COMPSCI 773 S1C

Dynamic Programming Stereo

Prof. Georgy Gimel'farb

1. Global optimisation
2. Profile model
3. GPV
4. Signal models
5. DP Stereo
6. Recurrent computation

Stereo Reconstruction: Practical Goal

Due to ill-posedness, precisely an original 3D scene simply cannot be reconstructed from its stereo pair

- Multiple-image 3D reconstruction is also ill-posed...

Goal is more limited and practical: to bring reconstructed surfaces close enough to visual (photogrammetic) perception

- Human visual perception combines binocular stereo with multiple monocular depth cues!
  - Human vision fails on uniform/ repetitive texture, large depth gradients, contrast differences, large y-disparities, etc.
- Computational stereo relies on regularising constraints to cope with ill-posedness (multiple equivalent solutions)
  - Due to a large variety of observed scenes, only very general prior knowledge is used to constrain optical surfaces
  - E.g. smoothness, depth gradient (curvature), continuity, ...

Stereo by Global Optimisation

Pros: low sensitivity to local errors

- Due to constraints on conjugate scanlines or entire images

Cons: generally, an NP-hard problem for 2D constraints on disparities in neighbouring points

- 2D constraints on corresponding signals make it harder!
- Feasible only in rare cases when direct exhaustion of surface variants (with exponential complexity) is avoided
    - Fast profile-wise 1D dynamic programming MAP / ML reconstruction takes no account of across-the-profiles constraints
    - Similar MAP/ML and MPM/MWM reconstruction by 1D belief propagation – the same drawback
- Most of approximate global optimisation algorithms are still too complex for large-size images of practical interest

Global Optimisation: Popular Tools

Exact 1D MAP/ML:

- Dynamic Programming; Belief Propagation (BP): max- or min-sum algorithm

Approximate 2D MAP/ML: Iterative Graph Cut; Loopy BP

- Exact minimum cut / maximum flow solution on networks with non-negative edge capacities and two special source and sink nodes
- Selecting a minimum subset of edges that separate the source from the sink and carry the maximum flow through the network
- Binary optimization on the lattice: reduces to the min-cut/max-flow for the cuts associated with 0–1 transitions

Exact 1D MPM/MWM: BP (sum-product algorithm)

Approximate 2D MPM/MWM: Loopy BP

No guaranteed convergence for loopy BP (but mostly exists in practice)

Markov Chain Profile Model

- Accounts for symmetry of stereo channels, visibility of 3D points and discontinuities due to occlusions
- A simplifying assumption: a single continuous surface only

Graph of profile variants (GPV)

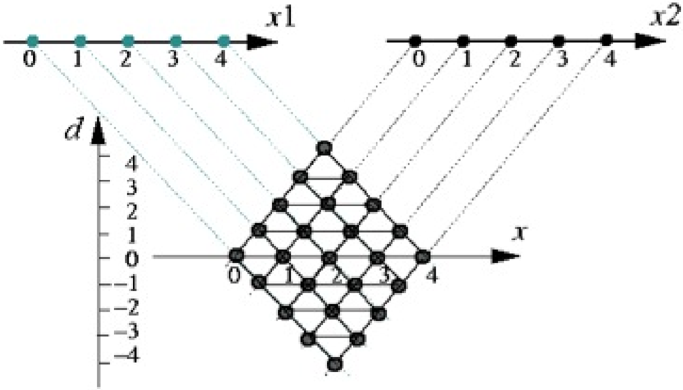

Markov Profile: Symmetric (x, d) Coordinates

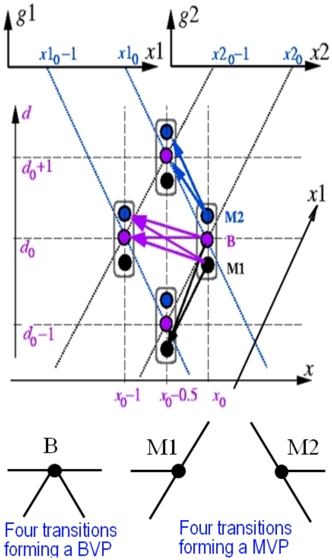

- Nodes along a continuous profile are subject to two constraints:
  1. Ordering constraint
  2. Visibility constraint
- $\{(x-1, d); (x-0.5, d \pm 1)\} \leftarrow (x, d)$
- Each node (x, d) has 3 states s
  - B – BVP (binocularly visible point): Point-wise matching score
  - M_i – MVP (monocularly visible point): Regularizing weight for partially occluded points with no correspondence
  - 8 allowable transitions
  - Popular constant MVP weight
    - More adequate MVP weights depend on related BVP scores

Graph of Profile Variants (GPV)

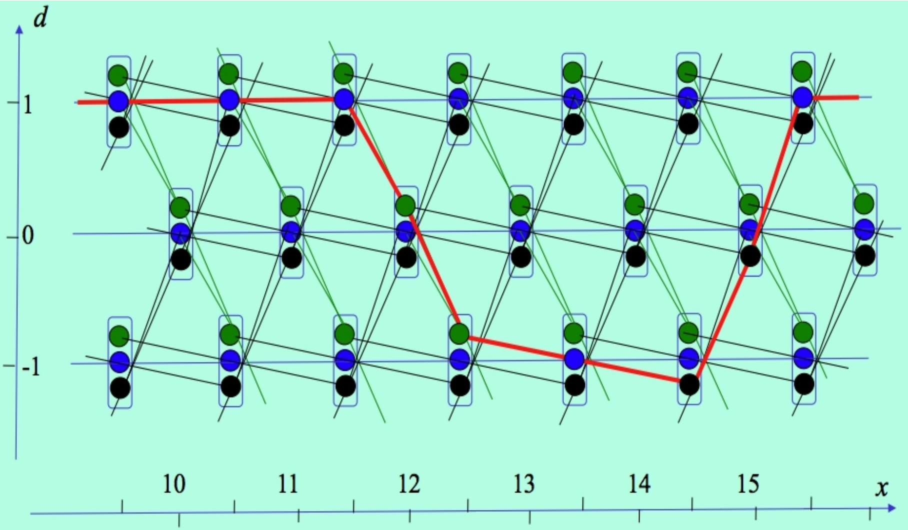

Allowable Transitions Along a Profile

- BVP (x, d, B):
  pixel-wise signal dissimilarity depending generally on the signals in the current and immediate preceding BVPs

$$\Rightarrow \left(g_{1:x+\frac{d}{2}};\ g_{2:x-\frac{d}{2}}\right);$$

$$\left(g_{1:x_{\mathrm{pr}}+\frac{d_{\mathrm{pr}}}{2}};\ g_{2:x_{\mathrm{pr}}-\frac{d_{\mathrm{pr}}}{2}}\right)$$

- MVP (x, d, M_1);
- MVP (x, d, M_2):
  fixed "dissimilarity" weight; generally, can depend on the related BVP dissimilarities

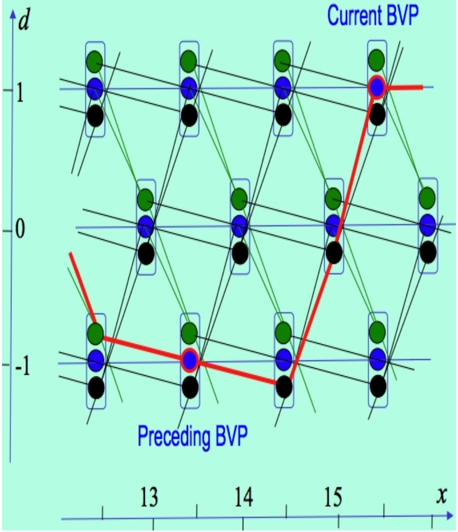

Simple Markov Chain x-Profile Model

Probability of a profile $\mathbf{d} = [(x_i, d_i, s_i) : \ i = 1, \ldots, n]$:

$$\Pr(\mathbf{d}|\mathbf{g}_1, \mathbf{g}_2) = p\left(x_1, d_1, s_1|\mathbf{g}_1, \mathbf{g}_2\right) \prod_{i=2}^{n} p\left(x_i, d_i, s_i|x_{i-1}, d_{i-1}, s_{i-1}; \mathbf{g}_1, \mathbf{g}_2\right)$$

- Each term depends on a transition from the GVP-node (x_{i-1}, d_{i-1}) in state s_{i-1} to the node (x_i, d_i) in state s_i along the profile
- Transitions are limited by the visibility states along a GVP
- The probability $p\left(x_i, d_i, \mathrm{B}|x_{i-1}, d_{i-1}, s_{i-1}; \mathbf{g}_1, \mathbf{g}_2\right)$ of a transition to state B depends on dissimilarity between the corresponding image signals for the current BVP on a profiles variant
    - Generally, it can also depend on the signals for the immediate preceding BVP along this variant
- Transition probabilities to the MVP can relate to those to the BVP
    - Typical simplification: a constant MVP probability

Simple Probability Models of Corresponding Signals

Symmetric model:

- Signal deviations in $\mathbf{g}_1$ and $\mathbf{g}_2$ w.r.t. an unobserved noiseless Cyclopean image (or ortho-image) $\mathbf{g}$ of a 3D scene

$$g_{1:x_1,y_1} = g_{x,y} + \nu_{1:x,y}; \quad g_{2:x_2,y_2} = g_{x,y} + \nu_{2:x,y}$$

- Independent central-symmetric random noise $\nu_{\dots}$: monotone decrease of the probability densities $p(\nu_{1:x,y}) \propto \exp(-\gamma\nu^2_{1:x,y})$ and $p(\nu_{2:x,y}) \propto \exp(-\gamma\nu^2_{2:x,y})$

Asymmetric models:

$$g_{1:x,y} = g_{2:x,y} + \nu_{x,y} \qquad \text{or} \qquad g_{2:x,y} = g_{1:x,y} + \nu_{x,y}$$

More Realistic Models of Corresponding Signals

Contrast deviation model – positive transfer factors, a, varying over the each camera field of view and pixel-wise independent random noise, n, of image sensors:

$$g_{1:x_1,y_1} = a_{1:x,y}g_{x,y} + \nu_{1:x,y}; \quad g_{2:x_2,y_2} = a_{2:x,y}g_{x,y} + \nu_{2:x,y}$$

- Transfer factors: strong interdependence for adjacent BVPs to account for visual resemblance of corresponding areas
- **Symmetric difference model** of the interdependence:
  - Limited direct proportion of the noiseless signal increments between the neighbouring BVPs (x, y) and (x', y) along the same epipolar profile: for $k = 1, 2$

$$\begin{aligned}\min_{e_k\in\varepsilon}\{e_k(g_{x,y} - g_{x',y})\} &\leq a_{k:x,y}g_{x,y} - a_{k:x',y}g_{x',y} \\ &\leq \max_{e_k\in\varepsilon}\{e_k(g_{x,y} - g_{x',y})\}\end{aligned}$$

  - $\varepsilon = [e_{\min}, e_{\max}]$ – a fixed range $0 < e_{\min} \leq e_{\max}$ of the difference factors e

Markov Chain Model of Contrast Deviations

Difference signal model → a **Markov chain of signals** for BVPs, which is mixed with independent signals for MVPs along a profile

- Assumption 1: Statistically independent ortho-image signals $\mathbf{g}$
- Assumption 2: Centre-symmetric independent random noise $\boldsymbol{\nu}$
- Assumption 3: Fixed range of allowable spatial contrast deviations
  - Analytic statistical estimates for an orthoimage $\mathbf{g}$ and transfer factors $\mathbf{a}_1, \mathbf{a}_2$ under a known 3-D profile
- Theoretically justified part of a (dis)similarity measure: from relationships between the corresponding signals for the BVPs
- Heuristic regularising part of the measure: for the MVPs
  - Fixed regularising values
  - Variable values derived from the assumed links between the MVPs and relevant BVPs

Dynamic Programming (DP) Stereo

Simplified notation with the omitted y-coordinate:

- For brevity:

$$g_i \equiv g_{x_i,y_i}; \quad g_{1:i} \equiv g_{1:x_i+\frac{d_i}{2},y}; \quad g_{2:i} \equiv g_{2:x_i-\frac{d_i}{2},y}$$

- $\mathbf{d} = ((x_i, d_i, s_i) : i = 1, 2, \ldots, N)$ – a digital profile (with allowable transitions between the adjacent GVP nodes)
- $\mathbf{g} = (g_i : i = 1, 2, \ldots, N)$ – the sequence of Cyclopean image signals along the profile $\mathbf{d}$
- $\mathbf{g}_1 = (g_{1:i} : i = 1, 2, \ldots, N_1)$ – the sequence of the corresponding left image signals for the profile $\mathbf{d}$
- $\mathbf{g}_2 = (g_{2:i} : i = 1, 2, \ldots, N_2)$ - the sequence of the corresponding right image signals for the profile $\mathbf{d}$

Dynamic Programming (DP) Stereo

Symmetric signal models with the omitted y-coordinate:

- Signal model with only random noise:

$$g_{1:i} = g_i + \nu_{1:i}; \quad g_{2:i} = g_i + \nu_{2:i}$$

- Signal model with the varying contrast:

$$g_{1:i} = a_{1,i} g_i + \nu_{1:i}; \quad g_{2:i} = a_{2:i} g_i + \nu_{2:i}$$

- Signal model with the varying contrast and offset (here, $\Delta\varphi_{i,i-1} \equiv \varphi_i - \varphi_{i-1}$):

$$\Delta g_{1:i,i-1} = e_{1:i} \Delta g_{i,i-1} + \Delta\nu_{1:i,i-1};$$

$$\Delta g_{2:i,i-1} = e_{2:i} \Delta g_{i,i-1} + \Delta\nu_{2:i,i-1}$$

Pixel-wise Signal Dissimilarity

The simplest symmetric signal model for BVPs:

- Noisy grayscale signals: $g_{1:i} = g_i + n_{1:i}$; $g_{2:i} = g_i + n_{2:i}$

$$\min_{g_i} \left\{ \max \left\{ (g_{1:i} - g_i)^2, (g_{2_i} - g_i)^2 \right\} \right\}$$
$$\Rightarrow \quad g_i = \frac{1}{2}\left(g_{1:i} + g_{2:i}\right) \Rightarrow D_i = \left(g_{1:i} - g_{2:i}\right)^2$$

- Noisy colour (RGB) signals: $D_i = \sum\limits_{c \in \{\mathrm{R,G,B}\}} (g_{c:1:i} - g_{c:2:i})^2$

The simplest dissimilarity for MVPs: $D_{\mathrm{occ}} = \mathrm{const}$

- The constant weight: an expected signal mismatch for partially occluded points observed only in one image
- A varying MVP weight depending on mismatches for the relevant BVPs might be more adequate

Dissimilarity for Varying Contrast Factors (optional)

Absolute signal dissimilarities for the BVPs ($s_i = \mathrm{B}$):

$$((|g_{1:i} - a_{1:i}g_i|;\ |g_{2:i} - a_{2:i}g_i|) : i = 1, 2, \ldots, N)$$

- $((a_{1:i}, a2 : i) : i = 1, \ldots, N)$ – sequences of the transfer factors
- Given the factors $\mathbf{a}$, an estimate for the unknown Cyclopean signal minimises the maximum of the two signal dissimilarities:

$$D_{i:a_{1:i},a_{2:i}} = \min_{g_i} \{\max \{|g_{1:i} - a_{1:i}g_i|;\ |g_{2:i} - a_{2:i}g_i|\}\}$$

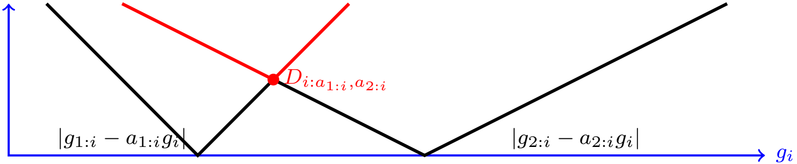

Minimax Parameter Estimates (optional)

Minimax estimate of the cyclopean signal

$$D_{i:a_{1:i},a_{2:i}} = \min_{g_i} \{\max\{|g_{1:i} - a_{1:i}g_i|;\ |g_{2:i} - a_{2:i}g_i|\}\}$$

$$\Rightarrow \quad \overbrace{g_{1:i} - a_{1:i}g_i = -g_{2:i} + a_{2:i}g_i}^{\text{Minimum by } g_i \text{ condition}}$$

$$\Rightarrow \quad g_i = \frac{g_{1:i}+g_{2:i}}{a_{1:i}+a_{2:i}}$$

$$\Rightarrow \quad \alpha_i = \frac{a_{1:i}}{a_{1:i}+a_{2:i}} \in [0,1]$$ Relative transfer factor

$$D_{i:\alpha_i} = |g_{1:i} - \alpha_i(g_{1:i} + g_{2:i})| \equiv |g_{2:i} - (1-\alpha_i)(g_{1:i} + g_{2:i})|$$

An allowable deviation range: $\alpha_i \in [\alpha_{\min}, \alpha_{\max}]$;

$$0 < \alpha_{\min} \leq 0.5 \leq \alpha_{\max} = 1 - \alpha_{\min} < 1$$

Point-wise Signal Dissimilarity (optional)

- Actual relative contrast deviation: $\alpha_i^\circ = \frac{g_{1:i}}{g_{1:i}+g_{2:i}}$
- Cyclopean signal adapted to the left signal:

$$\widehat{g}_{1:i} = \begin{cases} \alpha_{\min}(g_{1:i}+g_{2:i}) & \text{if} \quad \alpha_i^\circ < \alpha_{\min} \\ g_{1:i} & \text{if} \quad \alpha_i^\circ \in [\alpha_{\min}.\alpha_{\max}] \\ \alpha_{\max}(g_{1:i}+g_{2:i}) & \text{if} \quad \alpha_i^\circ > \alpha_{\max} \end{cases}$$

The same estimates for the Cyclopean signal and relative deviation factor are valid for the squared signal dissimilarity:

$$D_{i:a_{1:i},a_{2:i}} = \min_{g_i} \left\{ \max \left\{ (g_{1:i} - a_{1:i}g_i)^2 \,;\; (g_{2:i} - a_{2:i}g_i)^2 \right\} \right\}$$

$$\Rightarrow \quad D_i \equiv \min_{a_{1:i},a_{2:i}} D_{i:a_{1:i},a_{2:i}} = (g_{1:i} - \widehat{g}_{1:i})^2$$

Point-wise Signal Dissimilarity (optional)

For a BVP $(x_i, d_i, s_i = \mathrm{B})$ in a profile $\mathbf{d}$:

$$\alpha_i^\circ = \frac{g_{1:i}}{g_{1:i}+g_{2:i}}$$

$$\begin{aligned} D_{\mathrm{B}:i} &\equiv D_y(x_i, d_i, \mathrm{B}; x_{i-1}, d_{i-1}, s_{i-1} | \mathbf{g}_1, \mathbf{g}_2) \\ &= \begin{cases} \alpha_{\min}(g_{1:i} + g_{2:i}) - g_{1:i} & \text{if} \quad \alpha_i^\circ < \alpha_{\min} \\ 0 & \text{if} \quad \alpha_i^\circ \in [\alpha_{\min}, \alpha_{\max}] \\ g_{1:i} - \alpha_{\max}(g_{1:i} + g_{2:i}) & \text{if} \quad \alpha_i^\circ > \alpha_{\max} \end{cases} \end{aligned}$$

For an MVP $(x_i, d_i, s_i = \mathrm{M}_1 \text{ or } \mathrm{M}_2)$ in a profile $\mathbf{d}$:

a regularising constant "dissimilarity"

$$D_{\mathrm{M}:i} \equiv D_y(x_i, d_i, \mathrm{M}_k; x_{i-1}, d_{i-1}, s_{i-1} | \mathbf{g}_1, \mathbf{g}_2) = D_{\mathrm{occ}}; \quad k = 1, 2$$

to account for partially occluded points without stereo correspondence

Interdependent Contrast Factors (optional)

Floating ranges to account for spatial contrast dependencies

- "Difference" factor $\varepsilon \in [\varepsilon_{\min}, \varepsilon_{\max}]$;
  $0 < \varepsilon_{\min} \leq 1 \leq \varepsilon_{\max} = 2 - \varepsilon_{\min} < 2$
- Signal adaptation along a potentially optimal path:

$$\widehat{g}_{1:i} = \begin{cases} \widehat{g}_{1:i}^{\min} & \text{if} \quad g_{1:i} < \widehat{g}_{1:i}^{\min} \\ g_{1:i} & \text{if} \quad \widehat{g}_{1:i}^{\min} \leq g_{1:i} \leq \widehat{g}_{1:i}^{\max} \\ \widehat{g}_{1:i}^{\max} & \text{if} \quad g_{1:i} > \widehat{g}_{1:i}^{\max} \end{cases}$$

- Signal adaptation range along a potentially optimal path:

$$\begin{cases} \widehat{g}_{1:i}^{\min} & = & \min\limits_{\varepsilon} \left(\widehat{g}_{1:i_{\mathrm{pr:B}}} + \varepsilon \left(g_i - g_{i_{\mathrm{pr:B}}} \right) \right) \\ \widehat{g}_{1:i}^{\max} & = & \max\limits_{\varepsilon} \left(\widehat{g}_{1:i_{\mathrm{pr:B}}} + \varepsilon \left(g_i - g_{i_{\mathrm{pr:B}}} \right) \right) \end{cases}$$

  - $g_i = \frac{1}{2}\left(g_{1:i} + g_{2:i}\right)$ – the estimated Cyclopean signal
  - $i_{\mathrm{pr:B}}$ – the index of the BVP preceding the BVP i
  - Absolute dissimilarity $D_{\mathrm{B}:i} = |g_{1:i} - \widehat{g}_{1:i}|$

Constraints on Total Signal Dissimilarity

$$D_y(\mathbf{d}|\mathbf{g}_1,\mathbf{g}_2) = \sum_{i=1}^{N} D_{s_i:i} \equiv \sum_{i=1}^{N} D_y(x_i,d_i,s_i;x_{i-1},d_{i-1},s_{i-1}|\mathbf{g}_1,\mathbf{g}_2)$$

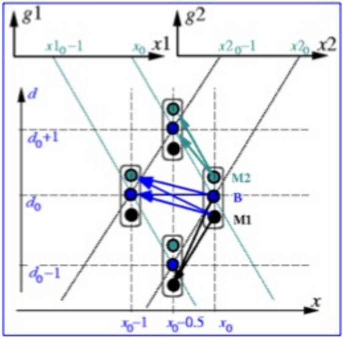

Visibility and ordering constraints on transitions between successive GPV-nodes along a continuous profile

$\Rightarrow$ Subsets Ω_{x_i,d_i,s_i} of GPV-nodes, which can precede any node $\mathbf{v}_i = (x_i,d_i,s_i)$ along a profile:

$$\Omega_{\mathbf{v}_i} \equiv \Omega_{x_i,d_i,s_i} =$$

$$\begin{cases} (x_i-\frac{1}{2},d_i+1,\mathrm{B});(x_i-\frac{1}{2},d_i+1,\mathrm{M}_2) & \text{if} \quad s_i=\mathrm{M}_2 \\ (x_i-\frac{1}{2},d_i-1,\mathrm{M}_1);(x_i-1,d_i,\mathrm{B});(x_i-1,d_i,\mathrm{M}_2 & \text{if} \quad s_i=\mathrm{B} \\ (x_i-\frac{1}{2},d_i-1,\mathrm{M}_1);(x_i-1,d_i,\mathrm{B});(x_i-1,d_i,\mathrm{M}_2 & \text{if} \quad s_i=\mathrm{M}_1 \end{cases}$$

DP to Minimise Signal Dissimilarity

Finds the profile for the globally minimal dissimilarity $D_y(\mathbf{d}|\mathbf{g}_1, \mathbf{g}_2)$

- DP exhausts all the profiles $\mathbf{d}$, which are possible in a GPV under the constrained transitions $\mathbf{\Omega}$
  - Forward pass along the x-axis of a GPV to find the minimum dissimilarity $D^* = \min\limits_{\mathbf{d}} D_y(\mathbf{d}|\mathbf{g}_1, \mathbf{g}_2)$
  - Backward pass to get the profile $\mathbf{d}^* = \arg\min\limits_{\mathbf{d}} D_y(\mathbf{d}|\mathbf{g}_1, \mathbf{g}_2)$
- At any current location, x_i, all GPV-nodes $\mathbf{v}_i = (x_i, d_i, s_i)$ are examined in order to calculate and store current **potentially optimal** total dissimilarities $D_{\mathrm{po}:i}(\mathbf{v}_i)$
  - Potential optimality: the stored dissimilarity is optimal if the node belongs to the globally optimal solution
  - $D_{\mathrm{po}:i}(\mathbf{v}_i)$ – the **minimal total signal dissimilarity** for the potentially optimal backward path from $\mathbf{v}_i = (x_i, d_i, s_i)$
- For each node, a potentially optimal backward transition $B_i(\mathbf{v}_i)$ to one of the preceding nodes $\mathbf{v}_{i-1}$ in $\mathbf{\Omega}$ is stored

Dynamic Programming (DP)

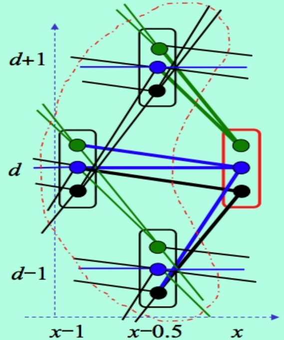

Recurrent DP computation:

$$
\begin{cases}
D_{\mathrm{po}:i}(\mathbf{v}_i) \\
\quad = \min\limits_{\mathbf{v}\in\Omega_{\mathbf{v}_i}} \{D_y(\mathbf{v}_i, \mathbf{v}|\mathbf{g}_1, \mathbf{g}_2) + D_{\mathrm{po}:i-1}(\mathbf{v})\} \\
\quad = D_y(\mathbf{v}_i, \mathbf{v}^*_{i-1}|\mathbf{g}_1, \mathbf{g}_2) + D_{\mathrm{po}:i-1}(\mathbf{v}^*_{i-1}) \\
\\
B_i(\mathbf{v}_i) \equiv \mathbf{v}^*_{i-1} \\
\quad = \arg\min\limits_{\mathbf{v}\in\Omega_{\mathbf{v}_i}} \{D_y(\mathbf{v}_i, \mathbf{v}|\mathbf{g}_1, \mathbf{g}_2) + D_{\mathrm{po}:i-1}(\mathbf{v})\}
\end{cases}
$$

$$
\begin{aligned}
D_{\mathrm{po}:i}(x,d,\mathrm{M}_2) &= D_{\mathrm{occ}} + \min\{D_{\mathrm{po}:i-1}(x-0.5,d+1,\mathrm{M}_2); D_{\mathrm{po}}(x-0.5,d+1,\mathrm{B})\} \\
D_{\mathrm{po}:i}(x,d,\mathrm{B}) &= \min\{D_{\mathrm{po}:i-1}(x-1,d,\mathrm{M}_2) + D_y(x,d;x_{\mathrm{prB:M}_2},d_{\mathrm{prB:M}_2}); \\
&\qquad D_{\mathrm{po}:i-1}(x-1,d+1,\mathrm{B}) + D_y(x,d;x-1,d); \\
&\qquad D_{\mathrm{po}:i-1}(x-0.5,d-1,\mathrm{M}_1) + D_y(x,d;x_{\mathrm{prB:M}_1},d_{\mathrm{prB:M}_1})\} \\
D_{\mathrm{po}:i}(x,d,\mathrm{M}_1) &= D_{\mathrm{occ}} + \min\{D_{\mathrm{po}:i-1}(x-0.5,d-1,\mathrm{M}_1); D_{\mathrm{po}:i-1}(x-1,d,\mathrm{B})\}
\end{aligned}
$$

Basic Recurrent DP Computation

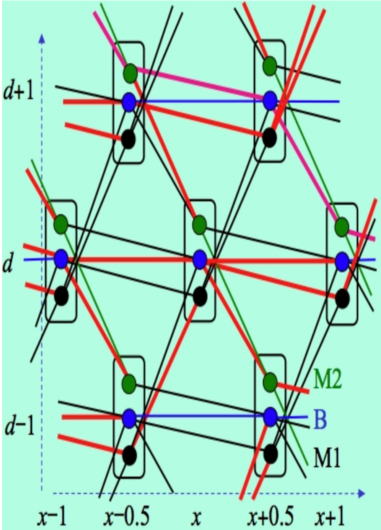

- $B_i(x_i, d_i, s_i)$ – a potentially optimal backward transition from a GPV-node
- Optimal profile: a sequence of potentially optimal backward transitions such that
  1. Begins (i.e. the profile ends) at point $x_N = x_{\max}$ or $x_{\max} - 0.5$ and
  2. Minimises the total dissimilarity $D_{\text{po}}(x_N, d_N, s_N)$ for all the GPV nodes x_N, d_N, s_N
- $D_{\text{po}}(x^*_N, d^*_N, s^*_N)$ – the minimal total signal dissimilarity

Basic Recurrent DP Computation

After the forward pass through a given x-coordinate range $[x_{\min}, x_{\max}]$, the optimal profile is recovered by the backward pass through the stored potentially optimal backward transitions:

$$D_y(\mathbf{d}^*|\mathbf{g}_1, \mathbf{g}_2) = D_{\mathrm{po}:N}(x_N^*, d_N^*, s_N^*)$$

$$(x_N^*, d_N^*, s_N^*) = \underset{\substack{x_N \in \{x_{\max} - 0.5, x_{\max}\} \\ d_N \in [d_{\min}, d_{\max}] \\ s_N \in \{\mathrm{M}_1, \mathrm{B}, \mathrm{M}_2\}}}{\arg\min} \{D_{\mathrm{po}:N}(x_N, y_N, s_N)\}$$

$$(x_{i-1}^*, d_{i-1}^*, s_{i-1}^*) = B_i\left(x_i^*, d_i^*, s_i^*\right); \quad i = N, N-1, \ldots, 2$$

Viterbi DP Algorithm: Probabilistic Model

DP search for the most likely sequence of unobserved (hidden) states from a sequence of observations (signals)

- Each observation corresponds to and depends on exactly one hidden state
- Hidden states are produced by a first-order Markov model:
  - Set of the hidden states $\mathbb{H} = \{h_1, \ldots, h_n\}$
  - Transitional probabilities $P_{\mathrm{t}}(h_i | h_j) \,:\, i, j \in \{1, \ldots, n\}$
- Given the states, the observations are statistically independent
  - Set of the signals $\mathbb{Q} = \{q_1, \ldots, q_m\}$
  - Observational probabilities $P_{\mathrm{o}}(q_j | h_i) \,:\, q_j \in \mathbb{Q};\; h_i \in \mathbb{H}$

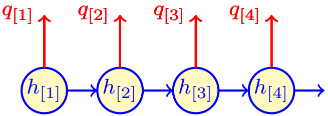

Log-likelihood of a sequence of states $\mathbf{h} = h_{[1]} h_{[2]} \ldots h_{[K]}$, given a sequence of signals $\mathbf{q} = q_{[1]} q_{[2]} \ldots q_{[K]}$:

$$L(\mathbf{h}|\mathbf{q}) \equiv \log \Pr(\mathbf{h}|\mathbf{q}) \Rightarrow \Pr(\mathbf{h}|\mathbf{q}) \sim \Pr(\mathbf{h}, \mathbf{q}) = \Pr_{\mathrm{t}}(\mathbf{h}) \Pr_{\mathrm{o}}(\mathbf{q}|\mathbf{h})$$

http://en.wikipedia.org/wiki/Viterbi_algorithm

Maximum (Log-)Likelihood $\mathbf{s}^* = \arg\max_{\mathbf{s}\in\mathbb{S}^K} L(\mathbf{s}|\mathbf{q})$

1. $\mathbf{h} = h_{[1]}h_{[2]}\dots h_{[K]}$ – a hidden (unobserved) Markov chain of states at steps $k = 1,\dots,K$ with joint probability
$$\Pr_{\mathrm{t}}(\mathbf{h}) = \pi\left(h_{[1]}\right)\prod_{k=2}^{K} P_{\mathrm{t}}\left(h_{[k]}|h_{[k-1]}\right)$$
   - $\pi(h)$ – prior probability of state $h \in \mathbb{H}$ at step $k = 1$
   - $P_{\mathrm{t}}(h|h')$ – probability of transition from a state h' to the next one, h
2. $\mathbf{q} = q_{[1]}q_{[2]}\dots q_{[K]}$ – an observed sequence of conditionally independent signals with probability $\Pr(\mathbf{q}|\mathbf{h}) = \prod_{k=1}^{K} P_{\mathrm{o}}\left(q_{[k]}|h_{[k]}\right)$
   - $P_{\mathrm{o}}(q|h)$ – probability of observing $q \in \mathbb{Q}$ in state $h \in \mathbb{H}$ at step k

$$\mathbf{h}^* = \arg\max_{\mathbf{h}\in\mathbb{H}^K} \sum_{k=1}^{K} \left(\psi_k\left(h_{[k]}\right) + \varphi\left(h_{[k]}|h_{[k-1]}\right)\right)$$

$$\psi_k(h) = \begin{cases} \log\pi(h) + \log P_{\mathrm{o}}\left(q_{[k]}|h\right) & k = 1;\ h \in \mathbb{H} \\ \log P_{\mathrm{o}}\left(q_{[k]}|h\right) & k > 1;\ h \in \mathbb{H} \end{cases}$$

$$\varphi(h|h') = \begin{cases} 0 & k = 1;\ h \in \mathbb{H} \\ \log P_{\mathrm{t}}(h|h') & k > 1;\ h \in \mathbb{H} \end{cases}$$

Graphical Model for $\mathbb{H} = \{a, b, c\}$ and Observed q

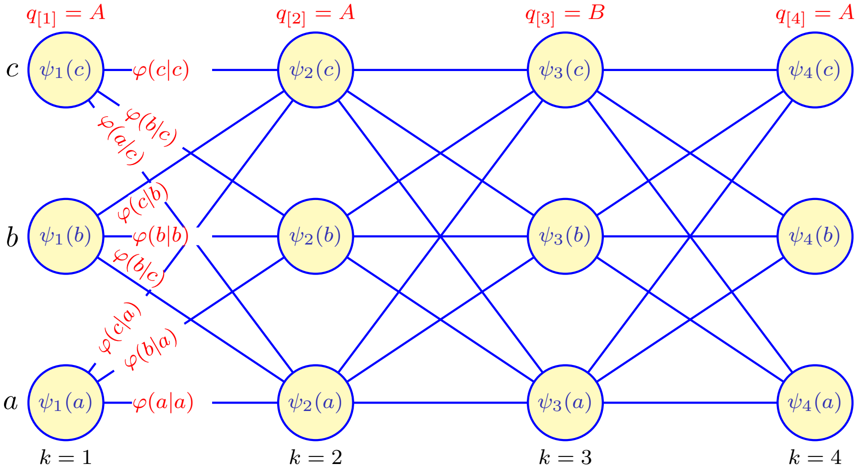

Maximum (Log-)Likelihood via Dynamic Programming

Viterbi DP algorithm:

1. Initialisation: $k = 1$; $\Phi_1(h_{[1]}) = \psi_1(h_{[1]})$ for all $h_{[1]} \in \mathbb{H}$
2. **Forward pass** for $k = 2, \ldots, K$ and all $h_{[k]} \in \mathbb{H}$:

$$\Phi_k\left(h_{[k]}\right) = \psi_k(h_{[k]}) + \max_{h_{[k-1]} \in \mathbb{H}} \left\{\varphi(h_{[k]}|h_{[k-1]}) + \Phi_{k-1}\left(h_{[k-1]}\right)\right\}$$

$$B_k\left(h_{[k]}\right) = \arg\max_{h_{[k-1]} \in \mathbb{H}} \left\{\varphi(h_{[k]}|h_{[k-1]}) + \Phi_{k-1}\left(h_{[k-1]}\right)\right\}$$

3. $k = K$: the maximum log-likelihood $\Phi_K(h^*_{[K]})$ in the state

$$h^*_{[K]} = \arg\max_{h_{[k]} \in \mathbb{S}} \Phi_K(h_{[K]})$$

4. **Backward pass** for $k = K, \ldots, 2$: $h^*_{[k-1]} = B_k\left(h^*_{[k]}\right)$

Example: $\mathbb{H} = \{a, b\}$; $\mathbb{Q} = \{A, B\}$; $\mathbf{q} = AABAB$

$$\varphi(h|h') = \begin{cases} 8 & h = h' \\ 1 & h \neq h' \end{cases}; \quad \text{for} \quad h, h' \in \mathbb{H}$$

$$\psi_k(h) = \begin{cases} 6 & v|s \in \{A|a, B|b\} \\ 2 & \text{otherwise} \end{cases}; \quad \text{for} \quad h \in \mathbb{H}$$

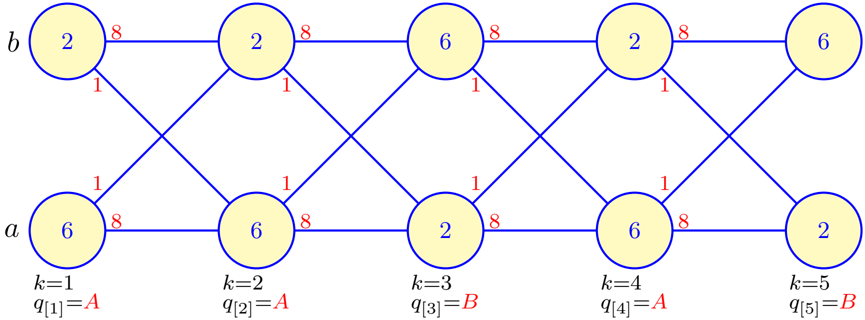

Example: $\mathbb{H} = \{a, b\}$; $\mathbb{Q} = \{A, B\}$; $\mathbf{q} = AABAB$

$$\varphi(h|h') = \begin{cases} 8 & h = h' \\ 1 & h \neq h' \end{cases}; \quad \text{for} \quad h, h' \in \mathbb{H}$$

$$\psi_k(h) = \begin{cases} 6 & v|s \in \{A|a, B|b\} \\ 2 & \text{otherwise} \end{cases}; \quad \text{for} \quad h \in \mathbb{H}$$

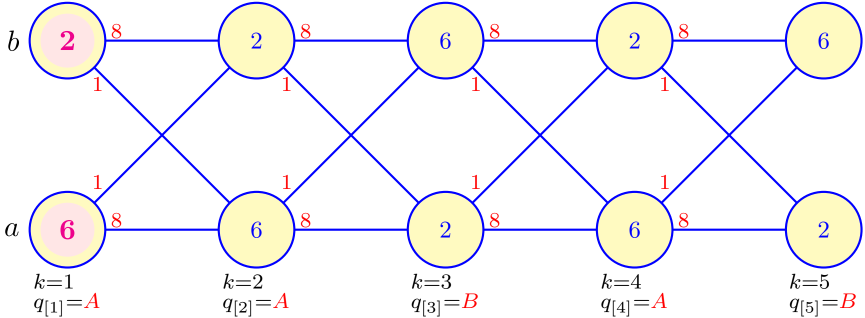

Step $k = 1$: Initialisation

Example: $\mathbb{H} = \{a, b\}$; $\mathbb{Q} = \{A, B\}$; $\mathbf{q} = AABAB$

$$\varphi(h|h') = \begin{cases} 8 & h = h' \\ 1 & h \neq h' \end{cases} ; \quad \text{for} \quad h, h' \in \mathbb{H}$$

$$\psi_k(h) = \begin{cases} 6 & v|s \in \{A|a, B|b\} \\ 2 & \text{otherwise} \end{cases} ; \quad \text{for} \quad h \in \mathbb{H}$$

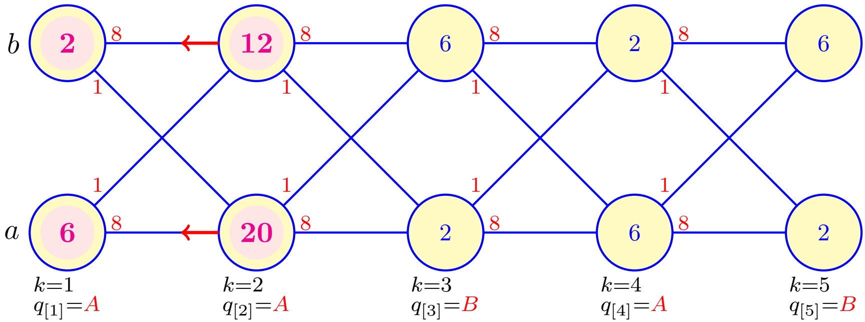

Step $k = 2$

Example: $\mathbb{H} = \{a, b\}$; $\mathbb{Q} = \{A, B\}$; $\mathbf{q} = AABAB$

$$\varphi(h|h') = \begin{cases} 8 & h = h' \\ 1 & h \neq h' \end{cases}; \quad \text{for} \quad h, h' \in \mathbb{H}$$

$$\psi_k(h) = \begin{cases} 6 & v|s \in \{A|a, B|b\} \\ 2 & \text{otherwise} \end{cases}; \quad \text{for} \quad h \in \mathbb{H}$$

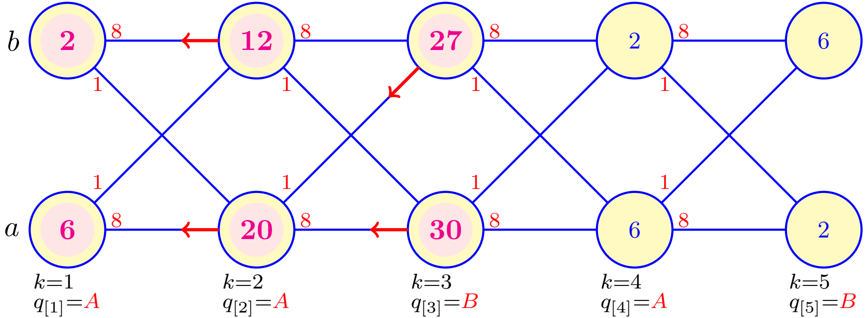

Step $k = 3$

Example: $\mathbb{H} = \{a, b\}$; $\mathbb{Q} = \{A, B\}$; $\mathbf{q} = AABAB$

$$\varphi(h|h') = \begin{cases} 8 & h = h' \\ 1 & h \neq h' \end{cases}; \quad \text{for} \quad h, h' \in \mathbb{H}$$

$$\psi_k(h) = \begin{cases} 6 & v|s \in \{A|a, B|b\} \\ 2 & \text{otherwise} \end{cases}; \quad \text{for} \quad h \in \mathbb{H}$$

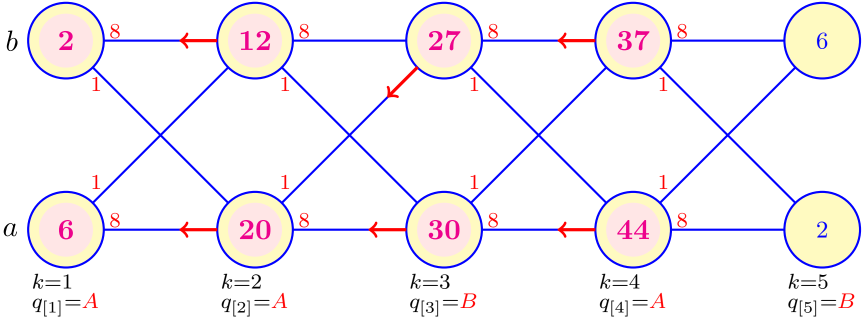

Step $k = 4$

Example: $\mathbb{H} = \{a, b\}$; $\mathbb{Q} = \{A, B\}$; $\mathbf{q} = AABAB$

$$\varphi(h|h') = \begin{cases} 8 & h = h' \\ 1 & h \neq h' \end{cases}; \quad \text{for} \quad h, h' \in \mathbb{H}$$

$$\psi_k(h) = \begin{cases} 6 & v|s \in \{A|a, B|b\} \\ 2 & \text{otherwise} \end{cases}; \quad \text{for} \quad h \in \mathbb{H}$$

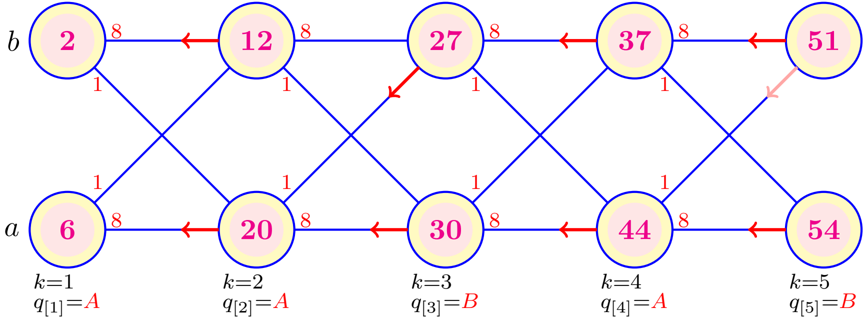

Step $k = 5$

Example: $\mathbb{H} = \{a, b\}$; $\mathbb{Q} = \{A, B\}$; $\mathbf{q} = AABAB$

$$\varphi(h|h') = \begin{cases} 8 & h = h' \\ 1 & h \neq h' \end{cases}; \quad \text{for} \quad h, h' \in \mathbb{H}$$

$$\psi_k(h) = \begin{cases} 6 & v|s \in \{A|a, B|b\} \\ 2 & \text{otherwise} \end{cases}; \quad \text{for} \quad h \in \mathbb{H}$$

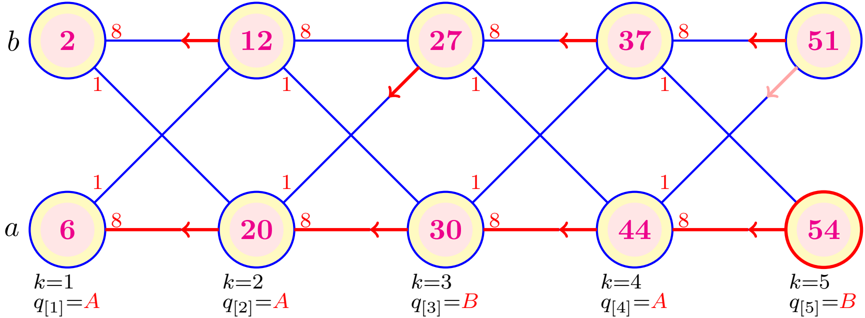

Backtracking: $\mathbf{h}^* = aaaaa$

DP stereo matching by energy minimization

- Markov models of surfaces / signals $\Longrightarrow$ Energy function $E(\mathbf{d}|\mathbf{g}_1, \mathbf{g}_2)$ as a matching score:

$$E(\mathbf{d}|\mathbf{g}_1, \mathbf{g}_2) = \sum_{i=1}^{n} \phi(x_i, d_i, s_i; x_{i-1}, d_{i-1}, s_{i-1}|\mathbf{g}_1, \mathbf{g}_2)$$

$$(x_{i-1}, d_{i-1}, s_{i-1}) \in \Omega(x_i, d_i, s_i) =$$

$$\begin{cases} \{(x_i - 0.5, d_i + 1, B); (x_i - 0.5, d_i + 1, M2)\} & \text{if} \quad s_i = M2 \\ \{(x_i - 0.5, d_i - 1, M1); (x_i - 1, d_i, B), (x_i - 1, d_i, M2)\} & \text{if} \quad s_i = B \\ \{(x_i - 0.5, d_i - 1, M1); (x_i - 1, d_i, B), (x_i - 1, d_i, M2)\} & \text{if} \quad s_i = M1 \end{cases}$$

  - Energy $=$ signal dissimilarity $+$ surface continuity $+$ surface smoothness $+$ occlusions $+ \ldots$
  - SDPS $\rightarrow$ energy accounts for different contrast and offset deviations along scanlines (approximate minimization by DP)
- DP: profile-wise global energy minima

Recurrent DP framework: Summary

Disparity "tube": $d_{\min} \leq d_i \leq d_{\max}$; $s_i \in \mathbb{S}$, for $i = 1, \ldots, n$

1. **Forward pass** along the x-axis of GPV
   - At each x_i, compute and store for each node $\mathbf{v}_i = (x_i, d_i, s_i)$
     - Potentially optimum partial total energy $E_i(\mathbf{v}_i)$
     - Potentially optimum backward transition $B_i(\mathbf{v}_i)$

$$
\begin{cases}
E_i(\mathbf{v}_i) & = \phi(\mathbf{v}_i; \mathbf{v}^{\circ}_{i-1} | \mathbf{g}_1, \mathbf{g}_2) + E_{i-1}(\mathbf{v}^{\circ}_{i-1}) \\
B_i(\mathbf{v}_i) \equiv \mathbf{v}^{\circ}_{i-1} & = \arg\min\limits_{\mathbf{v}_{i-1}} \{\phi(\mathbf{v}_i; \mathbf{v}_{i-1} | \mathbf{g}_1, \mathbf{g}_2) + E_{i-1}(\mathbf{v}_{i-1})\}
\end{cases}
$$

2. **Backward pass** along the x-axis of GPV
   - Optimal profile: a sequence of potentially optimum transitions:

$$
\begin{aligned}
\mathbf{v}^*_n &= \arg\min_{\mathbf{v}_n} E_n(\mathbf{v}_n) \\
\mathbf{v}^*_{i-1} &= B_n(\mathbf{v}^*_i);\; i = n, n-1, \ldots, 2
\end{aligned}
$$

Computational complexity $O(N\Delta)$ where $\Delta = d_{\max} - d_{\min} + 1$

Symmetric Dynamic Programming Stereo: An Example