Adversarial Principal Component Analysis

Daniel Pimentel-Alarcón,
Ari Biswas, Claudia Solís-Lemus

Wisconsin Institute for Discovery
UNIVERSITY *of* WISCONSIN-MADISON

ISIT 2017

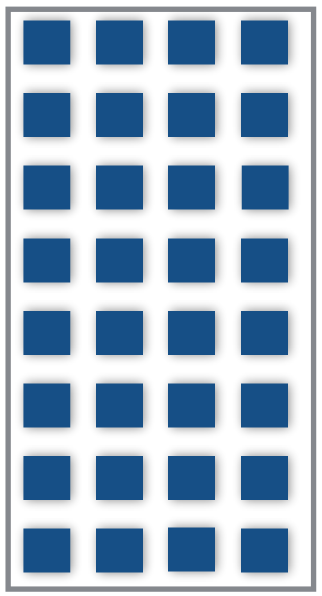

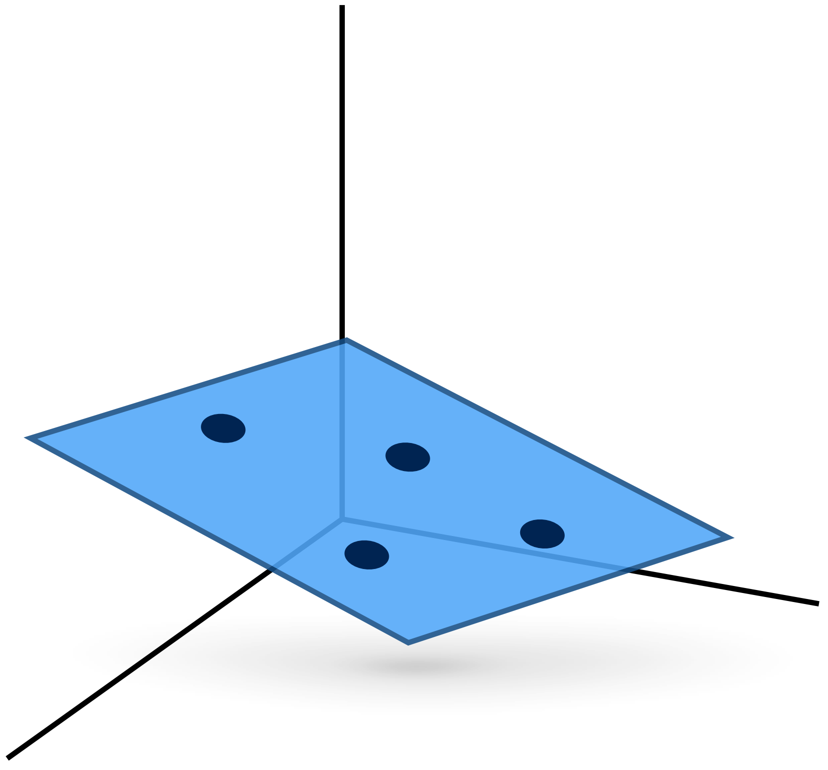

PCA

Finds a subspace that explains data

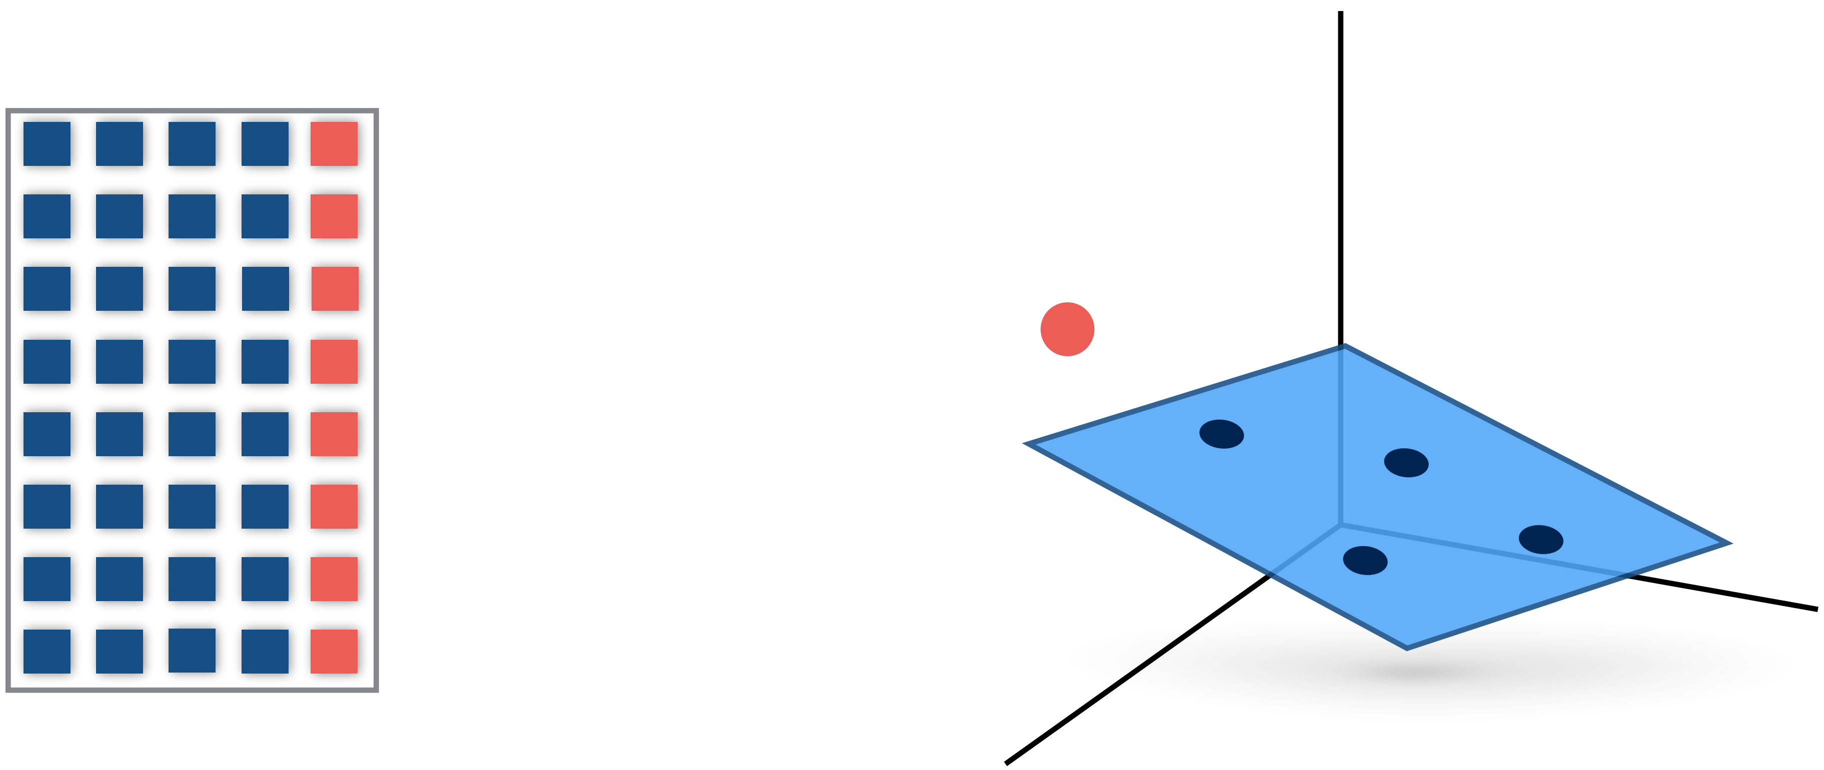

PCA

An outlier would *tilt* the subspace.

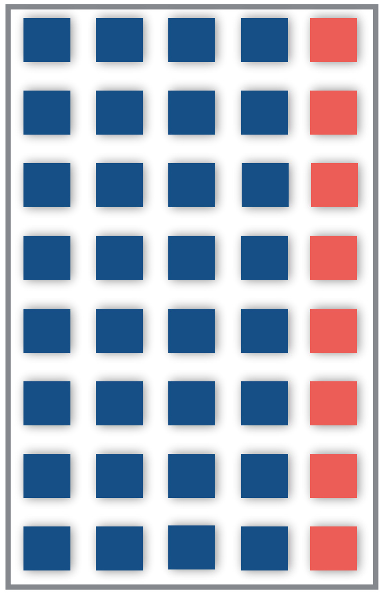

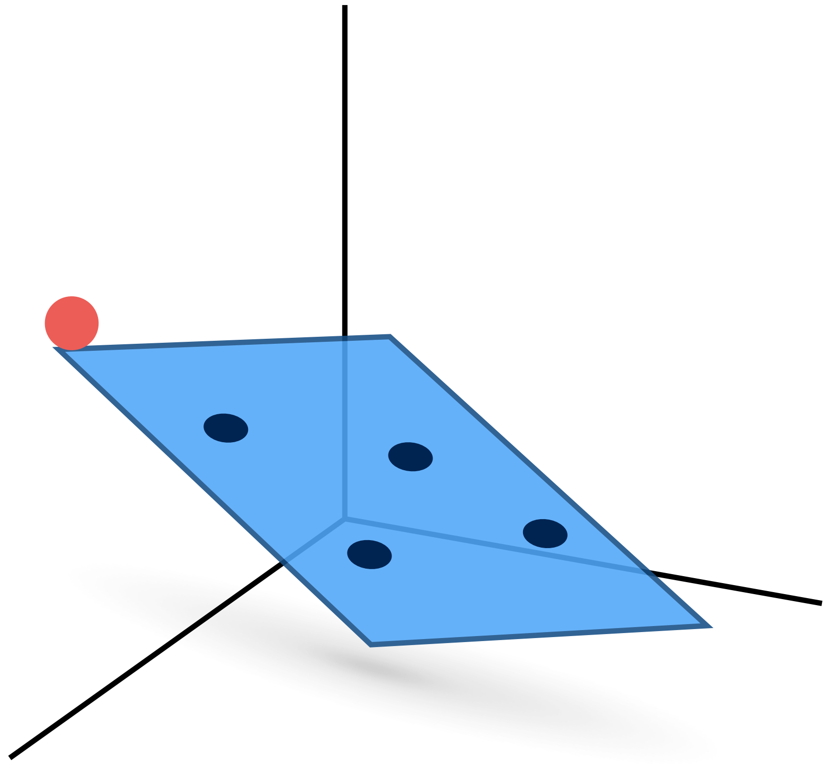

PCA

An outlier would *tilt* the subspace.

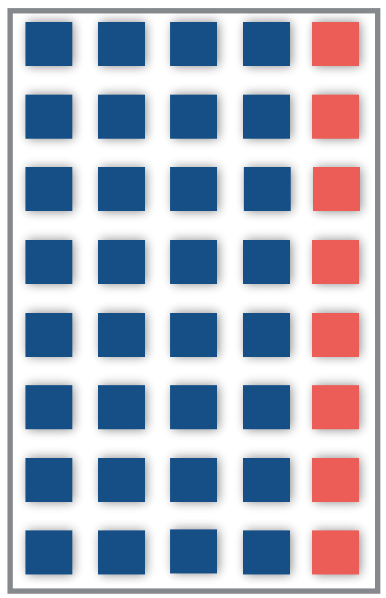

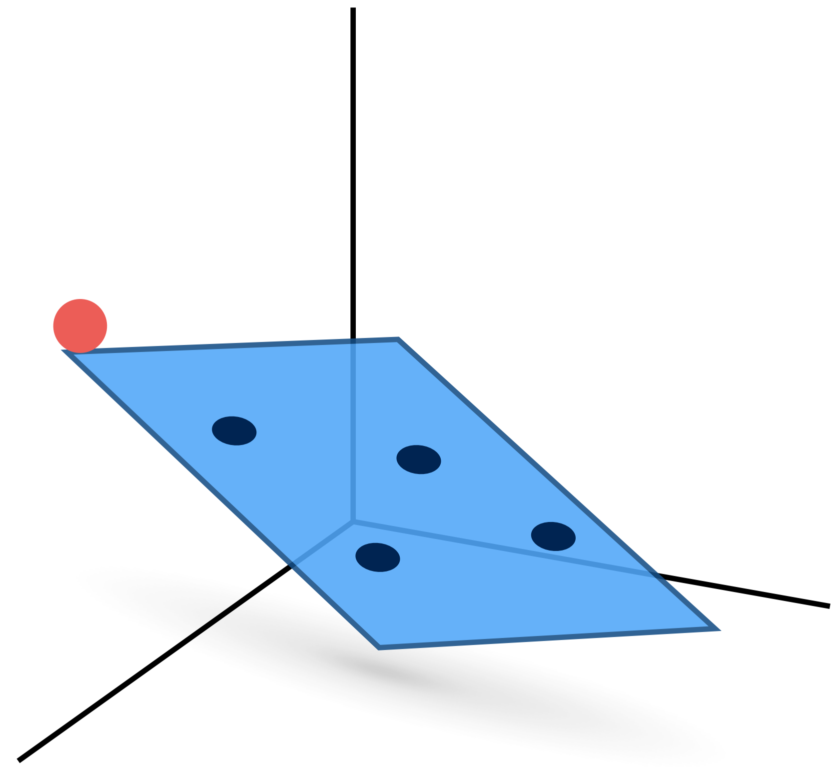

Adversarial PCA

Where should we put so that

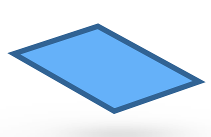 is tilted as much as possible?

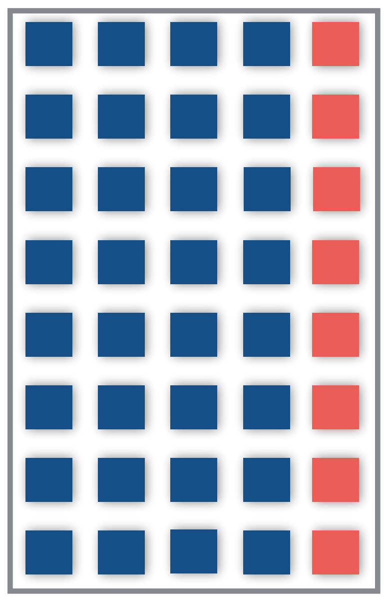

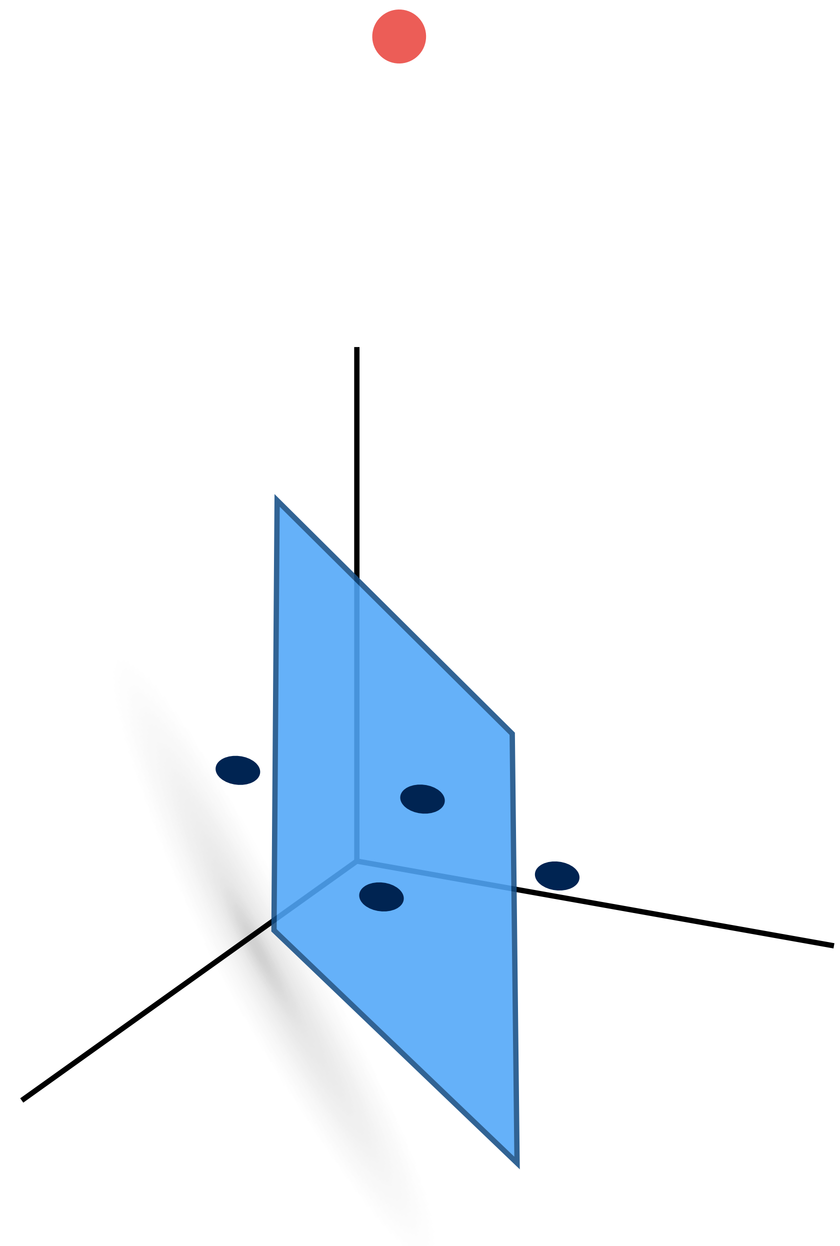

Adversarial PCA

Where should we put 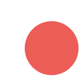 so that

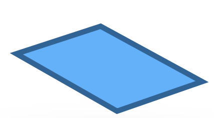 is tilted as much as possible?

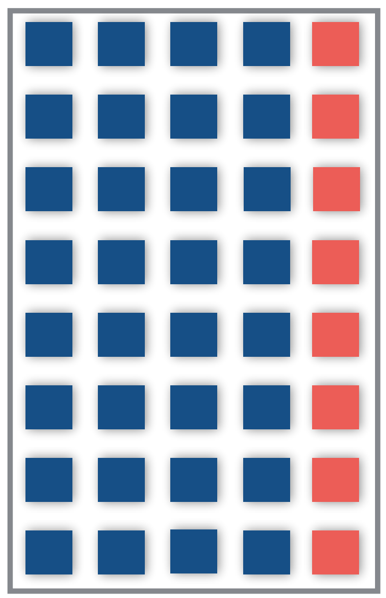

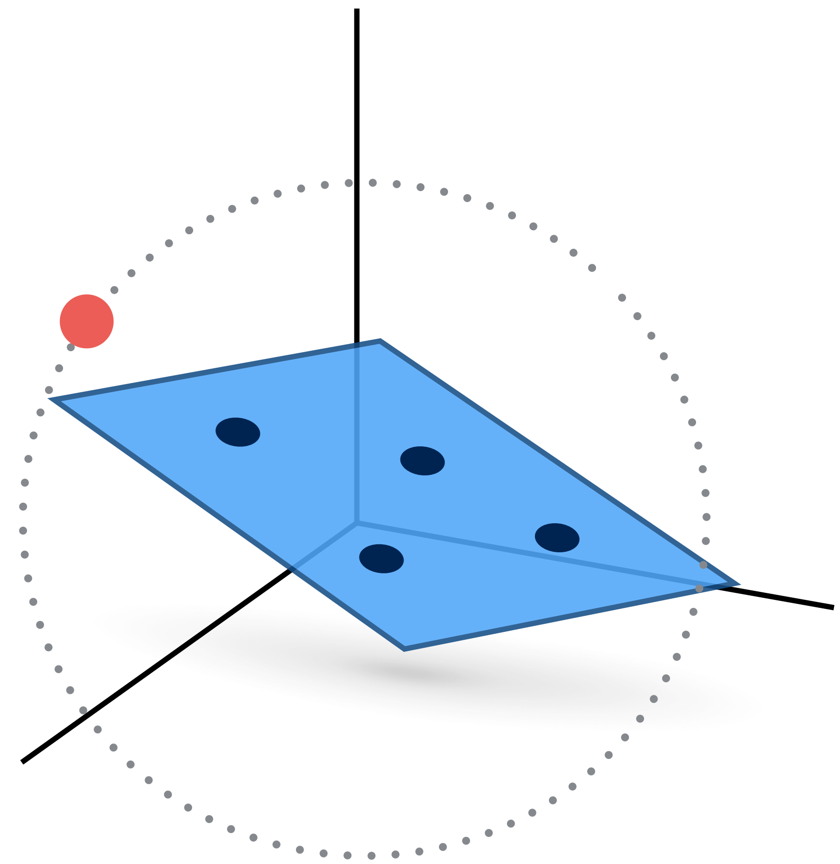

Adversarial PCA

Where should we put ● so that

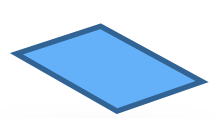 is tilted as much as possible?

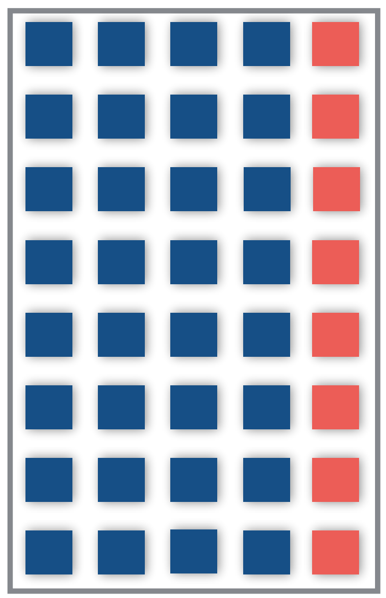

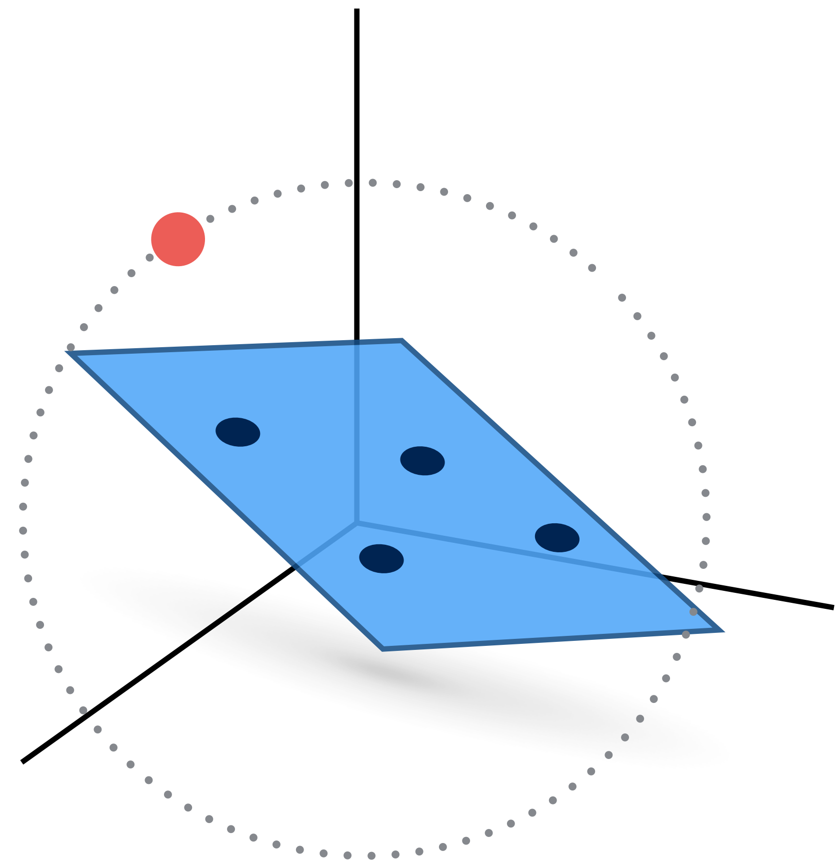

Adversarial PCA

Where should we put so that

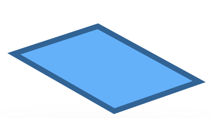 is tilted as much as possible?

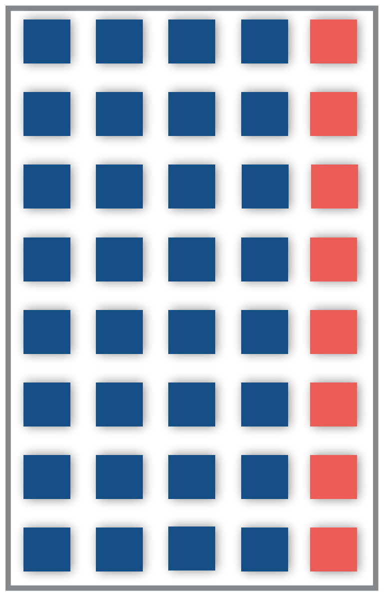

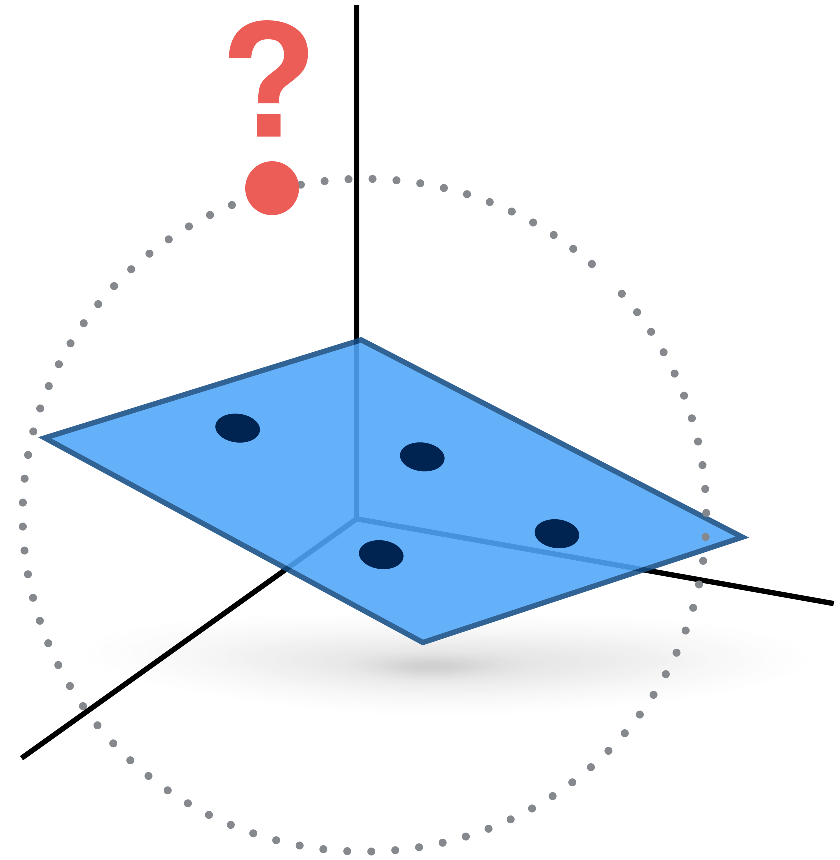

Adversarial PCA

Where should we put 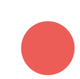 so that

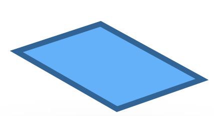 is tilted as much as possible?

Rob Nowak

Rob Nowak

Rob Nowak

Laura Balzano

Rob Nowak

Laura Balzano

John Lipor

Rob Nowak

Laura Balzano

John Lipor

Nigel Boston

Rob Nowak

Laura Balzano

John Lipor

Nigel Boston

Steve Wright

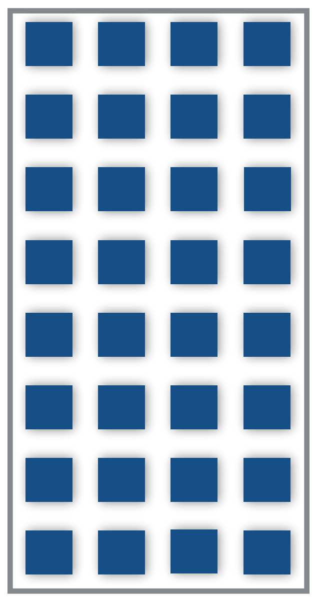

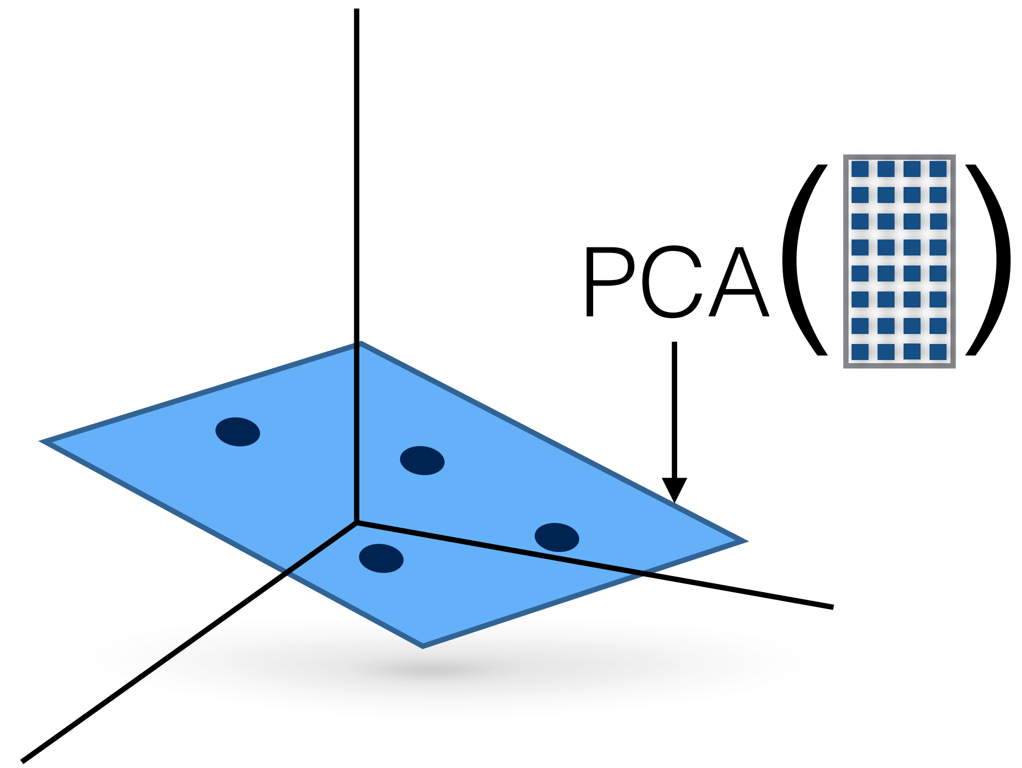

Rank-One Updates

Given a new point , how do
we compute new PCA efficiently?

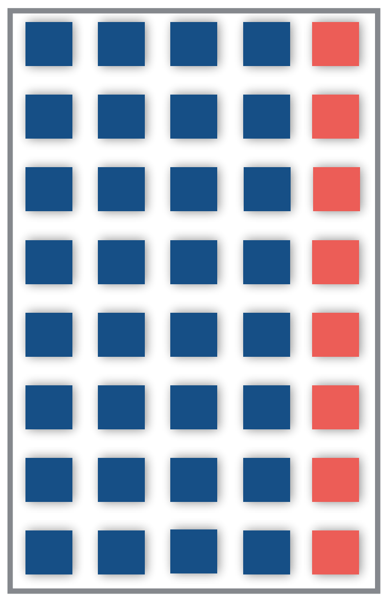

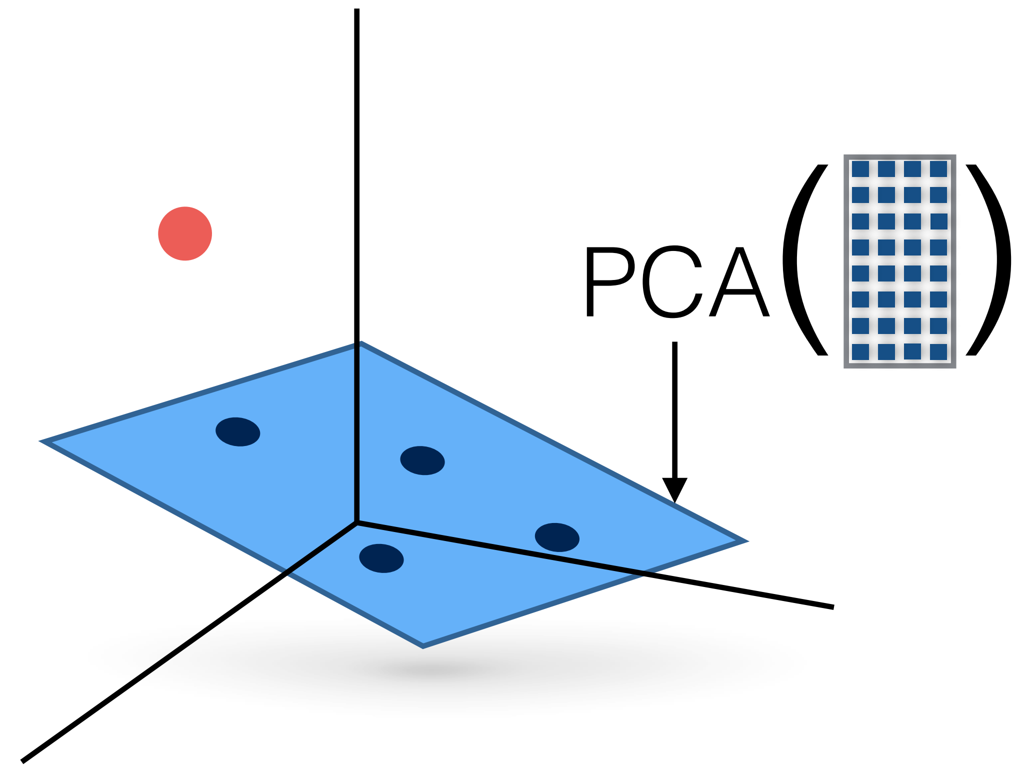

Rank-One Updates

Given a new point ●, how do
we compute new PCA efficiently?

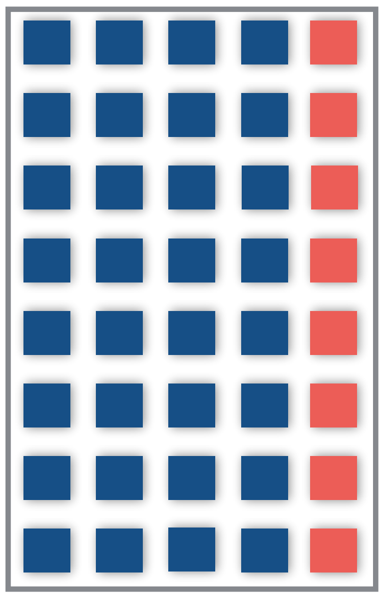

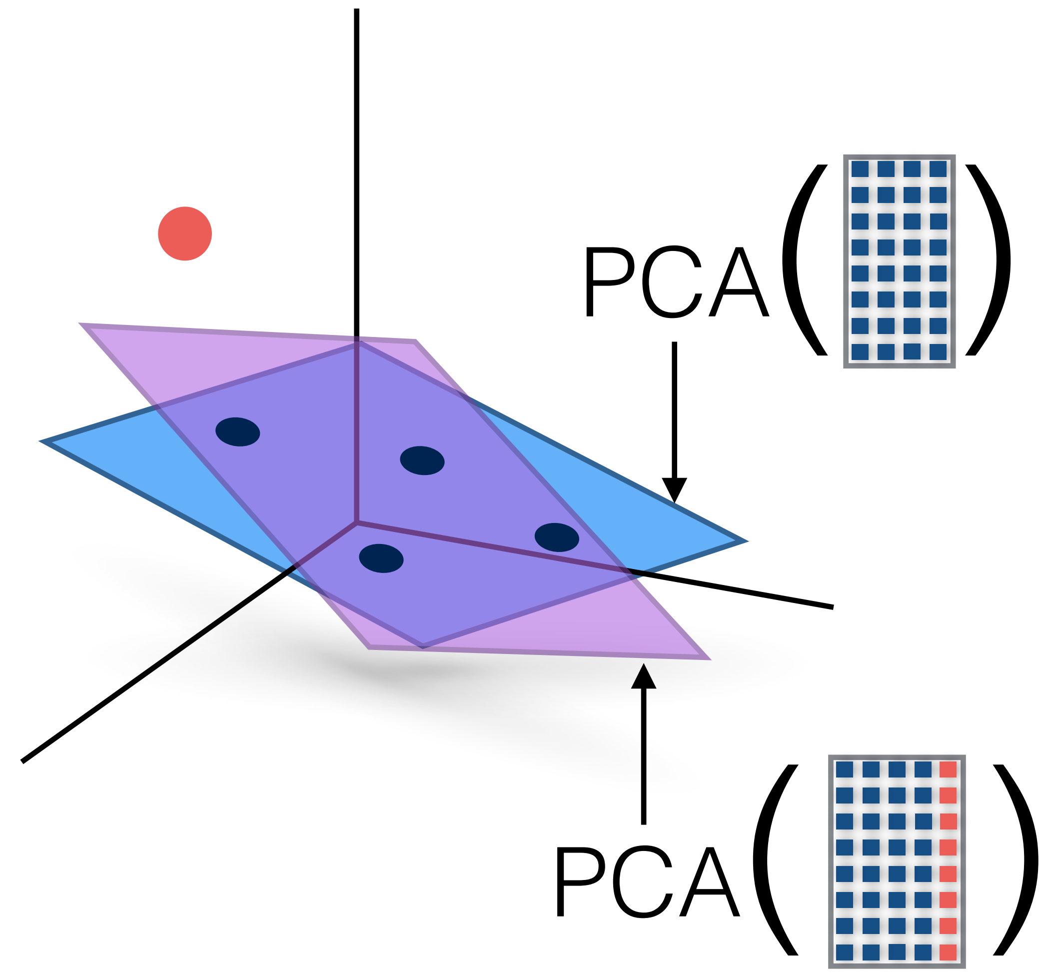

Rank-One Updates

Given a new point , how do
we compute new PCA efficiently?

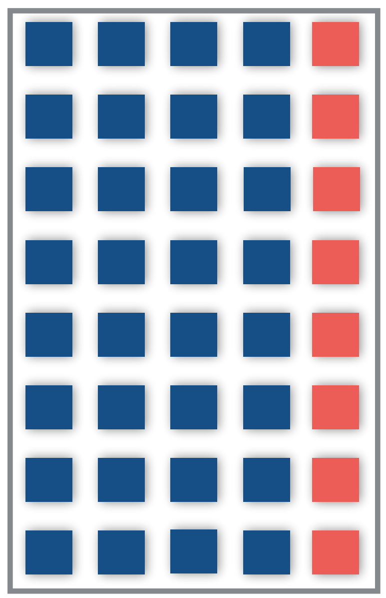

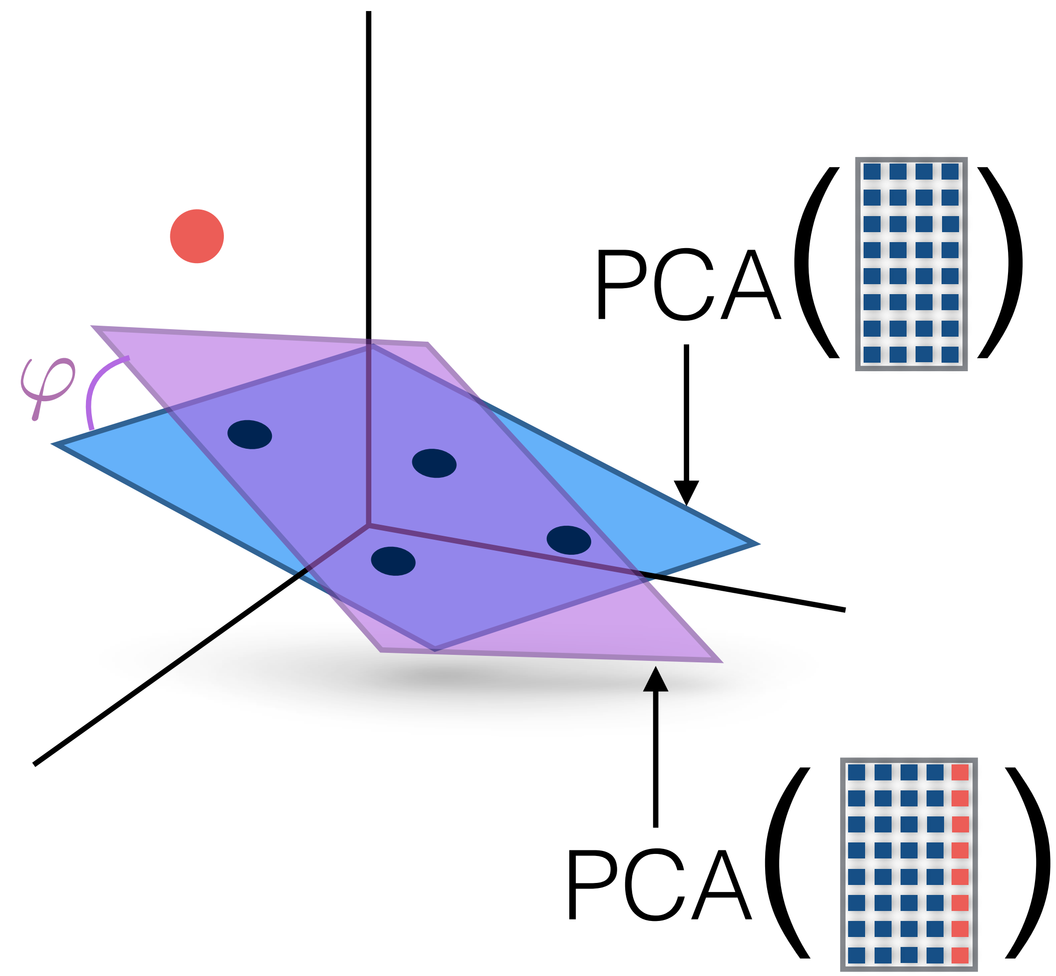

Adversarial PCA

Where should we put ● to maximize φ?


Why do you want to know?

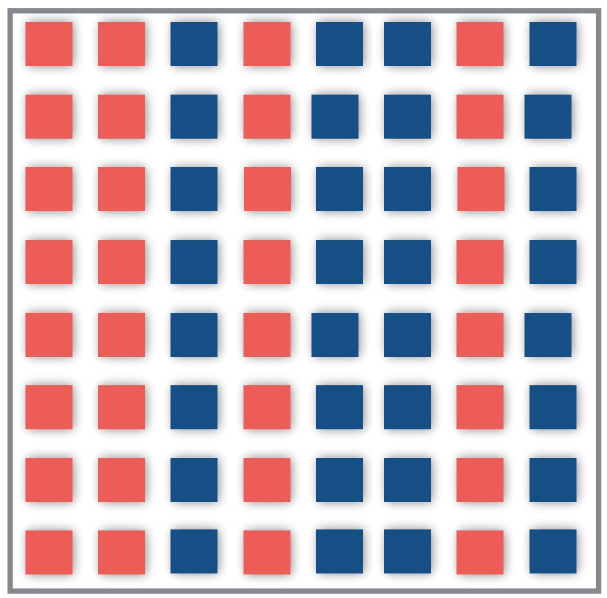

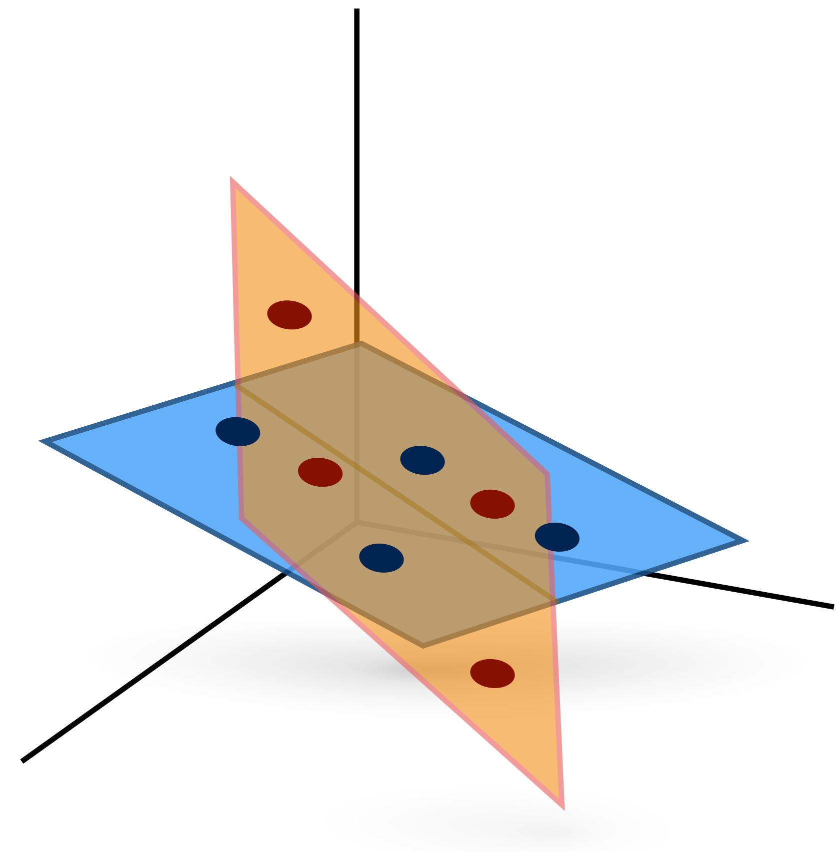

Subspace Clustering

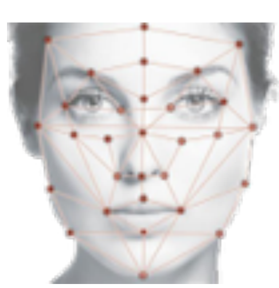

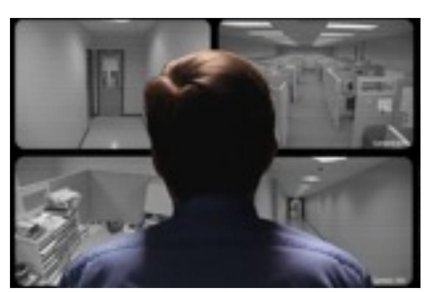

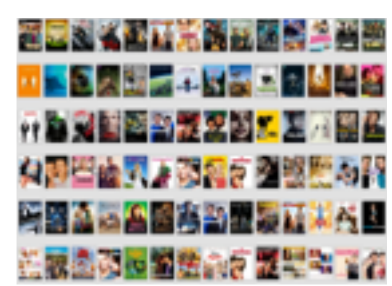

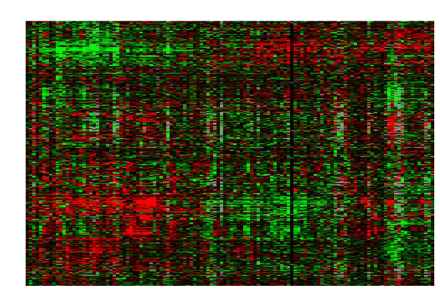

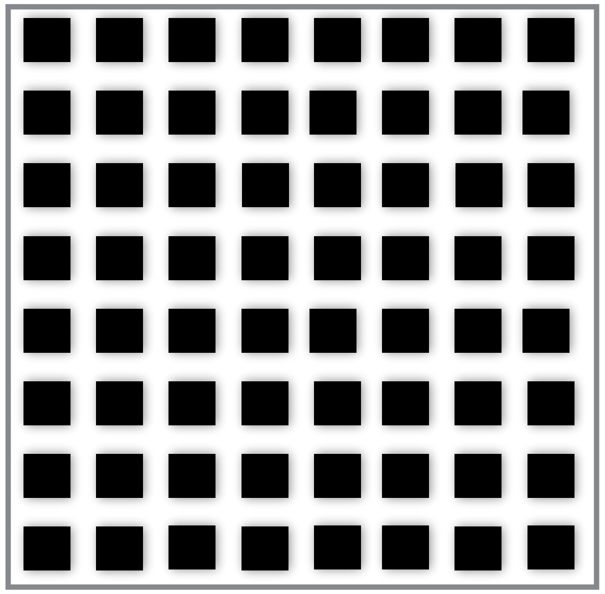

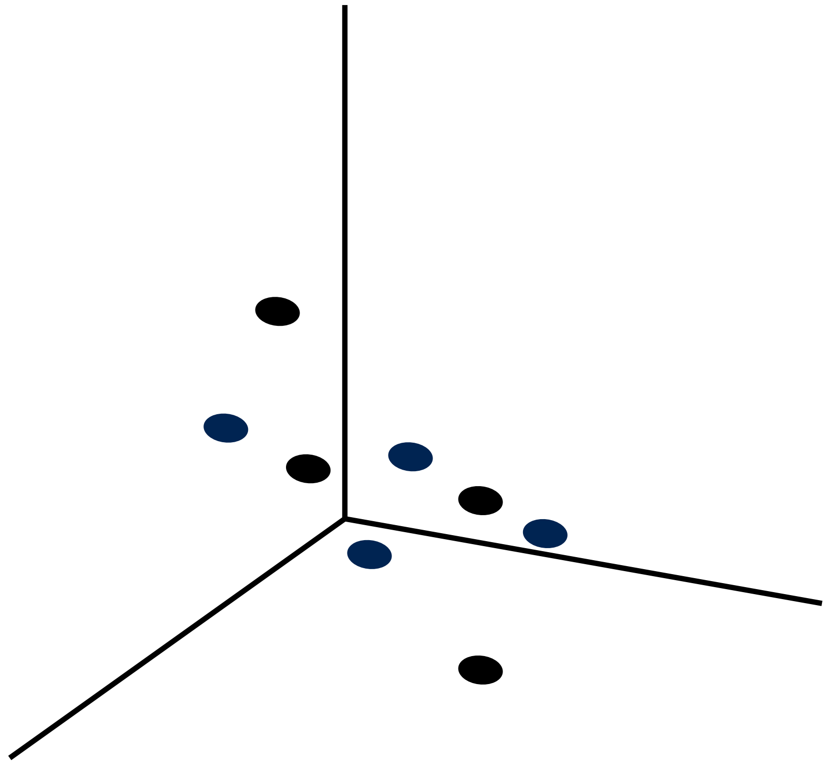

Subspace Clustering

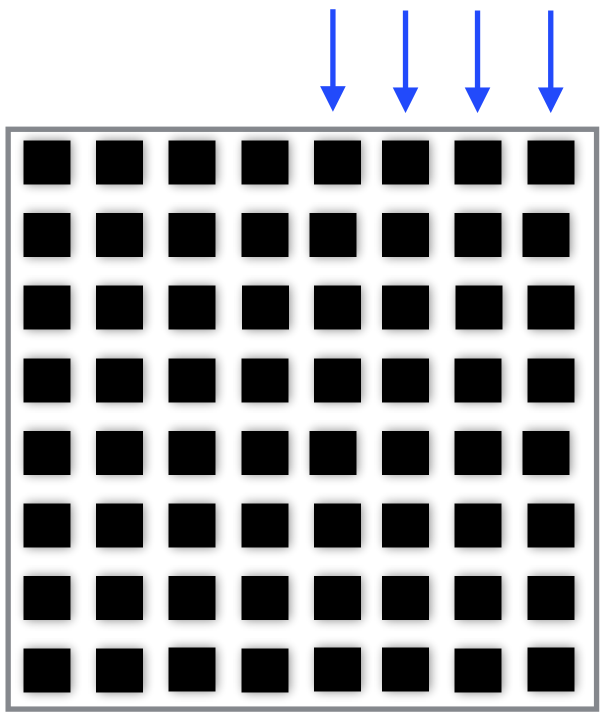

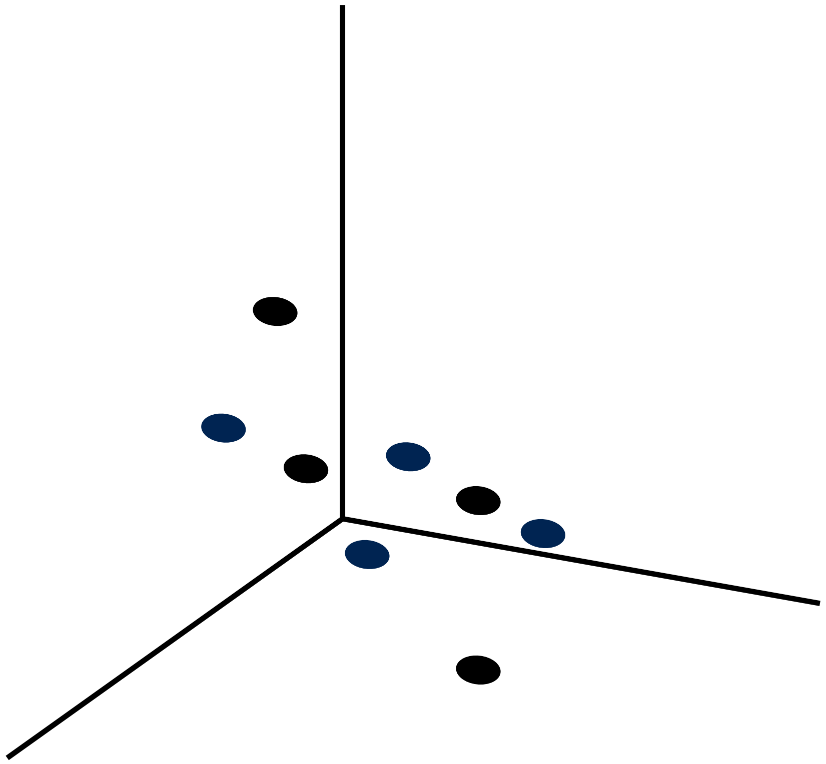

Subspace Clustering

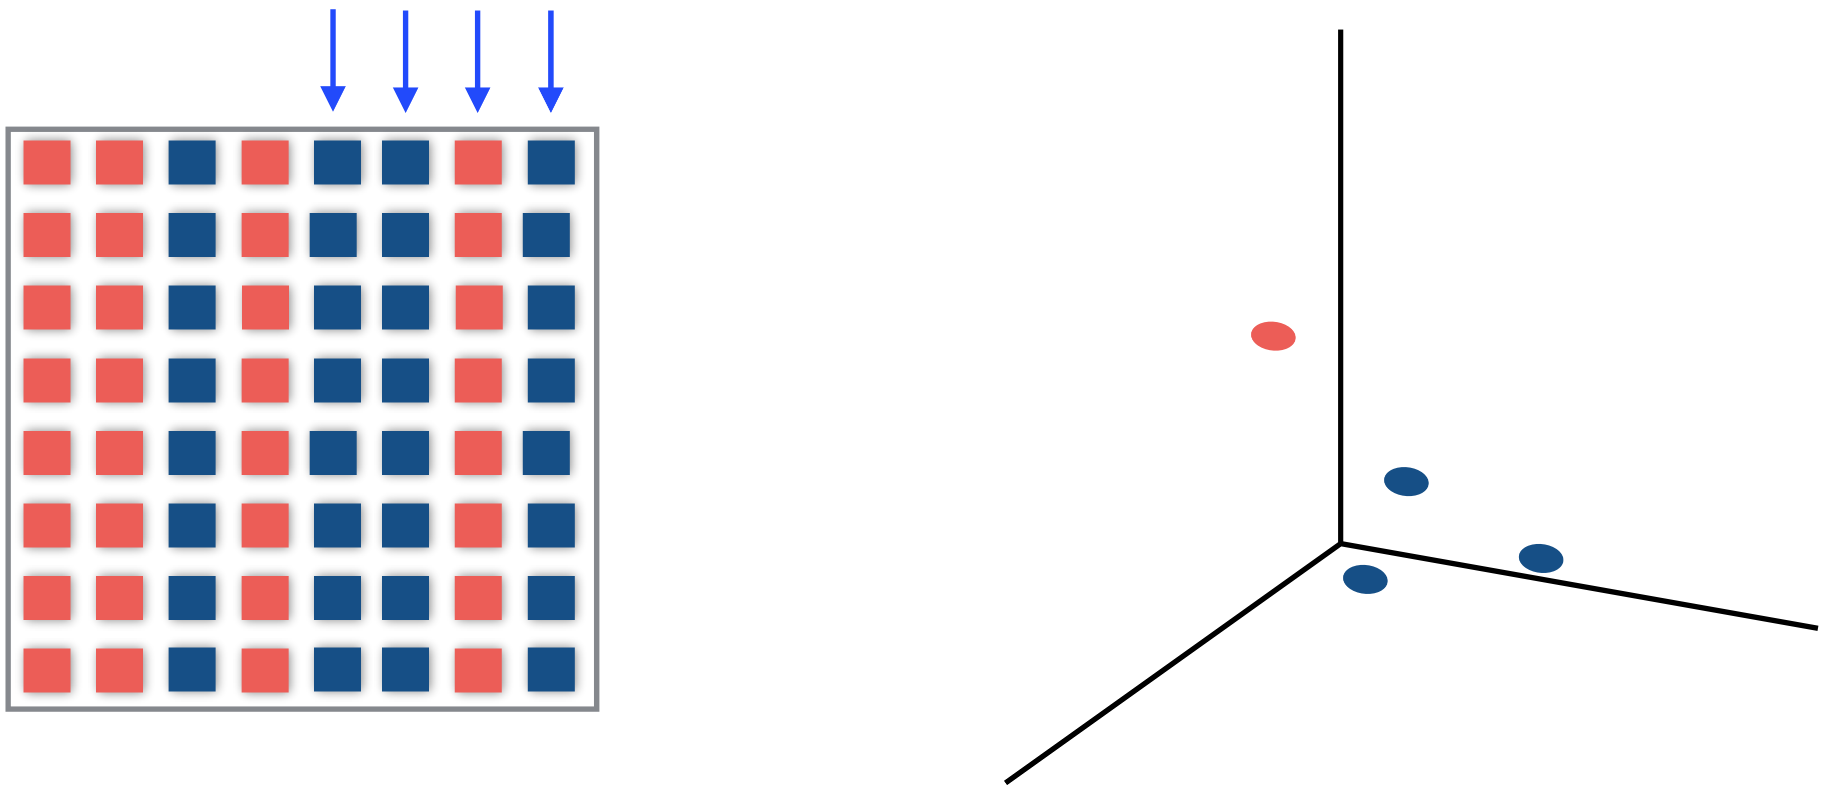

Subspace Clustering

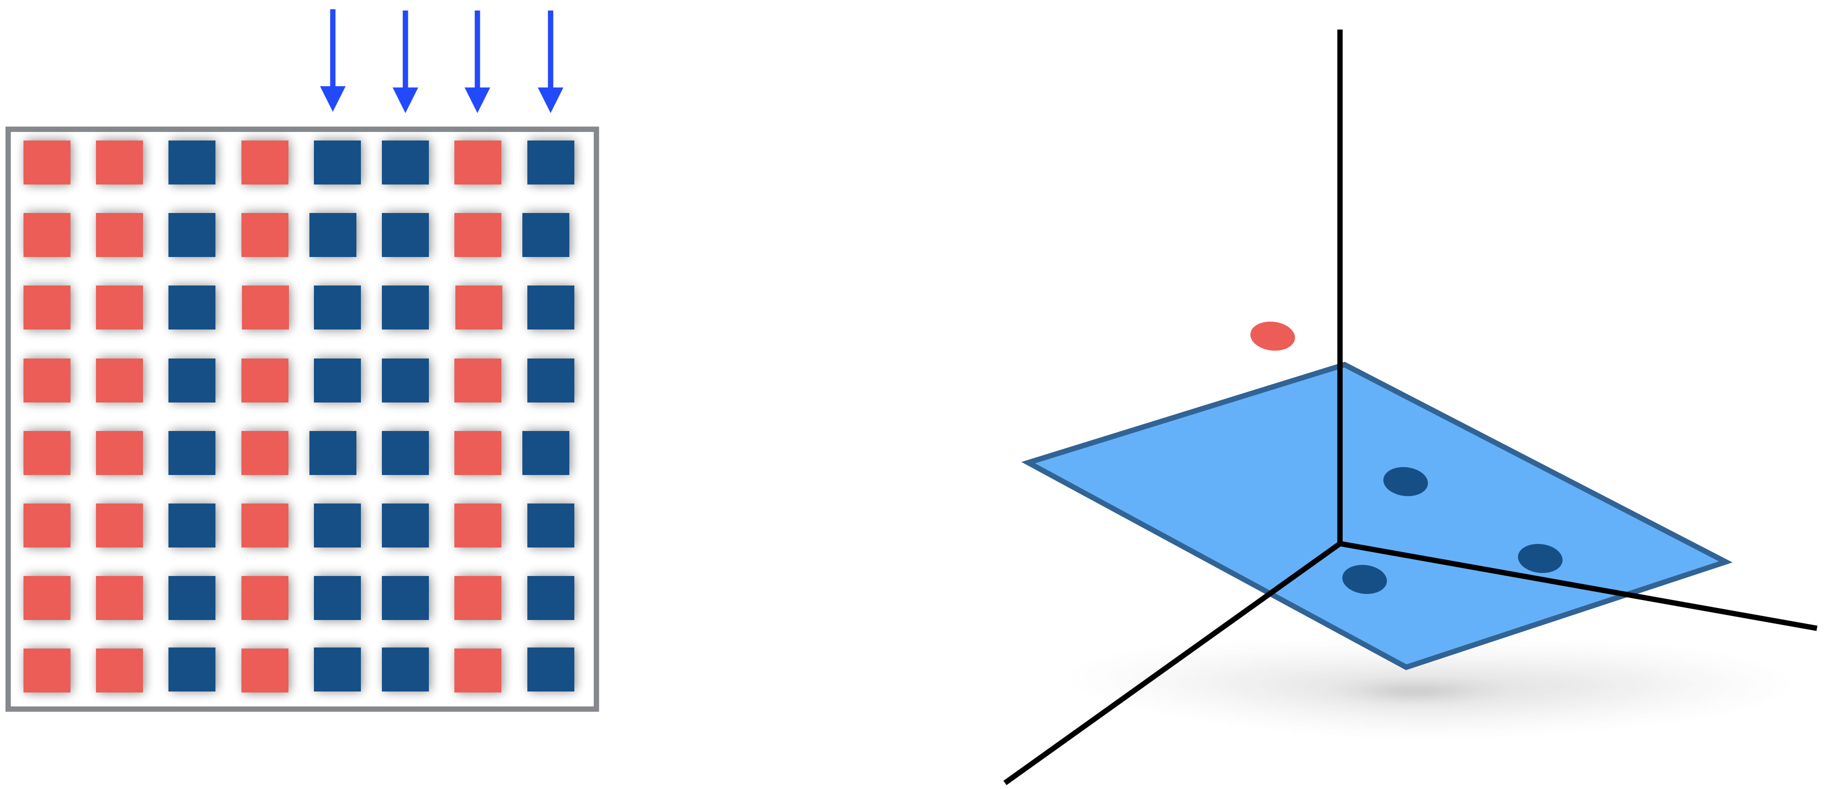

Subspace Clustering

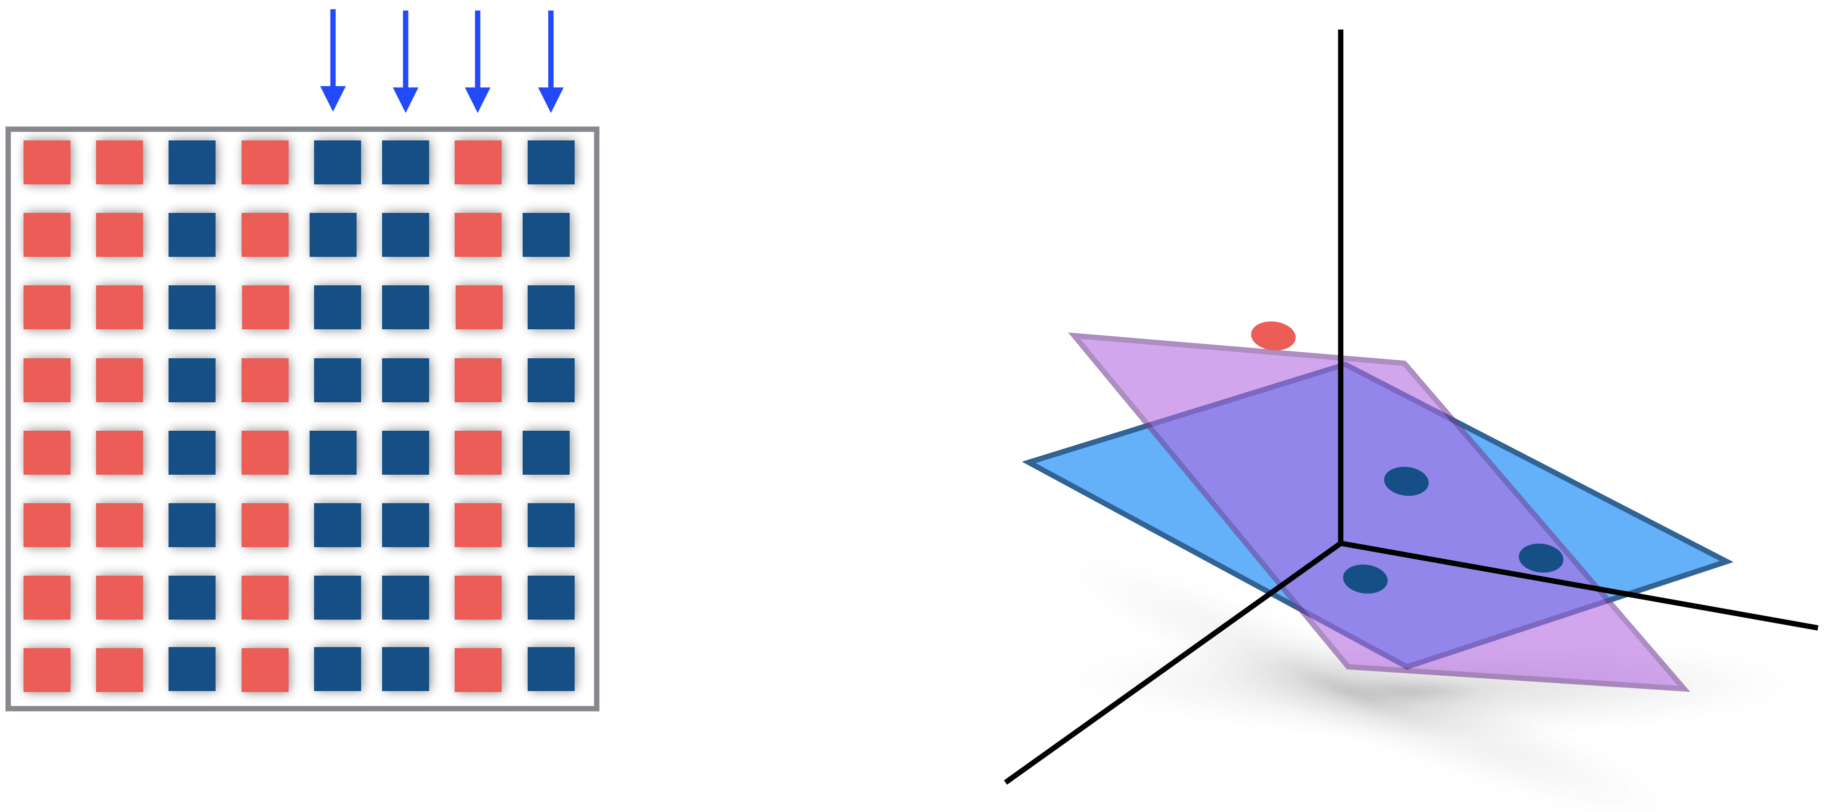

Subspace Clustering

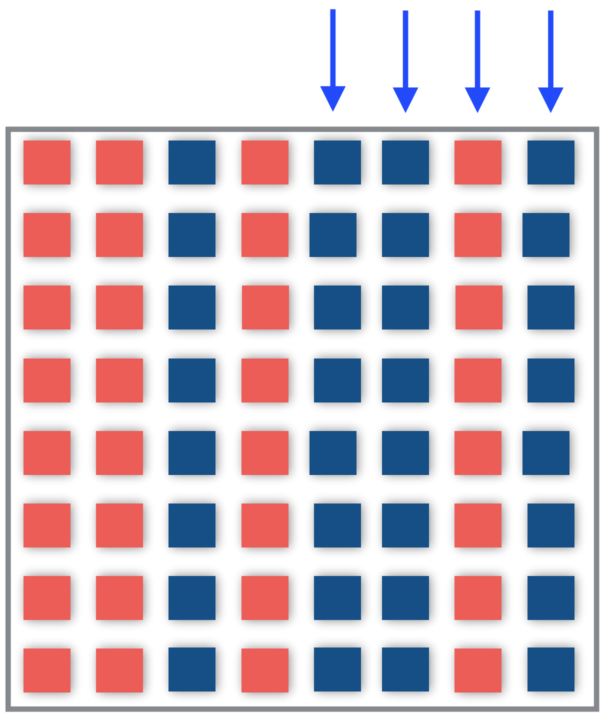

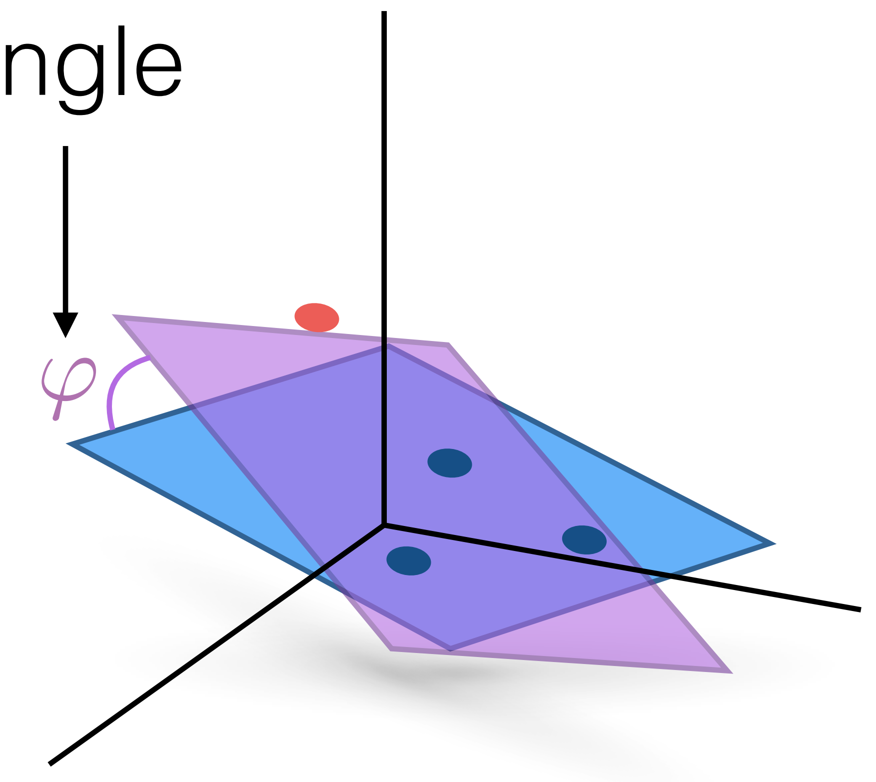

Subspace Clustering

We want to bound the *error* φ

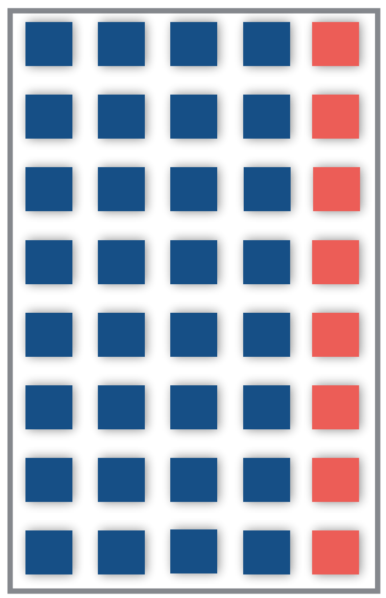

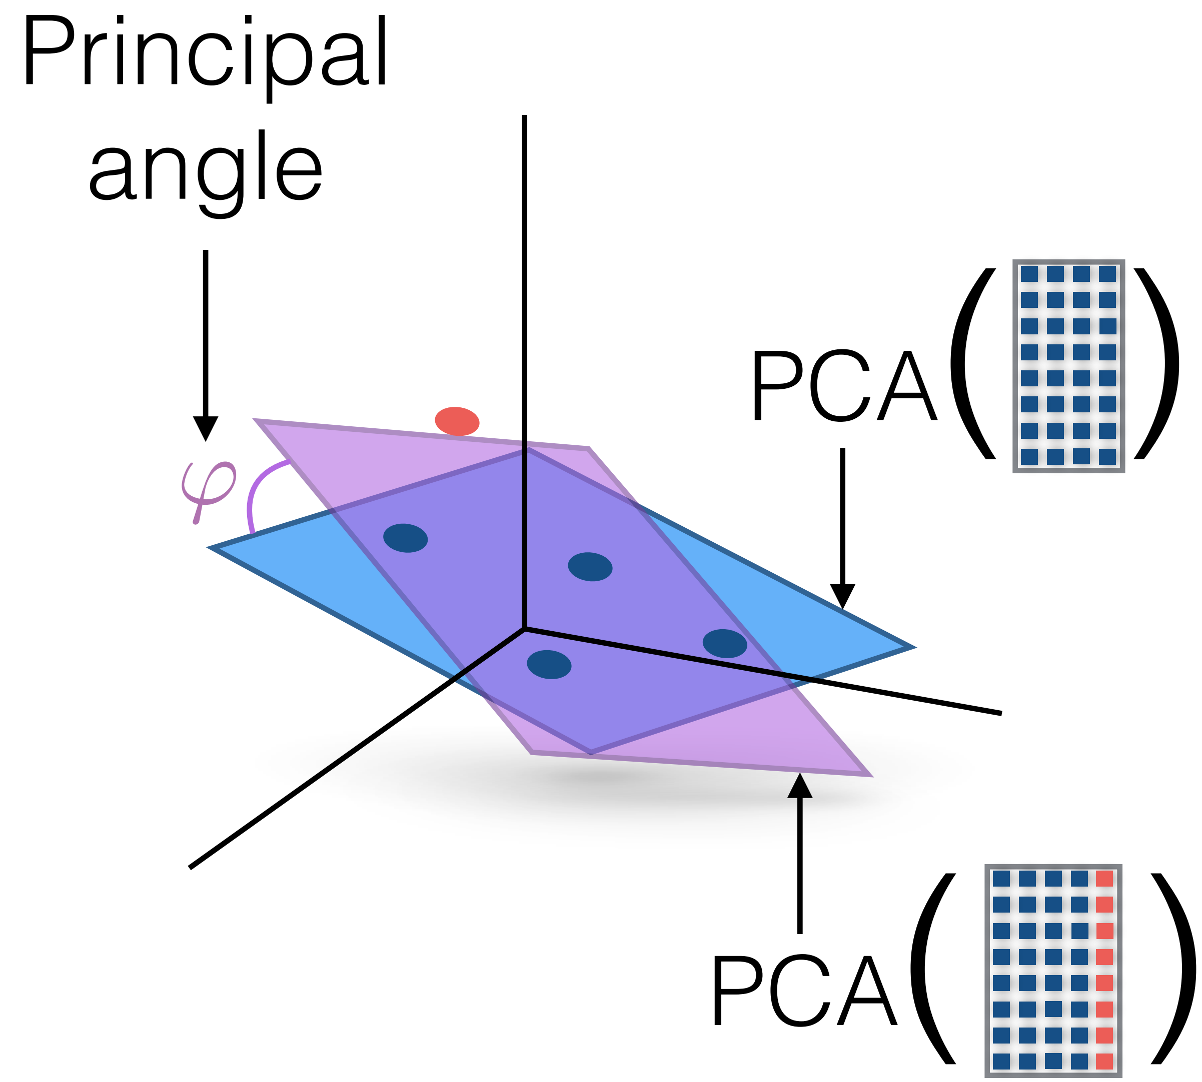

Adversarial PCA

Where should we put ● to maximize φ?

Our Main Theorem

Closed form solution

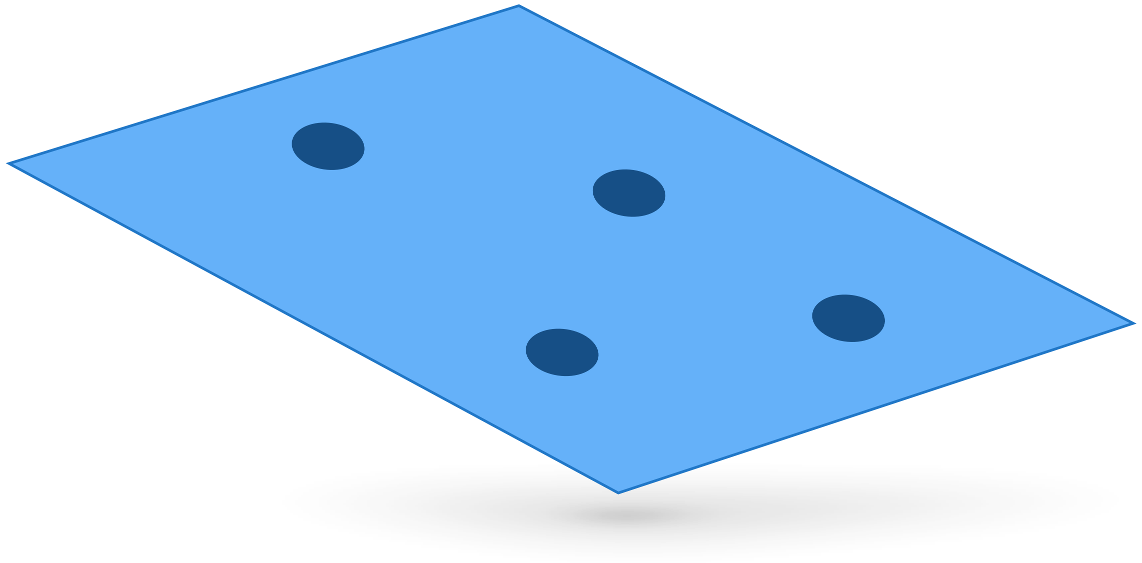

Our Main Theorem

Closed form solution

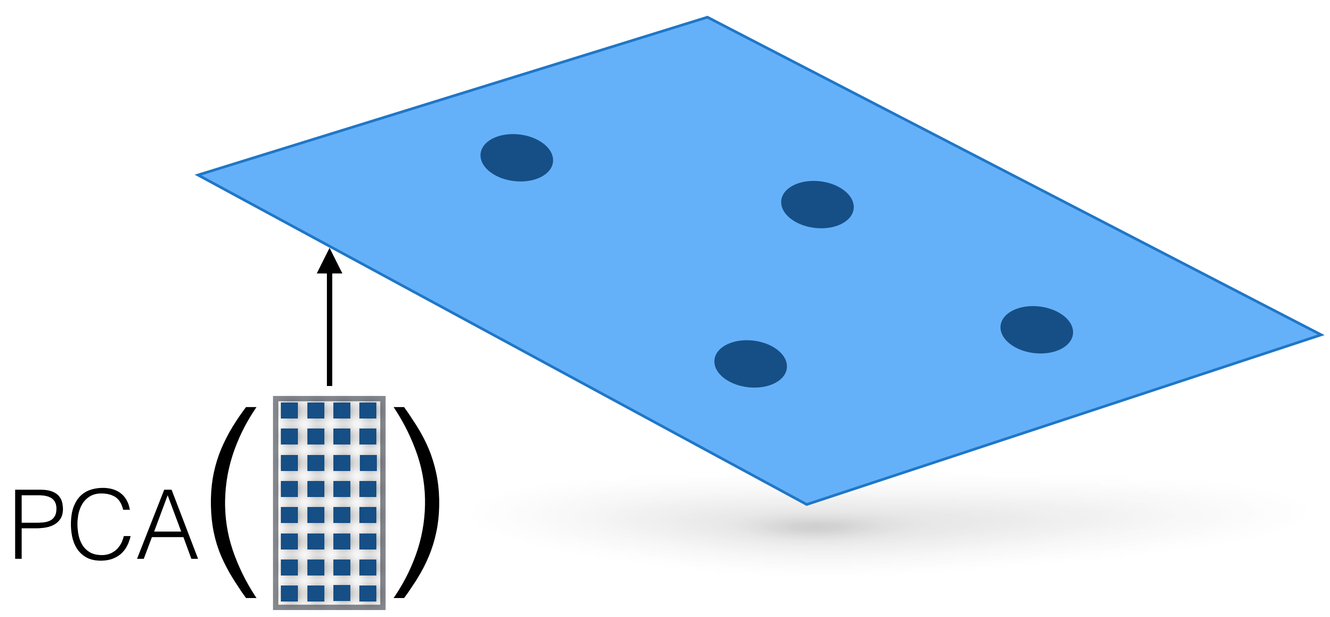

Our Main Theorem

Closed form solution

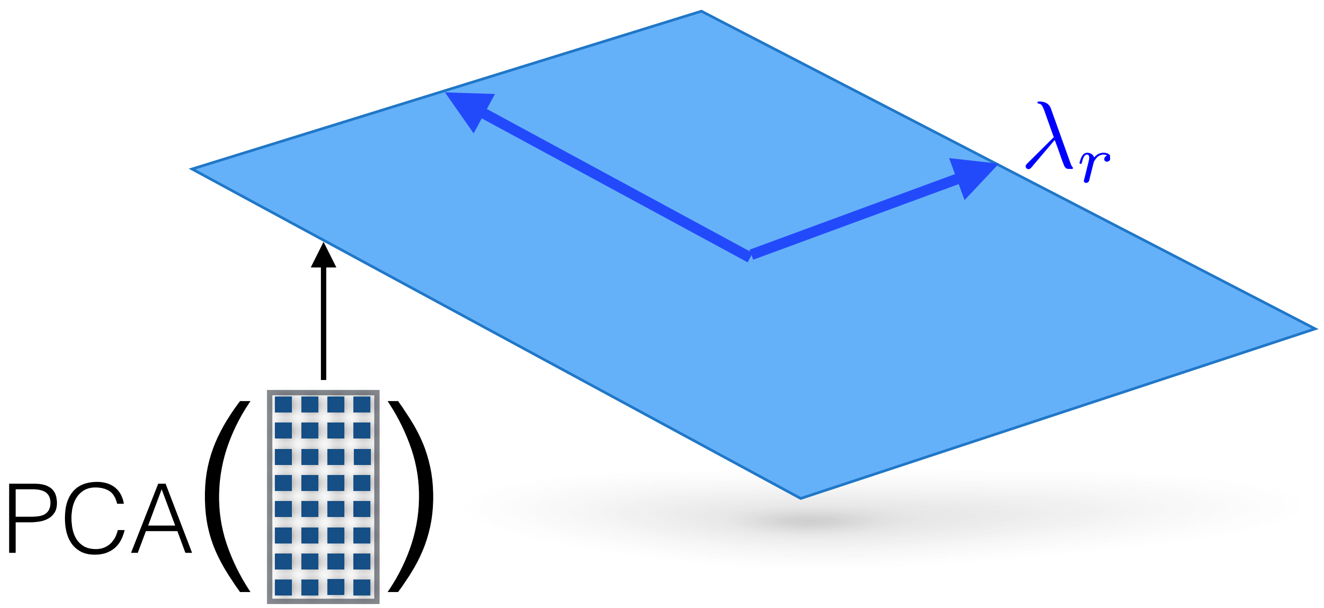

Our Main Theorem

Closed form solution

$$\theta^\star = \frac{1}{2} \arccos\left(-\frac{1}{{\lambda_r}^2}\right)$$

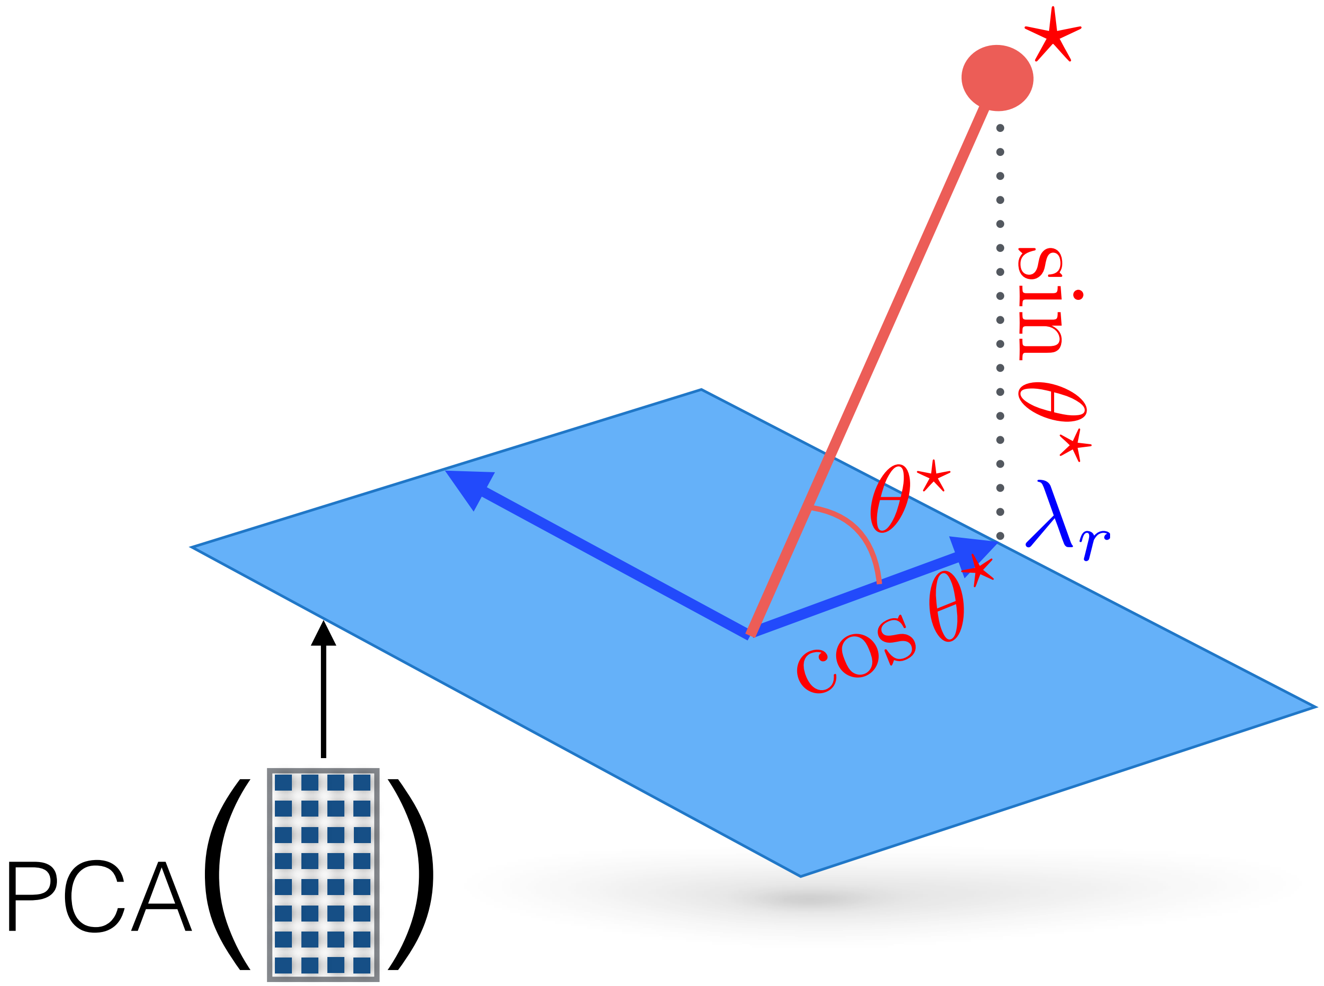

Our Main Theorem

Closed form solution

$$\theta^\star = \frac{1}{2}\arccos\left(-\frac{1}{{\lambda_r}^2}\right)$$

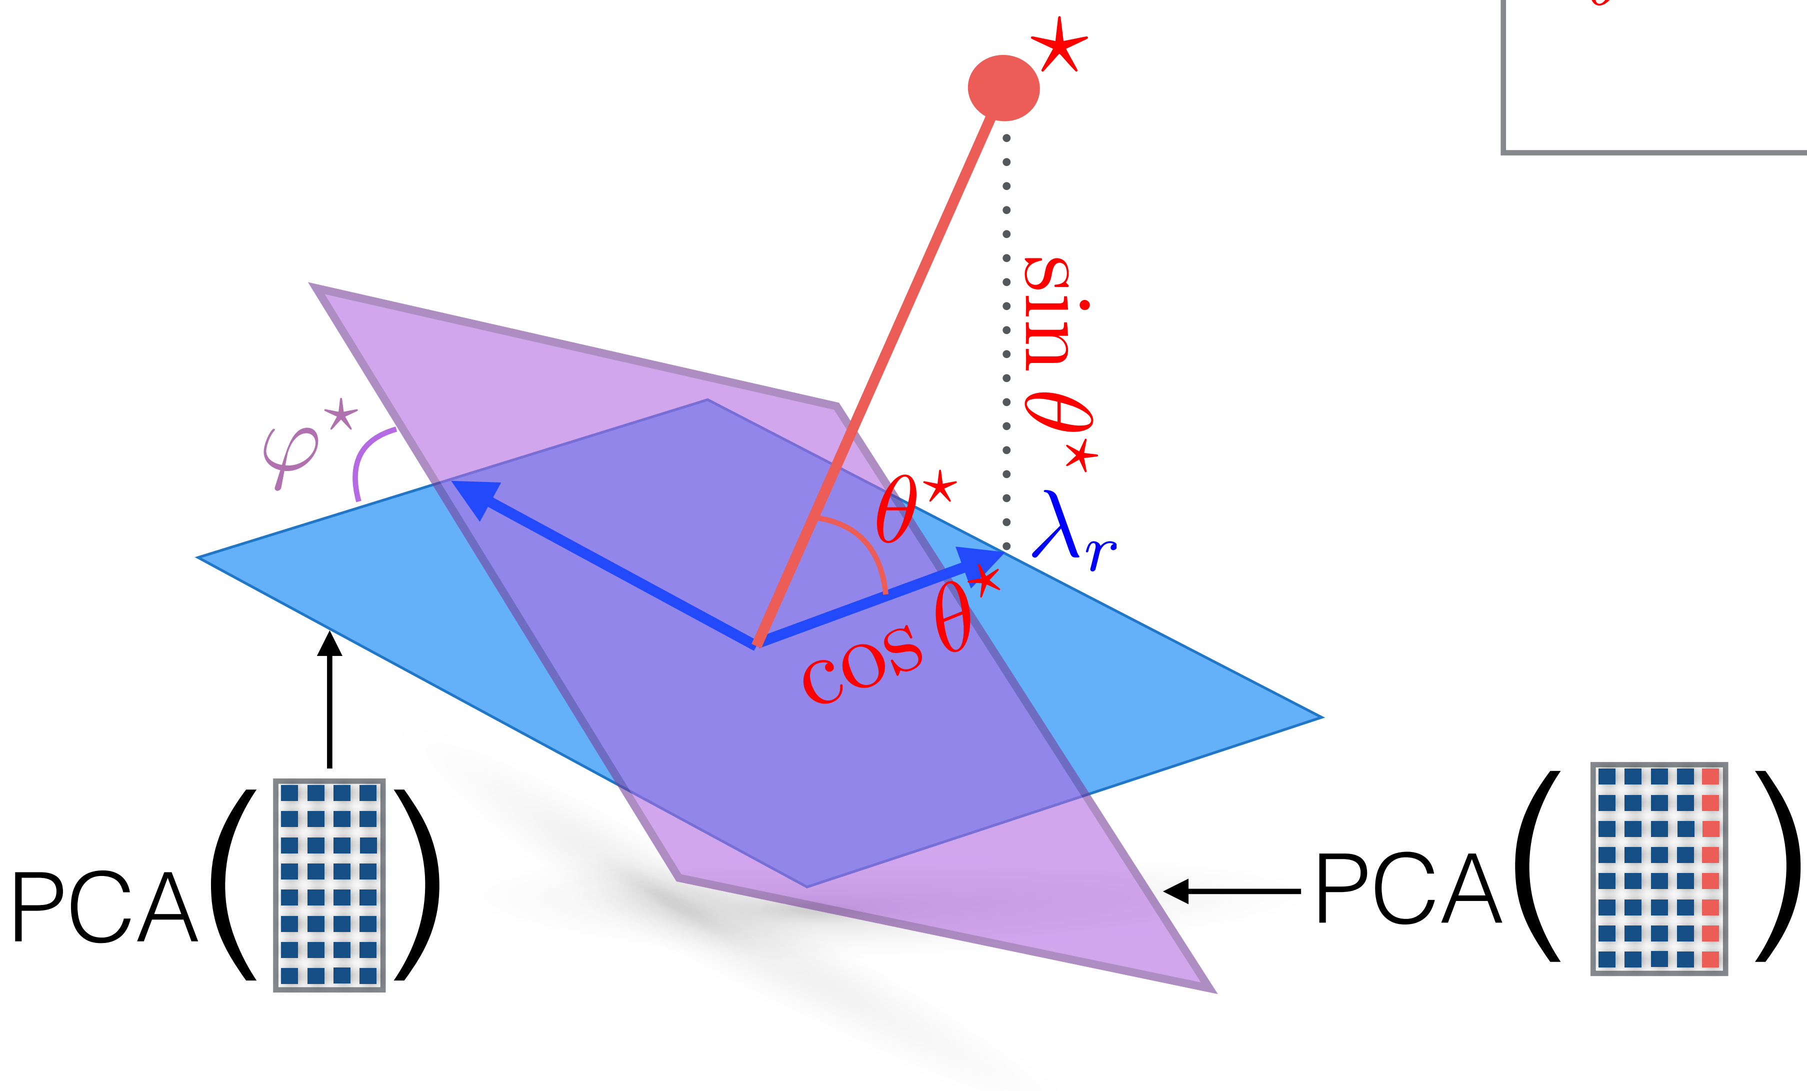

Our Main Theorem

Closed form solution

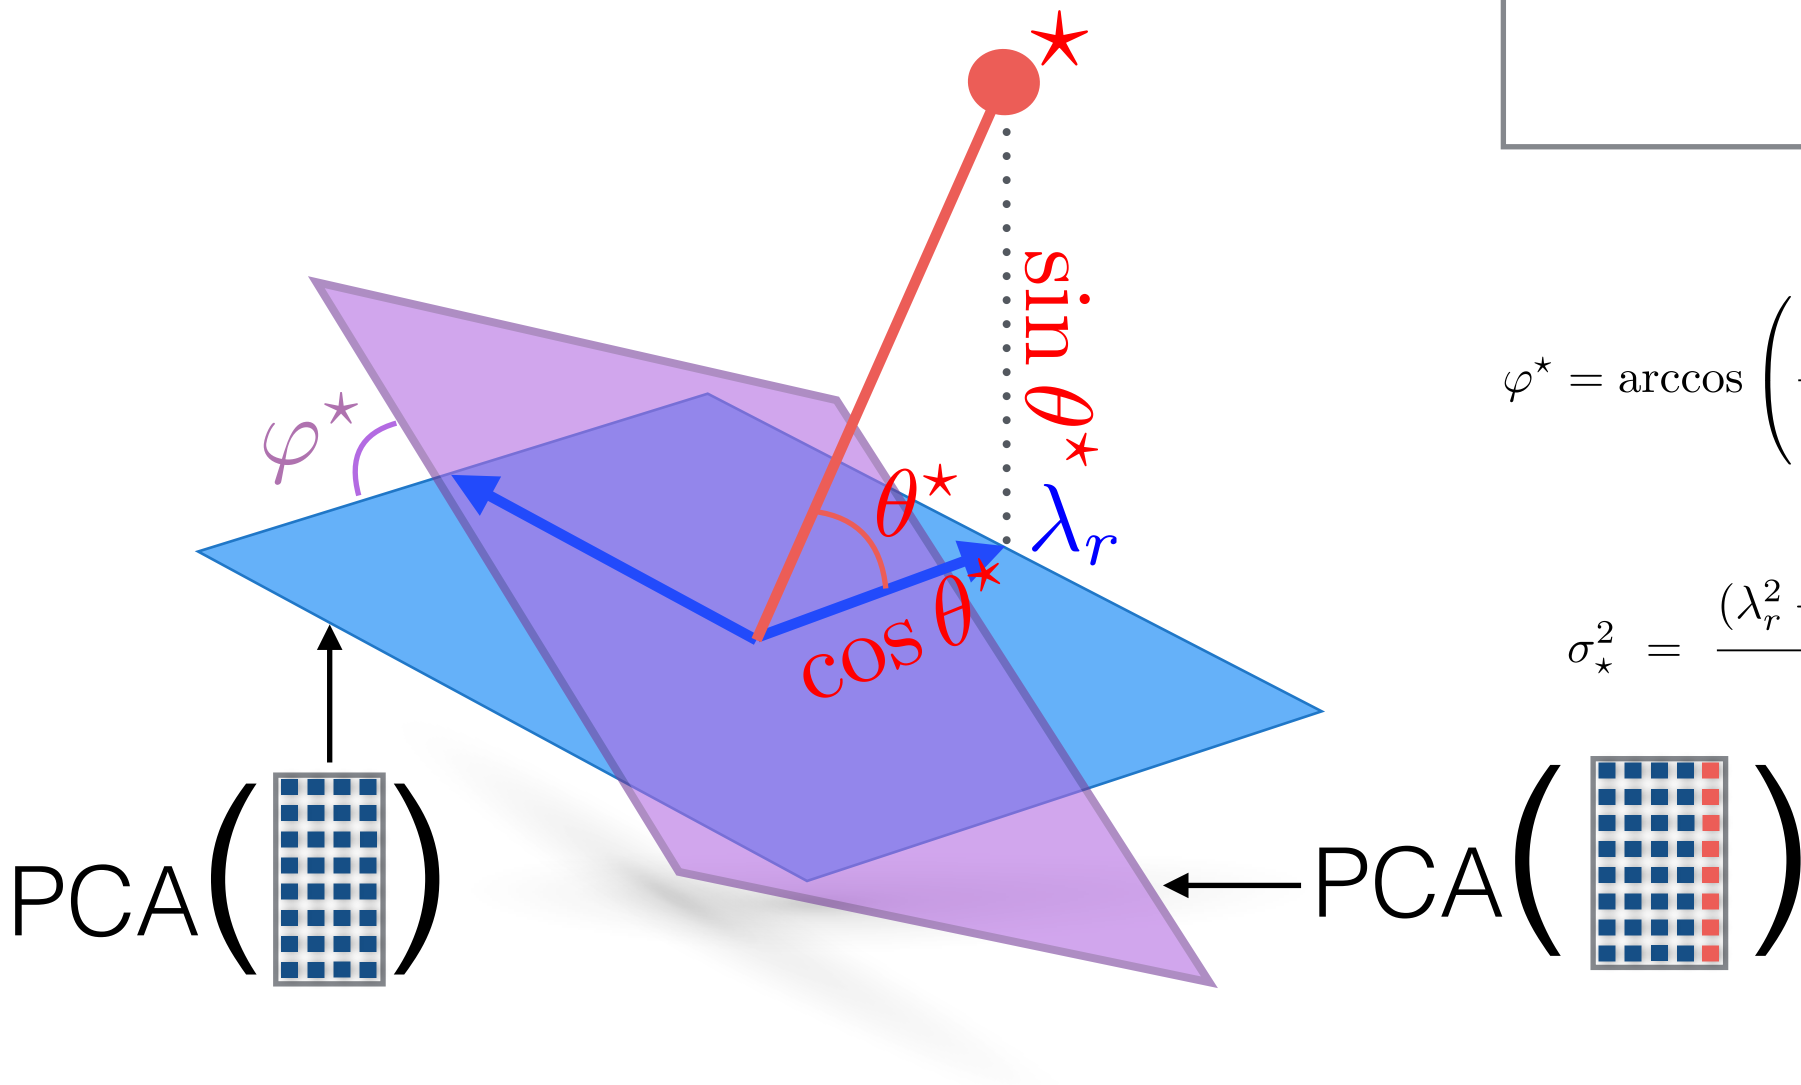

$$\theta^\star = \frac{1}{2}\arccos\left(-\frac{1}{{\lambda_r}^2}\right)$$

$$\varphi^\star = \arccos\left(\frac{\sin^2\theta^\star - \sigma_\star^2}{\sqrt{(\sin^2\theta^\star - \sigma_\star^2)^2 + (\sin\theta^\star\cos\theta^\star)^2}}\right)$$

$$\sigma_\star^2 = \frac{(\lambda_r^2+1) + \sqrt{(\lambda_r^2+1)^2 - 4\lambda_r^2\sin^2\theta^\star}}{2}.$$

Our Main Theorem

Closed form solution

THE FOLLOWING **PREVIEW** HAS BEEN APPROVED FOR

ALL AUDIENCES

BY THE MOTION PICTURE ASSOCIATION OF AMERICA INC.

THE FILM ADVERTISED HAS BEEN RATED

PG	GENERAL AUDIENCES
	All Ages Admitted
Linear Algebra, Geometry, Analysis	

www.filmratings.com

www.mpaa.org

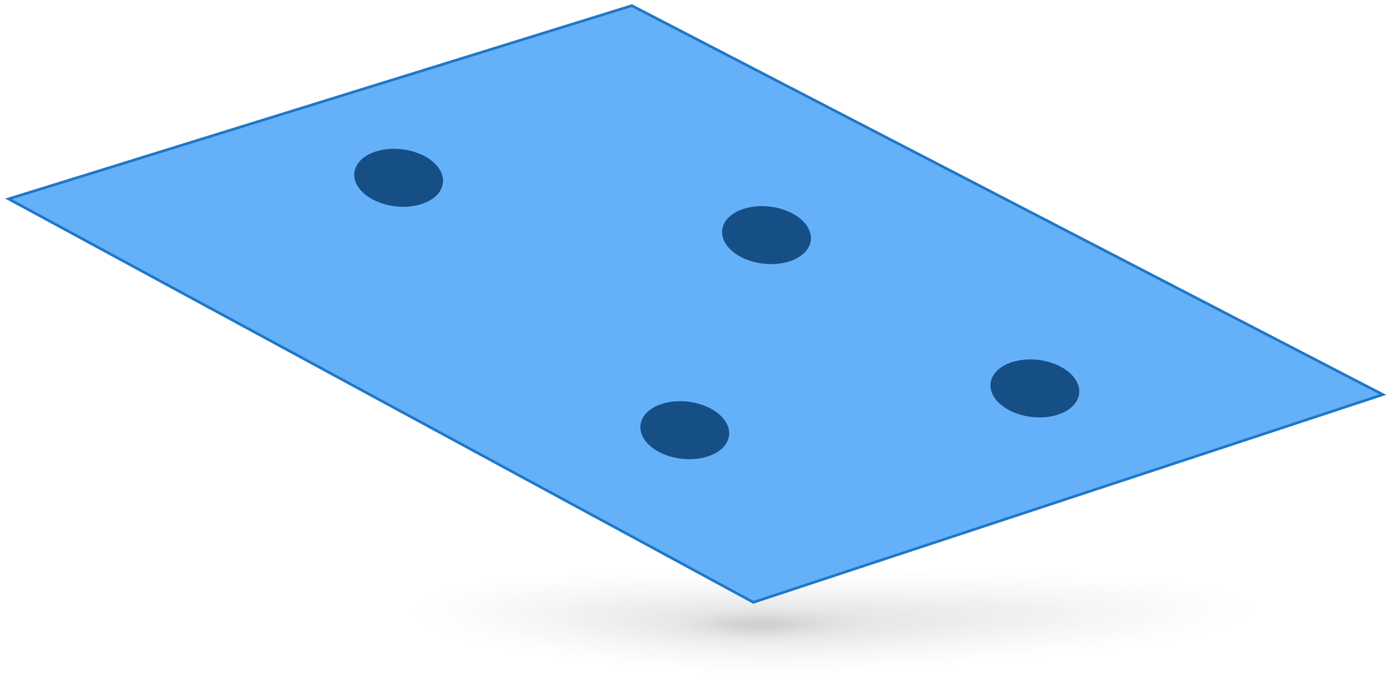

A flavor of the proof

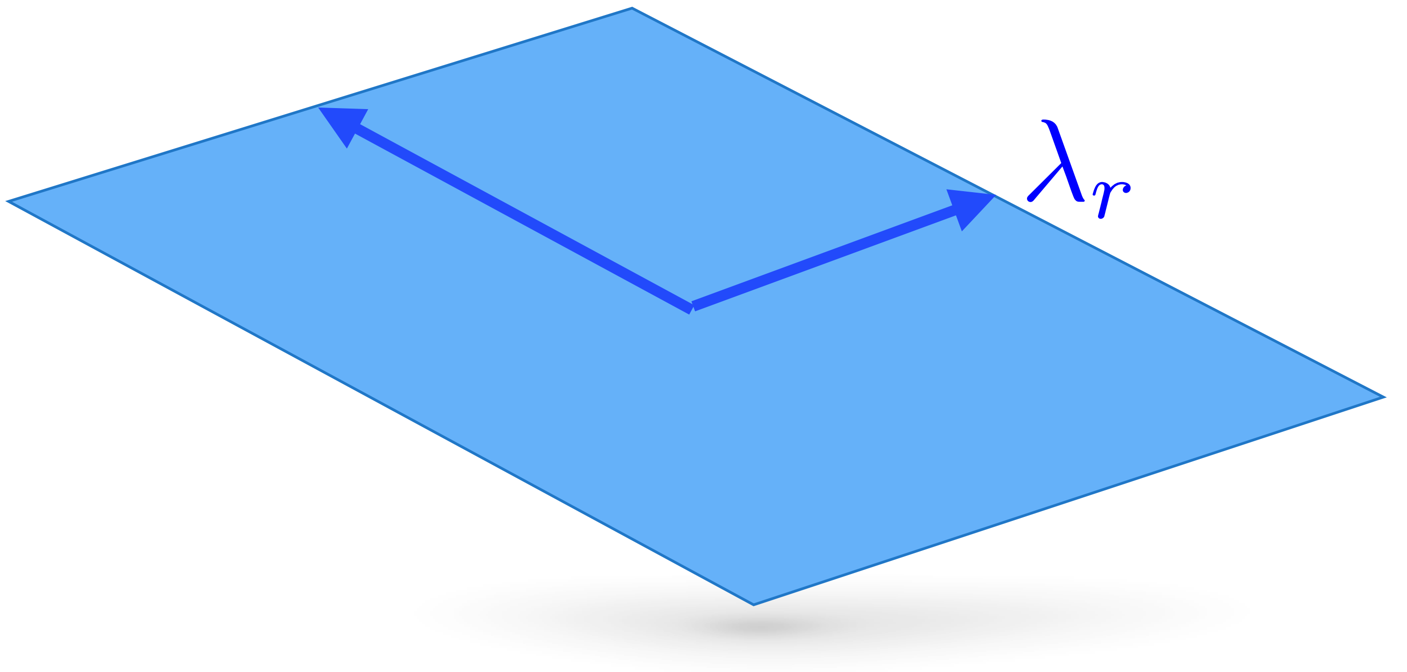

A flavor of the proof

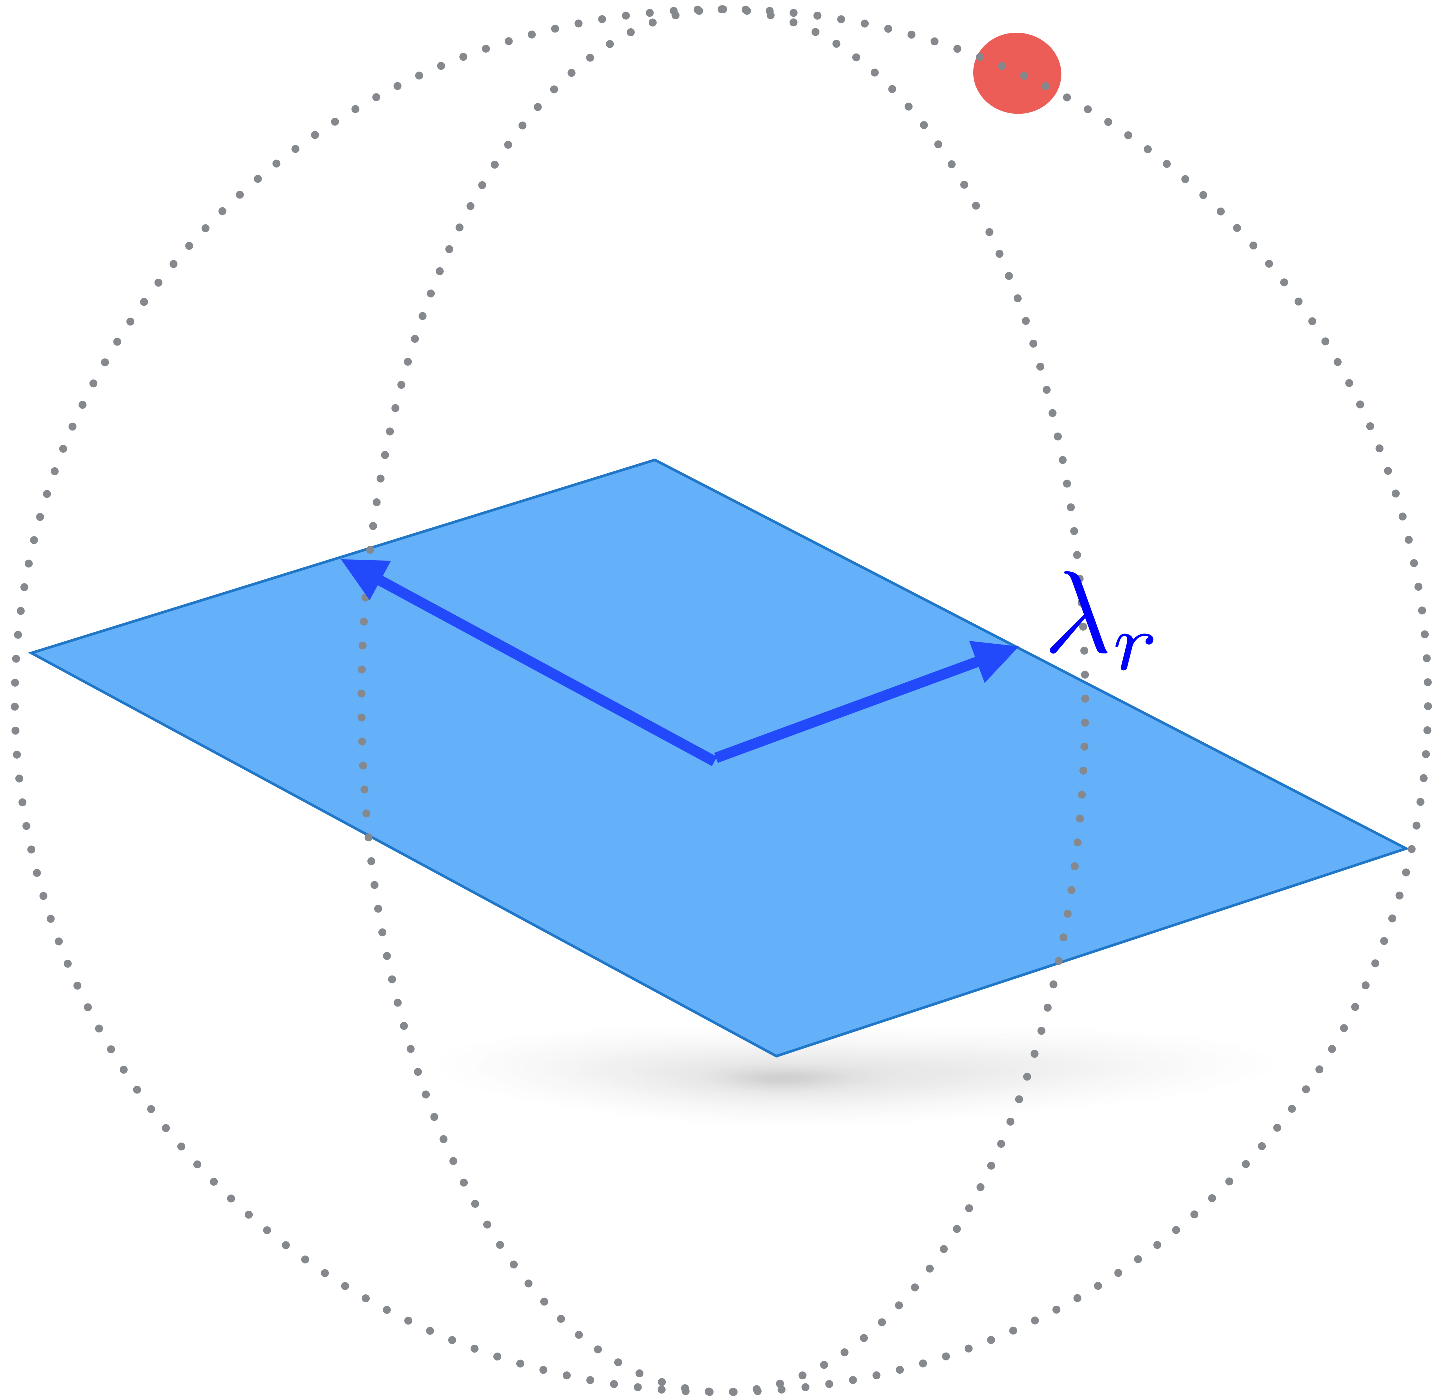

A flavor of the proof

Fix the magnitude of

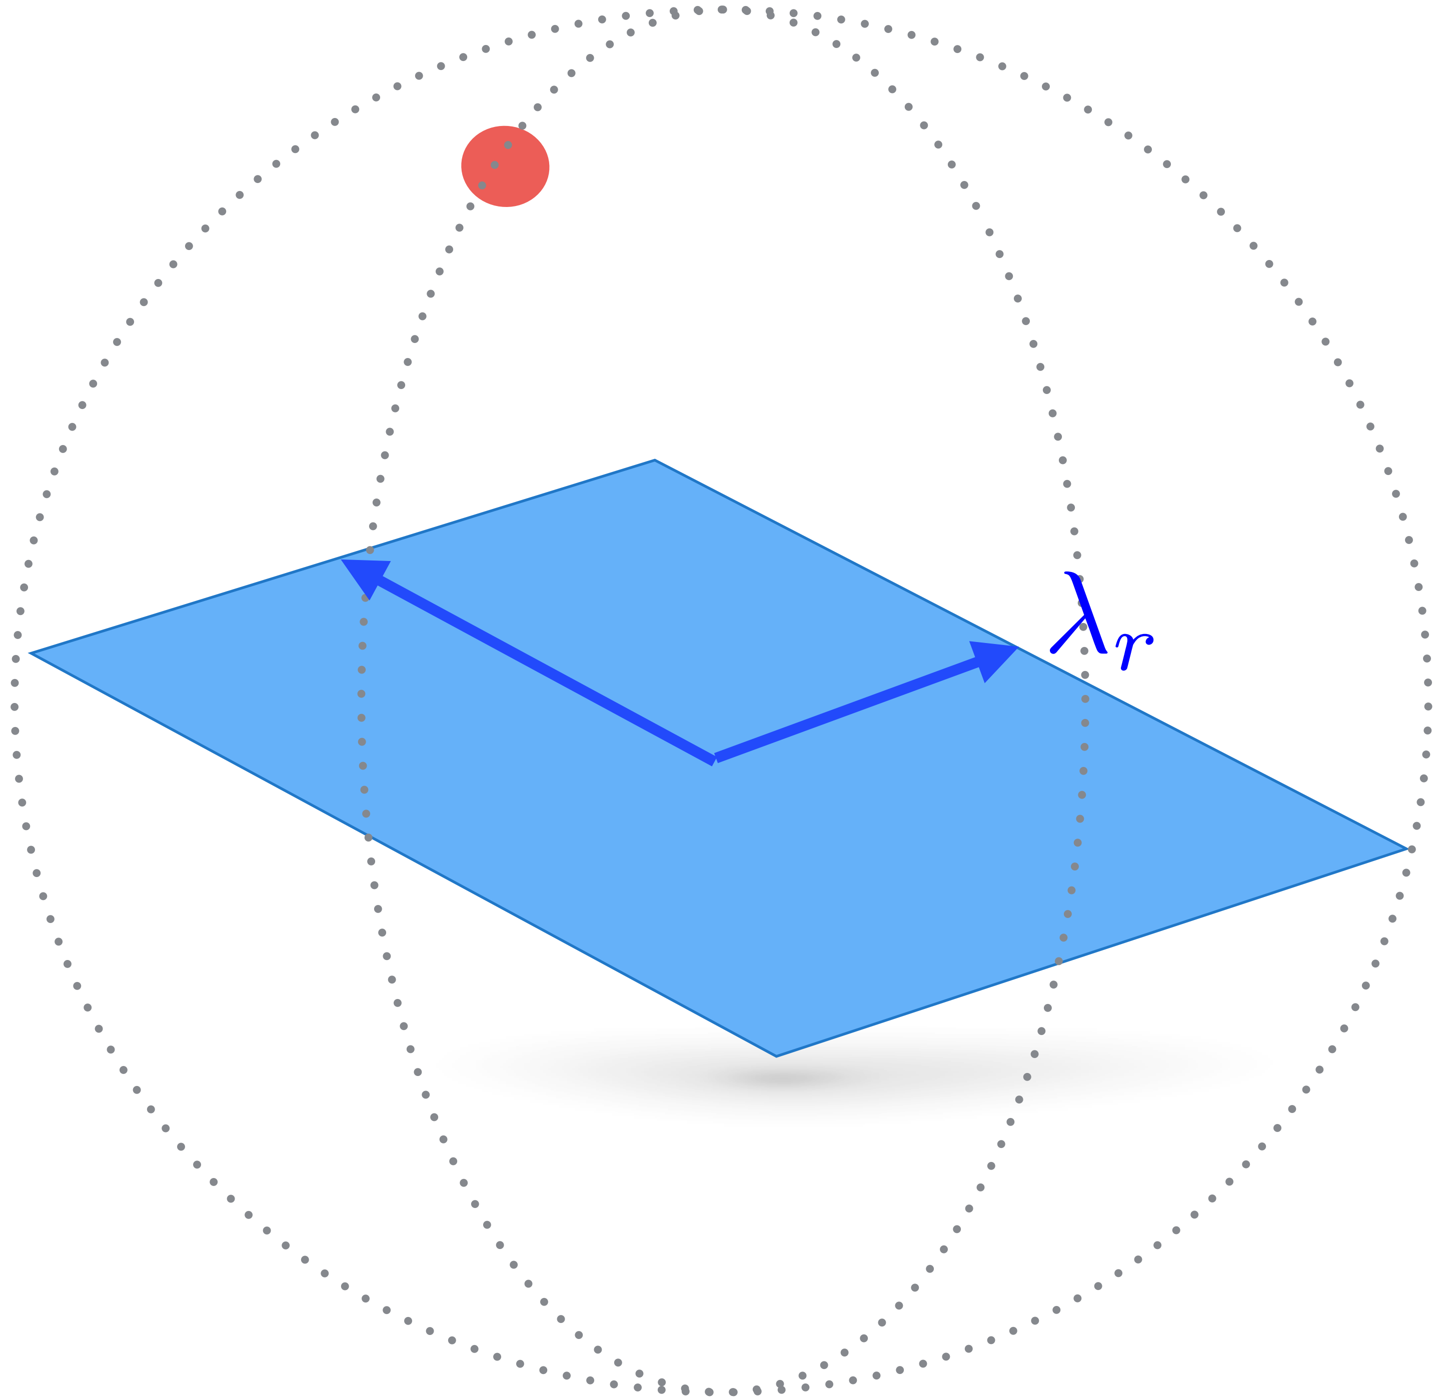

A flavor of the proof

Fix the magnitude of 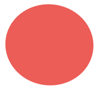

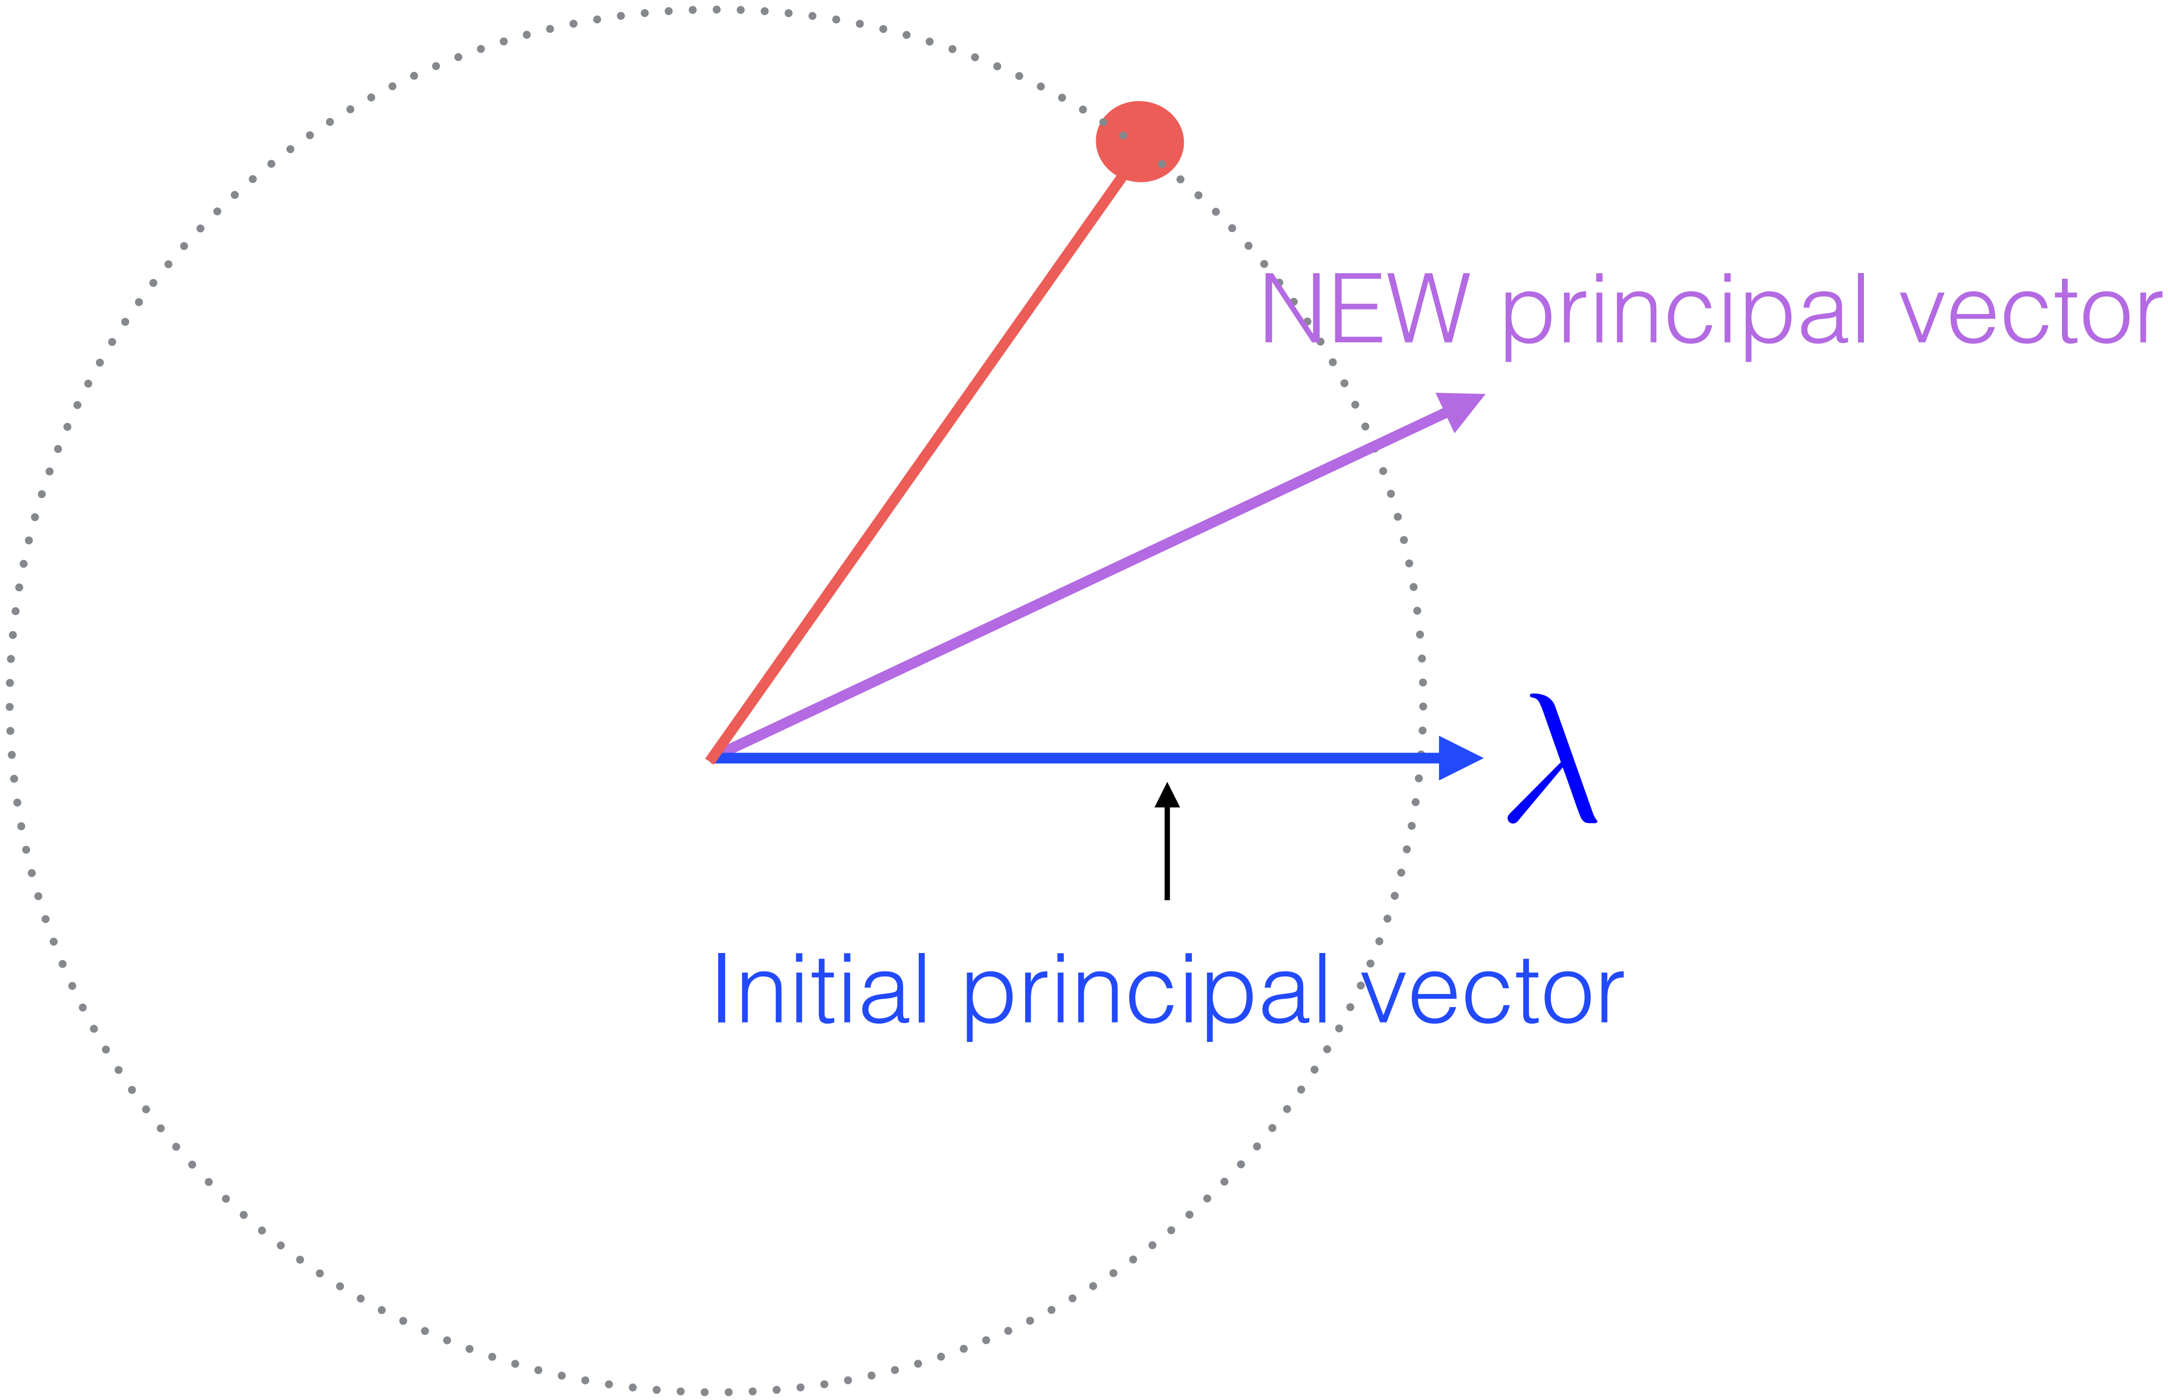

A flavor of the proof

Larger vectors are harder to tilt

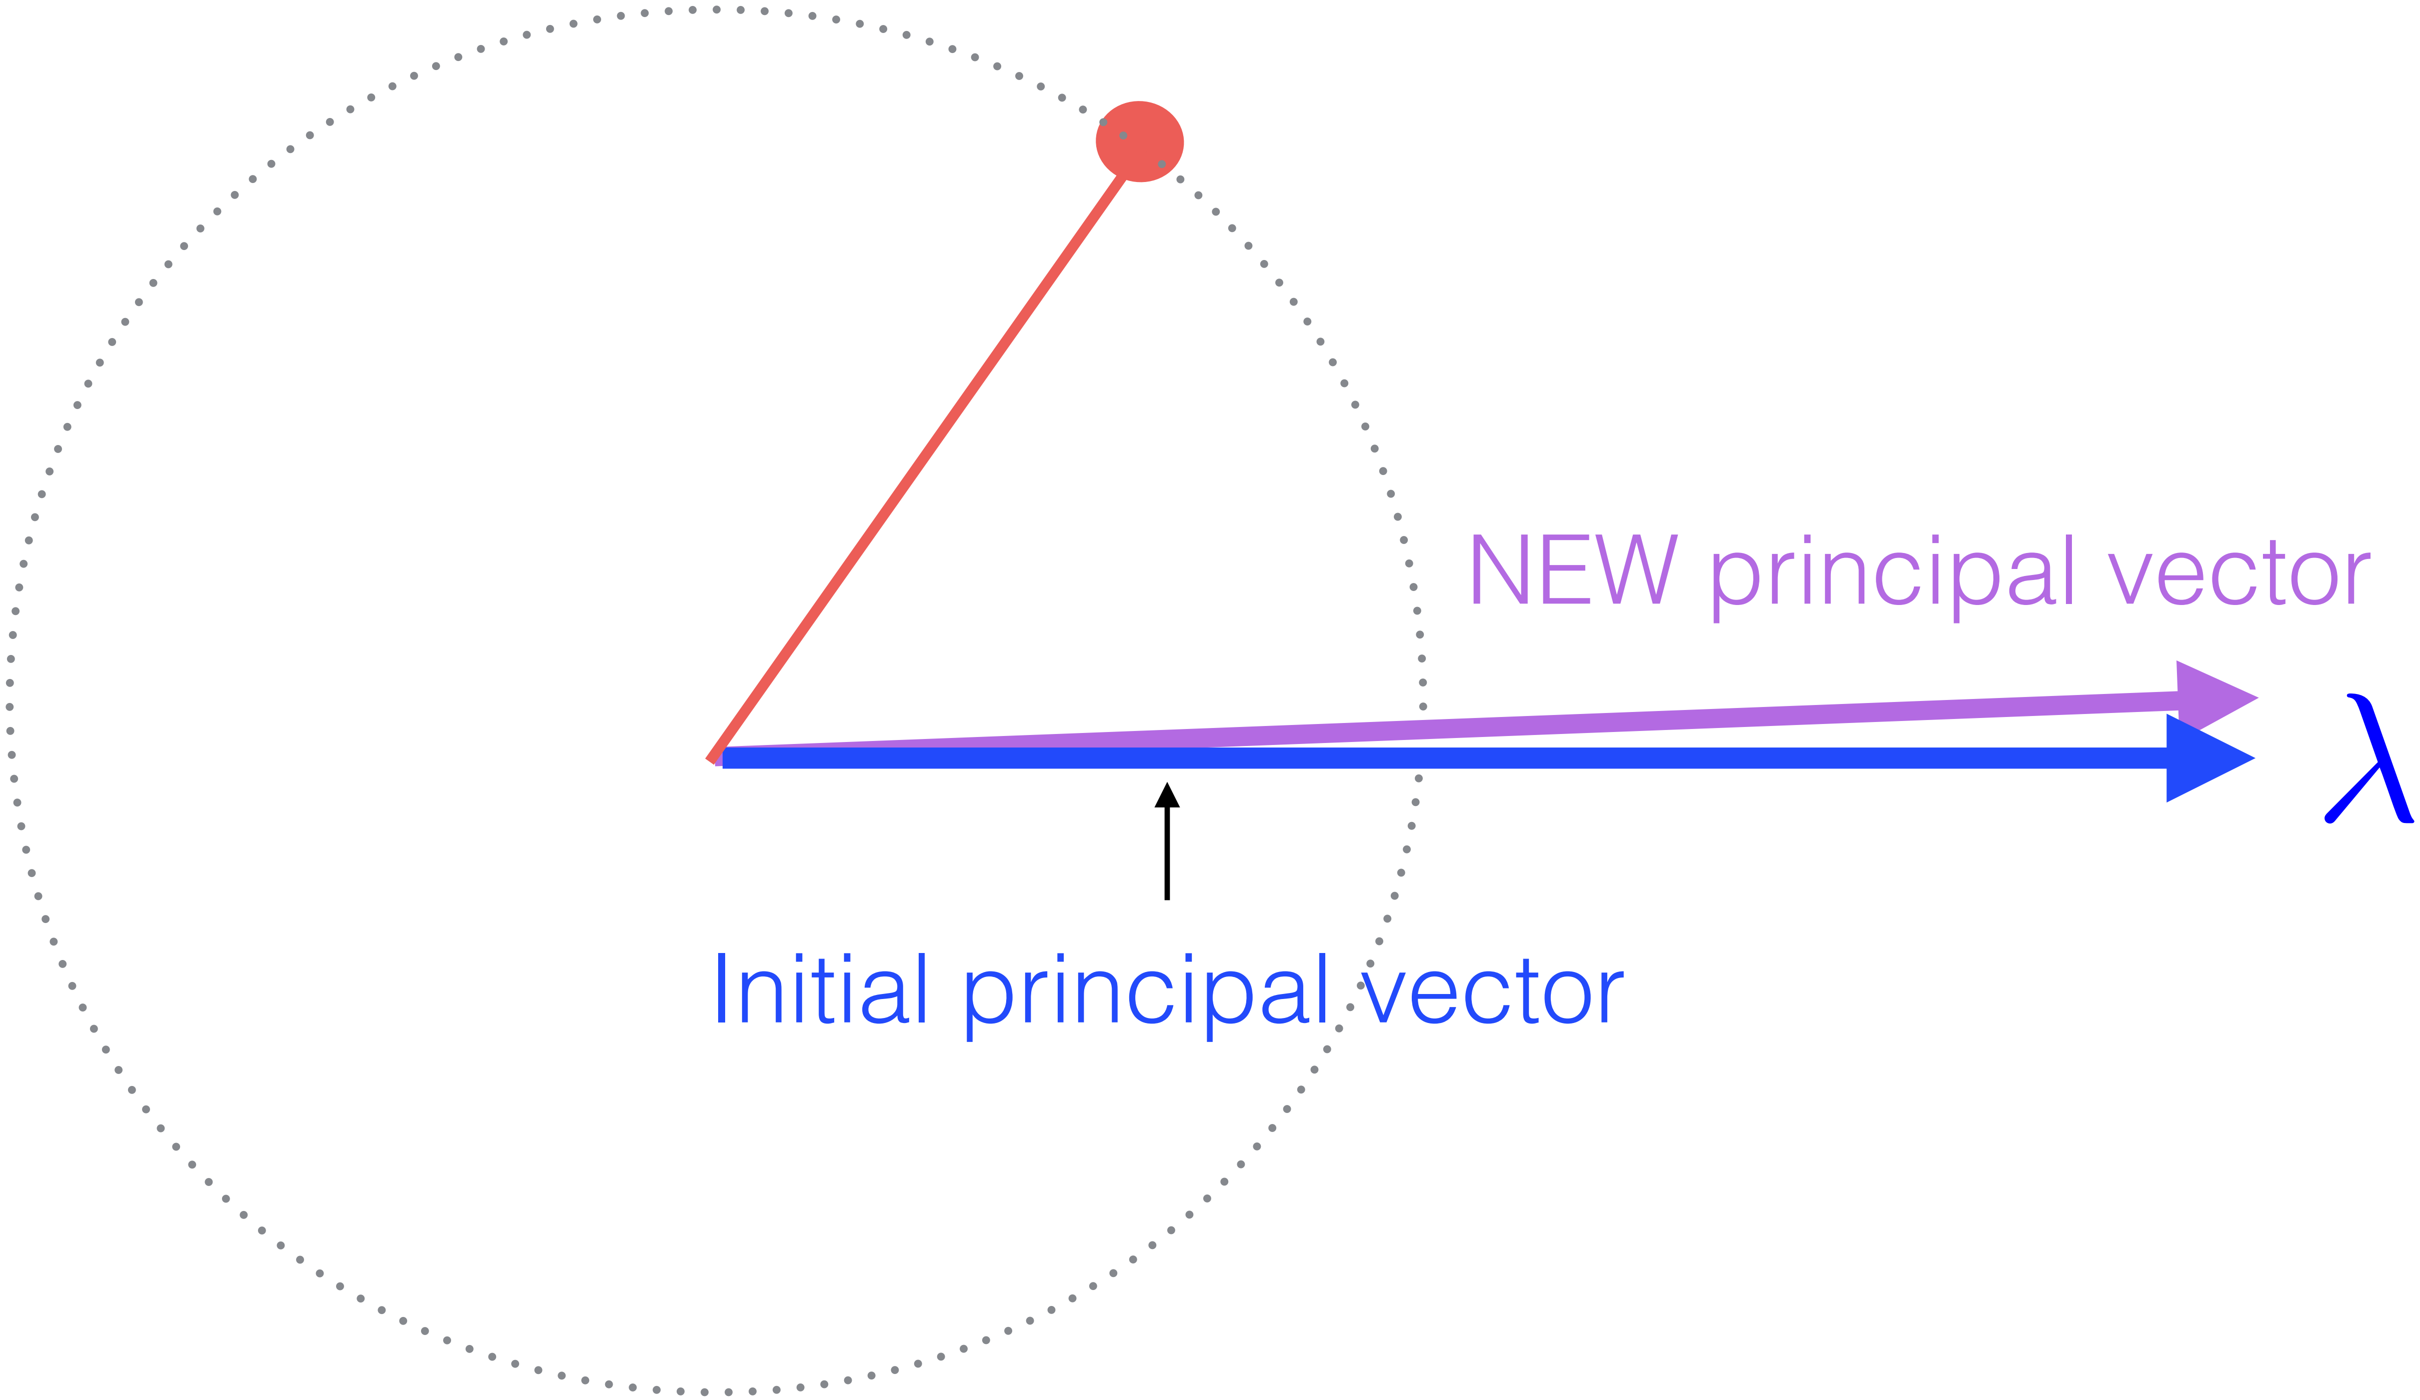

A flavor of the proof

Larger vectors are harder to tilt

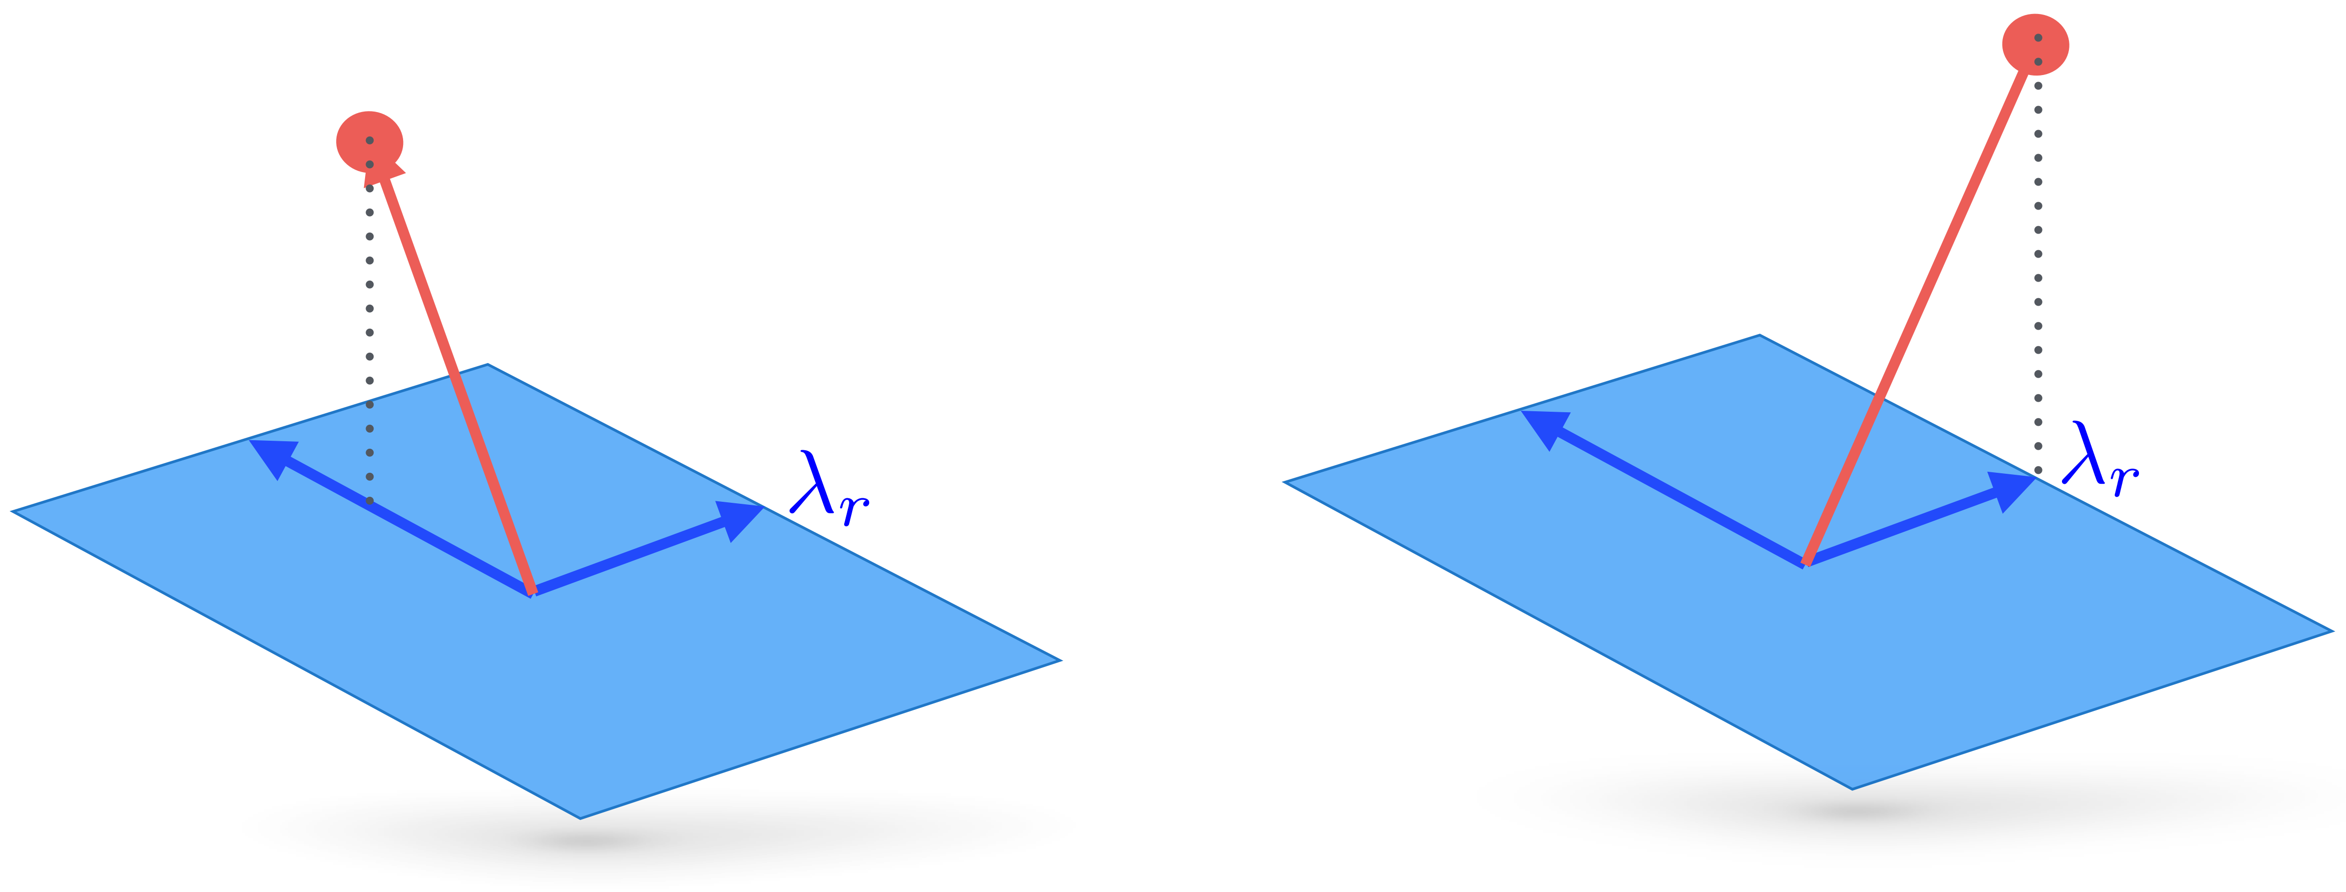

A flavor of the proof

Smaller directions can be tilted more

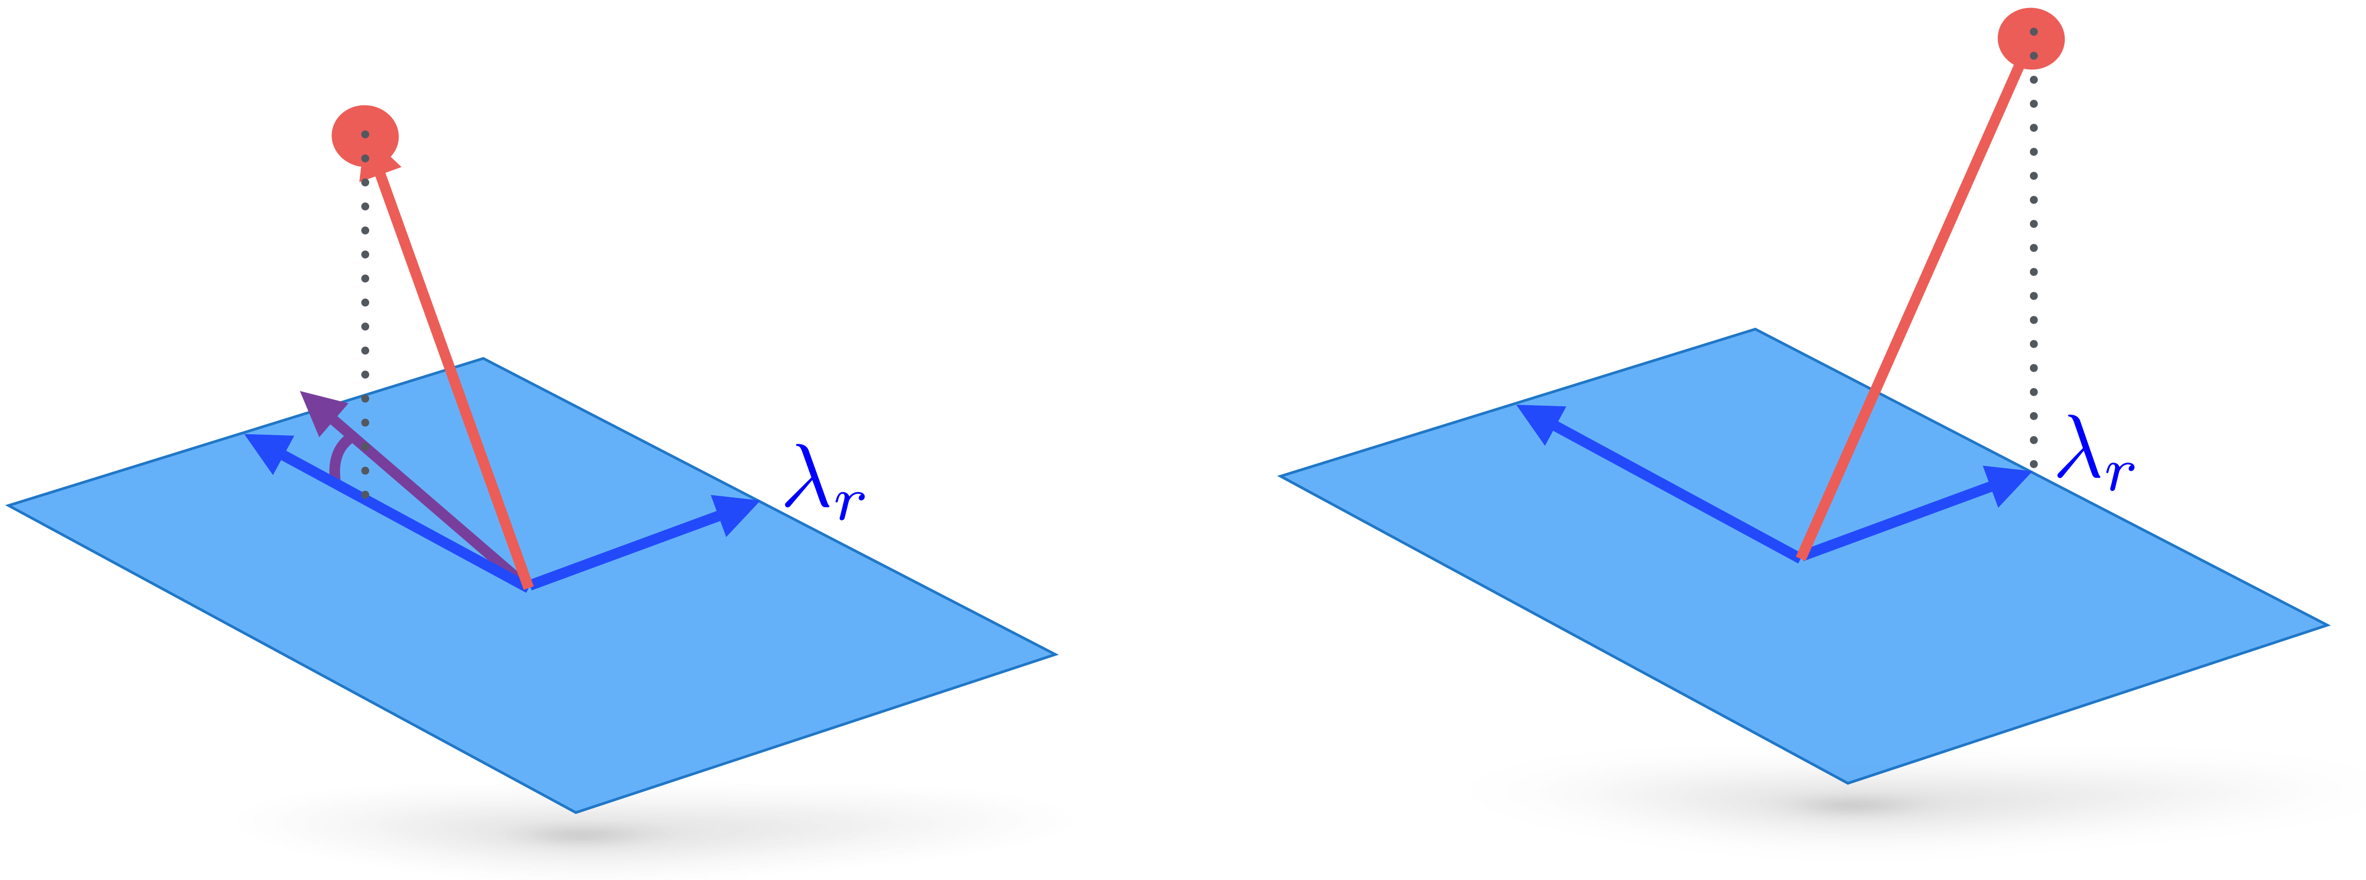

A flavor of the proof

Smaller directions can be tilted more

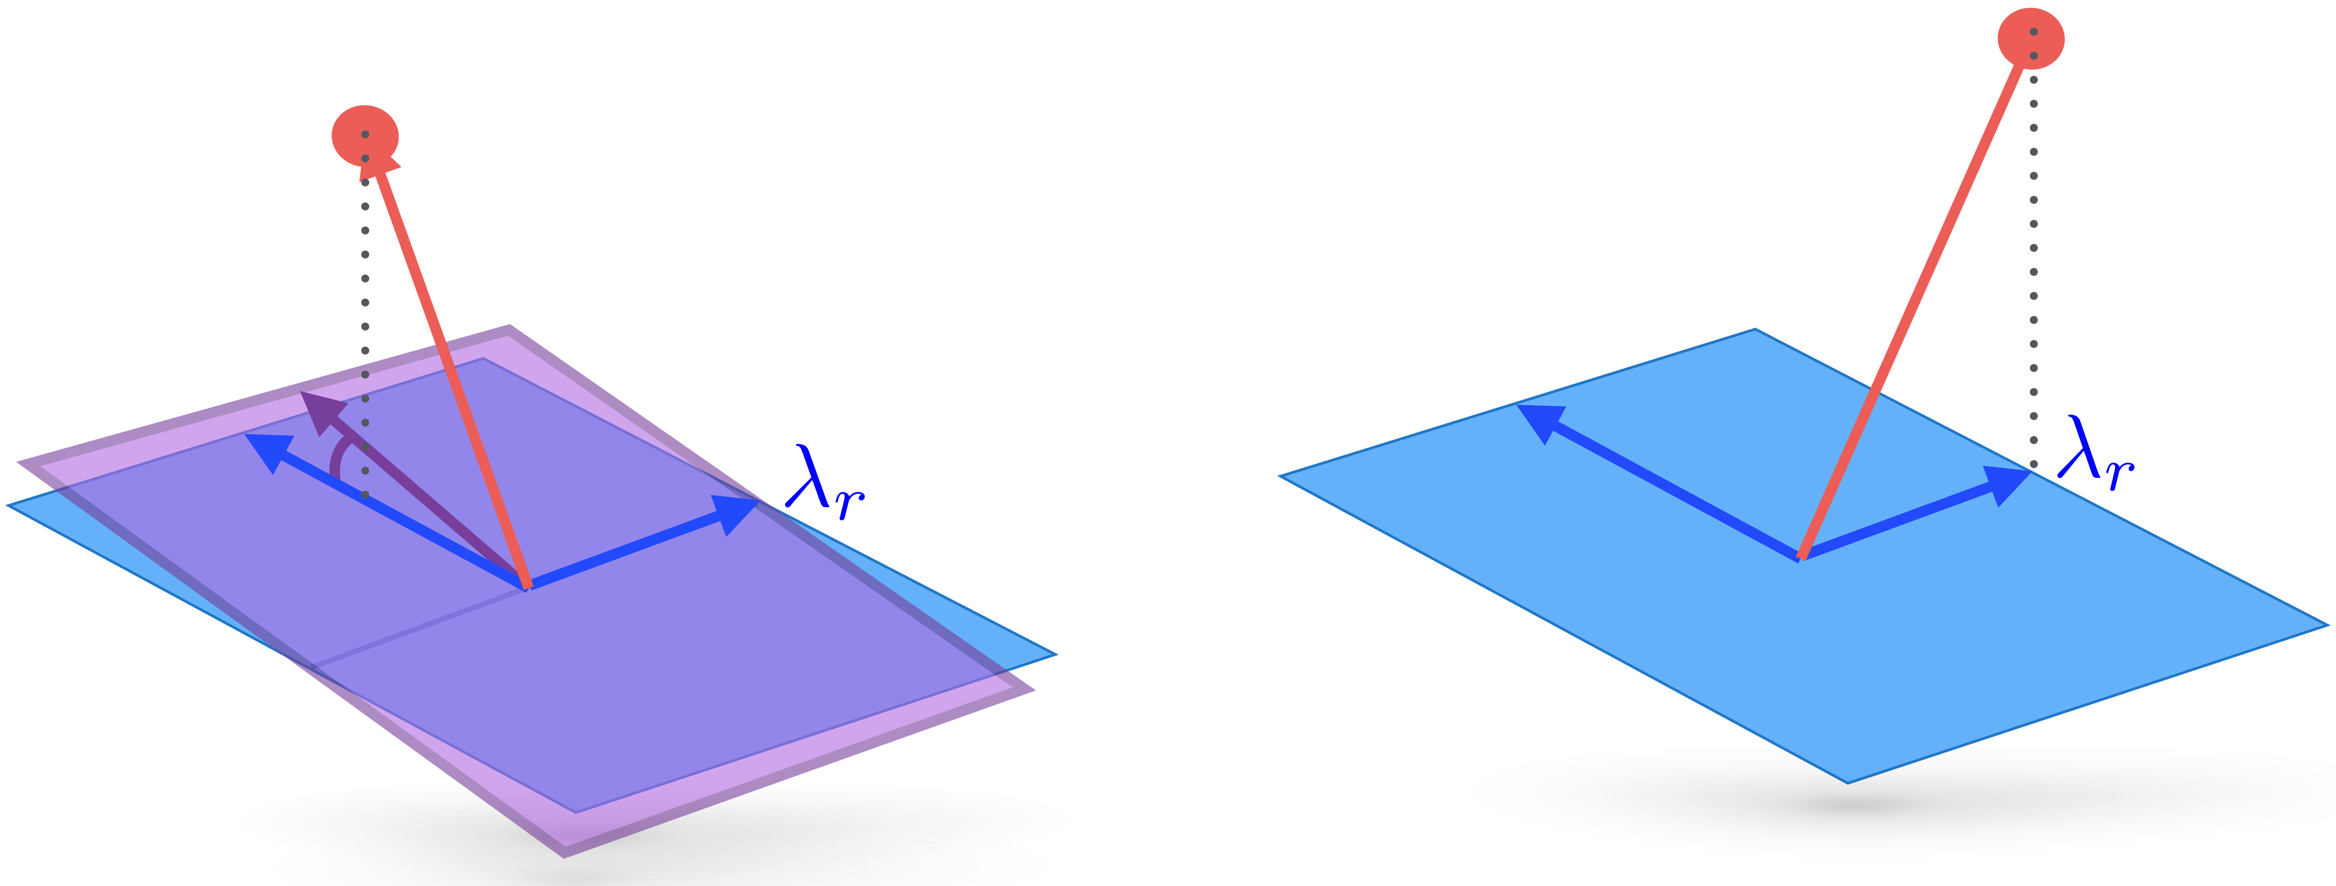

A flavor of the proof

Smaller directions can be tilted more

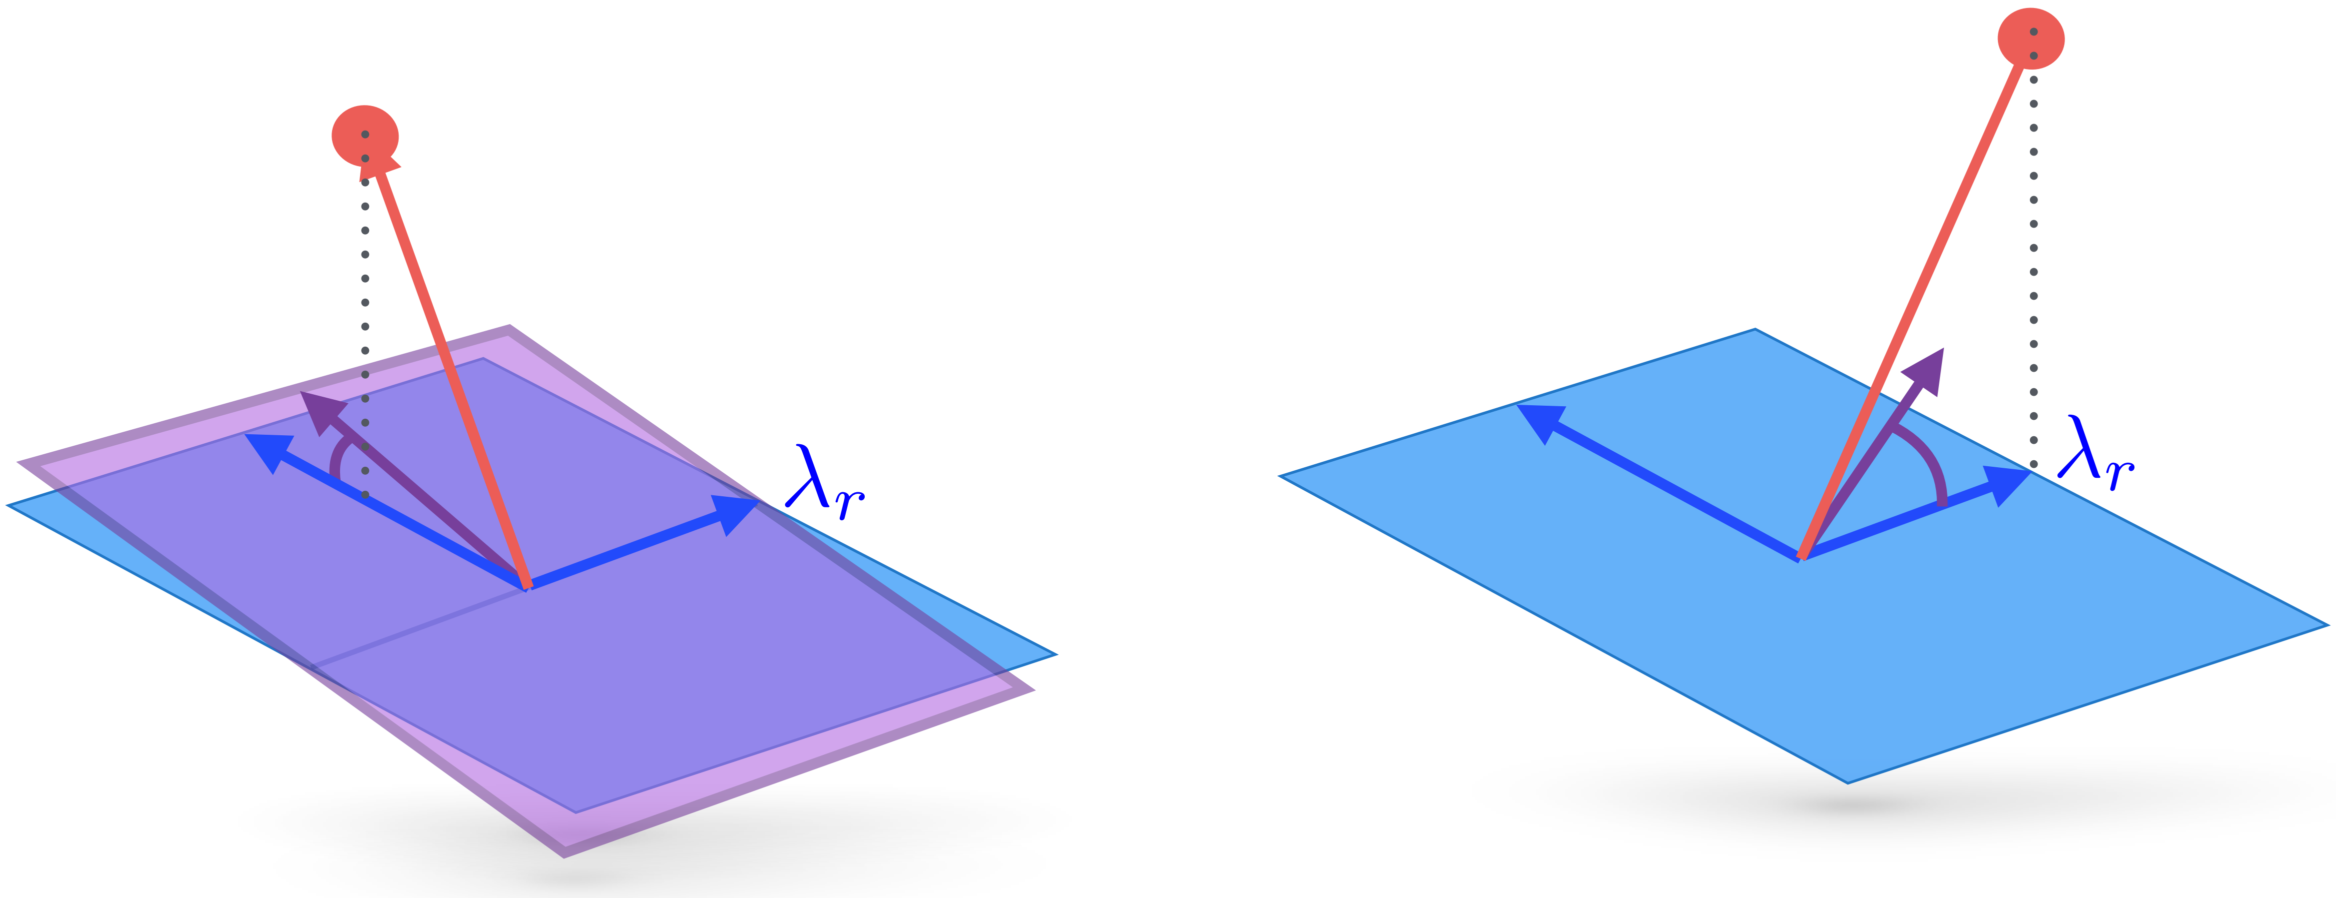

A flavor of the proof

Smaller directions can be tilted more

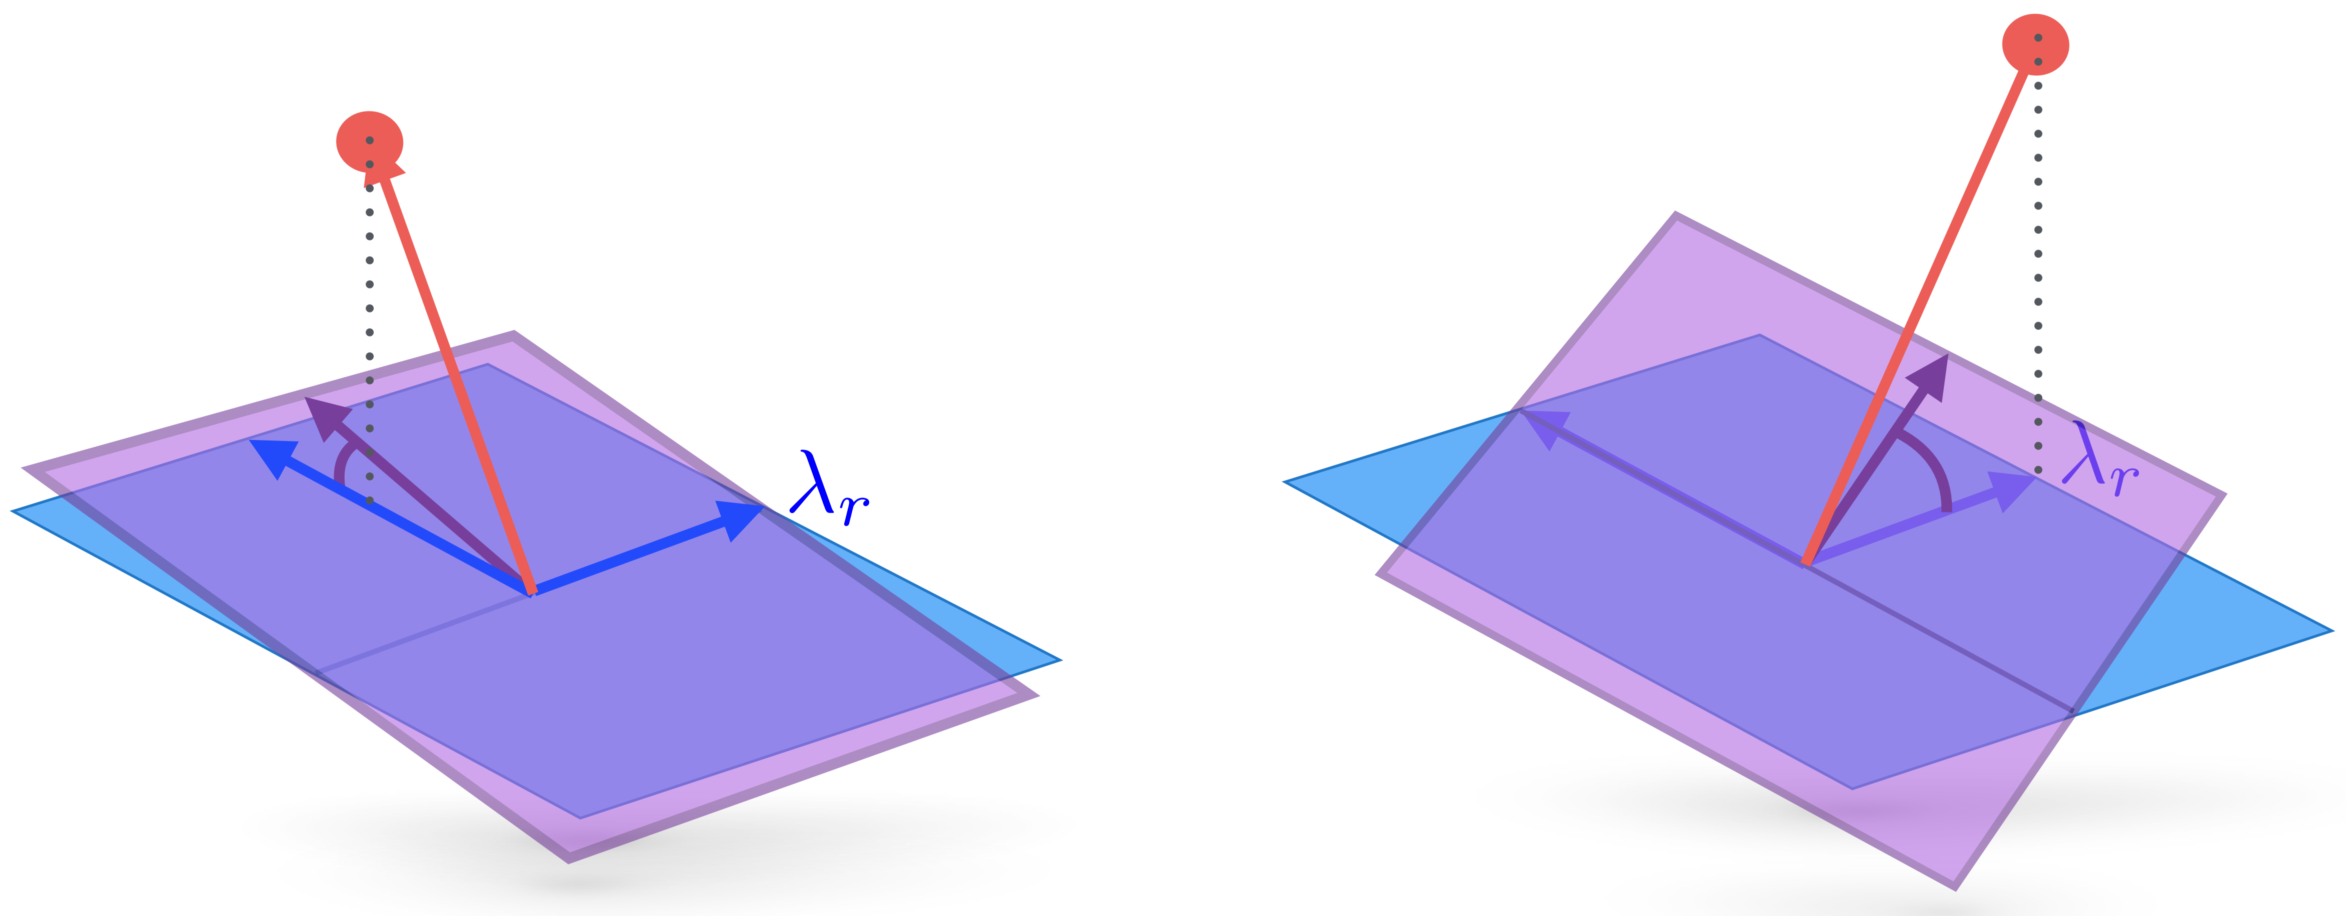

A flavor of the proof

Smaller directions can be tilted more

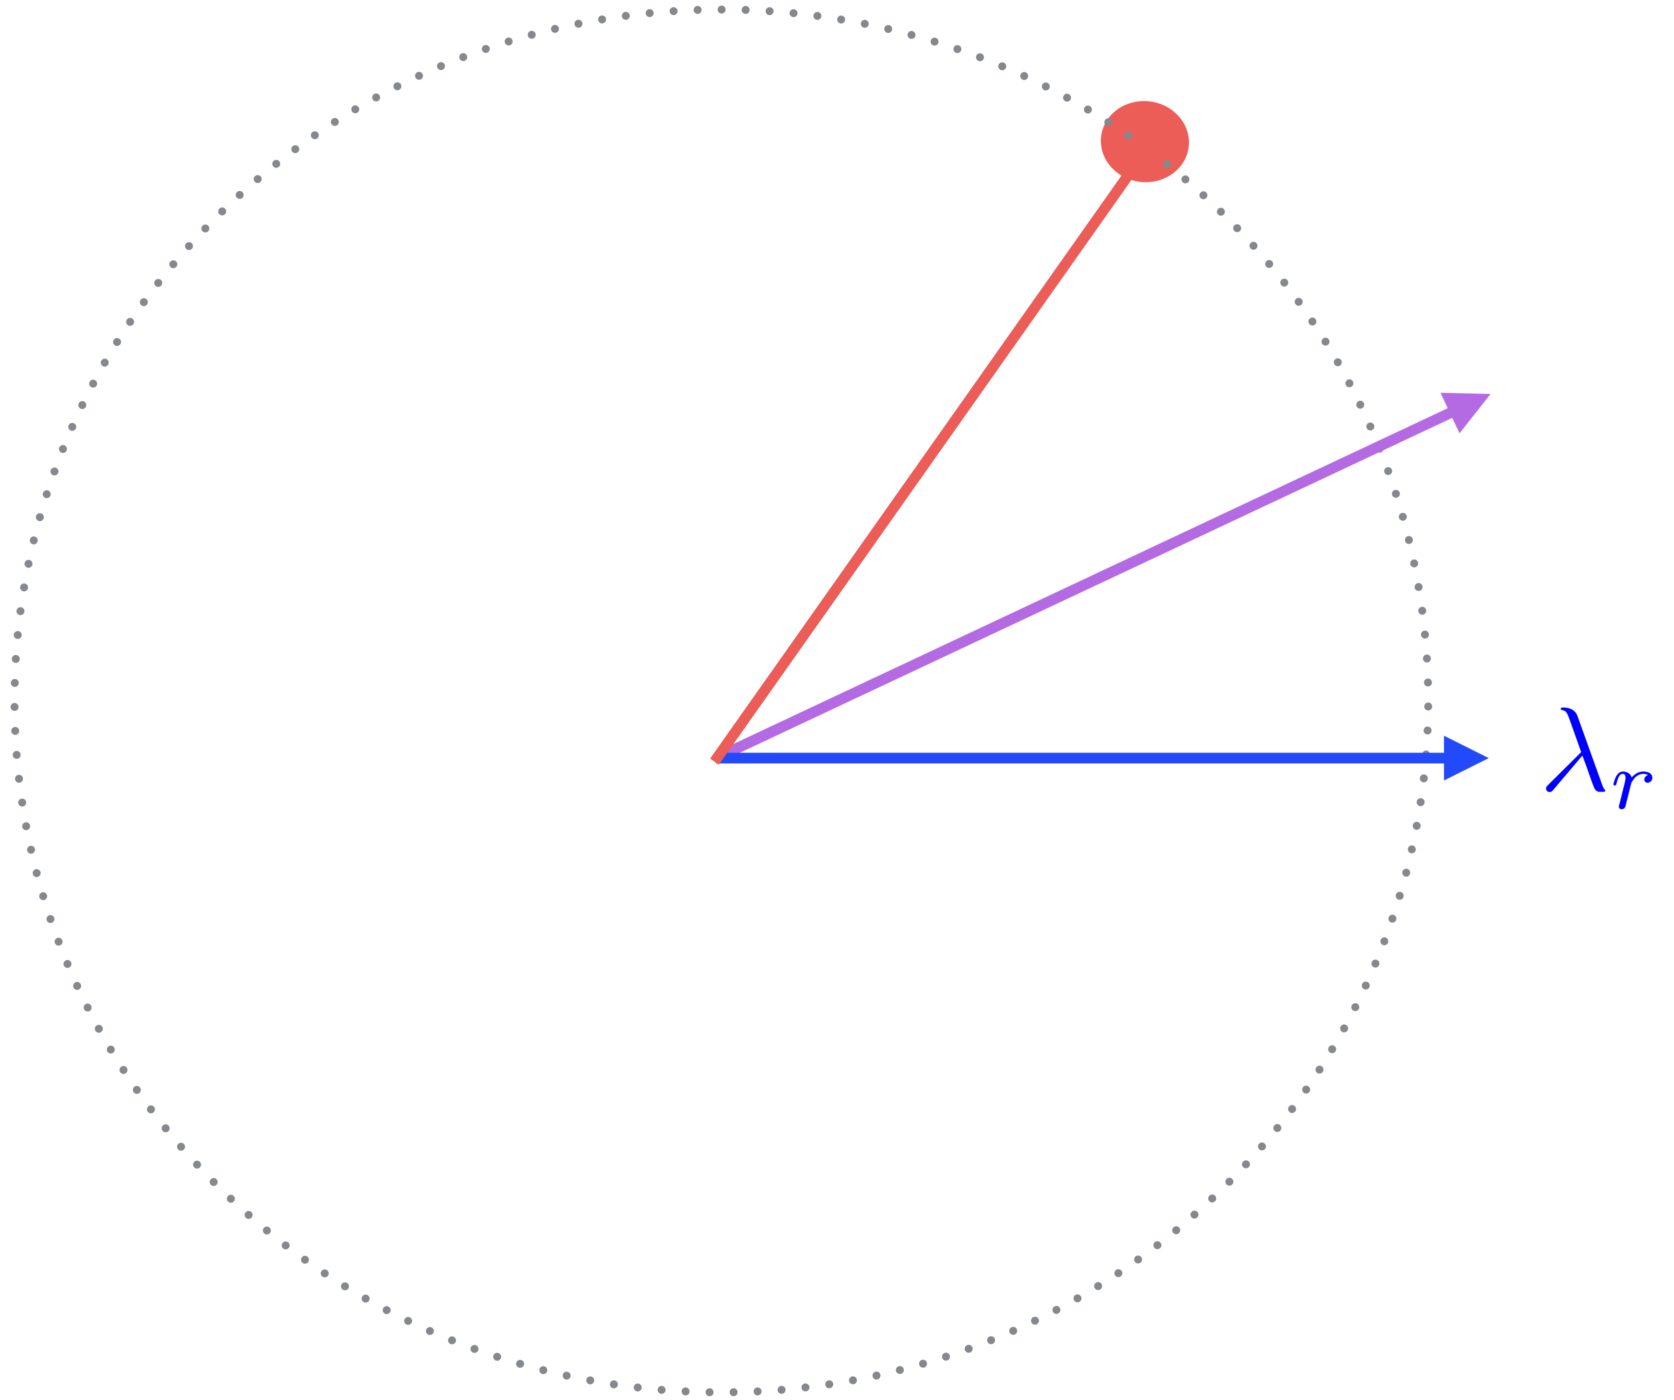

A flavor of the proof

How do we *tilt* maximally the smallest direction?

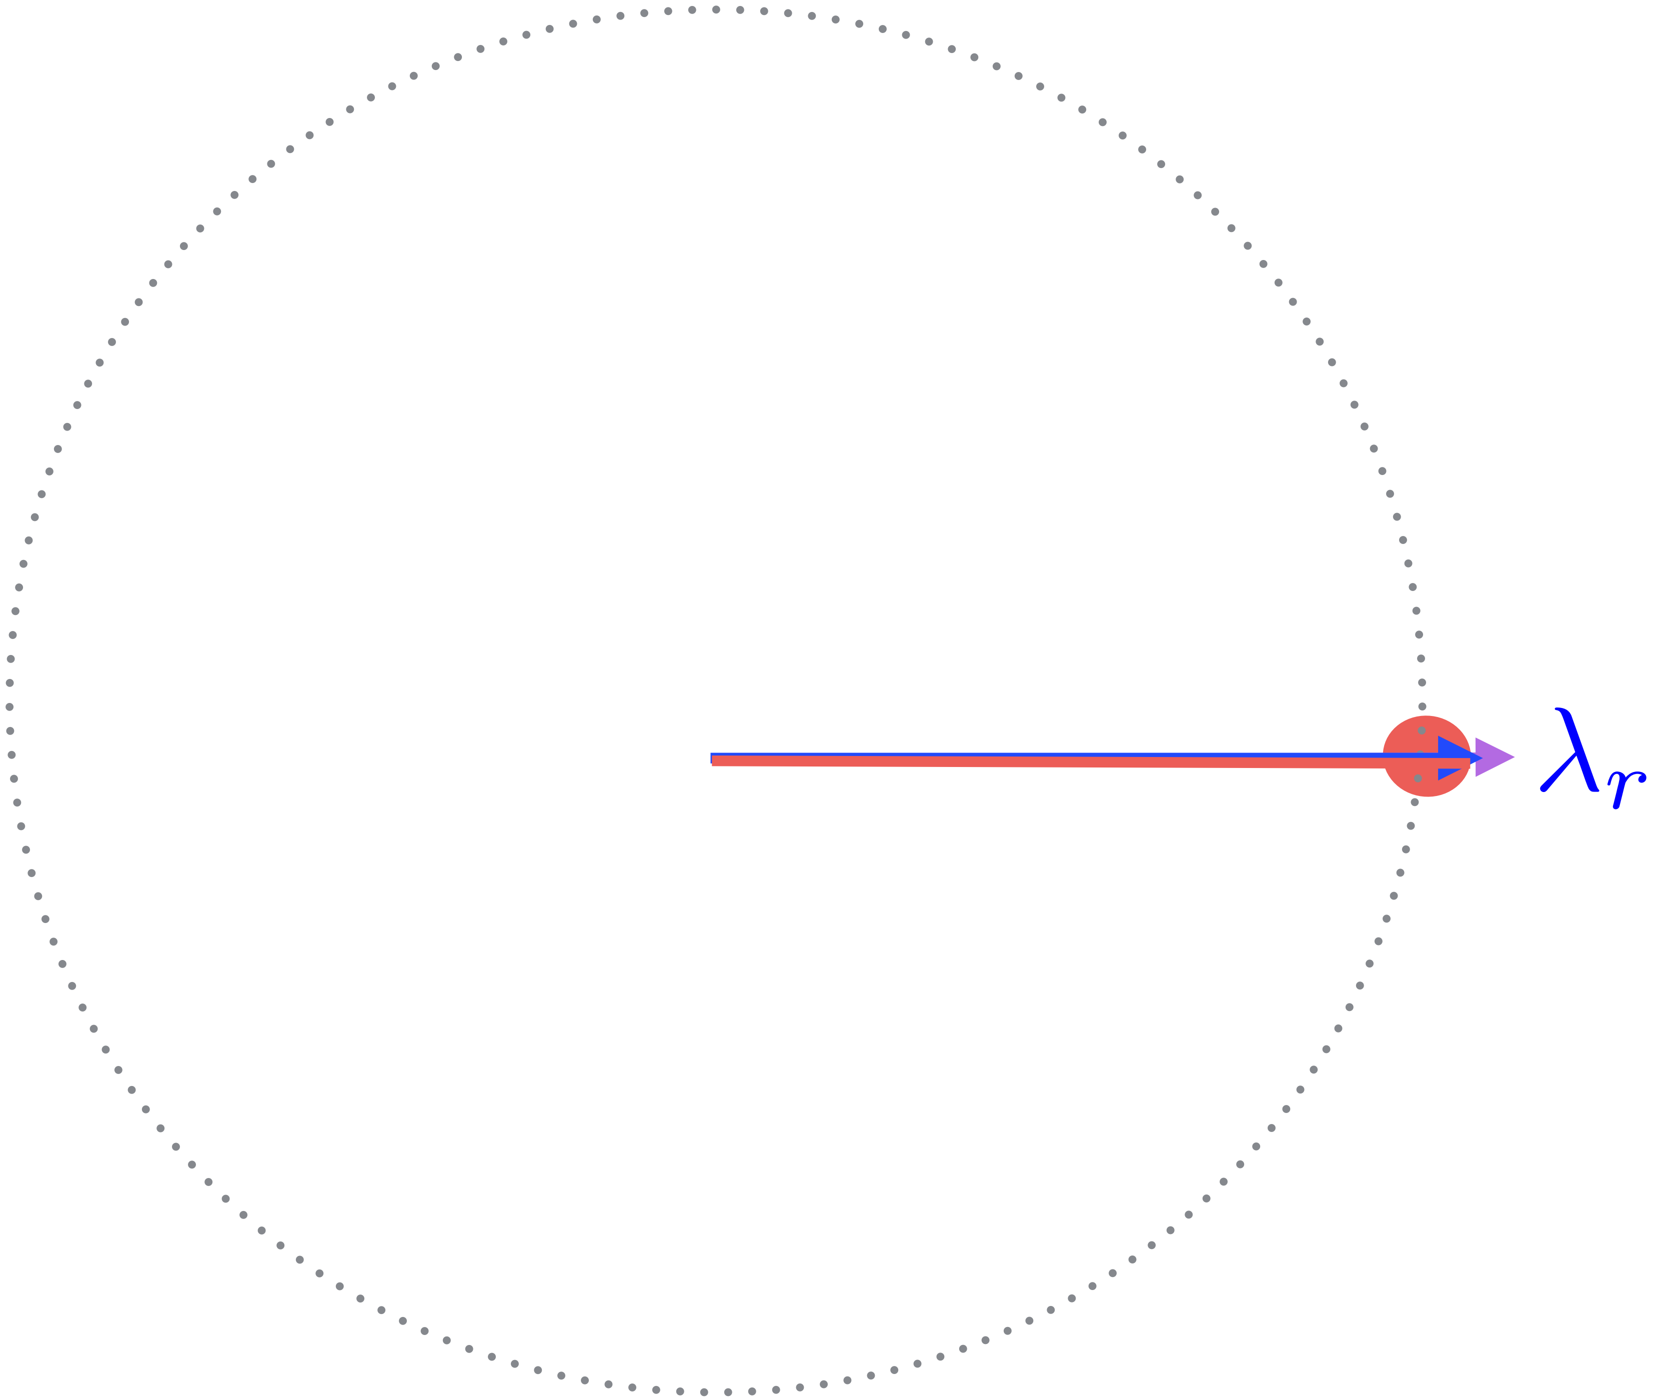

A flavor of the proof

How do we *tilt* maximally the smallest direction?

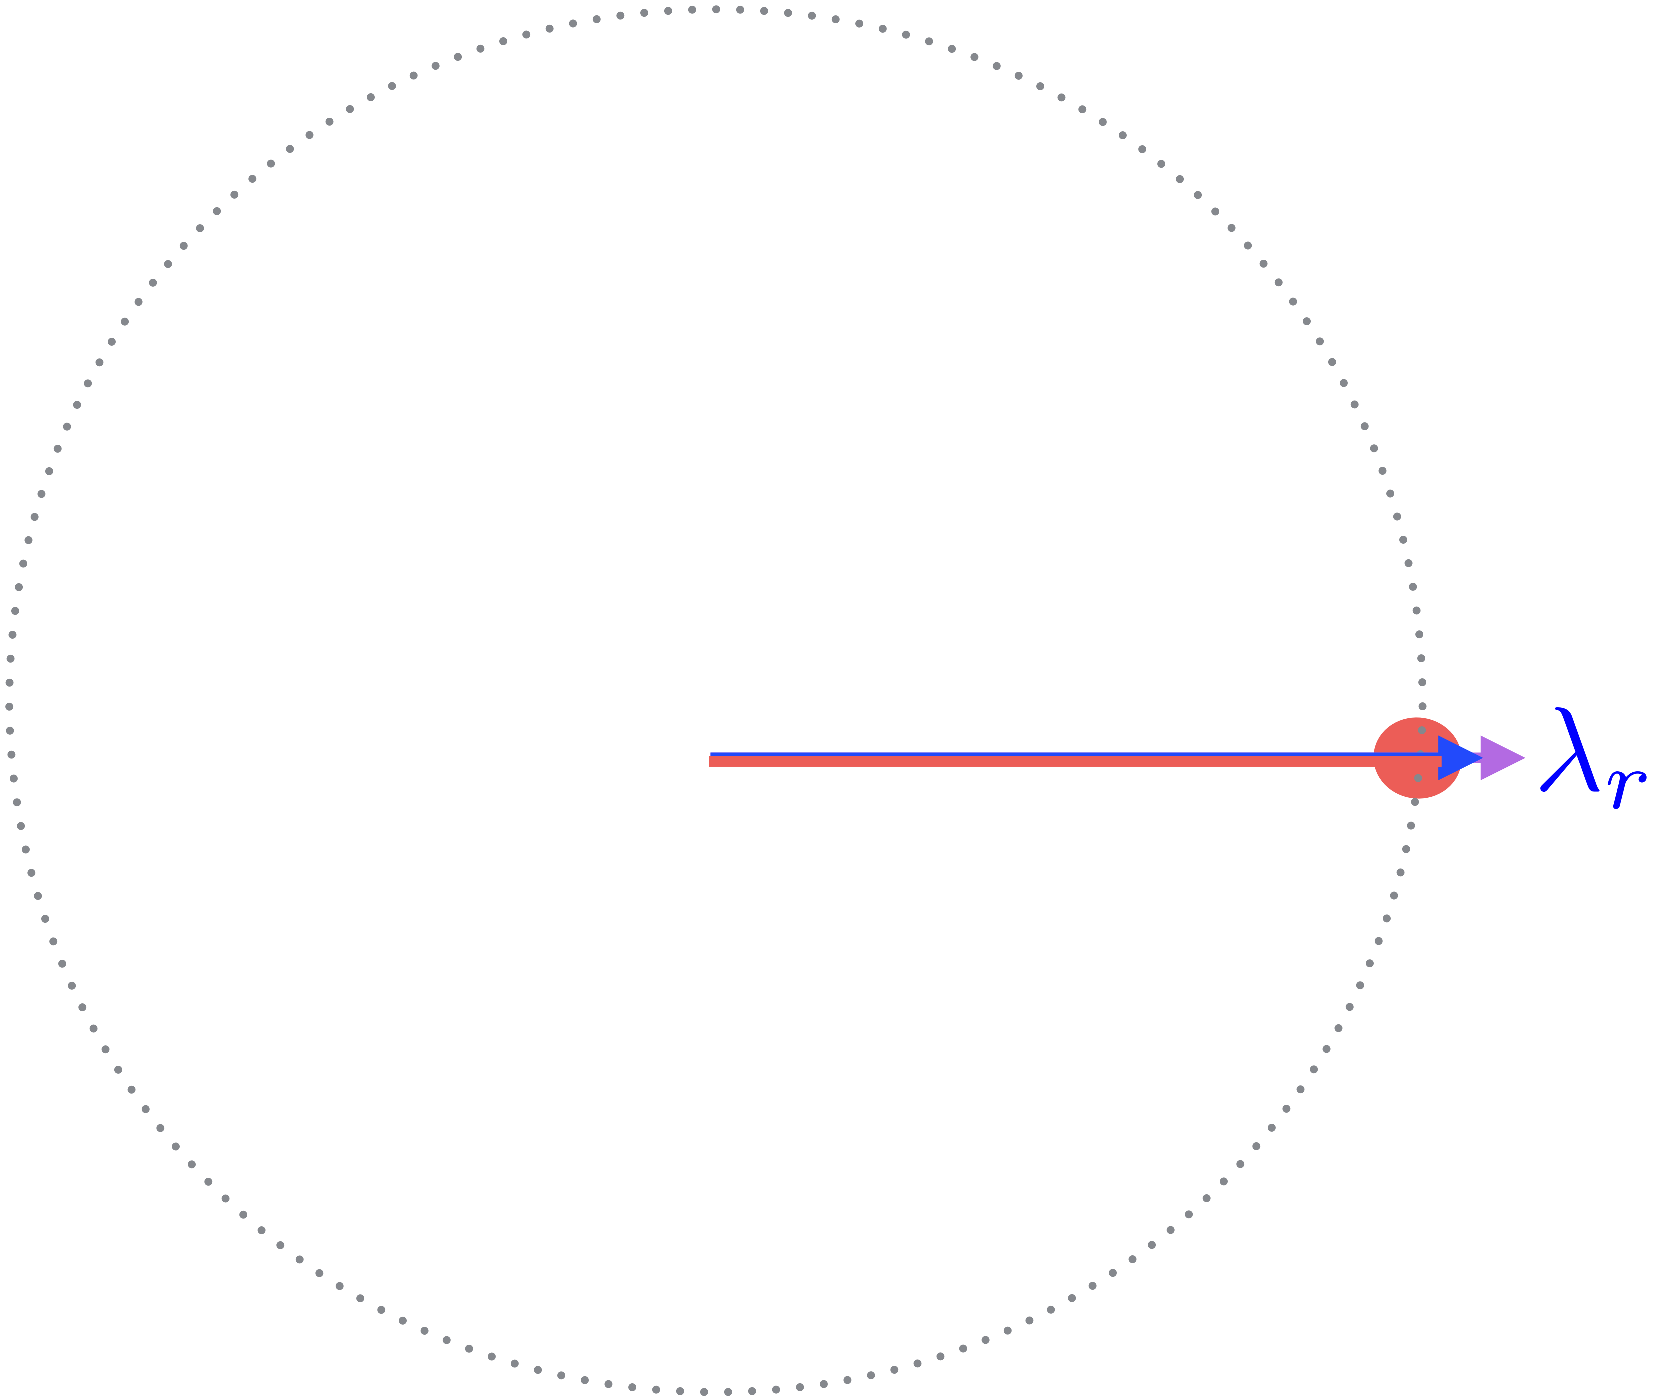

A flavor of the proof

How do we maximally *tilt* the smallest direction?

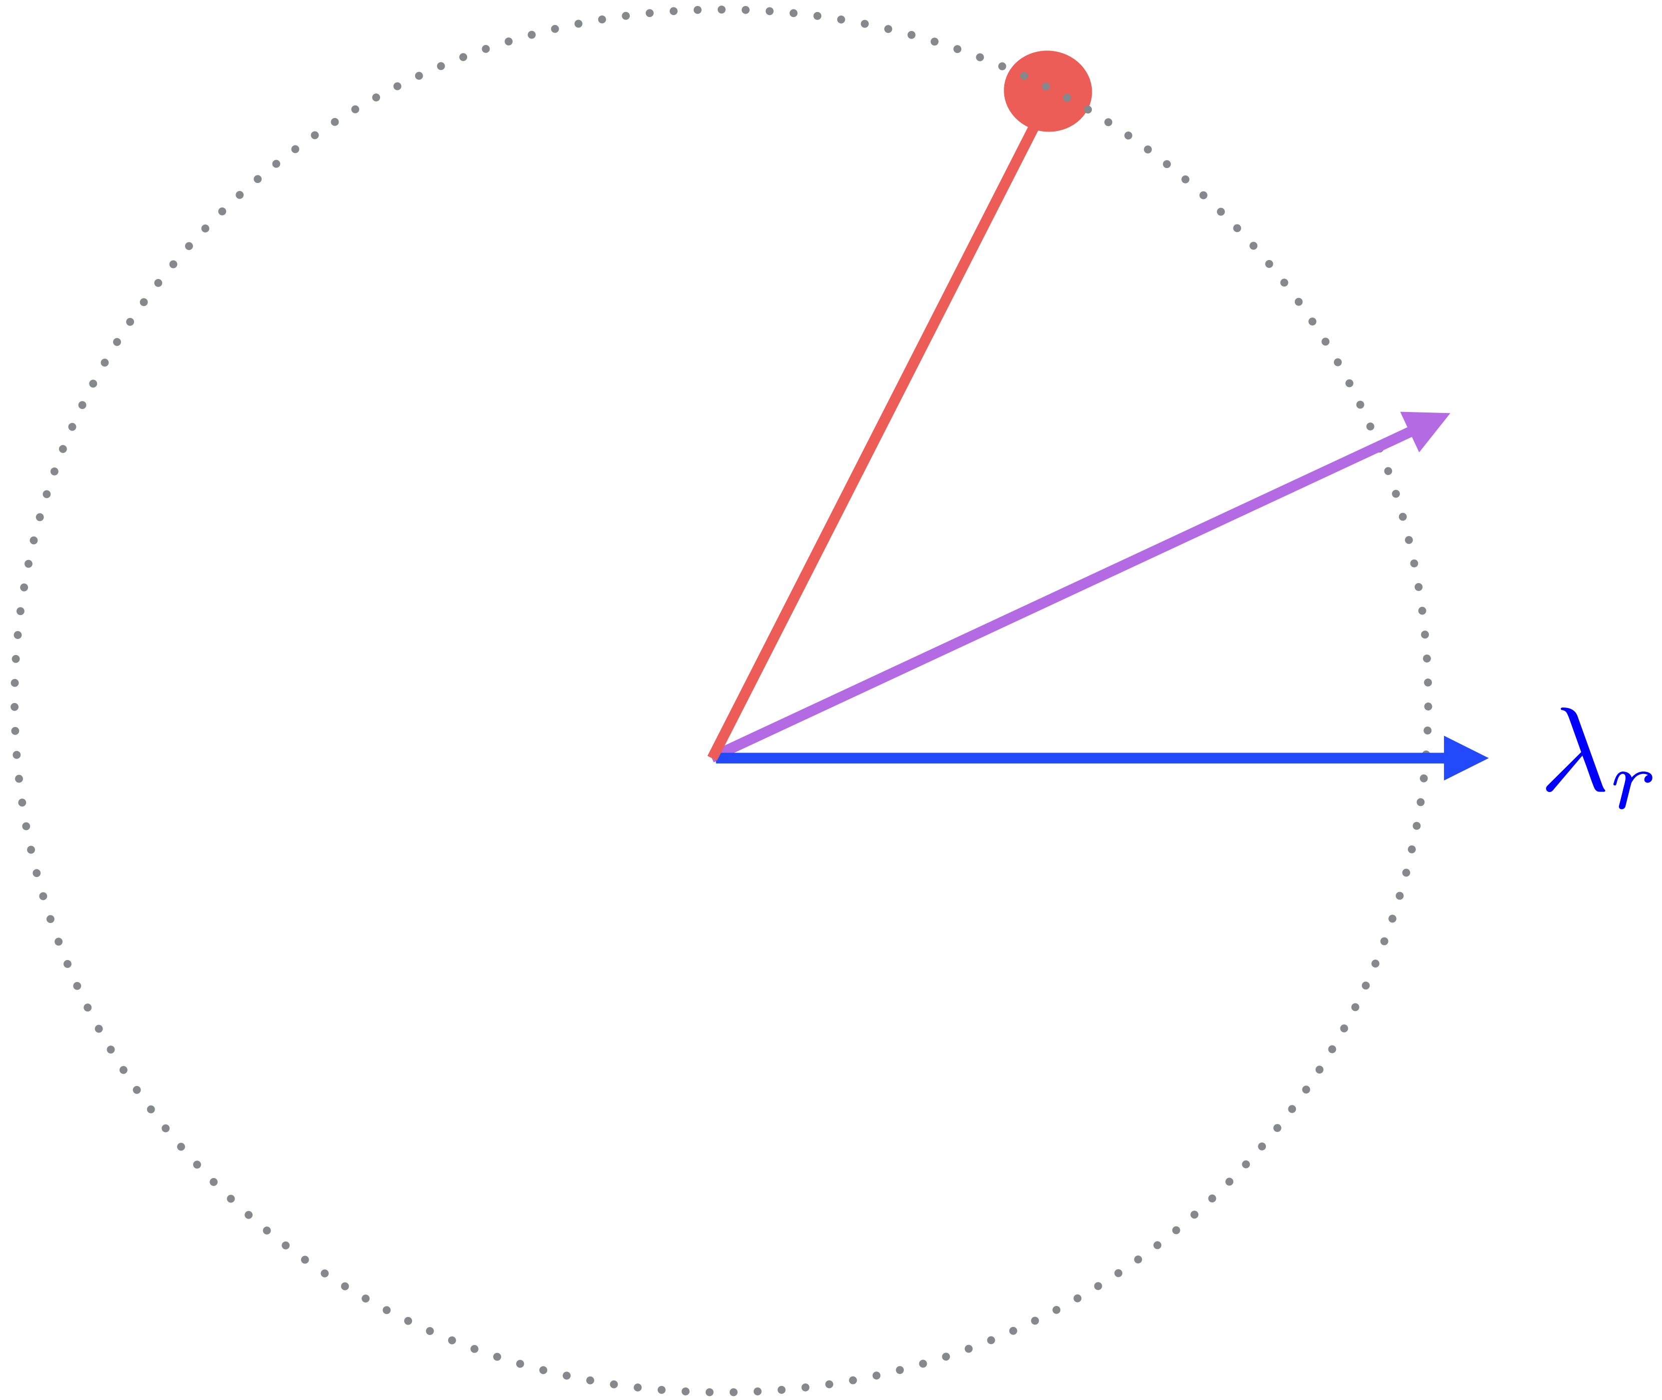

A flavor of the proof

How do we maximally *tilt* the smallest direction?

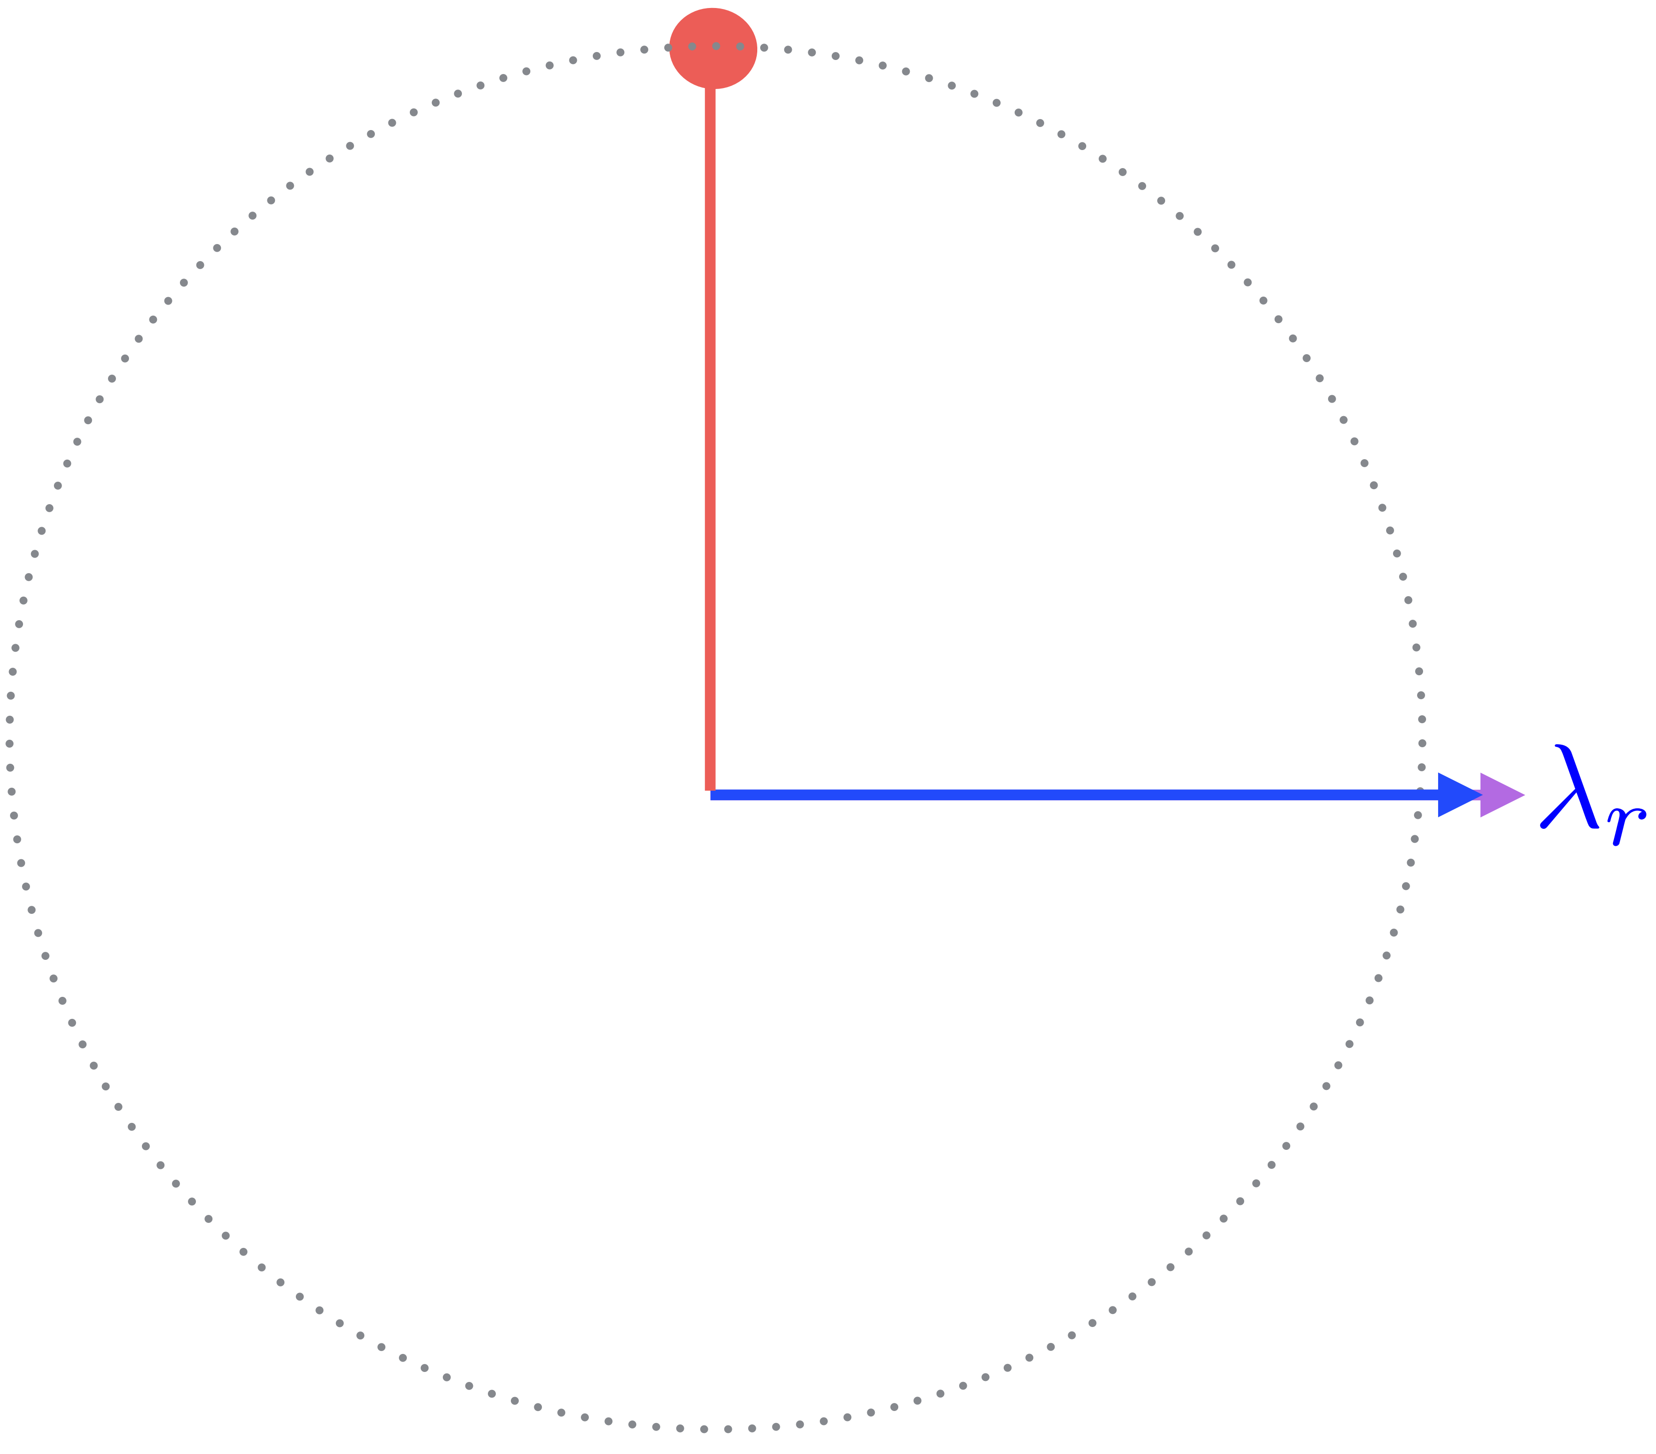

A flavor of the proof

How do we maximally *tilt* the smallest direction?

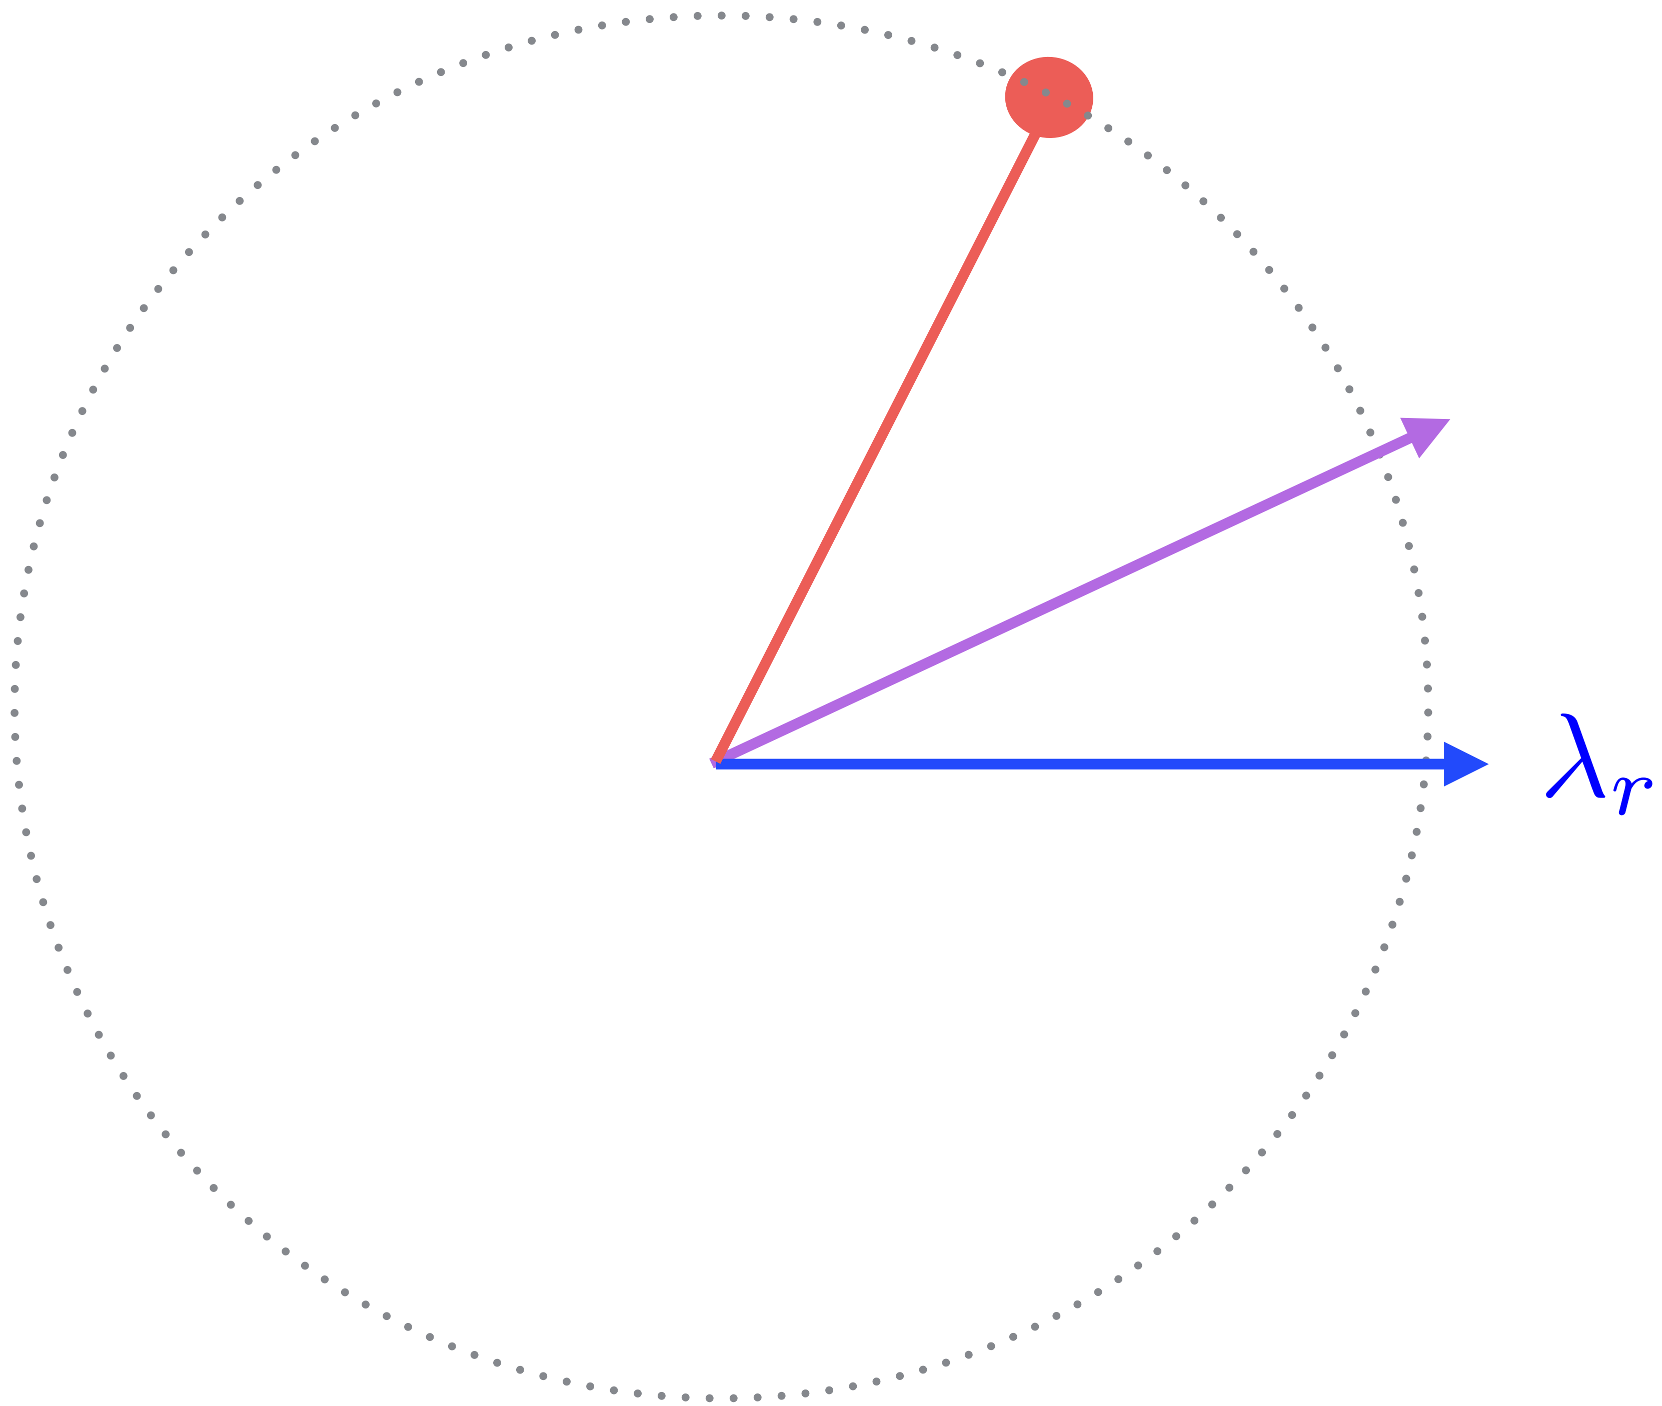

A flavor of the proof

How do we maximally *tilt* the smallest direction?

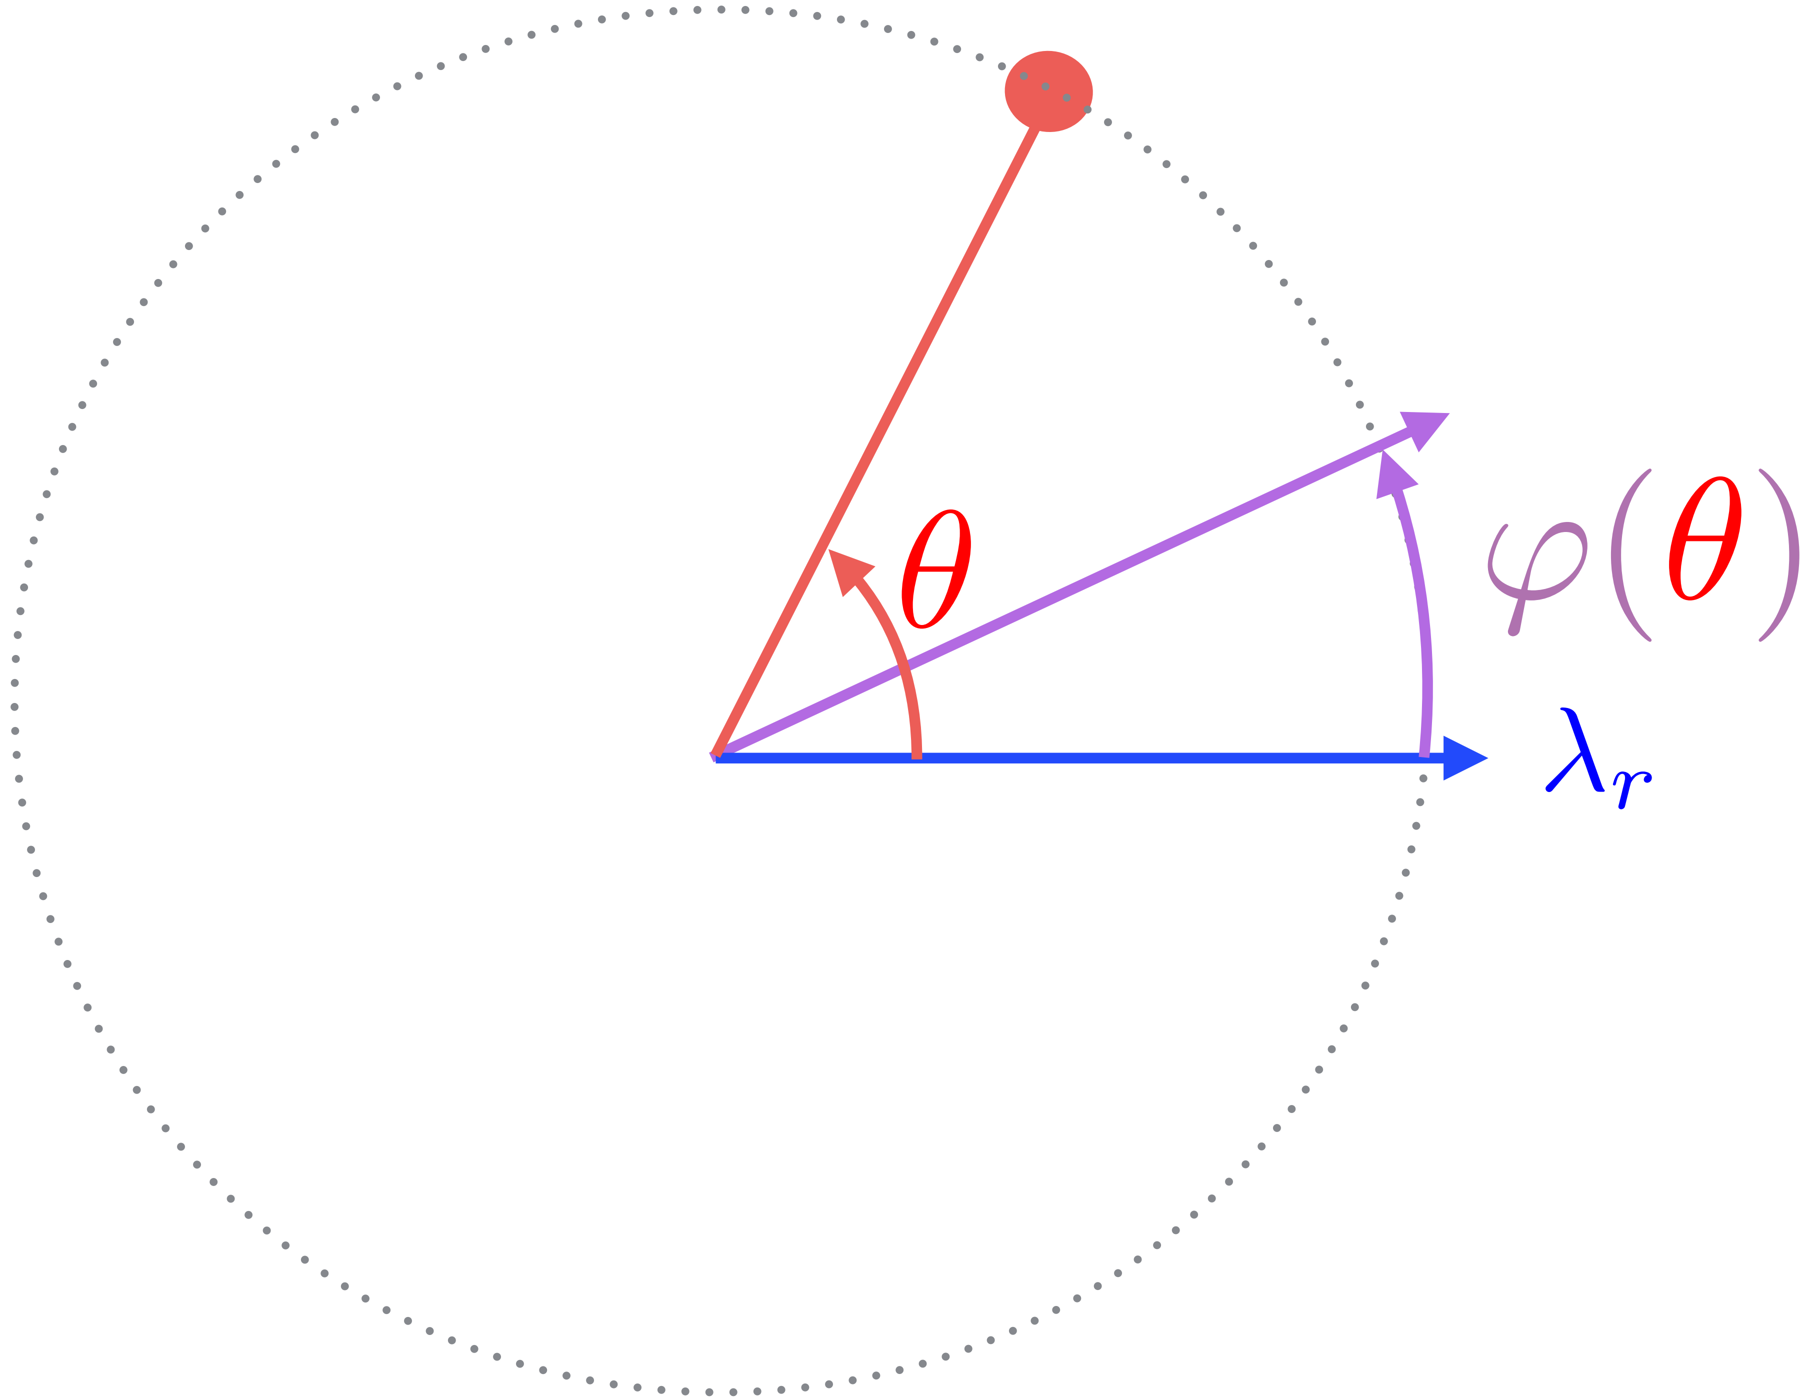

A flavor of the proof

How do we maximally *tilt* the smallest direction?

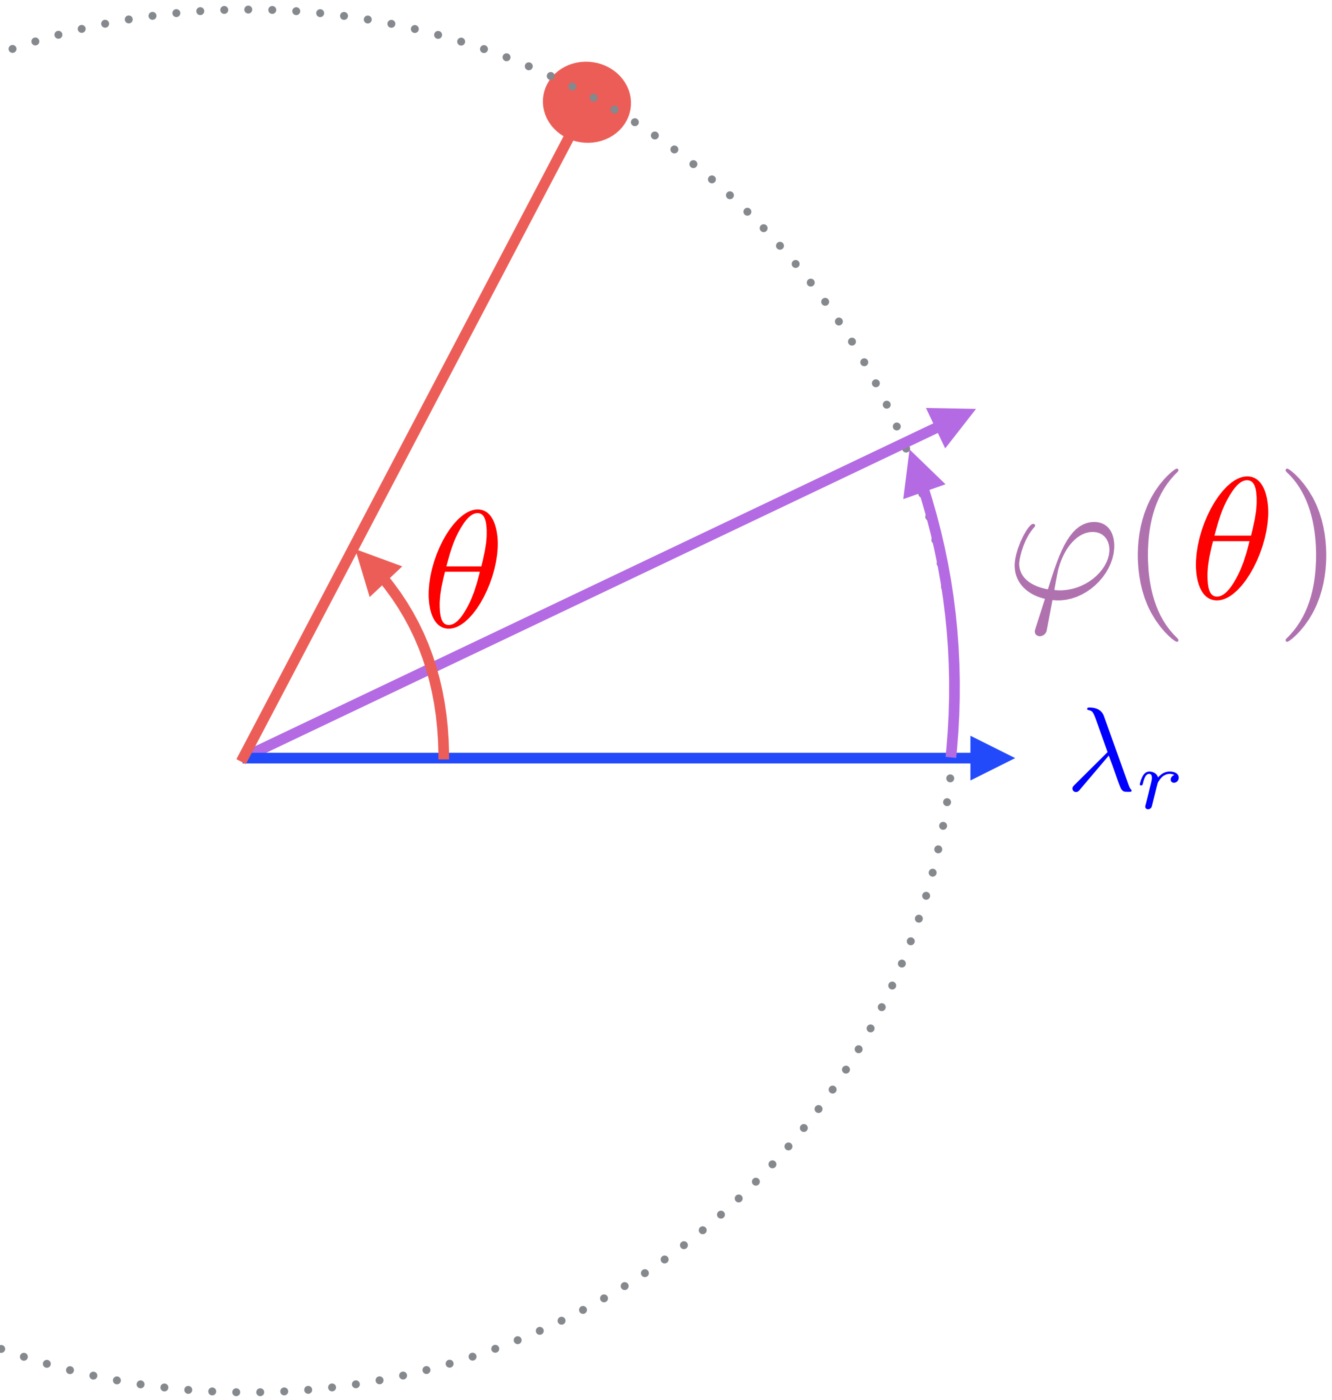

$$\boxed{\theta^\star = \arg\max_{\theta} \varphi(\theta)}$$

A flavor of the proof

How do we maximally *tilt* the smallest direction?

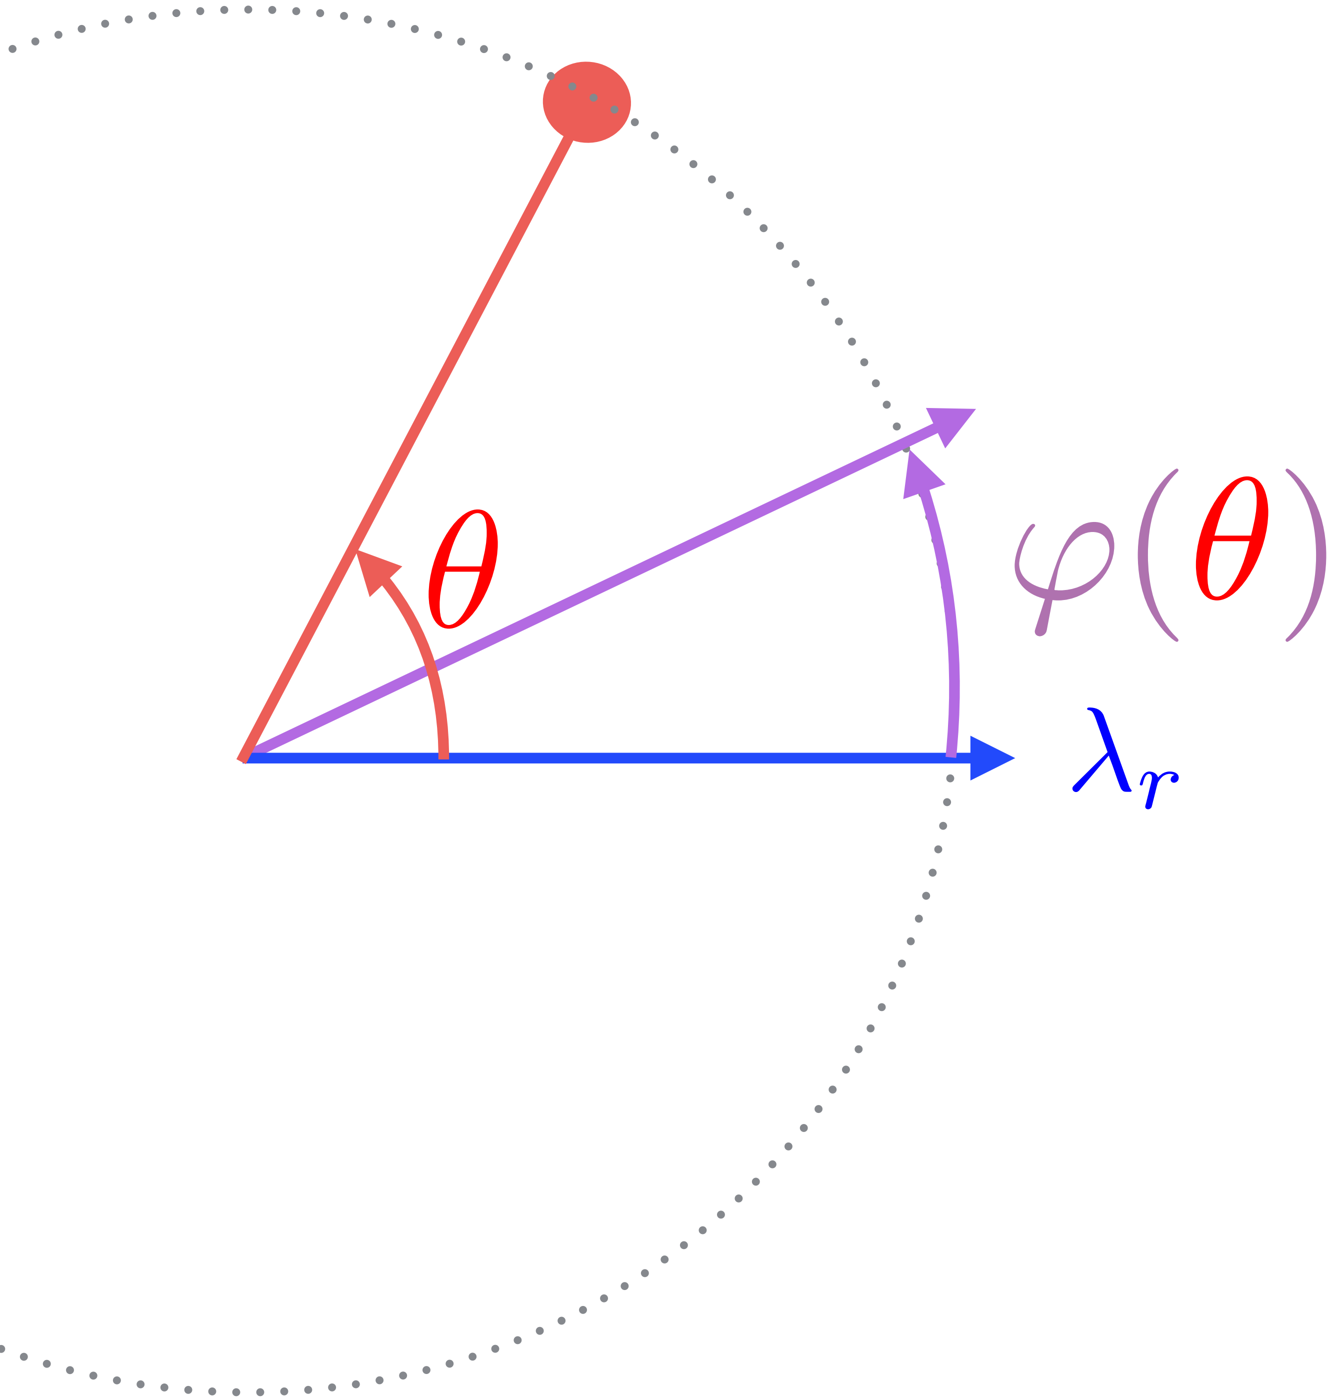

$$\boxed{\theta^\star = \arg\max_{\theta} \varphi(\theta)}$$

Usual tricks:

- Write in closed form.
- Take derivative.
- Set to zero.
- Solve.

(Easier said than done)

A flavor of the proof

How do we maximally *tilt* the smallest direction?

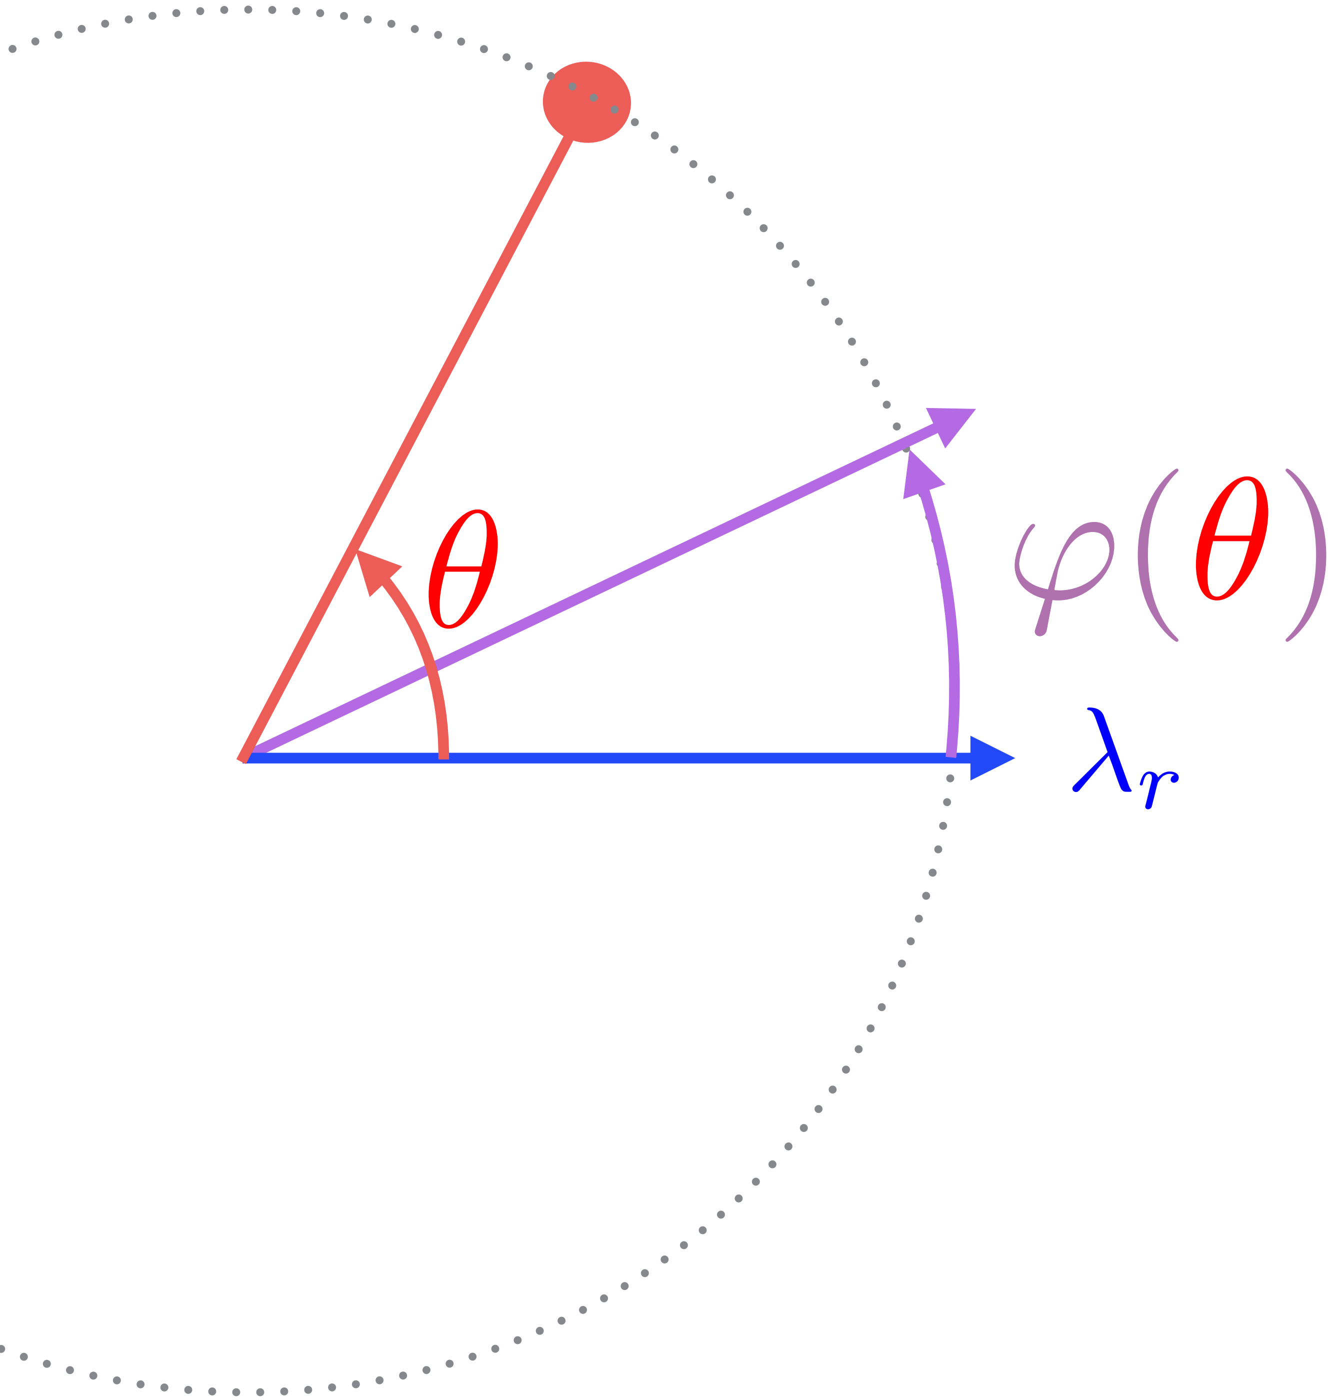

$$\boxed{\theta^\star = \arg\max_{\theta} \varphi(\theta)}$$

Usual tricks:

- Write in closed form.
- Take derivative.
- Set to zero.
- Solve.

(Easier said than done)

$$\boxed{\theta^\star = \frac{1}{2}\arccos\left(-\frac{1}{{\lambda_r}^2}\right)}$$

A flavor of the proof

How do we maximally *tilt* the smallest direction?

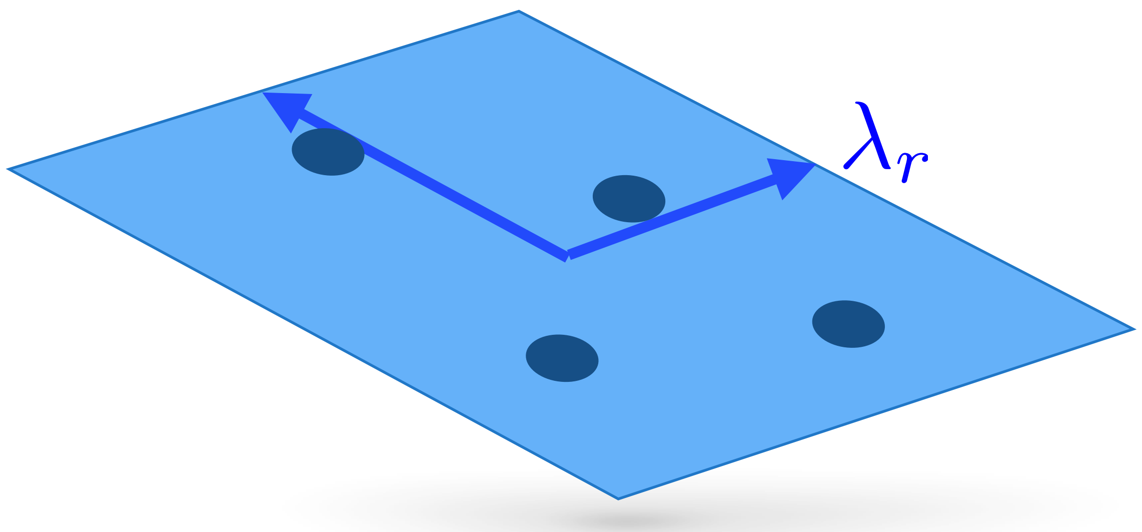

Putting everything together

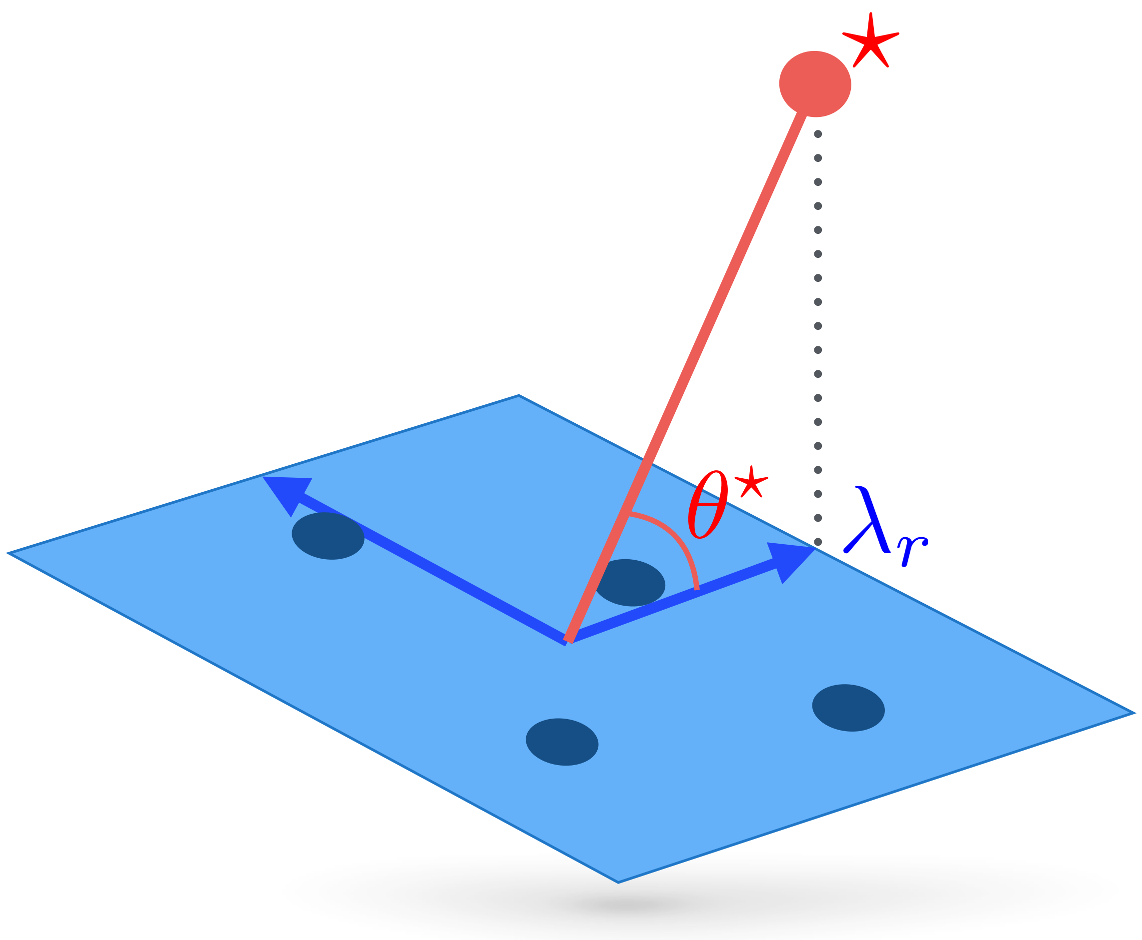

$$\theta^\star = \frac{1}{2}\arccos\left(-\frac{1}{{\lambda_r}^2}\right)$$

Putting everything together

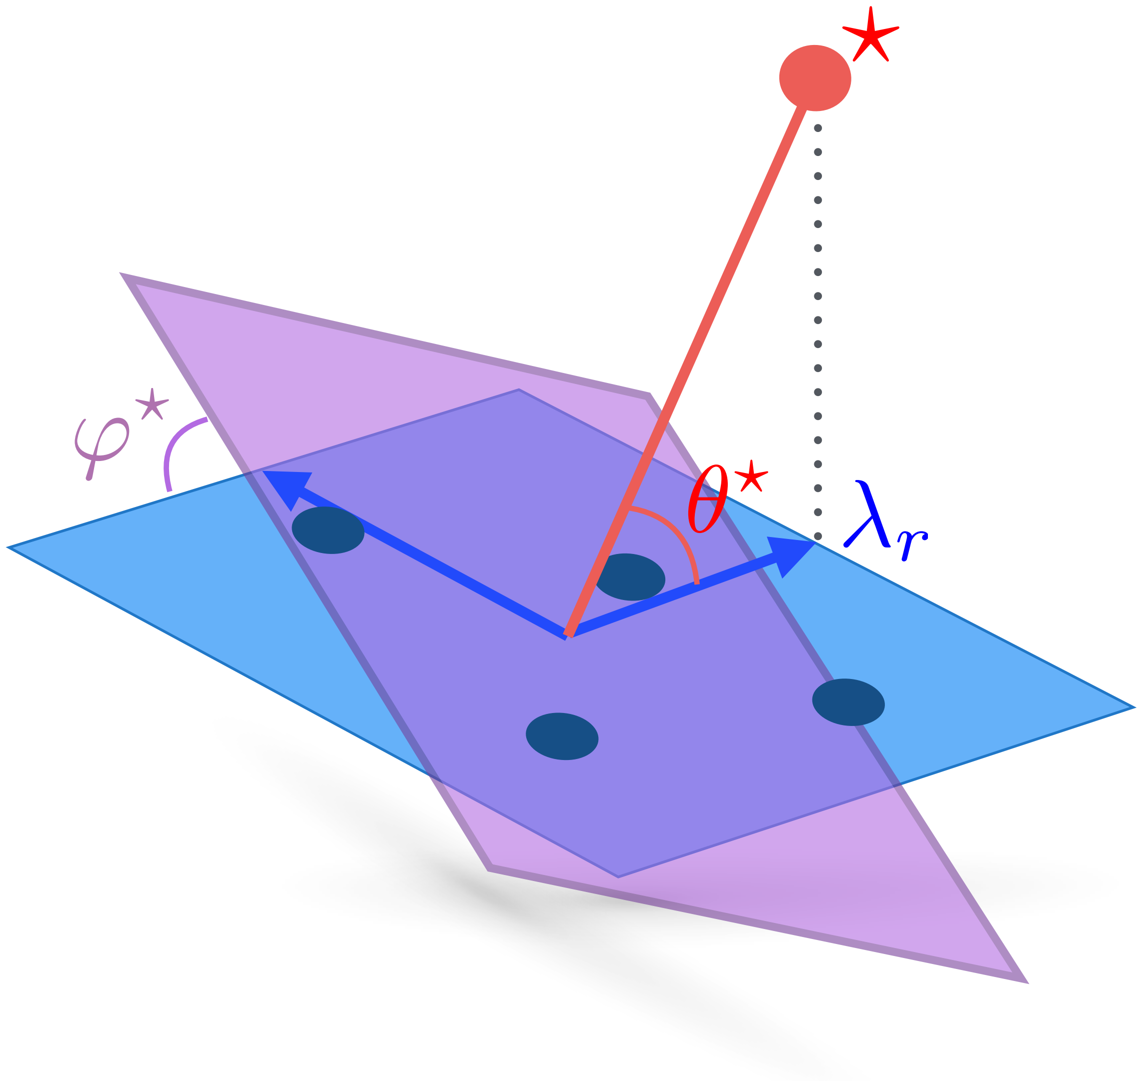

$$\theta^\star = \frac{1}{2}\arccos\left(-\frac{1}{{\lambda_r}^2}\right)$$

Putting everything together

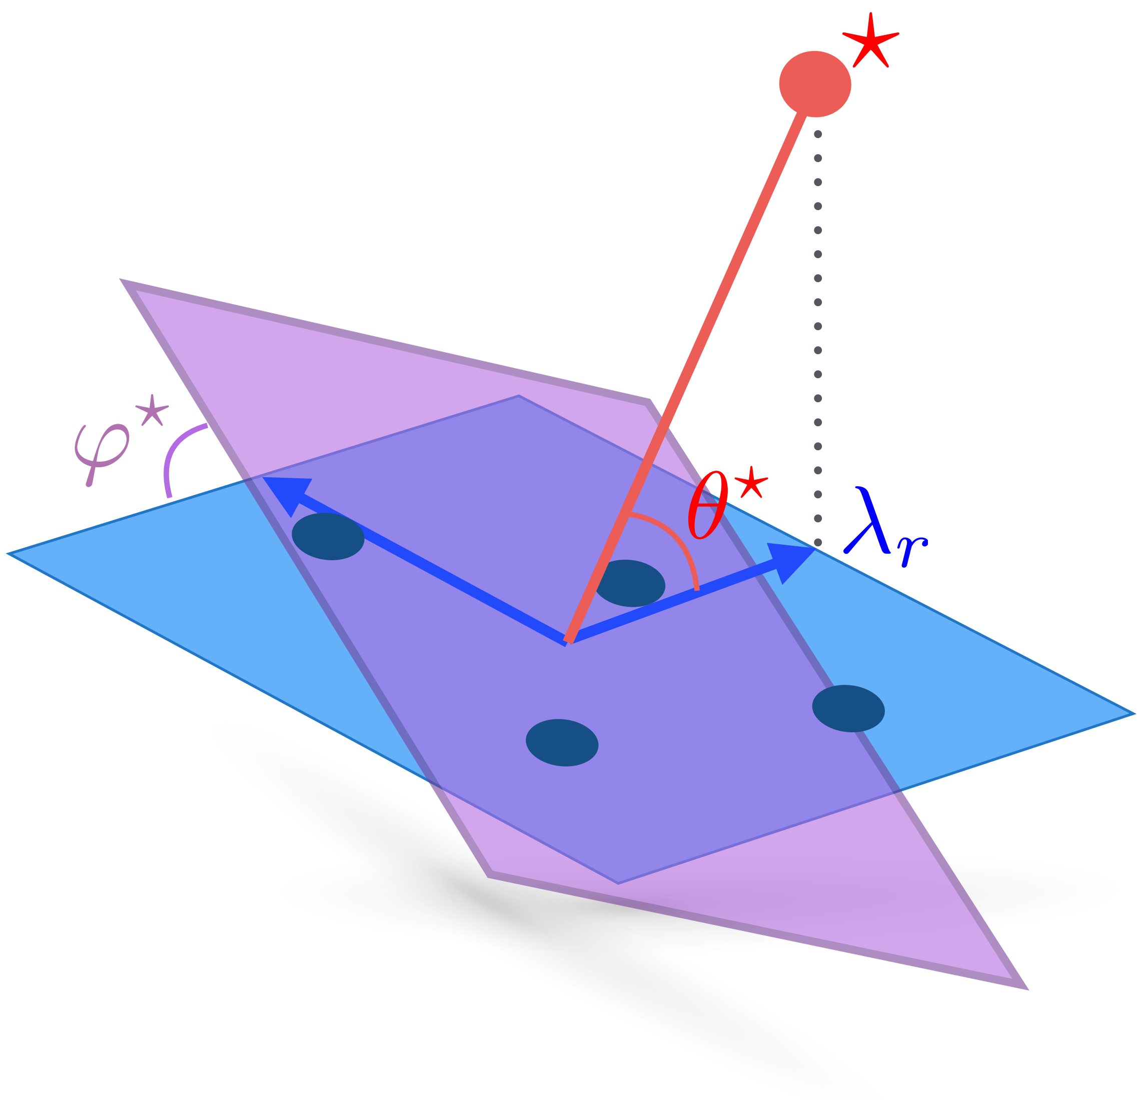

$$\theta^\star = \frac{1}{2}\arccos\left(-\frac{1}{{\lambda_r}^2}\right)$$

$$\varphi^\star = \arccos\left(\frac{\sin^2\theta^\star - \sigma_\star^2}{\sqrt{(\sin^2\theta^\star - \sigma_\star^2)^2 + (\sin\theta^\star\cos\theta^\star)^2}}\right)$$

$$\sigma_\star^2 = \frac{(\lambda_r^2+1) + \sqrt{(\lambda_r^2+1)^2 - 4\lambda_r^2\sin^2\theta^\star}}{2}.$$

Putting everything together

- Given a dataset 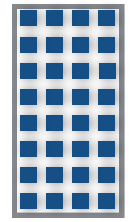

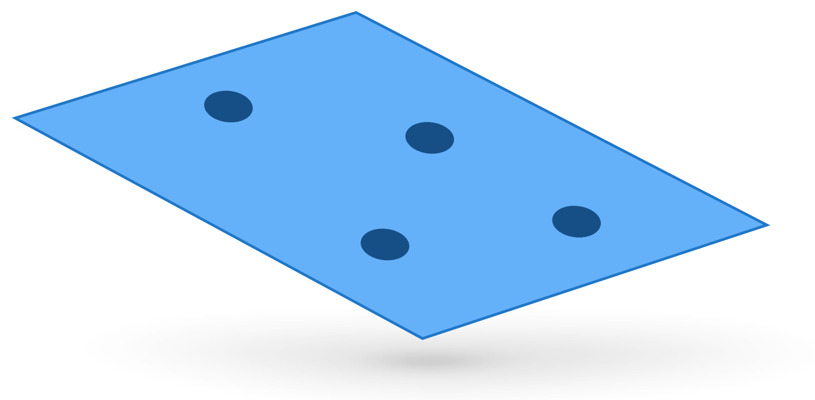

Take home message

- Given a dataset 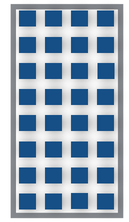, we know **exactly** what 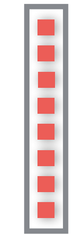 should be

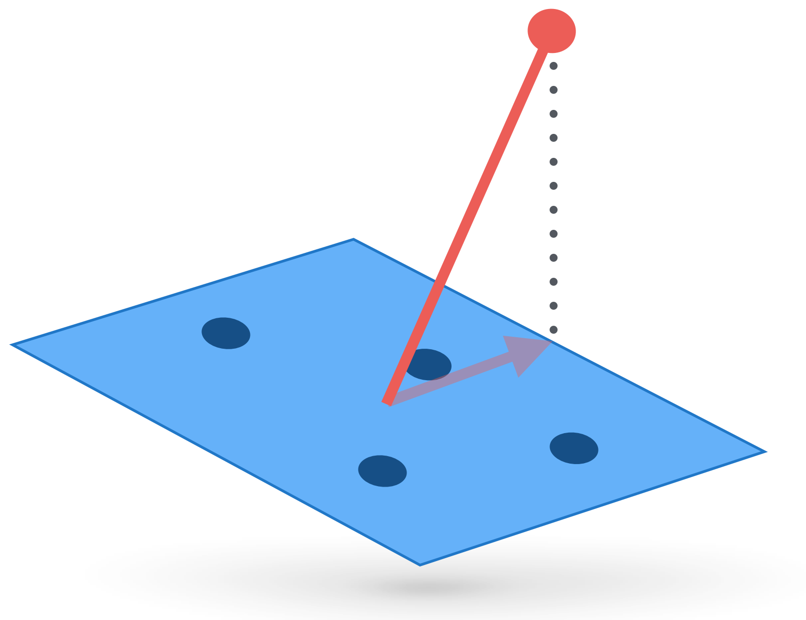

Take home message

- Given a dataset 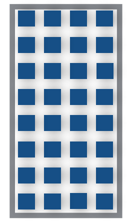, we know **exactly** what 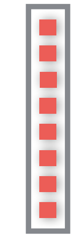 should be
  So that φ is maximal.

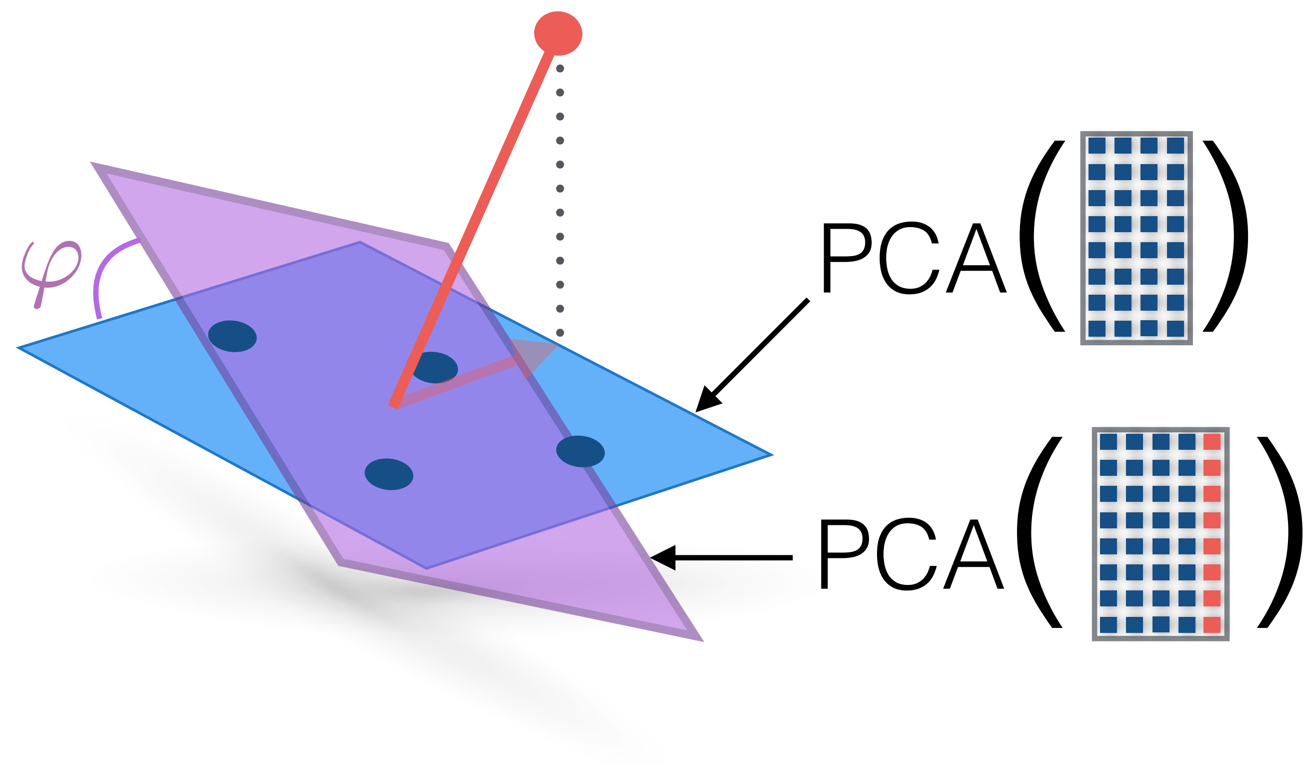

Take home message

- Given a dataset 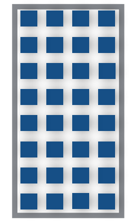, we know **exactly** what 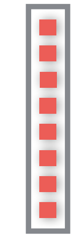 should be
  So that φ is maximal. (closed form)

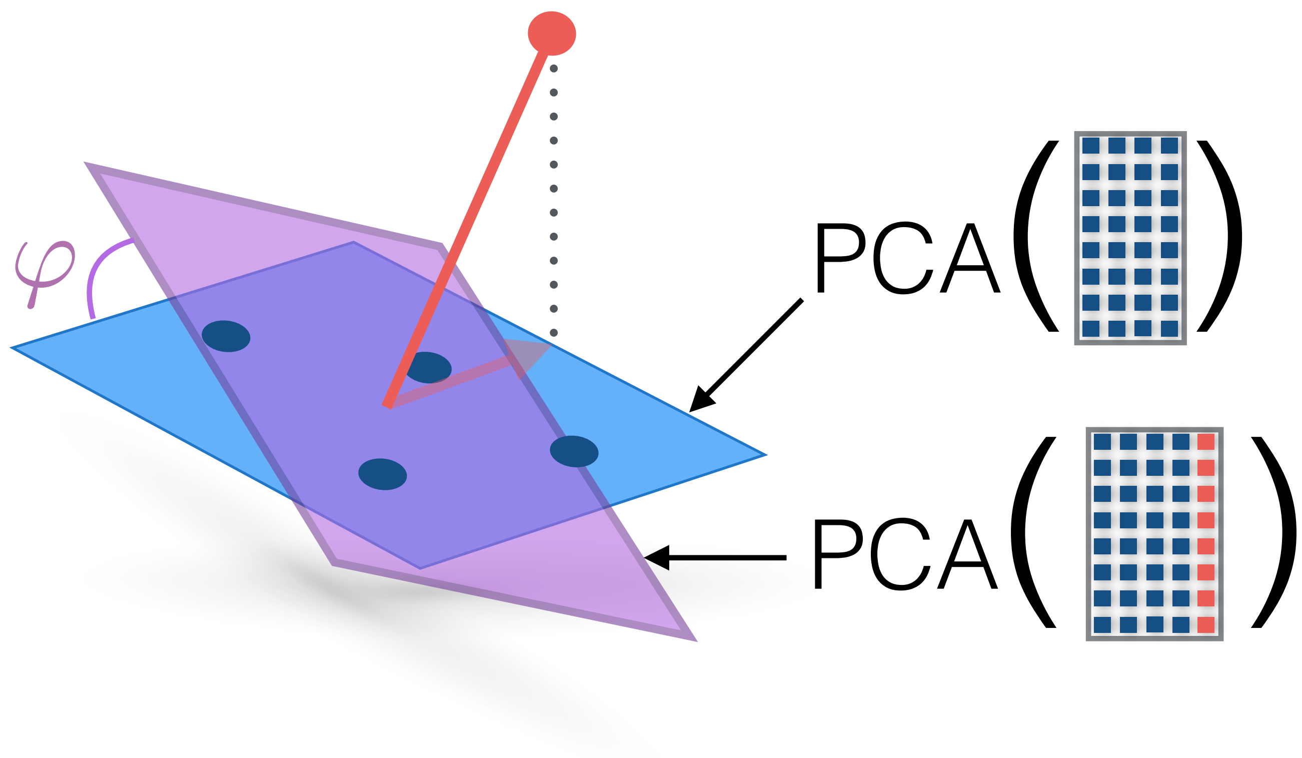

Take home message

- Given a dataset 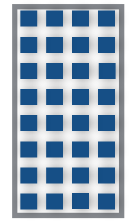, we know **exactly** what 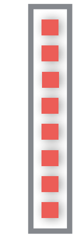 should be
  So that φ is maximal. (closed form)

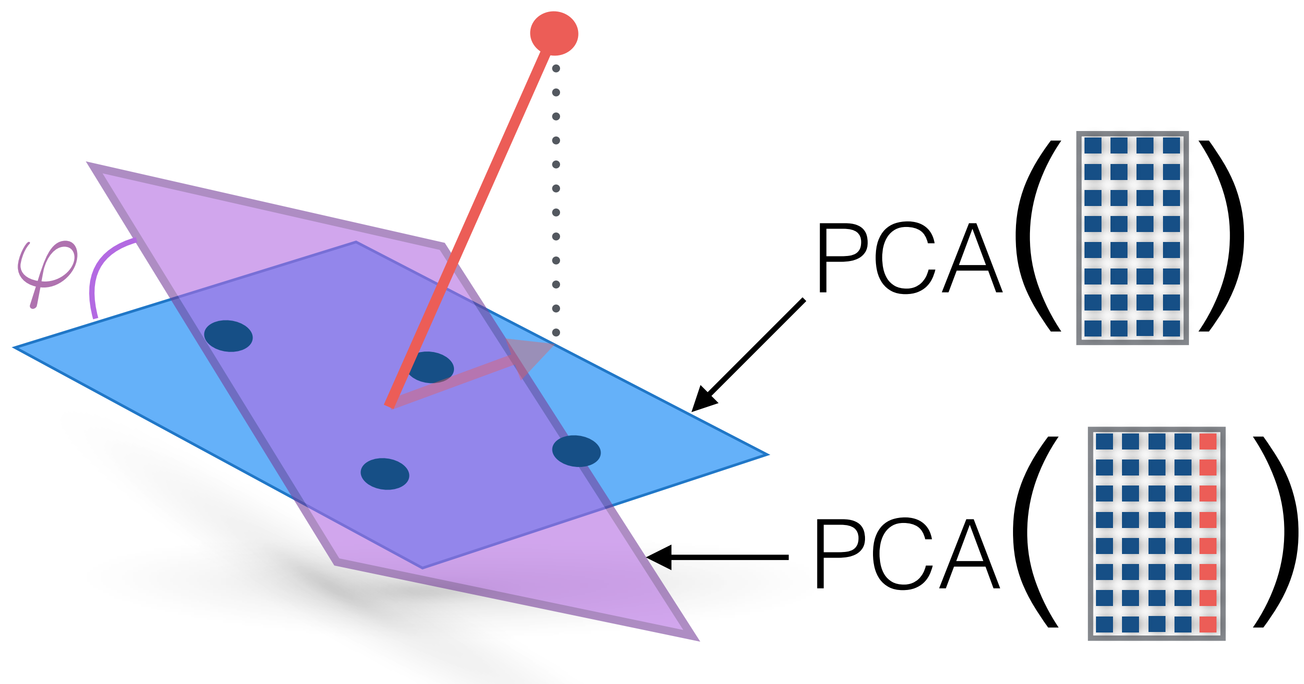

- Info-theory bound: how much one can *tilt* a subspace.
- Error bounds for Subspace Clustering.
- Applications in rank-one updates?
- Other applications?

Take home message

Dankeschön!