Meet-in-the-middle with Early and Efficient Termination (MEET)

Bidirectional Search while Ensuring Meet-In-The-Middle via Effective and Efficient-to-Compute Termination Conditions

Yi Wang[1], Eyal Weiss[3], Bingxian Mu[2], Oren Salzman[3]

[1]University of New Hampshire, [2]University of Prince Edward Island, [3]Technion-Israel Institute of Technology

yw1055@usnh.edu, bmu@upei.ca, {eweiss@campus,osalzman@cs}.technion.ac.il

Challenges, Open Questions, Contributions

Challenges:

- Advanced methodlogies, such as *Meet-in-the-middle* MM propety (MMP), *Must expand pairs* (MEP), and Errors Estimates (EE), terminate the search by exploring *all* possibly better solutions, which leads to many unnecessary expansions.
- MM and its successors algorithms terminate the search by requiring evaluating minimal g-values by iterating over *all* states in each open list at each iteration, which can lead to significant computational burden.

Open question:

Early termination while ensuring MMP and Optimality by no future generation can improve the current best solution in bidirectional heuristic search (Bi-HS), remains open.

Contributions:

- A tighter termination condition enables early stopping **while preserving** MMP and optimality.
- Eliminates the need to access minimal g-values in both open lists at each iteration.
- Runs at least two orders of magnitude over MM, and stays on par with or better than A* and BAE* even with weaker heuristics.

Illustration of Algorithms

Comparative illustration of A*, MM, BAE*, and MEET when finding an optimal solution in an instance on a grid map from MovingAI repository.

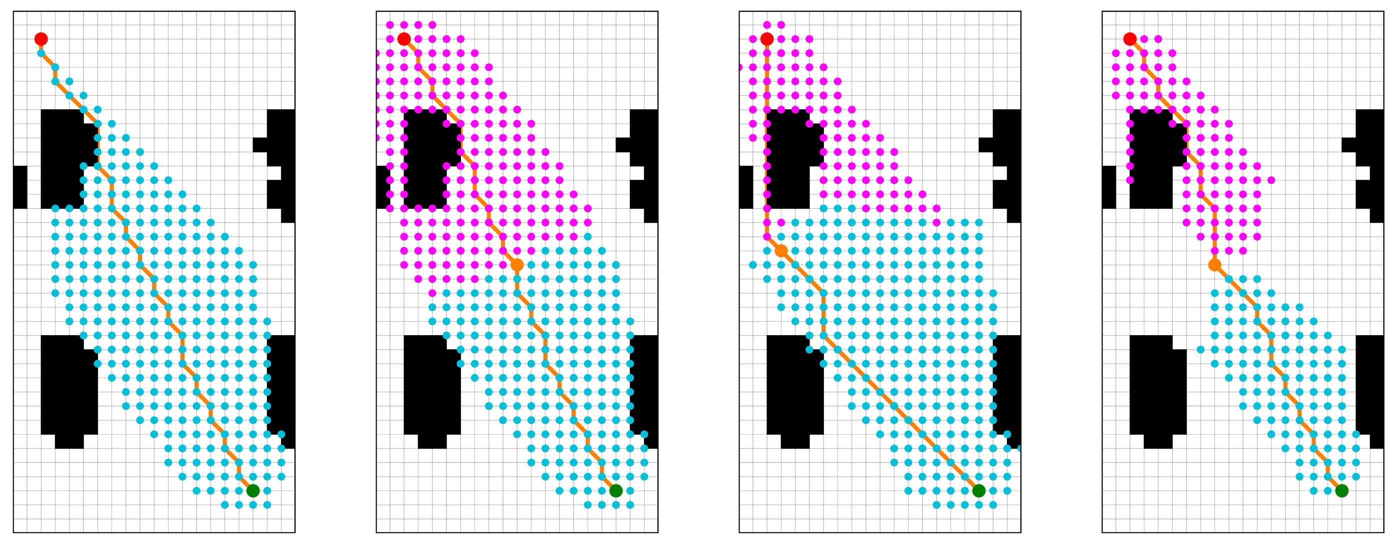

(c) A* (292) (d) MM (381) (e) BAE* (352) (f) MEET (202)

MEET ($t = 0.075$ ms), A* ($t = 0.106$ ms), MM ($t = 4.118$ ms), BAE* ($t = 0.232$ ms).

(•) and (•) states expanded in the forward and backward searches,
(•) start state, (•) goal state, — obstacle, (—) an optimal path.

Open Question

Why still open and so hard:

1. Achieving MMP took over 50 years due to key challenges such as proving optimality, frontier crossing, missing frontiers, etc.
2. **MM** algorithm incurs high computational cost due to requiring to access the minimal g-value and f-value at each iteration.
3. An in-exhaustive exploration of all potential better solutions can greatly risk losing optimality in a search algorithm.
4. Ensuring that "No unseen state can improve the current best solution" typically requires reasoning about all potential methods the frontiers might satisfy.
5. Unlike unidirectional heuristic search, bidirectional heuristic search lacks a tight bound on future improvements, which makes the search not terminate immediately, even when an optimal solution has been found.

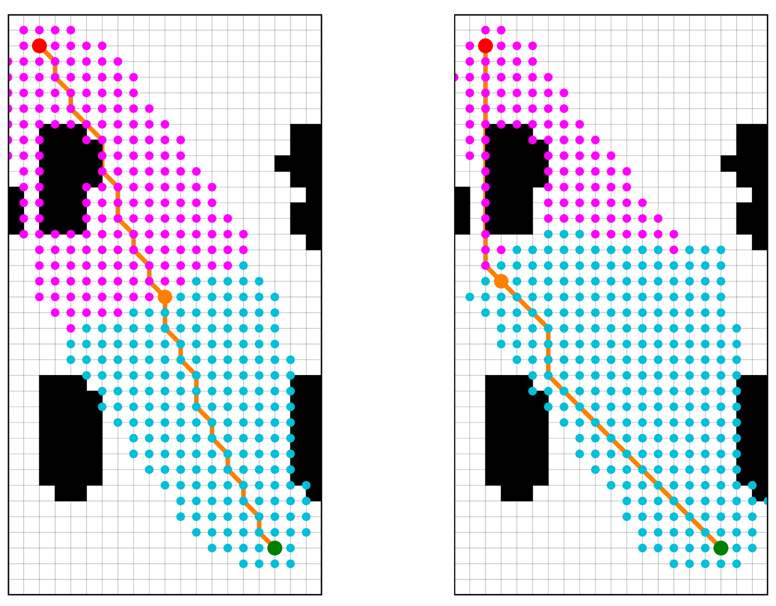

(a) MM-family (b) BAE*

(•) and (•) states expanded in the forward and backward searches,
(•) start state, (•) goal state, — obstacle, (—) an optimal path.

Comparison:

Concept	MM	MEET
State Pruning	No	Yes
State Priority Function	$\max(f, 2g)$	f with adaptive h
Admissible Heuristic	Yes	Yes
Termination Condition	Access g_{min}-and f_{min}-values	only f-value
MMP and Optimality	Yes	Yes

f, g, h: estimated total solution cost, cost-to-come, and cost-to-go, $f = g + h$.

Key Contributions:

Key Contributions:

A tighter termination condition (TC) that neither (i) requires accessing the minimal g-values in each open list at every iteration nor (ii) heavily relies on expanding all potential solution states.

MEET terminates when:

- TC1: the minimum priority exceeds the incumbent solution.
- TC2: the search fronts are too far apart to yield a better solution.
- TC3: the g-cost of a meeting state from its generating side is higher than from the other side, and nothing ahead can exceed it.
- TC4: the g-cost of a meeting state from its generating side is lower, but no further state could improve the incumbent solution.

.

Experimental Results

Comparative performance in Grid Maps and 10 Pancake Puzzle

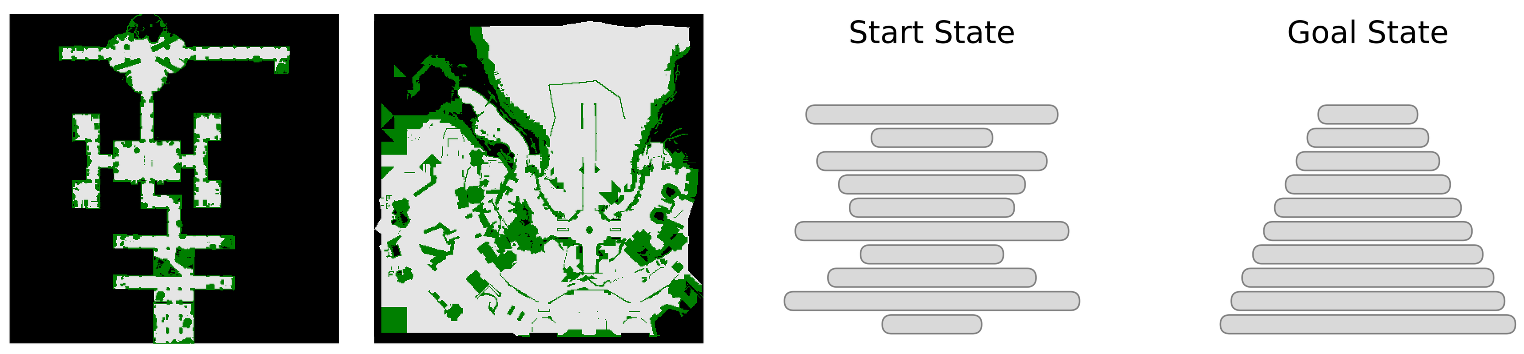

(g) Brc203d (h) Orz100d (i) 10 Pancake Puzzle

Domain	Heuristic	A*	MM	BAE*	MEET
brc203	Euclidean Distance	**1.4** (0.42, 2.4)	216 (52, 366)	3.9 (1.3, 6.1)	2.2 (0.65, 3.3)
	Octile Distance	**1.6** (0.36, 3.1)	406 (82, 879)	4.3 (1.3, 6.5)	2.4 (0.65, 3.8)
orz100d	Euclidean Distance	**8.5** (1.8, 14.5)	3.467×10^3 (588, 6.411×10^3)	23.8 (5.2, 34.3)	13.4 (2.7, 22.1)
	Octile Distance	**9.2** (1.4, 16.5)	5.214×10^3 (864, 11.339×10^3)	25 (4.6, 37.2)	14.7 (2.4, 24.5)
10-Pancake puzzle	GAP-1	**0.31** (0.08, 0.75)	7.1 (1.5, 22.3)	0.46 (0.2, 1.5)	0.48 (0.14, 0.74)
	GAP-2	5.6 (3.6, 17.9)	447 (110, 2.275×10^3)	1.3 (0.74, 3.1)	**1.0** (0.45, 2.1)
	GAP-3	81 (45.8, 291)	12.189×10^3 (403, 27.420×10^3)	8.2 (3.0, 24.8)	**3.9** (2.2, 11.1)
	GAP-4	510 (224, 1.605×10^3)	36.1326×10^3 (529, 40.419×10^3)	26.1 (13.7, 78.8)	**6.5** (3.8, 32.5)

Table: The median runtime and the inter-quartile range (IQR), indicated in parentheses as (25th percentile, 75th percentile).

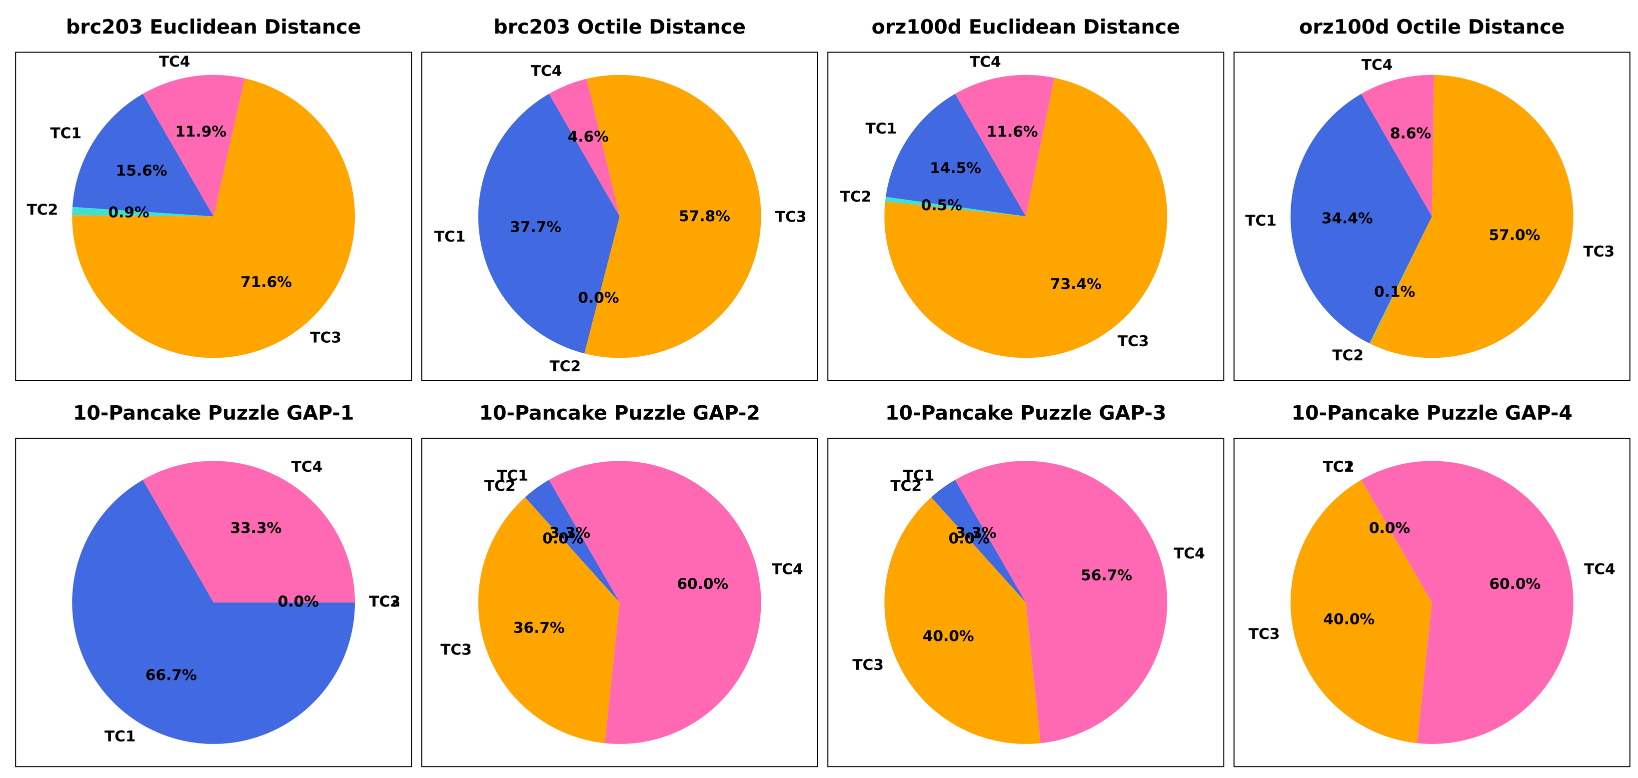

Figure: Triggered TC distributions. The colors (•), (•), (•), and (•) denote **TC1, TC2, TC3,** and **TC4**, respectively.

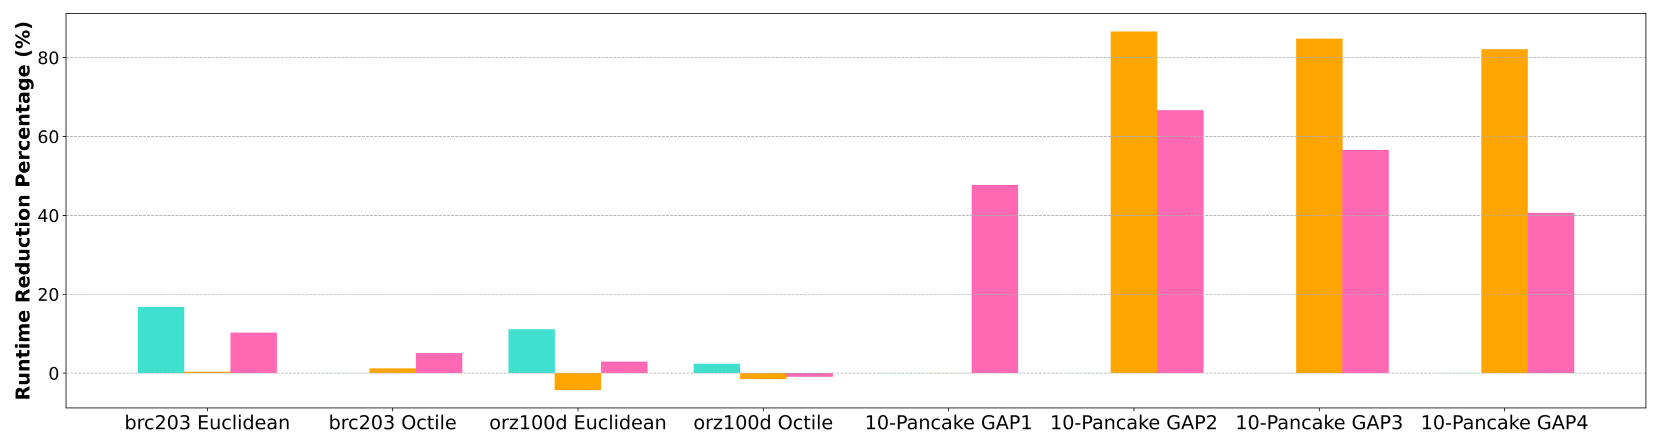

Figure: Average runtime reduction (%) vs. stopping on **TC1**. The bars (—), (—), and (—) denote **TC2, TC3,** and **TC4**, respectively.