PCRC Proposal Cover Sheet

Title of Study or Project:	Clustering Anesthesiology Case Data and Reinforcement Learning Decision Analysis
Primary Institution:	University of Michigan
Principal Investigator:	M. Burns
Co-Investigators:	L. Saager, D. Colquhoun, M. Mathis, J. Vandervest, Amy Shanks
Type of Study:	Retrospective Observational
IRB Number/Status:	HUM 00149171 (in process)
Hypothesis:	We hypothesize that anesthetic care varies widely across institutions and providers. We will study the heterogeneity of anesthetic practice and associated outcomes through clustering techniques, and in applying machine learning models we aim to understand and improve perioperative practice.
Number of Patients/Participants:	All valid cases from all sites contributing data to MPOG
Power Analysis:	None required
Proposed statistical test/analysis:	Machine learning techniques will use training and validation separation with 5-fold cross validation.
Resources (Brief summary of resources for data collection, personnel, financial):	Programmer to pull the data and statistical analysis

Clustering Anesthesiology Case Data and Reinforcement Learning Decision Analysis

M. Burns, L. Saager, D. Colquhoun, M. Mathis, J. Vandervest and colleagues

Background

Anesthesiologists frequently handle multiple cases in parallel, each requiring specific plans and modifications based on perioperative changes. Anesthesiologists make decisions regarding monitoring needs, airway management, IV access, fluids, blood products, laboratory draws, medication selection, route and dosing, among others. Throughout an operation, anesthesia providers make additional real-time decisions on a minute-by-minute basis. This amounts to over hundreds of decisions for each case and allowing for a wide range of variation in anesthetic practice. The decisions anesthesiologists make, both before and during a case, influence a patient's risk of complications (e.g., intraoperative blood loss, myocardial infarction (MI), stroke) and postoperative adverse outcomes (e.g., nausea/vomiting, pain control, respiratory distress)[1-4]. Plans and subsequent anesthetic execution are often provider specific and predicated on his/her training and the protocols of the institution under which they provide care. This allows for a large variation in how anesthesiology is performed, across patients, providers, and institutions. Additionally, while organization-based protocols exist, executed plans vary widely, with little agreement regarding what constitutes optimal. We plan to investigate this variation in anesthetic care.

While others have conducted observational studies regarding variation in anesthetic care and its impact on perioperative outcomes, these studies have been limited to either one particular operation or one particular decision (e.g., general anesthesia vs. neuraxial anesthesia)[5-7]. Moreover, prior work includes limited data regarding the patient receiving the care. In contrast, we plan to develop data-driven techniques that incorporate patient data available prior to and throughout their care to help anesthesiologists navigate the wide range of decisions they make. In doing so, we will learn patient-specific treatment protocols for improved outcomes, with a focus on the education of anesthesia providers as to the safe variability in patient care.

We will study the heterogeneity of delivered anesthesia and associated outcomes. This heterogeneity will be evaluated from several levels: surgical case, institutional, etc. From these data, we will gain insight into existing variation in the practice of anesthesiology and will develop a computational framework to learn from this heterogeneity and improve patient care. The results of this work aim to give anesthesiology providers deep insight into their practice. We will build a framework to identify important decisions made in anesthetic care and build a platform and to allow providers the ability to view variation in care. We have three specific aims of this project: (1) characterize heterogeneity in anesthesiology practice and associated outcomes, (2) create a visual tool for providers to better understand variations, (3) utilize aspects of this case organization within a specific case group to analyze intraoperative decisions made by anesthesiology providers as a method to gain insight into optimal anesthetic plans. In this work we will use clustering, natural language processing, and machine learning (including reinforcement learning techniques in time-series analysis) to both understand practice variation and to develop deeper understanding of the choices anesthesiologists make in their care for patients. We aim to investigate these choices in their variation, timing, quantity, quality, and surrounding informational and decisional structures. We aim to provide users with a web interface to view specific plan variation. Ultimately, this tool will be a valuable informational resource for providers at every level of their continued medical education. While this work will focus on the delivery of anesthesia, we expect that the theoretical and empirical results of our work in anesthesiology will generalize beyond this specific case, to other aspects of patient care involving sequential decision making.

Specific Aims

Aim 1. Characterize heterogeneity in anesthesiology provider decisions (1A) and associations with risk-adjusted outcomes (1B). Using case level, provider level, and institution level perioperative EHR data we will identify variables (factors) on which anesthesia providers make important decisions at a case-level. We will evaluate these factors as decision points with the intent to determine the critical features within care variation. Important risk-adjusted outcomes will be evaluated for the case heterogeneity as a means of clinical comparison.

Aim 2. Create a clinically usable visual tool to understand variation in decisions specific to cases and institution. This visualization tool will be constructed with the intended use in education for anesthesia providers to identify variation in anesthetic practice.

Aim 3. Analyze intraoperative decisions made by anesthesiology providers as a method to gain insight into optimal anesthetic decision plans. We will build on Aims 1-2 to investigate intraoperative decisions made my anesthesia providers on a minute-by-minute basis. For this aim, we will identify a cohort of cases in a defined patient subpopulation, specific outcomes, and target features after analysis of aims 1 and 2. With this aim we seek to provide near real-time prediction for common but complex treatment modalities.

Methods

Patient Population

The study population will consist of all procedural cases documented in the MPOG database including both adults and children. Not limited by age, surgical procedure, time period, or ASA status. An example of a cohort specific for Aim 3 of this project is adult patients undergoing major vascular procedures between the years 2010-2018.

Exclusion Criteria

We will exclude cases that have incomplete or duplicate records. An example exclusion criteria specific for Aim 3 of this project are patients with a status of ASA 6 or existing mechanical circulatory support.

Primary Objectives

Aim 1A: Identify and characterize by descriptive data within MPOG, as determined by clinical utility and usefulness in clinical practice.

Aim 1B: Establish and assess utility of risk-adjusted outcomes to be incorporated within a clinical tool for anesthesia providers. We will utilize phenotypes describing perioperative outcomes including 30-day in-hospital mortality, MI, acute kidney injury (AKI), estimated blood loss (EBL), stroke, pulmonary edema, case-specific information such as fluid and blood product administration, and provider quality metrics through the Anesthesiology Performance Improvement and Reporting Exchange (ASPIRE) such as blood-pressure and temperature monitoring and control.

Aim 2: Present the framework of important decisions made in anesthetic care within a visual interface to allow providers the ability to view variation in care.

Aim 3: Learn and evaluate intraoperative treatment utility to understand and improve care.

Data Organization and Clustering

Clustering is a critical component to all three aims of this project. We seek to characterize variation across treatment plans within a given patient subtype, based on case characteristics available prior to the operation. We will compare two different approaches to clustering: one based on a raw feature representation and another based on a learned latent representation that leverages the temporal structure of the medical plan. We can cluster patients into subtypes and executed plans and visualize empirical distributions. Each case will be represented by a high-dimensional feature vector.

We will start by investigating important decision points within anesthetic care. These decision points will help identify heterogeneity in anesthetic care and associated outcomes. We will cluster patients based on their pre-operative case characteristics including those found in Table 1.

Table 1: Patient and Anesthetic Management Characteristics

Categorization	**Features**	**Specified groupings**
Patient	Surgical procedure type	Anesthesia CPT codes
	Age	Years, grouped by number of decades over 30
	Gender	Male / female
	ASA status	1 / 1E / 2 / 2E / 3 / 3E / 4 / 4E / 5 / 5E / 6
	Comorbidities	Elixhauser comorbidities[8]
Case	Medications administered	Inhalational Anesthetics IV Sedative Hypnotics Analgesics Benzodiazepines
	Monitoring	Arterial line Central line PA catheter Echocardiography
	Airway management	MAC General - Supraglottic airway General - Endotracheal tube
	Anesthetic technique	Neuraxial Peripheral nerve block

We will first develop a vector representation for anesthetic plans. In addition, to a handcrafted representation, we will also consider representations learned via sequence embedding techniques (e.g., LSTM auto-encoder). Based on these representations, we will cluster anesthetic plans within a patient subtype cluster, and will visualize the different anesthetic plans based on the median of each cluster. After clustering and visualization, data will then be organized into hierarchical structuring after learning a mapping from the preoperative information to anesthetic execution. Clustering in this way allows us to partition decisions so we determine the common and the important features from each specific procedure. For example, important features having high heterogeneity in major abdominal surgery may include use of an arterial line and multimodal analgesia; when applied to cardiac surgery, however, these features may lack heterogeneity, and features such as selection of vasoactive medications and transfusion strategy may prevail as important features. An example can be seen in table 2.

Table 2: Simplified Hypothetical Cluster Results of Important Features for Specified Procedures

Procedure	Important features with high heterogeneity
Abdominal surgery	Use of arterial line, multimodal analgesia
Cardiac surgery	Vasoactive infusions, transfusion strategy
Total knee/hip surgery	Use of epidural, peripheral nerve block

Techniques that we will implement include spectral clustering, K-means, Density-based spatial clustering of applications with noise (DBSCAN), etc. Visualization of the clusters will be a key component in assessing clustering applicability and will be achieved using orthogonal and nonlinear dimensionality reduction with principal components analysis (PCA), T-distributed Stochastic Neighbor Embedding (t-SNE), among other methods (Figures 1-2). We will also select the number of clusters based on the change in eigenvalues when using a spectral clustering approach. Clusters will be validated based on stability in held-out data.

Variation in medication plans within a patient subtype may result from variation within the subtype itself. Since subtypes are identified in an unsupervised manner, leading to potentially irrelevant clusters, we will also explore a supervised learning approach. Specifically, we will learn a mapping from patient characteristics to a distribution over possible medication plans. Based on a global clustering of medication plans, we will label patients based on the cluster assignment of the executed medication plan. This results in a multi-class classification problem.

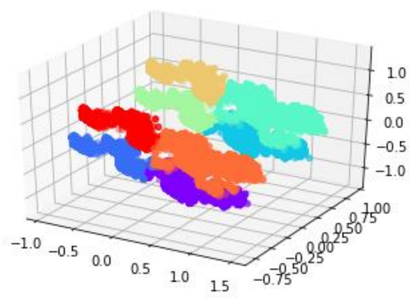

Figure 1: Clustering of case features using k-means and 8 clusters

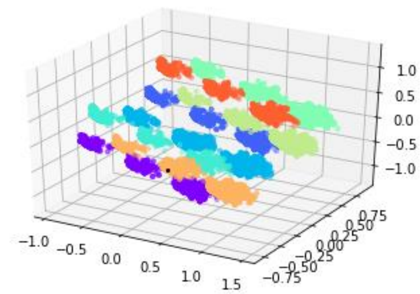

Figure 2: Clustering of case features using DBSCAN

Interactive Web Interface

We will create a web interface to allow users to explore and visualize variations in executed anesthetic care. The initial rendering of this interface can be seen in Figure 3. The user will login to the web interface which will limit views based on predetermined allowances for the user. For example, providers at a single institution will be able to see anesthetic paths specific to their institution and those in aggregate of the MPOG institutions, but will not be able to see information from a specific separate institution. After successful login, the user will be allowed to select case and patient information and view the details of anesthetic paths and their associated outcomes. We will design user inputs functions in the form of drop down menus, multiple selects, free entry, etc. To prevent specific case identification, views will be rendered only for a minimal number of cases, i.e. if the query does not render the number of cases to exceed the minimum threshold an error message will return. Successful queries will result in tabular and graphical displays that will themselves be interactive for the user. We anticipate the ability for the user to select single features (ex. a specific medication) to anchor or avoid within a query. Analysis of outcomes will be provided specific to the case query results. Outcomes and other features in the interface will be managed selectively by site champions, allowing the regulation of dissemination of information as to help limit misinterpretation.

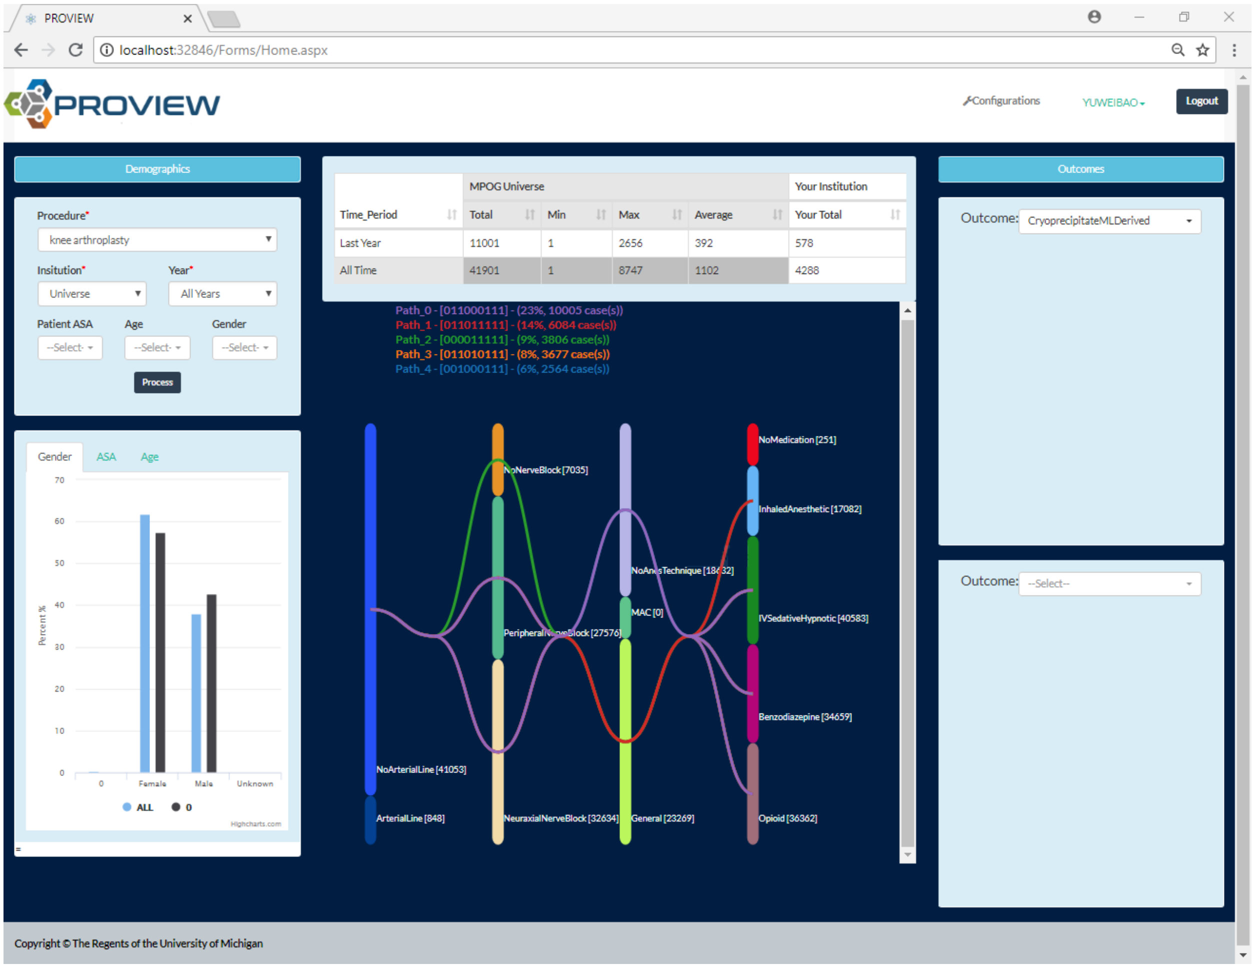

Figure 3: Knee arthroplasty anesthetic practice variation. The top five unique paths are illustrated by line tracing. Segments represent clustered anesthetic choices. Search criteria are represented as drop down menus in the upper left. Graphs at the lower left represent patient demographics within the returned query. The table at the top of the figure shows the selection highlighted, displaying the number of cases returned from the query.

Decision Analysis

Specific for Aim 3 of this project, we plan to investigate intraoperative decisions made by anesthesia care providers. Specific details into how the decision analysis will be performed can be found in the Machine Learning Methods and Implementation section of this document.

In addition to identifying potential management plans tailored to individual patients and procedures, the proposed approach will provide estimates of the probability that the anesthesia provider will have to deviate substantially from the proposed plan (i.e., estimates of uncertainty). This will be achieved by analyzing the variation in anesthetic path. This analysis will allow providers to further anticipate difficulties that may arise intraoperatively. Decisions made and their analysis will incorporate perioperative clinical features including vital signs in waveform and epoch forms. Additional decision analysis will incorporate actionable features and anesthesia quality metrics.

We will use blood pressure (BP) as an example. BP is often a frequent and challenging intraoperative vital sign to manage. Variation of BP within anesthetic care is a common research topic as well-established treatment guidelines to aid in decision making either do not exist or change frequently with changing practice[9]. As a result, anesthetic care varies widely across institutions and providers. We plan to learn and evaluate improved intraoperative treatment policies for blood pressure management. Using contextual data collected prior to the operation and detailed intraoperative data collected throughout the delivery of anesthesia, we will learn a mapping from context and patient state to action/decision that aims to optimally control blood pressure (e.g., provider decisions to administer fluids, administer vasopressor medications). We will develop and use batch off-policy reinforcement learning techniques to leverage the large amount of variation in current practice. Once developed, this framework can generalize to other perioperative outcomes.

Features and Feature Preparation

Features used within this study will include existing and any newly created phenotypes collected from the MPOG database as they are presented from the individual institutions including age, sex, weight, height, ASA status, emergent status, procedure text, procedure duration, medications administered (including dose, route, unit of measure, time), peripheral and neuraxial blocks performed, line information (IV, arterial, central, etc.), Elixhauser comorbidities, operative complications, current procedure terminology (CPT) codes, intraoperative oral morphine equivalents (OME), among others. Features may be combined into new features if necessary in model construction as a means of feature engineering. Vital signs will be processed in the form of epochs and as waveforms as deemed necessary by modeling techniques.

We are focused on providing clinically relevant outcomes. As an example, oral morphine equivalent (OME) is a common method to represent opioid utilization in patients[10]. We will utilize phenotypes around perioperative outcomes including 30-day in-hospital mortality, MI, AKI, EBL, stroke, pulmonary edema, case-specific information such as fluid and blood product administration, and provider quality metrics through the Anesthesiology Performance Improvement and Reporting Exchange (ASPIRE) such as blood-pressure and temperature monitoring and control.

Text-based features will be processed prior to use and will include custom spell check and correction, removal of punctuation and non-alpha characters, case-sensitive normalization, potential removal of single characters and identified words (ex. English StopWords). Medical abbreviations will be expanded to long form to aid in processing and information retention. Procedure text will be taken as “actual procedure text” uploaded by the institution. In the situation where “actual procedure text” is missing, “scheduled procedure text” will be.

Features may be constructed from text in a form the models will be able to properly assess. This will include n-gram creation where the text will be cut into smaller sized text phrases of “n” length. The cutting is in order such that “I went for a run today” would be processed as the following variables using bigram (i.e., n=2): “I went”, “went for”, “for a”, “a run”, and “run today.” Word correlations will be constructed as well, in which terms that are found together with the same text can be used to differentiate and assign value. We will use co-occurrence measures such as GloVe: Global Vectors for

Word Representation, which make use of global word-word co-occurrence counts to create word vectors[11]. Vector word representations are compressed representations of overall variable arrays constructed above, such as n-gram[12]. By compressing the data computational time and model construct time are reduced. Furthermore, these representations break words into indices, which allows for larger amounts of data to be processed and in doing so simpler models can have more data and outperform more complex models with less data[13]. Vector representations also allow for words to be expressed in relation to each other as opposed to simple atomic symbols. In this manner words can convey more representative meanings the way audio and image data can. We will use several different types of vector word representations.

We will use term frequency-inverse document frequency (tf-idf) as a numerical statistic and as a weighting method when comparing variables. Using tf-idf in addition to frequencies allow us to discover specific variables that could potentially be rare, but extremely useful for differentiation. A high tf-idf could indicate an important variable where as a low tf-idf could indicate common, less important variables.

We will use generative statistical models to increase the depth of representation of the text as to learn meaningful abstractions from an external dataset. We will do so using topic models, such as latent Dirichlet allocation (LDA). LDA views a document as a mixture of a range of topics in which each word's creation is generated by one of the document's topics. LDA works in an unsupervised way and is able to discover hidden topics from a large collection of documents. We propose to collect a much richer dataset using the Pubmed database and learn a topic model from this rich dataset. We will apply the model to procedural text and extract the hidden topics discovered from the Pubmed abstract to enrich the representation of features with shorter text.

Various additional natural language processes and techniques not listed above may be used in processing and model creation.

Machine Learning Methods and Implementation

The variation in care provides an opportunity from a machine learning perspective. We will use machine learning methods to both characterize existing heterogeneity in the data and develop frameworks with the goal of leveraging the observed treatment heterogeneity to learn patient-specific protocols for the optimal delivery of anesthesia and improved patient outcomes.

We will create various machine learning models, each with their own strengths and weaknesses in order to compare them against one another in decision and predictive analyses. Planned prediction models include linear support vector machine (SVM), long short-term memory neural network (LSTM), extreme gradient boosting (Xgboost), random forest (RF), etc.

One specific focus will be using a reinforcement learning (RL) framework in which we will focus on the sequence of decisions a provider makes, and where each patient trajectory can be modeled as a contextual Markov decision process (MDP). For example, we will use intraoperative blood pressure (BP) variation and resulting treatments in this model. We will use contextual MDPs in a 'batch' setting in which we are limited to learning from historical data to help understand BP treatment variations and their appropriateness.

We will consider techniques designed to explicitly model the temporal dependence between decisions (e.g., recurrent neural networks). Using a vector-to-sequence architecture, one can learn a representation of the case using an encoder and then learn a mapping from this representation to a sequence of distributions over possible decisions using a recurrent decoder. Such an approach generalizes to decisions regarding other aspects of the plan without requiring one to explicitly define a notion of similarity between plans. Once we have learned mappings from patient/procedure to executed anesthesia plan, we will measure expected outcomes associated with each anesthesia plan cluster. We will label plans according to our primary outcome of interest: for example, blood pressure management, specifically intraoperative hypotension, defined as 20, 30, or 40% decrease below the patient's preoperative baseline (for either SBP or MAP). Beyond blood pressure management, we will consider post-operative outcomes including but not limited to 30-day in-hospital mortality and acute kidney injury (AKI).

In reinforcement learning (RL) setting, an agent interacts sequentially with an environment, soliciting a reward. This is commonly modelled using a Markov decision process (MDP) M = (S, A, P, R, γ) where S is the state space, A is the action space, P is a transition probability function from state and action to the next state, R is a stochastic function from Q: S x A $\longrightarrow$ R, and $\gamma \in [0,1]$ is the discount factor for the reward[14]. "Q" stands for the "quality" of a state-action combination. In an episodic setting, an agent observes the current state $s_t \in S$, chooses an action $a_t \in A$ and then transitions to s_{t+1} according to some probability distribution P^a_{st}. In addition, the agent receives an instantaneous reward $r_t := R(s_t, a_t)$. This process continues until reaching the end of the episode at time step T. An agent behaves according to some policy Π where $\Pi(a|s) := P^a_{st}$. In our setting, each operation corresponds to an episode. We will sample trajectories on a minute by-minute basis. We will consider a continuous state space, S consisting of both preoperative and intraoperative variables. At each time step, patient state will be represented by a concatenation of time-invariant pre-operative and time-varying intra-operative variables. We will initially focus our analysis on blood pressure management. Thus, we limit the decision/actions to those that affect blood pressure management (e.g., delivery of medications and fluids).

All models will be constructed and optimized using the computational languages Python and R. Traditional ML models will be trained over a 5-fold cross validation and evaluated.

Evaluation Plan

We will validate our findings in a held-out dataset consisting of a subset of data from each MPOG institution and of data from MPOG institutions not present in the training data (to measure generalizability across institutions). We will measure the robustness/stability of the learned clusters, and we will evaluate discriminative performance in terms of class-specific accuracy for the multi-class classification task. When evaluating the vector-to-sequence model, we will cluster the actual and generated sequences in a post-processing step and evaluate class specific accuracy as above. While recurrent neural networks (RNNs) are a natural choice when modeling trajectories, they often lack the degree of interpretability required in clinical care. We will compare the tradeoffs between generalizability and interpretability.

For the RL framework, we will evaluate models on a held-out test set of patient trajectories by measuring the difference between current policies and the learned "optimal" policy in terms of patient outcomes (i.e., observed cumulative reward). We will bin Q-values and measure the relationship between the average Q-value and average outcome within each bin. This will provide a mapping from Q-value to reward (e.g., survival). Given that we are in a batch setting in which the data are sampled offline, we cannot test our policies prospectively, so we will consider batch off-policy evaluation techniques (e.g., doubly-robust off-policy evaluation[15] and MAGIC estimator[16]). We will compare the average expected reward associated with the learned policy, versus the average expected reward obtained by the anesthesia provider's policy (learned via SARSA).

Anticipated Limitations

Our proposed research study has anticipated limitations. One limitation is in the accuracy of the institutional data. We will hand audit outlier statistics and cases for validity. To prevent possible misinterpretation of data and its dissemination we will involve individual institution champions to aid in information release. Another limitation is case variability between institutions. There is likely a wide range in the number of specific cases at each institution. We will limit representation using a minimal threshold for each query result.

Statistical Analysis

Primary analysis

Primary analysis will assess variation in anesthetic execution. Overall accuracy of ML models will be calculated across clusters[17]. Precision (positive predictive value), recall (true positive rate, sensitivity), F1 scores (binary classification measure of accuracy), and Gini coefficients will be calculated within model. Standard error calculated from a 5-fold cross validation method in which the study population will be split into 5 distinct, but similar groupings. Models will be trained on 4 of the groupings with one group held out for testing. This will yield 5 distinct train/test scenarios. The average, precision, recall, and standard deviations will be calculated as comparison within these groupings. An additional, entirely novel set of data will be withheld for generalized testing.

Missing Data

We will use exclusion or imputation, as appropriate.

Human Subjects' Risks and Data Protection

Data analysis will be restricted to aggregated group data. Data will be de-identified regarding individual hospitals, unless specifically discussed and approved by individual hospitals for their own internal use.

While hospital and hospital characteristics might be part of the analysis to account for practice variation, no individual hospitals will be identifiable in the results or publication, again discussed and approved by individual hospitals for their own internal use. Each group will contain a sufficient number of hospitals and cases to ensure de-identification or no group analysis will be performed. Again, data analysis and results will not allow identification of individual contributing sites.

Data will be maintained on a password protected secure MPOG server hosted. The study data will be accessible only to the statistical team directly involved with analyzing the data. The system fully meets all applicable HIPAA privacy and security rules. Access to the database and backups are strictly monitored according to need.

The final dataset will contain no patient or caregiver identifier. No protected health information or identifying information about individual patients, caregivers or hospitals will be part of a publication.

Impact

Overall, this work will improve our current understanding of variation in the delivery of anesthetic care and will inform future work on the development of methodologies to automatically map patients to 'optimal' plans. The techniques developed here can apply more broadly as healthcare providers routinely develop and execute longitudinal treatment plans.

The RL framework developed in this research can be used to study other intraoperative actions/decisions and rewards, in which prior context is available. Since we consider the preparation plan as input to the model, in future work, we can build on the proposed work, jointly optimizing both the preparation plan and intraoperative plan. This work will lay the groundwork for future research directions in studying the limitations of learning from batch time-series data collected in a clinical setting.

The resulting visualization tool could be used by individual hospitals and their providers to aid in continued medical education and facilitate the preparation of anesthetic care plans.

Finally, longitudinal care/treatment analysis can be utilized across all medical specialties. This project, if successful, will provide a framework for retrospective and predictive analysis within future related projects.

References

1. Macario, A., et al., *Which Clinical Anesthesia Outcomes Are Important to Avoid? The Perspective of Patients.* Anesthesia & Analgesia, 1999. **89**(3).
2. Macario, A., et al., *Which clinical anesthesia outcomes are both common and important to avoid? The perspective of a panel of expert anesthesiologists.* Anesth Analg, 1999. **88**(5): p. 1085-91.
3. Myles, P.S., et al., *Systematic review and consensus definitions for the Standardised Endpoints in Perioperative Medicine (StEP) initiative: patient comfort.* British Journal of Anaesthesia, 2018. **120**(4): p. 705-711.
4. Jammer, I., et al., *Standards for definitions and use of outcome measures for clinical effectiveness research in perioperative medicine: European Perioperative Clinical Outcome (EPCO) definitions: a statement from the ESA-ESICM joint taskforce on perioperative outcome measures.* Eur J Anaesthesiol, 2015. **32**(2): p. 88-105.
5. Gabriel, R.A., et al., *Practice Variations in Anesthetic Care and Its Effect on Clinical Outcomes for Primary Total Hip Arthroplasties.* The Journal of Arthroplasty, 2016. **31**(4): p. 918-922.
6. Smith, L.M., et al., *Neuraxial and Combined Neuraxial/General Anesthesia Compared to General Anesthesia for Major Truncal and Lower Limb Surgery: A Systematic Review and Meta-analysis.* Anesth Analg, 2017. **125**(6): p. 1931-1945.
7. Giese, I. and C. Kloek, *Learning From the Past and Looking Toward the Future in Cataract Surgery: How to Evaluate Innovations and Incorporate Into Clinical Practice.* International Ophthalmology Clinics, 2017. **57**(4): p. 11-19.
8. Quan, H., et al., *Coding algorithms for defining comorbidities in ICD-9-CM and ICD-10 administrative data.* Med Care, 2005. **43**(11): p. 1130-9.
9. Li, D., C. Bohringer, and H. Liu, *What is "normal" intraoperative blood pressure and do deviations from it really affect postoperative outcome?* J Biomed Res, 2017. **31**(2): p. 79-81.
10. Nielsen, S., et al., *A synthesis of oral morphine equivalents (OME) for opioid utilisation studies.* Pharmacoepidemiol Drug Saf, 2016. **25**(6): p. 733-7.
11. Jeffrey Pennington, R.S., and Christopher Manning, *Glove: Global vectors for word representation.* Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2014: p. 1532–1543.
12. Sampo Pyysalo, F.G., Hans Moen, Tapio Salakoski, and Sophia Ananiadou, *Distributional semantic resources for biomedical text mining.* Proceedings of the 5th International Symposium on Languages in Biology and Medicine, 2013: p. 39-44.
13. Mikolov, T., Chen, K., Corrado, G., & Dean, J., *Efficient estimation of word representations in vector space.* ICLR, 2013.
14. Sutton, R.S. and A.G. Barto, *Reinforcement Learning: An Introduction.* IEEE Transactions on Neural Networks, 1998. **16**: p. 285-286.
15. Dudík, M., J. Langford, and L. Li, *Doubly robust policy evaluation and learning.* arXiv preprint arXiv:1103.4601, 2011.
16. Thomas, P. and E. Brunskill. *Data-efficient off-policy policy evaluation for reinforcement learning*. in *International Conference on Machine Learning*. 2016.
17. Lee, H.C., et al., *Prediction of Bispectral Index during Target-controlled Infusion of Propofol and Remifentanil: A Deep Learning Approach.* Anesthesiology, 2018. **128**(3): p. 492-501.