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

Guanqun Yang

Department of Electrical and Computer Engineering, University of California, Los Angeles
guanqun.yang@engineering.ucla.edu

Supervisor: Yoji Yamada

Graduate School of Engineering, Nagoya Universitys
yoji-yamada@mech.nagoya-u.ac.jp

ABSTRACT

The number and severity of elderly-involved accidents, especially falls, have posed great challenge to this aging society. Economical in cost and versatile in ability, mobile robots' role could not be neglected in elderly care service. In this paper, we put forward a novel architecture of autonomous motion planning which integrates object detection and reinforcement learning. This architecture makes possible secure path generation from environment prone to danger factors to elderly population in real time. Simulation results provides a visualization of each component and indicates effectiveness of this system.

1 INTRODUCTION

The United Nations data indicates that the number of elderly people could exceed 2.1 billion in 2050 and 3.1 billion in 2100 with growth rate higher than that of all other age groups [1]. At the same time, two tendencies are observed by academic community regarding aging population. One is the number of elderly-involved accidents, especially unexpected falls, is growing rapidly and could cause as many as 7 mortalities every single hour in United States by 2030 [2]. Another one is the steady increase of elderly-care service, which on average costs $6844 for a single month in United States each person [3].

Meanwhile, the versatility and cost-efficiency of mobile robots also draw researchers' attention. With their participation in elderly care service, the number of accidents and cost of service are expected to be considerably reduced. Furthermore, the novel deep learning and reinforcement learning theory could be directly applied to elderly care service through their implementation on these robotics platforms.

In this work, we present an architecture that detects safety route autonomously in hospital setting with mobile robots through deep reinforcement learning. This paper is organized as follows: part 1 shows the challenge and potential solution people have in elderly care service; part 2 provides an overview of our system through algorithmic diagram; part 3 shows the process we acquire the application-specific dataset; part 4 and part 5 briefly introduces ideas underlying and provides experimental results and corresponding evaluations.

2 SYSTEM ARCHITECTURE

From the introduction in part 1, the major objects of fall prevention is two-fold:

1. Detect a secure route from indoor map through deep reinforcement learning
2. Lead the user to his or her destination according to the route detected in step 1

The task in step 2 is largely a solved problem with multiple working implementations [4, 5]. Therefore, we mainly focus on step 1.

As is seen from the algorithm 1, when fed into the map of environment, the mobile robot continuously explore the entire environment and makes safety spot annotation or danger spot annotation through deep learning-based object detection and danger level evaluation. When all blocks are annotated, path planning based on reinforcement learning is enacted and thereby a safety route. Note multiple routes are possible after route detection so path could chosen by user based on their preferences.

3 DATASET PREPARATION

Preparing an dataset immediately related to our hospital scene setting is the requisite for utilizing out-of-state object detection

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();

Algorithm 1: Process of Secure Route Detection

Figure 1: The Hospital Scene Dataset With 29 Categories

framework for danger factor detection. However, to our best knowledge, such dataset is not available even though there does exist some datasets regarding surgery process and other medical activities where specific operations to patients are recorded [6, 7]. Our best hope is that common danger factors in hospitals and elderly care facilities could be included in our dataset.

By combining famous Places205 dataset and public domain Google image search results and after dataset preparation procedures described in algorithm 2[8], the dataset with 29 categories of danger factors is extracted. As is shown in Table 3, the number of each category of images is 100 and the image size is 256×256, which is suitable for the fine-tuning of convolution neural network. A snapshot of this dataset is shown in Figure 1.

Table 1: Statistics of Dataset

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

Result: Hospital Scene Dataset
Input: CategoryList, Places205 Dataset, Google Search Results
for $k = 1 : ImageCount$ **do**
    category = IdentifyingCategory(image);
    if $category \in CategoryList$ **then**
        labeling image with category identifier $(1, 2, \cdots, 29)$
    else
        lebeling image with 0;
    end
end
for $k = 1 : ImageCount$ **do**
    automatic resizing;
    automatic categorization with category identifier;
end

Algorithm 2: Procedures of Dataset Preparation

4 OBJECT DETECTION

YOLO is a object detection implementation that outperforms others regarding speed with comparable mean average precision (mAP). As is shown in Table 2, the platform could process video stream with as many as 91 frames per second (FPS) [9, 10, 11].

Table 2: Comparison of Out-of-State Object Detection Platform

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

However, the YOLO was original trained on PASCAL VOC 2007 dataset with limited number of categories [12]. Therefore, the fine-tuning of convolutional neural network (CNN) used in YOLO system is required.

Table 3: Categories in PASCAL VOC 2007 Dataset

Type	Category
Person	person
Animal	bird, cat, cow, dog, horse, sheep
Vehicle	aeroplane, bicycle, boat, bus, car, motorbike, train
Indoor	bottle, chair, dining table, potted plant, sofa, tv/monitor

Since our dataset is small and different from the PASCAL VOC 2007 dataset the YOLO was originally trained on, the fine-tuning of the network is allowed. Specifically, only layers that extract lower-level features including lines, arcs and circles are kept and the resulting network is trained on our dataset with linear classifier as output.

As is shown in Figure 2, the fine-tuned network could detect bed, which is previously unavailable.

Figure 2: Result of Fine-Tuned YOLO Detection System

5 MOTION PLANNING

Reinforcement learning is a process that enables agent to learn a sequence of actions through interaction with environment in a trial and error fashion. It could be formalized as a Markov decision process (MDP) and parameterized as a tuple $\langle S, A, P, R, s_0 \rangle$ [13]. Specifically, S is the state set where the agent could transition to; A is the action set including a list of allowable actions that could be taken by agent; $P(s_{t+1}|s_t, a_t)$ is the system dynamic that describes the probability of making transition from s_t to s_{t+1} by taking action $s_t \in S$; and $R(s_{t+1}|s_t, a_t)$ is the reward of making such transition.

The solution to the MDP could be described in algorithm 3, which is a dynamic programming-based approach. At each iteration, the agent chooses the action that generates most reward until there are no more rewards could be exploited from environment.

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

Algorithm 3: Algorithm to Solve MDP

In one complicated setting shown in Figure 3, where dark areas represent danger spots and bright area in down-right corner represent destination, the mobile robot could autonomously decide optimal action in each state, as is shown in Figure 4. More detailedly, the mobile robot tries to avoid the danger spots and the actions taken near these areas are all going down.

Based on the result provided by Figure 4, a secure route could be detected and the autonomous motion planning is made possible.

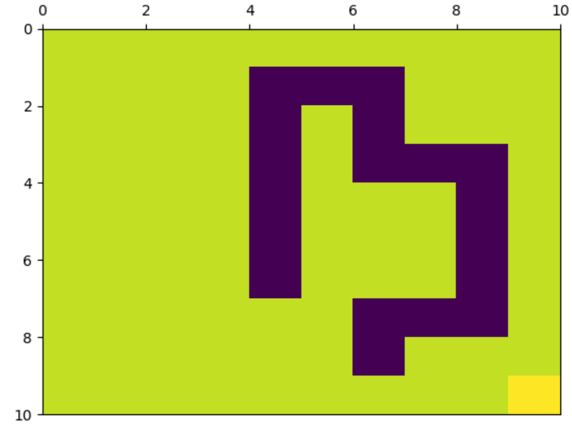

Figure 3: Simulation Setting

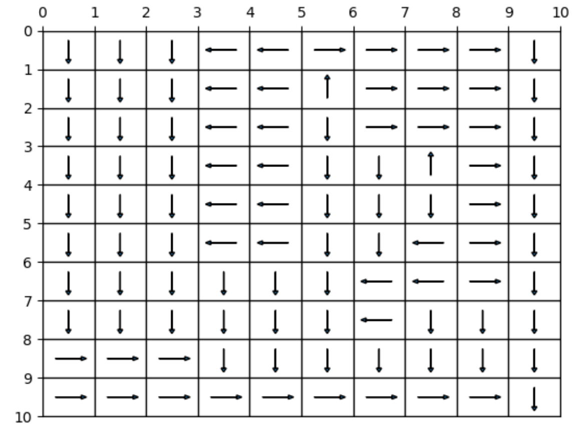

Figure 4: Simulation Result

6 CONCLUSION

In this work we put forward a novel architecture for secure route generation in the background of fall prevention in hospital or other elderly care facilities through deep reinforcement learning. A hospital scene dataset is first created to suit our purpose that applying our system in hospital setting. Then fine-tuned object detection system allows real time danger factor detection from live video stream. Finally, the reinforcement learning-based motion planning module makes possible the secure route generation in complicated hospital setting.

At the same time, we notice that the number of danger factors is still lacking when attempting to apply the system in real world. Therefore, more categories of images should be collected and added to our dataset. Additionally, the system integration of indoor mapping, object detection and motion planning is largely still a undone task and may involve numerous practical implementation issues. In the future, our work will continue in both of the two aspects.

ACKNOWLEDGEMENT

This work could not be made possible without the support coming from multiple sources. First we appreciate the JUACEP program and its coordination with several renowned institutions, including UCLA, in north America. The program and international cooperation is the requisite for this research experience. Then this work would be otherwise difficult, if not impossible, to complete in a smooth fashion without the invaluable support from Yamada lab, where Mr. Takaaki Namba, Mr. Eugene Kim and Prof. Yoji Yamada provides timely guidance in not only theoretical aspect but experimental aspect as well. Eventually, the generous financial support provided by JASSO guarantees a memorable research experience in Japan.

REFERENCES

[1] United Nations. World population prospects: The 2017 revision. *United Nations Econ Soc Aff*, 33(2):1–66, 2017.

[2] Centers for Disease Control and Prevention. Important facts about falls, 2016.

[3] Depart of Health and Human Services. Costs of care, 2016.

[4] Sebastian Thrun, Maren Bennewitz, Wolfram Burgard, Armin B Cremers, Frank Dellaert, Dieter Fox, Dirk Hahnel, Charles Rosenberg, Nicholas Roy, Jamieson Schulte, et al. Minerva: A second-generation museum tour-guide robot. In *Robotics and automation, 1999. Proceedings. 1999 IEEE international conference on*, volume 3. IEEE, 1999.

[5] Takayuki Kanda, Masahiro Shiomi, Zenta Miyashita, Hiroshi Ishiguro, and Norihiro Hagita. An affective guide robot in a shopping mall. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 173–180. ACM, 2009.

[6] Andru P Twinanda, Sherif Shehata, Didier Mutter, Jacques Marescaux, Michel De Mathelin, and Nicolas Padoy. Endonet: A deep architecture for recognition tasks on laparoscopic videos. *IEEE transactions on medical imaging*, 36(1):86–97, 2017.

[7] Ralf Stauder, Daniel Ostler, Michael Kranzfelder, Sebastian Koller, Hubertus Feußner, and Nassir Navab. The tum lapchole dataset for the m2cai 2016 workflow challenge. *arXiv preprint arXiv:1610.09278*, 2016.

[8] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. In *Advances in neural information processing systems*, pages 487–495, 2014.

[9] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.

[10] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. *arXiv preprint*, 2017.

[11] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.

[12] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010.

[13] Seyed Sajad Mousavi, Michael Schukat, and Enda Howley. Deep reinforcement learning: an overview. In *Proceedings of SAI Intelligent Systems Conference*, pages 426–440. Springer, 2016.