

Statement of Purpose

Carnegie Mellon University

Jiayi Weng (trinkle23897@gmail.com)

Contents

Ph.D. in Robotics	2
M.S. in Robotics	4
M.S. in Machine Learning	6
Master of Science in Computer Vision	8
M.S. in Computer Science	10
Master of Computational Data Science	12
Master of Science in Artificial Intelligence and Innovation	14
M.S. in Intelligent Information Systems	16

Statement of Purpose

Ph.D. in Robotics

Jiayi Weng (trinkle23897@gmail.com)

As an undergraduate, I was lucky to participate in several **reinforcement learning** projects at both Tsinghua University and MILA. This fascinating research area, with great potentials, deeply attracted me. To further pursue my interest, I am determined to apply for a Ph.D. program in computer science. I am confident that the Ph.D. training will prepare me to be an independent researcher, develop practical applications to solve challenging real-world problems, and ultimately pursue an academic career.

I first became interested in reinforcement learning when I worked under the supervision of **Professor Jun Zhu** on several VizDoom projects. I am most proud of the performance of our team in the *VizDoom AI Competition 2018*, a famous reinforcement learning challenge. Our agent **ranked 1/51**, far exceeding the score of second-place. This work¹ was ultimately **published in IJCAI 2019, of which I was the first author**. I confronted with several critical challenges in this competition: 1) Visual understanding from raw images; 2) Visual navigation in the 3D environment; 3) Sparse and delayed reward for the agent. First, I built up an object detection system aiming to efficiently identify the obstacles and enemies. Second, inspired by existing visual navigation methods, I incorporated the depth signal into the navigation network, which facilitated the agent to plot its paths. Third, I innovatively formulated the problem within a hierarchical reinforcement learning framework by dividing the challenge into subtasks. To make further improvements, I tuned the reward of the controller. Meanwhile, I accelerated both the training and inference phase and iterated the algorithm every day.

With these achievements on the VizDoom challenge, I was eager to start the independent project *Relational Markov Decision Process*, which intended to accelerate the deep reinforcement learning procedure by state abstraction (using first-order logic) and reward shaping. At first, I found the approach unsatisfying on some of VizDoom scenarios due to the instability of the training process. Therefore, I changed the network to double Deep-Q Network with duel training and adjusted the replay buffer size to improve stability. Next, I implemented parallel environment training to further speed up experiments, making the algorithm converge to the optima more rapidly. Finally, this method accelerated the convergence as expected, and was more robust and generalizable under harder games. The fantastic result far exceeded my expectations, and also strengthened my confidence in continuing research. We **submitted this work² to TNNLS in 2019**.

After these VizDoom projects, I was excited to apply what I had learned to create a conscious system, a topic I am especially interested in. This summer I led the *Rule-Transformer* project supervised by **Professor Yoshua Bengio**. We aimed to develop a

¹Project page: <https://trinkle23897.github.io/cv/viz2018.html>

²Project page: <https://trinkle23897.github.io/cv/rmdp.html>

system with *Consciousness Prior* (proposed by Professor Bengio), which integrated the advantage of both neural and symbolic models. By interpreting the attention architecture as applying rules to variables, we generalized attention to a novel mechanism called *Rule-Attention*. I designed and implemented this architecture in a short time, and performed a comprehensive investigation. The Rule-Transformer turned out to be satisfactory in a wide range of tasks from machine translation to multitask imitation learning. One of the interesting findings was that a decrease in dimension would encourage the Rule-Attention layer to activate more functions to deliver a similar performance. This is related to how we define a function and an argument to interpret a black concept box. The Rule-Transformer project gave me a glance at the fascinating nature of machine learning and the complicated challenges faced by the community. Meanwhile, the experience of working with such a passionate group of world-class researchers in MILA has equipped me with the necessary skills that will stand behind my future research.

Throughout my two-year immersion in reinforcement learning, I determined to pursue an academic career and focus on the essence and application in this field. From my understanding and experience, current reinforcement learning tasks, such as reasoning, planning and capturing causality, are conducted in an intuitive and unconscious manner with quite a high sample complexity and no interpretability, robustness, or generalization. To address these issues, I would like to research on modeling the dynamics of the world, leveraging and finding human priors for policy learning, and further enabling machine learning algorithms to learn in a sample-efficient manner. Also, based on my previous achievements in the VizDoom competition, I am interested in solving complex video games and real-world challenges with deep learning and reinforcement learning.

I am applying to the Ph.D. program at the Robotics Institute of Carnegie Mellon University due to its leading position in robotics research. There are four professors at CMU whose projects are particularly appealing to me. It is my dream to join **Professor Ruslan Salakhutdinov**'s research group. His team had a leading position in the previous VizDoom AI Competition. I especially admire his recent works of Transformer-XL and XLNet, and am thrilled in the MineRL challenge. I would like to incorporate language-based knowledge prior to tackle the aforementioned issues. **Professor David Held** and **Professor Deepak Pathak**'s projects on the intersection of robotics and computer vision fascinated me a lot. Based on robots, it would help me gain more insights into top-down visual understanding and robust reinforcement learning. Meanwhile, I am equally excited about the projects of Curiosity-driven learning and zero-shot imitation learning. I am keen to work under the supervision of **Professor Abhinav Gupta** who shares common research interest with me, since I have excellent engineering skills and have a literature view of reinforcement learning. It is of great significance to make related datasets and benchmark various existing methods in the field of robotics. I am convinced that with the help of the extraordinary minds and the top research environment at CMU, I can continue to contribute to the field of reinforcement learning.

Statement of Purpose

M.S. in Robotics

Jiayi Weng (trinkle23897@gmail.com)

As an undergraduate, I was lucky to participate in several **reinforcement learning** projects at both Tsinghua University and MILA. This fascinating research area, with great potentials, deeply attracted me. To further pursue my interest, I am determined to apply for an M.S. program in computer science. I am confident that the M.S. training will prepare me to be an independent researcher in the industry and develop practical applications to solve challenging real-world problems.

I first became interested in reinforcement learning when I worked under the supervision of **Professor Jun Zhu** on several VizDoom projects. I am most proud of the performance of our team in the *VizDoom AI Competition 2018*, a famous reinforcement learning challenge. Our agent **ranked 1/51**, far exceeding the score of second-place. This work¹ was ultimately **published in IJCAI 2019, of which I was the first author**. I confronted with several critical challenges in this competition: 1) Visual understanding from raw images; 2) Visual navigation in the 3D environment; 3) Sparse and delayed reward for the agent. First, I built up an object detection system aiming to efficiently identify the obstacles and enemies. Second, inspired by existing visual navigation methods, I incorporated the depth signal into the navigation network, which facilitated the agent to plot its paths. Third, I innovatively formulated the problem within a hierarchical reinforcement learning framework by dividing the challenge into subtasks. To make further improvements, I tuned the reward of the controller. Meanwhile, I accelerated both the training and inference phase and iterated the algorithm every day.

With these achievements on the VizDoom challenge, I was eager to start the independent project *Relational Markov Decision Process*, which intended to accelerate the deep reinforcement learning procedure by state abstraction (using first-order logic) and reward shaping. At first, I found the approach unsatisfying on some of VizDoom scenarios due to the instability of the training process. Therefore, I changed the network to double Deep-Q Network with duel training and adjusted the replay buffer size to improve stability. Next, I implemented parallel environment training to further speed up experiments, making the algorithm converge to the optima more rapidly. Finally, this method accelerated the convergence as expected, and was more robust and generalizable under harder games. The fantastic result far exceeded my expectations, and also strengthened my confidence in continuing research. We **submitted this work² to TNNLS in 2019**.

After these VizDoom projects, I was excited to apply what I had learned to create a conscious system, a topic I am especially interested in. This summer I led the *Rule-Transformer* project supervised by **Professor Yoshua Bengio**. We aimed to develop a

¹Project page: <https://trinkle23897.github.io/cv/viz2018.html>

²Project page: <https://trinkle23897.github.io/cv/rmdp.html>

system with *Consciousness Prior* (proposed by Professor Bengio), which integrated the advantage of both neural and symbolic models. By interpreting the attention architecture as applying rules to variables, we generalized attention to a novel mechanism called *Rule-Attention*. I designed and implemented this architecture in a short time, and performed a comprehensive investigation. The Rule-Transformer turned out to be satisfactory in a wide range of tasks from machine translation to multitask imitation learning. One of the interesting findings was that a decrease in dimension would encourage the Rule-Attention layer to activate more functions to deliver a similar performance. This is related to how we define a function and an argument to interpret a black concept box. The Rule-Transformer project gave me a glance at the fascinating nature of machine learning and the complicated challenges faced by the community. Meanwhile, the experience of working with such a passionate group of world-class researchers in MILA has equipped me with the necessary skills that will stand behind my future research.

Throughout my two-year immersion in reinforcement learning, I determined to pursue an academic career and focus on the essence and application in this field. From my understanding and experience, current reinforcement learning tasks, such as reasoning, planning and capturing causality, are conducted in an intuitive and unconscious manner with quite a high sample complexity and no interpretability, robustness, or generalization. To address these issues, I would like to research on modeling the dynamics of the world, leveraging and finding human priors for policy learning, and further enabling machine learning algorithms to learn in a sample-efficient manner. Also, based on my previous achievements in the VizDoom competition, I am interested in solving complex video games and real-world challenges with deep learning and reinforcement learning.

I am applying to the M.S. program at the Robotics Institute of Carnegie Mellon University due to its leading position in robotics research. There are four professors at CMU whose projects are particularly appealing to me. It is my dream to join **Professor Ruslan Salakhutdinov**'s research group. His team had a leading position in the previous VizDoom AI Competition. I especially admire his recent works of Transformer-XL and XLNet, and am thrilled in the MineRL challenge. I would like to incorporate language-based knowledge prior to tackle the aforementioned issues. **Professor David Held** and **Professor Deepak Pathak**'s projects on the intersection of robotics and computer vision fascinated me a lot. Based on robots, it would help me gain more insights into top-down visual understanding and robust reinforcement learning. Meanwhile, I am equally excited about the projects of Curiosity-driven learning and zero-shot imitation learning. I am keen to work under the supervision of **Professor Abhinav Gupta** who shares common research interest with me, since I have excellent engineering skills and have a literature view of reinforcement learning. It is of great significance to make related datasets and benchmark various existing methods in the field of robotics. I am convinced that with the help of the extraordinary minds and the top research environment at CMU, I can continue to contribute to the field of reinforcement learning.

Statement of Purpose

M.S. in Machine Learning

Jiayi Weng (trinkle23897@gmail.com)

As an undergraduate, I was lucky to participate in several **reinforcement learning** projects at both Tsinghua University and MILA. This fascinating research area, with great potentials, deeply attracted me. To further pursue my interest, I am determined to apply for an M.S. program in computer science. I am confident that the M.S. training will prepare me to be an independent researcher in the industry and develop practical applications to solve challenging real-world problems.

I first became interested in reinforcement learning when I worked under the supervision of **Professor Jun Zhu** on several VizDoom projects. I am most proud of the performance of our team in the *VizDoom AI Competition 2018*, a famous reinforcement learning challenge. Our agent **ranked 1/51**, far exceeding the score of second-place. This work¹ was ultimately **published in IJCAI 2019, of which I was the first author**. I confronted with several critical challenges in this competition: 1) Visual understanding from raw images; 2) Visual navigation in the 3D environment; 3) Sparse and delayed reward for the agent. First, I built up an object detection system aiming to efficiently identify the obstacles and enemies. Second, inspired by existing visual navigation methods, I incorporated the depth signal into the navigation network, which facilitated the agent to plot its paths. Third, I innovatively formulated the problem within a hierarchical reinforcement learning framework by dividing the challenge into subtasks. To make further improvements, I tuned the reward of the controller. Meanwhile, I accelerated both the training and inference phase and iterated the algorithm every day.

With these achievements on the VizDoom challenge, I was eager to start the independent project *Relational Markov Decision Process*, which intended to accelerate the deep reinforcement learning procedure by state abstraction (using first-order logic) and reward shaping. At first, I found the approach unsatisfying on some of VizDoom scenarios due to the instability of the training process. Therefore, I changed the network to double Deep-Q Network with duel training and adjusted the replay buffer size to improve stability. Next, I implemented parallel environment training to further speed up experiments, making the algorithm converge to the optima more rapidly. Finally, this method accelerated the convergence as expected, and was more robust and generalizable under harder games. The fantastic result far exceeded my expectations, and also strengthened my confidence in continuing research. We **submitted this work² to TNNLS in 2019**.

After these VizDoom projects, I was excited to apply what I had learned to create a conscious system, a topic I am especially interested in. This summer I led the *Rule-Transformer* project supervised by **Professor Yoshua Bengio**. We aimed to develop a

¹Project page: <https://trinkle23897.github.io/cv/viz2018.html>

²Project page: <https://trinkle23897.github.io/cv/rmdp.html>

system with *Consciousness Prior* (proposed by Professor Bengio), which integrated the advantage of both neural and symbolic models. By interpreting the attention architecture as applying rules to variables, we generalized attention to a novel mechanism called *Rule-Attention*. I designed and implemented this architecture in a short time, and performed a comprehensive investigation. The Rule-Transformer turned out to be satisfactory in a wide range of tasks from machine translation to multitask imitation learning. One of the interesting findings was that a decrease in dimension would encourage the Rule-Attention layer to activate more functions to deliver a similar performance. This is related to how we define a function and an argument to interpret a black concept box. The Rule-Transformer project gave me a glance at the fascinating nature of machine learning and the complicated challenges faced by the community. Meanwhile, the experience of working with such a passionate group of world-class researchers in MILA has equipped me with the necessary skills that will stand behind my future research.

Throughout my two-year immersion in reinforcement learning, I determined to pursue an academic career and focus on the essence and application in this field. From my understanding and experience, current reinforcement learning tasks, such as reasoning, planning and capturing causality, are conducted in an intuitive and unconscious manner with quite a high sample complexity and no interpretability, robustness, or generalization. To address these issues, I would like to research on modeling the dynamics of the world, leveraging and finding human priors for policy learning, and further enabling machine learning algorithms to learn in a sample-efficient manner. Also, based on my previous achievements in the VizDoom competition, I am interested in solving complex video games and real-world challenges with deep learning and reinforcement learning.

I am applying to M.S. in Machine Learning of Carnegie Mellon University due to its leading position in machine learning research. There are four professors at CMU whose projects are particularly appealing to me. It is my dream to join **Professor Ruslan Salakhutdinov**'s research group. His team had a leading position in the previous VizDoom AI Competition. I especially admire his recent works of Transformer-XL and XLNet, and am thrilled in the MineRL challenge. I would like to incorporate language-based knowledge prior to tackle the aforementioned issues. **Professor David Held** and **Professor Deepak Pathak**'s projects on the intersection of robotics and computer vision fascinated me a lot. Based on robots, it would help me gain more insights into top-down visual understanding and robust reinforcement learning. Meanwhile, I am equally excited about the projects of Curiosity-driven learning and zero-shot imitation learning. I am keen to work under the supervision of **Professor Abhinav Gupta** who shares common research interest with me, since I have excellent engineering skills and have a literature view of reinforcement learning. It is of great significance to make related datasets and benchmark various existing methods in the field of robotics. I am convinced that with the help of the extraordinary minds and the top research environment at CMU, I can continue to contribute to the field of reinforcement learning.

Statement of Purpose

Master of Science in Computer Vision

Jiayi Weng (trinkle23897@gmail.com)

As an undergraduate, I was lucky to participate in several **reinforcement learning** projects at both Tsinghua University and MILA. This fascinating research area, with great potentials, deeply attracted me. To further pursue my interest, I am determined to apply for an M.S. program in computer science. I am confident that the M.S. training will prepare me to be an independent researcher in the industry and develop practical applications to solve challenging real-world problems.

I first became interested in reinforcement learning when I worked under the supervision of **Professor Jun Zhu** on several VizDoom projects. I am most proud of the performance of our team in the *VizDoom AI Competition 2018*, a famous reinforcement learning challenge. Our agent **ranked 1/51**, far exceeding the score of second-place. This work¹ was ultimately **published in IJCAI 2019, of which I was the first author**. I confronted with several critical challenges in this competition: 1) Visual understanding from raw images; 2) Visual navigation in the 3D environment; 3) Sparse and delayed reward for the agent. First, I built up an object detection system aiming to efficiently identify the obstacles and enemies. Second, inspired by existing visual navigation methods, I incorporated the depth signal into the navigation network, which facilitated the agent to plot its paths. Third, I innovatively formulated the problem within a hierarchical reinforcement learning framework by dividing the challenge into subtasks. To make further improvements, I tuned the reward of the controller. Meanwhile, I accelerated both the training and inference phase and iterated the algorithm every day.

With these achievements on the VizDoom challenge, I was eager to start the independent project *Relational Markov Decision Process*, which intended to accelerate the deep reinforcement learning procedure by state abstraction (using first-order logic) and reward shaping. At first, I found the approach unsatisfying on some of VizDoom scenarios due to the instability of the training process. Therefore, I changed the network to double Deep-Q Network with duel training and adjusted the replay buffer size to improve stability. Next, I implemented parallel environment training to further speed up experiments, making the algorithm converge to the optima more rapidly. Finally, this method accelerated the convergence as expected, and was more robust and generalizable under harder games. The fantastic result far exceeded my expectations, and also strengthened my confidence in continuing research. We **submitted this work² to TNNLS in 2019**.

After these VizDoom projects, I was excited to apply what I had learned to create a conscious system, a topic I am especially interested in. This summer I led the *Rule-Transformer* project supervised by **Professor Yoshua Bengio**. We aimed to develop a

¹Project page: <https://trinkle23897.github.io/cv/viz2018.html>

²Project page: <https://trinkle23897.github.io/cv/rmdp.html>

system with *Consciousness Prior* (proposed by Professor Bengio), which integrated the advantage of both neural and symbolic models. By interpreting the attention architecture as applying rules to variables, we generalized attention to a novel mechanism called *Rule-Attention*. I designed and implemented this architecture in a short time, and performed a comprehensive investigation. The Rule-Transformer turned out to be satisfactory in a wide range of tasks from machine translation to multitask imitation learning. One of the interesting findings was that a decrease in dimension would encourage the Rule-Attention layer to activate more functions to deliver a similar performance. This is related to how we define a function and an argument to interpret a black concept box. The Rule-Transformer project gave me a glance at the fascinating nature of machine learning and the complicated challenges faced by the community. Meanwhile, the experience of working with such a passionate group of world-class researchers in MILA has equipped me with the necessary skills that will stand behind my future research.

I am applying to the Master of Science in Computer Vision at Carnegie Mellon University to master the skills needed to become a successful researcher in industry. Carnegie Mellon University has always been the school of my dreams because of its rigorous curriculum and exceptional research atmosphere. It is my dream to continue research on reinforcement learning at Professor **Ruslan Salakhutdinov, David Held, Deepak Pathak, and Abhinav Gupta**'s group. The advanced courses in your program also overlap with my interests and career plan, such as "Geometry-based Methods in Vision" and "Robot Localization and Mapping" that will enrich me in the intersection of computer vision and reinforcement learning. I believe it is the ideal place to reinforce my knowledge and prepare me to be an eminent industry researcher in the future.

Statement of Purpose

M.S. in Computer Science

Jiayi Weng (trinkle23897@gmail.com)

As an undergraduate, I was lucky to participate in several **reinforcement learning** projects at both Tsinghua University and MILA. This fascinating research area, with great potentials, deeply attracted me. To further pursue my interest, I am determined to apply for an M.S. program in computer science. I am confident that the M.S. training will prepare me to be an independent researcher in the industry and develop practical applications to solve challenging real-world problems.

I first became interested in reinforcement learning when I worked under the supervision of **Professor Jun Zhu** on several VizDoom projects. I am most proud of the performance of our team in the *VizDoom AI Competition 2018*, a famous reinforcement learning challenge. Our agent **ranked 1/51**, far exceeding the score of second-place. This work¹ was ultimately **published in IJCAI 2019, of which I was the first author**. I confronted with several critical challenges in this competition: 1) Visual understanding from raw images; 2) Visual navigation in the 3D environment; 3) Sparse and delayed reward for the agent. First, I built up an object detection system aiming to efficiently identify the obstacles and enemies. Second, inspired by existing visual navigation methods, I incorporated the depth signal into the navigation network, which facilitated the agent to plot its paths. Third, I innovatively formulated the problem within a hierarchical reinforcement learning framework by dividing the challenge into subtasks. To make further improvements, I tuned the reward of the controller. Meanwhile, I accelerated both the training and inference phase and iterated the algorithm every day.

With these achievements on the VizDoom challenge, I was eager to start the independent project *Relational Markov Decision Process*, which intended to accelerate the deep reinforcement learning procedure by state abstraction (using first-order logic) and reward shaping. At first, I found the approach unsatisfying on some of VizDoom scenarios due to the instability of the training process. Therefore, I changed the network to double Deep-Q Network with duel training and adjusted the replay buffer size to improve stability. Next, I implemented parallel environment training to further speed up experiments, making the algorithm converge to the optima more rapidly. Finally, this method accelerated the convergence as expected, and was more robust and generalizable under harder games. The fantastic result far exceeded my expectations, and also strengthened my confidence in continuing research. We **submitted this work² to TNNLS in 2019**.

After these VizDoom projects, I was excited to apply what I had learned to create a conscious system, a topic I am especially interested in. This summer I led the *Rule-Transformer* project supervised by **Professor Yoshua Bengio**. We aimed to develop a

¹Project page: <https://trinkle23897.github.io/cv/viz2018.html>

²Project page: <https://trinkle23897.github.io/cv/rmdp.html>

system with *Consciousness Prior* (proposed by Professor Bengio), which integrated the advantage of both neural and symbolic models. By interpreting the attention architecture as applying rules to variables, we generalized attention to a novel mechanism called *Rule-Attention*. I designed and implemented this architecture in a short time, and performed a comprehensive investigation. The Rule-Transformer turned out to be satisfactory in a wide range of tasks from machine translation to multitask imitation learning. One of the interesting findings was that a decrease in dimension would encourage the Rule-Attention layer to activate more functions to deliver a similar performance. This is related to how we define a function and an argument to interpret a black concept box. The Rule-Transformer project gave me a glance at the fascinating nature of machine learning and the complicated challenges faced by the community. Meanwhile, the experience of working with such a passionate group of world-class researchers in MILA has equipped me with the necessary skills that will stand behind my future research.

I am applying to the M.S. in Computer Science at Carnegie Mellon University to master the skills needed to become a successful researcher in industry. Carnegie Mellon University has always been the school of my dreams because of its rigorous curriculum and exceptional research atmosphere. It is my dream to continue research on reinforcement learning at Professor **Ruslan Salakhutdinov, David Held, Deepak Pathak, and Abhinav Gupta's** group. The three areas (AI, Systems, and Theory) offered by your honorable program will definitely equip me with a solid foundation and help me gain more insights in the field of reinforcement learning. I believe it is the ideal place to reinforce my knowledge and prepare me to be an eminent industry researcher in the future.

Statement of Purpose

Master of Computational Data Science

Jiayi Weng (trinkle23897@gmail.com)

As an undergraduate, I was lucky to participate in several **reinforcement learning** projects at both Tsinghua University and MILA. This fascinating research area, with great potentials, deeply attracted me. To further pursue my interest, I am determined to apply for an M.S. program in computer science. I am confident that the M.S. training will prepare me to be an independent researcher in the industry and develop practical applications to solve challenging real-world problems.

I first became interested in reinforcement learning when I worked under the supervision of **Professor Jun Zhu** on several VizDoom projects. I am most proud of the performance of our team in the *VizDoom AI Competition 2018*, a famous reinforcement learning challenge. Our agent **ranked 1/51**, far exceeding the score of second-place. This work¹ was ultimately **published in IJCAI 2019, of which I was the first author**. I confronted with several critical challenges in this competition: 1) Visual understanding from raw images; 2) Visual navigation in the 3D environment; 3) Sparse and delayed reward for the agent. First, I built up an object detection system aiming to efficiently identify the obstacles and enemies. Second, inspired by existing visual navigation methods, I incorporated the depth signal into the navigation network, which facilitated the agent to plot its paths. Third, I innovatively formulated the problem within a hierarchical reinforcement learning framework by dividing the challenge into subtasks. To make further improvements, I tuned the reward of the controller. Meanwhile, I accelerated both the training and inference phase and iterated the algorithm every day.

With these achievements on the VizDoom challenge, I was eager to start the independent project *Relational Markov Decision Process*, which intended to accelerate the deep reinforcement learning procedure by state abstraction (using first-order logic) and reward shaping. At first, I found the approach unsatisfying on some of VizDoom scenarios due to the instability of the training process. Therefore, I changed the network to double Deep-Q Network with duel training and adjusted the replay buffer size to improve stability. Next, I implemented parallel environment training to further speed up experiments, making the algorithm converge to the optima more rapidly. Finally, this method accelerated the convergence as expected, and was more robust and generalizable under harder games. The fantastic result far exceeded my expectations, and also strengthened my confidence in continuing research. We **submitted this work² to TNNLS in 2019**.

After these VizDoom projects, I was excited to apply what I had learned to create a conscious system, a topic I am especially interested in. This summer I led the *Rule-Transformer* project supervised by **Professor Yoshua Bengio**. We aimed to develop a

¹Project page: <https://trinkle23897.github.io/cv/viz2018.html>

²Project page: <https://trinkle23897.github.io/cv/rmdp.html>

system with *Consciousness Prior* (proposed by Professor Bengio), which integrated the advantage of both neural and symbolic models. By interpreting the attention architecture as applying rules to variables, we generalized attention to a novel mechanism called *Rule-Attention*. I designed and implemented this architecture in a short time, and performed a comprehensive investigation. The Rule-Transformer turned out to be satisfactory in a wide range of tasks from machine translation to multitask imitation learning. One of the interesting findings was that a decrease in dimension would encourage the Rule-Attention layer to activate more functions to deliver a similar performance. This is related to how we define a function and an argument to interpret a black concept box. The Rule-Transformer project gave me a glance at the fascinating nature of machine learning and the complicated challenges faced by the community. Meanwhile, the experience of working with such a passionate group of world-class researchers in MILA has equipped me with the necessary skills that will stand behind my future research.

I am applying to the Master of Computational Data Science at Carnegie Mellon University to master the skills needed to become a successful researcher in industry. Carnegie Mellon University has always been the school of my dreams because of its rigorous curriculum and exceptional research atmosphere. It is my dream to continue research on reinforcement learning at Professor **Ruslan Salakhutdinov, David Held, Deepak Pathak, and Abhinav Gupta**'s group. The advanced courses in your program also overlap with my interests and career plan, such as "Advanced Storage Systems" and "Design & Engineering of Intelligent Information Systems" that will enrich me in the system view of reinforcement learning. I believe it is the ideal place to reinforce my knowledge and prepare me to be an eminent industry researcher in the future.

Statement of Purpose

Master of Science in Artificial Intelligence and Innovation

Jiayi Weng (trinkle23897@gmail.com)

As an undergraduate, I was lucky to participate in several **reinforcement learning** projects at both Tsinghua University and MILA. This fascinating research area, with great potentials, deeply attracted me. To further pursue my interest, I am determined to apply for an M.S. program in computer science. I am confident that the M.S. training will prepare me to be an independent researcher in the industry and develop practical applications to solve challenging real-world problems.

I first became interested in reinforcement learning when I worked under the supervision of **Professor Jun Zhu** on several VizDoom projects. I am most proud of the performance of our team in the *VizDoom AI Competition 2018*, a famous reinforcement learning challenge. Our agent **ranked 1/51**, far exceeding the score of second-place. This work¹ was ultimately **published in IJCAI 2019, of which I was the first author**. I confronted with several critical challenges in this competition: 1) Visual understanding from raw images; 2) Visual navigation in the 3D environment; 3) Sparse and delayed reward for the agent. First, I built up an object detection system aiming to efficiently identify the obstacles and enemies. Second, inspired by existing visual navigation methods, I incorporated the depth signal into the navigation network, which facilitated the agent to plot its paths. Third, I innovatively formulated the problem within a hierarchical reinforcement learning framework by dividing the challenge into subtasks. To make further improvements, I tuned the reward of the controller. Meanwhile, I accelerated both the training and inference phase and iterated the algorithm every day.

With these achievements on the VizDoom challenge, I was eager to start the independent project *Relational Markov Decision Process*, which intended to accelerate the deep reinforcement learning procedure by state abstraction (using first-order logic) and reward shaping. At first, I found the approach unsatisfying on some of VizDoom scenarios due to the instability of the training process. Therefore, I changed the network to double Deep-Q Network with duel training and adjusted the replay buffer size to improve stability. Next, I implemented parallel environment training to further speed up experiments, making the algorithm converge to the optima more rapidly. Finally, this method accelerated the convergence as expected, and was more robust and generalizable under harder games. The fantastic result far exceeded my expectations, and also strengthened my confidence in continuing research. We **submitted this work² to TNNLS in 2019**.

After these VizDoom projects, I was excited to apply what I had learned to create a conscious system, a topic I am especially interested in. This summer I led the *Rule-Transformer* project supervised by **Professor Yoshua Bengio**. We aimed to develop a

¹Project page: <https://trinkle23897.github.io/cv/viz2018.html>

²Project page: <https://trinkle23897.github.io/cv/rmdp.html>

system with *Consciousness Prior* (proposed by Professor Bengio), which integrated the advantage of both neural and symbolic models. By interpreting the attention architecture as applying rules to variables, we generalized attention to a novel mechanism called *Rule-Attention*. I designed and implemented this architecture in a short time, and performed a comprehensive investigation. The Rule-Transformer turned out to be satisfactory in a wide range of tasks from machine translation to multitask imitation learning. One of the interesting findings was that a decrease in dimension would encourage the Rule-Attention layer to activate more functions to deliver a similar performance. This is related to how we define a function and an argument to interpret a black concept box. The Rule-Transformer project gave me a glance at the fascinating nature of machine learning and the complicated challenges faced by the community. Meanwhile, the experience of working with such a passionate group of world-class researchers in MILA has equipped me with the necessary skills that will stand behind my future research.

I am applying to the Master of Science in Artificial Intelligence and Innovation at Carnegie Mellon University to master the skills needed to become a successful researcher in industry. Carnegie Mellon University has always been the school of my dreams because of its rigorous curriculum and exceptional research atmosphere. It is my dream to continue research on reinforcement learning at Professor **Ruslan Salakhutdinov, David Held, Deepak Pathak, and Abhinav Gupta**'s group. The advanced courses in your program also overlap with my interests and career plan, such as "AI: Representation and Problem Solving" and "Design & Engineering of Intelligent Information Systems" that will give me a new perspective of the application of reinforcement learning. I believe it is the ideal place to reinforce my knowledge and prepare me to be an eminent industry researcher in the future.

Statement of Purpose

M.S. in Intelligent Information Systems

Jiayi Weng (trinkle23897@gmail.com)

As an undergraduate, I was lucky to participate in several **reinforcement learning** projects at both Tsinghua University and MILA. This fascinating research area, with great potentials, deeply attracted me. To further pursue my interest, I am determined to apply for an M.S. program in computer science. I am confident that the M.S. training will prepare me to be an independent researcher in the industry and develop practical applications to solve challenging real-world problems.

I first became interested in reinforcement learning when I worked under the supervision of **Professor Jun Zhu** on several VizDoom projects. I am most proud of the performance of our team in the *VizDoom AI Competition 2018*, a famous reinforcement learning challenge. Our agent **ranked 1/51**, far exceeding the score of second-place. This work¹ was ultimately **published in IJCAI 2019, of which I was the first author**. I confronted with several critical challenges in this competition: 1) Visual understanding from raw images; 2) Visual navigation in the 3D environment; 3) Sparse and delayed reward for the agent. First, I built up an object detection system aiming to efficiently identify the obstacles and enemies. Second, inspired by existing visual navigation methods, I incorporated the depth signal into the navigation network, which facilitated the agent to plot its paths. Third, I innovatively formulated the problem within a hierarchical reinforcement learning framework by dividing the challenge into subtasks. To make further improvements, I tuned the reward of the controller. Meanwhile, I accelerated both the training and inference phase and iterated the algorithm every day.

With these achievements on the VizDoom challenge, I was eager to start the independent project *Relational Markov Decision Process*, which intended to accelerate the deep reinforcement learning procedure by state abstraction (using first-order logic) and reward shaping. At first, I found the approach unsatisfying on some of VizDoom scenarios due to the instability of the training process. Therefore, I changed the network to double Deep-Q Network with duel training and adjusted the replay buffer size to improve stability. Next, I implemented parallel environment training to further speed up experiments, making the algorithm converge to the optima more rapidly. Finally, this method accelerated the convergence as expected, and was more robust and generalizable under harder games. The fantastic result far exceeded my expectations, and also strengthened my confidence in continuing research. We **submitted this work² to TNNLS in 2019**.

After these VizDoom projects, I was excited to apply what I had learned to create a conscious system, a topic I am especially interested in. This summer I led the *Rule-Transformer* project supervised by **Professor Yoshua Bengio**. We aimed to develop a

¹Project page: <https://trinkle23897.github.io/cv/viz2018.html>

²Project page: <https://trinkle23897.github.io/cv/rmdp.html>

system with *Consciousness Prior* (proposed by Professor Bengio), which integrated the advantage of both neural and symbolic models. By interpreting the attention architecture as applying rules to variables, we generalized attention to a novel mechanism called *Rule-Attention*. I designed and implemented this architecture in a short time, and performed a comprehensive investigation. The Rule-Transformer turned out to be satisfactory in a wide range of tasks from machine translation to multitask imitation learning. One of the interesting findings was that a decrease in dimension would encourage the Rule-Attention layer to activate more functions to deliver a similar performance. This is related to how we define a function and an argument to interpret a black concept box. The Rule-Transformer project gave me a glance at the fascinating nature of machine learning and the complicated challenges faced by the community. Meanwhile, the experience of working with such a passionate group of world-class researchers in MILA has equipped me with the necessary skills that will stand behind my future research.

I am applying to the M.S. in Intelligent Information Systems at Carnegie Mellon University to master the skills needed to become a successful researcher in industry. Carnegie Mellon University has always been the school of my dreams because of its rigorous curriculum and exceptional research atmosphere. It is my dream to continue research on reinforcement learning at Professor **Ruslan Salakhutdinov, David Held, Deepak Pathak, and Abhinav Gupta**'s group. Learning the advanced courses related to natural language processing in your honorable program, such as "Question Answering" and "Structured Prediction for Language and Other Discrete Data", will definitely give me a deeper understanding of incorporating human-knowledge prior into current machine learning and reinforcement learning framework. I believe it is the ideal place to reinforce my knowledge and prepare me to be an eminent industry researcher in the future.