One Model, Any CSP: Graph Neural Networks as Fast Global Search Heuristics for Constraint Satisfaction

IJCAI 2023

Jan Tönshoff, Berke Kisin, Jakob Lindner, Martin Grohe

[Stanford CS/MS&E 331]

Motivation

Constraint Satisfaction Problems (CSPs):

- Unify SAT, Graph Coloring, MAXCUT, and all NP-hard problems

Motivation for ANYCSP:

- Classical heuristics are hand-engineered, domain-specific
- Desire a single, general-purpose solver across CSPs

Goals:

- Learn *global* search heuristics through a shared GNN
- Train on small synthetic instances, generalize to large real ones

Key ideas

Unified View of CSPs

- Represent any CSP as a *Constraint–Value Graph (CVG)*

General search policy

- A single *Graph Neural Network (GNN)* operates on the CVG
- Learns to propose coordinated updates to all variables
  - Produces global actions rather than one-variable-at-a-time flips

Training objective

- Use RL (REINFORCE) to maximize solution quality
- No supervision from known solutions

CSPs: Recap

CSP instance $\mathcal{I} = (\mathcal{X}, \mathcal{D}, \mathcal{C})$:

- **Variables** $\mathcal{X}$
- $\mathcal{D}$ assigns to each variable $X \in \mathcal{X}$ a **domain** $\mathcal{D}(X)$
- **Assignment** α assigns each variable $\alpha(X) \in \mathcal{D}(X)$
- **Constraint** $C \in \mathcal{C}$
  - Defined by:
    1. Scope $s^C = (X_1, \dots, X_k)$
    2. Relation $R^C \subseteq \mathcal{D}(X_1) \times \cdots \times \mathcal{D}(X_k)$
  - Satisfied if $\big(\alpha(X_1), \dots, \alpha(X_k)\big) \in R^C$

Goal: find α that satisfies as many constraints as possible

Constraint-value graph (CVG)

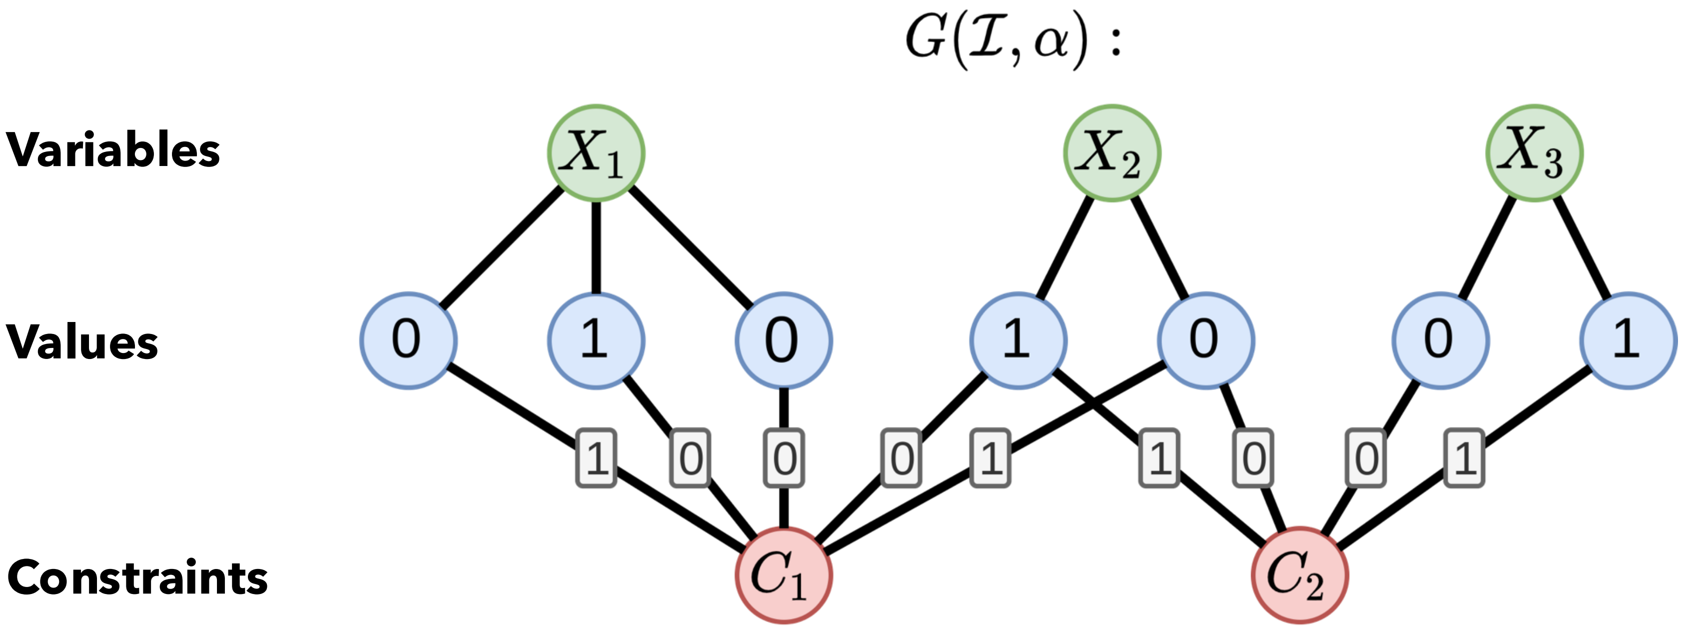

Constraint-value graph (CVG)

- $\mathcal{X} = \{X_1, X_2, X_3\}$
- $\mathcal{D}(X_1) = \{1,2,3\}$
- $\mathcal{D}(X_2) = \mathcal{D}(X_3) = \{1,2\}$
- $\alpha = (2,1,2)$
- $C_1 : X_1 \leq X_2$
- $C_2 : X_2 = X_3$
- For constraint C, variable X_i, value d, **label** is 1 iff $\left(\alpha(X_1), \dots, \alpha(X_{i-1}), d, \alpha(X_{i+1}), \dots, \alpha(X_k)\right) \in R^C$

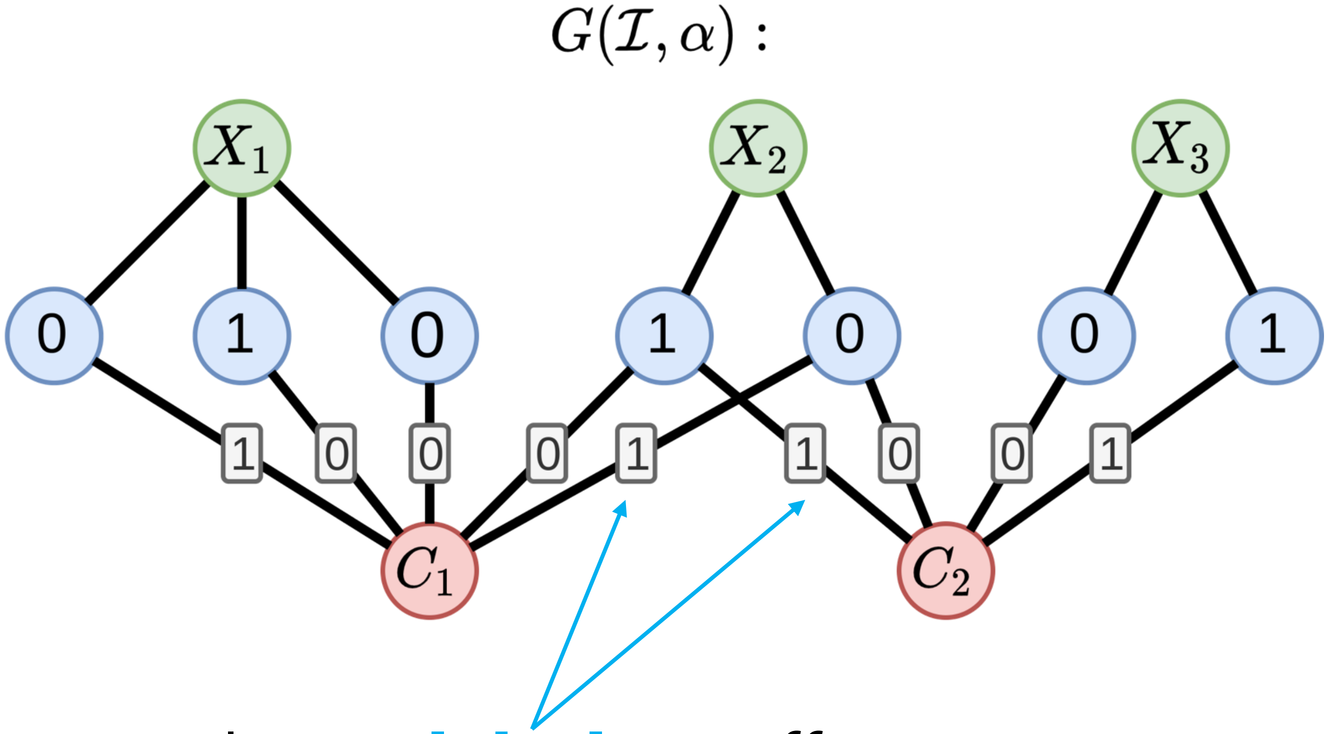

ANYCSP architecture

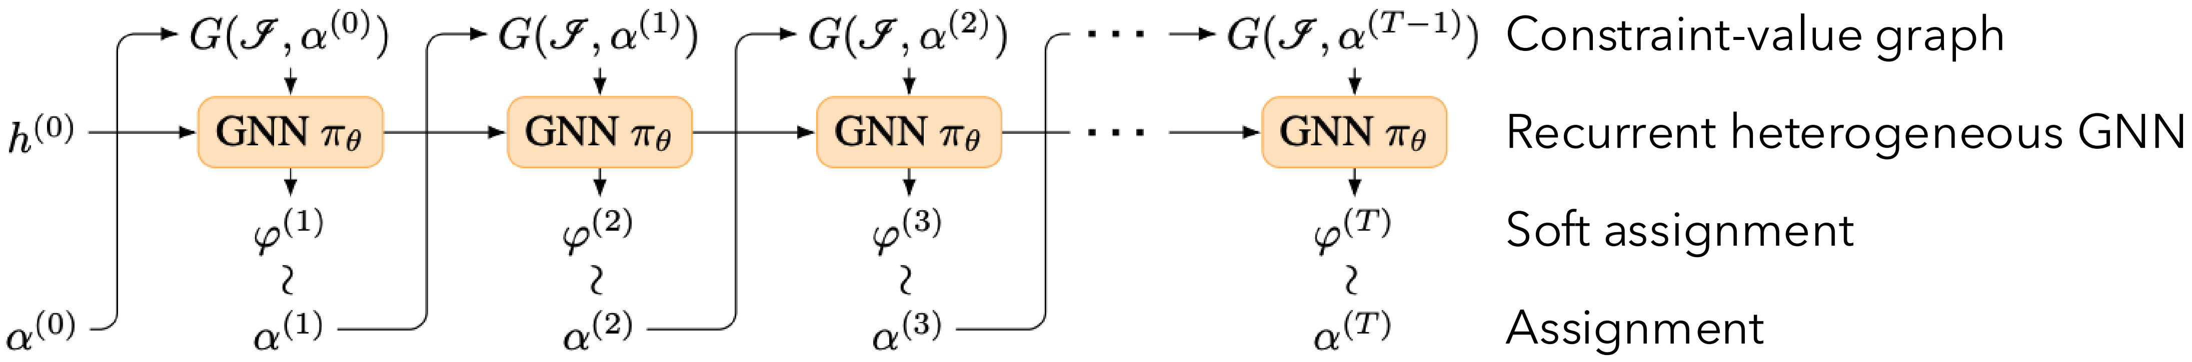

GNN has four directional layers:

- $V \to C$: inform constraints about tentative assignments
- $C \to V$: send satisfaction feedback to connected values
- $V \to X$: values aggregate into per-variable summaries
- $X \to V$: broadcast updated intent back to candidate values

Reward design

- **Goal:** reward policy iff it improves the best-so-far solution
- $Q_{\mathcal{J}}(\alpha)$ = fraction of satisfied constraints
- Naive reward $Q_{\mathcal{J}}\left(\alpha^{(t)}\right)$ caused stagnation at local maxima
- Define running best: $q^{(t)} = \max_{t'<t} Q_{\mathcal{J}}\left(\alpha^{(t')}\right)$
- Reward $r^{(t)} = \begin{cases} 0, & Q_{\mathcal{J}}\left(\alpha^{(t)}\right) \leq q^{(t)} \\ Q_{\mathcal{J}}\left(\alpha^{(t)}\right) - q^{(t)}, & Q_{\mathcal{J}}\left(\alpha^{(t)}\right) > q^{(t)} \end{cases}$
- No penalty for exploratory worse steps
- Training: vanilla REINFORCE, $T = 40$ iterations

Experimental setup

- Benchmark setup: MODEL RB [Xu, Li, '03]
  - Generates dense, random CSPs near satisfiability threshold
  - Training distribution: random MODEL RB instances
    - 30 variables and constraint arity 2
  - Test dataset (RB50): 50 satisfiable instances
    - 50 variables, domain size 23, ≈ 500 constraints
    - Standard benchmark in the XCSP Competition
- Baselines: ACE, CoSoCo, Picat
  - Picat: SAT-based solver, 2022 XCSP Competition winner
- Each solver runs once per instance with a 20-minute timeout
  - ANYCSP performs ≈ 500 k search iterations in this window

ModelRB results

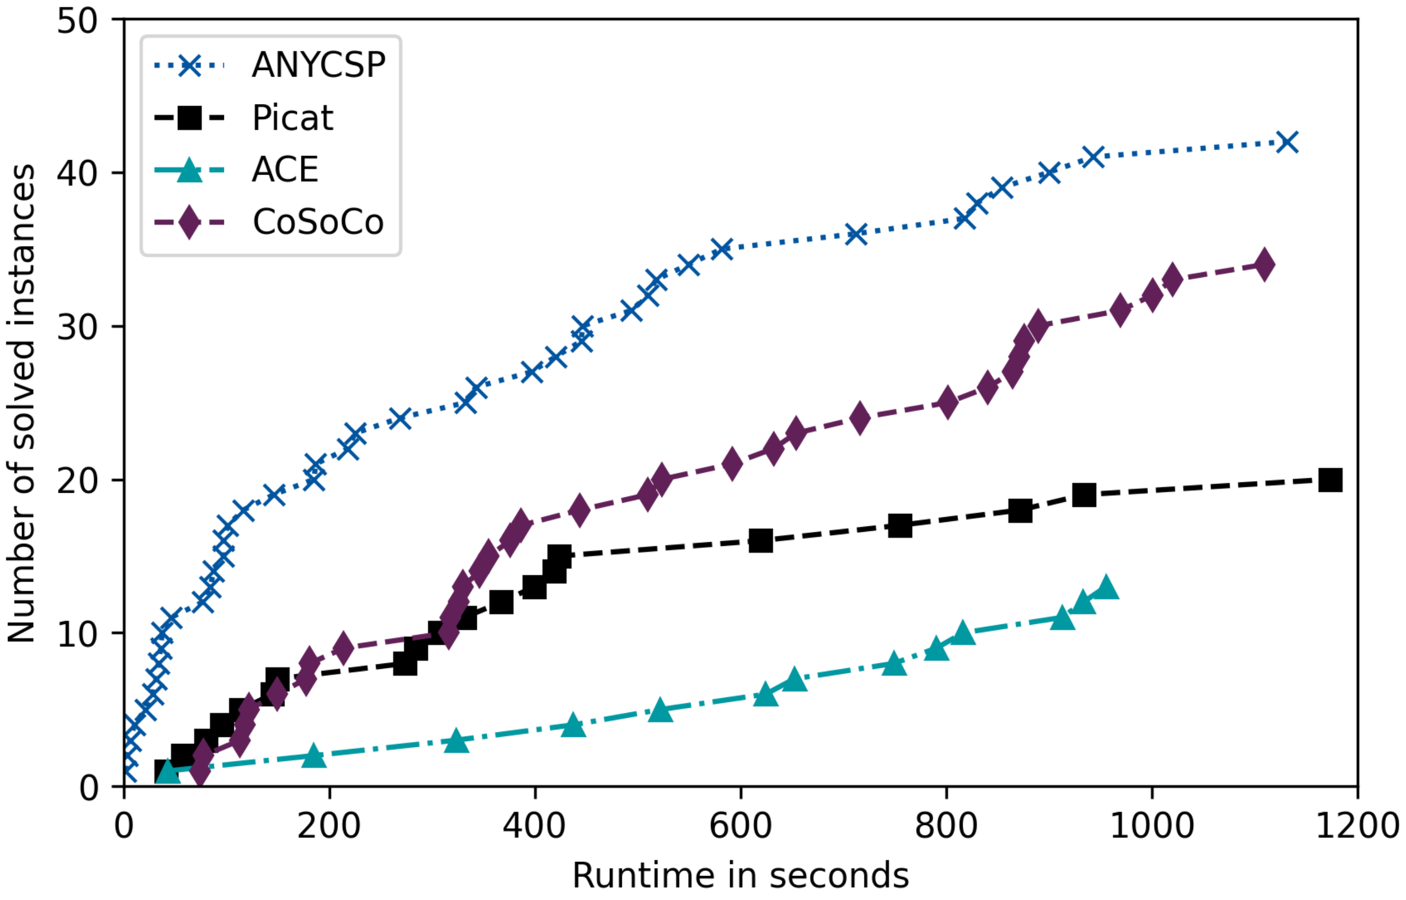

MaxCut: Experimental setup

- Training: unweighted Erdős–Rényi graphs with 100 vertices
- Testing: Gset [Ye'03]—diverse instances with 800–10k vertices
- Neural baselines
  - RUNCSP [Tönshoff, '21] (supervised)
  - ECO-DQN [Barrett et al., '20] (RL)
  - ECORD [Barrett et al., '22] (RL)
- Classical baselines: Greedy, Goemans–Williamson SDP
- Evaluation protocol
  - 20 parallel runs, 180-second timeout
  - Report mean deviation from best-known cuts [Benlic, Hao, '13]

MaxCut results

METHOD	$\|V\|$=800	$\|V\|$=1K	$\|V\|$=2K	$|V|\geq$3K
GREEDY	411.44	359.11	737.00	774.25
SDP	245.44	229.22	-	-
RUNCSP	185.89	156.56	357.33	401.00
ECO-DQN	65.11	54.67	157.00	428.25
ECORD	8.67	8.78	39.22	187.75
ANYCSP	**1.22**	**2.44**	**13.11**	**51.63**

Additional domains (see paper)

- Graph coloring:
  - Better than existing neural solvers; on par with best heuristic
- MAX-3-SAT
  - Baselines: neural approaches and conventional stochastic search
  - Neural approaches generally can't compete w/ CDCL solvers
- MAX-k-SET
  - Compares against SOTA local search algorithms
  - Global updates of ANYCSP beat local search in # iterations
    - Local-only variant of ANYCSP loses to strong heuristics
  - Classic baselines use CPU and have better runtimes

Motivation

Constraint Satisfaction Problems (CSPs):

- Unify SAT, Graph Coloring, MAXCUT, and all NP-hard problems

Motivation for ANYCSP:

- Classical heuristics are hand-engineered, domain-specific
- Desire a single, general-purpose solver across CSPs

Goals:

- Learn *global* search heuristics through a shared GNN
- Train on small synthetic instances, generalize to large real ones