A Study on Autonomous Motion Planning of Mobile Robot by Use of Deep Reinforcement Learning for Fall Prevention in Hospital

Guanqun Yang

University of California, Los Angeles
Department of Electrical and Computer Engineering

Japan-US-Canada Advanced Collaborative Education Program, 2018

UCLA

Outline

UCLA

Background

- Number of elderly-involved accidents is growing.
  - Most are related to unexpected falls.
- Elderly care service becomes unaffordable.
  - Average elderly care service could cost more than $6844 per month (as of 2016).

1. Suitable for repetitive duties
2. Higher efficiency and accuracy

Insight

Mobile robot could take a role in elderly care service, specially fall prevention.

UCLA

Background

- Number of elderly-involved accidents is growing.
  - Most are related to unexpected falls.
- Elderly care service becomes unaffordable.
  - Average elderly care service could cost more than $6844 per month (as of 2016).

1. Suitable for repetitive duties
2. Higher efficiency and accuracy

Insight

Mobile robot could take a role in elderly care service, specially fall prevention.

Background

- Number of elderly-involved accidents is growing.
  - Most are related to unexpected falls.
- Elderly care service becomes unaffordable.
  - Average elderly care service could cost more than \$6844 per month (as of 2016).

1. Suitable for repetitive duties
2. Higher efficiency and accuracy

Insight

Mobile robot could take a role in elderly care service, specially fall prevention.

UCLA

Background

- Number of elderly-involved accidents is growing.
  - Most are related to unexpected falls.
- Elderly care service becomes unaffordable.
  - Average elderly care service could cost more than \$6844 per month (as of 2016).

1. Suitable for repetitive duties
2. Higher efficiency and accuracy

Insight

Mobile robot could take a role in elderly care service, specially fall prevention.

UCLA

Background

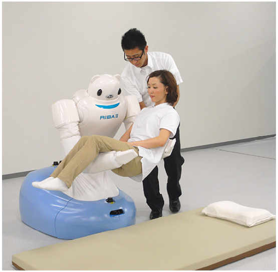

- Number of elderly-involved accidents is growing.
  - Most are related to unexpected falls.
- Elderly care service becomes unaffordable.
  - Average elderly care service could cost more than $6844 per month (as of 2016).

1. Suitable for repetitive duties
2. Higher efficiency and accuracy

Insight

Mobile robot could take a role in elderly care service, specially fall prevention.

UCLA

Background

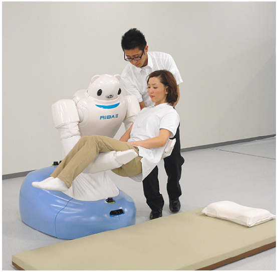

- Number of elderly-involved accidents is growing.
  - Most are related to unexpected falls.
- Elderly care service becomes unaffordable.
  - Average elderly care service could cost more than $6844 per month (as of 2016).

1. Suitable for repetitive duties
2. Higher efficiency and accuracy

Insight

Mobile robot could take a role in elderly care service, specially fall prevention.

UCLA

Objective

- Detect a safe route
- Lead the user to his/her destination

Result: Secure Path
Input: Environment Map
while *map not fully explored* **do**
    exploring environment and detecting dangers;
    if *danger detected* **then**
        DangerLevel = DangerEvaluation();
        DangerSpotAnnotation(DangerLevel);
    else
        SafetySpotAnnotation();
    end
end
SecurePath = MotionPlanning();

- Involving indoor mapping, object detection and path planning.
- Individually accomplished by DL, SLAM and RL.

UCLA

Objective

- Detect a safe route
- Lead the user to his/her destination

Result: Secure Path
Input: Environment Map
while *map not fully explored* **do**
    exploring environment and detecting dangers;
    if *danger detected* **then**
        DangerLevel = DangerEvaluation();
        DangerSpotAnnotation(DangerLevel);
    else
        SafetySpotAnnotation();
    end
end
SecurePath = MotionPlanning();

- Involving indoor mapping, object detection and path planning.
- Individually accomplished by DL, SLAM and RL.

Objective

- Detect a safe route
- Lead the user to his/her destination

Result: Secure Path
Input: Environment Map
while *map not fully explored* **do**
    exploring environment and detecting dangers;
    if *danger detected* **then**
        DangerLevel = DangerEvaluation();
        DangerSpotAnnotation(DangerLevel);
    else
        SafetySpotAnnotation();
    end
end
SecurePath = MotionPlanning();

- Involving indoor mapping, object detection and path planning.
- Individually accomplished by DL, SLAM and RL.

UCLA

Objective

- Detect a safe route
- Lead the user to his/her destination

Result: Secure Path
Input: Environment Map

while *map not fully explored* **do**
  exploring environment and detecting dangers;
  if *danger detected* **then**
    DangerLevel = DangerEvaluation();
    DangerSpotAnnotation(DangerLevel);
  else
    SafetySpotAnnotation();
  end
end
SecurePath = MotionPlanning();

- Involving indoor mapping, object detection and path planning.
- Individually accomplished by DL, SLAM and RL.

UCLA

Objective

- Detect a safe route
- Lead the user to his/her destination

```
Result: Secure Path
Input: Environment Map
while map not fully explored do
    exploring environment and detecting dangers;
    if danger detected then
        DangerLevel = DangerEvaluation();
        DangerSpotAnnotation(DangerLevel);
    else
        SafetySpotAnnotation();
    end
end
SecurePath = MotionPlanning();
```

- Involving indoor mapping, object detection and path planning.
- Individually accomplished by DL, SLAM and RL.

Hospital Scene Dataset Preparation

- The requisite of applying object detection network to our application.
- No application-specific dataset is available.
- Create dataset from scratch.

Table: Statistics of Dataset

Number of Categories	Number of Images	Image Size
29	2900	256×256

UCLA

Hospital Scene Dataset Preparation

- The requisite of applying object detection network to our application.
- No application-specific dataset is available.
- Create dataset from scratch.

Table: Statistics of Dataset

Number of Categories	Number of Images	Image Size
29	2900	256×256

UCLA

Hospital Scene Dataset Preparation

- The requisite of applying object detection network to our application.
- No application-specific dataset is available.
- Create dataset from scratch.

Table: Statistics of Dataset

Number of Categories	Number of Images	Image Size
29	2900	256×256

UCLA

Hospital Scene Dataset Preparation

- The requisite of applying object detection network to our application.
- No application-specific dataset is available.
- Create dataset from scratch.

Table: Statistics of Dataset

Number of Categories	Number of Images	Image Size
29	2900	256×256

UCLA

Hospital Scene Dataset Preparation

- The requisite of applying object detection network to our application.
- No application-specific dataset is available.
- Create dataset from scratch.

Table: Statistics of Dataset

Number of Categories	Number of Images	Image Size
29	2900	256×256

UCLA

Hospital Scene Dataset Preparation

- The requisite of applying object detection network to our application.
- No application-specific dataset is available.
- Create dataset from scratch.

Table: Statistics of Dataset

Number of Categories	Number of Images	Image Size
29	2900	256×256

Object Detection Based on YOLO

Detection Framework	mAP	FPS
Fast R-CNN	70.0	0.5
Faster R-CNN ResNet	76.4	5
Faster R-CNN VGG-16	73.2	7
SSD300	74.3	46
SSD500	76.8	19
YOLO 256×256	69.0	91

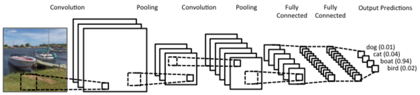

- A real-time detection system with high accuracy
- Fine tune the network to suit our purpose

Network Fine-tuning Procedure

1. Size and similarity of customized dataset
2. Large dataset: fine tuning entire network regardless of similarity
   Small dataset: only allowed when similarity is low

UCLA

Object Detection Based on YOLO

Detection Framework	mAP	FPS
Fast R-CNN	70.0	0.5
Faster R-CNN ResNet	76.4	5
Faster R-CNN VGG-16	73.2	7
SSD300	74.3	46
SSD500	76.8	19
YOLO 256 $\times$ 256	69.0	91

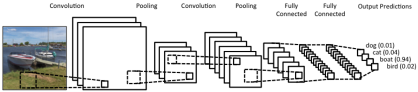

- A real-time detection system with high accuracy
- Fine tune the network to suit our purpose

Network Fine-tuning Procedure

1. Size and similarity of customized dataset
2. Large dataset: fine tuning entire network regardless of similarity
   Small dataset: only allowed when similarity is low

UCLA

Object Detection Based on YOLO

Detection Framework	mAP	FPS
Fast R-CNN	70.0	0.5
Faster R-CNN ResNet	76.4	5
Faster R-CNN VGG-16	73.2	7
SSD300	74.3	46
SSD500	76.8	19
YOLO 256×256	69.0	91

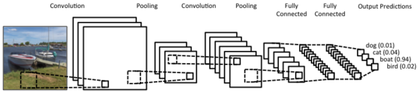

- A real-time detection system with high accuracy
- Fine tune the network to suit our purpose

Network Fine-tuning Procedure

1. Size and similarity of customized dataset
2. Large dataset: fine tuning entire network regardless of similarity
Small dataset: only allowed when similarity is low

Object Detection Based on YOLO

Detection Framework	mAP	FPS
Fast R-CNN	70.0	0.5
Faster R-CNN ResNet	76.4	5
Faster R-CNN VGG-16	73.2	7
SSD300	74.3	46
SSD500	76.8	19
YOLO 256×256	69.0	91

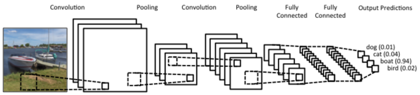

- A real-time detection system with high accuracy
- Fine tune the network to suit our purpose

Network Fine-tuning Procedure

1. Size and similarity of customized dataset
2. Large dataset: fine tuning entire network regardless of similarity
   Small dataset: only allowed when similarity is low

UCLA

Experimental Result

- Fine-tuned object detection system could detect items previously not applicable.
- Real time detection in video stream

UCLA

Experimental Result

- Fine-tuned object detection system could detect items previously not applicable.
- Real time detection in video stream

Experimental Result

- Fine-tuned object detection system could detect items previously not applicable.
- Real time detection in video stream

Reinforcement Learning-Based Motion Planning

- Trial and error learning through interaction to achieve optimal action sequence

Start	-1	0
0	-1	0
0	0	10

Problem State

- Markov Decision Process ($\langle S, A, P, R, s_0 \rangle$)
  1. $s_t \in S$, s_0 and $a_t \in A$: state, initial state and action
  2. $P(s_{t+1}|s_t, a_t)$: system dynamics
  3. $R(s_{t+1}|s_t, a_t)$: reward

- Could be solved with dynamic programming (DP)

Result: Optimal action $\pi(s)$ at each state s

for $k = 1 : \infty$ **do**

$V_k[s] = \max_a \sum_{s'} P(s'|s,a)R(s\prime|s,a) + \gamma V_{k-1}[s']$;

if $\forall s, |V_k(s) - V_{k-1}(s)| < \epsilon$ **then**

$\pi(s) = \text{argmax}_a \sum_{s'} P(s'|s,a)R(s\prime|s,a) + \gamma V_{k-1}[s']$;

end

end

UCLA

Reinforcement Learning-Based Motion Planning

- Trial and error learning through interaction to achieve optimal action sequence

Start	-1	0
0	-1	0
0	0	10

Problem State

- Markov Decision Process ($\langle S, A, P, R, s_0 \rangle$)
  1. $s_t \in S$, s_0 and $a_t \in A$: state, initial state and action
  2. $P(s_{t+1}|s_t, a_t)$: system dynamics
  3. $R(s_{t+1}|s_t, a_t)$: reward

- Could be solved with dynamic programming (DP)

Result: Optimal action $\pi(s)$ at each state s

for $k = 1 : \infty$ **do**

  $V_k[s] = \max_a \sum_{s'} P(s'|s, a)R(s\prime|s, a) + \gamma V_{k-1}[s']$;

  if $\forall s, |V_k(s) - V_{k-1}(s)| < \epsilon$ **then**

    $\pi(s) = \text{argmax}_a \sum_{s'} P(s'|s, a)R(s\prime|s, a) + \gamma V_{k-1}[s']$;

  end

end

UCLA

Reinforcement Learning-Based Motion Planning

- Trial and error learning through interaction to achieve optimal action sequence

Start	-1	0
0	-1	0
0	0	10

Problem State

- Markov Decision Process ($\langle S, A, P, R, s_0 \rangle$)
  1. $s_t \in S$, s_0 and $a_t \in A$: state, initial state and action
  2. $P(s_{t+1}|s_t, a_t)$: system dynamics
  3. $R(s_{t+1}|s_t, a_t)$: reward

- Could be solved with dynamic programming (DP)

Result: Optimal action $\pi(s)$ at each state s

for $k = 1 : \infty$ **do**

  $V_k[s] = \max_a \sum_{s'} P(s'|s,a)R(s'|s,a) + \gamma V_{k-1}[s']$;

  if $\forall s, |V_k(s) - V_{k-1}(s)| < \epsilon$ **then**

    $\pi(s) = \text{argmax}_a \sum_{s'} P(s'|s,a)R(s'|s,a) + \gamma V_{k-1}[s']$;

  end

end

UCLA

Reinforcement Learning-Based Motion Planning

- Trial and error learning through interaction to achieve optimal action sequence

Start	-1	0
0	-1	0
0	0	10

Problem State

- Markov Decision Process ($\langle S, A, P, R, s_0 \rangle$)
  1. $s_t \in S$, s_0 and $a_t \in A$: state, initial state and action
  2. $P(s_{t+1}|s_t, a_t)$: system dynamics
  3. $R(s_{t+1}|s_t, a_t)$: reward

- Could be solved with dynamic programming (DP)

Result: Optimal action $\pi(s)$ at each state s

for $k = 1 : \infty$ **do**

  $V_k[s] = \max_a \sum_{s'} P(s'|s,a)R(s\prime|s,a) + \gamma V_{k-1}[s']$;

  if $\forall s$, $|V_k(s) - V_{k-1}(s)| < \epsilon$ **then**

    $\pi(s) = \text{argmax}_a \sum_{s'} P(s'|s,a)R(s\prime|s,a) + \gamma V_{k-1}[s']$;

  end

end

UCLA

Experimental Result

- Providing an annotated map
  - Complicated danger distribution
- One secure route could be autonomously detected.

UCLA

Experimental Result

- Providing an annotated map
  - Complicated danger distribution
- One secure route could be autonomously detected.

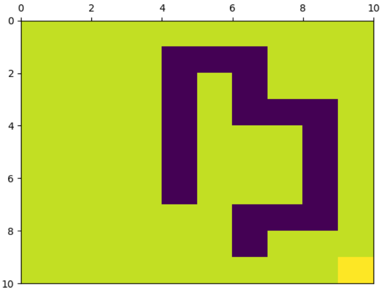

UCLA

Experimental Result

- Providing an annotated map
  - Complicated danger distribution
- One secure route could be autonomously detected.

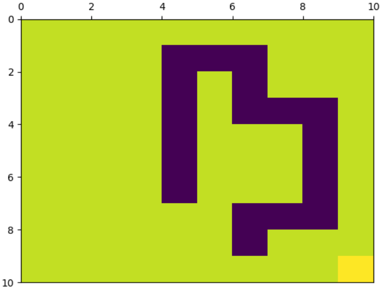

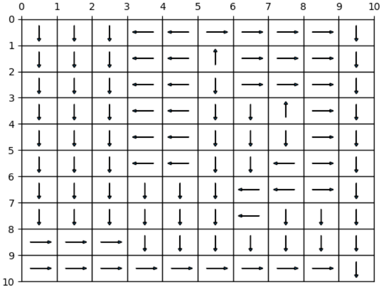

Summary

What Has Been Done

- A hospital scene-specific dataset
- Application of YOLO in hospital scene
- Previously environment from expert ⇒
  Autonomous and precise perception of environment

Future Work

- Collecting more data for dataset
- Integration of entire system

UCLA

Summary

What Has Been Done

- A hospital scene-specific dataset
- Application of YOLO in hospital scene
- Previously environment from expert $\Rightarrow$
  Autonomous and precise perception of environment

Future Work

- Collecting more data for dataset
- Integration of entire system

UCLA

Thank you for your listening!

Question and Answer