

1. Team Description

1.1. Team name

The name of our team is *TeleX Robotics*.

1.2. Team leader

Our team leader is Antonio Cao.

1.3. Primary and secondary points of contact

The primary point of contact is Antonio Cao. The secondary point of contact is Tim Foldy-Porto.

1.4. Primary location

Our team is primarily located in New Haven, CT, United States.

1.5. Secondary location

Currently, our team does not operate out of any secondary location.

1.6. Current expertise

There are two key components to our team: the core members and the advisory board. The core team is responsible for the design and development of the technology that will enable telepresence experiences. The advisory board is responsible for providing specialized knowledge that will help the core team solve field-specific challenges, and for critiquing the designs that the core team creates. Therefore, the current expertise of our team can similarly be divided into two components: the expertise of the core team members (discussed here) and the expertise of members of the advisory board (mentioned only in their biographies, as they are not full-time contributors to this project).

Antonio C. and Tim F-P. are undergraduate students with complimentary skill sets. Tim F-P. comes from a strong background of mechanical engineering, having designed and built many mechanical devices throughout high school and college. Through these activities, he became

versed in woodworking and metalworking. His knowledge in these fields was supplemented by college classes and projects, where he studied mechanical design and modern manufacturing methods. Concurrently with that coursework, he undertook a number of independent studies through which he designed novel embedded devices and learned about analog circuit design, radio technology, digital circuit design, microcontroller architecture, embedded programming, and software design for real-time operating systems. He developed this skill set throughout two summers of work as an embedded software engineer. Currently, he is performing research with his professors into the design of a biologically-inspired actuator and the development of novel neuromorphic computing architectures.

Antonio C. has come to the field of robotics with a slightly different set of knowledge; his expertise lies in the fields of software, video game design, algorithms, and machine learning. He is familiar with a number of programming languages—python (Tensorflow, PyTorch, SageMath), C, C++, C#, Java, Javascript—and is fluent in the use of Unity, OpenCV, Blender, and Photoshop, among others. In the past, he has designed applications for virtual reality and augmented reality systems. In addition to his more practical expertise, Antonio is knowledgeable in topics of advanced mathematics and theoretical physics. (345 words)

1.7. Needed expertise

While our core team brings a general knowledge of robotics engineering to the project, there is still a vast amount of expertise that will be needed to see the project to completion. To understand the extent of the expertise our team still needs, we have subdivided every aspect of the project into its most basic components. For example, we have identified for ourselves the three predominant features of telepresence technology: the robotic avatar, the controller system, and the system responsible for transmitting information between the two. Within the robotic avatar subsection, we have identified seven fundamental subsystems, each of which require a good deal of time and expertise to build. The seven systems are: the skeleton, the muscles (actuators), the sensors, the internal communication, the brain, the aesthetics/exterior, and the seventh system comprises all experimental features we might wish to include in the robot, such as flight. We have similarly identified all the salient subsystems required to build the controller device and the data transmission device. All that is to say that we understand, to a good extent, all that's required of our team in order to build telepresence technology.

A large portion of our missing expertise is filled in by our advisory board. So far, we have relied on them to answer questions such as “what are the characteristic features of muscles that our actuators must replicate?” (answered by Madhu V.), “how should we design interfaces between avatar subsystems to best prepare us for the manufacturing phase of development?” (answered by Joe Z.), and “what does it mean to create a believable and meaningful virtual experience for a person? How should they feel?” (answered by Justin B.). In the current stage of the project, our advisory board has been immensely helpful in guiding the direction of our research.

For the most part, the kinds of questions our core team has stumbled upon have been more along the lines of “what should we be building?” rather than “how do we build that?”. These types of questions have been more-or-less easily answered by our advisory board, of which all the members have extensive experience in exploring novel research areas and making informed decisions regarding the direction and the management of a project. As our project progresses however, we will likely begin encountering questions of the second kind; questions that require highly specific knowledge about a particular technology. We anticipate having questions about specific manufacturing processes or embedded protocols, to name a few. And most likely of all, we will have questions that we cannot currently anticipate, as they will be a direct result of whatever direction we decide to take the project. For those questions, we will rely on the extended network of both our advisory board and the larger ANA Avatar XPRIZE community. (463 words)

1.8. Team members, roles, and bios

Core team

1. Antonio Cao (put picture)

Antonio Cao is a junior at Yale University studying physics, mathematics, and computer science. In the summer of 2019, he held the position of research assistant at the GRASP Lab at the University of Pennsylvania, where he worked with Prof. Jianbo Shi to design a novel neural task embedding for robotic manipulations. Previously, he has researched Calabi-Yau manifolds with Prof. Shing-Tung Yau (Harvard University), implemented SLAM algorithms for self-driving vehicles in Prof. Zhong Shao’s (Yale University) research group, and designed autonomous underwater vehicles at the Hadal Science and Technology Research Center in Shanghai. He has published papers on neural networks and on augmented reality. His coursework has included: intensive algorithms, abstract algebra, deep learning, and physical simulations.

2. Tim Foldy-Porto (put picture)

Tim is a senior at Yale University studying physics and robotics engineering. He has been involved in several engineering projects on campus as a contributor and as a team leader. In 2017, he was selected as a summer fellow by the Yale Center for Engineering Innovation and Design. After that, he spent consecutive summers working as an embedded systems engineering at a robotics start up. In that role, he wrote custom device drivers and developed autonomous driving code for a microcontroller running TI-RTOS. His coursework has included mechanical design, analog and digital circuit design, mechatronics, and hardware development for neuromorphic computing.

Advisory board

1. Joe Zinter (put picture)

Dr. Joseph Zinter is the Assistant Director at the Yale Center for Engineering Innovation and Design (CEID), and a Lecturer and Associate Research Scientist in the Department of Mechanical Engineering and Materials Science. He specializes in teaching project-based courses that sit at the intersection of design, technology and entrepreneurship, which have led to dozens of student first authored publications and patents, national and international design awards, and several ventures. Prior to joining Yale, he acted as a Design Preceptor at the Harvard School of Engineering and Applied Sciences and as a Clinical Specialist in cardiac rhythm management for Medtronic. He holds a BA (Fairfield University) in Engineering, a BS (Columbia University) and MS (Cornell University) in Applied Physics, a MHS (Yale School of Medicine), and a PhD in Biomedical Engineering (Yale University).

2. Madhusudhan Venkadesan (put picture)

Professor Madhusudhan Venkadesan studies the biomechanics and control of animal movement, and develops mechanisms inspired by those studies. His lab uses human subject experiments, control theory, dynamics, and robotics to understand how evolution affects the way we control our bodies. Behaviors include walking, running, throwing, and grasping, in humans or other animals, and in biologically inspired machines built in the lab. He has won numerous awards, including the Human Frontier Science Program's Young Investigator Award (2013) and the Society for the Neural Control of Movement Scholarship (2010).

3. Justin Berry (put picture)

Justin Berry is a critic at the Yale School of Art, an interdisciplinary artist, and NYFA artist's fellowship recipient. His work has been exhibited internationally in various venues, with work recently on view at CAVE in Detroit, CUAC in Salt Lake City, and at the University of Richmond Art Museum. Recent issues of Frieze, Pin-up magazine, Media-N, and Prattfolio included features on his work and Bomb Magazine commissioned the piece *i-would.com* from him as part of their portfolio series in 2013. Berry currently serves as project lead for the center's Blended Reality collective, part of an applied research grant in mixed reality at the university. He is a member of the gallery collective Essex Flowers based in New York, and from 2007 to 2008 he was co-director of the artist run curatorial space Alogon, in Chicago, IL. He holds an MFA from the Art Institute of Chicago.

4. Brian Beitler (put picture)

Brian Beitler is a medical student at the Yale School of Medicine. He graduated from Yale in 2018 with a B.S. in mechanical engineering. While an undergraduate, he was involved in his college's aerospace club and worked as a research assistant. He, along with three classmates, founded a company (OnTrack Rehabilitation) that won the 2019 Rothberg Catalyzer start-up prize.

5. Jianbo Shi (put picture)

Jianbo Shi is a Professor of Computer and Information Science at the University of Pennsylvania. He received his Ph.D. degree in Computer Science from University of California at Berkeley in 1998. In 2007, he was awarded the Longuet-Higgins Prize for his work on Normalized Cuts. His current research focuses on first-person human behavior analysis and image recognition/segmentation. His other research interests include image/video retrieval, 3D vision, and vision-based desktop computing. His long-term interests center around a broader area of machine intelligence. He wishes to develop a "visual thinking" module that allows computers not only to understand the environment around us but also to achieve cognitive abilities such as machine memory and learning.

Bios (abbreviated, 200 words total)

1. Antonio Cao

Antonio Cao is a junior at Yale University studying physics, mathematics, and computer science. He has published papers on neural task embeddings for robotic manipulations and on augmented reality for use in surgery. His coursework has included: intensive algorithms, abstract algebra, deep learning, and physical simulations.

2. Tim Foldy-Porto

Tim is a senior at Yale University studying physics and robotics engineering. Most recently, he has worked as an embedded systems engineering at a robotics start up. His coursework has included mechanical design, analog and digital circuit design, mechatronics, and hardware development for neuromorphic computing.

3. Joe Zinter (adviser)

Dr. Joseph Zinter is the Assistant Director at the Yale Center for Engineering Innovation and Design (CEID), and a Lecturer and Associate Research Scientist in the Department of Mechanical Engineering.

4. Madhusudhan Venkadesan (adviser)

Madhusudhan Venkadesan is a professor in the Yale School of Engineering and Applied Sciences. He studies the biomechanics and control of animal movement.

5. Justin Berry (adviser)

Justin Berry is a critic at the Yale School of Art, an interdisciplinary artist, and NYFA artist's fellowship recipient.

6. Jianbo Shi (adviser)

Jianbo Shi is a Professor of Computer and Information Science at the University of Pennsylvania.

2. Team Objectives and Design Approach

A. Objectives

2.A.1. Given the vision of the Avatar competition, what are the specific goals of this team?

It has always been our dream to build technology that enables people from around the world to collaborate on creative projects. We envision a novel technology that allows people to interact with one another in the context of a fully immersive virtual environment. We imagine an architect taking her clients on a virtual tour of a proposed building; the architect walks them through the atrium, encourages them to look around and see the room, hear the sounds of the nearby traffic, and feel the building's material textures. We imagine engineers from around the world getting together in a virtual makerspace; they work with one another in real time, feeling the textures and the inertia of the tools as they use them.

The technology that we're building facilitates interactions not only in the virtual world, but in the real world as well. We imagine a doctor performing surgery—precise and with no latency—from a thousand miles away. We imagine a team of search and rescue operators saving lives in a hazardous environment, all without leaving the safety of their base. To many, these scenarios read like science-fiction, but to us, they are fully realizable goals. The technology that we envision is a biologically-inspired humanoid robot, operated remotely by a human who is wearing an immersive haptic feedback suit, capable of adding tactile and force-based sensations to the visual-audio stimuli of existing VR technology. These devices will support scenarios of real-world utility. (241 words)

2.A.2. How does the team composition contribute to achieving these goals?

Our team's composition will contribute to our success in two ways: the individual members of our team are highly knowledgeable in areas directly related to our goals; and the structure of our team will allow for fast-paced design and rapid iteration, as well as provide us with access to field-specific experts. Although our core team consists only of two people, each member provides a diverse set of technical knowledge that perfectly complements the other's. Tim's background in mechanical engineering, analog circuit design, power electronics, and embedded programming qualify him to handle most hardware challenges that arise in building our technology. Antonio's background in software development, machine learning, math, and industrial design allow him to tackle software challenges and to produce a cohesive design language for all our products.

Having a small core team allows for a quick, streamlined development and iteration process. We will be able to construct and evaluate prototypes much quicker than large teams, which are often

bogged down by bureaucratic and logistic constraints. To make up for the core team's knowledge deficit in highly specific areas, we will rely on our extensive advisory board. We have—on call to answer our questions and review our designs—a professor of engineering and product development, a professional art critic, an expert in computer vision and state-of-the-art neural networks, and a professor of control theory and biomechanics. Combined with our advisory board, we are prepared to tackle whatever challenges stand between us and our goals. (245 words)

2.A.3. How will the team demonstrate and measure success? What tools and approaches to you anticipate using for this?

Our internal measures of success will be drawn from our teams goals—success, in our minds, is creating the imagined experiences that we outlined in the first question of this section. At every turn, the decisions that we make regarding the technical development of our project are based on creating the most vivid, exciting, and realistic user experience possible. We entered the competition with a loose conception of the telepresence device we hope to build, but we made sure to keep an open mind throughout the research phase. This allowed us to perform an unbiased and comprehensive review of the literature regarding humanoid robotics and haptic feedback devices, which in turn let us select—for further exploration and development—the most promising and innovative technologies for producing believable virtual experiences.

We plan on taking a similar approach throughout the rest of the competition: we will constantly keep in mind the experiences we want to create for users and will remain open to new and innovative technologies as our project progresses. Our team is thrilled by the opportunity to provide people with exciting and meaningful virtual interactions. To measure success, we will periodically test our technology with a diverse set of potential users, from whom we will obtain qualitative and quantitative feedback. By comparing the users' perceived experiences with those we desire to create, we will be able to iterate on our designs and make more successful versions of our devices. (240 words)

2.A.4. What are the challenges you foresee your team having in achieving this success?

We anticipate that project management will be the biggest challenge that our team faces. As a small team, our core members will have to assume a variety of roles to achieve our goals. The technological scope of the project is massive—the design of a humanoid robot and the design of an immersive haptic feedback device are topics that would each fill multiple PhD theses. We are more susceptible than larger teams to the risks of wasting time on highly specific aspects of the project; where a large team could afford to have its members specialize in certain areas, we do not have such a luxury.

However, we see this as a strength so long as we remain focused: without the high overhead associated with managing many people, we are able to be ruthlessly efficient in our design process, constantly ensuring that our work time is optimally spent. We will rely heavily on our team of advisers for parts of the project where our core team lacks area-specific expertise. Additionally, by performing periodic literature reviews, we will be able to integrate existing cutting-edge technologies into our designs and cut down on the number of novel devices we have to develop. Collaboration with other teams competing in the competition will also help to mitigate this risk.

In addition to project management, we anticipate some difficulty finding funding for later rounds of the competition. Again, we plan to rely on our advisory board, of which several members have connections with venture capital. (250 words)

2.A.5. What impact will the team's success have on the future of robotic avatars?

We hope that the work we do for this competition affects the future of robotic avatars in two ways: that it solves challenges that have currently stalled progress in the field of telepresence; and that it provides our team and other teams around the world with a jumping-off-point for more advanced work in the field. Currently, robotic and VR technology is not good enough to support a realistic telepresence experience. While bipedal robots do exist, there is no robot that faithfully and successfully mimics the behaviors and attributes of the human body. And although VR technology is rapidly advancing, there does not yet exist a haptic feedback suit that creates realistic or intuitive tactile sensations. By combining cutting edge technologies in an innovative way, we hope to fill this technological void such that the average person is excited and optimistic about the field of telepresence.

Secondly, we want to kickstart the development of advanced and application-specific telepresence technology. The robot and haptic feedback suit that we create will enable a significant subset of telepresence activities, but it will not be sufficient or optimized for many others. For example, the devices we are building will most likely not allow for the remote operation of avatars in outer space, yet that is an important application of this technology. Future teams will be able to build upon our avatar system and expand telepresence robotics into niche markets. (234 words)

2.A.6. What are the overall capabilities and projected applications of your robotic avatar?

Our intention in creating a robotic avatar system is to flawlessly recreate the human experience for both the person controlling the avatar and the people interacting with the avatar. To achieve this goal, three things are necessary: 1) the humanoid robot must move indistinguishably from a real human, 2) the haptic feedback suit must provide highly realistic physical sensations to the wearer, and 3) the latency in the system must be low enough that the delay between action and


feedback, on the controller side and on the robot side, is shorter than that delay in the human body's own nervous systems (~200 milliseconds).

In order for our avatar system to successfully provide a telepresence experience, the operator of the system must not feel as if there are certain tasks that they can do in real life but that their robotic avatar cannot do. Many of the tasks that one might wish to do remotely are human centered, meaning that completing the task relies on using interfaces that are designed for human beings. Such interfaces include touch panels, wall-mounted buttons and levers, light switches, door knobs, serving utensils, etc. By nature of its humanoid form factor and biologically-motivated design, our robotic avatar system will be capable of using all of these interfaces. It will be able to walk, run, throw a football, fold a blanket, and manipulate objects in its hands.

The capability to perform these tasks will provide the operator with the sense that the avatar is an extension of their own body. To complete this experience the haptic feedback suit must supply the operator with sensations that are indistinguishable from the ones they experience in real life. The suit will be capable of communicating to the operator the weight of objects held in the robot's hands, the texture of surfaces the robot touches, and the smells and sounds of a room the robot is in. It will also be capable of extending the human experience by providing the operator with information that is not normally accessible to humans, such as the thermal profile of a room.

However, it is not enough to recreate a human's physical motions. To ensure people with whom the avatar interacts that they are in the presence of a human, not a machine, it is necessary to transport the operator's micro-expressions, subtle nods, and other cues that they use to navigate social interactions. To make the avatar's use of such cues convincing, the timing and delivery of them has to be exactly right. This is where the no-latency requirement comes in. If the avatar were to lag behind during conversations with real people, it would be an unpleasant experience for both the operator and their conversational partner. Our robot will have an unnoticeable level of lag, and will therefore find applications in a wide range of social activities, such as having conversations and demonstrating affection. (480 words)

B. Initial design concept



3. Technical Feasibility Assessment

3.1. Robot Implementation

What physical robotic avatar is the team planning to use?

We are planning on building a full-scale, biologically-inspired humanoid robot for use in the competition.

Describe the form of the robot.

Our robot will look and move like a human: it will comprise a head, two arms, and two legs; it will use underactuated, passively compliant hands for dexterous object manipulation; it will rely on a bipedal leg structure for fast, agile, and efficient locomotion. Our robot's joints will be milled from solid aluminum blocks, and its bones welded from hollow titanium tubing. Its actuators will be placed according to muscle placement in the human body, and then covered in a highly durable Kevlar skin to give the appearance of an organic form. In its movements, our robot will be deft, natural, and resilient to unexpected environmental perturbations. (107 words)

Is the robot bought, modified, self-made, other?

The robot will be designed by Antonio C. and Tim F-P, with supporting contributions from the advisory team, and will be manufactured by our partner company in China.

Does the robot have a name or nickname?

Not currently.


Please expand on current state of development, plans for enhancements to this robot etc.

In Q2 of 2019, we conducted a thorough review of the literature on humanoid robotics. Through this review, we identified the methods used and the features needed to build a state-of-the-art, innovative, and biologically-mimetic robot. During the summer of 2019, we used this information to guide the high-level design of our robot and its sub-components, including its actuation system, sensory system, electrical architecture, and control system. We spent Q3 generating CAD models for all subsystems of the robot, as well as building prototype versions of

select joints, limbs, actuators, and sensors. We are currently finishing up the initial CAD and prototyping processes and are meeting with our advisory team for design reviews. (112 words)

Provide an image or images of the robot (current and planned states)





Who on the team has expertise in this or will lead this effort?

- Tim F-P (mechanical design and control systems)
- Joe Z. (adviser, contributing knowledge in mechanical design, industrial design, and manufacturing methods)
- Madhu V. (adviser, contributing knowledge in biomechanics, the control of animal movement, and bio-inspired mechanisms)

3.2. Controller System

What type of controller system for the operator is the team planning to use? (VR, haptic suit, mixed, other? Describe as fully as possible.)

Our controller system will have two-fold functionality: it will capture movement data *from* and it will communicate haptic sensations *to* the wearer. To detect the wearer's motion, we plan to use an array of position sensors, accelerometers, and electromyography sensors placed around the body. This data will be aggregated, fused, and processed locally to produce an accurate model of the wearer's body position. To convey feedback to the wearer, we will use a virtual reality headset, headphones, and a custom-built haptic suit. Our haptic suit will provide the wearer with both tactile and force sensations, and will consist of a hard exoskeleton—for simulating interactions with rigid bodies—as well as a soft exoskeleton—for simulating interactions with compliant bodies. (120 words)

Who on the team has expertise in this?

- Tim F-P (hard and soft actuators, electrical architecture)
- Antonio C. (sensor fusion/machine learning, VR game design)
- Justin B. (adviser, contributing knowledge on input/output devices for VR, human-machine interactions)
- Brian B. (adviser, contributing knowledge on the sensory system in the human body)

3.3. What type of operating systems is the team planning to use?

Name and list reasons for choice

We have subdivided the choice of operating system into several categories: (1) controller-side operating system, responsible for collecting and processing movement data from the haptic suit,

as well as transforming robot sensor data into meaningful haptic feedback for the operator; (2) high-level robot operating system, responsible for general decision making and sensor-data processing; (3) low-level robot control system, which receives commands from (2) and in turn controls actuators and sensors. For (1), we are planning to use a Linux-based system. For (2), we will likely develop a custom computer architecture and OS. For (3), we will use an RTOS running on a secure embedded platform, such as the TI Hercules microcontroller. (111 words)

Who on the team has expertise in this?

- Tim F-P. (real-time operating systems, code for embedded systems)
- Antonio C. (Linux-based software development)

3.4. What types of sensory systems are the team planning to use?

We are planning to use a virtual reality headset, a surround-sound audio system, a custom-built haptic feedback suit, and an aroma device to convey sensory information to the operator.

Vision

We plan to use a virtual reality headset to convey visual feedback from the robot to the person controlling it. The VR system, most likely an HTC Vive or Oculus Rift, will receive video data from a variety of cameras mounted on the humanoid robot. These cameras will be able to capture a wide array of visual data—both visible light and infrared. The various video feeds will then be merged on the controller side and combined with telemetry data so that the wearer receives the most informationally dense VR experience possible.

Expertise provided by:

- Antonio C. (spatial experience design for VR)
- Jianbo S. (image/video transmission, 3D vision)

Hearing/sound

To communicate auditory information from the robot to the human controller, we use an array of speakers that are spaced evenly around the controller space. The human will move in a confined environment, so we will be able to mount speakers around the perimeter of that environment such that they don't interfere with the controller's movements. With this spatially diverse configuration, we will be able to faithfully reproduce the sounds that the robot perceives. To capture rich auditory information on the robot side, we will use an array of directional microphones mounted on the robot's head.

Expertise provided by:

- Antonio C. (spatial experience design for VR)
- Tim F-P. (audio electronics)

Haptics

We are building a custom haptic feedback suit for the purpose of communicating rich, intuitive, and life-like tactile and force-based sensations to the human controller. Our suit consists of three subsystems: a tactile feedback unit, capable of reproducing sensations such as skin stretch, skin deformation, and skin pressure; a hard exoskeleton, capable of simulating rigid contact forces; and a soft exoskeleton, capable of simulating interactions with compliant materials. These three systems will be merged into a single wearable device, mounted to the human rather than to an external frame. Mounting directly to the controller's body will allow for greater flexibility of movement, which will make the whole experience feel more life-like and less artificial.

Expertise provided by:

- Tim F-P. (mechanical / electrical engineering, embedded software)
- Antonio C. (industrial design, software to interface with VR)

Smell

To simulate smells, we plan on purchasing a third-party aroma device, such as the Aroma Shooter produced by AromaJoin. These devices are small, easily-controlled, and capable of producing a wide variety of smell sensations. By mounting them in a circle around the human controller—much like we plan to mount the speakers—we will be able to reproduce the exact spatial configuration of the smells that the robot perceives. Additionally, we will indicate smells to the wearer in an visual way, in case the robot detects odorless gases, such as carbon monoxide, or in case the wearer's own sense of smell is incapacitated.

Expertise provided by:

None; expertise needed in this field.

Other modalities

In addition to the haptic feedback suit, which provides the operator with tactile and force-based sensations, we are planning to influence the operator's sense of balance while they are

controlling the robot, thus providing them with a better sense of the robot's orientation in space. This will allow the operator to understand the robot's experience in dynamic situations, such as when the robot is jumping or running, or when the robot trips, stumbles, or is knocked off-balance by an external force. Understanding such experiences will enable the operator to help the robot recover from situations in which it has fallen or been knocked down.

Expertise provided by:

- Tim F-P. (mechanical / electrical engineering, embedded software)
- Antonio C. (software to interface with VR)
- Madhu V. (adviser, contributing knowledge in human biological systems)

3.5. Is your team planning to use AI technologies?

We will use AI in the following ways:

- Sensor fusion:
 - Tracking hand and body movement from sensor data (encoders, IMUs, etc.).
 - *Baseline data:* motion capture.
 - *Architecture:* CNN, Autoencoder with Kalman filters for noise reduction.
 - BMI: Understanding data from EMG sensors and EEG sensors.
- Navigation:
 - Object recognition: object segmentation and semantic understanding to help the robot interpret its surroundings.
 - *Architecture:* Mask-RCNN, PointNet++.
 - SLAM: Learning semantic landmarks to enable fast localization in a chaotic environment.
- Reinforcement learning for robot control: Learning dexterous control the body of a robot (e.g. object manipulation, self-balancing, and collision avoidance).
 - *Architecture:* PPO.
- AI-assisted rendering to reduce latency:
 - Foveated rendering: transmit sparse data points and use GAN to rapidly generate realistic images in the user's peripheral vision.

Who on the team has expertise in this?

- Antonio C. (implementation of machine learning algorithms)
- Tim F-P. (neuromorphic computing)

- Jianbo S. (adviser, contributing knowledge in machine learning, implementation in human-based environments)

3.6. Is your team planning to use a Brain-Machine Interface (BMI) or other sensor technology?

Define how these will be used and list reasons for choice.

Our current design for the operator suit incorporates EMG sensors along with accelerometers and position-based sensors. By fusing the biological sensor data with the physical sensor data, we will be able to obtain a higher level of motion-tracking accuracy than with either type of sensor alone. The EMG data will also be useful in training our latency-attenuating GAN, described in the following section. At the moment, we have no plans for using signals from the operator's brain to control the robot, but that is subject to change based on the results of our first prototyping round.

Who on the team has expertise in this?

- Antonio C. (machine learning)
- Tim F-P. (signal processing and filtering, sensor fusion)
- Madhu V. (adviser, contributing knowledge in EMG sensing procedures, EMG data processing)

3.7. How does your team plan to handle latency and lag challenges?

We will minimize latency and lag by implementing a robust continually-learning GAN between the operator and the robot to preemptively send data. Two of these networks will be implemented: one along the path of the control signal from the operator to the robot, and one along the path of the feedback signal from the robot's sensory systems to the haptic feedback suit. The job of these networks is to predict the input signal before it is actually generated. The networks will be pre-trained, but will also exhibit online learning to account for any given operator's unique movement characteristics. So long as the predictor generates highly accurate input signals, the effect on the output motion (on the robot side) or haptic sensations (on the operator side) will be imperceptible. (128 words)

3.8. Does your team plan to collect data on any aspects of the robot's operation?

Naturally, we will log data about the robots movements, its sensory perceptions, and its decision making process. This will allow us to refine the robot's control system, as well as identify any potential bugs or safety risks. By comparing this data to baseline data, obtained either through motion capture technology or an external position-detecting suit, we will be able to make informed adjustments to the robot's mechanical dynamics so as to mimic human motion as faithfully as possible. Additionally, storing this information will provide us with a large dataset with which to train our machine learning models. Logging video and audio feeds will allow us to tune our computer-vision and natural-language-processing algorithms for optimal performance. (115 words)

Who on the team has expertise in this?

- Antonio C. (data science and analysis)
- Tim F-P. (control system optimization)
- Madhu V. (adviser, contributing knowledge in control system tuning and optimization)
- Jianbo S. (adviser, contributing knowledge in ethical data collection procedures)

3.9. Does your system require use of the cloud?

We have identified three scenarios in which our system will use the cloud. During robot training, we will use cloud computing facilities (e.g. Amazon AWS) to train the sensor fusion and robot control networks. (e.g. create millions of simulations on the cloud to train the algorithms). Secondly, during remote operation of the avatar we will use the cloud to externally process high-level planning and navigation algorithms for the robot. Such algorithms include long-term motion planning, large-scale SLAM, and remote parsing of point clouds. Finally, we plan on using cloud services for data storage and post-processing. After testing and deployment of the avatar, we can use the cloud to the collected images and point clouds, automatically analyze robot performance, and produce artificial training data. (123 words)

3.10. Considering that the avatar system will be interacting with actual humans, what is your team's approach to mitigating potential safety issues?

We will be mitigating safety issues primarily through good design of the passive mechanics of our robot. Principally, we are designing a custom soft actuator based on cutting-edge research into the properties of human muscles (i.e. force controlled with variable impedance). By combining these actuators with intrinsically compliant joints and end-effectors, our robot will interact with the world much the same way a human would; it will be able to exert human-like forces with a finite bandwidth (as opposed to rigid, position based systems which exert arbitrarily high forces with theoretically infinite frequency). Additionally, we will have layers of

electrical- and software-based safety measures, which will constantly be working to ensure the robot is operating in a safe manner.

4. Work Plan

4.1. List each phase for the team effort from now until the end of the competition.

1. Prototyping round one

July 2019 – February 2020

In this phase, we will finish our initial prototypes of the humanoid robot and the haptic feedback suit so that we can present them to our advisory board for a design review. We will begin building relationships with potential investors / partners.

Key people: Tim F-P.
Antonio C.
Advisory board
Jiangsu Zhengyang Art and Communications Ltd.

2. Prototyping round two

March 2020 – September 2020

We will iterate on our initial prototypes and develop a full-scale prototype to present to our advisory board. On the business side, we will finalize our relationships with investors and begin contract negotiations.

Key people: Tim F-P.
Antonio C.
Advisory board
Investors

3. Production round one

Q4 2020 – Q1 2021

Based on feedback from prototyping round two, we will build a production-ready version of our avatar system for the purpose of competing in the semi-finals. At this stage, we hope to have secured the funding necessary for us to scale-up our manufacturing efforts.

Key people: Tim F-P.
Antonio C.
Manufacturers

4. Prototyping round three

Q2 2021 – Q3 2021

This will be the final round of prototyping, in which we will evaluate our performance in the semifinals and make the necessary adjustments to our designs. If needs be, we will have discussions during this time with other teams in the competition about potential collaboration.

Key people: Tim F-P.
Antonio C.
Advisory board

5. Production round two

Q4 2021 – Q1 2022

During this stage, we will produce the ultimate version of our avatar system that will compete in the competition finals.

Key people: Tim F-P.
Antonio C.
Manufacturers

(267 words)

4.2. Objectives and subtasks

4.2.1. Objective one

Description

Build a biomorphic, safe humanoid robot that faithfully mimics the behaviors and the attributes of human motion.

Technology approach

A key feature of our approach is that we will use cutting-edge research into the properties of human muscle to guide our design of actuators for our humanoid robot. Current research suggests that muscles respond to step input forces according to a stretched exponential function. This is a fundamentally different response than mainstream robotic actuators, which exhibit

theoretically infinite bandwidth and forces to enable position-based control. Generally, roboticists attempt to counteract the shortcomings of their position-based actuators through sophisticated control algorithms. However, this has traditionally led to large control complexities that result in lag, instabilities, and, in general, non-animal-like behaviors. By designing an actuator that passively recreates the salient qualities of human muscles, we will be able to reduce control complexity and accurately mimic animal motion.

Issues and challenges

We anticipate the fine tuning of the passive dynamics of our actuator to be a challenge. We expect the optimal parameters of the actuator to exist in a narrow window, where outside of that window the actuator exhibits unstable or unsafe behaviors. We will likely need to develop a novel robust tuning method to address this challenge. In terms of the robot in general, we expect challenges associated with aggregating, processing, and transmitting the massive amount of data that the robot will necessarily be collecting.

Desired impact

It is our hope that our work catalyzes development in the field of bio-inspired actuator design and enables the creation of other biologically-accurate robotic devices. (261 words)

4.2.2. Objective two

Description

Build a robust and portable haptic feedback suit—one that is able to integrate with existing VR technologies—for the purpose of providing fully immersive, life-like, and realistic virtual experiences to the wearer.

Technology approach

Our approach to this problem is to design a three-component haptic suit: one of the subsystems will handle tactile feedback while the other two subsystems handle force-based sensations. The tactile subsystem will consist of micro vibro-tactile actuators that will be able to apply vibration and light pressure at various points on the wearer's skin. One of the force subsystems will simulate interactions with soft or compliant materials, such as pillows or cushions. The other force subsystem will simulate interactions with hard and rigid materials, such as walls.

Issues and challenges

We expect the comfort of our haptic suit to present a challenge. Current VR technology suffers from inducing fatigue on its wearers. In order for telepresence to be truly viable, the devices that enable the virtual experience must be able to be worn for long amounts of time and be minimally noticeable by the wearers. We plan on performing multiple iterations of user studies and to work with professional clothing designers to perfect the aesthetics and ergonomics of our haptic suit.

Desired impact

We hope that a haptic suit that provides the wearer with realistic and intuitive tactile sensations finds applications beyond those of the competition. Such a device has the potential to have an impact on the VR gaming industry as well as other emerging VR fields. (254 words)

4.2.3. Objective three

Description

Design a system architecture that minimizes latency and maximizes noise-resilience for the connection between the avatar and the human operator.

Technology approach

Our approach to reducing latency and noise in the connection between the operator and the avatar is to develop a robust algorithm for predicting operator motion as well as robot sensory input. A network configuration without such an algorithm would yield a noticeable delay due to signal propagation time. By intelligently predicting the control signal and the feedback signal before it is actually received, we will be able to move the robot and to deliver haptic sensations to the operator with an imperceptible delay. Most likely, the operator side of this predictive technology will consist of a classifier that associates muscles impulses and previous movement with resultant movement, and a generative adversarial network that learns to rapidly generate realistic control commands to send to the robot. A similarly structured algorithm will be applied to the robot sensory data.

Issues and challenges

An issue naturally arises if there is a discrepancy between what our algorithms predict and what the human actually does or the robot actually senses. This will have to be solved by a comprehensive and strict training process for our predictors. The networks will have to continually learn the operator's movement habits so as to maximize the realism of the robot's movements and of the feedback sensations to the operator.

Desired impact

Continual learning in generative networks is a nascent topic in machine learning research. Implementing such a network into our robotic device will be a novel application of this technology, providing machine learning researchers with real-world data on the efficacy of such algorithms. (284 words)

4.3. Related work

Prior to beginning the competition, both members of the core team developed technologies that were related to the field of telepresence. Tim F-P. spent the summer of 2017 working on an exoskeleton robotic arm that was capable of providing the wearer with 100 lbs. of additional strength. Additionally, he researched formal verification algorithms for embedded systems during an independent study in the fall of 2017. During his time as an undergraduate researcher in the Yale autonomous driving lab, Y-Driving, Antonio C. developed an augmented reality app to teleoperate mobile robots and pose virtual obstacles. All of the work listed above has since been completed. All current efforts to develop telepresence technology are centered around the ANA Avatar XPRIZE competition.

4.4. Funding

- Current funding sources

Name of source: Jiangsu Zhengyang Art and Communications Ltd.
Type of funding: Payment in exchange for deliverables
Funding amount: \$20,000 USD
Type of contract: Project cooperation agreement
Window of validity: 1 September 2019 – 1 April 2020

- Expected funding sources

We are currently in the process of reaching out to venture capital firms, and expect to begin funding discussions in spring of 2020. By the fall of 2020, we hope to have secured the funding necessary for us to complete the development of our robotic avatar system.

- In-kind contributions

As part of our initial contract with Jiangsu Zhengyang Art and Communications Ltd., it was agreed that they would manufacture the humanoid robot free of charge (up to a certain monetary value, at which point their normal fees would take effect).

- Our team has not participated in any related events and is not planning on participating in related events during the timeframe of the ANA Avatar XPRIZE competition.

5. Executive summary

It has always been our dream to build technology that enables people from around the world to collaborate with one another and to push the limits of human exploration. We envision a novel technology that allows people to interact with one another in the context of a fully immersive virtual environment. We imagine an architect taking her clients on a virtual tour of a proposed building; the architect walks them through the atrium, encourages them to look around and see the room, hear the sounds of the nearby traffic, and feel the building's material textures. We imagine engineers from around the world getting together in a virtual makerspace; they work with one another in real time, feeling the textures and the inertia of the tools as they use them. We imagine humans going where no human has gone before: to a remote planet in outer space with a toxic atmosphere, to the bottom of the oceans, to the interior of a nuclear power plant during a meltdown.

We are proposing the development of a two-part system—a *robotic avatar* controlled by an *human operator* who is acting in a virtual space—that enables the above scenarios and more. One part of our system is a biologically-inspired humanoid robot, capable of moving with the speed, agility, and strength of a human. To facilitate social interactions, the robot will not only move like a human, but it will look like a human: it will comprise a head, two arms, and two legs; it will use underactuated, passively compliant hands for dexterous object manipulation; it will rely on a bipedal leg structure for fast, agile, and efficient locomotion. Our robot's joints will be milled from solid aluminum blocks, and its bones welded from hollow titanium tubing. Its actuators will be placed according to muscle placement in the human body, and then covered in a highly durable Kevlar skin to give the appearance of an organic form. In its movements, our robot will be deft, natural, and resilient to unexpected environmental perturbations. When interacting with people, our robot will faithfully replicate the operator's micro-expressions, subtle nods, and other social cues they use to navigate social interactions.

The other part of our system is an immersive virtual reality (VR) suit that senses the operator's movements and provides the operator with visual, aural, olfactory, tactile, and force-based sensations. The VR suit will comprise: a virtual reality headset to display visual data received from visible and infrared cameras on the robot; an array of speakers that spatially reproduce audio perceived by the robot; a custom haptic feedback bodysuit that renders tactile sensations as well as rigid and soft contact forces; an aroma device to replicate smells that the robot encounters; and a vestibular stimulation system, capable of influencing the operator's sense of balance for the purpose of communicating the robot's spatial orientation. Even if the operator is located thousands of miles away from the robot, the robot's experience will be transmitted to the VR suit and rendered for the operator with imperceptible latency. We attempt to achieve this by implementing a continually-learning generative adversarial network to preemptively send data between the operator and the robot.

We are aware that this project is a moonshot, but cutting-edge, never-been-done-before territory is where our team thrives. Our team is composed of two passionate core members and a diverse advisory board, consisting of professional engineers, artists, and professors, each an expert in their field. The small size of our core team allows us to move fast, quickly innovating and iterating on our designs. The intellectual breadth and depth of our advisory board allows us to tackle the most challenging problems standing in the way of developing full telepresence technology.

Why were suited to solve this problem