

# CS 260: Foundations of Data Science

Prof. Sara Mathieson

Fall 2021



- **Note-taker:** Trinity
- **Lab 6** due Thursday
  - Reminder to start early!
- Office hours TODAY: **3:30-5pm in H204**



# Outline for October 26

- Recap logistic regression and Handout 14
- Midterm 1: runtime (Part 5)
- Convert continuous features to discrete
- Introduction to visualization
- Begin: dimensionality reduction

# Outline for October 26


- Recap logistic regression and Handout 14
- Midterm 1: runtime (Part 5)
- Convert continuous features to discrete
- Introduction to visualization
- Begin: dimensionality reduction

# 3 important pieces to SGD

- Hypothesis function (prediction)

$$h_w(\mathbf{x}) = p(y = 1|\mathbf{x}) = \frac{1}{1 + e^{-w \cdot \mathbf{x}}}$$

# Logistic (sigmoid) function



# 3 important pieces to SGD

- Hypothesis function (prediction)

$$h_{\mathbf{w}}(\mathbf{x}) = p(y = 1|\mathbf{x}) = \frac{1}{1 + e^{-\mathbf{w} \cdot \mathbf{x}}}$$

- Cost function (want to minimize)

$$J(\mathbf{w}) = - \sum_{i=1}^n y_i \log h_{\mathbf{w}}(\mathbf{x}_i) + (1 - y_i) \log(1 - h_{\mathbf{w}}(\mathbf{x}_i))$$




# True and predicted y-values as probability distributions



true  
pred  $y_i = 1$

$$h_{\vec{w}}(\vec{x}_i) = \hat{y}_i = 0.6$$

$$1 - \hat{y}_i = 0.4$$



# 3 important pieces to SGD

- Hypothesis function (prediction)

$$h_{\mathbf{w}}(\mathbf{x}) = p(y = 1|\mathbf{x}) = \frac{1}{1 + e^{-\mathbf{w} \cdot \mathbf{x}}}$$

- Cost function (want to minimize)

$$J(\mathbf{w}) = - \sum_{i=1}^n y_i \log h_{\mathbf{w}}(\mathbf{x}_i) + (1 - y_i) \log(1 - h_{\mathbf{w}}(\mathbf{x}_i))$$

- Gradient of cost wrt single data point  $\mathbf{x}_i$

$$\nabla J_{\mathbf{x}_i}(\mathbf{w}) = (h_{\mathbf{w}}(\mathbf{x}_i) - y_i) \mathbf{x}_i$$

# Stochastic Gradient Descent for Logistic Regression (binary classification)

```
set  $w = 0$  vector
```

```
while cost  $J(w)$  still changing:
```

```
    shuffle data points
```

```
    for  $i = 1 \dots n$ :
```

```
         $w \leftarrow w - \alpha(\text{derivative of } J(w) \text{ wrt } x_i)$ 
```

```
    store  $J(w)$ 
```



Compare Handout 14 with a partner

# Handout 14

1. The output of logistic regression is a model that creates:
  - (a) a linear decision boundary
  - (b) a logistic decision boundary
  - (c) no decision boundary

# Handout 14

1. The output of logistic regression is a model that creates:

- (a) a linear decision boundary
- (b) a logistic decision boundary
- (c) no decision boundary

2. We use logistic regression for:

- (a) classification
- (b) regression
- (c) both

# Handout 14

1. The output of logistic regression is a model that creates:

- (a) a linear decision boundary
- (b) a logistic decision boundary
- (c) no decision boundary

2. We use logistic regression for:

- (a) classification
- (b) regression
- (c) both

3. Our hypothesis in logistic regression is:

$$h_{\mathbf{w}}(\mathbf{x}) = p(y = 1|\mathbf{x}) = \frac{1}{1 + e^{-\mathbf{w} \cdot \mathbf{x}}}$$

If  $\mathbf{w}$  is the zero vector (as it would be when starting SGD), what is the probability  $y = 1$ ?

# Handout 14

1. The output of logistic regression is a model that creates:

- (a) a linear decision boundary
- (b) a logistic decision boundary
- (c) no decision boundary

2. We use logistic regression for:

- (a) classification
- (b) regression
- (c) both

3. Our hypothesis in logistic regression is:

$$h_{\mathbf{w}}(\mathbf{x}) = p(y = 1|\mathbf{x}) = \frac{1}{1 + e^{-\mathbf{w} \cdot \mathbf{x}}}$$


If  $\mathbf{w}$  is the zero vector (as it would be when starting SGD), what is the probability  $y = 1$ ?

$\frac{1}{2}$

# Handout 14

Say I train a binary logistic regression model (i.e. outcomes  $\in \{0, 1\}$ ) and end up with  $\hat{\mathbf{w}} = [\hat{w}_0, \hat{w}_1]^T = [-4, -5]^T$ . What is the decision boundary? Sketch a graph of this logistic model and label the decision boundary. How would you classify a new point  $x_{\text{test}} = -2$ ?

$$\frac{1}{(1 + e^{-(-4 - 5x)})}$$




$$\begin{aligned} -4 - 5x &> 0 \\ -4 &> +5x \\ \boxed{x < -\frac{4}{5}} \end{aligned}$$

Decision boundary

# Handout 14

5. The graph below shows the cost for logistic regression as a function of the hypothesis  $h_w(\mathbf{x})$ , for one example  $\mathbf{x}$ . Which curve corresponds to the true label  $y = 0$  and which corresponds to  $y = 1$ ?



# Outline for October 26

- Recap logistic regression and Handout 14
- Midterm 1: runtime (Part 5)  
*Not posted online*
- Convert continuous features to discrete
- Introduction to visualization
- Begin: dimensionality reduction



# Outline for October 26

- Recap logistic regression and Handout 14
- Midterm 1: runtime (Part 5)
- **Convert continuous features to discrete**
- Introduction to visualization
- Begin: dimensionality reduction

# Continuous Features

(do this for the TRAIN only!)

1) Sort examples based on given feature

X	Y
10	Y
7	Y
8	N
3	Y
7	N
12	Y
2	Y

2	3	7	7	8	10	12
Y	Y	Y	N	N	Y	Y

# Continuous Features

(do this for the TRAIN only!)

X	Y
10	Y
7	Y
8	N
3	Y
7	N
12	Y
2	Y

1) Sort examples based on given feature

2	3	7	7	8	10	12
Y	Y	Y	N	N	Y	Y

2) Different label with same feature value, collapse to “None”

2	3	7	8	10	12
Y	Y	None	N	Y	Y

# Continuous Features

(do this for the TRAIN only!)

X	Y
10	Y
7	Y
8	N
3	Y
7	N
12	Y
2	Y


- 1) Sort examples based on given feature

2	3	7	7	8	10	12
Y	Y	Y	N	N	Y	Y

- 2) Different label with same feature value, collapse to "None"

2	3	7	8	10	12
Y	Y	None	N	Y	Y

- 3) Whenever label changes, make a feature (use avg)



.....  $x \geq 7.5$

$\Rightarrow O(n)$

		feature	label
		x	y
		<del>10</del>	Y
		<del>7</del>	Y
		<del>8</del>	N
		<del>3</del>	Y
		<del>7</del>	N
		<del>12</del>	Y
		<del>2</del>	Y

① sort

2	3	7	7	8	10	12
Y	Y	Y	N	N	Y	Y

② collapse features with more than one label

2	3	7	8	10	12
Y	Y	None	N	Y	Y

③

$x \geq 5$

$x \geq 7.5$

$x \geq 9$

change in label  $\Rightarrow$  make a feature!  
avg feature values

# Continuous Features (Handout 15)

(do this for the TRAIN only!)

temp	Y
80	Y
48	Y
60	N
48	N
40	N
48	Y
90	Y

- 1) Sort examples based on feature “temp”
- 2) Different label with same feature value, collapse to “None”
- 3) Whenever label changes, make a feature (use avg)



# Handout 15

$x \geq 44$

T

temp

~~80~~

label

Y

T

~~48~~

Y

T

~~60~~

N

T

~~48~~

N

F

~~40~~

N

T

~~48~~

Y

T

~~90~~

Y

①

40

N

48

Y

48

N

48

Y

60

N

80

Y

90

Y

②

40

N

48

None

60

N

80

Y

90

Y

$x \geq 44$

$x \geq 54$

$x \geq 70$

# Outline for October 26

- Recap logistic regression and Handout 14
- Midterm 1: runtime (Part 5)
- Convert continuous features to discrete
- Introduction to visualization
- Begin: dimensionality reduction



# Ugly, bad, wrong visualizations

- ugly**—A figure that has aesthetic problems but otherwise is clear and informative.
- bad**—A figure that has problems related to perception; it may be unclear, confusing, overly complicated, or deceiving.
- wrong**—A figure that has problems related to mathematics; it is objectively incorrect.

# Ugly, bad, wrong visualizations







Fig 1.1 from "Fundamentals of Data Visualization" by Claus Wilke

## Two different visualizations of the same data





Fig 2.3/2.4 from  
“Fundamentals of Data  
Visualization” by Claus Wilke



## All the same data – what do we want to convey?




Fig 3.2 from “Fundamentals of Data Visualization” by Claus Wilke


# Data Types

---

- Continuous vs. Discrete


# Data Types

## Continuous




# Data Types

## Continuous





# Data Types

## Discrete





# Overplotting



# Color


---

Why use color?

Which colors?


# Color

Individual vowel formant values  
non-normalized



# Color

## Sequential



# Sequential example

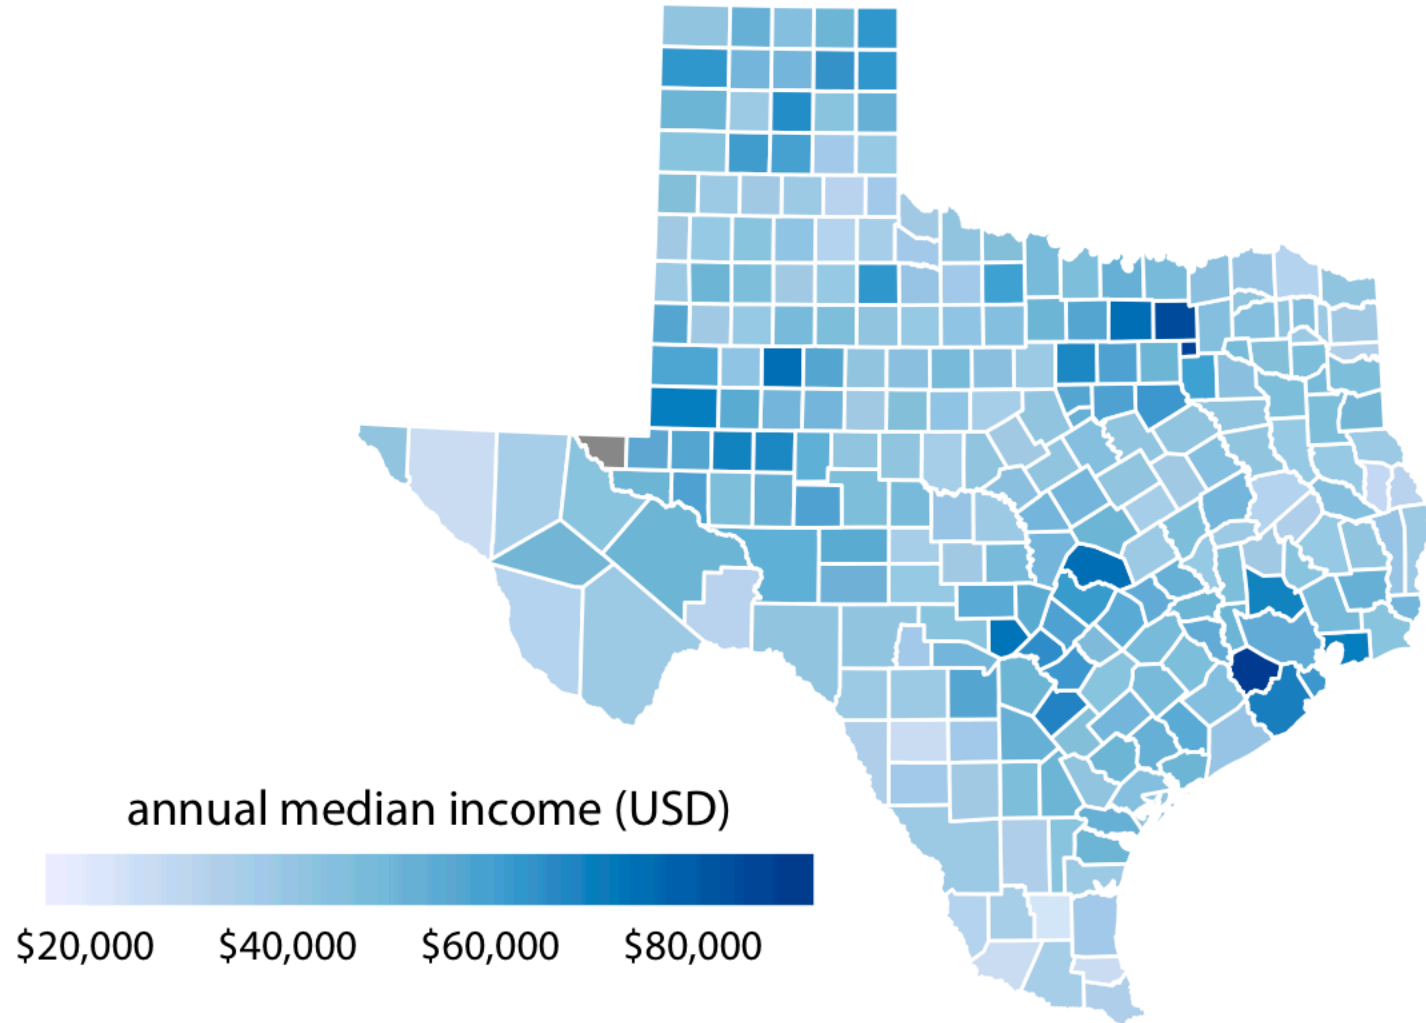



Fig 4.4 from "Fundamentals of Data Visualization" by Claus Wilke

# Color

## Diverging



# Diverging example

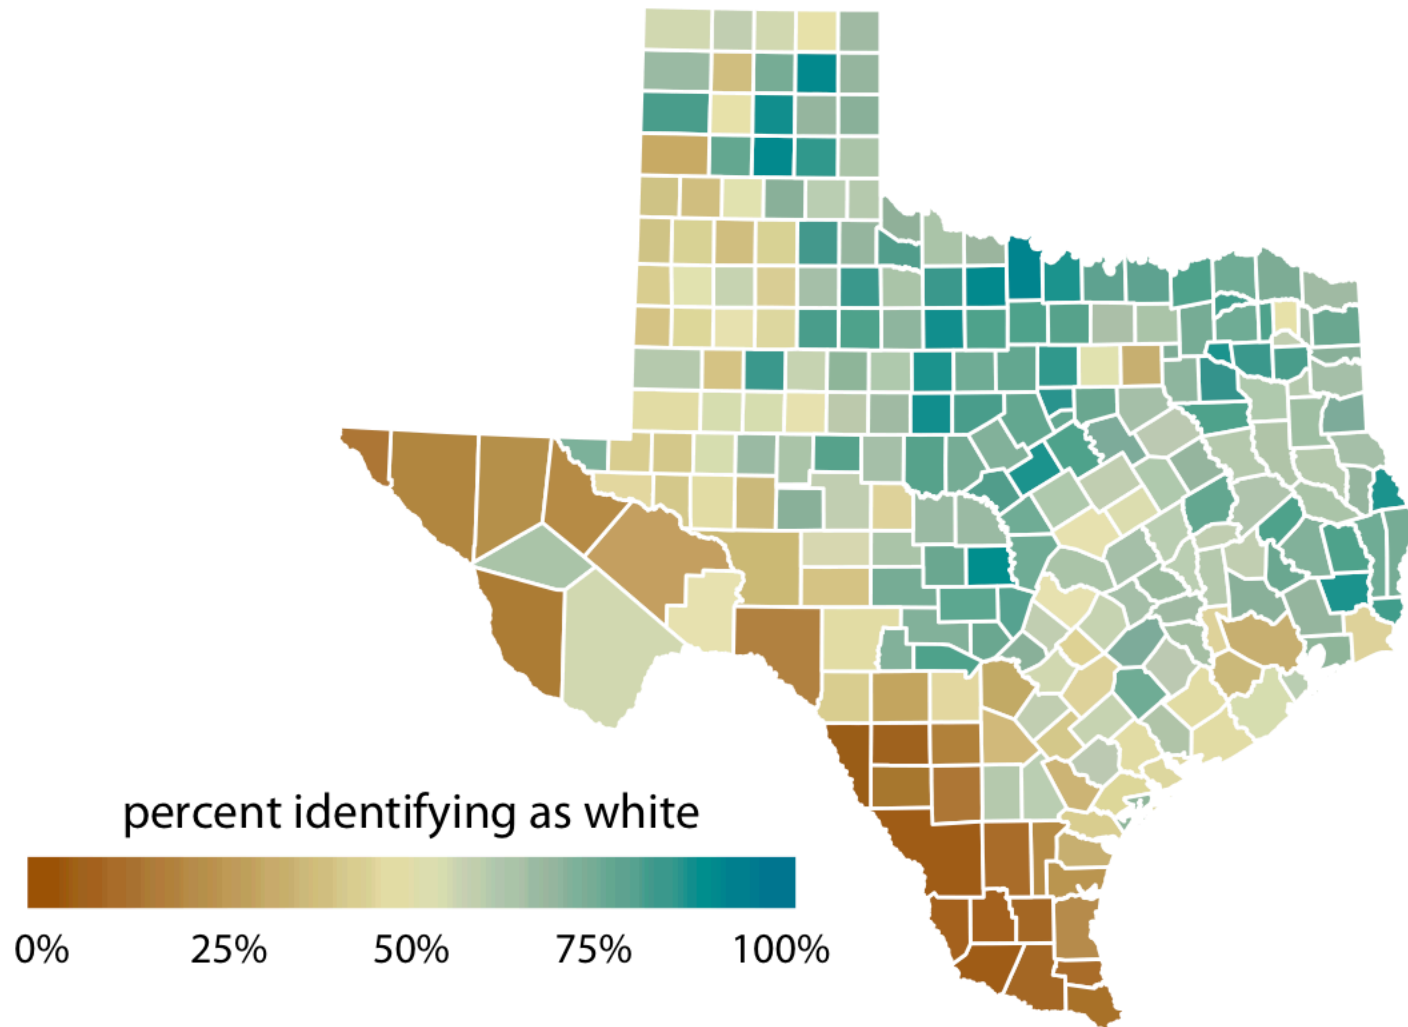



Fig 4.6 from "Fundamentals of Data Visualization" by Claus Wilke

# Color

## Qualitative





# Qualitative example







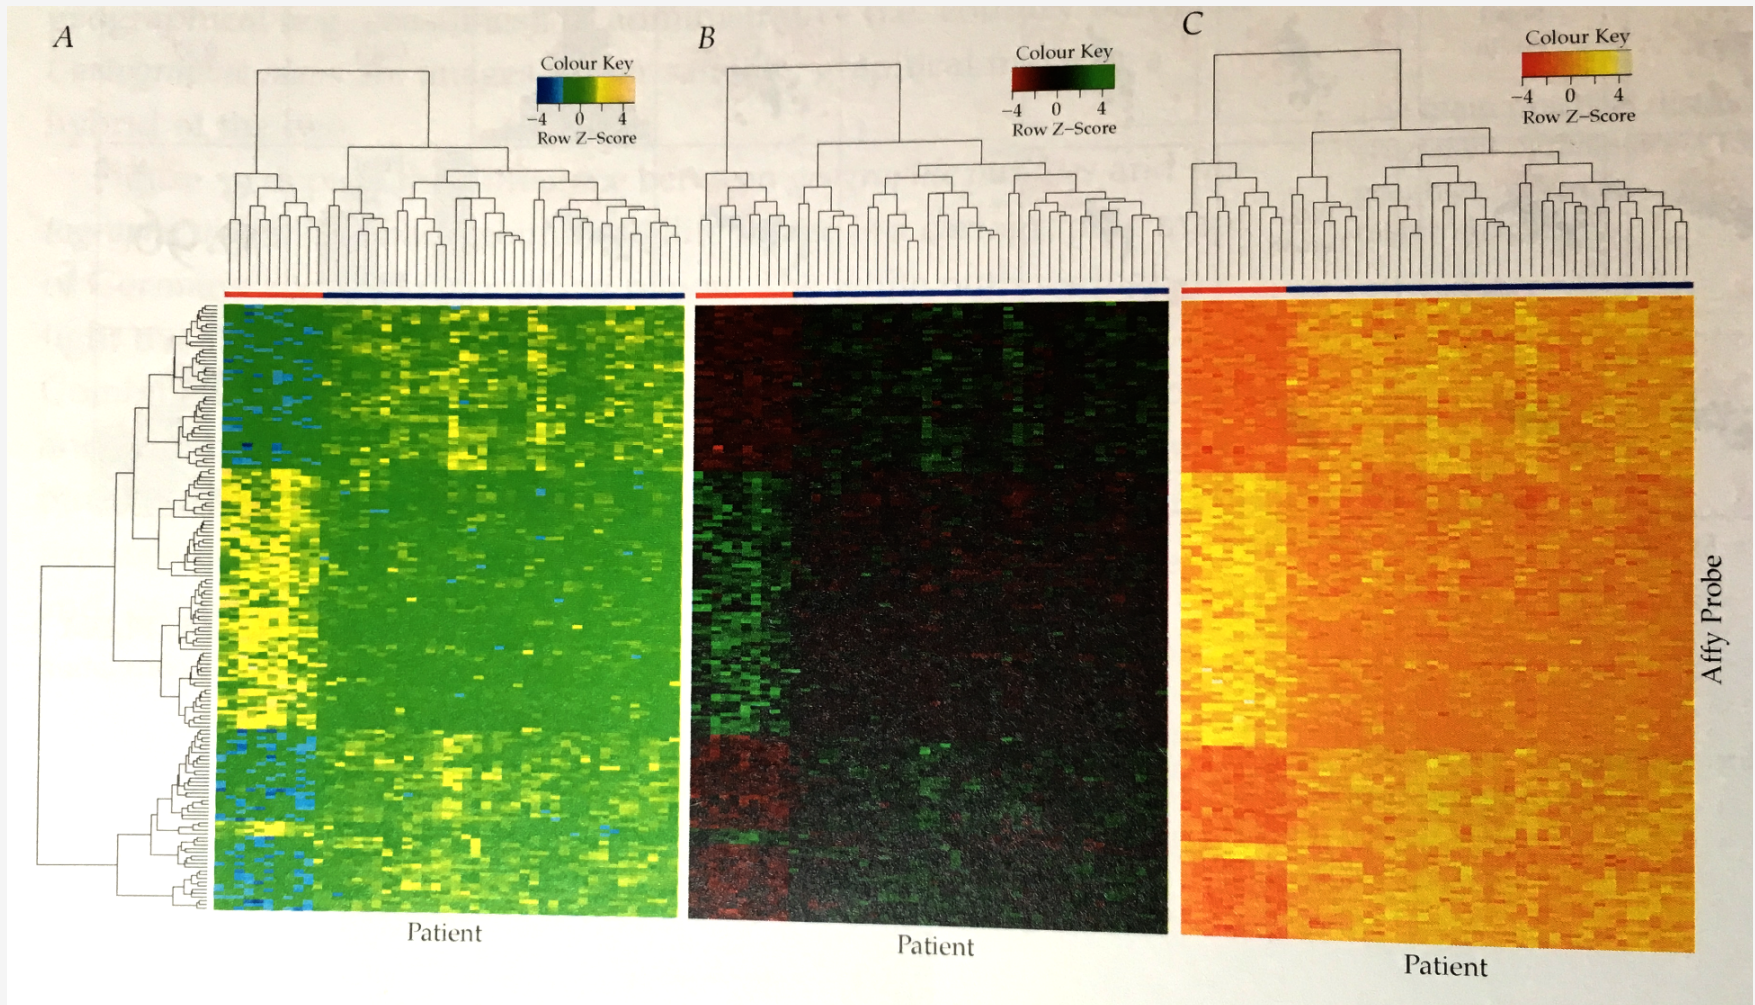
Fig 4.2 from "Fundamentals of Data Visualization"  
by Claus Wilke

# Color

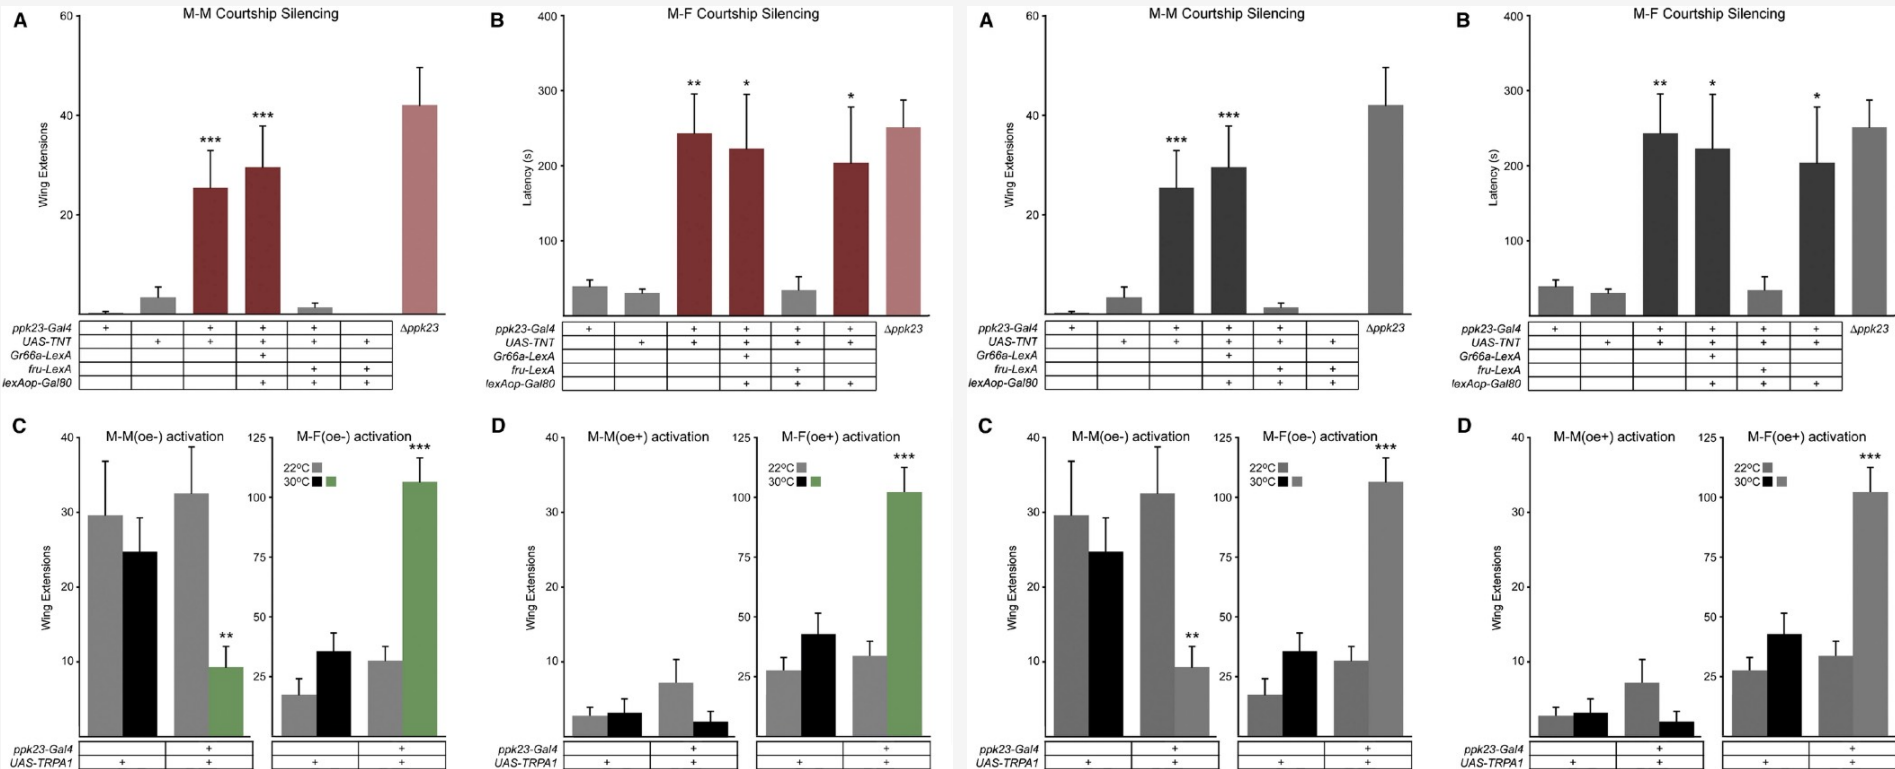
Color-blind  
friendly



# Color





# Color





A) Red bars denote responses different than ppk23-Gal4 control.  
 B) Red bars denote responses different than ppk23-Gal4  
 C)???  
 D)???

# Red/green vs. blue/orange



To black  
and white



# Color

---

- Avoid harsh, non-printable colors
- Use color with purpose: highlight data, or make it equally visible
- Do not encode continuous variables with color; categorical variables are best
- Avoid red/green; think about black/white printing
- Consistently use a color scheme throughout your figures