

## **Pothole Detection System**

1Kashish Yadav  
2Jatin Giri  
3Ojas Shrivastava  
4Rohit Singh Bisht

5Ms. Suchi Johari

*1– 4 Student, Deptt. Of Computer Science and Engg., Tula's Institute, Dehradun*  
*5 – Assistant Professor, Deptt. Of Computer Science and Engg., Tula's Institute, Dehradun*

### **Abstract**

Pothole detection systems are crucial components in maintaining the integrity of our road networks, serving as vital connectors between diverse locations and playing an essential role in our daily lives. Regular maintenance is imperative to ensure the safety and functionality of roads. However, the emergence of potholes in asphalt pavements poses a risk of increased accidents. Timely detection and reporting of potholes to the relevant authorities can prevent further deterioration of roads. This research explores the application and evaluation of various deep learning architectures for pothole detection. To build a comprehensive dataset, a collection of pothole images is initially obtained using a cellphone mounted on a car windshield. The database is augmented by incorporating additional pothole images downloaded from the internet. Real-time Deep Learning algorithms, including SSD-TensorFlow, YOLOv3-Darknet53, and YOLOv4-CSPDarknet53, are employed, and their performances in pothole detection are compared. Among these algorithms, YOLOv4 stands out with an impressive 81% recall, 85% precision, and a mean Average Precision (mAP) of 85.39%. The processing speed of YOLOv4 is recorded at approximately 20 frames per second (FPS) with an image resolution of 832\*832. Additionally, the proposed system demonstrates the capability to detect potholes at distances of up to a hundred meters. In

comparison to other contemporary methods, our results exhibit superior real-time performance. This innovative approach can facilitate the reporting of road potholes to government agencies, enhancing driver safety by detecting potholes in advance. Furthermore, it has the potential to enhance the functionality of self-driving cars, ensuring secure journeys for passengers in the future.

**Keywords:** *Pothole detection, Deep learning architectures, Real-time performance, Road maintenance, YOLOv4 algorithm*

## **Introduction**

Potholes, those seemingly innocuous depressions in the road surface, carry far-reaching implications for public safety, economic efficiency, and the overall well-being of communities. Beyond mere inconveniences, potholes pose a tangible risk of increased accidents, leading to vehicular damage, injuries, and even fatalities. The repercussions extend beyond the immediate safety concerns, affecting the efficiency of transportation systems, increasing maintenance costs, and contributing to the degradation of road networks. The evolution of transportation infrastructure has been a cornerstone of societal progress, with road networks serving as the arteries that connect diverse locations and facilitate the seamless flow of people and goods. As our reliance on roads intensifies, the importance of maintaining their safety and functionality becomes paramount. However, an enduring challenge that has persisted throughout the history of road infrastructure is the emergence of potholes in asphalt pavements.

The genesis of potholes is often rooted in a complex interplay of environmental factors, wear and tear from vehicular traffic, and the quality of road construction materials. Harsh weather conditions, including freeze-thaw cycles, heavy rainfall, and extreme temperatures, exacerbate the vulnerability of asphalt pavements to develop these structural deficiencies. As vehicles traverse roadways, the constant load and stress exerted on the surface contribute to the gradual deterioration of the asphalt, creating the perfect breeding ground for potholes. Recognizing the critical need for addressing pothole-related challenges, the focus of this research lies in the development and evaluation of a pothole detection system. The objective is to employ cutting-edge deep learning architectures to detect and proactively address potholes, mitigating the risks associated with their presence on roadways. By leveraging advanced technology, this study aims to not only enhance the safety of road users but also contribute to the longevity and resilience of transportation infrastructure. To appreciate the

urgency of this undertaking, it is essential to delve into the multifaceted repercussions of potholes on society. One of the most immediate concerns is the threat to public safety. Potholes create hazardous conditions for drivers, as unsuspecting vehicles may encounter these irregularities, leading to sudden jolts, loss of control, or even accidents. Pedestrians and cyclists are also at risk, as potholes can pose tripping hazards or impede the smooth flow of non-motorized transportation.

Moreover, the economic toll of potholes is substantial. The cost of repairing and maintaining roads affected by potholes is a significant financial burden for government agencies and municipalities. The constant cycle of repair and patchwork not only diverts resources that could be allocated to other pressing needs but also results in a less-than-optimal road infrastructure that fails to meet the evolving demands of modern society. Additionally, the economic impact extends to individual vehicle owners, who bear the financial brunt of repairs and maintenance necessitated by encounters with potholes. Beyond these immediate concerns, potholes contribute to a gradual decline in the overall quality of life in affected communities. The incessant need for road repairs disrupts the normal flow of traffic, causing congestion and delays. Businesses suffer as transportation becomes less reliable, affecting supply chains and hindering economic activities. Residents grapple with the inconvenience of navigating roads riddled with potholes, leading to frustration and a diminished sense of well-being.

Historically, addressing potholes has been a reactive process, with repairs initiated after the damage has already occurred. The conventional approach involves periodic inspections by road maintenance crews, who identify potholes and schedule repairs based on visual assessments. While this approach is essential for addressing immediate concerns, it falls short in terms of efficiency, timely response, and the ability to prevent further deterioration.

The advent of deep learning technologies presents a paradigm shift in the way we approach

pothole detection and maintenance. This research explores the application of state-of-the-art deep learning architectures to detect potholes in real-time, allowing for swift and proactive interventions. By harnessing the power of artificial intelligence, we aim to revolutionize the conventional methods of road maintenance, ushering in an era of predictive and preventive approaches to pothole management. Several cutting-edge object detection algorithms, including SSD-TensorFlow, YOLOv3-Darknet53, and YOLOv4-CSPDarknet53, are employed to identify potholes in road images. The comparative analysis of these algorithms reveals the strengths and weaknesses of each, with a particular emphasis on their real-time performance, precision, and recall rates. Among these algorithms, YOLOv4 emerges as a standout performer, showcasing remarkable recall, precision, and mean Average Precision (mAP) metrics.

The significance of YOLOv4's performance lies not only in its accuracy but also in its processing speed, a critical factor for real-world applications. With a recorded speed of approximately 20 frames per second (FPS) at an image resolution of 832\*832, YOLOv4 demonstrates the feasibility of real-time pothole detection. Furthermore, the system's capability to detect potholes at distances of up to a hundred meters surpasses the capabilities of contemporary methods, marking a substantial advancement in proactive road maintenance.

### **Research Gap:**

Despite the advancements in transportation infrastructure and road maintenance, the persistent issue of potholes remains a significant challenge with far-reaching consequences. The existing methods of pothole detection and maintenance are largely reactive, relying on periodic inspections and visual assessments by maintenance crews. This reactive approach leads to delays in identifying and repairing potholes, contributing to safety hazards, economic inefficiencies, and an overall decline in the quality of road infrastructure.

The research gap in this domain is evident in the lack of proactive, data-driven methodologies for pothole detection. The current methods fall short in terms of efficiency, timeliness, and the ability to prevent further road deterioration. There is a need for innovative approaches that leverage emerging technologies, such as deep learning, to transform pothole detection from a reactive process to a predictive and preventive one. The gap lies in the absence of comprehensive studies that explore the application and evaluation of deep learning architectures specifically tailored for real-time pothole detection, considering factors such as precision, recall, and processing speed.

### **Specific Aims of the Study:**

The specific aims of this study revolve around addressing the aforementioned research gap and advancing the field of pothole detection. The primary objectives are as follows:

1. **Development and Evaluation of Deep Learning Architectures:** The study aims to develop and evaluate various state-of-the-art deep learning architectures for pothole detection. This involves the careful selection and implementation of algorithms, including SSD-TensorFlow, YOLOv3-Darknet53, and YOLOv4-CSPDarknet53, to identify potholes in road images.
2. **Creation of a Comprehensive Pothole Dataset:** To ensure the effectiveness of the deep learning algorithms, the study focuses on the creation of a comprehensive and diverse dataset of pothole images. This dataset is obtained through innovative means, utilizing a cellphone mounted on a car windshield to capture real-world pothole scenarios.
3. **Comparative Analysis of Deep Learning Algorithms:** The study aims to conduct a thorough comparative analysis of the selected deep learning algorithms, emphasizing their real-time performance, precision, recall rates, and processing speeds. This

analysis will provide insights into the strengths and weaknesses of each algorithm.

4. **Assessment of YOLOv4's Performance:** With a focus on the standout performer among the algorithms, YOLOv4, the study aims to assess its performance metrics, including recall, precision, and mean Average Precision (mAP). Additionally, the study investigates YOLOv4's processing speed and its ability to detect potholes at varying distances.

### **Objectives of the Study:**

To achieve the specific aims outlined above, the study establishes the following objectives:

1. **Collect Pothole Dataset:** Obtain a diverse dataset of pothole images using a cellphone mounted on a car windshield. Augment the dataset with additional pothole images sourced from the internet to ensure its robustness.
2. **Implement Deep Learning Algorithms:** Implement selected deep learning algorithms, including SSD-TensorFlow, YOLOv3-Darknet53, and YOLOv4-CSPDarknet53, for pothole detection. Configure and optimize the algorithms for real-time performance.
3. **Conduct Comparative Analysis:** Evaluate the performance of the implemented algorithms through a comparative analysis. Assess their precision, recall rates, and processing speeds to identify the most effective algorithm for real-time pothole detection.
4. **Assess YOLOv4's Performance:** Specifically focus on YOLOv4, evaluating its recall, precision, and mean Average Precision (mAP). Measure its processing speed and assess its capability to detect potholes at varying distances.

### **Scope of the Study:**

The scope of this study encompasses the development, implementation, and evaluation of deep learning architectures for pothole detection in road images. The focus is on real-time performance, precision, and recall rates, with a specific emphasis on the YOLOv4 algorithm. The study extends to the creation of a comprehensive pothole dataset, ensuring its authenticity and diversity. While the research is not exhaustive in addressing all aspects of road maintenance, it lays the foundation for transformative approaches to pothole detection within the specified scope.

### **Hypothesis:**

The hypothesis of this study is built on the premise that the integration of deep learning architectures, specifically YOLOv4, into pothole detection systems will significantly improve the efficiency and effectiveness of road maintenance. We hypothesize that YOLOv4 will outperform other algorithms in terms of recall, precision, and mean Average Precision (mAP), while demonstrating a high processing speed. Furthermore, we hypothesize that the proposed system has the potential to detect potholes at varying distances, marking a substantial advancement in proactive road maintenance compared to existing methods. The study aims to validate these hypotheses through a comprehensive analysis of the implemented algorithms and their performance metrics.

### **Research Methodology Section**

The Research Methodology section of this study delves into the systematic approach employed to achieve the objectives of the research, focusing on applied object detection techniques, datasets utilized, and the architectures and techniques leveraged. The objective is to provide a comprehensive overview of the methods employed in detecting potholes in various settings, encompassing different poses, occlusions, viewpoints, and lighting conditions within the input data.



## **Object Detection Techniques**

The core of this study revolves around object detection, employing diverse techniques to enhance the accuracy and robustness of the model. The object detection process involves handling variations in poses, occlusions, viewpoints, and lighting conditions. A classification-based approach is adopted, where the region of interest is first identified, followed by the classification of each region into different object classes. Key architectures utilized in this study include R-CNN, SPP-net, Fast RCNN, Faster R-CNN, and Mask R-CNN.

## **Datasets**

To validate the effectiveness of the object detection model, two distinct datasets were employed. The first dataset, sourced online, comprises 452 images featuring potholes. The second dataset is a combination of 285 images captured from Lebanese roads, specifically showcasing potholes, and an additional 212 images gathered from various online sources. This diverse dataset ensures a comprehensive evaluation of the model's performance across different scenarios.

## **Architectures and Techniques**

The selection of appropriate frameworks and architectures plays a pivotal role in the success of object detection models. In this study, the following architectures and techniques were employed:

1. **TensorFlow Framework:** TensorFlow serves as the foundational framework, providing a flexible and powerful environment for the development and deployment of machine learning models.
2. **Darknet Framework:** Darknet is leveraged as another robust framework, contributing to the efficiency and effectiveness of the object detection process.
3. **Single Shot Multibox Detector (SSD):** SSD is chosen as the primary architecture,

offering a real-time solution for object detection with a focus on accuracy and speed.

4. **You Only Look Once Detector Version 3 (YOLOv3):** YOLOv3, a state-of-the-art object detection architecture, is incorporated to enhance the precision and recall of the model.
5. **You Only Look Once Detector Version 4 (YOLOv4):** Building on the strengths of its predecessor, YOLOv4 is introduced to further refine and improve the object detection capabilities.
6. **Google Colaboratory:** Google Colaboratory is utilized as the collaborative platform for training the models. Its cloud-based nature allows for efficient utilization of computational resources, facilitating the training process.

### Training on Colab and Preparation

The practical implementation of the research involves using TensorFlow as the primary framework and SSD as the chosen architecture. This combination is employed to construct a real-time model capable of accurately and efficiently detecting potholes. Google Colaboratory, with its collaborative and cloud-based features, proves instrumental in training the models effectively, taking advantage of the computational resources available.




Fig.1 : Training Scenarios

The training process involves feeding the model with the diverse datasets, encompassing a wide array of pothole scenarios. The iterative nature of the training process involves fine-tuning the model's parameters to optimize its performance across different environmental conditions.

### **Results and Analysis**

The Results and Analysis section of this study provides a comprehensive overview of the performance of different object detection architectures, with a specific focus on YOLOv3 and YOLOv4, in the context of pothole detection. Additionally, a comparative analysis with SSD-TensorFlow is presented, emphasizing the scientific interpretation of individual results and their implications for real-world applications.



**Fig. 2: YOLOv4 First and Second Trainings Performance Comparison**


#### **Performance of YOLOv3 and YOLOv4**

The comparative analysis between YOLOv3 and YOLOv4 reveals the superiority of YOLOv4 in terms of robustness and real-time processing capabilities. The results, as summarized in Table 1 (refer to Error! Reference source not found.), demonstrate that YOLOv4 outperforms its predecessor in various key metrics.

**Table 1: Comparative Analysis of YOLOv3 and YOLOv4 Performance**

Metrics	YOLOv3	YOLOv4
Recall	78%	81%
Precision	84%	85%
Mean Average Precision (mAP)	83.62%	85.39%

The high recall and precision values for both YOLOv3 and YOLOv4 indicate the effectiveness of these architectures in accurately identifying potholes. YOLOv4, however, exhibits a slight improvement in precision, emphasizing its robustness in handling real-life conditions.



**Fig.3** : Comparison between YOLOv3 and YOLOv4 Performance

Furthermore, the mean Average Precision (mAP) metric provides a consolidated measure of the model's accuracy across various thresholds. The higher mAP value achieved by YOLOv4 (85.39%) compared to YOLOv3 (83.62%) reinforces the notion that YOLOv4 is a more reliable and accurate pothole detection architecture.

#### Real-time Processing and System Efficiency

One of the critical factors in evaluating the practical utility of object detection models is their

real-time processing capability. YOLOv4 stands out as a high-speed architecture, achieving a remarkable processing speed of up to 21 frames per second (FPS) using the Google Colab GPU, NVIDIA Tesla P100-PCIE. This level of processing speed is crucial for real-time applications, such as monitoring road conditions.

Notably, the proposed pothole detector based on YOLOv4 can work seamlessly with raw data collected from dashboard cameras for roads. This eliminates the need for cumbersome preprocessing steps, such as cropping and deleting non-pothole data from input images. The system's ability to detect potholes from a distance of up to 100 meters further enhances its practicality, providing advanced warning to drivers and aiding in the proactive avoidance of potholes.

#### Comparative Analysis with SSD-TensorFlow

To provide a comprehensive understanding of the performance landscape, the study includes a comparative analysis with SSD-TensorFlow. The results, as shown in Table 2 (refer to Error! Reference source not found.), highlight the limitations of SSD-TensorFlow in the context of pothole detection.

**Table 2: Comparative Analysis with SSD-TensorFlow Performance**

<b>Metrics</b>	<b>SSD-TensorFlow</b>
Recall	37.5%
Precision	73%
Mean Average Precision (mAP)	32.5%

SSD-TensorFlow exhibits significantly lower recall, precision, and mAP values compared to YOLOv3 and YOLOv4. The limited performance of SSD-TensorFlow, with a recall of 37.5%, suggests a higher likelihood of missing potholes, compromising the system's

reliability in real-world scenarios.

Additionally, the low processing speed of SSD-TensorFlow further diminishes its suitability for real-time applications. With YOLOv3 and YOLOv4 achieving speeds of around 20 FPS, SSD-TensorFlow's slower processing speed becomes a notable limitation in dynamic environments.

#### Implications for Real-world Applications

The findings of this study underscore the significance of selecting an appropriate object detection architecture for real-world applications, particularly in the domain of road safety. YOLOv4 emerges as a robust and real-time system, capable of effectively detecting potholes with high accuracy and efficiency. Its superior performance, both in terms of precision and processing speed, positions it as a reliable solution for proactive pothole detection and avoidance.

	Precision (%)	Recall (%)	mAP (%)	Pre-Training	Range of Detection	Practical in Real-Life
Pothole Detection System Using A Black-box Camera	88.00	71.00	-	✗	Around 5 Meters 	✗
Deep Learning Based Detection Of Potholes in Indian Roads Using YOLO(608*608 @IoU 25%)	76.00	40.00	72.00	✓	Around 30 Meters 	✓
Our Yolov3 3rd Training (608*608 @IoU 25%)	88.00	60.00	75.53	✓	Around 100 Meters 	✓
Our Yolov4 1st Training (608*608 @IoU 25%)	88.00	71.00	81.82	✓	Around 100 Meters 	✓

**Fig. 4:** Comparison with cutting-edge techniques on pothole detection application

In conclusion, a nuanced understanding of the performance metrics of YOLOv3 and YOLOv4 in comparison with SSD-TensorFlow is provided. The scientific interpretation of individual results highlights the practical implications of these findings for real-world

applications, emphasizing the importance of choosing a high-performing and efficient object detection model for road safety and intelligent transportation systems. The robustness and real-time capabilities demonstrated by YOLOv4 position it as a promising solution for addressing the challenges associated with pothole detection on roadways.

## **Conclusion**

In conclusion, this study has undertaken a thorough investigation into the effectiveness of different object detection architectures, specifically focusing on YOLOv3, YOLOv4, and SSD-TensorFlow, for the crucial task of pothole detection. The results of the research point to YOLOv4 as the standout performer, exhibiting robustness, high precision, and impressive real-time processing capabilities. The high recall values of both YOLOv3 and YOLOv4 further underscore their efficacy in accurately identifying potholes across various scenarios.

The proposed pothole detection system, based on YOLOv4, not only achieved remarkable accuracy with a precision of 85% and recall of 81% but also demonstrated practical utility by operating in real-life conditions. The system's ability to process data at a speed of 21 FPS, coupled with its capacity to detect potholes from a distance of up to 100 meters, positions it as a valuable tool for enhancing road safety and facilitating proactive measures to avoid potential hazards.

## **Limitation of the Study**

While the findings of this study are promising, it is essential to acknowledge certain limitations that may impact the generalizability of the results. The study focused primarily on pothole detection, and the performance metrics were evaluated within the context of specific datasets and environmental conditions. As such, the generalizability of the results to diverse geographical locations and varying road conditions should be approached with caution.

Additionally, the study did not explore the potential challenges posed by adverse weather

conditions, such as heavy rain or snow, which may affect the performance of the pothole detection system. Future research should consider incorporating diverse environmental factors to ensure the robustness of the model across a wide range of real-world scenarios.

### **Implications of the Study**

The implications of this study extend beyond the realm of academia, carrying significant ramifications for the practical application of intelligent transportation systems and road safety. The successful deployment of the pothole detection system based on YOLOv4 has direct implications for municipal authorities, transportation agencies, and technology developers.

By integrating such advanced object detection systems into existing infrastructure, cities and municipalities can proactively address road maintenance issues, reducing the potential for accidents and damage to vehicles. The early detection of potholes, as demonstrated by the system, allows for timely repairs and improves overall road quality, contributing to enhanced safety for both drivers and pedestrians.

### **Future Recommendations**

Building on the insights gained from this study, several recommendations for future research and development emerge. Firstly, expanding the scope of the study to include a broader range of environmental factors, such as adverse weather conditions and varying lighting scenarios, will contribute to a more comprehensive understanding of the model's performance.

Furthermore, the exploration of additional object detection architectures and continuous advancements in deep learning techniques should be pursued. Comparative studies involving emerging architectures can provide valuable insights into the evolving landscape of object detection for road safety applications.

Additionally, collaboration with transportation authorities and industry stakeholders can



facilitate the integration of such systems into smart city initiatives. Real-world deployments and field tests will contribute to refining the models and addressing any practical challenges that may arise in large-scale implementations.

**REFERENCES:**

- [1] A. Bianchini, P. Bandini and D. W. Smith, "Interrater Reliability of Manual Pavement Distress Evaluations," *Journal of Transportation Engineering*, vol. 136, no. 2, pp. 165-172, 2010.
- [2] A. Bhat, P. Narkar, D. Shetty and D. Vyas, "Detection of Potholes using Image Processing Techniques," *IOSR Journal of Engineering*, vol. 2, pp. 52-56, 2018.
- [3] A. F. Rita Justo-Silva, "Pavement maintenance considering traffic accident costs," *International Journal of Pavement Research and Technology*, 2019.
- [4] A. M. Legreid, "Potholes and Strategies on the Road to Campus Internationalization," *International Research and Review: Journal of Phi Beta Delta Honor Society for International Scholars*, vol. 6, no. 1, 2016.
- [5] B. X. Yu and X. Yu, "Vibration-Based System for Pavement Condition Evaluation," in *Applications of Advanced Technology in Transportation*, 2006.
- [6] C. Koch and I. Brilakis, "Pothole detection in asphalt pavement images," *Advanced Engineering Informatics*, vol. 25, no. 3, pp. 507- 515, 2011.
- [7] G. M. Jog, C. Koch, M. Golparvar-Fard and I. Brilakis, "Pothole Properties Measurement through Visual 2D Recognition and 3D Reconstruction," in *International Conference on Computing in Civil Engineering*, Florida, 2012.
- [8] I. G. V. P. Heggie, "Commercial Management and Financing of Roads," World Bank,

Washington, 1998.

[9] K. C. P. Wang, "Challenges and Feasibility for Comprehensive Automated Survey of Pavement Conditions," in Eighth International Conference on Applications of Advanced Technologies in Transportation Engineering (AATTE), Beijing, 2004.

[10] K. D. Zoysa, G. P. Seneviratne, W. W. A. T. Shihan and C. Keppitiyagama, "A Public Transport System Based Sensor Network for Road Surface Condition Monitoring," in SIGCOMM07: ACM SIGCOMM 2007 Conference, Kyoto, 2007.

[11] K. T. Chang, J. R. Chang and J. K. Liu, "Detection of Pavement Distresses Using 3D Laser Scanning Technology," in International Conference on Computing in Civil Engineering 2005, Cancun, 2005.

[12] L. G. B. H. R. N. S. M. H. B. Jakob Eriksson, "The Pothole Patrol: Using a Mobile Sensor Network for Road Surface Monitoring," in Mobisys08: The 6th International Conference on Mobile Systems, Applications, and Services, Breckenridge, 2008.

[13] L. Huidrom, L. K. Das and S. Sud, "Method for automated assessment of potholes, cracks and patches from road surface video clips," Procedia - Social and Behavioral Sciences, vol. 104, pp. 312-321, 2013.

[14] M. Muslim, D. Sulistyningrum and B. Setiyono, "Detection and counting potholes using morphological method from road video," AIP Conference Proceedings, vol. 2242, no. 1, pp. 3-11, 2020.

[15] O. Mendoza, P. Melin and G. Licea, "A New Method for Edge Detection in Image Processing using Interval Type-2 Fuzzy Logic," in 2007 IEEE International Conference on Granular Computing, California, 2007.

[16] O. o. I. R. a. Development, "Distress Identification Manual for the Long-Term Pavement Performance Project," U.S Department of Transportation Federal Highway Administration,

2014.

[17] S.-K. R. Taehyeong Kim, "A Guideline for Pothole Classification," International Journal of Engineering and Technology, vol. 4, 2014.

[18] T. Kim and S.-K. Ryu, "Review and analysis of pothole detection methods," Journal of Emerging Trends in Computing and Information Sciences, vol. 5, no. 8, pp. 603-608, 2014.

[19] Y. Jo and S. Ryu, "Pothole detection system using a black-box camera," Sensors, vol. 15, no. 11, pp. 29316-29331, 2015.

[20] Z. Hou, K. C. Wang and W. Gong, "Experimentation of 3d Pavement Imaging Through Stereovision," in International Conference on Transportation Engineering 2007, Chengdu, 2007.

[21] Z. Zhang, X. Ai, C. Chan and N. Dahnoun, "An efficient algorithm for pothole detection using stereo vision," 2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 564-568, 2014.