Principles of Data Visualization I

Eamonn Maguire
CERN School of Computing, Israel
October 2018

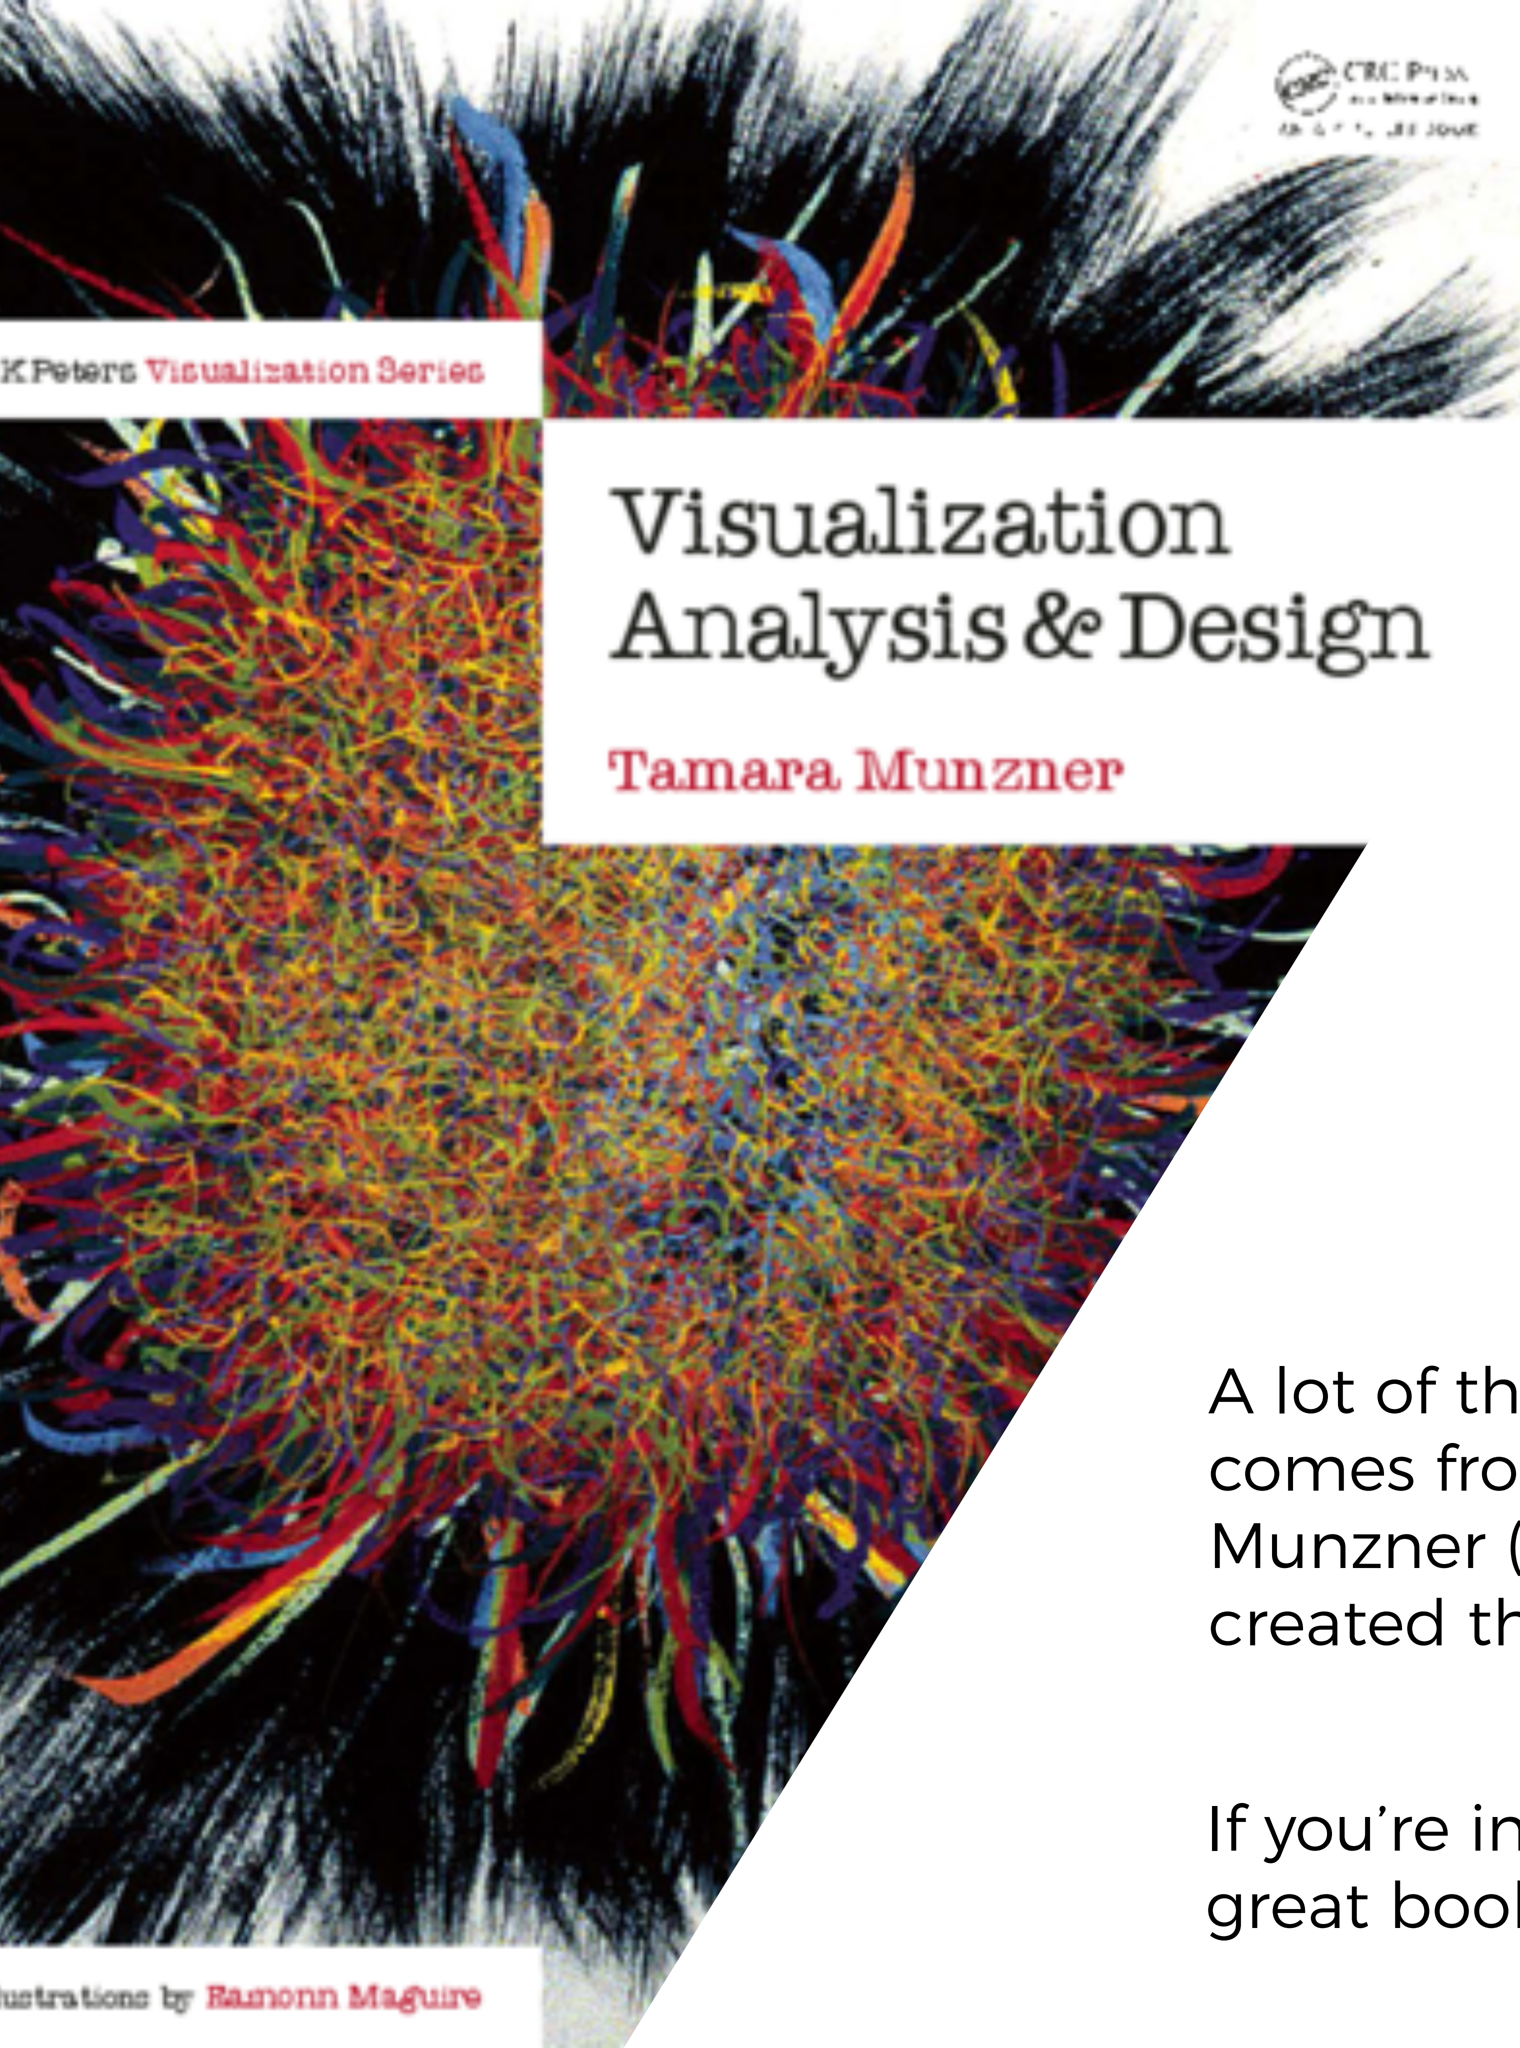

A lot of the content for this introduction comes from this book from Prof. Tamara Munzner (UBC, Vancouver, Canada) which I created the illustrations for.

If you're interested in learning more, it's a great book to check out :)

Visualization

The role of visualization systems is to provide visual representations of datasets that help people carry out tasks **more effectively**.

Tamara Munzner

A Visualization should:

1. Save time
2. Have a **clear purpose***
3. Include only the **relevant content***
4. **Encodes data/information** appropriately

* from Noel Illinsky, http://complexdiagrams.com/

Visualization

The role of visualization systems is to provide visual representations of datasets that help people carry out tasks **more effectively**.

Visualization is suitable when there is a need to augment human capabilities rather than replace people with computational decision-making methods.

Tamara Munzner

A Visualization should:

1. Save time
2. Have a **clear purpose***
3. Include only the **relevant content***
4. **Encodes data/information** appropriately

* from Noel Illinsky, http://complexdiagrams.com/

Visualization

The role of visualization systems is to provide visual representations of datasets that help people **carry out tasks more effectively**.

External representation:
replace cognition with perception

Visualization

The role of visualization systems is to provide visual representations of datasets that help people **carry out tasks more effectively**.

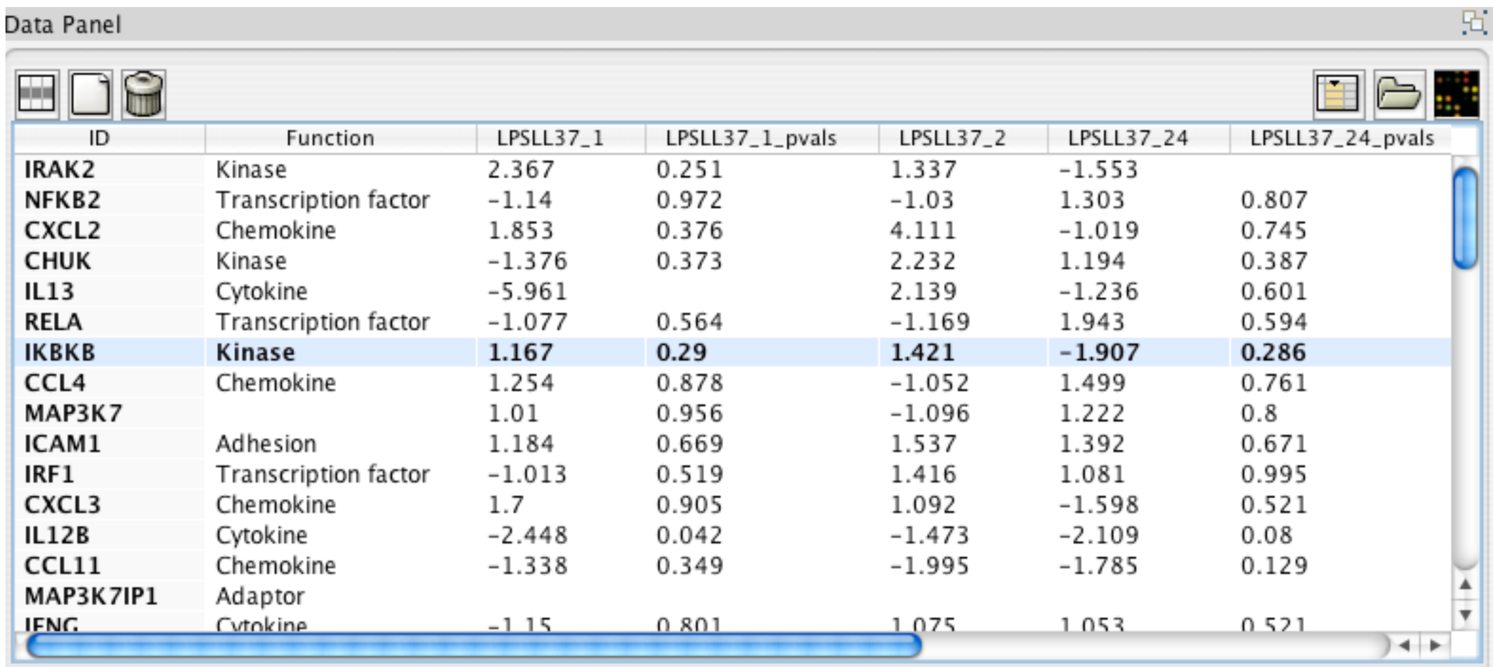

Data Panel

ID	Function	LPSLL37_1	LPSLL37_1_pvals	LPSLL37_2	LPSLL37_24	LPSLL37_24_pvals
IRAK2	Kinase	2.367	0.251	1.337	-1.553	
NFKB2	Transcription factor	-1.14	0.972	-1.03	1.303	0.807
CXCL2	Chemokine	1.853	0.376	4.111	-1.019	0.745
CHUK	Kinase	-1.376	0.373	2.232	1.194	0.387
IL13	Cytokine	-5.961		2.139	-1.236	0.601
RELA	Transcription factor	-1.077	0.564	-1.169	1.943	0.594
IKBKB	Kinase	1.167	0.29	1.421	-1.907	0.286
CCL4	Chemokine	1.254	0.878	-1.052	1.499	0.761
MAP3K7		1.01	0.956	-1.096	1.222	0.8
ICAM1	Adhesion	1.184	0.669	1.537	1.392	0.671
IRF1	Transcription factor	-1.013	0.519	1.416	1.081	0.995
CXCL3	Chemokine	1.7	0.905	1.092	-1.598	0.521
IL12B	Cytokine	-2.448	0.042	-1.473	-2.109	0.08
CCL11	Chemokine	-1.338	0.349	-1.995	-1.785	0.129
MAP3K7IP1	Adaptor					
IFNG	Cytokine	-1.15	0.801	1.075	1.053	0.521

External representation: replace cognition with perception

Visualization

The role of visualization systems is to provide visual representations of datasets that help people **carry out tasks more effectively**.

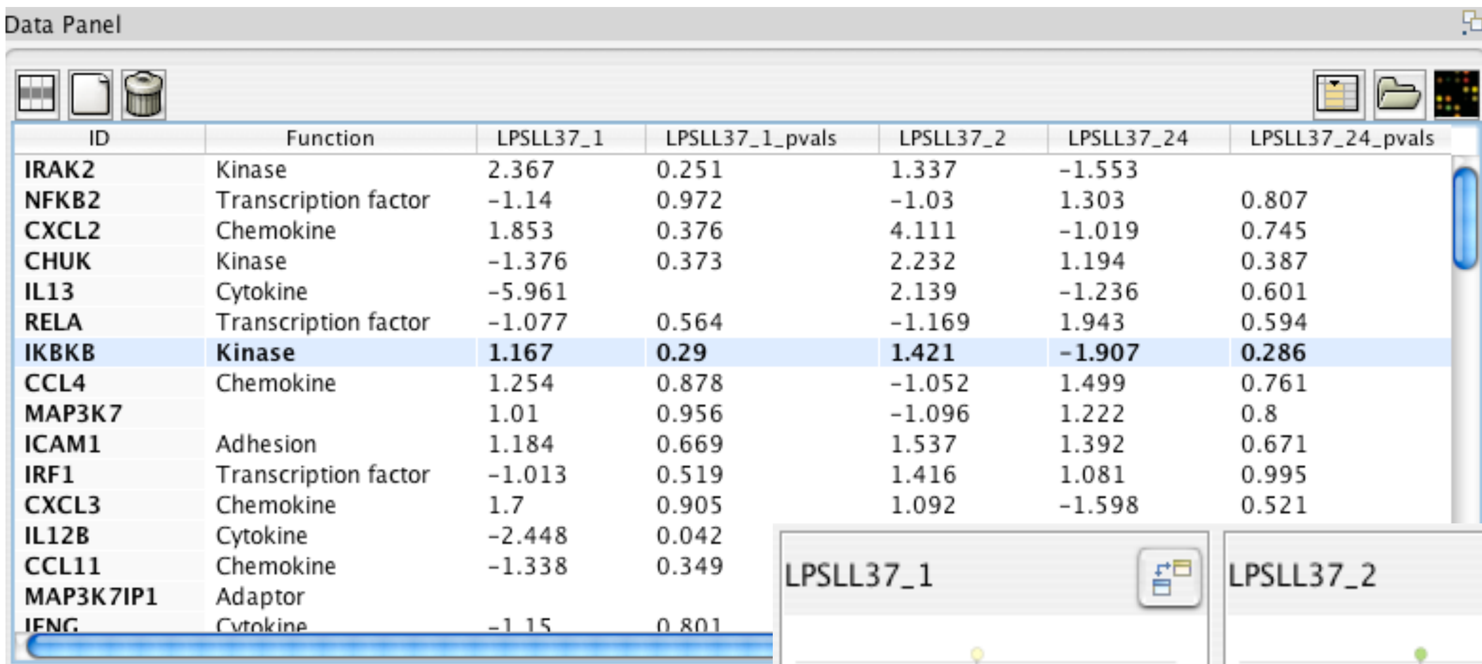

ID	Function	LPSLL37_1	LPSLL37_1_pvals	LPSLL37_2	LPSLL37_24	LPSLL37_24_pvals
IRAK2	Kinase	2.367	0.251	1.337	-1.553	
NFKB2	Transcription factor	-1.14	0.972	-1.03	1.303	0.807
CXCL2	Chemokine	1.853	0.376	4.111	-1.019	0.745
CHUK	Kinase	-1.376	0.373	2.232	1.194	0.387
IL13	Cytokine	-5.961		2.139	-1.236	0.601
RELA	Transcription factor	-1.077	0.564	-1.169	1.943	0.594
IKBKB	Kinase	1.167	0.29	1.421	-1.907	0.286
CCL4	Chemokine	1.254	0.878	-1.052	1.499	0.761
MAP3K7		1.01	0.956	-1.096	1.222	0.8
ICAM1	Adhesion	1.184	0.669	1.537	1.392	0.671
IRF1	Transcription factor	-1.013	0.519	1.416	1.081	0.995
CXCL3	Chemokine	1.7	0.905	1.092	-1.598	0.521
IL12B	Cytokine	-2.448	0.042			
CCL11	Chemokine	-1.338	0.349			
MAP3K7IP1	Adaptor					
IFNG	Cytokine	-1.15	0.801			

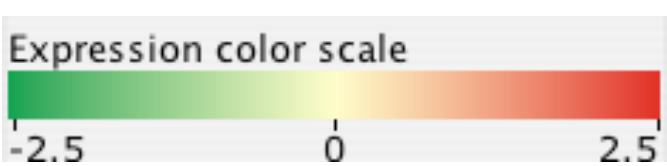

External representation: replace cognition with perception

Cerebral: Visualizing Multiple Experimental Conditions on a Graph with Biological Context. Barsky, Munzner, Gardy, and Kincaid. IEEE TVCG (Proc. InfoVis) 14(6): 1253-1260, 2008.]

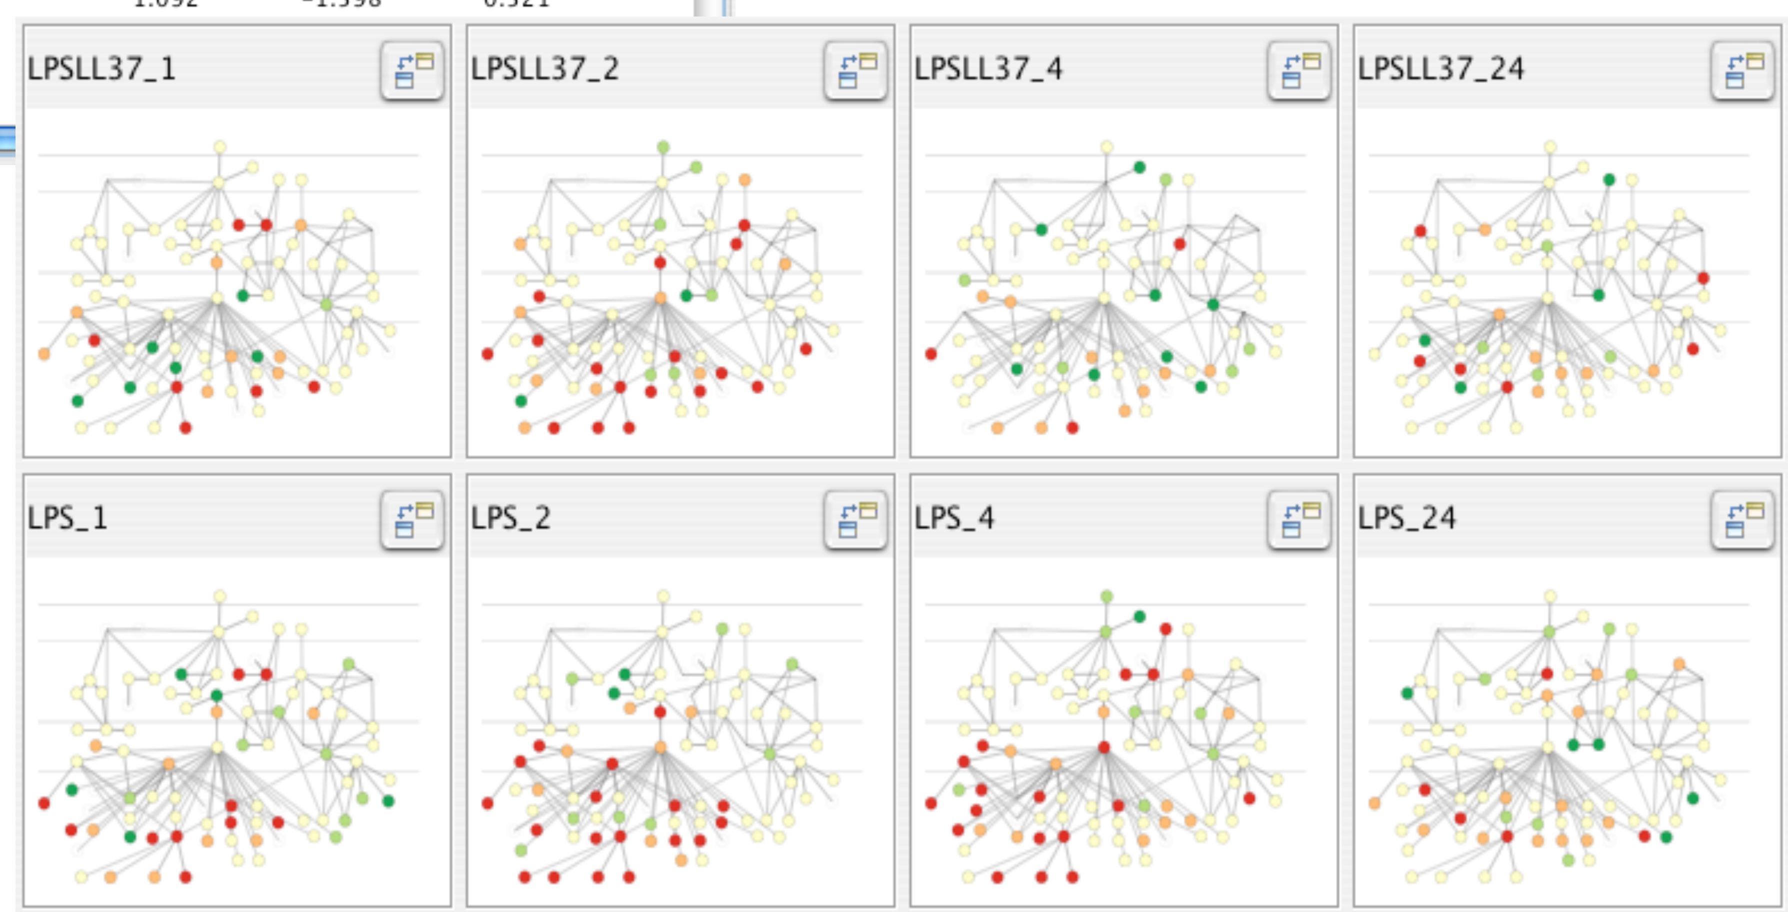

What are we visualising?

Major data types & classifications of them

Why are we visualising it?

What is the need for this visualization?

Why do the users need this, and what do they need to be able to do with it?

How can we visualise?

How can we visualize?

The components of a visualization.

Good and bad practices.

What are we visualising?

Major data types & classifications of them

Why are we visualising it?

What is the need for this visualization?

Why do the users need this, and what do they need to be able to do with it?

How can we visualise?

How can we visualize?

The components of a visualization.

Good and bad practices.

What are we visualising?

DATA TYPES

STATIC

DYNAMIC

DATASET TYPES

TABLES

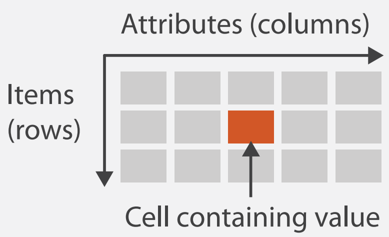

NETWORKS

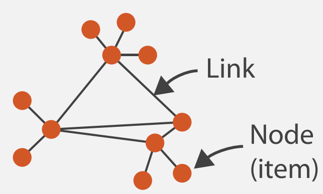

FIELDS (CONTINUOUS)

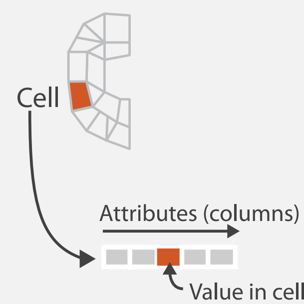

GEOMETRY (SPATIAL)

TEXT

- Prose Documents
- Document Collections
- Log Files
- Code
- Multimedia

- Multidimensional Table

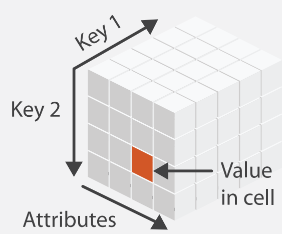

- Trees

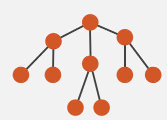

ATTRIBUTE TYPES

CATEGORICAL

ORDERED

- Ordinal

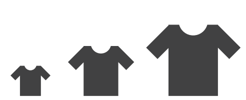

- Quantitative

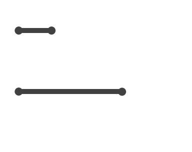

- Sequential

- Diverging

- Cyclic

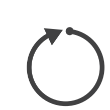

What are we visualising?

➔ STATIC

For static data, we have **fixed scales**.

We know our data range, therefore scales will not change.

What are we visualising?

STATIC

For static data, we have **fixed scales**.

We know our data range, therefore scales will not change.

DYNAMIC

For dynamic data, the observed min and max values can change, **therefore scales will change**.

This can have big consequences for the readability of our visualization.

What are you visualising?

The branches of data visualization

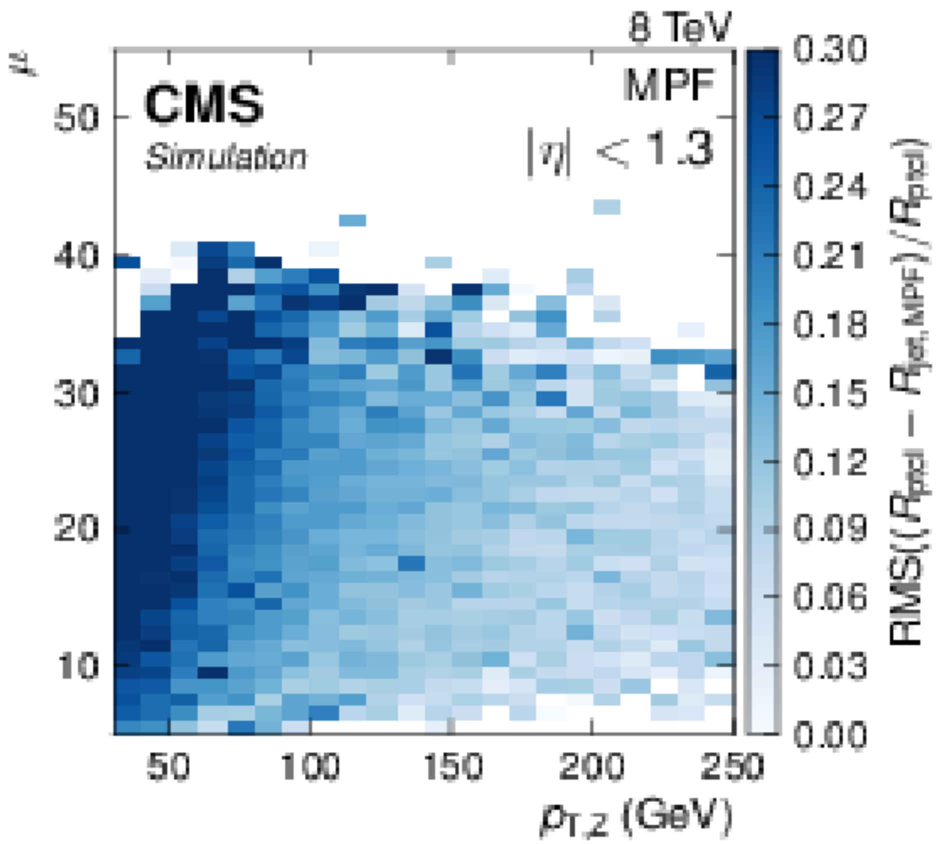

Information Visualization

Position is derived.
Incl. GeoVis

We decide what is on the X and Y axis, and what we do changes the information we extract.

Scientific Visualization

Position is given.
e.g. detector or medical visualizations

We have the X, Y, and Z coordinates of a cell in ATLAS, we show the energy deposit left here. We don't choose, the data tells us.

What are you visualising?

The branches of data visualization

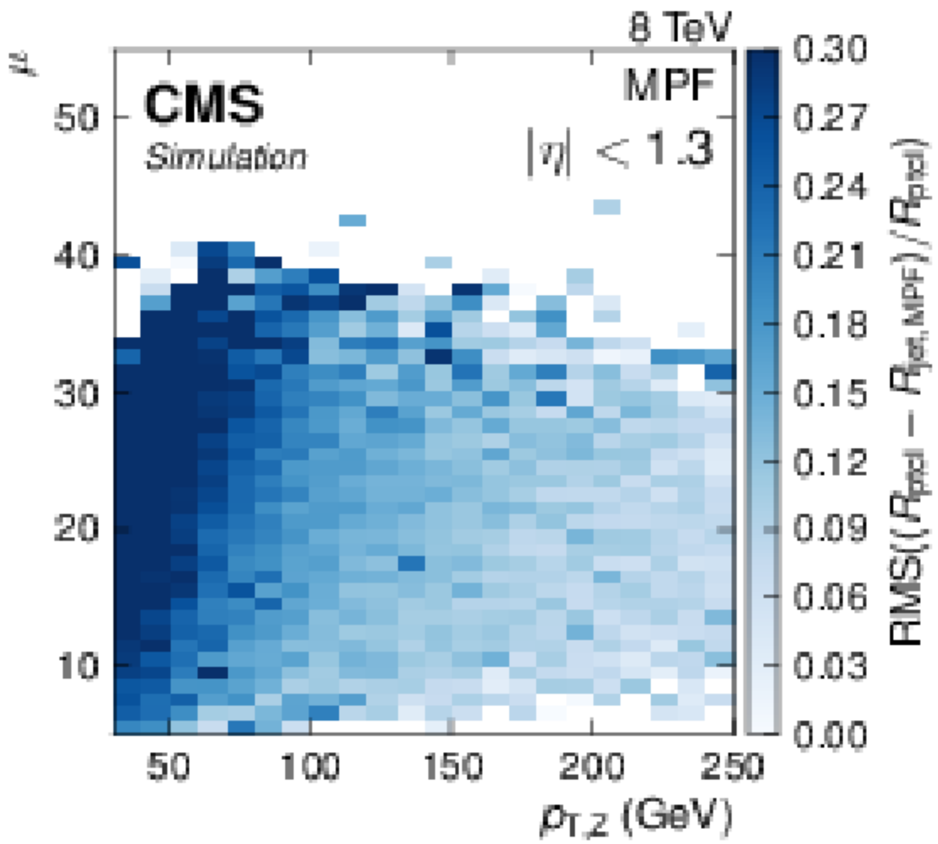

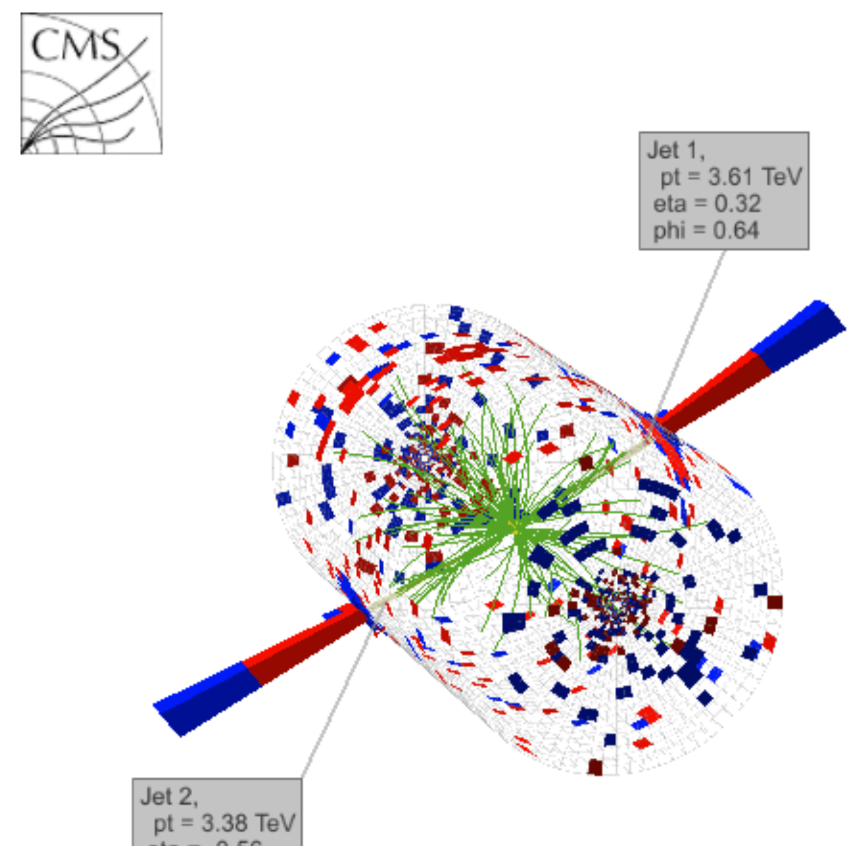

Information Visualization

Position is derived.
Incl. GeoVis

We decide what is on the X and Y axis, and what we do changes the information we extract.

Scientific Visualization

Position is given.
e.g. detector or medical visualizations

We have the X, Y, and Z coordinates of a cell in ATLAS, we show the energy deposit left here. We don't choose, the data tells us.

What are we visualising?

Major data types & classifications of them

Why are we visualising it?

What is the need for this visualization?

Why do the users need this, and what do they need to be able to do with it?

How can we visualise?

How can we visualize?

The components of a visualization.

Good and bad practices.

Why are we visualising?

The role of visualisation systems is to provide visual representations of datasets that help people **carry out tasks more effectively.**

Anscombe's Quartet: Raw Data

	1		2		3		4	
	X	Y	X	Y	X	Y	X	Y
	10.0	8.04	10.0	9.14	10.0	7.46	8.0	6.58
	8.0	6.95	8.0	8.14	8.0	6.77	8.0	5.76
	13.0	7.58	13.0	8.74	13.0	12.74	8.0	7.71
	9.0	8.81	9.0	8.77	9.0	7.11	8.0	8.84
	11.0	8.33	11.0	9.26	11.0	7.81	8.0	8.47
	14.0	9.96	14.0	8.10	14.0	8.84	8.0	7.04
	6.0	7.24	6.0	6.13	6.0	6.08	8.0	5.25
	4.0	4.26	4.0	3.10	4.0	5.39	19.0	12.50
	12.0	10.84	12.0	9.13	12.0	8.15	8.0	5.56
	7.0	4.82	7.0	7.26	7.0	6.42	8.0	7.91
	5.0	5.68	5.0	4.74	5.0	5.73	8.0	6.89
Mean	9.0	7.5	9.0	7.5	9.0	7.5	9.0	7.5
Variance	10.0	3.75	10.0	3.75	10.0	3.75	10.0	3.75
Correlation	0.816		0.816		0.816		0.816	

The statistics would lead us to believing that everything is the same

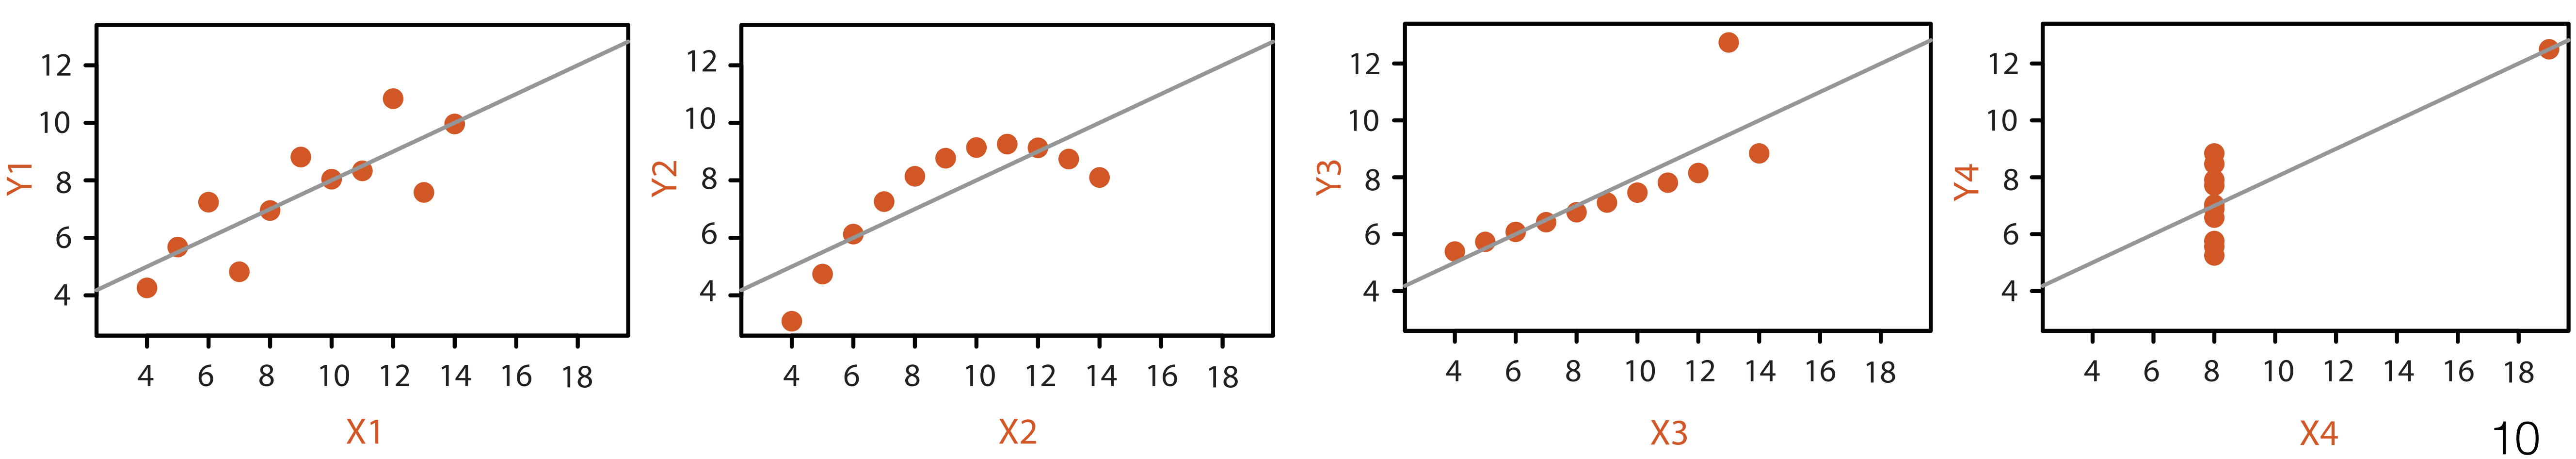

Why are we visualising?

Given a large matrix, or even a large series of numbers, it's difficult for humans to 'see' patterns in the data.

With a visualisation we want to transition a cognitively demanding task to a perceptual (less demanding) one.

Why are we visualising?

We can also use visualisation to better communicate concepts that aren't easily explained using text alone.

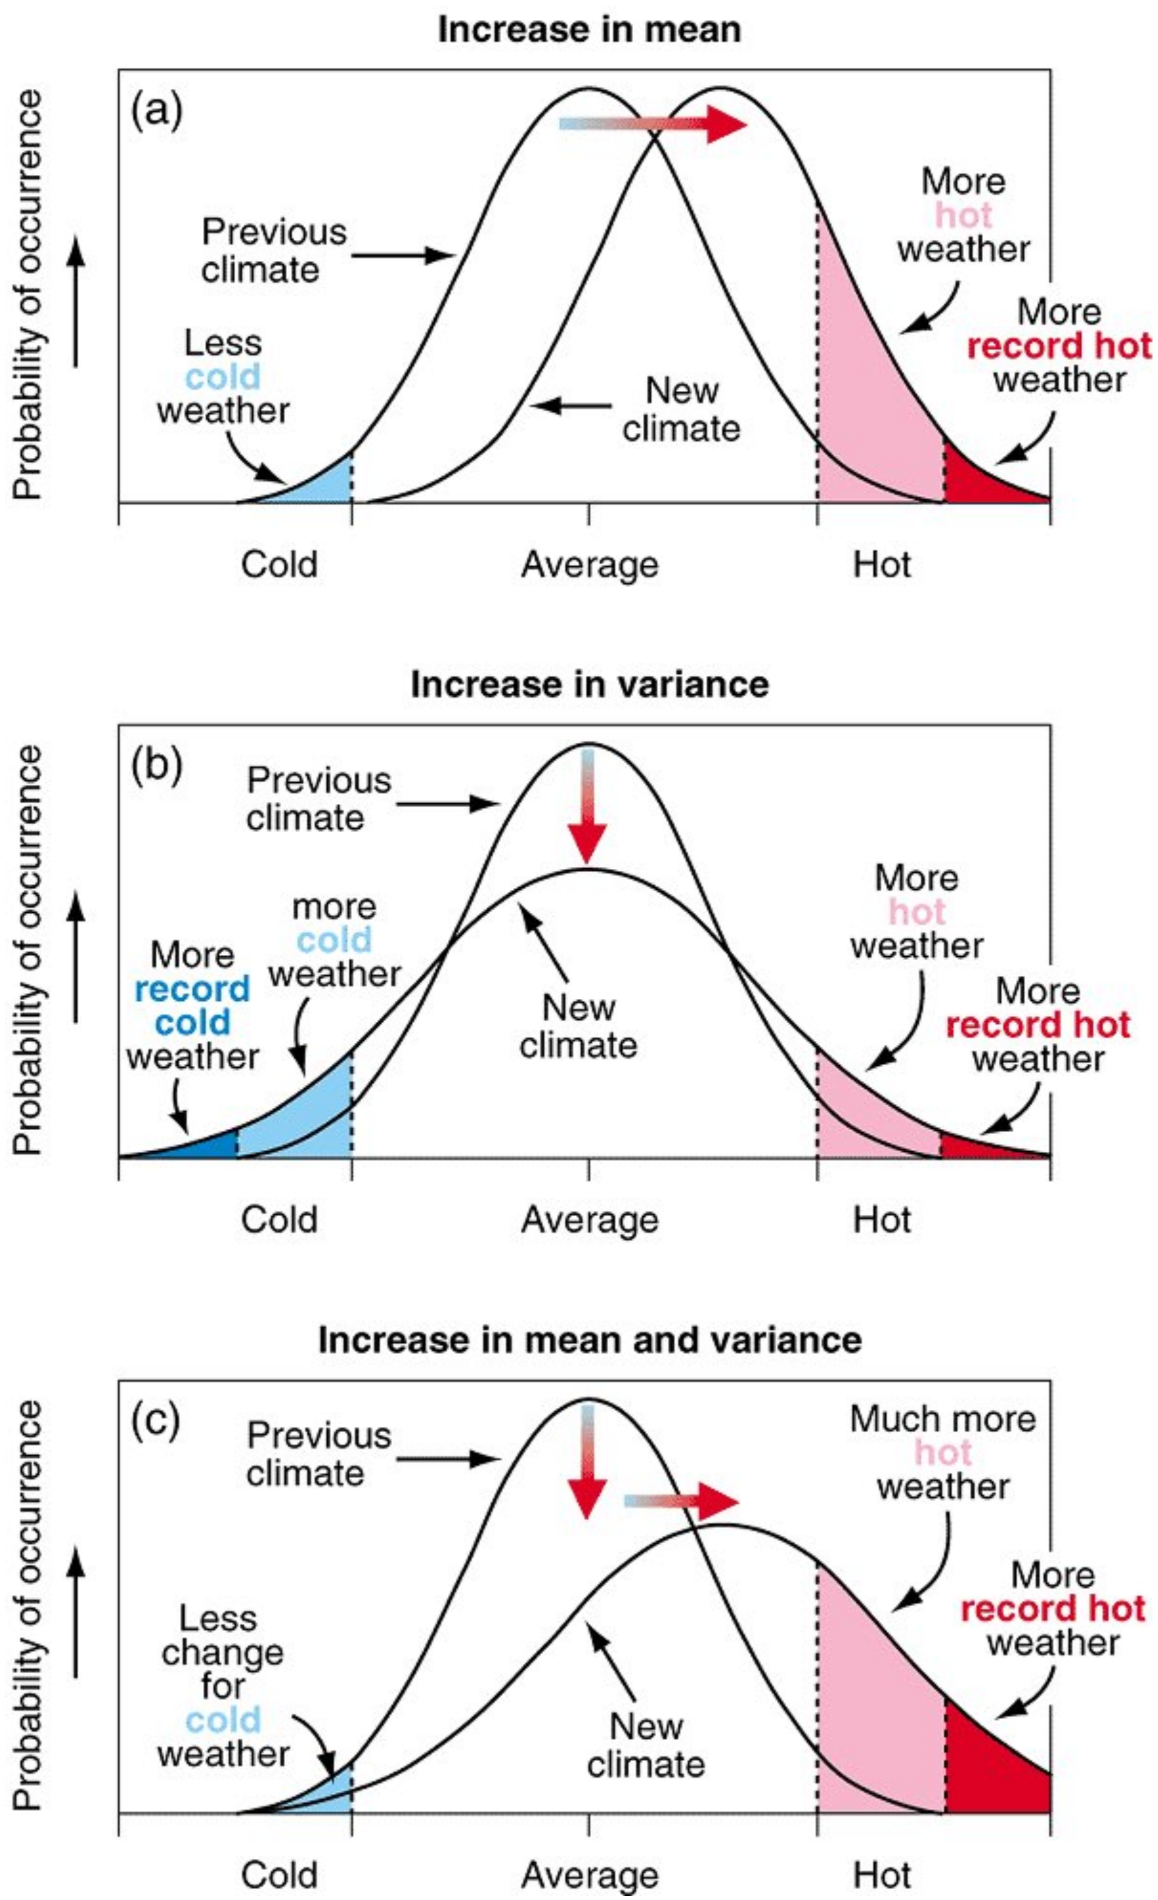

https://www.ipcc.ch/ipccreports/tar/wg1/fig2-32.htm

Why are we visualising?

Every visualisation should be thought of as a product of what actions the user needs to take to get to their objective (target)

Actions

Use

- Consume
  - *Discover*
  - *Present*
  - *Enjoy*
- Produce
  - *Annotate*
  - *Record*
  - *Derive*

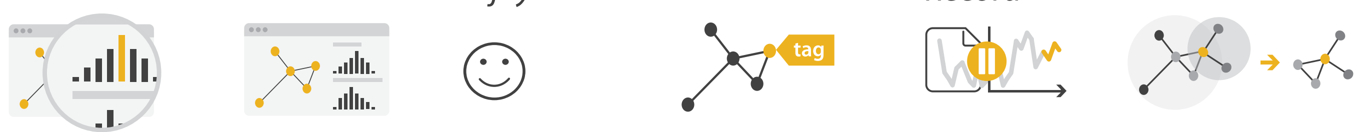

Search

	Target known	Target unknown
Location known	*Lookup*	*Browse*
Location unknown	*Locate*	*Explore*

Query

- Identify
- Compare
- Summarise

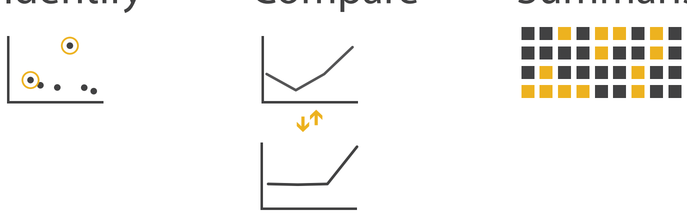

Why are we visualising?

Every visualisation should be thought of as a product of what actions the user needs to take to get to their objective (target)

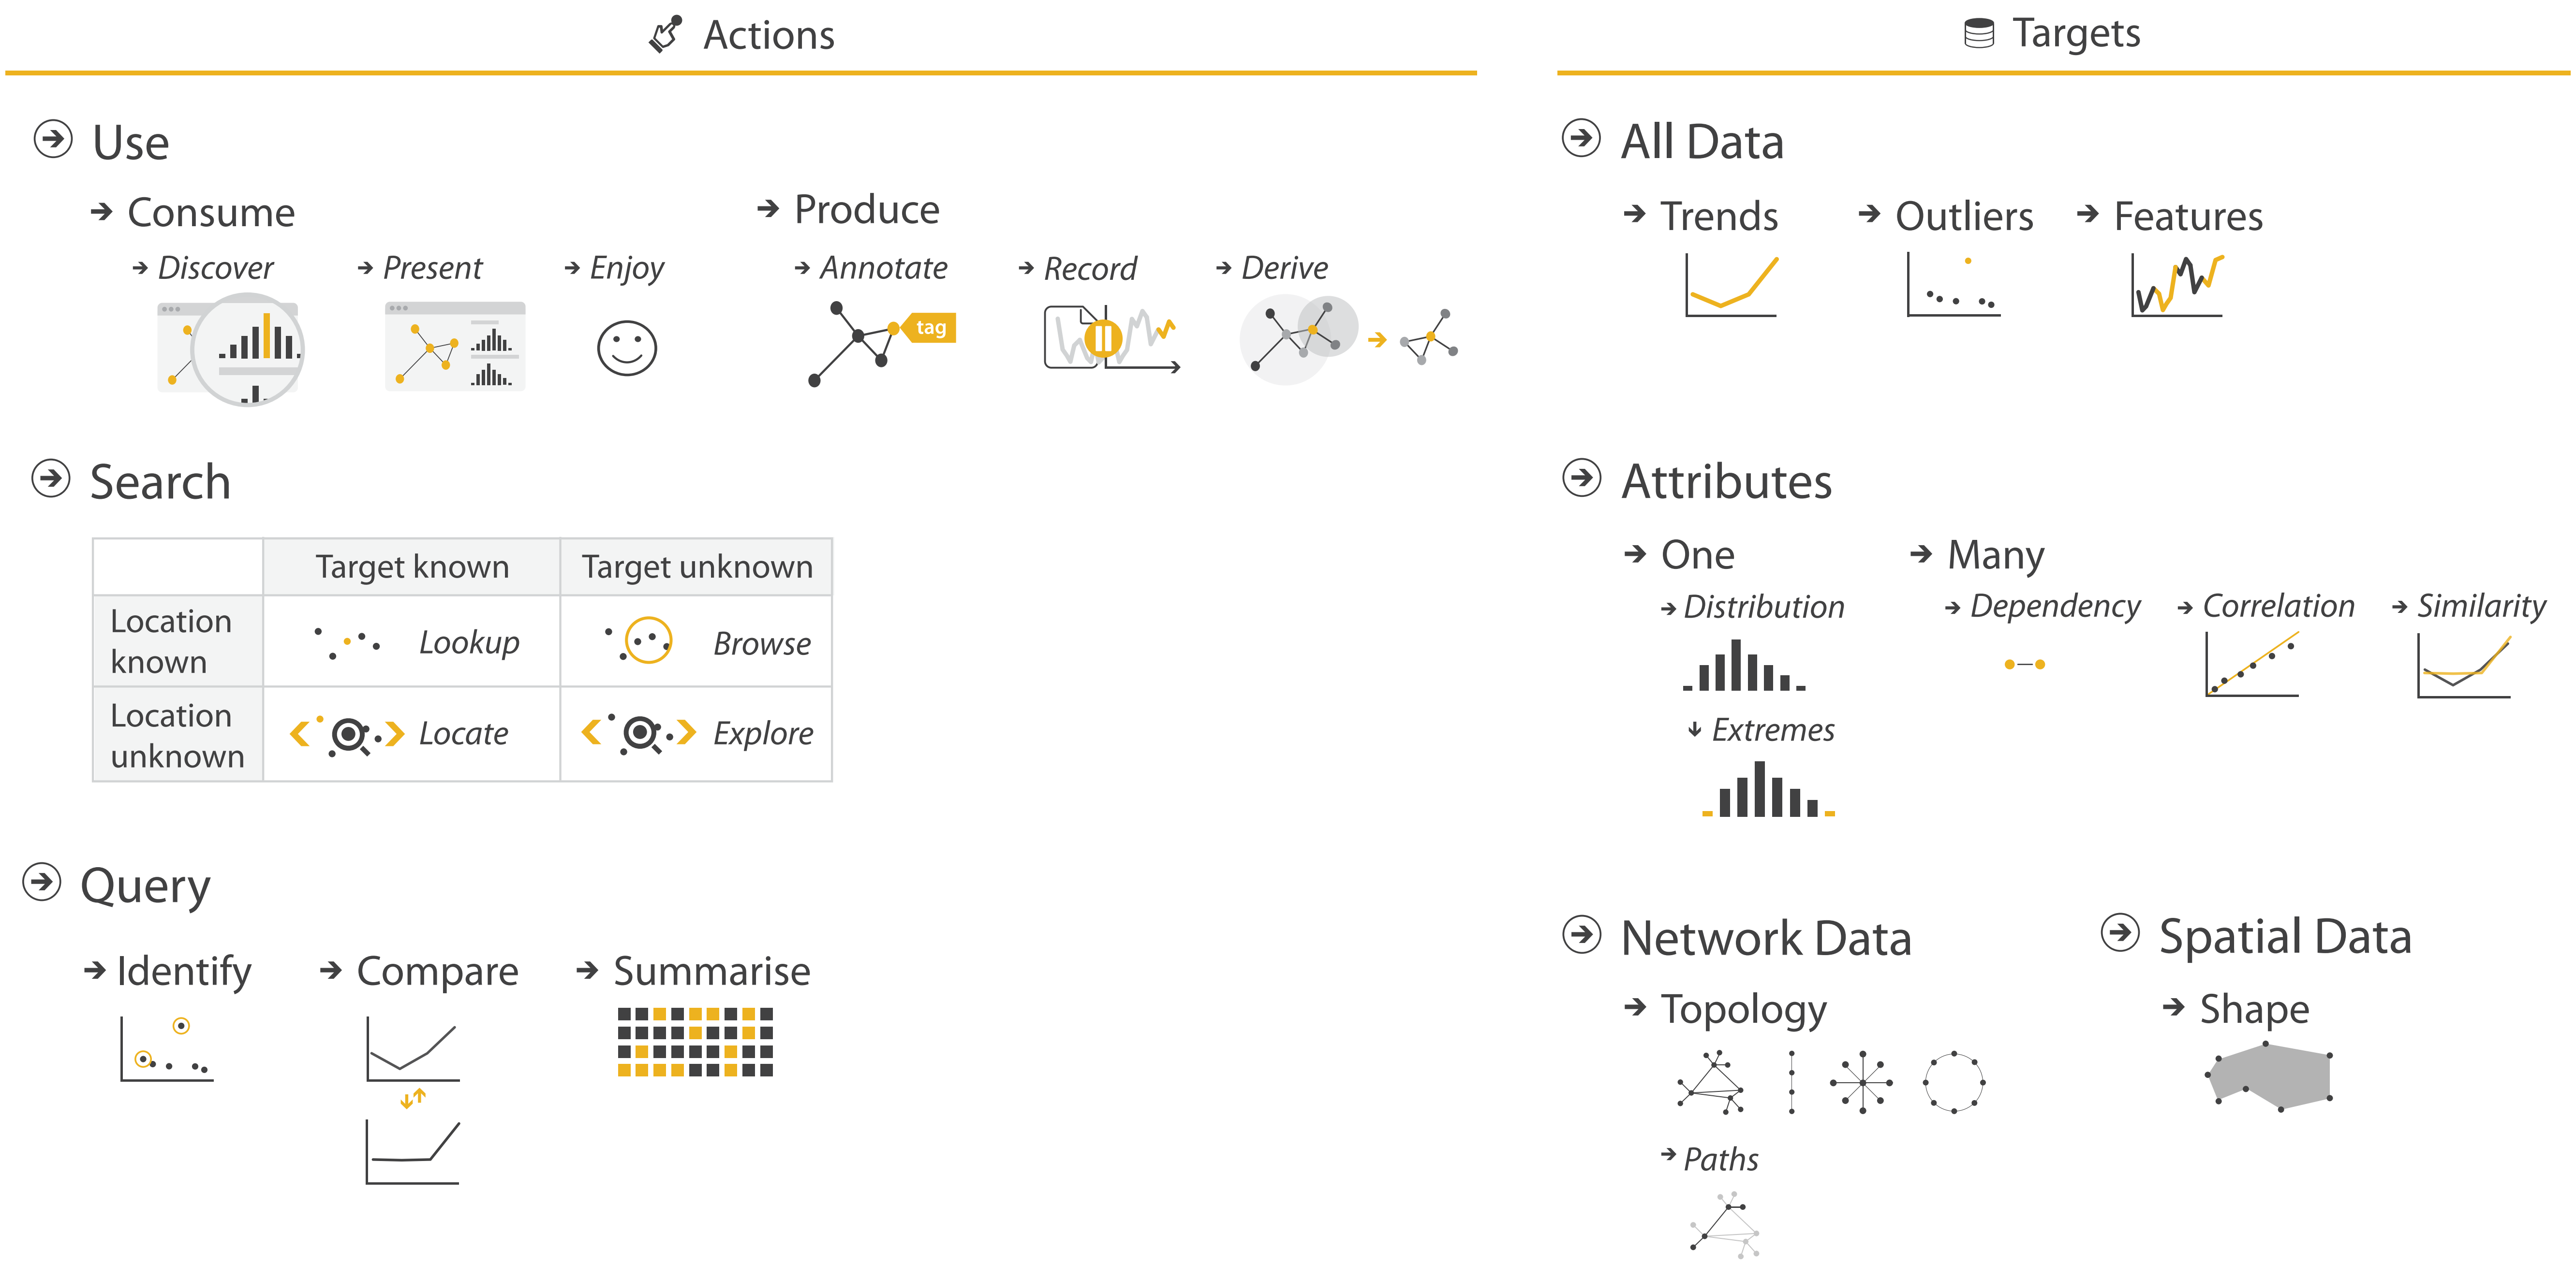

Why are we visualising?

Every visualisation should be thought of as a product of what actions the user needs to take to get to their objective (target)

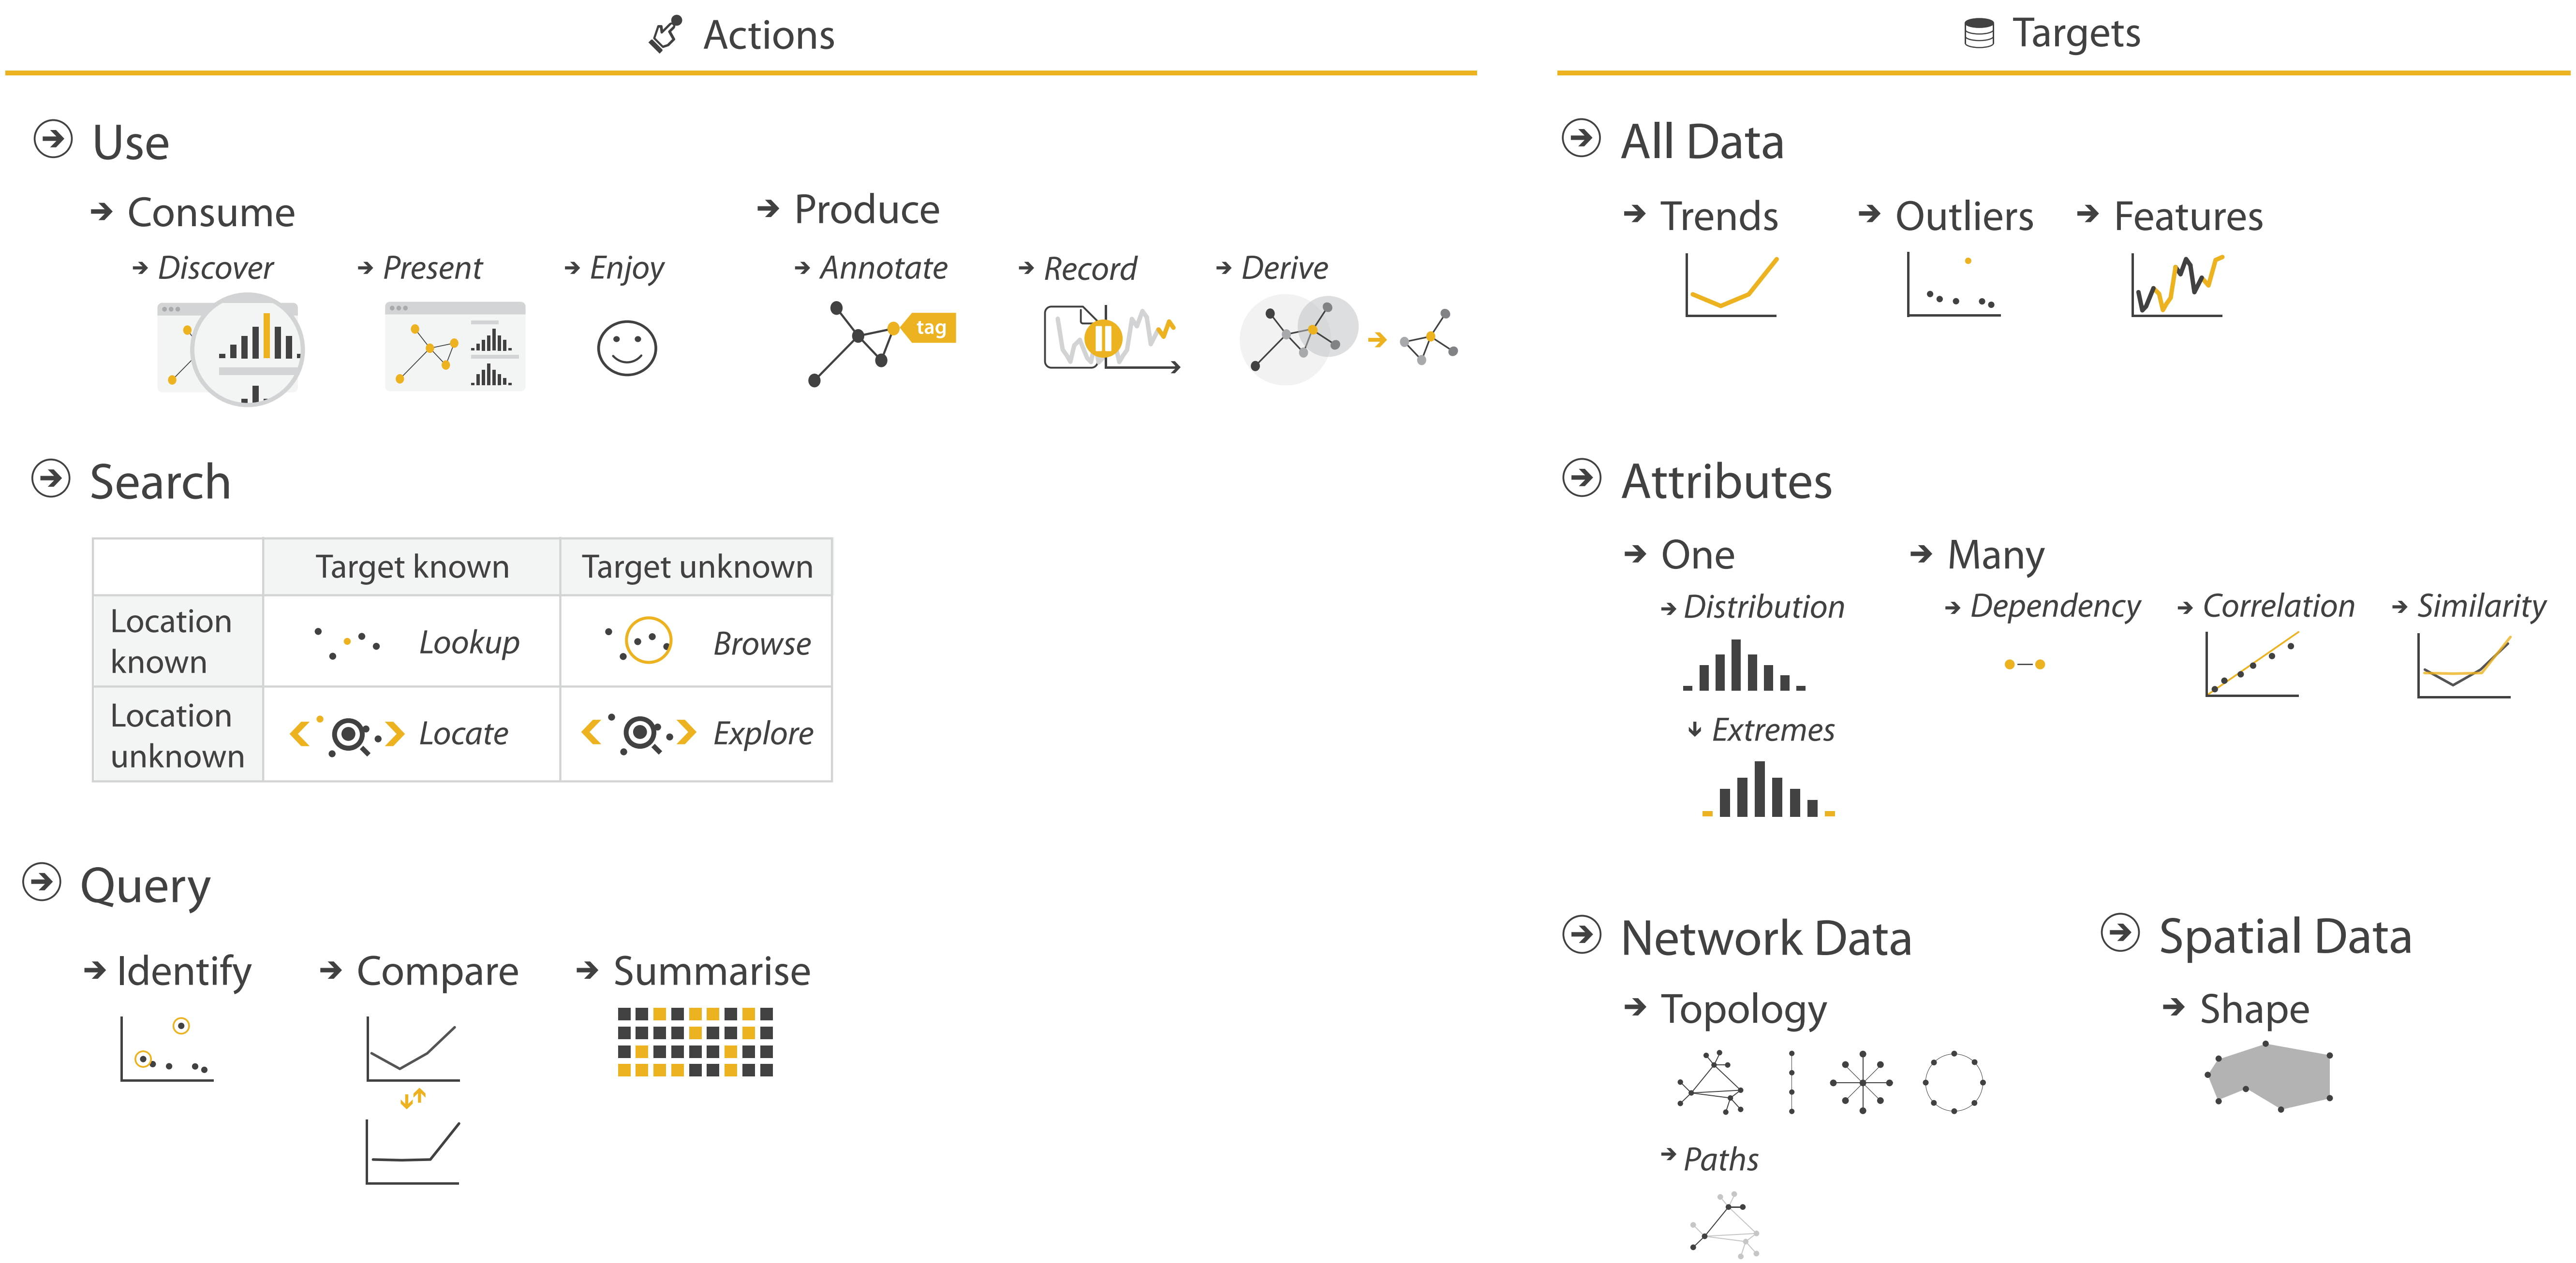

Always keep in mind why you're doing something. If what you create does not show what you intended, confuses, or misleads, it's time to rethink :)

Use

→ Consume

→ *Discover* → *Present* → *Enjoy*

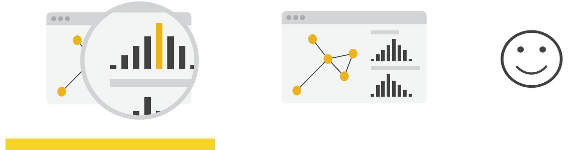

Discover

Finding new insights in your data

Implies a level of interactivity to query, compare, correlate etc.

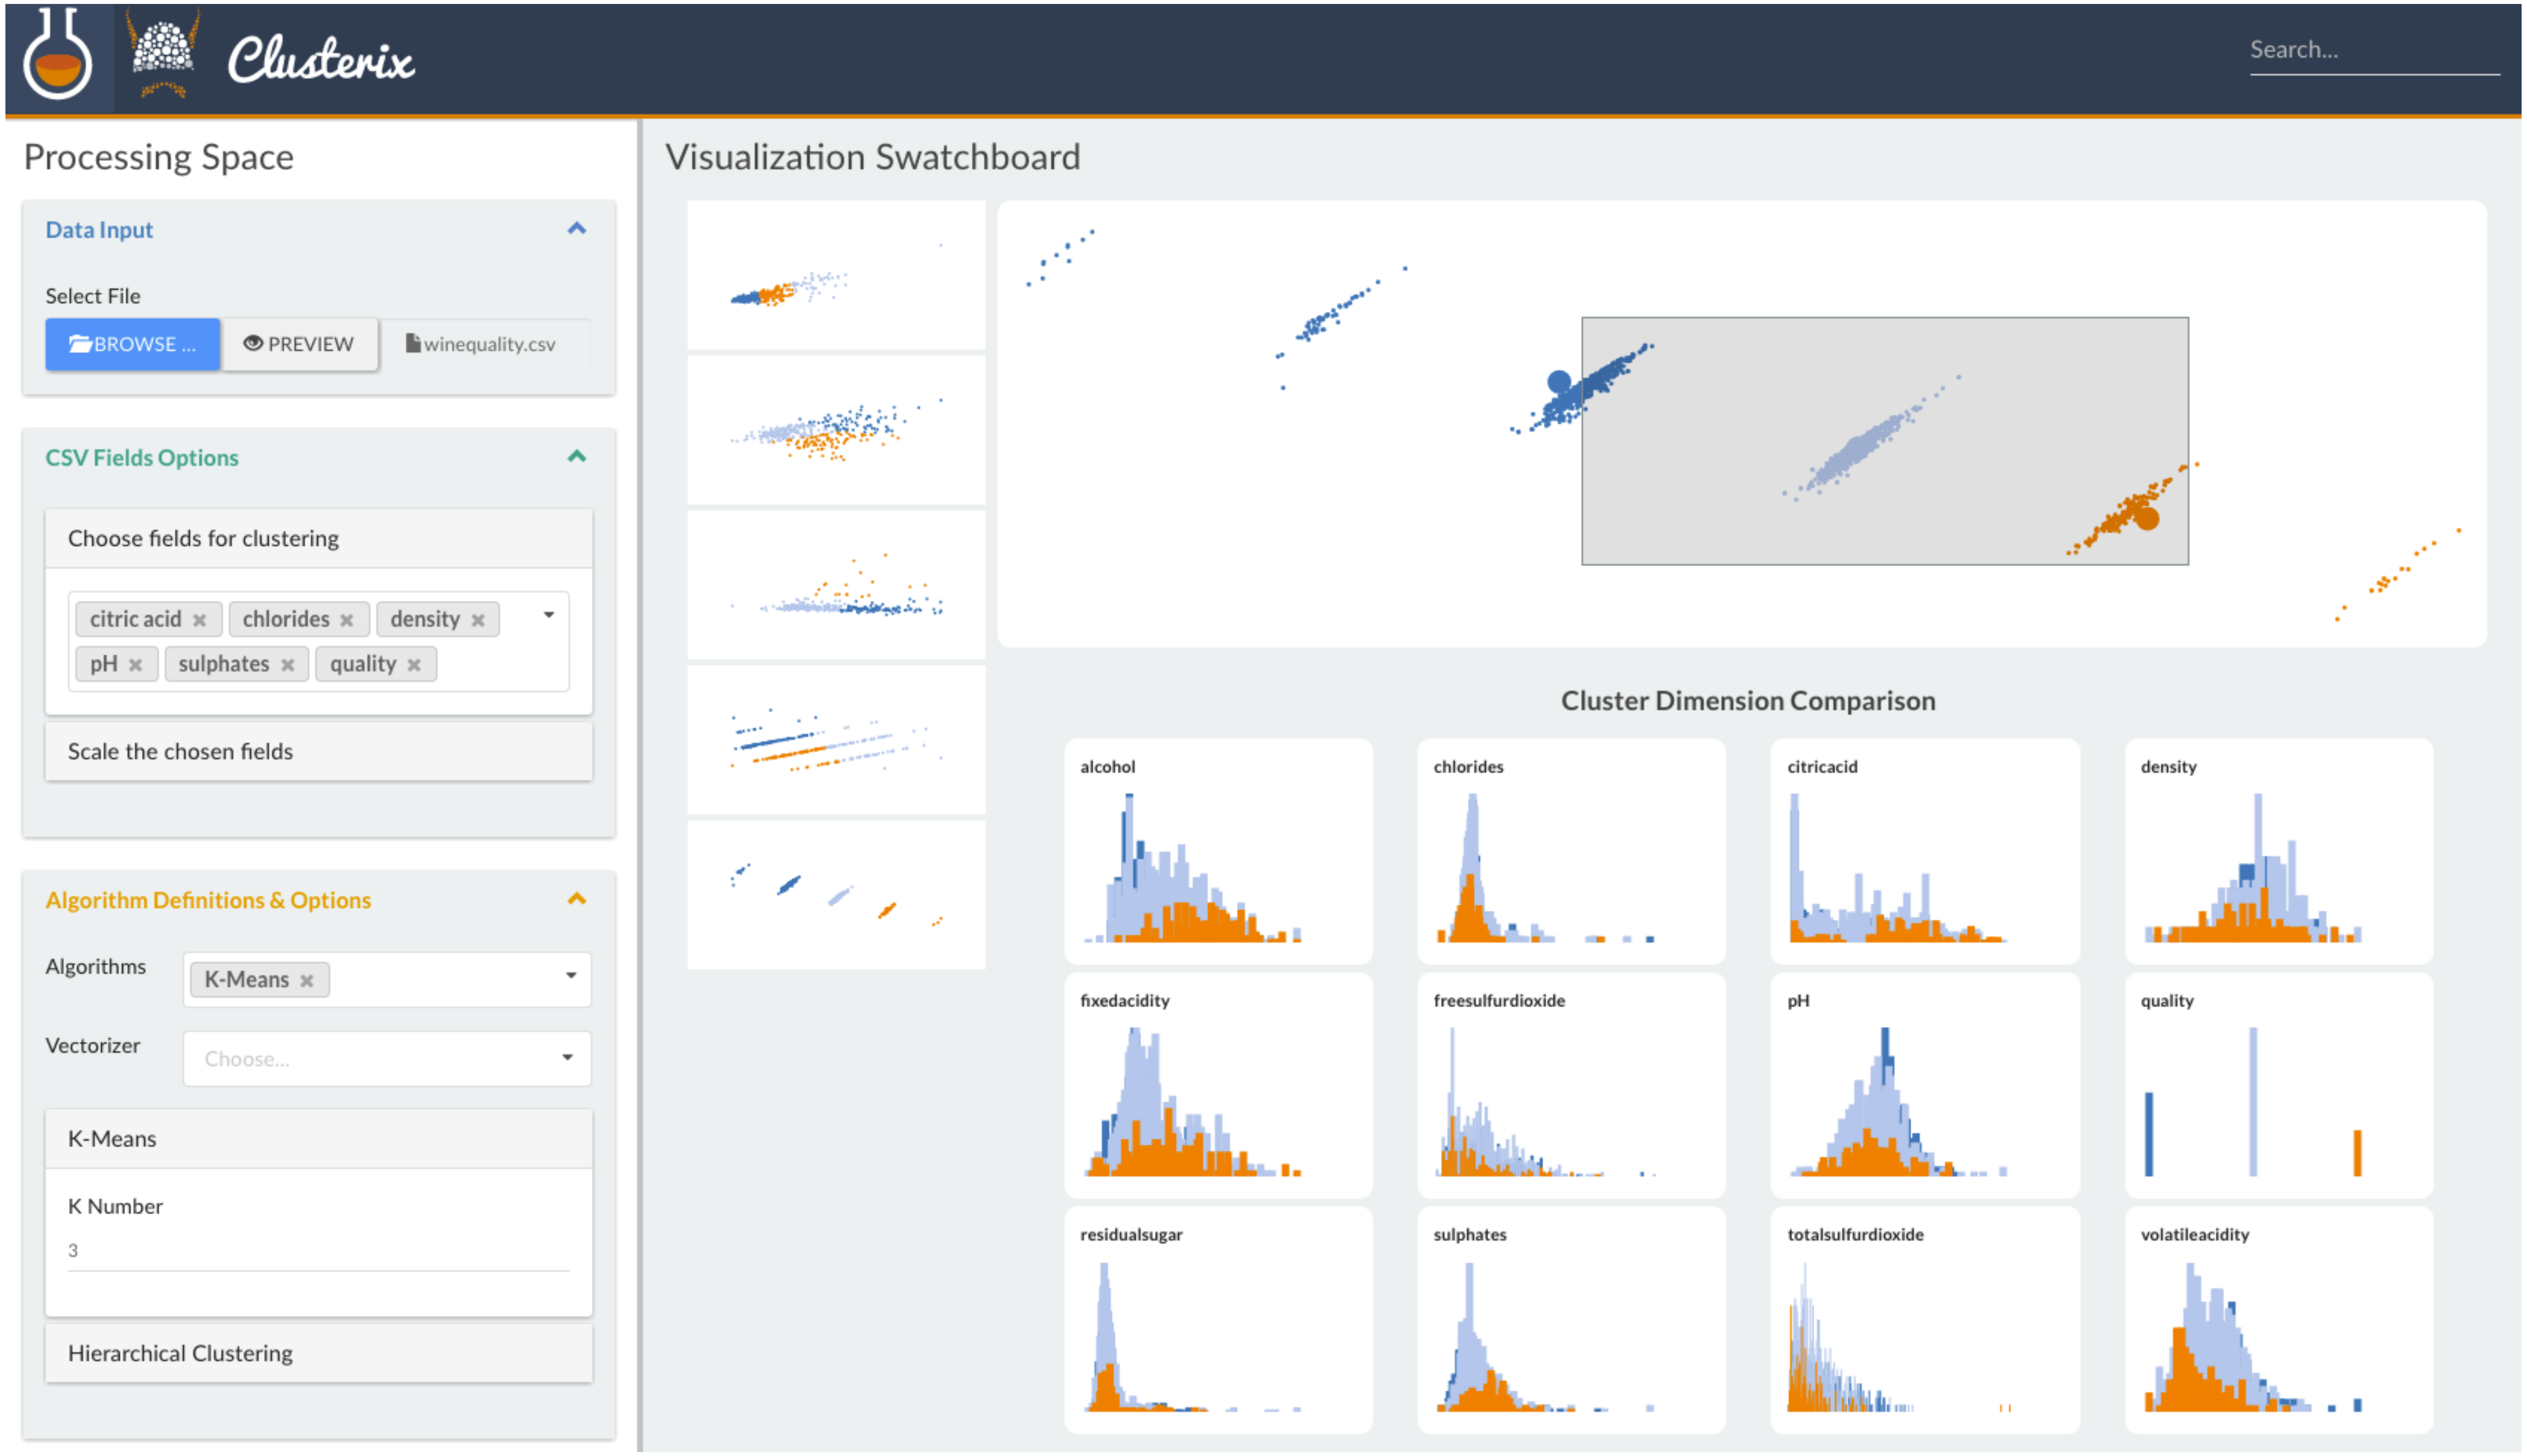

Work with Ilias Koutsakis and Gilles Louppe

What are we visualising?

Major data types & classifications of them

Why are we visualising it?

What is the need for this visualization?

Why do the users need this, and what do they need to be able to do with it?

How can we visualise?

How can we visualize?

The components of a visualization.

Good and bad practices.

How can you encode information optimally?

Encode	Manipulate	Facet	Reduce
	Change	Juxtapose	Filter
	Select	Partition	Aggregate
	Navigate	Superimpose	Embed

Arrange

- Express
- Separate
- Order
- Align
- Use

Map

from qualitative and quantitative attributes

- Color
  - *Hue*
  - *Saturation*
  - *Luminance*
  - *Transparency*
- Position, Size, Angle, Curvature, ...
- Region, Texture, Shape, ...
- Motion
  Direction, Rate, Frequency, ...

How can you encode information optimally?

Encode

Arrange
- Express
- Separate
- Order
- Align
- Use

Map
from qualitative and quantitative attributes

- Color
  - *Hue*
  - *Saturation*
  - *Luminance*
  - *Transparency*
- Position, Size, Angle, Curvature, ...
- Region, Texture, Shape, ...
- Motion
  Direction, Rate, Frequency, ...

Manipulate
- Change
- Select
- Navigate

Facet
- Juxtapose
- Partition
- Superimpose

Reduce
- Filter
- Aggregate
- Embed

If we don’t follow grammatical rules or spell correctly, the meaning of text can be lost.

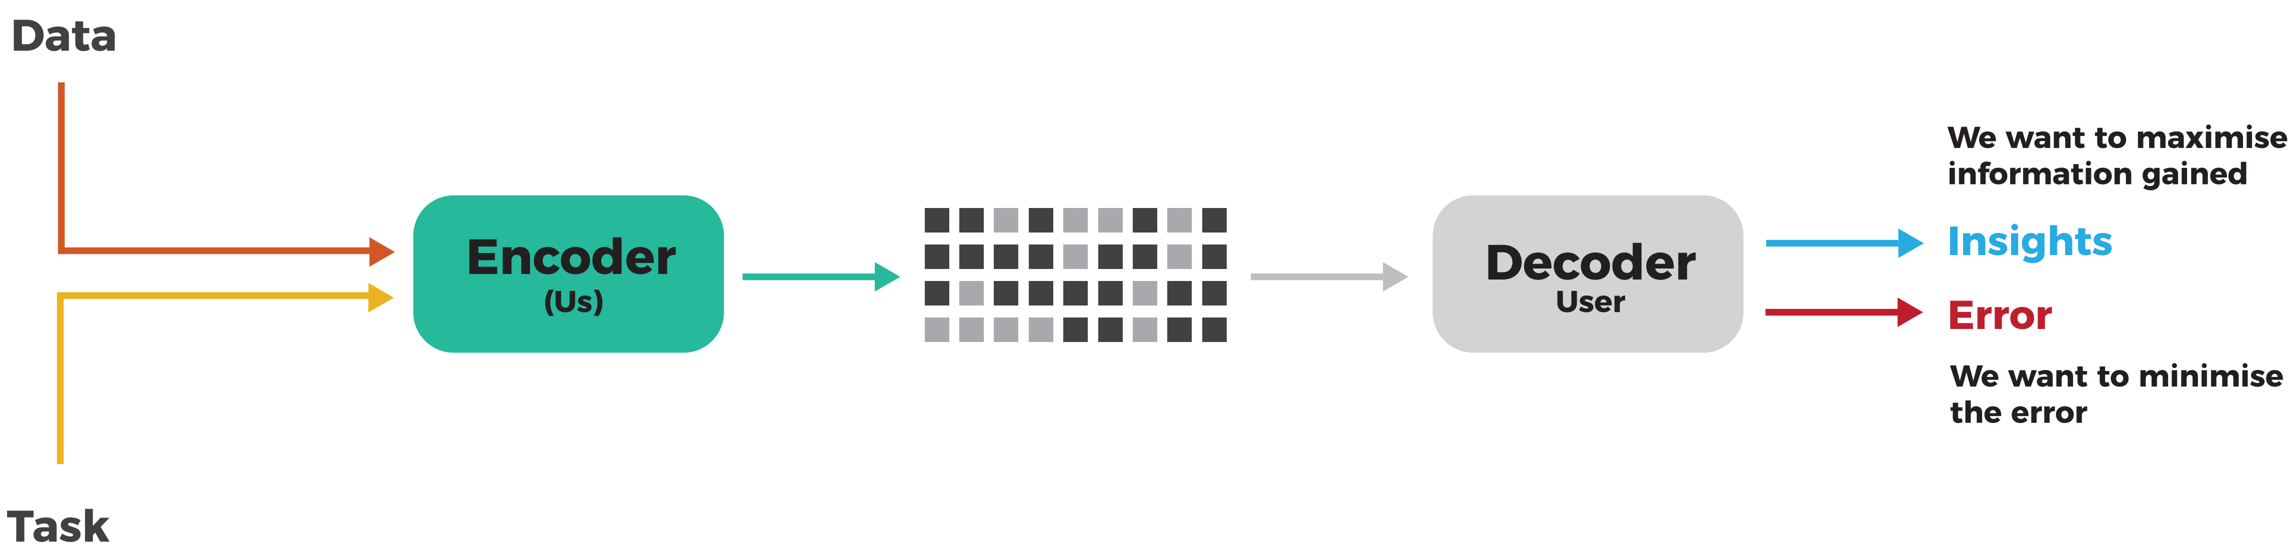

The same applies for visualisations. We can compose visualisations using a vocabulary (shapes, colour, texture,…), and a grammar. If we learn these, we can do better when it comes to communicating visually.

Graphs are like jokes.
If you have to explain them, they didn't work.

Anon.

How can you encode information optimally?

JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
2	10	4	5	6	9	1	3	5	3	4	7

How can you encode information optimally?

JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
2	10	4	5	6	9	1	3	5	3	4	7

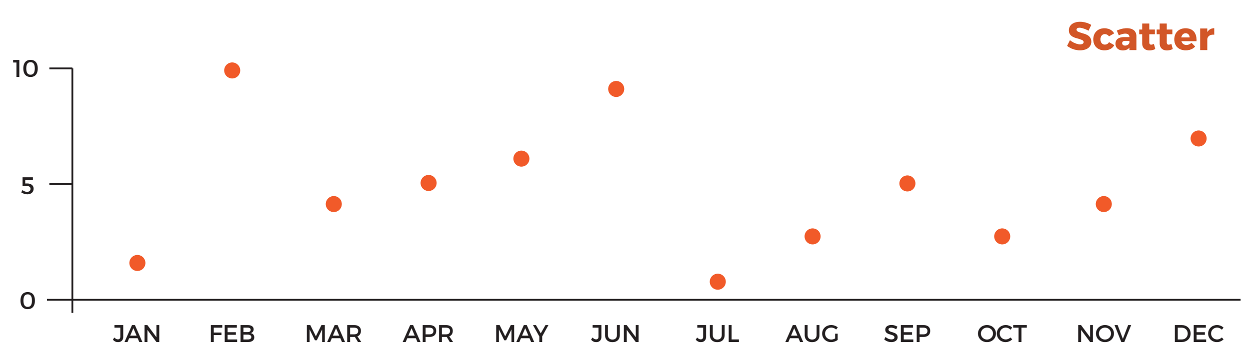

How can you encode information optimally?

JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
2	10	4	5	6	9	1	3	5	3	4	7

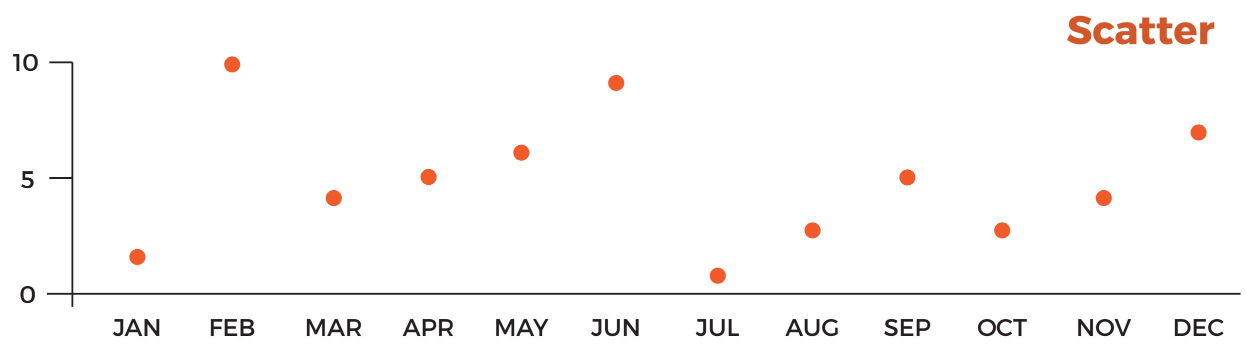

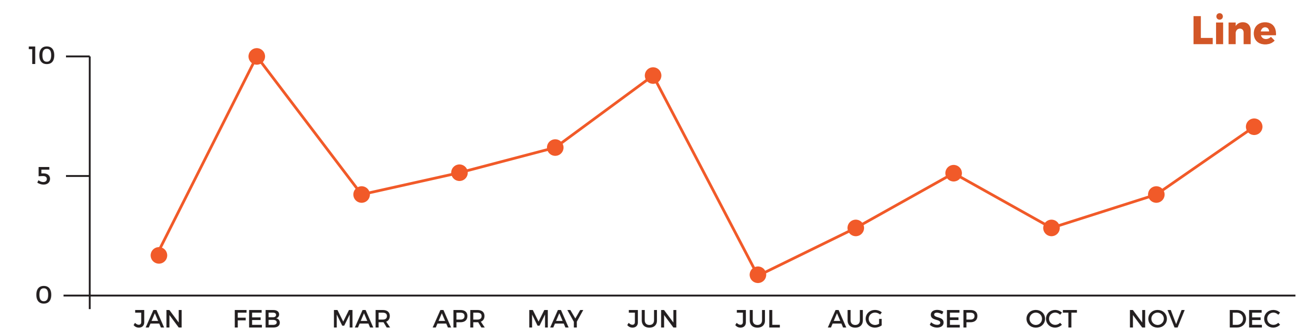

How can you encode information optimally?

JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
2	10	4	5	6	9	1	3	5	3	4	7

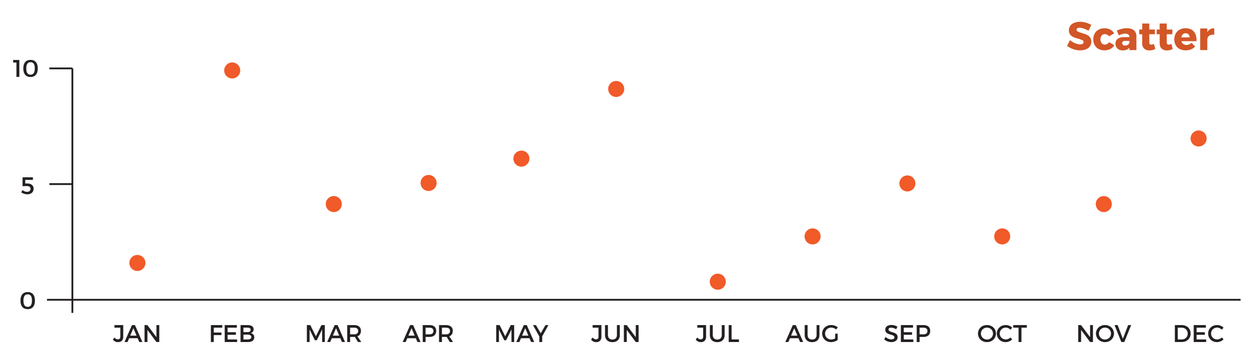

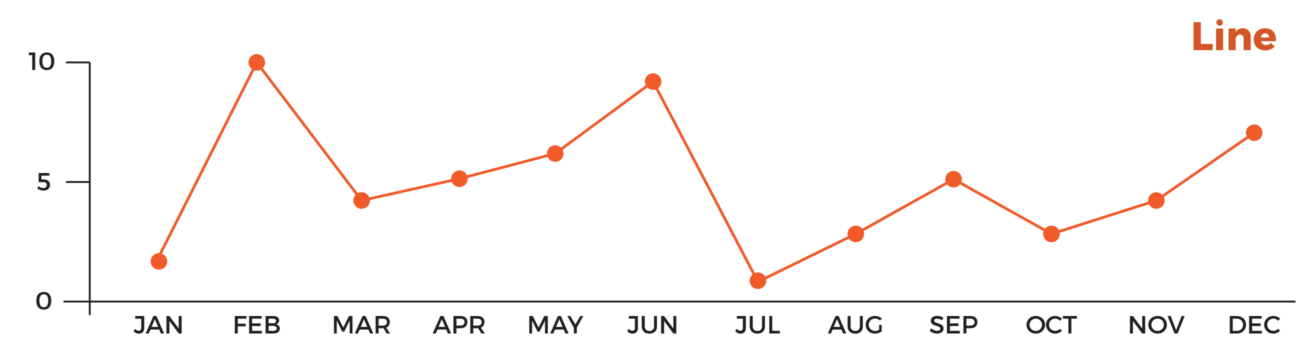

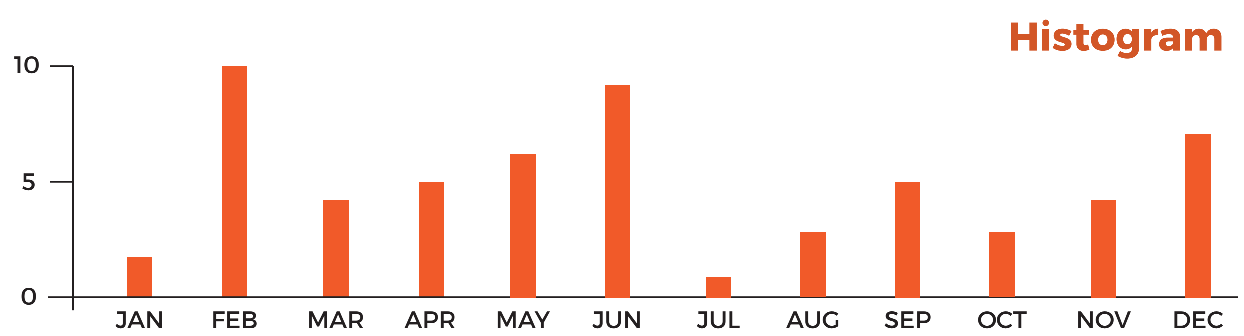

How can you encode information optimally?

JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
2	10	4	5	6	9	1	3	5	3	4	7

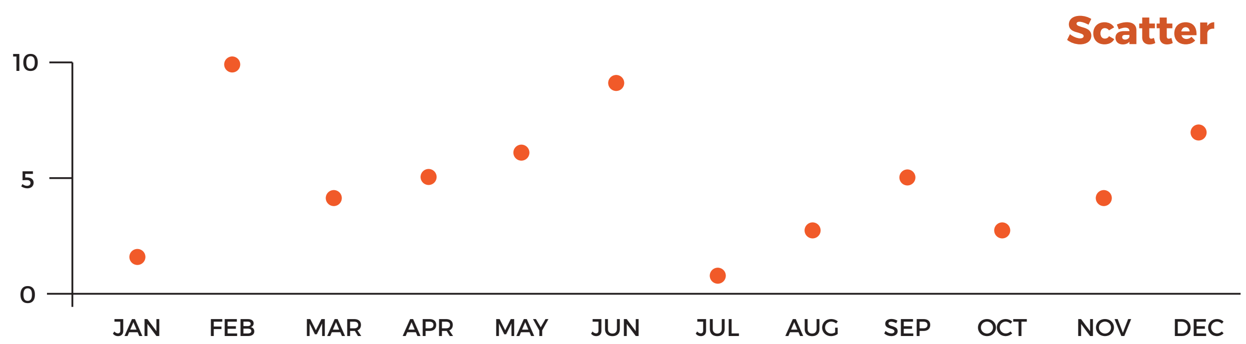

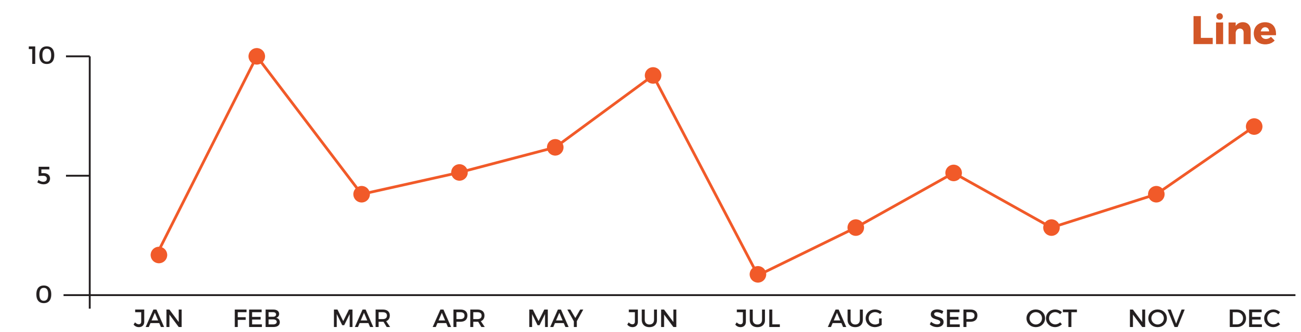

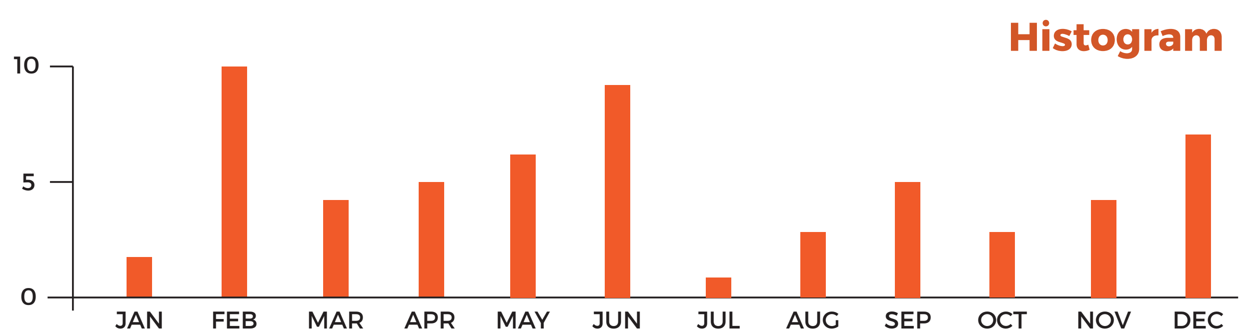

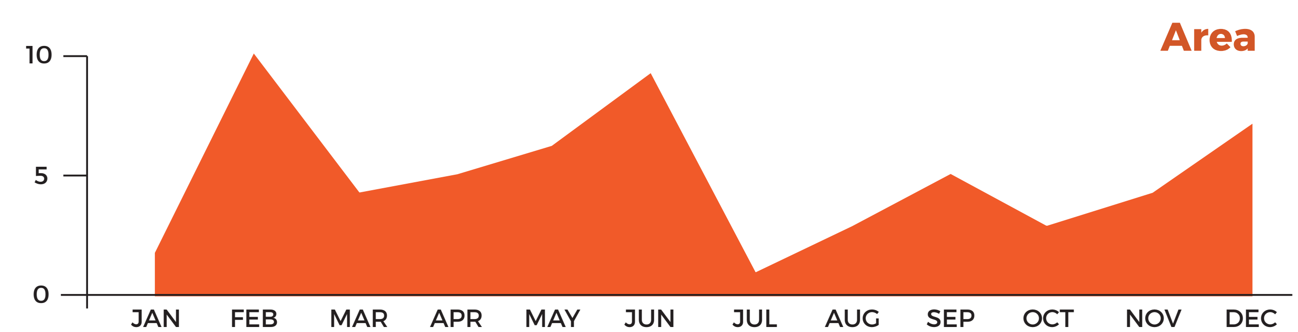

How can you encode information optimally?

JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
2	10	4	5	6	9	1	3	5	3	4	7

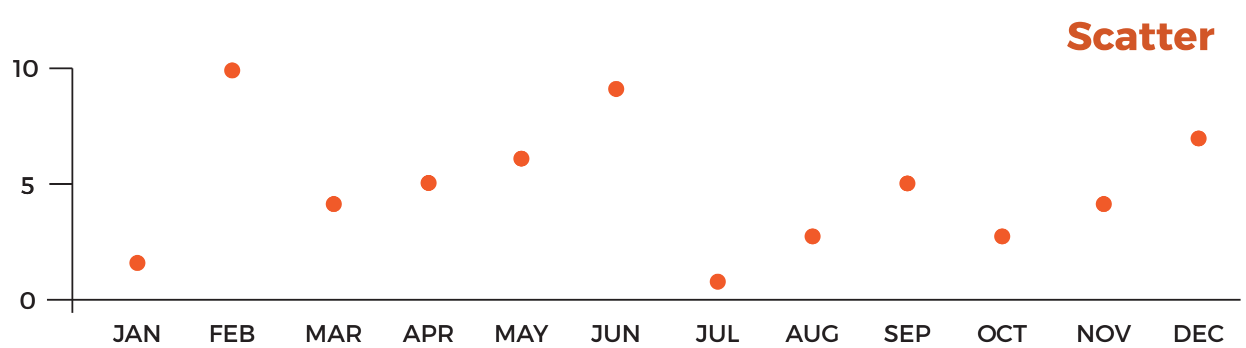

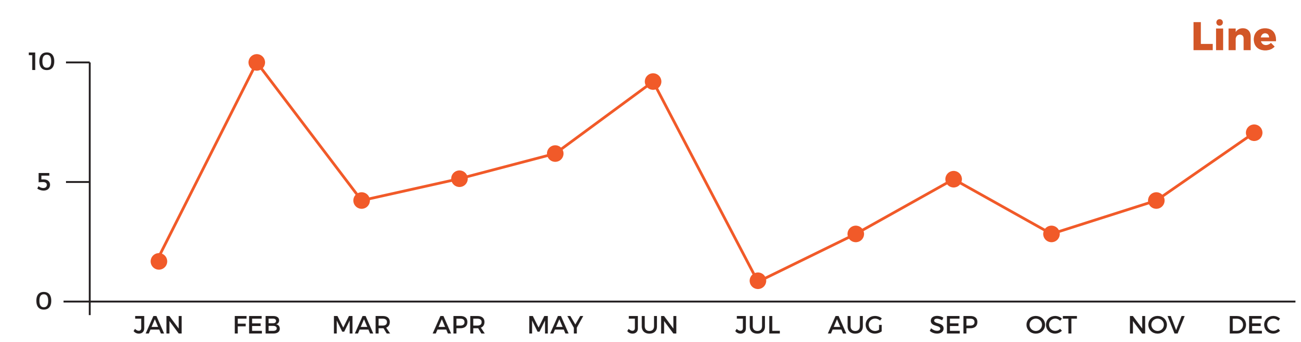

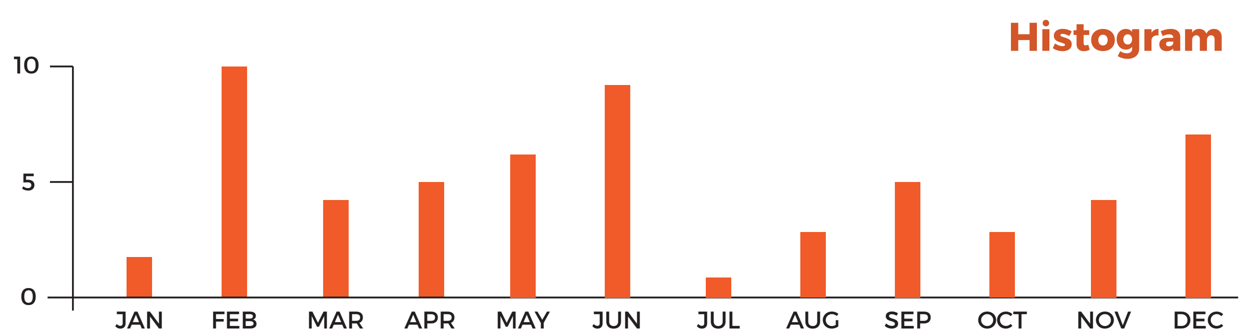

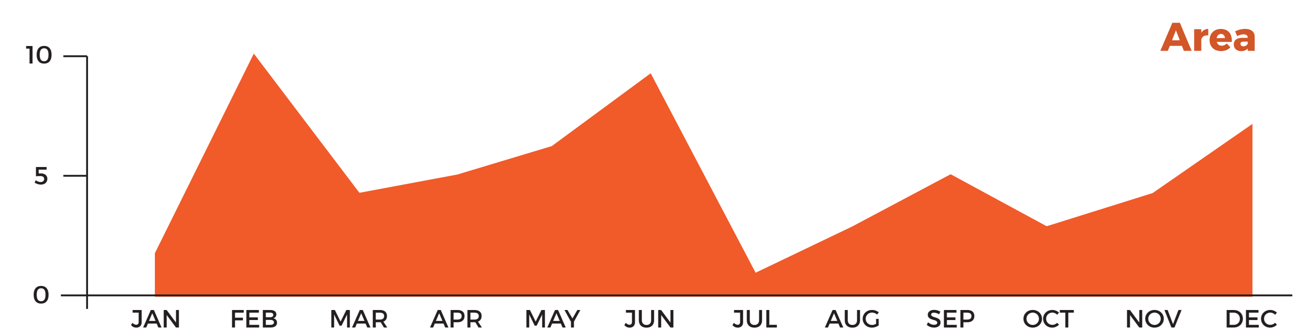

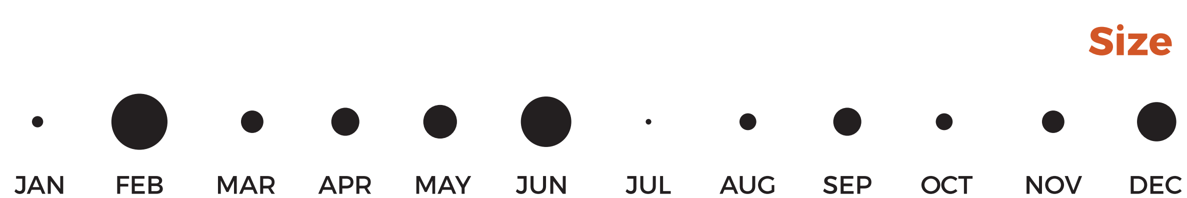

How can you encode information optimally?

JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
2	10	4	5	6	9	1	3	5	3	4	7

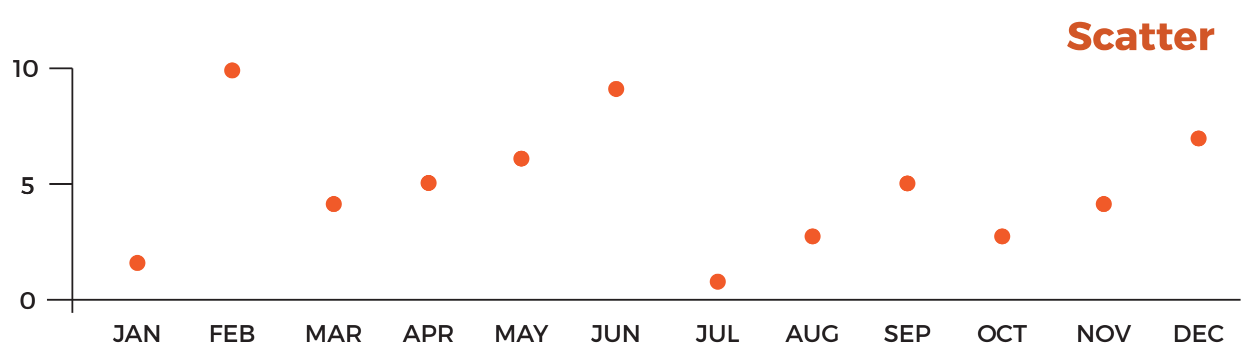

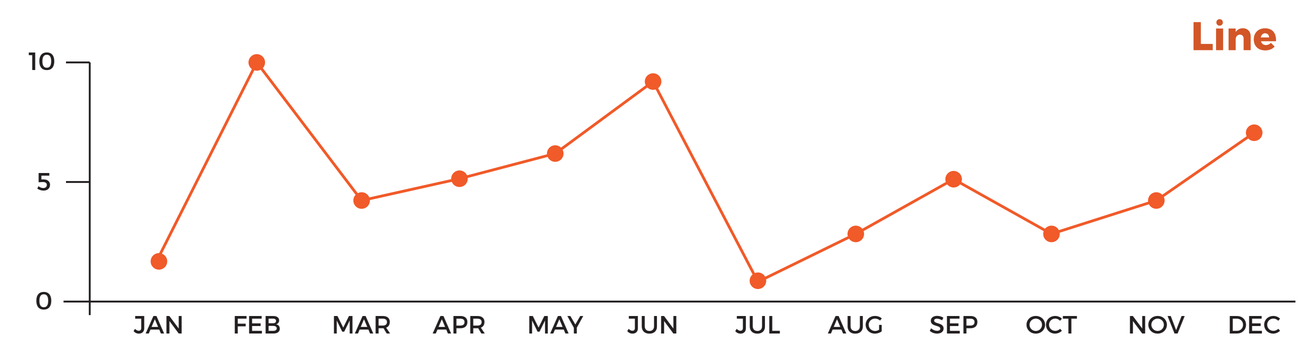

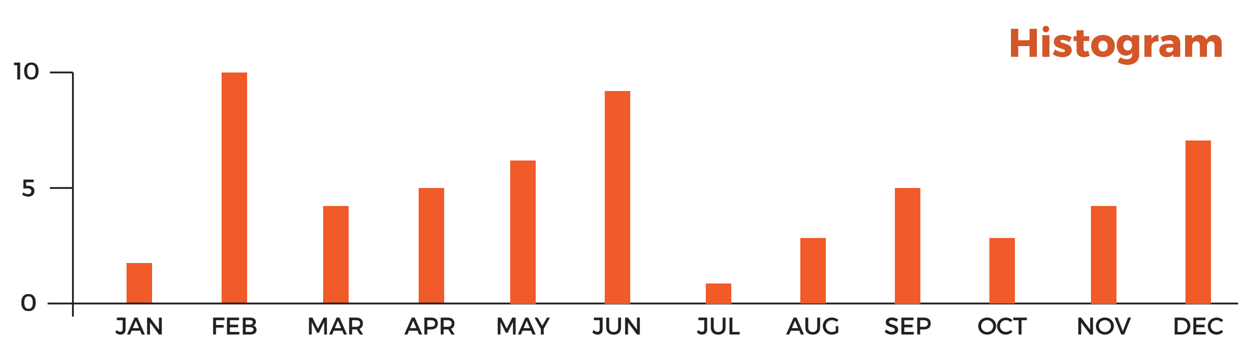

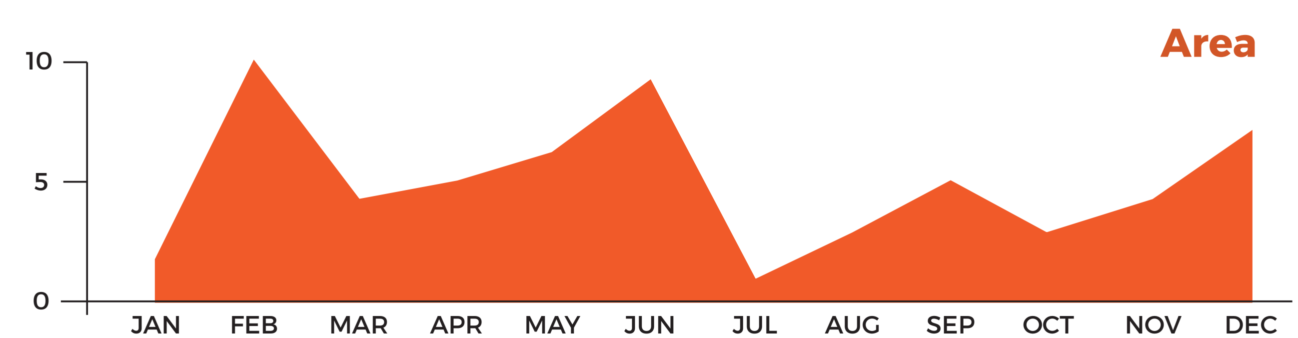

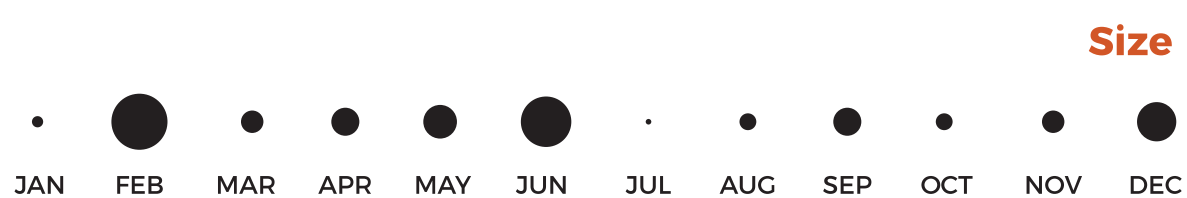

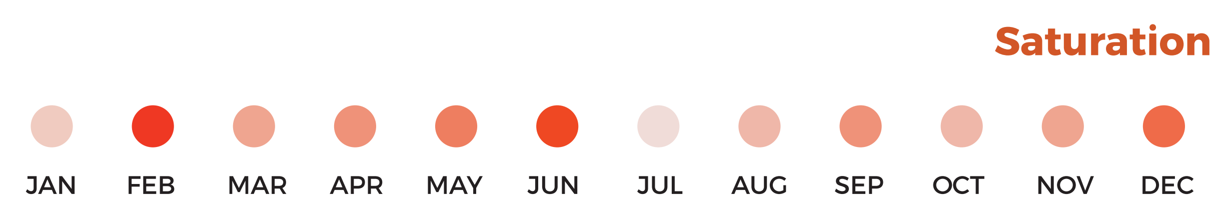

How can you encode information optimally?

JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
2	10	4	5	6	9	1	3	5	3	4	7

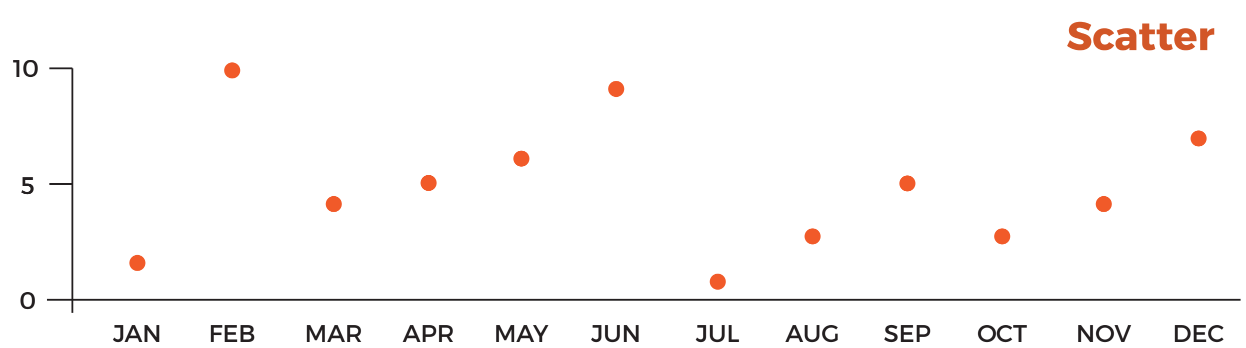

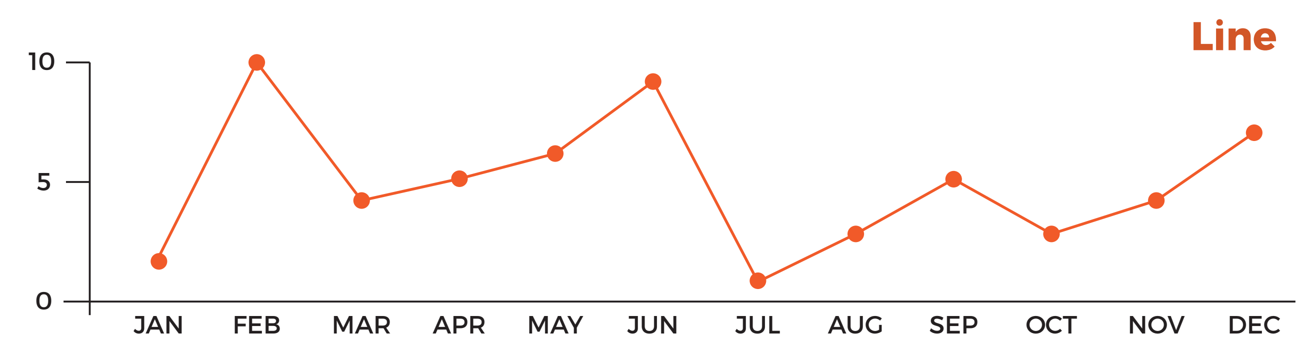

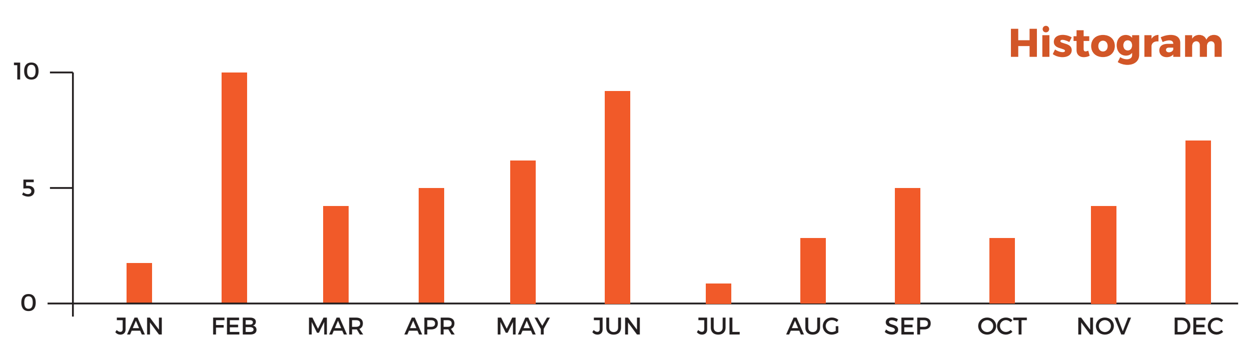

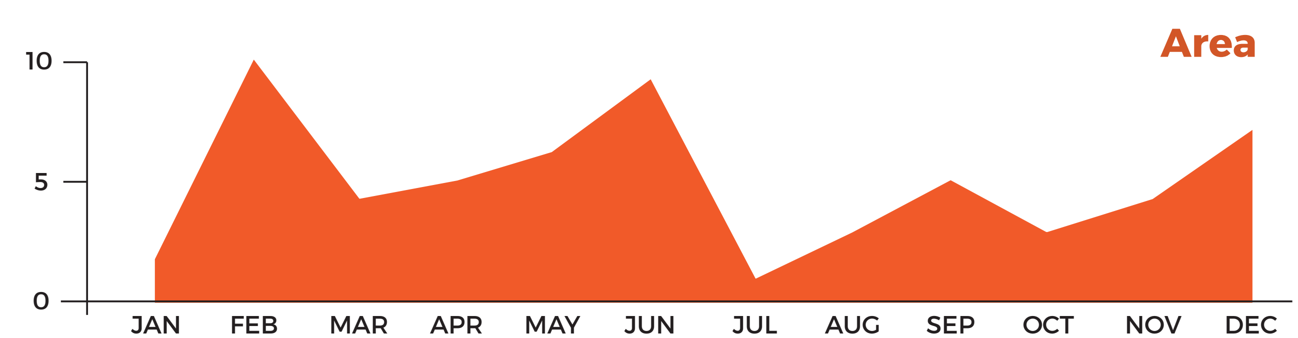

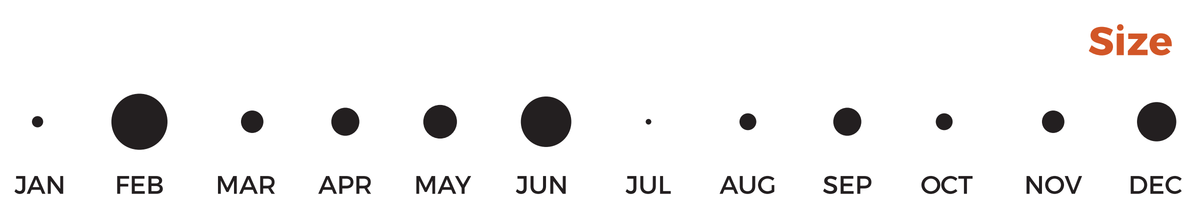

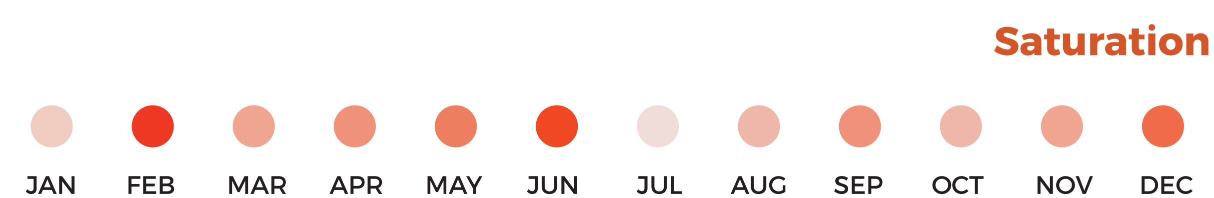

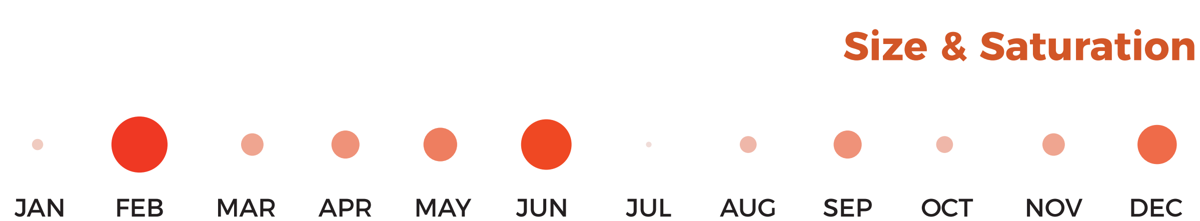

How can you encode information optimally?

JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
2	10	4	5	6	9	1	3	5	3	4	7

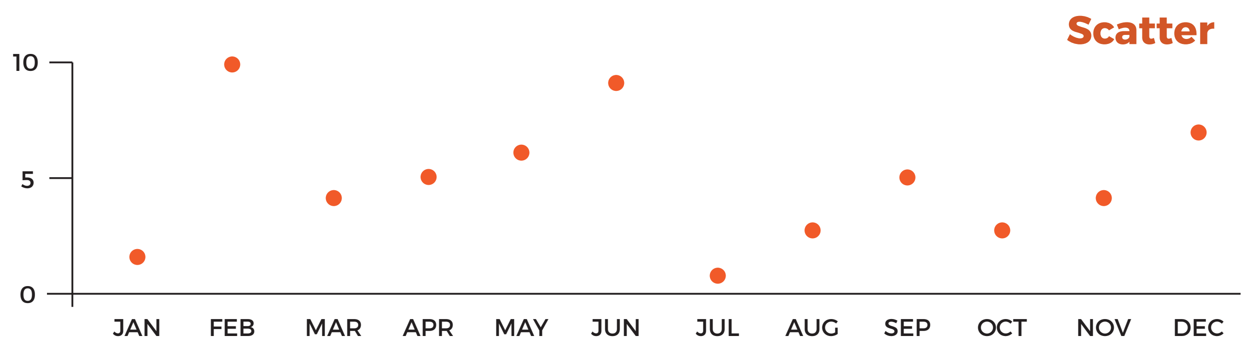

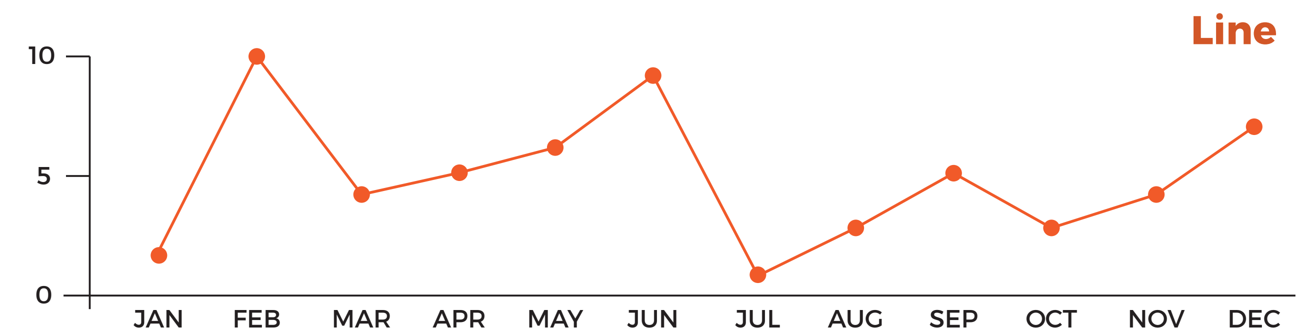

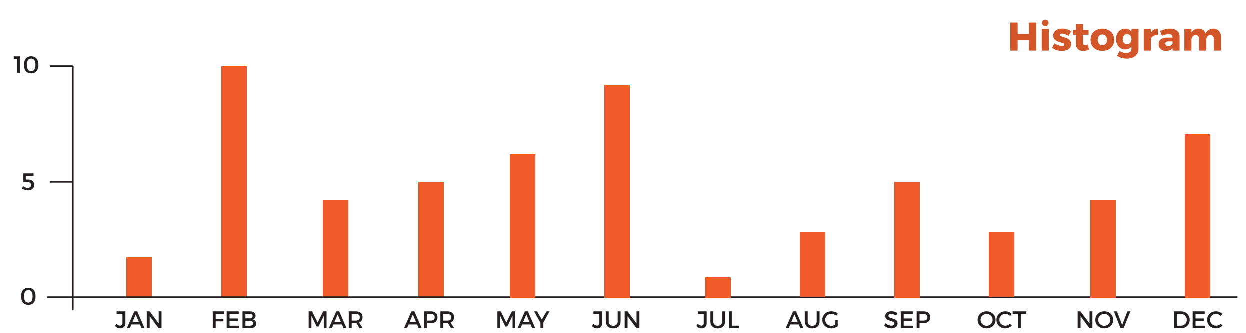

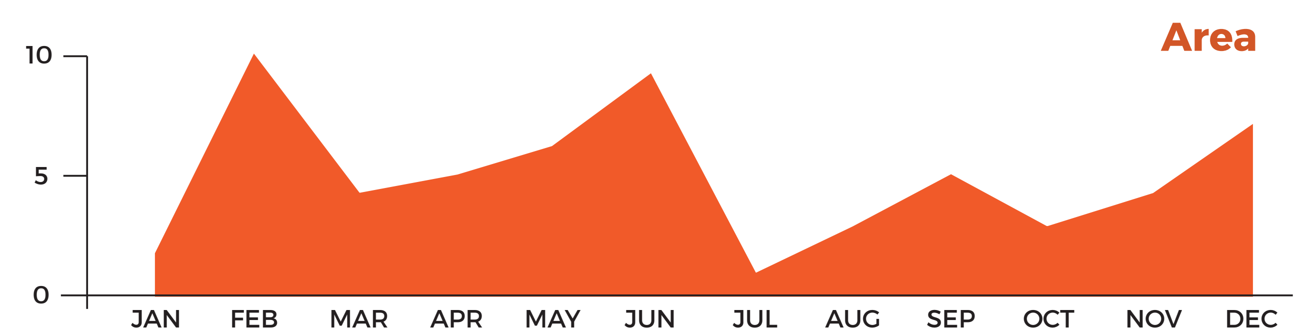

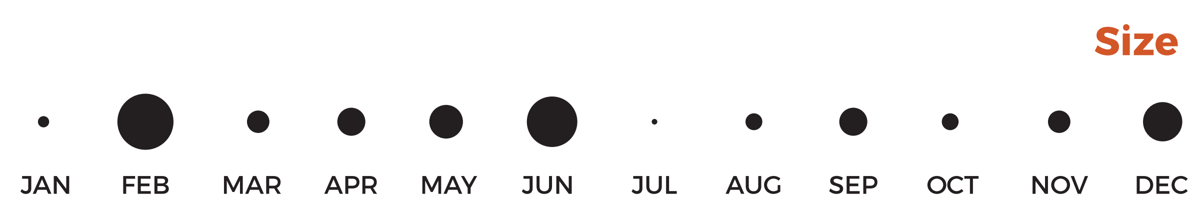

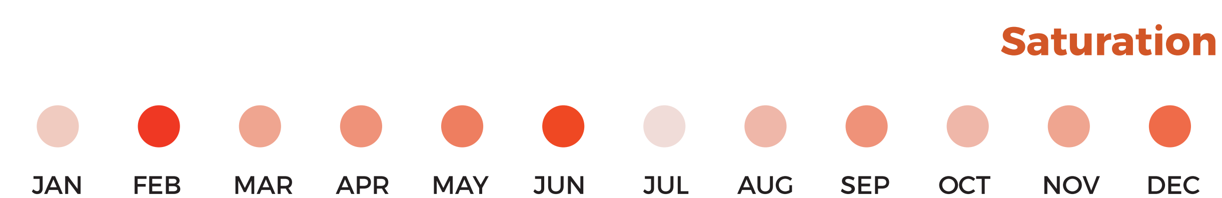

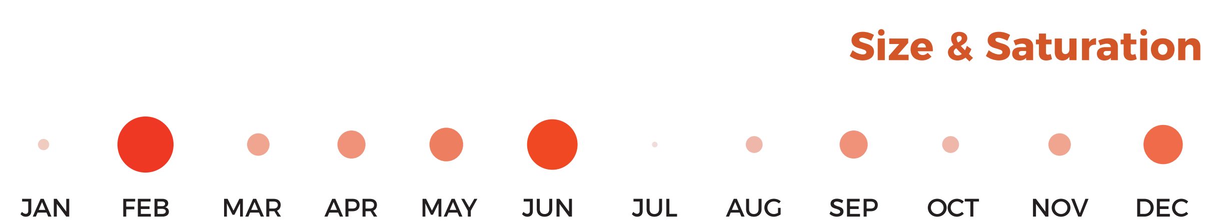

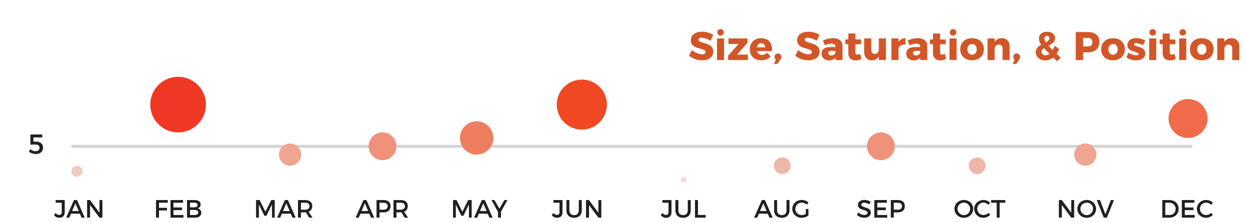

And that's just a really simple low dimensional example

Moreover, all of these visualizations encode the information,
but the decode error (interpreting, comparing, …) for each graph is different

But, why?

Our perception system does not behave linearly.

Some stimuli are perceived less or more than intended.

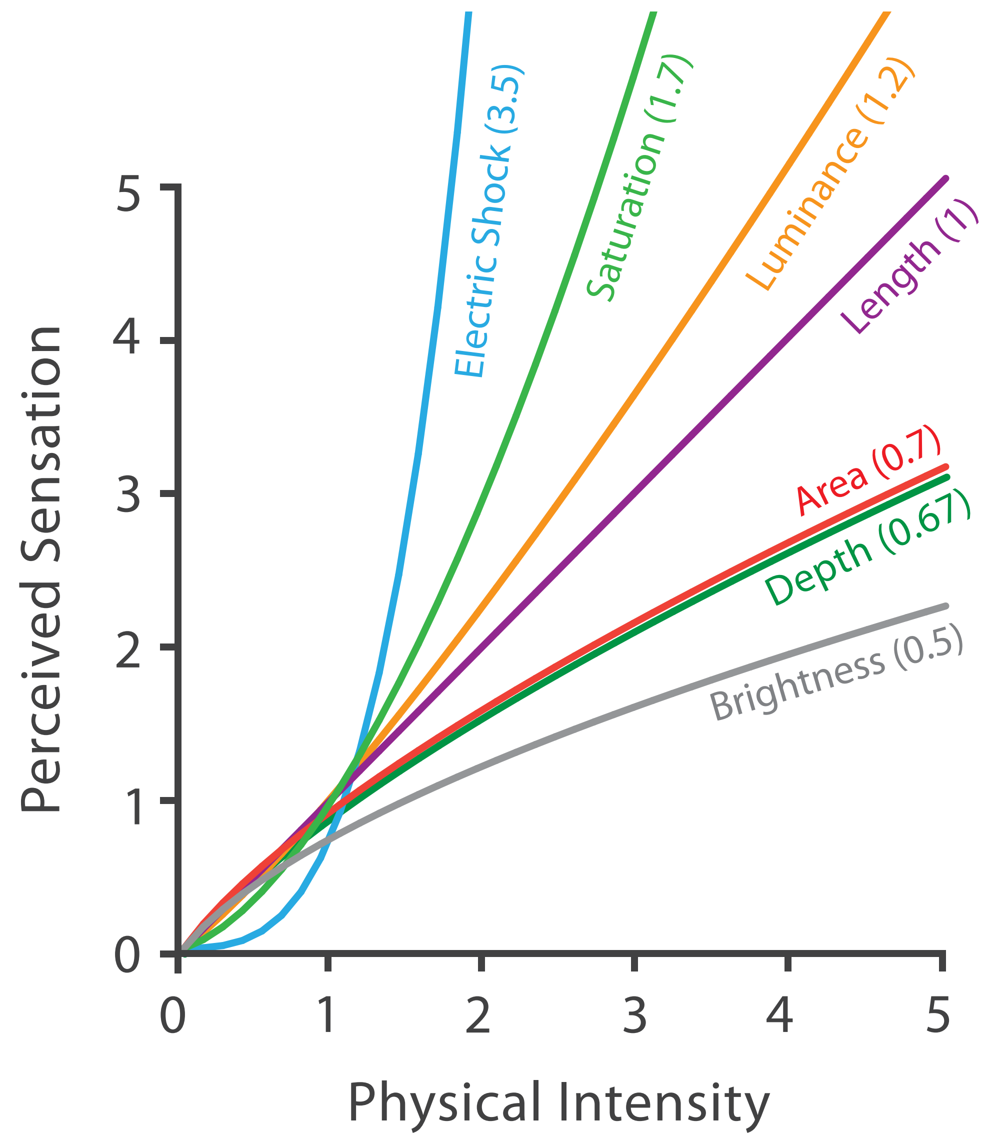

Steven's Psychophysical Power Law: $S = I^N$

Stevens, 1975

We have to be careful when mapping data to the visual world

Some visual channels are more effective for some data types over others.

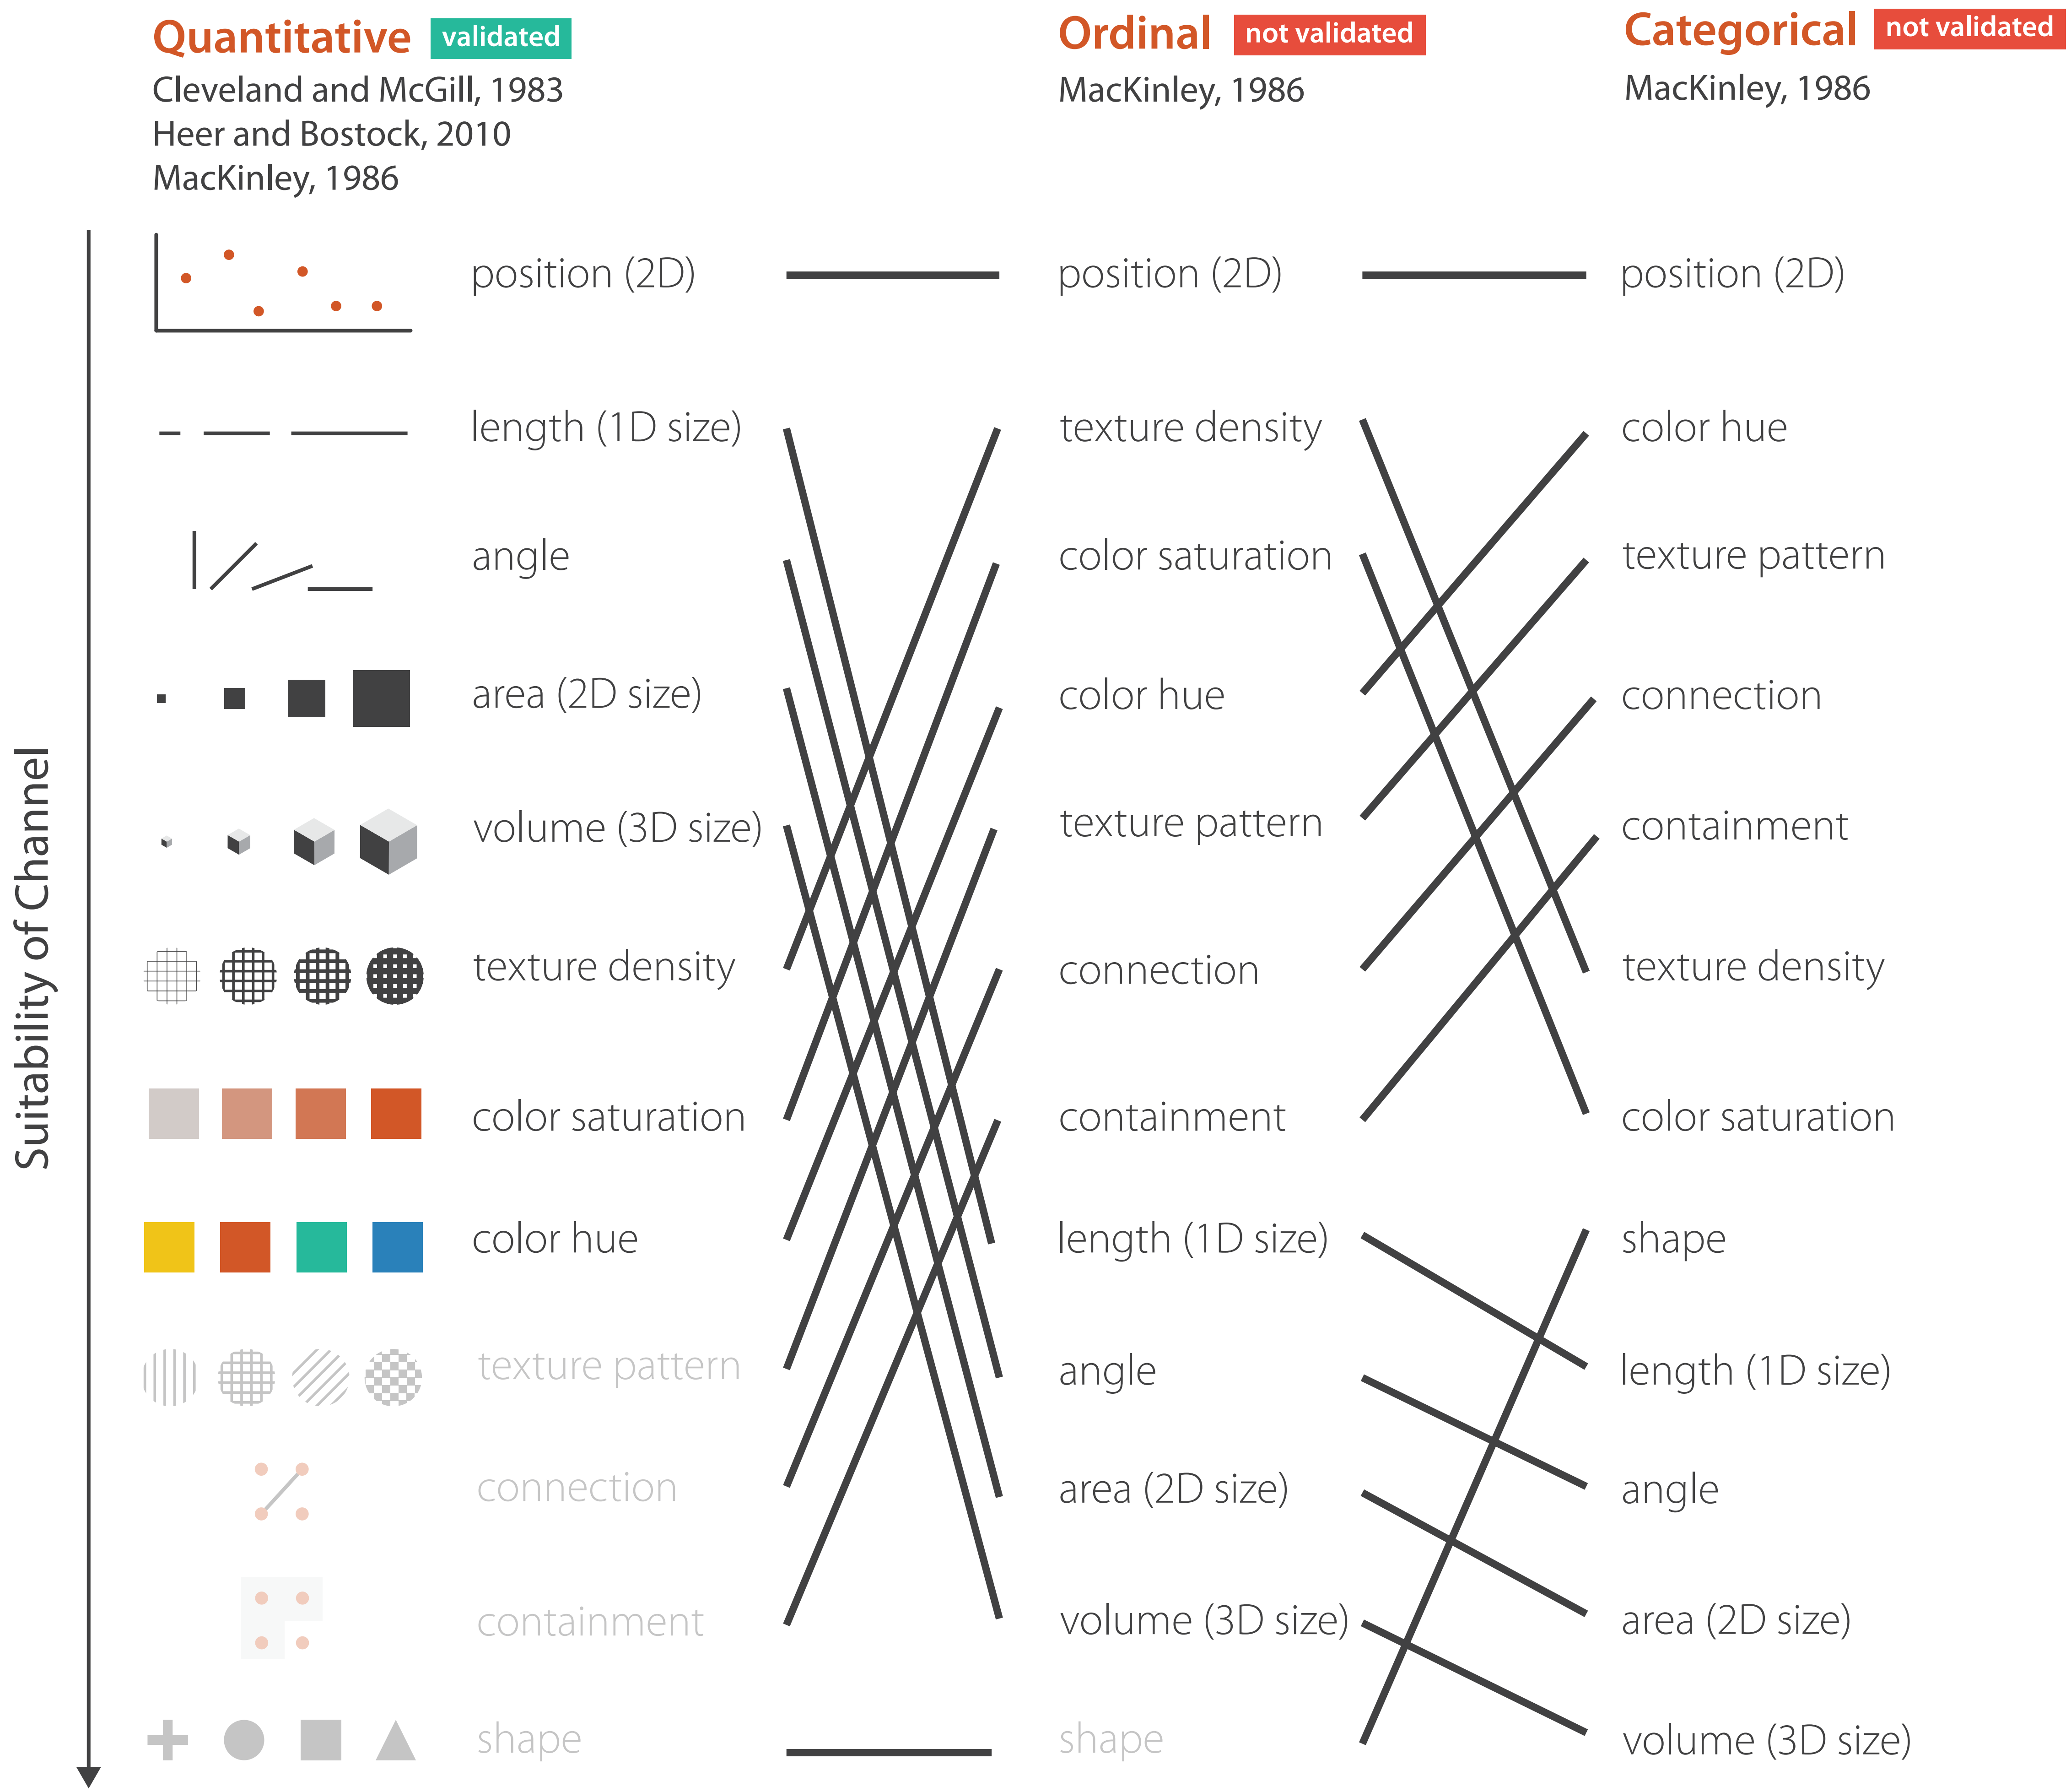

Quantitative
validated
Cleveland and McGill, 1983
Heer and Bostock, 2010
MacKinley, 1986
Ordinal
not validated
MacKinley, 1986
Categorical
not validated
MacKinley, 1986
Suitability of Channel
position (2D)
length (1D size)
angle
area (2D size)
volume (3D size)
texture density
color saturation
color hue
texture pattern
connection
containment
shape
position (2D)
texture density
color saturation
color hue
texture pattern
connection
containment
length (1D size)
angle
area (2D size)
volume (3D size)
shape
position (2D)
color hue
texture pattern
connection
containment
texture density
color saturation
shape
length (1D size)
angle
area (2D size)
volume (3D size)

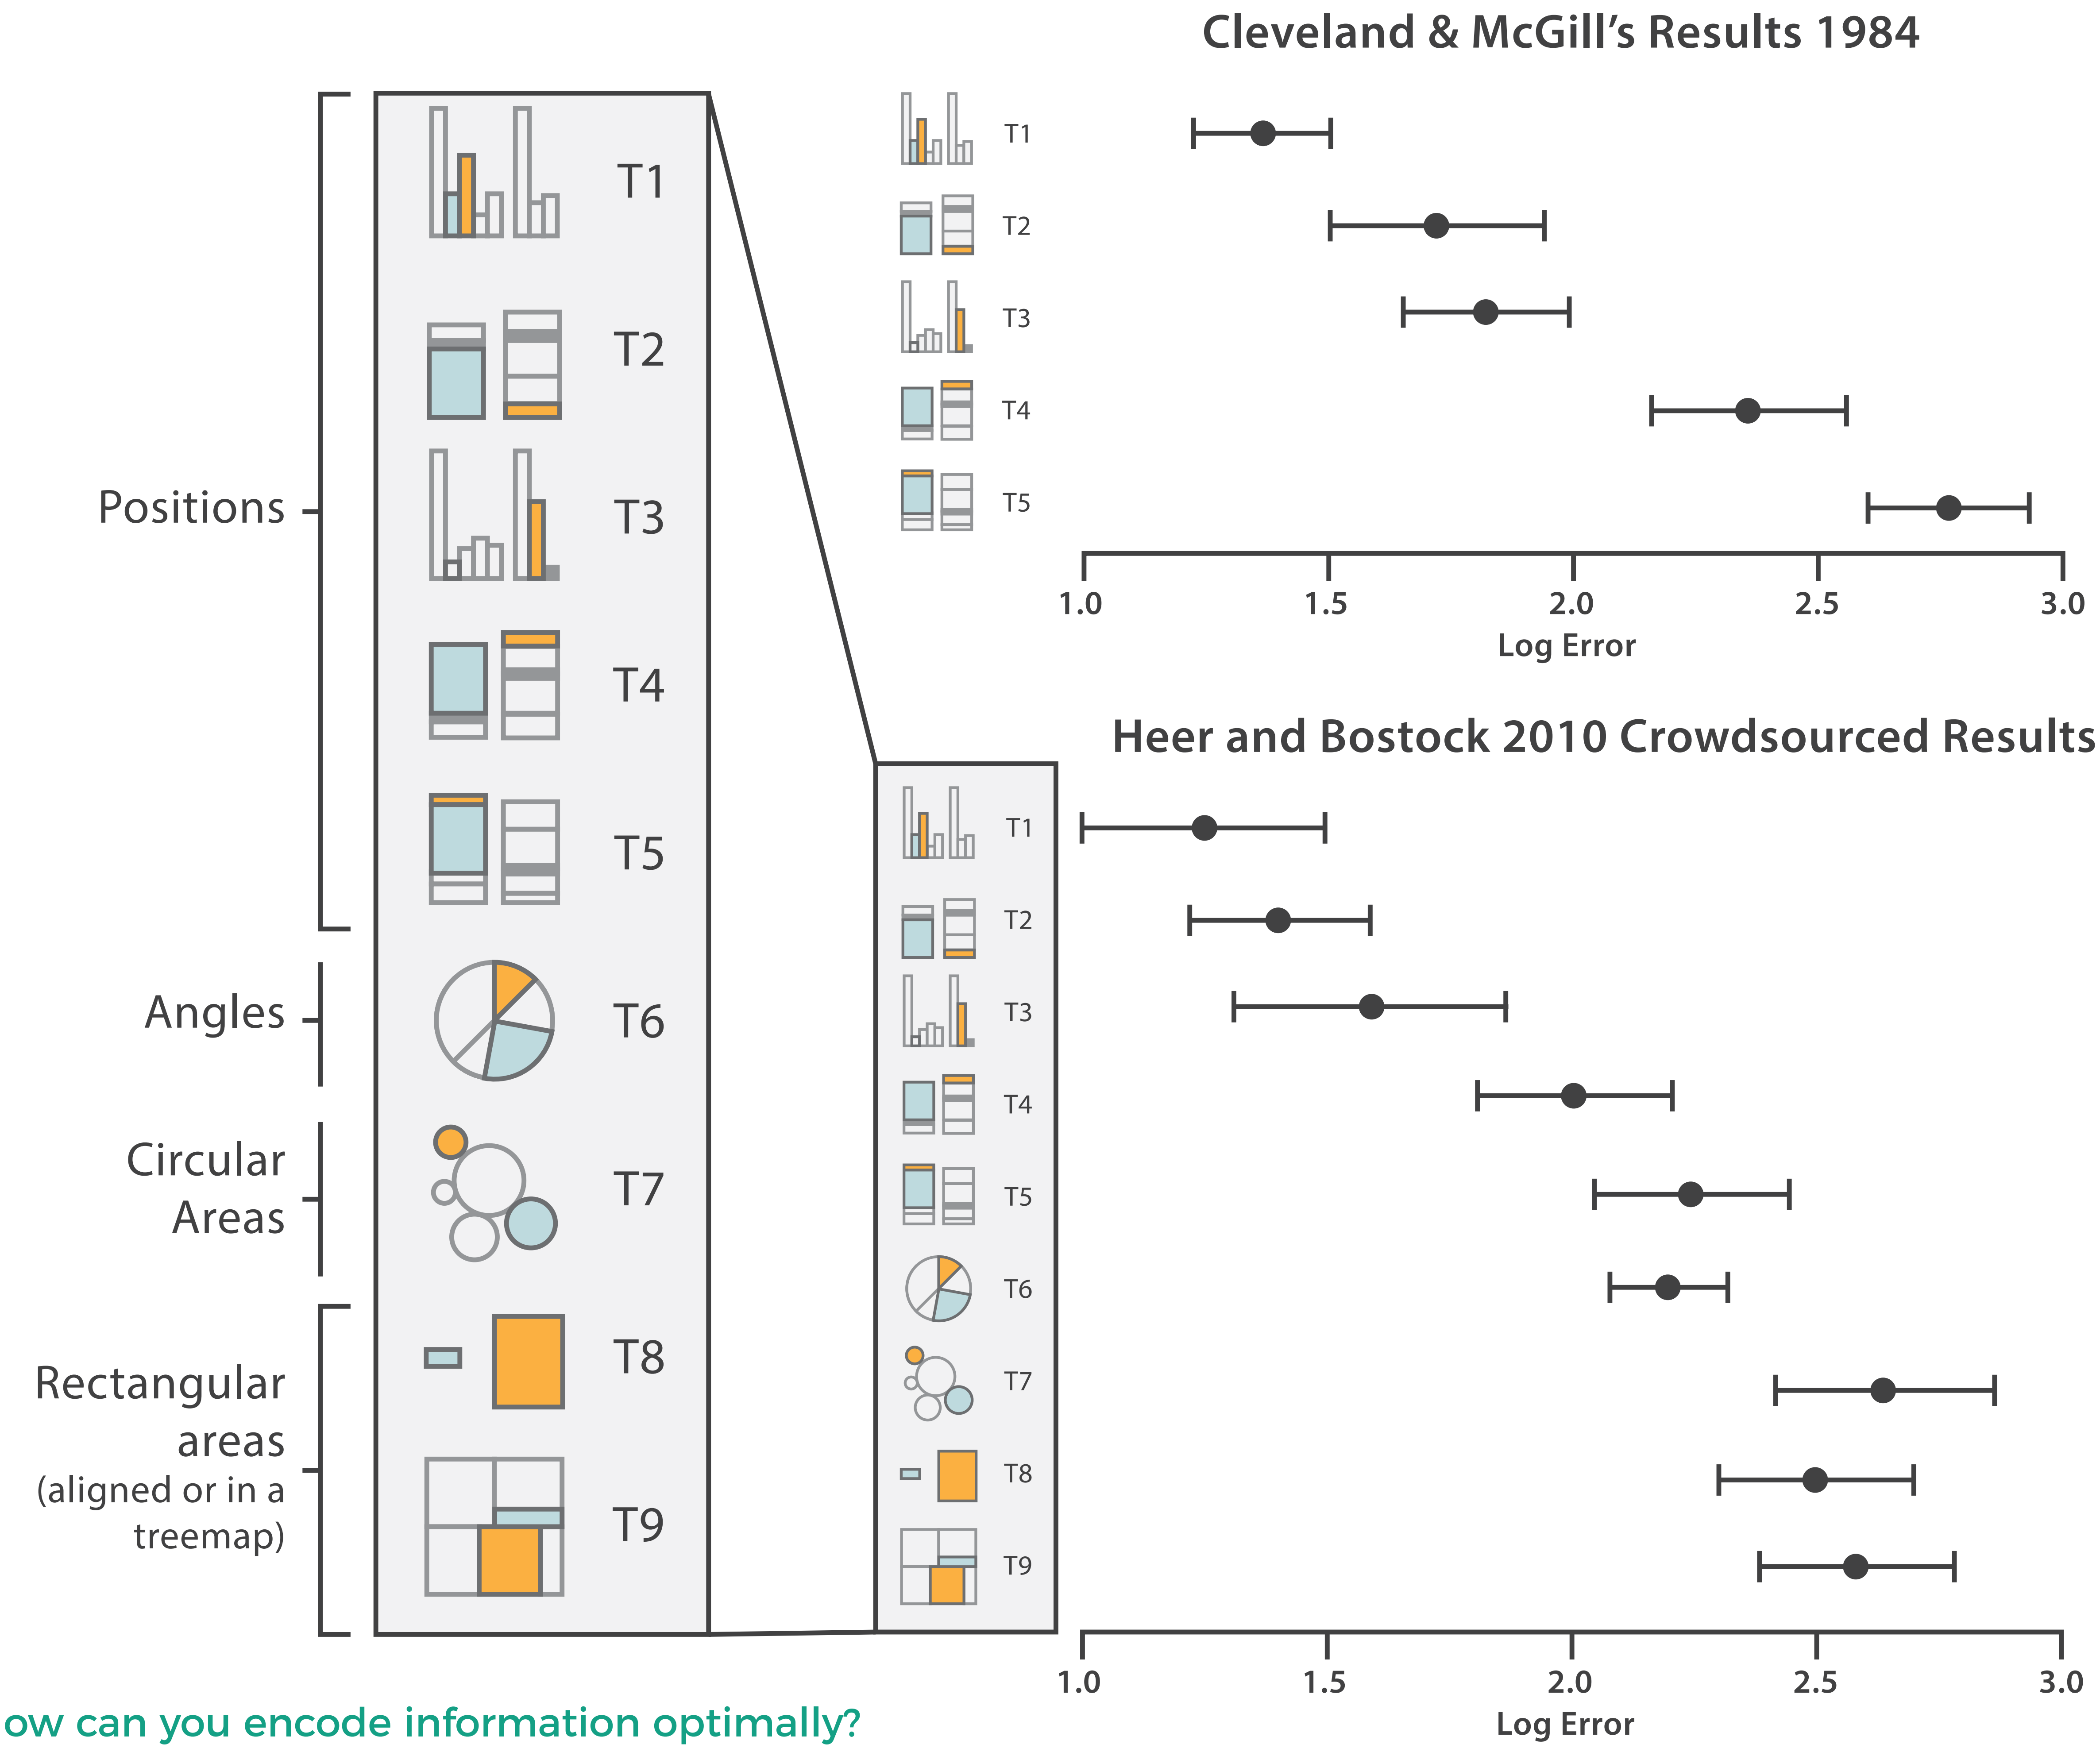

How can you encode information optimally?

T6: Pie charts have also been studied in more detail recently

It's quite clear that bar charts are a more effective visual encoding here than pie charts… our visual system is very good at judging lengths, but not so much at judging angles and areas.

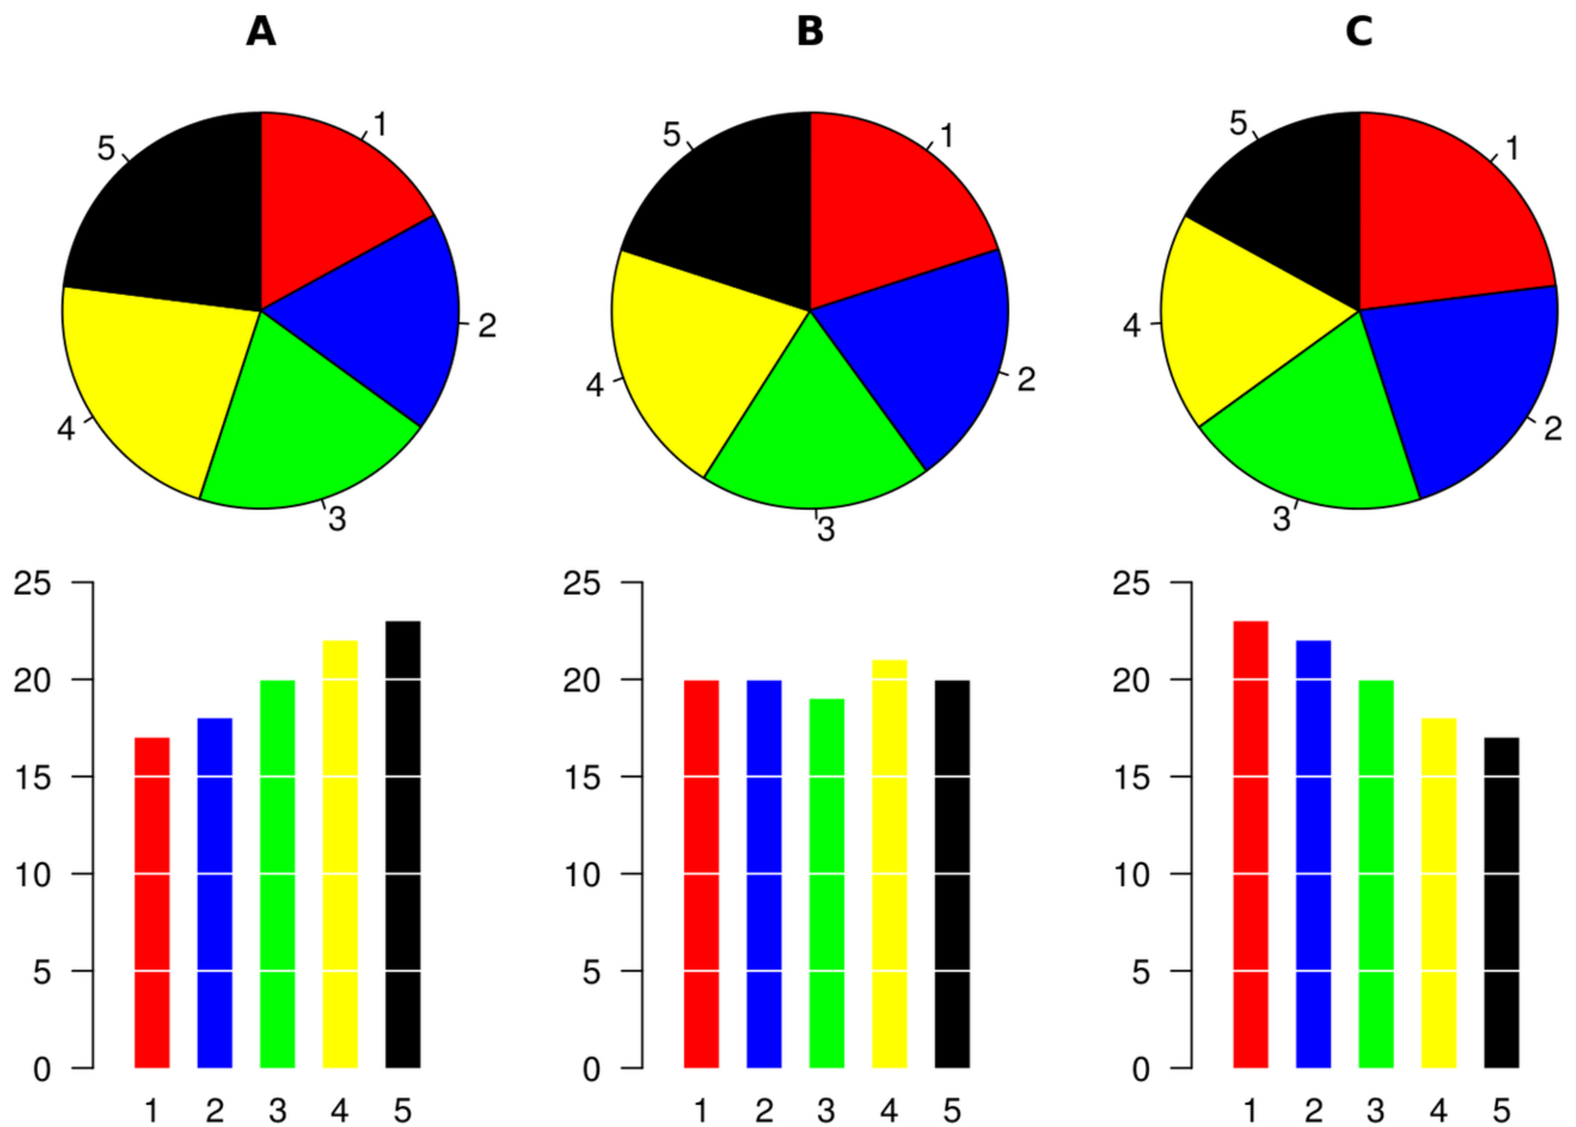

https://commons.wikimedia.org/wiki/File:Piecharts.svg

T6: Pie charts have also been studied in more detail recently

When someone reads or compares values in a pie chart, what are they doing? Comparing angles, areas, length of arc?

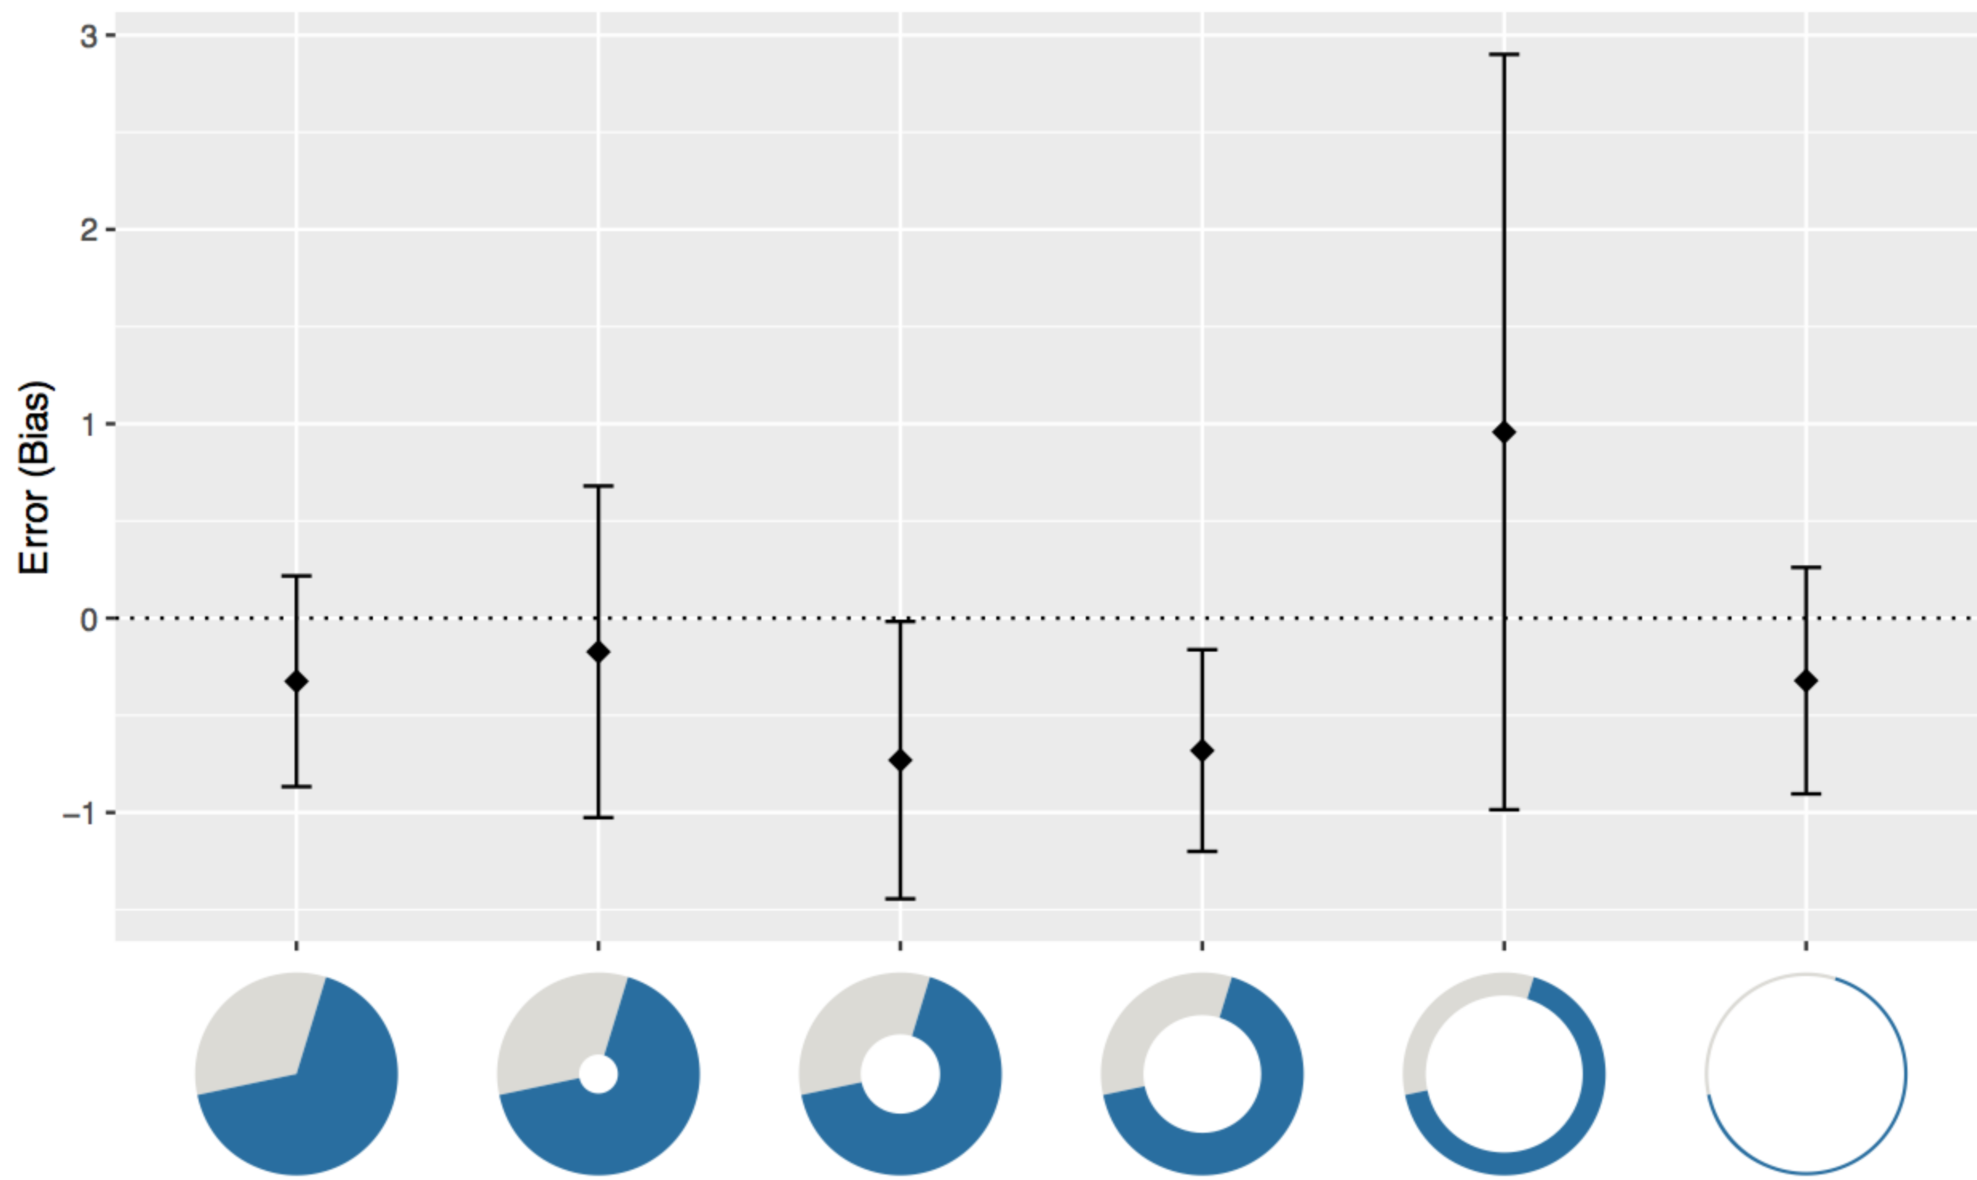

Robert Kosara and Drew Skau. 2016. **Judgment error in pie chart variations**. In Proceedings of the Eurographics: Short Papers (EuroVis '16). Eurographics Association, Goslar Germany, Germany, 91-95. DOI: https://doi.org/10.2312/eurovisshort.20161167

Drew Skau and Robert Kosara. 2016. **Arcs, Angles, or Areas: Individual Data Encodings in Pie and Donut Charts**. Comput. Graph. Forum 35, 3 (June 2016), 121-130. DOI: https://doi.org/10.1111/cgf.12888

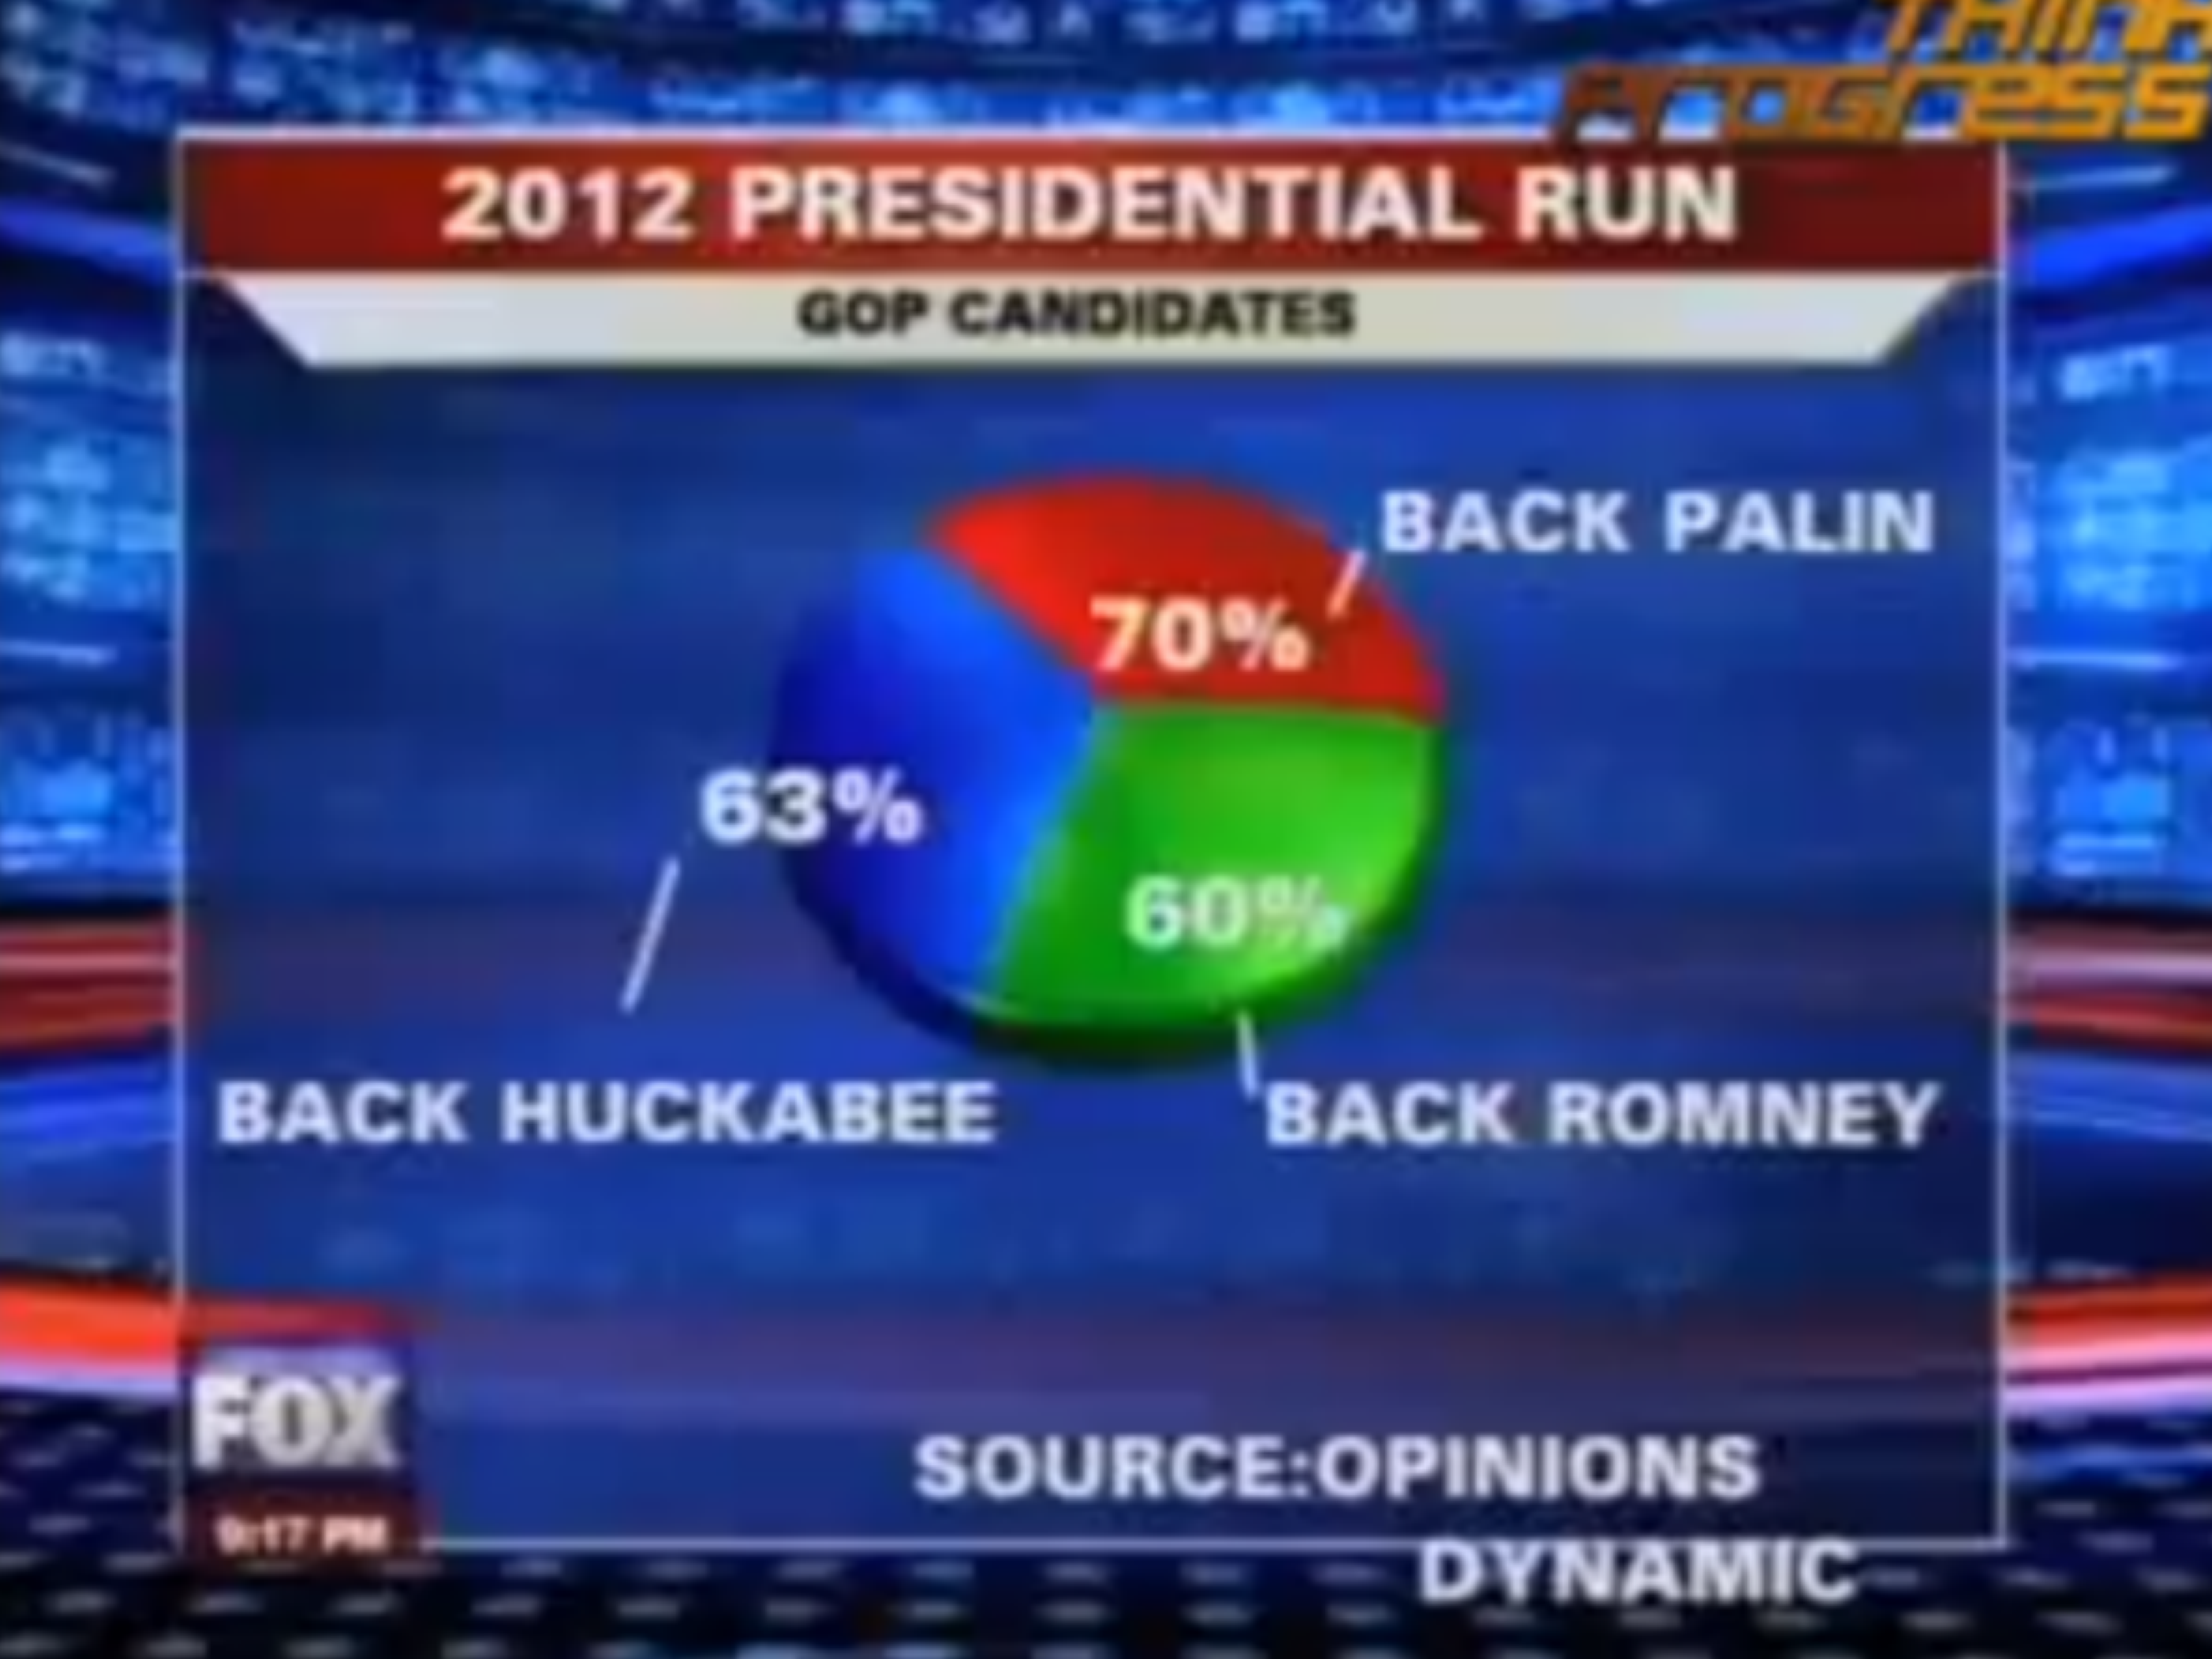
2012 PRESIDENTIAL RUN
GOP CANDIDATES
BACK PALIN
70%
63%
60%
BACK HUCKABEE
BACK ROMNEY
FOX
SOURCE:OPINIONS
DYNAMIC

T1/T7: Bar charts are better than areas...

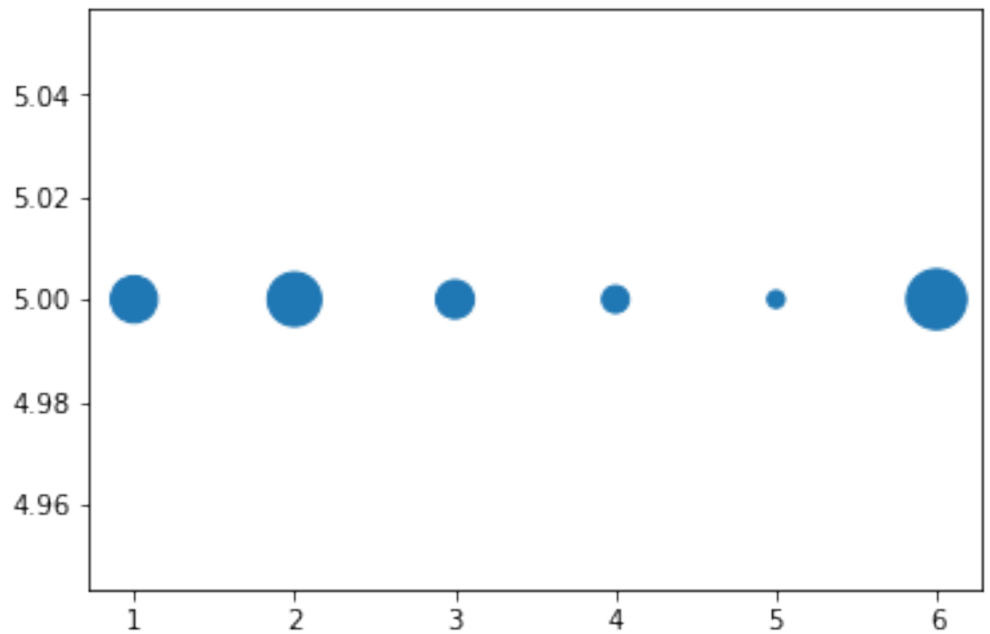

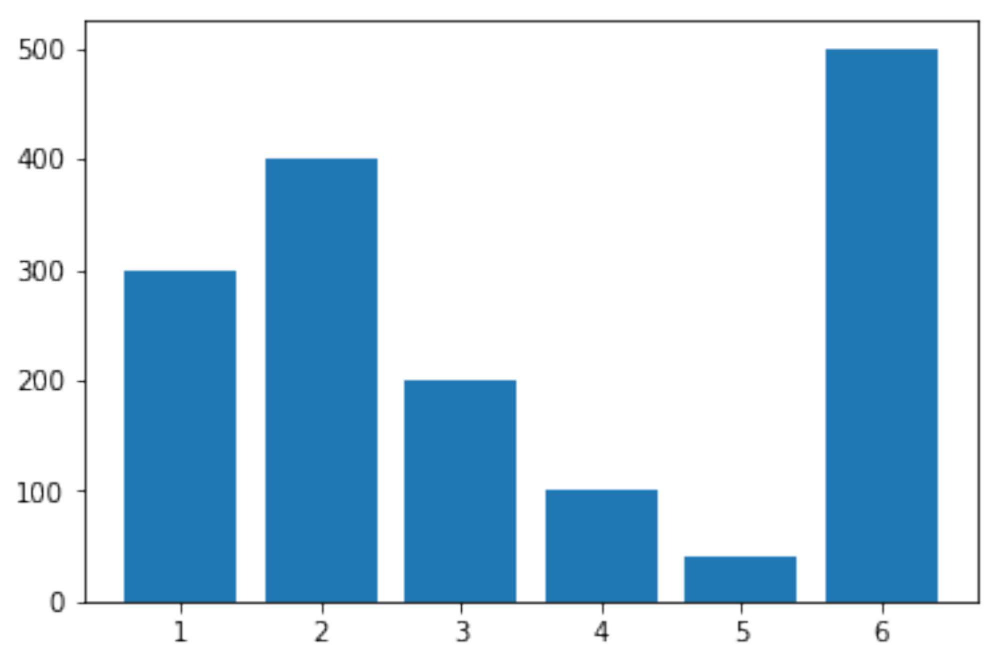

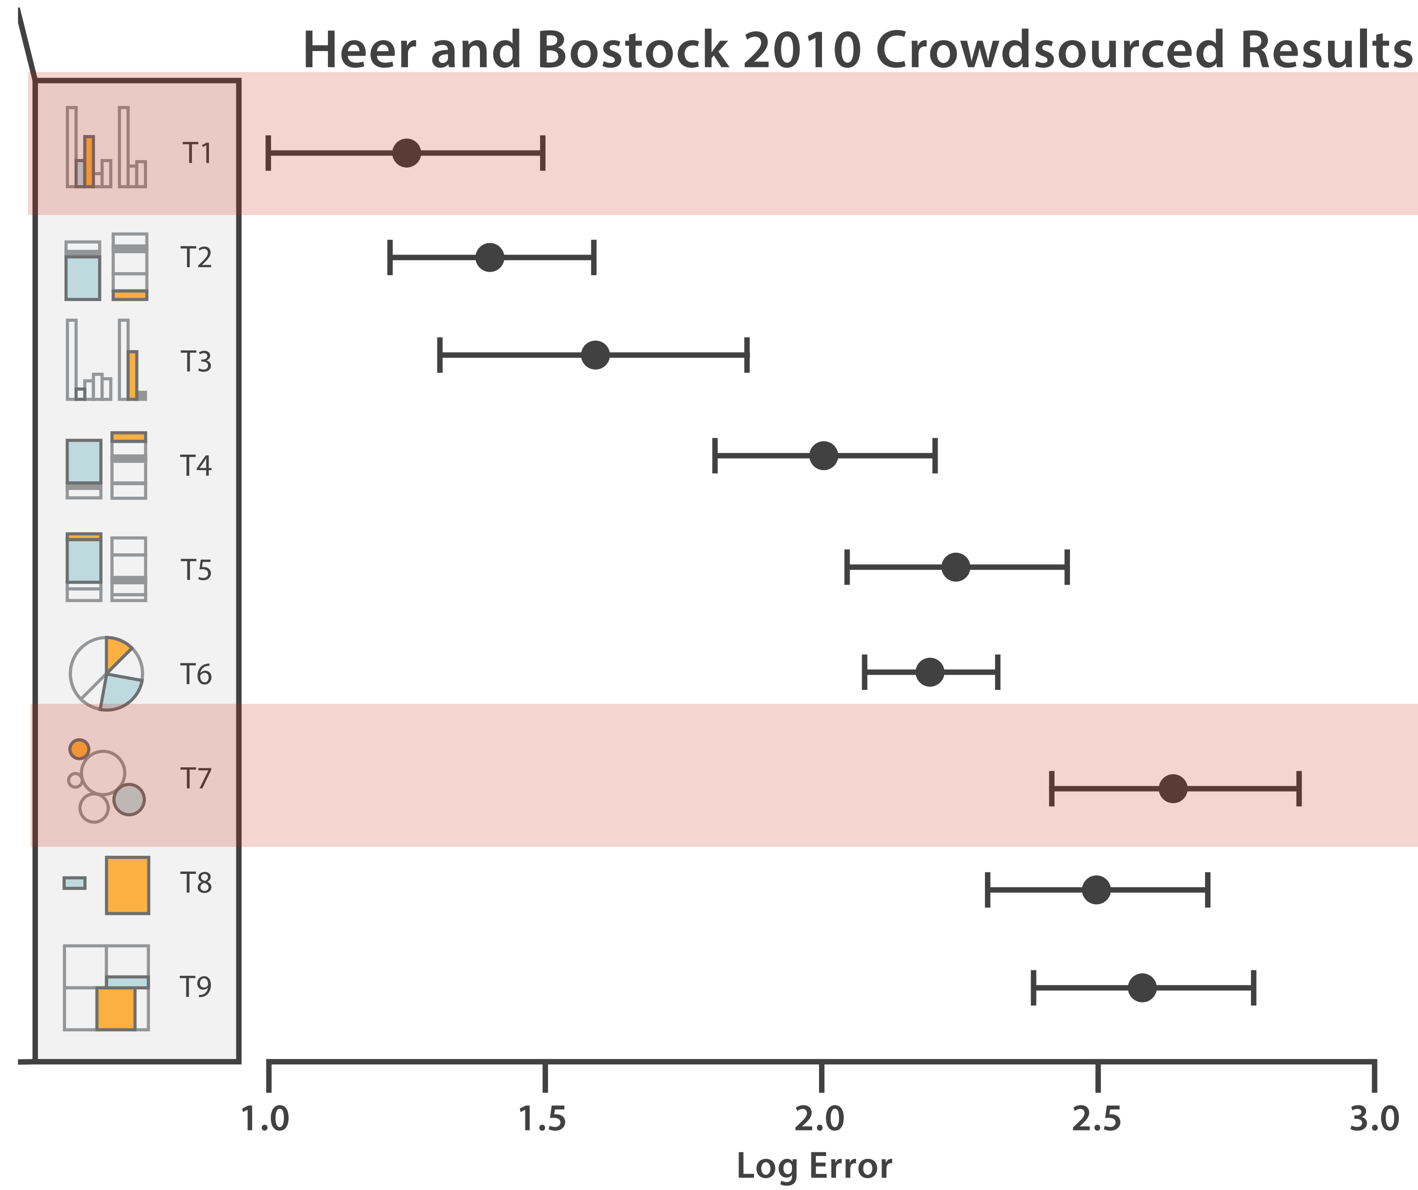

T1/T7: Bar charts are better than areas...

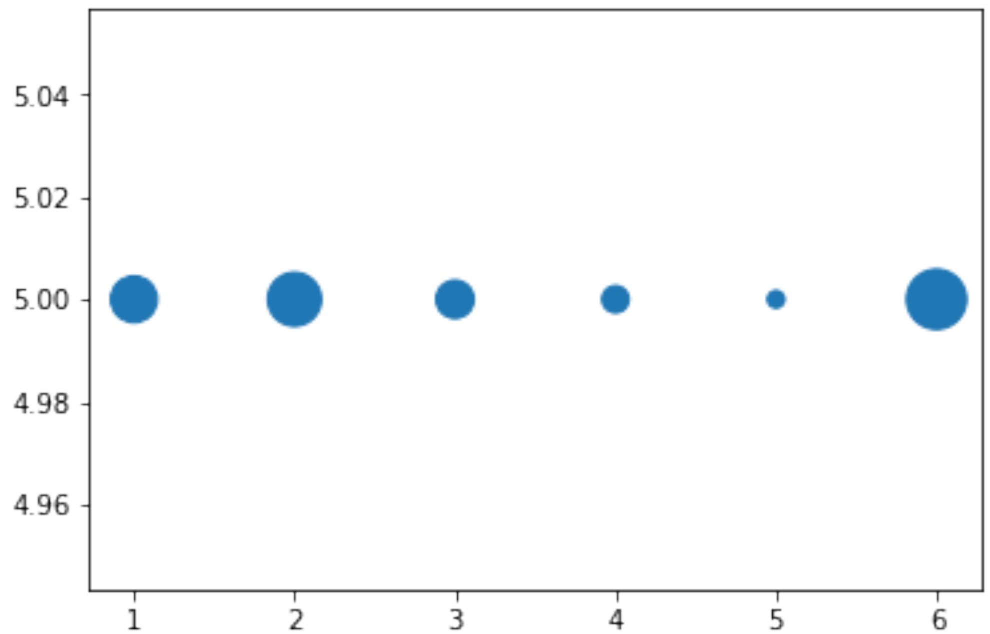

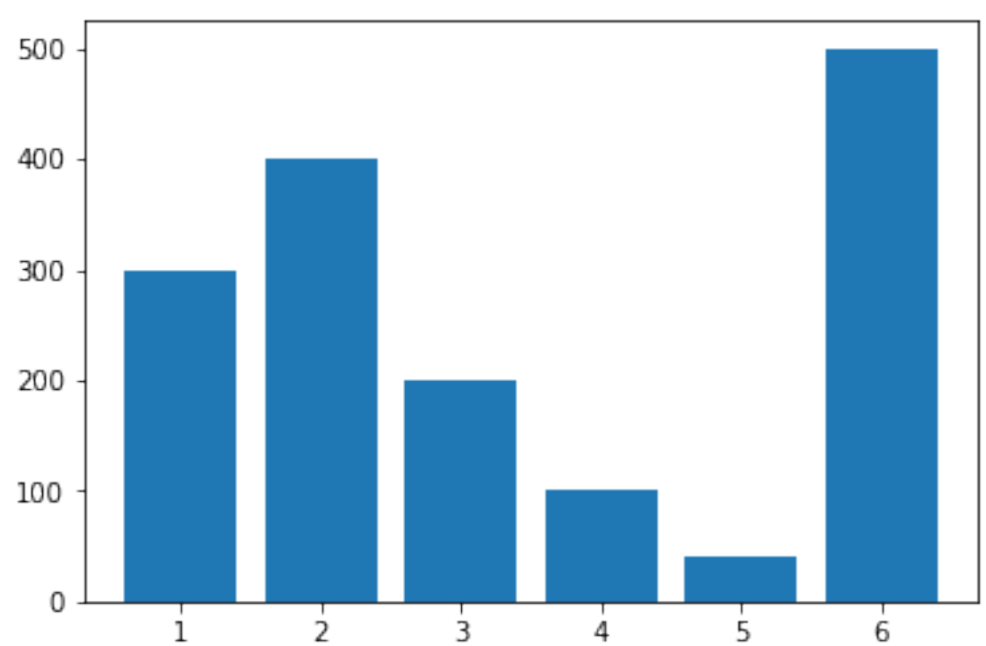

300 400 200 100 40 500

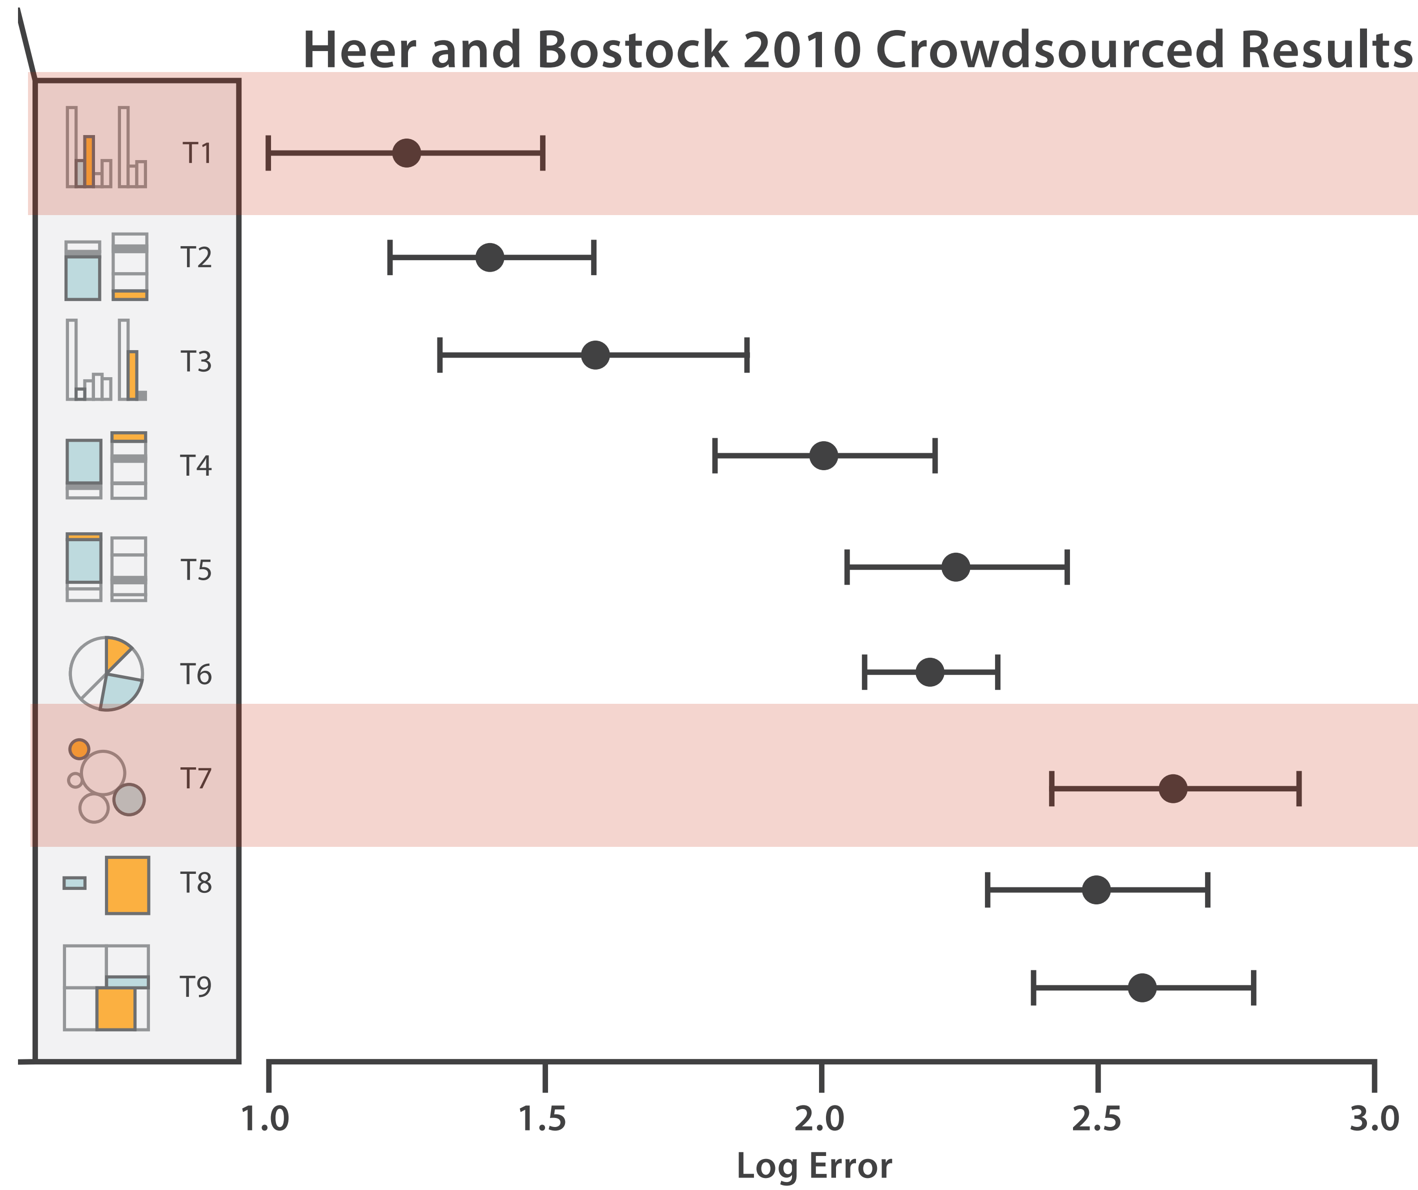

This is exactly the same data, at the right scaling.

T8/T9: Different aspect ratios for rectangles also result in greater or fewer errors

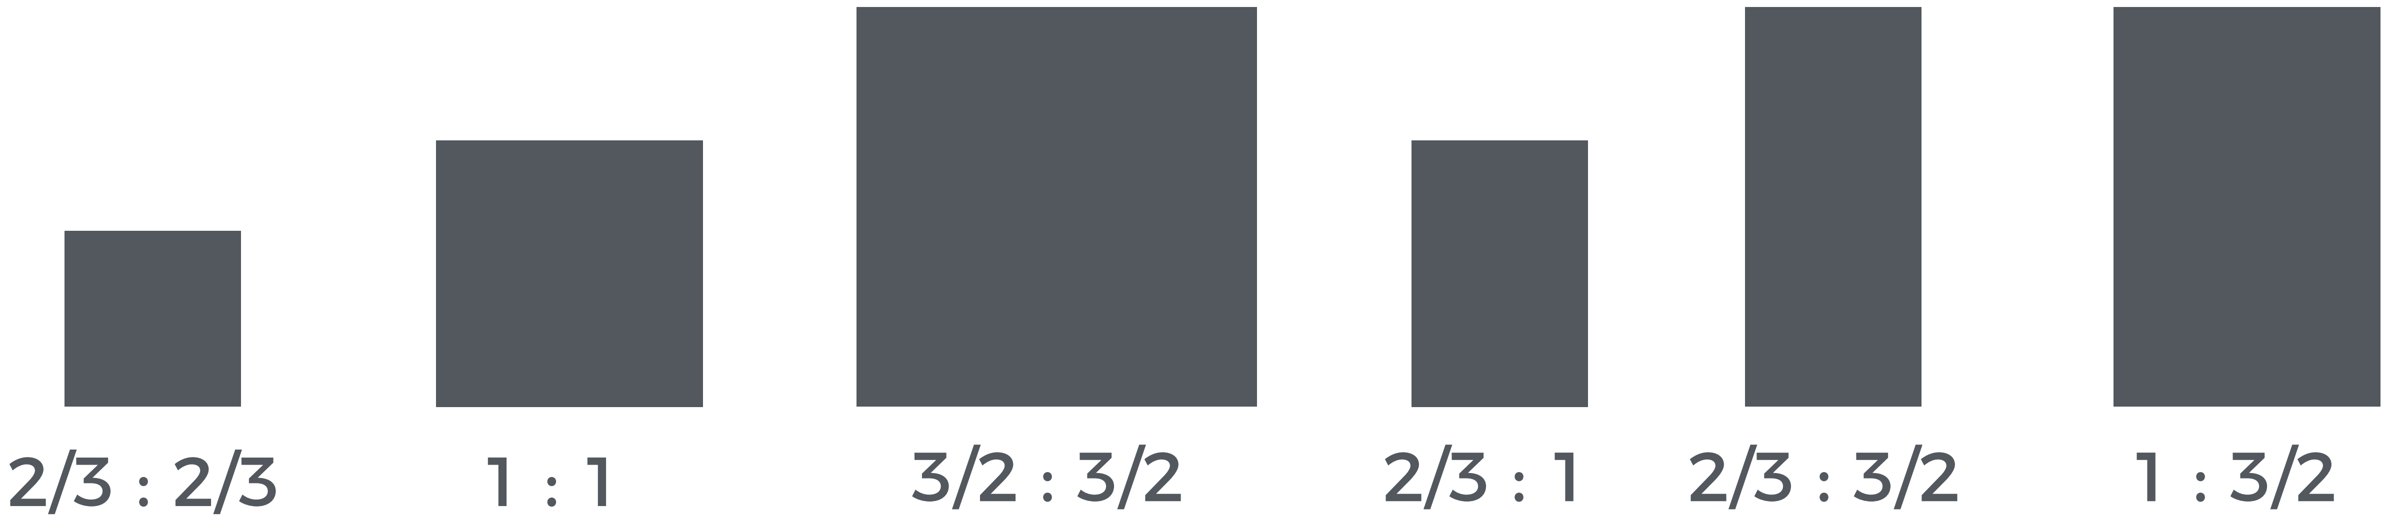

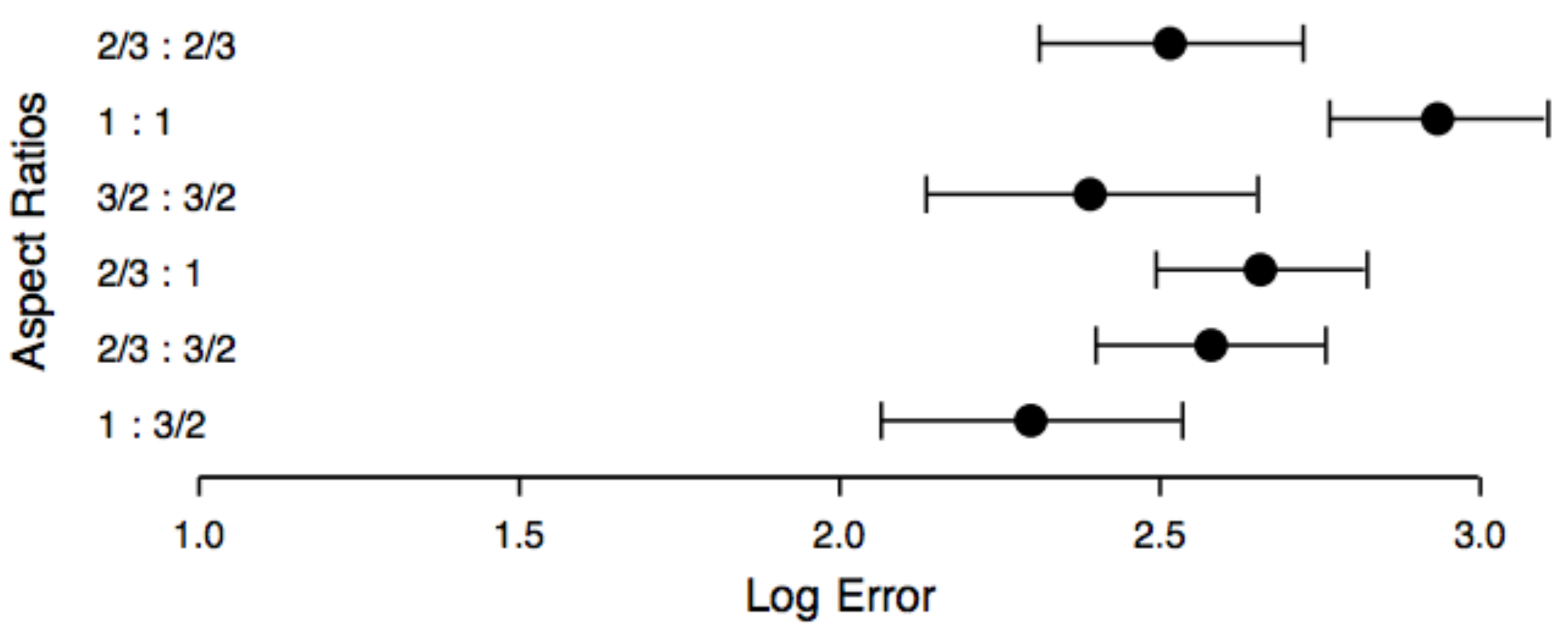

These results are directly relevant to treemap construction.

Aspect ratio is important!

For line charts there is a basic guideline on optimising plot aspect ratio to have an average angle of 45 degrees from Cleveland et al, 1988.

The Shape Parameter of a Two-Variable Graph
William Cleveland, Marylyn McGill, and Robert McGill
Journal of the American Statistical Association, 83, 289–300, 1988

Multi-Scale Banking to 45 Degrees
Jeffrey Heer, Maneesh Agrawala
IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis), 12(5), 701–708, 2006

An Empirical Model of Slope Ratio Comparisons
Justin Talbot, John Gerth, Pat Hanrahan
IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis), 2012

HOW

We have to be careful when mapping data to the visual world

Some visual channels are more effective for some data types over others.

Some data has a **natural mapping** that our brains expect given certain types of data

Natural Mappings

Graphical Code		Semantics
Small shapes defined by closed contour, texture, color.		Object, idea, entity, node.
Spatially ordered graphical objects.		Related information or a sequence. In a sequence the left-to-right ordering convention is borrowed from written language (English, French, etc.).
Graphical objects in proximity		Similar concepts
Graphical objects having the same shape color, or texture.	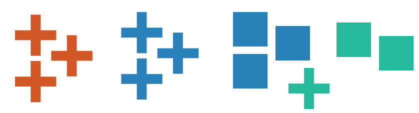	Similar concepts
Size, position or height of graphical object	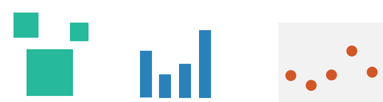	Size, quantity, importance, 2D location
Shapes connected by contour		Related entities, path between entities.
Thickness of connecting contour		Strength of relationship.
Color and texture of connecting contour	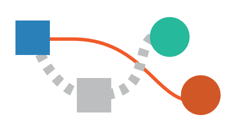	Type of relationship.
Shapes enclosed by a contour, a common texture or color	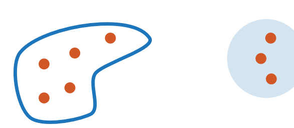	Contained/related entities.
Nested/partitioned regions		Hierarchical concepts.
Attached shapes	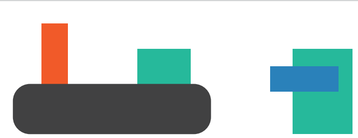	Parts of a conceptual structure.

HOW

We have to be careful when mapping data to the visual world

Some visual channels are more effective for some data types over others.

Some data has a **natural mapping** that our brains expect given certain types of data

There are many intricacies of the visual system that must be considered

The pop-out effect

We pre-attentively process a scene, and some visual elements stand out more than others.

- Parallel processing on many individual channels
  – speed independent of distractor count
  – speed depends on channel and amount of difference from distractors
- Serial search for (almost all) combinations
  – speed depends on number of distractors

The pop-out effect

We pre-attentively process a scene, and some visual elements stand out more than others.

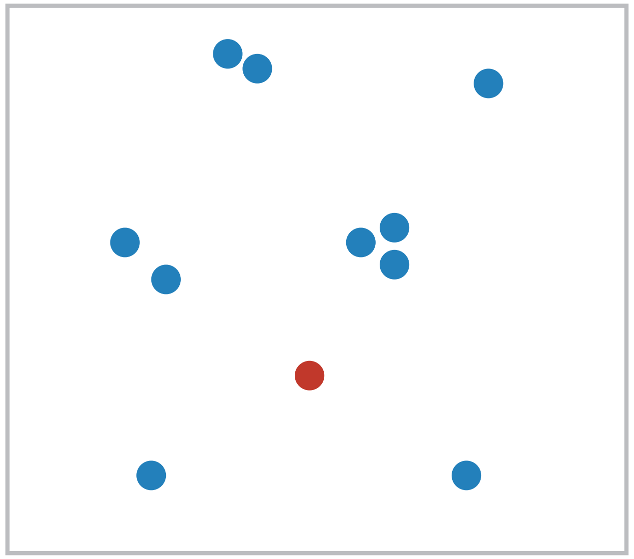

- Parallel processing on many individual channels
  - speed independent of distractor count
  - speed depends on channel and amount of difference from distractors
- Serial search for (almost all) combinations
  - speed depends on number of distractors

The pop-out effect

We pre-attentively process a scene, and some visual elements stand out more than others.

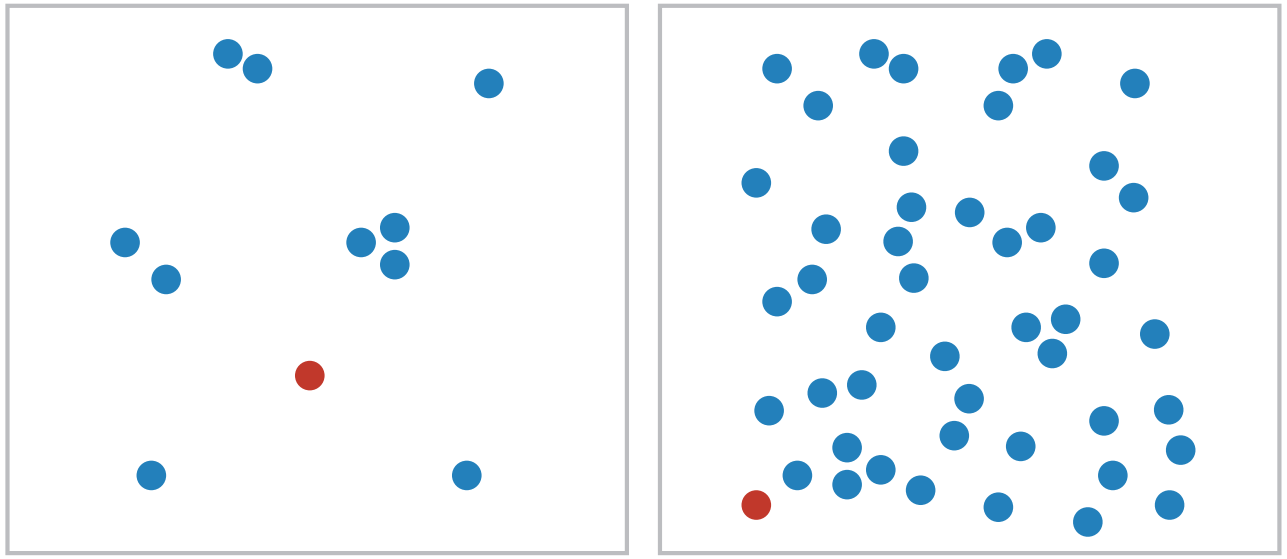

- Parallel processing on many individual channels
  – speed independent of distractor count
  – speed depends on channel and amount of difference from distractors
- Serial search for (almost all) combinations
  – speed depends on number of distractors

The pop-out effect

We pre-attentively process a scene, and some visual elements stand out more than others.

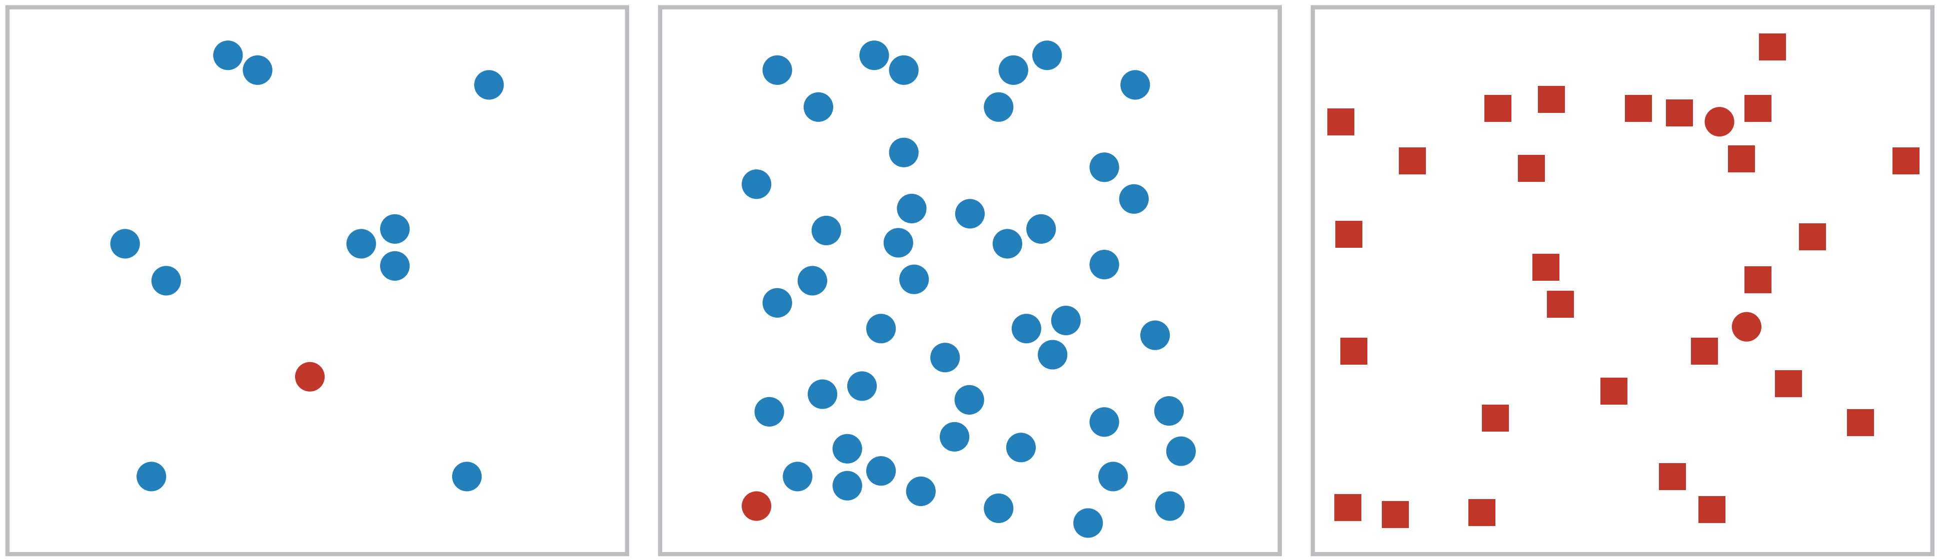

- Parallel processing on many individual channels
  - speed independent of distractor count
  - speed depends on channel and amount of difference from distractors
- Serial search for (almost all) combinations
  - speed depends on number of distractors

The pop-out effect

We pre-attentively process a scene, and some visual elements stand out more than others.

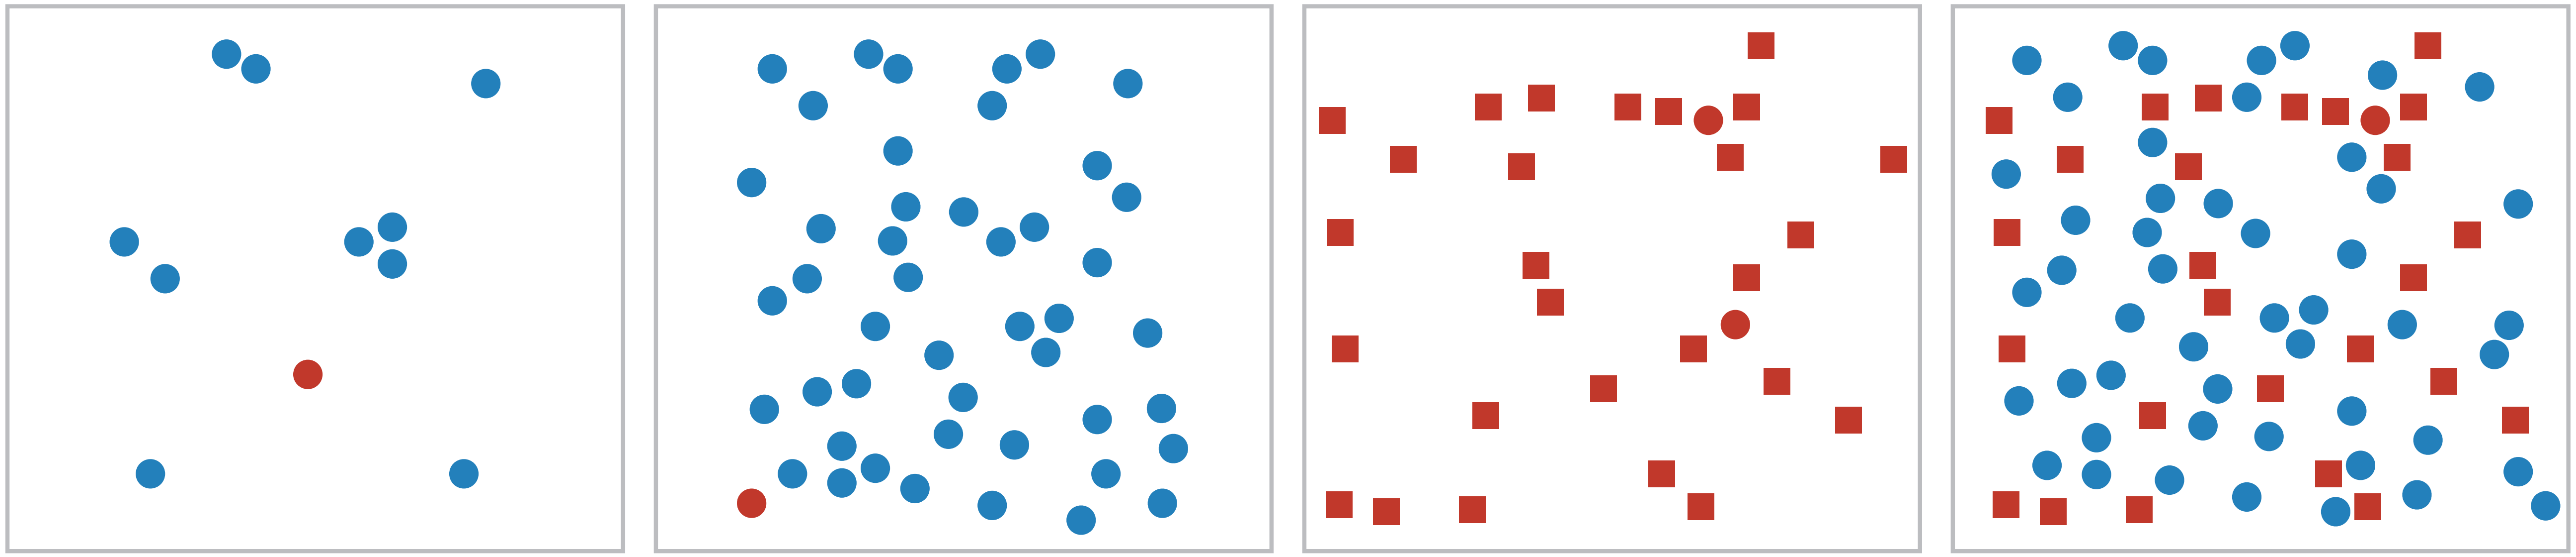

- Parallel processing on many individual channels
  - speed independent of distractor count
  - speed depends on channel and amount of difference from distractors
- Serial search for (almost all) combinations
  - speed depends on number of distractors

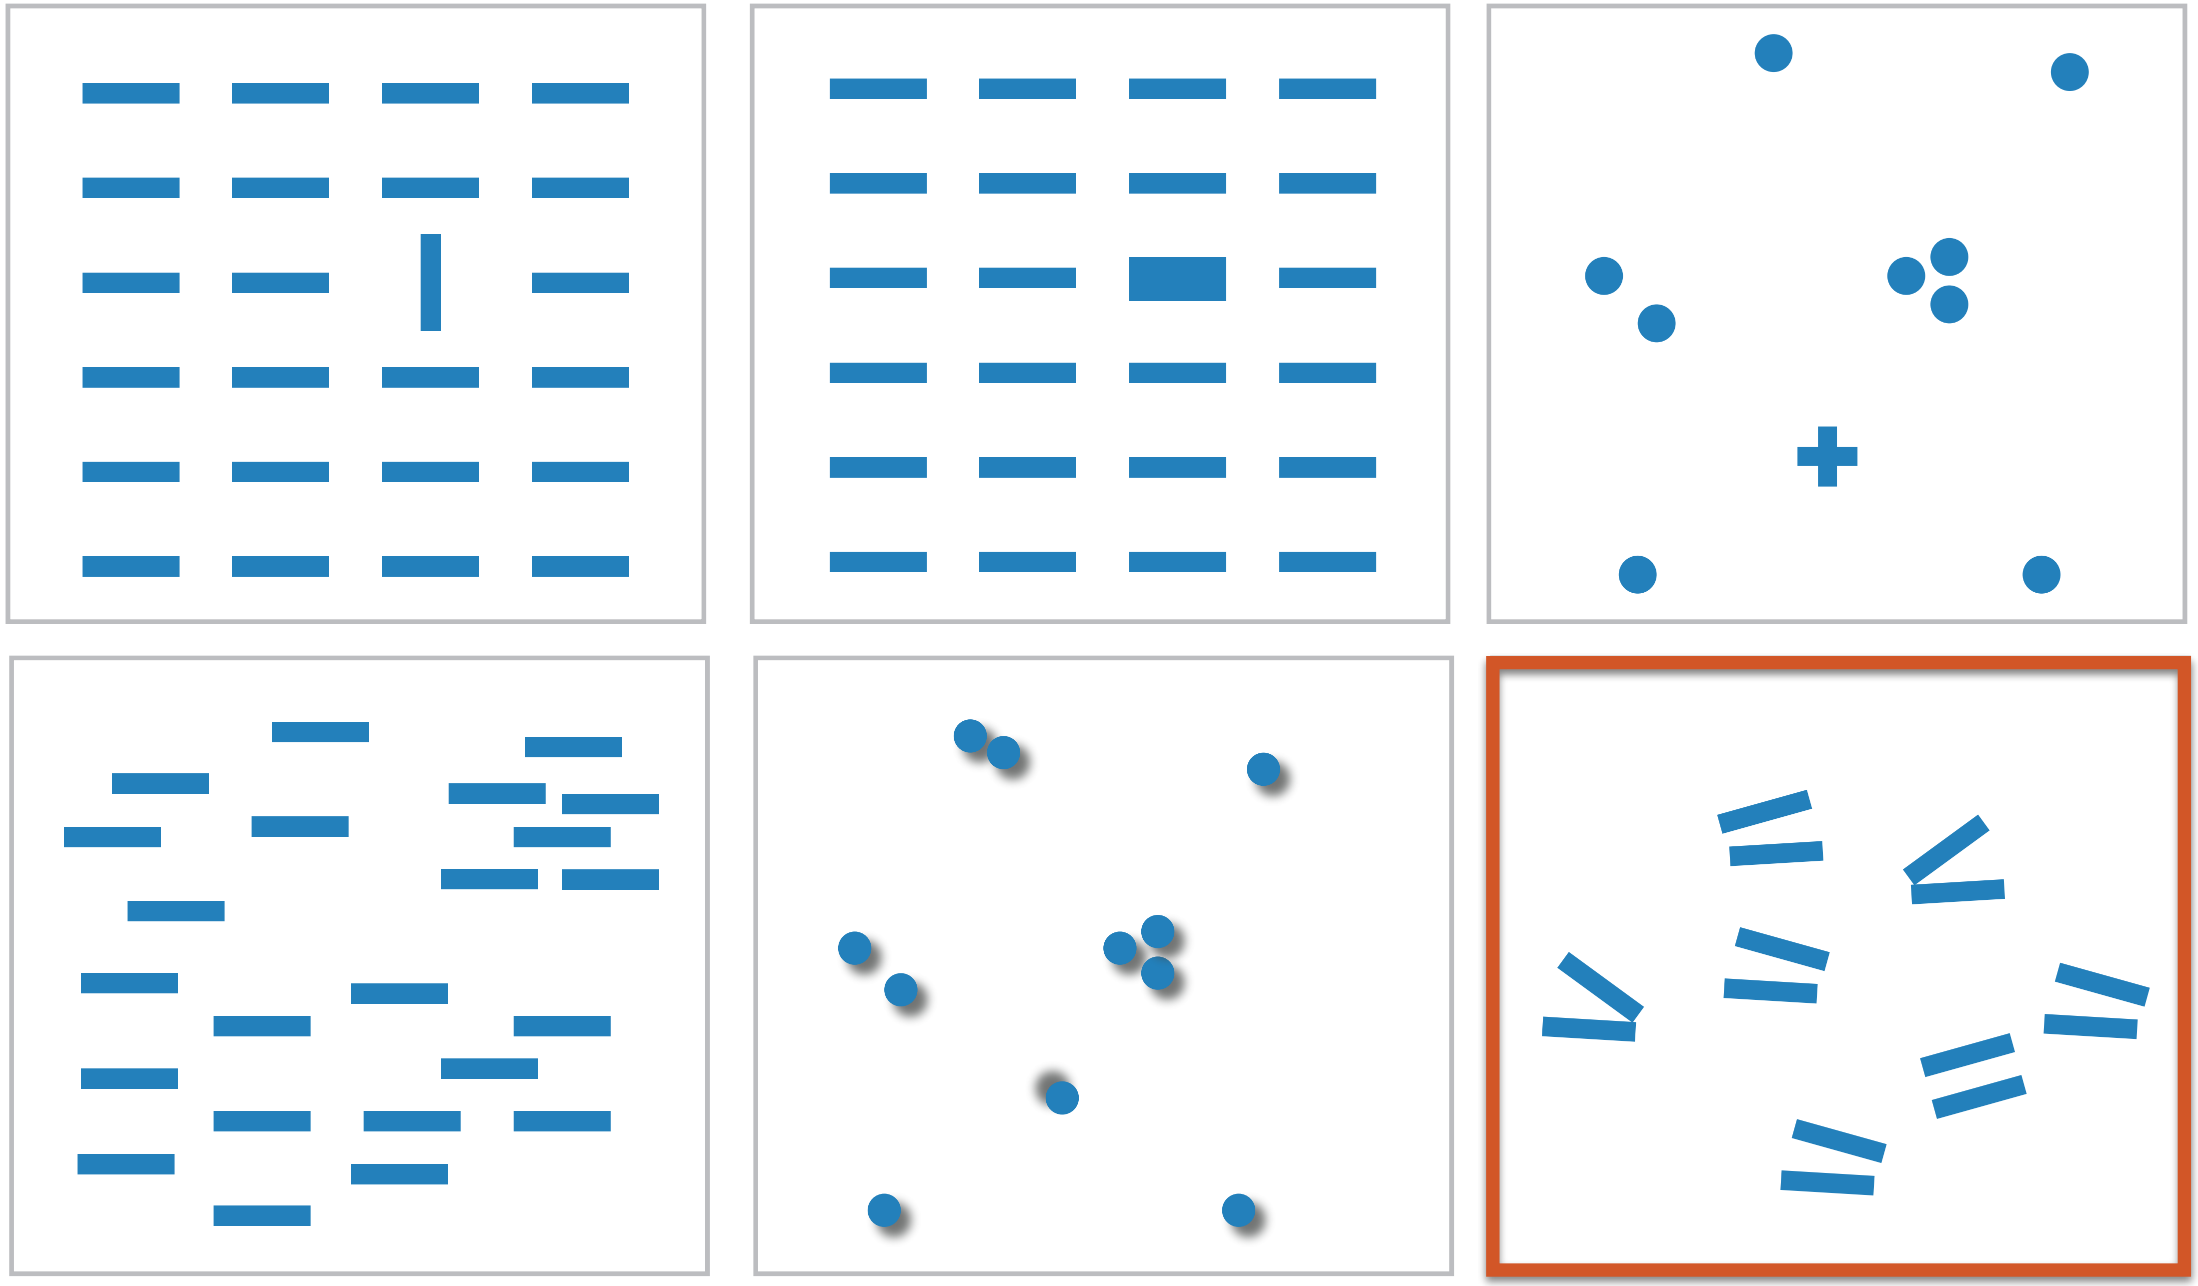

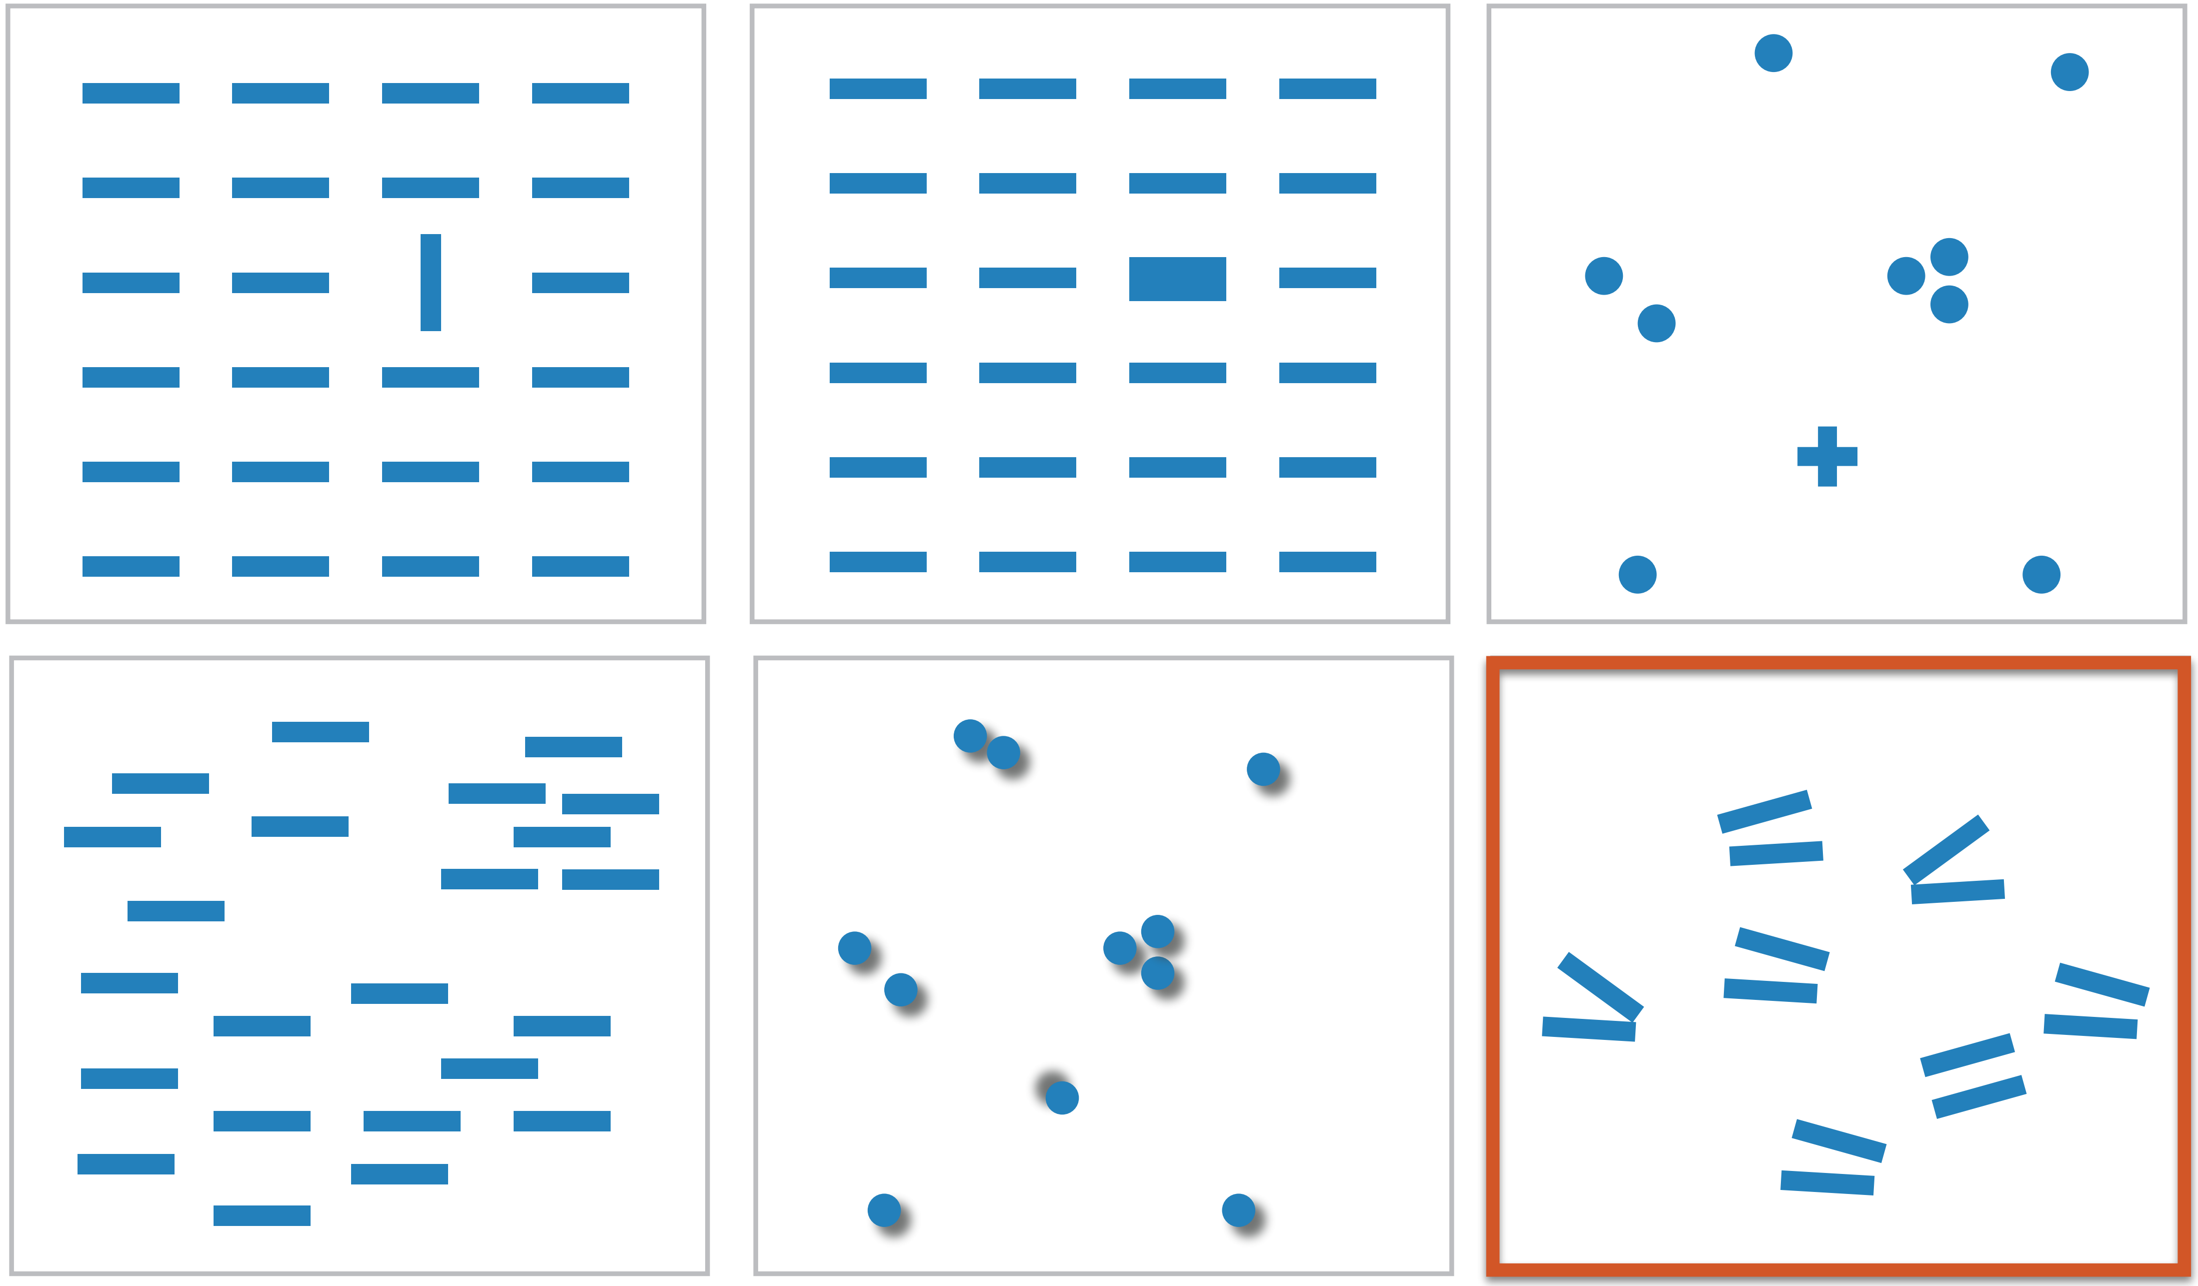

Not all exhibit the pop-out effect!

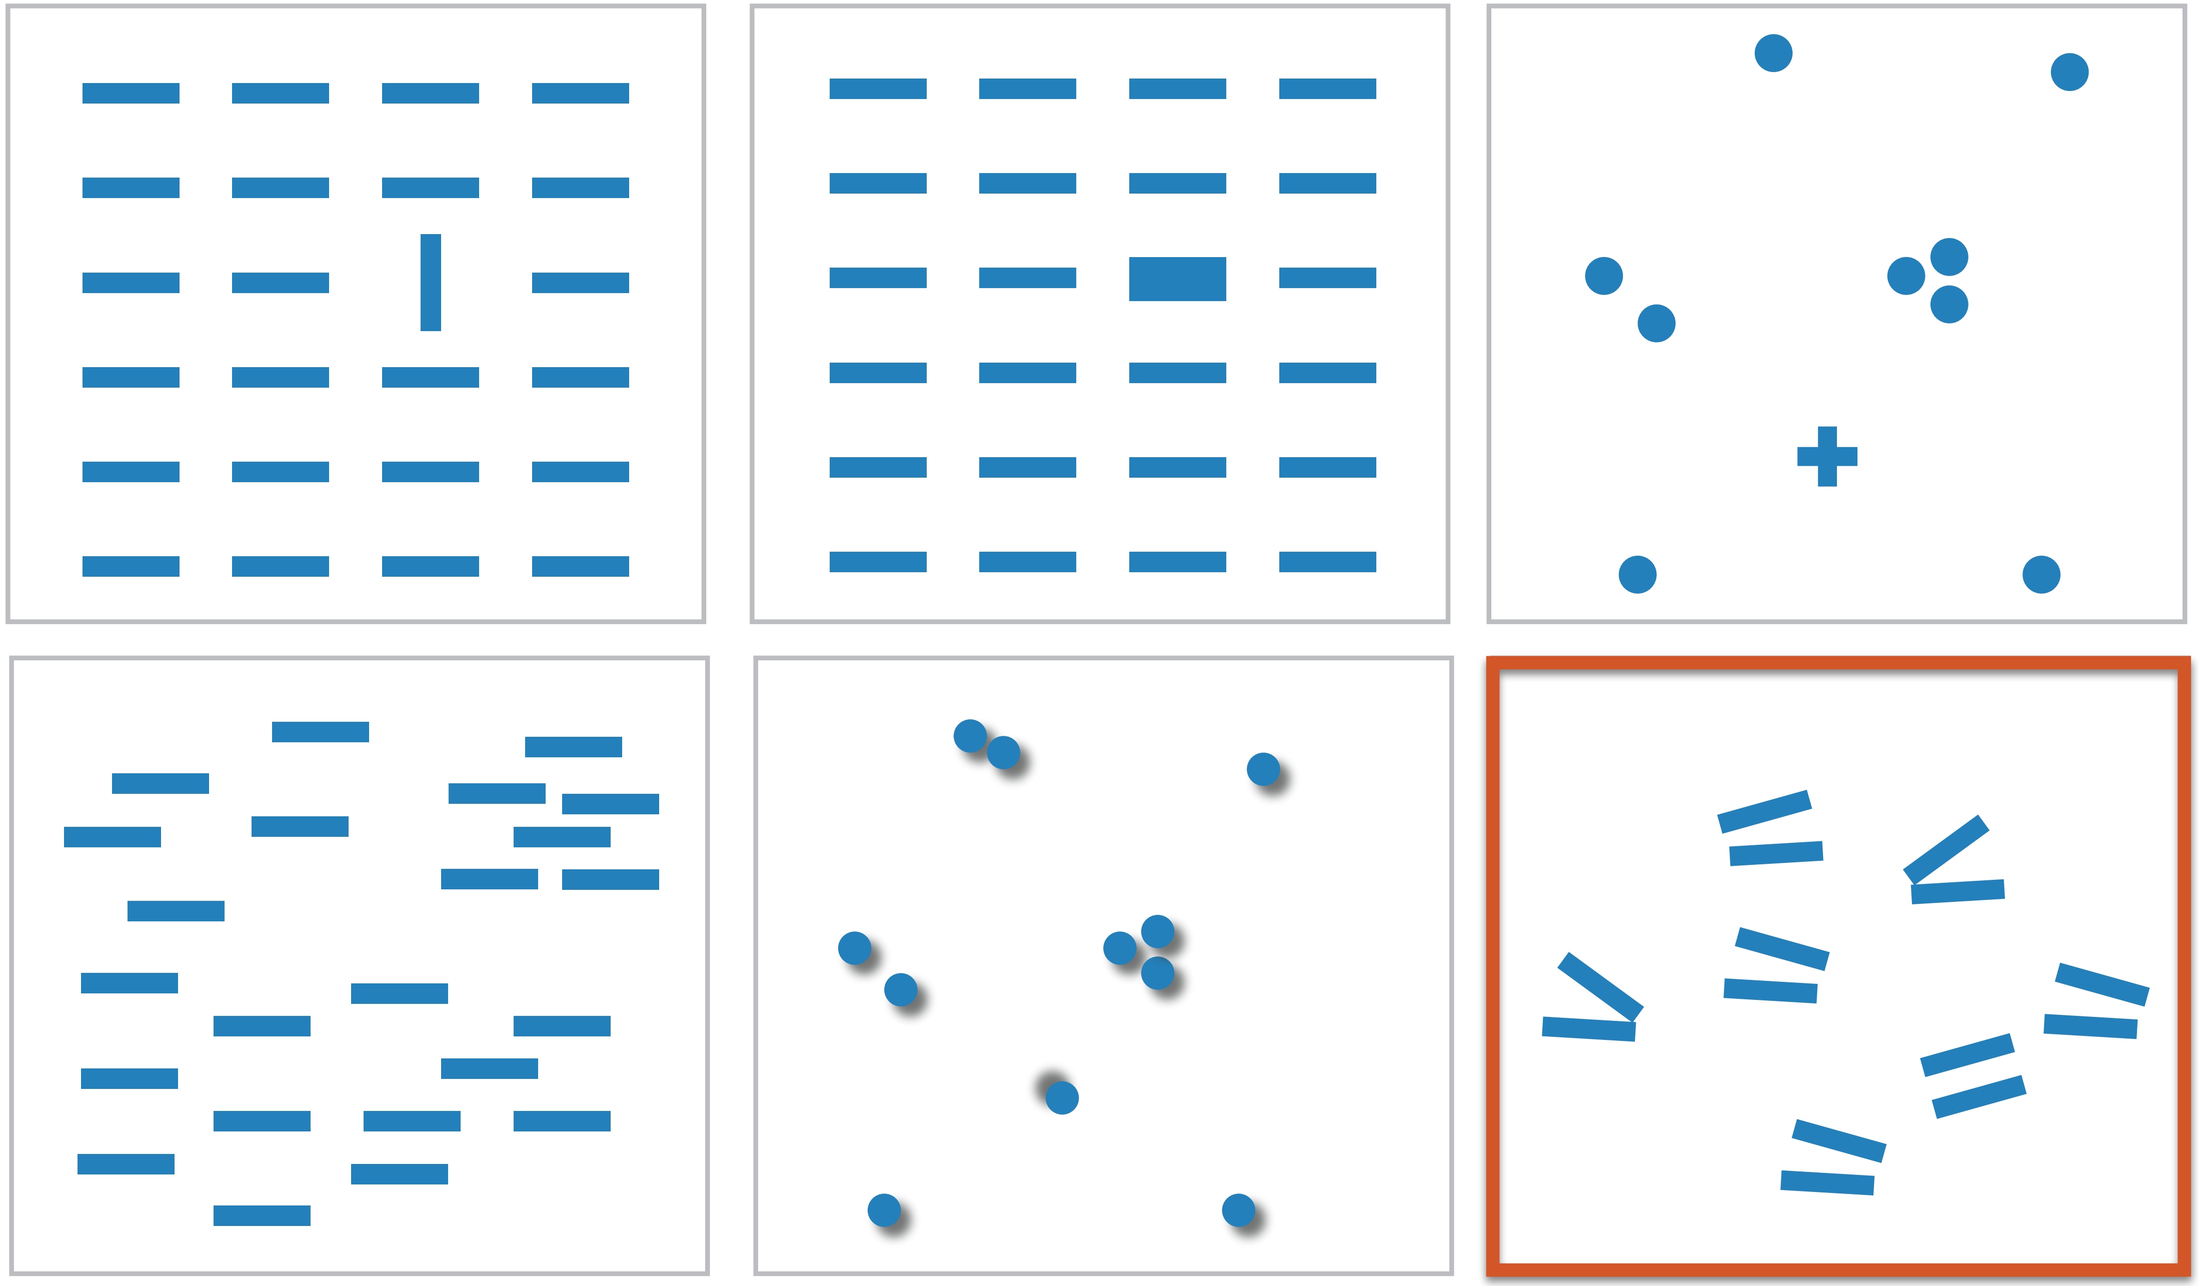

Not all exhibit the pop-out effect!

Parallel line pairs do not pop out from tilted pairs...

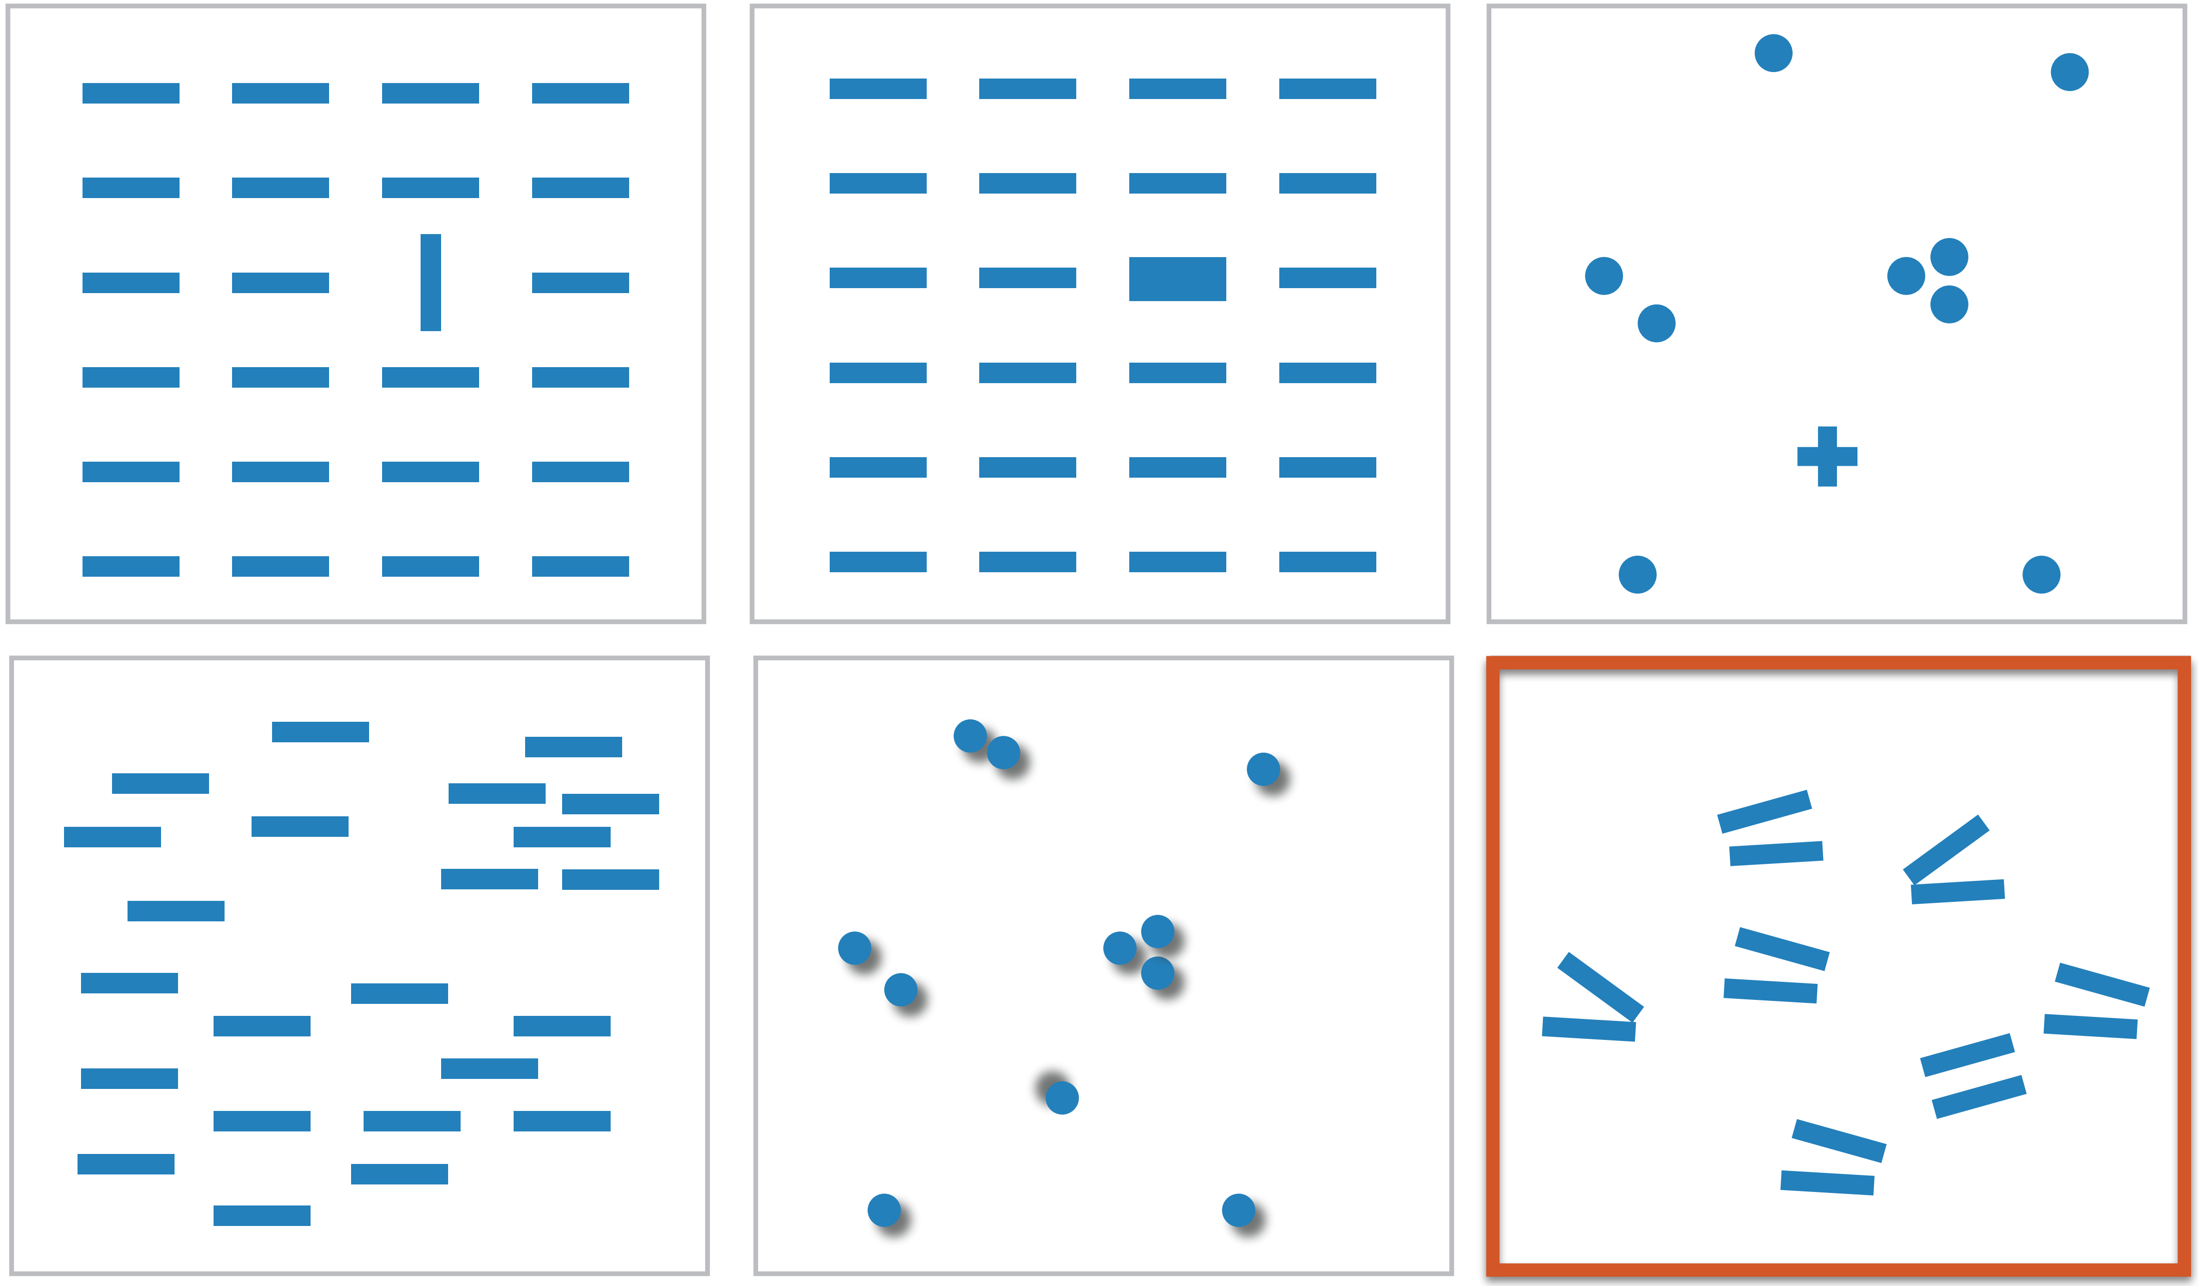

Not all exhibit the pop-out effect!

Parallel line pairs do not pop out from tilted pairs…

And not all visual channels pop out as quickly as other. E.g. colour is always on top.

Relative Comparison

Relative Comparison

Relative Comparison

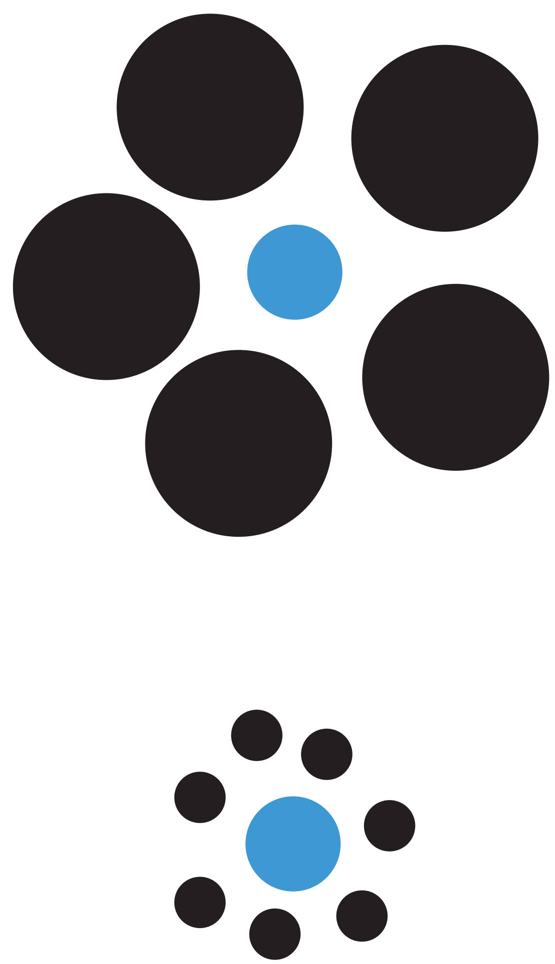

Relative Comparison

36px

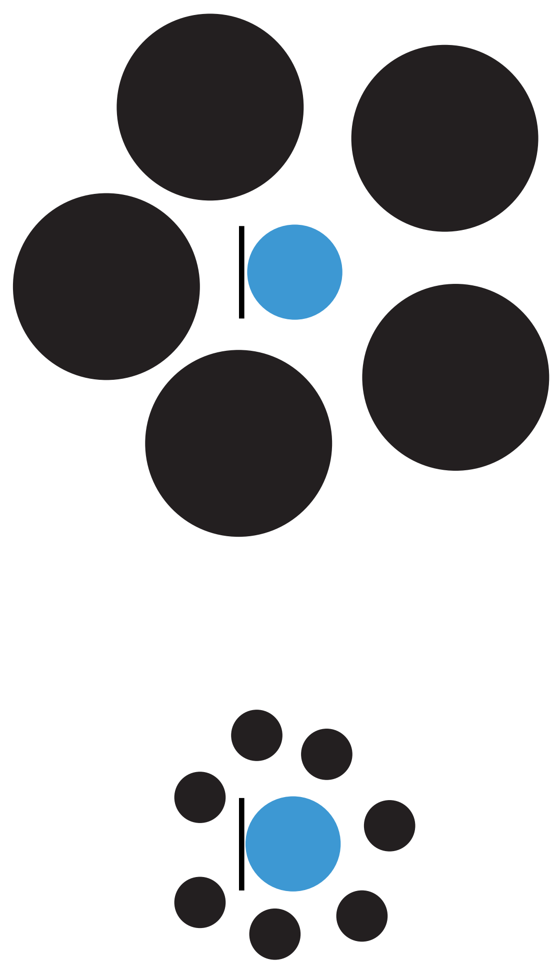

Relative Comparison

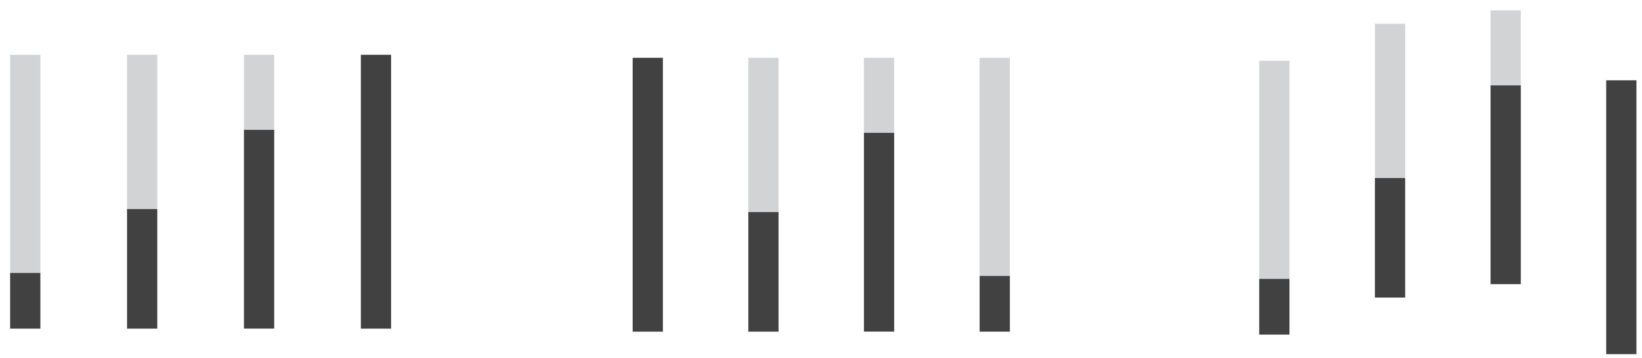

4 values Unordered Unaligned

Relative Comparison

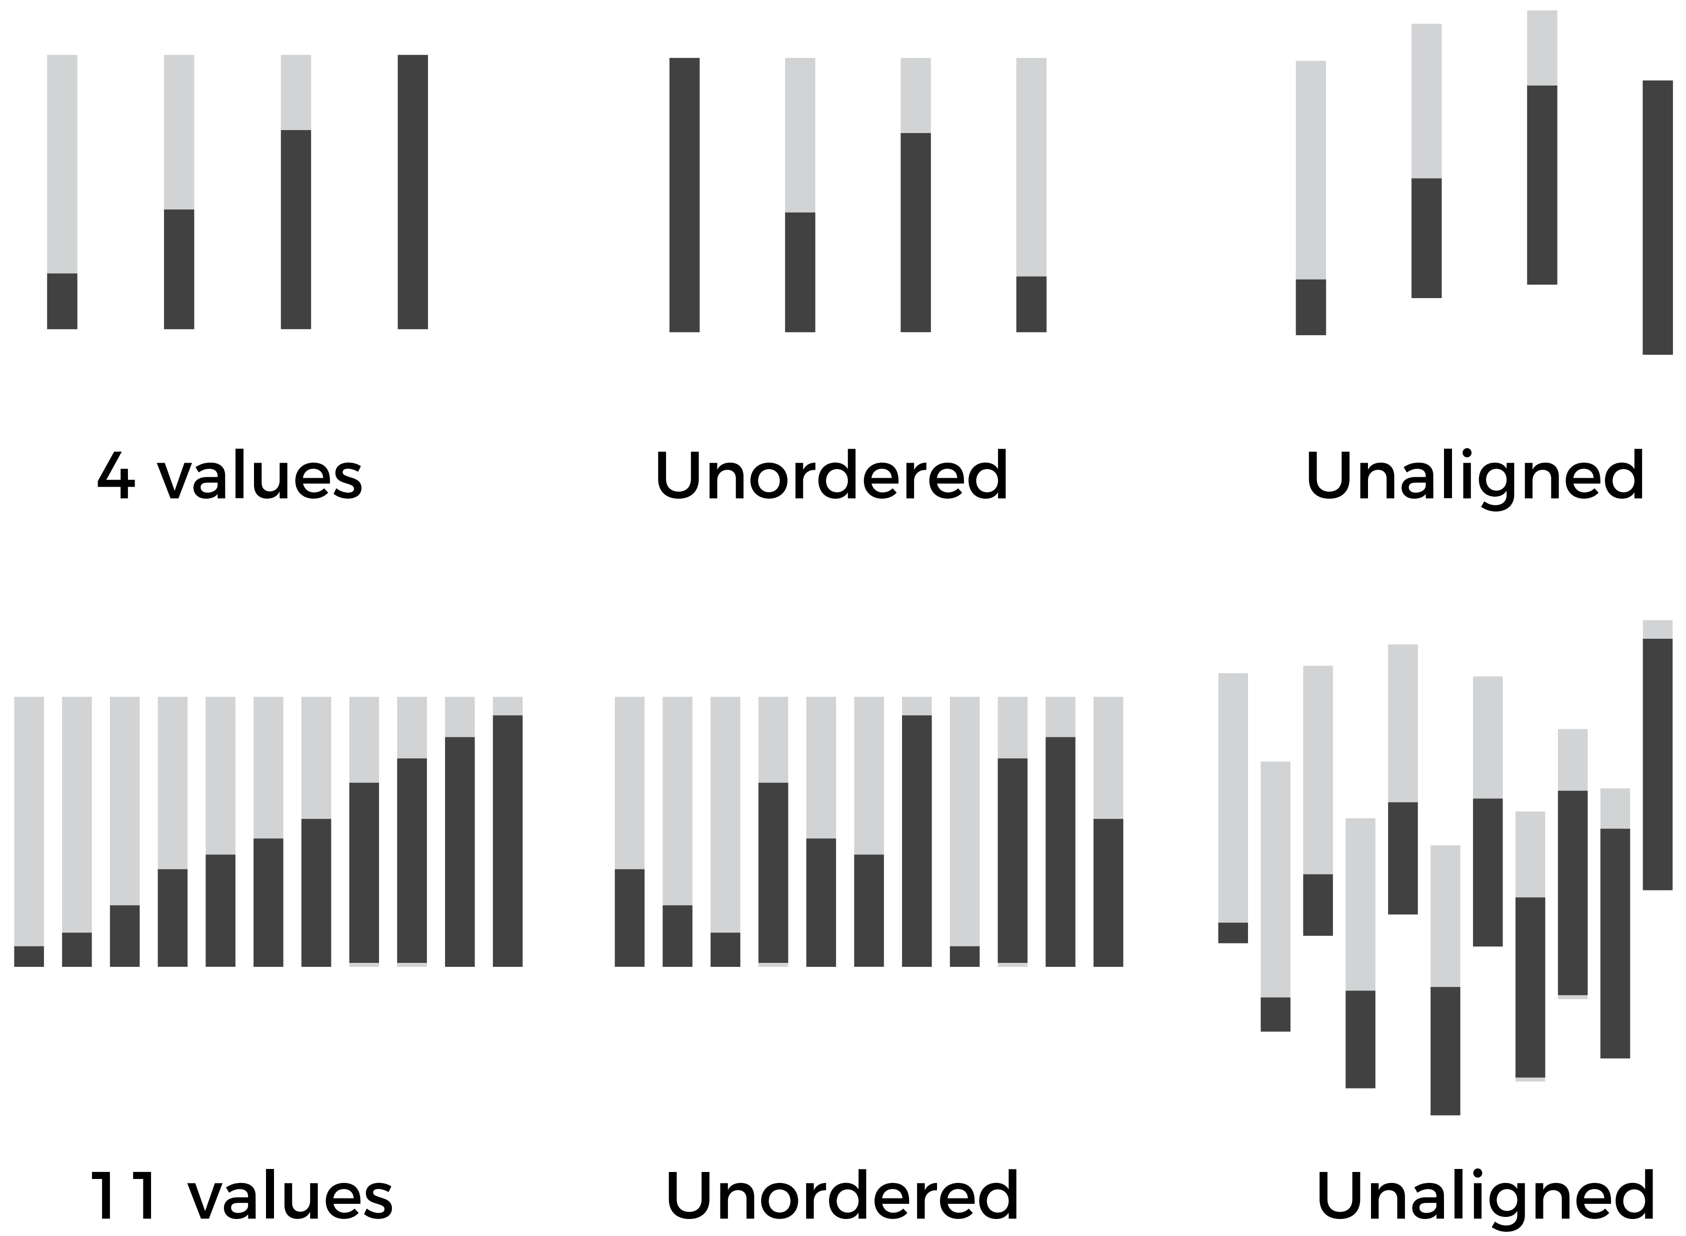

Relative Comparison

4 values

Aligned **Unordered**

Relative Comparison

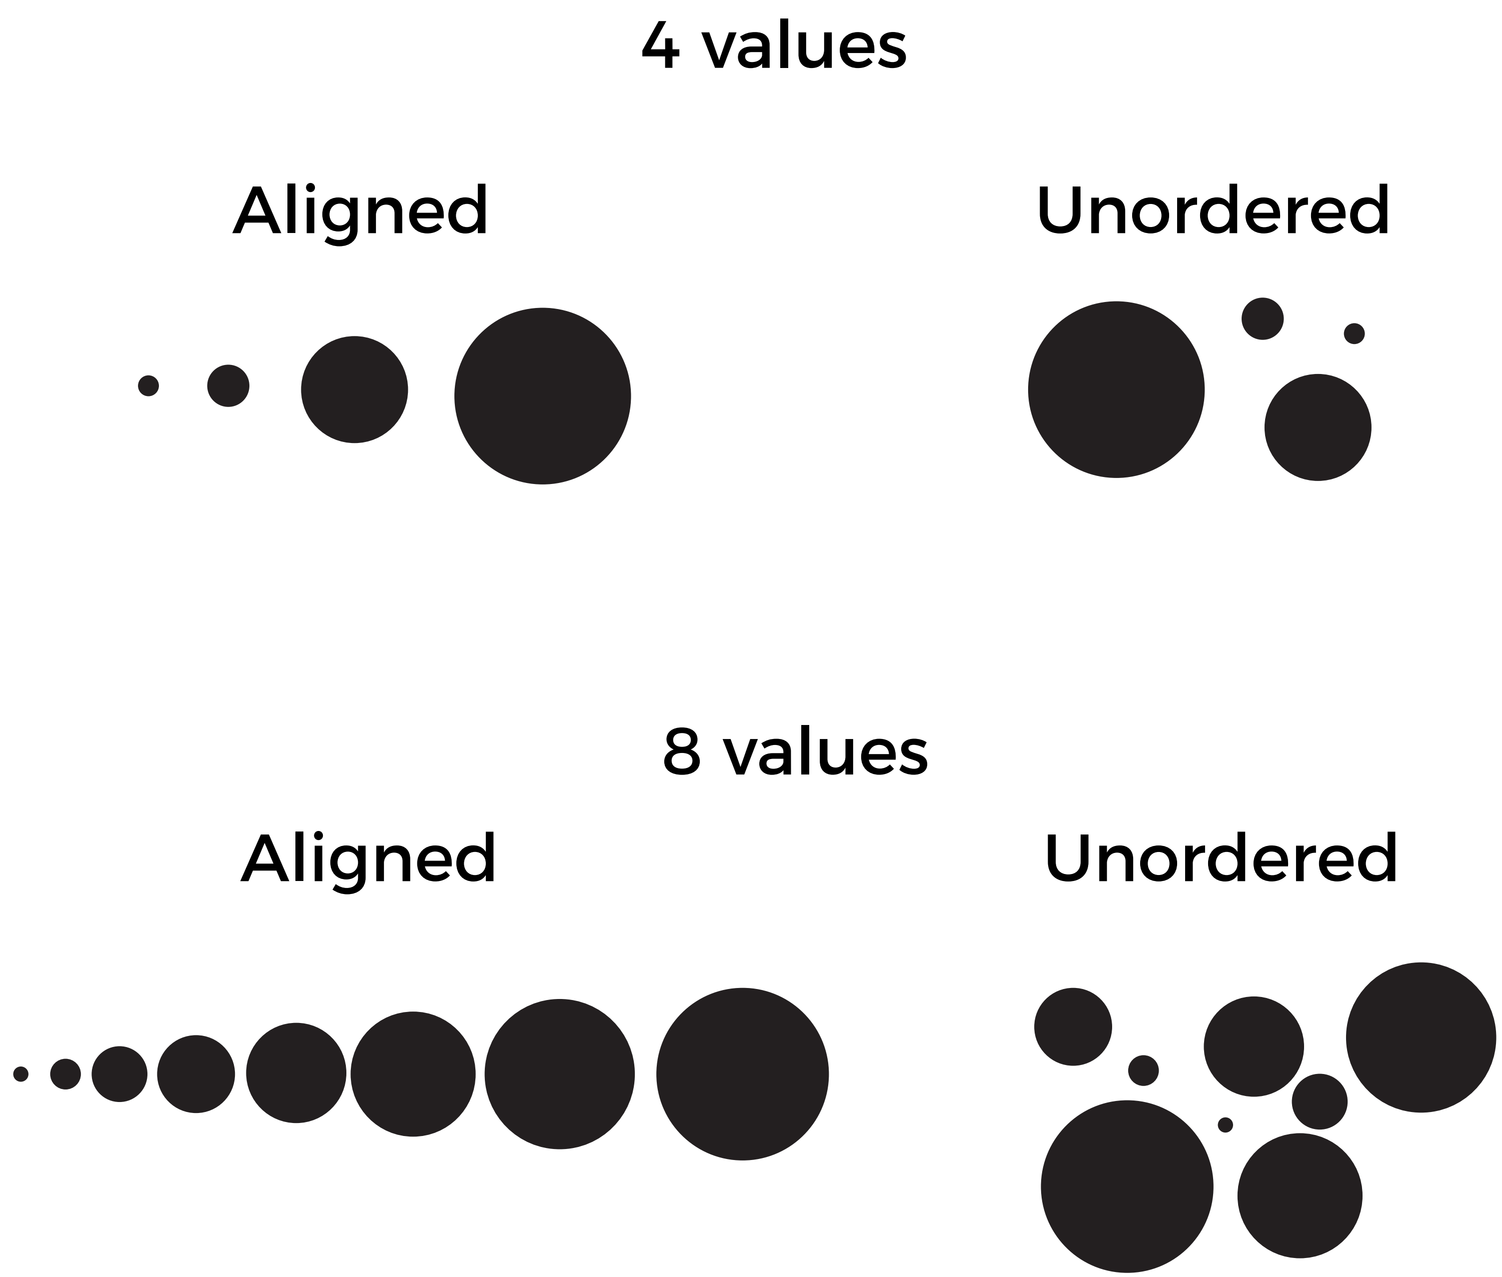

Relative Comparison

8 values

20 values

A) Known and Unknown Target Search

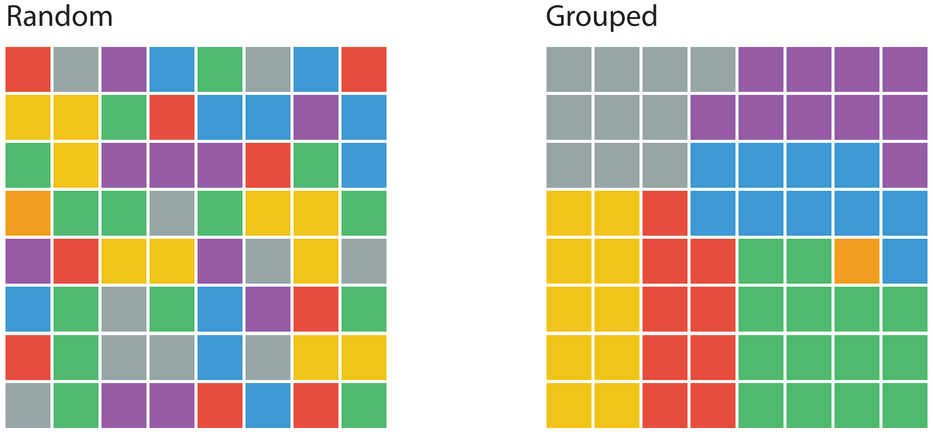

Target shown before hand (**known**) or **not shown** (**unknown**).
The unique colour here is the orange square.

B) Subitizing (how many colours?)

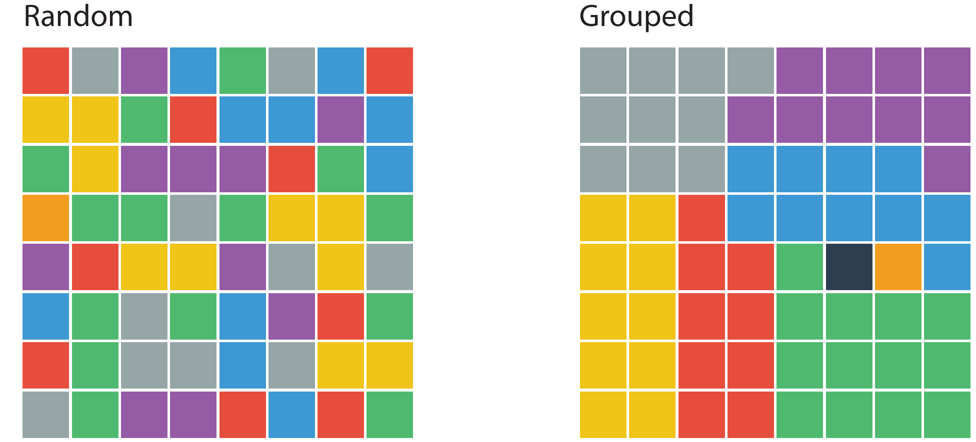

Which grid has more colours?

C) Response Time and Accuracy Results

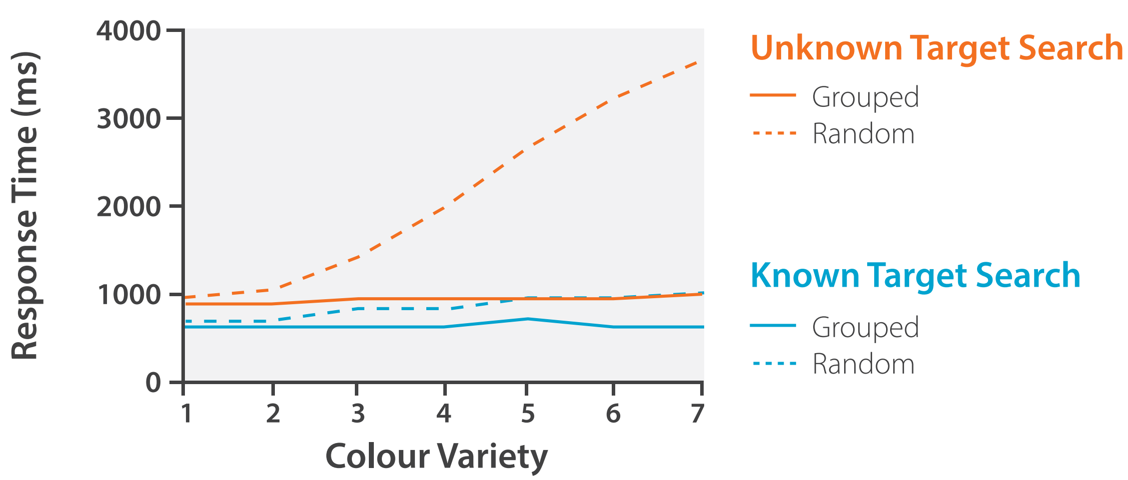

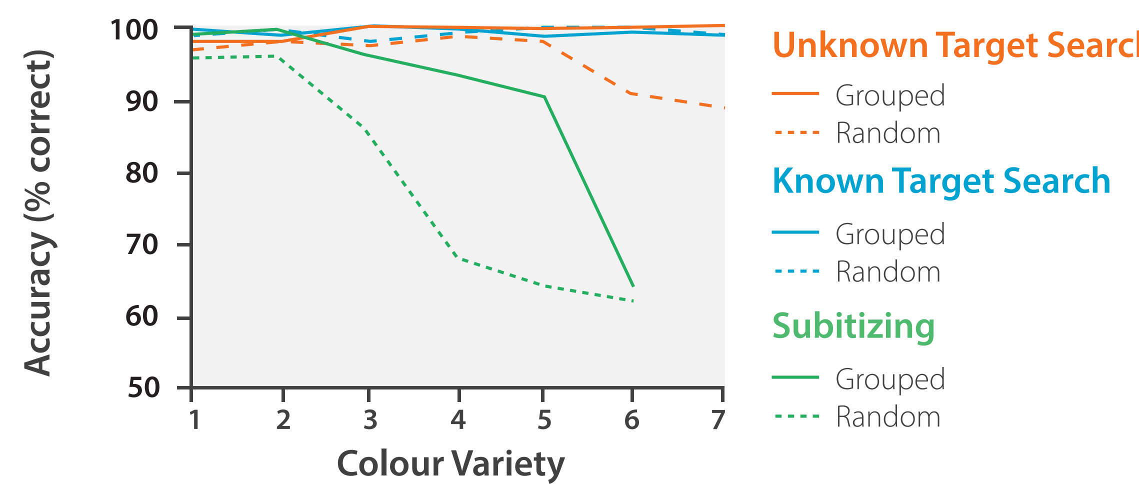

How Capacity Limits of Attention Influence Information Visualization Effectiveness.
Haroz S. and Whitney D., IEEE TVCG 2012

A) Known and Unknown Target Search

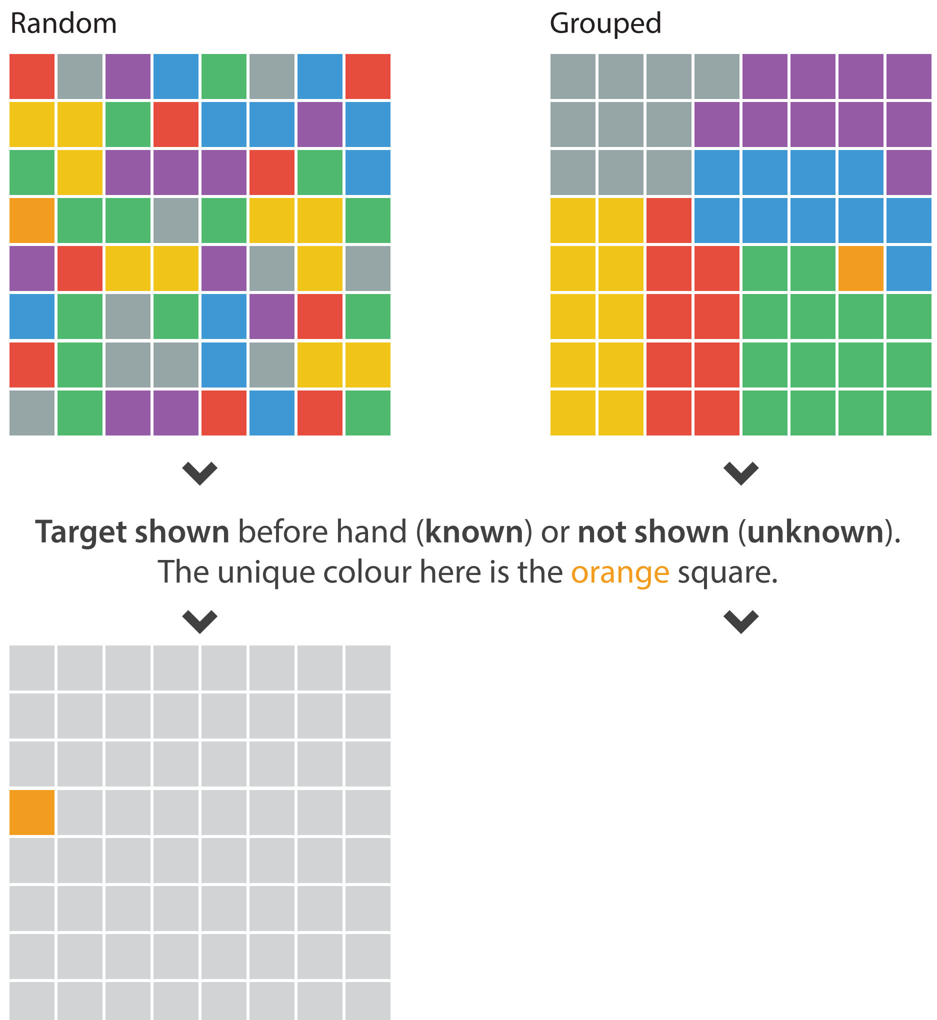

B) Subitizing (how many colours?)

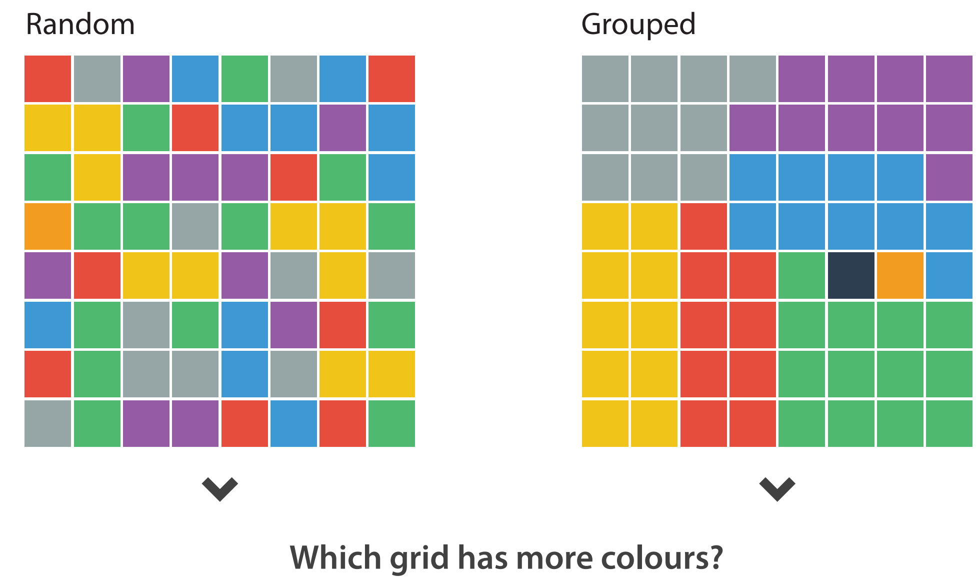

C) Response Time and Accuracy Results

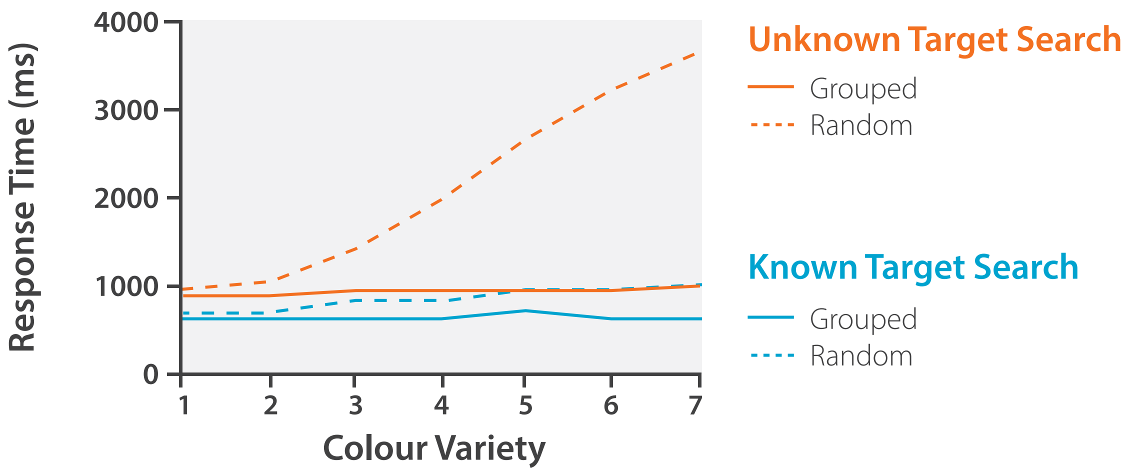

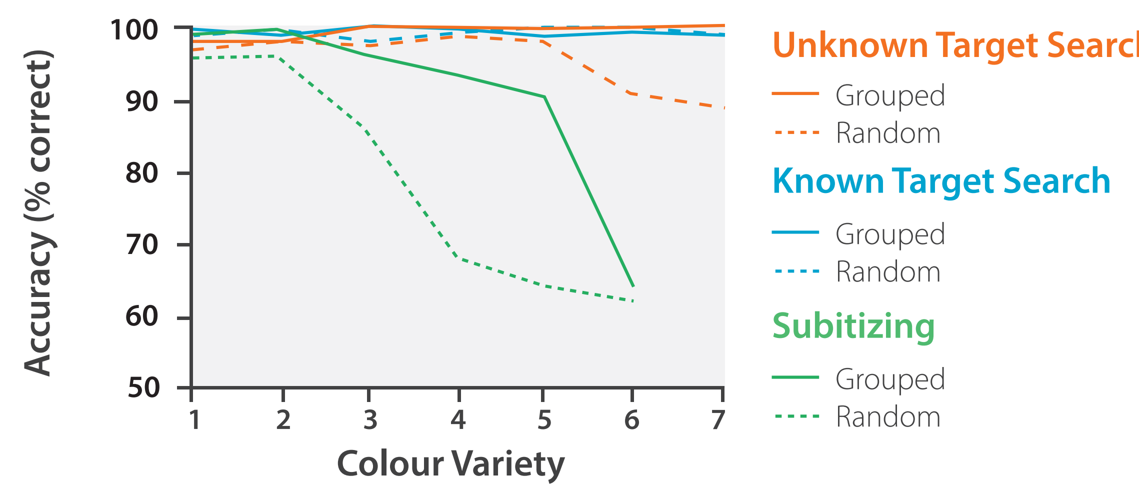

How Capacity Limits of Attention Influence Information Visualization Effectiveness.
Haroz S. and Whitney D., IEEE TVCG 2012

A) Known and Unknown Target Search

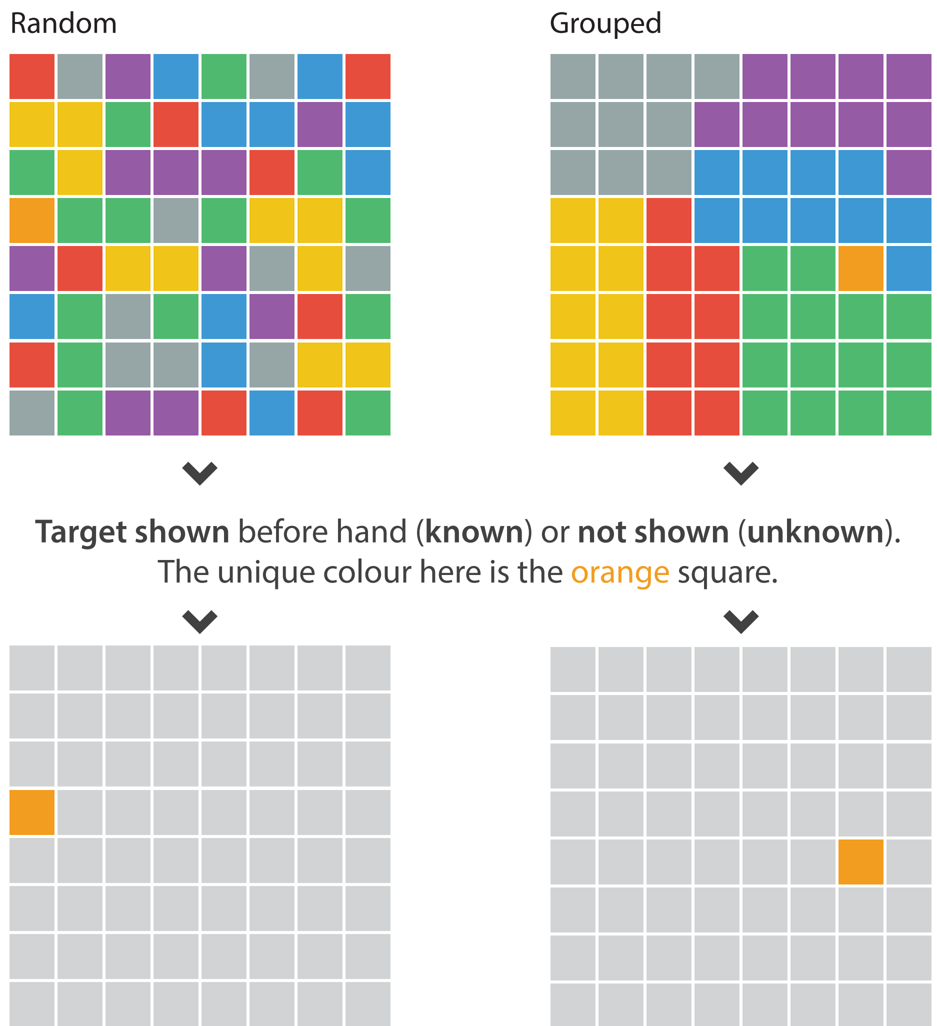

B) Subitizing (how many colours?)

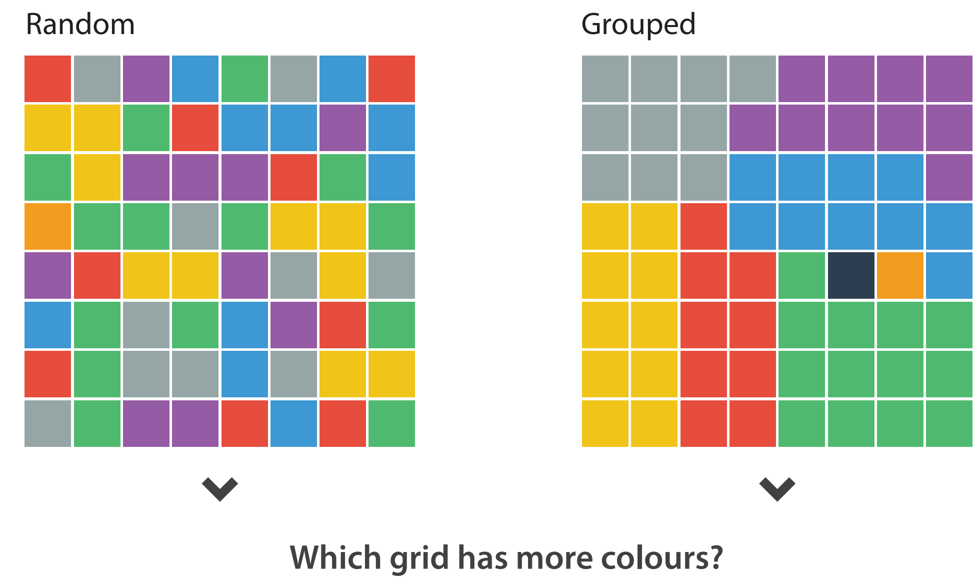

C) Response Time and Accuracy Results

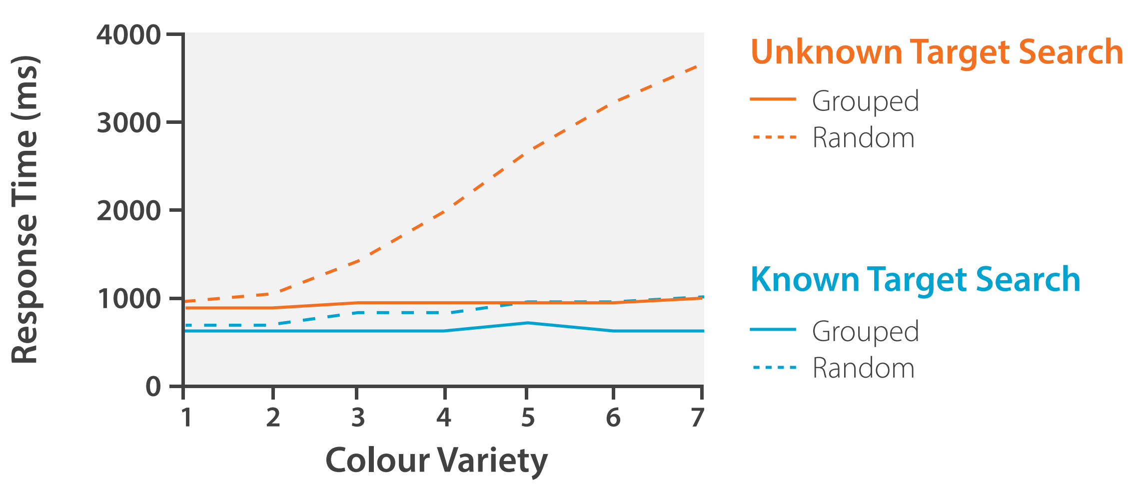

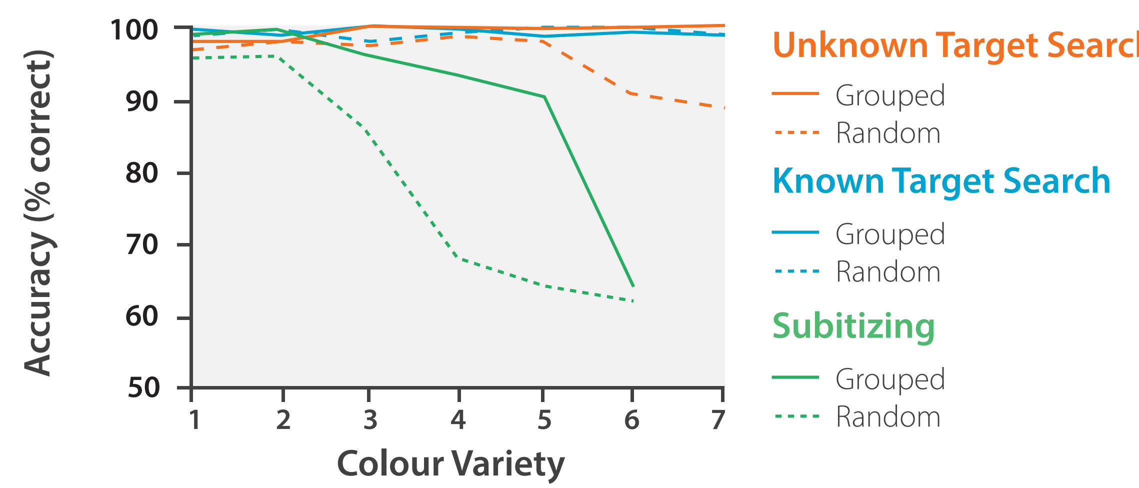

How Capacity Limits of Attention Influence Information Visualization Effectiveness.
Haroz S. and Whitney D., IEEE TVCG 2012

A) Known and Unknown Target Search

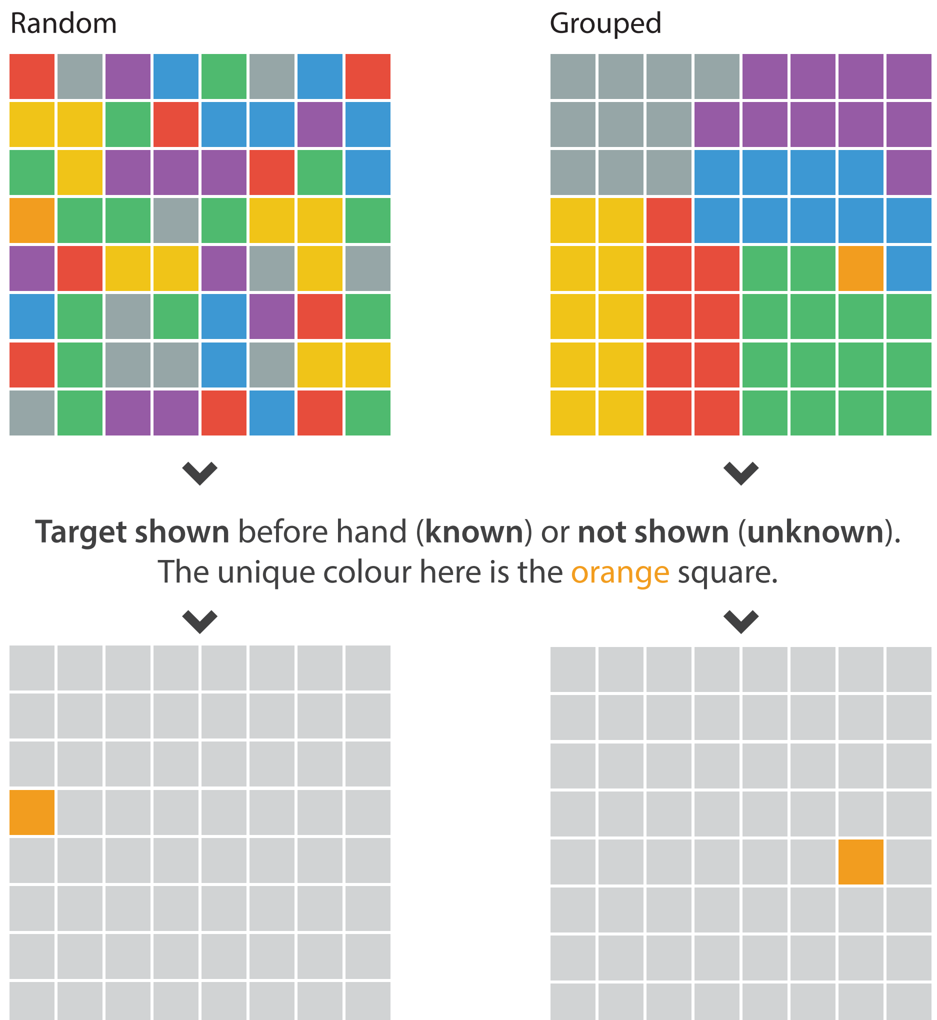

B) Subitizing (how many colours?)

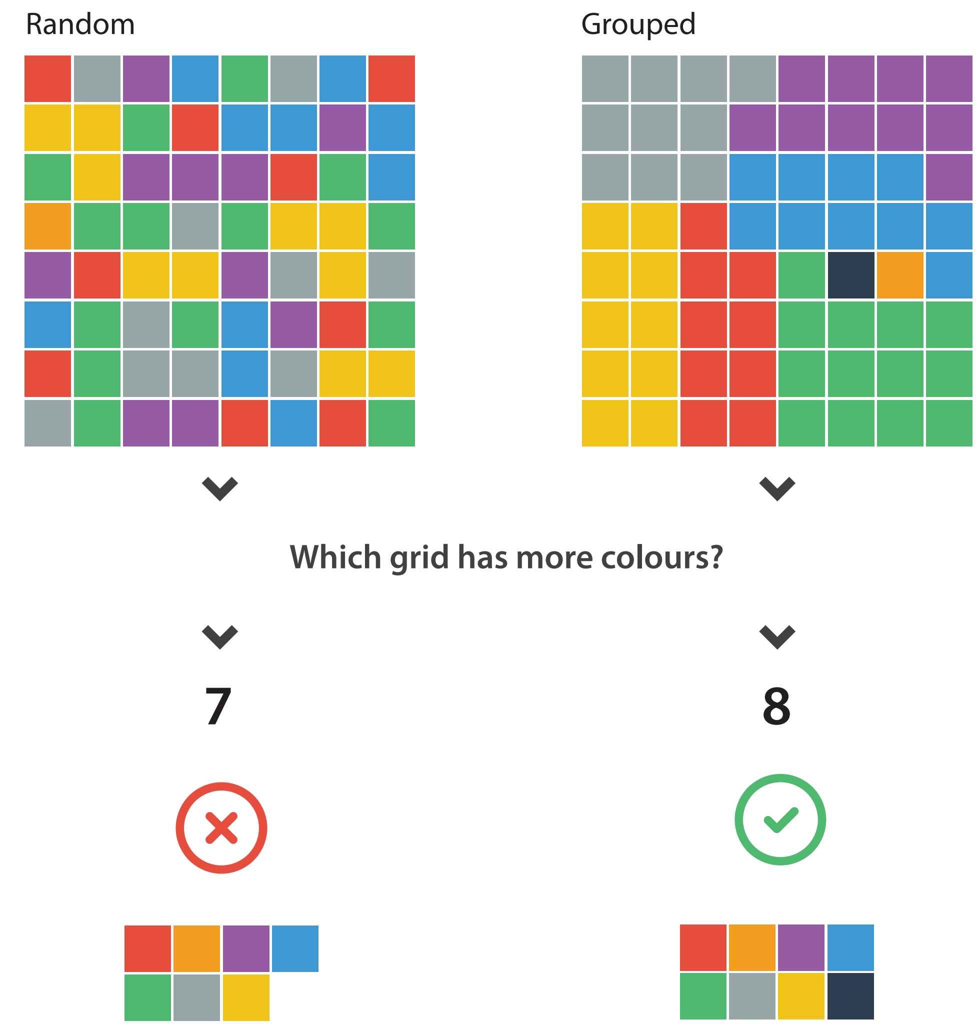

C) Response Time and Accuracy Results

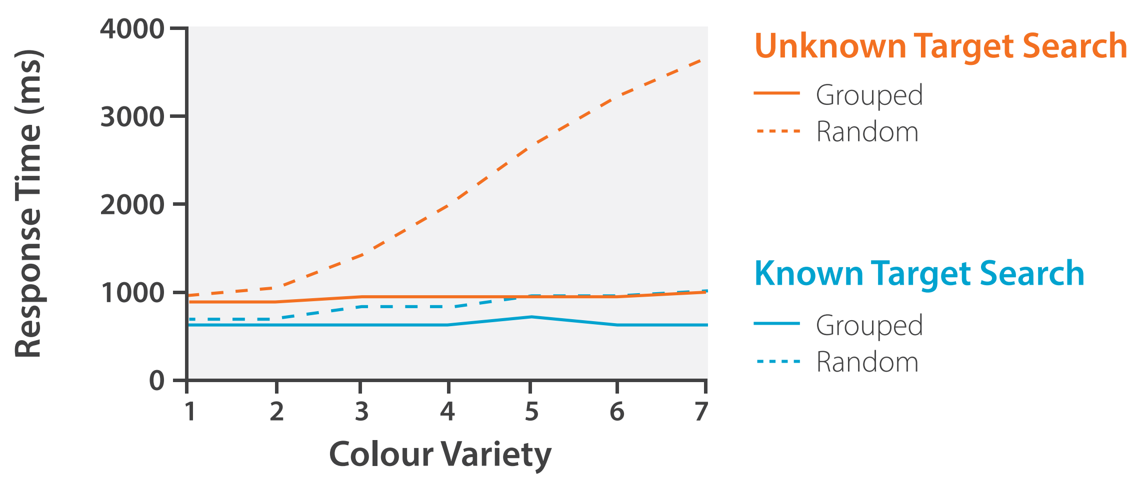

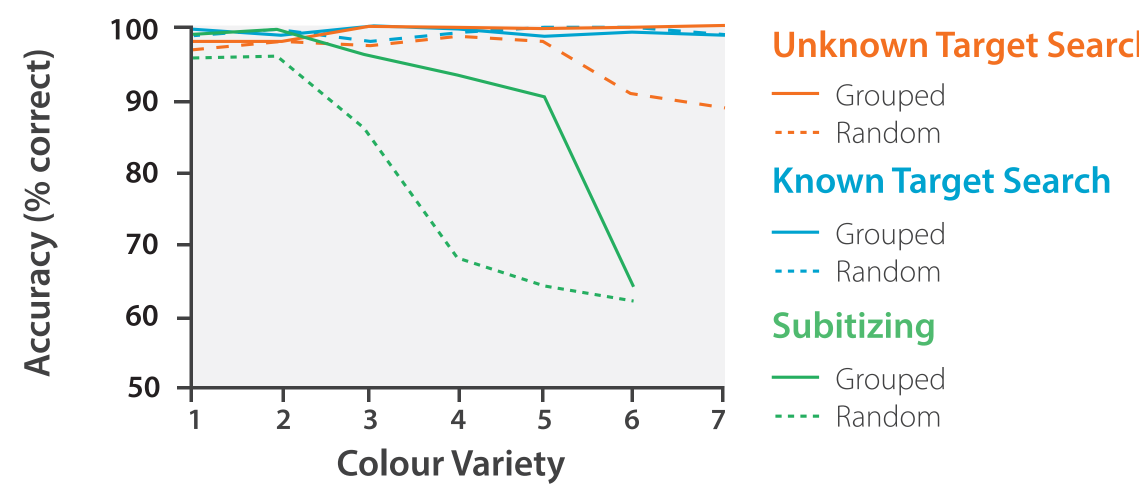

How Capacity Limits of Attention Influence Information Visualization Effectiveness.
Haroz S. and Whitney D., IEEE TVCG 2012

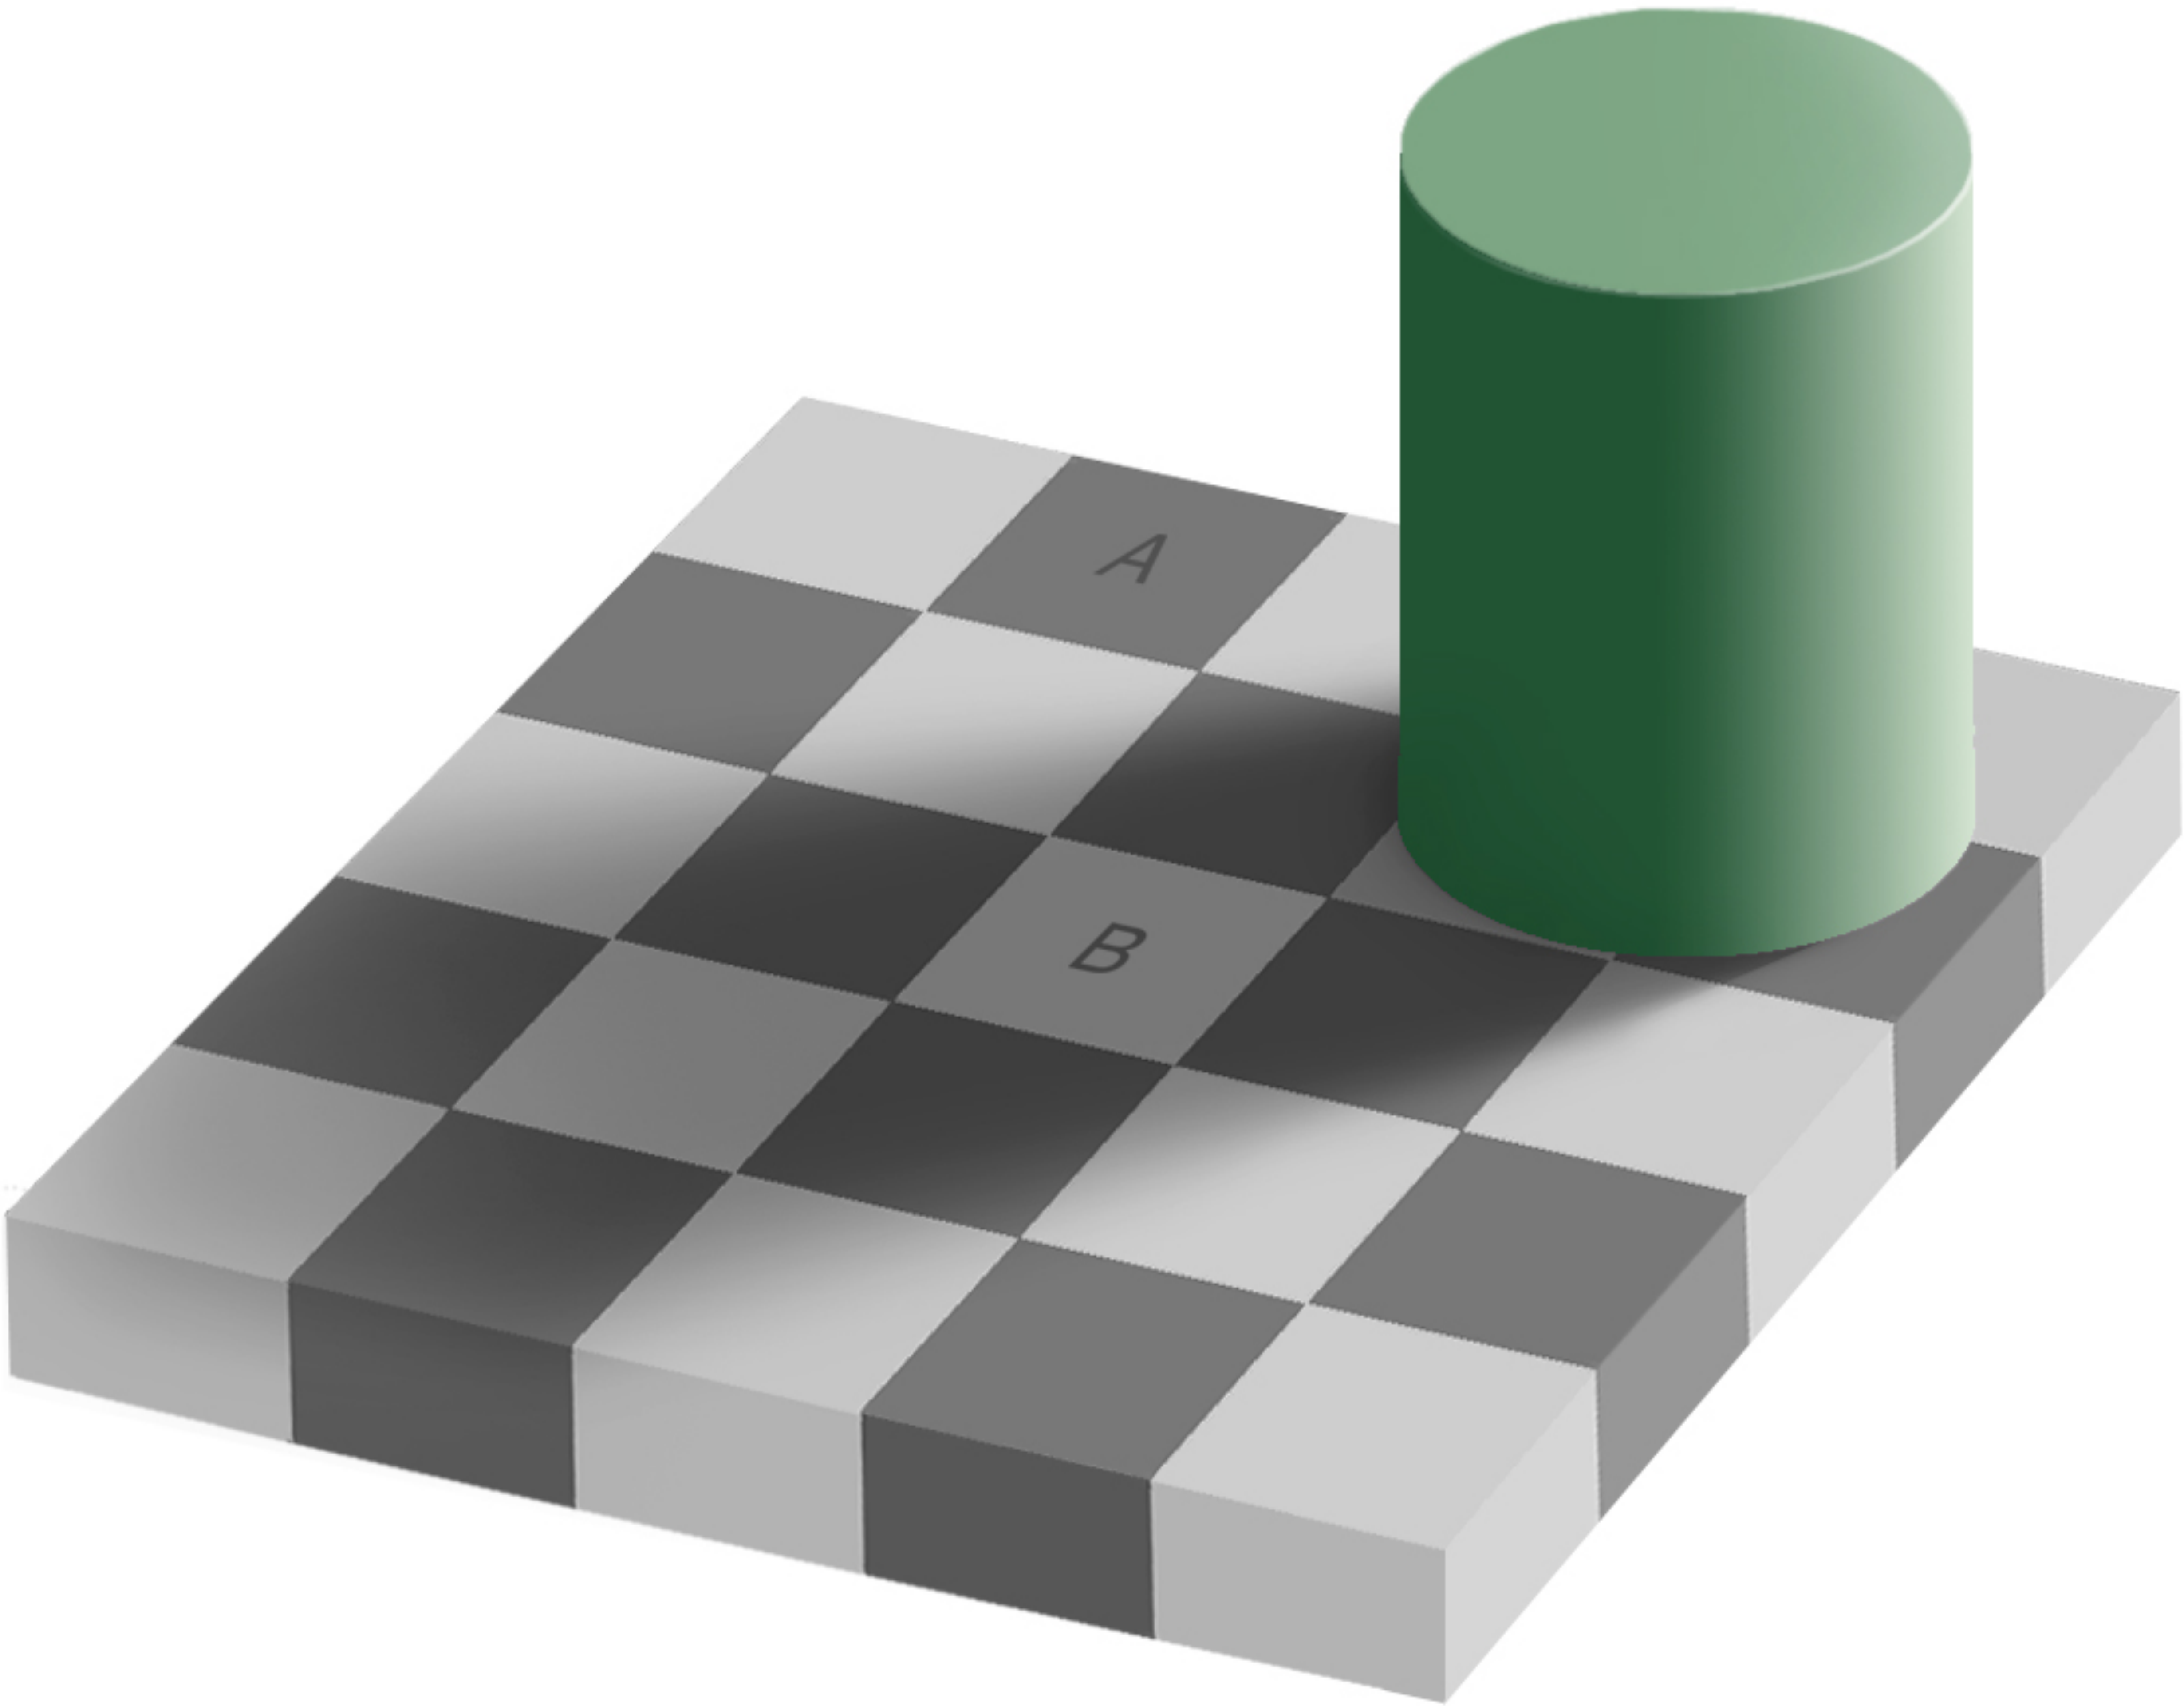
A
B

Gestalt Laws

A. Law of Closure

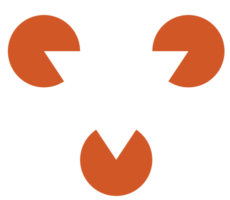

B. Law of Similarity

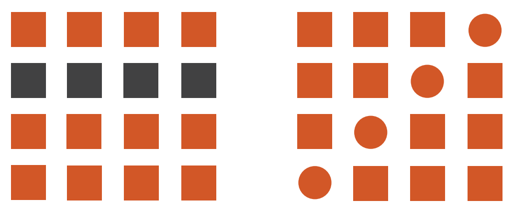

C. Law of Proximity

D. Law of Connectedness

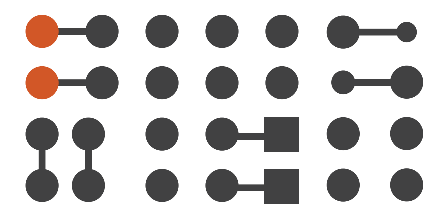

E. Law of Symmetry

[ ]{ }( )

F. Law of Good Continuation

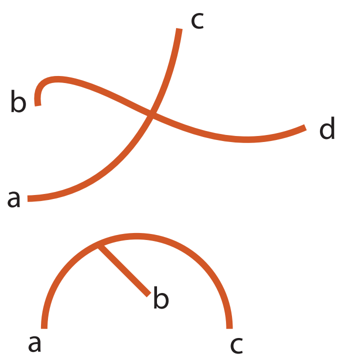

G. Contour Saliency

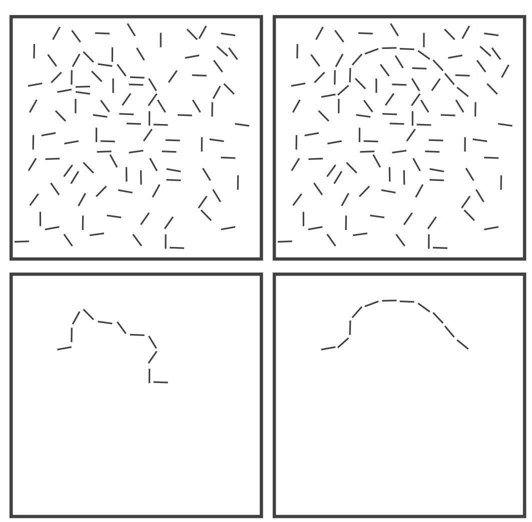

H. Law of Common Fate

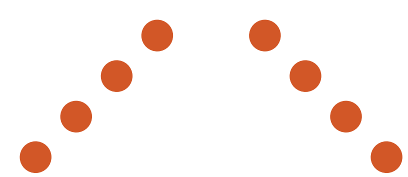

I. Law of Past Experience

J. Law of Pragnanz

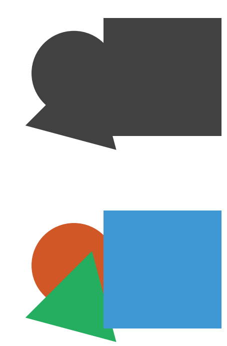

K. Figure/Ground

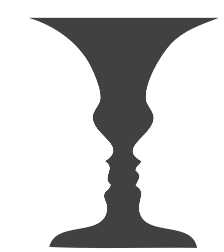

Integral/Separable Dimensions

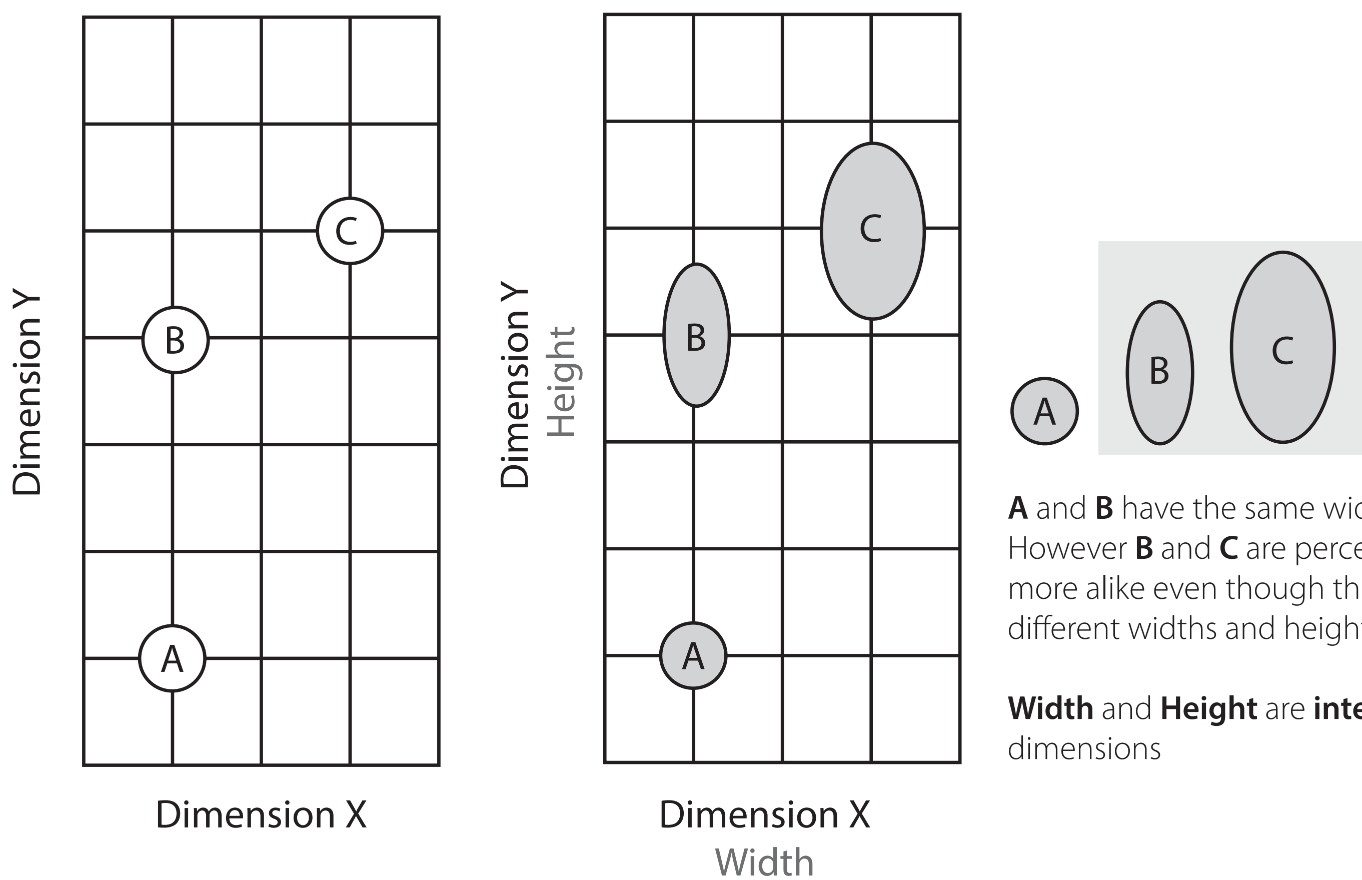

A and **B** have the same width. However **B** and **C** are perceived more alike even though they are different widths and heights.

Width and **Height** are **integral** dimensions

Integral/Separable Dimensions

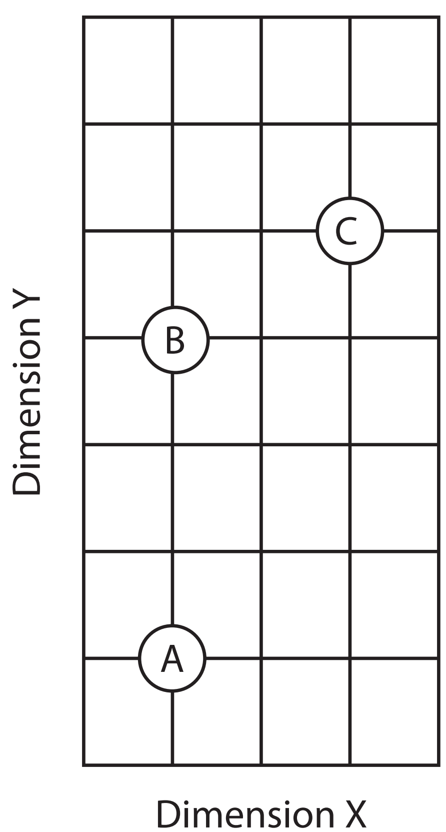

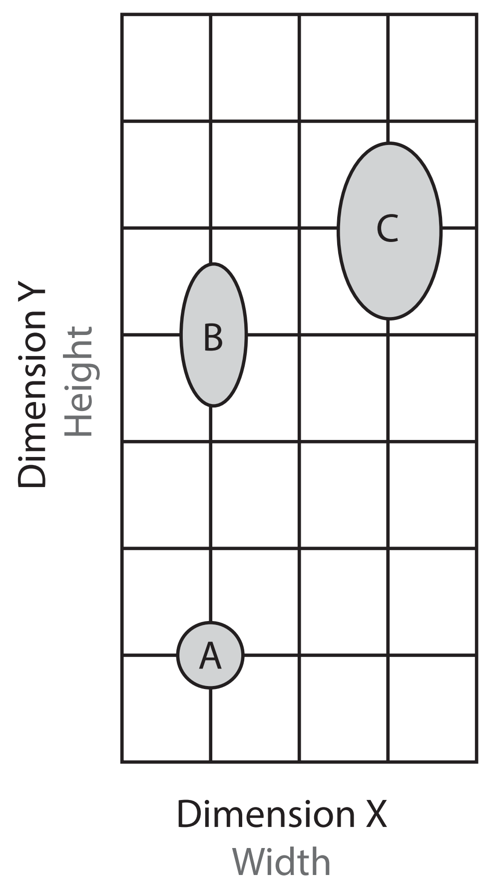

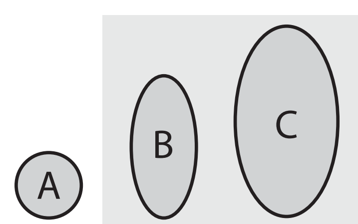

A and **B** have the same width. However **B** and **C** are perceived more alike even though they are different widths and heights.

Width and **Height** are **integral** dimensions

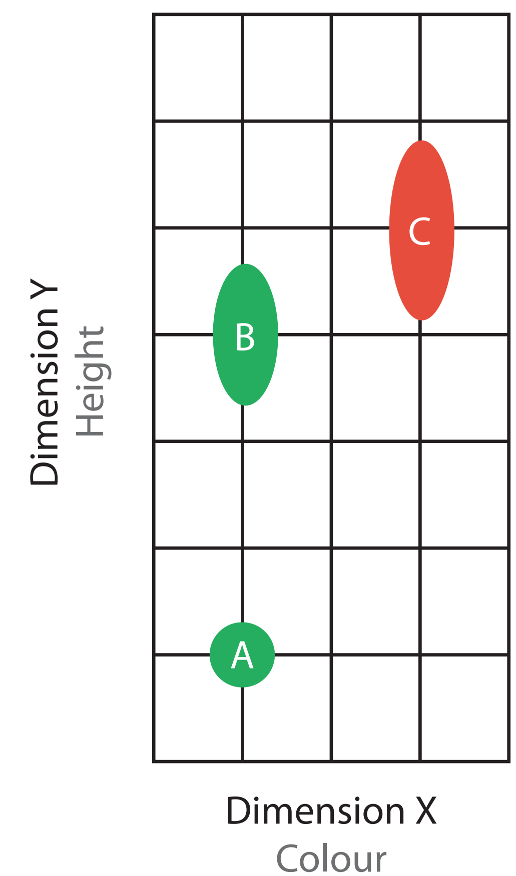

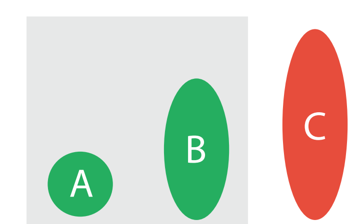

A and **B** have the same colour and are perceived more similar.

Colour and **Height** are **Separable** dimensions

Integral/Separable Dimensions

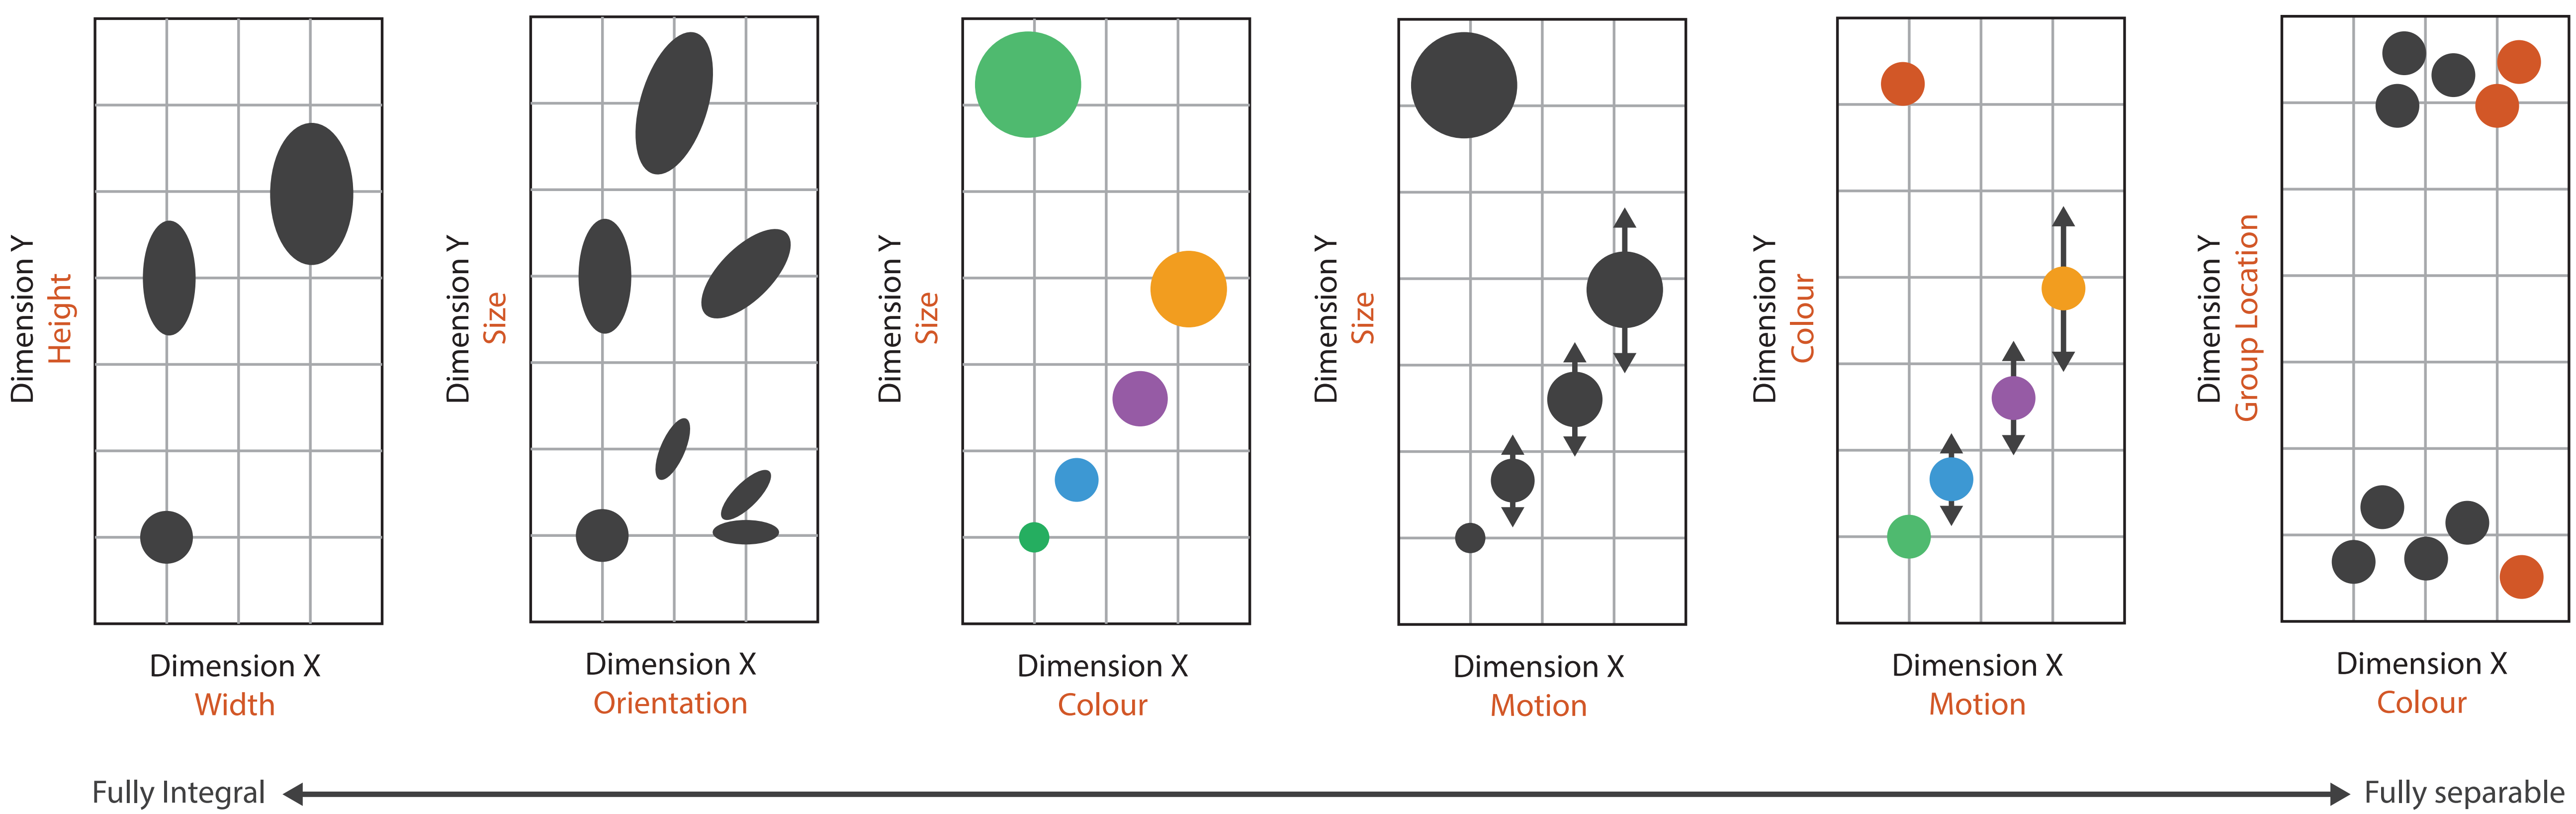

HOW

We have to be careful when mapping data to the visual world

Some visual channels are more effective for some data types over others.

Some data has a **natural mapping** that our brains expect given certain types of data

There are many visual tricks that can be observed due to how the visual system works

We don’t see in 3D, and we have difficulties interpreting information on the Z-axis.

2D always wins...

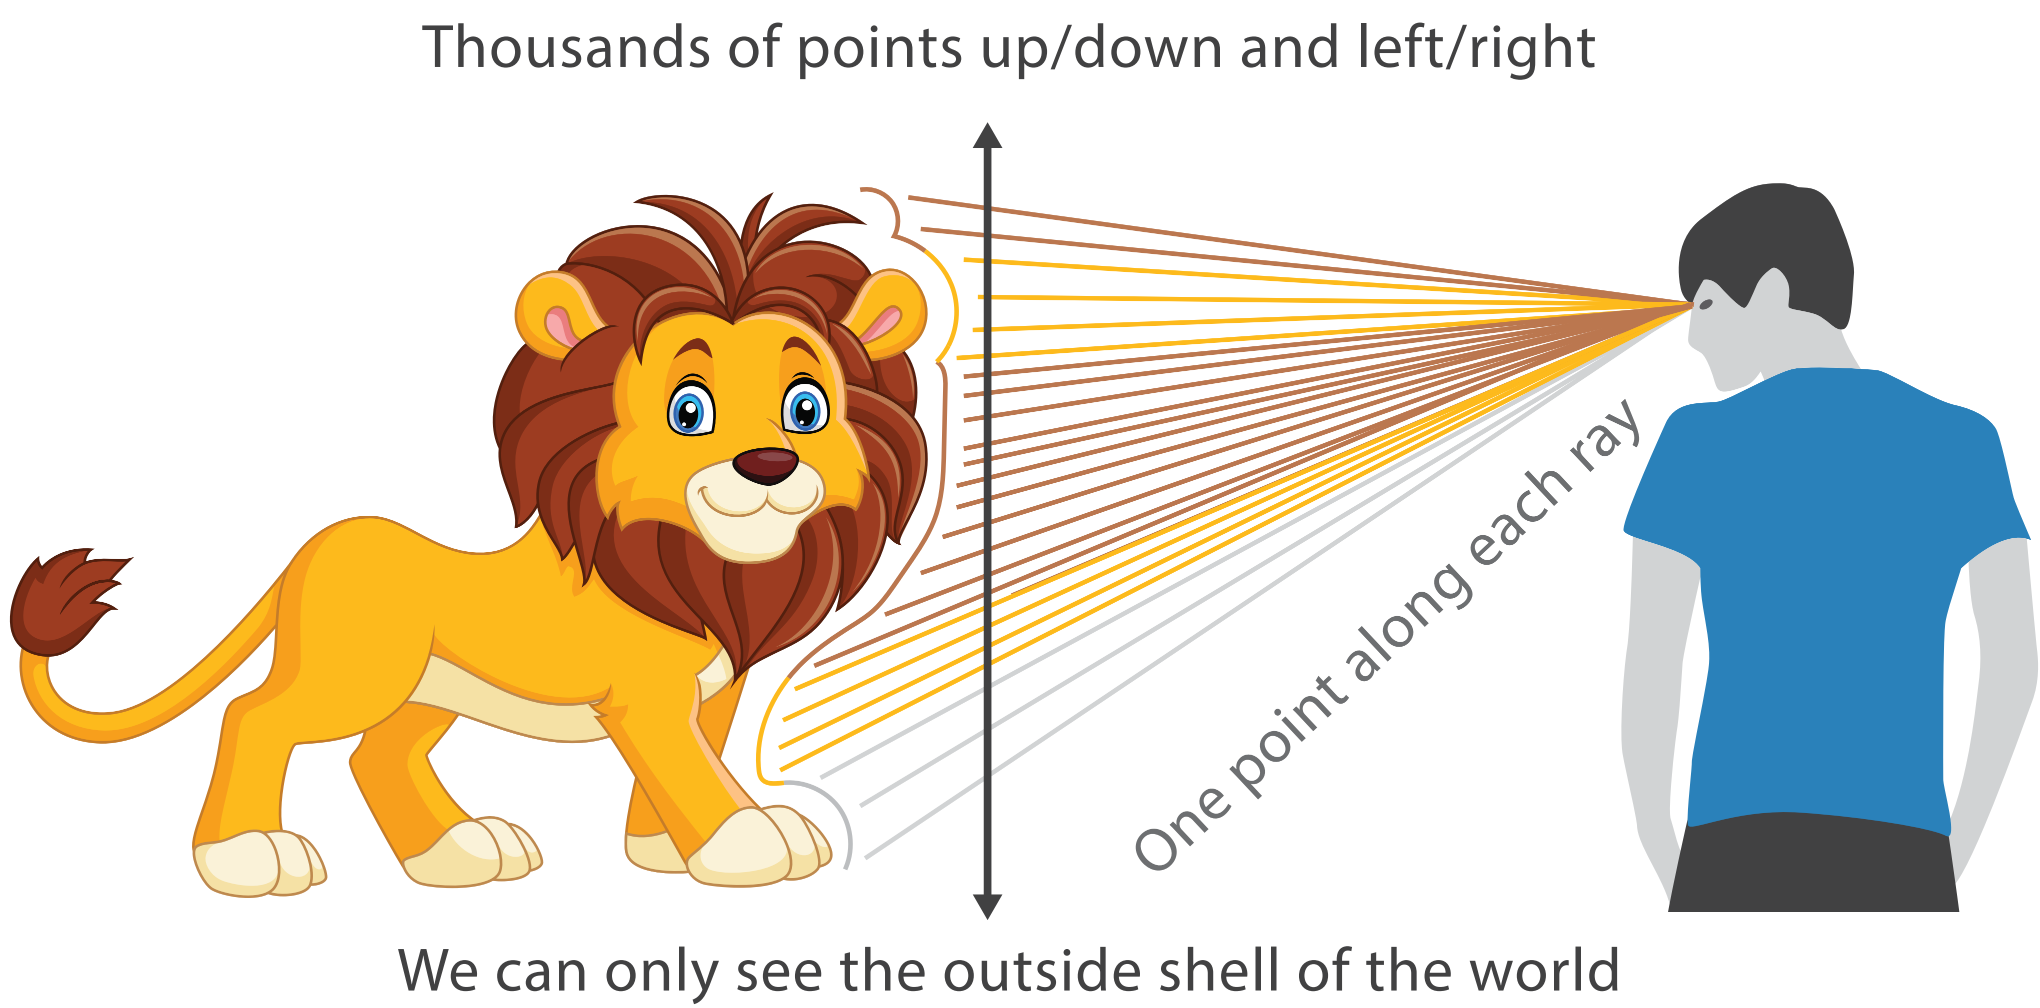

Our visual system is not good at interpreting information on the z-axis.

*3D is normally only used for exploration of inherently 3D information, such as medical imaging data…

2D always wins...

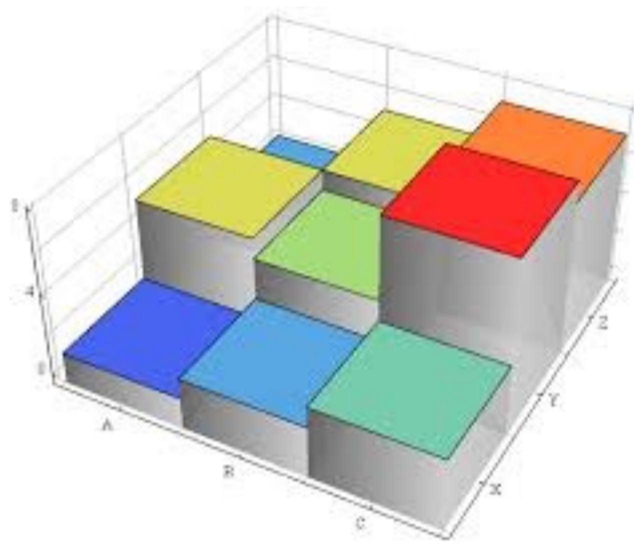

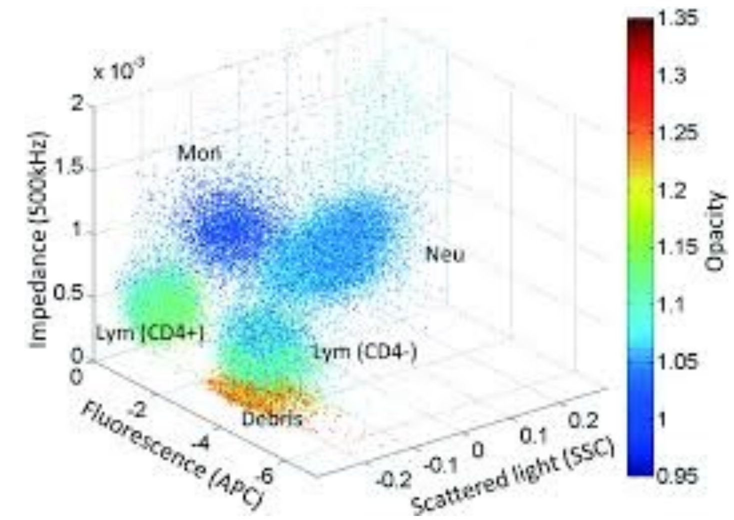

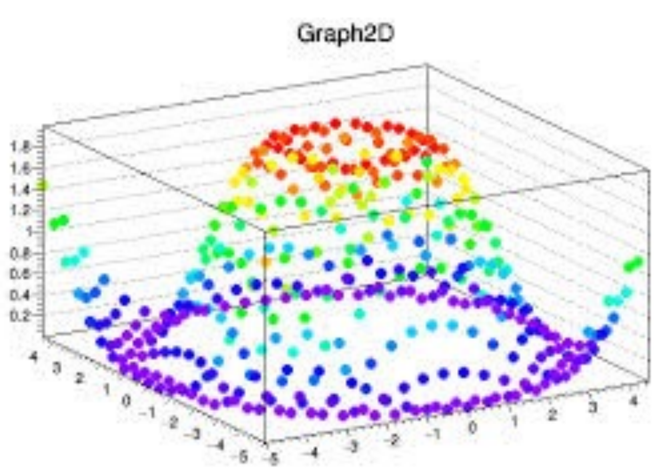

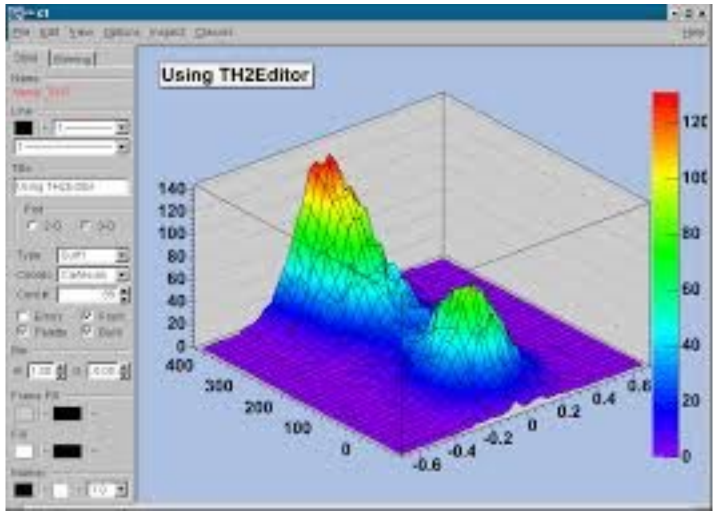

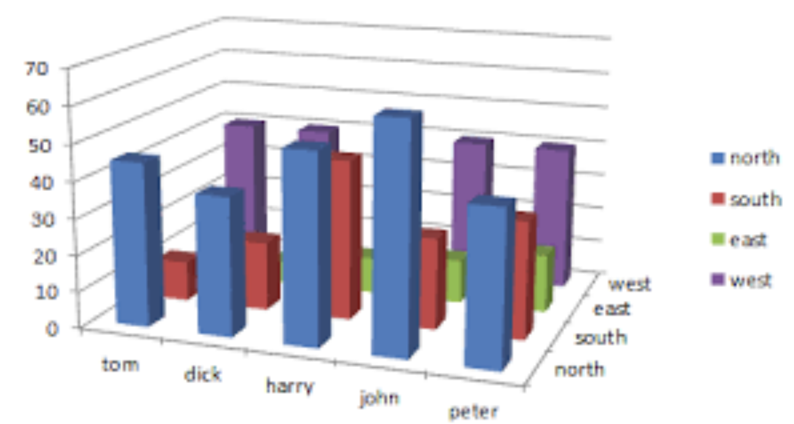

These options, taken randomly from google image searches so how widely 3D is abused in information visualisation. All of these charts are manipulating our perception of the data by using the Z axis to occlude information…it would be avoided in 2D.

2D always wins...

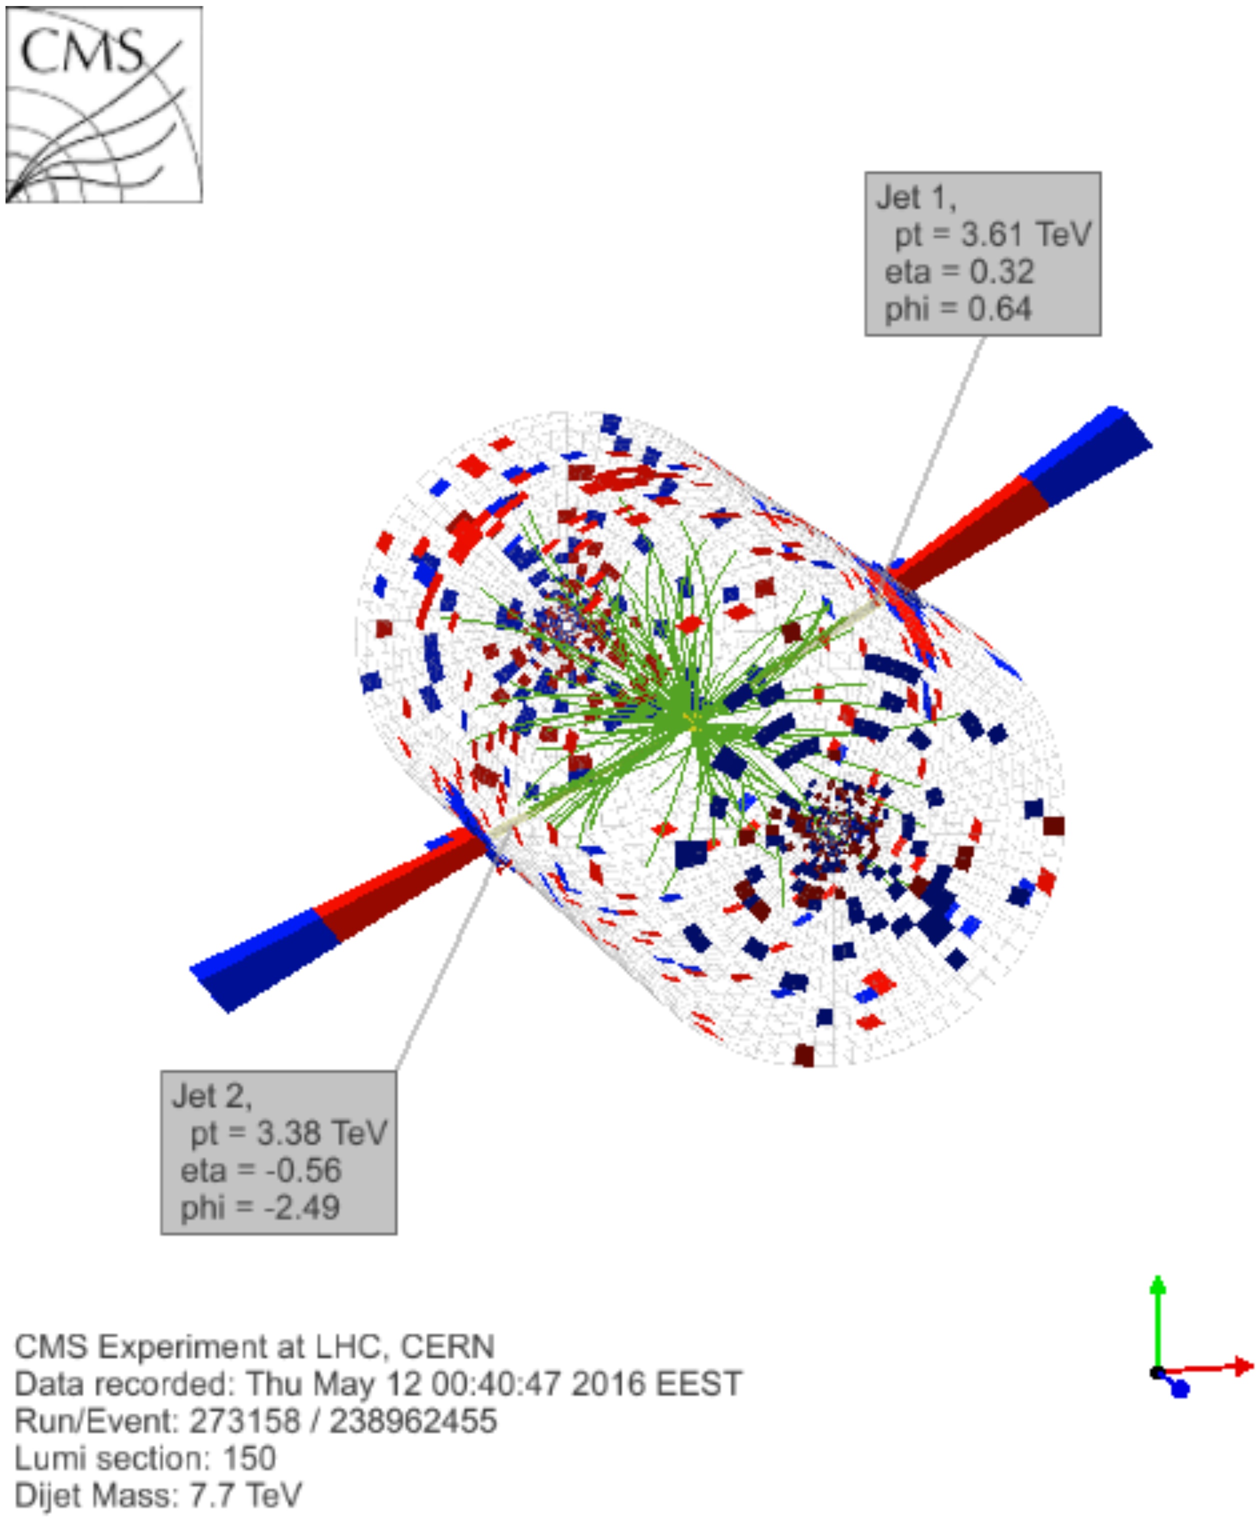

2D always wins...

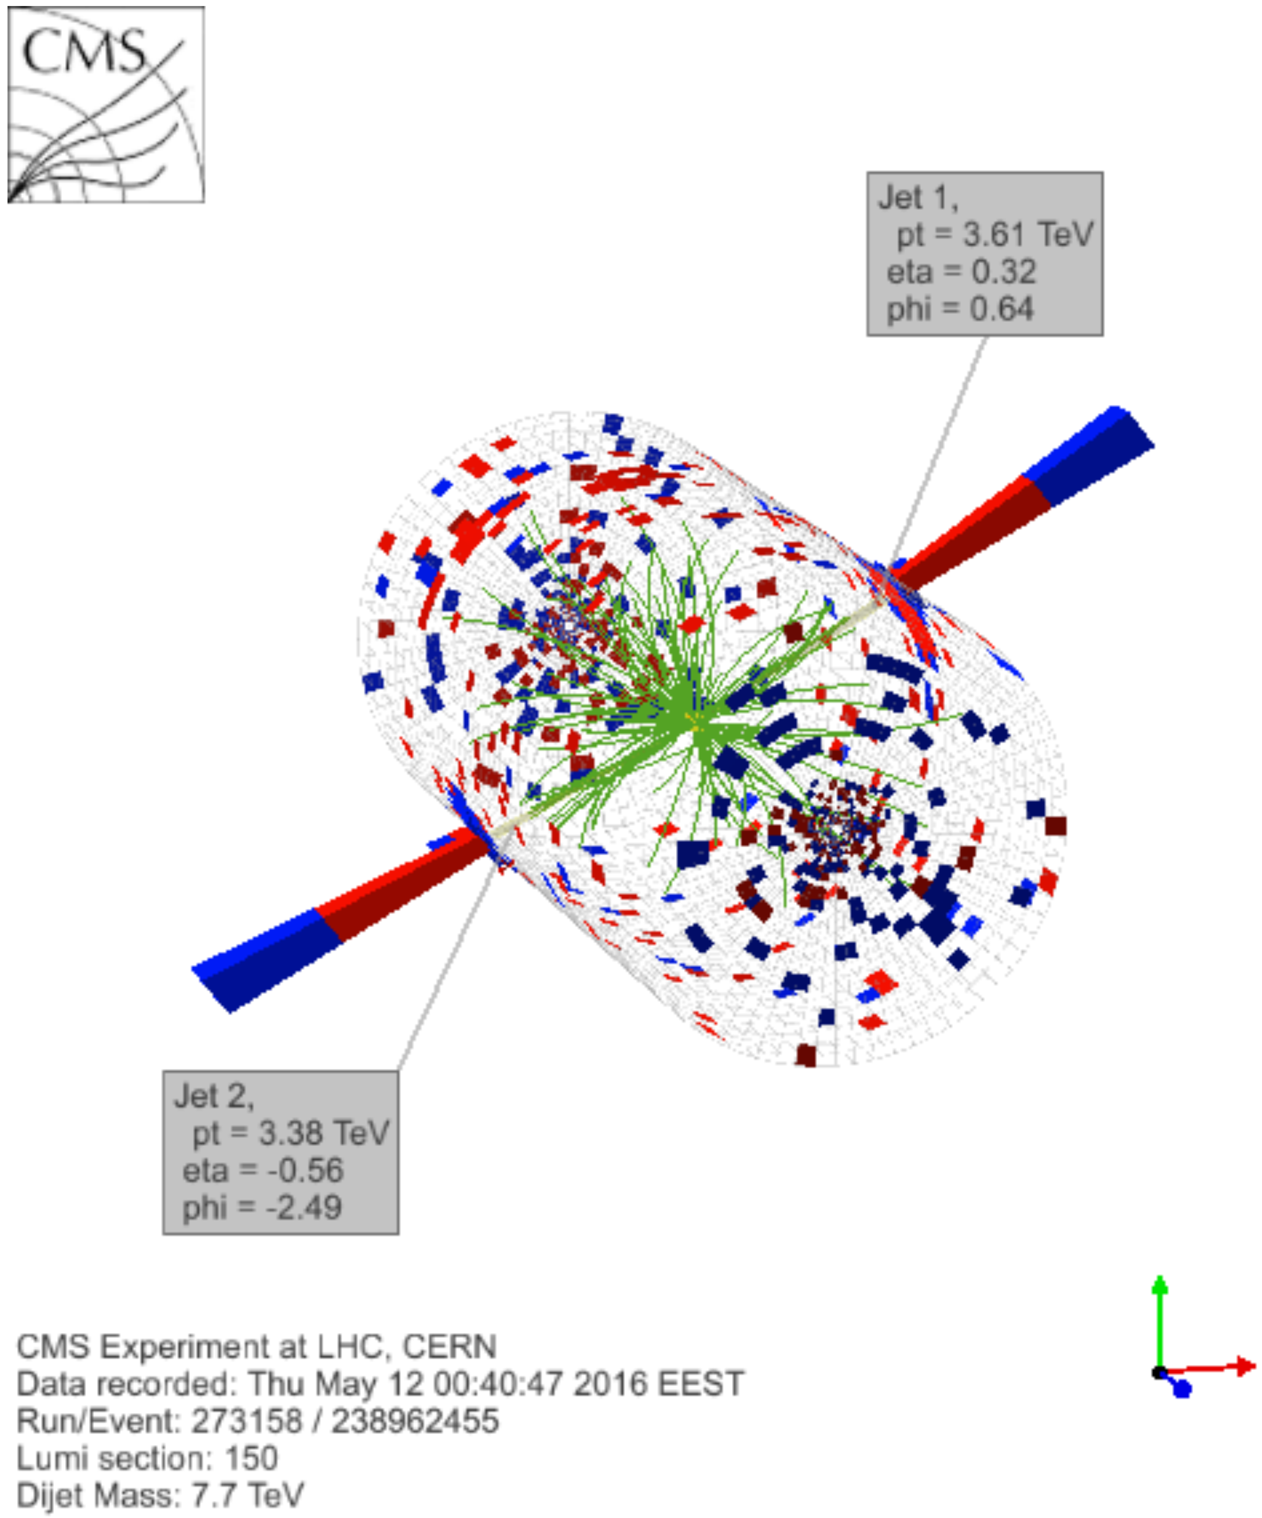

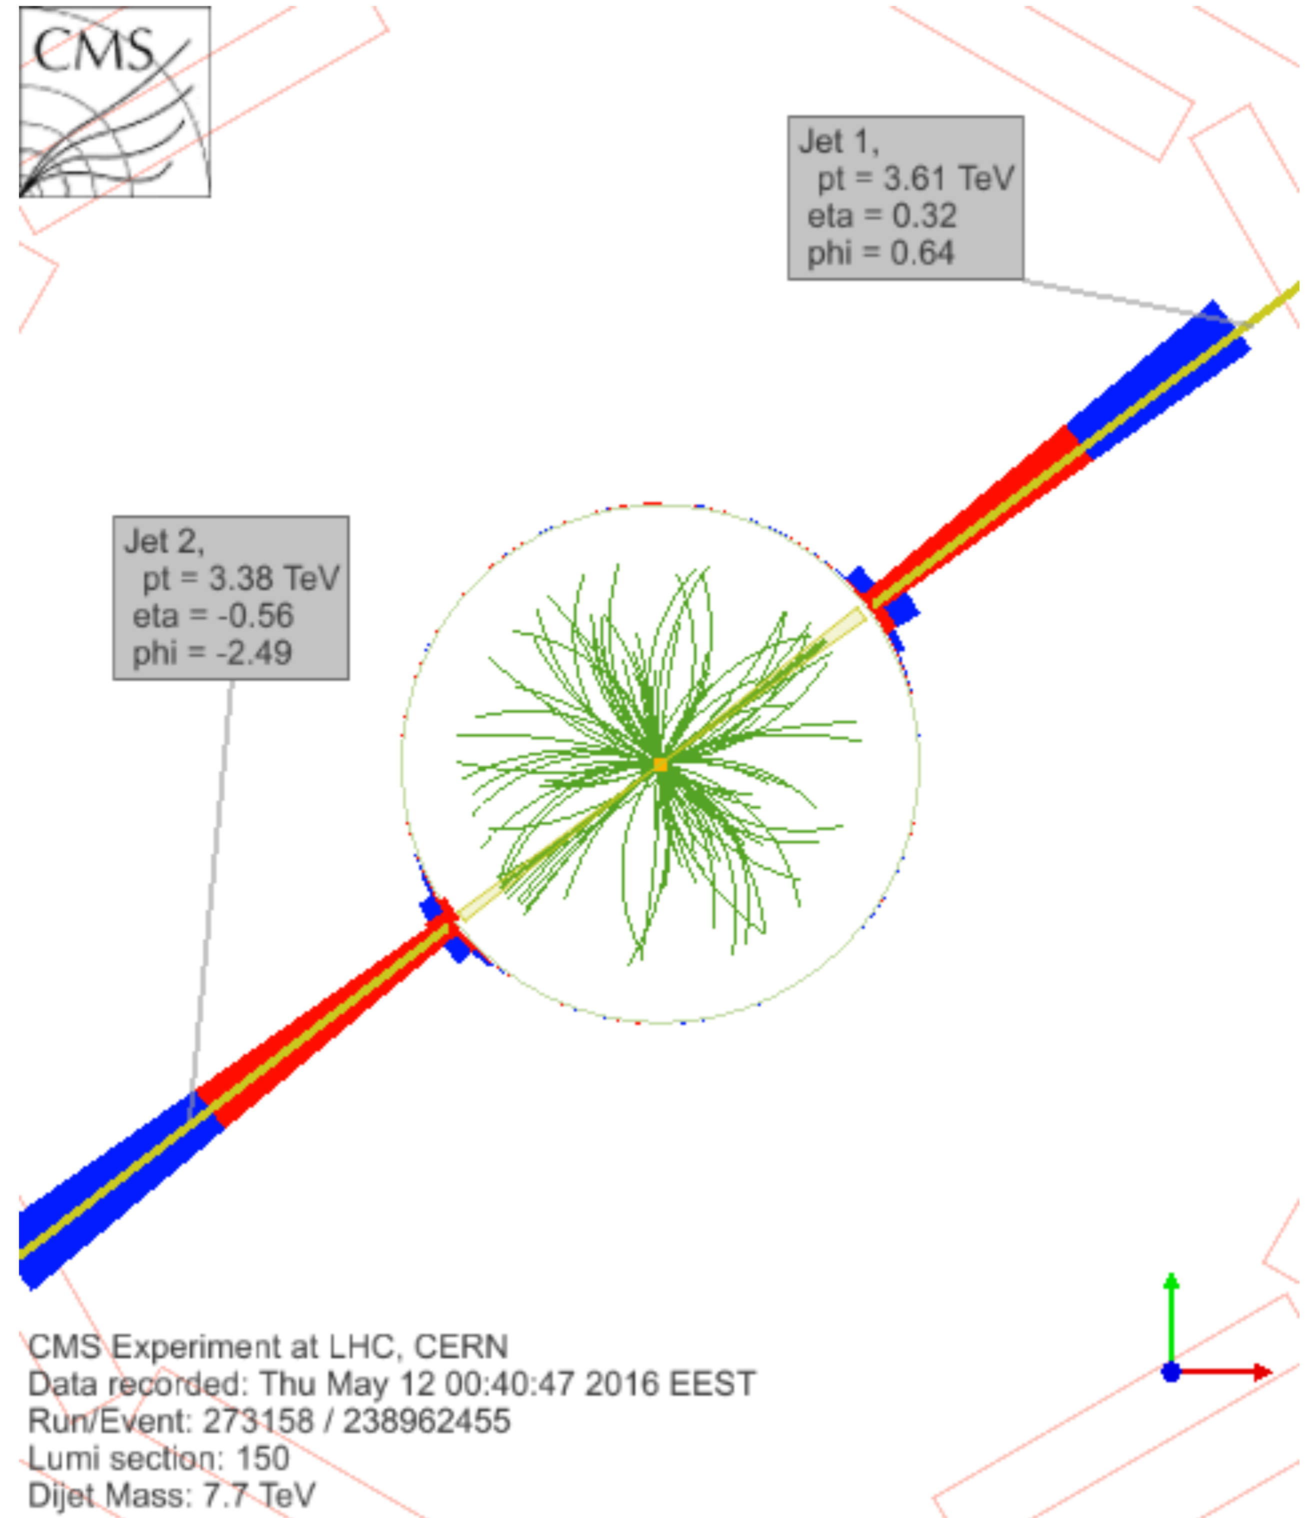

2D always wins…

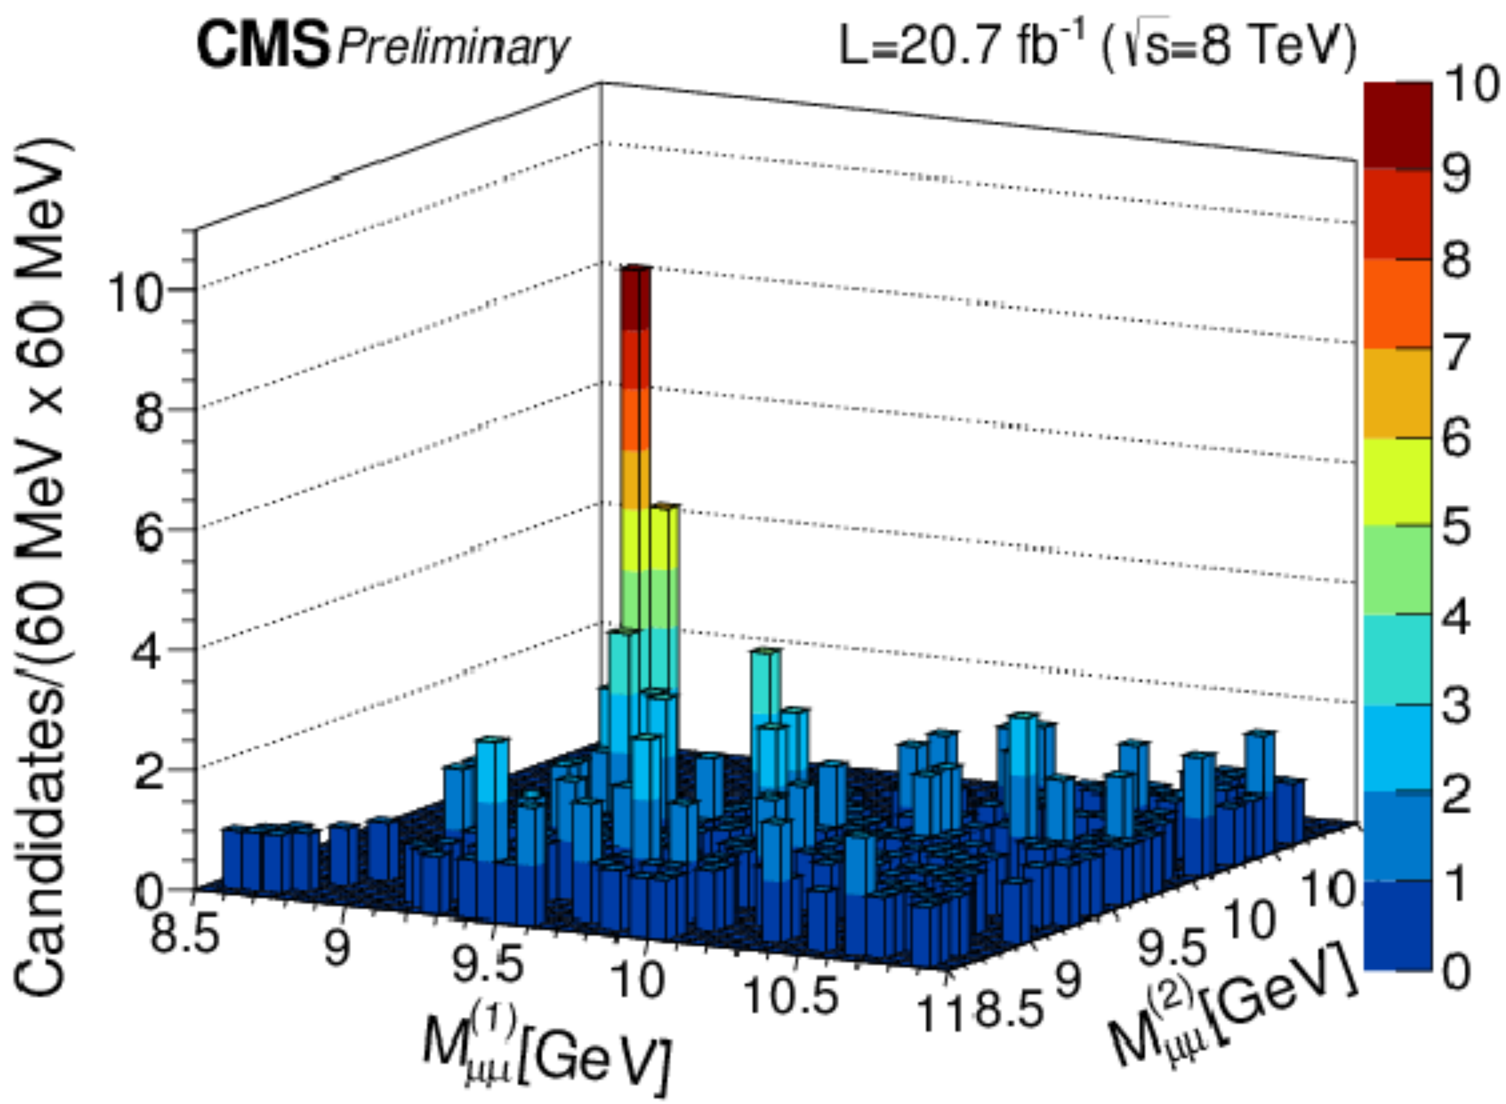

http://cms-results.web.cern.ch/cms-results/public-results/preliminary-results/BPH-14-008/index.html

HOW

We have to be careful when mapping data to the visual world

Some visual channels are more effective for some data types over others.

Some data has a **natural mapping** that our brains expect given certain types of data

There are many visual tricks that can be observed due to how the visual system works

We don’t see in 3D, and we have difficulties interpreting information on the Z-axis.

Colour

Colour

The simplest, yet most abused of all visual encodings.

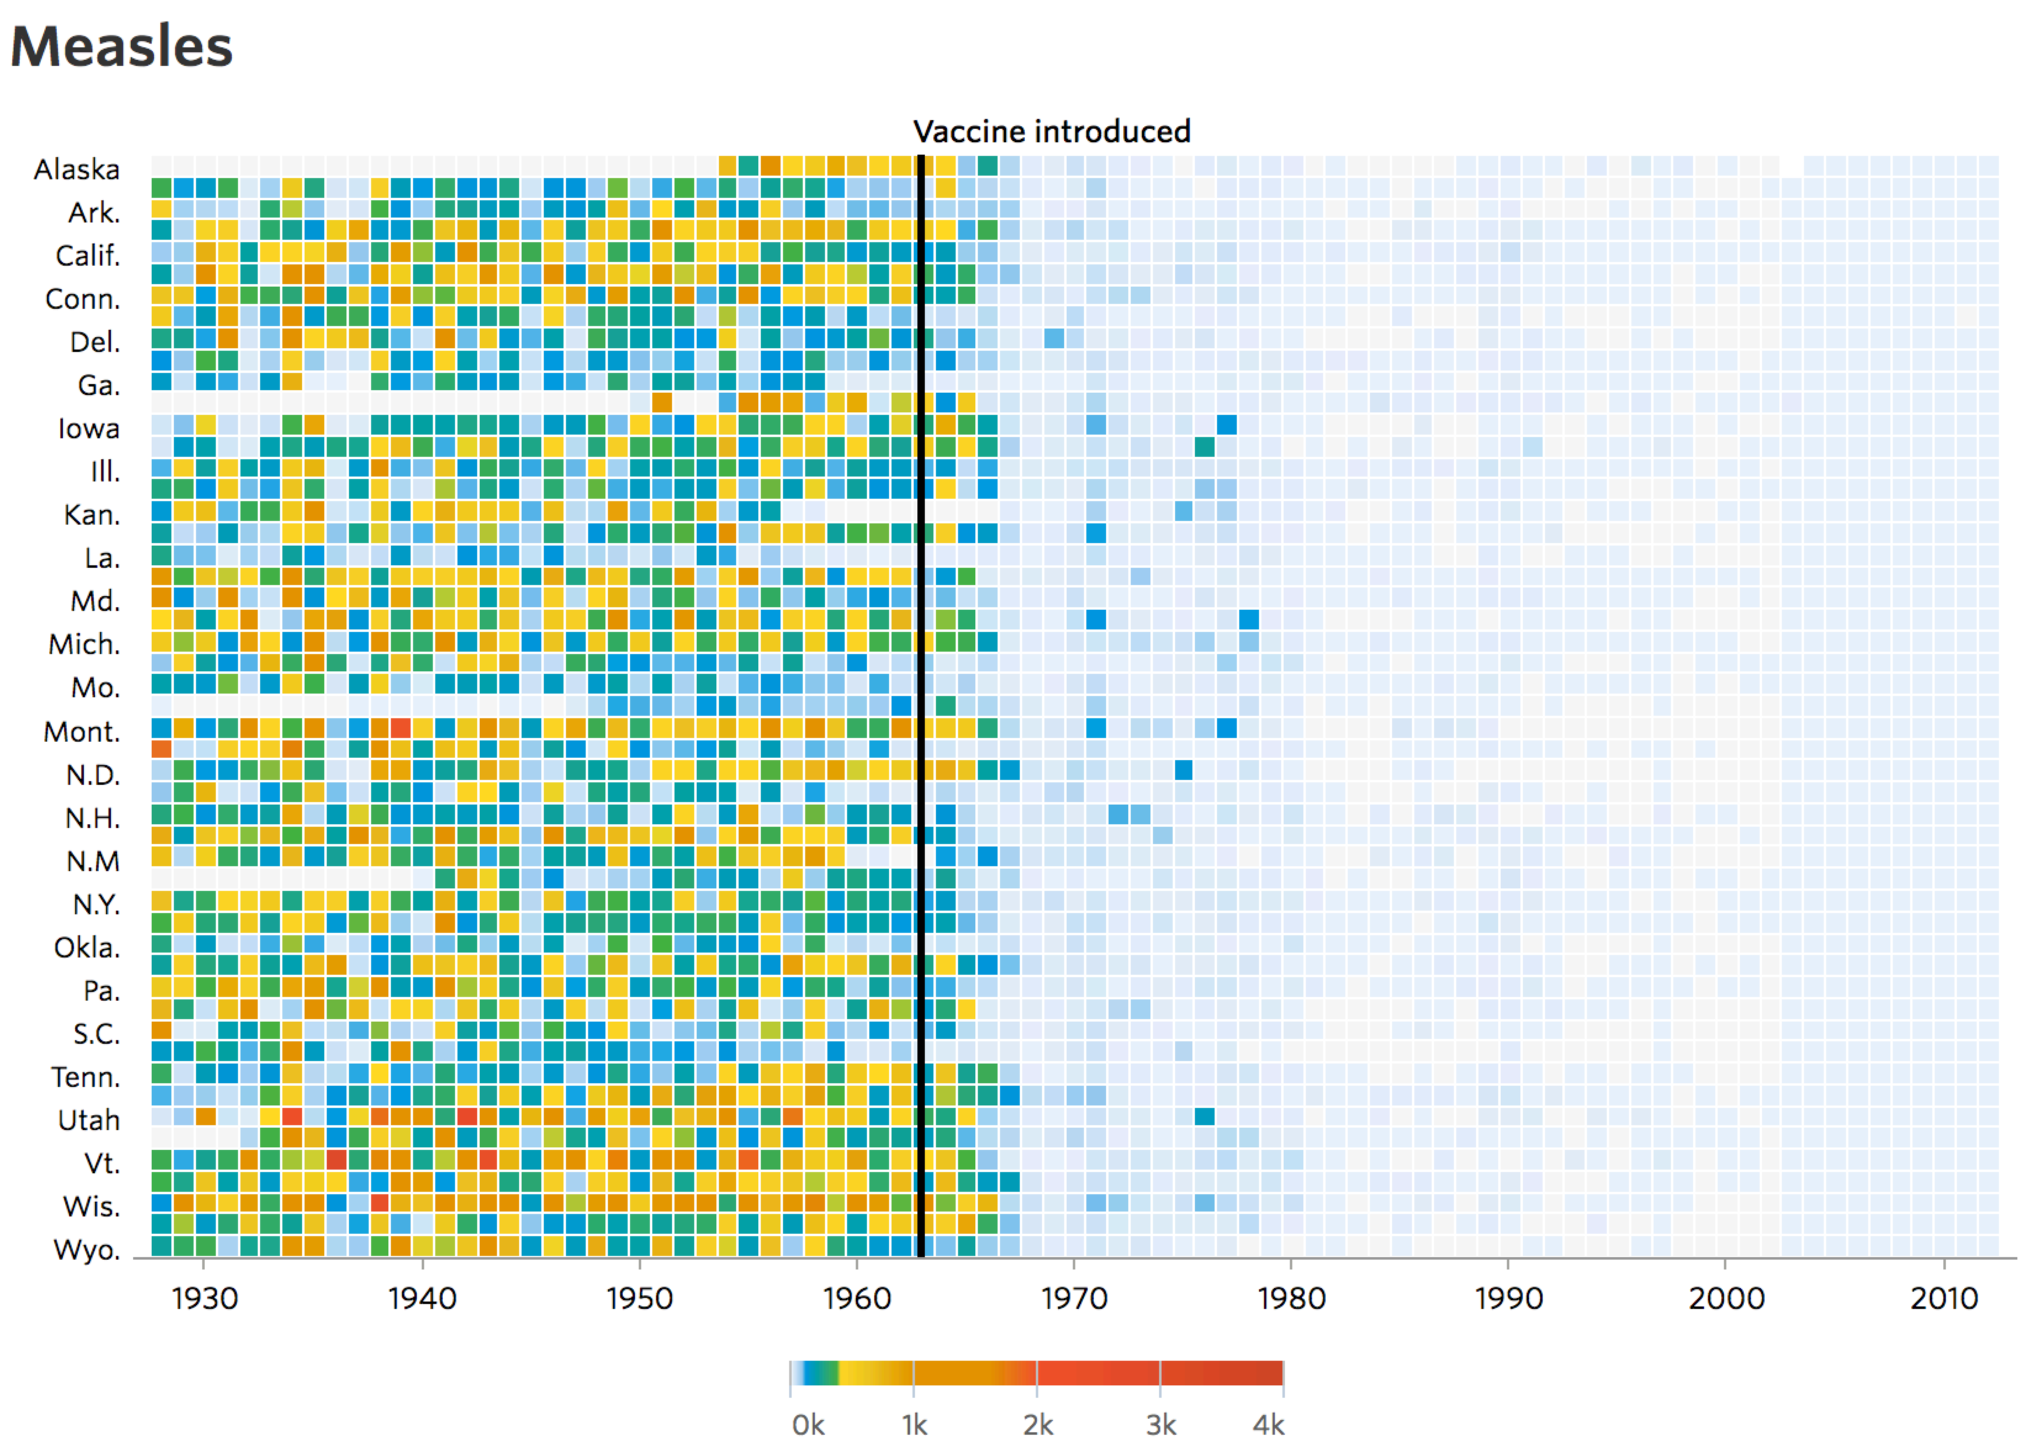

http://graphics.wsj.com/infectious-diseases-and-vaccines/

Colour

The simplest, yet most abused of all visual encodings.

SANFORD AND SELNICK

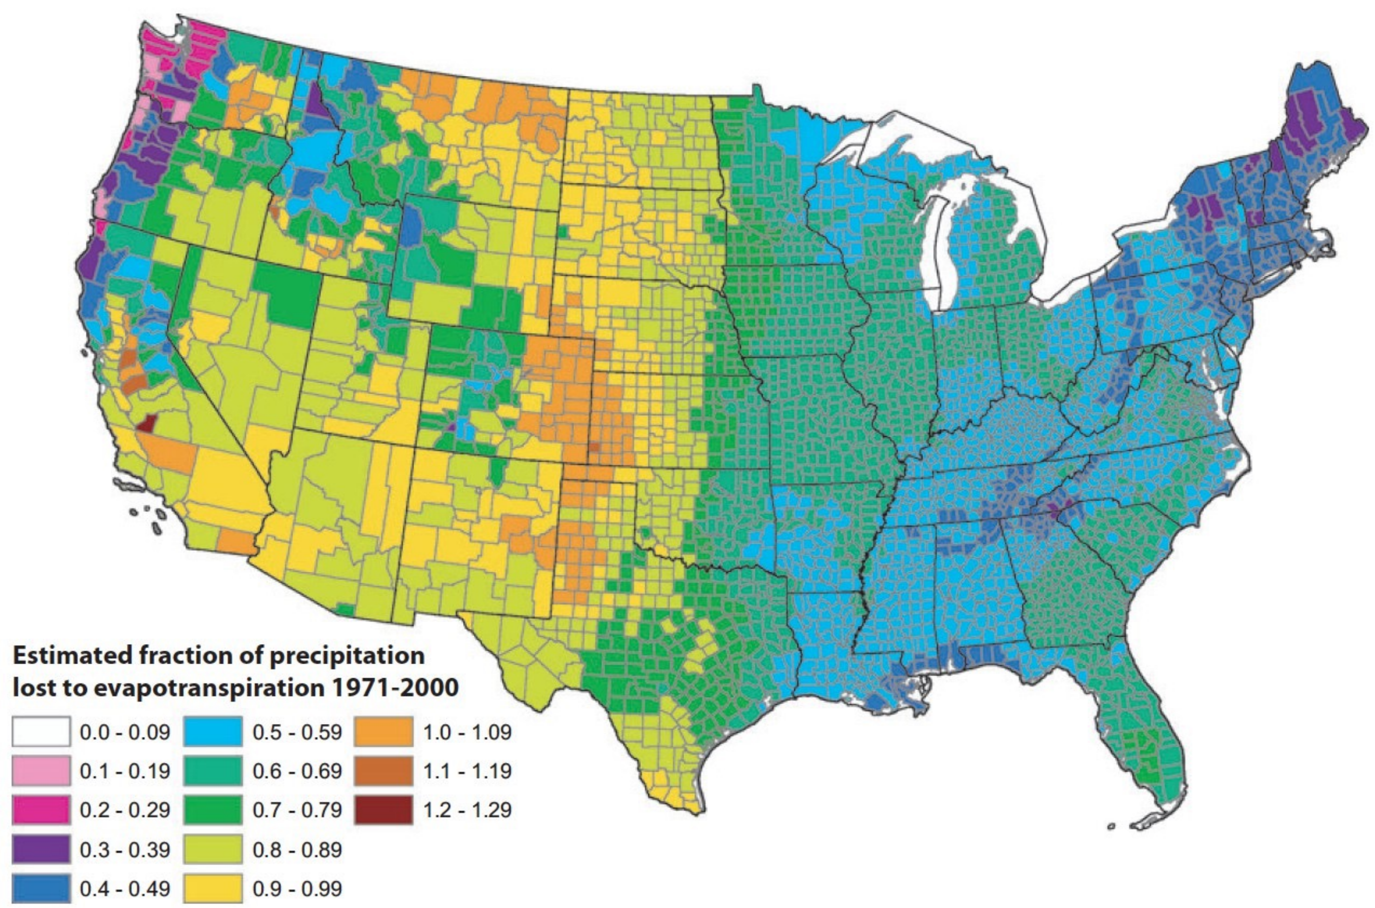

FIGURE 13. Estimated Mean Annual Ratio of Actual Evapotranspiration (ET) to Precipitation (P) for the Conterminous U.S. for the Period 1971-2000. Estimates are based on the regression equation in Table 1 that includes land cover. Calculations of ET/P were made first at the 800-m resolution of the PRISM climate data. The mean values for the counties (shown) were then calculated by averaging the 800-m values within each county. Areas with fractions >1 are agricultural counties that either import surface water or mine deep groundwater.

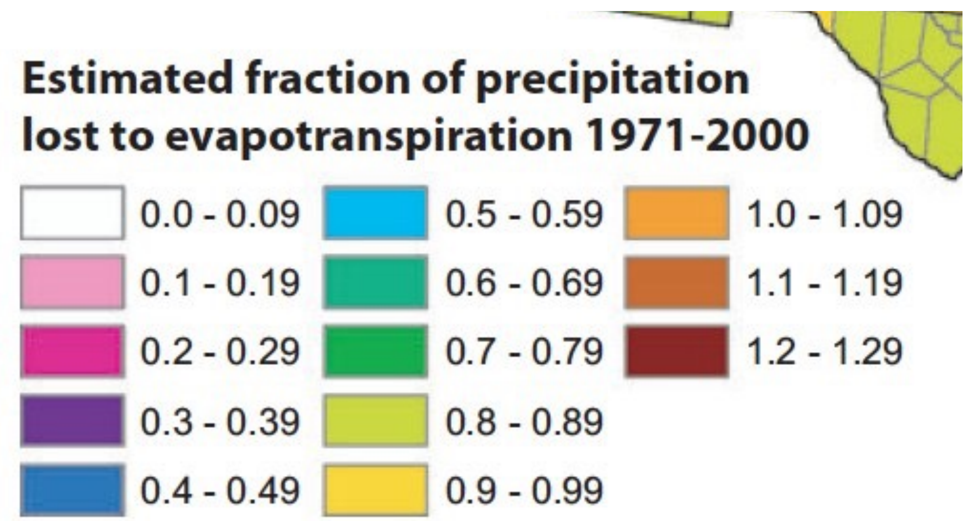

The problem is that a smooth step in a value does not equate to a smooth colour transition…

Colour

Additionally, colour is not equally binned in reality. We perceive colours differently due to an increased sensitivity to the yellow part of the spectrum…

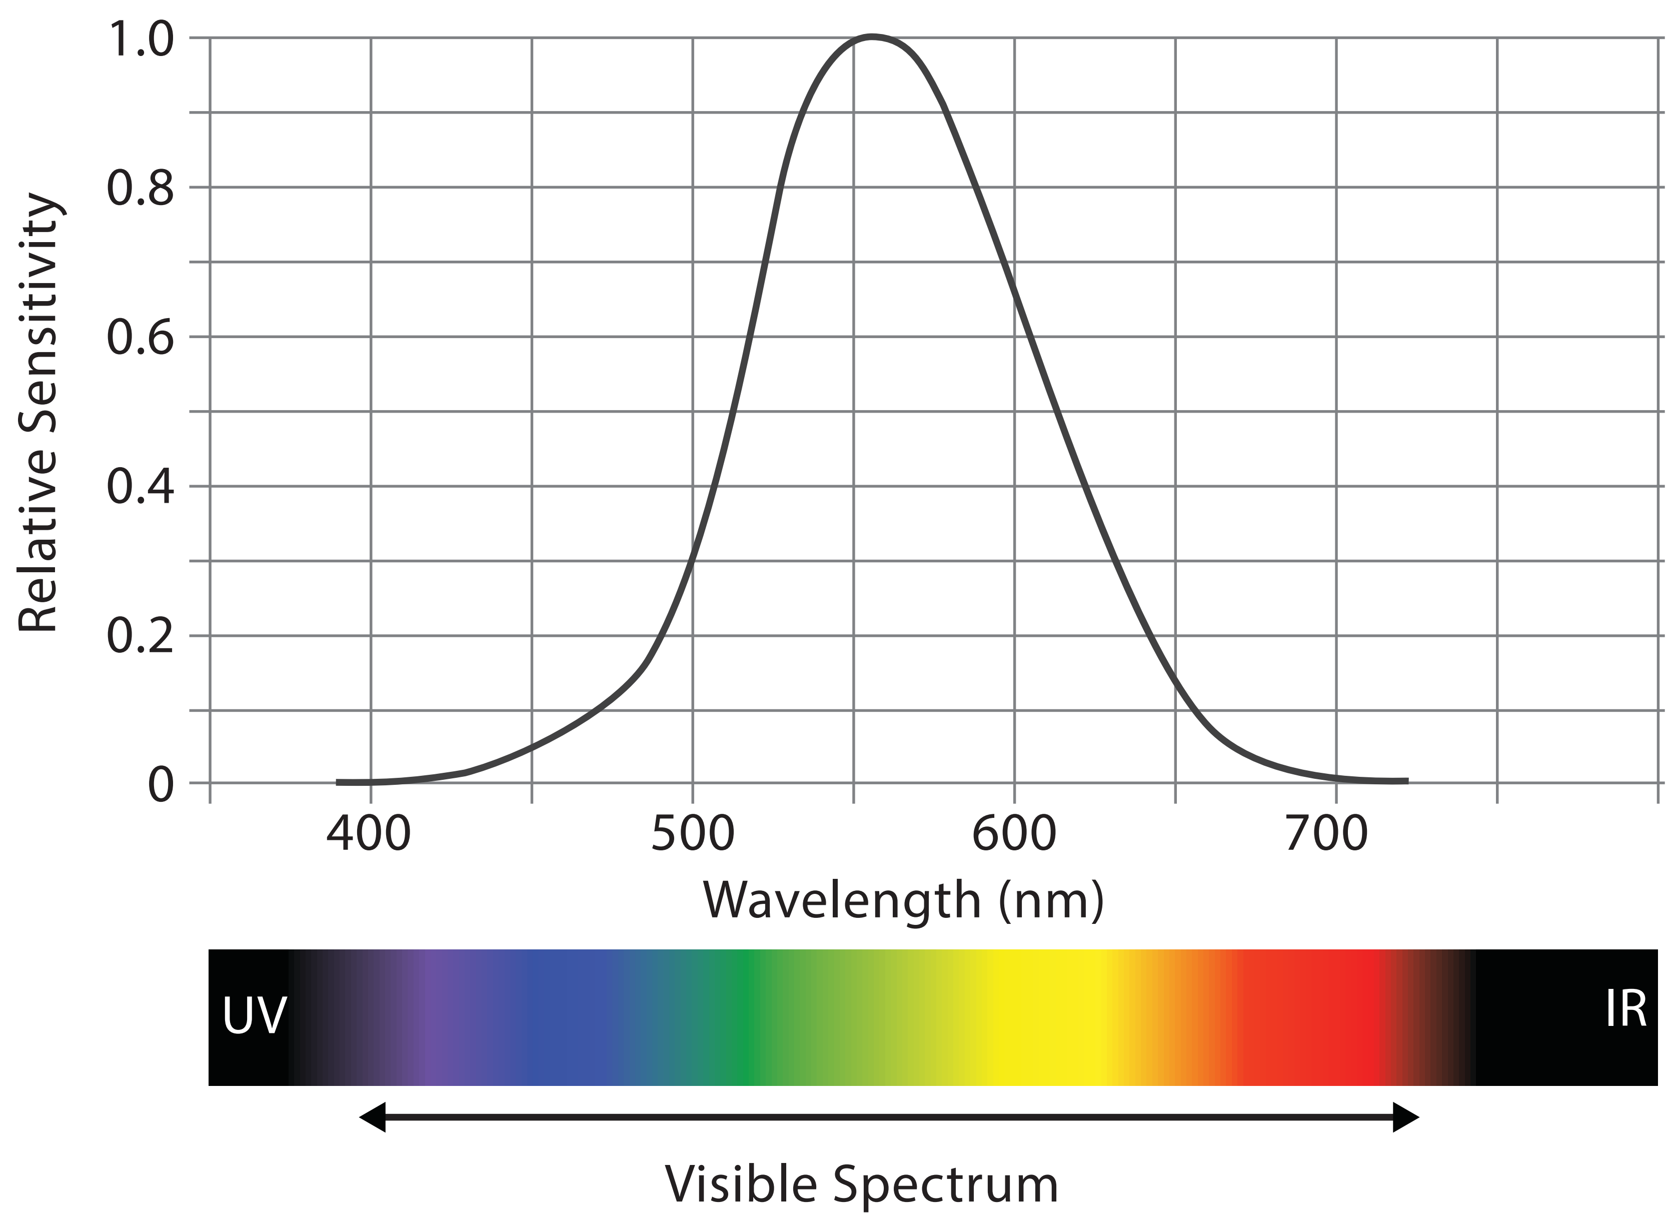

Colour

Luminosity is also not stable across the colours, meaning some colours will pop out more than others… and not always intentionally.

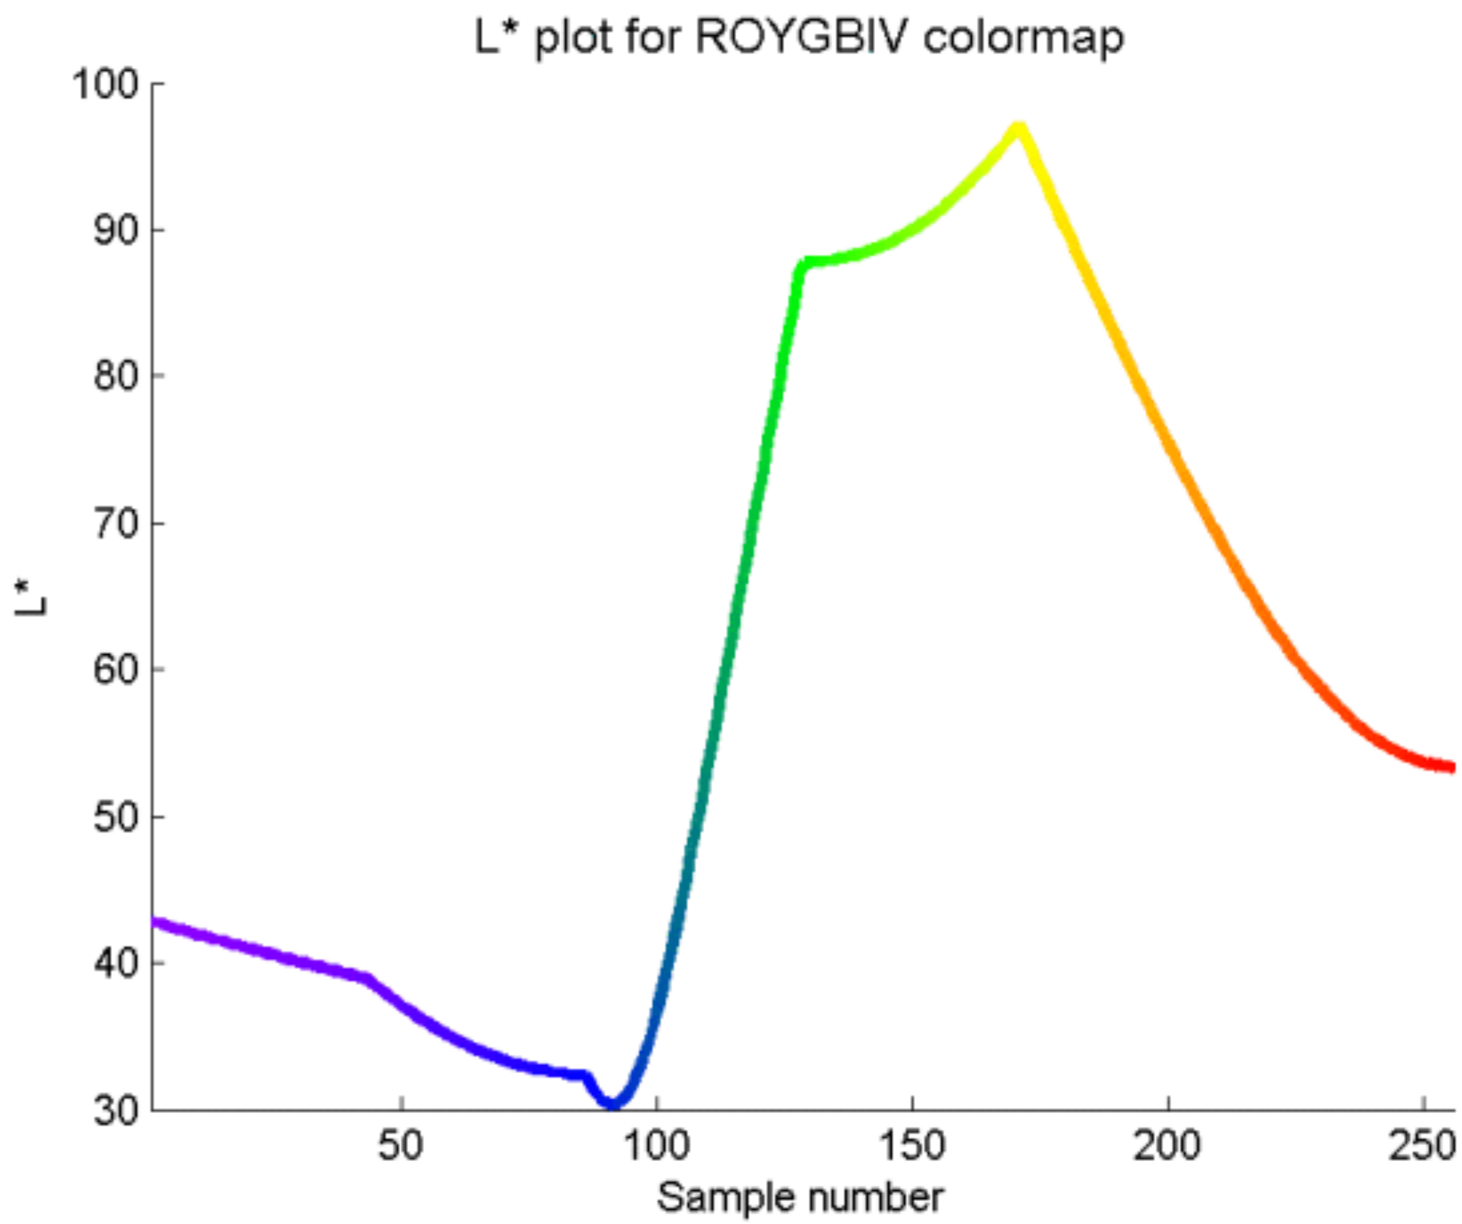

https://mycarta.wordpress.com/2012/10/06/the-rainbow-is-deadlong-live-the-rainbow-part-3/

Colour

Luminosity is also not stable across the colours, meaning some colours will pop out more than others… and not always intentionally.

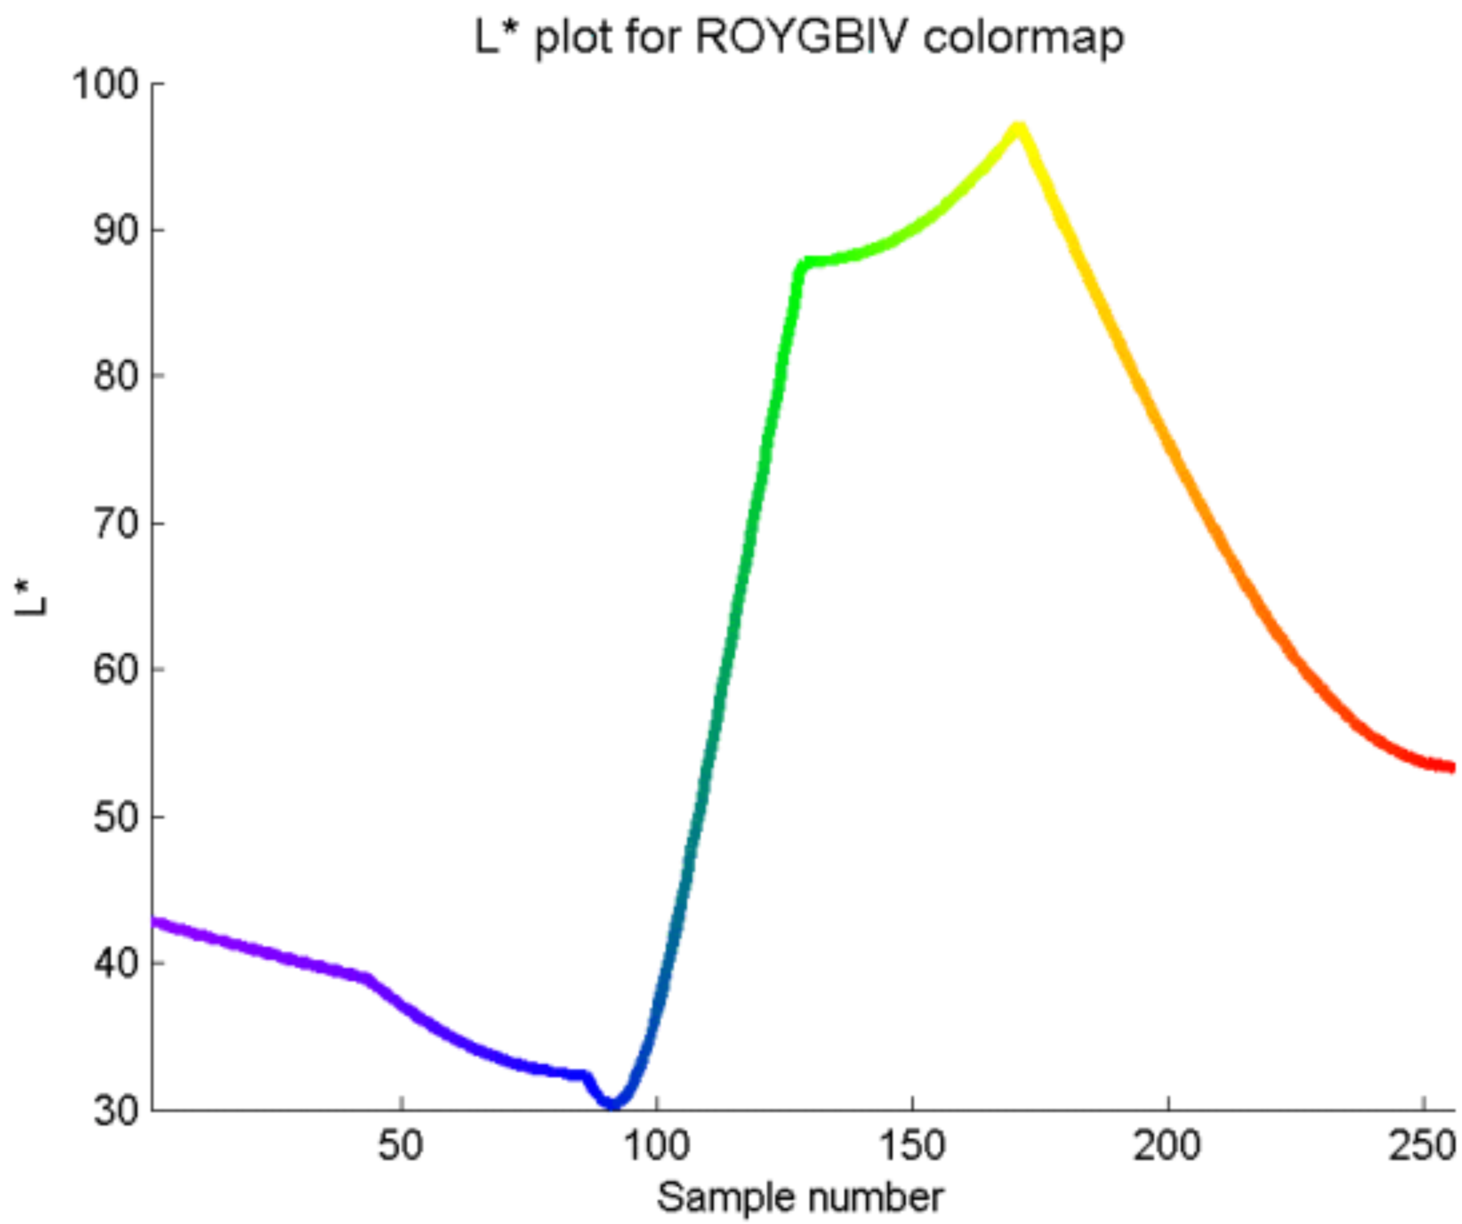

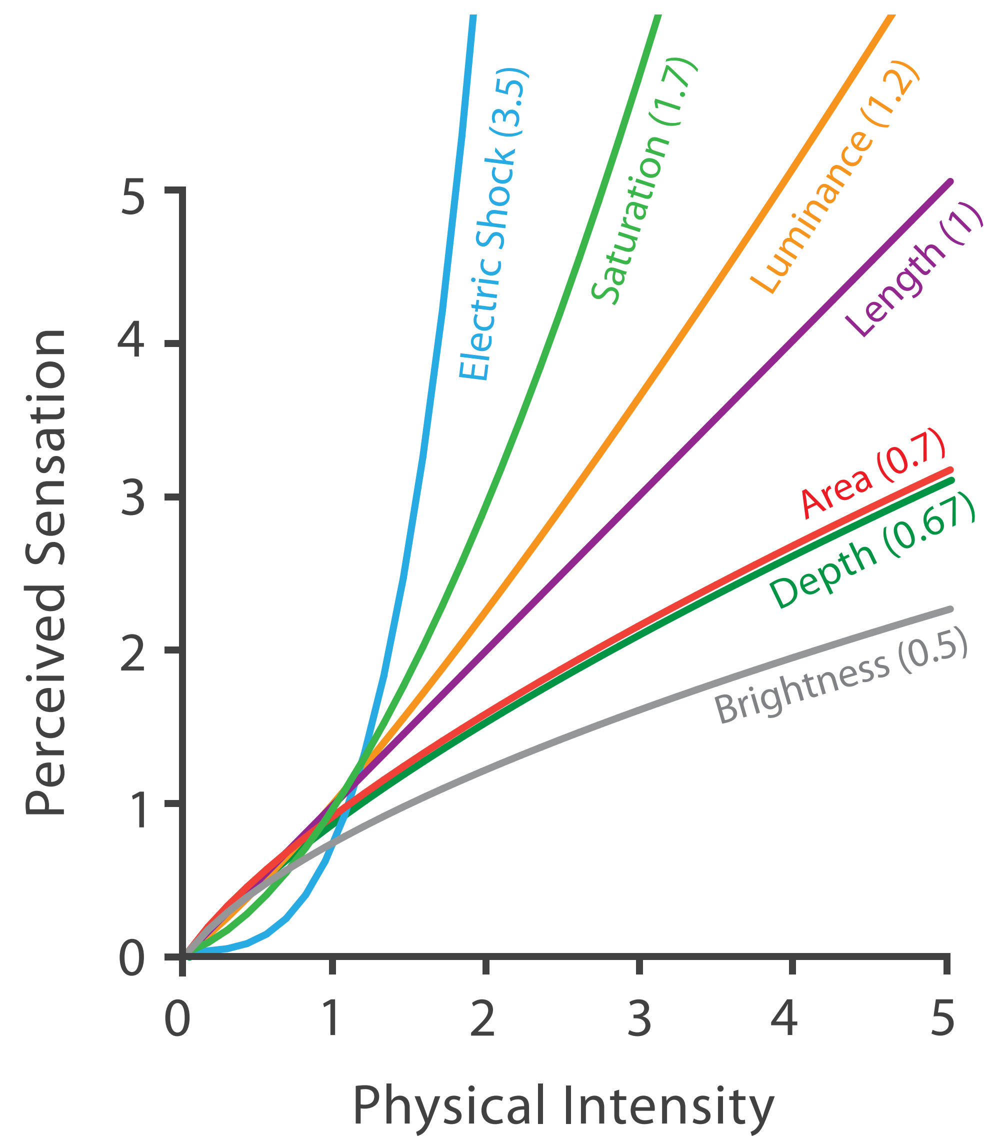

https://mycarta.wordpress.com/2012/10/06/the-rainbow-is-deadlong-live-the-rainbow-part-3/

Colour

And how we perceive changes in hue is also very different.

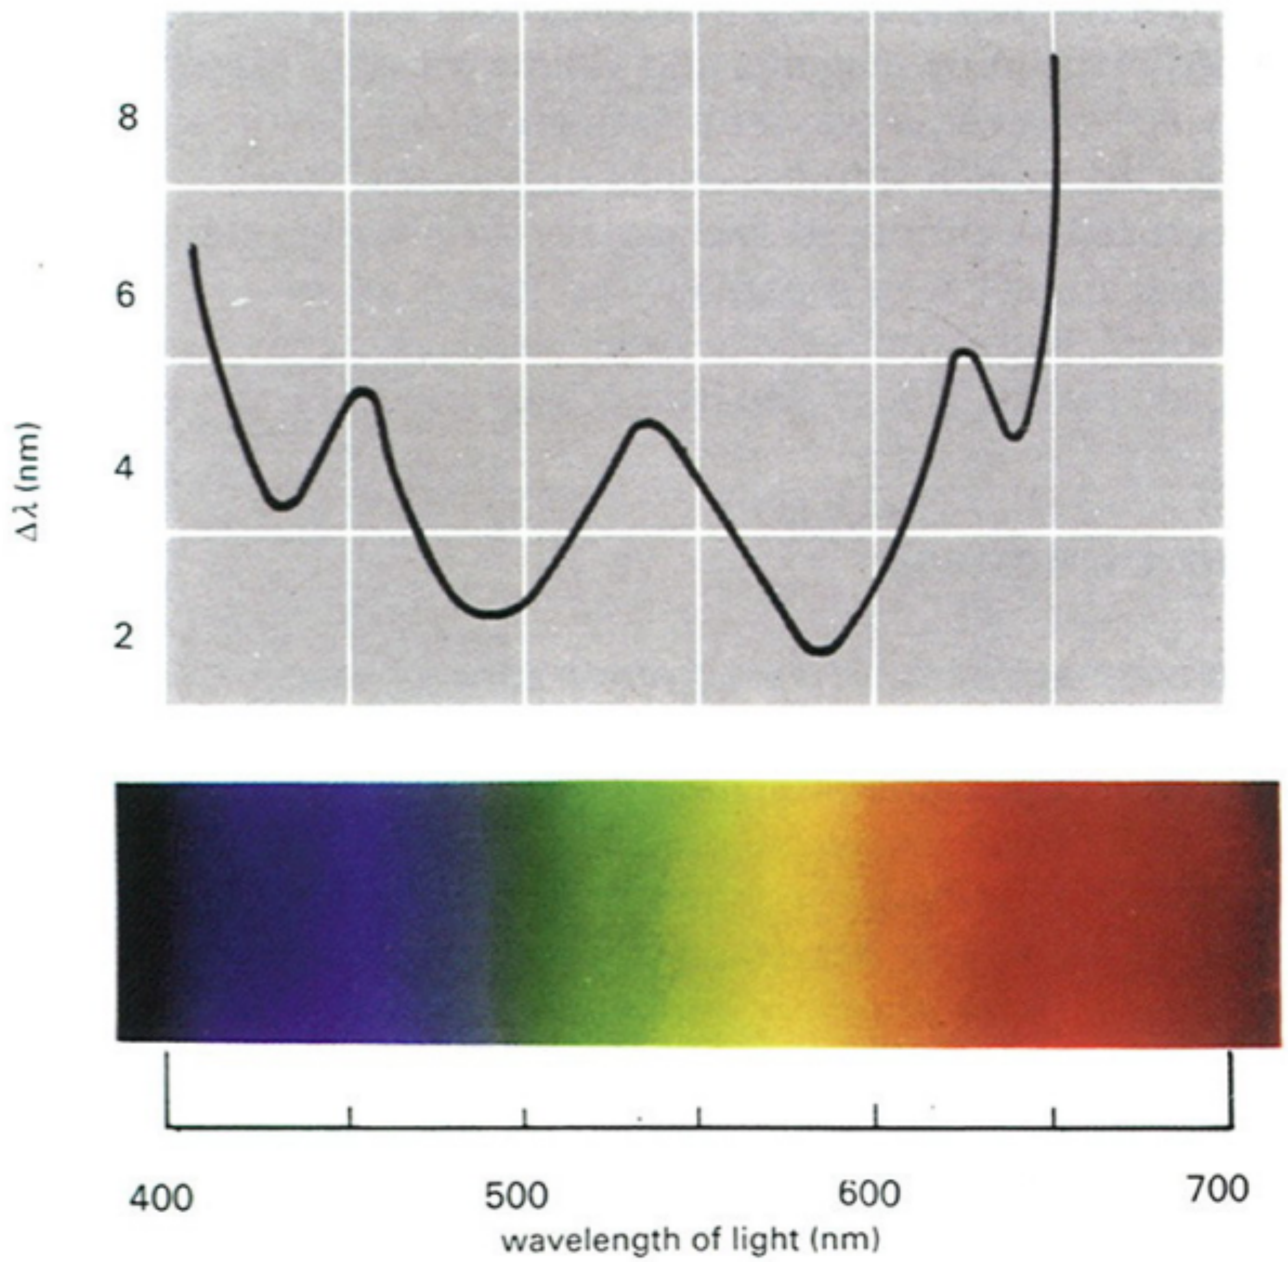

Gregory compared the wavelength of light with the smallest observable difference in hue (expressed as wavelength difference)

Is there a colour palette for scientific visualisation that works?

Colour
HSL linear L rainbow palette

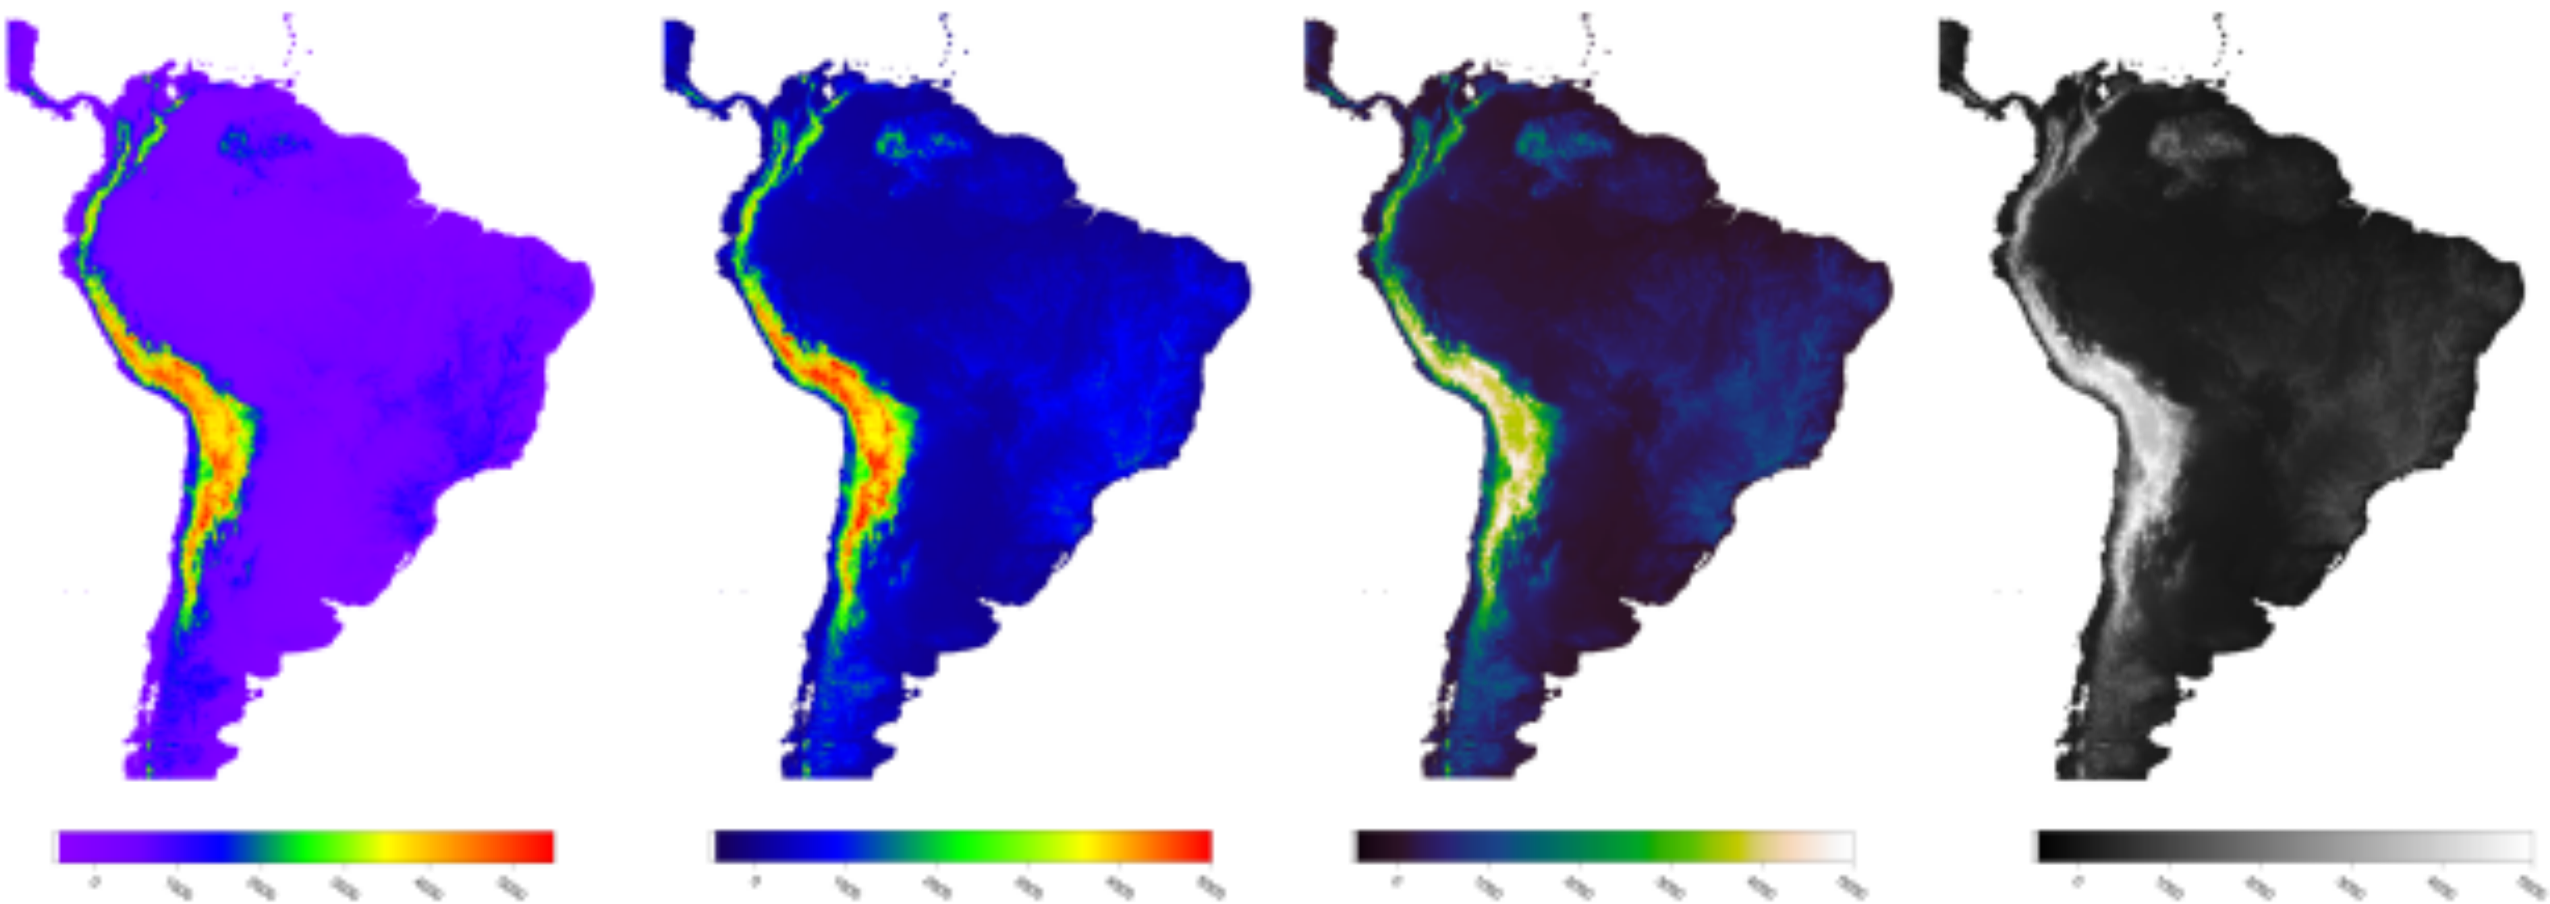

https://mycarta.wordpress.com/2012/10/06/the-rainbow-is-deadlong-live-the-rainbow-part-3/

Kindlmann, G. Reinhard, E. and Creem, S., 2002, Face-based Luminance Matching for Perceptual Colormap Generation, IEEE Proceedings of the conference on Visualization ’02

Colour

HSL linear L rainbow palette

These are available in matplotlib and therefore in seaborn, etc, so there’s no excuse :)

Colour

There are also lots of default colour maps that can be applied to particular data types.

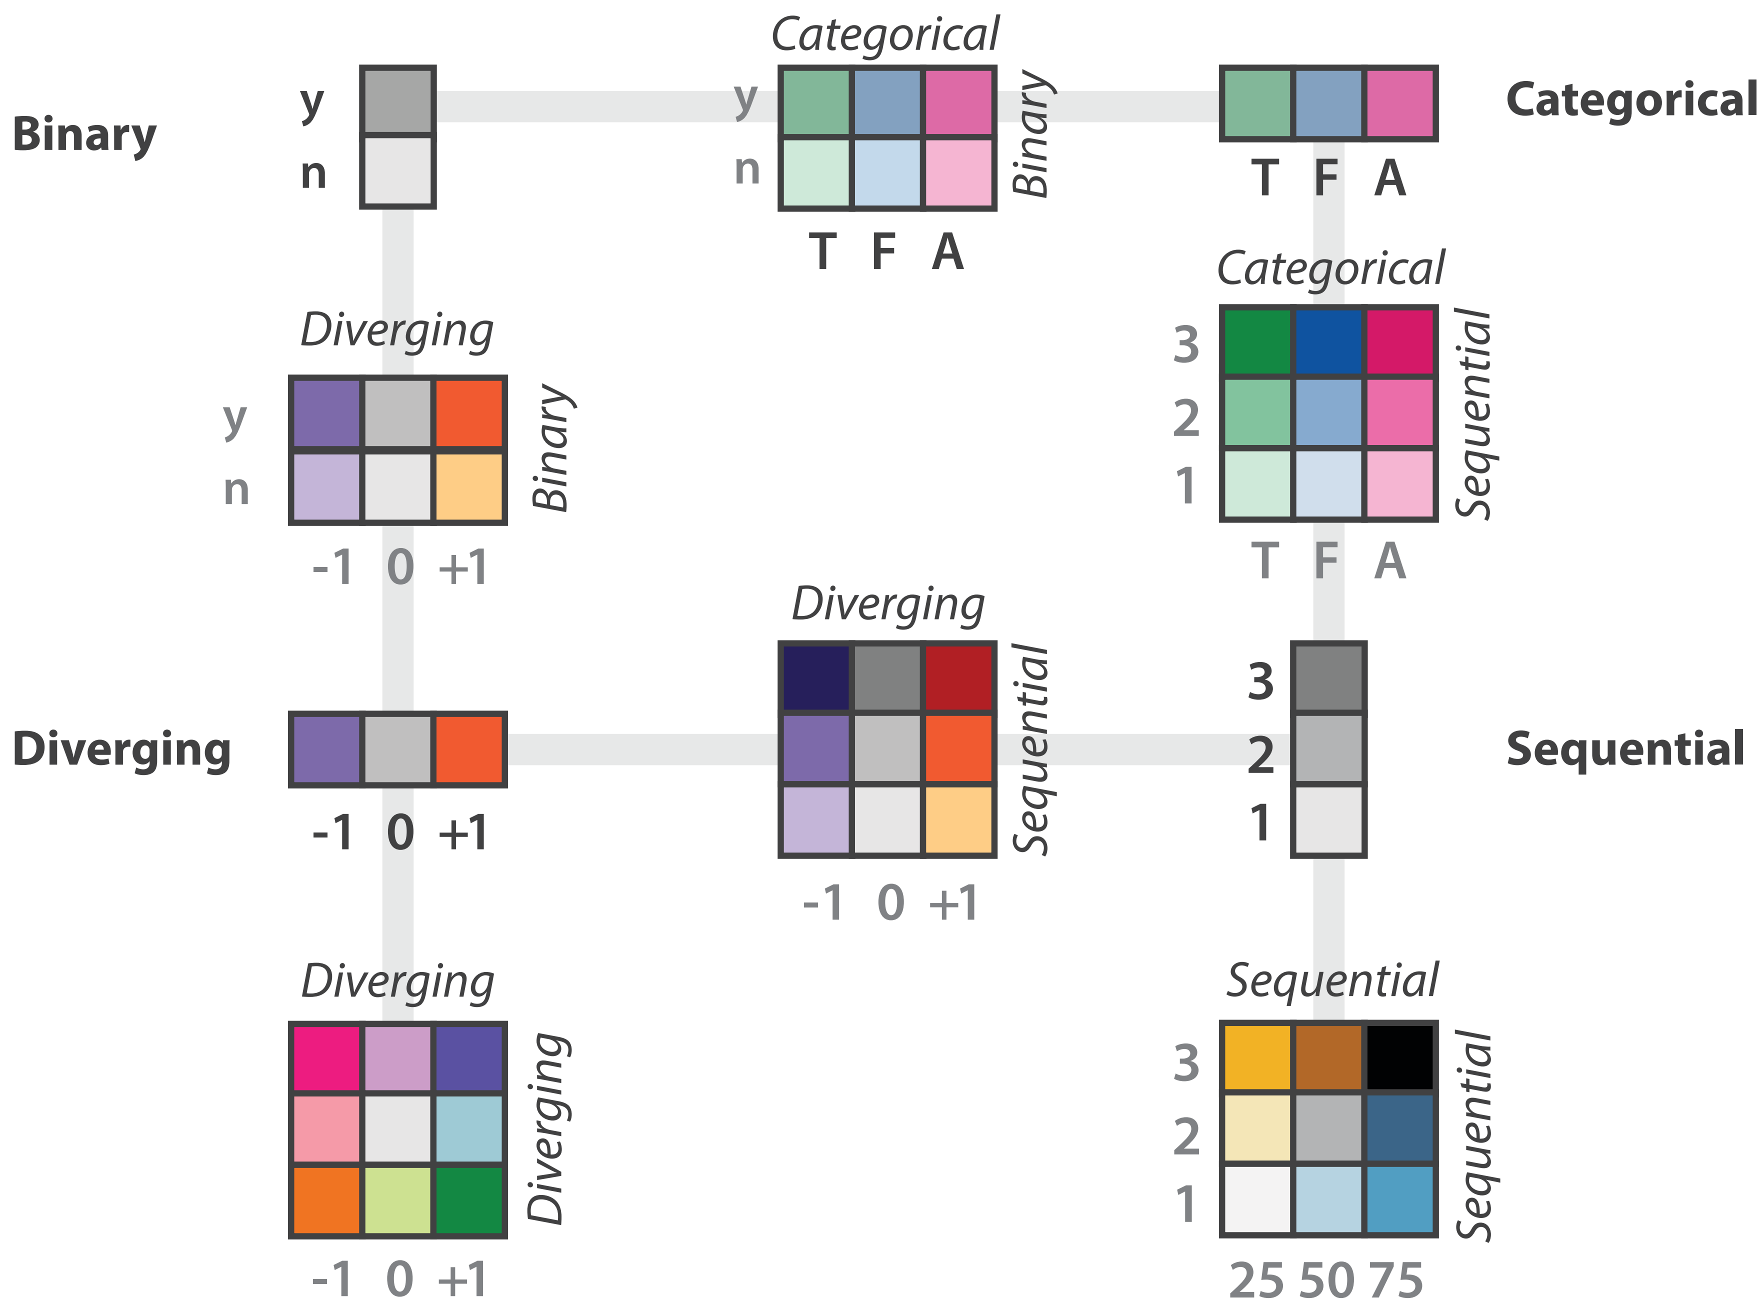

http://colorbrewer2.org/

Color

**Here I’m showing the correlation between football player attributes.
Is the choice of colour map helping this comparison?**

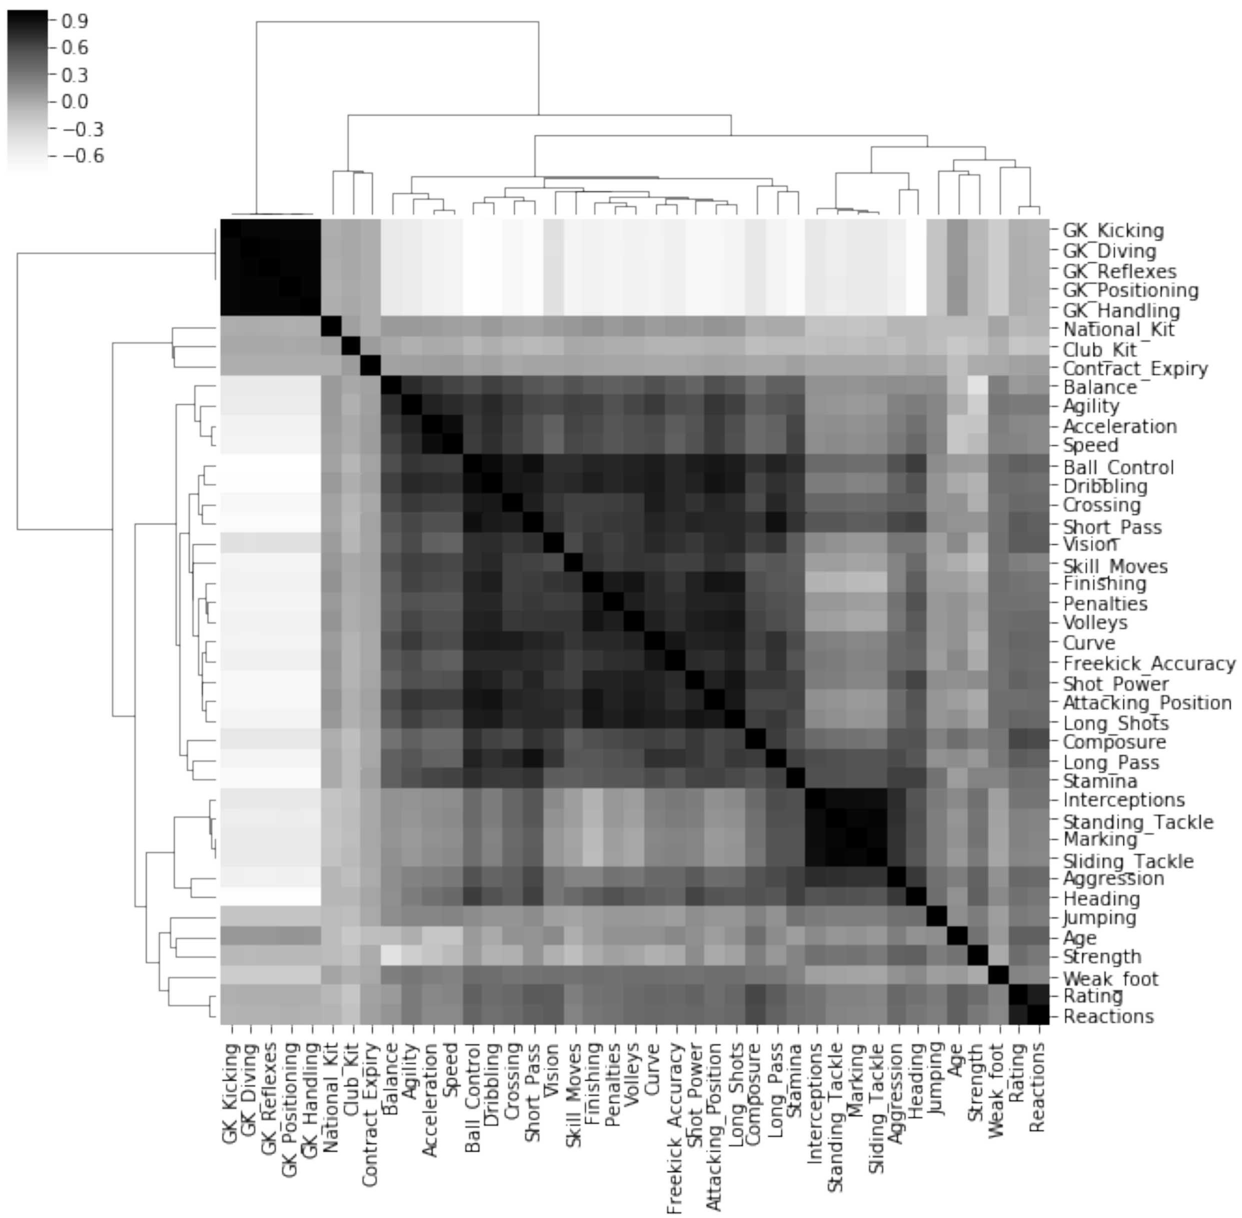

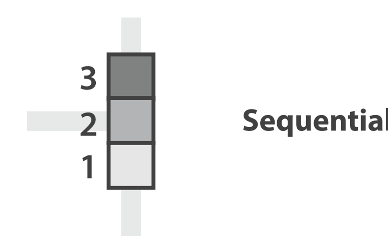


```
import seaborne as sns
sns.clustermap(fifa.corr(), cmap='Greys')
```

Color

Here I'm showing the correlation between football player attributes.
Is the choice of colour map helping this comparison?

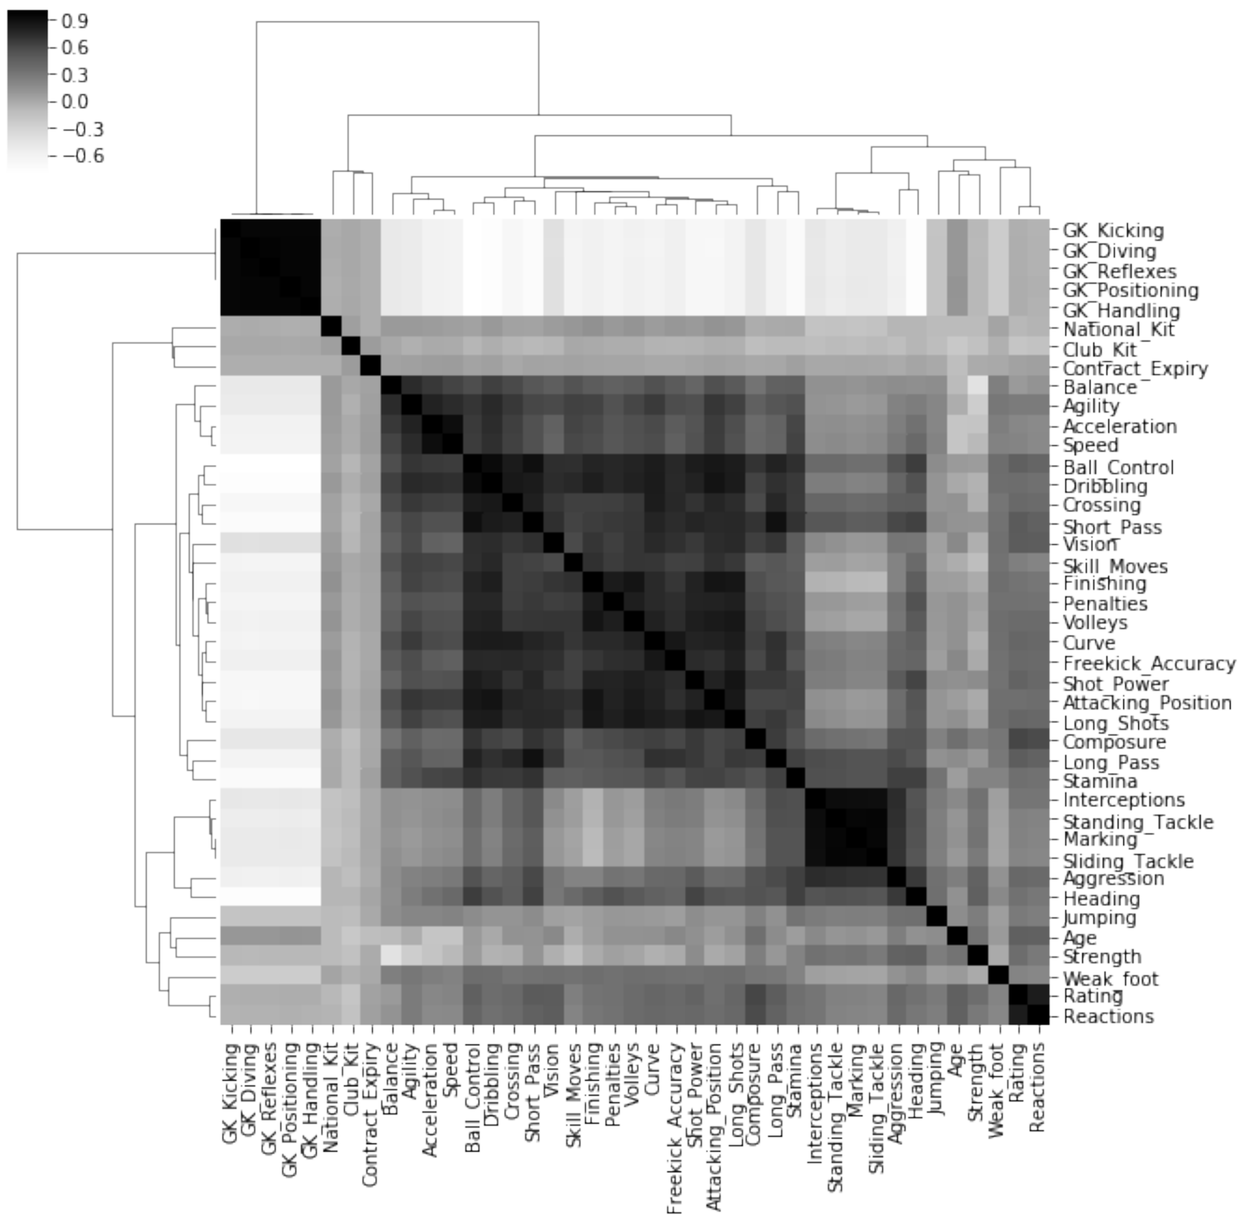

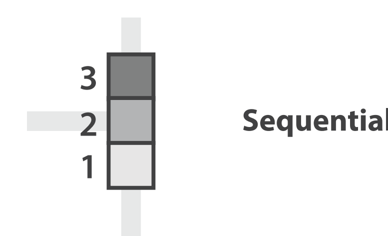

But our value pivots around 0, so the scale should be a diverging one.

```
import seaborne as sns
sns.clustermap(fifa.corr(), cmap='Greys')
```

Color

Here I'm showing the correlation between football player attributes.
Is the choice of colour map helping this comparison?

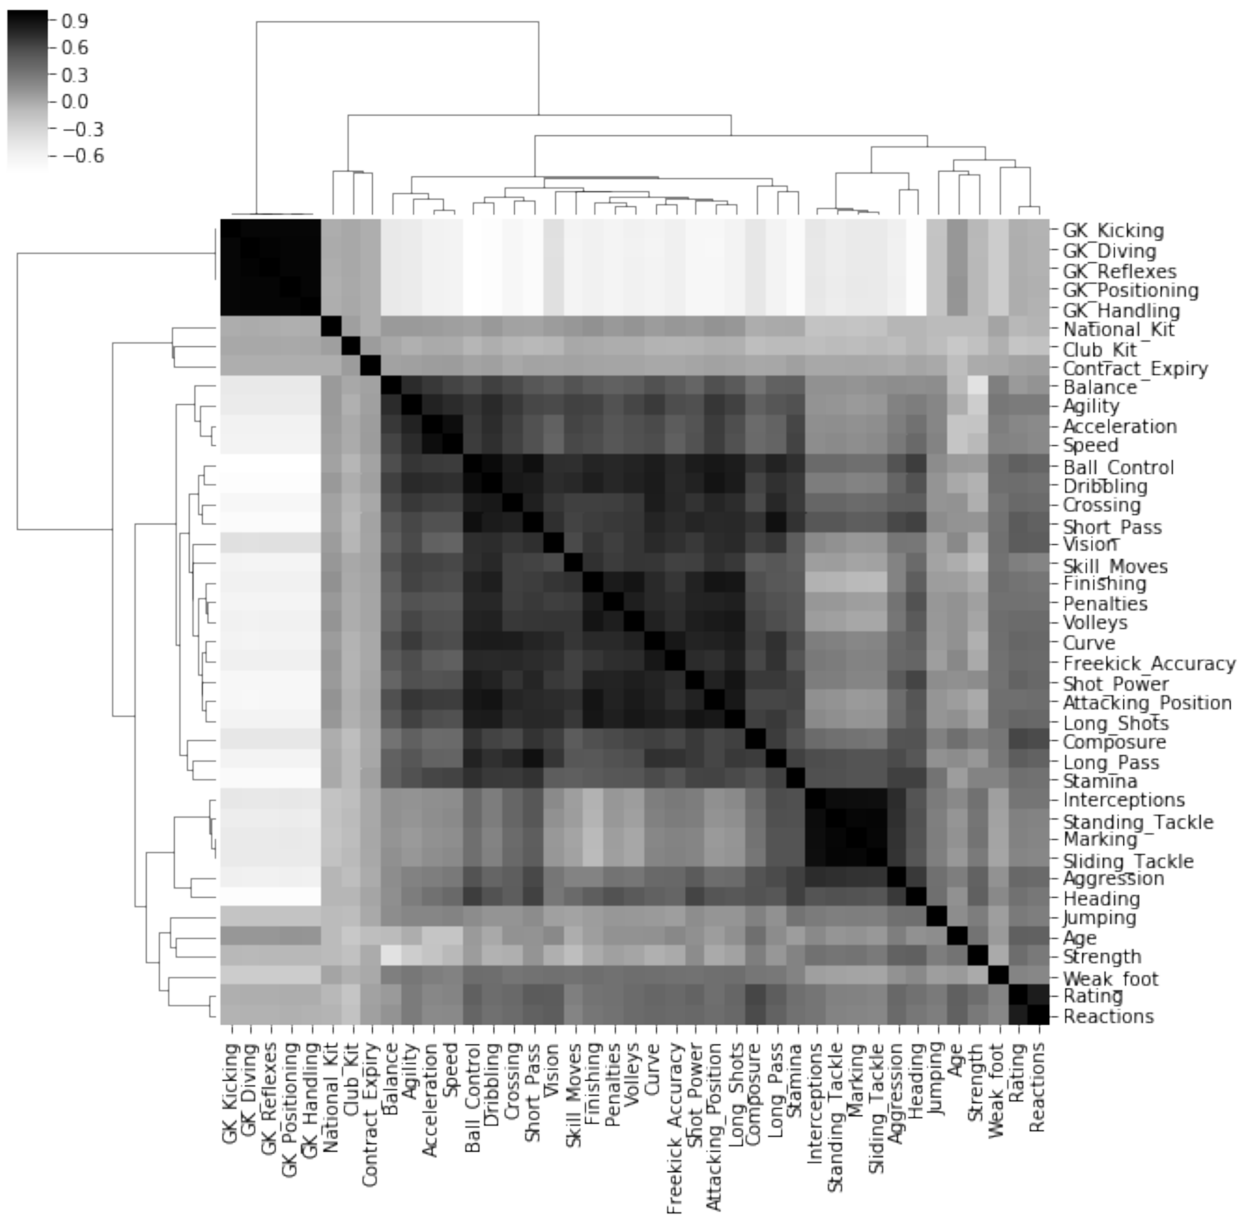

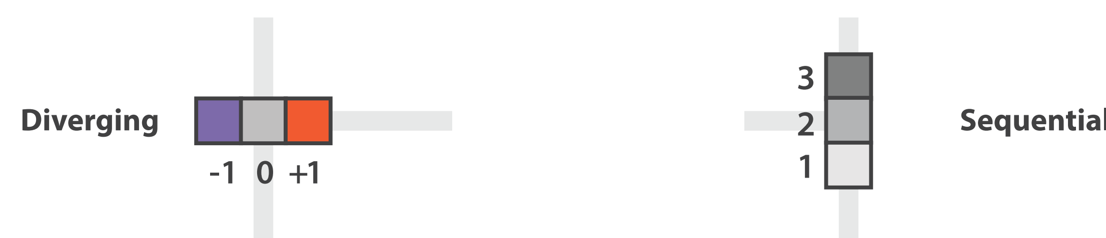

But our value pivots around 0, so the scale should be a diverging one.

```
import seaborne as sns
sns.clustermap(fifa.corr(), cmap='Greys')
```

Color

Here I'm showing the correlation between football player attributes.
Is the choice of colour map helping this comparison?

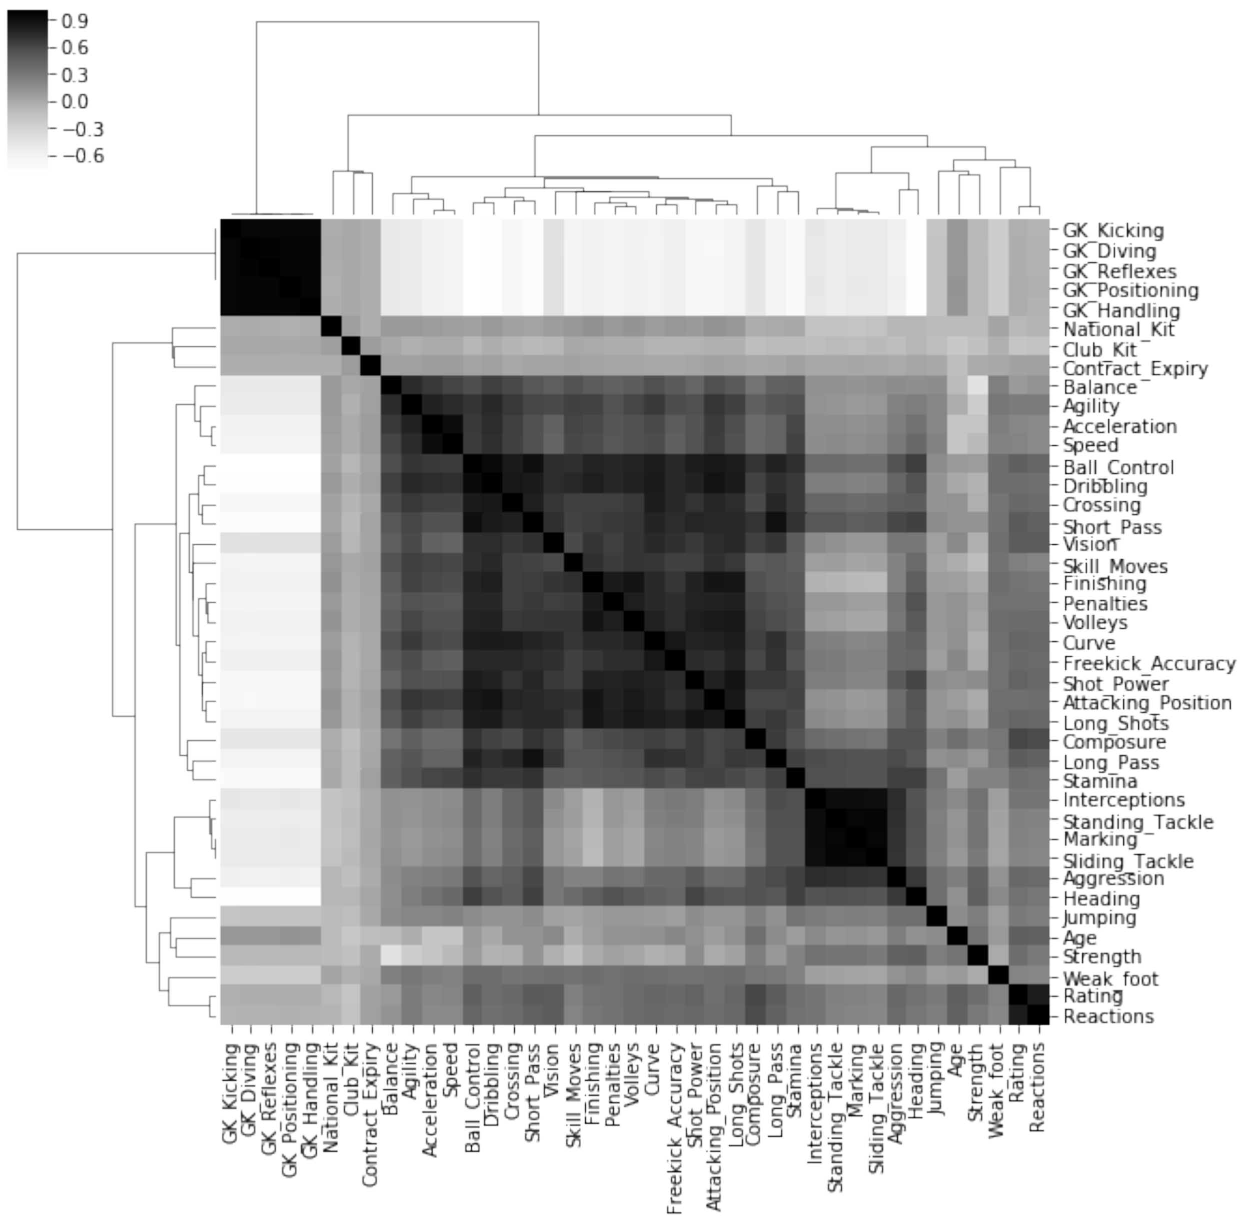

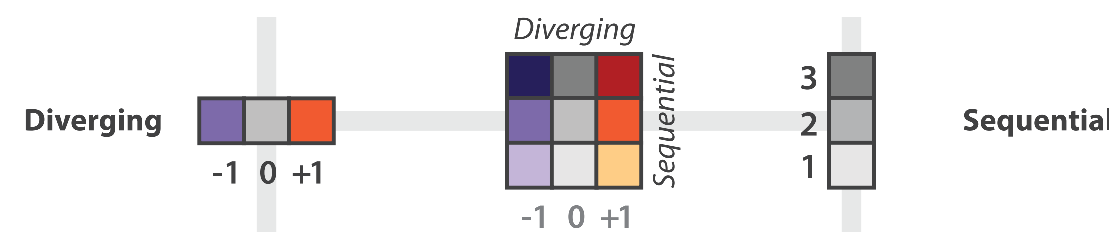

But our value pivots around 0, so the scale should be a diverging one.

```
import seaborne as sns
sns.clustermap(fifa.corr(), cmap='Greys')
```

Color

Here I'm showing the correlation between football player attributes. Is the choice of colour map helping this comparison?

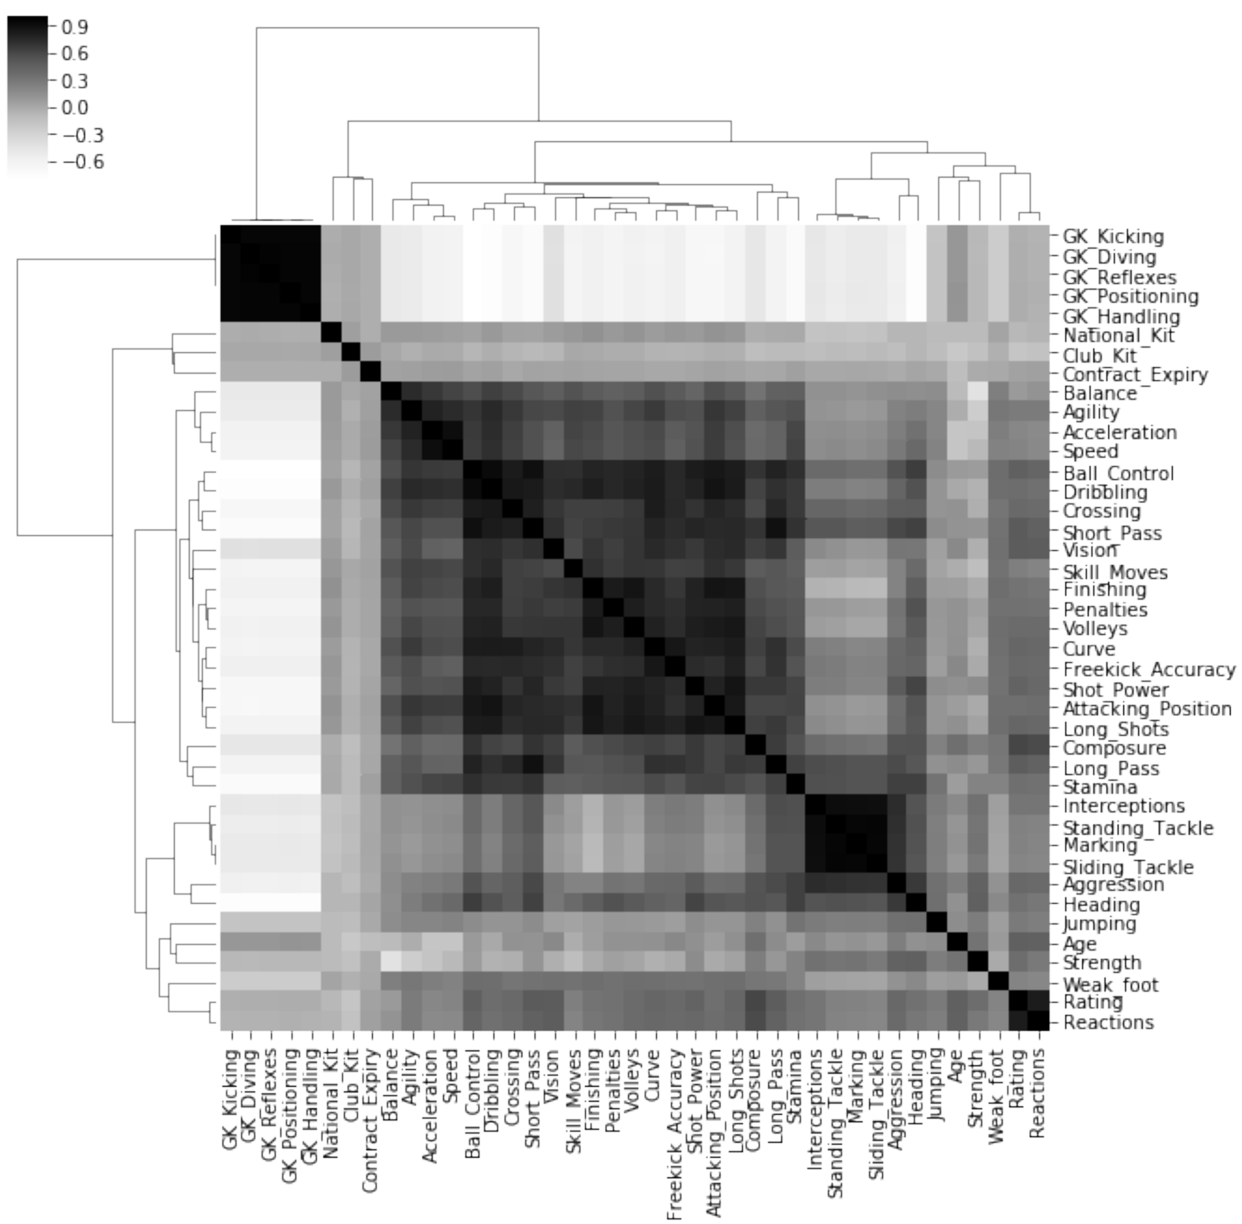

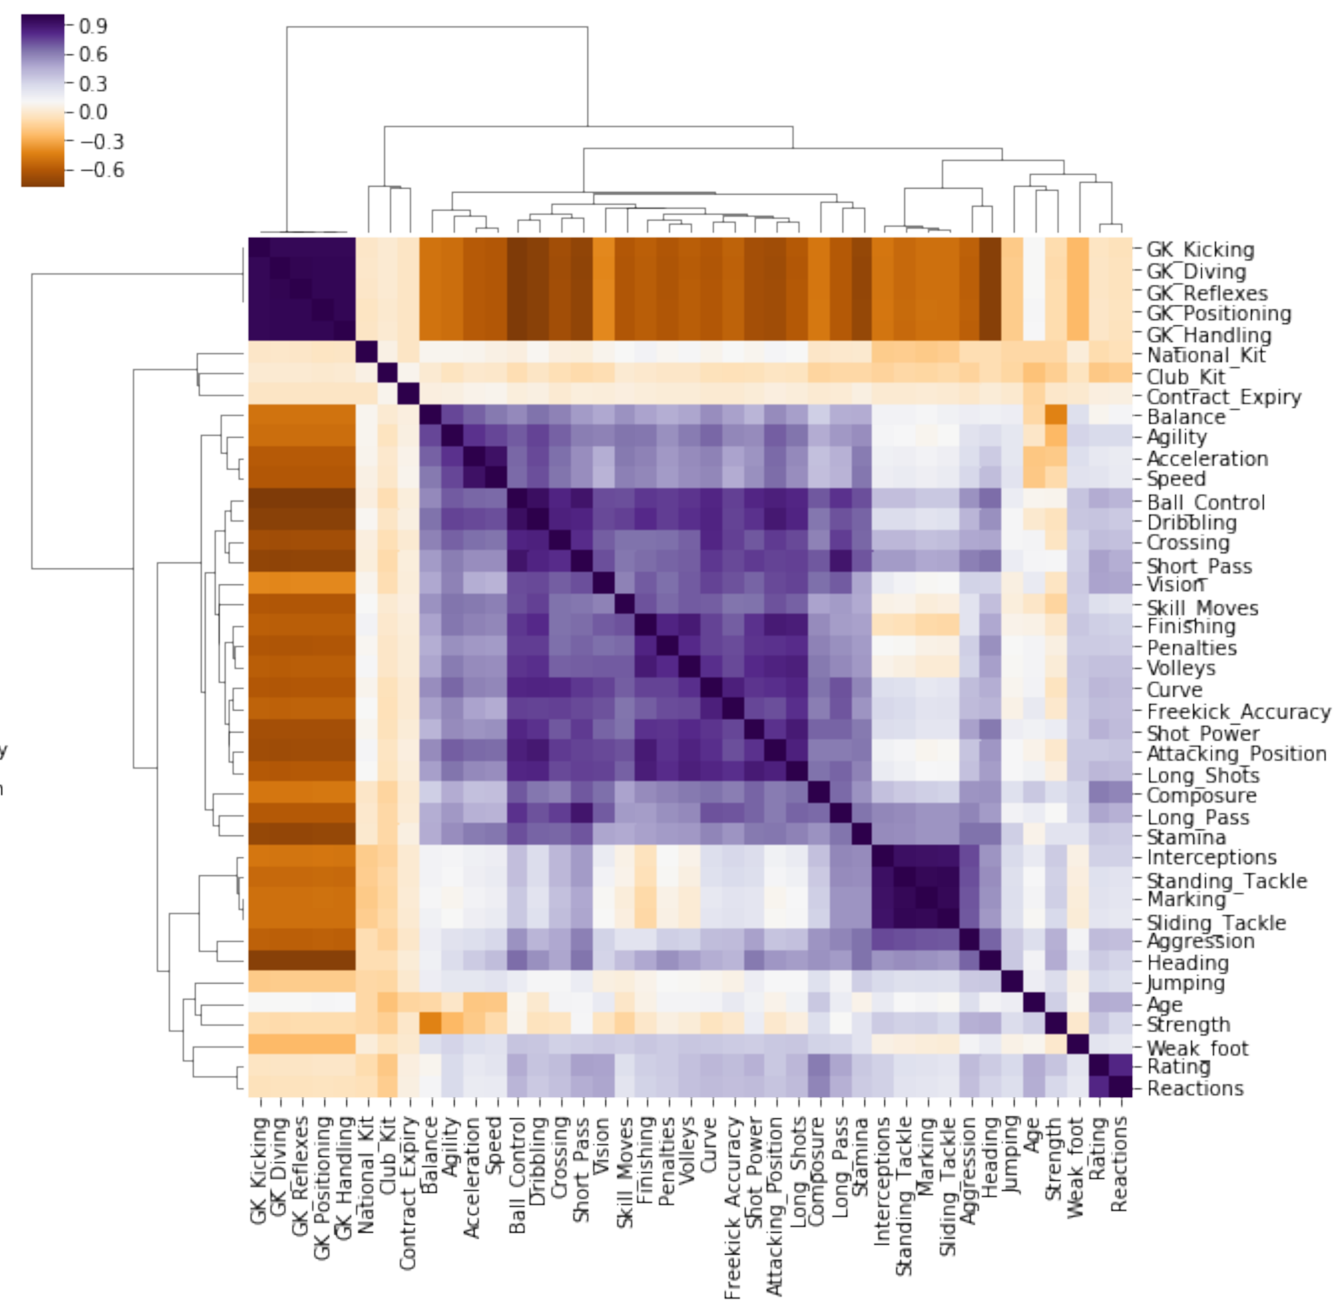


```
import seaborn  as sns
sns.clustermap(fifa.corr(), cmap='PuOr')
```

Color

You also don't want to have too many colours.
Too many colours means that users have to remember what a colour means. So a max of around 8 categories in a plot is recommended, otherwise the 'distance' between colours becomes too small.

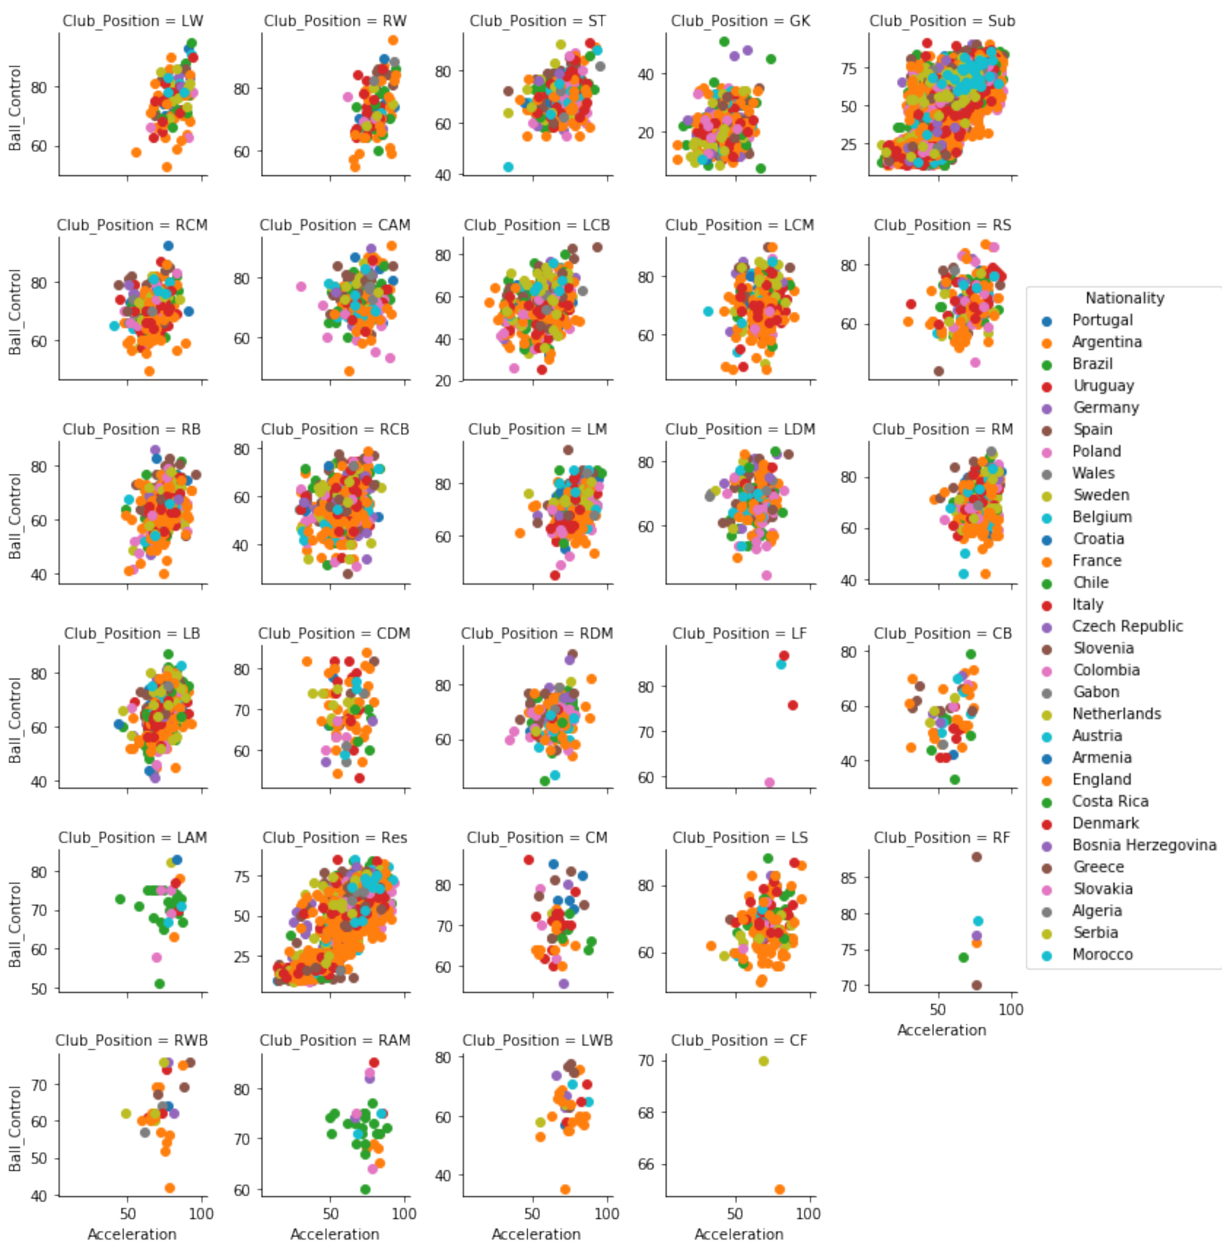

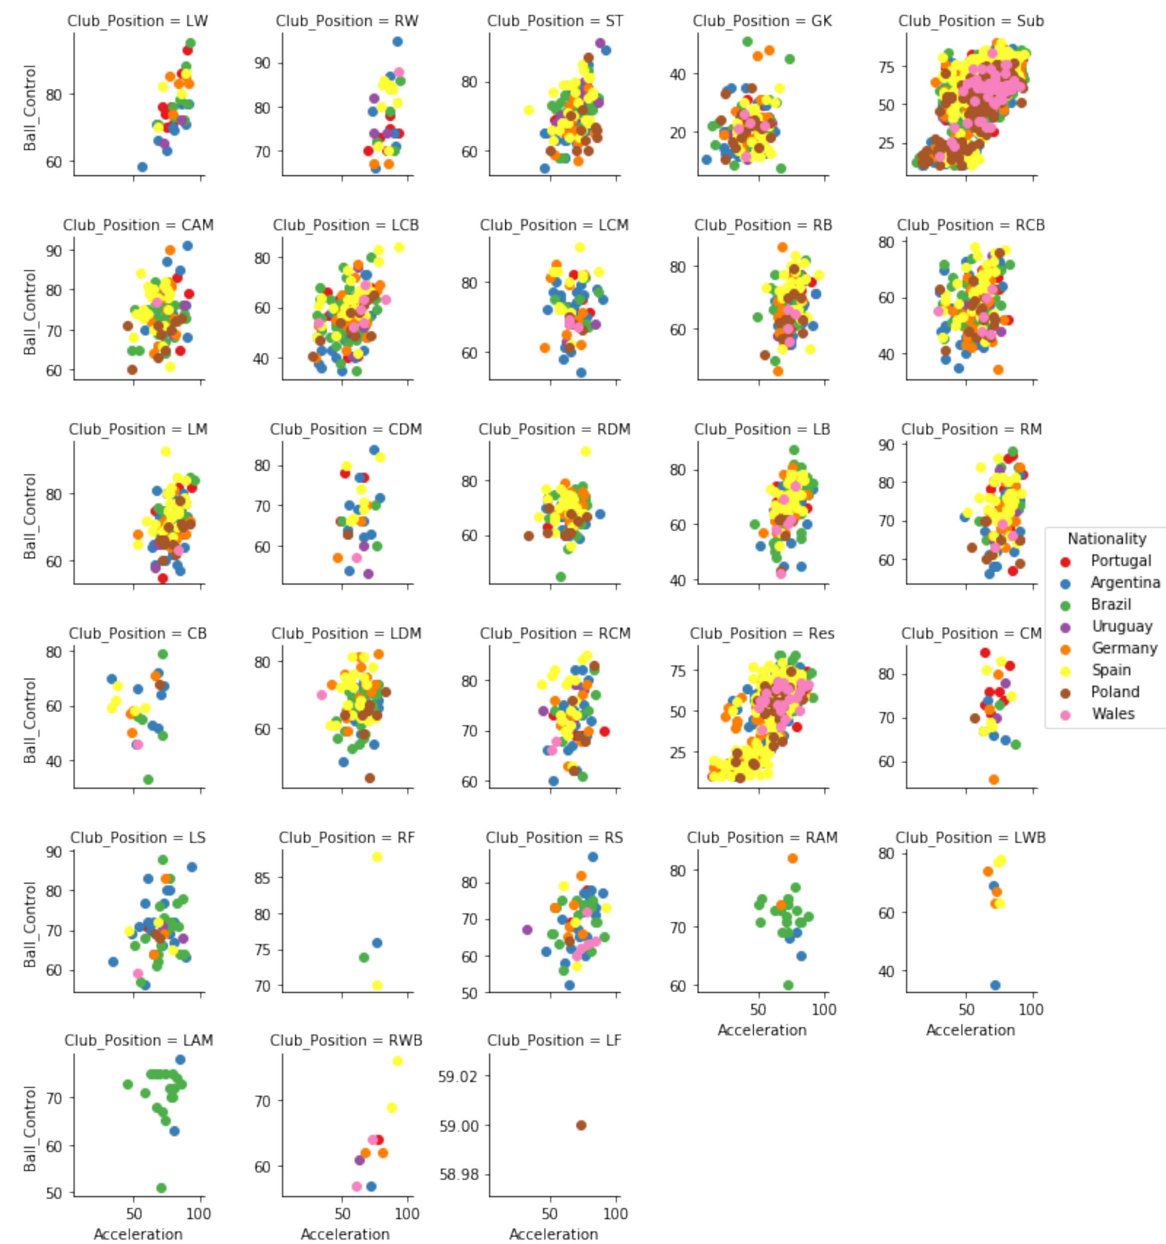

Semantic relevance

Or just consistency

When there are many colours for example, we find it difficult to remember abstract associations.

Color

What are semantically resonant colours?

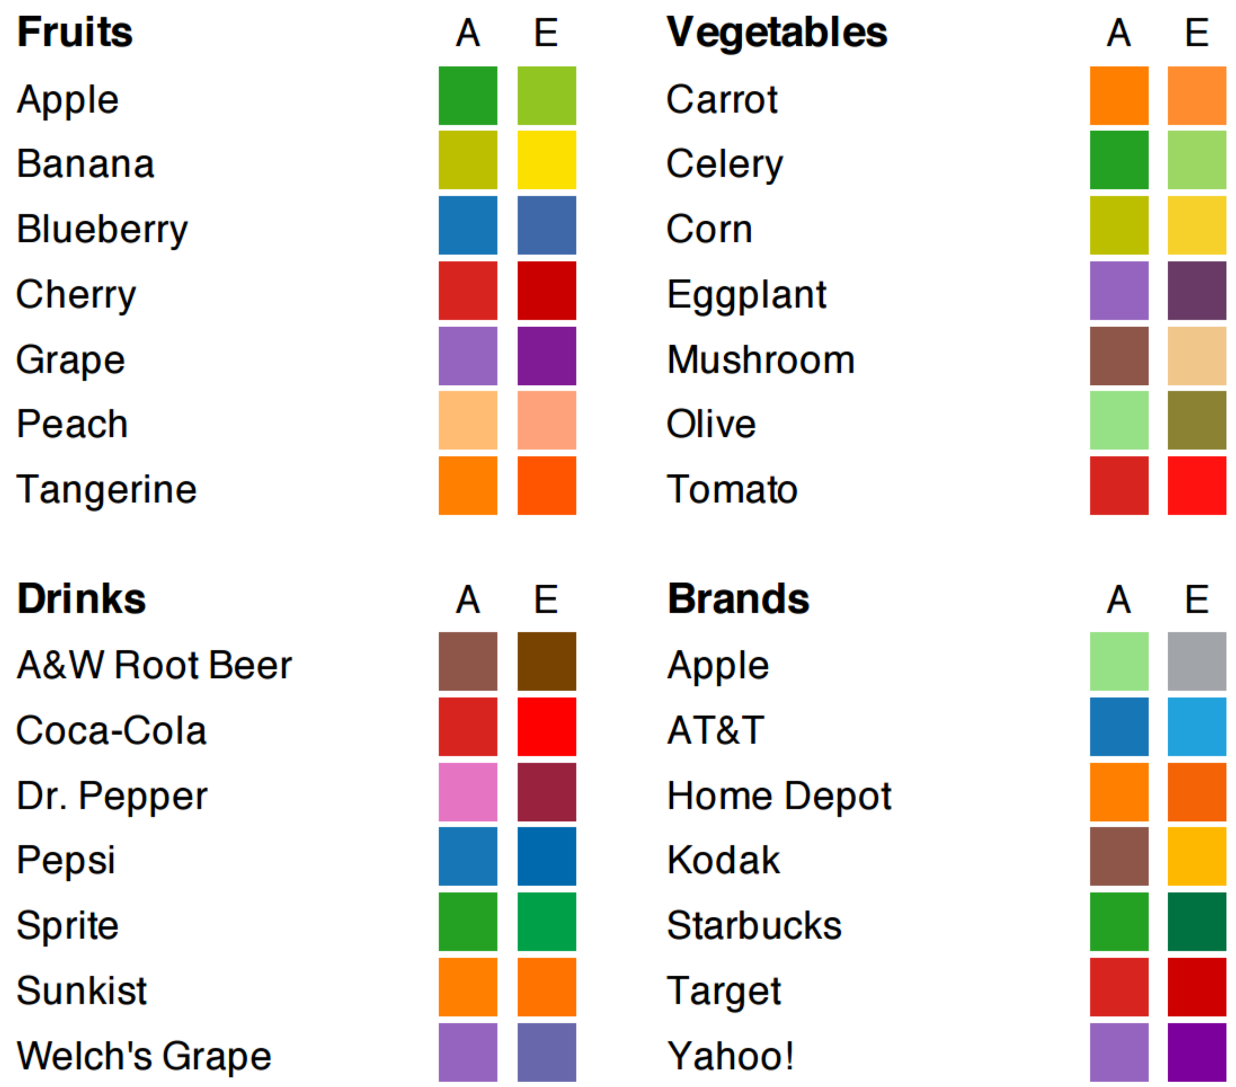

Figure 6: *Color assignments for categorical values in Experiment 1. (A = Algorithm, E = Expert)*

Selecting Semantically-Resonant Colors for Data Visualization
Sharon Lin, Julie Fortuna, Chinmay Kulkarni, Maureen Stone, Jeffrey Heer
Computer Graphics Forum (Proc. EuroVis), 2013

Color

What are semantically resonant colours?

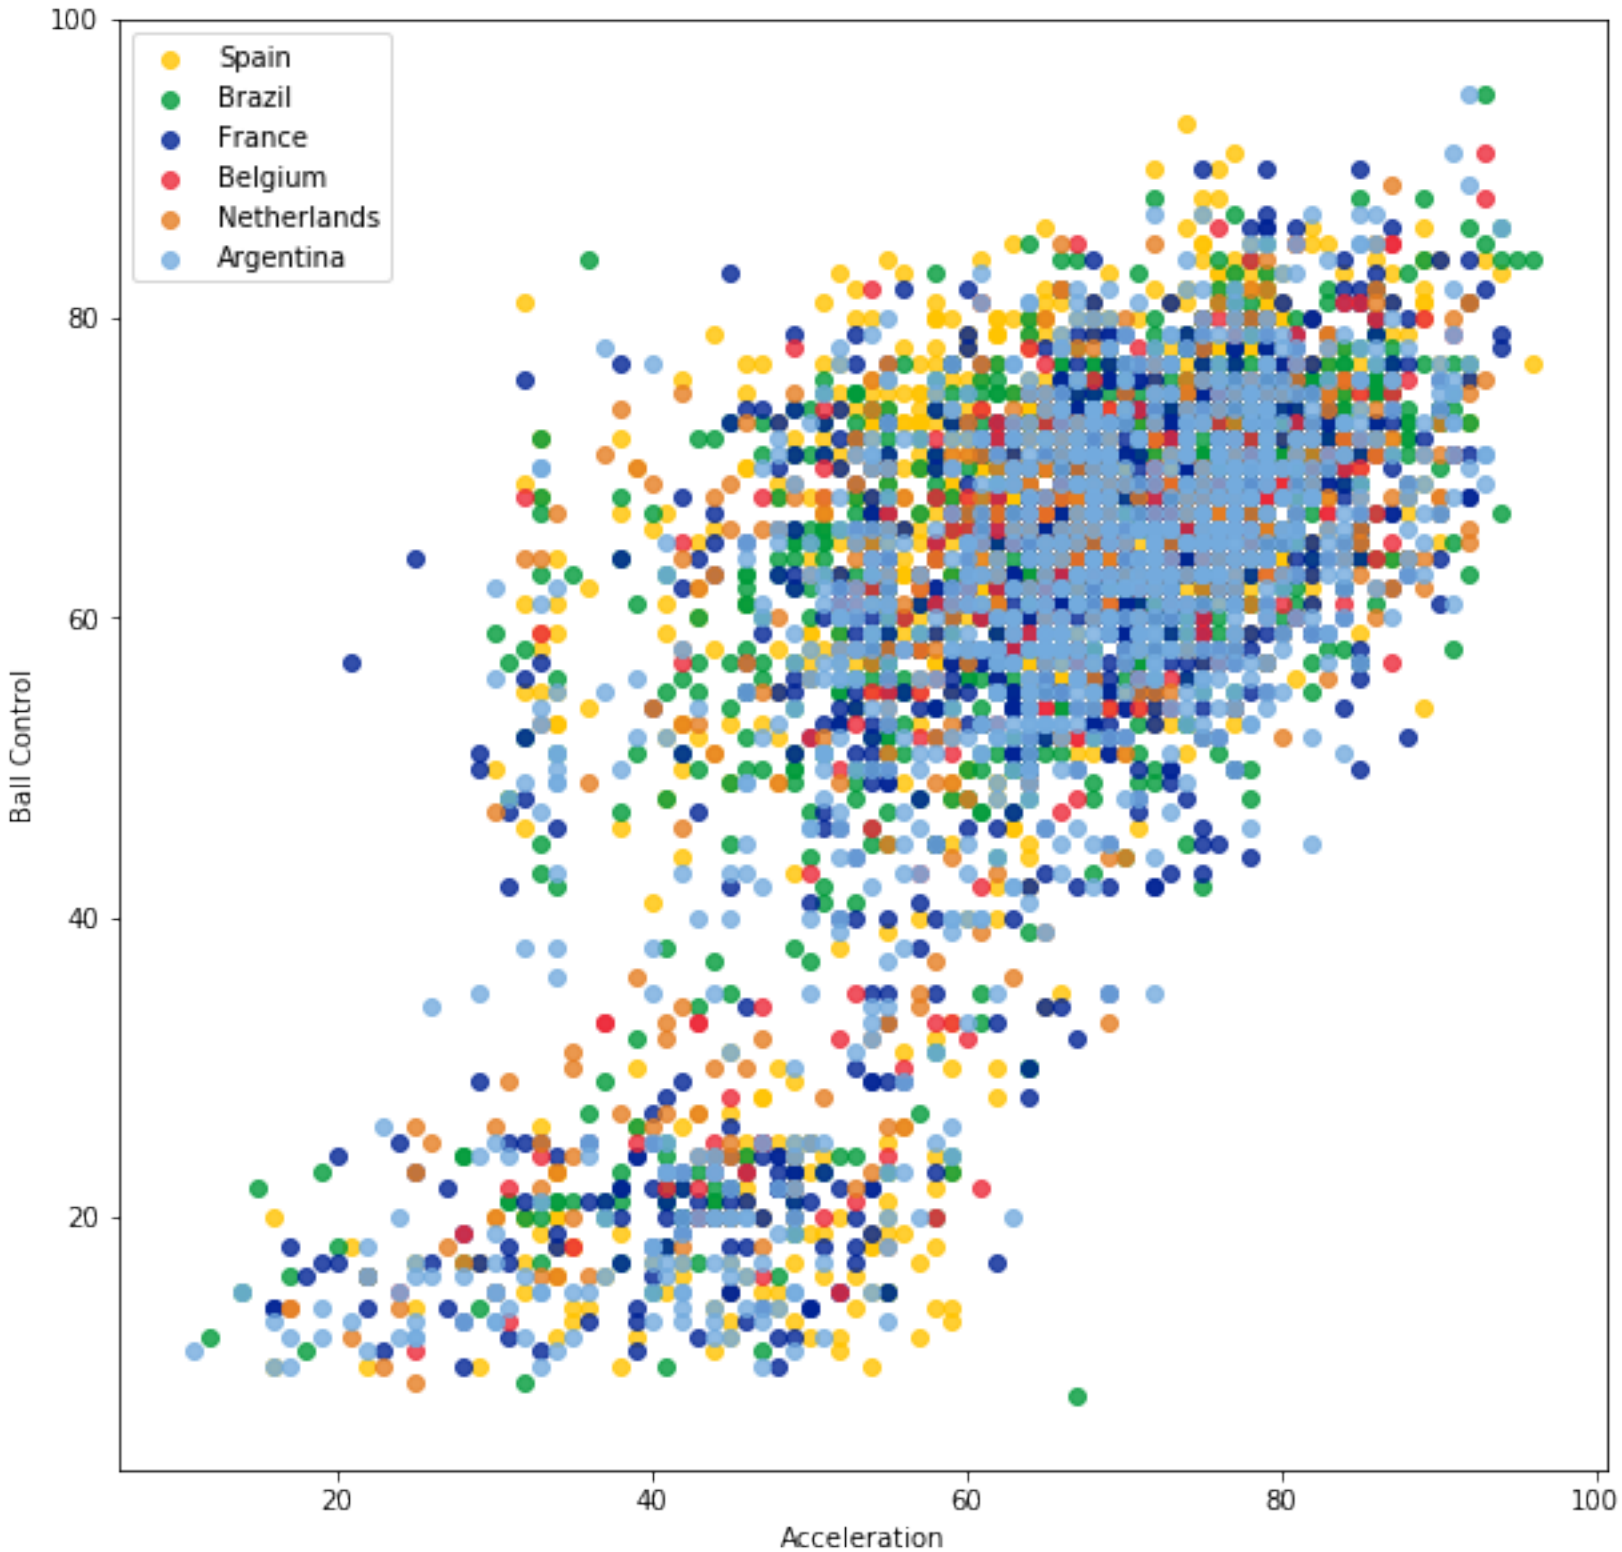

Semantic colouring is a good idea in theory, but there are limited areas where this really works.

But, if you are going to use colour, try to think how you can make it easier for users to decode the colour to the category without constantly having to look up a legend. That way, the decoding time is less.

Saving time…reducing cognitive load.