A Bayesian Ensemble Regression Framework on the Angry Birds Game

Nikolaos Tziortziotis, Georgios Papagiannis and Konstantinos Blekas

Department of Computer Science & Engineering, University of Ioannina, Greece

email: {ntziorzi, gpapagia kblekas}@cs.uoi.gr

Proposed Strategy

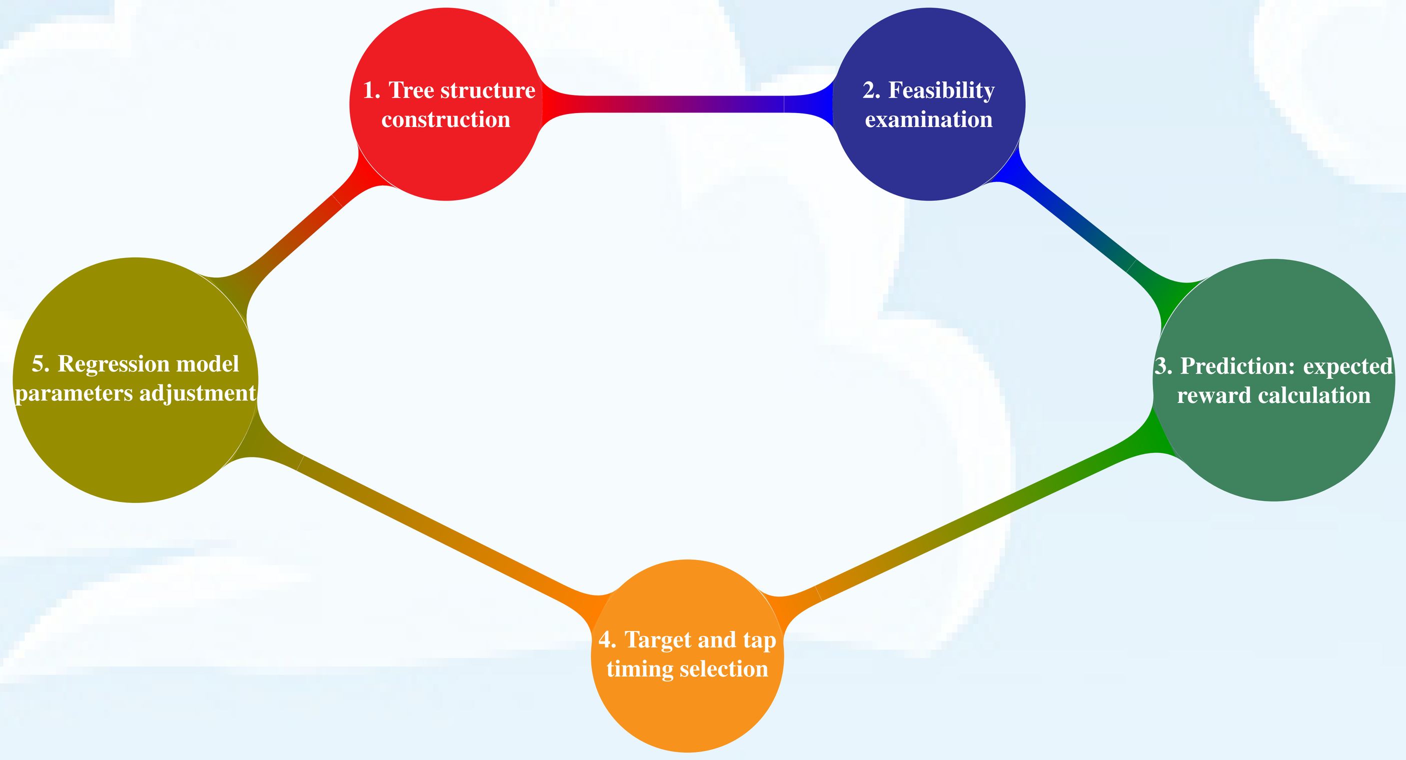

1. Tree Structure

Tree construction:

- Constructed in hierarchical fashion (*bottom-up*)
- Each node represents a union of adjacent objects of the same material.
- Root can be supposed as a virtual node

Three features per tree node are considered:

- $x_1(s)$: **Personal weight**, $x_1(s) = Area(s) \times c_s$
- $x_2(s)$: **Parents cumulative weight**, i.e. $x_2(s) = \sum_{s' \in \mathcal{P}(s)} x_1(s')$
- $x_3(s)$: **Distance** from the nearest pig (normalized to $[0, 1]$)

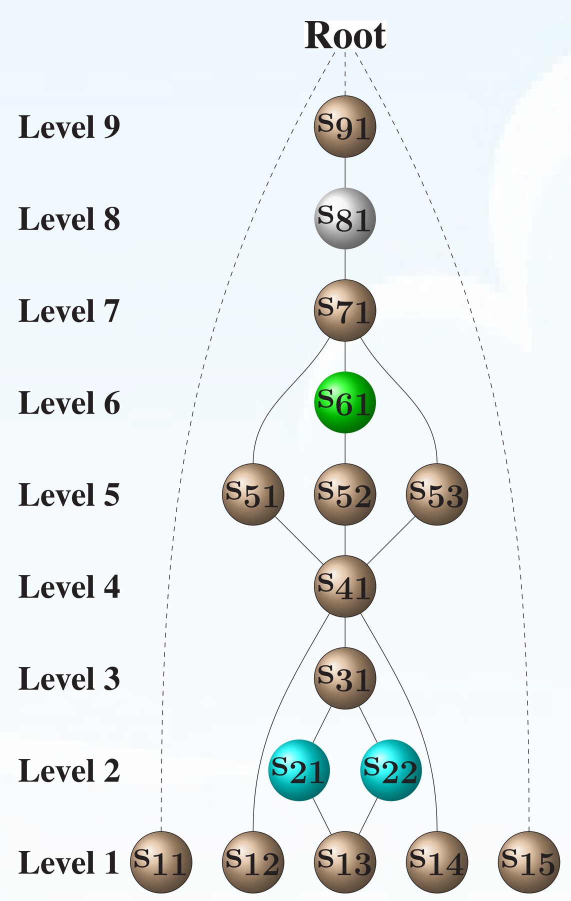

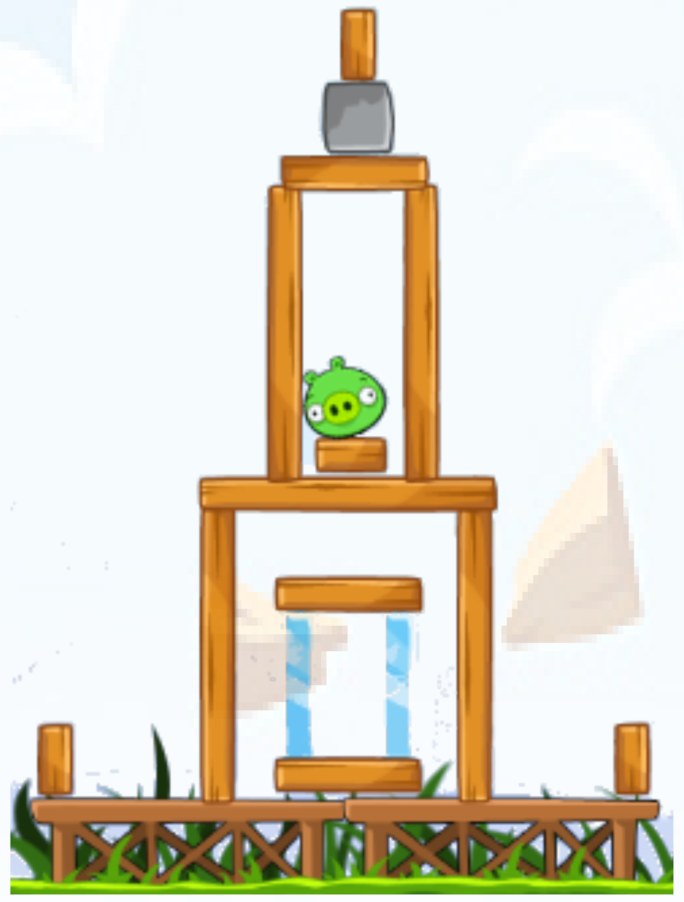

2. Feasibility notion

Two different trajectories are considered:

- Direct shot (angle $<= 45°$)
- High arching shot (angle $> 45°$)

A node is considered as:

- *feasible*, if can be reached directly by at least one shot[a]
- *infeasible*, otherwise

Feasible node

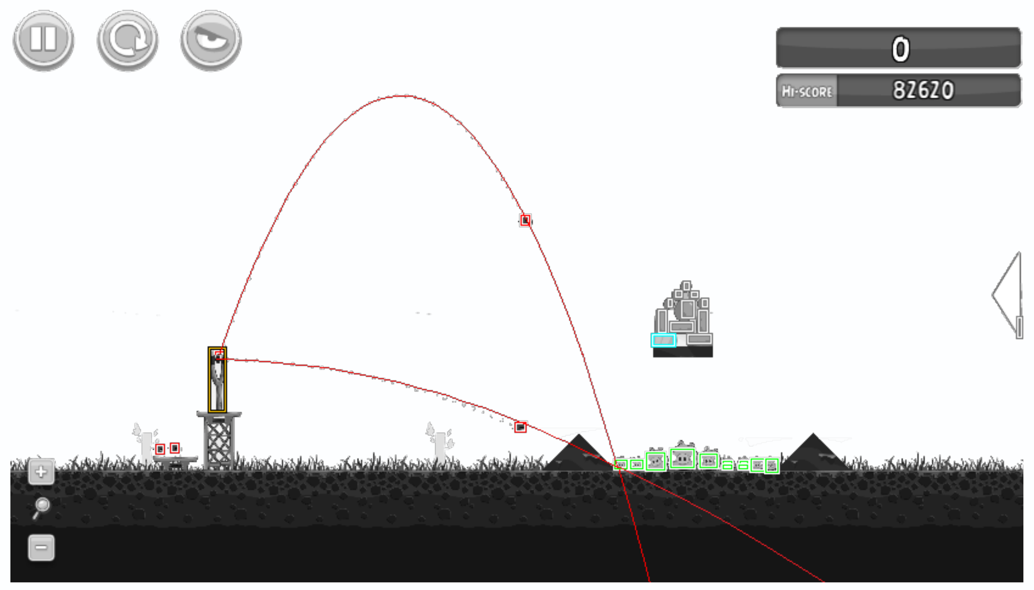

Pig is reachable by at least one trajectory

Infeasible node

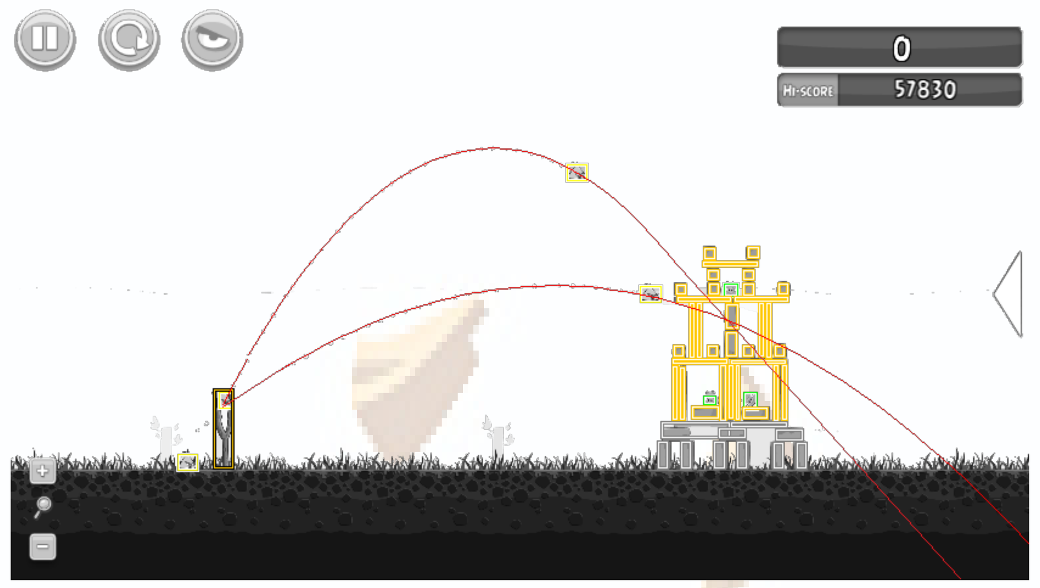

Wood is not directry reachable due to structure

[a] In the case of the white bird a node supposed as feasilbe if can be reached by bird's *egg*.

3. Ensemble linear regression models

A separate linear model is used for each (bird,object material) pair

Linear regression model

The rewards are considered as the target values:

$$t_n = \sum_{i=1}^{M} w_i\phi_i(x_n) + \epsilon_n = \boldsymbol{w}^\top \underbrace{\boldsymbol{\phi}(\boldsymbol{x}_n)}_{\text{Gaussian kernel}} + \underbrace{\epsilon_n}_{\text{noise}}$$

Supposing that $\epsilon_n \sim \mathcal{N}(0, \beta^{-1})$,

$$t_n \mid x = x_n \sim \mathcal{N}(\boldsymbol{w}^\top \boldsymbol{\phi}(\boldsymbol{x}), \beta^{-1})$$

Conditional probability:

$$p(t_{1:n} \mid \boldsymbol{w}, \beta) = \mathcal{N}(t_{1:n} \mid \underbrace{\Phi}_{\text{Design matrix}} \boldsymbol{w}, \beta^{-1} I_n), \quad t_{1:n} \triangleq \{t_k\}_{k=1}^{n}$$

Bayesian linear regression

Conjugate prior: $\boldsymbol{w}|\alpha \sim \mathcal{N}(\boldsymbol{w}|\boldsymbol{0}, \underbrace{a^{-1}}_{\text{prior parameter}} I)$.

Marginal distribution:

$$p(\boldsymbol{w}|t_{1:n}, \alpha, \beta) = \mathcal{N}(\boldsymbol{w}|\boldsymbol{\mu}_n, \Sigma_n),$$

where,

$$\boldsymbol{\mu}_n = \beta \Sigma_n \Phi_n^\top t_{1:n} \quad \text{and} \quad \Sigma_n = (\beta \Phi_n^\top \Phi_n + aI)^{-1}.$$

Predictive distribution:

$$p(t_*|t_{1:n}, \alpha, \beta) = \mathcal{N}(t_*|\boldsymbol{\mu}_n^\top \boldsymbol{\phi}(\boldsymbol{x}_*), \tfrac{1}{\beta} + \boldsymbol{\phi}(\boldsymbol{x}_*)^\top \Sigma_n \boldsymbol{\phi}(\boldsymbol{x}_*)).$$

4(a). Target selection mechanism

✓ Only the feasible nodes are examined

✓ The best arm is selected greedily according to:

$$j^* = \arg\max_q \left\{ \left(\boldsymbol{\mu}^{f(q)}_{n_{f(q)}}\right)^\top \boldsymbol{\phi}(\boldsymbol{x}_q) + C\sqrt{\frac{2\ln N}{n_{f(q)}}} \right\}$$

- $f(q)$: denotes the regression model for node, q
- $n_{f(q)}$: number of times where has been selected
- N: total number of plays
- C: has been selected equal to 3000

4(b). Tapping selection

No tapping

At the 75% of the trajectory from the slingshot to the first collision point

At the 85% of the trajectory(direct shot)
At the 90% of the trajectory(high-arching)

The tapping is performed when the bird lie above the target

No tapping

White bird tapping

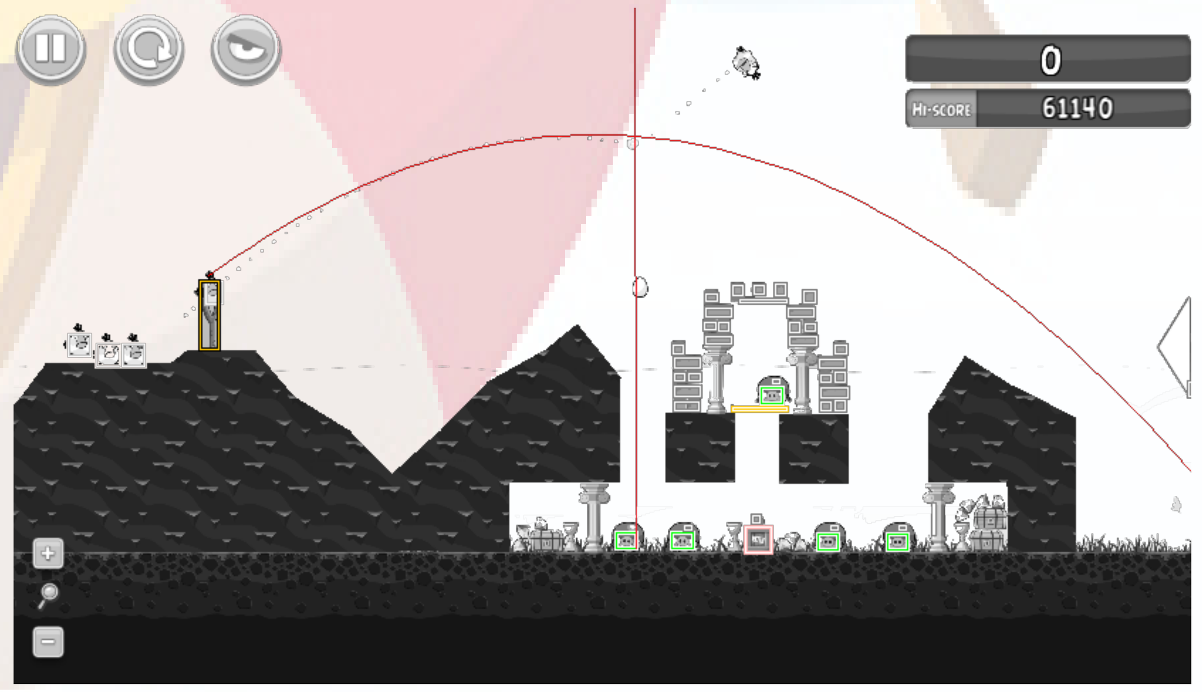

A node supposed as feasible if can be reached by bird's *egg*.

5. Online model's parameters learning

✓ Regressor $k \triangleq f(j^*)$ has been selected

The received observation (reward) t_{n_k+1} follows,

$$p(t_{n_k+1}|\boldsymbol{w}_k) = \mathcal{N}(t_{n_k+1}|\boldsymbol{w}_k^T \boldsymbol{\phi}(\boldsymbol{x}_{n_k+1}), \beta).$$

The weights' posterior distribution is given as:

$$p(\boldsymbol{w}_k|t_{1:n_k+1}) = p(t_{n_k+1}|\boldsymbol{w}_k)p(\boldsymbol{w}_k|t_{1:n_k}) = \mathcal{N}(\boldsymbol{w}_k|\boldsymbol{\mu}^k_{n_k+1}, \Sigma^k_{n_k+1}),$$

where the Gaussian parameters are given as:

$$\Sigma^k_{n_k+1} = \left[(\Sigma^k_{n_k})^{-1} + \beta \boldsymbol{\phi}(\boldsymbol{x}_{n_k+1})^T \boldsymbol{\phi}(\boldsymbol{x}_{n_k+1})\right]^{-1}$$

$$\boldsymbol{\mu}^k_{n_{j^*}+1} = \Sigma^k_{n_k+1}\left[\beta_k \boldsymbol{\phi}^T(\boldsymbol{x}_{n_k+1}) t_{n_k+1} + (\Sigma^k_{n_k})^{-1} \boldsymbol{\mu}^k_{n_k}\right].$$

Acknowledgment

This work was partially supported by a travel grant from the ECCAI society.