NICTA

Move-based Algorithms for the Optimization of an Isotropic Gradient MRF Model

Behrooz Nasihatkon   Richard Hartley

5 December 2012

Australian Government
Department of Broadband, Communications and the Digital Economy
Australian Research Council

NICTA Funding and Supporting Members and Partners

Outline

Total Variation

- good **regularizer**.

Total Variation

- good **regularizer**.
- For a function $x: \Omega \subseteq \mathbb{R}^n \to \mathbb{R}$ Total Variation is defined as

$$\mathrm{TV}(x) = \int_{\Omega} |\nabla_{\mathbf{t}} x(\mathbf{t})| \, d\mathbf{t}$$

Total Variation

- good **regularizer**.
- For a function $x : \Omega \subseteq \mathbb{R}^n \to \mathbb{R}$ Total Variation is defined as

$$\mathrm{TV}(x) = \int_{\Omega} |\nabla_{\mathbf{t}} x(\mathbf{t})| \, d\mathbf{t}$$

- **discontinuity preserving** (edge preserving for images).

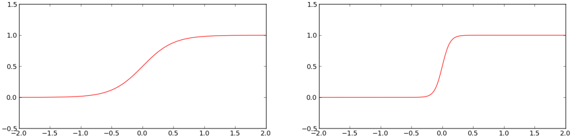

$x_1(t)$ $x_2(t)$

$$\mathrm{TV}(x_1) = \mathrm{TV}(x_2)$$

Modeling TV using MRFs

NICTA

- Approximate Total Variation using an MRF,
- A set of nodes $1, 2, \ldots, n$,
- A set of labels $\mathbf{x} = [x_1, x_2, \ldots, x_n]$, $x_i \in \mathcal{L}$.
- Energy function

$$E(\mathbf{x}) = \sum_i f_i(x_i) + \tilde{TV}(\mathbf{x}),$$

Current Models

- Approximate Magnitude of Gradient using **edge-based** potentials.

$$\tilde{TV}(\mathbf{x}) = \sum_{(i,j)\in\mathcal{C}^2} w_{ij} \left| x_i - x_j \right|$$

- Magnitude of Gradient (MoG) at each node i is approximated by

$$\mathrm{MoG}(i) = \sum_{j\in\mathcal{N}_i} w_{ij} \left| x_i - x_j \right|$$

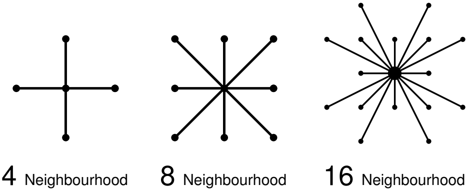

4-connected model

$$\mathrm{MoG}(i) = |x_i - x_j| + |x_i - x_k|$$

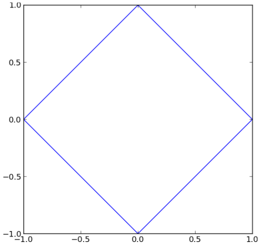

$$\mathrm{MoG}(i) = 1$$

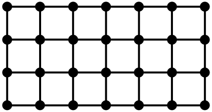

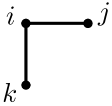

Diagonal Edges

$$\mathrm{MoG}(i) = |x_i - x_j| + |x_i - x_k| + \frac{\sqrt{2}}{2}|x_i - x_l|$$

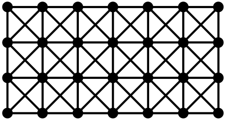

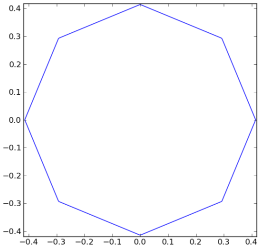

$\mathrm{MoG}(i) = 1$

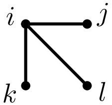

The 3-Clique Model

- The gradient vector $\approx \begin{pmatrix} x_j - x_i \\ x_k - x_i \end{pmatrix}$,

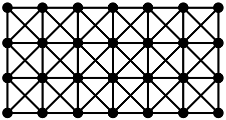

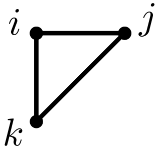

The 3-Clique Model

- The gradient vector $\approx \begin{pmatrix} x_j - x_i \\ x_k - x_i \end{pmatrix}$,
- For ordered labels $x_i \in \{1, 2, \ldots, M\}$

$$\mathrm{MoG}(i) = \sqrt{(x_i - x_j)^2 + (x_i - x_k)^2}$$

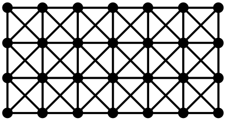

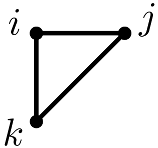

The 3-Clique Model

- The gradient vector $\approx \begin{pmatrix} x_j - x_i \\ x_k - x_i \end{pmatrix}$,
- For ordered labels $x_i \in \{1, 2, \ldots, M\}$

$$\mathrm{MoG}(i) = \sqrt{(x_i - x_j)^2 + (x_i - x_k)^2}$$

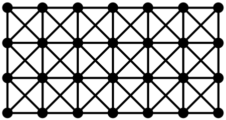

- For general labels $x_i \in \mathcal{L}$ with a semi-metric d

$$\mathrm{MoG}(i) = \sqrt{d(x_i, x_j)^2 + d(x_i, x_k)^2}$$

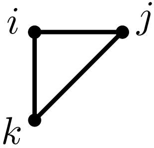

The 3-clique Model

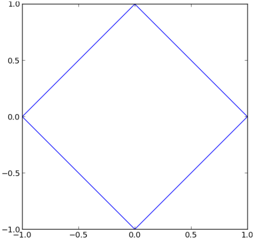

4 neighbours

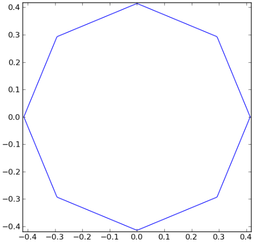

8 neighbours

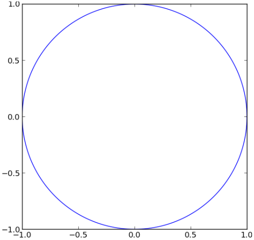

3-cliques

Move-Based Algorithms

$$\min_{\mathbf{x}\in\mathcal{L}} \sum_i f_i(x_i) + \gamma \sum_{(i,j,k)\in\mathcal{C}^3} \sqrt{d(x_i, x_j)^2 + d(x_i, x_k)^2}$$

- **Move-based approach** is a popular way of optimizing Multi-label MRFs.
- Optimizing the multi-label MRF iteratively by solving **a series of binary MRF optimizations**.

$x_1 \; x_2 \; x_3 \; x_4 \cdots$ $\qquad$ $u_1 \; u_2 \; u_3 \; u_4 \cdots$

$x_i \in \mathcal{L}$ $\qquad$ $u_i \in \{0, 1\}$

The Alpha-Expansion Algorithm

- Nodes have a choice to switch to α or stay unchanged:

$$l_\alpha^0(x_i) = x_i \qquad l_\alpha^1(x_i) = \alpha$$

- $\mathbf{l}_\alpha^{\mathbf{u}}(\mathbf{x}) = [l_\alpha^{u_1}(x_1), l_\alpha^{u_2}(x_2), \ldots, l_\alpha^{u_n}(x_n)]$

procedure ALPHA-EXPANSION($\mathbf{x}, \mathcal{L}$)
  repeat
    for each $\alpha \in \mathcal{L}$ **do**
      $\mathbf{u}^* \leftarrow \text{argmin}_{\mathbf{u}} E(\mathbf{l}_\alpha^{\mathbf{u}}(\mathbf{x}))$
      $\mathbf{x} \leftarrow \mathbf{l}_\alpha^{\mathbf{u}^*}(\mathbf{x})$
    end for
  until *convergence*
  return x
end procedure

The Alpha-Beta Swap Algorithm

- Nodes with labels α or β have a chance to swap.

$$I^0_{\alpha,\beta}(x_i) = \begin{cases} \alpha & \text{if } x_i \in \{\alpha, \beta\} \\ x_i & \text{otherwise} \end{cases} \qquad I^1_{\alpha,\beta}(x_i) = \begin{cases} \beta & \text{if } x_i \in \{\alpha, \beta\} \\ x_i & \text{otherwise} \end{cases}$$

procedure ALPHA-BETA-SWAP($\mathbf{x}, \mathcal{L}$)
  repeat
    for each $\alpha, \beta \in \mathcal{L} \times \mathcal{L}$ **do**
      $\mathbf{u}^* \leftarrow \operatorname{argmin}_{\mathbf{u}} E(\mathbf{I}^{\mathbf{u}}_{\alpha,\beta}(\mathbf{x}))$
      $\mathbf{x} \leftarrow \mathbf{I}^{\mathbf{u}^*}_{\alpha,\beta}(\mathbf{x})$
    end for
  until *convergence*
  return x
end procedure

General Move Algorithm

- Take arbitrary l^0 and l^1

$$l^0(x_i) = \text{arbitrary} \qquad l^1(x_i) = \text{arbitrary}$$

General Move Algorithm

NICTA

- Take arbitrary I^0 and I^1

$$I^0(x_i) = \text{arbitrary} \qquad\qquad I^1(x_i) = \text{arbitrary}$$

- The pair of functions (I^0, I^1) is called the **update policy**.
- **State Preservation Property**

$$\forall x \in \mathcal{L} \quad I^0(x) = x \quad \text{or} \quad I^1(x) = x$$

Solving the Binary Problem

- How to solve

$$\mathbf{u}^* \leftarrow \mathrm{argmin}_{\mathbf{u}} E(\mathbf{I}^{\mathbf{u}}(\mathbf{x}))$$

- Energy functions consisting of **quadratic** and **cubic** terms are solvable by graph-cuts if and only if they are **submodular**[1].
- The function $f\colon \{0,1\}\times\{0,1\} \to \mathbb{R}$ is **submodular** if

$$f(0,1) + f(1,0) \geq f(0,0) + f(1,1).$$

- A pseudo-Boolean function of n variables is submodular if **any restriction** to any pair of variables is submodular.

[1] Kolmogorov and Zabih 2004.

Central Question

NICTA

- **Main Question:** Given

$$E(\mathbf{x}) = \sum_{(i,j,k)\in\mathcal{C}^3} \sqrt{d(x_i, x_j)^2 + d(x_i, x_k)^2}$$

what choice of policy (l^0, l^1) results in a submodular $E(\mathbf{l^u}(\mathbf{x}))$ as a function of $\mathbf{u}$, so we can solve

$$\mathbf{u}^* \leftarrow \mathrm{argmin}_{\mathbf{u}} E(\mathbf{l^u}(\mathbf{x}))$$

Main Theorem (General Case)

$$E(\mathbf{x}) = \sum_{(i,j,k)\in\mathcal{C}^3} \sqrt{d(x_i, x_j)^2 + d(x_i, x_k)^2}$$

Theorem

Assume $d\colon \mathcal{L} \times \mathcal{L} \to \mathbb{R}$ is a **semi-metric** *and the update policy has the* **state preservation property**, *the energy function $E'_{\mathbf{x}}(\mathbf{u}) = E(\mathbf{l}^{\mathbf{u}}(\mathbf{x}))$ is submodular for all* **x**, **if and only if** *for any three labels $x, y, z \in \mathcal{L}$*

$$\left(d(x, y^1) - d(x, y^0)\right)\left(d(x, z^1) - d(x, z^0)\right) \geq 0,$$

where x^u is a compact form for $l^u(x)$.

Main Theorem (Ordered Labels)

NICTA

$$E(\mathbf{x}) = \sum_{(i,j,k)\in\mathcal{C}^3} \sqrt{(x_i - x_j)^2 + (x_i - x_k)^2}$$

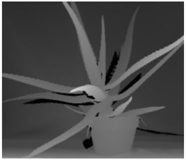

(Middlebury Dataset)

Proposition

With $\mathcal{L} = \{0, 1, \ldots, M-1\}$ and $d(x, y) = |x - y|$ (ordered labels), and having **state preservation property** *for the update policy, the energy function $E'_{\mathbf{x}}(\mathbf{u}) = E(\mathbf{l}^{\mathbf{u}}(\mathbf{x}))$ is submodular for all $\mathbf{x}$ if and only if (l^0, l^1) is a* **mirrored update policy**.

Mirrored Policy

Definition

An update policy (l^0, l^1) is called **mirrored** if

(i) $\forall x \in \mathcal{A} \quad l^0(x) < l^1(x)$ or $\forall x \in \mathcal{A} \quad l^0(x) > l^1(x)$,

(ii) $\exists \mu \in \mathcal{L}$ such that $\forall x \in \mathcal{A}$

$$\frac{l^0(x)+l^1(x)}{2} \in \{\mu,\ \mu+\frac{1}{2},\ \mu+1\}.$$

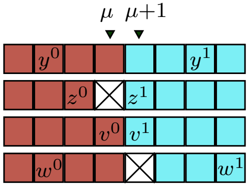

Main Theorem (Unordered Labels)

NICTA

$$E(\mathbf{x}) = \sum_{(i,j,k)\in\mathcal{C}^3} \sqrt{1(x_i \neq x_j) + 1(x_i \neq x_k)}$$

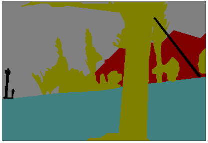

(Buffalo-Xiph.org)

Proposition

With $d(x,y) = 1(x \neq y)$ *(unordered labels), and assuming the* **state preservation property** *for the update policy, the energy function* $E'_{\mathbf{x}}(\mathbf{u}) = E(\mathbf{l}^{\mathbf{u}}(\mathbf{x}))$ *is submodular for all* $\mathbf{x}$ *if and only if for any pair of active labels* $y, z \in \mathcal{A}$*, we have* $l^0(y) \neq l^1(z)$*.*

Mirrored Swap

$$I_s^0(x) = \begin{cases} \min(x, s-x) & 0 \le s-x < M, \\ x & \text{otherwise,} \end{cases} \qquad I_s^1(x) = \begin{cases} \max(x, s-x) & 0 \le s-x < M, \\ x & \text{otherwise.} \end{cases}$$

procedure MIRRORED-SWAP($\mathbf{x}$, M)
    repeat
        for each $s \in \{1, 2, \ldots, 2n-3\}$ **do**
            $\mathbf{u}^* \leftarrow \mathrm{argmin}_{\mathbf{u}} E(\mathbf{I}_s^{\mathbf{u}}(\mathbf{x}))$
            $\mathbf{x} \leftarrow \mathbf{I}_s^{\mathbf{u}}(\mathbf{x})$
        end for
    until *convergence*
    return x
end procedure

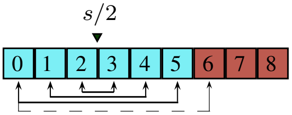

Conclusion

- Suits MRFs with the **continuous** and **ordered** labels,
- Alpha-expansion cannot be applied to **ordered labels**,
- Mirrored Swap algorithm for the **ordered labels**,
- For **unordered labels**, the submodularity holds for vaster types of binary moves, including alpha-expansion and alpha-beta swap.

Thanks

Questions?

? ? ? ?