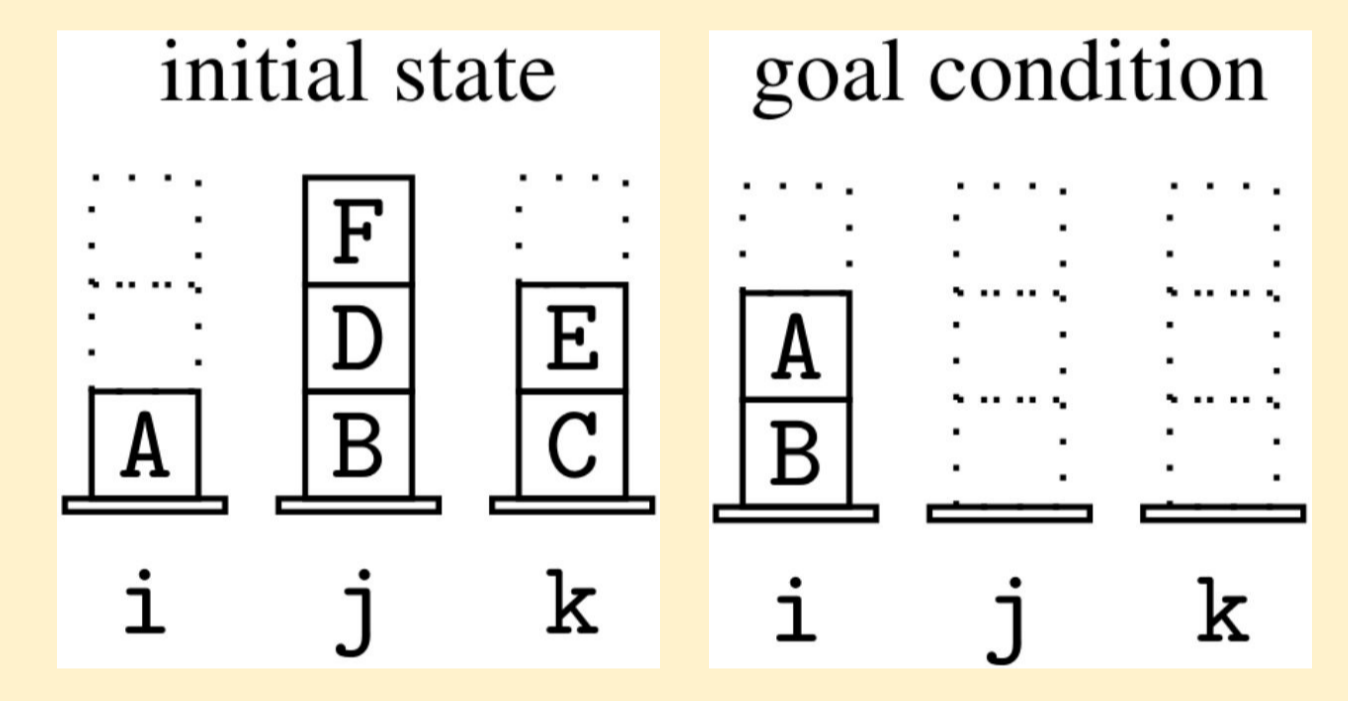




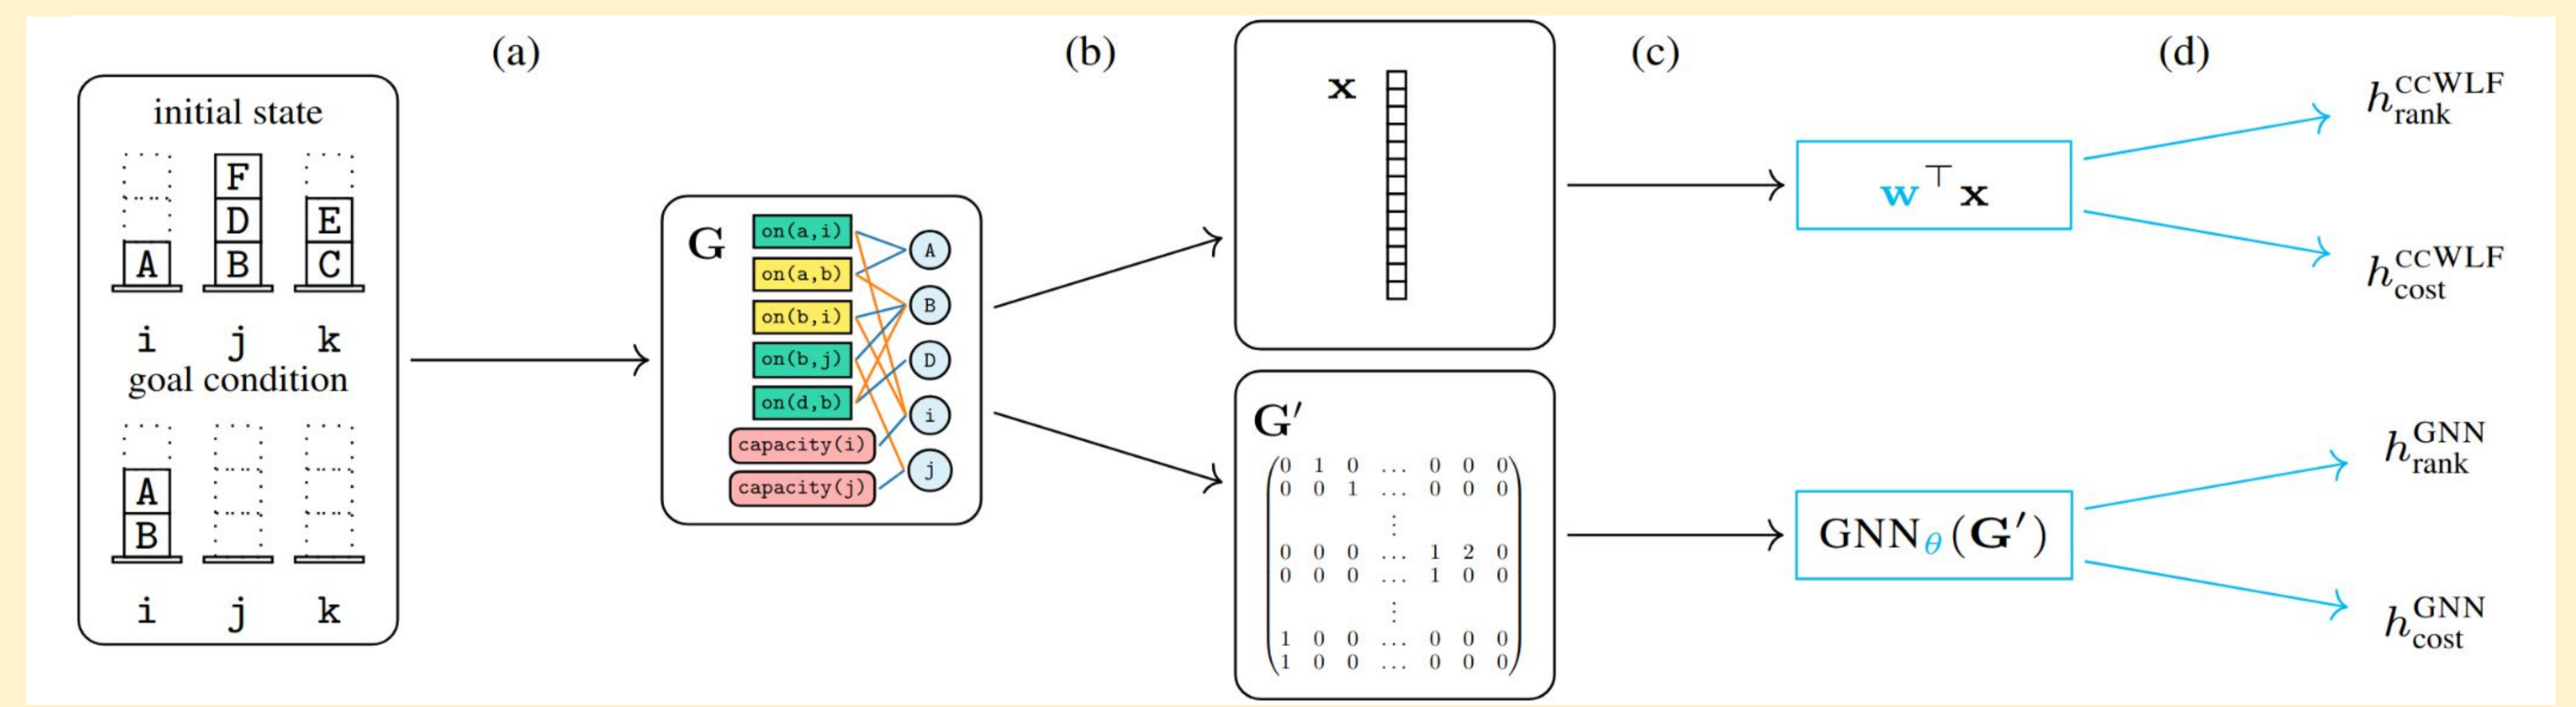
## Numeric Planning

- factored representations of states
  - predicates** for Boolean values:
    - $on(A, i), (F, D), on(D, B), on(B, j), on(E, C), on(C, k)$
  - functions** for numeric values
    - $capacity(i) = 2, capacity(j) = 0, capacity(k) = 1$
- transition system and goals specified
  - by Boolean assignments
    - $on(A, B), on(B, i)$
  - numeric conditions
    - $capacity(i) \geq 1$
- solution:
  - goal-reaching sequence of actions



## Learning Heuristic Functions

- (a) numeric planning tasks are **transformed into graphs** with edge labels, and both categorical and continuous node features
- (b) graphs are fed into graph kernels (top) or transformed into matrix representations (bottom)
- (c) graph kernel features use a linear kernel; matrix representations are fed into GNNs
- (d) models **learn a heuristic function** as either cost-to-go estimators or ranking functions used for search



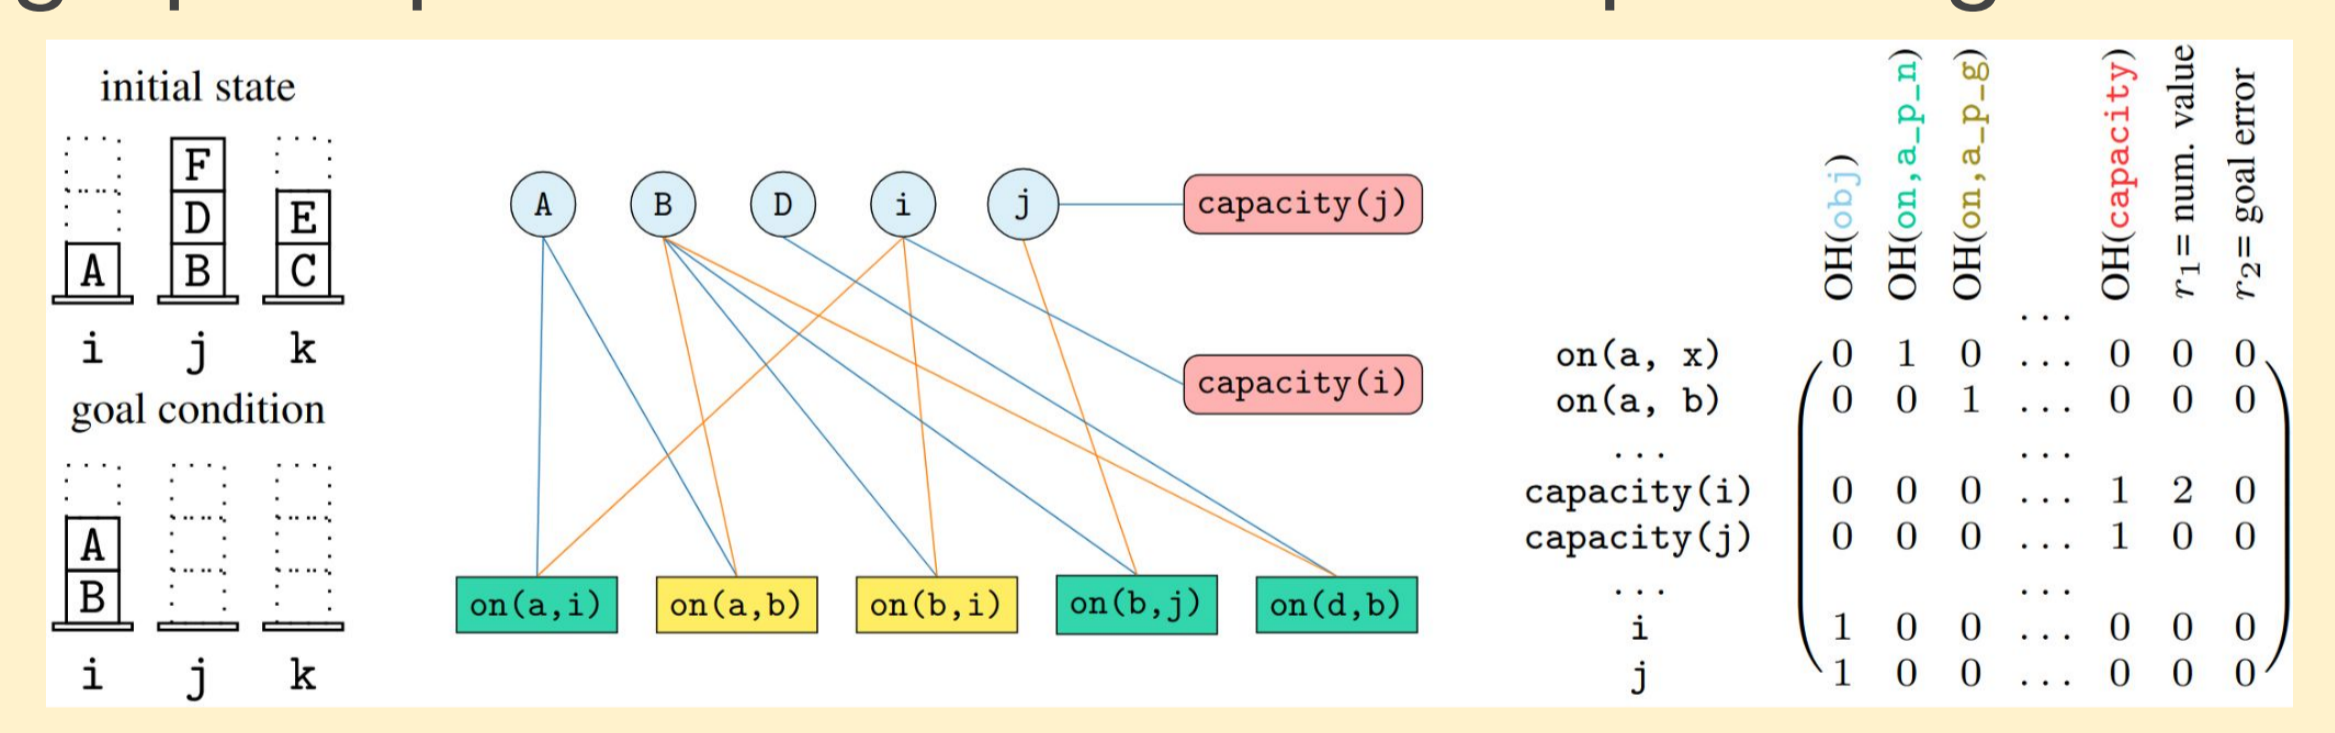
## Planning and RL: What's the Difference?

Both Solve MDPs!

Planning	Reinforcement Learning
model-known	model-free or model-based
goal-conditioned, minimise cost	maximise reward
transitions modelled symbolically	transitions modelled as distributions
search algorithms: A*, GBFS, iLAO*, LRTDP	search algorithms: MTCS, UCT, TD( $\lambda$ )
heuristic functions (cost-to-go estimator)	value functions (expected reward)
algorithms guided by models	algorithms guided by rewards

## New Contributions

- graph representation for numeric planning



- graph kernel for categorical-continuous attributes


**Algorithm 1: CCWL algorithm**

**Data:** A graph  $G = \langle V, E, F_{cat}, F_{con}, L \rangle$  with continuous and categorical attributes, a deterministic HASH function, allowed colours  $C = |C|$ , a pooling function POOL, and number of CCWL iterations  $L$ .

**Result:** Feature vector of size  $\mathbb{R}^{(1+d)|C|}$ .

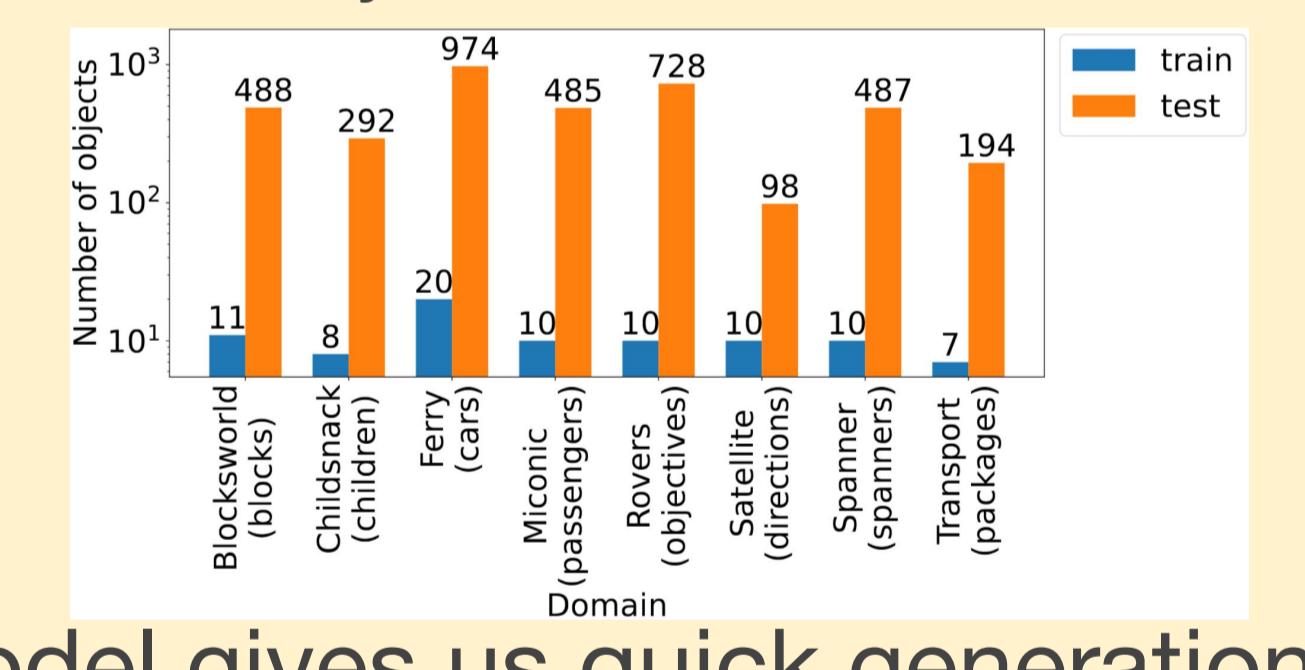
- $\kappa^0(v) \leftarrow F_{cat}(v), \forall v \in V$
- for**  $j \in [L]$  **do for**  $v \in V$  **do**
- $\kappa^j(v) \leftarrow HASH(\kappa^{j-1}(v), \bigcup_{u \in N_i(v)} \{\kappa^{j-1}(u), \iota\} \mid u \in N_i(v)\})$
- $M \leftarrow \bigcup_{j=0, \dots, L} \{\kappa^j(v) \mid v \in V\}$
- $\vec{v}_{cat} \leftarrow [COUNT(M, c_1), \dots, COUNT(M, c_{|C|})]$
- $S_i = \{v \in V \text{ s.t. } \exists j \in \{0\} \cup [L], \kappa^j(v) = c_i\}, \forall i \in C$
- $\vec{v}_{con} \leftarrow [con(1) \parallel \dots \parallel con(|C|)], con(i) = POOL_{v \in S_i}(F_{con}(v))$
- return**  $\vec{v}_{cat} \parallel \vec{v}_{con}$

- LP ranking formulation for learning heuristic

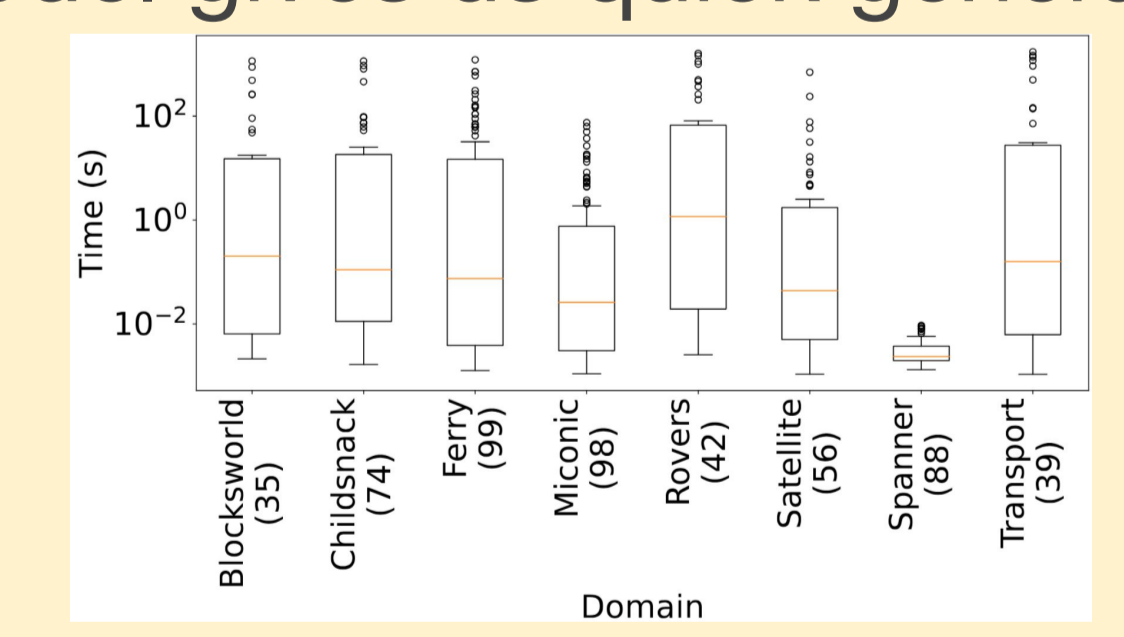


## Experiments


- Generalise across object sizes  $\rightarrow$  out-of-distribution learning



- access to model gives us quick generation of training data



- competitive with domain-independent numeric planners
  - purple = numeric planners
  - orange = learning planners



## Graph Learning

- in the top 3 keywords of the past 3 ICLR conferences
- handles relational data and arbitrarily sized inputs = planning
- deep learning
  - Graph Neural Networks (GNNs)
- classical ML
  - graph kernels