Robotic Comprehension of Viscosity

JosieKate Cooley and Jacqueline Gibson

Department of Computer Science, The University of Texas at Austin
cooley.jk8@gmail.com, jacqueline.gibson@utexas.edu

Abstract

To the best of our knowledge, at the time of this report's writing there are no published attempts at teaching a robot to distinguish substances by their current material state (e.g. liquid or solid) or to classify distinct liquids based solely on haptic data. This project aims to create a reliable and customizable data collection program as a base for future work in machine learning. The program created performs a series of behaviors at different speeds and is successful in collecting large amounts of effort, force, and position data from a Kinova robotic arm, with four different stirring implements used to create motions in four distinct substances. More testing, including the implementation of a learning algorithm, should be done to fully realize the goals of this project. This research represents an intersection between the fields of personal robotics and autonomous intelligence, with potential applications in such disparate areas as accessibility robotics, food preparation, and automotive technology, among others.

Introduction and Related Work

Liquids are an integral part of everyday life. All living organisms rely on liquids, especially water, to function and thrive. That said, the majority of autonomous intelligence research does not concern robot interaction with liquid. As robotics becomes more incorporated with everyday society, it is important that these systems exhibit comprehensive knowledge of the world of which they are a part. Similarly, if autonomous robotics researchers want to expand their reach or collaborate with other fields, it is important that their work too features a broader set of applications.

Identifying and interacting with liquids is important in everything from food production to automotive care, yet this is a virtually unexplored region of robotics, with the exception of some niche work (e.g. work done with submerging robots in water completely and using them to collect data in underwater conditions). Current research efforts include training robots to categorize objects as containers to hold water and utilizing neural networks to detect and track liquids [5]. Work has been done by research collaborators such as Schenk and Fox to use machine learning techniques to get robots to understand the movement of liquids, but these movements are so complex that so far only computer-simulated liquids have been analyzed, and these necessarily differ from real-world liquids [6, 7]. It is agreed that robots need to be able to reason about liquids in order to function in any useful way, but this reasoning is difficult, not only because robots working directly with liquids is a challenge, but because liquids themselves present a unique physical challenge that roboticists are still looking for algorithms to solve. Teams like Klapfer, Kunze, and Beetz have proposed such programs, but even the simulations of liquids often have problems that must be addressed before a simulated robot can attempt to analyze them [6]. Groups like Griffith and Shane et al. have done experiments with robotic comprehension objects, but most of these objects are either solid or empty and none contained liquids [5]. In a highly related experiment, Elbrechter, Maycock, Haschke, and Ritter used 3D video feed to teach a robot about the viscosities of different liquids [3]. This team had a robot shake containers of various liquids, and had success in getting the robot to perceive differences among their viscosities. The robot did not directly interact with the liquids, and the hundreds of tests conducted were per-formed manually. The liquid set used was relatively small, and consisted of relatively similar liquids, such as milk and buttermilk. This was largely because the robot used was a kitchen assistant robot, so the experiment focused on kitchen-related liquids and motions humans might use to determine viscosity when preparing food. The motion used was a back-and-forth motion designed to create movement on the surface of the liquid being tested, as in food preparation less direct contact is considered ideal for sanitary purposes. Alternative approaches to liquid classification are non-invasive. For instance, placing specially designed ultrasound sensors within a container, and measuring "acoustic impedance," as well as other calibrations to characterize liquids, specifically through comprehending their density; this proved beneficial for systems developed with the culinary industry in mind [1][2]. Similarly, Hara et al, utilized light refraction and manipulation combined with depth cameras to aid robots in detecting the presence of a liquid in small containers, such as a cup; this approach, however proved limited in scope as it was only effective for clear liquids and not for substances with lower rates of transparency [8].

Our goal is to build upon the breakthroughs made with water-related robotics research and expand the working substances to include additional liquids as well as solids. Ultimately, our project aims to explore the cross section of personal robotics and autonomous intelligence, representing an opportunity to create systems which are able to learn about all elements of their environments and better equipped to perform a greater variety of tasks more efficiently and robustly.

Overview of the Concept of Viscosity

Viscosity is measured in pascal seconds (Pa s) or poise (dyne second per square centimeter [dyne s/cm2]), with ten poise being equal to one pascal second. A higher viscosity corresponds to a thicker (more viscous) fluid, with the viscosities of semisolids and solids being the highest - for example, tar at room temperature [20^{o} C] has a viscosity of 30,000, and water at the same temperature has a viscosity of 1. Viscosity is highly dependent on temperature, with a higher temperature yielding more viscous gases and less viscous liquids [2]. This is an important consideration for this project, as varying temperatures will yield inconsistent results. We solely conduct tests with substances generally stores at room temperature. This is to prevent skewing that could occur due to liquids exhibiting an unusual viscosity as a result of ecologically invalid circumstances (e.g. as milk is often refrigerated, it potentially curdles when left in warmer temperatures and this would affect its viscosity).

Table 1: Courtesy of Physics.info [4]

Substances	Viscosity at Room Temperature(in Pa s)
Blood (Human)	3-4
Canola Oil	33
Caulk	1000
Chocolate Syrup	10-25
Corn Syrup	2-3
Honey	10
Ketchup	50
Lard	1000
Machine Oil (Light)	102
Machine Oil (Heavy)	233
Maple Syrup	2-3
Milk	3
Molasses	5
Mustard	70
Olive Oil	84
Peanut Butter	150-250
Rubbing alcohol	2.4
Sour Cream	100
Soybean Oil	69
Tar	20
Vegetable Shortening	1200
Water	1

Methodology

Technical Approach: Software

Our program is a ROS node written in C++ using the Ubuntu operating system. It utilizes ROS messages (specifically from the topics `sensor_msgs::JointState`, `geometry_msgs::WrenchStamped`, and `geometry_msgs::PoseStamped`) to gather data on the position, effort, force, and torque from the eight joints of the Kinova robotic arm. We create global vectors to store the messages received from these topics, as well as a global flag to indicate when to store the data to avoid recording when the motions of the arm are not producing relevant data points. Data is pushed to the vectors in callback functions when applicable during trials, and the vectors created are passed to a file writing method once a behavior is completed. Using this, we are able to write all the data for a given trial, using the following naming convention: `path_behavior_stirringImplement_stir_substanceStirred_dataMessage_trial_trialNumber_run_runNumber.csv`.

The messages in the vector are passed through loops and are broken into component parts; these components are stored in the aforementioned .csv files. Each file also stores the timestamp from the message header, as well as data specific to the message. The `geometry_msgs::WrenchStamped` message stores the x, y, and z force and torque from `/mico_arm_driver/out/tool_wrench`, which is important for determining the effort the arm uses to move through different substances. The `geometry_msgs::PoseStamped` message gives the 3D position coordinates and orientation quaternion of the arm's end effector, which helps determine how the resistance of each substance affects the position of the arm, especially when compared to the resistance results from the control run using air. The `sensor_msgs::JointState` message is used twice, yielding both the position and velocity information of the eight joints (six arm joints and two fingers) and the efforts of the joints, all of which are presented as arrays of float64 objects. These arrays are used to determine the impact the liquid has on the movement and force output of the entire arm. Combining this information with the wrench and pose data can be used to aid the robot in developing a more comprehensive picture of the way its motions affect each of the stirred substances; this understanding represents a rudimentary understanding of viscosity and serves as the foundation for the future training.

Figure 1: A screen capture of the C++ code for the beginning of a run

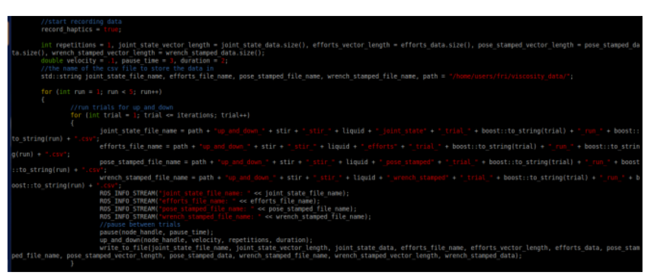

Figure 2: A screen capture of an example call to rosrun robotic_comprehension_of_viscosity behavior node showing user input and the subsequent file nomenclature

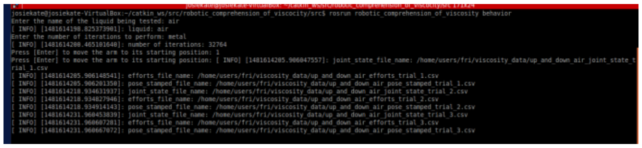

In a typical run (using the command `rosrun robotic_comprehension_of_viscosity behavior`), the program asks the user to input the name of the material used to stir, the name of the substance being tested, and the number of iterations to run, all of which are used to give the .csv files specific and meaningful names. The user would also be asked to press enter to begin the run, to give the researchers time to ensure everything has been properly set up before trials begin: the stirring implement needed to be manually placed in the robot's grasp before testing could begin, and the container of the substance being stirred needed to be in a particular orientation at a particular height, dependent on the stirring implement being used. Once this is verified, the robot would move to a preprogrammed position (found by echoing the robot's Cartesian velocity at the preferred position) using the `segbot_arm_manipulation::moveToPoseMoveIt` method in an attempt to create consistency of starting position across trials. The robot begins its movements at the slowest velocity (.1) with the longest duration (2 seconds); these values are doubled and halved respectively after each set of three movements is completed. Please see the appendix for images of the robot in the starting position and various behaviors.

The first motion to be called is the up-and-down behavior, during which the arm rises for the duration to exit the liquid completely, then lowers (running at the speed of velocity * -1) back into the liquid, returning to the starting position. During this and every behavior, data is published at 40Hz, ultimately creating thousands of data points for future analysis. The data gathered from the up-and-down behavior provides valuable knowledge about the entry of the stirring implement into the substance (not recorded during the initial entry when the robot is assuming the starting position), and also helped to distinguish between solid substances (such as beans), less viscous liquids (such as water), and thicker liquids (such as shampoo), as thicker liquids generally adhered to the stirring implement, while solid substances and less viscous liquids did not. The next motion is the back-and-forth behavior, which, similarly to the up-and-down behavior, travels backwards at negative velocity for the duration, then forwards at positive velocity for the same amount of time. This is again useful for distinguishing between liquid substances and solid substances, as the solid substances tended to stay where they had been pushed, providing less resistance during the forwards part of the motion, as opposed to the liquid substances, which would move back to equilibrium immediately after being pushed. This motion also provides data about the linear movement of the substance being tested, important in studying the change in position of the joints. The final motion is the circle behavior, which moves the stirring implement in a full counterclockwise circle, most reminiscent of traditional human stirring. Unlike the previous two behaviors, the circle maintained a constant velocity, as the motion is difficult to execute on a scale small enough to fit in the container and thus difficult to modify. The circular motion provides the most ecologically valid data for robots attempting to recreate human motions such as mixing cake batter, although the unchanging velocity means this motion pro-vided different data than the other behaviors. There was only data gathered for one-fourth of the velocities of the other behaviors, as this behavior ran at one speed rather than four as the others did. However, since the circular behavior was repeated in every trial (staying the same as the velocities of the other behaviors increased), there was four times as much data gathered for this single velocity. The decision to repeat the circular motion with every set of behaviors was made based on the large and varied amount of movement data the behavior provides.

After each behavior, the data collected is written to a file with the aforementioned naming convention, the velocity is doubled, and the trial duration halved; the velocity and the duration are at most increased and decreased respectively by a factor of 8, with the final velocity being .8 and the final duration being 1/8th of the original time (equivalent to ¼ second). To show progress and allow error-checking, each file name is printed to the console as it is created. This resulted in a total of four runs per trial, each with three behaviors, each with four .csv files (one for each of the four messages). With the four substances and four stirring materials, this amounted to 786 .csv files, each with thousands of data points (as large amounts of data are pushed each 1/40th of a second).

Technical Approach: Hardware

The hardware used for the experiments is the two fingered, Kinova robotic arm. To implement behaviors, motions are first tested manually on the arm using both the provided joystick and the arm's ability to be manipulated with human force. The arm's Cartesian velocity is echoed often to ensure that the correct, relevant joint movements are being implemented, and the message `geometry_msgs::TwistStamped` is used to publish the linear movement of the desired joint at the given velocity. To prevent the motions from overlapping and obfuscating one another, a pause behavior was written that simply published a velocity of 0 for all linear and angular velocities for the given duration, which is set to 3 seconds for all trials used in this paper. The pause behavior is implemented between all behaviors, and at the beginning of the program to differentiate the up-and-down motion from the arm's movement to the start position.

Figure 3: Photo of the joint positions used for stirring trials (L); Close-up of the finger positioning for gripping the stirring implements, shown with the wooden paint stirrer (R)

The fingers of the arm are of particular interest in this project, as they not only grasped the stirring implement for the entirety of each trial, but also provided more than half of the data in the form of position, torque, and force information. While there is no code to open or close the fingers, this is a source of much manual effort both before and after the program ran. Getting a firm grip on the stirring implement is vitally important, and a consistent grip is important for consistent data collection. Another consideration is the direction of the grasp - because the arm used only has two fingers, the implement could fairly easily be knocked loose by moving in a direction parallel to that of the grip. This is solved fairly easily in the back-and-forth behavior by moving in a direction perpendicular to the grip, but is an unavoidable issue in the circular behavior. This made it even more important to establish a good grip on the stirring implement, a procedure that is practiced often by the research team but is unfortunately not able to be automated for the data collection utilized for this project. An initial concern for this project is whether the Kinova arm is sensitive enough to detect the differences in the forces and positions used to stir different substances. For this reason, a disparate set of substances was chosen, which seems to have resulted in viable data. However, more testing is needed to more clearly determine the sensitivity of the arm, and whether or not it can distinguish the difference between more similar substances (e.g. the difference between juice and soda) with a reasonable degree of accuracy.

Materials

This project uses a preliminary set of substances able to be stirred by the robot, consisting of water, shampoo, dry Anasazi beans, and air. Air is chosen as the control, as it is easily accessible and provides limited resistance. The other materials were selected for their convenience and their frequency of use in everyday capacity, as well as their differences in viscosity. In the case of the beans, it is decided that since the robot's haptic sensitivity is relatively unknown, an extremely different substance should be used to maximize the possibility of viable data. The four stirring

Figure 4: The Tupperware container used to hold the test substances (L), including dry Anasazi beans (R)

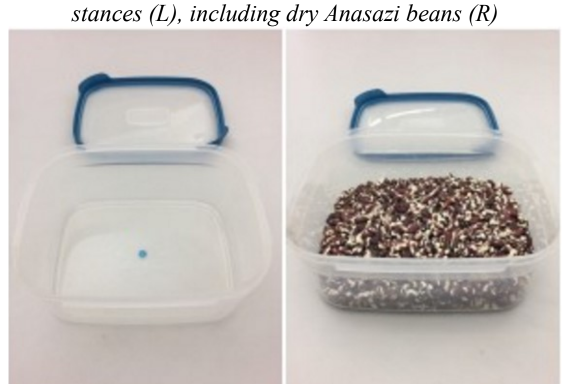

implements chosen were a wooden paint stick, a metal icing spatula with a silicon handle, a hard plastic spatula for mixers, and a soft plastic spatula with a hard plastic handle. These implements were selected due to their combination of availability, common usage in the real world, and their

Figure 5: The stirring implements used for the trials, from left to right: wooden paint stirrer, metal icing spreader, plastic spatula and soft plastic/silicone spatula

basic difference in material composition. In future trials, it would be useful to have an expanded set of both substances and stirring implements, as the end goal is to have the robot distinguish between very similar liquids, such as milk and water. It should be noted that a more comprehensive way of gathering data would also be useful here, as milk and water are typically distinguished by characteristics such as color, and such comprehension could be implemented via visual data on the robot.

Challenges

This project presented many challenges, many of which were overcome and some of which will need to be addressed in the future. Several proved unavoidable, such as the aforementioned gripping issues. The Kinova arm had a tendency to shake when moving very slowly or very quickly, which is consistent across trials but not necessarily within trials. This is especially concerning in this project, as position data is heavily used. The arm also had some issues moving to the programmed start position. When moved from the "home" position, the arm would choose one of a number of ways to move to the given start position, some of which triggered an obscure bug that prevented the behaviors from running at all. To solve this issue, the arm was moved close to the desired starting position before each run; this allowed the program to successfully move to the starting position without human intervention. The initial movement is not always consistent, however, with the arm often moved to a position near the programmed starting point, but stopping at an angle that prevented successful data collection. This is mostly remedied by calling the function responsible for moving the arm into the starting position twice in succession, although the exactness of the arm's starting position remains a confounding variable.

Another confounding variable is the position of the container, both in terms of exact placement on the table and in terms of height. Although tape was used to mark the container's position, the necessary act of moving the container to switch the substances being tested resulted in some shifting, which is somewhat problematic. The motions are tailored to fit in the container used, a process that is relatively simple but was ultimately deemed too time-consuming to reconfigure for every trial of every substance and stirring material for this dataset; as such, some movement naturally occurred. To compensate for the height differences in the stirring materials, the container of the substance being stirred was either lowered or raised. The height aids were kept consistent across trials of each implement (e.g. the number of books and the size of the books remained the same for the tool regardless of testing substance), but placement consistency across implements could not be achieved due to the nature of the program. Data from trials during which the stirring implement noticeable ran into the sides of the container for extended periods of time were thrown out from the final collection, but it is possible slight bumps or other such issues are not noticed by the researchers and may have affected the data.

Figure 6: Image of the height adjustments required for trials with shorter stirrers. Pictured is the shampoo run with the plastic spatula; two books were used to raise the container.

Results

This project is successful in terms of data collection, but unfortunately fell short in the goal of applying a machine learning algorithm to the data; the latter would have allowed the team to train the robot to recognize substances based on differences in viscosity. The large amount of data collected made analysis somewhat difficult and time-consuming, but it was ultimately determined that the robotic arm used is sensitive enough to detect differences in the haptic and positional manipulation required to stir substances of disparate viscosities.

Figure 7: The data points shown represent, from left to right (in columns of x, y, and z): force used when completing the circular motion with hard plastic in air, beans, shampoo, and water.

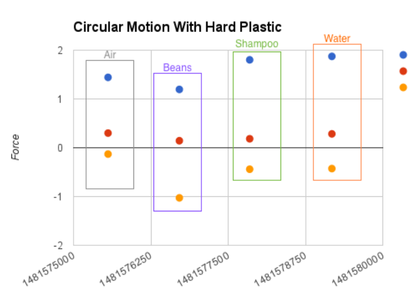

Brief analysis of the data collected proved that the robot would be able to successfully determine the difference between a solid and a liquid. See the appendix for an interpretation of some of the findings. Further testing is needed to prove whether or not the current setup is qualified to help the robot determine between liquids. This indicates that future work with this program on this arm is viable, and further data and learning can be pursued as is a goal of this project. The data collection program created can be adjusted to include new behaviors, or modify existing behaviors, with minimal difficulty. Additionally, the trials and behavior runs can be made to iterate any number of times via user input, and velocity and duration can be similarly easily changed to accommodate a wide range of substance containers. The velocity and duration of the behaviors can also be changed relatively simply (keeping the constraints of the arm, container, and stirring implement in mind), making mimicking common human stirring motions another way to continue research in this area. The dataset gathered is valuable, but more valuable is what it represents - a way to easily collect and store viscosity data for a large variety of substances, containers, and stirring implements.

Conclusion and Future Work

Our ultimate goal for this project is to teach the arm to correctly classify substances (either as liquid or solid or to distinguish between liquids) as a result of haptic sensation.

If successful, this project could be used as a springboard for research on robotic interaction with and comprehension of liquids, which could in turn be used to create robots that can work with liquids intelligently. This would include such activities as helping with cleaning (distinguishing between cleaning liquids as well as personal care items like shampoo and conditioner), cooking (determining the quality of ingredients and utilizing the viscosity of cake batter to gauge how much longer mixing needs to occur and at what point the batter is ready to be placed into the pan for baking), construction (testing the viscosity of time-sensitive materials like concrete and tar, specifically through stirring the mixture to prevent it from solidifying too early), and even automotive care (to optimize the viscosity of fluids important to car performance, such as freon, motor oil, and windshield washer fluid). The research also poses as an opportunity to advance the field of accessibility robots, specifically in the area of food preparation and aid. Ultimately, the advances made with this research can open doors in several fields, and will hopefully lead to changes that will improve the quality of life for millions.

As mentioned, the main goal for the continuation of this project is further data collection, with the eventual goal of implementing a machine learning algorithm that allows the robot to accurately distinguish between substances using the haptic data as a means to interpret viscosity. To this end, it would be useful to implement more user input to make updating variables, for instance velocity or duration of behaviors and starting position of the arm easier. The obstacles presented by the rigidity of these variables is one of the main challenges of the data collection for this project, and in order to meet future goals such challenges must be overcome. The data collection implemented in future iterations of this project should also feature greater breadth of substances tested, in order to more fully establish the usefulness of the data collection program and to make the robot's understanding more valuable. If a more representative dataset could be gathered, the robot could apply its knowledge of viscosity and substance manipulation to complete a variety of tasks as previously mentioned.

Additionally, another way to increase the accuracy of measurement as well as increase the number of trials completed would be to place the liquids in a closed container lined with a viscometer; the robot could then interact with the container through a set of hard-coded actions (e.g. push, shake, move in a circle, etc.), and based on the feedback recorded attempt to classify the liquids. The data collected from these trials could prove more comprehensive than attempting to convert the haptic data from the arm's sensors into viscosity interpretations. Additionally, to aid the robot in comprehending distinct albeit similarly viscous liquids, other forms of feedback could be employed (e.g. through using a Kinect camera as well as audio feedback), increasing the interaction between the robot and the substances, similar to the Elbrechter et al experiment [3]. If this could be successfully implemented, the arm's learning would more closely mirror the proprioceptive process infants use to learn to distinguish between substances, and future training could continue autonomously.

Acknowledgements

Thank you Jivko Sinapov for your guidance in defining the scope of this project as well as your willingness to answer questions during the research process.

References

1. Bamberger, Judith A and Margaret S. Greenwood "Non-invasive characterization of fluid foodstuffs based on ultrasonic measurements" Food Research International; Volume 37, Number 6, pp. 621-625, 2004.
2. Bamberger, Judith A and Margaret S. Greenwood "Measuring fluid and slurry density and solids concentration non-invasively," Ultrasonics; Volume 42, Number 1-9, pp. 563-567, 2004.
3. C. Elbrechter, J. Maycock, R. Haschke and H. Ritter, "Discriminating liquids using a robotic kitch-

en assistant," 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Hamburg, 2015, pp. 703-708.
4. Elert, Glenn. "Viscosity." The Physics Hypertextbook, 2016, physics.info/viscosity/. Accessed 1 Oct. 2016.
5. Griffith, Shane, et al. "Object Categorization in the Sink: Learning Behavior–Grounded Object Categories with Water." Object Categorization in the Sink: Learning Behavior–Grounded Object Categories with Water, 2012, www.vladsukhoy.com/papers/Griffith_Water_SPME2012.pdf. Accessed 27 Sept. 2016.
6. Reinhard Klapfer, Lars Kunze, Michael Beetz, "Pouring and Mixing Liquids --- Understanding the Physical Effects of Everyday Robot Manipulation Actions", In 35th German Conference on Artificial Intelligence (KI-2012), Workshop on Human Reasoning and Automated Deduction, Saarbrücken, Germany, 2012.
7. Schenck, Connor, and Dieter Fox. "Detection and Tracking of Liquids with Fully Convolutional Networks." arXiv.org, Cornell University Library, 20 June 2016, arxiv.org/pdf/1606.06266v1.pdf. Accessed 29 Sept. 2016.
8. Yoshitaka Hara, Fuhito Honda, Takashio Tsubouchi, and Aakhisa Ohya, "Detection of liquids in cups based on the refraction of light with a depth camera using triangulation." 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Tsukuba, 2014.

Appendix

An example of the full sequence of movements using the wooden paint stirrer, L to R, top to bottom: Beginning in starting position; the highest point of the up-down movement; coming back down from the up-down movement; completely down (returning to the starting position); beginning to move forward for the back-and-forth movement (changed for actual trials); furthest forward; beginning to move backwards; back in starting position; beginning to move in a counterclockwise direction; furthest to the left in the circular behavior (note that the fingers do not twist); past the furthest backwards point, beginning to move back towards the starting position; furthest point to the right; returning to the starting position