Austin Pets Alive! & Maddie's Fund

Canine Parvovirus Research Project

Treatment Study Final Report

Author: Dr. Kevin Horecka, June 25th, 2018

Introduction

The canine parvovirus is a viral illness known for its incredible infectiousness and infliction primarily on young dog populations. Although a vaccine exists which is highly effective in preventing infection, this vaccine takes time to generate immunity, and animal populations lacking access to the vaccine will, obviously, not be protected. Because the canine parvovirus (parvo, for short) can live on surfaces for up to a year, transmitted via the fecal-oral route, it can easily infect virtually any young dog who is unprotected. Moreover, households which are infected may incorrectly assume particular cleaning methods normally utilized for infectious disease control (i.e. soap, all-purpose cleaners, etc) will kill parvo; however, the only household cleaner which kills parvo is bleach. If bleach is not used, it is likely a household will remain contaminated. All of that said, parvo is a very treatable illness. In many private practice environments, the treatment can cost more than 1,000 to 2,000 USD in the United States, and shelters are often hesitant to treat it due to the highly infectious nature of the illness - instead, choosing euthanasia to protect the rest of the shelter population.

Austin Pets Alive! has been successfully treating the canine parvovirus in a quarantine environment since 2009, treating anywhere from 200 to 800 dogs in any given year in a shelter environment with save rates consistently above 80% and often as high as 90-95%. The treatments can be performed economically due to the isolation of all animals within a section of the shelter exclusively designated for the treatment of parvo, the presence of a dedicated volunteer staff to perform treatments twice a day, and the donations of the public of blankets, food, and funds to pay for medicines and other equipment necessary for treatment.

In order to evaluate the efficacy of this program, this report outlines research into the Economic Cost associated with treatment (including volunteer hours) as well as modeling of the elements which predict survival (via a Survival/Hazard model). This work can serve as a starting point for future experimental treatments of Parvo as its results can be used to identify which animals require special intervention. Moreover, this work can be used to set expectations with shelters that want to implement their own Parvo programs.

Methods

Two primary sets of methods were used to evaluate the economic costs and survival rates of Parvo animals in this study. First, the economics of treatment are evaluated via a breakdown of the costs of the medications, equipment, and volunteer time for each animal. A further analysis breaks down how different forms of treatment may have different associated costs. Next, a breakdown of the volunteer hours needed to treat certain numbers of dogs is evaluated via both a linear and logarithmic statistical model. Finally, a linear model is used to predict the cost of treating any particular dog using their weight as a predictor variable.

In evaluating the success of treatment, several evaluations of severity of condition and probability of recovery are employed. First, simple metrics of severity are evaluated to observe the time course of the disease, identifying critical windows during which treatment must be modified if survival rates are to be increased. Next, a Survival Model is used to form a high-level predictor variable which can be used to identify which dogs are not expected to recover given their current state and short-term medical history. This model has the added effect of determining which of the observation variables available are significant predictors of the treatment being ineffective. Finally, this model is used in linear classification and the specificity/sensitivity is evaluated.

Subjects

796 dogs were evaluated with complete data present only for 710 of these dogs. 9403 total treatments were evaluated across all animals with an average number of treatments of 11.39 per animal. Of the 710 dogs with complete data who are used in the primary analyses, 609 survived, resulting in an overall survival rate of 85.77% (slightly lower than typical, potentially due to higher population counts due to hurricane Harvey or due to the approximately 10% of animals being excluded due to inconsistent data). The treatments were performed between January 3rd, 2017 and February 18th, 2018. The treatments were recorded on standard paper treatment sheets and later transcribed into a Google Sheets document and cleaned both by hand and by outlier detection.

Statistical Analyses

All statistical analyses were conducted using Scipy, Numpy, Pandas, and Seaborn packages in Python 3.6. Linear regression and logarithmic regression were performed using a significance value of 0.05 to determine model significance. Survival analyses were performed using Lifelines in Python and classification was performed using Scikit-Learn. All classifiers were evaluated using 100 fold cross validation with F1 score as the primary evaluation metric. Balanced class weightings were employed to account for class imbalance. Logistic Regression classifiers were used with a 0.2/0.8 test/train split.

Results

The results from all of the analyses are described below.

Economic Analysis

Economic analyses were performed to determine the various costs associated with treatment.

Medication Costs

When we examine the use of different medications over the whole data set by time, we see that certain medications consistently account for a large portion of the medication costs.

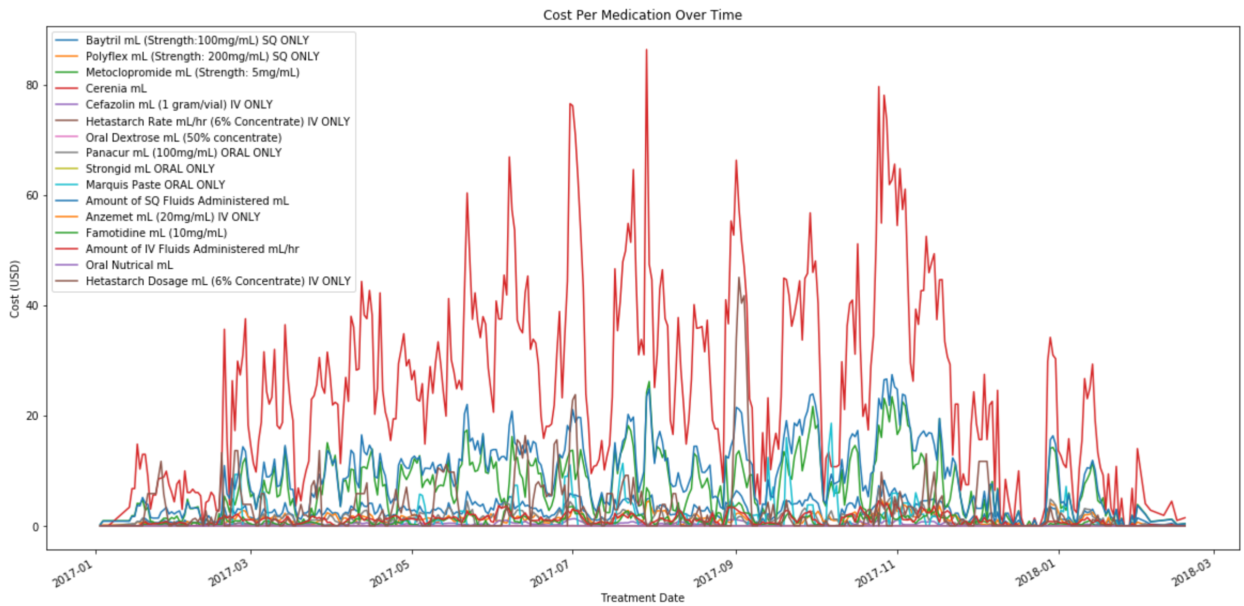

Figure 1: A timeline of the medications used in treating parvo dogs. Note: Cerenia is the most expensive medication.

Cerenia accounts for a huge portion of the overall costs. We can break down this per-medication cost and see the impact of different medications more clearly.

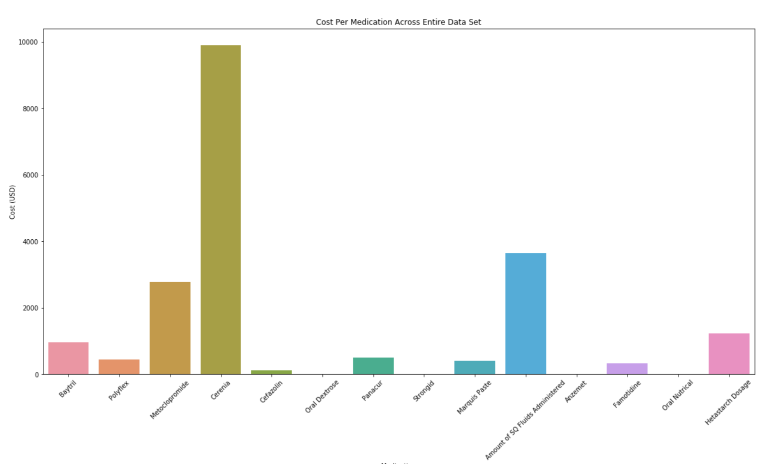

Figure 2: The total cost of each medication across the entire treatment period of 1 year.

Clearly, the amount of fluids, metoclopramide, hetastarch and baytril account for a large amount of the treatment costs together, but Cerenia accounts for almost 50% of the total costs. We can further examine these costs associated with different treatment categories (animals who received IV fluids, animals who received SQ fluids but not IV fluids, and animals who did not receive IV or SQ fluids).

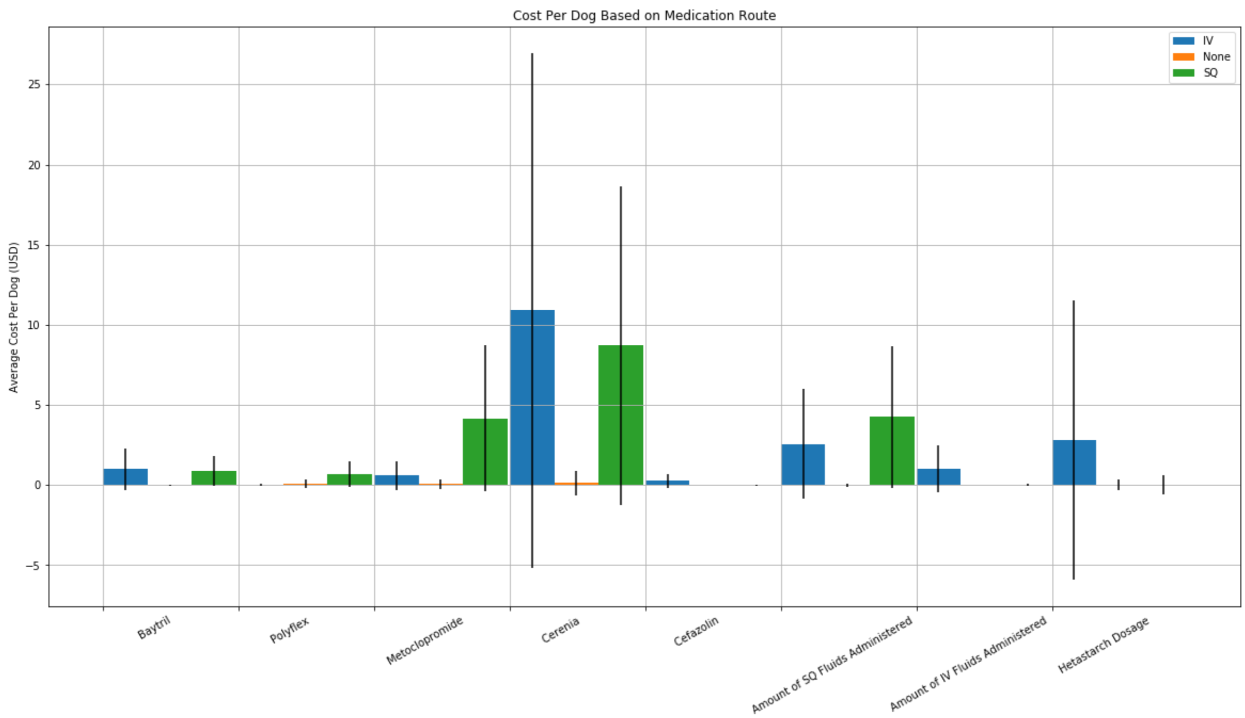

Figure 3: A breakdown of the costs of different medications and their use by IV, SQ, and Oral-only animals.

We can see that the costs associated with Cerenia in the SQ and IV groups are similar, suggesting that there may be room for optimization of usage of that resource. It is also interesting to note that Cerenia has not always been used in the ICU as part of standard treatment, so other ICUs may consider only adding it once the program is stable. It is unclear as to whether or not Cerenia improves outcomes.

Per Dog Treatment Costs

The average cost of treatment of animals is 56.08 USD with approximately 30 USD being cost of syringes, food, and cleaning supplies. This total does not include labor (which can, in principle, be done almost entirely with volunteers) or facilities costs (which are highly variable depending on the area but could, in principle, be run out of a home or other donated space).

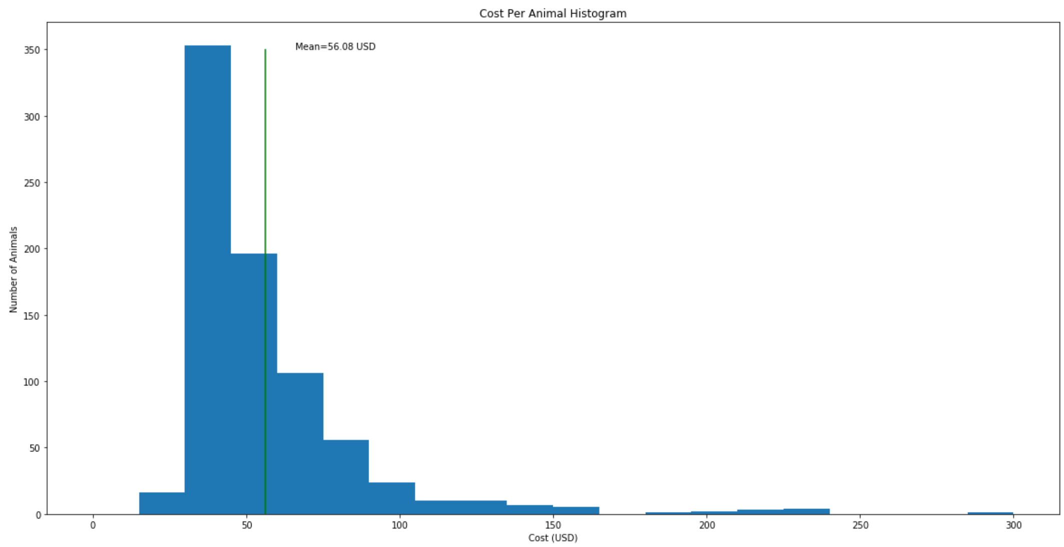

Figure 4: A histogram of the average cost to treat each animal. The mean is shown in green.

One outlier animal existed in the data set that cost 454.88 to treat as this dog was in the ICU for the upper end of the duration distribution as well as in the upper end of the weight distribution. This means it received an incredible amount of SQ fluids, accounting for most of its high cost. In this cases, outlier animals can be avoided by carefully determining when to stop giving SQ fluids.

It is important to note, as is implied by the previous outlier example, that weight is the most critical predictor of cost as weight is used to determine the medication dosage for each animal. When a linear model is trained on weight vs. cost, we find an r squared value of 0.42, suggesting a fairly strong relationship between weight and cost.

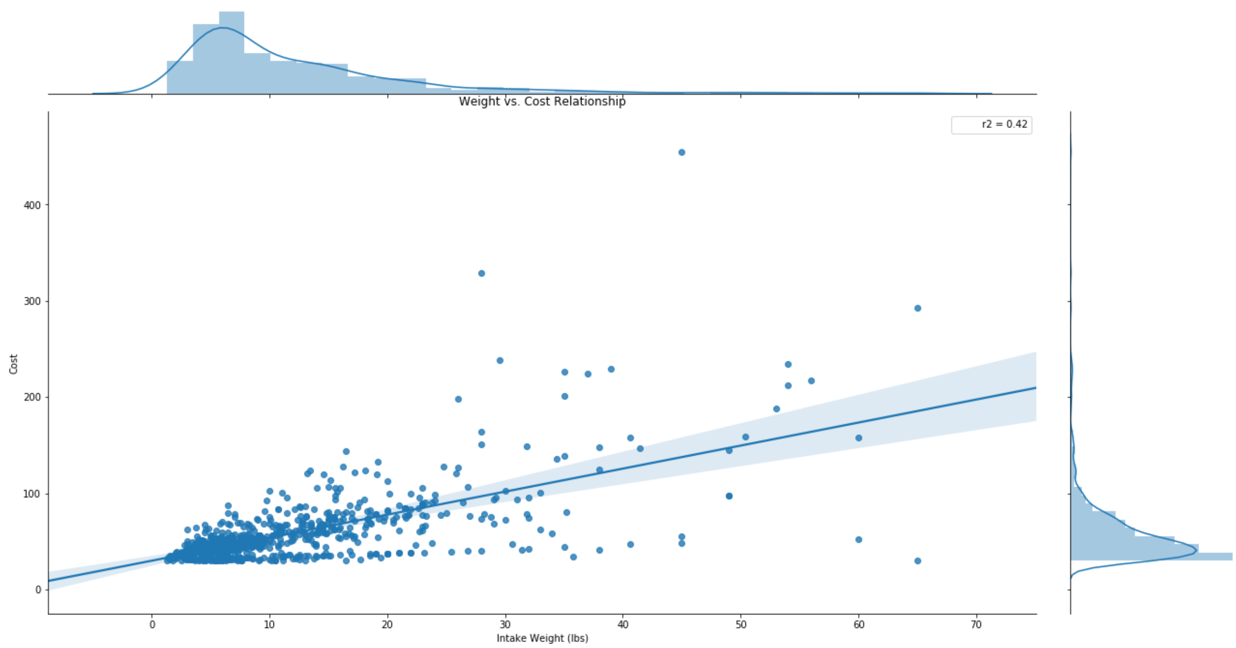

Figure 4: The relationship between cost and weight of the animal at intake. Note that there is a strong positive relationship between these variables primarily because medication dosage is determined by weight.

Volunteer Time

Finally, on the economic analysis, the amount of volunteer time for treatment was evaluated. In the case of this data set, it is critical to note that kittens, in addition to puppies, were being treated in the ICU. Therefore, it is useful to compare models of volunteer time which involve both of these factors. First, we can evaluate the average amount of volunteer time per shift. Note that this number is computed by taking the shift time (typically 1-3 hours) and multiply it by the number of volunteers present during that shift (typically 2-4). This may include care volunteers (who do not perform treatment and focus on cleaning and feeding the animals) along with medical volunteers (who perform medical treatments) and thus provides an overestimate of the amount of time required for treatment.

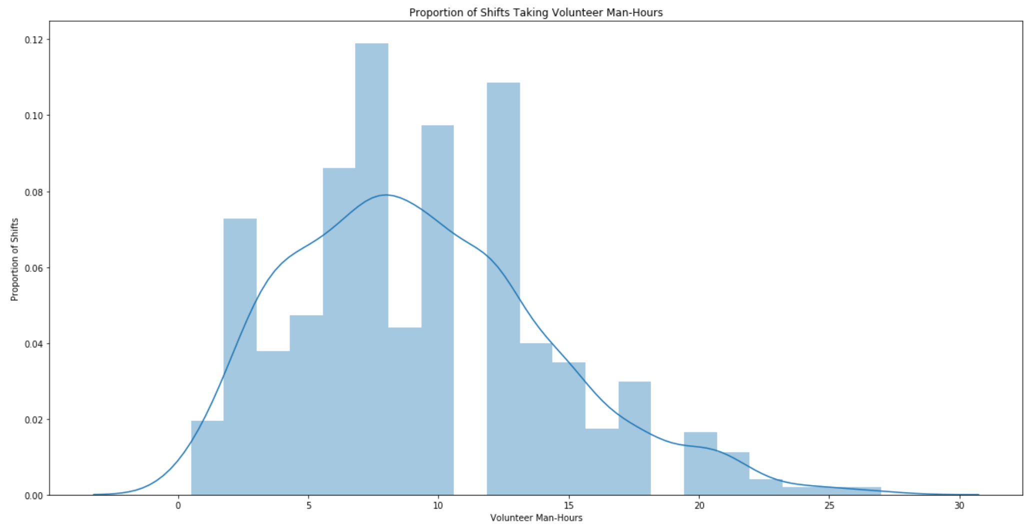

Figure 5: A histogram of the number of hours contributed by all volunteers as a proportion of the number of shifts in the data set. Note that the average is approximately 9 hours (i.e. a 3 hour shift with 3 volunteers) or 9 hours per animal across their entire stay. Critically, this is an average given the 700+ animals treated in the ICU during the year. Smaller shelters that treat fewer animals will likely have far smaller average volunteer hours.

We we can see, the average shift time is approximately 9 work hours (i.e. 3 hours with 3 volunteers). A significant proportion of shifts obviously lie below this mean, and a left skew can be seen in the kernel density estimation, suggesting that the mean is being distorted by occasional artificially large shift durations. Some of these shifts were during Hurricane Harvey when a large number of volunteers were needed to process the influx of animals.

A more granular analysis can be performed by looking at the number of animals present (dogs and cats) as well as the number of volunteers present and the associated shift time.

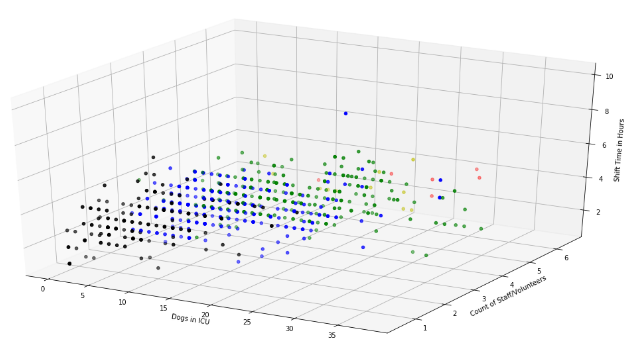

Figure 6: Animals in the ICU vs. Volunteer Count vs. Shift Time. Note that the number of volunteers is roughly scaled with the number of animals in order to keep the maximum shift time under 4 hours.

Here, we can see that the number of volunteers is increased with the associated number of animals in order to keep the shift times below approximately 4 hours. This load balancing is a critical part of maintaining an active volunteer group as consistently long shifts can lead to volunteer burn-out.

Finally, it is interesting to note that although there is a linear relationship between hours of work and animal counts ($p<0.05$, $r=0.69$), a logarithmic model provides a slightly better fit, suggesting that volunteers are able to optimize treatment, to some small degree, when the number of dogs increases.

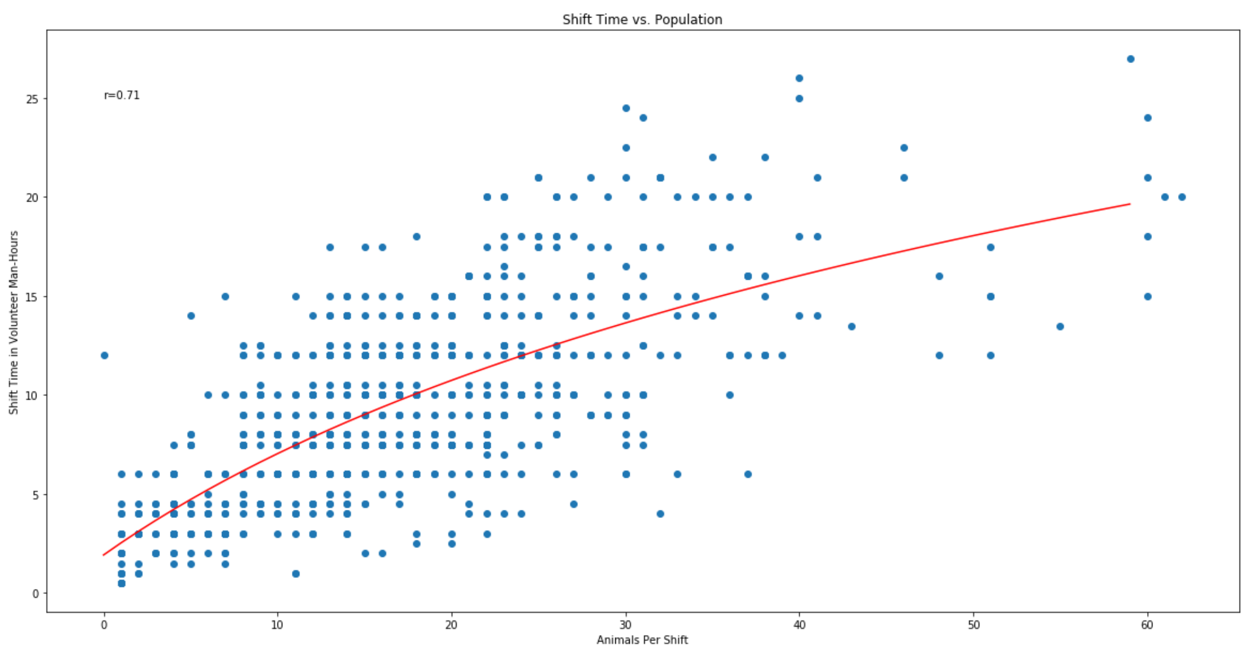

Figure 7: Animals per shift vs. volunteer hours fit to a logarithmic curve. Note that this significant fit shows that volunteers may be optimizing treatment times when higher number of animals per shift are present (i.e. >30 animals).

Staff/Veterinarian Time

Although this component of cost is going to be entirely dependent upon the structure of the organization that runs the ICU, in this case, because of the very large volume of animals treated and the requirement to do data gathering/data entry, 1 veterinarian and 3 staff members participated. The staff members each cost approximately 97 dollars per day (averaged across the full 365 day year) while the veterinarian cost approximately 40 dollars per one hour of evaluation (where 1 hour of evaluation is typically performed per day). Note that this adds approximately 151.32 dollars per animal across the entire year. Prior to 2013, the ICU ran with no paid staff and only a veterinarian. With only a veterinarian's time, the cost increases by 18.34 dollars per dog on average, bringing the total cost to treat per dog to 74.42 USD per dog.

Survival Analyses

In addition to evaluating the costs (and predictors there-in) associated with treatment, we evaluate the predictors for survival which might be useful to determining which elements of the population require additional (potentially experimental) treatments or other alterations to the standard protocol.

Models of Severity

The development of models of severity of condition is a complicated issue. At the simplest level, the indicator variable "Attitude" in the data set represents the simple, behavioral correlate of severity. It does not account for fecal state or vomiting, but it provides an indirect measure of the condition of the animal. When we look at how this variable evolves over time, we see that, for healthy animals, this variable generally only improves. For animals that do not survive, this variable gets worse rapidly.

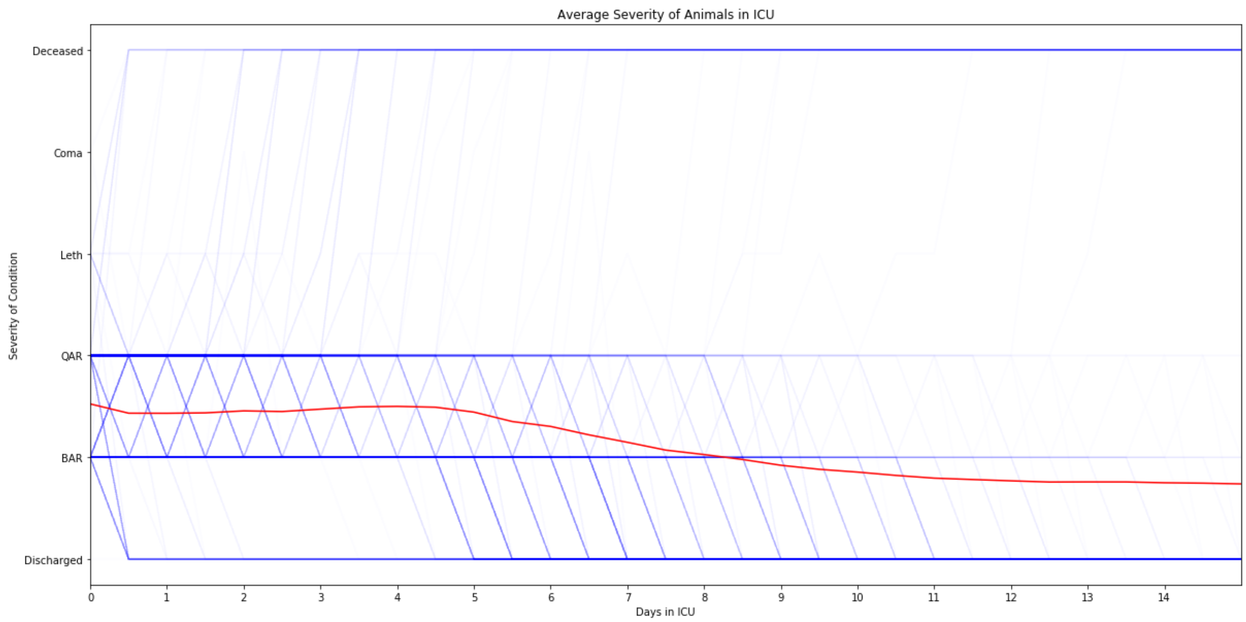

Figure 8: The average severity as evaluated by the Attitude of an animal given the number of days they've been in the ICU. BAR (Bright and Responsive), QAR (Quiet and Responsive), Leth (Lethargic), and Coma (Comatose) values are assigned to each animal during each treatment, providing a simple metric of severity of condition.

When we look at this variable in relation to one of our primary animal descriptor variables (Intake Weight) we find that there is a particular period during which no animals which survive are being discharged (as per the protocol), but animals are still dying.

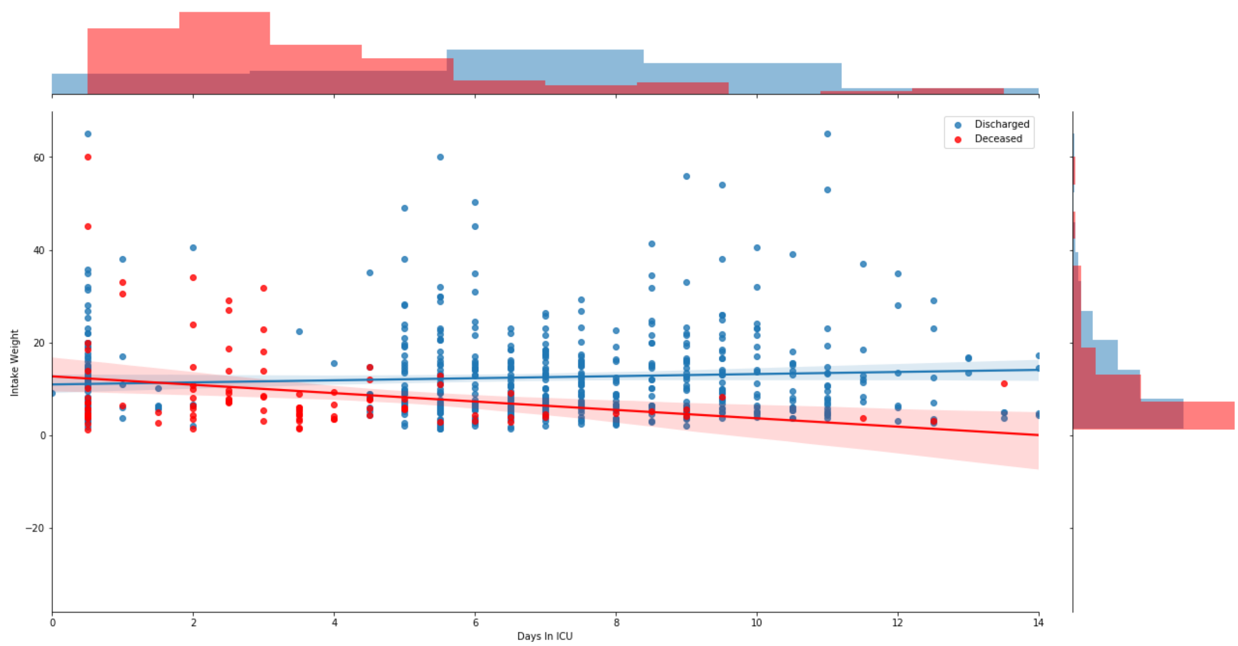

Figure 9: The Weight vs. Days In ICU for all of the animals, separated by outcome (with red animals dying and blue surviving). Note that a significant portion of the animals that die do so during the first 5 days in the ICU. These animals are generally lighter weight than their surviving counterparts.

Additionally, we see a pattern of lighter weight dogs being more prone to dying early, though this pattern is occasionally violated by large dogs dying on intake.

Beyond this single measure of severity (i.e. “Attitude”) several attempts were made to create aggregate measures of severity (Principal Component Analysis [PCA], Linear Discriminant Analysis [LDA], Regression Modeling, and Weighted Sum models) which account for other variables (such as fecal type, eating and drinking behavior, and vomiting). Although some of these severity types point to potential bimodal distributions in the data, all have one critical failing: they do not account for the temporal progression of the illness.

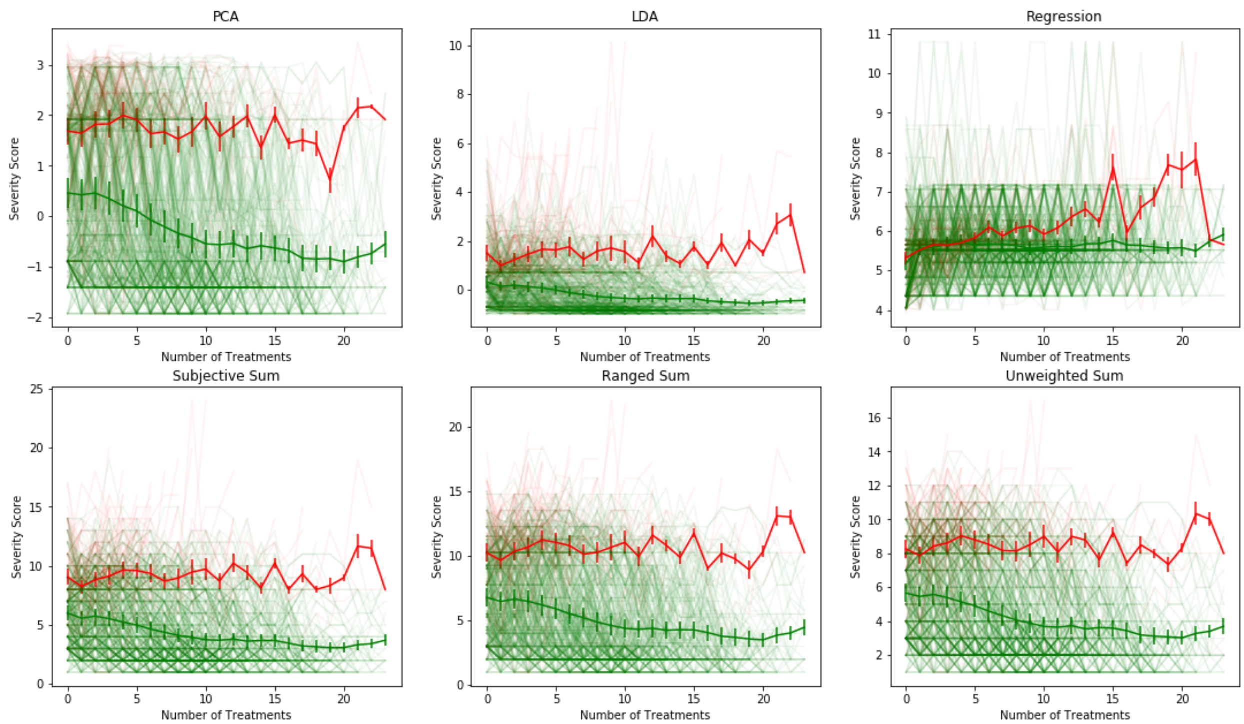

Figure 10: Evaluation of 6 different models of severity based on PCA (Principal Component Analysis), LDA (Linear Discriminant Analysis), Regression (Logistic Regression), Subjective Sum (a weighted sum based on the subjective importance of each variable as evaluated by a veterinarian), Ranged Sum (a weighted sum based on the range of values for each variable), Unweighted Sum (a straight sum of the ordinal values). Note that in all cases, animals that die are generally evaluated as more severe. This suggest there are many ways to generate a severity metric, but none of these methods include temporal information.

Note that many of these methods did separate animals according to outcome (with Red lines representing animals which ultimately died and Green lines representing animals that survived).

Hazard Modeling

A more advanced model which accounts for the temporal information within the data is known as the Cox Proportional Hazard Model. This model is within a class of models known as survival models which attempt to model the probability of an event occuring (in this case, death) given a set of variables which may or may not change over time leading up to the event.

When one of these models is trained on the data, we find only some of the behavioral variables are significant predictors of death at any given moment in time.

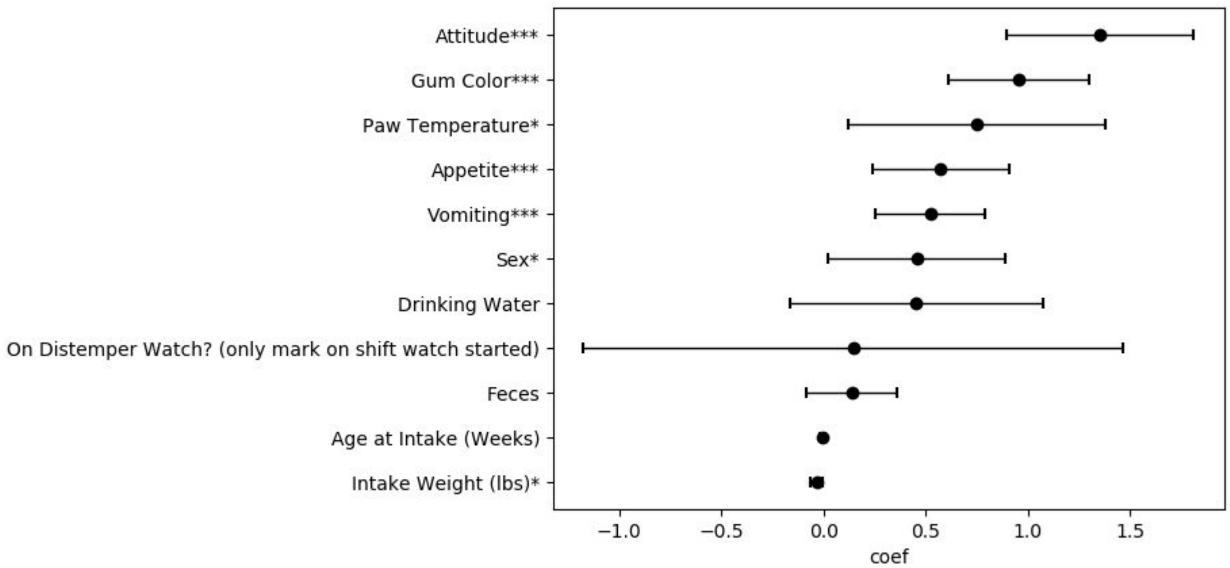

Figure 11: The Range and Significance of the coefficients in the hazard model. * - p<0.05, *** - p<0.001.

In particular, Attitude (from the previous sections), Gum Color, Paw Temperature, Appetite, Vomiting, Sex, and Intake Weight (lbs) are significant predictors, while Drinking Water, Distemper Watch, Feces and Age are not.

As an example of how this model produces a hazard score, we can take a sample of animals from the data set who stayed in the ICU, with 3 members of the sample ultimately surviving to be discharged after the 8th treatment and 3 members of the sample passing away after the 8th treatment.

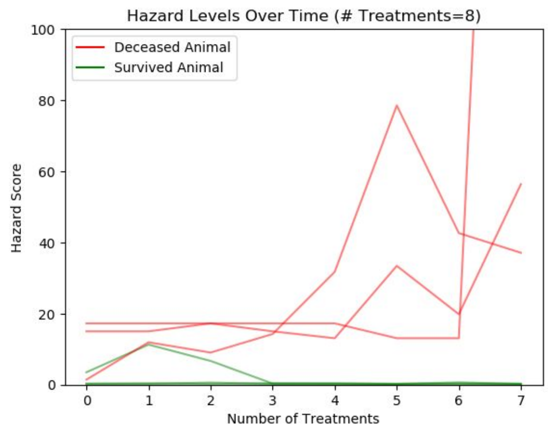

Figure 12: An example of 3 animals that survived and 3 that died and their hazard scores over the 8 treatments they received.

Note that the animals that ultimately died generally began with a higher hazard score from the first treatment (though this was not consistently true until the 4th treatment). When we look at this across the entire data set, we find that no animal that reaches a hazard score of 896 ever survives, making this threshold a useful threshold for determining when to perform experimental treatments.

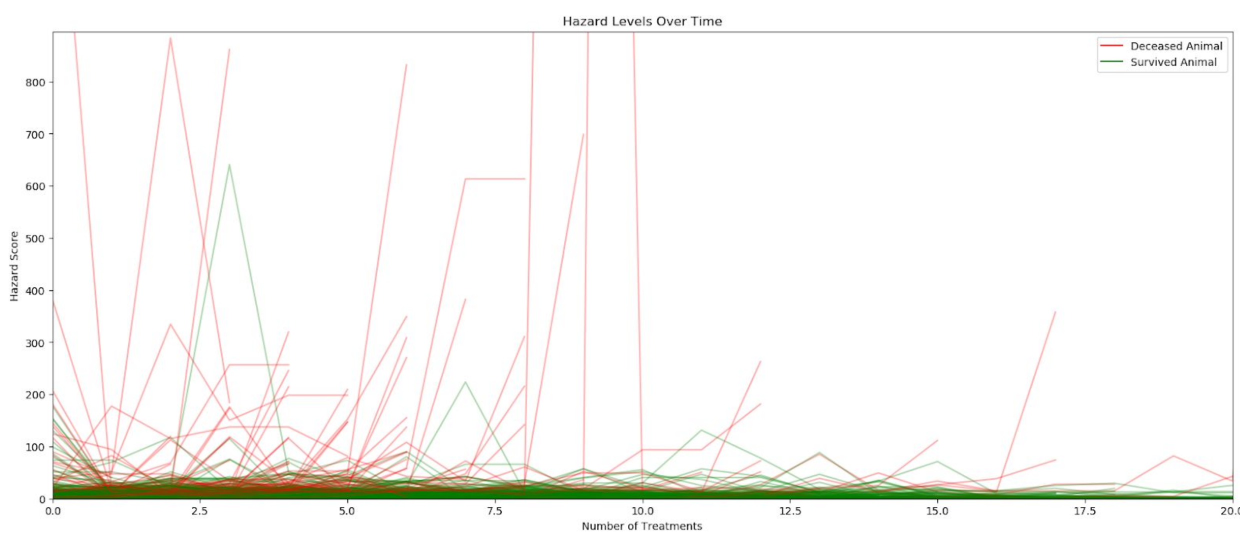

Figure 13: The hazard scores of all animals in the data set. Note that although there are outliers, the hazard scores of dogs that die generally spike during the pre-death treatment.

When we aggregate this data, we find some interesting patterns which might inform when experimental treatment should begin.

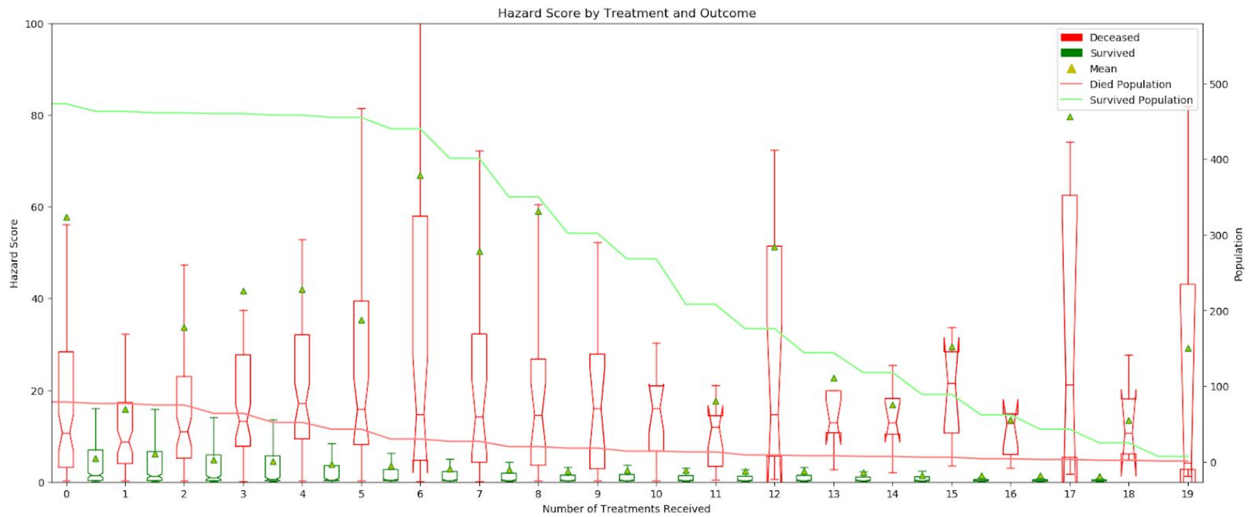

Figure 14: A statistical breakdown of the status of all survived and died dogs across each treatment they receive. The solid lines represent the populations within each group across treatment days (it can only decreased monotonically). The box-and-whisker plots show the quartiles and median for each treatment's hazard score for each group. The triangle represents the associated means. Note that even on the first day, many dogs can be clearly identified as more severe and unlikely to survive.

Note that for many animals which die, their hazard score is high from the first day they arrive. As animals begin to die more rapidly around treatment 2, it is critical that these animals be identified as early as possible. After the 8th treatment, most of the animals who are going to die already have.

One advantage to using a score like this is that it can be trained and cross validated to improve predictive power with more data, ultimately allowing it to be used as a classifier. This means its false positive and negative rates can be directly evaluated. Although this work is ongoing, the currently trained linear model performs moderately well with approximately 25% false positive and negative rates. This performance will likely improve as we continue to evaluate the hazard score and measures of interest with more complex models.

Discussion

This work provides several critical analyses which may serve to aid new clinics intending to treat the canine parvovirus en masse as well as clinics which may want to perform experimental treatments to boost the standard survival rate for parvo from ~90% to 95% or even 100%. First, an economic breakdown of the costs involved in mass treatment of the disease is provided. This included a base cost of 30 dollars per dog for food, syringes, and cleaning supplies but did not include facilities costs or volunteer time. On average, animals cost just over 53 dollars to treat, making treatment incredibly affordable. Moreover, the largest contributing factor to the medical cost was a medication which is not required for treatment (but improves quality of life during treatment). Thus, costs could potentially be even lower than this while still maintaining a satisfactory save rate. Volunteer time is the most critical component of treatment, with an average of just over 8 hours per animal of volunteer time and 18 volunteer hours (computed by multiplying the number of volunteers by the time spent in the ICU) per day. This average is, of course, an average given the extreme volume of dogs seen in this data set (over 700 in a year). The value of 8 hours per dog (spread out over the average stay in the ICU of 10 days) should be considered a more reasonable estimate for clinics beginning treatments. Additionally, we see that the critical driver of cost is the weight of the animal. This provides another potential avenue for cost savings by identifying the animals which are in better condition (i.e. eating and drinking on their own) and reducing more costly medications and treatments (with fluids being a prime candidate). Finally, we find that volunteers can potentially optimize treatments when larger numbers of animals are present, but critically, management of any ICU seeing large volunteer animals should balance the number of volunteers per shift against the number of animals to keep the total amount of time any given volunteer is working under 4 hours per shift, max.

In addition to the economic breakdown provided, an analysis of factors contributing to survival given the currently presented protocol is examined. Several critical factors are found to be related to survival rates. First, lighter dogs are generally more susceptible to dying from the disease. These deaths generally occur before the 4th or 5th day or treatment, thus requiring early identification of critical animals for alternate interventions to be available. The use of hazard/survival modeling presents one potential avenue for this early identification, showing false positive and negative rates of approximately 75% in the first day of treatment at identifying animals which will go on to die. For animals that exceed a particular threshold of hazard score, experimental treatments may be warranted such as fecal transplants. These, of course, should be at the discretion of the veterinarian overseeing the treatment.

In conclusion, the canine parvovirus is an imminently treatable illness. Not only can nearly 90% of treated animals survive with supportive care, but this care can be given with fairly minimal economic requirements (~53 dollars and ~8 hours of care per animal). Future work should attempt to determine if animals for whom the treatment will not succeed can be identified early in order to try alternate interventions.