Socially Cognizant Industrial Waste Classification Robot (SC-IWCR)

Robert Contofalsky, Joe Doerr

May 13, 2025

1 Introduction

1.1 Landscape of recycling in the United States

Since 1970, the United States (US) has been enhancing its recycling capabilities to offset its issues with waste generation and landfill capacities [17]. In 2018, the US recycled nearly 25% of its municipal solid waste (MSW; 69 million tons), and the material that was recycled varied substantially by material type [36]. While much higher than the 1960 recycling rate of 10%, these numbers are far from ideal and have much room for improvement, as 50% of the 2018 MSW (146.1 million tons) still ended up in landfills. While almost a quarter of the landfill was food, much of the material was recyclable. Plastics, paper & paperboard, rubber, glass, metal and textiles made up just over half of the material that ended up in landfills, meaning that there is large room for improvement in this area [36].[1]

Efficient recycling sorting is imperative for the recycling economy to survive. In the US, recycling is sorted through material recovery facilities (MRFs), which are central for segregating waste from valuable recyclables [28]. Typically, MRFs receive recyclables from “single-stream collection in a mixed material stream that contains metals, plastics, glass, paper, and cardboard”. Much of the sorting is either automated or semi-automated, and workers remove problematic materials that could damage machinery, and handle the pre-sorting of incoming waste streams where human judgment still outperforms automated systems [28, 11, 18]. Additionally, contamination (non-recyclable products) is a major issue in single stream recyclables, as up 50% of loads can be contaminated and on average, 16% of loads are contaminated [31, 3]. If not removed from the stream, contaminated loads can damage machinery, which is estimated to cost at least $300 million a year for the US recycling system, or $140/ton of recycling [31, 29]. These expenses as a result of inefficient sorting drives up the cost of recycling, and shifts that burden to the consumers; making the economic benefits of recycling ultimately less appealing to governments and MRFs [28, 31]

Efficient recycling sorting is imperative for the recycling economy to survive. In the US, recycling is sorted through material recovery facilities (MRFs), which are central for segregating waste from valuable recyclables. Typically, MRFs receive recyclables from “single-stream collection in a mixed material stream that contains metals, plastics, glass, paper, and cardboard” [28]. Much of the sorting is either automated or semi-automated, and workers remove problematic materials that could damage machinery, and handle the pre-sorting of incoming waste streams where human judgment still outperforms automated systems [28, 11, 18]. Additionally, contamination (non-recyclable products) is a major issue in single stream recyclables, as up to 50% of loads can be contaminated and on average, 16% of loads are contaminated [31, 3]. If not removed from the stream, contaminated loads can damage machinery, which is estimated to cost at least $300 million a year for the US recycling system, or $140/ton of recycling [31, 29]. These expenses as a result of inefficient sorting increases the cost of recycling, and shifts that burden to the consumers; making the economic benefits of recycling ultimately less appealing to governments and MRFs [28, 31].

1.2 Impact on Workers

Employees at MRF facilities face difficult working conditions. For starters, they have higher injury rates of 4.4 incidents per 100 workers compared to 2.7 in all industries [7, 8, 34]. Second, they are exposed to dangerous chemicals by virtue of having to sort contaminated objects that increase their

---

[1] Though it is important to keep in mind that holistically analyzing these numbers is challenging due to the varying nature of local and state laws/public attitudes towards recycling in the US [17, 2].

risk for potentially untreatable illnesses. Third, the volatility of scrap market prices, changing recycling quality requirements, and low wages, lead to high turnover rates and recruitment difficulties for MRF employees and managers [28]. It is for these reasons that there is currently a push for sorting robots to enter these facilities, to decrease this manual burden.

1.3 Robots in MRFs

Robots provide a way to increase recycling rates without putting recycling literacy pressure on citizens, and reduce sorting errors that lead to contamination. As mentioned before, there is already a push for MRFs to rely less on human labour, with robot sorters being the best alternative [10]. Currently, there already exist robot systems that can pick between 2-10 times faster than humans. Many companies are increasingly interested in automating recycling sorting to save costs that result from the low retention rate of MRFs employees [10].

The current state of the art robot sorter are the Centralized and Fixed (C-Fixed) robots and Modular and Flexible (M-Flex) systems. C-Fixed solutions, can achieve up to 100 picks per minute but operate within safety cages isolated from workers, while M-Flex solutions enable greater worker collaboration and adaptability at 10-20 picks per minute. Companies such as Waste Management have already documented significant operational improvements with automation, including approximately 30% lower labor costs, 18% lower operational costs, and a 40% improvement in key safety metrics at fully automated facilities. Robotic sorting technologies are particularly promising for improving recovery of recyclables from residue streams, with case studies showing robots recovering nearly 600,000 additional recyclable items per month that would otherwise be landfilled, providing a return on investment within the first year of operation through additional commodity revenue and avoided disposal costs [10].

1.4 Our Approach

Building a low cost and highly efficient robotic recycling sorter is an important challenge in today's environment. To this end, this paper presents the Socially Cognizant Industrial Waste Classification Robot (SC-IWCR) for sorting recyclables in MRFs. This robot arm is positioned next to a conveyor belt of trash that needs to be sorted. From RGBD cameras the robot fuses together a 3D representation of the scene. Utilizing fine-tuned state of the art foundation models from computer vision [24][30], recyclable objects can be panoptically segmented with labels. Using the 3D segmented geometries, grasp poses can be extracted using grasp planning networks [6] that continually learn. Given these grasp poses, motion planning can generate the full robot motion to grasp the object, retrieve it, and place it in the proper bin. We provide code for the dataset generation component of grasp planning. We first introduce the technical approach then devise socially cognizant adjustments instructed by questions from [12]. The socially cognizant updates include robot supervision, a screen showing the robot's perception and planning for explainability, a programming interface for workers to ensure they maintain autonomy, and a washing station for the robot. We additionally provide a thorough analysis of the technical and socially cognizant designs. Finally, we refine the SC-IWCR by explicitly considering and answering the 10 questions for socially cognizant designs from [12].

2 Related Work

2.1 Grasping

Finding high quality grasp configurations in highly occluded and cluttered scenes is a difficult task for robots. Grasping tasks are trivial to human beings as we implicitly use shape completion, rich tactile information, and very capable end effectors. Humans can intuitively assess friction and slippage of a grasp, reactively apply forces, and estimate full object shapes, not requiring extensive training on how to do so. Conversely, the most common form of robot grasping is to find an SE(3) two-finger end effector gripper pose as a grasp target. This is used for its simplicity in grasp generation and end effector design. Uncertainty of the object geometry is a primary difficulty inhibiting this approach. Grasp planning has additional difficulties in inference time where sampled grasps need IK solutions and to be collision checked.

Given a full object geometry, grasps can be sampled and analytically evaluated using metrics such as force closure [23] and GWS [23]. Grasps can also be generated through simulating sampled grasps and keeping successes [6]. The issue with these methods is that it is not fast enough for online inference and requires a full object geometry. Due to these drawbacks, many approaches have been centered around preparing grasps offline with privileged object geometries, then training networks to recreate these grasps given partial observability of the object geometries [23][33][6][20]. This can be in the form of matching geometries to a library of known objects with associated grasps [26]. More commonly, point clouds with color are be directly processed with a learned network to output grasp candidates [23][33]. Grasping can also use a truncated signed distance field as input [6][20]. The ACRONYM [13] dataset contains about $8,000$ objects and about 17 million grasps which works such as [33] use in place of generating their own grasps. There also have been attempts to analytically find grasps directly on the partial geometries [26]. GPD receives a set of desired points to sample grasps around then uses a network to score the grasps [35].

One major drawback of learned methods is their reliance on the training input type. Models such as [33] having been trained on partial point clouds from one view means the model has inconsistent performance on a fused point cloud from multiple views. Additionally, these learned grasp planners do not have a robust way to filter grasps for a target object that is not sensitive to an imperfect mask. The use of deep learning leads to unexplainability for grasp failures. Furthermore, these methods fail to be adaptive, something that this work aims to remedy through lifelong learning on new data.

2.2 Distinguishing recyclables

In order to grasp and sort objects properly, they must have accurate semantic segmentation given a text description of the desired object. The state of the art method GroundedSAM [30] first uses an object detection method, Grounding DINO [24], to create bounding boxes on semantically segmented objects. Then they use the segment anything model (SAM) [22] to create pixel-level segmentations of the objects within the bounding boxes. Yolov8 [37] does instance segmentation but cannot segment objects conditioned on language. Florence2 [41] is a large vision language model (VLM) that can segment objects under text conditions but has lower precision than GroundedSAM and is more difficult to finetune given its non-modular structure. In this work we aim to finetune the object detection portion of GroundedSAM in order to properly distinguish between the different types of recyclable objects.

Computer vision is not the only way to distinguish between types of recyclable objects. Spectroscopic techniques such as near infrared [5], X-ray fluorescence [39], and color analysis [25] can be used to distinguish the type of recyclable plastic. These techniques are good for a set of materials but not for everything, making them less applicable for this application where we want to sort all recyclables in one system [25]. Near infrared spectroscopy alone can be noisy and inaccurate in practice, X-ray fluorescence detects PVC, and color analysis works for different colors of plastic [25]. It has been shown that with the help of deep learning, near infrared spectroscopy can classify plastics extremely accurately [9]. Spectroscopy is one of the industry standard methods for sorting plastics [25].

2.3 Garbage sorting

State of the art garbage sorting comes out of the company Pellenc ST which can distinguish recyclables on a conveyor belt using machine learning and spectroscopy. The objects are sorted into separate groups using well-timed air bursts. They specialize machines for objects of different types such as plastics versus metals which require unique spectroscopic techniques to sort. They are effective on sparsely occupied conveyor belts with few item types, trash that does not roll, and trash that is not overlapping. Their products have the most use through finely sorting already separated recyclables. The advantage of using a robot arm in place of a specialized machine is that the robot arm can do any given task. It is flexible to whatever setup the recycling factory already has for easier integration. Additionally, the robot arm can work with any amount and type of trash coming down the conveyor.

3 Method

3.1 Technical-only design

Our overall approach is to have a set of robotic arms that pick out different recyclable materials and place them into sorted bins from a starting and stopping conveyor belt of trash. The robot will have different chutes for different recyclable materials surrounding it that it may place trash into.

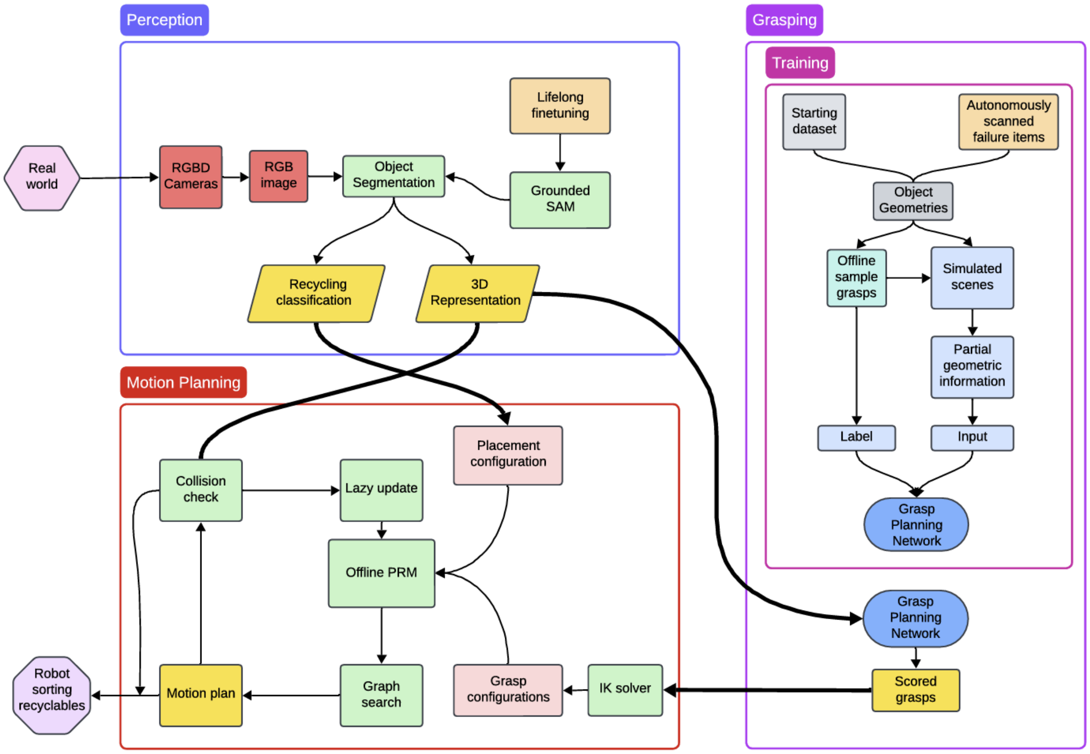

Figure 1: Technical Flowchart

3.1.1 Perception

With many cameras in stable positions, when the conveyor belt stops, the views can be segmented for each desired object and fused into a labeled truncated signed distance field (TSDF) which represents the 3D geometry of the scene. Each camera can be a RGBD camera that we can get a point cloud from. An alternative way to get point clouds is to use multiple cameras to get depth information where models such as Foundation Stereo [40] can help with accuracy. Given precisely measured camera placements, these point clouds can be fused together. To create 3D segmentation of the desired objects to pick, every RGB image can get a 2D panoptic segmentation mask that can be projected to the image's respective point cloud before fusion. This can be done with algorithms such as Grounded SAM 2 [30] which require a text prompt alongside an RGB image. Other options include Detic [42] and Florence2 [41].

3.1.2 Robot architecture

There are a variety of choices for robot architecture such as a simpler and cheaper 3 dof top down gripper similar to a 3d printer, or a more involved and capable $>$ 6 degree-of-freedom (dof) redundant robot arm to control an end effector in SE(3) space. We choose the ladder here for the increased flexibility to different objects and capability to pick up harder objects. For the end effector, options such as two finger grippers, suction, and robot hands can be used for grasping. We choose to use the two finger grippers for their simplicity, cost, and body of research around grasping with them. Suction does poorly when dealing with very unstructured objects and might suck up smaller particles and liquids. Robot hands are very expensive and while more capable, learning new grasps with robot hands requires more careful engineering. We will use the Franka Research 3 [19], 7-dof arm for this task.

3.1.3 Grasping

Given a segmented geometric representation of the scene, a grasping procedure must be completed to sort the recyclables from the rest of the garbage. A common way to grasp objects with a robot arm is to use a grasp planner. Grasps can be generated analytically or sampled then tested in simulation, but this process is lengthy and requires the full object model. Therefore, state of the art approaches gather high quality grasps for a large variety of objects offline, and train a model to recreate these grasps given partial geometric information [23][33][6]. Taking an existing grasp planner out of the box, such as point-net GPD [23] or VGN [6], would have a low success rate on arbitrary piles of garbage coming on a conveyor belt. Additionally, trash is constantly changing as different products and societal trends emerge, which would make the distribution of grasping targets constantly changing.

Therefore, we adopt an approach of lifelong learning to continually learn to grasp new objects that the robot is unable to generalize to initially. The robot can autonomously detect ungraspable objects or grasping failures and record the partial geometry of the given target object into a buffer. Similar to point-net GPD [23] and VGN [6], we can generate grasps offline and train on simulated partial point clouds of the object geometry. One issue with this approach is that in cluttered and unstructured trash, a wide variety of partial views of the object will exist due to occlusion which will commonly fall out of training distribution. Additionally, grasps generated on noisy partial views of an object may be very low quality. A potentially more effective human-in-the-loop approach is explained in section 3.2.3.

We will train the grasp planner initially on both a large labeled dataset Graspnet-1billion [14] and from our own labeled grasps to get the benefits for general grasping alongside data closer to our deployment domain. Our approach to grasp creation is similar to [6]. Given a set of object geometries, we can generate random scenes. We can sample surface points and for each surface point sample grasps with an orientation at some offset from the normal of the surface point. For evaluation, these grasps can be fully executed in simulation which gives richer information than analytical metrics. Data points can be stored as the input point cloud with all sampled grasps and their associated scores. Our object geometries will include a starting set of objects and our newly added geometries from continual learning. There are a variety of object geometry datasets available. We will choose EDAG [27] which includes a diverse range of object shapes unlike other datasets that contain mostly household objects. This will be important for picking trash as trash will commonly be deformed and thus in very unusual shapes. We will train a fully convolutional network with this dataset that outputs a set of grasps poses, grasp scores, and gripper closing widths.

3.1.4 Motion planning

In order to generate motion to grasp objects, our approach is to use motion planning to the grasp, closing the gripper, then motion planning to the object placement. For motion quality and speed, we can compute a dense probabilistic roadmap (PRM) [21] offline in an area around the stationary robot arm and do lazy collision evaluation [4] of created trajectories with the fused TSDF. This is especially effective as few configurations should be in collision as the top-down setup will cause the arm to mostly move in free space. For a given grasp, *n inverse* kinematic solutions can be solved as configurations to add to the PRM. Then with a graph search algorithm such as A*, trajectories can be found and collision-checked. Other motion planning options that one can potentially explore include a rapidly exploring random tree (RRT) and trajectory optimization.

3.1.5 Recyclable type detection

Recyclables of different materials need to be sorted properly into separate groups. As listed in [25], the main groups of plastics to recycle are PET, high-density polyethylene (HDPE), low density polyethylene, PVC, polypropylene, and polystyrene. The most valuable to recycle are PET and HDPE [25]. Semantic object detection models such as Grounding DINO [groundingdino] used in grounded SAM [30] is not trained on distinguishing between different types of recyclable materials which is essential for proper sorting. Different material types can look virtually identical, so additional background knowledge such as what logo on a piece of garbage corresponds to what materials will be essential. Due to the consistently changing stream of products within garbage over time, the semantic object detection model should be continually learning on new data. A labeled dataset of objects with their corresponding recyclable properties will need to be consistently updated. While finetuning, we can continue to sample datapoints from the datasets that Grounding DINO was originally trained on such as O365 [32].

Grounding DINO uses pre-trained networks for the text and image backbone. For finetuning, we will freeze the text and image backbone and train the feature enhancer and the cross-modality decoder. We will train in batches of 64 following the Grounding DINO paper [24] using datapoints of (image, text prompt, bounding box). To gather this dataset, we will use human annotators. Recycling sorting workers will already have the required knowledge to discriminate between different plastic types and can directly annotate images from the stream of garbage at their workplace. Concretely, there will be a screen with a given image of garbage on the conveyor belt with a set of discrete categories of recyclables to choose from. One click will be needed to choose the type of recyclable, and another click on the desired object will be used to create the bounding box. SAM [22] can be prompted with just a point, so we will extract the object segmentation from that point and create a bounding box around the object segmentation. This annotation can happen directly in the workplace with the same conveyor belt it will be used on, making the data very in-distribution.

3.1.6 Evaluation

We will evaluate the modules separately and as a whole. For grasping, we will use success rate and incorrect object grasp rate. We will evaluate this by counting every attempted grasp and if the segmented object was grasped, if no object was grasped, or if the wrong object was grasped. For motion planning the planning time, rate of failure, and collision avoidance will be evaluated. The planning time is essential for the efficiency of the robot. It should be very fast using a precomputed roadmap. If the robot consistently cannot find a plan to a desired grasp, the roadmap is likely too sparse or we are not giving the newly sampled point a large enough number of neighbors to attach to. These parameters of the motion planning can be tuned to fix this problem. We will also ensure that the collision checking is working properly and the robot avoids collision with its surroundings. Collisions could arise from bugs or faulty cameras. For object detection, there is not an easy way to automate evaluation. Since there are workers around, they can approximate the success rate of segmentation by watching the explanatory screen on the back of the robot as it works. As a whole, the robot will be evaluated on sorting accuracy and sorting speed. Sorting accuracy will be counted as what percentage of objects that the robot has picked are sorted correctly. Sorting speeding with evaluate the average time it takes the robot to pick an object starting at segmentation and ending at placing the object in a chute.

3.2 Socially Cognizant adjustments

For our design to be socially cognizant, we consider four overarching goals of socially cognizant robotics from [12]. These are to improve the quality of life of people, avoid unintended consequences, be adaptive to people and their needs, and maintain people's autonomy. These goals exist at the human-robot level and the human-society level [12]. We use these goals to adjust the original technical design in the following ways.

3.2.1 Quality of Life

From a robot-society perspective, our goal is to help the environment by recycling more material which is a general benefit for society. From a robot-human perspective, the workers that will be interacting with the robot daily should have the most positive experience possible. One frustration with robots is their difficult-to-understand failures. To mitigate this, this work aims to make this system explainable. One advantage of a modular system like ours is that you can clearly see which part is outputting what. This can be printed on a screen nearby the robot, showing outputs of the 3D geometry reconstruction, 3D segmentation, options for grasping, and animations of its upcoming trajectory to move to a given grasp target. While there still might be frustration when one of the component's is inexplicably not working, at least the workers can which part is failing and how it is failing. This can add additional layers of trust and understanding of the robot. Workers can build a mental model of how the robot thinks and behaves intuitively.

3.2.2 Unintended Consequences

From a robot-society perspective, a careless implementation of this robot might lead to wasted recyclable material. There could be segmentation or grasping issues with a certain type of object which could lead to a large waste of recyclable material. This can be mitigated with weekly monitoring of the robot performance by a human being and a few human sorters that viewing the downstream waste to see if the robot missed anything.

Another unintended consequence could be contamination of recyclables due to the robot's end effector getting various substances on it from the trash and continuing to sort recyclables without cleaning it off. To solve this, the robot can periodically have the conveyor stop so it can clean its end effector. The end effector can be dipped into nearby soap water then use a pre-trained motion to wipe itself off on a static dry cloth. This cleaning station can exist at a nearby position to the robot. The soap water and the dry cloth will need to be changed periodically by the workers.

From a human-robot perspective, working with a brittle robot that requires frequent maintenance can be painful. The RGBD cameras around the robot might get bumped by workers moving around the robot, messing up the camera extrinsics. Additionally, RGBD cameras get worn down over time and need to be replaced. When camera poses are perturbed, the extrinsics need to carefully be recalculated by a professional to avoid noisy fusions. This means that the workers would have a robot that would commonly be out of commission, adding stress and unpredictability to the job. To mitigate this, there can be a metal mount rigidly attached to the robot base with camera-holding slots that cameras tightly slot into. The mount and the robot will be rigidly connected, ensuring that the camera are always in the same position in the robot frame.

3.2.3 Adaptive

From a robot-society perspective, the robot should be adaptive to the changing types of recyclable materials and the changing look of recyclable materials. Eventually we will find ways to recycle things previously unrecyclable and new materials will be used that might be recyclable. It is important the the robot can adapt to these changing demands. The technical design already proposed continual learning, but the efficacy can be improved through human intervention. With a few humans double checking the robot's work, the items the robot missed can be used to create new full geometries for the robot to use to train with and semantic labels to train the segmentation model with. This is adaptive at the robot-human level as well because the workers can choose which items to scan and show so that the robot adapts to their needs. The full object geometry data will lead to much stronger improvements for grasping than the robot learning off of only partial geometries. If we had the full object mesh of objects that the robot fails to grasp, we can create many diverse and realistic trash piles offline to train on.

3.2.4 Autonomy and agency

For the people that work with the robot, the way they do their jobs should not be controlled by the robot. They should have the power to control the robot within reason to fit how they want their job to be completed. We make this happen by creating a simple programming interface for the robot where people can set discrete parameters concerning which objects to keep going for, how fast to move, the gripper squeeze strength, what grouping of recyclables to do, and how often to clean the end effector if at all. These parameters will ensure that the workers can see a better way to do the job for themselves and make it happen without being restrained by the robot. Additionally, the workers should be able to turn it on and off easily.

3.2.5 Cognitive Models

A key component of our waste classification robot is to have an accurate model of Theory of Mind (ToM) working in MRFs. TOM is defined as understanding the intentions, beliefs, and desires of other people. This can be understood as a *model* of human goals and desires [15]. This can be in the form of treating people as obstacles to avoid collisions with. We can use a conservative estimate of their collision geometries to check the motion plan against closed loop. There will be RGBD cameras that cover the reachable space of the robot arm that are consistently computing the geometry of the space around the robot. If collisions are expected with the current plan, this triggers a re-planning phase to the current grasp that the robot is moving towards or current bin to place the recyclable object. The ToM comes in with how we consider collisions of people. Through sequential frames of geometries, the direction and velocity of the person can be estimated and the person's collision model can be expanded larger in their areas.

3.2.6 Social Science Evaluation Metrics

To evaluate our socially cognizant solution, we propose a mixed-methods evaluation framework that integrates quantitative performance metrics with qualitative social science assessments:

1. **Worker Agency Assessment** The first metric that we would want to regularly evaluate is how do employees about their sense of agency at work. The primary reason for this is because knowing their sense of agency (and if it is increasing or decreasing) will be a key way to help employees

not feel alienated by their robot companions at work, but also help us as designers empirically validate how socially cognizant our robot is. If our robot significantly decreases employee sense of agency, then we know our robots' socially cognizant design is flawed. Measuring perceived agency would be done through the perceived autonomy scale developed by Ryan & Deci'[16]). This could be administered as a survey on a monthly basis that would anonymous and be later analyzed by social scientists that would be consultants.

2. **Technological Mediation Analysis** The second metric that we would want to measure are related to how our SC-IWRC transforms the relationship between workers and materials. This would be done through Verbeek's mediation theory [38]). This theory argues that technology doesn't simply facilitate work but *transforms* the work experience because of the new technology adds a layer of complexity in the workspace. The most ideal way to do this would be through an ethnographic study/analysis where researchers would observe how employees work alongside the SC-IWRC, however, this would be time-consuming and costly. A more flexible, and inexpensive approach would be giving monthly surveys that would ask the employees how they feel the robots have impacted their work, improved/worsened their sorting capabilities, etc. This would be more of a performative/empirical metric compared to the subjective metric offered by Deci's questionnaire.

3. **Public Policy Implications Assessment** The third metric is about evaluating the broader policy implications of implementing our SC-IWRC system. This assessment would track quantifiable improvements in recycling rates and reductions in contamination. We would compile this data bi-annually to provide policymakers and consultants with empirical evidence of how our socially cognizant design can advance sustainability goals without sacrificing worker well-being. Additionally, we would monitor regulatory compliance across different jurisdictions as robotic sorting becomes more prevalent, identifying potential regulatory gaps or needed adjustments. The results obtained from these evaluations would hopefully inform future iterations in public policy when developing future policy recommendations that balance technological advancement with worker protections, potentially establishing best practices for future implementations. Unlike the previous metrics which focus on individual experiences, this assessment considers system-level impacts that could inform municipal waste management policies and future automation regulations.

3.2.7 Evaluation Timeline

Our evaluation approach integrates technical and social assessment in three phases:

Phase 1: Baseline and Initial Deployment (Months 1-3) To start, we will establish pre-implementation baselines for all social science evaluation metrics at the start of the project. Initial worker interviews will gauge expectations and concerns before deployment. The SC-IWRC will then be deployed to collect early technical performance data.

Phase 2: Adaptive Implementation (Months 4-9) After Phase 1 is completed, worker feedback from Phase 1 will be analyzed and implemented to improve the system. Various social scientists will be consulted to analyze this data, as there will be a lot of it. During this phase we will conduct monthly technical evaluations alongside bi-monthly social impact assessments (worker agency and technological mediation). Interface and parameter options will be refined based on observed worker usage patterns. If possible, we will try and co-analyze this data with the MRF first quarterly report with our SC-IWRCs deployed.

Phase 3: Longitudinal Assessment (Months 10-24) For Phase 3, we will attempt to have a comprehensive evaluation across all technical and social dimensions. We will perform comparative analysis with non-socially cognizant implementations at other facilities across all metrics that we have collected in Phase 2, and present these findings to policymakers and MRF facility managers. We aim to be as transparent as possible with our metrics so as to increase collaboration among MRFs to deploy our SC-IWRCs.

This structured evaluation approach ensures our waste classification robot not only performs technically but creates positive social outcomes aligned with the principles of socially cognizant robotics. By systematically measuring both performance and social impact through established social science

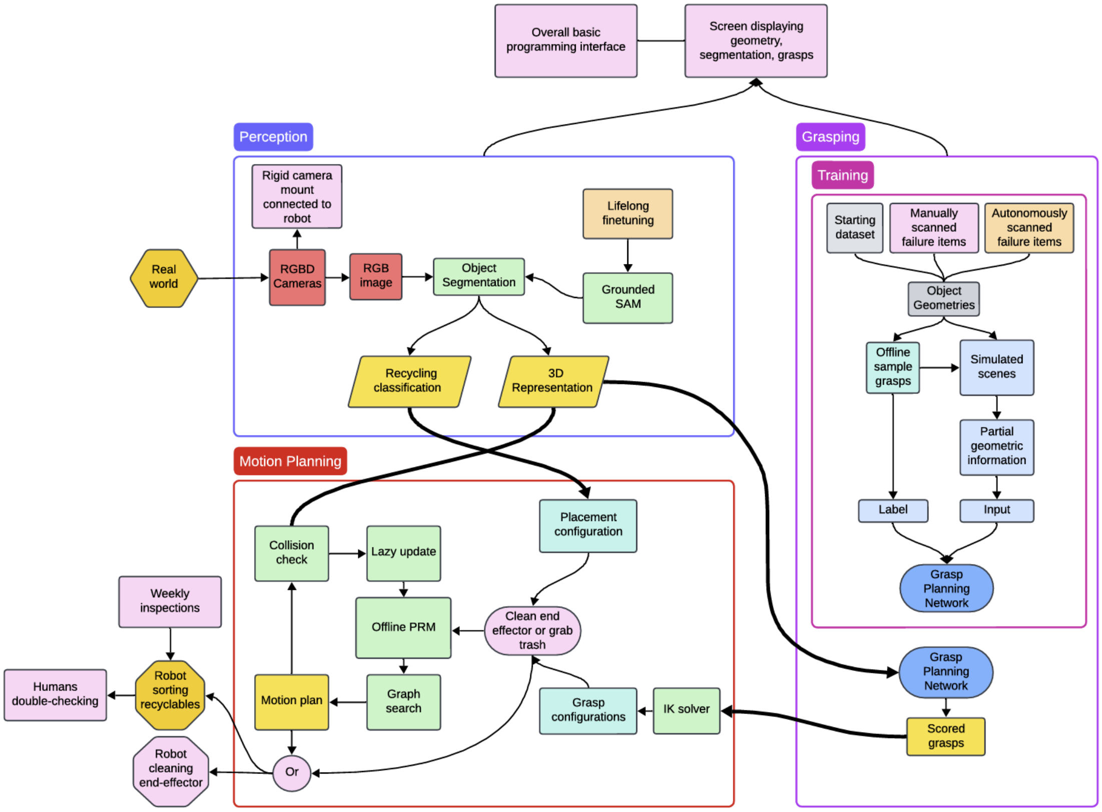

Figure 2: Socially Cognizant Flowchart. The socially cognizant adaptations are shown in light pink.

methodologies, we can demonstrate how technical innovation and social responsibility can be successfully integrated.

3.3 Toy coded version

One of the essential components of our method is grasp planning which necessitates grasp dataset generation. Since we are adopting a similar grasp dataset generation structure to VGN [6], we used their open source code as a base. They generate grasps in simulation rather than analytically, producing more robust grasps at a slower rate. This is done through random scene generation of piles or tightly packed groups of objects. An object surface point is randomly sampled and 6 grasps are generated at the position of the point with rotations that are offset from the normal of the given point. These grasps are fully executed in the scene, resetting the scene each time. The simulation is through pybullet.

To fit our SC-IWCR, the types of scenes to generate need to include recyclable objects and be in the style of conveyor belt piles that are seen in recycling plants. For data generation we added recyclable objects such as aluminum cans, plastic shampoo bottles, and condiment bottles. For the conveyor belt style scene generation, we added a version that randomly samples objects at different orientations and lets them fall freely onto a flat table surface, creating conveyor-belt-style piles. Additionally, our code includes an easy way to add new objects to the scene generation which only requires an obj and urdf file for each new object. Moreover, the garbage at a recycling plant will have a variety of substances on the objects in addition to the objects being composed of a variety of materials. This motivates finding grasps that are robust to unknown friction coefficients. For this we calculate the expected value of the success of a given grasp on a given scene over n trials of adding uniform noise to the friction coefficients for all objects in the scene. The expected value is kept as the label for the grasp. To speed up dataset generation time, we only calculate expected value for grasps that had initial success.

Our code is within this repository: https://github.com/JoeDoerr/ExpectedGrasps

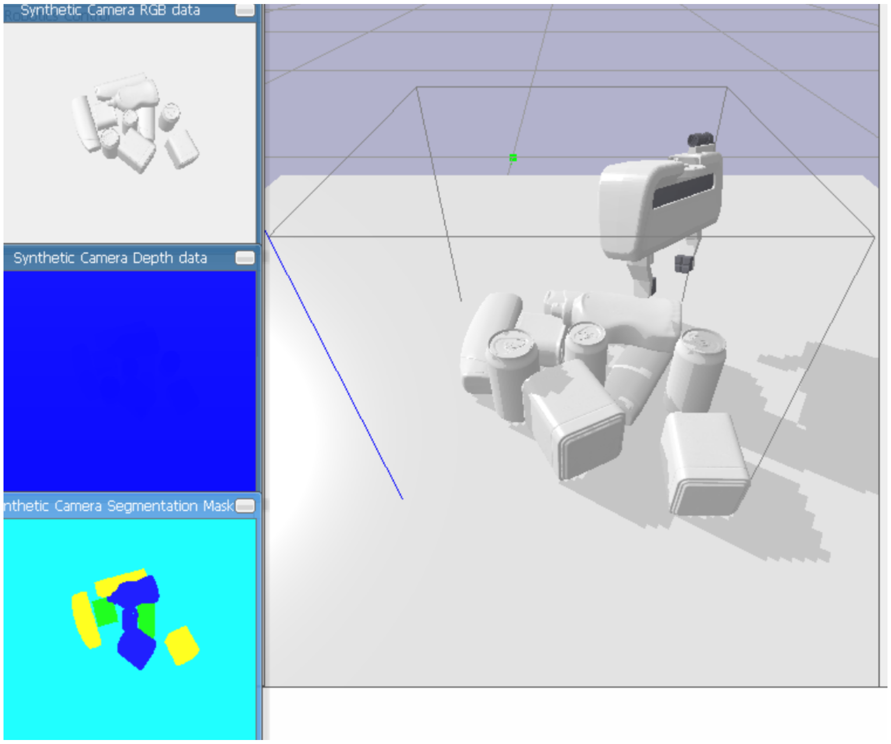

Figure 3: Toy example of the grasp dataset generation pipeline

4 Analysis

4.1 Technology-Only vs. Socially Cognizant Waste Classification Robot

In this section, we will provide a framework for anticipating unintended consequences of technology-only versus socially cognizant approaches to waste classification. The following subsections will discuss 1) Potential Problems and Unintended Consequences, 2) Insufficiency of the Technology-Only Approach, and 3) Measuring the Impact of these robots.

4.2 Comparative Analysis of Solutions

Aspect	**Technology-Only Solution**	**Socially Cognizant Solution**
Primary Goal	Maximize efficiency & Shareholder Value	Balance efficiency with human well-being
Control System	Fully automated without human involvement	Collaborative with adjustable parameters
Interface	Technical readouts only	Explanatory screen showing decisions
Adaptation	Autonomous learning	Worker-guided learning
Worker Role	Minimal maintenance	Collaborative programming and training

Table 1: Comparative Analysis of Technology-Only and Socially Cognizant Waste Classification Robot Solutions

4.3 Potential Problems and Unintended Consequences

Robot-Society Perspective:

1. **Job Displacement Without Transition:** Automated facilities experience approximately 30% lower labor costs (section 1.3), which translates to workforce reduction without creating alternative pathways for workers.

2. **Material-Specific Biases:** Without human oversight, the AI system would develop biases toward easily identifiable materials (section 5.4.1), systematically missing valuable recyclables, undermining environmental goals, and allowing contaminated recycling to ruin recycling batches.

3. **Economic Inequity:** The 18% operational cost savings would primarily would be directed to facility owners without including any savings distributions with the workers they replaced.

4. **Consumer Detachment:** By removing human involvement, the technology-only approach risks disconnecting consumers from waste consequences, reinforcing the perception that improper waste management techniques can be ”fixed” by having ”better” robots.

5. **Cross-Contamination:** Without self-cleaning protocols (section 5.4.3), the technology-only solution risks contaminating recyclables, potentially causing rejection of entire batches and undermining recycling goals. This would be an easy oversight in a factory that is focused on mainly robots working by themselves over working alongside humans.

Human-Robot Perspective:

1. **Diminished Worker Agency:** The technology-only approach relegates humans to reactive maintenance roles rather than leveraging their expertise and judgment and offers no upward mobility in this area beyond this role.

2. **Technological Distancing:** Following Prof. Hill's concept of ”*actio in distans,*” the technology-only solution creates a separation between human decisions and outcomes without establishing clear accountability structures.

3. **Adaptation Failures:** Without worker-guided learning mechanisms, the system would struggle to recognize new packaging materials, creating a growing gap between system capabilities and waste-stream realities.

4. **Maintenance Burden:** As noted in section 5.4.5, the technology-only solution ignores practical maintenance challenges around camera calibration and equipment reliability, creating stress and unpredictability for workers, on top of increased stress around robot costs.

4.4 Insufficiency of the Technology-Only Approach

The technology-only approach is fundamentally insufficient when measured against both theoretical frameworks:

1. **“Ought” vs. “Can”:** As Prof. Hill emphasizes, technological capability must be guided by ethical imperatives. The technology-only approach focuses on technical feasibility, and shareholder, value without addressing ethical dimensions. Technology first; ethics later.

2. **Procedural Rationality:** Following Prof. Andrews' distinction, the technology-only approach may achieve technical efficiency (substantive rationality) but lacks the “appropriate deliberation” required for procedural rationality [1].

3. **Facility Resilience:** A technology-only approach runs the real risk of over-reliance on automation; creating high vulnerability to system failures, where facilities would experience complete shutdowns rather than being able to shift to human-guided sorting.

4. **Early Learning Opportunities:** The technology-only solution lacks the ”learning by doing” feedback mechanisms that Prof. Andrews identifies as crucial for adapting quickly to unexpected outcomes through observation and testing.

4.5 Measuring Impact

Our measurement framework incorporates both substantive and procedural dimensions:

1. **Technical Performance Metrics:** Our framework will measure sorting accuracy across different material types and throughput relative to industry standards. We'll track contamination reduction and improvements in material recovery rates. The system's speed in adapting to novel materials will also be evaluated.

2. **Worker Impact Metrics:** We'll monitor safety improvements from the baseline 4.4 incidents per 100 workers. Job quality metrics will include autonomy, skill utilization, and overall satisfaction. Worker retention rates and technical skill development will indicate long-term workforce benefits.
3. **Economic and Societal Metrics:** The broader impact will be measured through benefit distribution and community employment effects. We'll quantify environmental benefits such as reduced landfill usage. The overall viability and sustainability of the recycling economy will be assessed by measuring how much more profitable the MRF become.

By implementing the iterative schedule outlined in section 5.6, our socially cognizant approach integrates Prof. Andrews' methods (analogizing, interpolation, projection) with procedural approaches (reflection, reasoning, discourse) to enable continuous learning and adaptation. This aims to transform what would be unanticipated consequences into anticipated ones that can be effectively managed, while preserving human dignity and agency in the recycling process.

5 Socially Cognizant Robot Design: Best Practices

5.1 Does The Design Improve Quality of Life

Our waste classification robot directly improves quality of life for multiple stakeholders. First, our robot will alleviate household recycling literacy, as their increased sorting efficiency will shift the burden of inconsistent recycling practices to the MRF facility robots as opposed to the average American. Second, our design will reduce the health burden of MRF workers. By automating recycling sorting, MRF employee exposure to hazardous working conditions and potentially reducing their high injury rate closer to the national average or below.

Our socially cognizant robot design will have downstream effects of improving the environment and recycling economy. Increased recycling rates will result in less landfill usage, and less damage to recycling equipment due to decreasing contamination rates. These benefits will extend to the entire community, and not just MRF facilities, and their workers.

5.2 Does The Technolgy Address a Critical Societal Need?

The technology addresses the critical societal need for improved recycling efficiency and sustainability. With only 32.1% of municipal solid waste being recycled or composted in the U.S., and plastics recycling rates at a troubling 8.7%, there is an urgent need for solutions that can improve these metrics. The high contamination rates (15-25% on average) further underscore the necessity of better sorting technologies. Our technology responds directly to the economic challenges facing recycling facilities, where high processing costs often make recycling financially unviable. By reducing contamination rates and increasing the value of recovered materials, our robot helps make recycling economically sustainable—a prerequisite for addressing the environmental need for waste diversion.

5.3 Have discussions occurred with a diverse set of stakeholders prior to deployment?

Our socially cognizant design process includes extensive stakeholder consultation beginning in the planning phase. We specifically identify MRF workers, facility managers, recycling industry experts (e.g., Waste Management), community representatives (local governments), and potential end users of recycled materials (manufacturing companies using recycled content) as key stakeholders. The worker-guided learning approach in our system ensures that recycling sorting workers who have specialized knowledge about material types are integrated into the annotation process for training our model(s). This not only improves technical performance but gives workers direct input into how the robot identifies and categorizes materials. Additionally, our design includes mechanisms for community members to identify systematic errors in sorting, creating multiple channels for stakeholder feedback both before and after deployment.

5.4 What are the potential unintended consequences?

We've identified several potential unintended consequences and built mitigation strategies into our design:

From a robot-society perspective:

1. **Material-Specific Biases:** The learning system might develop biases toward easily identifiable materials at the expense of others. One good example of this is with the varying types of plastic that are incredibly difficult for a robot to identify with high accuracy. Our socially cognizant design addresses this through worker-guided training and continuous monitoring of material recovery rates.
2. **Wasted Recyclable Material:** A careless implementation of this robot might lead to wasted recyclable material. There could be segmentation or grasping issues with certain types of objects, which could lead to substantial waste of recyclable material, and ultimately not saving the MRF any money in the process. Making recycling more expensive could yield serious blowback in the implementation of this robot on a grander scale.
3. **Cross-Contamination:** Another unintended consequence could be contamination of recyclables due to the robot's end effector getting various substances on it from the trash and continuing to sort recyclables without cleaning it off. To solve this, the robot can periodically have the conveyor stop so it can clean its end effector. The end effector can be dipped into nearby soap water then use a pre-trained motion to wipe itself off on a static dry cloth. This cleaning station can exist at a nearby position to the robot. The soap water and the dry cloth will need to be changed periodically by the workers.

From a human-robot perspective:

4. **Job Displacement:** Perhaps the most significant potential consequence is worker displacement. While our socially cognizant design deliberately creates complementary roles rather than replacement roles, this does not mean that facilities that employ these robots would not seek to replace human workers in totality. Indeed, the economic opportunities that these robots could yield could seriously increase the demand for them at a larger scale that could wipe out the need for human workers, ultimately causing societal stress that cannot be repaired later down the line. This was placed in "human-robot perspective" as this immediate impact would be troubling for the workers that got replaced and only as a collective impact would this be deemed "robot-society perspective". However, this process of job displacement would start more locally than broadly.
5. **Maintenance Burden:** Working with a robot that requires frequent maintenance can be problematic. The RGBD cameras around the robot might get moved by workers moving around the robot, messing up the camera calibration which can require an expert to fix. Additionally, RGBD cameras worsen in quality over time and need to be replaced. To avoid the robot from being commonly out of commission, there can be a mount that is rigidly attached to the robot base with camera-holding slots. This ensures that the cameras are always in the same position in the robot frame regardless of being bumped or needing to put in new cameras.

5.5 Are there safety issues in the application domain and have the appropriate permissions and inspections been completed?

Our design addresses several safety concerns inherent to MRF environments:

1. **Worker Safety:** By physically separating the robotic arm from worker areas and incorporating cameras for area monitoring, we reduce collision risks.
2. **Hazardous Material Handling:** The robot can identify and properly sort hazardous materials that would otherwise pose health risks to human workers.
3. **Mechanical Safety:** Our implementation includes emergency stop mechanisms accessible to workers and automated fault detection systems.

Prior to deployment, our design would require inspection and certification under OSHA regulations for industrial robots, as well as compliance with specific waste handling regulations in the deployment jurisdiction. As such, we would work in collaboration with several law firms to ensure that we exceed the legal standards required for robotic deployment in MRFs.

5.6 What is the schedule for iteration?

Our design incorporates a deliberate iteration schedule with three phases:

1. **Initial Deployment and Baseline Assessment (Months 1-3):** Establish performance baselines for sorting accuracy, contamination rates, and worker feedback.
2. **First Iteration Cycle (Months 4-6):** Implement adjustments based on initial feedback, focusing on technical improvements and worker interface refinements.
3. **Second Iteration Cycle (Months 7-12):** Address more complex issues identified during extended operation, including adapting to seasonal changes in waste stream composition.

Each iteration includes collection of quantitative metrics (sorting accuracy, throughput, contamination rates, etc) and qualitative feedback from workers and facility managers (trust, safety, comfort, usability, feedback, etc). The schedule deliberately balances the tension between rapid deployment and thorough evaluation.

5.7 How will the technology adapt to human desires and needs of people once deployed?

Our waste classification robot incorporates multiple adaptation mechanisms:

1. **Continual Learning System:** The robot continuously updates its model to recognize new packaging types and materials as they enter the waste stream.
2. **Customizable Parameters:** Workers can adjust sorting priorities, gripper strength, and cleaning frequency through a simple programming interface, giving them control over how the robot operates.
3. **Worker-Guided Learning:** The annotation system allows recycling sorting workers to teach the robot about new or challenging materials by clicking on images from the waste stream.
4. **Performance Monitoring and Feedback Loop:** The explanatory screen shows the robot's decision process, allowing workers to identify and correct errors in real-time.

These features ensure that the robot can adapt to both the evolving needs of workers and changes in the recycling stream over time.

5.8 Is there a mechanism for users to provide feedback and is there a plan to act on that feedback?

Our design incorporates several feedback mechanisms:

1. **Real-Time Interface:** The explanatory screen showing the robot's perception, grasping options, and planned movements allows workers to identify issues immediately. This interface is also user friendly for anyone to use and not just experts, effectively democratizing the robot in the facility.
2. **Annotation System:** Workers can directly label new materials or correct misclassifications through the touchscreen interface.
3. **Parameter Adjustment:** The programming interface allows workers to adjust robot behavior based on their experience and preferences.
4. **Scheduled Review Sessions:** Weekly performance reviews examine missed recyclables and systematic errors, with explicit processes for incorporating feedback into software updates.

The feedback cycle is completed through scheduled system updates that implement changes based on worker input, creating a closed loop between observation, feedback, and adaptation.

5.9 What are the ethical considerations?

Our waste classification robot design addresses several key ethical considerations, following Professor Hill's framework of moving beyond "what can we do?" to the essential question of "what ought we to do?":

1. **Responsibility in Technological Distancing:** We counter Professor Hill's concept of "actio in distans" (action from a distance) by maintaining clear accountability through transparent decision processes, audit trails, and meaningful human oversight. Maintaining proximity between workers, managers, and robots.
2. **Human-Centered Values:** Our design aligns with core human values rather than merely maximizing efficiency, preserving the dignity of labor by creating complementary roles that respect worker expertise and autonomy. No where in our design do we prioritize profit maximization over human agency or other similar metrics.
3. **Congruence with Human Capabilities:** Following Prof. Hill's framework, our technology remains congruent with human reasoning and nature through explainable decisions, adaptation to human expertise, and enhancement rather than replacement of human capabilities. Clarity is our number one focus.
4. **Balanced Innovation and Caution:** Addressing the "ought" question requires careful consideration of potential harms, implemented through regular monitoring, contamination prevention, and human verification systems. Distinguishing "oughts" from "is" is a difficult challenge that we actively tackle when designing our robot.
5. **Ethical Evaluation Framework:** We ensure ongoing ethical alignment by evaluating worker impact, environmental benefits, economic fairness, autonomy preservation, and clarity of responsibility once per quarter.

This ethical framework ensures our technology embodies Professor Hill's principle that technological development must be guided by ethical imperatives rather than merely technical possibilities.

5.10 Does the design respect human agency and autonomy?

Our socially cognizant design deliberately preserves and enhances human agency through:

1. **Worker Control Interface:** The programming interface gives workers direct control over robot parameters, allowing them to adjust operations based on their expertise and preferences.
2. **Override Capability:** Workers can manually override sorting decisions when they observe errors, maintaining human judgment as the final authority.
3. **Decision Transparency:** The explanatory screen makes the robot's perception and decision-making visible, allowing workers to understand and predict system behavior.
4. **Complementary Task Design:** Rather than replacing human judgment, our system complements it by handling repetitive, dangerous tasks while preserving human control over complex decisions.
5. **Skill Development:** The design creates opportunities for workers to develop new technical skills in robot supervision and maintenance, expanding rather than diminishing their agency within the workplace.

By addressing these ten questions throughout our design process, we've created a waste classification robot that exemplifies socially cognizant robotics principles. Our system balances technological capability with human needs, enhances worker dignity while improving safety, and contributes to environmental sustainability while maintaining economic viability. Most importantly, it places human agency at the center of the design, ensuring that technology serves human needs rather than subordinating them to technical imperatives.

6 Conclusion

Our Socially Cognizant Industrial Waste Classification Robot (SC-IWCR) represents a significant advancement in addressing the critical challenges facing material recovery facilities today. It uses state of the art computer vision, robot grasping, and robot motion planning algorithms to create a robust trash sorting solution. By balancing technical capabilities with human-centered design principles, we've created a system that not only improves recycling efficiency but does so while enhancing worker experiences and addressing potential unintended consequences.

When scaled to numerous implementations, we expect our socially cognizant approach to provide several long-term impacts:

1. **Environmental Benefits:** Increased sorting efficiency translates into increased recycling rates, decreased landfill usage, decreased contamination, leading to more efficient resource utilization.
2. **Worker Benefits:** Rather than displacing workers, our system will transform recycling jobs into higher-skilled positions focused on robot supervision, maintenance, and continuous improvement, in addition to decreasing MRF employee injury rates. The efficiency of the SC-IWRC can cause MRFs to be more profitable and thus create more MRFs, introducing more jobs. Moreover, workers will maintain their autonomy rather than working around the robot.
3. **Adaptive Resilience:** Our continual learning framework, guided by human expertise, will allow recycling systems to evolve with changing materials and societal needs.

By embedding socially cognizant considerations at the core of our robotic design, the SC-IWCR provides an opportunity to benefit the world on the scale of the environment, the society, and the individuals that work with the robot.

References

[1] Clinton J Andrews. Better anticipating unintended consequences. *IEEE Transactions on Technology and Society*, 2024. Accessed: May 12, 2025.

[2] Anonymous. *Quantifying Barriers to Recycling Plastic Packaging in the United States.* PhD thesis, ProQuest, 2025. Accessed: May 10, 2025.

[3] Catalina Blanco, Caitlin Spanbauer, and Stephanie Stienecker. America's broken recycling system. *California Management Review*, May 2023. Accessed: May 10, 2025.

[4] Robert Bohlin and Lydia E Kavraki. Path planning using lazy prm. In *Proceedings 2000 ICRA. Millennium conference. IEEE international conference on robotics and automation. Symposia proceedings (Cat. No. 00CH37065)*, volume 1, pages 521–528. IEEE, 2000.

[5] Giuseppe Bonifazi and Silvia Serranti. Quality control by hyperspectral imaging (hsi) in solid waste recycling: logics, algorithms and procedures. In *Image Processing: Machine Vision Applications VII*, volume 9024, pages 189–203. SPIE, 2014.

[6] Michel Breyer, Jen Jen Chung, Lionel Ott, Roland Siegwart, and Juan Nieto. Volumetric grasping network: Real-time 6 dof grasp detection in clutter. In *Conference on Robot Learning*, pages 1602–1611. PMLR, 2021.

[7] Bureau of Labor Statistics. Table 1. incidence rates of nonfatal occupational injuries and illnesses by industry and case types, 2022. Website, 2022. Accessed: May 12, 2025.

[8] Bureau of Labor Statistics. Employer-reported workplace injuries and illnesses, 2023 - 2023 a01 results. Website, 2023. Accessed: May 12, 2025.

[9] Berny Carrera, Victor Luid Piñol, Judit Bazin Mata, and Kwanho Kim. A machine learning based classification models for plastic recycling using different wavelength range spectrums. *Journal of Cleaner Production*, 374:133883, 2022.

[10] Marian Chertow, Barbara K Reck, Amy Wrzesniewski, and Berk Calli. Outlook on the future role of robots and ai in material recovery facilities: Implications for u.s. recycling and the workforce. *Journal of Cleaner Production*, 470:143234, Sep 2024.

[11] Ciprian Cimpan, Anja Maul, Michael Jansen, Thomas Pretz, and Henrik Wenzel. Central sorting and recovery of msw recyclable materials: A review of technological state-of-the-art, cases, practice and implications for materials recycling. *Journal of Environmental Management*, 156:181–199, Jun 2015.

[12] Kristin J Dana, Clinton Andrews, Kostas Bekris, Jacob Feldman, Matthew Stone, Pernille Hemmer, Aaron Mazzeo, Hal Salzman, and Jingang Yi. Socially cognizant robotics for a technology enhanced society. *arXiv preprint arXiv:2310.18303*, 2023.

[13] Clemens Eppner, Arsalan Mousavian, and Dieter Fox. Acronym: A large-scale grasp dataset based on simulation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6222–6227. IEEE, 2021.

[14] Hao-Shu Fang, Chenxi Wang, Minghao Gou, and Cewu Lu. Graspnet-1billion: A large-scale benchmark for general object grasping. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11444–11453, 2020.

[15] Jacob Feldman. Lecture 3: Theory of mind. Lecture, February 2025. Course: Socially Cognizant Robotics.

[16] Marylène Gagné and Edward L Deci. Self-determination theory and work motivation. *Journal of Organizational Behavior*, 26(4):331–362, 2005.

[17] Krista L Greene and David J Tonjes. Quantitative assessments of municipal waste management systems: Using different indicators to compare and rank programs in new york state. *Waste Management*, 34(4):825–836, 2014.

[18] Sathish Paulraj Gundupalli, Subrata Hait, and Atul Thakur. A review on automated sorting of source-separated municipal solid waste for recycling. *Waste Management*, 60:56–74, Feb 2017.

[19] Sami Haddadin. The franka emika robot: A standard platform in robotics research. *IEEE Robotics & Automation Magazine*, 2024.

[20] Zhenyu Jiang, Yifeng Zhu, Maxwell Svetlik, Kuan Fang, and Yuke Zhu. Synergies between affordance and geometry: 6-dof grasp detection via implicit representations. *arXiv preprint arXiv:2104.01542*, 2021.

[21] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE transactions on Robotics and Automation*, 12(4):566–580, 1996.

[22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4015–4026, 2023.

[23] Hongzhuo Liang, Xiaojian Ma, Shuang Li, Michael Görner, Song Tang, Bin Fang, Fuchun Sun, and Jianwei Zhang. Pointnetgpd: Detecting grasp configurations from point sets. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3629–3635. IEEE, 2019.

[24] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, pages 38–55. Springer, 2024.

[25] Cesar Lubongo, Mohammed AA Bin Daej, and Paschalis Alexandridis. Recent developments in technology for sorting plastic for recycling: The emergence of artificial intelligence and the rise of the robots. *Recycling*, 9(4):59, 2024.

[26] Jeffrey Mahler, Florian T Pokorny, Brian Hou, Melrose Roderick, Michael Laskey, Mathieu Aubry, Kai Kohlhoff, Torsten Kröger, James Kuffner, and Ken Goldberg. Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 1957–1964. IEEE, 2016.

[27] Douglas Morrison, Peter Corke, and Jürgen Leitner. Egad! an evolved grasping analysis dataset for diversity and reproducibility in robotic manipulation. *IEEE Robotics and Automation Letters*, 5(3):4368–4375, 2020.

[28] S M M Rahman and B K Reck. Techno-economic analysis of a material recovery facility employing robotic sorting technology. In N Nasr, editor, *Technology Innovation for the Circular Economy*, pages 269–278. Wiley, 1st edition, 2024.

[29] Recycling Partnership. 2019 west coast contamination initiative research report. Technical report, Recycling Partnership, May 2020.

[30] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024.

[31] Susan Robinson. The changing waste stream. Technical report, EPA: Houston, TX, USA, 2014.

[32] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8430–8439, 2019.

[33] Martin Sundermeyer, Arsalan Mousavian, Rudolph Triebel, and Dieter Fox. Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13438–13444. IEEE, 2021.

[34] Talent.com. Recycling sorter salary in usa - average salary. Website, 2025. Accessed: May 12, 2025.

[35] Andreas Ten Pas, Marcus Gualtieri, Kate Saenko, and Robert Platt. Grasp pose detection in point clouds. *The International Journal of Robotics Research*, 36(13-14):1455–1473, 2017.

[36] US EPA. National overview: Facts and figures on materials, wastes and recycling. Website, 2025. Accessed: May 10, 2025.

[37] Rejin Varghese and M Sambath. Yolov8: A novel object detection algorithm with enhanced performance and robustness. In *2024 International Conference on Advances in Data Engineering and Intelligent Computing Systems (ADICS)*, pages 1–6. IEEE, 2024.

[38] Peter-Paul Verbeek. Toward a theory of technological mediation: A program for postphenomenological research. In *Technoscience and postphenomenology: The Manhattan papers*, page 189. Lexington Books, 2016. Accessed: May 12, 2025.

[39] Bryan D Vogt, Kristoffer K Stokes, and Sanat K Kumar. Why is recycling of postconsumer plastics so challenging? *ACS Applied Polymer Materials*, 3(9):4325–4346, 2021.

[40] Bowen Wen, Matthew Trepte, Joseph Aribido, Jan Kautz, Orazio Gallo, and Stan Birchfield. Foundationstereo: Zero-shot stereo matching. *arXiv preprint arXiv:2501.09898*, 2025.

[41] Bin Xiao, Haiping Wu, Weijian Xu, Xiyang Dai, Houdong Hu, Yumao Lu, Michael Zeng, Ce Liu, and Lu Yuan. Florence-2: Advancing a unified representation for a variety of vision tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4818–4829, 2024.

[42] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *European conference on computer vision*, pages 350–368. Springer, 2022.