Magic: The Sortening

Henry Powell, EE, Malcolm Okaya, CE, Joe Maloyan, CE, Zalman Lipschitz, ME, and Liam Rees, ME

***Abstract*— The secondhand market for Magic cards is very large, but dealing with cards in bulk requires significant time, manual effort, and money to sort and check for damage. Current automatic sorting solutions, which are themselves quite limited, can't check for damage properly. Our proposed solution is able to take several stacks of cards as an input, detect and grade each card's condition, identify the cards, check for forgeries, and sort them into several configurable bins. This allows users to properly value and sort large quantities of cards. We did not have time to complete all parts of the project, but we were able to create a proof-of-concept that demonstrated the potential of such a device.**

I. Introduction

There is a significant secondhand market for collectible trading cards in general and *Magic: The Gathering* cards in particular. The value of any given card is a function of its condition and its rarity. This project aims to create a solution for sorting and grading bulk cards.

A. Significance

The secondhand *Magic: The Gathering* (MTG) market is estimated to be worth over $10 billion [1], and the market for trading cards in general is estimated to be worth $67 billion by 2027 [2]

B. Context and Competing Solutions in Marketplace

There are a number of card sorting machines currently on the market, and likewise many grading and verification services. The *Roca Sorter* by TCGplayer is one such sorting machine, it features configurable sorting bins for up to 1000 cards and automatic digital catalog but no grading or verification features [3]. Currently, if one wants their card verified and graded, they will have to send it to a professional grading service. Even the most basic grading can cost $20 per card and take over a year to turn around; a grading with a more reasonable turnaround, and with a breakdown over different parts of a card can cost upwards of $200 each [4]. Our solution pairs the bulk sorting capabilities of machines like the Roca, with the value-add of the grade and validation offered by grading services.

C. Societal Impacts

The prototype machine is aimed at secondhand card sellers such as brick-and-mortar stores and online distributors. The sorting functionality reduces the labor costs and time required for simply sorting and cataloging cards. The grading functionality reduces or eliminates the significant costs associated with generating accurate grades (and thus values) for cards through 3rd party services. The forgery detection functionality helps protect sellers and their customers from fraud.

D. System Requirements and Specifications

The specifications of the system requirements shown below in table 1 reflect the emphasized desires of a potential client. They may be summarized as: a high quality, high accuracy, and multifaceted analysis of each card's condition and veracity with minimal potential wear or damage to cards as they are handled. Desired speed and capability is based on the approximate speed and capability of a human grader, and as many output bins as are feasible is desirable.

Item	Description & Spec
Card Throughput	>2 cards/min, >1000 card capacity
Condition Grading	Catch damage visible to the human eye. Provide a valid grading based on that damage based on client guidelines.
Card ID	Solution must find the name of the card as well as the set it comes from. >95% accuracy
Forgery Detection	Perform weight check, glossiness check, printing verification check. 90% of forgeries fail at least one.
Card Damaging	No noticeable damage to cards.
Output Bins	At least 15 software-configurable output bins

Table 1: Requirements and Specifications

II. Design

A. Overview

Our solution is formed around a 3-axis gantry with a pick-and-place system. Cards are taken from an input bin, brought to each grading and verification step in turn, then placed in their appropriate user-configured output bin. See section IIB for discussion at-length of the mechanical system. Originally, the prototype mechanical design utilized rollers and a number of chutes in series, though this route was abandoned due to customer feedback and concerns regarding damaging cards and system reliability; See appendix A for discussion of design alternatives at length.

The condition grading, card identification, and forgery detection are each performed over several separate stages, one for each side of the card and also for the sides. See section IIC for discussion at-length of the damage detection system, utilizing a technique known as microscopy to highlight damage. See section IID for discussion at-length of the forgery-detection system, using 4 different methods to determine the authenticity of each card. See section IIE for discussion at-length of the card identification system utilizing both feature matching and perceptual hashing. The system block diagrams are provided below.

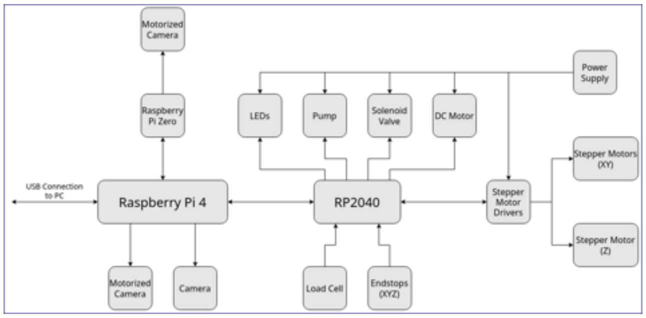

Hardware Block Diagram

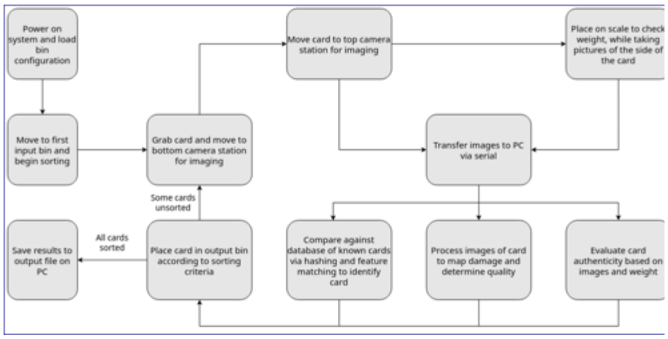

Software Block Diagram

B. Mechanical Pick & Place System

The mechanical pick-and-place system is build off of a 3-Axis gantry system that we built from scratch. The gantry system is modeled after 3D printer and CNC-style grantries. A vacuum pump and solenoid is used to generate a vacuum, which is used with a suction cup to pick up and release cards. The use of a suction cup took influence from the Roca sorter [3], a similar design, though not as ambitious (as it did not attempt grading or fraud-detection). 3 Stepper motors drive the gantry in its 3 axes; roller wheels sized to extruded aluminum sections are what allow for smooth and precise motion. Small endstop switches are used for homing and ensuring an accurate idea of the position in the XY direction, and they are used every few cards to ensure that the position does not slip too much. In the Z direction of motion, which is done independently by means of a vertically moving shaft fixed to the head of the gantry, is monitored by both a stepper motor and an IR distance sensor to ensure that the cards are picked up precisely regardless of how full any in/output bin might be.

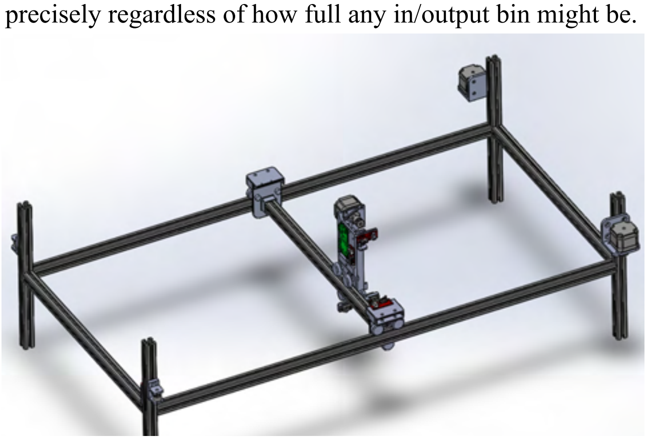

CAD Model of the gantry

C. Damage Detection, Cropping, and Rotating

Damage detection is performed in the topside evaluator, bottom side evaluator, and edge evaluator utilizing a technique known as microscopy.

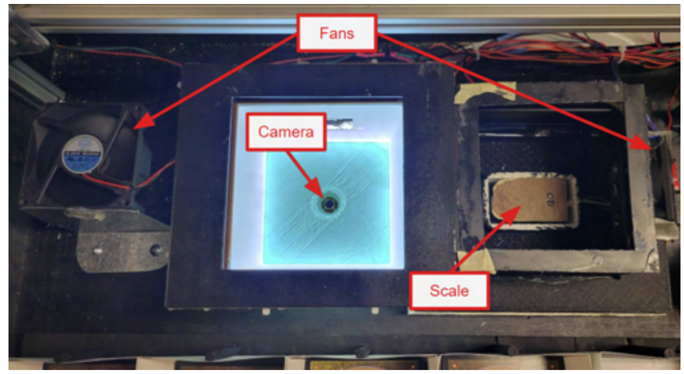

The Evaluator, with key features noted.

i. Cropping and Rotating

Here, we explain briefly the act of cropping and rotating the card.

It is in the first stage of image processing we already see the benefits of having multiple illuminations on the card…

Raw imaging of the cards

As shown above, the side lighting(right) clearly outlines the border of the card, which can be leveraged by the algorithm for efficient and accurate detection of the edge. From this, we can derive a contour of the card using an algorithm supplied by OpenCV. From that, sample a subset of the contour points to derive a median slope which can be used to calculate how much the card should rotate…

Median Slope

From this, we redraw the contour and find the tightest fitting points to get the crop and thus have completed the crop and initial rotation. The CardID algorithm expects the card to be upright when it receives it, so the pard is first rotated 90 degrees, and then the std of pixel values is taken for the top half and bottom half. Whichever half has a higher std is considered to be the true top half, and the card is re-rotated as needed. That information used to rotate the card with side lighting can then be used to identically rotate the card illuminated from above, resulting in a properly cropped and rotated image with minimal visible damage on it, which is ideal for Card ID.\

Cropped and Rotated Images

ii. *Edge Damage Detection*

Though edge damage and surface damage are handled in separate algorithms, they still share the same general principle of comparing pixel brightness between images. We discuss edge damage detection first. First, we detect the edge on the side-illuminated card using the contour algorithm. This acts as a path for a sliding window algorithm. From this, the window slides over the edge of the above-illuminated card, collecting “scores” of damage, which is a function of the amount of whiteness in the image. The higher the score, the more likely there is damage. The above-illuminated card is used, as the only scenario where white could be visible on the edge when illuminated from above is if it is some form of damage (or dust, for which a fan attempts to blow it away before imaging takes place).

From this, we get a signal of the whiteness inside the window versus the position of the window along the path. This signal can then be parsed according to spread and amplitude according to user preferences to determine edge damage. The picture below shows highlighted damaged reasons based on a set of user preferences.

Above images are examples of the edge damage map generated on several different cards.

iii. *Surface Damage Detection*

It becomes intuitive to solve for surface detection when describing variable illumination as a transformation on a vector space. We first start with describing a card, C, as a set of vectors in a color space, here we use RGB, though we have also explored color constancy spaces such as YCbCr and HSV. We assume a fixed position between the card and the camera to let T be the set of all possible illuminations in the device. So we have

$$C \subset R \cup G \cup B \text{ and } T(C) \subset R \cup G \cup B$$

It can be very difficult to model the difference between two transformations in a predictable manner, so instead we rely on statistics to separate vectors of damaged and undamaged pixels on the card. We assume that the set of damaged and undamaged pixels follow different paths on RGB space when undergoing a transformation (i.e switching from above lighting to side lighting). However, we must organize our pixels by similarity, so when calculating an expected path, we consider the location of the pixel on the card - due to uneven lighting conditions - by considering all pixels under a sliding window. We also only consider pixels that are close in RGB space. The illustrations below summarize this process (for simplicity's sake, we explain using 2d spaces)…

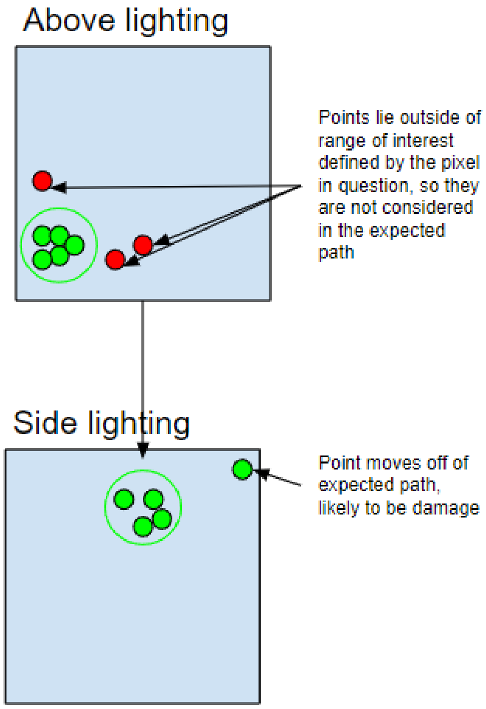

From this, we generate actual and expected trajectory arrays which are used to separate damage and dust from the card.

D. Forgery Detection

The MTG secondhand industry has several standard methodologies for detecting forgeries, so this project aims to replicate several of them. The first was that some fake cards have a mass of over 1.8g, whereas all real cards had a mass under 1.8g. A load cell connected to an (differential) amplifier allowed for a precise measurement of the mass of the card. A version of this was implemented for MDR, but there wasn't time to put it into the system for CDR or later.

A second metric was the existence of 4 red ink dots in the small white portion of the green circle on the back of cards. This is a feature impossible to detect with the human eye, but with a camera that could zoom in to look at cards more closely the problem became more manageable. An RPi camera was purchased that had a focusing motor on the lens, allowing for closeup images to be taken. An algorithm was used to find the green circle on the back, from which the white portion and the red dots could just barely be detected. However, there wasn't time to implement this in the final design either.

The glossiness of the card was also something that could be measured by the level of glare reflected from a point source, with many fraudulent cards being either too dull or too glossy under light. Because we only have 2 light sources, and one camera, this did not work on the prototype. A future, improved version may be able to accomplish this with additional cameras and lighting sources.

Finally, legitimate cards have a specific pattern to the printing of their art— with rings of dots and circles about them. The system is capable of seeing this, though software wasn't written to detect this pattern consistently.

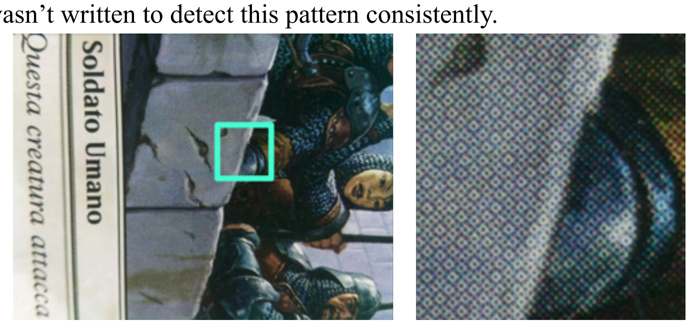

The print pattern for a legitimate card

E. Card Identification

The card identification is performed by an algorithm featuring both perceptual hashing and feature matching on a large database of reference photos of every card ever printed. The perceptual hashing is used to narrow possible card matches to only the few most likely options. Then, feature matching is used to determine which card most closely matches. This 2-step process is used as perceptual hashing is very fast and efficient to search through the database, but is unable to distinguish between similar cards, such as a card that has been reprinted in a later edition. Feature matching has sufficient accuracy to provide this final narrowing step, but is too slow to run over the entire card database.

III. The Refined Prototype

[1] *A. Prototype Overview*

The FPR-ready prototype features the ability to continuously pick up cards from a bin, move them to all required stations, transfer and process all images, place the card in another bin, and then repeat with another card, all without human intervention.

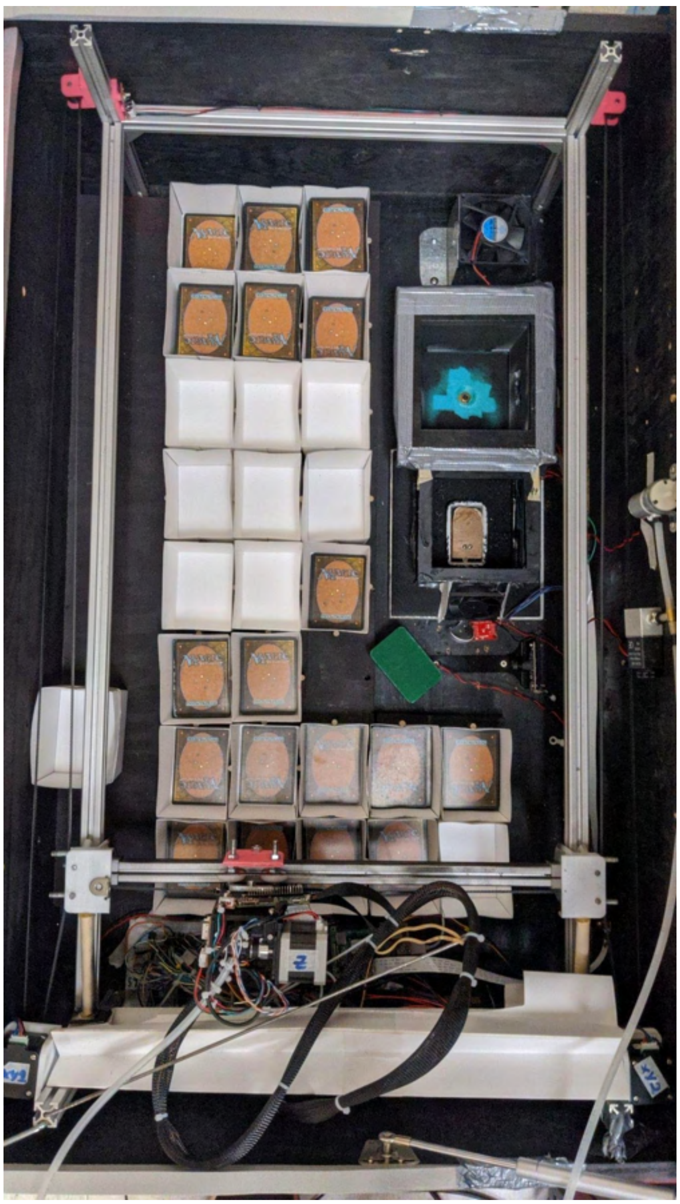

Figure : The prototype design.

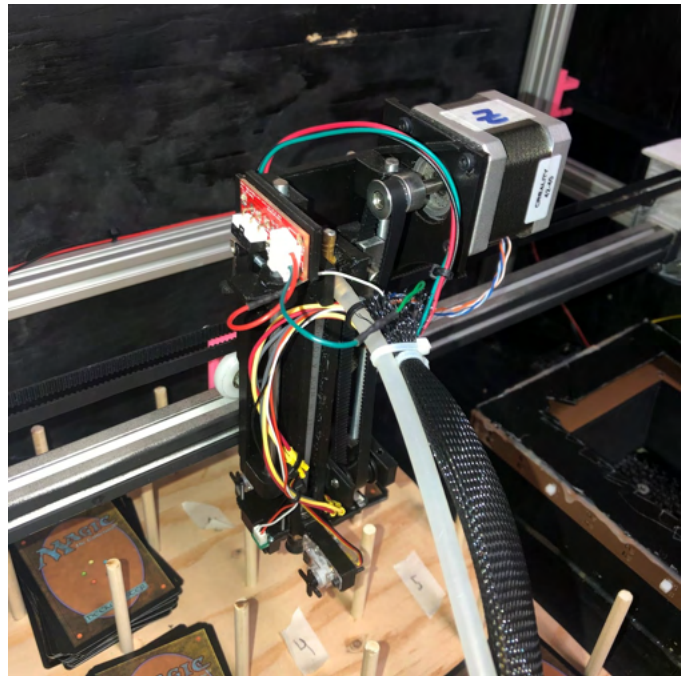

Figure : The gantry head featuring camera, endstops, vacuum hose, pickup system.

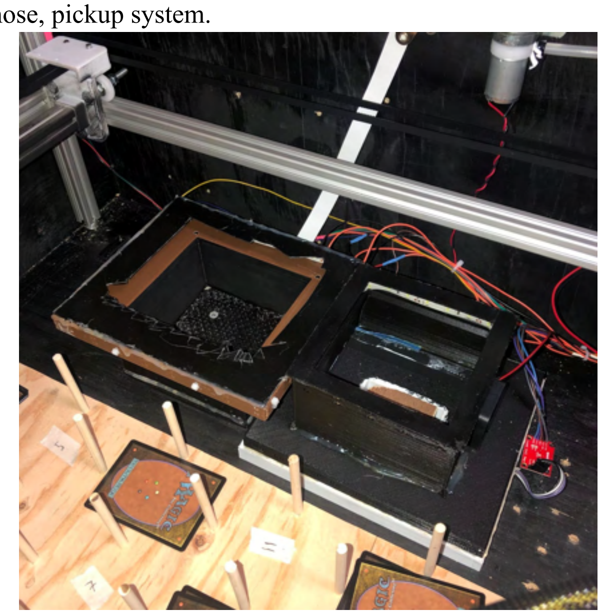

Figure : Top and Bottom Evaluators

B. List of Hardware and Software

The prototype is a pick and place system based around core-XY gantry, constructed from aluminum extrusion. 3 stepper motors allow for the 3 degrees of freedom necessary to transport cards to all required locations. A pump-actuated suction cup on the head is used to pick up and carry cards. The microcontroller is an RP2040, and the firmware for it is coded in C. The image processing is all done in Python.

[2] *C. Custom Hardware*

The main board PCB features an RP2040 for controlling motors, toggling LEDs, and receiving input from sensors. It also takes a 24V input and steps that down to 12V and 5V so that almost all of the system can be powered from one source. Another PCB was designed to hold the three stepper motor drivers used by the system, though it is not used in the final prototype.

The mechanical system is based around a gantry system built of extruded aluminum and using stepper motors for XYZ motion. The overall enclosure is a wooden frame.

D. Prototype Functionality and Performance

The prototype used for Demo Day picked up cards, brought them over to a fan to attempt to remove large particles of dust, and then brought the card to the first imaging station where pictures were taken at the two lighting conditions, and a closeup was taken at a different focus to allow for pattern-detect, though that algorithm was not implemented. With the images captured, the cards were cropped for cardID, the effectiveness of which is discussed in appendix C. Additionally the damage on the edge was detected and used to find a grade for the card, which in turn was used to sort the cards into output bins.

IV. Conclusion

The prototype has demonstrated all functionality necessary to justify additional funding for a future group to take it to the next step. Current competitors overcharge and underdeliver, and this prototype undercharges and underdelivers

Acknowledgment

Professor Parente demonstrated great interest in the development of the project, and encouraged us to stay on track.

Our Evaluators, Professor Gummeson and Professor Polizzi, offered good feedback and seemed to be genuinely interested in our project, which we greatly appreciated.

M5 and the All-Campus Makerspace were vital for this project. The resources, expertise, and hospitality we received were very helpful.

References

[1] "Magic Card Market." *PW Consulting*, 7 Mar. 2022, https://pmarketresearch.com/magic-card-market/.

[2] Markets, Research and. "U.S. Sports Trading Card Market 2021-2027: Digital Integration of Sports Trading Card Games Drives Growth." 13 Dec. 2021

[3] https://www.rocarobotics.com/

[4] https://www.beckett.com/grading/

Appendix

Sections *A - F* are required appendix sections. You may refer to these materials in the report body as needed. An appendix is a useful catchall when there is too much detail to include in the report body, but where this content is still useful to be communicated. Include such details in Section *G, H, …* and so on.

A. Design Alternatives

We originally considered a design featuring a roller-based setup for handling cards. Rollers offer significantly improved speeds over the pick-and-place style system the final prototype uses, and allows grading and sorting activities in parallel. These benefits are overshadowed by what we felt were significant challenges with the mechanical design of a roller system, the increased chance of causing damage to cards as they are handled, and feedback we received from a potential customer de-emphasizing the speed of operation over fail-safe operation and a greater number of configurable outputs. A pick-and-place system permitted many more outputs, as it more readily allows bins to be placed in 2 dimensions.

B. Technical Standards

UART was a standard used for communication with the RP2040, and serial on C and python was used to send commands and images between the RPi 4 and the external PC. The USB-USB PCB was designed to try to fit the standards for USB C connectivity. This meant ground vias were necessary next to high frequency vias, and differential traces had to be routed right next to each other while trying to be close to the same length.

C. *Testing Methods*

Card throughput & output verification: We wanted >2 cards/min, and >1000 card capacity. For the process used on demo day that excluded taking pictures of the back of the card, it still took 70 seconds to process a card. There were 28 bins that could each easily hold 200 cards, so all capacity specifications were met.

Condition Grading verification: Find >95% of damage visible to the human eye with 90% confidence. Manually map all damage on the digital image of the card. Compare it with the computer-generated damage map. Do this for front/back/sides. 19/20 pieces of damage visually identified should also be identified by the computer. Do this 10 times. We did not have time to actually implement this.

Card Identification Verification: Correctly identify cards 95% of the time. 95% certainty. We ran 134 cards through, and found 123 correctly, which came out to a success rate of 92% with an std of 2%, so we didn't quite reach that mark.

Forgery Detection Verification: perform Weight check, Glossiness check, Printing verification checks. 90% of legitimate cards pass all 4 tests. 90% of fake cards fail at least one. 90% confidence: Run through 10 of each, 10 times. No more than 1 false positive/negative each run. We did not have time for any of these tests.

No card damaging verification: Cause minimal damage to cards through processing. Run 1 card through 5 times. Damage detection software should not detect more than 5% additional damage. Visual inspection should also show no damage. Demonstrate fail-safe by unplugging mid-operation and showing no fires. For this we were able to test that cards were not damaged, and this proved successful. As for sudden stops, this happened frequently and the system was never harmed..

C. *Project Expenditures*

Being a large and hardware-intensive project, the expenditures were significant. There were a number of pieces of hardware that group members already owned, but often a quick order of parts would be necessary, as it was unclear what additional funding would be available until later on. In mid March, a budget extension of $350 was approved, allowing for the purchase of additional PCB designs as well as electronic and mechanical parts needed to finalize the prototype.

D. *Project Management*

The team consisted of two computer engineers, an electrical engineer, and two mechanical engineers. The project was divided into a number of broad tasks assigned for an individual member to be responsible for. Zalman Lipschitz was responsible for the overall gantry design, and was considered the general MechE lead. Liam Rees was responsible for the head on the gantry, which consisted largely of designing an intricate and robust set of 3D printed parts to mount all necessary hardware. Henry Powell was the team lead, and was responsible for the optics design as well as the PCB design. Malcolm Okaya was responsible for embedded system design as well as card ID. Joseph Maloyan was responsible for image processing with respect to damage and fraud detection. This organization of responsibilities usually worked well, but the size of the project meant that every miscommunication was one that couldn't be afforded, and resulted in lost time that couldn't be regained, resulting in an incomplete project. The best that could be done was finding what miscommunications had occurred and figuring out what the best design choice was from that point. The project was one that required significant individual accomplishment, and every time this was achieved it became an act of leadership. Malcolm got a prototype of cardID working very early on, which encouraged the rest of the team to put extra work in. Zalman built a full model of the design in Solidworks quickly, which had a similar effect. One area of struggle occurred when the diodes on the PCB were designed to point the wrong direction, which resulted in a significant loss of hours when traces appeared to be shorting out. Once the problem was detected it was resolved, but the lost time could not be returned.

E. *Beyond the Classroom*

Henry: For my PCB design, I had used KiCad briefly in the past, but for this project I tried learning Altium as well, since it was considered more professional. Having designed PCBs with both for this project, I cannot dispute that Altium has superior trace routing, but the rest of the software seemed unnecessarily convoluted.

Liam: I spent a lot of time working on the mechanical design and on the wiring for the system. While I had had some mechanical design experience in the past, it hadn't been of such a level of complexity or with so many moving parts. Likewise, I had had a limited amount of electrical working experience, which I got to work on and grow through this project.

Zalman: I've learned a ton throughout the course of this project. I primarily worked on the gantry and the enclosure, but my touches can be seen throughout the whole system. Much of the work I did was fine-tuning the system and modifying it to meet the shifting needs. I went into the design with about a year of internship experience and several projects under my belt, but I've never worked on something with the format we had (small-scale, highly organized). Just like many full-scale engineering projects, there are inherently many facets of this project: mechanical, electrical, and software design, budget distribution, customer requirements, and course-defined milestones to meet. There was plenty of cross-discipline work, high level project coordination, and major decision making that we all took part in. This project really embodied what I feel a capstone project should – it was a scaled-down engineering challenge, and it really taught me a lot about the overall design process, the other engineering disciplines, and how to work well as part of a team.

Malcolm: This project has taught me a lot about writing code for embedded systems and working with custom PCBs. I have worked on embedded systems projects before, but not one this complex. The stepper drivers, microcontroller, and single-board computers all communicate with each other to some degree, so there are many moving parts. This also means the PCB was very complicated, so assembly and testing was a long process. While I have done through-hole soldering before, I had limited experience with SMD parts, and learning to solder with hot air has opened up a world of possibilities for future projects.

Joseph: I was able to apply what I had done at previous occupations and in classes in a theoretical manner to the image processing and optics components of the project. I read through many academic papers on optics and computer vision theory to develop a methodology to detect damage on the card without the reliance of machine learning. I believe at the end of the day we have explored a potentially novel way to process images by manipulating the illumination that could have potentially useful applications outside of automatic trading card grading.