Gaussian noise stability

Elchanan Mossel[1]   Joe Neeman[2]

[1]UC Berkeley

[2]UT Austin

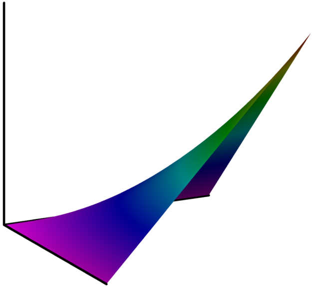

Gaussian noise stability

Fix a parameter $0 < \rho < 1$. Take

$$(X, Y) \sim \mathcal{N}\left(0, \begin{pmatrix} I_n & \rho I_n \\ \rho I_n & I_n \end{pmatrix}\right)$$

Gaussian noise stability

Fix a parameter $0 < \rho < 1$. Take

$$(X, Y) \sim \mathcal{N}\left(0, \begin{pmatrix} I_n & \rho I_n \\ \rho I_n & I_n \end{pmatrix}\right)$$

The *Gaussian noise stability* of $A \subset \mathbb{R}^n$ is $\Pr(X \in A, Y \in A)$.

Gaussian noise stability

Fix a parameter $0 < \rho < 1$. Take

$$(X, Y) \sim \mathcal{N}\left(0, \begin{pmatrix} I_n & \rho I_n \\ \rho I_n & I_n \end{pmatrix}\right)$$

The *Gaussian noise stability* of $A \subset \mathbb{R}^n$ is $\Pr(X \in A, Y \in A)$.
Applications in

- approximability (e.g., optimal UGC hardness of Max-Cut, KKMO '05)
- testing (e.g., testing half-spaces, MORS '09)

Borell's theorem

What sets have high noise stability?

Borell's theorem

What sets have high noise stability?
Half-spaces maximize the noise stability (among all sets of a given volume):

Theorem (Borell '85)

For any $A \subset \mathbb{R}^n$*, if* $A' \subset \mathbb{R}^n$ *is a half-space with* $\Pr(A') = \Pr(A)$ *then*

$$\Pr(X \in A, Y \in A) \leq \Pr(X \in A', Y \in A').$$

A *half-space* is a set of the form $\{x \in \mathbb{R}^n : x \cdot a \leq b\}$.

Borell's theorem

Define $\Phi(x) = \Pr(X_1 \leq x)$. Then $\{x \in \mathbb{R}^n : x_1 \leq \Phi^{-1}(a)\}$ is a half-space of volume a. Define

$$J(a, b) = \Pr(X_1 \leq \Phi^{-1}(a), Y_1 \leq \Phi^{-1}(b)).$$

Borell's theorem

Define $\Phi(x) = \Pr(X_1 \le x)$. Then $\{x \in \mathbb{R}^n : x_1 \le \Phi^{-1}(a)\}$ is a half-space of volume a. Define

$$J(a,b) = \Pr(X_1 \le \Phi^{-1}(a), Y_1 \le \Phi^{-1}(b)).$$

Since the Gaussian measure is rotationally invariant, Borell's theorem is equivalent to

$$\Pr(X \in A, Y \in A) \le J(\Pr(A), \Pr(A)).$$

Borell's theorem

Define $\Phi(x) = \Pr(X_1 \leq x)$. Then $\{x \in \mathbb{R}^n : x_1 \leq \Phi^{-1}(a)\}$ is a half-space of volume a. Define

$$J(a, b) = \Pr(X_1 \leq \Phi^{-1}(a), Y_1 \leq \Phi^{-1}(b)).$$

Since the Gaussian measure is rotationally invariant, Borell's theorem is equivalent to

$$\Pr(X \in A, Y \in A) \leq J(\Pr(A), \Pr(A)).$$

Theorem (Mossel, N. '12)

If $\Pr(X, Y \in A) = J(\Pr(A), \Pr(A))$ *then* A *is a.s. equal to a half-space.*

Borell's theorem

Define $\Phi(x) = \Pr(X_1 \le x)$. Then $\{x \in \mathbb{R}^n : x_1 \le \Phi^{-1}(a)\}$ is a half-space of volume a. Define

$$J(a, b) = \Pr(X_1 \le \Phi^{-1}(a), Y_1 \le \Phi^{-1}(b)).$$

Since the Gaussian measure is rotationally invariant, Borell's theorem is equivalent to

$$\Pr(X \in A, Y \in A) \le J(\Pr(A), \Pr(A)).$$

Theorem (Mossel, N. '12)

If $\Pr(X, Y \in A) = J(\Pr(A), \Pr(A))$ *then* A *is a.s. equal to a half-space.*

If $\Pr(X, Y \in A) \ge J(\Pr(A), \Pr(A)) - \delta$ *then there is a half-space* B *with*

$$\Pr(A \Delta B) \le C(\rho, \Pr(A)) \delta^{c(\rho)}.$$

Borell's theorem

Define $\Phi(x) = \Pr(X_1 \le x)$. Then $\{x \in \mathbb{R}^n : x_1 \le \Phi^{-1}(a)\}$ is a half-space of volume a. Define

$$J(a, b) = \Pr(X_1 \le \Phi^{-1}(a), Y_1 \le \Phi^{-1}(b)).$$

Since the Gaussian measure is rotationally invariant, Borell's theorem is equivalent to

$$\Pr(X \in A, Y \in A) \le J(\Pr(A), \Pr(A)).$$

Theorem (Mossel, N. '12, Eldan '13)

If $\Pr(X, Y \in A) = J(\Pr(A), \Pr(A))$ *then* A *is a.s. equal to a half-space.*

If $\Pr(X, Y \in A) \ge J(\Pr(A), \Pr(A)) - \delta$ *then there is a half-space* B *with*

$$\Pr(A \Delta B) \le \frac{C(\Pr(A))}{\sqrt{1-\rho}} \sqrt{\delta \log(1/\delta)}.$$

Borell's theorem: previous proofs

- Borell's original proof, using Ehrhard symmetrization.

Borell's theorem: previous proofs

- Borell's original proof, using Ehrhard symmetrization.
- Burchard-Schmuckenschlager and Issakson-Mossel, using spherical symmetrization.

Borell's theorem: previous proofs

- Borell's original proof, using Ehrhard symmetrization.
- Burchard-Schmuckenschlager and Issakson-Mossel, using spherical symmetrization.
- Kindler-O'Donnell (when $\Pr(X \in A) = \frac{1}{2}$, and for certain values of ρ), using subadditivity.

An application: half-space testing

An application: half-space testing

Suppose we have query access to some unknown $A \subset \mathbb{R}^n$ (ie. we can ask whether $x \in A$) and we want to check if A is a half-space.

An application: half-space testing

Suppose we have query access to some unknown $A \subset \mathbb{R}^n$ (ie. we can ask whether $x \in A$) and we want to check if A is a half-space.

1. Sample $Z_1, \ldots, Z_m \sim \mathcal{N}(0, I_n)$ and let $\hat{p} = \frac{\#\{Z_i \in A\}}{m}$.
2. Sample $(X_1, Y_1), \ldots, (X_m, Y_m) \sim \Pr_\rho$. Answer "yes" if

$$\frac{\#\{i : X_i \in A, Y_i \in A\}}{m} \geq J(\hat{p}, \hat{p}) - \tilde{O}(\epsilon^2)$$

and "no" otherwise.

An application: half-space testing

Suppose we have query access to some unknown $A \subset \mathbb{R}^n$ (ie. we can ask whether $x \in A$) and we want to check if A is a half-space.

1. Sample $Z_1, \ldots, Z_m \sim \mathcal{N}(0, I_n)$ and let $\hat{p} = \frac{\#\{Z_i \in A\}}{m}$.
2. Sample $(X_1, Y_1), \ldots, (X_m, Y_m) \sim \Pr_\rho$. Answer "yes" if

$$\frac{\#\{i : X_i \in A, Y_i \in A\}}{m} \geq J(\hat{p}, \hat{p}) - \tilde{O}(\epsilon^2)$$

   and "no" otherwise.

Theorem (Mossel, N. '12, Eldan '13)

If A is a half-space, then the algorithm above answers "yes" w.h.p.
If A is ϵ-far from a half-space and $m \geq \tilde{O}(\epsilon^{-4})$ then the algorithm answers "no" w.h.p.

MORS '09 showed that a similar algorithm works if $m \geq \epsilon^{-6}$.

Proof of Borell's theorem

Proof of Borell's theorem

Recall $J(a, b) = \Pr(X_1 \leq \Phi^{-1}(a), Y_1 \leq \Phi^{-1}(b))$.

Theorem

For any $f : \mathbb{R}^n \to [0, 1]$,

$$\mathbb{E} J(f(X), f(Y)) \leq J(\mathbb{E} f, \mathbb{E} f).$$

Proof of Borell's theorem

Recall $J(a, b) = \Pr(X_1 \le \Phi^{-1}(a), Y_1 \le \Phi^{-1}(b))$.

Theorem

For any $f : \mathbb{R}^n \to [0, 1]$,

$$\mathbb{E} J(f(X), f(Y)) \le J(\mathbb{E} f, \mathbb{E} f).$$

To get the original statement,

$$\Pr(X \in A, Y \in A) \le J(\Pr(A), \Pr(A)),$$

set $f = 1_A$.
(Note that $J(1, 1) = 1$ and $J(0, 1) = J(1, 0) = J(0, 0) = 0$.)

Proof of Borell's theorem

Want to show $\mathbb{E}J(f(X), f(Y)) \leq J(\mathbb{E}f, \mathbb{E}f)$.

Define the operator P_t by

$$(P_t f)(x) = \mathbb{E}f(e^{-t}x + \sqrt{1 - e^{-2t}}X).$$

Note that $P_0 f = f$ and $P_\infty f = \mathbb{E}f$.

Proof of Borell's theorem

Want to show $\mathbb{E} J(f(X), f(Y)) \leq J(\mathbb{E} f, \mathbb{E} f)$.

Define the operator P_t by

$$(P_t f)(x) = \mathbb{E} f(e^{-t} x + \sqrt{1 - e^{-2t}} X).$$

Note that $P_0 f = f$ and $P_\infty f = \mathbb{E} f$.
Consider $\mathbb{E} J(P_t f(X), P_t f(Y))$.

Proof of Borell's theorem

Want to show $\mathbb{E}J(f(X), f(Y)) \leq J(\mathbb{E}f, \mathbb{E}f)$.

Define the operator P_t by

$$(P_t f)(x) = \mathbb{E}f(e^{-t}x + \sqrt{1 - e^{-2t}}X).$$

Note that $P_0 f = f$ and $P_\infty f = \mathbb{E}f$.
Consider $\mathbb{E}J(P_t f(X), P_t f(Y))$.

The punchline: this is an increasing function of t.

Proof of Borell's theorem

Let

$$
\begin{aligned}
v_t &= v_t(X) = \Phi^{-1}(P_t f(X)) \\
w_t &= w_t(Y) = \Phi^{-1}(P_t f(Y)).
\end{aligned}
$$

Proof of Borell's theorem

Let

$$v_t = v_t(X) = \Phi^{-1}(P_t f(X))$$
$$w_t = w_t(Y) = \Phi^{-1}(P_t f(Y)).$$

$$\frac{d}{dt}\mathbb{E}J(P_t f(X), P_t f(Y))$$

Proof of Borell's theorem

Let

$$v_t = v_t(X) = \Phi^{-1}(P_t f(X))$$
$$w_t = w_t(Y) = \Phi^{-1}(P_t f(Y)).$$

$$\frac{d}{dt}\mathbb{E}J(P_t f(X), P_t f(Y))$$
$$= \ \ldots \text{chain rule } (\times 8) \ \ldots$$

Proof of Borell's theorem

Let

$$v_t = v_t(X) = \Phi^{-1}(P_t f(X))$$
$$w_t = w_t(Y) = \Phi^{-1}(P_t f(Y)).$$

$$\frac{d}{dt}\mathbb{E}J(P_t f(X), P_t f(Y))$$
$$= \ldots \text{chain rule } (\times 8) \ldots$$
$$= \ldots \text{integrate by parts} \ldots$$

Proof of Borell's theorem

Let

$$v_t = v_t(X) = \Phi^{-1}(P_t f(X))$$
$$w_t = w_t(Y) = \Phi^{-1}(P_t f(Y)).$$

$$\begin{aligned}
&\frac{d}{dt}\mathbb{E}J(P_t f(X), P_t f(Y)) \\
&= \ \ldots \text{chain rule } (\times 8) \ \ldots \\
&= \ \ldots \text{integrate by parts} \ \ldots \\
&= \frac{\rho}{2\pi\sqrt{1-\rho^2}}\mathbb{E}e^{-(v_t^2+w_t^2-2\rho v_t w_t)/(1-\rho^2)}|\nabla v_t - \nabla w_t|^2
\end{aligned}$$

Proof of Borell's theorem

Let

$$v_t = v_t(X) = \Phi^{-1}(P_t f(X))$$
$$w_t = w_t(Y) = \Phi^{-1}(P_t f(Y)).$$

$$\begin{aligned}
&\frac{d}{dt}\mathbb{E}J(P_t f(X), P_t f(Y)) \\
&= \ \ldots \text{chain rule } (\times 8) \ \ldots \\
&= \ \ldots \text{integrate by parts} \ \ldots \\
&= \frac{\rho}{2\pi\sqrt{1-\rho^2}}\mathbb{E}e^{-(v_t^2+w_t^2-2\rho v_t w_t)/(1-\rho^2)}|\nabla v_t - \nabla w_t|^2 \\
&\geq 0 \qquad \square
\end{aligned}$$

What's going on?

Why consider $\mathbb{E}J(f(X), f(Y))$? Why does the proof work?

What's going on?

Why consider $\mathbb{E}J(f(X), f(Y))$?

What's going on?

Why consider $\mathbb{E}J(f(X), f(Y))$?

Given $f : \mathbb{R}^n \to [0,1]$, define $A_f \subset \mathbb{R}^{n+1}$ by

$$A_f = \{(x, x_{n+1}) \in \mathbb{R}^{n+1} : x_{n+1} \le \Phi^{-1}(f(x))\}.$$

What's going on?

Why consider $\mathbb{E}J(f(X), f(Y))$?
Given $f : \mathbb{R}^n \to [0,1]$, define $A_f \subset \mathbb{R}^{n+1}$ by

$$A_f = \{(x, x_{n+1}) \in \mathbb{R}^{n+1} : x_{n+1} \le \Phi^{-1}(f(x))\}.$$

Then

$$\begin{aligned}
&\Pr((X, X_{n+1}) \in A_f) \\
&= \Pr(X_{n+1} \le \Phi^{-1}(f(X))) \\
&= \mathbb{E}f(X)
\end{aligned}$$

What's going on?

Why consider $\mathbb{E}J(f(X), f(Y))$?
Given $f : \mathbb{R}^n \to [0,1]$, define $A_f \subset \mathbb{R}^{n+1}$ by

$$A_f = \{(x, x_{n+1}) \in \mathbb{R}^{n+1} : x_{n+1} \le \Phi^{-1}(f(x))\}.$$

Then

$$\begin{aligned}
&\Pr((X, X_{n+1}) \in A_f) \\
&= \Pr(X_{n+1} \le \Phi^{-1}(f(X))) \\
&= \mathbb{E}f(X)
\end{aligned}$$

and

$$\begin{aligned}
&\Pr((X, X_{n+1}) \in A_f, (Y, Y_{n+1}) \in A_f) \\
&= \Pr(X_{n+1} \le \Phi^{-1}(f(X)), Y_{n+1} \le \Phi^{-1}(f(Y))) \\
&= \mathbb{E}J(f(X), f(Y)).
\end{aligned}$$

What's going on?

Why consider $\mathbb{E}J(f(X), f(Y))$?
Given $f : \mathbb{R}^n \to [0,1]$, define $A_f \subset \mathbb{R}^{n+1}$ by

$$A_f = \{(x, x_{n+1}) \in \mathbb{R}^{n+1} : x_{n+1} \leq \Phi^{-1}(f(x))\}.$$

Then

$$\Pr((X, X_{n+1}) \in A_f) = \mathbb{E}f(X)$$
$$\Pr((X, X_{n+1}) \in A_f, (Y, Y_{n+1}) \in A_f) = \mathbb{E}J(f(X), f(Y)).$$

What's going on?

Why consider $\mathbb{E}J(f(X), f(Y))$?
Given $f : \mathbb{R}^n \to [0,1]$, define $A_f \subset \mathbb{R}^{n+1}$ by

$$A_f = \{(x, x_{n+1}) \in \mathbb{R}^{n+1} : x_{n+1} \le \Phi^{-1}(f(x))\}.$$

Then

$$\Pr((X, X_{n+1}) \in A_f) = \mathbb{E}f(X)$$
$$\Pr((X, X_{n+1}) \in A_f, (Y, Y_{n+1}) \in A_f) = \mathbb{E}J(f(X), f(Y)).$$

and so Borell's theorem (in $\mathbb{R}^{n+1}$) applied to A_f gives

$$\begin{aligned} &\mathbb{E}J(f(X), f(Y)) \\ &= \Pr((X, X_{n+1}) \in A_f, (Y, Y_{n+1}) \in A_f) \\ &\le J(\Pr(A_f), \Pr(A_f)) \\ &= J(\mathbb{E}f, \mathbb{E}f). \end{aligned}$$

What's going on?

Why does the proof work?

What's going on?

Why does the proof work?

A_f:

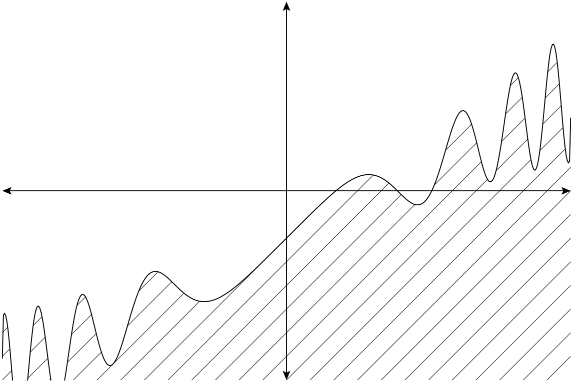

What's going on?

Why does the proof work?

$A_{P_{0.01}f}$:

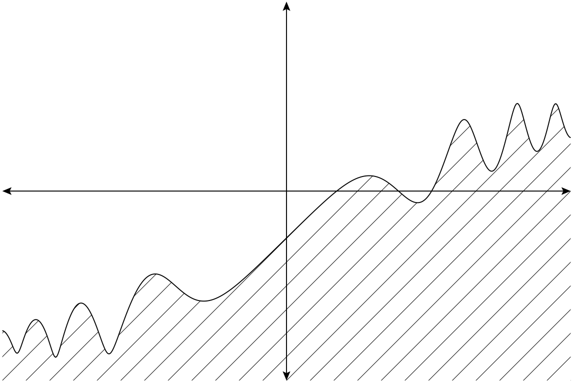

What's going on?

Why does the proof work?

$A_{P_{0.02}f}$:

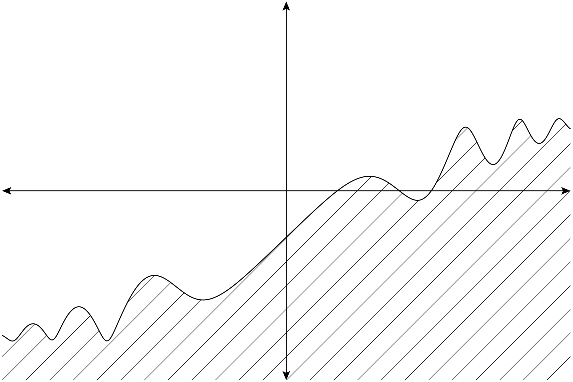

What's going on?

Why does the proof work?

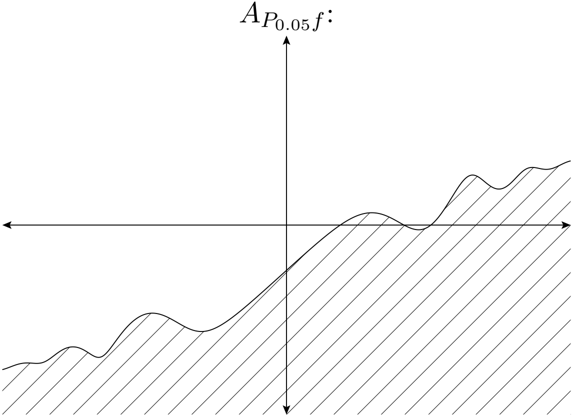

What's going on?

Why does the proof work?

$A_{P_{0.1}f}$:

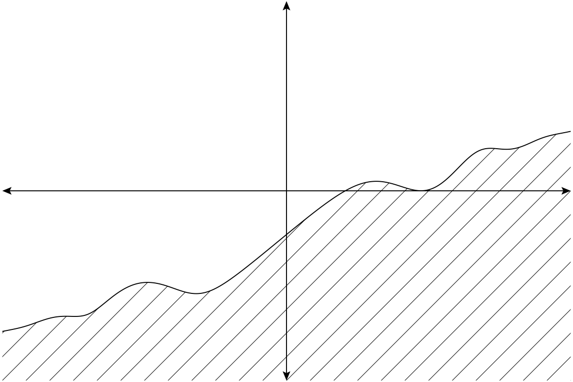

What's going on?

Why does the proof work?

$A_{P_{0.2}f}$:

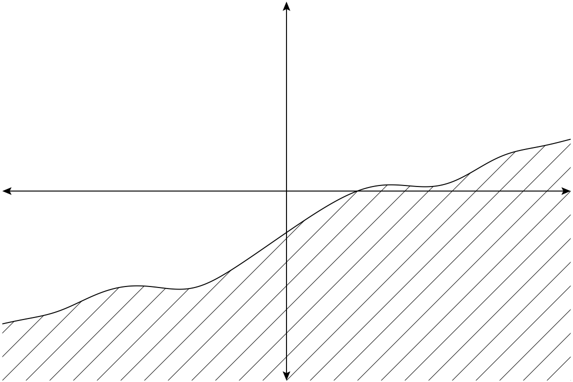

What's going on?

Why does the proof work?

$A_{P_{0.5}f}$:

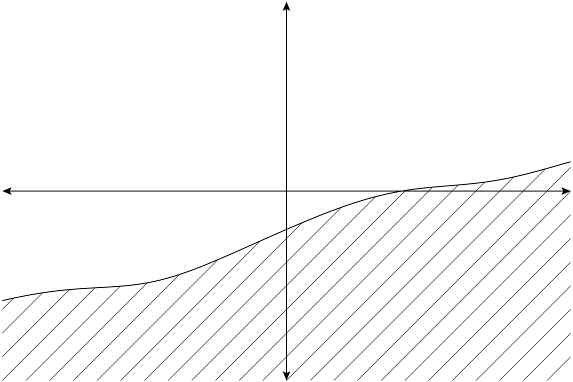

What's going on?

Why does the proof work?

$A_{P_1 f}$:

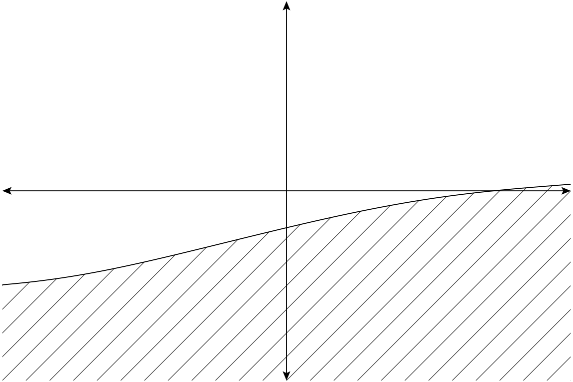

What's going on?

Why does the proof work?

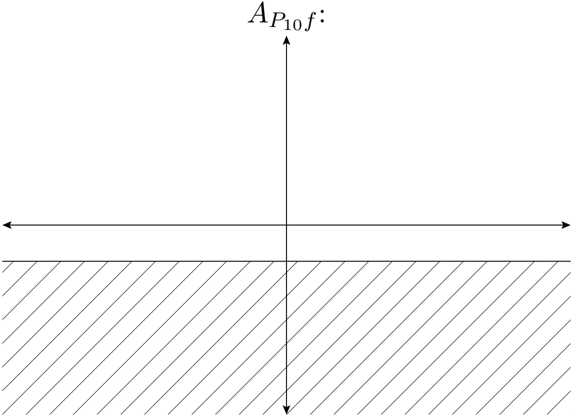

What's going on?

Why does the proof work?

$A_{P_{10} f}$:

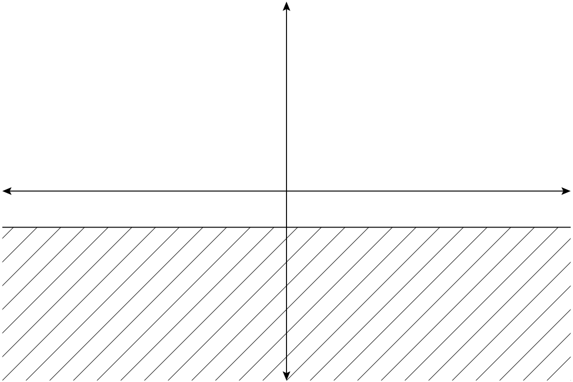

We showed that this transformation only increases the noise stability.

This idea has been used before: Bakry and Ledoux '96 used it to prove the Gaussian isoperimetric inequality.

Borell's theorem vs. Jensen's inequality

Theorem (Mossel, N. '12)

If $J : [0,1] \times [0,1] \to \mathbb{R}$ *satisfies* $\begin{pmatrix} \frac{\partial^2 J(x,y)}{\partial x^2} & \rho \frac{\partial^2 J(x,y)}{\partial x \partial y} \\ \rho \frac{\partial^2 J(x,y)}{\partial x \partial y} & \frac{\partial^2 J(x,y)}{\partial y^2} \end{pmatrix} \leq 0$ *then*

$$\mathbb{E} J(f(X), f(Y)) \leq J(\mathbb{E} f, \mathbb{E} f)$$

whenever X *and* Y *are* ρ*-correlated Gaussians.*

Borell's theorem vs. Jensen's inequality

Theorem (~~Mossel, N. '12~~ Jensen 1906)

If $J : [0,1] \times [0,1] \to \mathbb{R}$ *satisfies* $\begin{pmatrix} \frac{\partial^2 J(x,y)}{\partial x^2} & \cancel{\rho} \frac{\partial^2 J(x,y)}{\partial x \partial y} \\ \cancel{\rho} \frac{\partial^2 J(x,y)}{\partial x \partial y} & \frac{\partial^2 J(x,y)}{\partial y^2} \end{pmatrix} \leq 0$ *then*

$$\mathbb{E} J(f(X), f(Y)) \leq J(\mathbb{E} f, \mathbb{E} f)$$

whenever X *and* Y *are* ~~*ρ-correlated Gaussians*~~ *any random variables.*

Borell's theorem vs. Jensen's inequality

Theorem (Mossel, N. '12)

If $J : [0,1] \times [0,1] \to \mathbb{R}$ *satisfies* $\begin{pmatrix} \frac{\partial^2 J(x,y)}{\partial x^2} & \rho \frac{\partial^2 J(x,y)}{\partial x \partial y} \\ \rho \frac{\partial^2 J(x,y)}{\partial x \partial y} & \frac{\partial^2 J(x,y)}{\partial y^2} \end{pmatrix} \le 0$ *then*

$$\mathbb{E} J(f(X), f(Y)) \le J(\mathbb{E} f, \mathbb{E} f)$$

whenever X *and* Y *are* ρ*-correlated Gaussians.*

Borell's theorem vs. Jensen's inequality

Theorem (Mossel, N. '12)

If $J : [0,1] \times [0,1] \to \mathbb{R}$ *satisfies* $\begin{pmatrix} \frac{\partial^2 J(x,y)}{\partial x^2} & \rho \frac{\partial^2 J(x,y)}{\partial x \partial y} \\ \rho \frac{\partial^2 J(x,y)}{\partial x \partial y} & \frac{\partial^2 J(x,y)}{\partial y^2} \end{pmatrix} \le 0$ *then*

$$\mathbb{E} J(f(X), f(Y)) \le J(\mathbb{E} f, \mathbb{E} f)$$

whenever X *and* Y *are* ρ*-correlated Gaussians.*

Does the condition mean anything? Our J is the smallest one satisfying it.

This is what J looks like ($\rho = 0.1$)

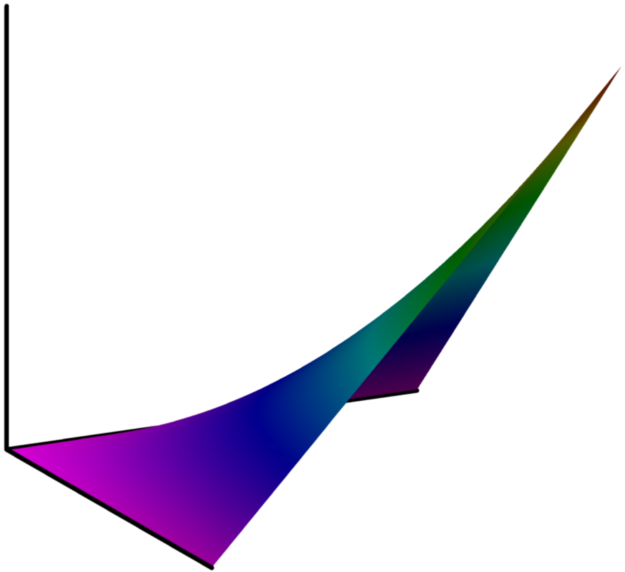

This is what J looks like ($\rho = 0.3$)

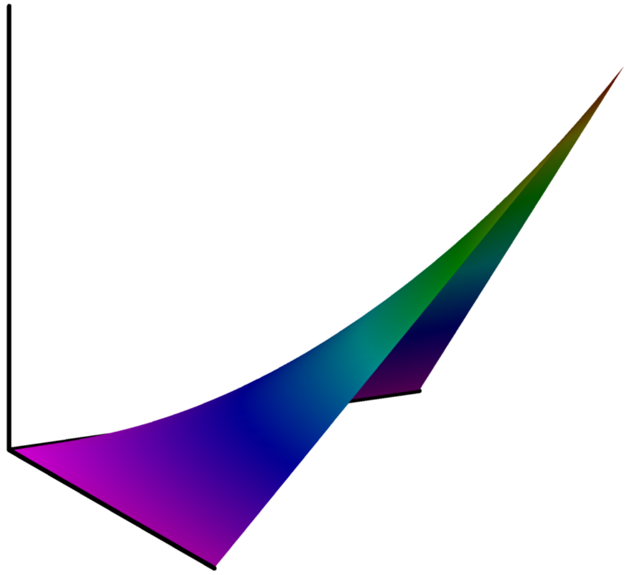

This is what J looks like ($\rho = 0.5$)

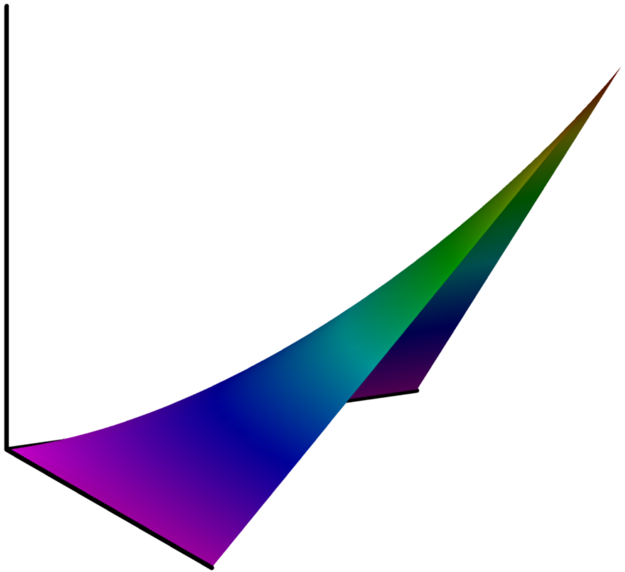

This is what J looks like ($\rho = 0.7$)

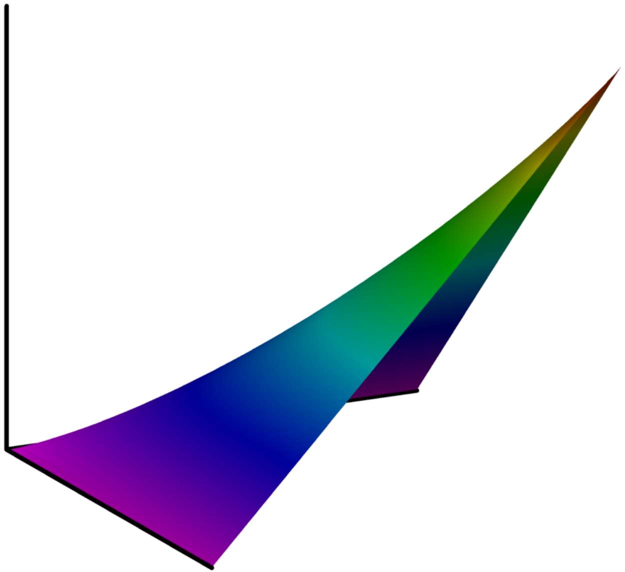

This is what J looks like ($\rho = 0.9$)

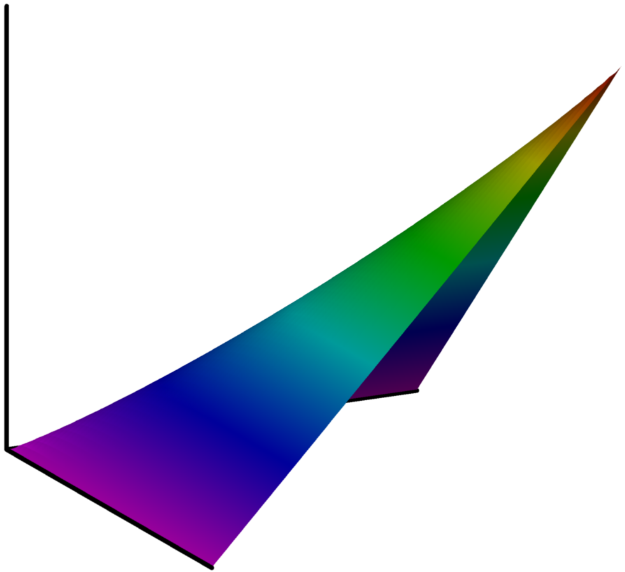

Proof: the equality case

Claim: if $f = 1_A$ and $\mathbb{E}J(f(X), f(Y)) = J(\mathbb{E}f, \mathbb{E}f)$ then A is a half-space.

Proof: the equality case

Claim: if $f = 1_A$ and $\mathbb{E}J(f(X), f(Y)) = J(\mathbb{E}f, \mathbb{E}f)$ then A is a half-space.

Recall that

$$\frac{d}{dt}\mathbb{E}J(P_t f(X), P_t f(Y)) = \frac{\rho}{2\pi\sqrt{1-\rho^2}}\mathbb{E}e^{-(v_t^2+w_t^2-2\rho v_t w_t)}|\nabla v_t - \nabla w_t|^2$$

where

$$v_t = v_t(X) = \Phi^{-1}(P_t f(X))$$
$$w_t = w_t(Y) = \Phi^{-1}(P_t f(Y))$$

Proof: the equality case

Claim: if $f = 1_A$ and $\mathbb{E}J(f(X), f(Y)) = J(\mathbb{E}f, \mathbb{E}f)$ then A is a half-space.
Recall that

$$\frac{d}{dt}\mathbb{E}J(P_tf(X), P_tf(Y)) = \frac{\rho}{2\pi\sqrt{1-\rho^2}}\mathbb{E}e^{-(v_t^2+w_t^2-2\rho v_t w_t)}|\nabla v_t - \nabla w_t|^2$$

where

$$v_t = v_t(X) = \Phi^{-1}(P_tf(X))$$
$$w_t = w_t(Y) = \Phi^{-1}(P_tf(Y))$$

$$\mathbb{E}J(f, f) = J(\mathbb{E}f, \mathbb{E}f) \iff \forall t \; \nabla v_t(X) = \nabla w_t(Y) = \text{constant}$$
$$\iff P_tf(x) = \Phi(a(t) \cdot x + b(t))$$

Proof: the equality case

Claim: if $f = 1_A$ and $\mathbb{E}J(f(X), f(Y)) = J(\mathbb{E}f, \mathbb{E}f)$ then A is a half-space.
Recall that

$$\frac{d}{dt}\mathbb{E}J(P_t f(X), P_t f(Y)) = \frac{\rho}{2\pi\sqrt{1-\rho^2}}\mathbb{E}e^{-(v_t^2+w_t^2-2\rho v_t w_t)}|\nabla v_t - \nabla w_t|^2$$

where

$$v_t = v_t(X) = \Phi^{-1}(P_t f(X))$$
$$w_t = w_t(Y) = \Phi^{-1}(P_t f(Y))$$

$$\begin{aligned}\mathbb{E}J(f,f) = J(\mathbb{E}f, \mathbb{E}f) &\iff \forall t \; \nabla v_t(X) = \nabla w_t(Y) = \text{constant} \\ &\iff P_t f(x) = \Phi(a(t)\cdot x + b(t)) \\ &\iff \text{if } f = 1_A \text{ then } A \text{ is a half-space.} \quad \square\end{aligned}$$

Proof: robustness

Claim: if $f = 1_A$ and $\mathbb{E}J(f(X), f(Y)) \geq J(\mathbb{E}f, \mathbb{E}f) - \delta$ then A is almost a half-space.

Proof: robustness

Claim: if $f = 1_A$ and $\mathbb{E} J(f(X), f(Y)) \geq J(\mathbb{E} f, \mathbb{E} f) - \delta$ then A is almost a half-space.

Recall that

$$
\begin{aligned}
&J(\mathbb{E} f, \mathbb{E} f) - \mathbb{E} J(f(X), f(Y)) \\
&\qquad = \frac{\rho}{2\pi\sqrt{1-\rho^2}} \int_0^\infty \mathbb{E} e^{-(v_t^2 + w_t^2 - 2\rho v_t w_t)} |\nabla v_t - \nabla w_t|^2 \, dt.
\end{aligned}
$$

Proof: robustness

Claim: if $f = 1_A$ and $\mathbb{E}J(f(X), f(Y)) \geq J(\mathbb{E}f, \mathbb{E}f) - \delta$ then A is almost a half-space.
Recall that

$$J(\mathbb{E}f, \mathbb{E}f) - \mathbb{E}J(f(X), f(Y)) \\ = \frac{\rho}{2\pi\sqrt{1-\rho^2}} \int_0^\infty \mathbb{E}e^{-(v_t^2 + w_t^2 - 2\rho v_t w_t)} |\nabla v_t - \nabla w_t|^2 \, dt.$$

Lemma

For any $t > 0$, $P_t f$ is close to a function of the form $\Phi(a \cdot x + b)$.

Proof: robustness

Claim: if $f = 1_A$ and $\mathbb{E}J(f(X), f(Y)) \geq J(\mathbb{E}f, \mathbb{E}f) - \delta$ then A is almost a half-space.
Recall that

$$J(\mathbb{E}f, \mathbb{E}f) - \mathbb{E}J(f(X), f(Y)) = \frac{\rho}{2\pi\sqrt{1-\rho^2}} \int_0^\infty \mathbb{E}e^{-(v_t^2 + w_t^2 - 2\rho v_t w_t)} |\nabla v_t - \nabla w_t|^2 \, dt.$$

Lemma

For any $t > 0$, $P_t f$ is close to a function of the form $\Phi(a \cdot x + b)$.

Lemma

If $P_t f$ is close to a function of the form $\Phi(a \cdot x + b)$ then f is also close to a function of the same form.