25th INTERNATIONAL CONFERENCE ON PATTERN RECOGNITION

Milan, Italy 10|15 January 2021

Deep Next-Best-View Planner for Cross-Season Visual Route Classification

Kurauchi Kanya, Tanaka Kanji

University of Fukui, Japan

Outline

- Introduction
- Method
- Implementation
- Experiments
- Conclusions

Long-term visual place recognition (VPR)

Goal: classify ego-centric view images into pre-defined place classes

Challenge: domain-shifts (e.g., changes in viewpoints, illumination, weather conditions, and seasons)

Common approach: fine-tuning DNN (deep convolutional neural network) as VPC (visual place classifier)

Training data: action-observation sequences

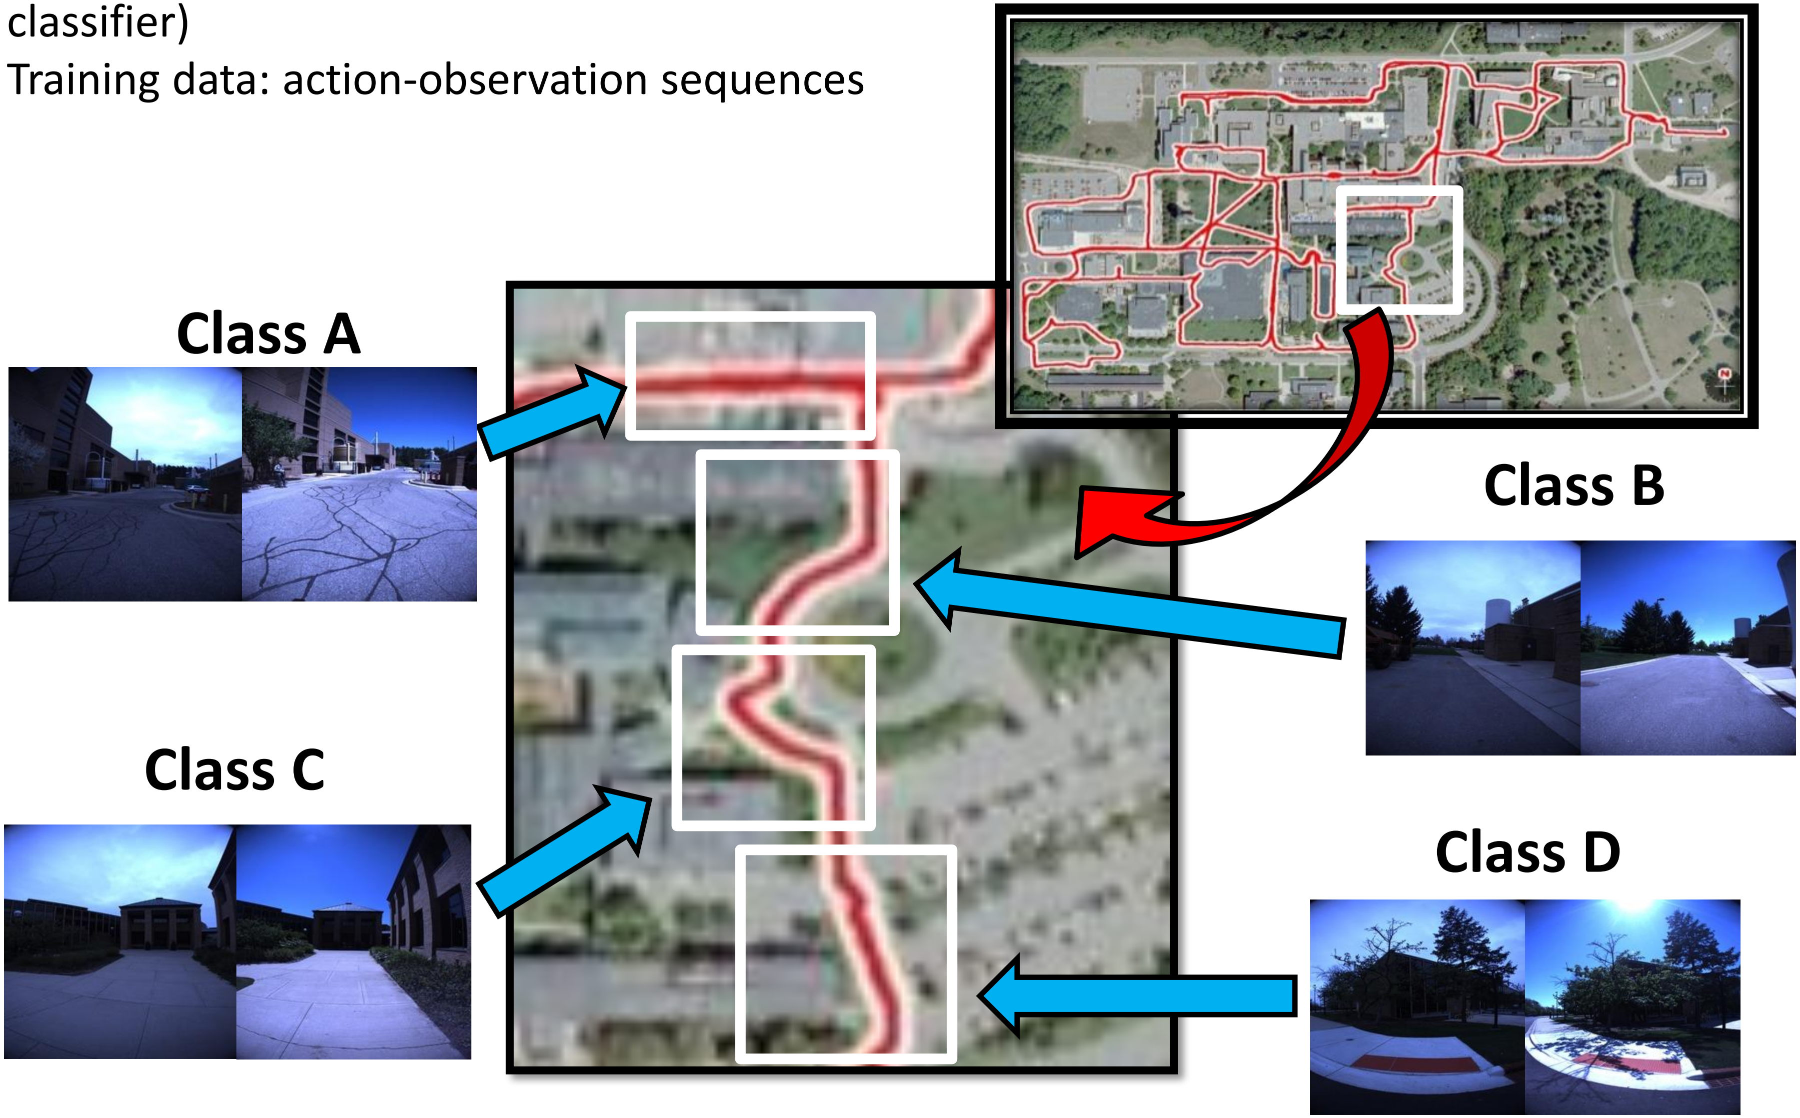

Common assumption: passive setting

Assumption: the robot's action is determined by a predefined control rule, such as a constant speed motion rule.

Limitations:

1. viewpoints are not necessarily optimized for the VPR task i.e., less informative non-discriminative images, sub-optimal VPR performance
2. produce an unnecessarily large number of redundant observations, i.e., significant decrease in the cost performance of VPR

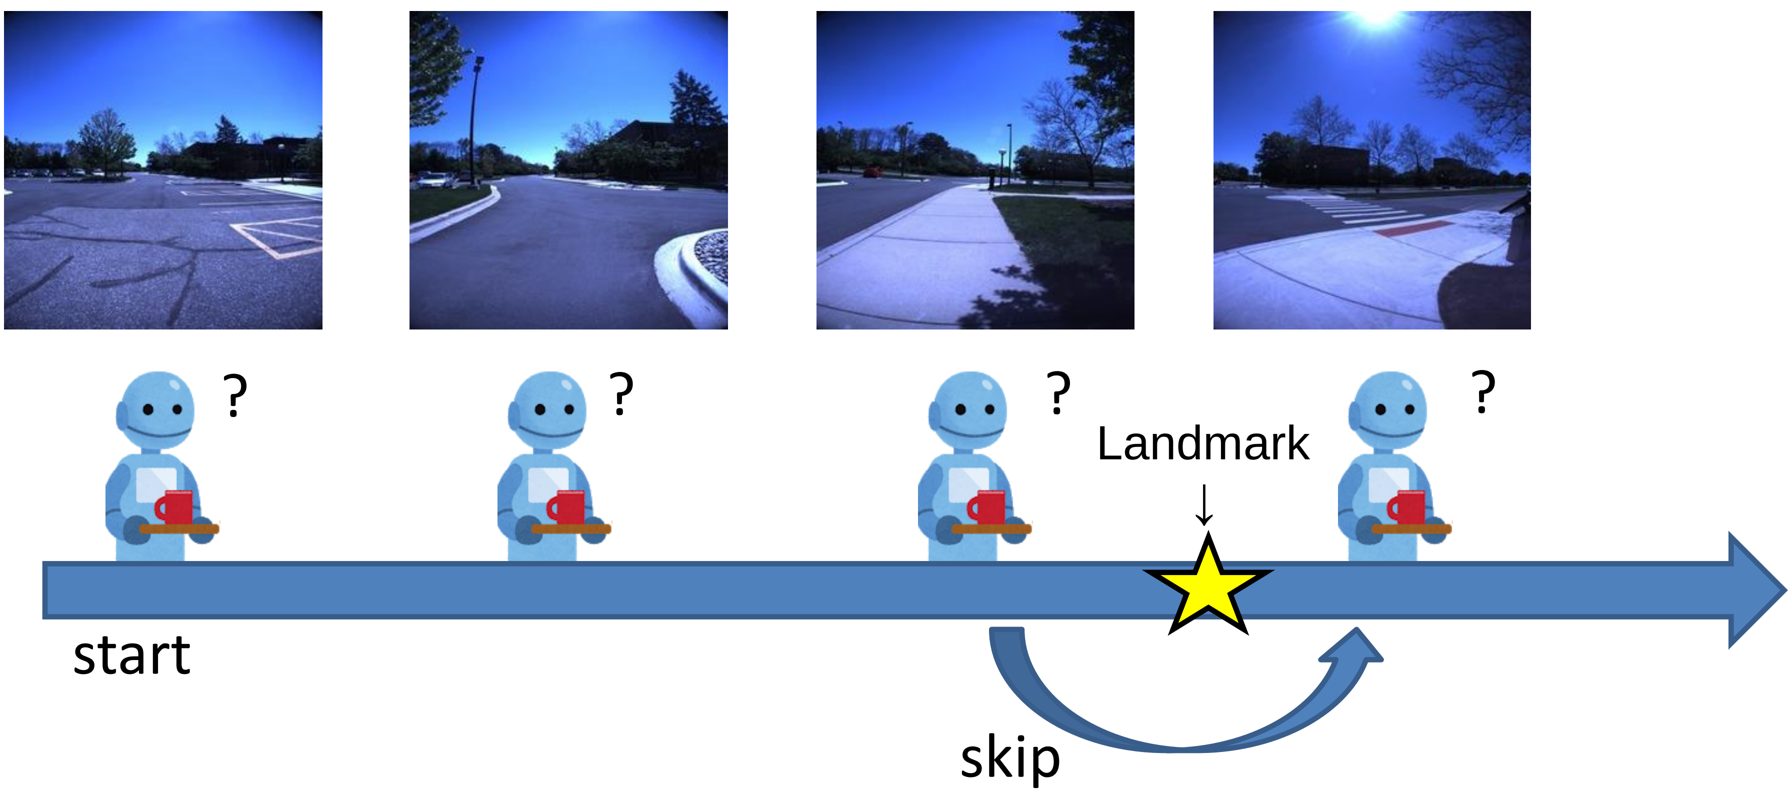

Our approach: active VPR in visual route classification

Goal: classify a view sequence into predefined route classes.

Method: train a state-to-action mapping function, i.e., next-best-view (NBV) planner

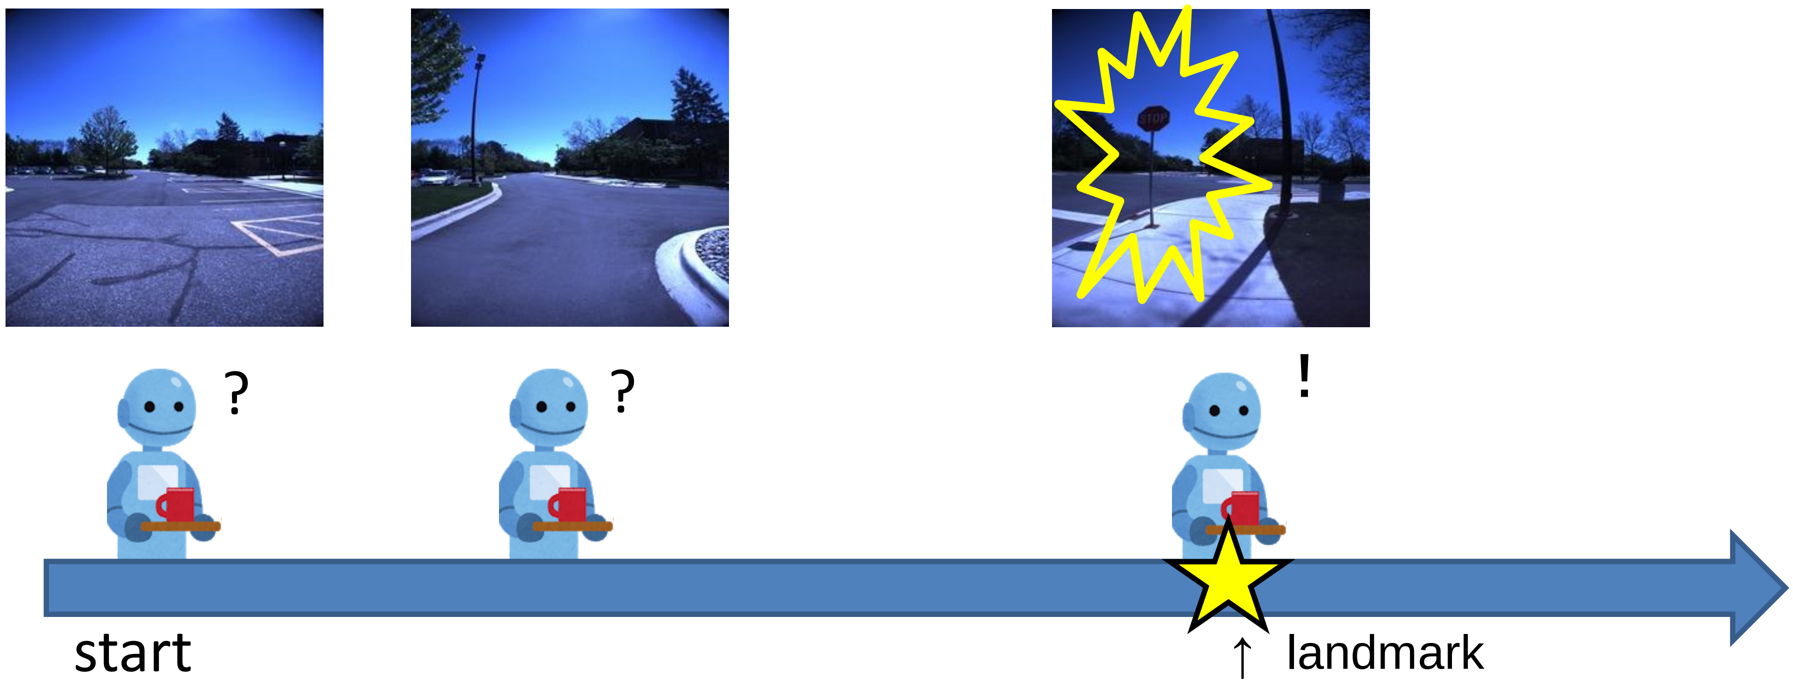

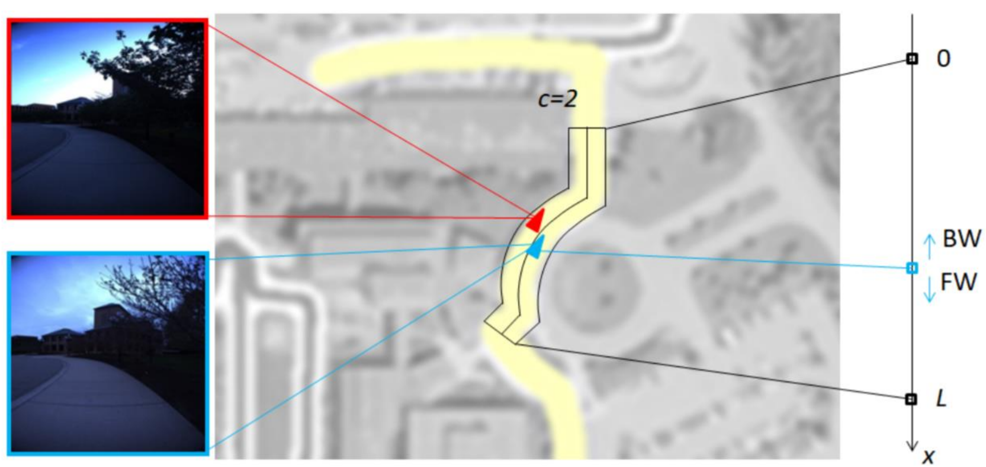

Difficulty: domain shift
Typical planners may be confused and deteriorated
low-level image features (e.g., color, texture, shape) -based planenr
mid-level image features (e.g., objects, landmarks, GIST) -based planner
Adapting these planners to a new domain
requires significant efforts
collect training data
retrain the classifier

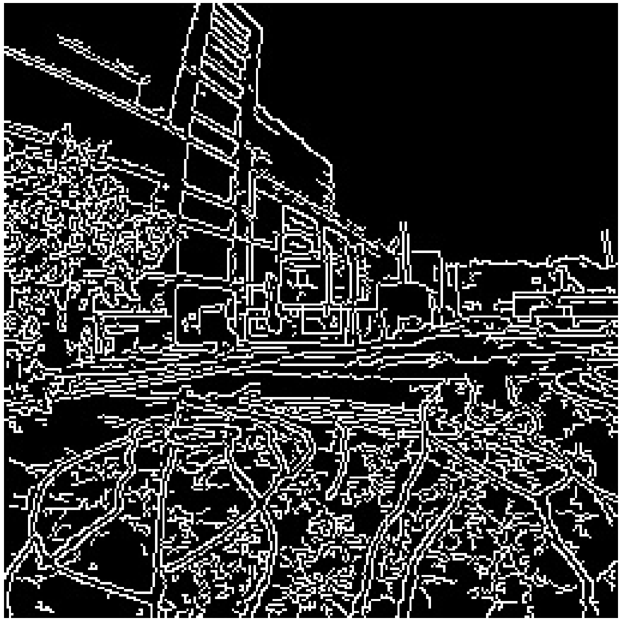

low-level image feature

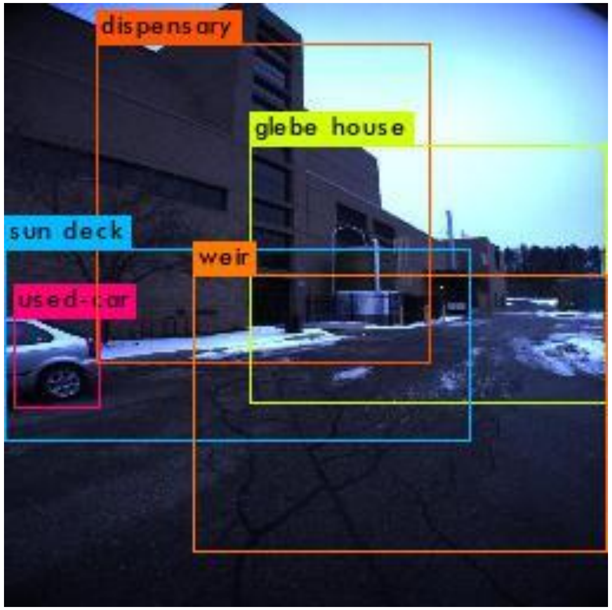

middle-level image feature

Our approach: domain-invariant NBV planner

Idea: Use output signal of DNN as planner input

Motivation: Our previous work of knowledge distillation (KD)

DNN's output provides abundant information that can be used even for another DNN

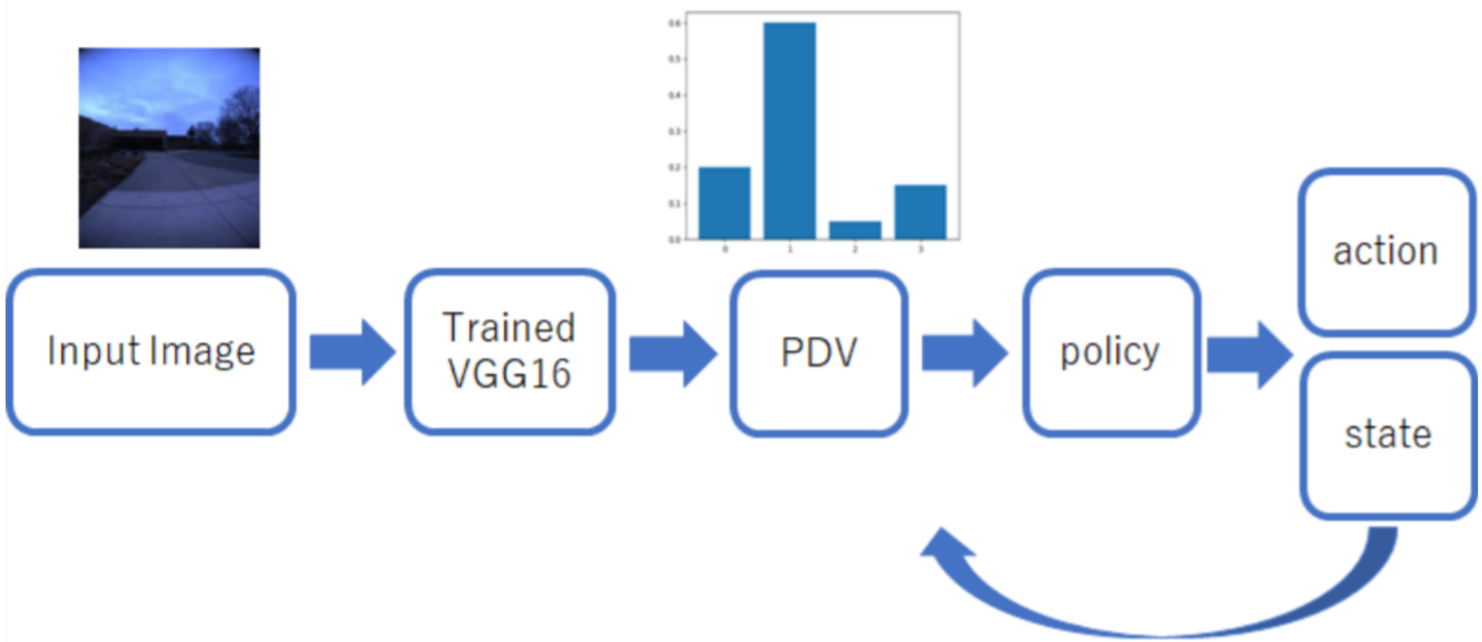

Proposal: NBV as POMDP

partially observable decision process

State: s=(x,c)

c: hidden place class

x: viewpoint wrt the place c

Main contributions: domain-invariant NBV planner in POMDP formulation

1. We address the problem of active VPR from a novel perspective of long-term VPR, in which the effect of domain shift is explicitly addressed by introducing a domain-invariant NBV planner.
2. We formulate the active VPR task as a POMDP problem and provide a feasible solution to address the inherent intractability.
3. We verify the efficacy of the proposed approach through challenging cross-season VPR experiments, focusing on outdoor scenes with no distinctive landmark objects, using the publicly available NCLT dataset.

(a) success rate [%]

	ws1	ws2	ws3
single-view	86.7	64.5	77.1
multi-view	84.5	64.5	77.2
proposed	87.8	66.7	82.8

(b) average number of observations

	ws1	ws2	ws3
single-view	1.00	1.00	1.00
multi-view	2.57	2.67	2.66
proposed	2.14	1.99	2.51

(c) average travel distance

	ws1	ws2	ws3
single-view	0.0	0.0	0.0
multi-view	79.2	89.3	78.1
proposed	34.3	34.4	40.0

Outline

- Introduction
- Method
- Implementation
- Experiments
- Conclusions

The goal of an NBV planner is to plan an optimal action-observation-sequence that is expected to maximally improve the cost performance of VPR.

(1) We assume the availability of a DNN-based VPR that is pretrained in a self-supervised manner.

(2) We also assume that, during an active VPR task, the robot is located in one of the routes or places.

(3) Finally, we assume the availability of a scoring function, which takes as input an image sequence produced by an action-observation sequence, applies each t-th image to the DNN, and aggregates the PDV sequence {PDVt}t into a final decision in the form of a K-dimensional score vector.

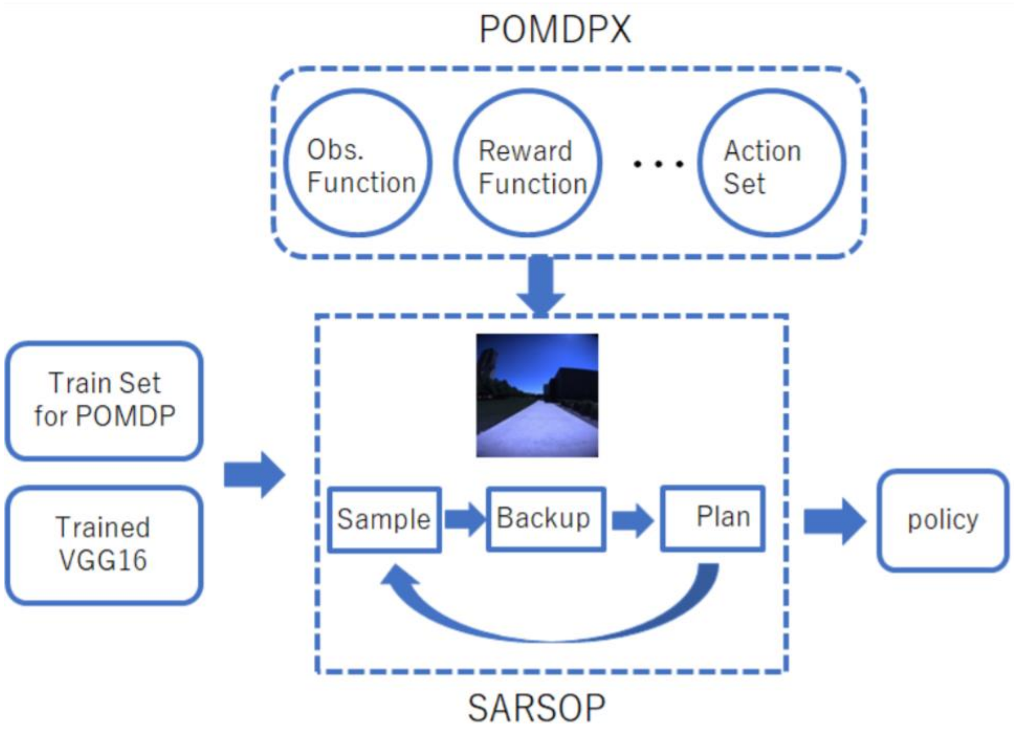

Outline

- Introduction
- Method
- Implementation
- Experiments
- Conclusions

Active VPR Formulation

Following the definition by Kaelbling [12], we define POMDP as a six-tuple (S,A,T,R,W,O),

S : a finite set of states,

A : a finite set of actions,

T : S × A × S → [0,1] is the transition function defining the probability of state change caused by an action,

R : S × A → R is the reward function that represents the reward granted after having reached the new state with the given action,

W : a finite set of observations,

O : S × A × W → [0,1] is the probability distribution of the observations according to the states and the actions.

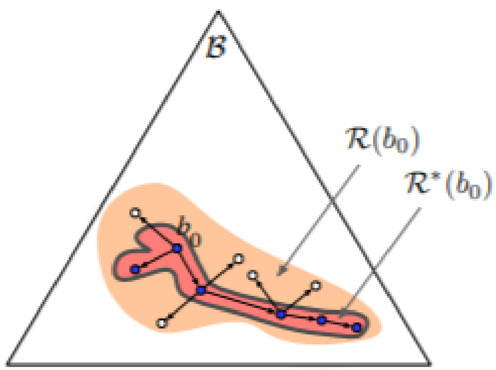

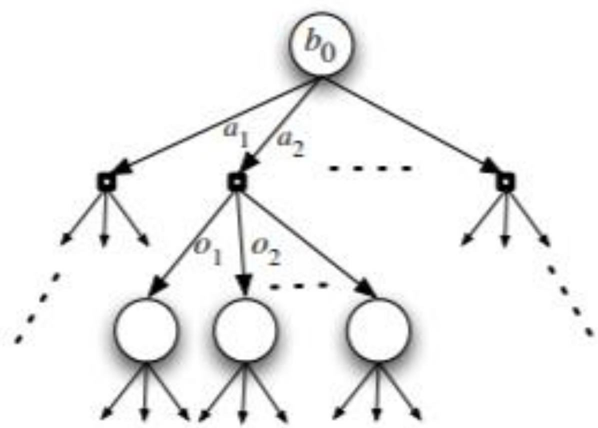

REFERENCES↓

SARSOP: Efficient Point-Based POMDP Planning by Approximating Optimally Reachable Belief Spaces

Action Set

This includes three forward moves (FWs, i.e., “a1,” “a3,” and “a5”),
three backward moves (BWs, i.e., “a2,” “a4,” and “a6”),
and an additional special action (TN, i.e., “a7”) that terminates the active VPR and outputs the most likely place class as the final answer

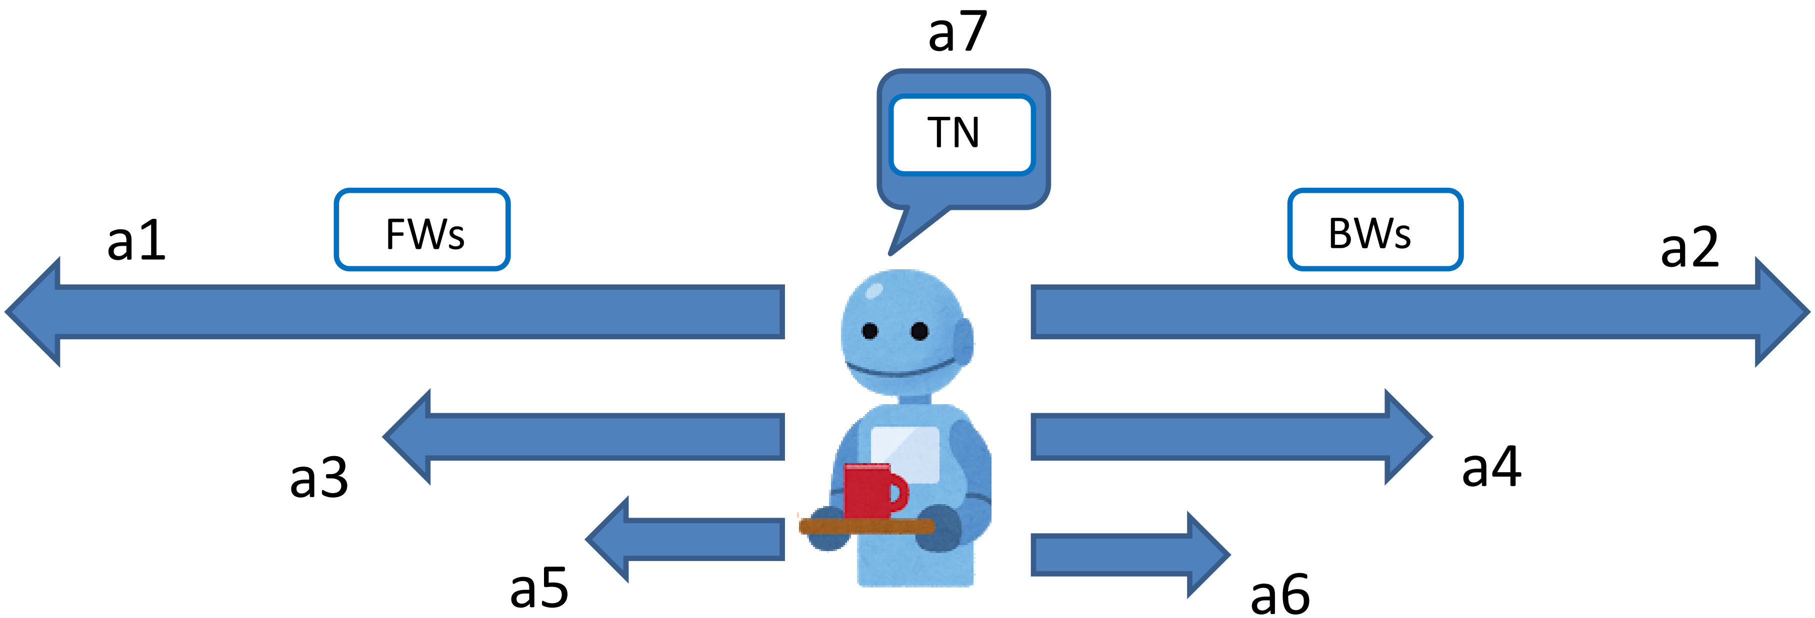

Reward Function

The reward function R(s,a) represents the reward for reaching a new state s with a given action a.

We define rewards for several different events.

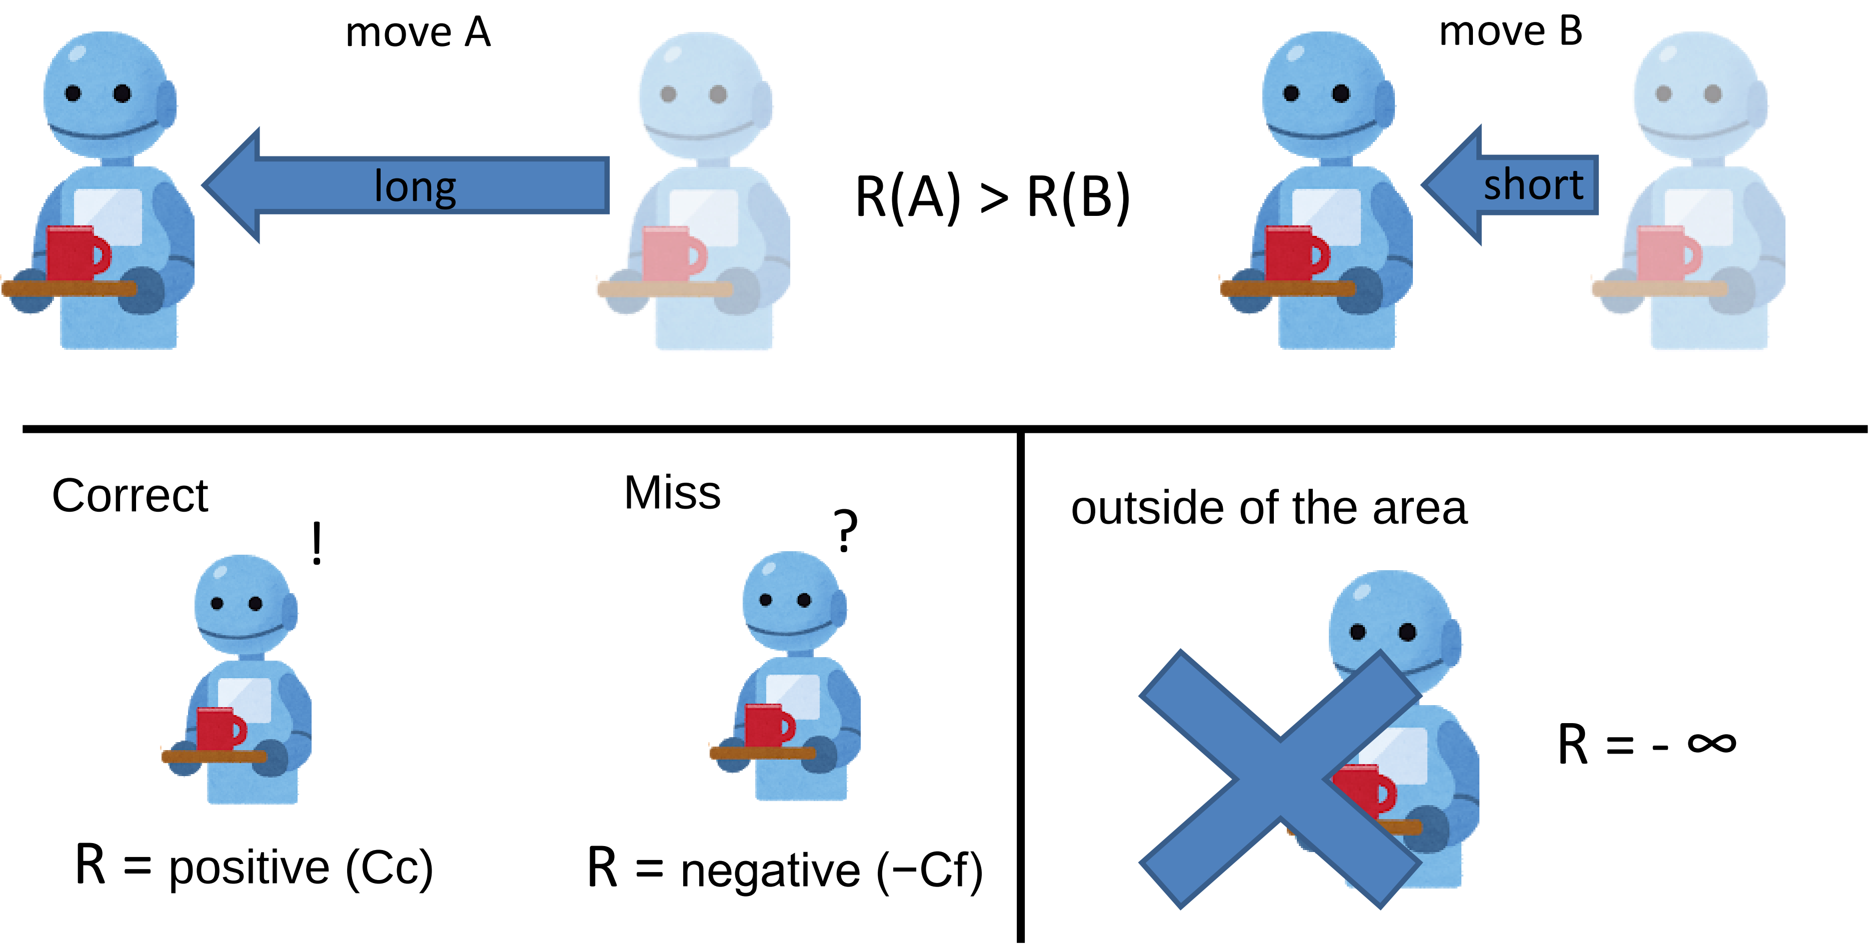

Observation Function

The observation set is represented as W.

As mentioned previously, we are motivated to use the PDV output by a DNN as a domain-invariant observation model.

A PDV is a K-dimensional real-valued vector, and its elements sum up to 1.

For the DNN, we use a convolutional neural network with the VGG16 architecture that has been fine-tuned on our K-class VPR task in the training domain.

We relax the assumption that the viewpoint x is measurable and consider multiple hypotheses for x.
In our approach, this is simply realized by hypothesizing the initial viewpoint.
At each time step t, the H different actions outputted by the H planners are aggregated into a single action plan, which is the final decision that is executed by the real robot controller.

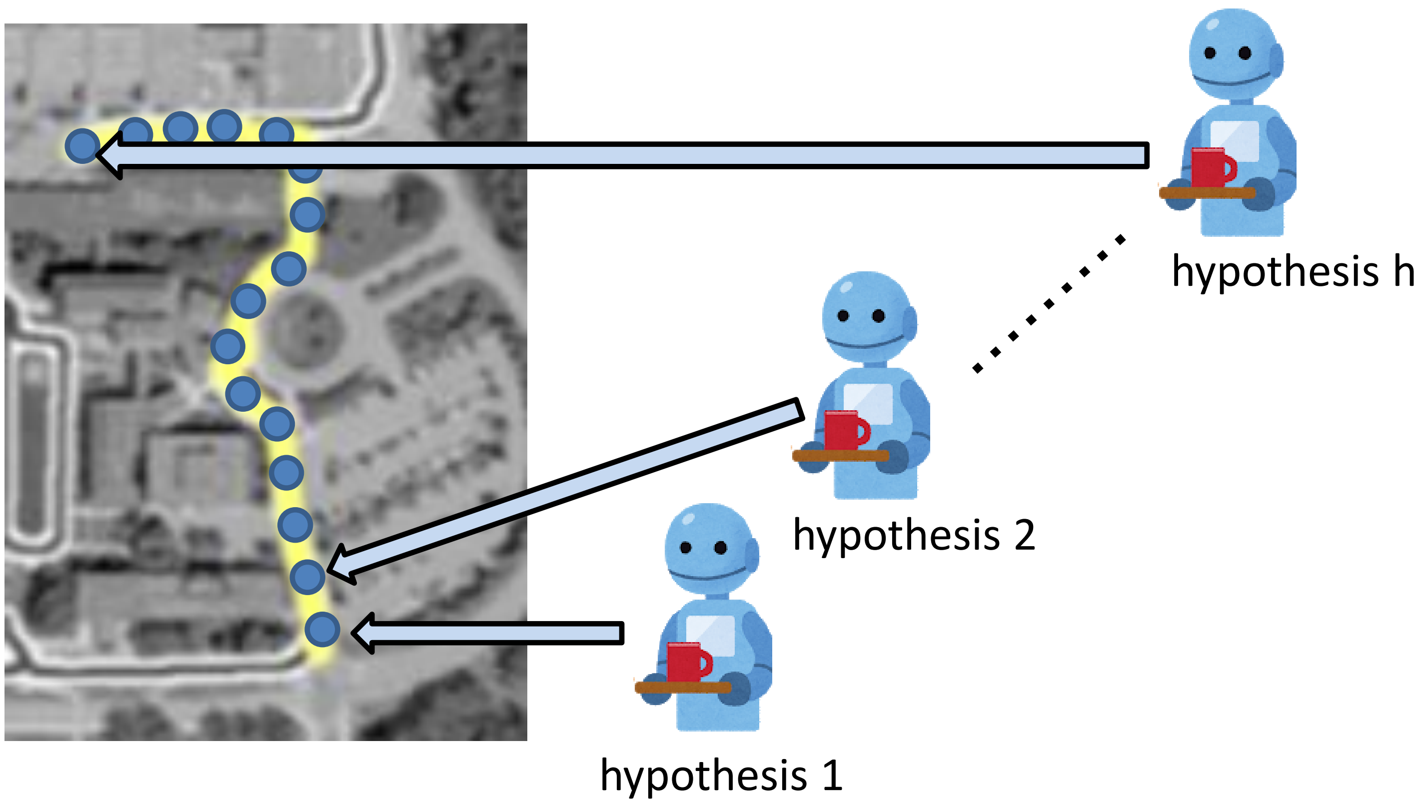

Outline

- Introduction
- Method
- Implementation
- <u>Experiments</u>
- Conclusions

Datasets

We verified the efficacy of the proposed method through active VPR experiments.

For the dataset, we considered the NCLT dataset.

We considered three independent workspaces (WSs), as shown in Fig. 4, using the outdoor parts of the datasets that contain no distinctive landmark objects.

All the three WSs were almost of the same length in terms of travel distance, and each WS was divided into four equal-length subsequences, each of which defines a route or a place class.

In this way, we considered three independent sets of four-class classification problems.

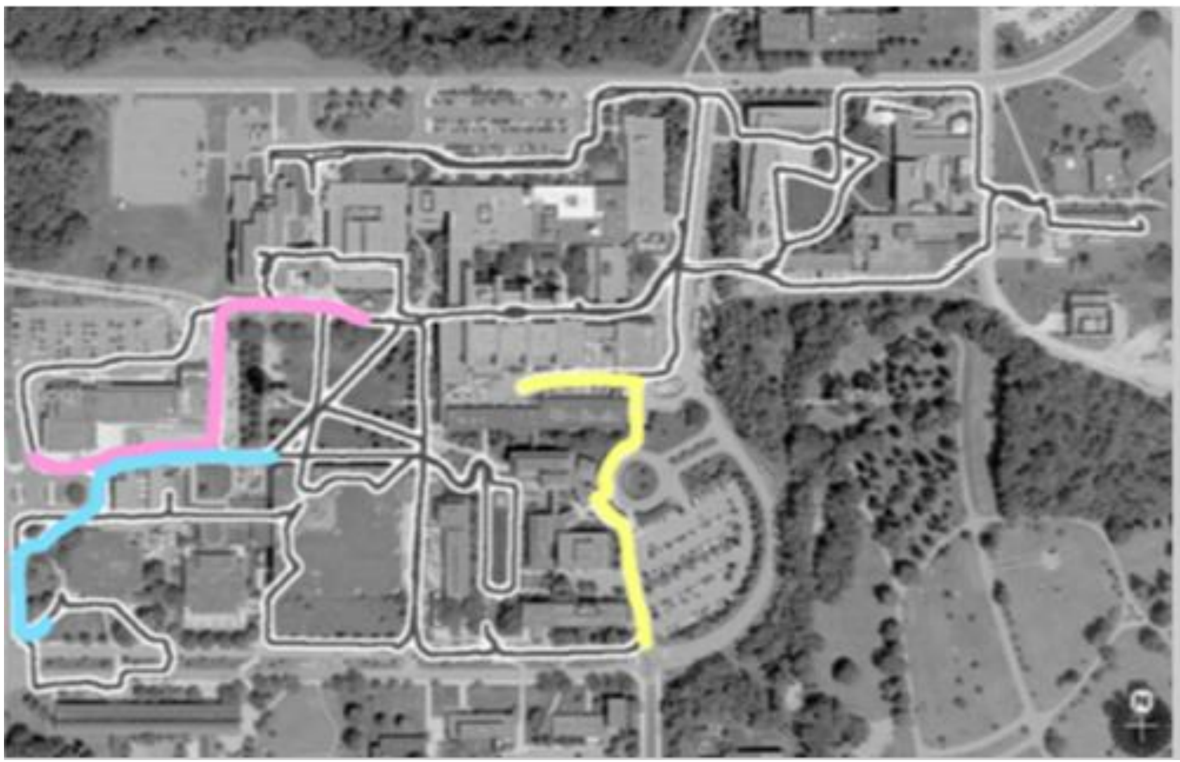

Fig. 4

Experiments

For performance comparison, we considered three comparison methods.

The first method (i.e., “single-view”) :

the conventional passive single-view VPR, which uses just a single image at the initial viewpoint as the sole query input.

The second method (i.e., “multi-view”) :

a multi-view VPR, which randomly selects the next action among the seven possible actions.

The last one (i.e., “proposed”) :

the proposed multi-view VPR method, which is based on the pretrained POMDP model and the PDV-based observation model.

(a) success rate [%]

	ws1	ws2	ws3
single-view	86.7	64.5	77.1
multi-view	84.5	64.5	77.2
proposed	87.8	66.7	82.8

(b) average number of observations

	ws1	ws2	ws3
single-view	1.00	1.00	1.00
multi-view	2.57	2.67	2.66
proposed	2.14	1.99	2.51

(c) average travel distance

	ws1	ws2	ws3
single-view	0.0	0.0	0.0
multi-view	79.2	89.3	78.1
proposed	34.3	34.4	40.0

Fig. 5 shows examples of the image sequences acquired at the planned viewpoints.
Note that, for the single-view method, only the first (i.e., leftmost) images were used as the input to the DNN. Comparing the two multi-view methods, we can see that the proposed method successfully selected visual salient scenes that could help distinguish the correct class from the other ones.

For the Fig. 5 (a) example, a relatively distinctive scene with a mountain in the background, trees, leftmost bench in the foreground, and relatively broad part of the road was selected as the second view by the proposed NBV planner.

Fig. 5 (a)

Outline

- Introduction
- Method
- Implementation
- Experiments
- <u>Conclusions</u>

We have presented a DNN-based NBV planner that incorporates the output PDVs as a domain-invariant observation model.
An algorithm is developed for planning an optimal action observation-sequence to maximize the expected cost performance of a visual route classification task.
We also formulate the active VPR as a POMDP problem and present a feasible solution to address the inherent intractability.