Reverse Engineering a PDE from an Image Inpainting Algorithm

Rob Hocking

University of Cambridge

Supervisor: Carola-Bibiane Schönlieb

December 6, 2015

Outline of Talk

- Problem of Interest

Outline of Talk

- Problem of Interest
  - Image inpainting

Outline of Talk

- Problem of Interest
  - Image inpainting
  - Class of Non-iterative, Geometric Methods

Outline of Talk

- Problem of Interest
  - Image inpainting
  - Class of Non-iterative, Geometric Methods
- Motivation

Outline of Talk

- Problem of Interest
  - Image inpainting
  - Class of Non-iterative, Geometric Methods
- Motivation
  - Target application: 3D Conversion

Outline of Talk

- Problem of Interest
  - Image inpainting
  - Class of Non-iterative, Geometric Methods
- Motivation
  - Target application: 3D Conversion
  - Example usage in Industry

Outline of Talk

- Problem of Interest
  - Image inpainting
  - Class of Non-iterative, Geometric Methods
- Motivation
  - Target application: 3D Conversion
  - Example usage in Industry
  - Not mathematically motivated

Outline of Talk

- Problem of Interest
  - Image inpainting
  - Class of Non-iterative, Geometric Methods
- Motivation
  - Target application: 3D Conversion
  - Example usage in Industry
  - Not mathematically motivated
- Maths

Outline of Talk

- Problem of Interest
  - Image inpainting
  - Class of Non-iterative, Geometric Methods
- Motivation
  - Target application: 3D Conversion
  - Example usage in Industry
  - Not mathematically motivated
- Maths
  - Can we use analysis to understand these methods?

Outline of Talk

- Problem of Interest
  - Image inpainting
  - Class of Non-iterative, Geometric Methods
- Motivation
  - Target application: 3D Conversion
  - Example usage in Industry
  - Not mathematically motivated
- Maths
  - Can we use analysis to understand these methods?
  - Yes. In fact analysis of (a) continuum limit already exists.

Outline of Talk

- Problem of Interest
  - Image inpainting
  - Class of Non-iterative, Geometric Methods
- Motivation
  - Target application: 3D Conversion
  - Example usage in Industry
  - Not mathematically motivated
- Maths
  - Can we use analysis to understand these methods?
  - Yes. In fact analysis of (a) continuum limit already exists.
  - Actually, there is a second possible continuum limit...

Outline of Talk

- Problem of Interest
  - Image inpainting
  - Class of Non-iterative, Geometric Methods
- Motivation
  - Target application: 3D Conversion
  - Example usage in Industry
  - Not mathematically motivated
- Maths
  - Can we use analysis to understand these methods?
  - Yes. In fact analysis of (a) continuum limit already exists.
  - Actually, there is a second possible continuum limit...
  - ...which is "closer" to the discrete solution.

Outline of Talk

- Problem of Interest
    - Image inpainting
    - Class of Non-iterative, Geometric Methods
- Motivation
    - Target application: 3D Conversion
    - Example usage in Industry
    - Not mathematically motivated
- Maths
    - Can we use analysis to understand these methods?
    - Yes. In fact analysis of (a) continuum limit already exists.
    - Actually, there is a second possible continuum limit...
    - ...which is “closer” to the discrete solution.
    - analyzing the *closer* limit facilitates the design of better algorithms.

Part I

The Problem

Image Inpainting

Filling holes / removing objects from images.

Image Inpainting

Filling holes / removing objects from images.

Figure

Image Inpainting

Filling holes / removing objects from images.

Figure

Approach of Interest

Fill in Shells using by taking weighted averages

Approach of Interest

Fill in Shells using by taking weighted averages

(Denote the hole to be filled by O, pixel *coordinates* by x, pixel *color* by $u(x)$, weights $w(x, y)$)

Approach of Interest

Fill in Shells using by taking weighted averages

(Denote the hole to be filled by O, pixel *coordinates* by x, pixel *color* by $u(x)$, weights $w(x, y)$)

while $O \neq \emptyset$

  for pixels $x \in \partial O$

    $u(x) \leftarrow \frac{\sum_{y \in B_\epsilon(x) \cap O^c} w(x,y)u(y)}{\sum_{y \in B_\epsilon(x) \cap O^c} w(x,y)}$

  end

  $O \leftarrow O \backslash \partial O$.

end

Example: Uniform Weights

Figure

Example: Uniform Weights

Figure

Example: Uniform Weights

Figure

Example: Uniform Weights

Figure

Example: Uniform Weights

Figure

Example: Uniform Weights

Figure

Example: Uniform Weights

Figure

Example: Uniform Weights

Figure

Immediate Observations

"Kinking "of extrapolated isophotes.

Immediate Observations

“Kinking ”of extrapolated isophotes.

Can be explained by proposing a continuum limit [Marz2007].

Immediate Observations

“Kinking ”of extrapolated isophotes.

Can be explained by proposing a continuum limit [Marz2007]. However,

- This limit does not account for other behaviour we will see later.

Immediate Observations

“Kinking ”of extrapolated isophotes.

Can be explained by proposing a continuum limit [Marz2007]. However,

- This limit does not account for other behaviour we will see later.
- It turns out another, *closer* limit also exists.

Immediate Observations

“Kinking ”of extrapolated isophotes.

Can be explained by proposing a continuum limit [Marz2007]. However,

- This limit does not account for other behaviour we will see later.
- It turns out another, *closer* limit also exists.
- It does.

Part II

Motivation

3D Conversion

- Given only the left eye view of the world, can we construct what the right eye sees?

3D Conversion

- Given only the left eye view of the world, can we construct what the right eye sees?
- Yes - in fact, in the movie industry entire films are routinely converted.

3D Conversion

- Given only the left eye view of the world, can we construct what the right eye sees?
- Yes - in fact, in the movie industry entire films are routinely converted.
- There are companies that exist solely for this purpose.

3D Conversion

- Given only the left eye view of the world, can we construct what the right eye sees?
- Yes - in fact, in the movie industry entire films are routinely converted.
- There are companies that exist solely for this purpose.
- One of them is Gener8, a company I worked for at the start of my PhD.

Figure

“Recent” Converted Films

(a)

(b)

Render the Scene from a new viewpoint

A depth map is used to decide how much to shift pixels left or right.

Render the Scene from a new viewpoint

A depth map is used to decide how much to shift pixels left or right.
"Gaps" are created at depth discontinuities.

Render the Scene from a new viewpoint

A depth map is used to decide how much to shift pixels left or right.
“Gaps” are created at depth discontinuities.

Figure

Incomplete Disocclusion

In classical inpainting, our goal is to fill the hole in an image *in its entirety*. But in the present case, we only have to fill *part* of the hole.

Figure

Classical Inpainting vs. Present Case

- Filling the entire hole is *allowed*, but it is potentially very wasteful.

Classical Inpainting vs. Present Case

- Filling the entire hole is *allowed*, but it is potentially very wasteful.
- Shell based approaches are natural, as they can be “stopped early”.

Classical Inpainting vs. Present Case

- Filling the entire hole is *allowed*, but it is potentially very wasteful.
- Shell based approaches are natural, as they can be "stopped early".
- GPU implementation is extremely fast.

ErodeFill and Compass Fill

- When I arrived at Gener8, the company was using two inpainting schemes, both special cases of the approach described earlier.

ErodeFill and Compass Fill

- When I arrived at Gener8, the company was using two inpainting schemes, both special cases of the approach described earlier.
- Both suffered from the "kinking" phenomena we saw earlier.

ErodeFill and Compass Fill

- When I arrived at Gener8, the company was using two inpainting schemes, both special cases of the approach described earlier.
- Both suffered from the "kinking" phenomena we saw earlier.
- My Job: "Can you make that kinking go away?"

Coherence Transport

Idea: Instead of uniform weights, assign higher weights to pixels on edges.

Coherence Transport

Idea: Instead of uniform weights, assign higher weights to pixels on edges.

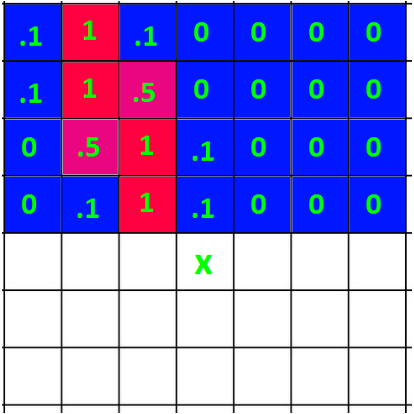

Figure

Coherence Transport

For each pixel x due to be filled, estimate local edge magnitude $\mu(x) \geq 0$ and direction $\mathbf{g}(x) \in S^1$.

Coherence Transport

For each pixel x due to be filled, estimate local edge magnitude $\mu(x) \geq 0$ and direction $\mathbf{g}(x) \in S^1$.
Adapt weights accordingly:

$$w(x, y) = \frac{\exp\left(-\frac{\mu(x)^2}{2\epsilon^2}\left(\mathbf{g}^\perp(x) \cdot (y - x)\right)^2\right)}{\|x - y\|}.$$

Coherence Transport

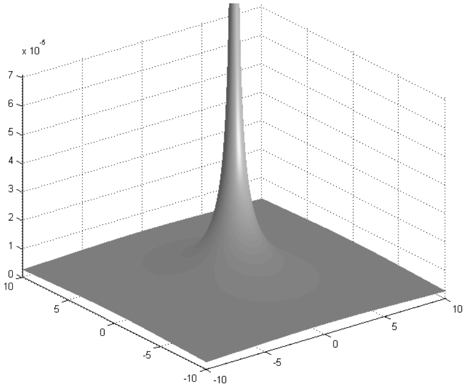

Figure: $\mu = 0$

Coherence Transport

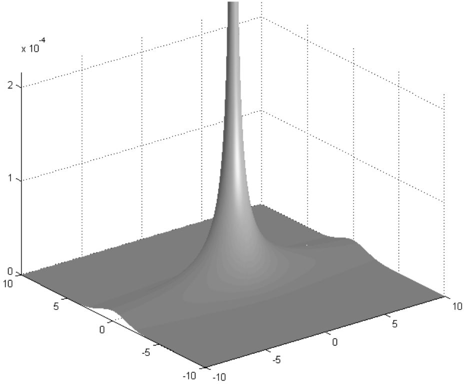

Figure: $\mu = 1$

Coherence Transport

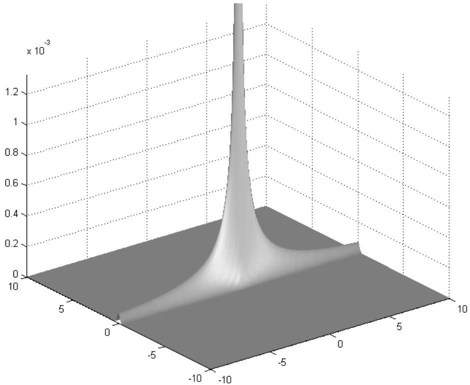

Figure: $\mu = 10$

Is the kinking fixed?

Test the method on a toy problem, feeding in the correct $\mathbf{g}$ by hand and set $\mu \gg 1$.

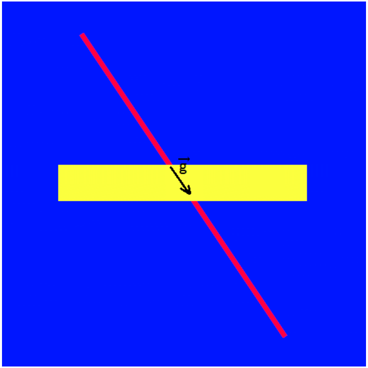

Figure

Actually, no

$\epsilon = 3\text{px},\ \mu = 50.$

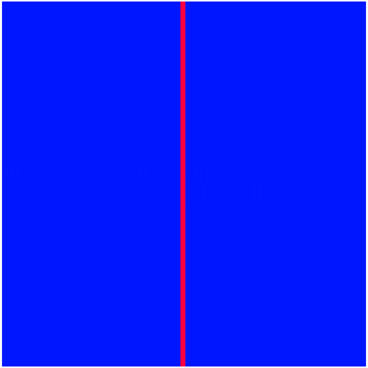

Figure

Actually, no

$\epsilon = 3\text{px}$, $\mu = 50$.

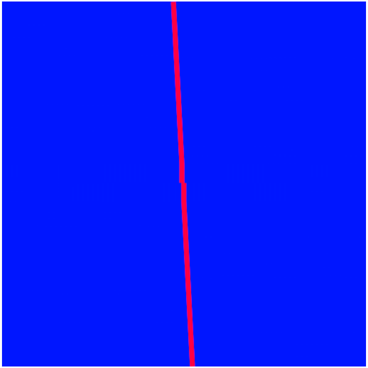

Figure

Actually, no

$\epsilon = 3\text{px}$, $\mu = 50$.

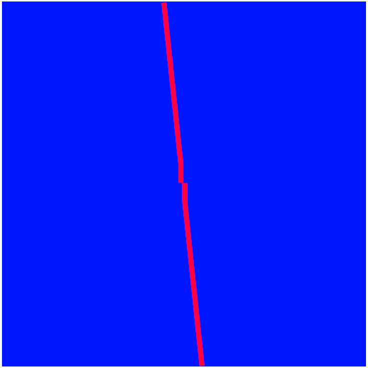

Figure

Actually, no

$\epsilon = 3\text{px}$, $\mu = 50$.

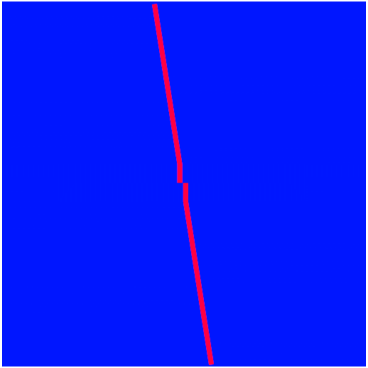

Figure

Actually, no

$\epsilon = 3\text{px}$, $\mu = 50$.

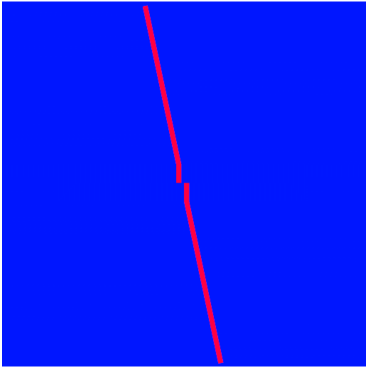

Figure

Actually, no

$\epsilon = 3\text{px}$, $\mu = 50$.

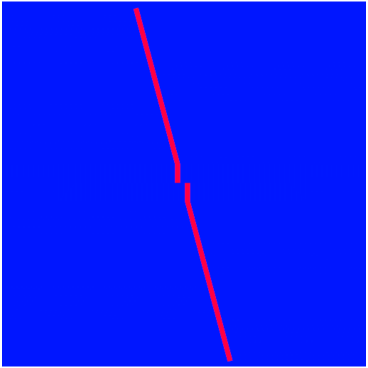

Figure

Actually, no

$\epsilon = 3\text{px}$, $\mu = 50$.

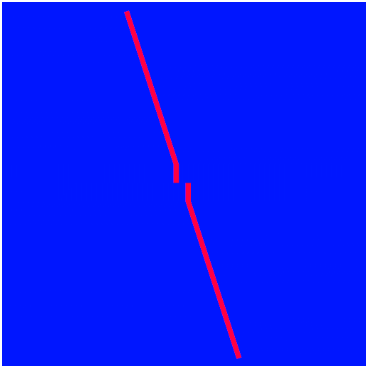

Figure

Actually, no

$\epsilon = 3\text{px}$, $\mu = 50$.

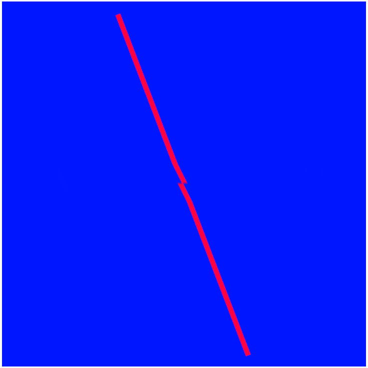

Figure

Actually, no

$\epsilon = 3\text{px}$, $\mu = 50$.

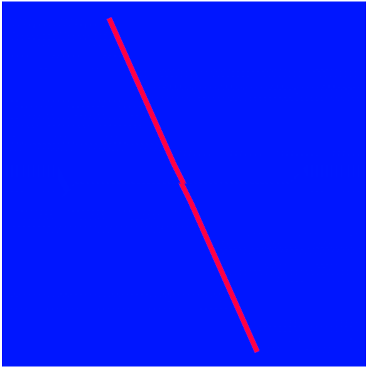

Figure

Actually, no

$\epsilon = 3\text{px}$, $\mu = 50$.

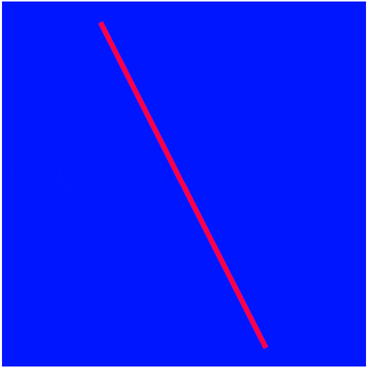

Figure

Actually, no

$\epsilon = 3\text{px},\ \mu = 50.$

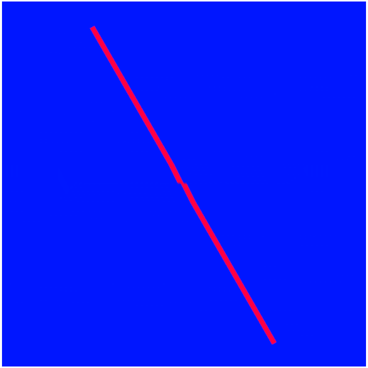

Figure

Intuitive Fix

- weights are highest for pixels on the line passing through $\mathbf{x}$ parallel to $\mathbf{g}$.

Intuitive Fix

- weights are highest for pixels on the line passing through $\mathbf{x}$ parallel to $\mathbf{g}$.
- Problems come up when this line "misses" pixel centers in $B_{\epsilon,h}(\mathbf{x})$.

Intuitive Fix

- weights are highest for pixels on the line passing through $\mathbf{x}$ parallel to $\mathbf{g}$.
- Problems come up when this line "misses" pixel centers in $B_{\epsilon,h}(\mathbf{x})$.

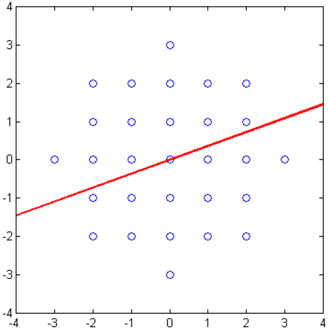

Figure

Intuitive Fix

Sum over a rotated ball of “ghost pixels” instead.

Intuitive Fix

Sum over a rotated ball of "ghost pixels" instead.

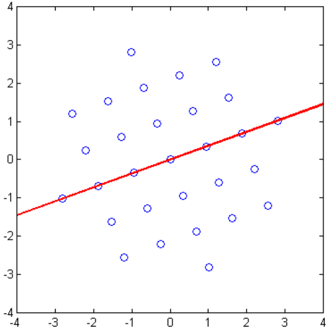

Figure

Define "ghost pixels" using bilinear interpolation.

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

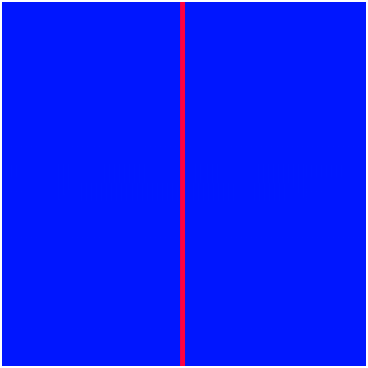

Figure

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

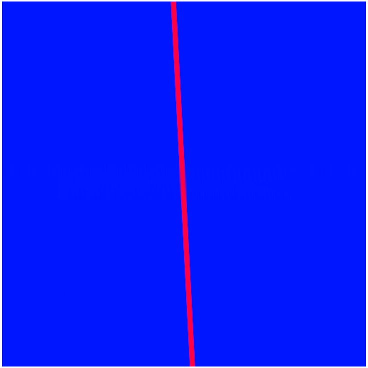

Figure

Lines do not bend

$\epsilon = 3\text{px},\ \mu = 50.$

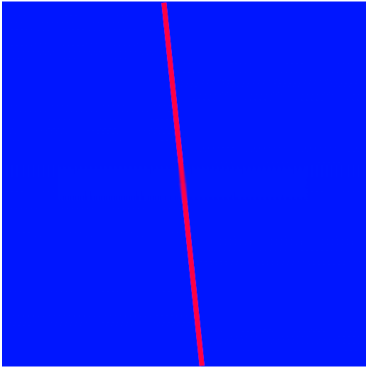

Figure

Lines do not bend

$\epsilon = 3\text{px},\ \mu = 50.$

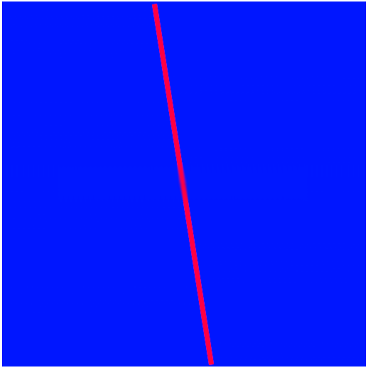

Figure

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

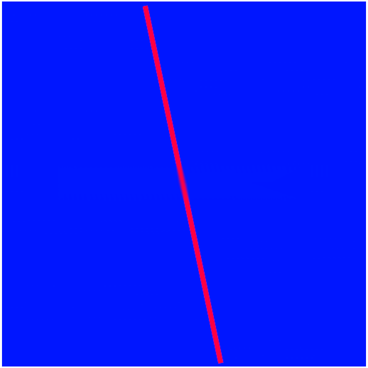

Figure

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

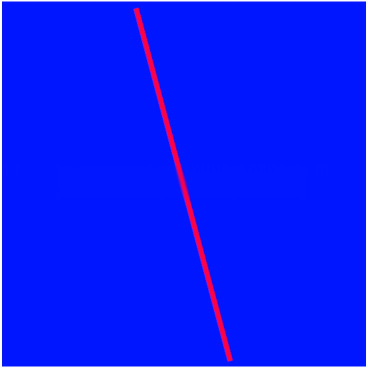

Figure

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

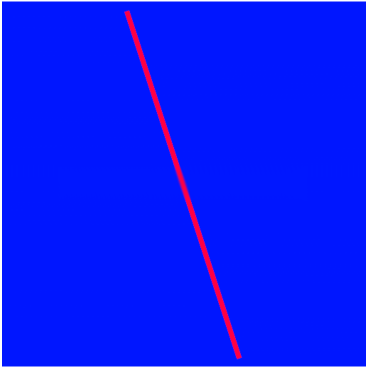

Figure

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

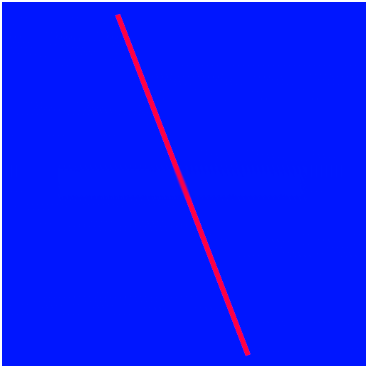

Figure

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

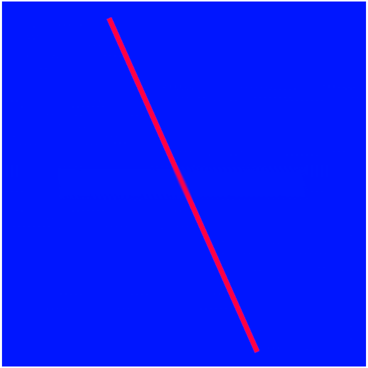

Figure

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

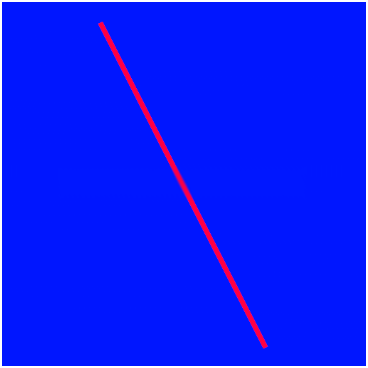

Figure

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

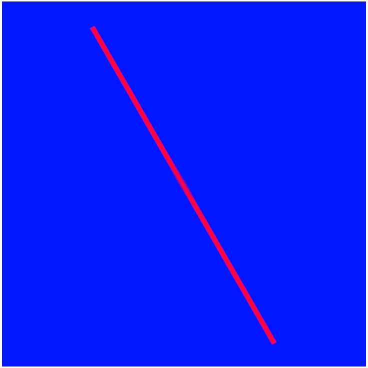

Figure

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

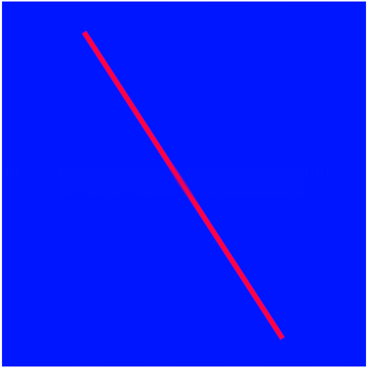

Figure

Lines do not bend

$\epsilon = 3\text{px}$, $\mu = 50$.

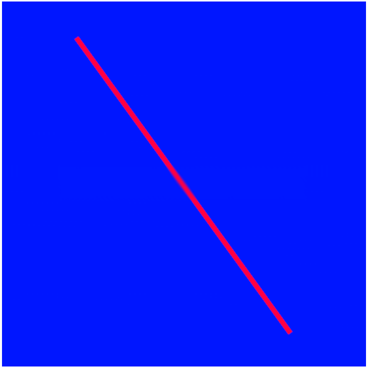

Figure

Part III

Maths

What's going on?

- What caused the kinking observed earlier?

What's going on?

- What caused the kinking observed earlier?
- Why did ghost pixels fix it?

Assumptions

From now on, assuming inpainting domain $O = (0, 1]^2$, with periodic boundary conditions at $x = 0$ and $x = 1$.

Boundary data is on the strip $(0, 1] \times (-\delta, 0]$.

Assumptions

From now on, assuming inpainting domain $O = (0,1]^2$, with periodic boundary conditions at $x = 0$ and $x = 1$.

Boundary data is on the strip $(0,1] \times (-\delta, 0]$.

Also some restrictions on the weights.

Existing Theory

März justifies Coherence Transport by arguing that in the *double limit* $h \to 0$ and then $\epsilon \to 0$, the method behaves like the transport PDE

$$\mathbf{g}_\mu^* \cdot \nabla u_{\text{märz}} = 0$$

Existing Theory

März justifies Coherence Transport by arguing that in the *double limit* $h \to 0$ and then $\epsilon \to 0$, the method behaves like the transport PDE

$$\mathbf{g}_\mu^* \cdot \nabla u_{\text{märz}} = 0$$

where,

- $\mathbf{g}_\mu^* := \int_{y \in B_1^-} w_{\mu,1}(0,y) y dy \longrightarrow \mathbf{g}$ as $\mu \to \infty$,

Existing Theory

März justifies Coherence Transport by arguing that in the *double limit* $h \to 0$ and then $\epsilon \to 0$, the method behaves like the transport PDE

$$\mathbf{g}_\mu^* \cdot \nabla u_{\text{märz}} = 0$$

where,

- $\mathbf{g}_\mu^* := \int_{y \in B_1^-} w_{\mu,1}(0,y) y dy \longrightarrow \mathbf{g}$ as $\mu \to \infty$,
- $B_1^- := \{(y_1, y_2) \in \mathbb{R}^2 : y_1^2 + y_2^2 \leq 1$ and $y_2 \leq 0\}$,

Existing Theory

Call the continuum solution $u_{\text{märz}}$, discrete solution u_h.

Existing Theory

Call the continuum solution $u_{\text{märz}}$, discrete solution u_h.
We know that the double limit

$$u_{\text{märz}} = \lim_{\epsilon \to 0} \left[\lim_{h \to 0} u_h \right]$$

does what we want.

Existing Theory

Call the continuum solution $u_{\text{märz}}$, discrete solution u_h.
We know that the double limit

$$u_{\text{märz}} = \lim_{\epsilon \to 0} \left[\lim_{h \to 0} u_h \right]$$

does what we want.

But does $h \approx 0$ and $\epsilon \approx 0$ mean $u_h \approx u_{\text{märz}}$?

Existing Theory

Call the continuum solution $u_{\text{märz}}$, discrete solution u_h.
We know that the double limit

$$u_{\text{märz}} = \lim_{\epsilon \to 0} \left[\lim_{h \to 0} u_h \right]$$

does what we want.

But does $h \approx 0$ and $\epsilon \approx 0$ mean $u_h \approx u_{\text{märz}}$?

Actually, no.

Closer Look

First, let's take $h \to 0$...

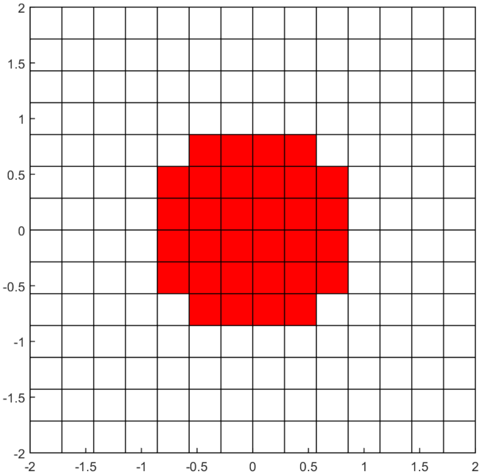

Figure

Closer Look

First, let's take $h \to 0$...

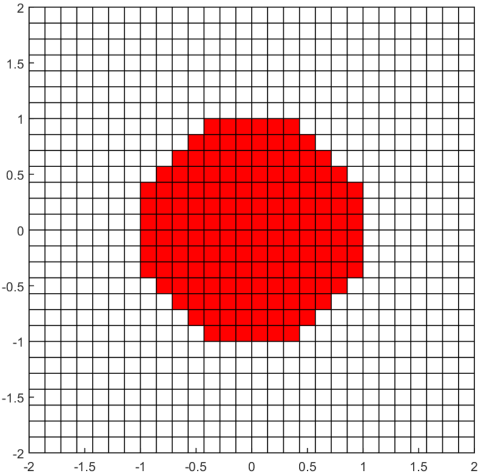

Figure

Closer Look

First, let's take $h \to 0$...

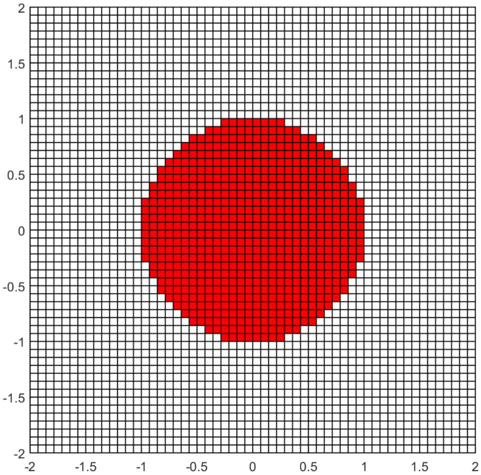

Figure

Closer Look

First, let's take $h \to 0$...

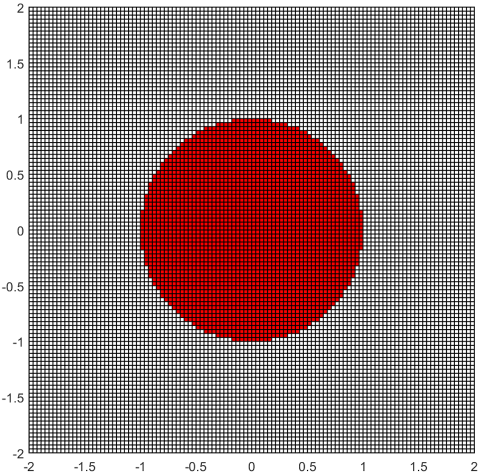

Figure

Closer Look

First, let's take $h \to 0$...

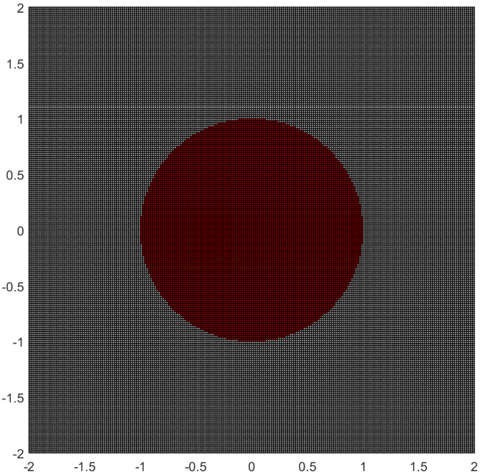

Figure

Closer Look

First, let's take $h \to 0$...

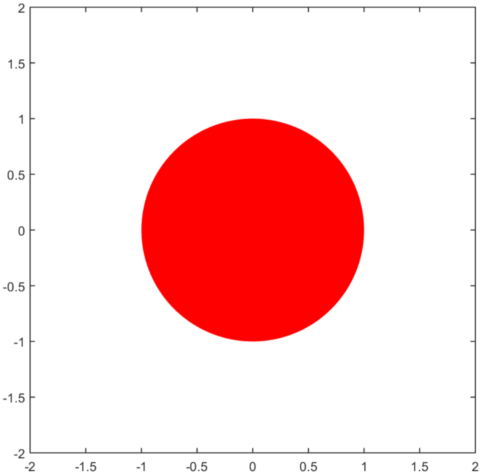

Figure

Closer Look

Next, take $\epsilon \to 0$...

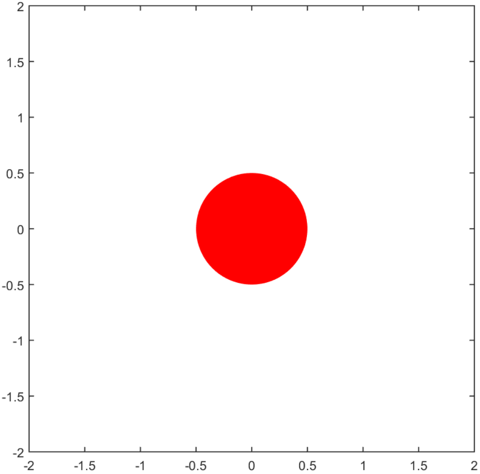

Figure

Closer Look

Next, take $\epsilon \to 0$...

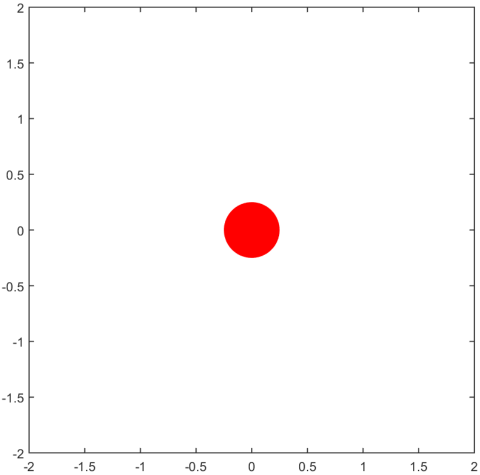

Figure

Closer Look

Next, take $\epsilon \to 0$...

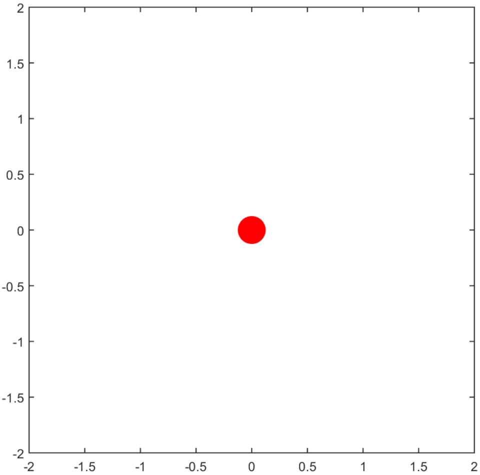

Figure

Closer Look

Next, take $\epsilon \to 0$...

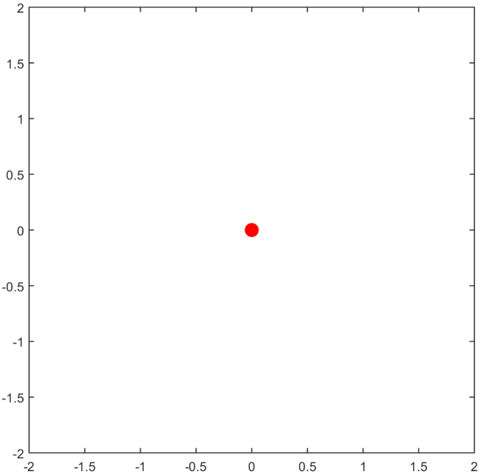

Figure

Closer Look

Next, take $\epsilon \to 0$...

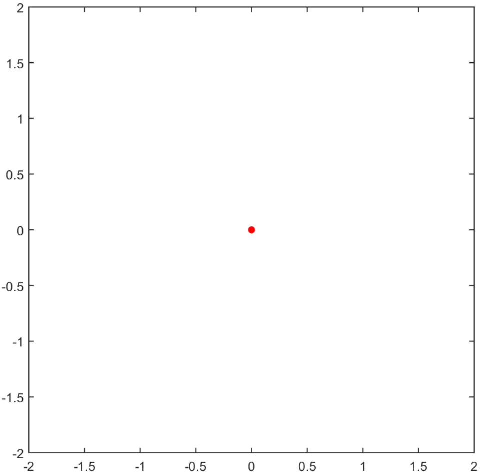

Figure

Closer Look

Next, take $\epsilon \to 0$...

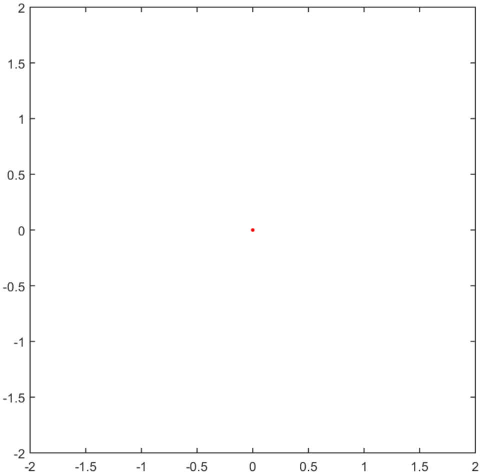

Figure

Closer Look

Next, take $\epsilon \to 0$...

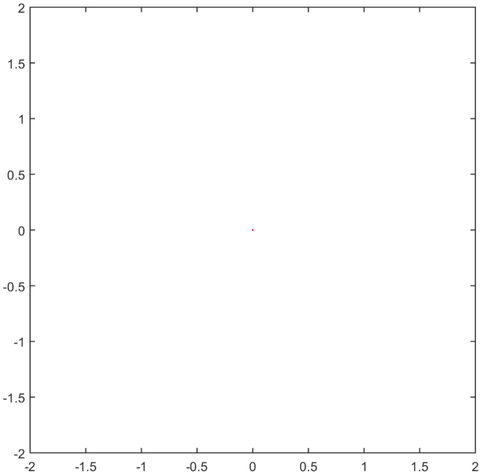

Figure

Closer Look

Next, take $\epsilon \to 0$...

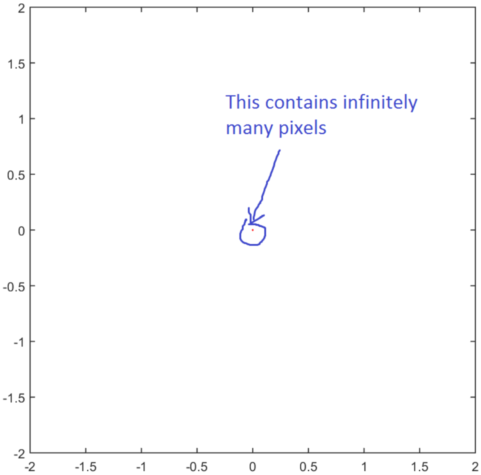

Figure

Closer Look

First, let's make h "pretty small".

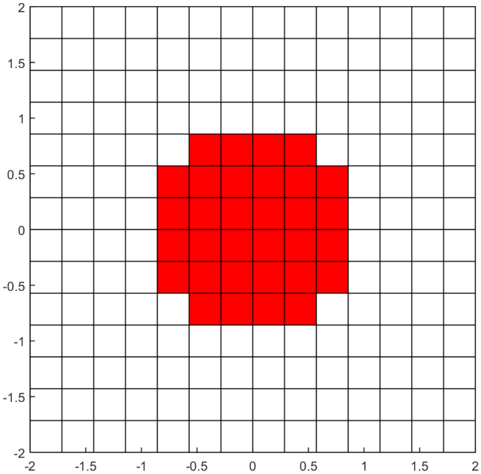

Figure

Closer Look

First, let's make h "pretty small".

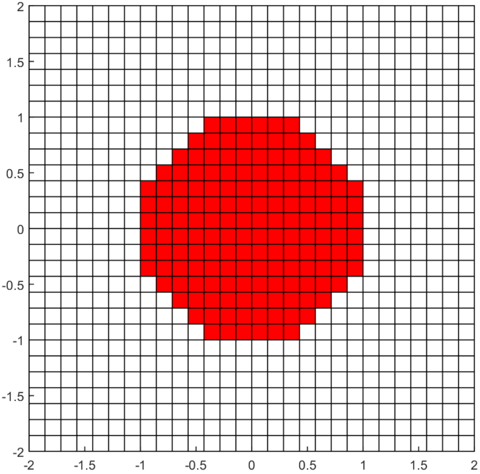

Figure

Closer Look

First, let's make h "pretty small".

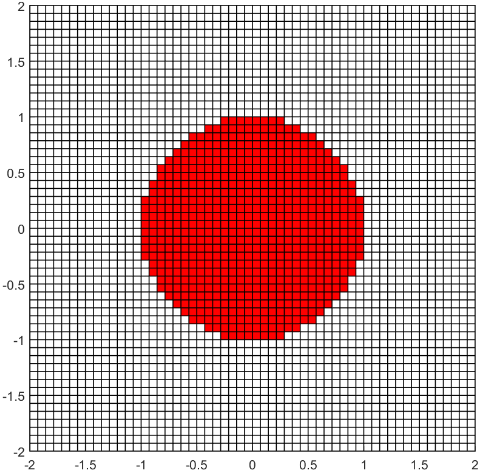

Figure

Closer Look

First, let's make h "pretty small".

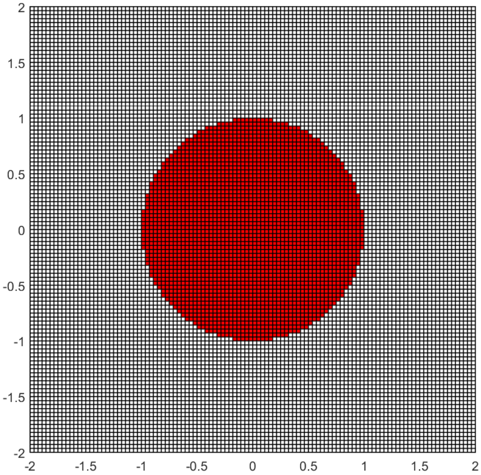

Figure

Closer Look

Now let's make ϵ "pretty small" as well.

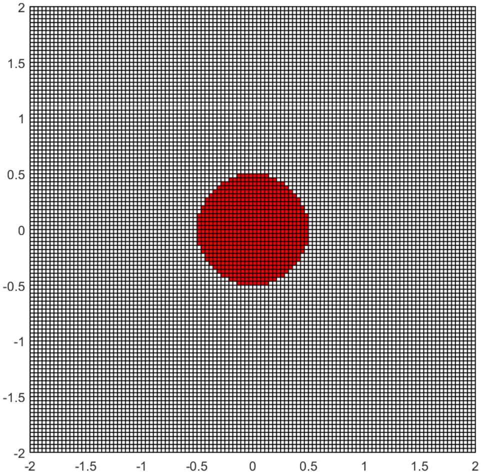

Figure

Closer Look

Now let's make ϵ "pretty small" as well.

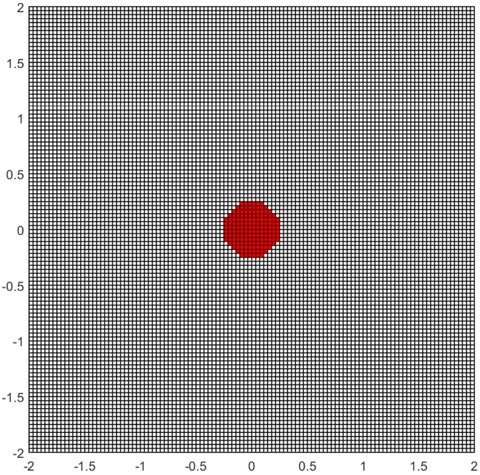

Figure

Closer Look

Now let's make ϵ "pretty small" as well.

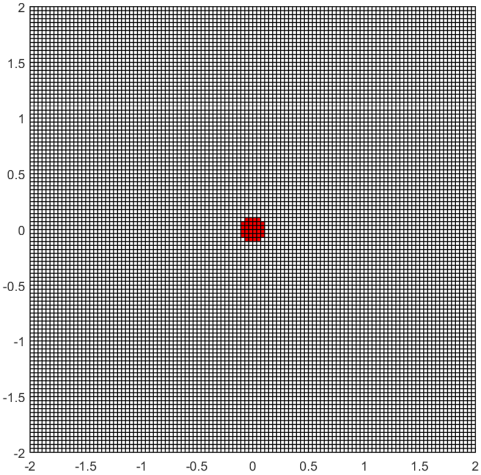

Figure

Closer Look

Now let's make ϵ "pretty small" as well.

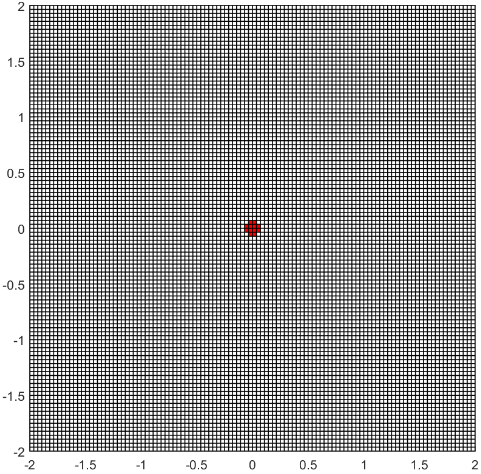

Figure

Closer Look

Now let's make ϵ "pretty small" as well.

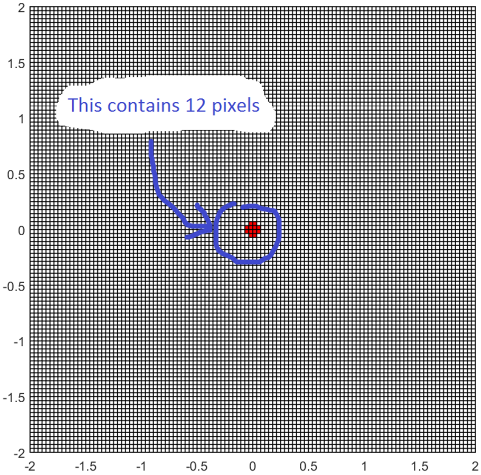

Figure

Closer Look

Define $r = \epsilon/h$ (the radius of the ball measured in pixels).

Closer Look

Define $r = \epsilon/h$ (the radius of the ball measured in pixels).

For $u_h \approx u_{\text{märz}}$, we need not only $h \approx 0$ and $\epsilon \approx 0$, but also $r \gg 1$.

Closer Look

Define $r = \epsilon/h$ (the radius of the ball measured in pixels).

For $u_h \approx u_{\text{märz}}$, we need not only $h \approx 0$ and $\epsilon \approx 0$, but also $r \gg 1$.

But in practice, one fixes $r = 3$px to $r = 5$px.

Closer Look

Define $r = \epsilon/h$ (the radius of the ball measured in pixels).

For $u_h \approx u_{\text{märz}}$, we need not only $h \approx 0$ and $\epsilon \approx 0$, but also $r \gg 1$.

But in practice, one fixes $r = 3$px to $r = 5$px.

Makes more sense to study the limit $h \to 0$ with r *fixed*.

New Limit

$h \to 0$, $r = \epsilon/h$ constant.

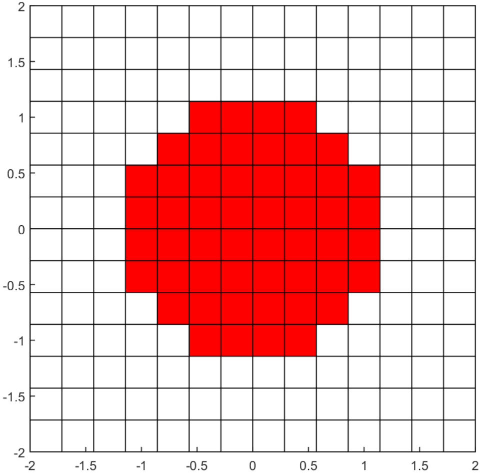

Figure

New Limit

$h \to 0$, $r = \epsilon/h$ constant.

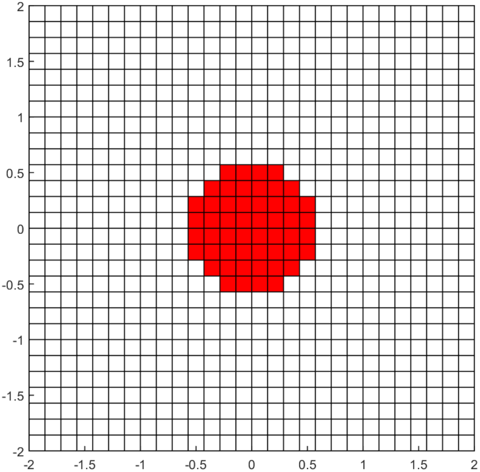

Figure

New Limit

$h \to 0$, $r = \epsilon/h$ constant.

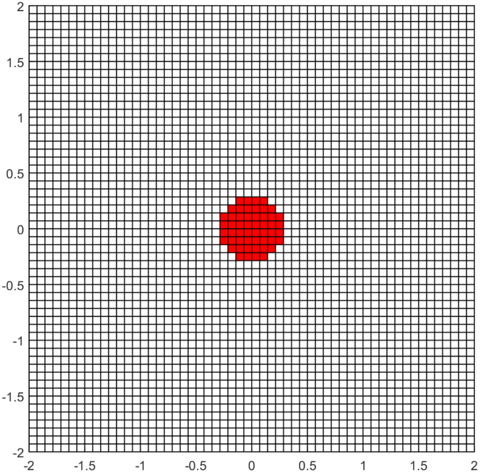

Figure

New Limit

$h \to 0$, $r = \epsilon/h$ constant.

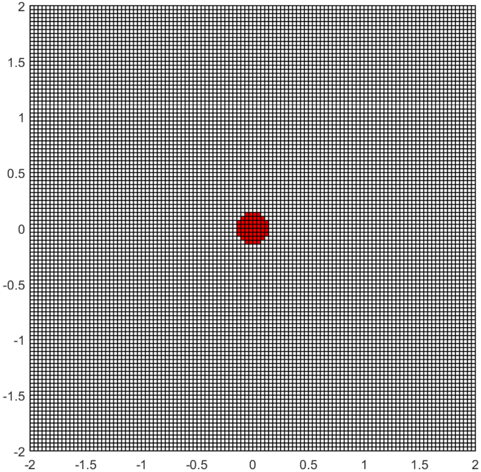

Figure

$h \to 0$, $r = \epsilon/h$ constant.

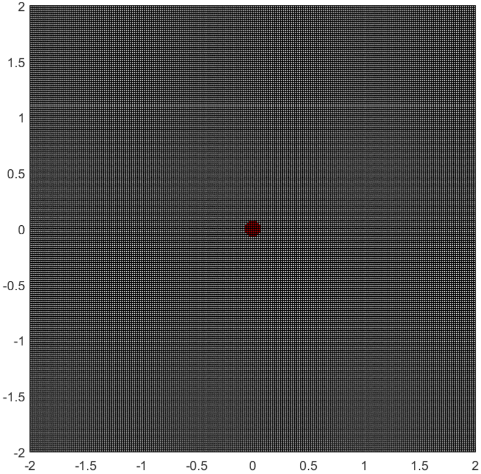

Figure

Which Limit is Closer?

It is possible to prove convergence of u_h to *either* limit:

Which Limit is Closer?

It is possible to prove convergence of u_h to *either* limit:

$$\|u_h - u_r\|_p \quad \to \quad 0 \qquad \text{as } h \to 0 \text{ with } r \text{ fixed.}$$

Which Limit is Closer?

It is possible to prove convergence of u_h to *either* limit:

$$\begin{aligned} \|u_h - u_r\|_p &\to 0 \quad \text{as } h \to 0 \text{ with } r \text{ fixed.} \\ \|u_h - u_{\text{märz}}\|_p &\to 0 \quad \text{as } h \to 0 \text{ and } r \to \infty \text{ but } r^2 h \to 0. \end{aligned}$$

Which Limit is Closer?

It is possible to prove convergence of u_h to *either* limit:

$$\begin{aligned} \|u_h - u_r\|_p &\to 0 \quad \text{as } h \to 0 \text{ with } r \text{ fixed.} \\ \|u_h - u_{\text{märz}}\|_p &\to 0 \quad \text{as } h \to 0 \text{ and } r \to \infty \text{ but } r^2 h \to 0. \end{aligned}$$

- true for $p < \infty$ for boundary data with finitely many jump discontinuities.

Which Limit is Closer?

It is possible to prove convergence of u_h to *either* limit:

$$\begin{aligned} \|u_h - u_r\|_p &\to 0 \quad \text{as } h \to 0 \text{ with } r \text{ fixed.} \\ \|u_h - u_{\text{märz}}\|_p &\to 0 \quad \text{as } h \to 0 \text{ and } r \to \infty \text{ but } r^2 h \to 0. \end{aligned}$$

- true for $p < \infty$ for boundary data with finitely many jump discontinuities.
- true for all $1 \leq p \leq \infty$ if there are no jumps.

Which Limit is Closer?

But for *fixed* h, one finds:

$$\|u_h - u_r\|_p \lessapprox C_r \|u_h - u_{\text{märz}}\|_p^2.$$

New Limit

Under this limit we still get a transport equation

$$\mathbf{g}^*_{\mu,r} \cdot \nabla u_r = 0,$$

New Limit

Under this limit we still get a transport equation

$$\mathbf{g}^*_{\mu,r} \cdot \nabla u_r = 0,$$

but the transport direction is different:

$$\mathbf{g}^*_{\mu,r} := \sum_{\mathbf{j} \in b_r^-} w_{\mu,r}(0, \mathbf{j})\mathbf{j}.$$

where

$$b_r^- := \{(i,j) \in \mathbb{Z}^2 : i^2 + j^2 \leq r^2 \text{ and } j < 0\}$$

New Limit

Under this limit we still get a transport equation

$$\mathbf{g}^*_{\mu,r} \cdot \nabla u_r = 0,$$

but the transport direction is different:

$$\mathbf{g}^*_{\mu,r} := \sum_{\mathbf{j} \in b_r^-} w_{\mu,r}(0,\mathbf{j})\mathbf{j}.$$

where

$$b_r^- := \{(i,j) \in \mathbb{Z}^2 : i^2 + j^2 \le r^2 \text{ and } j < 0\}$$

vs

$$\mathbf{g}^*_{\mu} := \int_{y \in B_1^-} w_{\mu,1}(0,y) y dy$$

where

$$B_1^- := \{(y_1,y_2) \in \mathbb{R}^2 : y_1^2 + y_2^2 \le 1 \text{ and } y_2 \le 0\}.$$

Explanation of Clamping

In this case we have

$$\mathbf{g}^*_{r,\mu} := \sum_{\mathbf{j}\in b_r^-} e^{-\frac{\mu^2}{2r^2}(\mathbf{j}\cdot\mathbf{g}^\perp)^2} \frac{\mathbf{j}}{\|\mathbf{j}\|},$$

Explanation of Clamping

In this case we have

$$\mathbf{g}^*_{r,\mu} := \sum_{\mathbf{j} \in b_r^-} e^{-\frac{\mu^2}{2r^2}(\mathbf{j}\cdot\mathbf{g}^\perp)^2} \frac{\mathbf{j}}{\|\mathbf{j}\|},$$

- Suppose $\mathbf{j}^*$ is the *unique* minimizer of $|\mathbf{j} \cdot \mathbf{g}^\perp|$ for $\mathbf{j} \in b_r^-$

Explanation of Clamping

In this case we have

$$\mathbf{g}^*_{r,\mu} := \sum_{\mathbf{j}\in b_r^-} e^{-\frac{\mu^2}{2r^2}(\mathbf{j}\cdot\mathbf{g}^\perp)^2} \frac{\mathbf{j}}{\|\mathbf{j}\|},$$

- Suppose $\mathbf{j}^*$ is the *unique* minimizer of $|\mathbf{j}\cdot\mathbf{g}^\perp|$ for $\mathbf{j}\in b_r^-$

Then, rescaling by $e^{\frac{\mu^2}{2r^2}(\mathbf{j}^*\cdot\mathbf{g}^\perp)^2}$ we have

$$\mathbf{g}^*_{r,\mu} = \frac{\mathbf{j}^*}{\|\mathbf{j}^*\|} + \sum_{\mathbf{j}\in b_r^-\setminus\{\mathbf{j}^*\}} e^{-\frac{\mu^2}{2r^2}\{(\mathbf{j}\cdot\mathbf{g}^\perp)^2-(\mathbf{j}^*\cdot\mathbf{g}^\perp)^2\}} \frac{\mathbf{j}}{\|\mathbf{j}\|}$$

Explanation of Clamping

In this case we have

$$\mathbf{g}^*_{r,\mu} := \sum_{\mathbf{j} \in b_r^-} e^{-\frac{\mu^2}{2r^2}(\mathbf{j}\cdot\mathbf{g}^\perp)^2} \frac{\mathbf{j}}{\|\mathbf{j}\|},$$

- Suppose $\mathbf{j}^*$ is the *unique* minimizer of $|\mathbf{j} \cdot \mathbf{g}^\perp|$ for $\mathbf{j} \in b_r^-$

Then, rescaling by $e^{\frac{\mu^2}{2r^2}(\mathbf{j}^*\cdot\mathbf{g}^\perp)^2}$ we have

$$\begin{aligned} \mathbf{g}^*_{r,\mu} &= \frac{\mathbf{j}^*}{\|\mathbf{j}^*\|} + \sum_{\mathbf{j} \in b_r^- \setminus \{\mathbf{j}^*\}} e^{-\frac{\mu^2}{2r^2}\left\{(\mathbf{j}\cdot\mathbf{g}^\perp)^2 - (\mathbf{j}^*\cdot\mathbf{g}^\perp)^2\right\}} \frac{\mathbf{j}}{\|\mathbf{j}\|} \\ &\to \frac{\mathbf{j}^*}{\|\mathbf{j}^*\|} \qquad \text{as } \mu \to \infty. \end{aligned}$$

Comparison with real Life

Define $\theta = \angle \mathbf{g}$, $\theta_r^* = \angle \mathbf{g}_r^*$, consider $\theta_r^* = \Theta(\theta)$.

Comparison with real Life

Theoretical Curve $(r = 3)$

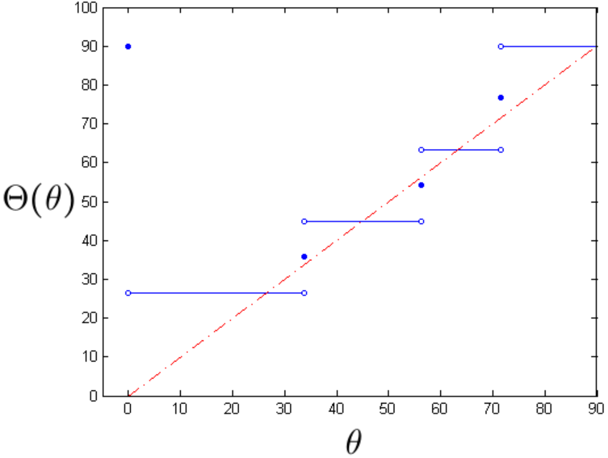

Figure

Comparison with real Life

Real Curve ($r = 3,\ \mu = 40$)

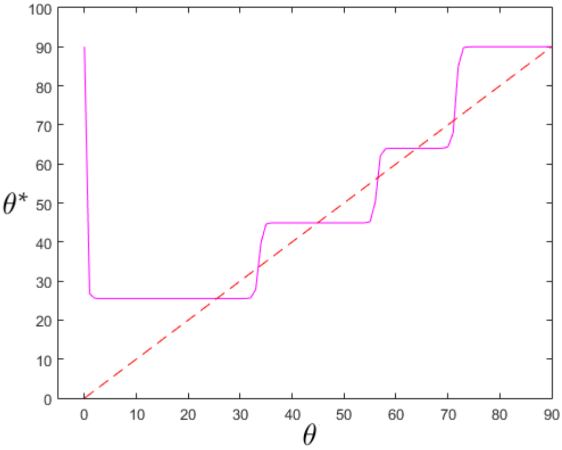

Figure

Comparison with real Life

Theoretical Curve $(r = 5)$

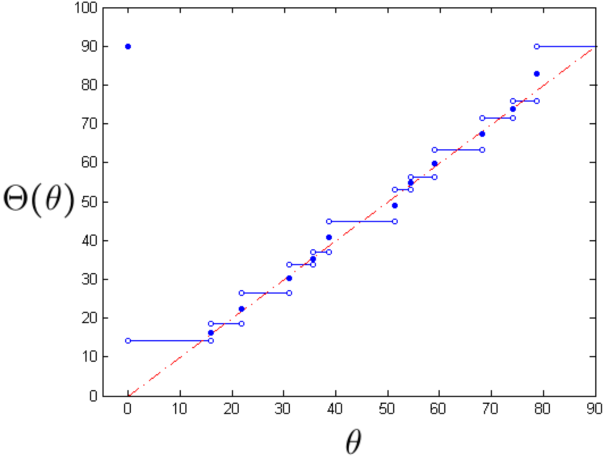

Figure

Comparison with real Life

Real Curve ($r = 5$, $\mu = 40$)

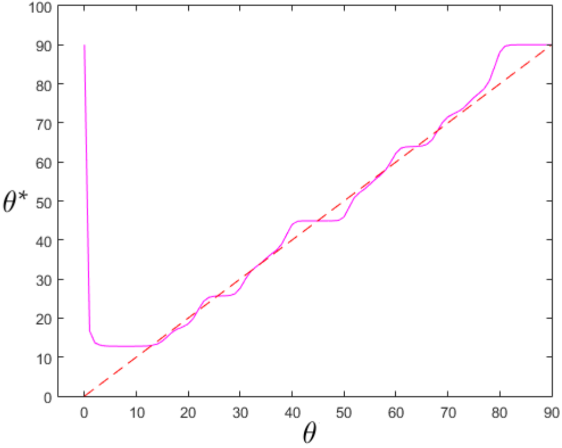

Figure

Explaining our Fix

- This explains earlier kinking.

Explaining our Fix

- This explains earlier kinking.
- But why were we able to make it go away?

Explaining our Fix

Earlier, we proposed a continuum limit u_r with transport direction

$$\mathbf{g}_r^* = \sum_{\mathbf{j} \in b_r^-} w(0, \mathbf{j})\mathbf{j}.$$

Explaining our Fix

Earlier, we proposed a continuum limit u_r with transport direction

$$\mathbf{g}_r^* = \sum_{\mathbf{j} \in b_r^-} w(0, \mathbf{j})\mathbf{j}.$$

But actually this assumed that we use no ghost pixels.

Back to the Fix

Now assume we sum over a rotated ball $\tilde{B}_{\epsilon,h}(\mathbf{x})$.

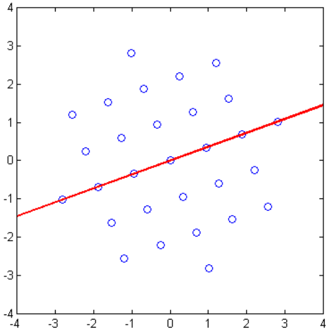

Figure

Back to the Fix

We get a remarkably similar formula:

$$\mathbf{g}_r^* = \sum_{\mathbf{j} \in \tilde{b}_r^-} w(0, \mathbf{j})\mathbf{j} \qquad \text{vs.} \qquad \mathbf{g}_r^* = \sum_{\mathbf{j} \in b_r^-} w(0, \mathbf{j})\mathbf{j}$$

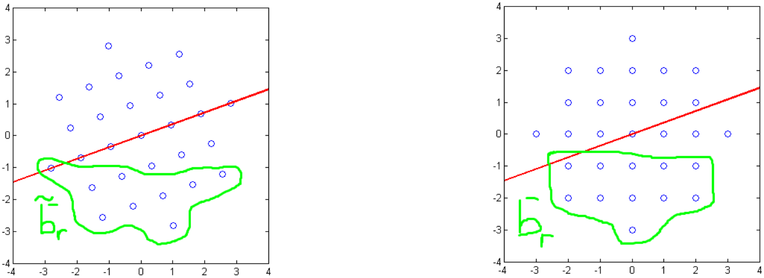

Figure

Back to the Fix

- This simple formula is a consequence of our choice to define ghost pixels via bilinear interpolation.

Back to the Fix

- This simple formula is a consequence of our choice to define ghost pixels via bilinear interpolation.
- It is *not* true for generic interpolants.

Back to the Fix

- This simple formula is a consequence of our choice to define ghost pixels via bilinear interpolation.
- It is *not* true for generic interpolants.
- Key property is that the bilinear interpolant of a linear function is the function.

Back to the Fix

- This simple formula is a consequence of our choice to define ghost pixels via bilinear interpolation.
- It is *not* true for generic interpolants.
- Key property is that the bilinear interpolant of a linear function is the function.
- Now let's see why everything is fixed.

Back to the Fix

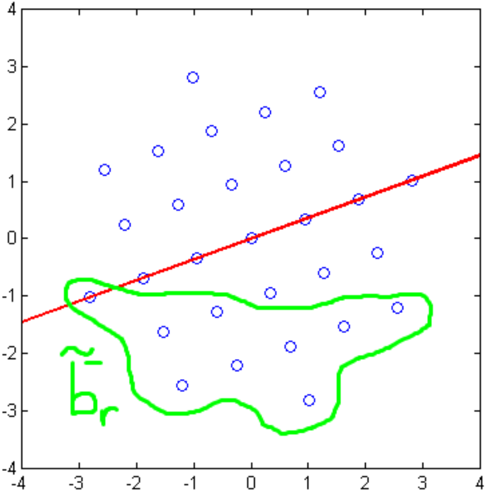

Figure

Back to the Fix

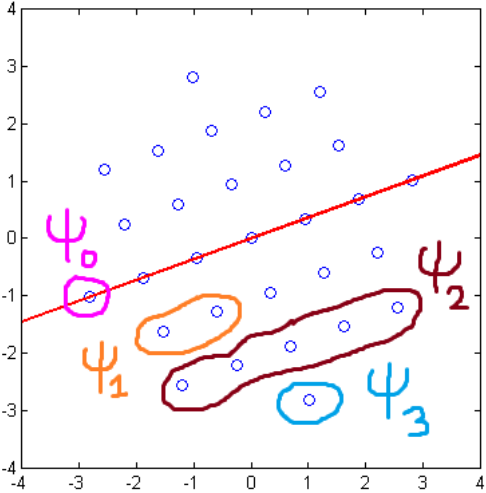

Figure

Back to the Fix

Assume $\Psi_0 \neq \emptyset$.

Back to the Fix

Assume $\Psi_0 \neq \emptyset$.

$$\mathbf{g}^*_{r,\mu} = \sum_{\mathbf{j}\in\Psi_0} \frac{\mathbf{j}}{\|\mathbf{j}\|} + \sum_{k=1}^{r} \sum_{\mathbf{j}\in\Psi_k} e^{-\frac{\mu^2}{2r^2}k^2\|\mathbf{g}\|^2} \frac{\mathbf{j}}{\|\mathbf{j}\|}$$

Back to the Fix

Assume $\Psi_0 \neq \emptyset$.

$$\begin{aligned}\mathbf{g}^*_{r,\mu} &= \sum_{\mathbf{j}\in\Psi_0} \frac{\mathbf{j}}{\|\mathbf{j}\|} + \sum_{k=1}^{r} \sum_{\mathbf{j}\in\Psi_k} e^{-\frac{\mu^2}{2r^2}k^2\|\mathbf{g}\|^2} \frac{\mathbf{j}}{\|\mathbf{j}\|} \\ &\to \sum_{\mathbf{j}\in\Psi_0} \frac{\mathbf{j}}{\|\mathbf{j}\|} \qquad \text{as } \mu \to \infty.\end{aligned}$$

Back to the Fix

Assume $\Psi_0 \neq \emptyset$.

$$\begin{aligned}
\mathbf{g}^*_{r,\mu} &= \sum_{\mathbf{j}\in\Psi_0} \frac{\mathbf{j}}{\|\mathbf{j}\|} + \sum_{k=1}^{r} \sum_{\mathbf{j}\in\Psi_k} e^{-\frac{\mu^2}{2r^2}k^2\|\mathbf{g}\|^2} \frac{\mathbf{j}}{\|\mathbf{j}\|} \\
&\to \sum_{\mathbf{j}\in\Psi_0} \frac{\mathbf{j}}{\|\mathbf{j}\|} \qquad \text{as } \mu \to \infty. \\
&= \mathbf{g} \qquad \text{after rescaling.}
\end{aligned}$$

Back to the Fix

Assume $\Psi_0 \neq \emptyset$.

$$\begin{aligned} \mathbf{g}^*_{r,\mu} &= \sum_{\mathbf{j}\in\Psi_0} \frac{\mathbf{j}}{\|\mathbf{j}\|} + \sum_{k=1}^{r} \sum_{\mathbf{j}\in\Psi_k} e^{-\frac{\mu^2}{2r^2}k^2\|\mathbf{g}\|^2} \frac{\mathbf{j}}{\|\mathbf{j}\|} \\ &\to \sum_{\mathbf{j}\in\Psi_0} \frac{\mathbf{j}}{\|\mathbf{j}\|} \qquad \text{as } \mu \to \infty. \\ &= \mathbf{g} \qquad \text{after rescaling.} \end{aligned}$$

Can show $\Psi_0 \neq \emptyset$ if $\theta > \theta_c(r)$.

Comparison with real Life

Theoretical Curve $(r = 3)$

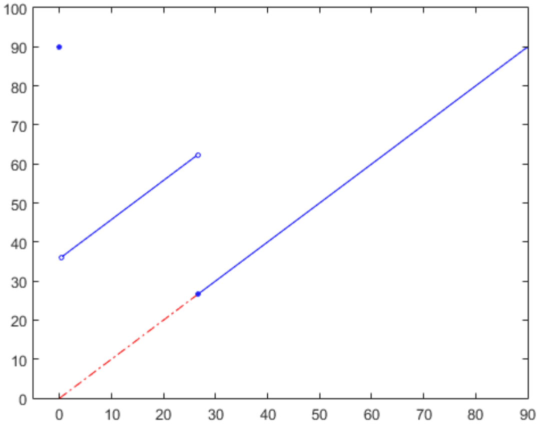

Figure

Comparison with real Life

Real Curve $(r = 3,\ \mu = 40)$

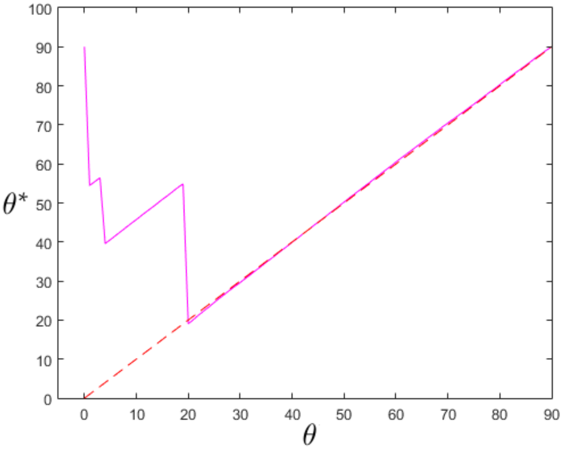

Figure

Comparison with real Life

Theoretical Curve $(r = 5)$

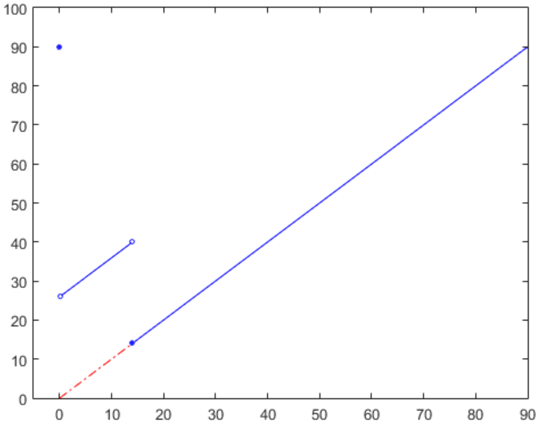

Figure

Comparison with real Life

Real Curve $(r = 5,\ \mu = 40)$

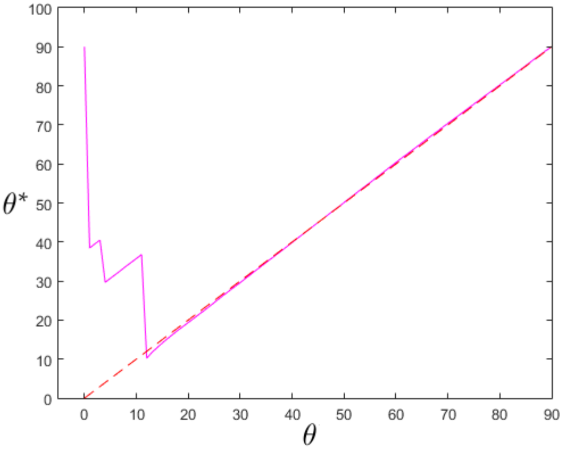

Figure

Convergence

- For *smooth* boundary data, proving convergence is routine.

Convergence

- For *smooth* boundary data, proving convergence is routine.
- However, *nonsmooth* boundary data (e.g. images) is much more challenging.

Convergence

- In this case, the fact that our weights are non-negative and sum to one means they can be interpretted as a probability density.

Convergence

- In this case, the fact that our weights are non-negative and sum to one means they can be interpretted as a probability density.
- This opens the door to a probabilistic line of attack based on martingales.

Convergence

- In this case, the fact that our weights are non-negative and sum to one means they can be interpretted as a probability density.
- This opens the door to a probabilistic line of attack based on martingales.
- Enables us to prove convergence even for data with jump discontinuities.

Better weights

Figure

Better weights

Figure

Better weights

Figure

Better weights

Figure

Better weights

Figure

Better weights

Figure

Better weights

Figure

Better weights

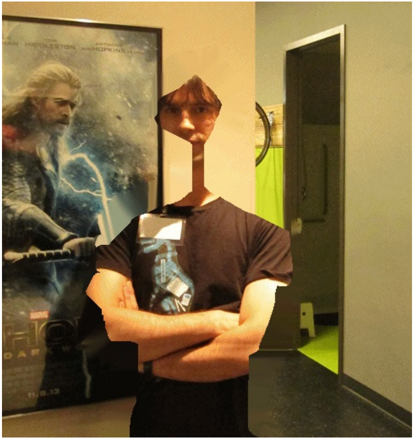

Figure

Better weights

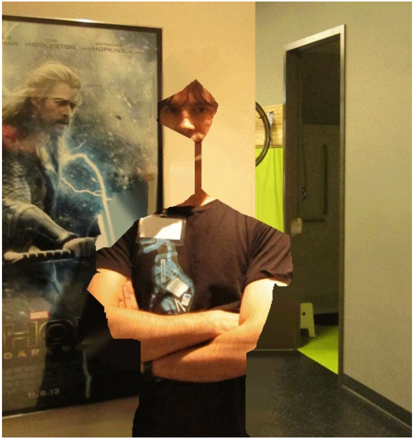

Figure

Better weights

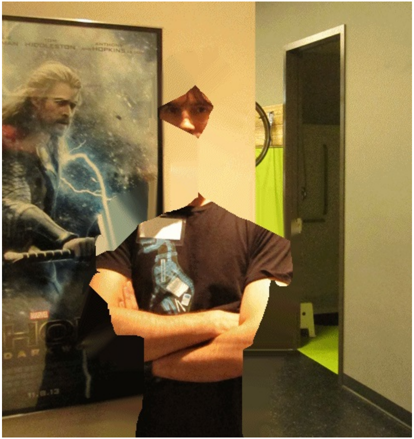

Figure

Better weights

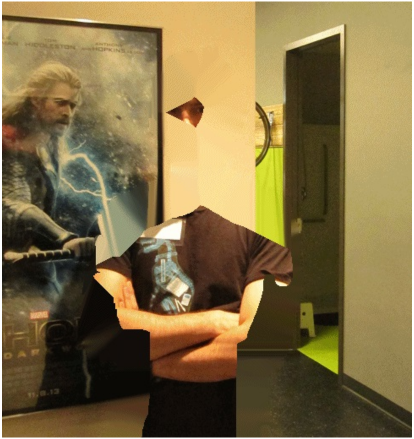

Figure

Better weights

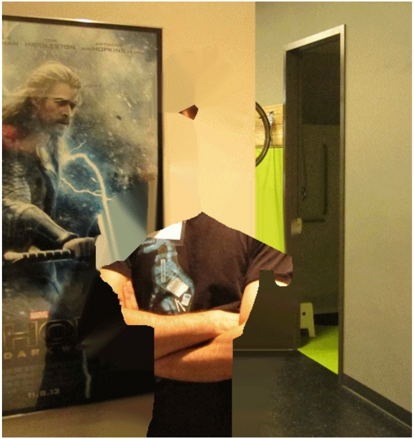

Figure

Better weights

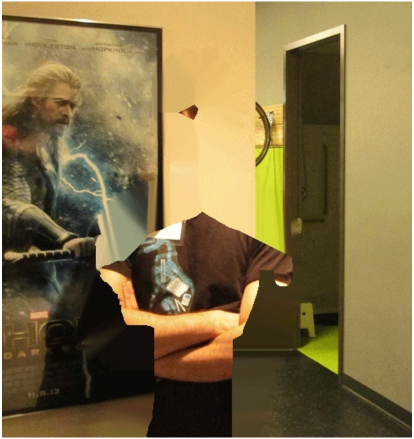

Figure

The End