Robotic Manipulation with a Human in the Loop - Accuracy Estimation of the Baxter Robot

Jiewen Sun
Sebastian Schweigert
James Su
Sunil Srinivasan
Mark Jouppi

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2015-67
http://www.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-67.html

May 13, 2015

University of California, Berkeley College of Engineering

MASTER OF ENGINEERING - SPRING 2015

Electrical Engineering and Computer Sciences

Robotics and Embedded Software

ROBOTIC MANIPULATION WITH A HUMAN IN THE LOOP – ACCURACY ESTIMATION OF THE BAXTER ROBOT

JIEWEN SUN

This **Masters Project Paper** fulfills the Master of Engineering degree requirement.

Approved by:

1. Capstone Project Advisor:

Signature: ___________________________ Date _____________

Print Name/Department: PROFESSOR RUZENA BAJCSY, ELECTRICAL ENGINEERING & COMPUTER SCIENCES

2. Faculty Committee Member #2:

Signature: ___________________________ Date _____________

Print Name/Department: PROFESSOR RONALD S. FEARING, ELECTRICAL ENGINEERING & COMPUTER SCIENCES

University of California, Berkeley College of Engineering

MASTER OF ENGINEERING - SPRING 2015

Electrical Engineering and Computer Sciences

Robotics and Embedded Software

ROBOTIC MANIPULATION WITH A HUMAN IN THE LOOP – ACCURACY ESTIMATION OF THE BAXTER ROBOT

JIEWEN SUN

ABSTRACT

This paper illustrates the design and implementation of a robotic software system based on robot-human interactions, which increases the applications of low-cost, inaccurate robots such as Baxter. The first four parts of the paper discuss the motivation, market & industry, IP strategy and system workflow. Then the paper explains a specific technical contribution of the project in detail – estimating Baxter's positioning accuracy. Based on forward kinematics and Jacobian matrix respectively, two models are built to convert the error tolerance of the joints to the error tolerance of the end effector in the Cartesian space. The second model based on Jacobian matrix is chosen as the final approach since it has linear complexity and it is more computationally efficient. The result gives a reasonable estimation of Baxter's accuracy, which allows valid comparison with the task's accuracy requirements. Thus we are able to determine how the human can interact with the robot and improve its performance.

Robotic Manipulation with a Human in the Loop

Jiewen Sun

Final Capstone Report – Accuracy Estimation of the Baxter Robot

1 PROBLEM STATEMENT

(Authors: Schweigert, Srinivasan, Sun, Su, Jouppi)

We live in an age of increasing automation, but while we have machines that can open a can, pour a glass of water, or assemble a piece of furniture, the world does not have a machine that is versatile enough to do *all* of these tasks.

Normally, when people think of automation, they think of robots designed to accomplish a very specific task, such as lifting a car-door into a car-frame, over and over again. The newer generation of robots though is the class of general-purpose robots. While such robots have yet to materialize commercially, general-purpose is a great concept. Imagine if families could have a robotic assistant to take care of household tasks or run daily errands. In short, human life would be much more convenient and efficient.

Unfortunately, a major limitation towards reaching such a milestone is the engineering trade-off between cost and performance: with a limited budget of resources, it is almost impossible to add additional levels of complexity without decreasing performance. As such, it has traditionally been challenging to use robots that are both low-cost and versatile in domestic environments because the applications of these robots are limited by their low performance – specifically, their inaccuracy. This is where we decided that the human should come in.

To address this barrier of limited resources, our capstone team has developed a system that is designed around robot-human interaction, where human instructors train and work with cost-effective robots to accomplish a broad range of tasks with high accuracy. Using a set of algorithms that we have developed, the robot learns how to perform a task from a human who teaches it through a series of demonstrations. Following this learning process, the robot evaluates the task and identifies the precision requirements using a mathematical model. And when the robot detects that it is unable to achieve the accuracy required for a certain portion of the task, it requests human assistance. The final outcome is a system that excels across a vast range of duties, due to the combination of both the efficiency of robots working on a large-scale and the precision of humans working on a small-scale.

This revolutionary design of cooperation between man and machine succeeds at tasks that are otherwise impossible for the machine to accomplish alone. In essence, we added in the human as an additional resource to improve the overall performance of the system. This was the rationale behind our capstone project, for we saw an opportunity here to make an enormous technical stride in society's current usage of commercial robots: we took an otherwise unimpressive commodity – the low-cost and inaccurate robot – and engineered commercial value from it in the form of robotic adaptability.

2 INDUSTRY AND MARKET TRENDS

(Authors: Schweigert, Srinivasan, Sun, Su, Jouppi)

Before examining any technical details though, we first wanted to scope out the business potential of our project. Consequently, in an attempt to analyze our strategic position in the market, we evaluated the competitive forces outlined by Porter (Porter, 2008) because we felt that an in-depth analysis of the intensity of these forces will influence our marketing strategies. In other words, analyzing these five forces enabled us to have a better understanding of our industry and shaped our strategies to sustain long-term profitability. Before we begin our analysis however, let us first clearly define both our market and our product.

2.1 Market and Product

We defined our market to be consumer households with the intent that our algorithms accomplish household tasks, such as assembling furniture. We chose to target the elderly and the disabled as buyers of our product because this is a large, growing population with critical and largely unmet needs. Simply put, the elderly population in the United States is growing. While the current number of senior citizens in the US is roughly 40 million, that number is expected to grow to over 80 million by 2050 (Ortman et al., 2014). Additionally, according to US 2010 census data, about 30%, 38%, and 56% of the population aged 70 to 74, 75 to 79, and 80 or over, respectively, live with severe disabilities (Brault, 2012). To further narrow our market though, we chose to focus specifically on affluent elderly-and-disabled individuals as our target customers. This is a reasonable objective because many elderly people have amassed a wealth savings and investments cultivated over their lifetimes. Indeed, according to a Pew Research study, the median net worth of senior citizens in the US is $170,000, which is 47 times greater than that of people aged 35 and younger (Censky, 2011).

The definition of our product is a more complex matter because, at its core, our capstone project involved the research and development of an algorithm that allows a robot to learn a task and cooperate with a human to perform that task; it is not a complete software – or hardware – solution. Unfortunately, while software solutions usually have commercialization potential, algorithms alone do not. In order to take our robot-learning algorithm and relate it to a commercial application, we had to decide what form that application should take and how to take such a product to market. One option was to simply license out our algorithm for others to utilize; we would receive royalties as a result of these sold licenses, and companies could make products or provide services using our algorithm. One major caveat, though, is that our algorithm incorporates ideas presented in externally-published research, so the intellectual property for this method may not lie entirely with us. We therefore chose not to investigate this option any further. Our next option to consider was to sell a software solution for users to install on devices that they already own. However, the “device” in this case would be a full-fledged robot, where, as a point of reference, a Baxter robot from Rethink Robotics – our current hardware-platform of choice – has a set price of approximately $35,000 (Rethink Robotics, 2015). Clearly, it would be ludicrous for people to purchase such costly technology without ensuring that it already comes with the necessary software to function. This left us with our final choice: a “full package”, in which we offer a robotic apparatus preloaded and set up with our software such that a consumer only needs to buy one product, with installation services if necessary. This way, we can market our product directly to our target consumers and eliminate the customer’s barrier-to-purchase that comes from setting up the technology. Thus, we decided on this “full package” as the form for our product: a physical robot bundled with software algorithms that we implement.

We must consider several factors with the decision to market this “full package”. The first is price, and this is largely influenced by the suppliers since we must obtain the proper robotic hardware or components externally. After all, according to an IBISWorld report, the cost of mechanical manufacturing is increasing as the expenditure of raw materials increases, so we opt to purchase a whole robot setup instead of building our own robot from basic components (Crompton, 2014:25). As a result, we would look to Rethink Robotics as a supplier of our Baxter robot, a hardware platform. With a markup from our software and services, selling our product at around $40,000, or at about a 15% markup, is not an unreasonable price point – especially if we were to get an Original Equipment Manufacturer (OEM) discount for Baxter. This provides us with a defined pricing model.

Lastly, we must discuss promotion and place/distribution. As O'Donnell points out, 50% of seniors are not using the internet, so marketing is better achieved through conventional channels such as mail, television, and newspapers (O'Donnell, 2015). Interestingly, O'Donnell also predicts an increased use of social media by seniors in 2020, making social-media campaigns a possibility in the near future (O'Donnell, 2015). Distribution of this product, however, is complicated; while we would like to be able to sell our product online, providing setup services would require a trained professional to be present. As such, we will most likely have to either distribute through local partners that provide such services or create a local presence ourselves, incurring additional costs. With our product, price, promotion, and place now defined, we have all the significant facets of a commercialization strategy. Note that we do not analyze the minimum viable product (MVP) in detail. This is because our research specifically investigates the Baxter robot's ability to learn the task of assembling a coffee table, at which point we will have a decent MVP that performs table assembly. Thus, we have established a viable (if hypothetical) commercialization strategy for our research efforts.

2.2 Competitive Forces Analysis

2.2.1 Power of Buyers

With the market and product definition out of the way, we can begin to evaluate Porter's five forces, the first of which is the power of buyers (Porter, 2008). We deduce this force to be relatively weak, since the large population of potential buyers means that individual buyers do not have much leverage or bargaining power with us in our product offering. Moreover, as we will address later on, there are few – if any – direct rivals in our industry. Thus, a scarcity of competing products only elevates our power, as options are limited for the buyer. Furthermore, the switching costs for complex, robotic solutions would be high; given that the price of these robots with our software would be roughly $40,000, it is not an expense to be made frequently. We imagine that a typical customer will only purchase one such robotic system in their life. Thus, it is not of great concern that customers would switch to using a competitor's domestic robot solution after purchasing our product. All in all, the power of buyers is assessed to be fairly weak, and we do not concern ourselves in mitigating this force.

2.2.2 Power of Competitors

Regarding rivalry within our industry, there are two main classifications of competitors: robotics companies and robotics research institutions. Some of these competitors offer products that are mildly similar to our envisioned product, and they also target similar markets. For example, Clearpath Robotics, a robotics company (Hoover's, Inc. "Clearpath Robotics Inc.," n.d.), offers support to the PR2 robot to perform household chores like fetching food from the refrigerator and cooking it. Alternatively, there are research institutions like the Information and Robot Technology Research Initiative (IRT) at the University of Tokyo working on developing software that allows the AR robot to accomplish household assignments such as cleaning floors, washing laundry, and so on. Fortunately, companies and research institutions like these will only indirectly compete with us because our product differs from theirs in the extent that humans are involved. The robotic systems these competitors are developing are meant to be fully autonomous – the robots execute their tasks independent of any human interaction – while our system is meant to be semi-autonomous, enabling a human to both work with and teach a robot to perform various tasks. This is an advantageously superior method because now the scope of the system is not limited to what the robot can accomplish independently; the scope is broadened to what the robot and human can accomplish together synergistically. Simply put, the generality of our method enhances a robot's utility and flexibility. Apart from offering a unique product though, we also have some advantages over our competitors in terms of hardware costs. To illustrate, a two-arm PR2 robot is priced around $400,000 (Owano, 2011) while a Baxter robot, as mentioned previously, is priced around only $35,000 (Rethink Robotics, 2015), a relatively far-cheaper option. To summarize, since we are working in a fairly new field, there are no true established rivals in this specific area yet. Thus, we can conclude that the force of competitors is weak.

2.2.3 Power of Substitutes

Moving onto the next force listed by Porter, we realize that significant attention needs to be given to the force of substitutes since there are, broadly speaking, quite a number of substitutes to our product. For instance, alternative technologies, like the iRobot Roomba (Biesada, n.d.) – a popular floor cleaning robot, have existed in the consumer market for many years, and these established technologies have a large customer base. Customers are more comfortable with familiar products, so it will not be easy to encourage customers to migrate to a substitute product. Moreover, if we look past the technological substitutes, there are a variety of human-labor alternatives in regards to accomplishing household tasks, such as employing a live-in caretaker or residing in a nursing home. However, similar to our stance against the competitor force, we again have some advantages due to our functionality and low cost. Addressing the concern of alternative technologies, even though products like the iRobot Roomba are popular and functional, they tend to have a limited set of features, such as floor cleaning. Our product, on the other hand, is a more general solution which can be used to tackle a variety of household chores. Along that same line, for many tasks in this set, our robot can be more efficient than a human caretaker due to its autonomous nature. Furthermore, as mentioned previously, our pricing model markets our product at a cost of about $40,000 with an extensive lifespan, while most nursing homes cost up to $80,000 – and that is per year (Ellis, 2013). All of these arguments make our product competitive to existing substitutes, motivating us to divert attention from this force and concentrate on more pressing ones.

2.2.4 Power of New Entrants

In contrast to the mild nature of the forces mentioned previously, new-entrant competition looming over the horizon should be of great concern. For instance, some of the heavy-hitters in robotic research include companies like Clearpath Robotics (McMillan, 2015) and 3D Robotics (Bi, 2015), both of which were founded only six years ago in 2009. It seems that, unlike the issue of existing rivals and possible substitutes, there is indeed a strong force in regards to new entrants. To further illustrate this fact, large corporations with broader goals in the technological field can certainly seep into our industry, such as Amazon with its Amazon Prime Air drones or Google with its autonomous cars. Big players such as these would certainly have the resources to quickly create a new division within their company and fund research in alternative robotic avenues. Furthermore, even our suppliers can be considered possible new entrants, since they both already posses their own hardware and can additionally reverse-engineer our software algorithm that was, in large-part, acquired from public research papers. All in all, to summarize, we see that dangerous incoming players in this industry are: either startups or big companies with other additional focuses, or suppliers that provide our hardware. When combined with the fact that there are no true established rivals yet as mentioned previously, this danger reinforces both the notion that robotics is a relatively new field and that the threat of new entrants is high.

2.2.5 Power of Suppliers

The last of Porter's five forces to address is the threat of suppliers (Porter, 2008). This threat is a complex point that requires careful analysis in our business strategy. To first clarify, we envision robotic-hardware-platform manufacturers as our suppliers. As per our product description, we would take the robotic hardware platforms from companies like Rethink Robotics and Universal Robotics, customize the robots with our specialized software that gives them practical intelligence to work alongside humans, and then sell them to customers. In particular, we would purchase from companies that produce innovative, low-cost robotic hardware platforms upon which we can then build our solution. Our smart software would make up for the inaccuracies in the cheaper hardware with better algorithms and human-in-the-loop collaboration. Since there are currently only a few firms producing such low-cost platforms, these few suppliers have high bargaining power, as we are left with fewer alternate firms from which to choose.

2.3 Market Strategy

We see that presently, of Porter's five forces, both new entrants and existing suppliers hold strong power (Porter, 2008). Knowing this, we can establish our market strategy to mitigate these two forces, strategically positioning ourselves in a superior situation.

To mitigate the threat of new entrants from the suppliers themselves (see Section 2.2.4), we can generalize our software to work across multiple platforms and disincentivize suppliers to enter the market, as they would only be encouraged to produce software across their own single platform. Additionally, to discourage new and small startups from forming, we can both establish strong relationships with suppliers to gain a leg up on others looking to pursue our method of utilizing existing hardware and maintain a high fixed cost – such as a high R&D cost by developing more proprietary algorithms – to deter incomers that have a small amount of seed funding. Finally, we can address the threat of entry from large corporations by realizing that these companies have more overarching goals, so focus on their robotics branch will not be as heavy as on their other branches. As such, we can capture a niche market to detour focus and attention away from us. Fortunately, we have already positioned ourselves in such a situation, in which we target a niche group of customers – the elderly and the disabled. As a result, we see that our competitive landscape as it applies to new entrants can be classified as quite aggressive, but there are indeed routes we can take to dodge much of this aggression.

To mitigate the issue of being locked into a single supplier (see Section 2.2.5), the core strategy is still to generalize our software. This would considerably increase our power, since we would no longer be dependent on any one supplier. Note that as a trend, robotics startups are becoming increasingly common (Tobe, 2013), and we thus anticipate more suppliers coming into the market in the future. As of right now though, suppliers are a strong force that must be considered carefully in our strategy, and we must route efforts to ease this force.

2.4 Market Trends

With an evaluation of competitive forces complete, we end with a discussion of the major trends that influence our project strategy. Aside from the trends of both the changing age demographic of the US – affecting the power of both buyers and substitutes – and the increased interest in the robotic industry – affecting the power of both substitutes and new entrants, another trend to consider is the recent advancement in integrated circuit (IC) technology that has resulted in improved computing performance and reduced cost, resulting in reduced barriers-to-entry and thus further enhancing the threat of new entrants. IC technology has seen consistent improvements in computing power and consistent reductions in cost since their inception. Our industry is directly affected by these advancements; in recent years, more powerful computational devices have generated more robotic technology in the household arena, for engineers are allowed to easily incorporate computing power into the chassis of the robot. This design contrasts with industrial robots, where the computational power is often located in an external computer. The trend is summarized with a concept known as Moore's law, stating that the computational power of the average IC doubles nearly every two years. This trend has been relatively consistent since the early history of ICs. However, there is disagreement among analysts about how much longer this trend will continue (Hoover's, Inc. "Home Health Care Services.," n.d.). The trend has the effect of making our products more functionally efficient and versatile, which reduces the power of substitutes. However, the lower cost of computing technology also reduces the barriers-to-entry in the industry, which increases the power of rivals. Only time will reveal the overall impact that this trend will have.

To summarize, from our strategy analysis, we have deduced that while some competitive forces are certainly in our favor, a few forces bring cause-for-concern and need to be addressed. With adequate industry analysis, we can plan our strategy in order to leverage ourselves into a better position within the market. Summarizing our findings, we have identified within the market both the power of new entrants and the power of suppliers to be strong forces. Consequently, to dampen these threats, we would generalize our software to work across multiple platforms, disincentivizing suppliers from entering the market as well as taking away supplier bargaining power. We would also encourage people to use our product instead of substitutes by having features and functionality that other products do not, at a price point that is not prohibitively expensive.

3 IP STRATEGY

(Authors: Schweigert, Srinivasan, Sun, Su, Jouppi)

Aside from a business standpoint though, we must also consider which legal avenues to take in order to protect our intellectual property (IP): in particular, whether or not our idea is patentable. After all, in many research scenarios such as ours, a patent is the most feasible way to safeguard any IP that is developed. Unfortunately, as this section will argue, patenting our work may not be the most practical path to pursue; however, we do have an alternative strategy better suited to our purposes, in the form of copyright.

We feel that in our more specific situation, the costs of attempting to obtain and enforce a patent far outweigh the benefits, for a number of reasons. One consideration is that the mathematics behind the algorithms we employ are pulled from published research papers, particularly those that deal with robot learning-by-demonstration (Billard et al., 2008). Therefore, the proprietary essence of such research is not ours to claim. By the same token, we cannot patent the ROS (Robot Operating System) software platform upon which we develop because it is open-source and thus, once again, publically available. Most importantly, we do not feel that it is pragmatic to patent the software code itself. This is because software, at its core, is the manifestation of logical deductions, and another group or individual may take a different route of logical deductions to arrive at the same conclusion. Following this train of thought, it is ordinarily quite difficult to obtain and/or protect a patent when the end result can be reached in various ways. As explained by Anthony Klein, an IP attorney at Latham & Watkins LLP, pure software patents remain controversial since “what would constitute patentable subject matter is unclear” (Klein, 2015).

Before investigating an alternative means at protecting our ideas however, it is important that we have the foresight to research whether existing patents overlap with our results. We discovered that the closest patent to our project is entitled: “Method and system for training a robot using human-assisted task demonstration” (Bajaras, 2014). It describes a system for humans to train robots for pick-and-place tasks by moving the robot’s arm while recording trajectory-and-perception data through the robot’s sensory systems. At first glance, it may appear that our project directly infringes upon this patent. However, after delving into the details, this is not the case due to the limited scope of this patent. To give some background on the nature of patents, a patent consists of independent claims and dependent claims; if one does not violate the independent claims, then by definition, one does not violate the dependent claims (Brown & Michaels, PC 2006). Now, many of our project’s similarities with this patent lie in the dependent claims. However, if we can argue that our capstone project does not infringe upon any of the independent claims, then we can legally claim that we do not infringe upon the dependent claims as well – and thus the patent as a whole.

There are two independent claims mentioned in this patent. To quote the first independent claim (claim 1):

> A method for training a robot to execute a robotic task in a work environment, the method comprising: moving the robot across its configuration space … assigning, via the ECU, virtual deictic markers to the detected perceptual features (Bajaras, 2014).

We argue that we do not infringe this claim because our project does not use “virtual deictic markers” – markers based on a representational paradigm that use “selective attention and pointers … to learn and reason about rich complex environments” (Ravindran, 2007). As for the second independent patent claim (claim 2):

> The method of claim 1, wherein moving the robot across its configuration space includes moving at least one of a robot arm and a robot manipulator attached to the robot arm (Bajaras, 2014).

Our project does not use a single “arm and a manipulator”, but rather a dual-armed Baxter-robot. Hence, our project does not violate any of the independent claims and thus none of the dependent claims. Therefore, while this is the closest patent to our idea, we do not infringe upon it and are therefore not required to license from it. Since other existing patents are even less related, a breach of IP is of no worry to us.

With the threat of similar, existing IP out of the way, we can now begin to pursue an alternative strategy of IP protection. After much consideration, we believe that copyright is the most appropriate option – in fact, this happens to be the choice for many software companies. Of course, copyright does indeed present a few risks since, in general, patents protect ideas while copyright only protects the expressions of ideas.

The first risk is the risk of knock-offs: there are ways around copyright such that people can make products very similar to ours but are not in violation of copyright law. This includes implementing our algorithm through a different tactic – one example is converting our code to a different programming language – as well as merely adapting functions from our program. The point is that copyright does not protect our ideas, making it incredibly easy for others to take our ideas and tweak them to look slightly different in their end product. We would need to mitigate this issue by implementing our algorithm across multiple programming languages to prevent the scenarios where someone claims credit on our ideas based on simple modifications.

The second risk is the risk of undetected duplication. It is the first risk in reverse, where certain competitors are indeed copying our code directly, but we have no way of detecting that they are doing so. The reason for this is that we will generally not have the source code of our competitors to compare to our own; all we will have is the compiled functionality that their code is capable of demonstrating. In that sense, it is near impossible to identify specifically if they have violated copyright. Consequently, it is quite difficult to mitigate this risk.

While copyright does offer less protection than patents, it is nonetheless more feasible and realistic to acquire. For instance, copyright is granted automatically when an original work is created, so registration is not required. This simplified procedure immediately eliminates the time and money that we would otherwise need to spend to obtain a patent. Moreover, the duration of copyright is the life of the author(s) plus 70 years, which is plenty of time for us given the short life cycle of software. Furthermore, copyright offers authors the exclusive rights to reproduction, protects against public displays and derivatives of their work, and establishes a public credibility that can attract investment and customers. Licensing can also present itself as a way to increase profit and expand a business.

All in all, it appears that pursuing a patent is not the route for us to go. Instead, a more practical approach at protecting our IP is for us to pursue copyright, due to both the more lenient restrictions and more efficient timeline at obtaining copyright.

4 SYSTEM ORGANIZATION

(Authors: Schweigert, Srinivasan, Sun, Su, Jouppi)

As mentioned before, we are developing a general-purpose robotic software system that incorporates robot-human interactions. In order to implement and test our system, we choose to work on a specific task: having Baxter assembling a coffee table with the help of human. Figure 1 shows the workflow. Figure 2 shows the organization of system components.

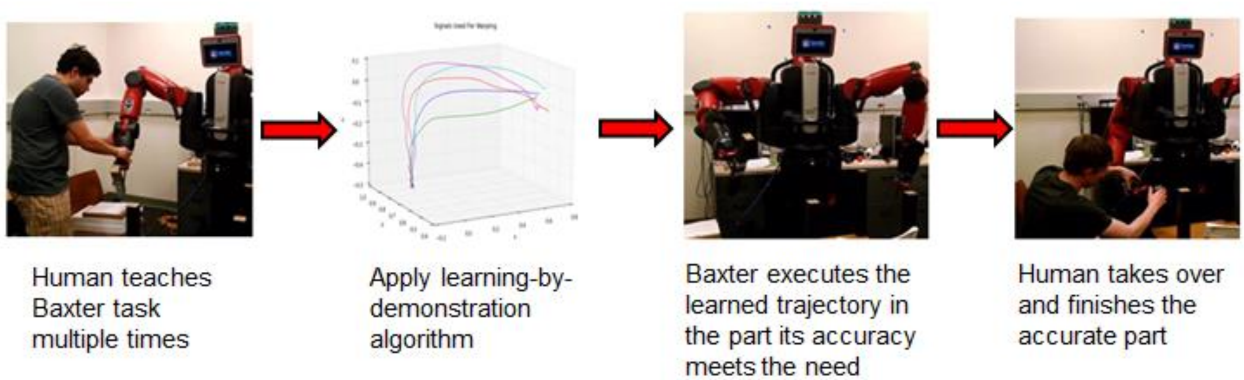

Figure 1: System Workflow

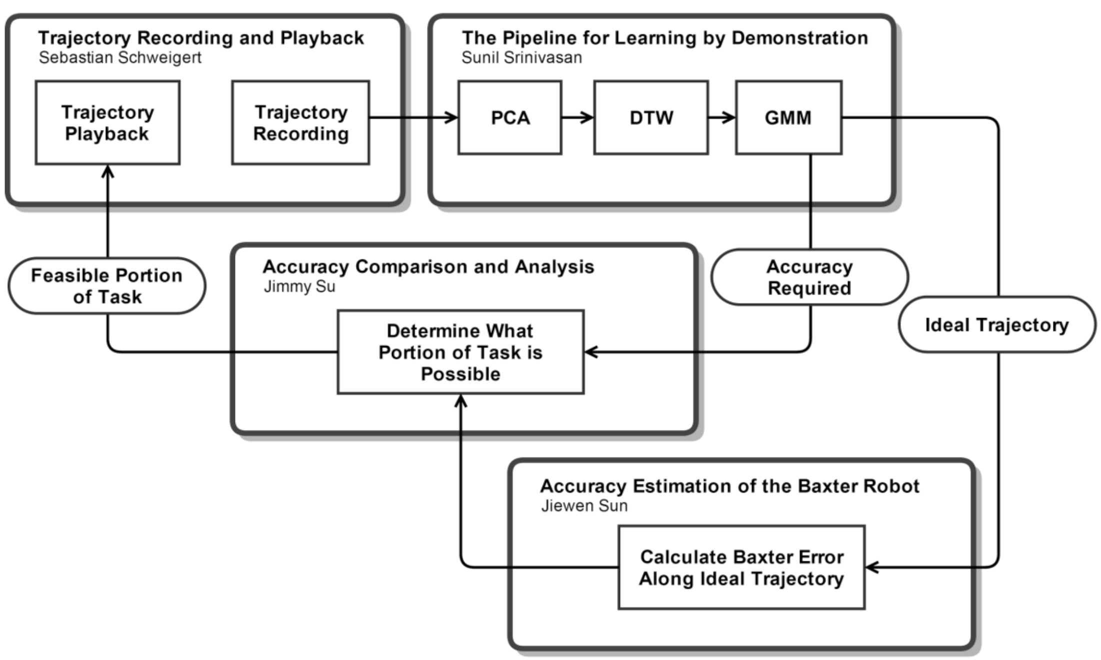

Figure 2: System Component Organization

5 TECHNICAL CONTRIBUTIONS

(Author: Sun)

5.1 Overview

The goal of our project is to design a software system around robot-human interaction, where a human can train and work with cost-effective robots to accomplish a broad range of tasks, including those which require high accuracy. To achieve this goal, we addressed a series of technical issues, which can be divided into four parts: (1) operating the Baxter robot based on ROS; (2) implementing learning-by-demonstration algorithms; (3) estimating the accuracy of the Baxter robot; and (4) integrating the human-robot interactions.

In this paper, I present several methods I adopted to estimate the accuracy of Baxter and analyze the results. More specifically, I addressed the following issue: given a certain joint configuration of Baxter, what is the positioning accuracy of its gripper? Answering this question is an intermediate part of our project: it requires knowledge of robotic operation in ROS, and it plays a key role in the ultimate integration phase, when it comes to compare the robot's accuracy with the task's accuracy requirements.

5.2 Literature Review

The accuracy of a robotic system depends on many different factors, including manufacturing tolerance, sensor resolution, and various environmental conditions (e.g. temperature, humidity). There have been some successful attempts in robot-accuracy estimation, both in academia and industry. In this section, I introduce two main approaches in this field, discuss their pros and cons, and choose one that fits our project.

First of all, the definition of robot positioning accuracy needs to be clarified. A widely accepted definition is: "the ability of the robot to precisely move to a desired position in 3-D space" (Conrad, 2000). In our case, I further refine the definition of accuracy as the ability of Baxter to position its gripper to a desired position in the 3-D Cartesian coordinate system.

The first approach for accuracy estimation is called metrology, which is essentially using sensing equipment to directly track the actions of a robot in repeated experiments. For example, an LED strip can be attached on a robot's gripper, with a camera fixed at a certain point (Alici, 2005). Through calibration and image processing, we can estimate the accuracy statistically. Of course, there are also other sensing methods, like optical tracking, inertial measurement, and trilateration probing. In general, the advantage of the metrological approach is that it is explicit and precise. However, it is not widely utilized due to its high cost and selectivity.

Another approach for accuracy estimation is called accuracy parameterization. That is to say, a model is built upon a set of parameters and it calculates the accuracy estimation based on these parameters. Accuracy parameterization models are greatly based on robot kinematics and dynamics, more specifically, on forward kinematics and the Jacobian matrix. In general, this approach is not as accurate as the metrology approach. However, it is more widely adopted because of its versatility, simplicity, and low cost.

In our project, we don't require a highly precise estimation of the accuracy. Therefore, the second approach, accuracy parameterization, is a better fit for us. Our accuracy parameterization model is based on the common robot-kinematics models, with some customized modifications to make it more effective. First, I made it as simple as possible. Since the data size of our project is quite large, simple model enables more efficient computation. Next, I tailored the outputs to make sure that the data format is consistent with the rest of our system, which facilitates the effectiveness of the integration.

5.3 Implementation

The implementation of accuracy parameterization is based on many mathematical tools, especially linear algebra and probability theory. First of all, I implemented the kinematics model for Baxter. Then I decided on the related parameters. Lastly, I built two models for accuracy parameterization and applied different ways to quantify the results. All code is written in Python, which is widely used in robotic systems nowadays.

5.3.1 Baxter Robot Model

Baxter is a humanoid robot with two arms. Each arm has seven joints, that is, seven degrees of freedom, as illustrated in Figure 3. In the case of Baxter, we consider a manipulator to be one of its arms. In this paper, I only analyze one arm because the other arm is symmetric. In the following parts, when I say manipulator, it refers to the left arm of Baxter. Additionally, where there is mention of an end effector, it refers to the gripper.

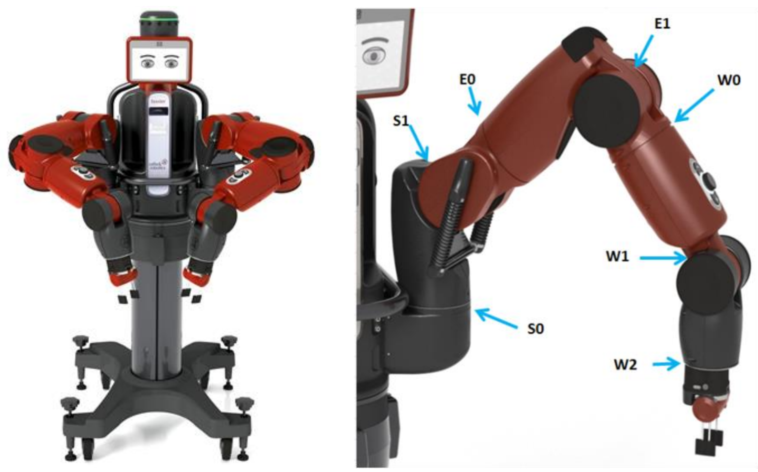

Figure 3: Baxter Robot (Rethink Robotics 2015)

5.3.1.1 Forward Kinematics

Coordinate transformation is a fundamental mathematical tool in analyzing robotic manipulators. The most common application in coordinate transformation is forward kinematics. Forward kinematics answers the following question: given the angle of each joint in the manipulator, what is the transformation (rotation and translation) of the tool frame with respect to the base frame? I used the following steps to write the forward kinematics map for Baxter's serial chain manipulator (Bestick 2014).

First of all, we need to define a zero configuration for the manipulator, as shown in Figure 4. That is to say, in this configuration: $\theta = [\theta_1, \theta_2, \theta_3, \theta_4, \theta_5, \theta_6, \theta_7] = 0$. In this way, we are able to set values for the seven joints. The base frame (S) is at the center of Baxter's body base while the tool frame (T) is at the center of its gripper.

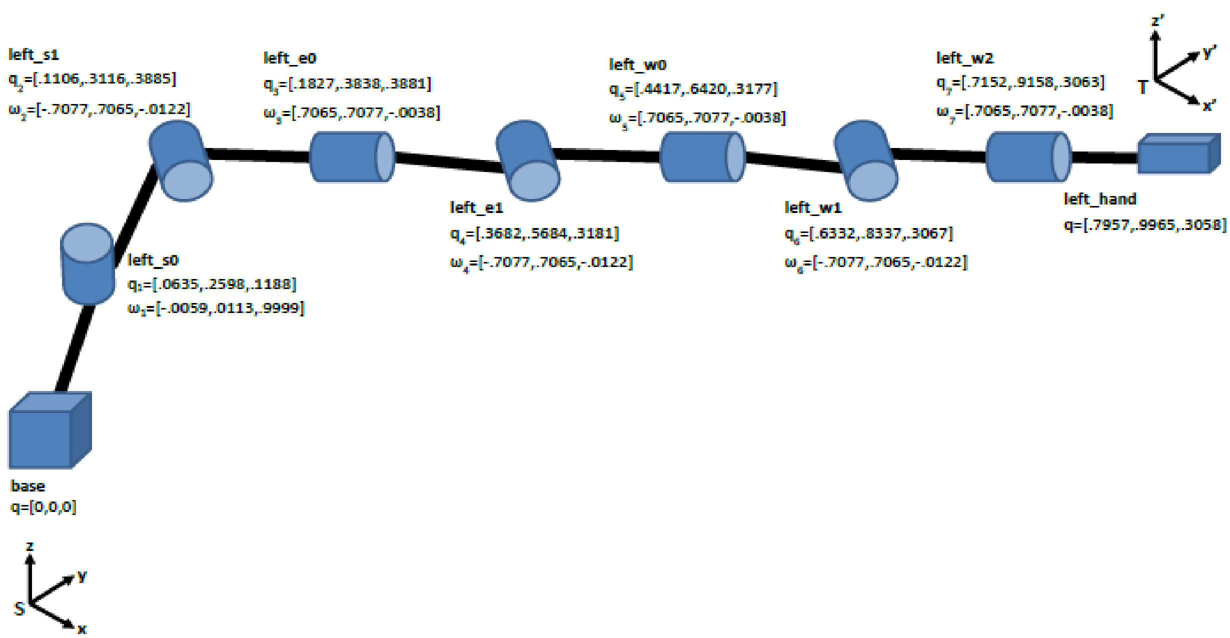

Figure 4: A Zero Configuration of Baxter's Left Arm (Bestick 2014)

The transformation of this zero configuration is in the form:

$$g_{ST}(0) = \begin{bmatrix} R_{ST}(0) & q \\ 0 & 0 \end{bmatrix}$$

We calculate the twist for each joint in the manipulator accordingly:

$$\xi_i = \begin{bmatrix} -\omega_i \times q_i \\ q_i \end{bmatrix}$$

Then we can calculate the product-of-exponential map for the complete manipulator, or rather, the homogenous coordinate transformation, which is a 4×4 matrix:

$$g_{ST}(\theta_1, \theta_2, \theta_3, \theta_4, \theta_5, \theta_6, \theta_7) = e^{\widehat{\xi_1}\theta_1} e^{\widehat{\xi_2}2} \cdots e^{\widehat{\xi_7}\theta_7} g_{ST}(0)$$

Following these steps, I built a forward kinematics function which inputs seven joint angles and outputs the transformation of the tool frame with respect to the base frame.

5.3.1.2 Jacobian Matrix

In addition to the relationship between joint angles and end-effector transformation, one often makes use of the relationship between joint velocities and end-effector velocities (Murray 1994), which can be expressed in a Jacobian matrix.

The Jacobian of a manipulator is traditionally described as the differentiation of the forward-kinematics map.

If we represent forward kinematics as a mapping:

$$g: R^n \to R^p$$

The Jacobian is the linear mapping:

$$\frac{\partial g}{\partial \theta}(\theta): R^n \to R^p$$

There are two kinds of Jacobian matrices: the spatial Jacobian matrix and the body Jacobian matrix. I am using the body Jacobian matrix because it measures the velocity of the tool frame with respect to the base frame, which is the actual velocity we are looking for.

The body Jacobian matrix is calculated according to the following formulas:

$$V_{ST}^{b} = J_{ST}^{b}\dot{\theta}$$
$$J_{ST}^{b} = \begin{bmatrix} \xi_1^{\dagger} & \xi_2^{\dagger} & \dots & \xi_n^{\dagger} \end{bmatrix}$$
$$\xi_i^{\dagger} = Ad^{-1}_{\left(e^{\hat{\xi}_i \theta_i} \dots e^{\hat{\xi}_n \theta_n} g_{ST}(0)\right)} \xi_i$$

In our case, Baxter has seven joints: thus the body Jacobian matrix is a 6×7 matrix. In the following sections, when mentioning the Jacobian, it refers to the body Jacobian.

5.3.2 Joint Error

Given the above models of the Baxter robot, we can use the angle error of each joint to calculate the error of the end-effector, so as to analyze the end-effector positioning accuracy. As stated in the hardware specification of Baxter (Rethink Robotics, 2014), the joint error is influenced by both the sensor resolution and the controller tolerance. The typical sensor resolution is ± 0.1 degrees while the default controller tolerance is ± 5 degrees. Therefore, we can estimate that the accumulated error for each joint is ± 5.1 degrees $= \pm 0.089$ rads.

After deciding on the value of joint error, I further assumed that it follows a Gaussian (Normal) distribution, which is a common distribution for a robotic system. It was also reasonable to assume that the confidence level is 95%. That is to say, the joint error is 95% certain to be in the range [– joint_error, + joint_error]. In other words, the joint angle is 95% certain to be in the range [expected_joint_angle – joint_error, expected_joint_angle + joint_error]. Due to the linear mapping of Jacobian, these assumptions also guarantee that, in common cases (when the robot is not in singularity and it doesn't reach its joint limit), the positioning error of the end effector is also in the form of a Gaussian distribution (with a confidence level of 95%). This matches the outputs of the Gaussian Mixture Model (Calinon 2007) used in our learning-by-demonstration algorithms.

5.3.3 Accuracy Parameterization Model

Given the foregoing analysis of the Baxter model and its joint errors, there are different ways to build the accuracy parameterization model. In this part, I present two kinds of models: a model based on forward kinematics and a model based on the Jacobian.

5.3.3.1 Based on Forward Kinematics

An intuitive way of parameterizing accuracy is using forward kinematics. To illustrate, assume that, given a target joint configuration, each joint jitters in a certain range according to the joint error determined above. By generating sets of joint configurations (let us call them testing points), we can use forward kinematics to calculate the scattered positions of the end effector. This is a purely statistical approach.

The way to generate the testing points is quite explicit. For example, given an expected set of joint angles (listed in Table 1) and the joint error (± 0.089 rads), each joint has three testing values: a lower bound, an expected angle and an upper bound. A set of testing points is a combination of seven joint angles (e.g., $\theta =$[-0.769, -0.123, 0.00422, 1.51, -0.437, -0.0349, -0.381]). The total number of combination is $3^7 = 2187$. That is to say, there are 2187 sets of testing points in total.

Table 1: Explanation of Testing Points

Unit: rads	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7
Expected Angles	-0.769	-0.0337	0.00422	1.60	-0.526	-0.0349	-0.292
Lower Bound (Expected Angles - 0.089)	-0.858	-0.123	-0.0848	1.51	-0.615	-0.124	-0.381
Upper Bound (Expected Angles + 0.089)	-0.680	0.0553	0.0932	1.69	-0.437	0.0541	-0.203

After generating testing points, we can use forward kinematics to calculate the corresponding positions of the end effector. One set of testing points uniquely determines an end effector position, thus there are 2187 resulting points. The expected resulting point is the position of the end effector when each joint is in its expected angle. The scatter plot is shown in Figure 5, with the expected resulting point (0.5017, 0.2647, -0.2703).

From the graph, we can see that this result has large biases in the X and Y directions but has relatively little bias in Z direction. This is likely due to the fact that, at this configuration, the shoulder joint (s0, whose rotation axis is approximately in the Z axis) has the greatest effect to the position of the end effector. To clarify, we project this scatter plot into X-Y plane, as shown in Figure 6. The red point in the center is the expected resulting point.

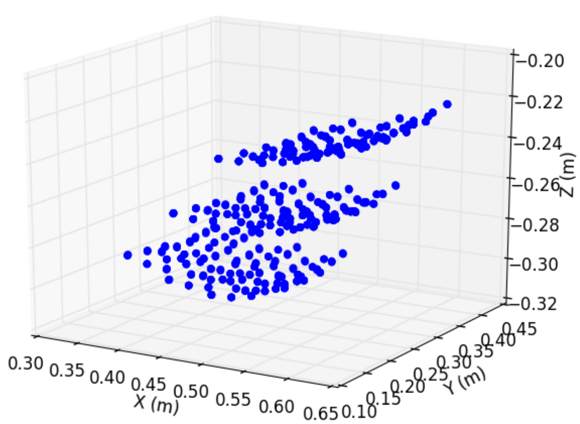

Figure 5: 3-D Scatter Resulting Points

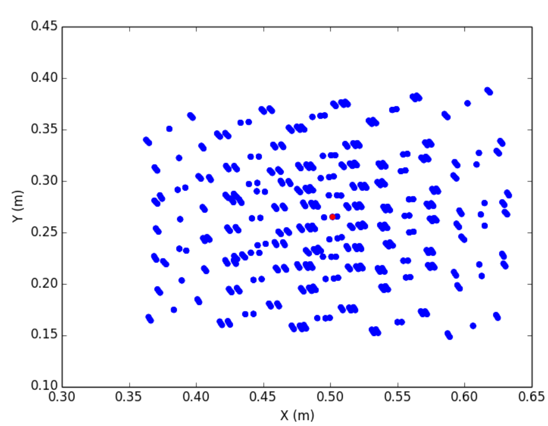

Figure 6: 2-D Scatter Resulting Points

Subsequently, we can perform density analysis on this data. To accomplish this, I use a class called "gaussian_kde" in SciPy (the scientific tools library in Python) to achieve kernel distribution estimation based on Gaussian kernels. In other words, given a set of data, we can calculate the density in a certain dimensional space. Figure 7 shows the density map at resulting points. The density is represented by color: the warmer (redder) the color, the denser the data. Generally speaking, the data is denser near the center point, which follows a Gaussian distribution. This further verifies that our assumption of a Gaussian distribution is reasonable.

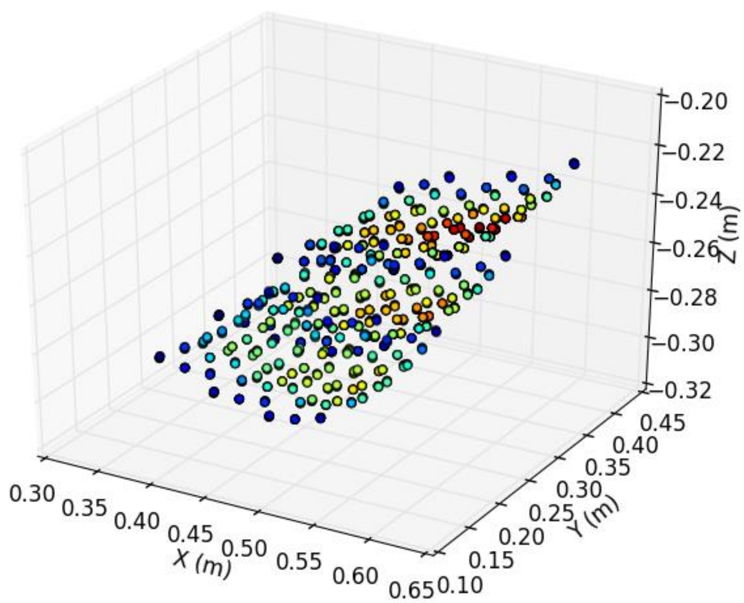

Figure 7: Density Map of Resulting Points

Given these results, we still need a way to quantify the accuracy. An intuitive solution is to use the volume of a convex hull. The convex hull is the smallest convex that contains a set of data points in the Euclidean space. Therefore, it describes the bound of scatter points. The smaller the convex hull, the more accurate Baxter is. There is a class in SciPy that can be used to calculate a convex hull. Functions return simplexes, vertices, volume, etc. The convex hull of this example is shown in Figure 8, with a volume of 0.00229 m^3.

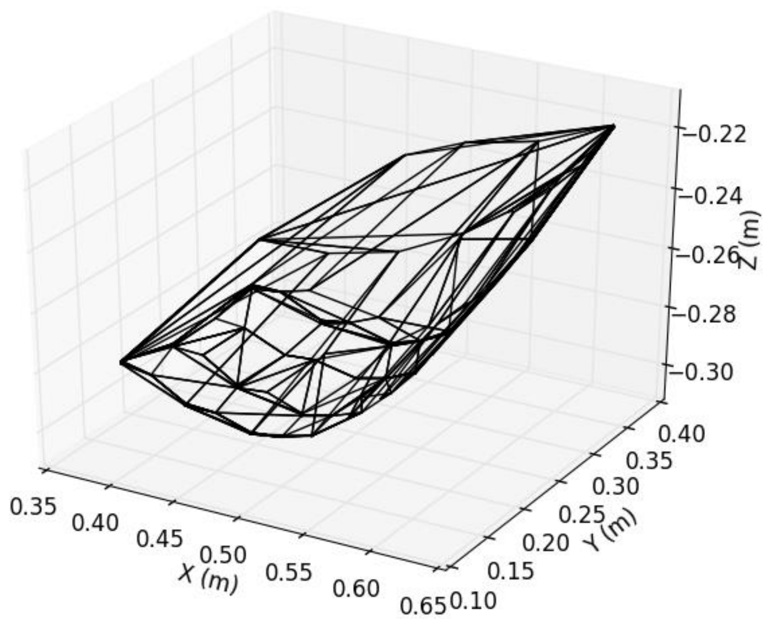

Figure 8: Convex Hull of Resulting Points

This accuracy parameterization model, based on forward kinematics and statistics, is accurate and complete. It offers a large amount of data which allows for ample analyses. However, it is computationally expensive. Calculating accuracy for just a single example joint configuration takes about 5 seconds. In our project, we ultimately needed to calculate approximately 1000 joint configurations, and that calculation would have taken quite a long time. Therefore, this approach was not feasible. Moreover, in more general cases, where a robot has N joints, the complexity of this algorithm is $O(3^N)$, which is exponential; and an algorithm with exponential complexity should be used with caution in practice.

5.3.3.2 Based on Jacobian Matrix

An alternative model of accuracy parameterization is based on the Jacobian matrix. As mentioned before, the Jacobian matrix reflects the relationship between joint velocities and end-effector velocities:

$$V_{ST}^{b} = J_{ST}^{b}\dot{\theta}$$

According to the Euler Derivative Method, the above equation can be rewritten as:

$$\begin{bmatrix}\Delta q\\ \Delta\omega\end{bmatrix} = J_{ST}^{b}\Delta\theta, \text{ where } \Delta q = \begin{bmatrix}\Delta x\\ \Delta y\\ \Delta z\end{bmatrix}$$

Therefore, by multiplying the Jacobian matrix with the joint error, we can know the estimated bias of the end effector in the X, Y and Z directions. To get a bound of the bias, simply accumulate the absolute value of bias generated by each joint error, which represents the worst case in the current configuration.

Taking the same joint configuration from Table 1 as an example, the Jacobian matrix is then:

```
[[-0.2071  0.8708 -0.9719 -0.2858 -1.0637 -1.3459 -1.6625]
 [ 0.2218  1.0139 -0.4278 -0.8549  1.2624  0.4668  1.5563]
 [ 0.313  -0.0509  0.3958  0.3772 -0.1881 -1.1954  1.5937]
 [-0.6988 -0.5112 -0.4991 -0.9563  0.5398 -0.66    0.7951]
 [-0.714   0.475   0.4468  0.266   0.5498  0.0205  0.3251]
 [ 0.0434  0.7164 -0.7425 -0.1216  0.6375  0.7511  0.5119]]
```

By multiplying it with the joint error 0.089, we can then estimate the biases in the X, Y and Z directions. This tells us the largest distance the end-effector can vary in the three Cartesian directions, which can be represented as a bounding box in Figure 9. The red point is the expected resulting point, which is at the center of the box. The result is in accordance with the result of the scatter plot from the previous approach.

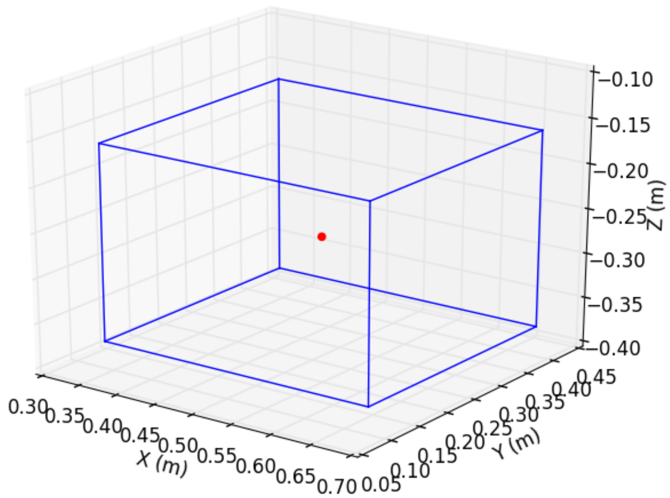

Figure 9: Bound of the End Effector Position

This accuracy parameterization model, based on the Jacobian, is also an effective approach. It is less precise than the former model, but it is computationally cheaper, with linear complexity. More importantly, this approach works well to maintain consistency in our overall system. That is to say, the results of this accuracy parameterization model are comparable to the results of the learning-by-demonstration portion of our capstone project. Under the assumption of a Gaussian distribution, the learning-by-demonstration algorithms can also generate 95% confidence intervals in Cartesian coordinates. At a certain point in the learned trajectory, this confidence interval can also be interpreted as a bounding box, which represents the accuracy requirements of the task at this point.

5.4 Results and Discussion

As discussed above, I have built two models for accuracy parameterization. The first model is based on forward kinematics, where the accuracy is quantified by the volume of the convex hull. The second model is based on the Jacobian matrix, where the accuracy is quantified by a bounding box in 3-D Cartesian space. Note that these two models are both effective and can be applied to other robots as well. However, considering the limitation of computation time as well as the consistency of data in our project, we decided to use the bounding-box model based on the Jacobian matrix as our final approach for Baxter accuracy estimation. We use the first model based on forward kinematics for verification only.

Figure 10 shows an offline workflow of our system. First of all, a set of trajectories are recorded and saved in csv files. Then, the means (representing the learned trajectories) and standard deviations (representing the accuracy requirements of the task) are calculated by our learning-by-demonstration algorithms, which are based on Gaussian Mixture Model (GMM) with a pre-processing of Principal Component Analysis (PCA) and Dynamic Time Warping (DTW). Estimating Baxter's inherent accuracy requires the joint angle configurations of the learned trajectories as inputs, so I use MoveIt! tools to compute a plan. Joint angles are read from the plan and then fed to the accuracy estimation part to calculate the bounding boxes. By comparing the bounding boxes with the confidence intervals, we will know when Baxter is accurate enough to accomplish the task and when it is not, which is the core idea of our project.

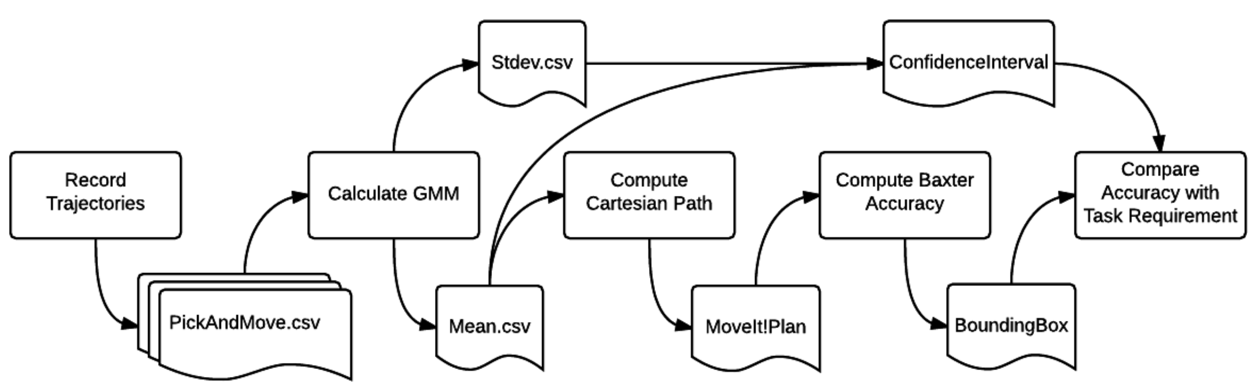

Figure 10: Offline Workflow

In our demo, Baxter learns the task of assembling a coffee table. A human holds one of Baxter's arms and demonstrates the following action: pick up a table leg from a random starting point, move it close to a screw hole on the table surface, and plug it in. While doing this, record both the position and the orientation of the end effector. Repeat this for five times and save these trajectories in separate csv files. Process the data using the learning-by-demonstration algorithms, and a mean trajectory is learned. The original trajectories (dash lines) and the learned trajectory (full line) are shown in Figure 11.

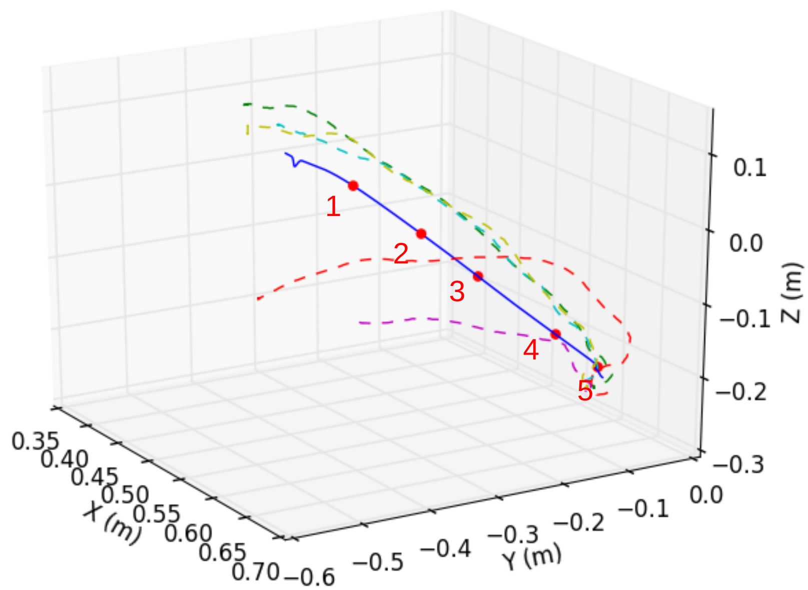

Figure 11: Trajectories and Example Points

Pick five points on the learned trajectory as example points and calculate Baxter's positioning accuracy at these points. Example results are listed in Table 2. For example, for the first point, Baxter's end effector may be off by 0.190 m in the X direction, 0.206 m in Y direction and 0.297 in Z direction, in the worst case. This also shows that Baxter's accuracy varies at different configurations.

Table 2: Baxter Accuracy Estimation at Example Points

Number	Position (m)	Bias in X (m)	Bias in Y (m)	Bias in Z (m)
1	(0.508, -0.310, 0.024)	0.190	0.206	0.297
2	(0.548, -0.249, -0.034)	0.150	0.230	0.229
3	(0.580, -0.195, -0.086)	0.121	0.233	0.112
4	(0.623, -0.122, -0.157)	0.099	0.113	0.098
5	(0.644, -0.079, -0.199)	0.071	0.103	0.129

In conclusion, this accuracy estimation approach works quite well. We can use it to calculate Baxter's positioning accuracy at any configuration. It is effective, computationally efficient, and can be integrated into the system easily.

6 CONCLUDING REFLECTIONS

(Author: Sun)

In general, we accomplished the majority of our project plan. We accumulated a sizeable knowledge base of using both the Baxter robot and the ROS platform. We successfully implemented a set of algorithms for learning-by-demonstration, adopted different approaches for accuracy parameterization, and presented a demo with these parts fully integrated. In our demo, Baxter is able to learn a simple action from a human, perform the learned trajectory, and identify whether its accuracy meets the task's accuracy requirements, so as to segment the task. Additionally, we made our code general and modularized so that it can be applied to other robotic systems with good scalability.

A part of the project that we have not fully worked out is how the human should help the robot when the robot is not capable of finishing a task. For example, when Baxter is not accurate enough to finish the peg-in-hole and screwing segments of the task, what is the best way for a human to help? We considered three possible solutions: (1) a human takes over the task; (2) a human offers a small amount of fixture guidance and Baxter finishes the rest; or (3) a human enables the force sensor on Baxter's gripper that can be used to detect the position of the hole during the adjustment process. All in all, robot-human interaction is a fairly new area, and there are various interesting related topics for future research.

Apart from our technical accomplishments, we also gained experience in project management. Our team had good cooperation: we split tasks properly according to our different expertise, and we used GitHub to manage our source code to achieve effective collaboration. It is inevitable to have encountered some form of difficulty, for example, setting up and debugging took more time than we had anticipated. However, in the end, we successfully adjusted our schedules based on a limited time frame to deliver the best results that we could.

REFERENCES

Alici, G., & Shirinzadeh, B. (2005). A systematic technique to estimate positioning errors for robot accuracy improvement using laser interferometry based sensing. Mechanism and Machine Theory, 40(8), 879-906.

Barajas, L. G., Martinson, E., Payton, D. W., & Uhlenbrock, R. M. (2014). Method and system for training a robot using human-assisted task demonstration. *U.S. Patent No. 8,843,236*. Retrieved from https://www.google.com/patents/US8843236

Bestick, A., Buchan, A. (2014). Forward Kinematics/Coordinate Transformations EE125 Manual, Lab 3.

Bi, F. & Mac, R. (2015). Drone Maker 3D Robotics Raises $50 Million in Latest Round. *Forbes*. Retrieved from http://www.forbes.com/sites/frankbi/2015/02/26/drone-maker-3d-robotics-raises-50-million-in-latest-round/

Biesada, A. (n.d.). Hoover's Company Profiles: iRobot Corporation. *Hoover's, Inc.* Retrieved from http://subscriber.hoovers.com/H/company360/fulldescription.html?companyId=132607000000000

Billard, A., Calinon, S., Dillmann, R., & Schaal, S. (2008). Robot Programming by Demonstration. *Springer Handbook of Robotics*, 1371-1394.

Brault, M. (2012). Americans With Disabilities: 2010 - Household Economic Studies. *Economics and Statistics Administration, US Department of Commerce*. Retrieved from http://www.census.gov/prod/2012pubs/p70-131.pdf

Brown & Michaels, PC. (2006). How do I read a patent? - the Claims. *Brown & Michaels*. Retrieved from http://www.bpmlegal.com/howtopat5.html

Calinon, S., Guenter, F., & Billard, A. (2007). On learning, representing, and generalizing a task in a humanoid robot. Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on, 37(2), 286-298.

Censky, A. (2011). Older Americans are 47 times richer than young. *CNN Money - The New American Dream*. Retrieved from http://money.cnn.com/2011/11/07/news/economy/wealth_gap_age/

Conrad, K. L., Shiakolas, P. S., & Yih, T. C. (2000). Robotic calibration issues: Accuracy, repeatability and calibration. In Proceedings of the 8th Mediterranean Conference on Control and Automation (MED2000), Rio, Patras, Greece.

Crompton, J. (2014). IBISWorld Industry Report 33399: Power Tools & Other General Purpose Machinery Manufacturing in the US. *IBISWorld.* Retrieved from http://www.ibis.com

Ellis, B. (2013). Nursing home costs top $80,000 a year. *CNNMoney (New York)*. Retrieved from http://money.cnn.com/2013/04/09/retirement/nursing-home-costs/

Hoover's, Inc. (n.d.). Hoover's Company Profiles: Clearpath Robotics Inc. *Hoover's, Inc.* Retrieved from http://subscriber.hoovers.com/H/company360/overview.html?companyId=244260399

Hoover's, Inc. (n.d.). Industry Report: Home Health Care Services. *Hoover's, Inc*. Retrieved from http://subscriber.hoovers.com/H/industry360/overview.html?industryId=1383

Klein, A. R. (2015). Intellectual Property Basics for Technology Companies. *Latham & Watkins*.

McMillan, R. (2015). Now We Can Build Autonomous Killing Machines. *Wired*. Retrieved from http://www.wired.com/2015/02/can-now-build-autonomous-killing-machines-thats-bad-idea/

Murray, R. M., Li, Z., Sastry, S. S., & Sastry, S. S. (1994). A mathematical introduction to robotic manipulation. CRC press.

Owano, N. (2011). Willow Garage slashes price (and arm) of PR2 robot for research. *PhysOrg*. Retrieved from http://phys.org/news/2011-08-willow-garage-slashes-price-arm.html

O'Donnell, F. (2015). Issues and Insights. *Mintel: Senior Lifestyles - US - December 2013*. Retrieved from http://academic.mintel.com/display/689700/

Ortman, J. M., Velkoff, V. A., & Hogan, H. (2014). An Aging Nation: The Older Population in the United States. *Economics and Statistics Administration, US Department of Commerce*. Retrieved from http://www.census.gov/prod/2014pubs/p25-1140.pdf

Porter, M. E. (2008). The Five Competitive Forces That Shape Strategy. *Harvard Business Review 86*(1), 25-40.

Ravindran, B., Barto, A. G., & Mathew, V. (2007, January). Deictic Option Schemas. In *IJCAI* (pp. 1023-1028).

Rethink Robotics (2015). Build a Baxter. *Rethink Robotics.* Retrieved from

Rethink Robotics. (2015). http://www.rethinkrobotics.com/baxter/.

Rethink Robotics. (2014). Hardware Specifications.

Robotic Operating System. (2015). http://www.ros.org/.

Tobe, F. (2013). A glimpse at our robotic future: 235 start-ups reviewed. *Robohub*. Retrieved from http://robohub.org/a-glimpse-at-our-robotic-future-235-start-ups-reviewed/