A Bayesian Framework for Statistical, Multi-Modal Sensor Fusion

LTC Mick Smith

Department of Mathematical Sciences, United States Military Academy

Anuj Srivastava

Co-Author & Ph.D. Major Professor, Florida State University

20 October 2004

Army Conference on Applied Statistics

Acknowledgements

J. Sethuraman, FSU; Gary Krahn, USMA

ARO DAAD 19-99-1-0267; NMA 201-01-2010; USMA AN40

Outline of Presentation

- Bayesian Sensor Fusion: Definition & Motivation
- Statistical Models for Scene & Sensors
- Implementing a Metropolis-Hastings Algorithm
- Example Markov Chain with Application to Tactical Questions
- Directions for Further Work

Definition

Bayesian sensor fusion is a methodology for scene inference that:

- Formulates a prior distribution for the scene.
- Constructs probability models or likelihood functions for sensor data conditioned on the scene.
- Conducts unified inference about the scene using the posterior distibution of the scene given the sensor data.

Motivation

Our Bayesian methodology for sensor fusion:

- Recognizes that sensors are partial observers of the scene.
- Exploits the complementary nature of the multi-sensor suite by merging joint probabilities.
- Affords inclusion of the commander's estimate by way of a prior distribution.

Sensors Detect Different Aspects of the Scene

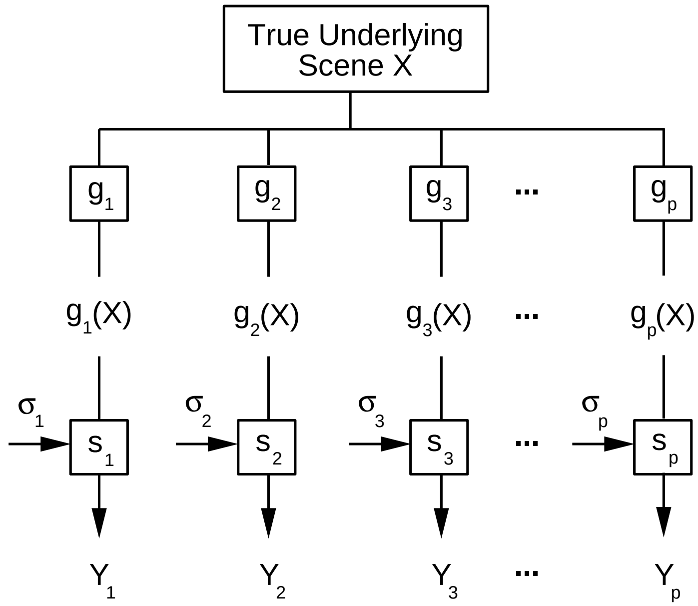

- Sensors $s_1, \ldots, s_p$ observe the projections g_i of the scene X and generate data vectors $Y_1, \ldots, Y_p$.
- Data vectors are corrupted by sensor noise σ_i.

Simulated Battlefield Scene

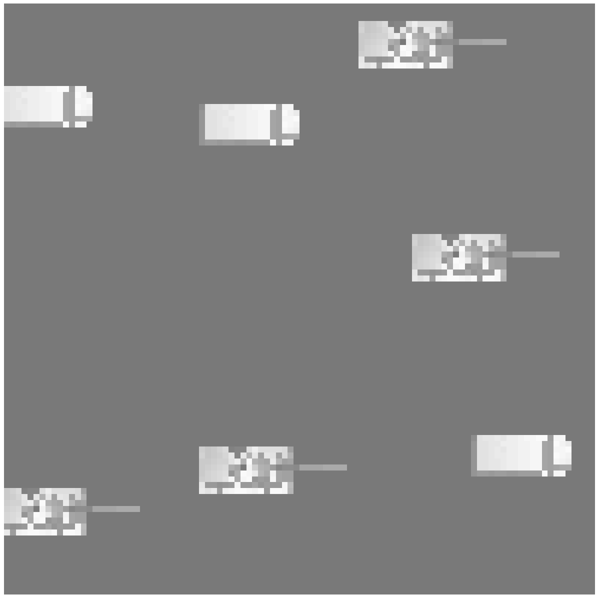

Types of combat vehicles limited to tanks and trucks with fixed orientations.

Mathematical Description of the Scene

- We take the scene X to be a point in the space $\mathcal{X} = \bigcup_{n=0}^{\infty} (\mathfrak{D} \times \mathcal{A})^n$ where:
  - $\mathfrak{D} \subset \mathbb{R}^2$ is a battlefield region of interest;
  - $\mathcal{A} = \{\alpha_1, \ldots, \alpha_M\}$ is a set of M possible target types;
  - n is the number of targets present.
- After discretizing & truncating $\mathcal{X}$, a typical state in our Markov chain is a matrix with columns corresponding to target vehicles:

$$X = \begin{bmatrix} r_1 & r_2 & \cdots & r_n \\ c_1 & c_2 & \cdots & c_n \\ \alpha_1 & \alpha_2 & \cdots & \alpha_n \end{bmatrix}.$$

Prior Distribution: The Commander's Experience

- X is a realization of a marked homogeneous Poisson spatial point process:
  - $N \sim \mathrm{Poisson}(\lambda|\mathcal{D}|)$ for some $\lambda > 0$.
  - Given $\{N = n\}$, let the locations $q_1, \ldots, q_n$ of targets be distributed independently and uniformly in $\mathcal{D}$.
- More realistic prior distributions ν_0 on $\mathcal{X}$ are desirable.

Multi-Modal, Multi-Sensor Environment

Table 1: Sensors Considered in the Paper

Label	*Sensor*	*Nature of Operation*	*Detected Aspects*	*Data Output* (Y_i)
s_1	Infrared Camera	Low-Resolution Imager	Target Location & ID	2D Image Array
s_2	Acoustic Array	Audio Signal Receiver	Direction Only; No ID	1D Signal Vector
s_3	Scout	Human Vision	Rough Location; ID	Categorical Data
s_4	Seismic Array	Wave Receiver	Rough Location; Partial ID	Local Detection

- Likelihood functions for s_1 and s_2 are adopted from published research.
- Probability models for the scout's spot report and for the seismic sensor array are newly proposed in this work & are described on the following slides.

Scout's Spot Report

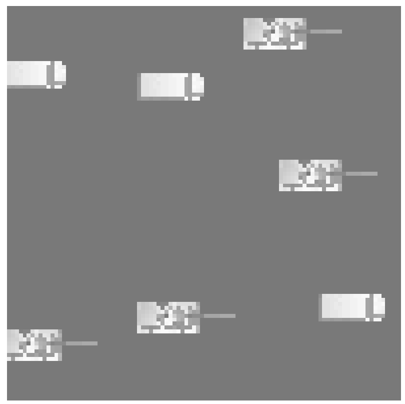

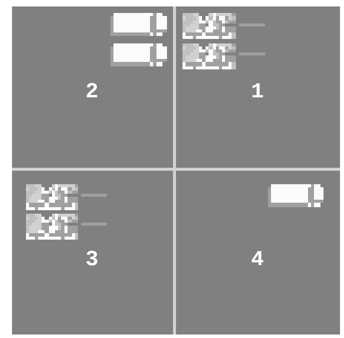

- Suppose that a scout reports target counts by quadrant & by type.
- To construct a likelihood function conditioned on the scene, we imagine the scout asking & answering three questions:
- How many targets? Where are they? What are they?

Scout Likelihood: How Many Targets? Where?

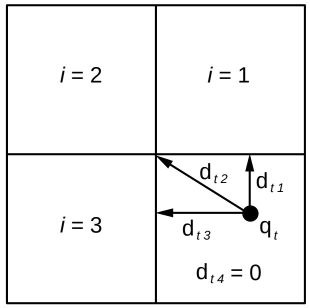

- Number of targets observed: $N_S \sim$ discretized Gaussian with mean n.
- Given $\{N_S = n_0\}$, require that the spot report Y_3 satisfies $\sum_{j=1}^{4M} (Y_3)_j = n_0$.
- Let d_{ti} denote distance from location q_t to quadrant i.
- Define the probability that the scout reports quadrant i as location for target t:

$$\tilde{p}_{ti} = \frac{\exp(-d_{ti}/a)}{\sum_{j=1}^{4} \exp(-d_{tj}/a)}, \quad i = 1, 2, 3, 4; \quad a > 0.$$

Scout Likelihood: What are They?

- Let $I\{\alpha_t = j\}$ indicate that α_j is the identity of target t.
- Define generalized Bernoulli parameters $\{p_{tj}\}_{j=1}^{4M}$ for the t^{th} target:

$$p_{tj} \;=\; (1-\sigma_3)\,\tilde{p}_{ti}\,I\{\alpha_t = j_i\} \;+\; \frac{\sigma_3}{M-1}\,\tilde{p}_{ti}\,I\{\alpha_t \neq j_i\}.$$

- Parameter σ_3 is the classification error.
- In words, the scout correctly reports the target type w.p. $(1-\sigma_3)$ and he is equally likely to report any of the incorrect target types.
- Likelihood function:

$$L_3(Y_3 \,|\, X) = \frac{\mathcal{G}(n_0)}{n_0\,!} \sum_{T_o \in \mathcal{P}(T)} \prod_{t=(T_o)_1}^{(T_o)_{n_0}} {p_{t1}}^{(Y_3)_1}\, {p_{t2}}^{(Y_3)_2} \,\cdots\, {p_{t,4M}}^{(Y_3)_{4M}}.$$

Seismic Sensor Array

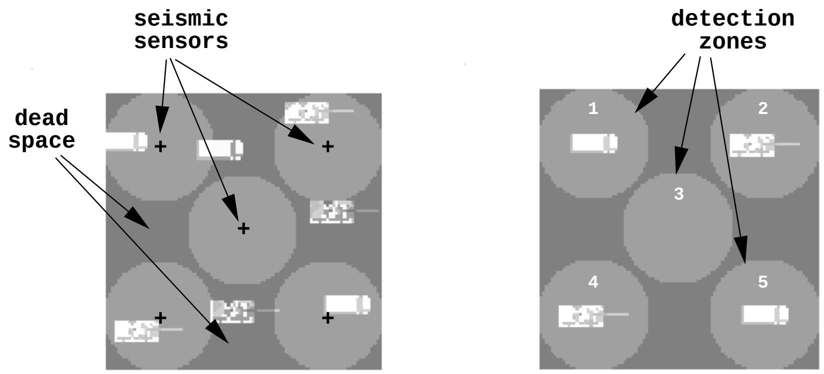

- Seismic sensor detects & classifies targets but does not count them.
- Sensors are deployed in an array; each has a known detection-zone radius.
- Array may admit gaps of “dead space.”

Seismic Sensor Behavior

- Case 1: Zone j is devoid of targets:

$$P\{(Y_4)_j = y \,|\, n_{1j} = \cdots = n_{Mj} = 0\} = \begin{cases} 1 - \sigma_4, & y = \alpha_\emptyset; \\ \frac{\sigma_4}{M}, & y \in \mathcal{A}; \\ 0, & \text{otherwise.} \end{cases}$$

- Case 2: Zone j contains exactly 1 target type:

$$P\{(Y_4)_j = y \,|\, n_{ij} > 0 \text{ for } i = i_0 \text{ only}\} = \begin{cases} \frac{\sigma_4}{4}, & y = \alpha_\emptyset; \\ 1 - \sigma_4, & y = \alpha_{i_0}; \\ \frac{3\sigma_4}{4(M-1)}, & y \in \mathcal{A} \backslash \{\alpha_{i_0}\}; \\ 0, & \text{otherwise.} \end{cases}$$

Seismic Sensor Behavior

- Case 3: Sensor must "decide" among competing target types in Zone j:

$$P\big\{(Y_4)_j = y \,\big|\, 2 \leq |\{i : n_{ij} > 0\}| \,\big\} =$$

$$(1-\sigma_4)\frac{\sum_{t=1}^{n_{\cdot j}} I\{\alpha_t = i\} e^{-d_t/a}}{\sum_{t=1}^{n_{\cdot j}} e^{-d_t/a}}, \quad y = \alpha_i \in \mathcal{A}.$$

- Likelihood function:

$$L_4(Y_4 \,|\, X) = \prod_{j=1}^{k} P\big\{(Y_4)_j = y \,|\, X\big\}.$$

Posterior Distribution of the Scene

- We assume that, given the scene X, the sensor data vectors Y_i are conditionally independent.
- Applying Bayes' rule, we obtain an expression for the posterior distribution:

$$\nu(X) \equiv \nu(X \,|\, Y_1, \ldots, Y_p) \propto L_1(Y_1 \,|\, X) \cdots L_p(Y_p \,|\, X)\, \nu_0(X).$$

- To conduct inference, we must generate samples from ν.
- We do this via Metropolis-Hastings: we construct an ergodic Markov chain on $\mathcal{X}$ having stationary distibution ν.

Metropolis-Hastings Algorithm

Given the current state $X^{(t)} \in \mathcal{X}$,

1. Generate $Y_t \sim G(y|X^{(t)})$. G is called the proposal distribution.
2. Set $X^{(t+1)} = \begin{cases} Y_t & \text{w.p. } \gamma(X^{(t)}, Y_t); \\ X^{(t)} & \text{w.p. } 1 - \gamma(X^{(t)}, Y_t), \end{cases}$ where

$$\gamma(x, y) = \min\left\{1, \frac{\nu(y)\, G(x|y)}{\nu(x)\, G(y|x)}\right\}.$$

For a large class of proposal distributions G and for $X^{(1)} \sim F$ where F is an arbitrary probability distribution on $\mathcal{X}$, this algorithm is known to generate a Markov chain with unique stationary distribution ν.

Proposal Distribution: “Simple Moves”

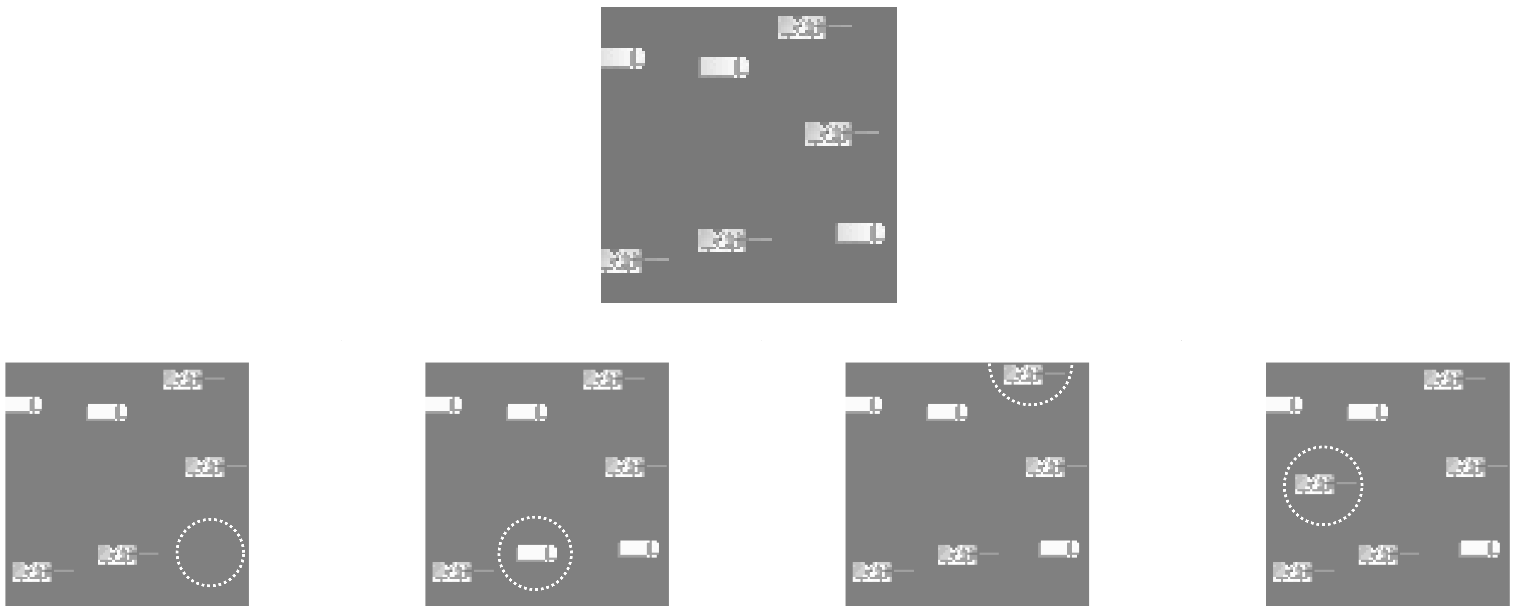

Death Change ID Adjust Birth

$G(y \mid X^{(t)})$ nominates a state from one of the indicated *neighborhoods*.

Example Scene & Sensor Data

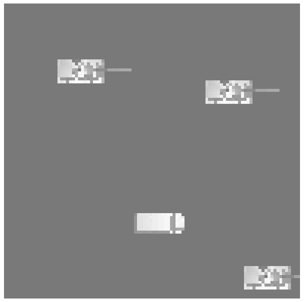

Original Scene

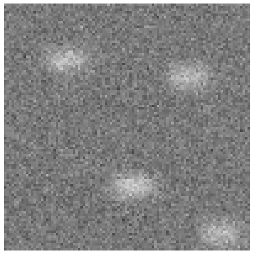

Infrared Image

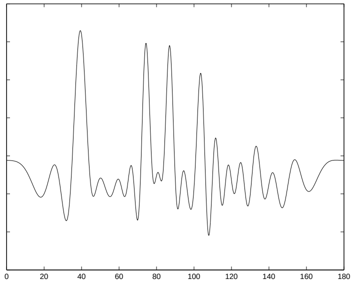

Acoustic Signal

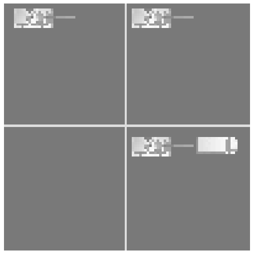

Scout’s Report

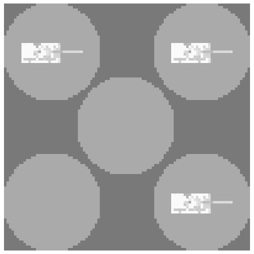

Seismic Data

Evolution of the Markov Chain

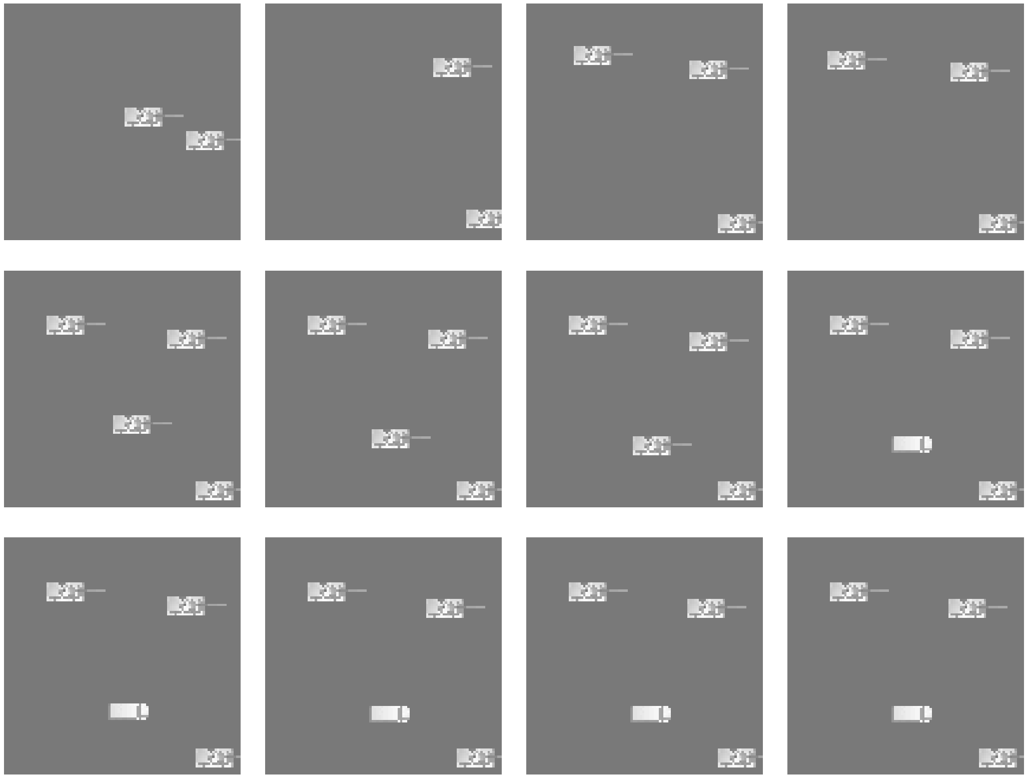

Solution   Original

Answering Tactical Questions

- Discard the first B states (*burn-in period*) and retain, for purposes of inference,

$$\{X^{(B+1)}, X^{(B+2)}, \dots, X^{(B+R)}\}.$$

- Typical commander's question: *How many tanks are out there?*
- Let $A = \{X \in \mathcal{X} : \text{number of enemy tanks} \geq k\}$.
- The ergodic property of our Markov chain allows us to estimate the posterior probability of this event by $\frac{1}{R} \sum_{j=1}^{R} \mathbf{1}_A(X_j)$.
- If the commander requires this probability to be at least 0.95 (say), we may construct a simple rule based on our sample:

$$\frac{1}{R} \sum_{j=1}^{R} \mathbf{1}_A(X_j) \geq 0.95 \quad \Rightarrow \quad \text{Respond.}$$

Directions for Future Work

- Models for Additional Sensors — e.g., Magnetic Sensors
- Improved M-H Proposal Distribution G to Increase Acceptance Rate
- Designed Experiment to Estimate Parameters for Scout Likelihood
- Validation Using Real Data
- Recoding the Algorithm to Achieve Fast Execution

Likelihood for IR Image

$$L_1(Y_1 \mid X) = \prod_{i=1}^{rc} \frac{\big((I_0 * h)(z_i)\big)^{Y_1(z_i)} \, \mathrm{e}^{-(I_0 * h)(z_i)}}{Y_1(z_i)!}.$$

$$\hat{L}_1(Y_1 \mid X) = \frac{1}{Z} \exp\Big(-\frac{1}{2\sigma_1^2} \, \|Y_1 - I_0 * h\|_F^2\Big).$$

Likelihood for Acoustic Signal

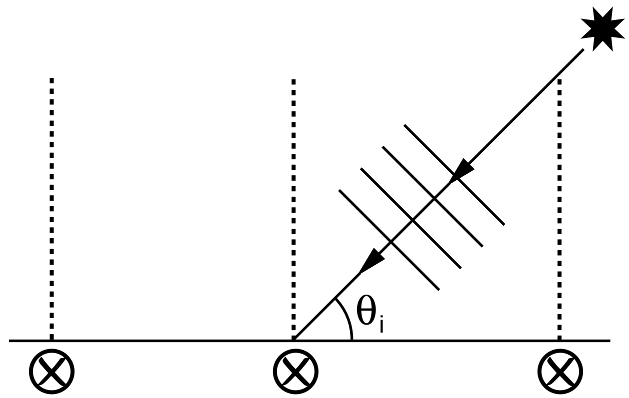

$$L_2(Y_2 \mid X) = \frac{1}{Z} \exp\Big(-\frac{1}{\sigma_2^2} \|Y_2 - \sum_{i=1}^{n} d(\theta_i)\|^2\Big).$$

$$d(\theta_i) = \big[\, 1, \exp\{-j\pi\cos(\theta_i)\}, \ldots, \exp\{-(m-1)j\pi\cos(\theta_i)\} \,\big]' \quad (j^2 = -1).$$

Neighboring States for DEATH and CHANGE ID

$$\mathcal{N}_D(X^{(t)}) = \begin{cases} \{X_{-j}^{(t)} : j = 1, \ldots, n\}, & \text{if } \|X^{(t)}\| \geq 1; \\ \{X^{(t)}\}, & \text{if } \|X^{(t)}\| = 0, \end{cases}$$

where $X_{-j}^{(t)}$ denotes the matrix $X^{(t)}$ after removing column j.

$$\mathcal{N}_C(X^{(t)}) = \begin{cases} \{X_{\Delta j}^{(t)} : j = 1, \ldots, n\}, & \text{if } \|X^{(t)}\| \geq 1; \\ \{X^{(t)}\}, & \text{if } \|X^{(t)}\| = 0, \end{cases}$$

where $X_{\Delta j}^{(t)}$ denotes the matrix $X^{(t)}$ after changing the identity component of column j.

Neighboring States for ADJUST and BIRTH

$$\mathcal{N}_A(X^{(t)}) = \begin{cases} \{X^{(t)}_{\oplus j} : j = 1, \ldots, n\}, & \text{if } \|X^{(t)}\| \geq 1; \\ \{X^{(t)}\}, & \text{if } \|X^{(t)}\| = 0, \end{cases}$$

where each $X^{(t)}_{\oplus j}$ denotes as many as eight perturbations to the location components of $X^{(t)}_j$.

$$\mathcal{N}_B(X^{(t)}) = \{X^{(t)}_\tau : \tau \in (\mathfrak{D} \times \mathcal{A}) \setminus T_{X^{(t)}}\},$$

where $X^{(t)}_\tau$ is the augmentation of the matrix $X^{(t)}$ by one additional column τ corresponding to any "legal" target not already present: $\|X^{(t)}_\tau\| = \|X^{(t)}\| + 1$.

Proposal Distribution for Metropolis-Hastings

$$
\begin{aligned}
G(y \,|\, X^{(t)}) \;\; = \;\; & w_D \, \frac{1}{|\mathcal{N}_D(X^{(t)})|} \, \mathbf{1}_{\mathcal{N}_D(X^{(t)})}(y) \\
& + \; w_C \, \frac{1}{|\mathcal{N}_C(X^{(t)})|} \, \mathbf{1}_{\mathcal{N}_C(X^{(t)})}(y) \; + \; w_A \, \frac{1}{|\mathcal{N}_A(X^{(t)})|} \, \mathbf{1}_{\mathcal{N}_A(X^{(t)})}(y) \\
& + \; w_B \, \mathbb{P}_{T_{X^{(t)}}}(\tau) \, \mathbf{1}_{\mathcal{N}_B(X^{(t)})}(y),
\end{aligned}
$$

where $\mathbb{P}_{T_{X^{(t)}}}(\cdot)$ is a probability mass function on $(\mathcal{D} \times \mathcal{A}) \setminus T_{X^{(t)}}$ and where we introduce fixed positive weights satisfying $w_D + w_C + w_A + w_B = 1$.