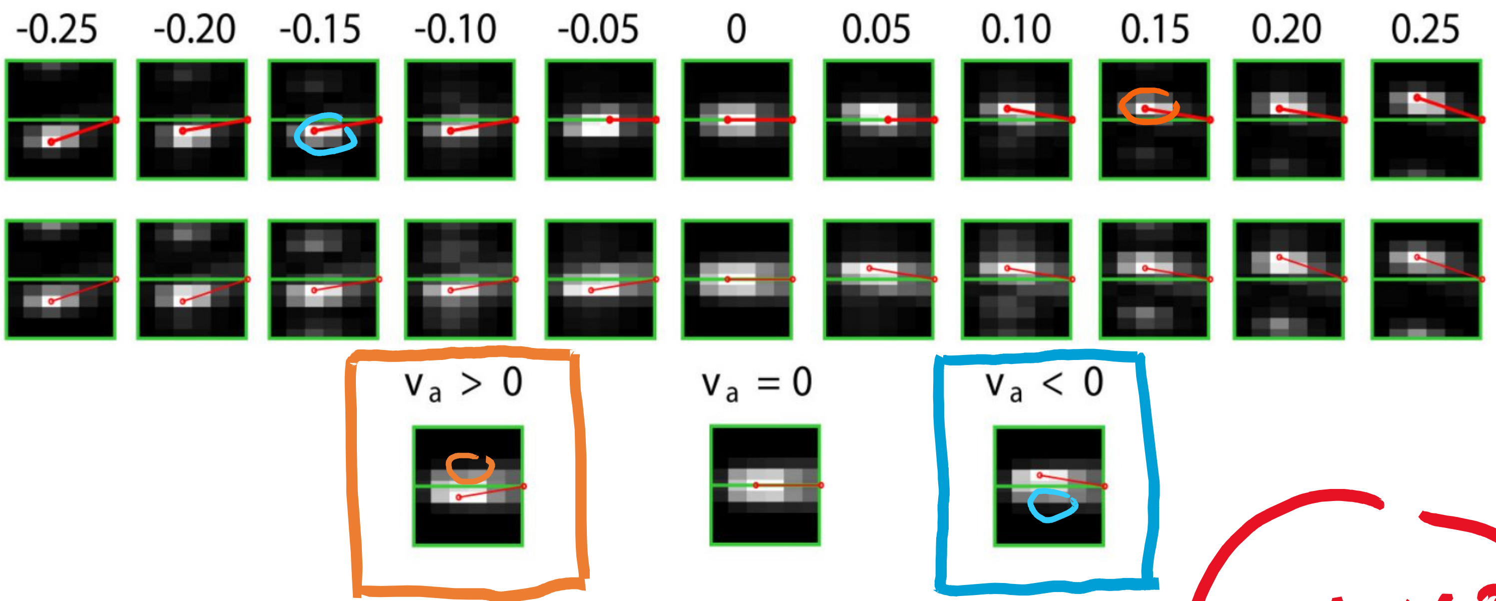

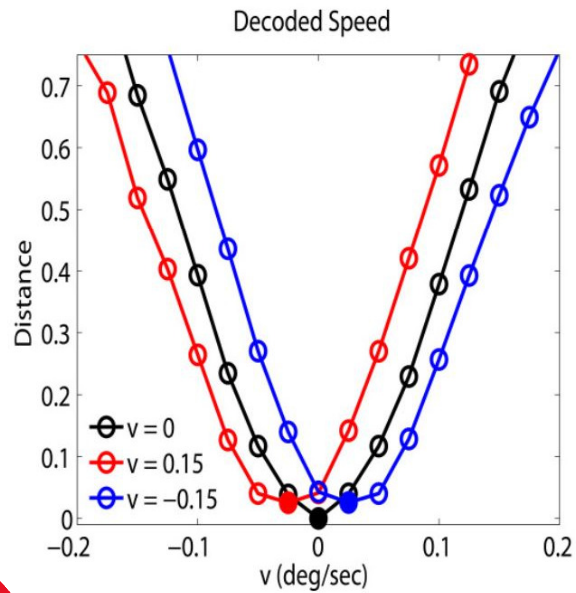

WHY?

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN

- Images & manifolds
- Uniformization idea
- SPCA-uniformization works!
- Results
- Summary

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN

IMAGES & MANIFOLDS

COLOR

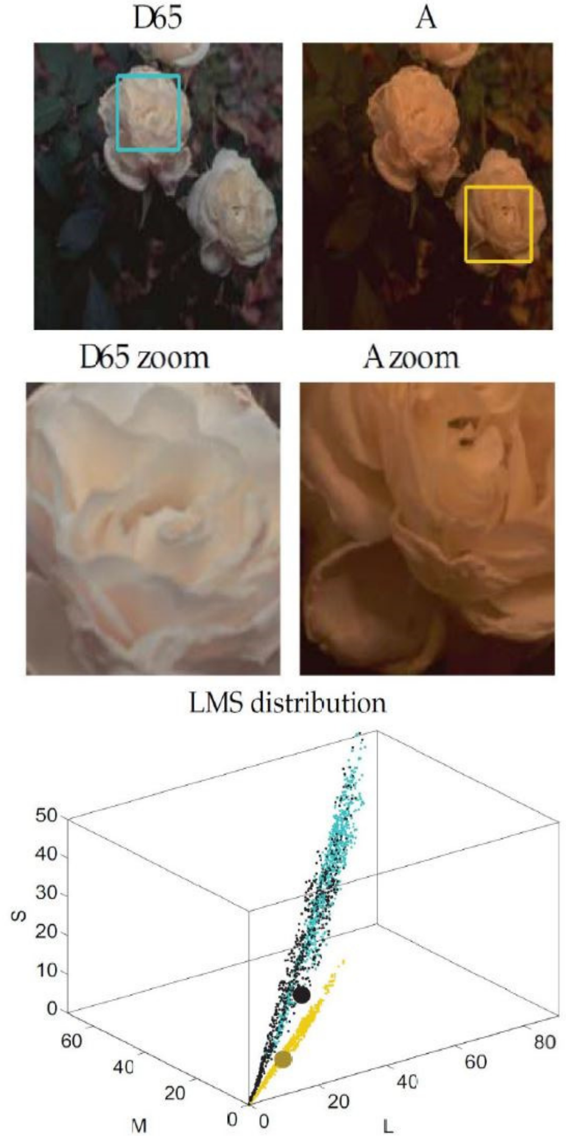

SPATIAL TEXTURE

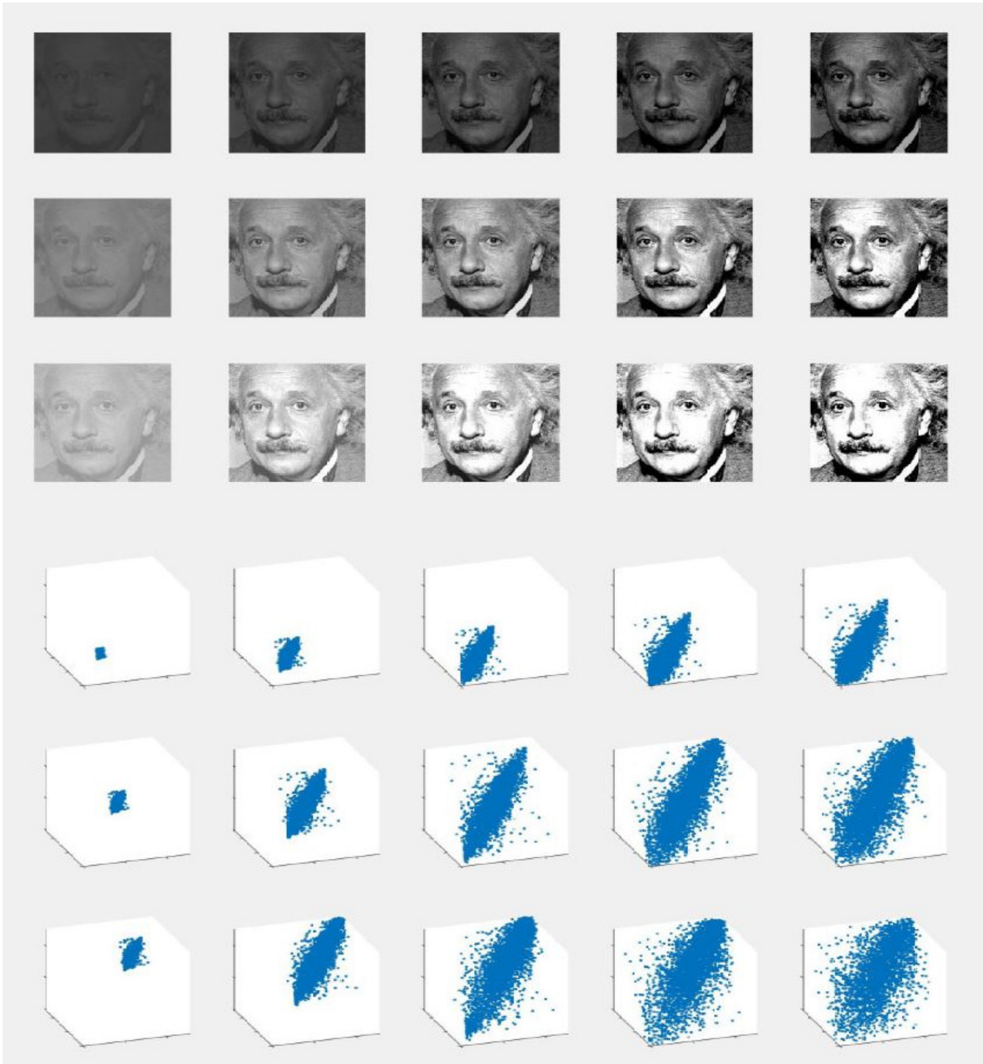

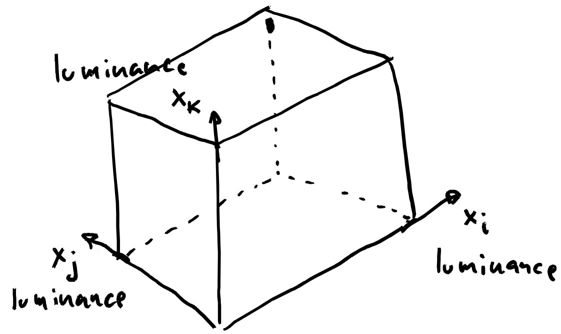

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN IMAGES & MANIFOLDS

COLOR

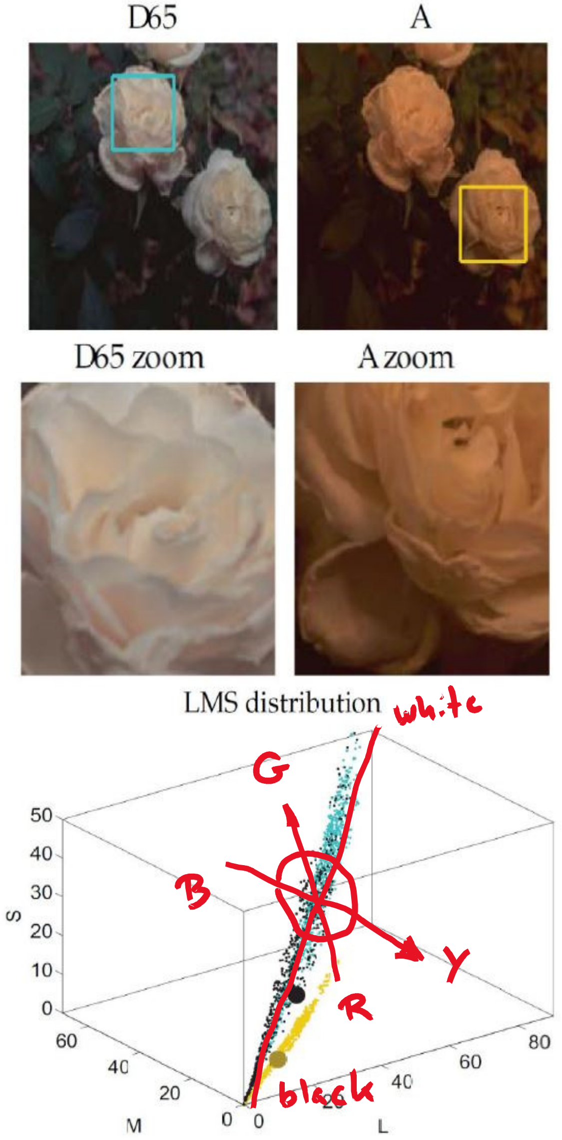

SPATIAL TEXTURE

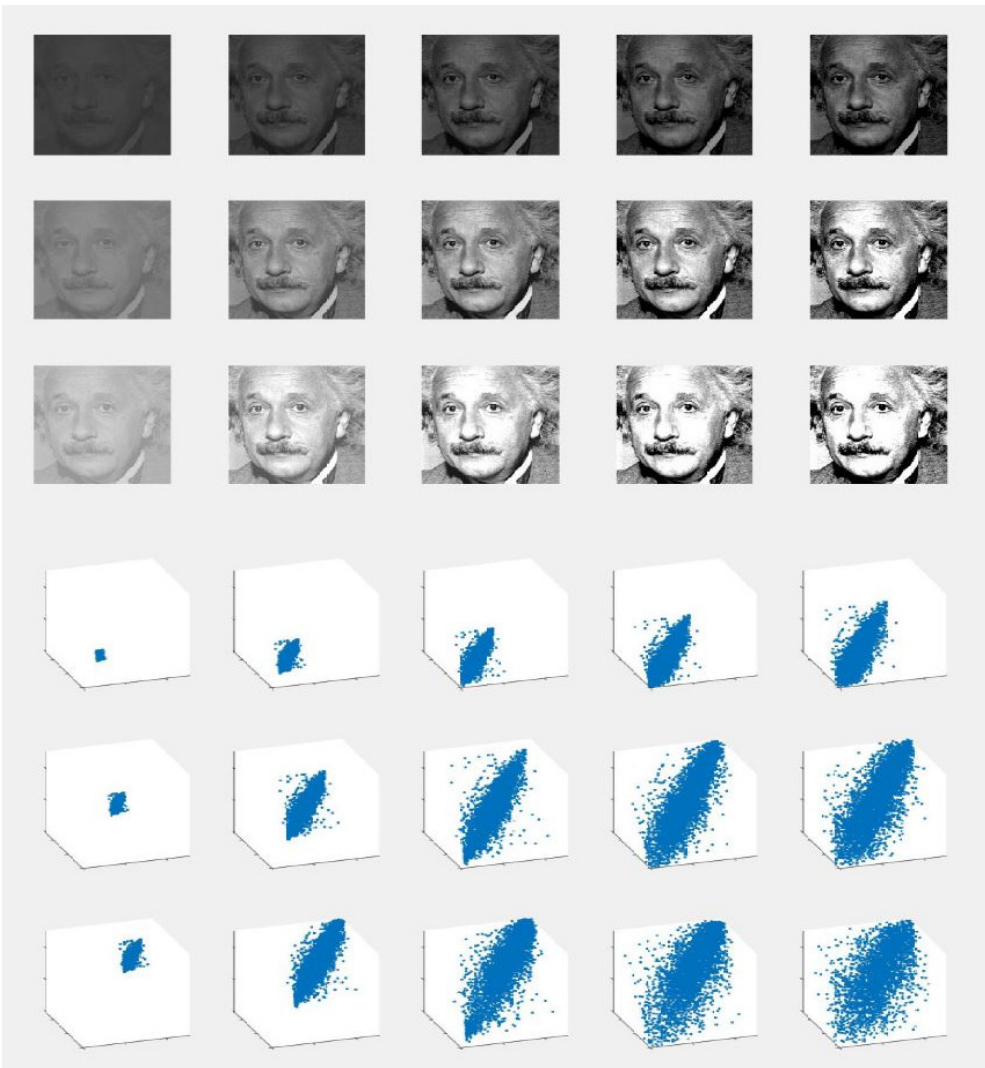

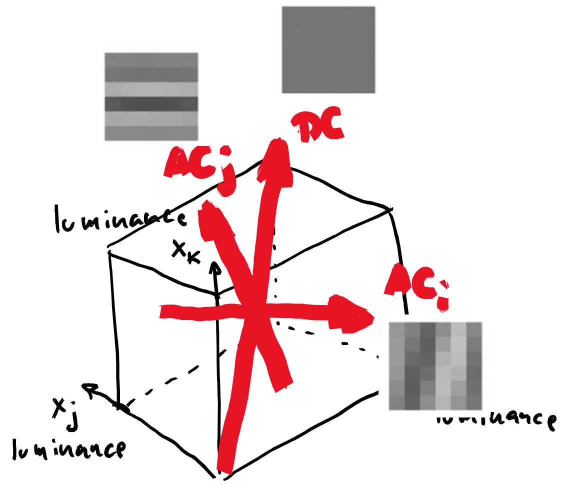

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN   THE UNIFORMIZ. IDEA

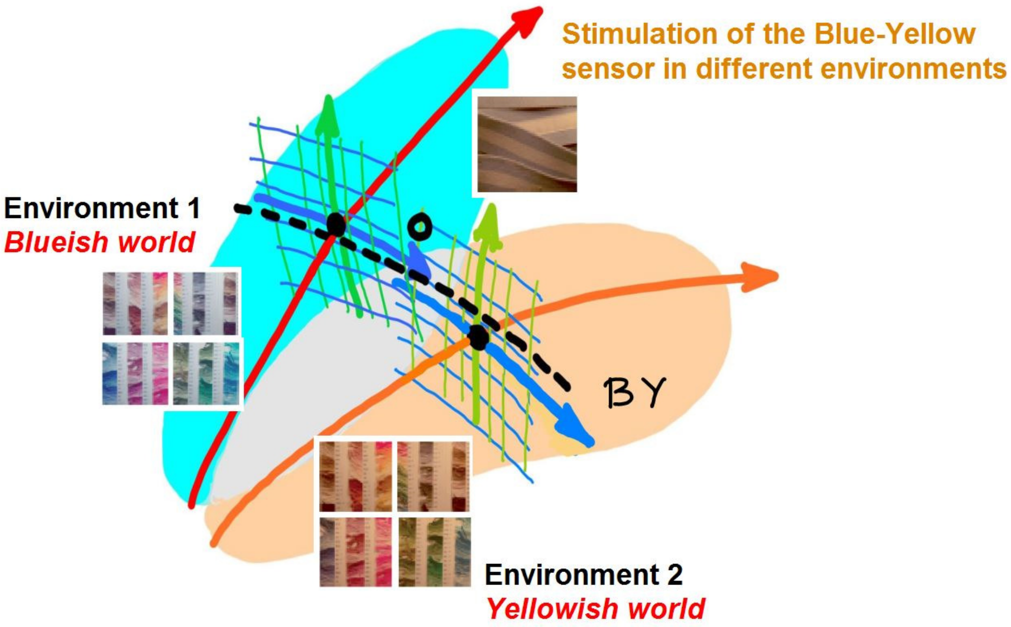

(a)

③ RESULTS I: Human illusions from INFOMAX and ERRORMIN THE UNIFORMIZ. IDEA

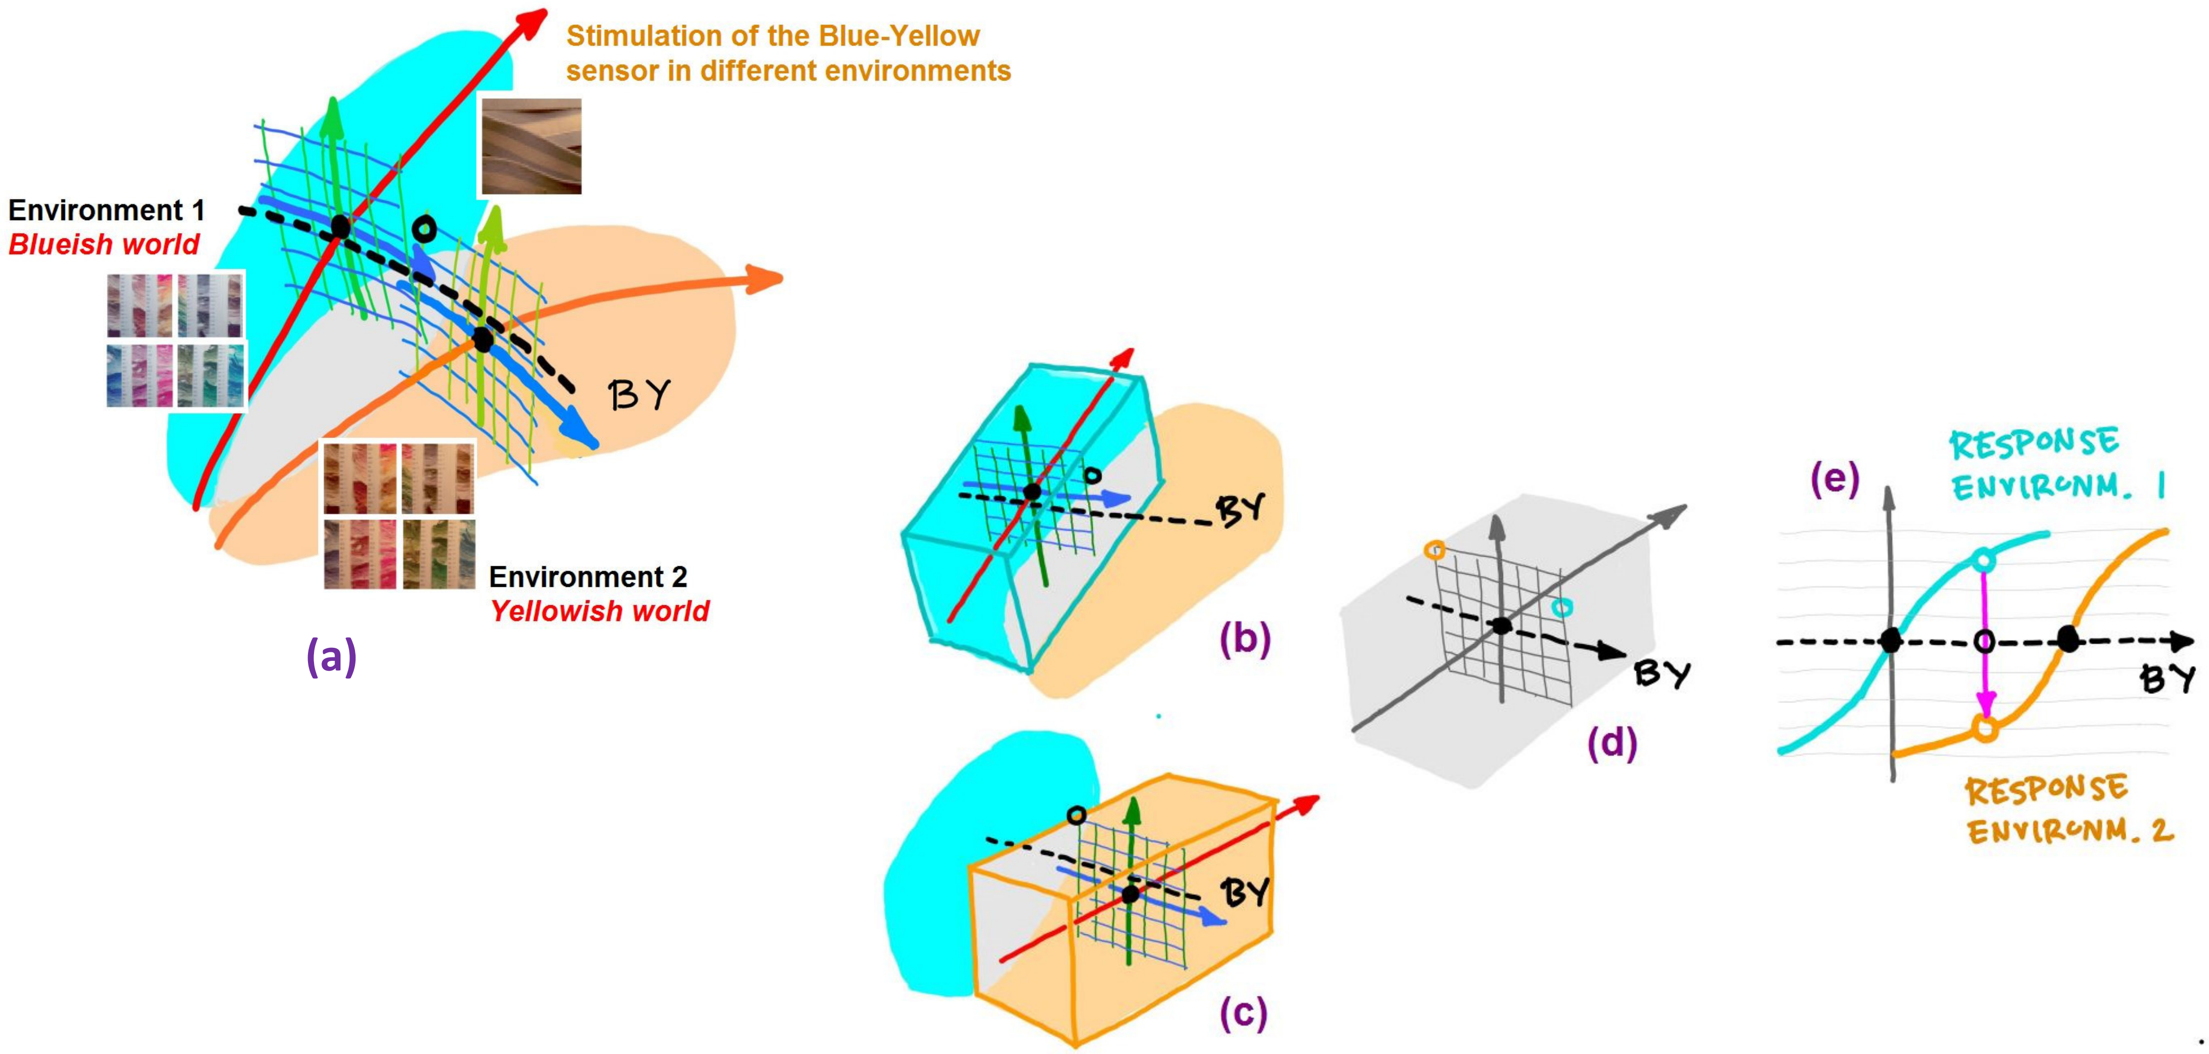

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN THE UNIFORMIZ. IDEA

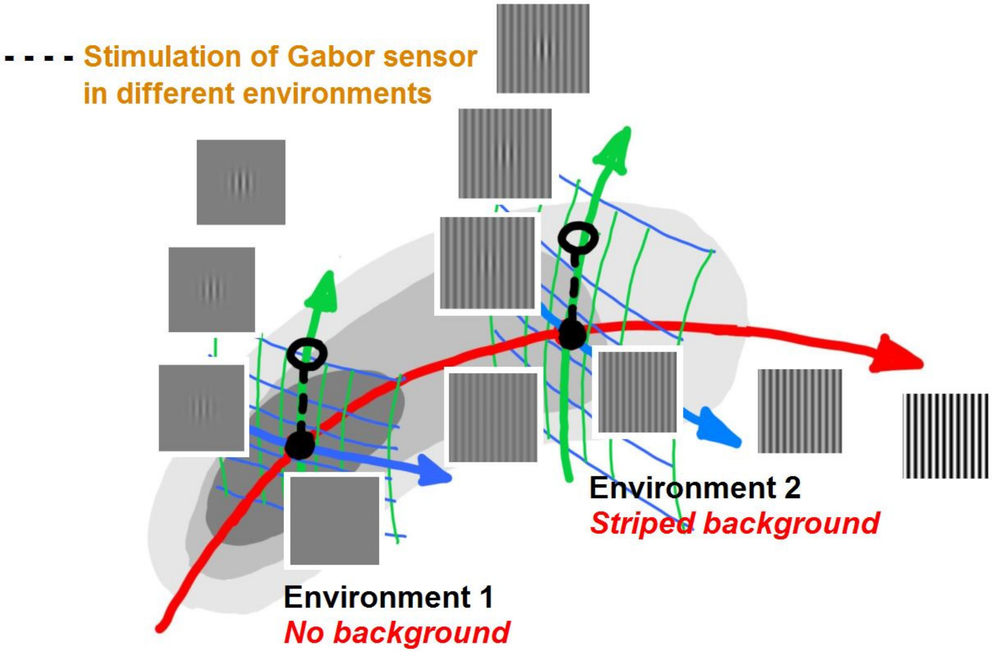

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN THE UNIFORMIZ. IDEA

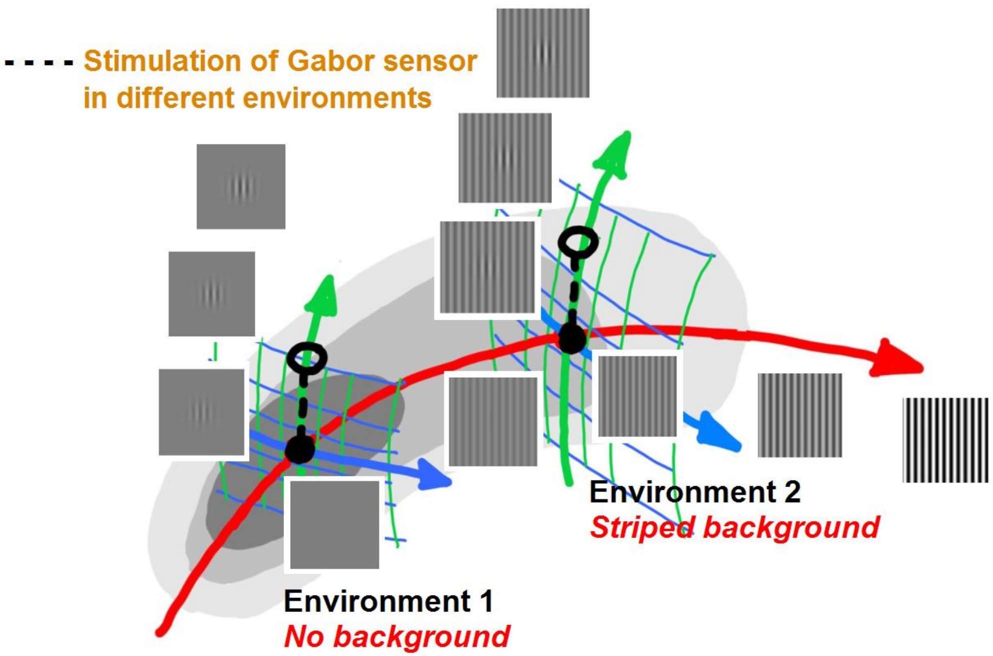

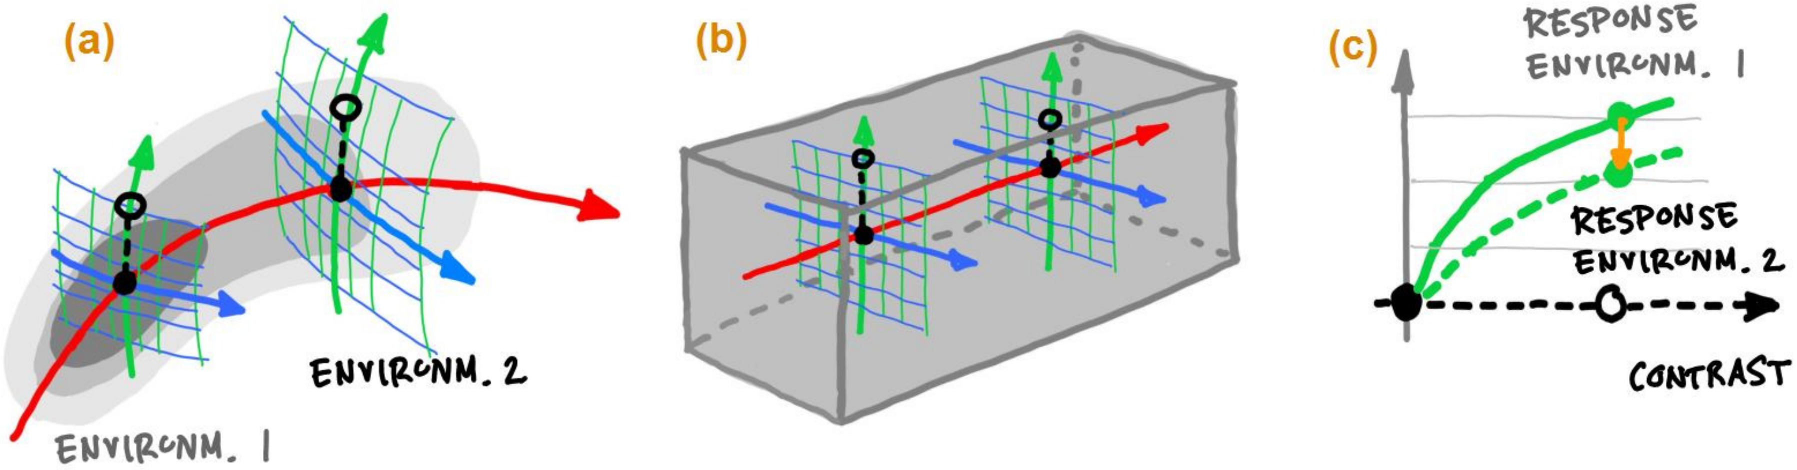

③ RESULTS I: Human illusions from INFOMAX and ERRORMIN UNIFORMIZ. WORKS!

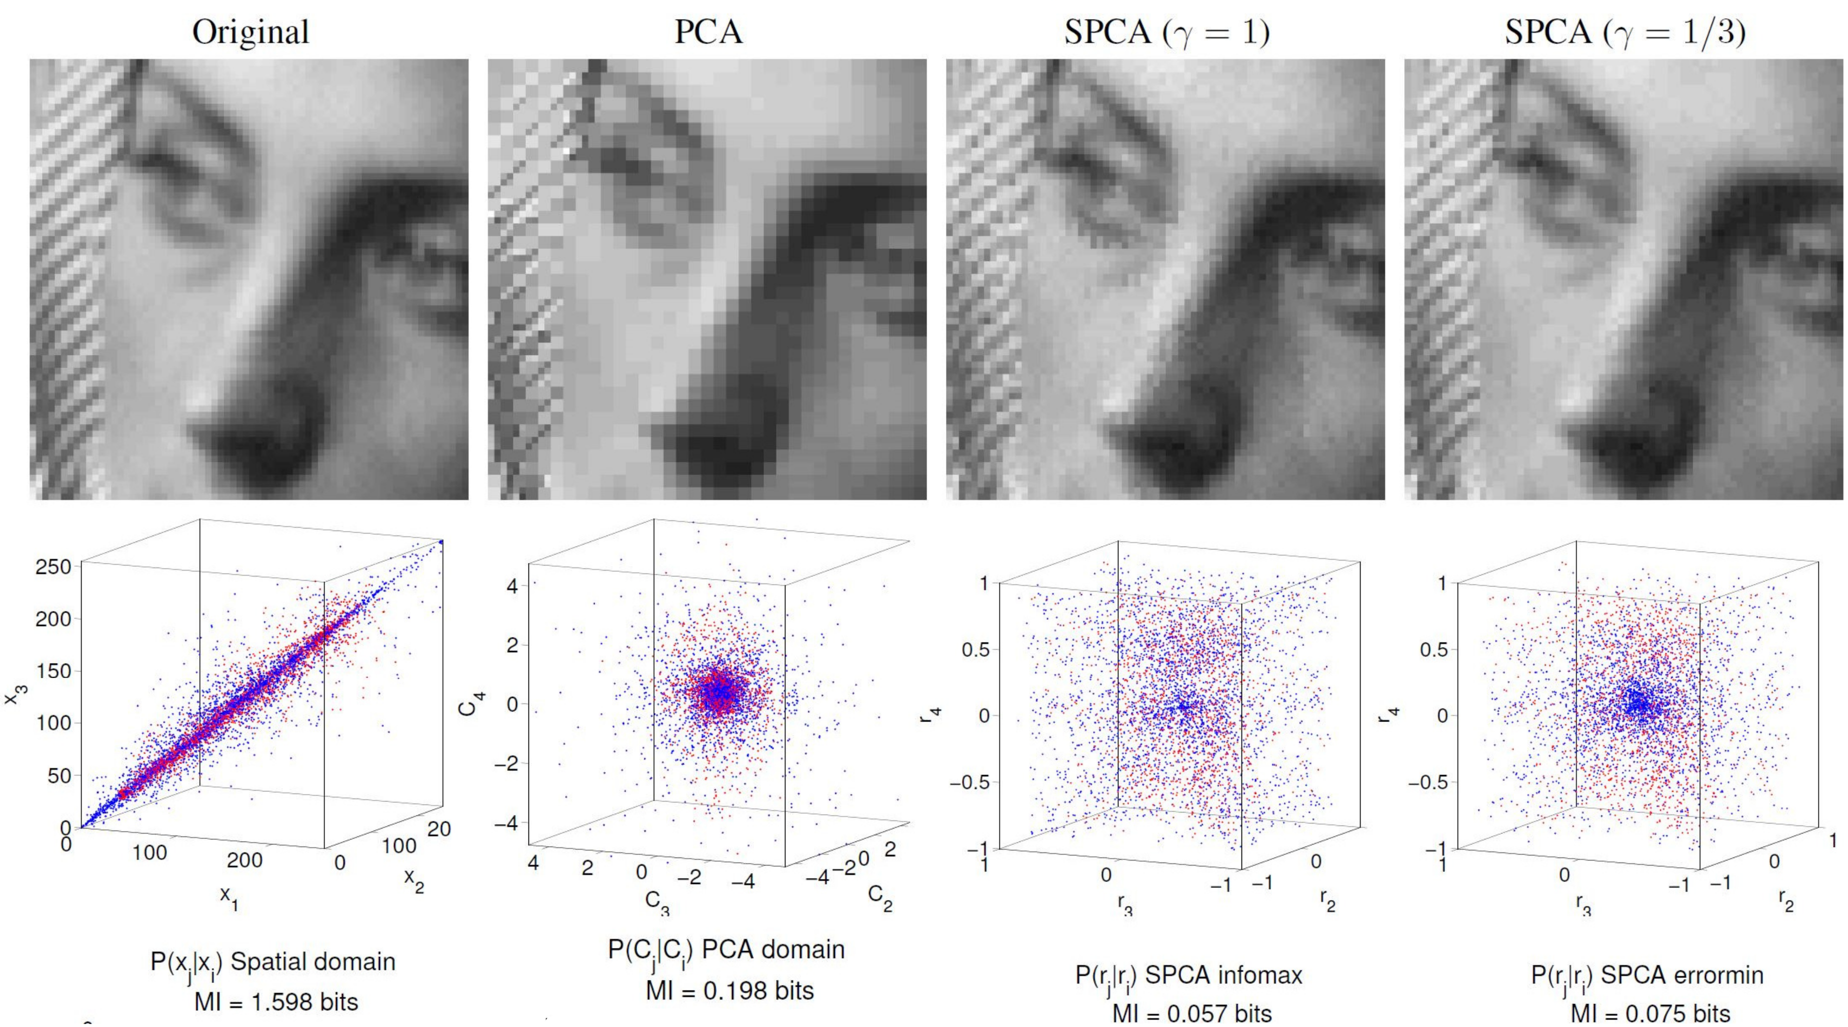

(3) RESULTS I: Human illusions from INFOMAX and ERRORMIN UNIFORMIZ. WORKS!

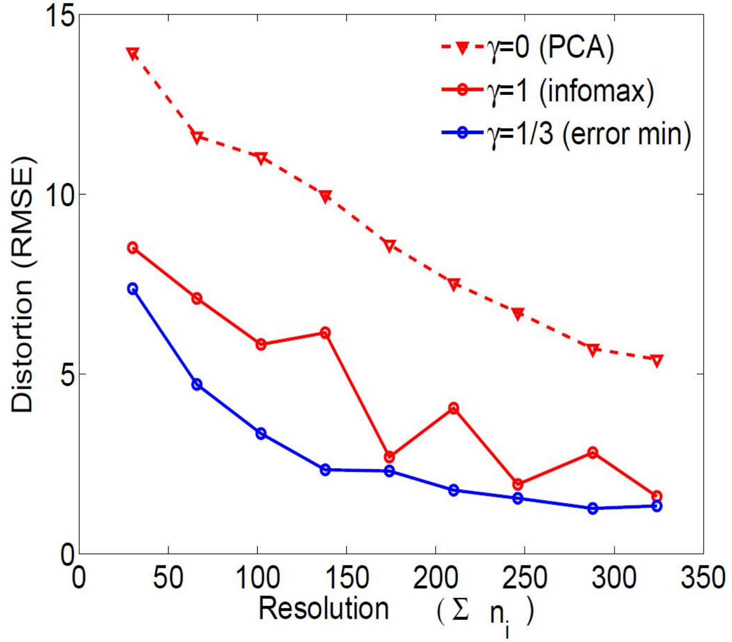

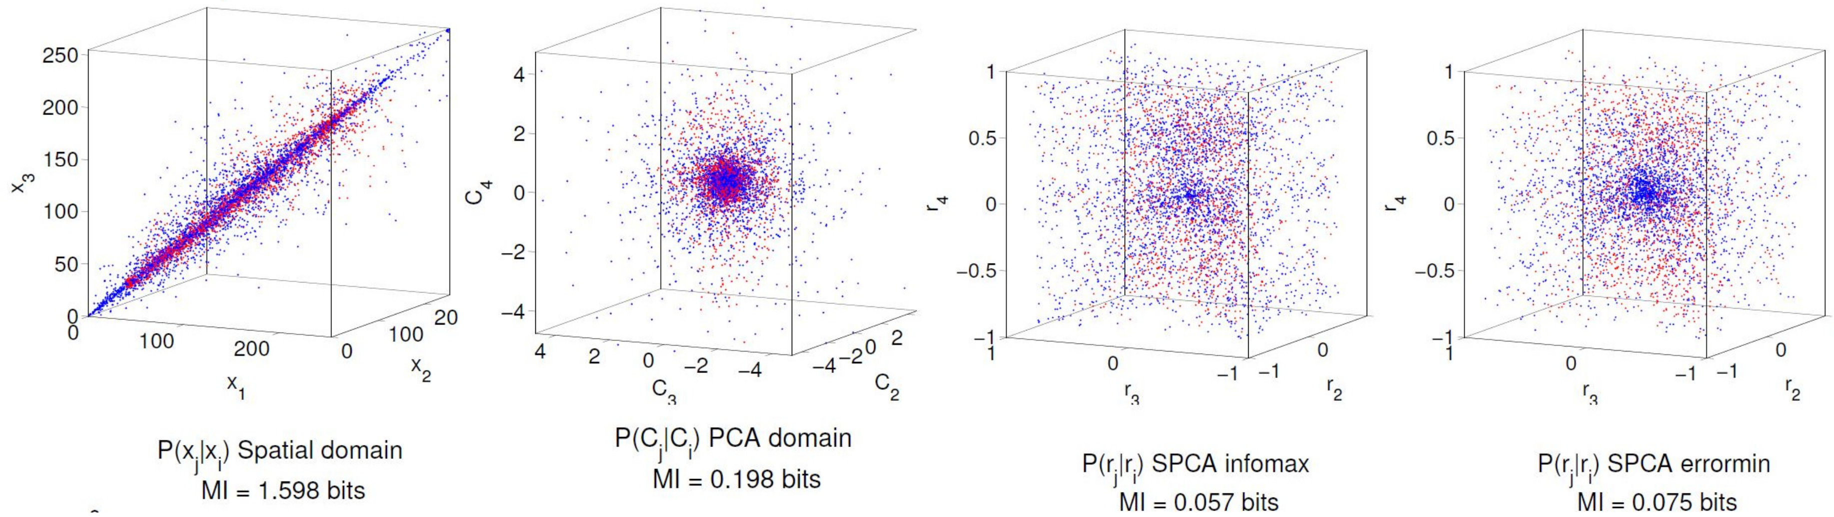

③ RESULTS I: Human illusions from INFOMAX and ERRORMIN RESULTS: TEXTURE

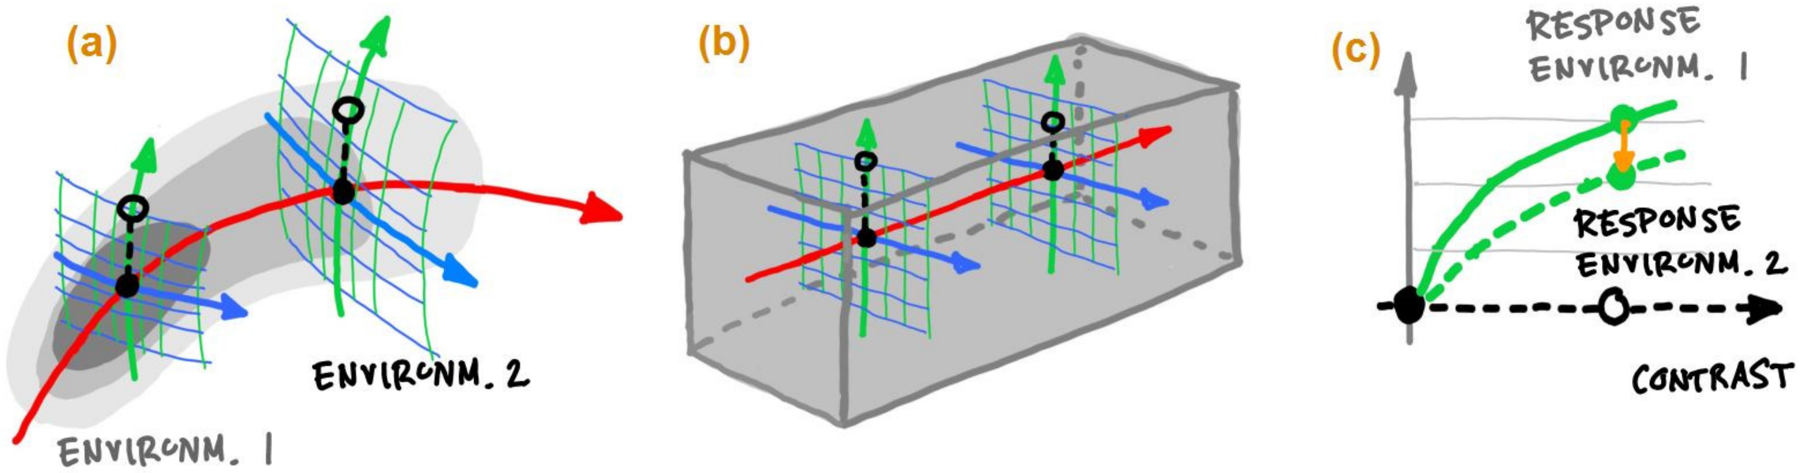

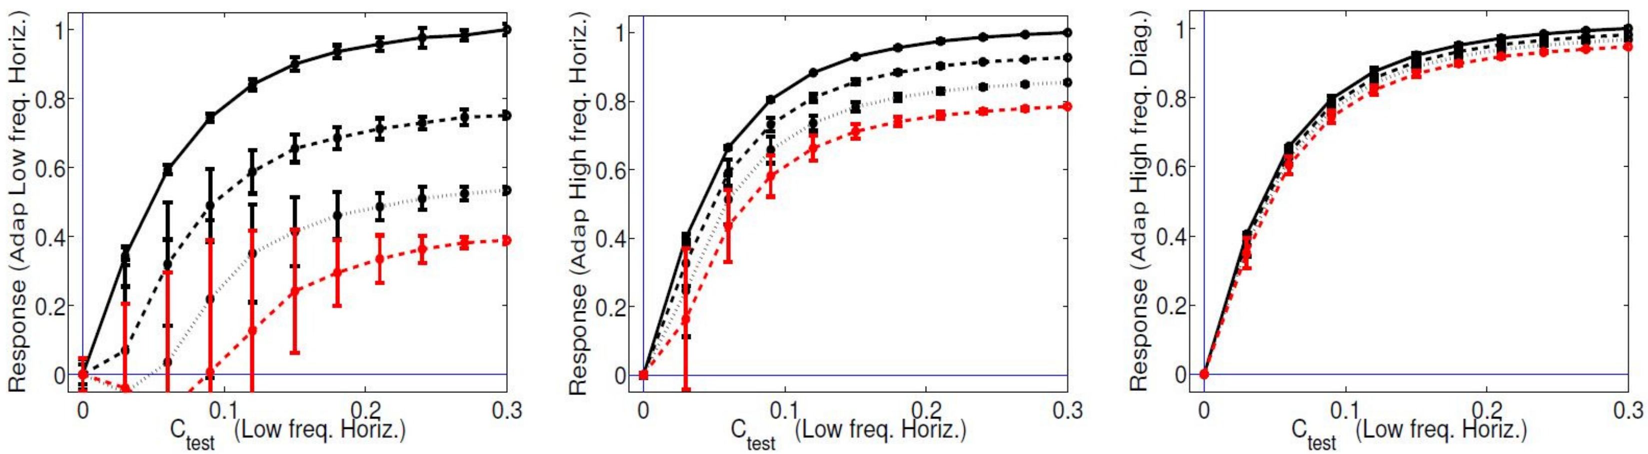

③ RESULTS I: Human illusions from INFOMAX and ERRORMIN RESULTS: MOTION

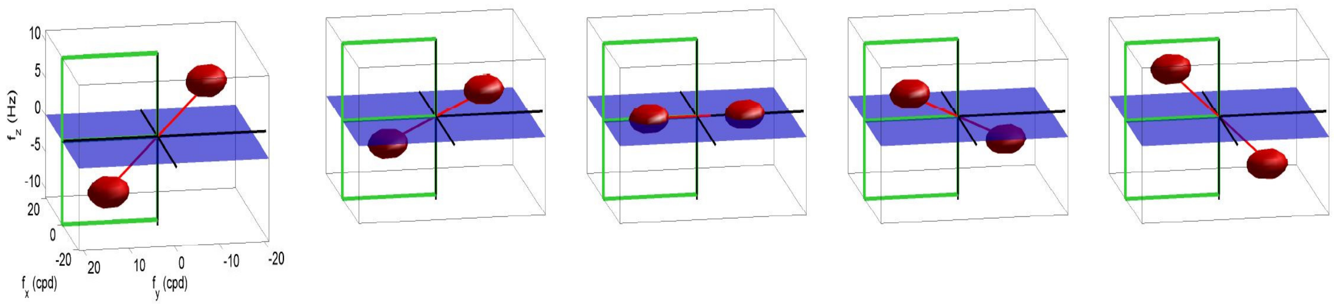

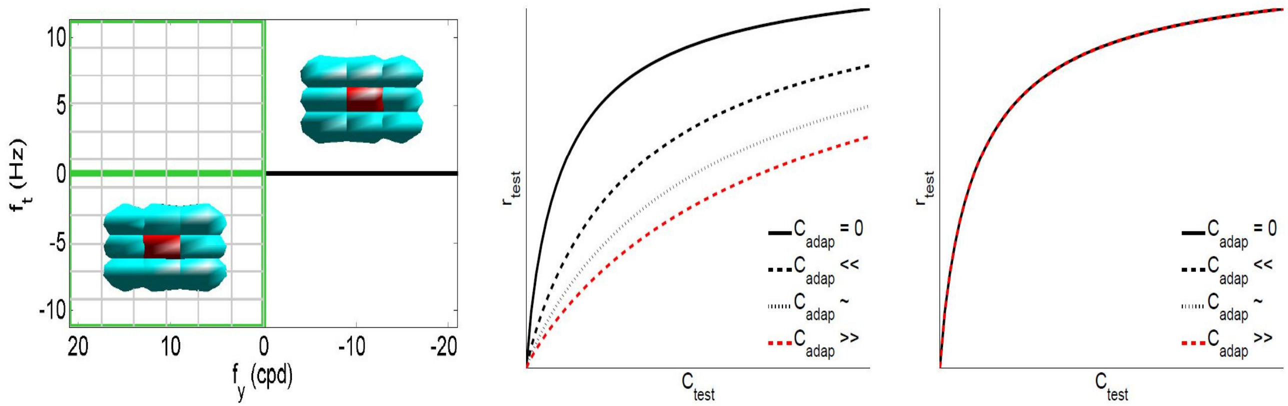

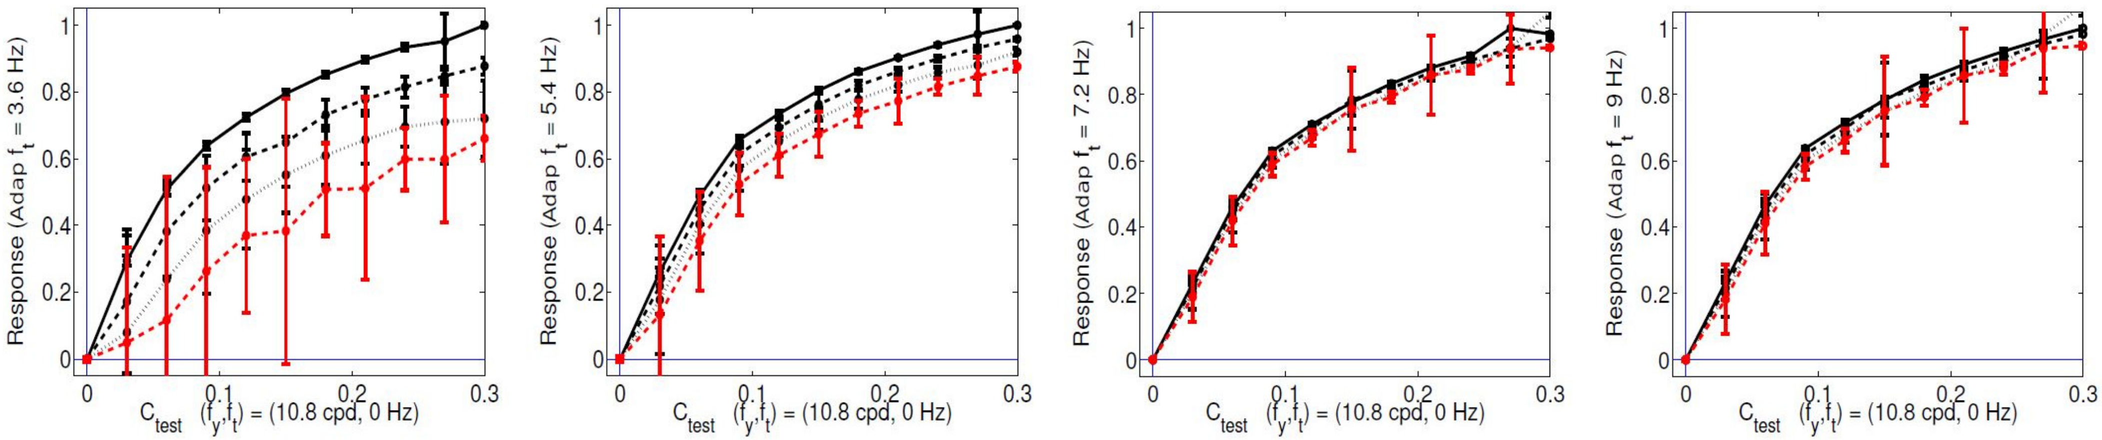

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN RESULTS: COLOR

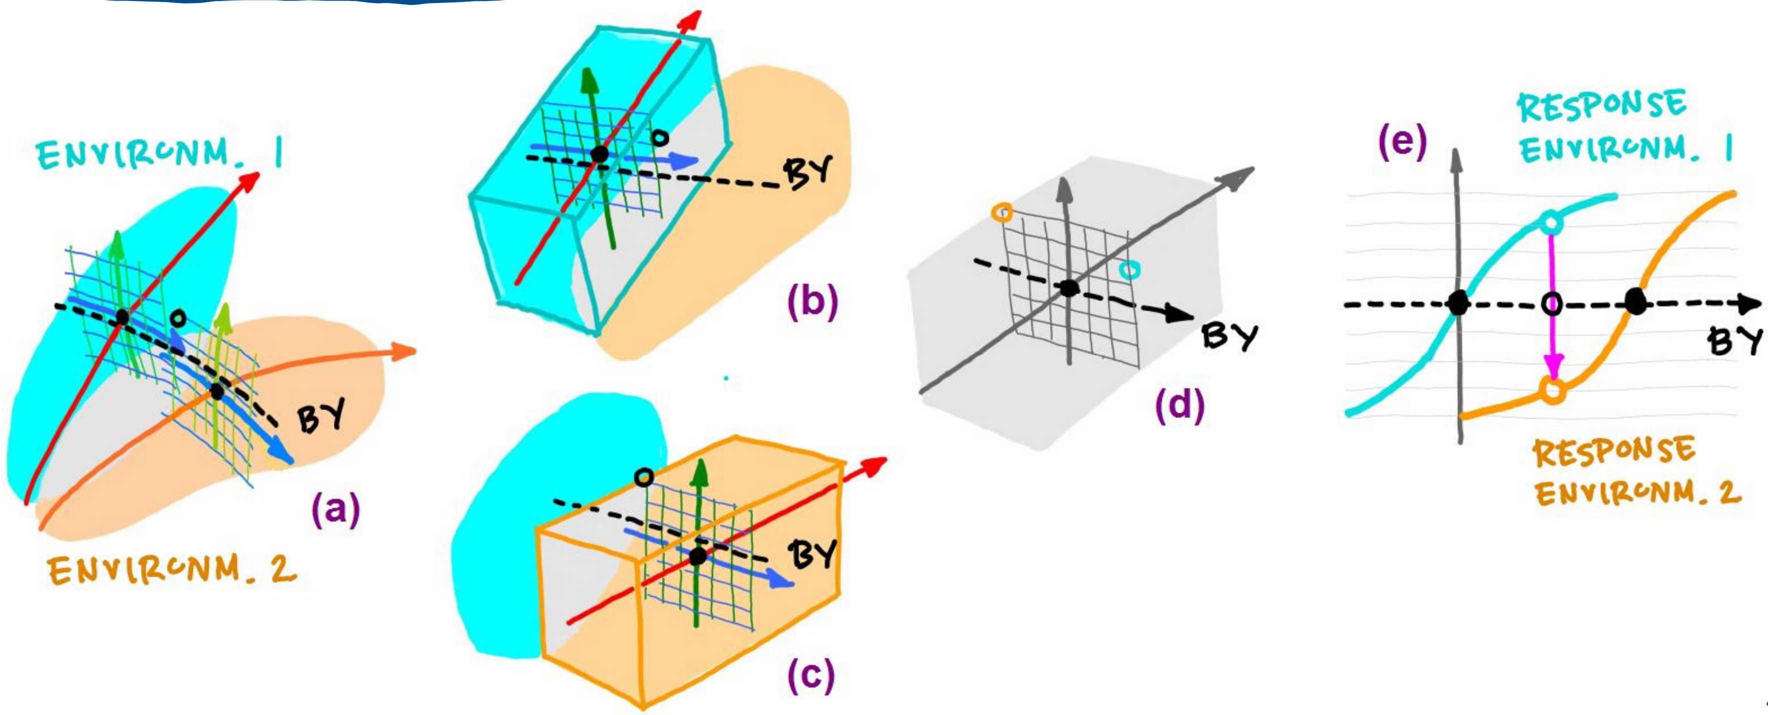

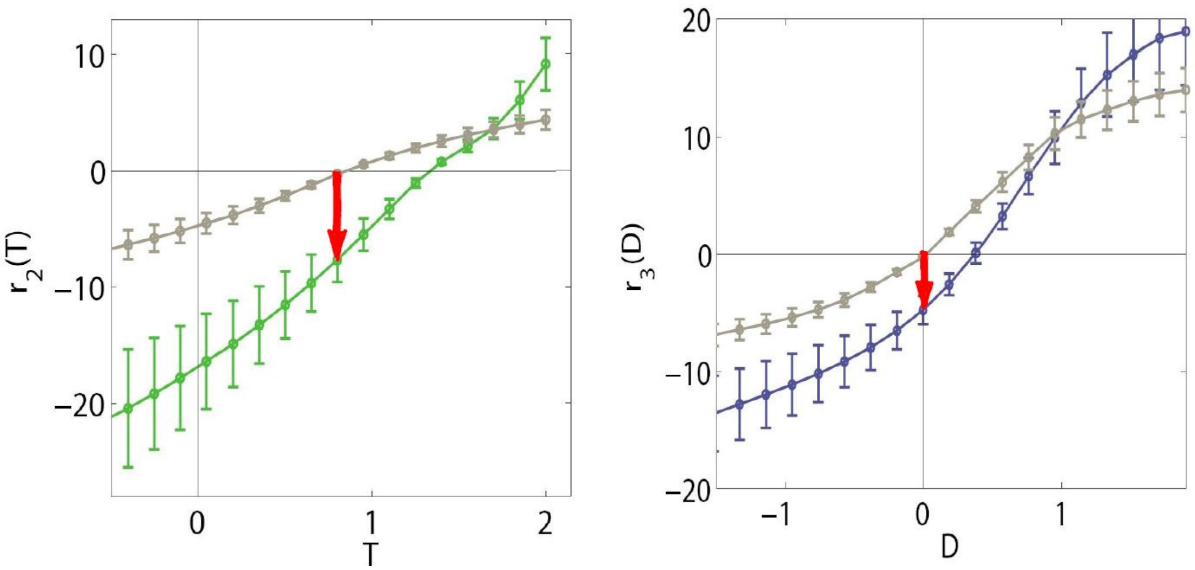

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN

SUMMARY RESULTS I

* Uniformization via SPCA explains aftereffects
  - Texture
  - Motion
  - Color

V. Laparra & J. Malo (2015) Visual aftereffects and sensory nonlinearities from a single statistical framework
Front. Hum. Neurosci., https://doi.org/10.3389/fnhum.2015.00557

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN

SUMMARY RESULTS I

* Uniformization via SPCA explains aftereffects
  - Texture
  - Motion
  - Color (ERRORMIN)

V. Laparra & J. Malo (2015) Visual aftereffects and sensory nonlinearities from a single statistical framework **Front. Hum. Neurosci.**, https://doi.org/10.3389/fnhum.2015.00557

V. Laparra, J. Malo et al. (2012) Nonlinearities and adaptation in color vision from Sequential Principal Curves Analysis **Neural Comput.** Vol. 24, 2751–2788. doi: 10.1162/NECO_a_00342

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN

SUMMARY RESULTS I

* Uniformization via SPCA explains aftereffects
  - Texture
  - Motion
  - Color (ERRORMIN)

* Uniformization goal is consistent with I and ΔT in psychophysical models

V. Laparra & J. Malo (2015) Visual aftereffects and sensory nonlinearities from a single statistical framework **Front. Hum. Neurosci.**, https://doi.org/10.3389/fnhum.2015.00557

V. Laparra, J. Malo et al. (2012) Nonlinearities and adaptation in color vision from Sequential Principal Curves Analysis **Neural Comput.** Vol. 24, 2751–2788. doi: 10.1162/NECO_a_00342

J. Malo & V. Laparra (2010) "Psychophysically-tuned Divisive Normalization factorizes the PDF of natural images" **Neural Comput.**, Vol. 22, 12:3179-3206

A. Gómez, M. Bertalmío & J. Malo (2019) Information Flow in Wilson-Cowan Networks Accepted in **J. Neurosphysiol**. Preprint in https://arxiv.org/abs/1907.13046

J. Malo (2019) Spatio-chromatic information available from different neural layers via Gaussianization Submitted to **J. Math. Neurosci.** https://arxiv.org/abs/1910.01559

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN

SUMMARY RESULTS I

* Uniformization via SPCA explains aftereffects
  - Texture
  - Motion
  - Color (ERRORMIN)

* Uniformization goal is consistent with I and ΔT in psychophysical models

Information available from different layers

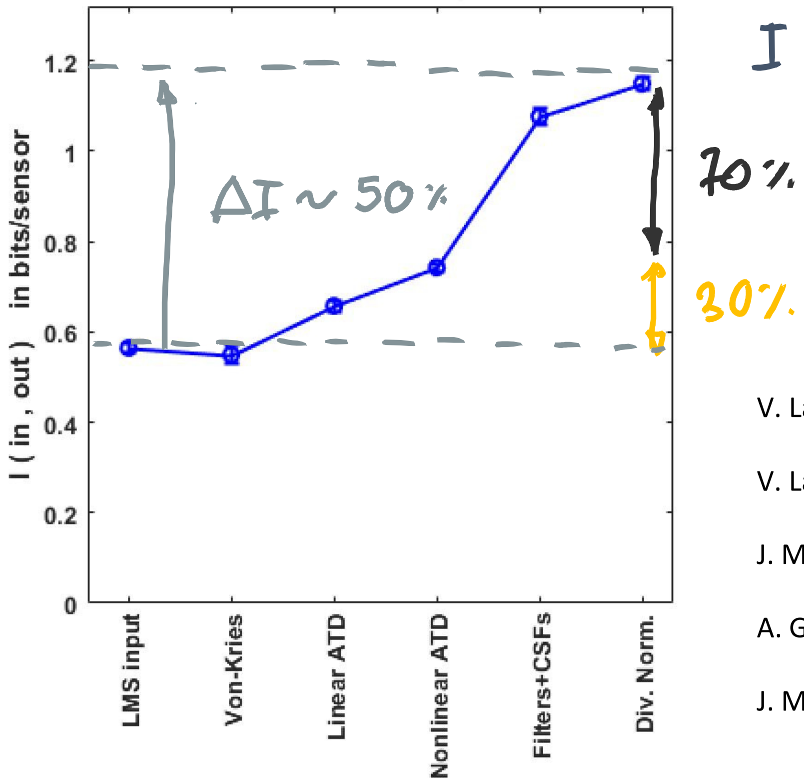

I and ΔT computed through Gaussianization!

	NO-SUBSAMPL.		
	$\Delta x = 0.24$	$\Delta x = 0.12$	$\Delta x = 0.06$
Divisive Norm.	$\mathbf{84 \pm 3}$	67 ± 3	54 ± 2
Wilson-Cowan	$\mathbf{85 \pm 4}$	67 ± 3	54 ± 1

V. Laparra & J. Malo (2015) Visual aftereffects and sensory nonlinearities from a single statistical framework **Front. Hum. Neurosci.**, https://doi.org/10.3389/fnhum.2015.00557

V. Laparra, J. Malo et al. (2012) Nonlinearities and adaptation in color vision from Sequential Principal Curves Analysis **Neural Comput.** Vol. 24, 2751–2788. doi: 10.1162/NECO_a_00342

J. Malo & V. Laparra (2010) "Psychophysically-tuned Divisive Normalization factorizes the PDF of natural images" **Neural Comput.**, Vol. 22, 12:3179-3206

A. Gómez, M. Bertalmío & J. Malo (2019) Information Flow in Wilson-Cowan Networks Accepted in **J. Neurosphysiol**. Preprint in https://arxiv.org/abs/1907.13046

J. Malo (2019) Spatio-chromatic information available from different neural layers via Gaussianization Submitted to **J. Math. Neurosci.** https://arxiv.org/abs/1910.01559

3 RESULTS I: Human illusions from INFOMAX and ERRORMIN

SUMMARY RESULTS I

* Uniformization via SPCA explains aftereffects { · Texture · Motion · Color (ERRORMIN)

* Uniformization goal is consistent with I and ΔT in psychophysical models

* INFOMAX & ERRORMIN GOALS EXPLAIN { NON-LINEARITIES / VISUAL ILLUSIONS

Information available
from different layers

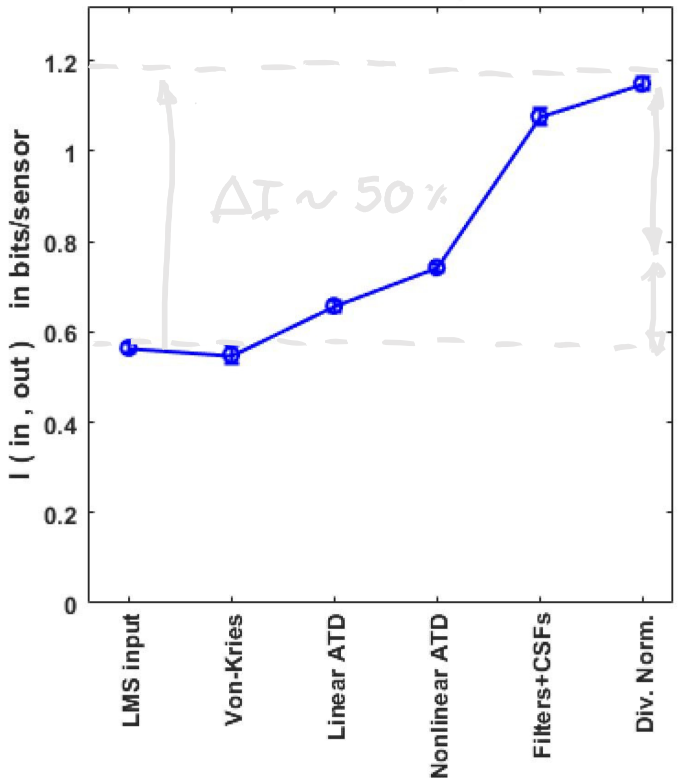

I and ΔT computed through Gaussianization!

	NO-SUBSAMPL.		
	$\Delta x = 0.24$	$\Delta x = 0.12$	$\Delta x = 0.06$
Divisive Norm.	$\mathbf{84 \pm 3}$	67 ± 3	54 ± 2
Wilson-Cowan	$\mathbf{85 \pm 4}$	67 ± 3	54 ± 1

V. Laparra & J. Malo (2015) Visual aftereffects and sensory nonlinearities from a single statistical framework
Front. Hum. Neurosci., https://doi.org/10.3389/fnhum.2015.00557

V. Laparra, J. Malo et al. (2012) Nonlinearities and adaptation in color vision from Sequential Principal Curves Analysis
Neural Comput. Vol. 24, 2751–2788. doi: 10.1162/NECO_a_00342

J. Malo & V. Laparra (2010) "Psychophysically-tuned Divisive Normalization factorizes the PDF of natural images"
Neural Comput., Vol. 22, 12:3179-3206

A. Gómez, M. Bertalmío & J. Malo (2019) Information Flow in Wilson-Cowan Networks
Accepted in **J. Neurosphysiol**. Preprint in https://arxiv.org/abs/1907.13046

J. Malo (2019) Spatio-chromatic information available from different neural layers via Gaussianization
Submitted to **J. Math. Neurosci.** https://arxiv.org/abs/1910.01559

4) Illusions in Convolutional Neural Networks

4.1 – Stimuli for experiments

4.2 – Learning low-level vision tasks

4.3 – Results { · Artificial Physiology
· Artificial Psychophysics

A. Gómez et al. (2018) "Convolutional Neural Networks Can Be Deceived by Visual Illusions".
Proc. IEEE Conf. Comp. Vision Patt. Recog. CVPR18, Prerpint in https://arxiv.org/abs/1811.10565
A. Gómez, J. Malo, et al. (2019) "Visual Illusions also deceive Convolutional Neural Networks: Analysis and Implications".
Submitted to **Vision Research** https://arxiv.org/abs/1912.01643

(4) Illusions in Convolutional Neural Networks   STIMULI FOR EXPERIMENTS

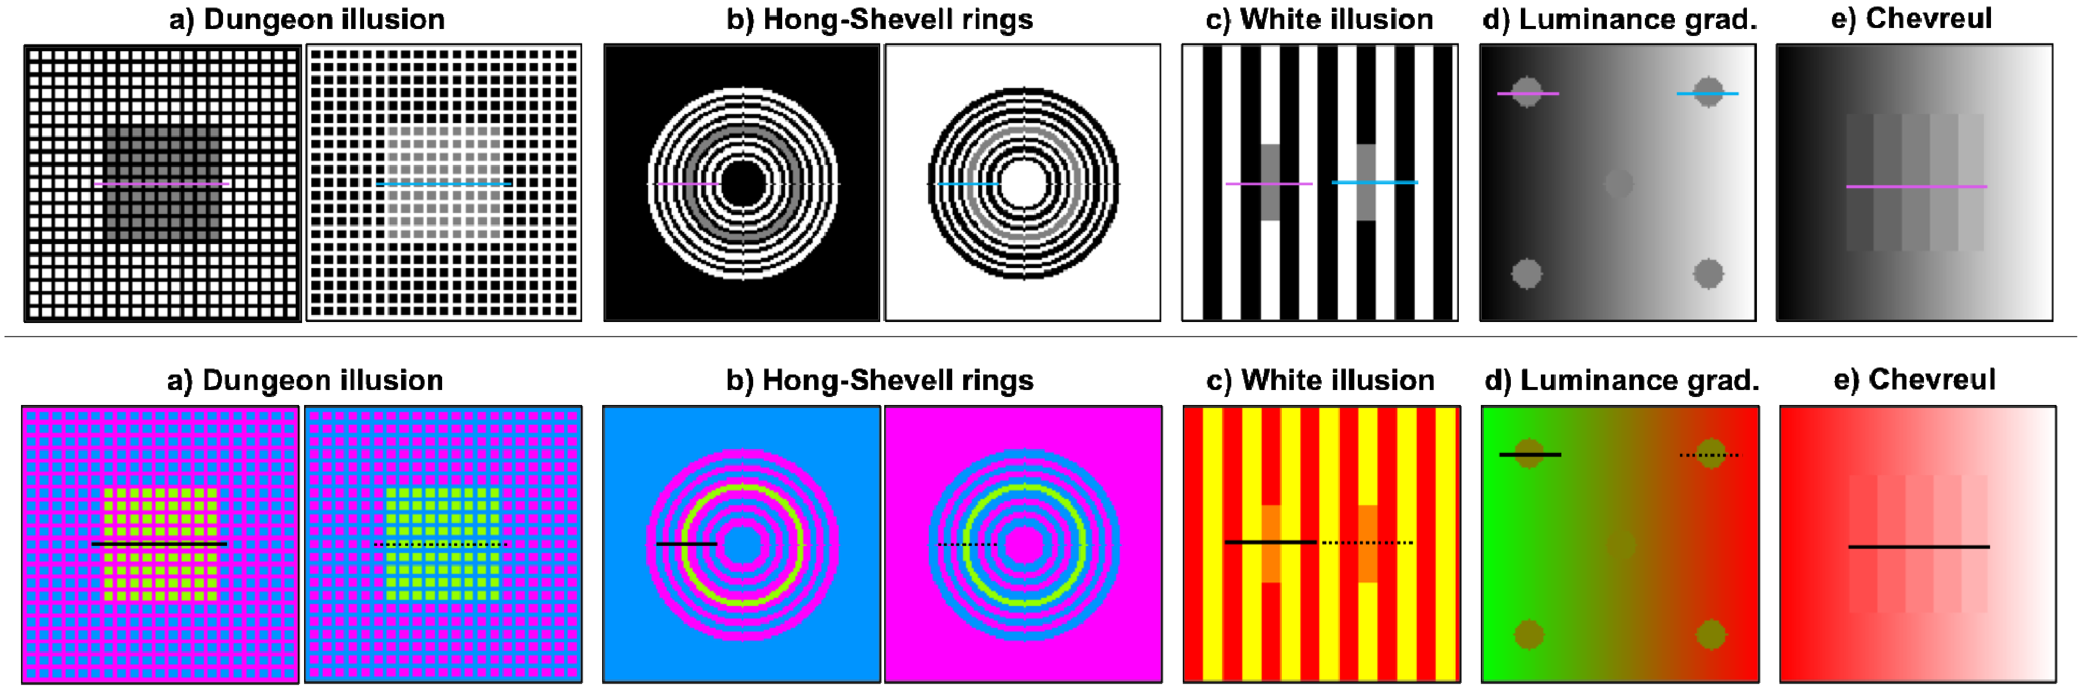

4 Illusions in Convolutional Neural Networks   STIMULI FOR EXPERIMENTS

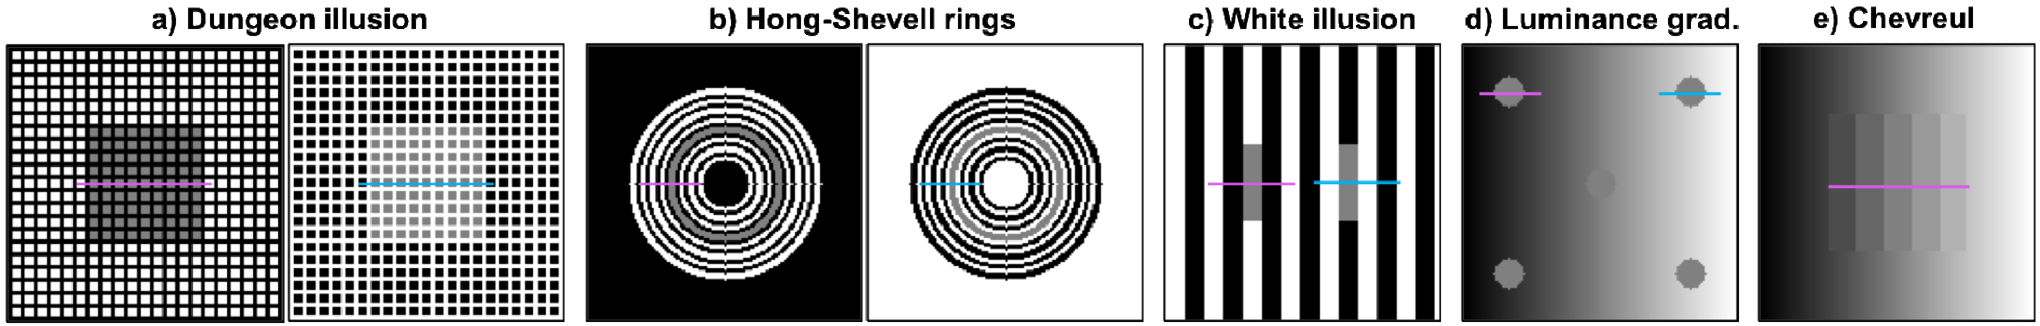

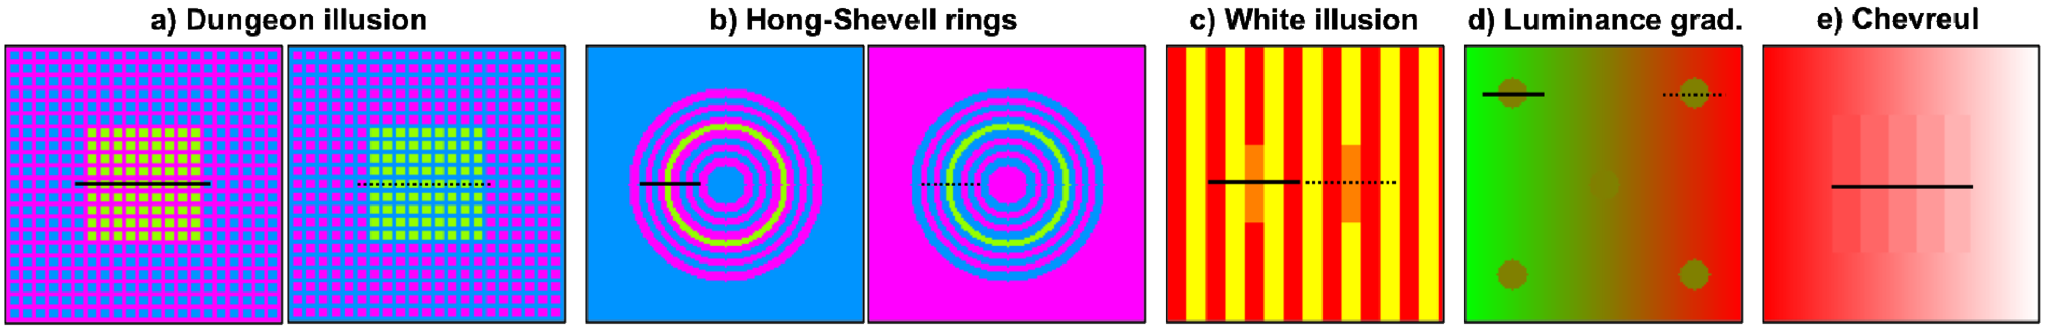

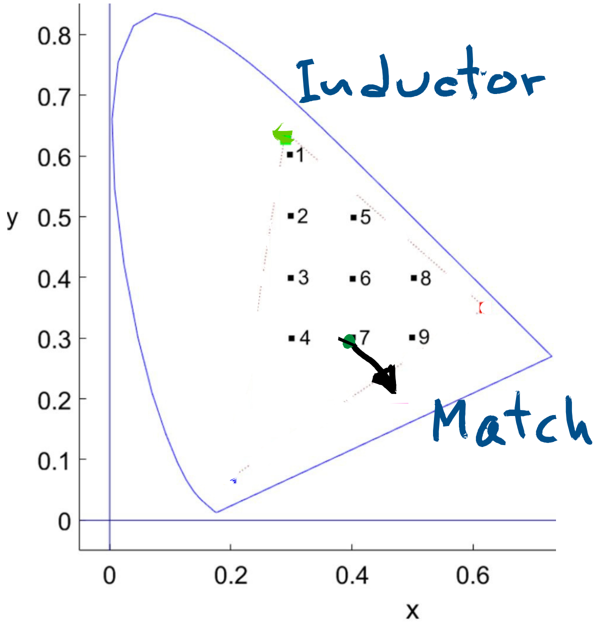

Ware & Cowan Vis. Res. 82
ASYMMETRIC COLOR MATCHING

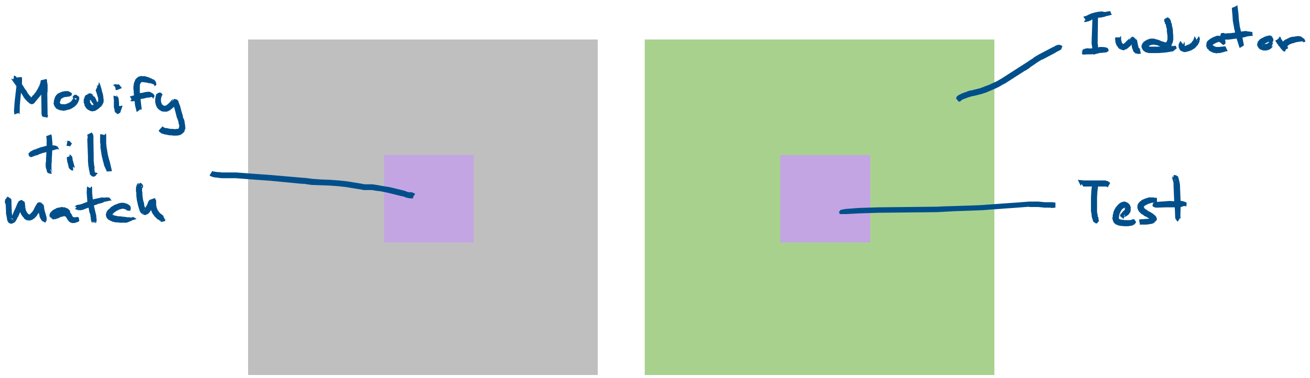

4 Illusions in Convolutional Neural Networks

CNNs LEARNING LOW-LEVEL VISION TASKS

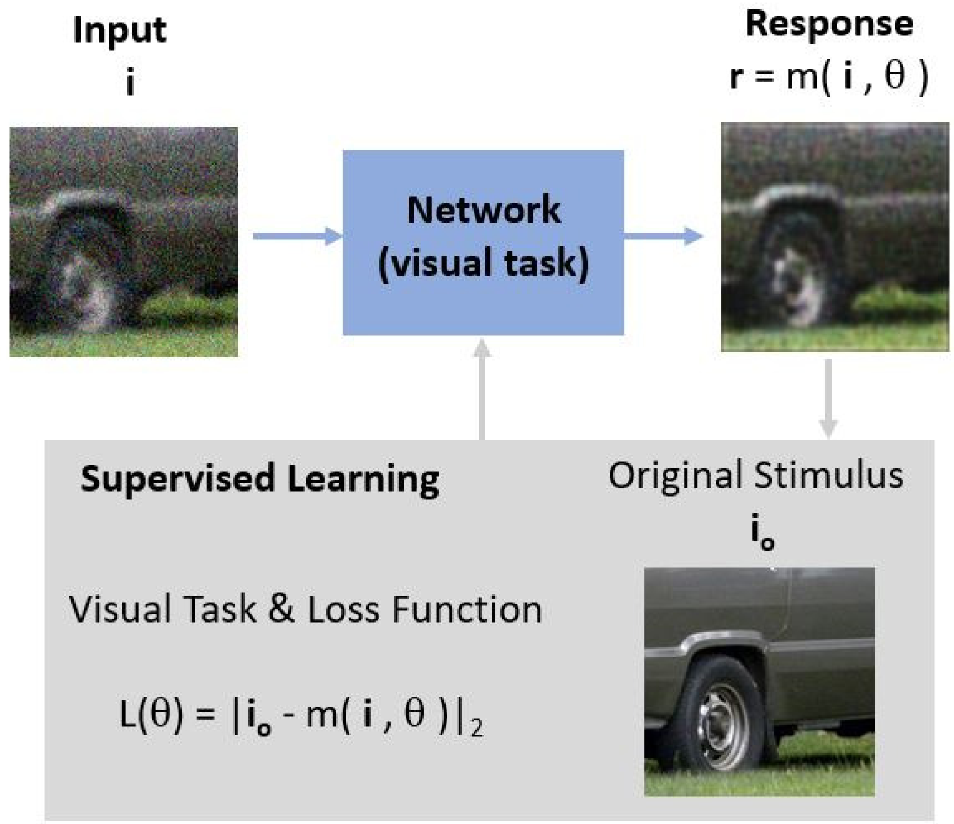

Learning
- Denoising
- Deblurring
- Restoration

Data sets
- Russlowsky IJCV 15
- Vazquez et al. Percept. 09
- Malo et al. Neur. Comp. 12

Architectures
- Shallow 2-layers 5×5 kernels
- Deep 4-layers 10×10 kernels
- Very Deep 20-layers [Zhang et al. CVPR 17]

(4) Illusions in Convolutional Neural Networks ARTIFICIAL PHYSIOLOGY VS PSYCHOPHYSICS

Artificial PHYSIOLOGY: Measure responses

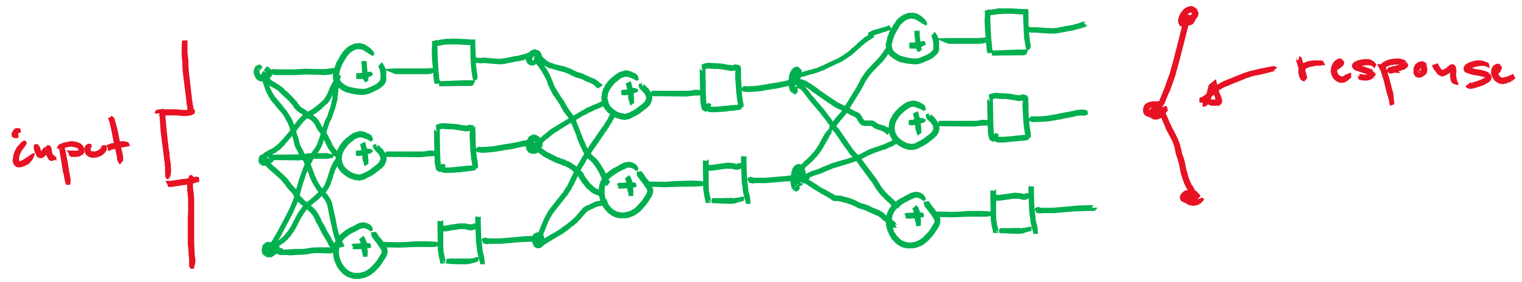

Artificial PSYCHOPHYSICS: Modify input to achieve match in response

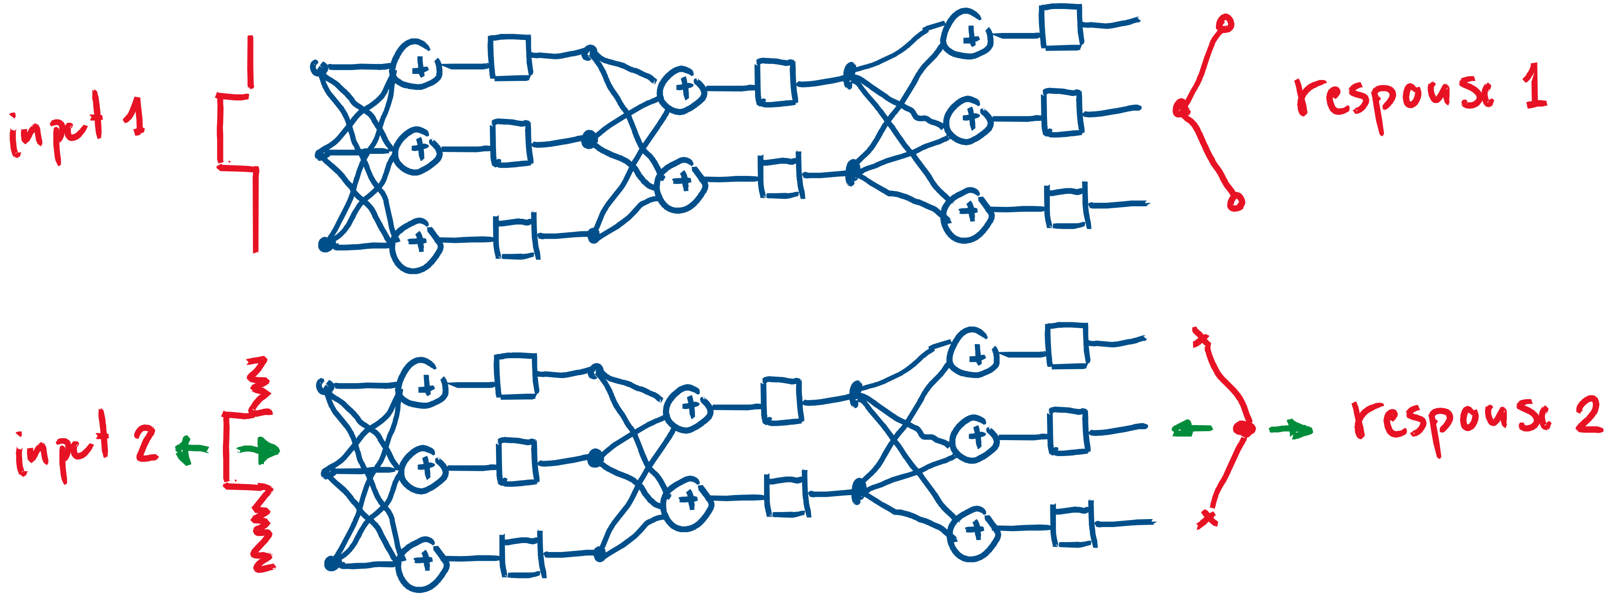

(4) Illusions in Convolutional Neural Networks

Results ARTIFICIAL PHYSIOLOGY

SHALLOW

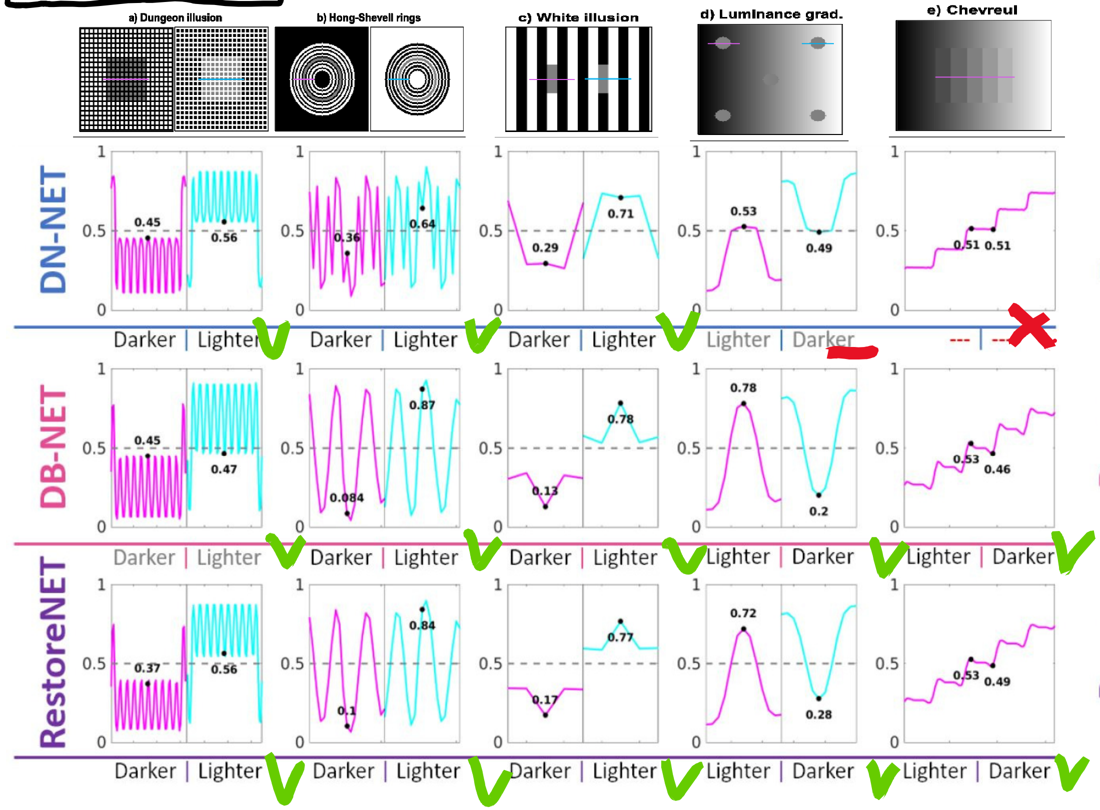

DEEP

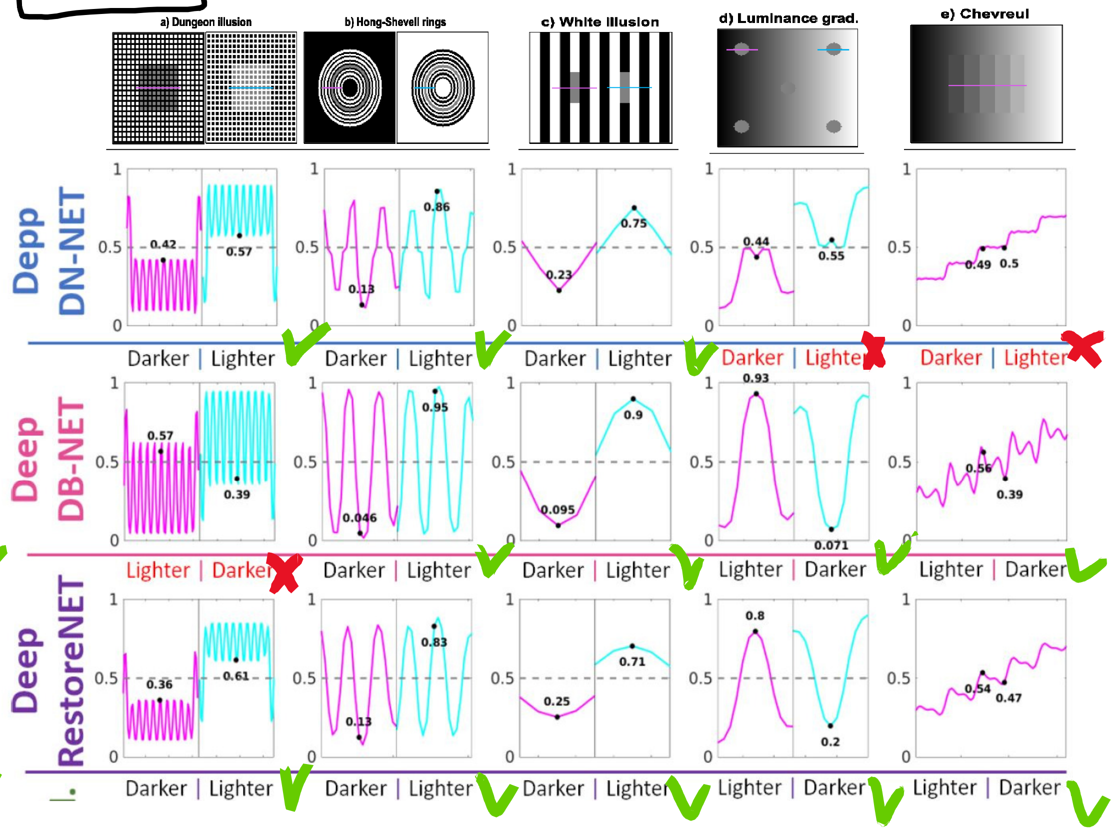

VERY DEEP

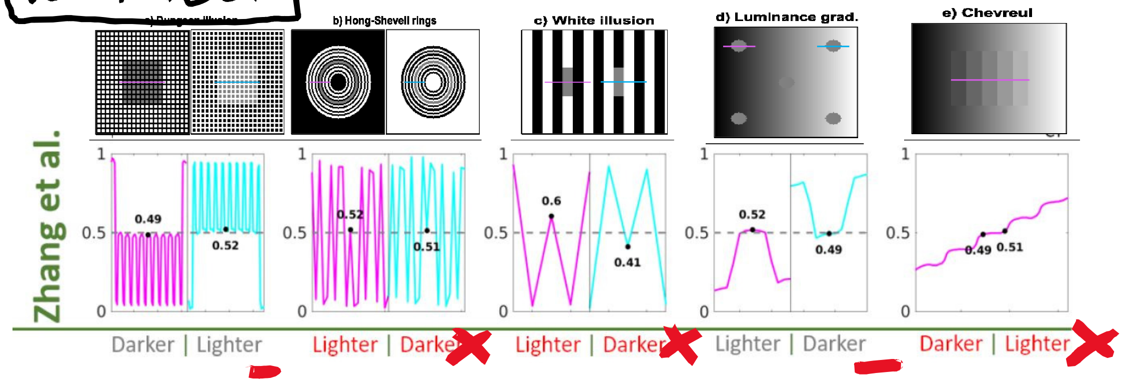

4 Illusions in Convolutional Neural Networks

Results ARTIFICIAL PHYSIOLOGY

SHALLOW

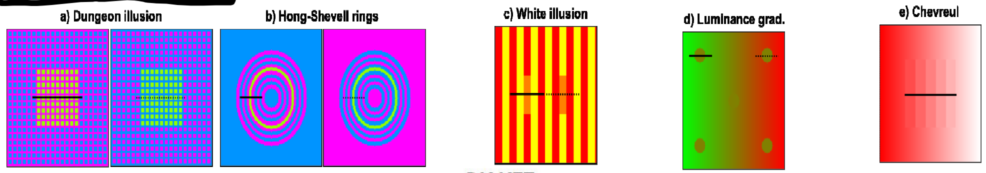

DN-NET

	Dungeon			Hong-Shevell			White			Gradient			Chevreul		
	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R
R	0.58	0.77	0.27	0.58	0.39	0.72	1	0.85	0.92	0.5	0.43	0.61	1	0.92	0.91
G	1	0.38	0.76	1	0.74	0.53	0.5	0.64	0.36	0.5	0.63	0.36	0.5	0.5	0.51
B	0	0.47	0.51	0	0.44	0.37	0	0.074	0.13	0	0.079	0.079	0.5	0.49	0.51

DB-NET

	Dungeon			Hong-Shevell			White			Gradient			Chevreul		
	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R
R	0.58	0.68	0.22	0.58	0.21	0.44	1	0.81	0.94	0.5	0.37	0.64	1	0.92	0.9
G	1	0.36	0.75	1	0.62	0.39	0.5	0.55	0.39	0.5	0.6	0.38	0.5	0.47	0.45
B	0	0.34	0.44	0	0.25	0.21	0	0.058	0.12	0	0.078	0.075	0.5	0.48	0.46

RestoreNET

	Dungeon			Hong-Shevell			White			Gradient			Chevreul		
	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R
R	0.58	0.77	0.24	0.58	0.3	0.67	1	0.85	0.93	0.5	0.36	0.64	1	0.92	0.9
G	1	0.36	0.76	1	0.64	0.41	0.5	0.55	0.39	0.5	0.6	0.37	0.5	0.49	0.48
B	0	0.43	0.53	0	0.38	0.34	0	0.057	0.11	0	0.073	0.075	0.5	0.48	0.47

DEEP

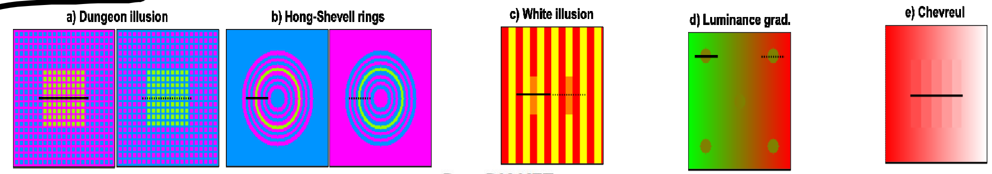

Deep DN-NET

	Dungeon			Hong-Shevell			White			Gradient			Chevreul		
	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R
R	0.58	0.77	0.27	0.58	0.42	0.63	1	0.94	0.97	0.5	0.47	0.49	1	0.96	0.95
G	1	0.38	0.72	1	0.57	0.4	0.5	0.61	0.38	0.5	0.54	0.48	0.5	0.49	0.49
B	0	0.5	0.57	0	0.56	0.49	0	0.028	0.056	0	0.027	0.03	0.5	0.52	0.52

Deep DB-NET

	Dungeon			Hong-Shevell			White			Gradient			Chevreul		
	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R
R	0.58	0.67	0.29	0.58	0.36	0.45	1	0.89	0.97	0.5	0.5	0.45	1	0.95	0.92
G	1	0.27	0.72	1	0.51	0.25	0.5	0.48	0.47	0.5	0.56	0.41	0.5	0.52	0.45
B	0	0.41	0.57	0	0.52	0.36	0	0.024	0.094	0	0.068	0.052	0.5	0.53	0.46

Deep RestoreNET

	Dungeon			Hong-Shevell			White			Gradient			Chevreul		
	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R
R	0.58	0.74	0.28	0.58	0.31	0.58	1	0.93	0.97	0.5	0.48	0.55	1	0.93	0.91
G	1	0.39	0.7	1	0.5	0.42	0.5	0.51	0.39	0.5	0.48	0.46	0.5	0.53	0.5
B	0	0.49	0.56	0	0.49	0.53	0	0.036	0.075	0	0.074	0.073	0.5	0.5	0.48

VERY DEEP

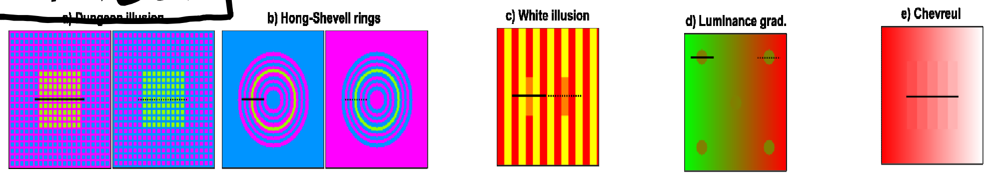

Zhang et al.

	Dungeon			Hong-Shevell			White			Gradient			Chevreul		
	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R	In	Out-L	Out-R
R	0.58	0.59	0.56	0.58	0.58	0.59	1	1	1	0.5	0.52	0.49	1	1	1
G	1	0.98	0.98	1	0.99	0.98	0.5	0.5	0.49	0.5	0.49	0.51	0.5	0.48	0.51
B	0	0.027	0.027	0	0.02	0.02	0	0.012	0.012	0	0.012	0.012	0.5	0.48	0.51

4) Illusions in Convolutional Neural Networks

Results ARTIFICIAL PSYCHOPHYSICS

SHALLOW

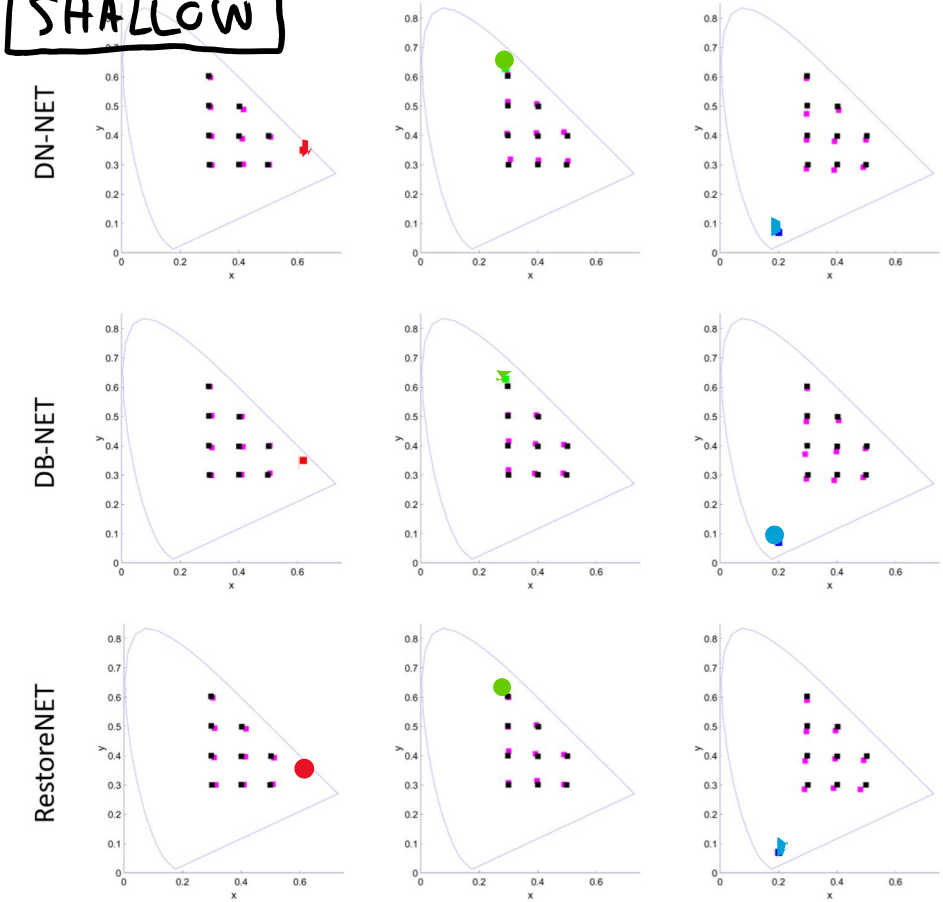

DEEP

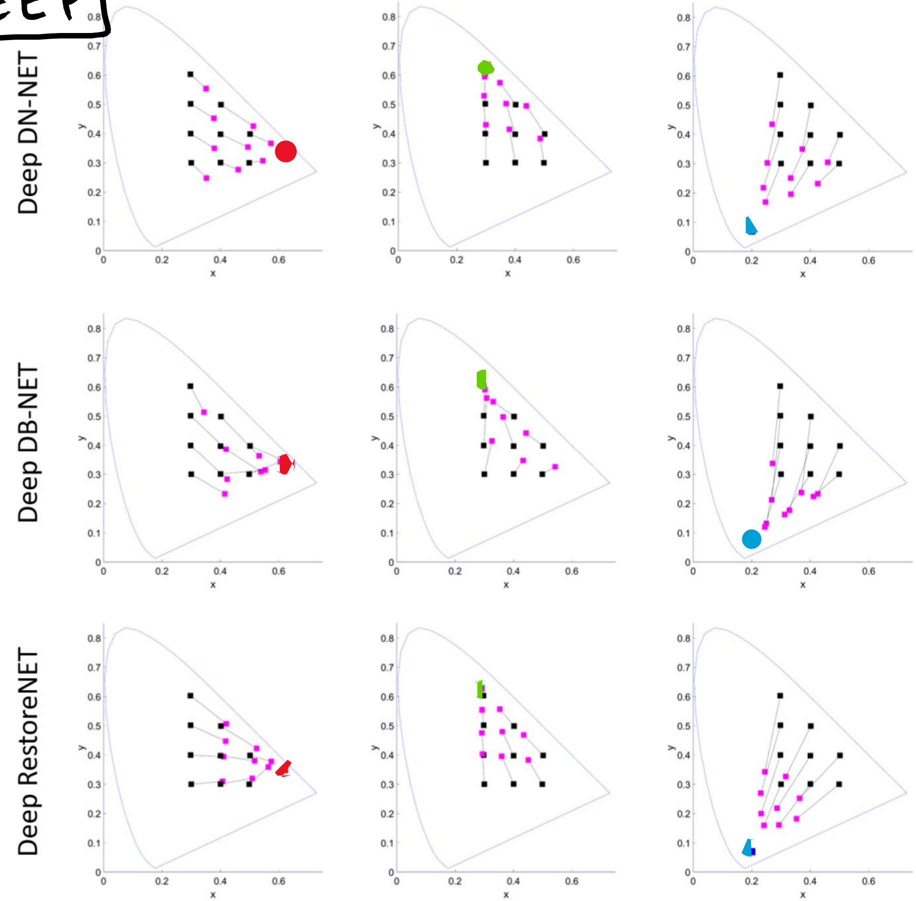

VERY DEEP

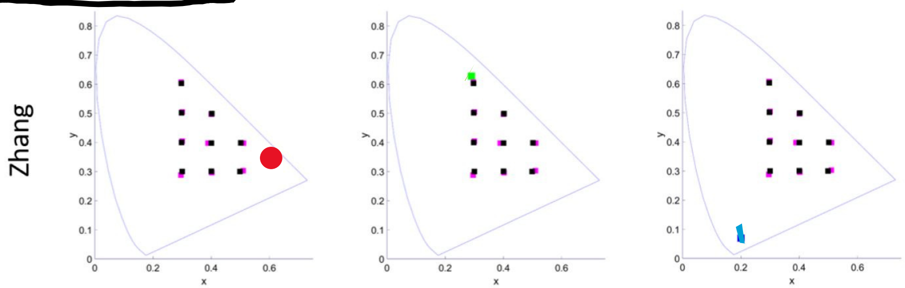

HUMANS

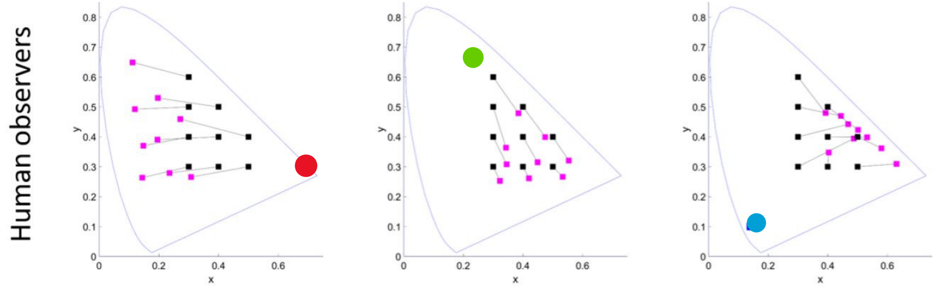

4 Illusions in Convolutional Neural Networks

SUMMARY { ARTIFICIAL PHYSIOLOGY
ARTIFICIAL PSYCHOPHYSICS

* CNNs do have visual illusions

* Visual illusions of CNNs are similar but not like ours

  - Artificial physiol. ⇒ shifts in resp. ~ ok 75%
  - Artificial psychophys. ⇒ Assimilation vs contrast

* Complexity of architecture is an issue

4 Illusions in Convolutional Neural Networks

SUMMARY { ARTIFICIAL PHYSIOLOGY
ARTIFICIAL PSYCHOPHYSICS

* CNNs do have visual illusions

* Visual illusions of CNNs are similar but not like ours

  - Artificial physiol. ⇒ shifts in resp. ~ ok 75%
  - Artificial psychophys. ⇒ Assimilation vs contrast

* Complexity of architecture is an issue

WHY?

5 RESULTS II: Artificial illusions from "artificial" CSFs

5.1 - Linearization analysis

5.2 - Jacobian, eigenvalues & eigenvectors

5.3 - Artificial CSFs

5 RESULTS II: Artificial illusions from "artificial" CSFs

5.1 - Linearization analysis

$$\underset{\text{input}}{i} \xrightarrow{m(i,\theta)} \underset{\text{response}}{r}$$

$$r = m(i, \theta)$$

5 RESULTS II: Artificial illusions from "artificial" CSFs   5.1 - Linearization analysis

$$\underset{\text{input}}{i} \xrightarrow{m(i,\theta)} \underset{\text{response}}{r}$$

$$r = m(i, \theta)$$

$$r = m(0 + i, \theta) \simeq \cancelto{0}{m(0, \theta)} + \nabla_i m(0) \cdot i$$

$$r = \underbrace{\nabla_i m(0)}_{\text{Jacobian at } 0} \cdot i$$

5 RESULTS II: Artificial illusions from "artificial" CSFs

5.1 – Linearization analysis

$$\underset{\text{input}}{i} \xrightarrow{m(i,\theta)} \underset{\text{response}}{r}$$

$$r = m(i, \theta)$$

$$r = m(0 + i, \theta) \simeq \cancel{m(0, \theta)}^{\,0} + \nabla_i m(0) \cdot i$$

$$r = \underbrace{\nabla_i m(0)}_{\text{Jacobian at } 0} \cdot i$$

- Analytic Martinez, Malo et al. PLOS 18
- Autograd
- Linear regression $\nabla_i m(0) = R \cdot I^{+}$

$$R = [\, r^{(1)} \; r^{(2)} \; \dots \; r^{(n)} \,] \quad I = [\, i^{(1)} \; i^{(2)} \; \dots \; i^{(n)} \,]$$

5 RESULTS II: Artificial illusions from "artificial" CSFs

5.1 - Linearization analysis

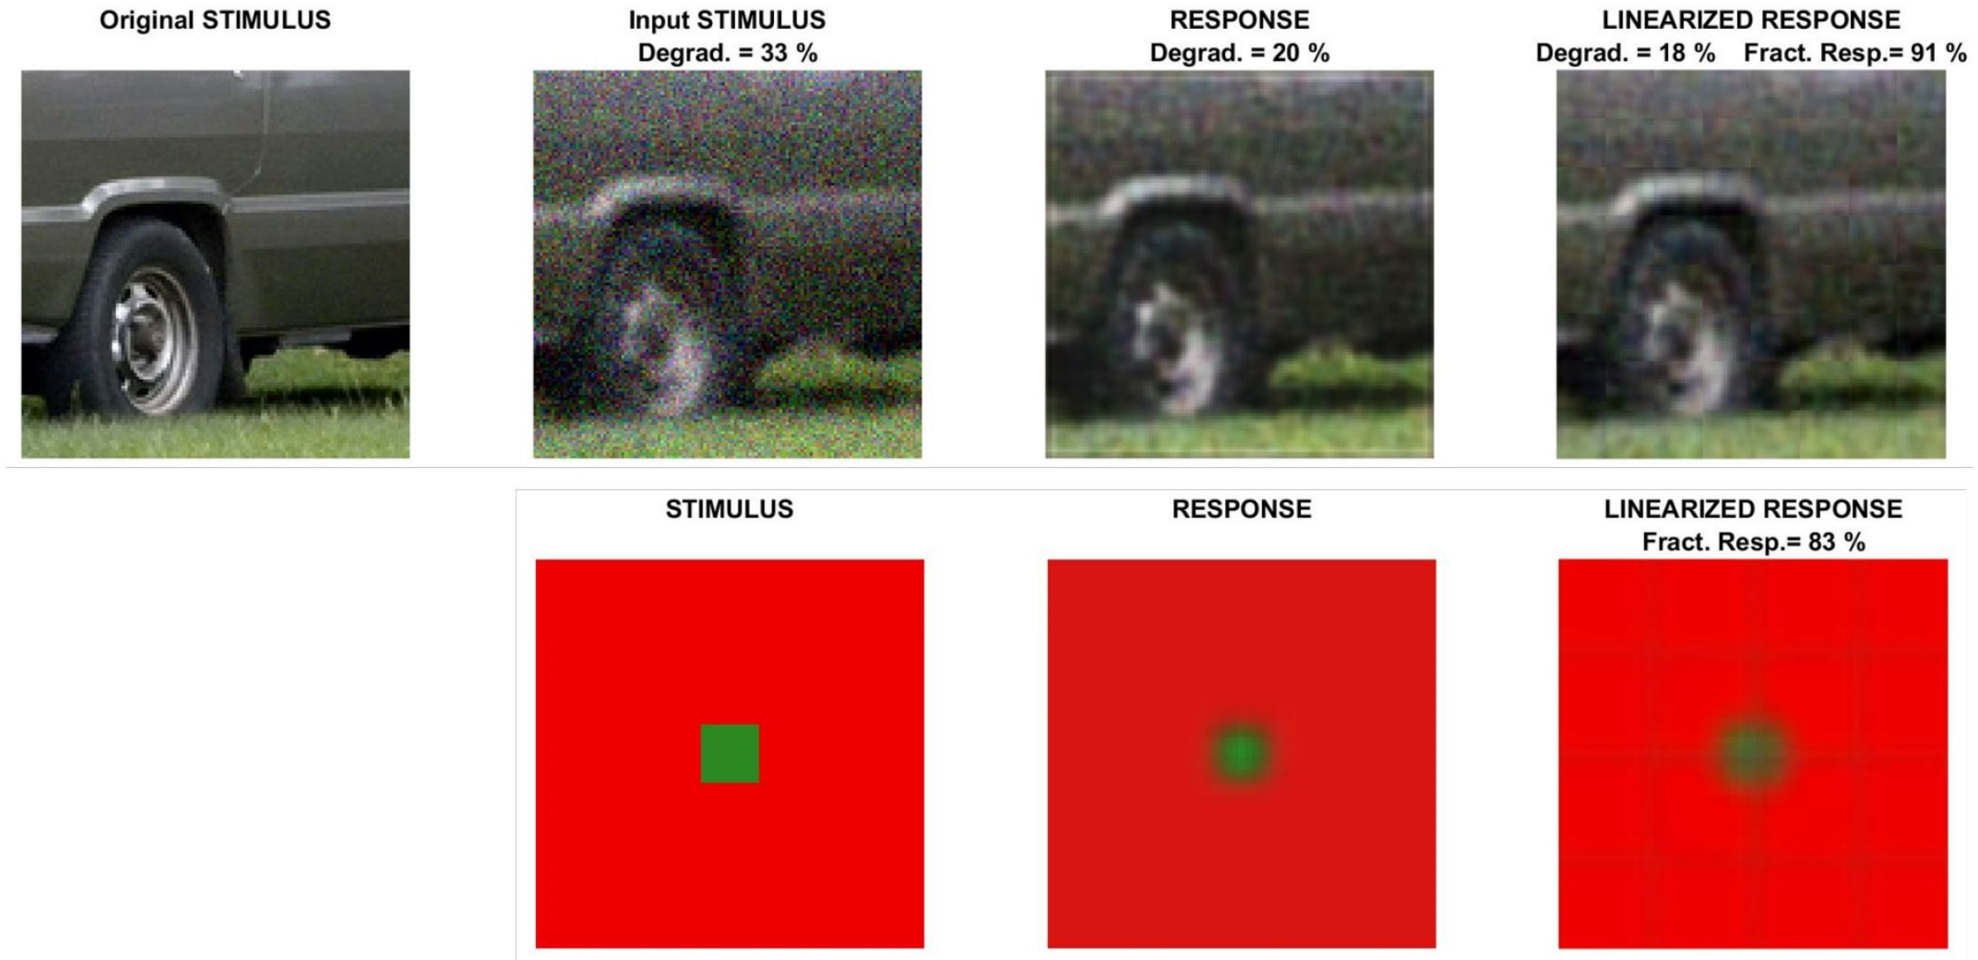

5 RESULTS II: Artificial illusions from "artificial" CSFs 5.2 - Jacobian, eigenvalues & eigenvectors

Jacobian

$D_i m =$

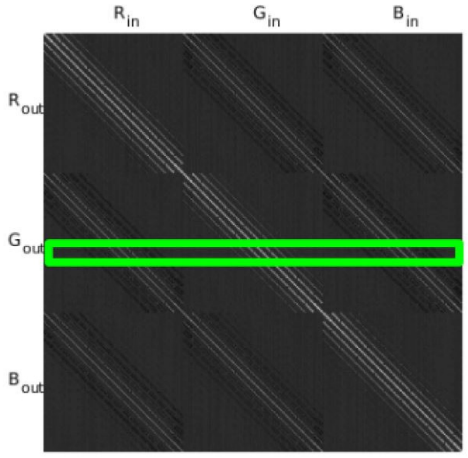

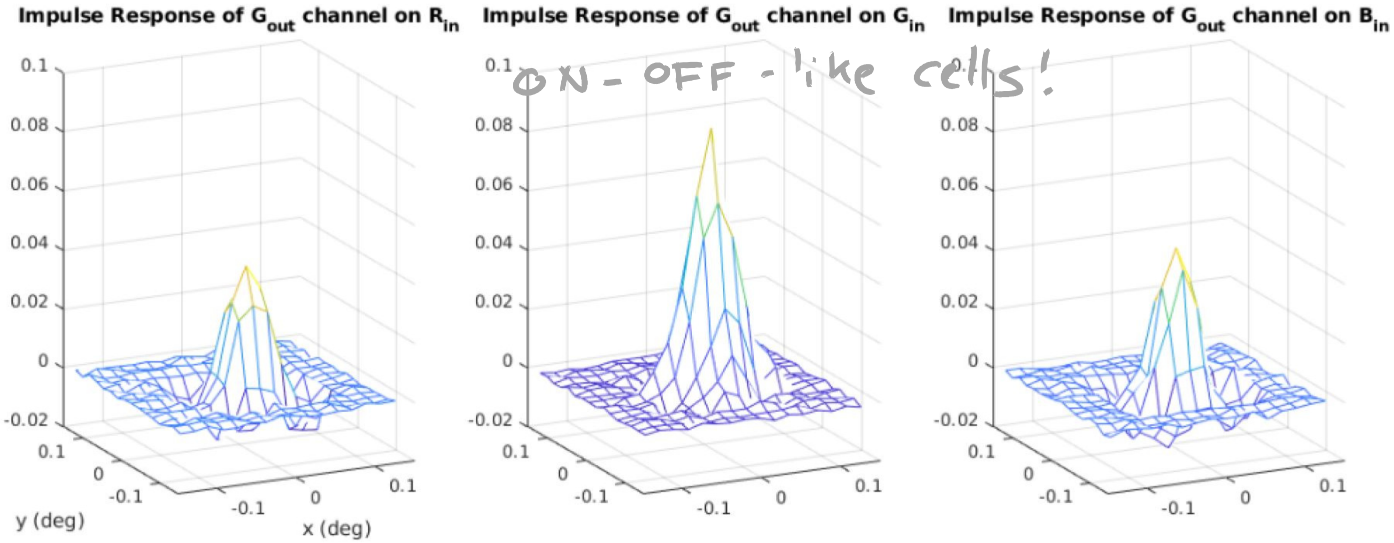

ON - OFF - like cells!

Eigen vectors

(~ Fourier & opponents)

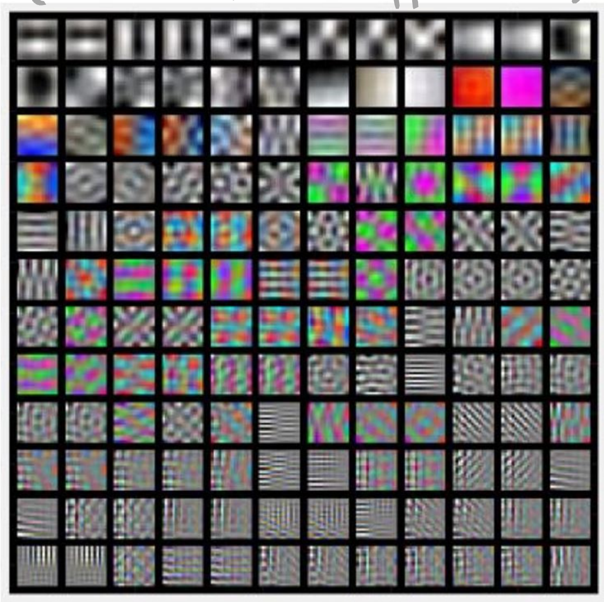

Eigen values

(~ low-pass)

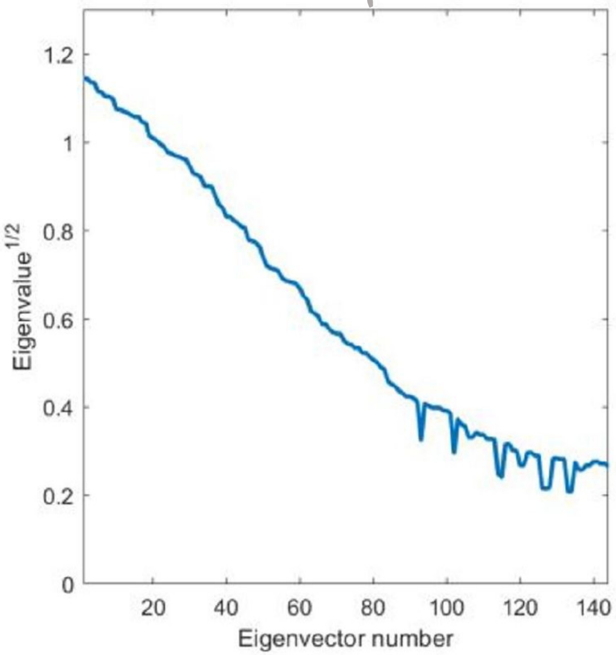

Color eigen-directions

RG - YB !

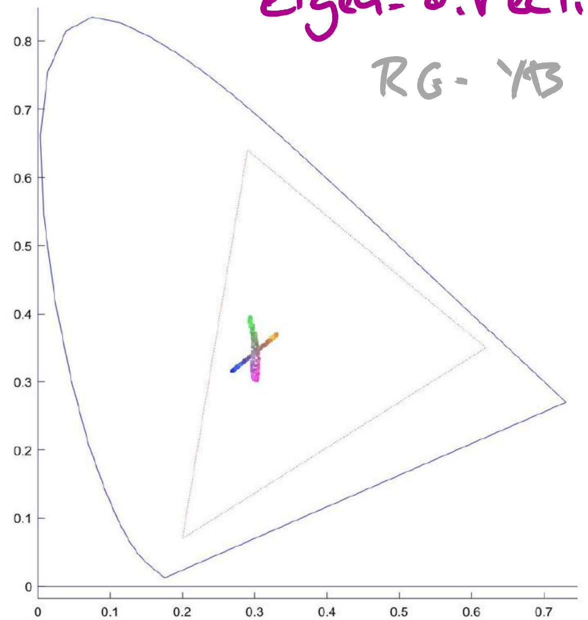

5 RESULTS II: Artificial illusions from "artificial" CSFs

5.3 – Artificial CSFs

ARTIFICIAL SPECTRAL SENSITIV.

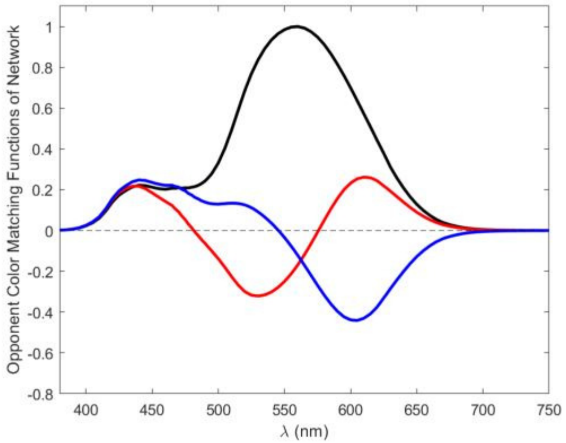

HUMAN SPECTRAL SENSITIV.

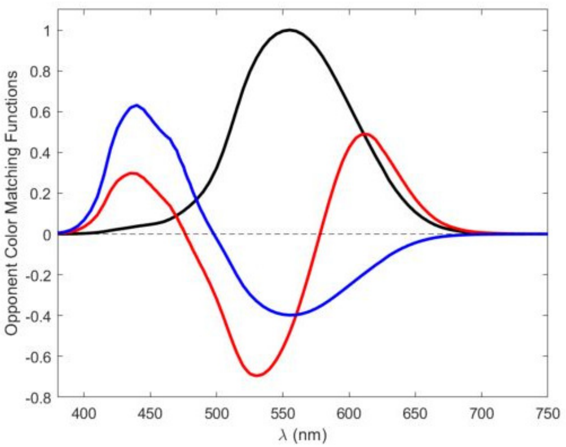

ARTIFICIAL SPATIAL FILTERS

Denoise

Deblur

Restore

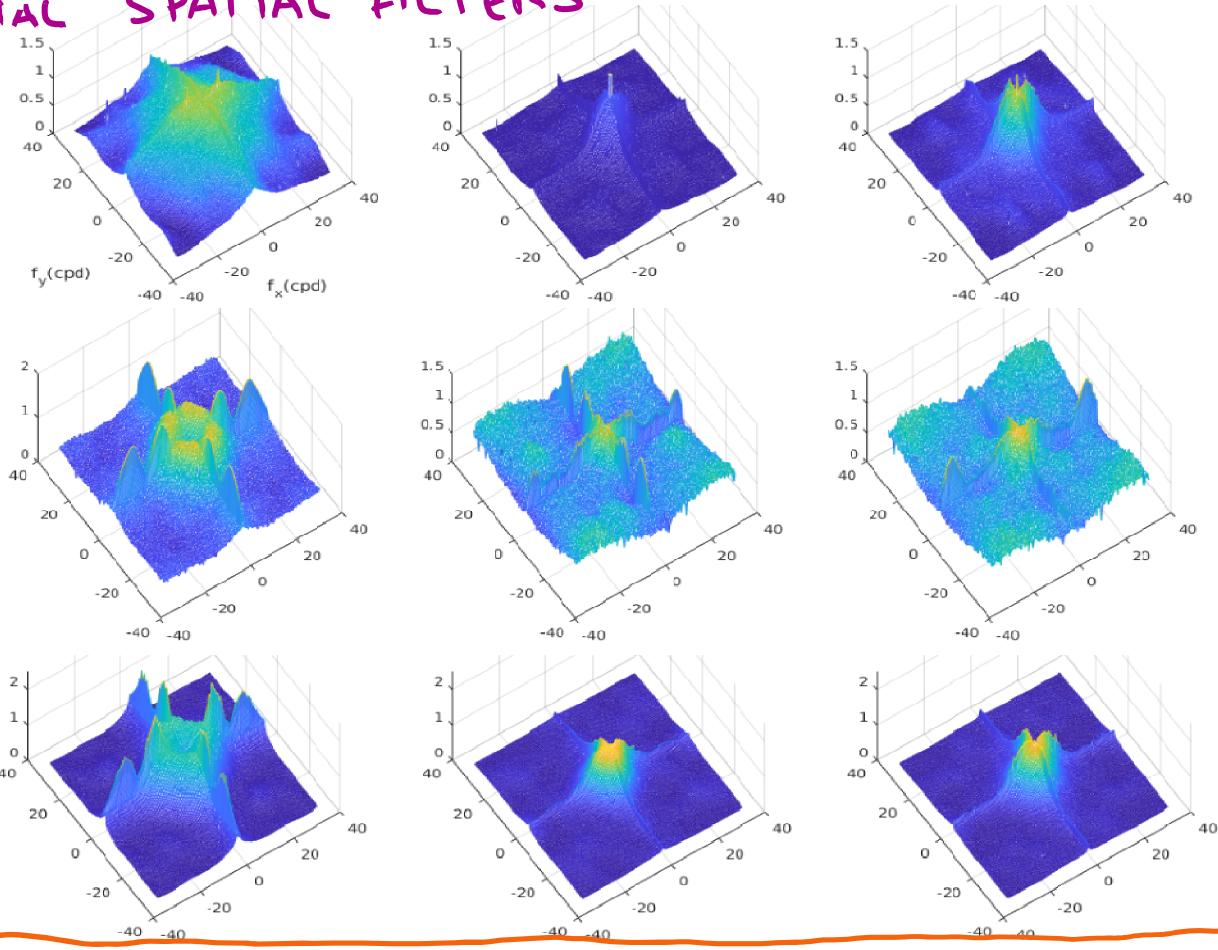

HUMAN CONTRAST SENSITIV. (CSFs)

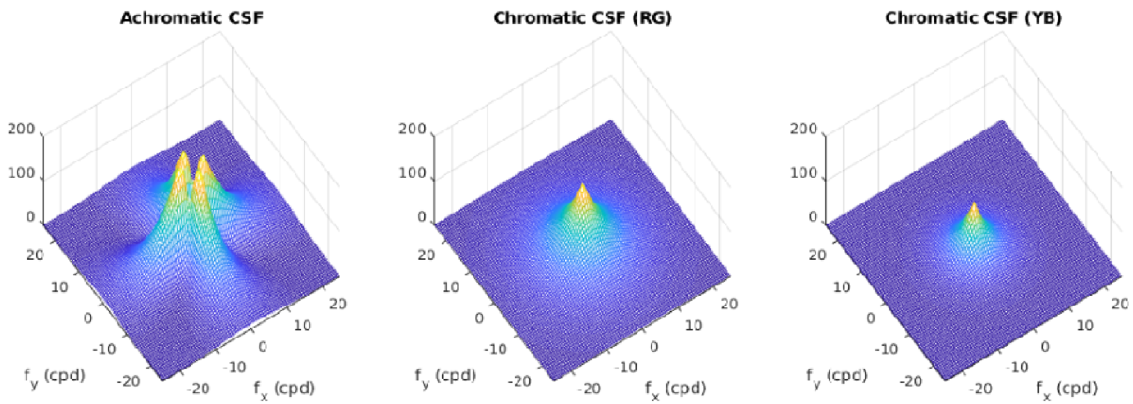

6 DISCUSSION & CONCLUSIONS

* CNNs for low-level vision develop a number of human-like features
  - On-off cells
  - Opponent color channels RG, YB
  - Human-like spatial filters (CSFs) in Achrom. RG, YB

A. Gómez, J. Malo, et al. (2019) "Visual Illusions also deceive Convolutional Neural Networks: Analysis and Implications". Submitted to **Vision Research** https://arxiv.org/abs/1912.01643

6 DISCUSSION & CONCLUSIONS

* CNNs for low-level vision develop a number of human-like features
  - On-off cells
  - Opponent color channels RG, YB
  - Human-like spatial filters (CSFs) in Achrom. RG, YB

* Why? ≡ error-minimization goal (Wiener restorat.) ⇒ fit to spectrum of natural images

  Atick Li Neur. Comp. 92
  Malo et al. IEEE TIP 06

A. Gómez, J. Malo, et al. (2019) "Visual Illusions also deceive Convolutional Neural Networks: Analysis and Implications". Submitted to **Vision Research** https://arxiv.org/abs/1912.01643

6 DISCUSSION & CONCLUSIONS

* CNNs for low-level vision develop a number of human-like features
  - On-off cells
  - Opponent color channels RG, YB
  - Human-like spatial filters (CSFs) in Achrom. RG, YB

* Why? = error-minimization goal (Wiener restorat.) ⇒ fit to spectrum of natural images — Atick Li Neur. Comp. 92; Malo et al. IEEE TIP 06

* (Generality/illusions, simple architect.) vs (Specificity/no-illusions, complex architect.)

	Shallow	Deep	Zhang et al.
Fract. Lin. Resp.	94%	95%	40%
Performance (error)	11%	10%	2%
Illusion Strength	++	+++	-

A. Gómez, J. Malo, et al. (2019) "Visual Illusions also deceive Convolutional Neural Networks: Analysis and Implications".
Submitted to **Vision Research** https://arxiv.org/abs/1912.01643

C) DISCUSSION & CONCLUSIONS

* CNNs for low-level vision develop a number of human-like features

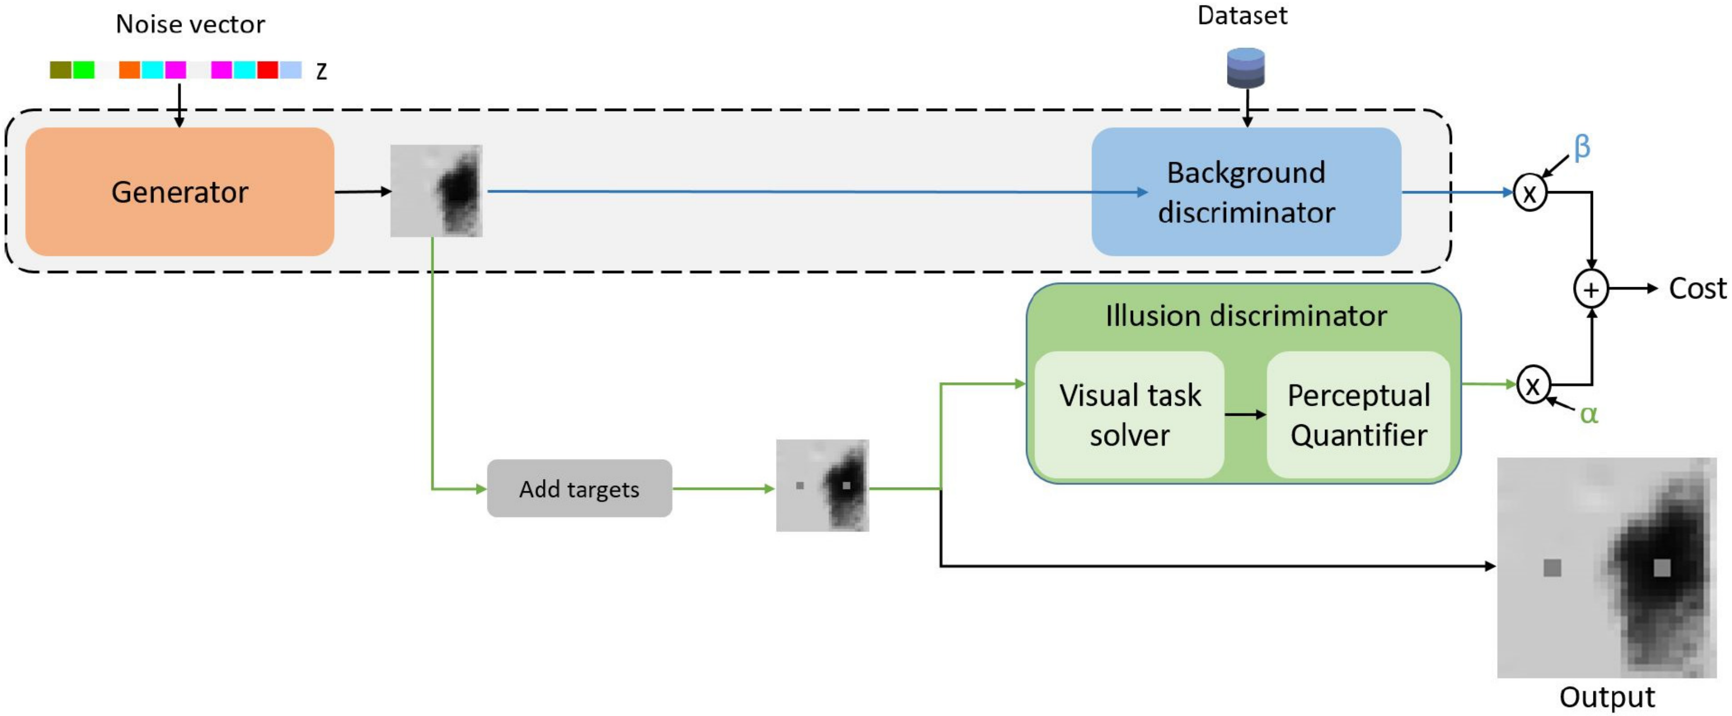

* Vision models & img. stats. in synthesizing visual illusions with GANs

A. Gómez, J. Malo, et al. (2019) "Visual Illusions also deceive Convolutional Neural Networks: Analysis and Implications". Submitted to **Vision Research** https://arxiv.org/abs/1912.01643

A. Gómez, J. Malo, et al. (2019) "Synthesizing Visual Illusions Using Generative Adversarial Networks" Submitted to **IEEE CVPR 19'** https://arxiv.org/abs/1911.09599

6 DISCUSSION & CONCLUSIONS

* Focus on the objective functions! (the "why" question)

nature neuroscience

FOCUS | PERSPECTIVE

https://doi.org/10.1038/s41593-019-0520-2

A deep learning framework for neuroscience

Blake A. Richards[1,2,3,4,42*], Timothy P. Lillicrap[5,6,42], Philippe Beaudoin[7], Yoshua Bengio[1,4,8],

Systems neuroscience seeks explanations for how the brain implements a wide variety of perceptual, cognitive and motor tasks. Conversely, artificial intelligence attempts to design computational systems based on the tasks they will have to solve. In artificial neural networks, the three components specified by design are the objective functions, the learning rules and the architectures. With the growing success of deep learning, which utilizes brain-inspired architectures, these three designed components have increasingly become central to how we model, engineer and optimize complex artificial learning systems. Here we argue that a greater focus on these components would also benefit systems neuroscience. We give examples of how this optimization-based framework can drive theoretical and experimental progress in neuroscience. We contend that this principled perspective on systems neuroscience will help to generate more rapid progress.

NATURE NEUROSCIENCE | VOL 22 | NOVEMBER 2019 | 1761–1770 | www.nature.com/natureneuroscience

⑥ DISCUSSION & CONCLUSIONS

* Focus on the objective functions! (the "why" question)

nature neuroscience

FOCUS | PERSPECTIVE

https://doi.org/10.1038/s41593-019-0520-2

A deep learning framework for neuroscience

Blake A. Richards[1,2,3,4,42*], Timothy P. Lillicrap[5,6,42], Philippe Beaudoin[7], Yoshua Bengio[1,4,8],

Systems neuroscience seeks explanations for how the brain implements a wide variety of perceptual, cognitive and motor tasks. Conversely, artificial intelligence attempts to design computational systems based on the tasks they will have to solve. In artificial neural networks, the three components specified by design are the objective functions, the learning rules and the architectures. With the growing success of deep learning, which utilizes brain-inspired architectures, these three designed components have increasingly become central to how we model, engineer and optimize complex artificial learning systems. Here we argue that a greater focus on these components would also benefit systems neuroscience. We give examples of how this optimization-based framework can drive theoretical and experimental progress in neuroscience. We contend that this principled perspective on systems neuroscience will help to generate more rapid progress.

NATURE NEUROSCIENCE | VOL 22 | NOVEMBER 2019 | 1761–1770 | www.nature.com/natureneuroscience

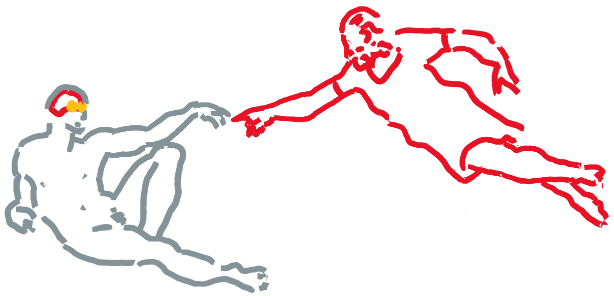

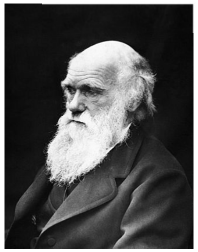

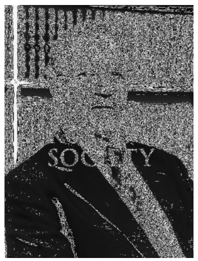

* Human illusions may come from infomax / error min

* CNN illusions come from error minimization

7. Other stuff we do

- Psychophysics - MAD
- Physiology - fMRI