# Tight Bounds for Learning a Mixture of Two Gaussians 

Moritz Hardt Eric Price<br>Google Research UT Austin

## 2015-06-17

## Problem



- Height distribution of American 20 year olds.


## Problem



- Height distribution of American 20 year olds.
- Male/female heights are very close to Gaussian distribution.


## Problem



- Height distribution of American 20 year olds.
- Male/female heights are very close to Gaussian distribution.
- Can we learn the average male and female heights from unlabeled population data?


## Problem



- Height distribution of American 20 year olds.
- Male/female heights are very close to Gaussian distribution.
- Can we learn the average male and female heights from unlabeled population data?
- How many samples to learn $\mu_{1}, \mu_{2}$ to $\pm \epsilon \sigma$ ?


## Problem



- Height distribution of American 20 year olds.
- Male/female heights are very close to Gaussian distribution.
- Can we learn the average male and female heights from unlabeled population data?
- How many samples to learn $\mu_{1}, \mu_{2}$ to $\pm \epsilon \sigma$ ?
- d-dimensional setting: also learn weight, shoe size, ...


## Gaussian Mixtures: Origins

> III. Contributions to the Mathematical Theory of Erolution.
> By Kart, Prarson, University College, London.
> Communicated by Professor. Henrict, Ir.R.S.

Received October 18,-Read November 16, 1893.

> [Plates 1-5.]

## Contents.

Page.
I.-On the Dissection of Asymmetrical Frequency-Curves. Gcueral Thoory, §\$ 1-8. ..... 71-85
Example: Professor Wemosis mensmrements of the "Forchead" of Crabs.§§ $9-10$$85-90$
IL.-On the Dissoction of Symmetrical Trequency-Curves. General Theory, §§ 11-12 Application. Crabs "No. 4," §§ 13-1" ..... $90-100$
11L.-Investigation of an Asymmetrical Frequency, Curve repmesenting Mr. H. Thomsor's measurcments of the Carapaco of Prawns. §§ 16-18 ..... 100-106
Table J. First Six Powers of First Thirty Natural Nambers ..... 106
Table LLL. Ordinates of Normal Frequency-Curve ..... 107
Note alded February 10, 1894 ..... 107-11.0

## Gaussian Mixtures: Origins

Contributions to the Mathematical Theory of Evolution, Karl Pearson, 1894


- Pearson's naturalist buddy measured lots of crab body parts.


## Gaussian Mixtures: Origins

Contributions to the Mathematical Theory of Evolution, Karl Pearson, 1894


- Pearson's naturalist buddy measured lots of crab body parts.
- Most lengths seemed to follow the "normal" distribution (a recently coined name)


## Gaussian Mixtures: Origins

Contributions to the Mathematical Theory of Evolution, Karl Pearson, 1894


- Pearson's naturalist buddy measured lots of crab body parts.
- Most lengths seemed to follow the "normal" distribution (a recently coined name)
- But the "forehead" size wasn't symmetric.


## Gaussian Mixtures: Origins

Contributions to the Mathematical Theory of Evolution, Karl Pearson, 1894


- Pearson's naturalist buddy measured lots of crab body parts.
- Most lengths seemed to follow the "normal" distribution (a recently coined name)
- But the "forehead" size wasn't symmetric.
- Maybe there were actually two species of crabs?


## More previous work

- Pearson 1894: proposed method for 2 Gaussians


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:
- Royce '58, Gridgeman '70, Gupta-Huang '80


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:
- Royce '58, Gridgeman '70, Gupta-Huang '80
- Provable results assuming the components are well-separated:


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:
- Royce '58, Gridgeman '70, Gupta-Huang '80
- Provable results assuming the components are well-separated:
- Clustering: Dasgupta '99, DA '00


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:
- Royce '58, Gridgeman '70, Gupta-Huang '80
- Provable results assuming the components are well-separated:
- Clustering: Dasgupta '99, DA '00
- Spectral methods: VW '04, AK '05, KSV '05, AM '05, VW '05


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:
- Royce '58, Gridgeman '70, Gupta-Huang '80
- Provable results assuming the components are well-separated:
- Clustering: Dasgupta '99, DA '00
- Spectral methods: VW '04, AK '05, KSV '05, AM '05, VW '05
- Kalai-Moitra-Valiant 2010: first general polynomial bound.


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:
- Royce '58, Gridgeman '70, Gupta-Huang '80
- Provable results assuming the components are well-separated:
- Clustering: Dasgupta '99, DA '00
- Spectral methods: VW '04, AK '05, KSV '05, AM '05, VW '05
- Kalai-Moitra-Valiant 2010: first general polynomial bound.
- Extended to general $k$ mixtures: Moitra-Valiant '10, Belkin-Sinha '10


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:
- Royce '58, Gridgeman '70, Gupta-Huang '80
- Provable results assuming the components are well-separated:
- Clustering: Dasgupta '99, DA '00
- Spectral methods: VW '04, AK '05, KSV '05, AM '05, VW '05
- Kalai-Moitra-Valiant 2010: first general polynomial bound.
- Extended to general $k$ mixtures: Moitra-Valiant '10, Belkin-Sinha '10
- The KMV polynomial is very large.


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:
- Royce '58, Gridgeman '70, Gupta-Huang '80
- Provable results assuming the components are well-separated:
- Clustering: Dasgupta '99, DA '00
- Spectral methods: VW '04, AK '05, KSV '05, AM '05, VW '05
- Kalai-Moitra-Valiant 2010: first general polynomial bound.
- Extended to general $k$ mixtures: Moitra-Valiant '10, Belkin-Sinha '10
- The KMV polynomial is very large.
- Our result: tight upper and lower bounds for the sample complexity.


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:
- Royce '58, Gridgeman '70, Gupta-Huang '80
- Provable results assuming the components are well-separated:
- Clustering: Dasgupta '99, DA '00
- Spectral methods: VW '04, AK '05, KSV '05, AM '05, VW '05
- Kalai-Moitra-Valiant 2010: first general polynomial bound.
- Extended to general $k$ mixtures: Moitra-Valiant '10, Belkin-Sinha '10
- The KMV polynomial is very large.
- Our result: tight upper and lower bounds for the sample complexity.
- For $k=2$ mixtures, arbitrary $d$ dimensions.


## More previous work

- Pearson 1894: proposed method for 2 Gaussians
- "Method of moments"
- Other empirical papers over the years:
- Royce '58, Gridgeman '70, Gupta-Huang '80
- Provable results assuming the components are well-separated:
- Clustering: Dasgupta '99, DA '00
- Spectral methods: VW '04, AK '05, KSV '05, AM '05, VW '05
- Kalai-Moitra-Valiant 2010: first general polynomial bound.
- Extended to general $k$ mixtures: Moitra-Valiant '10, Belkin-Sinha '10
- The KMV polynomial is very large.
- Our result: tight upper and lower bounds for the sample complexity.
- For $k=2$ mixtures, arbitrary $d$ dimensions.
- Lower bound extends to larger $k$.


## Learning the components vs. learning the sum



- It's important that we want to learn the individual components:


## Learning the components vs. learning the sum



- It's important that we want to learn the individual components:
- Male/female average heights, std. deviations.


## Learning the components vs. learning the sum



- It's important that we want to learn the individual components:
- Male/female average heights, std. deviations.
- Getting $\epsilon$ approximation in TV norm to overall distribution takes $\widetilde{\Theta}\left(1 / \epsilon^{2}\right)$ samples from black box techniques.


## Learning the components vs. learning the sum



- It's important that we want to learn the individual components:
- Male/female average heights, std. deviations.
- Getting $\epsilon$ approximation in TV norm to overall distribution takes $\widetilde{\Theta}\left(1 / \epsilon^{2}\right)$ samples from black box techniques.
- Quite general: non-properly for any mixture of known unimodal distributions. [Chan, Diakonikolas, Servedio, Sun '13]


## Learning the components vs. learning the sum



- It's important that we want to learn the individual components:
- Male/female average heights, std. deviations.
- Getting $\epsilon$ approximation in TV norm to overall distribution takes $\widetilde{\Theta}\left(1 / \epsilon^{2}\right)$ samples from black box techniques.
- Quite general: non-properly for any mixture of known unimodal distributions. [Chan, Diakonikolas, Servedio, Sun '13]
- Proper learning: [Daskalakis-Kamath '14]


## Learning the components vs. learning the sum



- It's important that we want to learn the individual components:
- Male/female average heights, std. deviations.
- Getting $\epsilon$ approximation in TV norm to overall distribution takes $\widetilde{\Theta}\left(1 / \epsilon^{2}\right)$ samples from black box techniques.
- Quite general: non-properly for any mixture of known unimodal distributions. [Chan, Diakonikolas, Servedio, Sun '13]
- Proper learning: [Daskalakis-Kamath '14]
- But only in low dimensions.


## Learning the components vs. learning the sum



- It's important that we want to learn the individual components:
- Male/female average heights, std. deviations.
- Getting $\epsilon$ approximation in TV norm to overall distribution takes $\widetilde{\Theta}\left(1 / \epsilon^{2}\right)$ samples from black box techniques.
- Quite general: non-properly for any mixture of known unimodal distributions. [Chan, Diakonikolas, Servedio, Sun '13]
- Proper learning: [Daskalakis-Kamath '14]
- But only in low dimensions.
- Generic high-d TV estimation algs use 1d parameter estimation.


## Our result

- A variant of Pearson's 1894 method is optimal!


## Our result

- A variant of Pearson's 1894 method is optimal!
- Suppose we want means and variances to $\epsilon$ accuracy:


## Our result

- A variant of Pearson's 1894 method is optimal!
- Suppose we want means and variances to $\epsilon$ accuracy:
- $\mu_{i}$ to $\pm \epsilon \sigma$


## Our result

- A variant of Pearson's 1894 method is optimal!
- Suppose we want means and variances to $\epsilon$ accuracy:
- $\mu_{i}$ to $\pm \epsilon \sigma$
- $\sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$


## Our result

- A variant of Pearson's 1894 method is optimal!
- Suppose we want means and variances to $\epsilon$ accuracy:
- $\mu_{i}$ to $\pm \epsilon \sigma$
- $\sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$
- In one dimension: $\Theta\left(1 / \epsilon^{12}\right)$ samples necessary and sufficient.


## Our result

- A variant of Pearson's 1894 method is optimal!
- Suppose we want means and variances to $\epsilon$ accuracy:
- $\mu_{i}$ to $\pm \epsilon \sigma$
- $\sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$
- In one dimension: $\Theta\left(1 / \epsilon^{12}\right)$ samples necessary and sufficient.
- Previously: $1 / \epsilon \approx 300$, no lower bound.


## Our result

- A variant of Pearson's 1894 method is optimal!
- Suppose we want means and variances to $\epsilon$ accuracy:
- $\mu_{i}$ to $\pm \epsilon \sigma$
- $\sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$
- In one dimension: $\Theta\left(1 / \epsilon^{12}\right)$ samples necessary and sufficient.
- Previously: $1 / \epsilon^{\approx 300}$, no lower bound.
- Moreover: algorithm is almost the same as Pearson (1894).


## Our result

- A variant of Pearson's 1894 method is optimal!
- Suppose we want means and variances to $\epsilon$ accuracy:
- $\mu_{i}$ to $\pm \epsilon \sigma$
- $\sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$
- In one dimension: $\Theta\left(1 / \epsilon^{12}\right)$ samples necessary and sufficient.
- Previously: $1 / \epsilon^{\approx 300}$, no lower bound.
- Moreover: algorithm is almost the same as Pearson (1894).

- More precisely: if two gaussians are $\alpha$ standard deviations apart, getting $\epsilon \alpha$ precision takes $\Theta\left(\frac{1}{\alpha^{12} \epsilon^{2}}\right)$ samples.


## Our result: higher dimensions

- In dimensions, $\Theta\left(1 / \epsilon^{12} \log d\right)$ samples for parameter distance.


## Our result: higher dimensions

- In $d$ dimensions, $\Theta\left(1 / \epsilon^{12} \log d\right)$ samples for parameter distance.
- " $\sigma^{2}$ " is max variance in any coordinate.


## Our result: higher dimensions

- In $d$ dimensions, $\Theta\left(1 / \epsilon^{12} \log d\right)$ samples for parameter distance.
- " $\sigma^{2}$ " is max variance in any coordinate.
- Get each entry of covariance matrix to $\pm \epsilon^{2} \sigma^{2}$.


## Our result: higher dimensions

- In $d$ dimensions, $\Theta\left(1 / \epsilon^{12} \log d\right)$ samples for parameter distance.
- " $\sigma^{2}$ " is max variance in any coordinate.
- Get each entry of covariance matrix to $\pm \epsilon^{2} \sigma^{2}$.
- Useful when covariance matrix is sparse.


## Our result: higher dimensions

- In $d$ dimensions, $\Theta\left(1 / \epsilon^{12} \log d\right)$ samples for parameter distance.
- " $\sigma^{2}$ " is max variance in any coordinate.
- Get each entry of covariance matrix to $\pm \epsilon^{2} \sigma^{2}$.
- Useful when covariance matrix is sparse.
- Also gives an improved bound in TV error of each component:


## Our result: higher dimensions



- In $d$ dimensions, $\Theta\left(1 / \epsilon^{12} \log d\right)$ samples for parameter distance.
- " $\sigma^{2}$ " is max variance in any coordinate.
- Get each entry of covariance matrix to $\pm \epsilon^{2} \sigma^{2}$.
- Useful when covariance matrix is sparse.
- Also gives an improved bound in TV error of each component:
- If components overlap, then parameter distance $\approx$ TV.


## Our result: higher dimensions



- In $d$ dimensions, $\Theta\left(1 / \epsilon^{12} \log d\right)$ samples for parameter distance.
- " $\sigma^{2}$ " is max variance in any coordinate.
- Get each entry of covariance matrix to $\pm \epsilon^{2} \sigma^{2}$.
- Useful when covariance matrix is sparse.
- Also gives an improved bound in TV error of each component:
- If components overlap, then parameter distance $\approx$ TV.
- If components don't overlap, then clustering is trivial.


## Our result: higher dimensions



- In $d$ dimensions, $\Theta\left(1 / \epsilon^{12} \log d\right)$ samples for parameter distance.
- " $\sigma^{2}$ " is max variance in any coordinate.
- Get each entry of covariance matrix to $\pm \epsilon^{2} \sigma^{2}$.
- Useful when covariance matrix is sparse.
- Also gives an improved bound in TV error of each component:
- If components overlap, then parameter distance $\approx$ TV.
- If components don't overlap, then clustering is trivial.
- Straightforwardly gives $\widetilde{O}\left(d^{30} / \epsilon^{36}\right)$ samples.


## Our result: higher dimensions



- In $d$ dimensions, $\Theta\left(1 / \epsilon^{12} \log d\right)$ samples for parameter distance.
- " $\sigma^{2}$ " is max variance in any coordinate.
- Get each entry of covariance matrix to $\pm \epsilon^{2} \sigma^{2}$.
- Useful when covariance matrix is sparse.
- Also gives an improved bound in TV error of each component:
- If components overlap, then parameter distance $\approx$ TV.
- If components don't overlap, then clustering is trivial.
- Straightforwardly gives $\widetilde{O}\left(d^{30} / \epsilon^{36}\right)$ samples.
- Best known, but not the $\widetilde{O}\left(d / \epsilon^{c}\right)$ we want.


## Our result: higher dimensions



- In $d$ dimensions, $\Theta\left(1 / \epsilon^{12} \log d\right)$ samples for parameter distance.
- " $\sigma^{2}$ " is max variance in any coordinate.
- Get each entry of covariance matrix to $\pm \epsilon^{2} \sigma^{2}$.
- Useful when covariance matrix is sparse.
- Also gives an improved bound in TV error of each component:
- If components overlap, then parameter distance $\approx$ TV.
- If components don't overlap, then clustering is trivial.
- Straightforwardly gives $\widetilde{O}\left(d^{30} / \epsilon^{36}\right)$ samples.
- Best known, but not the $\widetilde{O}\left(d / \epsilon^{c}\right)$ we want.
- Caveat: assume $p_{1}, p_{2}$ are bounded away from zero throughout.


## Outline

## (1) Algorithm in One Dimension

## Outline

# (1) Algorithm in One Dimension 

(2) Lower Bound

## Outline

# (1) Algorithm in One Dimension 

(2) Lower Bound
(3) Algorithm in $d$ Dimensions

## Outline

# (1) Algorithm in One Dimension 

(2) Lower Bound

## (3) Algorithm in $d$ Dimensions

## Method of Moments



- We want to learn five parameters: $\mu_{1}, \mu_{2}, \sigma_{1}, \sigma_{2}, p_{1}, p_{2}$ with $p_{1}+p_{2}=1$.


## Method of Moments



- We want to learn five parameters: $\mu_{1}, \mu_{2}, \sigma_{1}, \sigma_{2}, p_{1}, p_{2}$ with $p_{1}+p_{2}=1$.
- Moments give polynomial equations in parameters:

$$
\begin{aligned}
& M_{1}:=\mathbb{E}\left[x^{1}\right]=p_{1} \mu_{1}+p_{2} \mu_{2} \\
& M_{2}:=\mathbb{E}\left[x^{2}\right] \\
&=p_{1} \mu_{1}^{2}+p_{2} \mu_{2}^{2}+p_{1} \sigma_{1}^{2}+p_{2} \sigma_{2}^{2} \\
& M_{3}, M_{4}, M_{5}, M_{6}=[\ldots]
\end{aligned}
$$

## Method of Moments



- We want to learn five parameters: $\mu_{1}, \mu_{2}, \sigma_{1}, \sigma_{2}, p_{1}, p_{2}$ with $p_{1}+p_{2}=1$.
- Moments give polynomial equations in parameters:

$$
\begin{aligned}
M_{1} & :=\mathbb{E}\left[x^{1}\right] \\
M_{2} & :=p_{1} \mu_{1}+p_{2} \mu_{2} \\
M_{3}, M_{4}, M_{5}, M_{6} & =[\ldots]
\end{aligned}
$$

- Use our samples to estimate the moments.


## Method of Moments



- We want to learn five parameters: $\mu_{1}, \mu_{2}, \sigma_{1}, \sigma_{2}, p_{1}, p_{2}$ with $p_{1}+p_{2}=1$.
- Moments give polynomial equations in parameters:

$$
\begin{aligned}
& M_{1}:=\mathbb{E}\left[x^{1}\right]=p_{1} \mu_{1}+p_{2} \mu_{2} \\
& M_{2}:=\mathbb{E}\left[x^{2}\right] \\
&=p_{1} \mu_{1}^{2}+p_{2} \mu_{2}^{2}+p_{1} \sigma_{1}^{2}+p_{2} \sigma_{2}^{2} \\
& M_{3}, M_{4}, M_{5}, M_{6}=[\ldots]
\end{aligned}
$$

- Use our samples to estimate the moments.
- Solve the system of equations to find the parameters.


## Method of Moments

Solving the system

- Start with five parameters.


## Method of Moments

Solving the system

- Start with five parameters.
- First, can assume mean zero:
- Convert to "central moments"


## Method of Moments

Solving the system

- Start with five parameters.
- First, can assume mean zero:
- Convert to "central moments"
- $M_{2}^{\prime}=M_{2}-M_{1}^{2}$ is independent of translation.


## Method of Moments

## Solving the system

- Start with five parameters.
- First, can assume mean zero:
- Convert to "central moments"
- $M_{2}^{\prime}=M_{2}-M_{1}^{2}$ is independent of translation.
- Analogously, can assume $\min \left(\sigma_{1}, \sigma_{2}\right)=0$ by converting to "excess moments"


## Method of Moments

## Solving the system

- Start with five parameters.
- First, can assume mean zero:
- Convert to "central moments"
- $M_{2}^{\prime}=M_{2}-M_{1}^{2}$ is independent of translation.
- Analogously, can assume $\min \left(\sigma_{1}, \sigma_{2}\right)=0$ by converting to "excess moments"
- $X_{4}=M_{4}-3 M_{2}^{2}$ is independent of adding $N\left(0, \sigma^{2}\right)$.


## Method of Moments

Solving the system

- Start with five parameters.
- First, can assume mean zero:

| Parameters | $\lambda>0$ rate, or inverse scale |
| :--- | :--- |
| Support | $x \in[0, \infty)$ |
| pdf | $\lambda e^{-\lambda x}$ |
| CDF | $1-e^{-\lambda x}$ |
| Mean | $\lambda^{-1}$ |
| Median | $\lambda^{-1} \ln (2)$ |
| Mode | 0 |
| Variance | $\lambda^{-2}$ |
| Skewness | 2 |
| Ex. kurtosis | $1-\ln (\lambda)$ |
| Entropy | $\left(1-\frac{t}{\lambda}\right)^{-1}$ for $t<\lambda$ |
| MGF | $\left(1-\frac{i t}{\lambda}\right)^{-1}$ |
| CF | $1-2$ |
| Cichar infarmation |  |

- Convert to "central moments"
- $M_{2}^{\prime}=M_{2}-M_{1}^{2}$ is independent of translation.
- Analogously, can assume $\min \left(\sigma_{1}, \sigma_{2}\right)=0$ by converting to "excess moments"
- $X_{4}=M_{4}-3 M_{2}^{2}$ is independent of adding $N\left(0, \sigma^{2}\right)$.
- "Excess kurtosis" coined by Pearson, appearing in every Wikipedia probability distribution infobox.


## Method of Moments

Solving the system

- Start with five parameters.
- First, can assume mean zero:

| Parameters | $\lambda>0$ rate, or inverse scale |
| :--- | :--- |
| Support | $x \in[0, \infty)$ |
| pdf | $\lambda e^{-\lambda x}$ |
| CDF | $1-e^{-\lambda x}$ |
| Mean | $\lambda^{-1}$ |
| Median | $\lambda^{-1} \ln (2)$ |
| Mode | 0 |
| Variance | $\lambda^{-2}$ |
| Skewness | 2 |
| Ex. kurtosis | 6 |
| Entropy | $1-\ln (\lambda)$ |
| MGF | $\left(1-\frac{t}{\lambda}\right)^{-1}$ for $t<\lambda$ |
| CF | $\left(1-\frac{i t}{\lambda}\right)^{-1}$ |
| Fichar infarmation | $1-2$ |

- Convert to "central moments"
- $M_{2}^{\prime}=M_{2}-M_{1}^{2}$ is independent of translation.
- Analogously, can assume $\min \left(\sigma_{1}, \sigma_{2}\right)=0$ by converting to "excess moments"
- $X_{4}=M_{4}-3 M_{2}^{2}$ is independent of adding $N\left(0, \sigma^{2}\right)$.
- "Excess kurtosis" coined by Pearson, appearing in every Wikipedia probability distribution infobox.
- Leaves three free parameters.


## Method of Moments: system of equations

- Convenient to reparameterize by

$$
\alpha=-\mu_{1} \mu_{2}, \beta=\mu_{1}+\mu_{2}, \gamma=\frac{\sigma_{2}^{2}-\sigma_{1}^{2}}{\mu_{2}-\mu_{1}}
$$

## Method of Moments: system of equations

- Convenient to reparameterize by

$$
\alpha=-\mu_{1} \mu_{2}, \beta=\mu_{1}+\mu_{2}, \gamma=\frac{\sigma_{2}^{2}-\sigma_{1}^{2}}{\mu_{2}-\mu_{1}}
$$

- Gives that

$$
\begin{aligned}
& X_{3}=\alpha(\beta+3 \gamma) \\
& X_{4}=\alpha\left(-2 \alpha+\beta^{2}+6 \beta \gamma+3 \gamma^{2}\right) \\
& X_{5}=\alpha\left(\beta^{3}-8 \alpha \beta+10 \beta^{2} \gamma+15 \gamma^{2} \beta-20 \alpha \gamma\right) \\
& X_{6}=\alpha\left(16 \alpha^{2}-12 \alpha \beta^{2}-60 \alpha \beta \gamma+\beta^{4}+15 \beta^{3} \gamma+45 \beta^{2} \gamma^{2}+15 \beta \gamma^{3}\right)
\end{aligned}
$$

## Method of Moments: system of equations

- Convenient to reparameterize by

$$
\alpha=-\mu_{1} \mu_{2}, \beta=\mu_{1}+\mu_{2}, \gamma=\frac{\sigma_{2}^{2}-\sigma_{1}^{2}}{\mu_{2}-\mu_{1}}
$$

- Gives that

$$
\begin{aligned}
& X_{3}=\alpha(\beta+3 \gamma) \\
& X_{4}=\alpha\left(-2 \alpha+\beta^{2}+6 \beta \gamma+3 \gamma^{2}\right) \\
& X_{5}=\alpha\left(\beta^{3}-8 \alpha \beta+10 \beta^{2} \gamma+15 \gamma^{2} \beta-20 \alpha \gamma\right) \\
& X_{6}=\alpha\left(16 \alpha^{2}-12 \alpha \beta^{2}-60 \alpha \beta \gamma+\beta^{4}+15 \beta^{3} \gamma+45 \beta^{2} \gamma^{2}+15 \beta \gamma^{3}\right)
\end{aligned}
$$

All my attempts to obtain a simpler set have failed... It is possible, however, that some other ... equations of a less complex kind may ultimately be found.

## Pearson's Polynomial

- Chug chug chug...


## Pearson's Polynomial

- Chug chug chug...
- Get a 9th degree polynomial in the excess moments $X_{3}, X_{4}, X_{5}$ :

$$
\begin{aligned}
p(\alpha)= & 8 \alpha^{9}+28 X_{4} \alpha^{7}-12 X_{3}^{2} \alpha^{6}+\left(24 X_{3} X_{5}+30 X_{4}^{2}\right) \alpha^{5} \\
& \quad+\left(6 X_{5}^{2}-148 X_{3}^{2} X_{4}\right) \alpha^{4}+\left(96 X_{3}^{4}-36 X_{3} X_{4} X_{5}+9 X_{4}^{3}\right) \alpha^{3} \\
& \quad+\left(24 X_{3}^{3} X_{5}+21 X_{3}^{2} X_{4}^{2}\right) \alpha^{2}-32 X_{3}^{4} X_{4} \alpha+8 X_{3}^{6} \\
= & 0
\end{aligned}
$$

## Pearson's Polynomial

- Chug chug chug...
- Get a 9th degree polynomial in the excess moments $X_{3}, X_{4}, X_{5}$ :

$$
\begin{aligned}
p(\alpha)= & 8 \alpha^{9}+28 X_{4} \alpha^{7}-12 X_{3}^{2} \alpha^{6}+\left(24 X_{3} X_{5}+30 X_{4}^{2}\right) \alpha^{5} \\
& +\left(6 X_{5}^{2}-148 X_{3}^{2} X_{4}\right) \alpha^{4}+\left(96 X_{3}^{4}-36 X_{3} X_{4} X_{5}+9 X_{4}^{3}\right) \alpha^{3} \\
& \quad+\left(24 X_{3}^{3} X_{5}+21 X_{3}^{2} X_{4}^{2}\right) \alpha^{2}-32 X_{3}^{4} X_{4} \alpha+8 X_{3}^{6} \\
= & 0
\end{aligned}
$$

- Easy to go from solutions $\alpha=-\mu_{1} \mu_{2}$ to mixtures $\mu_{i}, \sigma_{i}, p_{i}$.


## Pearson's Polynomial



- Get a 9th degree polynomial in the excess moments $X_{3}, X_{4}, X_{5}$.


## Pearson's Polynomial



- Get a 9th degree polynomial in the excess moments $X_{3}, X_{4}, X_{5}$.
- Positive roots correspond to mixtures that match on five moments.


## Pearson's Polynomial



- Get a 9th degree polynomial in the excess moments $X_{3}, X_{4}, X_{5}$.
- Positive roots correspond to mixtures that match on five moments.
- Pearson's proposal: choose root with closer 6th moment.


## Pearson's Polynomial



- Get a 9th degree polynomial in the excess moments $X_{3}, X_{4}, X_{5}$.
- Positive roots correspond to mixtures that match on five moments.
- Pearson's proposal: choose root with closer 6th moment.
- Works because six moments uniquely identify mixture [KMV]


## Pearson's Polynomial



- Get a 9th degree polynomial in the excess moments $X_{3}, X_{4}, X_{5}$.
- Positive roots correspond to mixtures that match on five moments.
- Pearson's proposal: choose root with closer 6th moment.
- Works because six moments uniquely identify mixture [KMV]
- How robust to moment estimation error?


## Pearson's Polynomial



- Get a 9th degree polynomial in the excess moments $X_{3}, X_{4}, X_{5}$.
- Positive roots correspond to mixtures that match on five moments.
- Pearson's proposal: choose root with closer 6th moment.
- Works because six moments uniquely identify mixture [KMV]
- How robust to moment estimation error?
- Usually works well


## Pearson's Polynomial



- Get a 9th degree polynomial in the excess moments $X_{3}, X_{4}, X_{5}$.
- Positive roots correspond to mixtures that match on five moments.
- Pearson's proposal: choose root with closer 6th moment.
- Works because six moments uniquely identify mixture [KMV]
- How robust to moment estimation error?
- Usually works well
- Not when there's a double root.


## Making it robust in all cases

- Can create another ninth degree polynomial $p_{6}$ from $X_{3}, X_{4}, X_{5}, X_{6}$.


## Making it robust in all cases

- Can create another ninth degree polynomial $p_{6}$ from $X_{3}, X_{4}, X_{5}, X_{6}$.
- Then $\alpha$ is the unique positive root of

$$
r(\alpha):=p_{5}(\alpha)^{2}+p_{6}(\alpha)^{2}=0
$$

## Making it robust in all cases

- Can create another ninth degree polynomial $p_{6}$ from $X_{3}, X_{4}, X_{5}, X_{6}$.
- Then $\alpha$ is the unique positive root of

$$
r(\alpha):=p_{5}(\alpha)^{2}+p_{6}(\alpha)^{2}=0
$$

- How robust is the solution to perturbations of $X_{3}, \ldots, X_{6}$ ?


## Making it robust in all cases

- Can create another ninth degree polynomial $p_{6}$ from $X_{3}, X_{4}, X_{5}, X_{6}$.
- Then $\alpha$ is the unique positive root of

$$
r(\alpha):=p_{5}(\alpha)^{2}+p_{6}(\alpha)^{2}=0
$$

- How robust is the solution to perturbations of $X_{3}, \ldots, X_{6}$ ?
- We know $q(x):=r /(x-\alpha)^{2}$ has no positive roots.


## Making it robust in all cases

- Can create another ninth degree polynomial $p_{6}$ from $X_{3}, X_{4}, X_{5}, X_{6}$.
- Then $\alpha$ is the unique positive root of

$$
r(\alpha):=p_{5}(\alpha)^{2}+p_{6}(\alpha)^{2}=0
$$

- How robust is the solution to perturbations of $X_{3}, \ldots, X_{6}$ ?
- We know $q(x):=r /(x-\alpha)^{2}$ has no positive roots.
- By compactness: $q(x) \geq c>0$ for some constant $c$.


## Making it robust in all cases

- Can create another ninth degree polynomial $p_{6}$ from $X_{3}, X_{4}, X_{5}, X_{6}$.
- Then $\alpha$ is the unique positive root of

$$
r(\alpha):=p_{5}(\alpha)^{2}+p_{6}(\alpha)^{2}=0
$$

- How robust is the solution to perturbations of $X_{3}, \ldots, X_{6}$ ?
- We know $q(x):=r /(x-\alpha)^{2}$ has no positive roots.
- By compactness: $q(x) \geq c>0$ for some constant $c$.
- Therefore plugging in empirical moments $\widetilde{X}_{i}$ to estimate polynomials $p_{5}, p_{6}$ is robust:


## Making it robust in all cases

- Can create another ninth degree polynomial $p_{6}$ from $X_{3}, X_{4}, X_{5}, X_{6}$.
- Then $\alpha$ is the unique positive root of

$$
r(\alpha):=p_{5}(\alpha)^{2}+p_{6}(\alpha)^{2}=0
$$

- How robust is the solution to perturbations of $X_{3}, \ldots, X_{6}$ ?
- We know $q(x):=r /(x-\alpha)^{2}$ has no positive roots.
- By compactness: $q(x) \geq c>0$ for some constant $c$.
- Therefore plugging in empirical moments $\widetilde{X}_{i}$ to estimate polynomials $p_{5}, p_{6}$ is robust:
- Given approximations $\left|\widetilde{p}_{5}-p_{5}\right|,\left|\widetilde{p}_{6}-p_{6}\right| \leq \epsilon$,

$$
|\alpha-\arg \min \widetilde{r}(x)| \lesssim \epsilon .
$$

## Making it robust in all cases

- Can create another ninth degree polynomial $p_{6}$ from $X_{3}, X_{4}, X_{5}, X_{6}$.
- Then $\alpha$ is the unique positive root of

$$
r(\alpha):=p_{5}(\alpha)^{2}+p_{6}(\alpha)^{2}=0
$$

- How robust is the solution to perturbations of $X_{3}, \ldots, X_{6}$ ?
- We know $q(x):=r /(x-\alpha)^{2}$ has no positive roots.
- By compactness: $q(x) \geq c>0$ for some constant $c$.
- Therefore plugging in empirical moments $\widetilde{X}_{i}$ to estimate polynomials $p_{5}, p_{6}$ is robust:
- Given approximations $\left|\widetilde{p}_{5}-p_{5}\right|,\left|\widetilde{p}_{6}-p_{6}\right| \leq \epsilon$,

$$
|\alpha-\arg \min \widetilde{r}(x)| \lesssim \epsilon .
$$

- Getting $\alpha$ lets us estimate means, variances.


## Result



- Scale so the excess moments are $O(1): \mu_{i}$ are $\pm O(1)$.


## Result



- Scale so the excess moments are $O(1): \mu_{i}$ are $\pm O(1)$.
- Getting the $\widetilde{p}_{i}$ to $O(\epsilon)$ requires getting the first six moments to $\pm O(\epsilon)$.


## Result



- Scale so the excess moments are $O(1): \mu_{i}$ are $\pm O(1)$.
- Getting the $\widetilde{p}_{i}$ to $O(\epsilon)$ requires getting the first six moments to $\pm O(\epsilon)$.
- If the variance is $\sigma^{2}$, then $M_{i}$ has variance $O\left(\sigma^{2 i}\right)$.


## Result



- Scale so the excess moments are $O(1): \mu_{i}$ are $\pm O(1)$.
- Getting the $\widetilde{p}_{i}$ to $O(\epsilon)$ requires getting the first six moments to $\pm O(\epsilon)$.
- If the variance is $\sigma^{2}$, then $M_{i}$ has variance $O\left(\sigma^{2 i}\right)$.
- Thus $O\left(\sigma^{12} / \epsilon^{2}\right)$ samples to learn the $\mu_{i}$ to $\pm \epsilon$.


## Result



- Scale so the excess moments are $O(1): \mu_{i}$ are $\pm O(1)$.
- Getting the $\widetilde{p}_{i}$ to $O(\epsilon)$ requires getting the first six moments to $\pm O(\epsilon)$.
- If the variance is $\sigma^{2}$, then $M_{i}$ has variance $O\left(\sigma^{2 i}\right)$.
- Thus $O\left(\sigma^{12} / \epsilon^{2}\right)$ samples to learn the $\mu_{i}$ to $\pm \epsilon$.
- If components are $\Omega(1)$ standard deviations apart, $O\left(1 / \epsilon^{2}\right)$ samples suffice.


## Result



- Scale so the excess moments are $O(1): \mu_{i}$ are $\pm O(1)$.
- Getting the $\widetilde{p}_{i}$ to $O(\epsilon)$ requires getting the first six moments to $\pm O(\epsilon)$.
- If the variance is $\sigma^{2}$, then $M_{i}$ has variance $O\left(\sigma^{2 i}\right)$.
- Thus $O\left(\sigma^{12} / \epsilon^{2}\right)$ samples to learn the $\mu_{i}$ to $\pm \epsilon$.
- If components are $\Omega(1)$ standard deviations apart, $O\left(1 / \epsilon^{2}\right)$ samples suffice.
- In general, $O\left(1 / \epsilon^{12}\right)$ samples suffice to get $\epsilon \sigma$ accuracy.


## Outline

## (1) Algorithm in One Dimension

## 3) Algorithm in $d$ Dimensions

## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments


## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments
- Necessary to get sixth moment to $\pm(\epsilon \sigma)^{6}$.


## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments
- Necessary to get sixth moment to $\pm(\epsilon \sigma)^{6}$.
- Let $F, F^{\prime}$ be any two mixtures with five matching moments:

- Constant means and variances.


## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments
- Necessary to get sixth moment to $\pm(\epsilon \sigma)^{6}$.
- Let $F, F^{\prime}$ be any two mixtures with five matching moments:

- Constant means and variances.
- Add $N\left(0, \sigma^{2}\right)$ to each mixture for growing $\sigma$.


## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments
- Necessary to get sixth moment to $\pm(\epsilon \sigma)^{6}$.
- Let $F, F^{\prime}$ be any two mixtures with five matching moments:

- Constant means and variances.
- Add $N\left(0, \sigma^{2}\right)$ to each mixture for growing $\sigma$.


## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments
- Necessary to get sixth moment to $\pm(\epsilon \sigma)^{6}$.
- Let $F, F^{\prime}$ be any two mixtures with five matching moments:

- Constant means and variances.
- Add $N\left(0, \sigma^{2}\right)$ to each mixture for growing $\sigma$.


## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments
- Necessary to get sixth moment to $\pm(\epsilon \sigma)^{6}$.
- Let $F, F^{\prime}$ be any two mixtures with five matching moments:

- Constant means and variances.
- Add $N\left(0, \sigma^{2}\right)$ to each mixture for growing $\sigma$.


## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments
- Necessary to get sixth moment to $\pm(\epsilon \sigma)^{6}$.
- Let $F, F^{\prime}$ be any two mixtures with five matching moments:

- Constant means and variances.
- Add $N\left(0, \sigma^{2}\right)$ to each mixture for growing $\sigma$.


## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments
- Necessary to get sixth moment to $\pm(\epsilon \sigma)^{6}$.
- Let $F, F^{\prime}$ be any two mixtures with five matching moments:

- Constant means and variances.
- Add $N\left(0, \sigma^{2}\right)$ to each mixture for growing $\sigma$.


## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments
- Necessary to get sixth moment to $\pm(\epsilon \sigma)^{6}$.
- Let $F, F^{\prime}$ be any two mixtures with five matching moments:

- Constant means and variances.
- Add $N\left(0, \sigma^{2}\right)$ to each mixture for growing $\sigma$.


## Lower bound in one dimension

- The algorithm takes $O\left(\epsilon^{-12}\right)$ samples because it uses six moments
- Necessary to get sixth moment to $\pm(\epsilon \sigma)^{6}$.
- Let $F, F^{\prime}$ be any two mixtures with five matching moments:

- Constant means and variances.
- Add $N\left(0, \sigma^{2}\right)$ to each mixture for growing $\sigma$.
- Claim: $\Omega\left(\sigma^{12}\right)$ samples necessary to distinguish the distributions.


## Lower bound in one dimension

- Two mixtures $F, F^{\prime}$ with $F \approx F^{\prime}$.



## Lower bound in one dimension

- Two mixtures $F, F^{\prime}$ with $F \approx F^{\prime}$.
- Have $\operatorname{TV}\left(F, F^{\prime}\right) \approx 1 / \sigma^{6}$.



## Lower bound in one dimension

- Two mixtures $F, F^{\prime}$ with $F \approx F^{\prime}$.
- Have $\operatorname{TV}\left(F, F^{\prime}\right) \approx 1 / \sigma^{6}$.
- Shows $\Omega\left(\sigma^{6}\right)$ samples, $O\left(\sigma^{12}\right)$ samples.



## Lower bound in one dimension

- Two mixtures $F, F^{\prime}$ with $F \approx F^{\prime}$.
- Have $\operatorname{TV}\left(F, F^{\prime}\right) \approx 1 / \sigma^{6}$.
- Shows $\Omega\left(\sigma^{6}\right)$ samples, $O\left(\sigma^{12}\right)$ samples.

- Improve using squared Hellinger distance.


## Lower bound in one dimension

- Two mixtures $F, F^{\prime}$ with $F \approx F^{\prime}$.
- Have $\operatorname{TV}\left(F, F^{\prime}\right) \approx 1 / \sigma^{6}$.
- Shows $\Omega\left(\sigma^{6}\right)$ samples, $O\left(\sigma^{12}\right)$ samples.

- Improve using squared Hellinger distance.
- $H^{2}(P, Q):=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x$


## Lower bound in one dimension

- Two mixtures $F, F^{\prime}$ with $F \approx F^{\prime}$.
- Have TV $\left(F, F^{\prime}\right) \approx 1 / \sigma^{6}$.
- Shows $\Omega\left(\sigma^{6}\right)$ samples, $O\left(\sigma^{12}\right)$ samples.

- Improve using squared Hellinger distance.
- $H^{2}(P, Q):=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x$
- $H^{2}$ is subadditive on product measures:


## Lower bound in one dimension

- Two mixtures $F, F^{\prime}$ with $F \approx F^{\prime}$.
- Have TV $\left(F, F^{\prime}\right) \approx 1 / \sigma^{6}$.
- Shows $\Omega\left(\sigma^{6}\right)$ samples, $O\left(\sigma^{12}\right)$ samples.

- Improve using squared Hellinger distance.
- $H^{2}(P, Q):=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x$
- $H^{2}$ is subadditive on product measures:

$$
\star H^{2}\left(\left(x_{1}, \ldots, x_{m}\right),\left(x_{1}^{\prime}, \ldots, x_{m}^{\prime}\right)\right) \leq m H^{2}\left(x, x^{\prime}\right)
$$

## Lower bound in one dimension

- Two mixtures $F, F^{\prime}$ with $F \approx F^{\prime}$.
- Have TV $\left(F, F^{\prime}\right) \approx 1 / \sigma^{6}$.
- Shows $\Omega\left(\sigma^{6}\right)$ samples, $O\left(\sigma^{12}\right)$ samples.

- Improve using squared Hellinger distance.
- $H^{2}(P, Q):=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x$
- $H^{2}$ is subadditive on product measures:

$$
\star H^{2}\left(\left(x_{1}, \ldots, x_{m}\right),\left(x_{1}^{\prime}, \ldots, x_{m}^{\prime}\right)\right) \leq m H^{2}\left(x, x^{\prime}\right) .
$$

- Sample complexity is $\Omega\left(1 / H^{2}\left(F, F^{\prime}\right)\right)$


## Lower bound in one dimension

- Two mixtures $F, F^{\prime}$ with $F \approx F^{\prime}$.
- Have TV $\left(F, F^{\prime}\right) \approx 1 / \sigma^{6}$.
- Shows $\Omega\left(\sigma^{6}\right)$ samples, $O\left(\sigma^{12}\right)$ samples.

- Improve using squared Hellinger distance.
- $H^{2}(P, Q):=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x$
- $H^{2}$ is subadditive on product measures:

$$
\star H^{2}\left(\left(x_{1}, \ldots, x_{m}\right),\left(x_{1}^{\prime}, \ldots, x_{m}^{\prime}\right)\right) \leq m H^{2}\left(x, x^{\prime}\right) .
$$

- Sample complexity is $\Omega\left(1 / H^{2}\left(F, F^{\prime}\right)\right)$
- $H^{2} \lesssim T V \lesssim H$, but often $H \approx T V$.


## Bounding the Hellinger distance: general idea

## Definition

$$
H^{2}(P, Q)=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x
$$

## Bounding the Hellinger distance: general idea

## Definition

$$
H^{2}(P, Q)=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x=1-\int \sqrt{p(x) q(x)} d x
$$

## Bounding the Hellinger distance: general idea

## Definition

$$
H^{2}(P, Q)=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x=1-\int \sqrt{p(x) q(x)} d x
$$

- If $q(x)=(1+\Delta(x)) p(x)$ for some small $\Delta$, then [Pollard '00]


## Bounding the Hellinger distance: general idea

## Definition

$$
H^{2}(P, Q)=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x=1-\int \sqrt{p(x) q(x)} d x
$$

- If $q(x)=(1+\Delta(x)) p(x)$ for some small $\Delta$, then [Pollard '00]

$$
H^{2}(p, q)=1-\int \sqrt{1+\Delta(x)} p(x) d x
$$

## Bounding the Hellinger distance: general idea

## Definition

$$
H^{2}(P, Q)=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x=1-\int \sqrt{p(x) q(x)} d x
$$

- If $q(x)=(1+\Delta(x)) p(x)$ for some small $\Delta$, then [Pollard '00]

$$
\begin{aligned}
H^{2}(p, q) & =1-\int \sqrt{1+\Delta(x)} p(x) d x \\
& =1-\underset{x \sim p}{\mathbb{E}}[\sqrt{1+\Delta(x)}]
\end{aligned}
$$

## Bounding the Hellinger distance: general idea

## Definition

$$
H^{2}(P, Q)=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x=1-\int \sqrt{p(x) q(x)} d x
$$

- If $q(x)=(1+\Delta(x)) p(x)$ for some small $\Delta$, then [Pollard '00]

$$
\begin{aligned}
H^{2}(p, q) & =1-\int \sqrt{1+\Delta(x)} p(x) d x \\
& =1-\underset{x \sim p}{\mathbb{E}}[\sqrt{1+\Delta(x)}] \\
& =1-\underset{x \sim p}{\mathbb{E}}\left[1+\Delta(x) / 2-O\left(\Delta^{2}(x)\right)\right]
\end{aligned}
$$

## Bounding the Hellinger distance: general idea

## Definition

$$
H^{2}(P, Q)=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x=1-\int \sqrt{p(x) q(x)} d x
$$

- If $q(x)=(1+\Delta(x)) p(x)$ for some small $\Delta$, then [Pollard '00]

$$
\begin{aligned}
H^{2}(p, q) & =1-\int \sqrt{1+\Delta(x)} p(x) d x \\
& =1-\underset{x \sim p}{\mathbb{E}}[\sqrt{1+\Delta(x)}] \\
& =1-\underset{x \sim p}{\mathbb{E}}\left[1+\Delta(x) / 2-O\left(\Delta^{2}(x)\right)\right]
\end{aligned}
$$

## Bounding the Hellinger distance: general idea

## Definition

$$
H^{2}(P, Q)=\frac{1}{2} \int(\sqrt{p(x)}-\sqrt{q(x)})^{2} d x=1-\int \sqrt{p(x) q(x)} d x
$$

- If $q(x)=(1+\Delta(x)) p(x)$ for some small $\Delta$, then [Pollard '00]

$$
\begin{aligned}
H^{2}(p, q) & =1-\int \sqrt{1+\Delta(x)} p(x) d x \\
& =1-\underset{x \sim p}{\mathbb{E}}[\sqrt{1+\Delta(x)}] \\
& =1-\underset{x \sim p}{\mathbb{E}}[1+\underbrace{\Delta(x)}_{\int q(x)-p(x)=0} / 2-O\left(\Delta^{2}(x)\right)]
\end{aligned}
$$

## Bounding the Hellinger distance: general idea

## Definition

$$
H^{2}(P, Q)=\frac{1}{2} \int\left(\sqrt{p(x)}-\sqrt{q(x))^{2}} d x=1-\int \sqrt{p(x) q(x)} d x\right.
$$

- If $q(x)=(1+\Delta(x)) p(x)$ for some small $\Delta$, then [Pollard '00]

$$
\begin{aligned}
H^{2}(p, q) & =1-\int \sqrt{1+\Delta(x)} p(x) d x \\
& =1-\underset{x \sim p}{\mathbb{E}}[\sqrt{1+\Delta(x)}] \\
& =1-\underset{x \sim p}{\mathbb{E}}[1+\underbrace{1+(x)} / 2-O\left(\Delta^{2}(x)\right)] \\
& \lesssim \underset{x \sim p}{\mathbb{E}}\left[\Delta^{2}(x)\right]
\end{aligned}
$$

## Bounding the Hellinger distance: general idea

## Definition

$$
H^{2}(P, Q)=\frac{1}{2} \int\left(\sqrt{p(x)}-\sqrt{q(x))^{2}} d x=1-\int \sqrt{p(x) q(x)} d x\right.
$$

- If $q(x)=(1+\Delta(x)) p(x)$ for some small $\Delta$, then [Pollard '00]

$$
\begin{aligned}
H^{2}(p, q) & =1-\int \sqrt{1+\Delta(x)} p(x) d x \\
& =1-\underset{x \sim p}{\mathbb{E}}[\sqrt{1+\Delta(x)}] \\
& =1-\underset{x \sim p}{\mathbb{E}}[1+\underbrace{1+(x)} / 2-O\left(\Delta^{2}(x)\right)] \\
& \lesssim \underset{x \sim p}{\mathbb{E}}\left[\Delta^{2}(x)\right]
\end{aligned}
$$

- Compare to $T V(p, q)=\frac{1}{2} \mathbb{E}_{x \sim p}[|\Delta(x)|]$


## Bounding the Hellinger distance: our setting

## Lemma

Let $F, F^{\prime}$ be two subgaussian distributions with $k$ matching moments and constant parameters. Then for $G, G^{\prime}=F+N\left(0, \sigma^{2}\right), F^{\prime}+N\left(0, \sigma^{2}\right)$,

$$
H^{2}\left(G, G^{\prime}\right) \lesssim 1 / \sigma^{2 k+2} .
$$

## Bounding the Hellinger distance: our setting

## Lemma

Let $F, F^{\prime}$ be two subgaussian distributions with $k$ matching moments and constant parameters. Then for $G, G^{\prime}=F+N\left(0, \sigma^{2}\right), F^{\prime}+N\left(0, \sigma^{2}\right)$,

$$
H^{2}\left(G, G^{\prime}\right) \lesssim 1 / \sigma^{2 k+2} .
$$

- Power series expansion of $\mathbb{E}\left[\Delta^{2}\right]=\mathbb{E}\left[\left(\frac{G^{\prime}(x)-G(x)}{G(x)}\right)^{2}\right]$.


## Bounding the Hellinger distance: our setting

## Lemma

Let $F, F^{\prime}$ be two subgaussian distributions with $k$ matching moments and constant parameters. Then for $G, G^{\prime}=F+N\left(0, \sigma^{2}\right), F^{\prime}+N\left(0, \sigma^{2}\right)$,

$$
H^{2}\left(G, G^{\prime}\right) \lesssim 1 / \sigma^{2 k+2} .
$$

- Power series expansion of $\mathbb{E}\left[\Delta^{2}\right]=\mathbb{E}\left[\left(\frac{G^{\prime}(x)-G(x)}{G(x)}\right)^{2}\right]$.
- Matching moments make the first $k$ terms zero.


## Bounding the Hellinger distance: our setting

## Lemma

Let $F, F^{\prime}$ be two subgaussian distributions with $k$ matching moments and constant parameters. Then for $G, G^{\prime}=F+N\left(0, \sigma^{2}\right), F^{\prime}+N\left(0, \sigma^{2}\right)$,

$$
H^{2}\left(G, G^{\prime}\right) \lesssim 1 / \sigma^{2 k+2} .
$$

- Power series expansion of $\mathbb{E}\left[\Delta^{2}\right]=\mathbb{E}\left[\left(\frac{G^{\prime}(x)-G(x)}{G(x)}\right)^{2}\right]$.
- Matching moments make the first $k$ terms zero.
- Leaves $\left(1 / \sigma^{k+1}\right)^{2}$ as largest remaining term.


## Lower bound in one dimension

- Add $N\left(0, \sigma^{2}\right)$ to two mixtures with five matching moments.



## Lower bound in one dimension

- Add $N\left(0, \sigma^{2}\right)$ to two mixtures with five matching moments.



## Lower bound in one dimension

- Add $N\left(0, \sigma^{2}\right)$ to two mixtures with five matching moments.



## Lower bound in one dimension

- Add $N\left(0, \sigma^{2}\right)$ to two mixtures with five matching moments.



## Lower bound in one dimension

- Add $N\left(0, \sigma^{2}\right)$ to two mixtures with five matching moments.



## Lower bound in one dimension

- Add $N\left(0, \sigma^{2}\right)$ to two mixtures with five matching moments.



## Lower bound in one dimension

- Add $N\left(0, \sigma^{2}\right)$ to two mixtures with five matching moments.




## Lower bound in one dimension

- Add $N\left(0, \sigma^{2}\right)$ to two mixtures with five matching moments.


- For

$$
\begin{aligned}
G & =\frac{1}{2} N\left(-1,1+\sigma^{2}\right)+\frac{1}{2} N\left(1,2+\sigma^{2}\right) \\
G^{\prime} & \approx 0.297 N\left(-1.226,0.610+\sigma^{2}\right)+0.703 N\left(0.517,2.396+\sigma^{2}\right)
\end{aligned}
$$

have $H^{2}\left(G, G^{\prime}\right) \lesssim 1 / \sigma^{12}$.

## Lower bound in one dimension

- Add $N\left(0, \sigma^{2}\right)$ to two mixtures with five matching moments.


- For

$$
\begin{aligned}
G & =\frac{1}{2} N\left(-1,1+\sigma^{2}\right)+\frac{1}{2} N\left(1,2+\sigma^{2}\right) \\
G^{\prime} & \approx 0.297 N\left(-1.226,0.610+\sigma^{2}\right)+0.703 N\left(0.517,2.396+\sigma^{2}\right)
\end{aligned}
$$

have $H^{2}\left(G, G^{\prime}\right) \lesssim 1 / \sigma^{12}$.

- Therefore distinguishing $G$ from $G^{\prime}$ takes $\Omega\left(\sigma^{12}\right)$ samples.


## Lower bound in one dimension

- Add $N\left(0, \sigma^{2}\right)$ to two mixtures with five matching moments.

- For

$$
\begin{aligned}
G & =\frac{1}{2} N\left(-1,1+\sigma^{2}\right)+\frac{1}{2} N\left(1,2+\sigma^{2}\right) \\
G^{\prime} & \approx 0.297 N\left(-1.226,0.610+\sigma^{2}\right)+0.703 N\left(0.517,2.396+\sigma^{2}\right)
\end{aligned}
$$

have $H^{2}\left(G, G^{\prime}\right) \lesssim 1 / \sigma^{12}$.

- Therefore distinguishing $G$ from $G^{\prime}$ takes $\Omega\left(\sigma^{12}\right)$ samples.
- Cannot learn either means to $\pm \epsilon \sigma$ or variance to $\pm \epsilon^{2} \sigma^{2}$ with $o\left(1 / \epsilon^{12}\right)$ samples.


## Lower bound in $d$ dimensions

- Trivial based on the Hellinger distance bound.


## Lower bound in $d$ dimensions

- Trivial based on the Hellinger distance bound.
- Place the "hard" instance independently in all $d$ coordinates.


## Lower bound in $d$ dimensions

- Trivial based on the Hellinger distance bound.
- Place the "hard" instance independently in all $d$ coordinates.
- Solution must solve all $d$ instances.


## Lower bound in $d$ dimensions

- Trivial based on the Hellinger distance bound.
- Place the "hard" instance independently in all $d$ coordinates.
- Solution must solve all $d$ instances.
- Each instance has Hellinger distance $O\left(\epsilon^{12}\right)$.


## Lower bound in $d$ dimensions

- Trivial based on the Hellinger distance bound.
- Place the "hard" instance independently in all $d$ coordinates.
- Solution must solve all $d$ instances.
- Each instance has Hellinger distance $O\left(\epsilon^{12}\right)$.
- Therefore $\Omega\left(\epsilon^{-12} \log (d / \delta)\right)$ samples are necessary to succeed with probability $1-\delta$ :


## Lower bound in $d$ dimensions

- Trivial based on the Hellinger distance bound.
- Place the "hard" instance independently in all $d$ coordinates.
- Solution must solve all $d$ instances.
- Each instance has Hellinger distance $O\left(\epsilon^{12}\right)$.
- Therefore $\Omega\left(\epsilon^{-12} \log (d / \delta)\right)$ samples are necessary to succeed with probability $1-\delta$ :
- Each set of $\epsilon^{-12}$ samples has a constant chance of giving no information about each coordinate.


## Lower bound in $d$ dimensions

- Trivial based on the Hellinger distance bound.
- Place the "hard" instance independently in all $d$ coordinates.
- Solution must solve all $d$ instances.
- Each instance has Hellinger distance $O\left(\epsilon^{12}\right)$.
- Therefore $\Omega\left(\epsilon^{-12} \log (d / \delta)\right)$ samples are necessary to succeed with probability $1-\delta$ :
- Each set of $\epsilon^{-12}$ samples has a constant chance of giving no information about each coordinate.
- With $o\left(\epsilon^{-12} \log d\right)$ samples, some coordinate will be independent of all the samples.


## Outline

## (1) Algorithm in One Dimension

(2) Lower Bound
(3) Algorithm in $d$ Dimensions

## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...


## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...
- (And covariance matrix)


## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...
- (And covariance matrix)
- Look at individual attributes to get all these.


## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...
- (And covariance matrix)
- Look at individual attributes to get all these.
- Just need to know: is the taller group also heavier or lighter?


## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...
- (And covariance matrix)
- Look at individual attributes to get all these.
- Just need to know: is the taller group also heavier or lighter?
- Suffices to consider $d=2$ :


## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...
- (And covariance matrix)
- Look at individual attributes to get all these.
- Just need to know: is the taller group also heavier or lighter?
- Suffices to consider $d=2$ :
- Does $\mu_{i}$ go with $\mu_{j}$ or $\mu_{j}^{\prime}$ ?


## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...
- (And covariance matrix)
- Look at individual attributes to get all these.
- Just need to know: is the taller group also heavier or lighter?
- Suffices to consider $d=2$ :
- Does $\mu_{i}$ go with $\mu_{j}$ or $\mu_{j}^{\prime}$ ?
- Project onto a random direction $e_{i} \sin \theta+e_{j} \cos \theta$.


## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...
- (And covariance matrix)
- Look at individual attributes to get all these.
- Just need to know: is the taller group also heavier or lighter?
- Suffices to consider $d=2$ :
- Does $\mu_{i}$ go with $\mu_{j}$ or $\mu_{j}^{\prime}$ ?
- Project onto a random direction $e_{i} \sin \theta+e_{j} \cos \theta$.
- $\left(\mu_{i}, \mu_{j}\right)$ usually has a significantly different projection from $\left(\mu_{i}, \mu_{j}^{\prime}\right)$.


## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...
- (And covariance matrix)
- Look at individual attributes to get all these.
- Just need to know: is the taller group also heavier or lighter?
- Suffices to consider $d=2$ :
- Does $\mu_{i}$ go with $\mu_{j}$ or $\mu_{j}^{\prime}$ ?
- Project onto a random direction $e_{i} \sin \theta+e_{j} \cos \theta$.
- $\left(\mu_{i}, \mu_{j}\right)$ usually has a significantly different projection from $\left(\mu_{i}, \mu_{j}^{\prime}\right)$.
- Thus we can piece them together by solving the $O\left(d^{2}\right)$ one dimensional problems.


## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...
- (And covariance matrix)
- Look at individual attributes to get all these.
- Just need to know: is the taller group also heavier or lighter?
- Suffices to consider $d=2$ :
- Does $\mu_{i}$ go with $\mu_{j}$ or $\mu_{j}^{\prime}$ ?
- Project onto a random direction $e_{i} \sin \theta+e_{j} \cos \theta$.
- $\left(\mu_{i}, \mu_{j}\right)$ usually has a significantly different projection from $\left(\mu_{i}, \mu_{j}^{\prime}\right)$.
- Thus we can piece them together by solving the $O\left(d^{2}\right)$ one dimensional problems.
- For covariances: reduce to $d=4$, so $O\left(d^{4}\right)$ one dimensional problems.


## Algorithm in $d$ dimensions

- Want to learn average male/female height, weight, shoe size, ...
- (And covariance matrix)
- Look at individual attributes to get all these.
- Just need to know: is the taller group also heavier or lighter?
- Suffices to consider $d=2$ :
- Does $\mu_{i}$ go with $\mu_{j}$ or $\mu_{j}^{\prime}$ ?
- Project onto a random direction $e_{i} \sin \theta+e_{j} \cos \theta$.
- $\left(\mu_{i}, \mu_{j}\right)$ usually has a significantly different projection from $\left(\mu_{i}, \mu_{j}^{\prime}\right)$.
- Thus we can piece them together by solving the $O\left(d^{2}\right)$ one dimensional problems.
- For covariances: reduce to $d=4$, so $O\left(d^{4}\right)$ one dimensional problems.
- Only loss is $\log (1 / \delta) \rightarrow \log (d / \delta):$

$$
\Theta\left(1 / \epsilon^{12} \log (d / \delta)\right) \text { samples }
$$

## Recap and open questions

- Our result:
- $\Theta\left(\epsilon^{-12} \log d\right)$ samples necessary and sufficient to estimate $\mu_{i}$ to $\pm \epsilon \sigma, \sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$.


## Recap and open questions

- Our result:
- $\Theta\left(\epsilon^{-12} \log d\right)$ samples necessary and sufficient to estimate $\mu_{i}$ to $\pm \epsilon \sigma, \sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$.
- If the means have $\alpha \sigma$ separation, just $O\left(\epsilon^{-2} \alpha^{-12}\right)$ for $\epsilon \alpha \sigma$ accuracy.


## Recap and open questions

- Our result:
- $\Theta\left(\epsilon^{-12} \log d\right)$ samples necessary and sufficient to estimate $\mu_{i}$ to $\pm \epsilon \sigma, \sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$.
- If the means have $\alpha \sigma$ separation, just $O\left(\epsilon^{-2} \alpha^{-12}\right)$ for $\epsilon \alpha \sigma$ accuracy.
- Extend to $k>2$ ?


## Recap and open questions

- Our result:
- $\Theta\left(\epsilon^{-12} \log d\right)$ samples necessary and sufficient to estimate $\mu_{i}$ to $\pm \epsilon \sigma, \sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$.
- If the means have $\alpha \sigma$ separation, just $O\left(\epsilon^{-2} \alpha^{-12}\right)$ for $\epsilon \alpha \sigma$ accuracy.
- Extend to $k>2$ ?
- Lower bound extends, at least to $\Omega\left(\epsilon^{-6 k-2}\right)$.


## Recap and open questions

- Our result:
- $\Theta\left(\epsilon^{-12} \log d\right)$ samples necessary and sufficient to estimate $\mu_{i}$ to $\pm \epsilon \sigma, \sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$.
- If the means have $\alpha \sigma$ separation, just $O\left(\epsilon^{-2} \alpha^{-12}\right)$ for $\epsilon \alpha \sigma$ accuracy.
- Extend to $k>2$ ?
- Lower bound extends, at least to $\Omega\left(\epsilon^{-6 k-2}\right)$.
- Do we really care about finding an $O\left(\epsilon^{-22}\right)$ algorithm?


## Recap and open questions

- Our result:
- $\Theta\left(\epsilon^{-12} \log d\right)$ samples necessary and sufficient to estimate $\mu_{i}$ to $\pm \epsilon \sigma, \sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$.
- If the means have $\alpha \sigma$ separation, just $O\left(\epsilon^{-2} \alpha^{-12}\right)$ for $\epsilon \alpha \sigma$ accuracy.
- Extend to $k>2$ ?
- Lower bound extends, at least to $\Omega\left(\epsilon^{-6 k-2}\right)$.
- Do we really care about finding an $O\left(\epsilon^{-22}\right)$ algorithm?
- Solving the system of equations gets nasty.


## Recap and open questions

- Our result:
- $\Theta\left(\epsilon^{-12} \log d\right)$ samples necessary and sufficient to estimate $\mu_{i}$ to $\pm \epsilon \sigma, \sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$.
- If the means have $\alpha \sigma$ separation, just $O\left(\epsilon^{-2} \alpha^{-12}\right)$ for $\epsilon \alpha \sigma$ accuracy.
- Extend to $k>2$ ?
- Lower bound extends, at least to $\Omega\left(\epsilon^{-6 k-2}\right)$.
- Do we really care about finding an $O\left(\epsilon^{-22}\right)$ algorithm?
- Solving the system of equations gets nasty.
- [Next talk: Ge-Huang-Kakade avoid this for smoothed instances]


## Recap and open questions

- Our result:
- $\Theta\left(\epsilon^{-12} \log d\right)$ samples necessary and sufficient to estimate $\mu_{i}$ to $\pm \epsilon \sigma, \sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$.
- If the means have $\alpha \sigma$ separation, just $O\left(\epsilon^{-2} \alpha^{-12}\right)$ for $\epsilon \alpha \sigma$ accuracy.
- Extend to $k>2$ ?
- Lower bound extends, at least to $\Omega\left(\epsilon^{-6 k-2}\right)$.
- Do we really care about finding an $O\left(\epsilon^{-22}\right)$ algorithm?
- Solving the system of equations gets nasty.
- [Next talk: Ge-Huang-Kakade avoid this for smoothed instances]
- Automated way of figuring out whether solution to system of polynomial equations is robust?


## Recap and open questions

- Our result:
- $\Theta\left(\epsilon^{-12} \log d\right)$ samples necessary and sufficient to estimate $\mu_{i}$ to $\pm \epsilon \sigma, \sigma_{i}^{2}$ to $\pm \epsilon^{2} \sigma^{2}$.
- If the means have $\alpha \sigma$ separation, just $O\left(\epsilon^{-2} \alpha^{-12}\right)$ for $\epsilon \alpha \sigma$ accuracy.
- Extend to $k>2$ ?
- Lower bound extends, at least to $\Omega\left(\epsilon^{-6 k-2}\right)$.
- Do we really care about finding an $O\left(\epsilon^{-22}\right)$ algorithm?
- Solving the system of equations gets nasty.
- [Next talk: Ge-Huang-Kakade avoid this for smoothed instances]
- Automated way of figuring out whether solution to system of polynomial equations is robust?
- TV estimation in $d$ dimensions with $d / \epsilon^{c}$ rather than $d^{30} / \epsilon^{c}$ ?

