Image inpainting and a geometric model of the visual cortex.

Anton Lukyanenko

UIUC Mathematics Department

November 30, 2011

Outline

Outline

Outline

Outline

Inpainting and disocclusion: examples

Inpainting and disocclusion: examples

Inpainting and disocclusion: examples

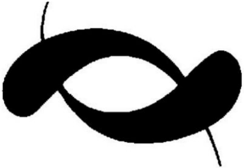

Inpainting and disocclusion: examples



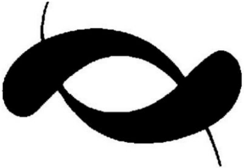
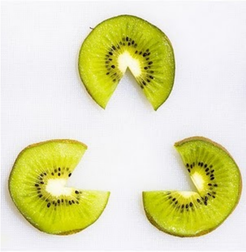

Inpainting and disocclusion: curve completion

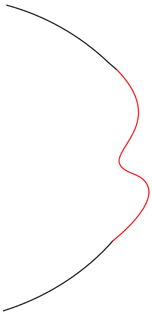

Curve completion should have:

1. Isotropy
2. Scale invariance
3. Smoothness
4. Extensibility
5. Roundness
6. Total minimum curvature

Problem

The curve completion axioms cannot all be satisfied simultaneously.

Inpainting and disocclusion: curve completion

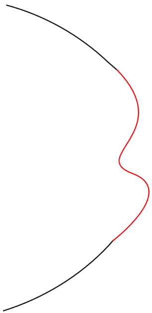

Curve completion should have:

1. Isotropy
2. Scale invariance
3. Smoothness
4. Extensibility
5. Roundness
6. Total minimum curvature

Problem

The curve completion axioms cannot all be satisfied simultaneously.

Inpainting and disocclusion: curve completion

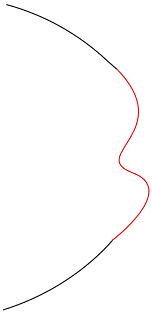

Curve completion should have:

1. Isotropy
2. Scale invariance
3. Smoothness
4. Extensibility
5. Roundness
6. Total minimum curvature

Problem

The curve completion axioms cannot all be satisfied simultaneously.

Inpainting and disocclusion: image restoration

Key methods for inpainting:

1. Averaging:
   Laplace equation
2. Prolongation of contour lines
3. Diffusion

Inpainting and disocclusion: image restoration

Key methods for inpainting:

1. Averaging:
   Laplace equation
2. Prolongation of contour lines
3. Diffusion

Inpainting and disocclusion: image restoration

Key methods for inpainting:

1. Averaging:
   Laplace equation
2. Prolongation of contour lines
3. Diffusion

Inpainting and disocclusion: image restoration

Key methods for inpainting:

1. Averaging:
   Laplace equation
2. Prolongation of contour lines
3. Diffusion

Inpainting and disocclusion: image restoration

Key methods for inpainting:

1. Averaging:
   Laplace equation
2. Prolongation of contour lines
3. Diffusion

Inpainting and disocclusion: image restoration

Key methods for inpainting:

1. Averaging:
   Laplace equation
2. Prolongation of contour lines
3. Diffusion

Question

What about disocclusion?

Inpainting and disocclusion: image restoration

Key methods for inpainting:

1. Averaging:
   Laplace equation
2. Prolongation of contour lines
3. Diffusion

Question

What about disocclusion?

Question

Is this what the brain does?

Biological inspiration: V1

Early visual processing:

1. Retina cells react to light intensity.
2. Information sent to striate cortex V1.

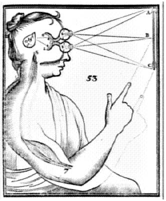

Biological inspiration: V1

Early visual processing:

1. Retina cells react to light intensity.
2. Information sent to striate cortex V1.

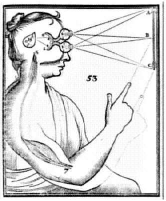

3. Each V1 cell responds to a region of retinal data.
4. Cell responses are modeled by Gabor filters.

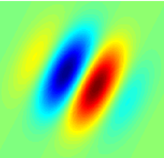

Biological inspiration: natural images

Olshausen, Field 1996:
Gabor filters are good for encoding real-world images. A learning algorithm decided the following 400 filters are the best.

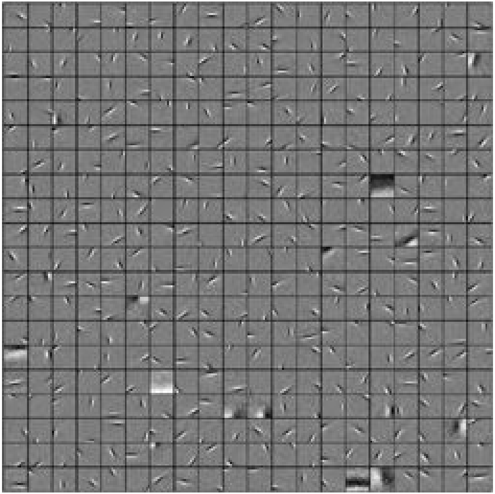

Biological inspiration: V1 structure

1. V1 cells have:
   - Receptive field
   - Direction preference
2. Columns and hypercolumns
3. Long-range connections
4. Association field

Biological inspiration: V1 structure

1. V1 cells have:
    - Receptive field
    - Direction preference
2. Columns and hypercolumns
3. Long-range connections
4. Association field

Biological inspiration: V1 structure

1. V1 cells have:
   - Receptive field
   - Direction preference
2. Columns and hypercolumns
3. Long-range connections
4. Association field

Biological inspiration: V1 structure

1. V1 cells have:
   - Receptive field
   - Direction preference
2. Columns and hypercolumns
3. Long-range connections
4. Association field

Biological inspiration: V1 structure

1. V1 cells have:
   - Receptive field
   - Direction preference
2. Columns and hypercolumns
3. Long-range connections
4. Association field

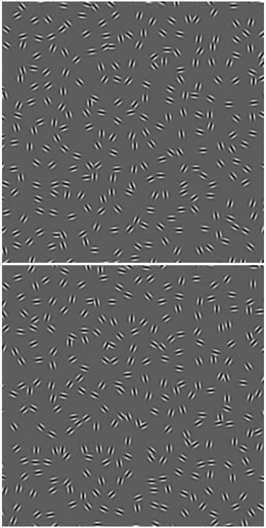

Biological inspiration: V1 structure

1. V1 cells have:
   - Receptive field
   - Direction preference
2. Columns and hypercolumns
3. Long-range connections
4. Association field

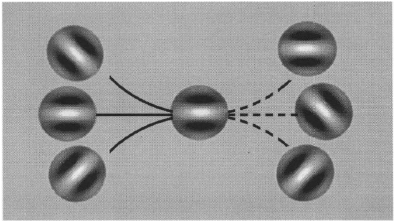

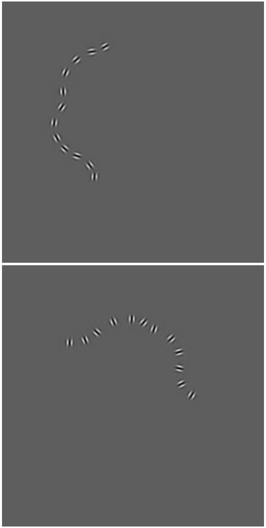

Mathematical filling: modeling V1 structure

We will use the following model of V1:

1. Each cell has coordinates (x, y, θ).
2. (x, y) represent the center of the receptive field.
3. θ represents the orientation preference.

Neural connections are encoded in the geometry.

1. "Allowed" curves represent signal propagation.
2. An allowed curve starting at (x, y, θ) can move *only*:
   - Along θ.
   - Along (x, y) in the θ direction.
   - A combination of these.
3. Note: this model also describes wheelbarrow motion.
4. Distance measured using *only* allowed curves.
5. Distance is the length of the shortest allowed curve.
6. (Chow's Theorem) An allowed curve connects any two cells.

Mathematical filling: modeling V1 structure

We will use the following model of V1:

1. Each cell has coordinates (x, y, θ).
2. (x, y) represent the center of the receptive field.
3. θ represents the orientation preference.

Neural connections are encoded in the geometry.

1. "Allowed" curves represent signal propagation.
2. An allowed curve starting at (x, y, θ) can move *only*:
   - Along θ.
   - Along (x, y) in the θ direction.
   - A combination of these.
3. Note: this model also describes wheelbarrow motion.
4. Distance measured using *only* allowed curves.
5. Distance is the length of the shortest allowed curve.
6. (Chow's Theorem) An allowed curve connects any two cells.

Mathematical filling: modeling V1 structure

We will use the following model of V1:

1. Each cell has coordinates (x, y, θ).
2. (x, y) represent the center of the receptive field.
3. θ represents the orientation preference.

Neural connections are encoded in the geometry.

1. "Allowed" curves represent signal propagation.
2. An allowed curve starting at (x, y, θ) can move *only*:
   - Along θ.
   - Along (x, y) in the θ direction.
   - A combination of these.
3. Note: this model also describes wheelbarrow motion.
4. Distance measured using *only* allowed curves.
5. Distance is the length of the shortest allowed curve.
6. (Chow's Theorem) An allowed curve connects any two cells.

Mathematical filling: modeling V1 structure

We will use the following model of V1:

1. Each cell has coordinates (x, y, θ).
2. (x, y) represent the center of the receptive field.
3. θ represents the orientation preference.

Neural connections are encoded in the geometry.

1. "Allowed" curves represent signal propagation.
2. An allowed curve starting at (x, y, θ) can move *only*:
    - Along θ.
    - Along (x, y) in the θ direction.
    - A combination of these.
3. Note: this model also describes wheelbarrow motion.
4. Distance measured using *only* allowed curves.
5. Distance is the length of the shortest allowed curve.
6. (Chow's Theorem) An allowed curve connects any two cells.

Mathematical filling: modeling V1 structure

We will use the following model of V1:

1. Each cell has coordinates (x, y, θ).
2. (x, y) represent the center of the receptive field.
3. θ represents the orientation preference.

Neural connections are encoded in the geometry.

1. "Allowed" curves represent signal propagation.
2. An allowed curve starting at (x, y, θ) can move *only*:
   - Along θ.
   - Along (x, y) in the θ direction.
   - A combination of these.
3. Note: this model also describes wheelbarrow motion.
4. Distance measured using *only* allowed curves.
5. Distance is the length of the shortest allowed curve.
6. (Chow's Theorem) An allowed curve connects any two cells.

Mathematical filling: modeling V1 structure

We will use the following model of V1:

1. Each cell has coordinates (x, y, θ).
2. (x, y) represent the center of the receptive field.
3. θ represents the orientation preference.

Neural connections are encoded in the geometry.

1. "Allowed" curves represent signal propagation.
2. An allowed curve starting at (x, y, θ) can move *only*:
   - Along θ.
   - Along (x, y) in the θ direction.
   - A combination of these.
3. Note: this model also describes wheelbarrow motion.
4. Distance measured using *only* allowed curves.
5. Distance is the length of the shortest allowed curve.
6. (Chow's Theorem) An allowed curve connects any two cells.

Mathematical filling: modeling V1 structure

We will use the following model of V1:

1. Each cell has coordinates (x, y, θ).
2. (x, y) represent the center of the receptive field.
3. θ represents the orientation preference.

Neural connections are encoded in the geometry.

1. "Allowed" curves represent signal propagation.
2. An allowed curve starting at (x, y, θ) can move *only*:
   - Along θ.
   - Along (x, y) in the θ direction.
   - A combination of these.
3. Note: this model also describes wheelbarrow motion.
4. Distance measured using *only* allowed curves.
5. Distance is the length of the shortest allowed curve.
6. (Chow's Theorem) An allowed curve connects any two cells.

Mathematical filling: modeling V1 structure

We will use the following model of V1:

1. Each cell has coordinates (x, y, θ).
2. (x, y) represent the center of the receptive field.
3. θ represents the orientation preference.

Neural connections are encoded in the geometry.

1. "Allowed" curves represent signal propagation.
2. An allowed curve starting at (x, y, θ) can move *only*:
   - Along θ.
   - Along (x, y) in the θ direction.
   - A combination of these.
3. Note: this model also describes wheelbarrow motion.
4. Distance measured using *only* allowed curves.
5. Distance is the length of the shortest allowed curve.
6. (Chow's Theorem) An allowed curve connects any two cells.

Mathematical filling: modeling V1 structure

We will use the following model of V1:

1. Each cell has coordinates (x, y, θ).
2. (x, y) represent the center of the receptive field.
3. θ represents the orientation preference.

Neural connections are encoded in the geometry.

1. "Allowed" curves represent signal propagation.
2. An allowed curve starting at (x, y, θ) can move *only*:
   - Along θ.
   - Along (x, y) in the θ direction.
   - A combination of these.
3. Note: this model also describes wheelbarrow motion.
4. Distance measured using *only* allowed curves.
5. Distance is the length of the shortest allowed curve.
6. (Chow's Theorem) An allowed curve connects any two cells.

Mathematical filling: modeling V1 structure

We will use the following model of V1:

1. Each cell has coordinates (x, y, θ).
2. (x, y) represent the center of the receptive field.
3. θ represents the orientation preference.

Neural connections are encoded in the geometry.

1. "Allowed" curves represent signal propagation.
2. An allowed curve starting at (x, y, θ) can move *only*:
   - Along θ.
   - Along (x, y) in the θ direction.
   - A combination of these.
3. Note: this model also describes wheelbarrow motion.
4. Distance measured using *only* allowed curves.
5. Distance is the length of the shortest allowed curve.
6. (Chow's Theorem) An allowed curve connects any two cells.

Mathematical filling: modeling V1 activity

Processing of visual data is encoded by a "lifting":

1. Retinal data is encoded as an intensity function $I(x, y)$.
2. The direction of a contour at (x, y) is given by $\nabla I(x, y)$.
3. Cells near $(x, y, \theta(\nabla I(x, y)))$ get excited.
4. $I(x, y)$ is replaced by a surface S.

Mathematical filling: modeling V1 activity

Processing of visual data is encoded by a "lifting":

1. Retinal data is encoded as an intensity function $I(x, y)$.
2. The direction of a contour at (x, y) is given by $\nabla I(x, y)$.
3. Cells near $(x, y, \theta(\nabla I(x, y)))$ get excited.
4. $I(x, y)$ is replaced by a surface S.

Mathematical filling: modeling V1 activity

Processing of visual data is encoded by a "lifting":

1. Retinal data is encoded as an intensity function $I(x, y)$.
2. The direction of a contour at (x, y) is given by $\nabla I(x, y)$.
3. Cells near $(x, y, \theta(\nabla I(x, y)))$ get excited.
4. $I(x, y)$ is replaced by a surface S.

Mathematical filling: modeling V1 activity

Processing of visual data is encoded by a "lifting":

1. Retinal data is encoded as an intensity function $I(x, y)$.
2. The direction of a contour at (x, y) is given by $\nabla I(x, y)$.
3. Cells near $(x, y, \theta(\nabla I(x, y)))$ get excited.
4. $I(x, y)$ is replaced by a surface S.

Mathematical filling: modeling V1 activity

Processing of visual data is encoded by a "lifting":

1. Retinal data is encoded as an intensity function $I(x, y)$.
2. The direction of a contour at (x, y) is given by $\nabla I(x, y)$.
3. Cells near $(x, y, \theta(\nabla I(x, y)))$ get excited.
4. $I(x, y)$ is replaced by a surface S.

Original data:

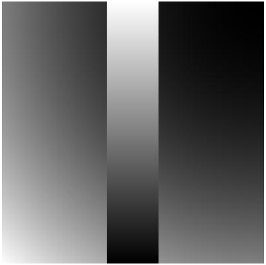

Mathematical filling: modeling V1 activity

Processing of visual data is encoded by a "lifting":

1. Retinal data is encoded as an intensity function $I(x, y)$.
2. The direction of a contour at (x, y) is given by $\nabla I(x, y)$.
3. Cells near $(x, y, \theta(\nabla I(x, y)))$ get excited.
4. $I(x, y)$ is replaced by a surface S.

Original data:

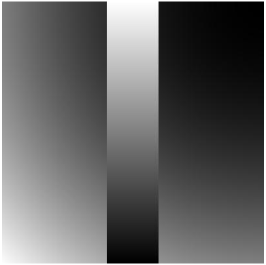

Associated surface in V1:

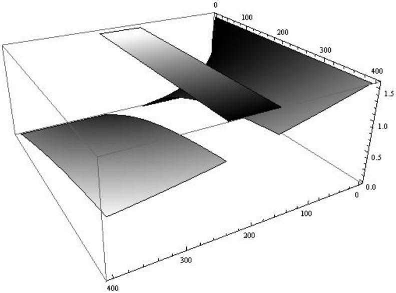

Mathematical filling: soap bubbles

Even in Euclidean space, filling surfaces is hard.

Original surface:

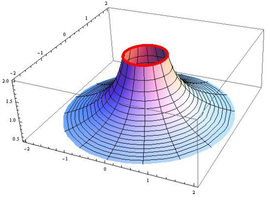

Mathematical filling: soap bubbles

Even in Euclidean space, filling surfaces is hard.

Original surface:

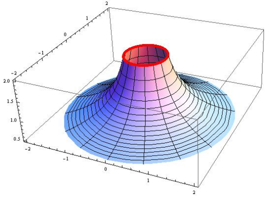

One potential filling:

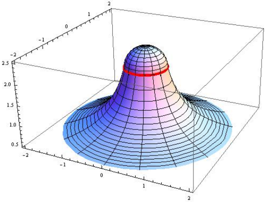

Mathematical filling: soap bubbles

Even in Euclidean space, filling surfaces is hard.

Original surface:

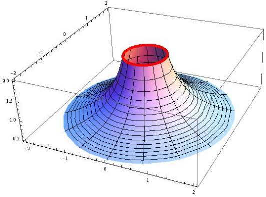

One potential filling:

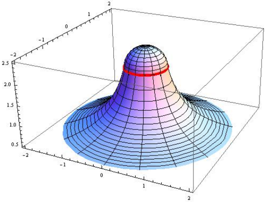

Least surface area,
officially the best:

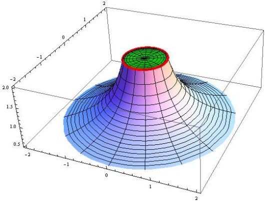

Mathematical filling: soap bubbles

Even in Euclidean space, filling surfaces is hard.

Original surface:

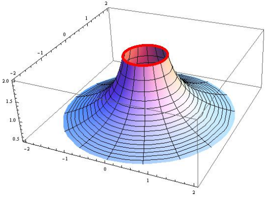

One potential filling:

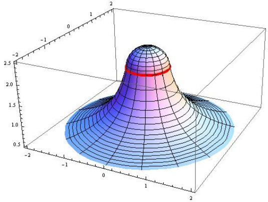

Least surface area,
officially the best:

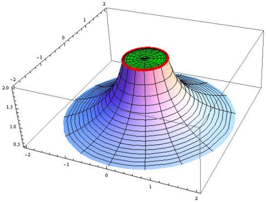

Small surface area,
but problems appearing:

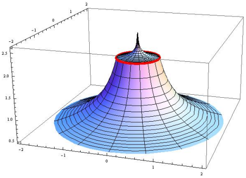

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along "allowed directions".
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces "best" any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford '92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti '06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls '08,'09) In certain cases, minimal surfaces are extremely easy to find.

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along "allowed directions".
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces "best" any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford '92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti '06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls '08,'09) In certain cases, minimal surfaces are extremely easy to find.

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along "allowed directions".
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces "best" any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford '92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti '06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls '08,'09) In certain cases, minimal surfaces are extremely easy to find.

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along "allowed directions".
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces "best" any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford '92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti '06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls '08,'09) In certain cases, minimal surfaces are extremely easy to find.

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along “allowed directions”.
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces “best” any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford '92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti '06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls '08,'09) In certain cases, minimal surfaces are extremely easy to find.

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along "allowed directions".
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces "best" any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford '92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti '06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls '08,'09) In certain cases, minimal surfaces are extremely easy to find.

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along "allowed directions".
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces "best" any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford '92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti '06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls '08,'09) In certain cases, minimal surfaces are extremely easy to find.

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along "allowed directions".
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces "best" any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford '92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti '06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls '08,'09) In certain cases, minimal surfaces are extremely easy to find.

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along “allowed directions”.
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces “best” any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford ’92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti ’06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls ’08,’09) In certain cases, minimal surfaces are extremely easy to find.

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along "allowed directions".
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces "best" any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford '92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti '06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls '08,'09) In certain cases, minimal surfaces are extremely easy to find.

Mathematical filling: minimal surfaces in V1 geometry

Our model of V1 is not Euclidean, which causes many issues:

1. Information moves only along “allowed directions”.
2. Propagation in other directions is possible, but indirect.
3. Notion of surface area needs adjustment.
4. Finding minimal surfaces becomes even harder.
5. Are minimal surfaces “best” any more?

These problems have been resolved:

1. V1 is a *Sub-Riemannian manifold*.
2. Surface area, etc. of SR manifolds is now well-studied.
3. (Mumford ’92) Brownian motion in V1 leads to minimal surfaces.
4. (Citti-Sarti ’06) Minimal surfaces useful for image completion.
5. (Hladky-Pauls ’08,’09) In certain cases, minimal surfaces are extremely easy to find.

Conclusion: examples

Citti-Sarti example:

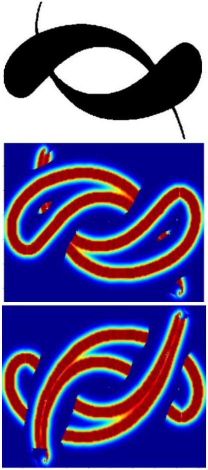

Hladky-Pauls examples:

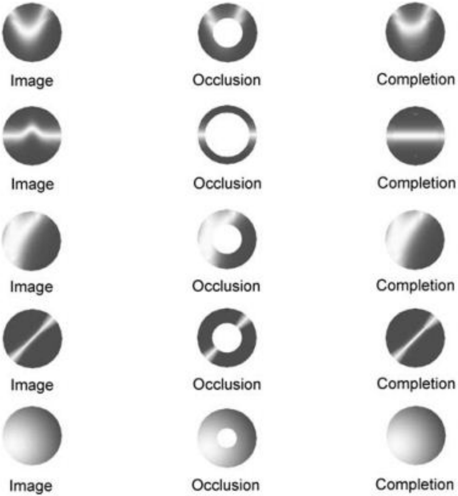

Conclusion: questions

Question

How useful and accurate is this model of V1?

Question

How useful are the inpainting and disocclusion methods?

Question

How can the method be generalized to include extra dimensions (e.g. video correction) or color?

Conclusion: questions

Question

How useful and accurate is this model of V1?

Question

How useful are the inpainting and disocclusion methods?

Question

How can the method be generalized to include extra dimensions (e.g. video correction) or color?

Conclusion: questions

Question

How useful and accurate is this model of V1?

Question

How useful are the inpainting and disocclusion methods?

Question

How can the method be generalized to include extra dimensions (e.g. video correction) or color?