Summary Doctoral Thesis – Bavo Kempen

ISBN 978 94 6165 403 8

D/2026/1869/24

Title: Modelling Cerebrovascular Autoregulation in Traumatic Brain Injury: Evaluation of Established Methods and Development of Novel Machine Learning Models

Chapter 1: General Introduction

Severe traumatic brain injury (TBI) is a life-threatening condition requiring specialized neurocritical care. Following the initial structural damage caused by primary injury, a cascade of secondary pathophysiological events may occur, and because these secondary insults can significantly worsen both short- and long-term outcomes, their prevention or early detection is central to current TBI management.

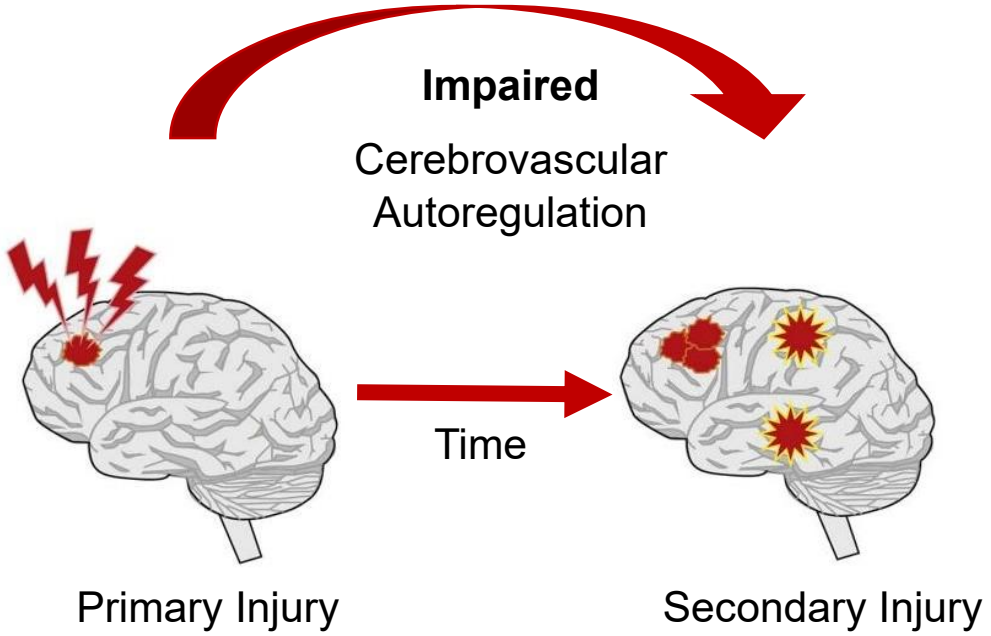

Figure 1. Transition from the primary injury to secondary injuries over time. The primary injury initiates pathophysiological processes which can develop into secondary injuries. A pivot mechanism that can be disturbed by the initial primary injury is cerebrovascular autoregulation.

Cerebrovascular autoregulation (CA) plays a pivotal role in preserving brain health by maintaining stable cerebral blood flow (CBF) despite fluctuating cerebral perfusion pressures (CPP). However, CA can be dynamically impaired following the primary injury, which can further compromise cerebral perfusion and oxygenation, exacerbating tissue injury. Maintaining an active CA by continuously monitoring CA

using neuromonitoring in TBI patients is therefore essential to prevent the transition from primary to secondary injuries.

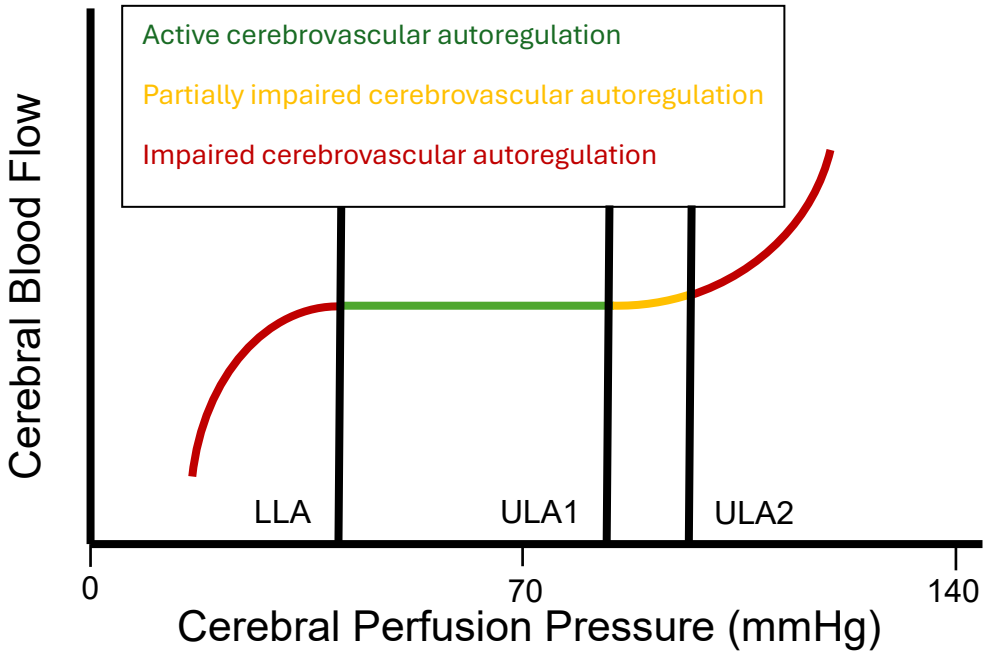

Figure 2. The static cerebrovascular autoregulation curve. The green plateau region reflects the cerebral perfusion pressures at which cerebral blood flow remains adequate and approximately constant due to the active cerebrovascular autoregulation mechanism. This plateau regions is edged by the lower limit of autoregulation (LLA) and first upper limit of autoregulation (ULA1). Between ULA1 and ULA2 autoregulation is partially impaired. Below the LLA and above the ULA2 autoregulation is considered completely impaired and pressure passive, i.e. cerebral blood flow will increase -or decrease with increases or decreases in cerebral perfusion pressure.

Chapter 2: Exploration of associations between vascular reactivity in acute traumatic brain injury and variation in selected candidate genes: A pilot study

Genetic variation may influence baseline CA capacity by altering the myogenic, metabolic, endothelial, and neurogenic pathways regulating CBF, including variants in calcium-signaling (DUSP5, TRPM2, TRPM4, TRPM8) and renin-angiotensin system genes (ACE, AGTR2, AGT, ENPEP, ATP2B1), endothelial mediators such as NOS isoforms (NOS1, NOS2, NOS3), adenosine receptors (ADORA1, ADORA2A), endothelin genes (EDN1-3, EDNRA, EDNRB), prostaglandin/eicosanoid-related genes (PHACTR1, PTGER2), and CYP4A11 involved in 20-HETE production.

The exploratory study outlined here found several physiologically plausible associations between these genes - both coding and common non-coding variants -

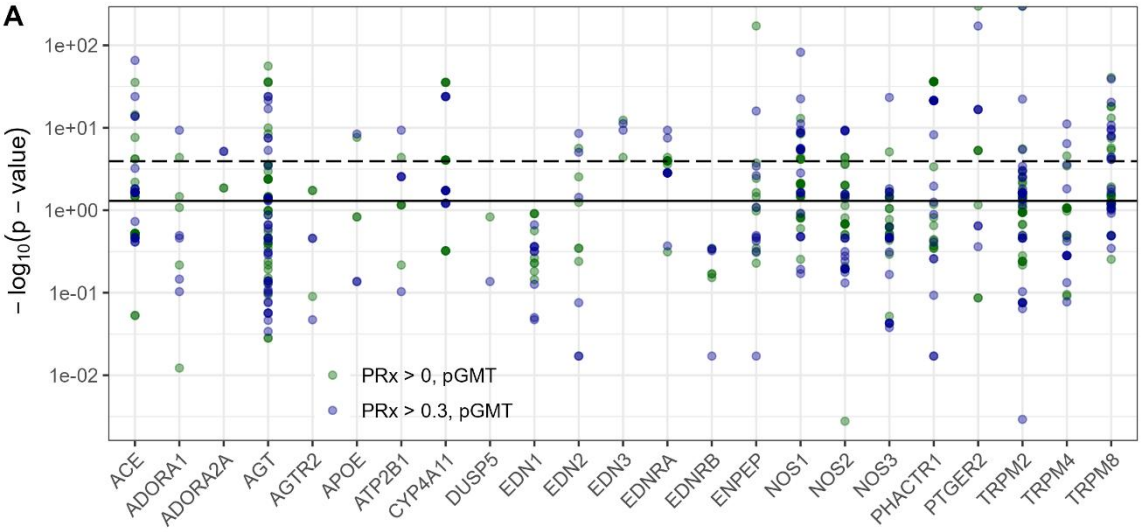

Figure 3. Association of individual common variants with cerebrovascular reactivity status, adjusted for CRASH covariates. Manhattan style plot which highlights significance for each of the variants. Each dot represents one variant, sorted by gene. The black horizontal line indicates p < 0.05, while the striped horizontal line reflects the Bonferroni-corrected significance level with 0.05/428=0.000117. Each variant (dot) above the horizontal line was found to be associated with cerebrovascular reactivity, as defined by the pressure-reactivity index (PRx).

and cerebrovascular reactivity (CVR) after TBI. Most variants showed the expected direction of effect with respect to what was previously found in the literature in relation to associated biological mechanisms. Most significant associations involved non-coding variants. Future studies should aim to gather larger cohorts of patients and improved CVR definitions, which will be essential to validate these findings and enable investigation of rare variants. Integrating genetic data into TBI research may help identify key molecular pathways affecting CA, enhance prediction models, and support development of targeted therapies for prevention of failing CA after TBI.

Chapter 3: Visualization of the intracranial pressure and time burden in childhood brain trauma: What we have learned one decade on with KidsBrainIT

&

Chapter 4: KidsBrainIT: Visualization of the impact of cerebral perfusion pressure insult intensity and duration on childhood brain trauma outcome

An important concept in TBI research is the dose-response concept, visualizing the relationship of sustained physiological events for a certain duration with outcome using important physiological parameters such as intracranial pressure (ICP) and CPP. Here this concept was explored using a novel prospectively collected pediatric TBI dataset from the multi-center, multi-national KidsBrainIT consortium (n=104). Minute-by-minute ICP and CPP time series were transformed into intensity-duration episodes and linked to 6-month outcome using dose-response visualizations.

The analysis validated the previously published pediatric ICP dose-response plot, demonstrating that higher ICP levels were tolerated only for short periods along an exponential transition curve and that any ICP exceeding 20 mmHg was associated with poorer outcome.

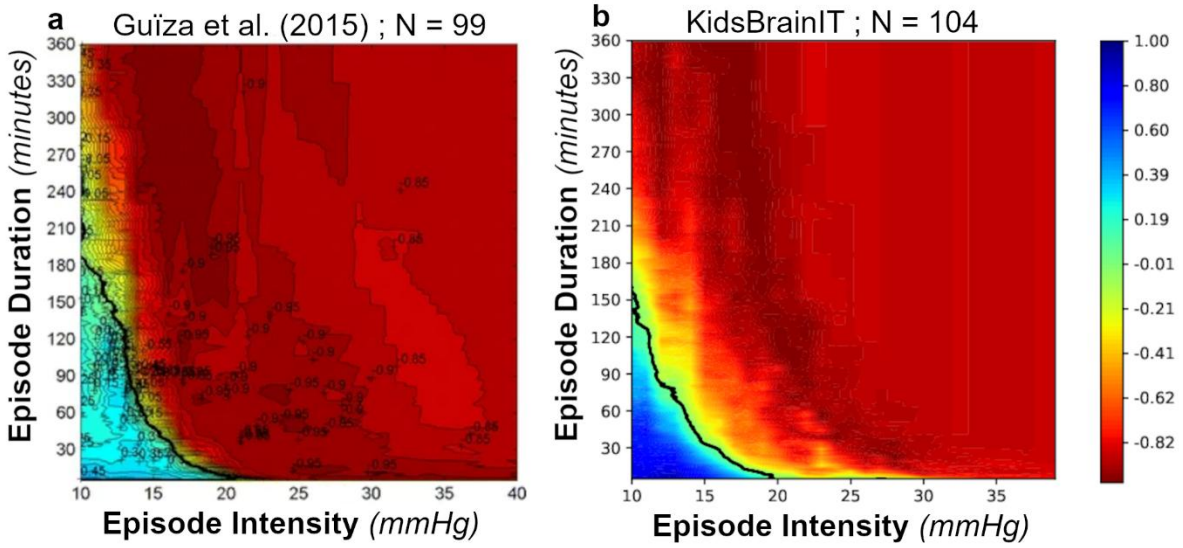

Figure 4. Visualization of the univariate correlation, expressed in color code, between the Glasgow Outcome Score at 6 months (GOS-6) and the average number of intracranial pressure (ICP) events per GOS-6 category, n = 104. Each dot in the graph represents an intracranial pressure (ICP) event. A blue color depicts a positive association between the number of such events and GOS, while a red color depicts a negative association. The color bar on the right-hand side of the figure displays the continuous transition of the process from -1 to 1 with its respective color code. The transition curve (ie., zero correlation line) is highlighted in black. A) The Original ICP dose-response visualization (produced in Matlab) published in 2015. B) The KidsBrainIT ICP dose-response validation plot (produced in Python) based on our novel dataset.

Moreover, the first pediatric CPP dose-response plots were generated, revealing a similar exponential separation between good and poor outcomes as in adult TBI CPP dose-response plots.

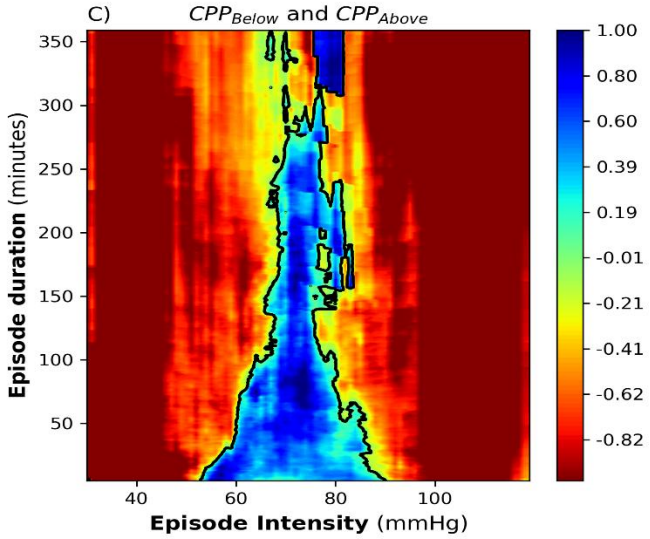

Figure 5. Visualization of the univariate correlation between the Glasgow Outcome Score (GOS) and the average amount of cerebral perfusion pressure (CPP) events per GOS category, N = 104. CPP events are denoted by the colored dots in the graph, defined by an intensity threshold (x-axis) maintained for a certain duration (y-axis). A negative association of a CPP event with GOS is color coded as dark red, which gradually transitions toward dark blue for positive correlations. The zero correlation line, that is, the transition curve, is highlighted in black.

Together, these findings validate pediatric ICP dose-response relationships, establish the first pediatric CPP dose-response thresholds, and emphasize the need to reconsider current treatment targets in childhood TBI. However, unlike in adult TBI cohorts, "normal" ICP and CPP values in children vary across age groups because both baseline arterial blood pressure (ABP) and ICP change with development, with CPP = ABP - ICP. Hence, future studies should increase pediatric datasets to further investigate the effect of age by stratifying along age bands.

Chapter 5: Optimal cerebral perfusion pressure estimation in adult and pediatric TBI patients: insights from a multiverse analysis

Optimal CPP (CPPopt) represents the CPP at which CVR - as surrogate for CA - is most effective and CBF remains stable.

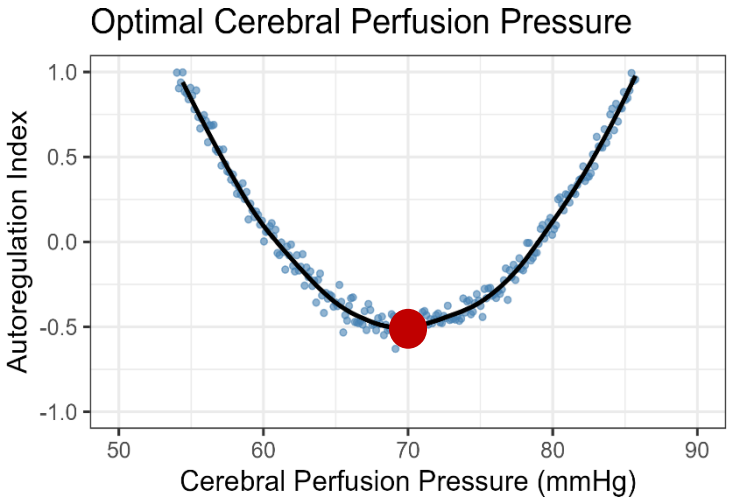

Figure 6. Theoretical representation of optimal cerebral perfusion pressure (CPPopt). The autoregulation index - for instance the pressure-reactivity index (PRx) - plotted against cerebral perfusion pressure (CPP). CPPopt reflects the CPP value at which the autoregulation index was minimal.

Although numerous algorithms have been proposed to compute CPPopt and linked to functional outcomes, studies vary widely in methodological choices, including inclusion criteria, CPPopt variable definitions, preprocessing, autoregulation indices, calculation methods, and statistical modeling. Here, a comprehensive literature review was performed followed by a large-scale multiverse analysis on a novel high-resolution adult (n=57) and low-resolution pediatric (n=202) severe TBI datasets. We systematically applied 7,497,520,044 plausible analytical combinations to the adult cohort and a subset of these combinations on the pediatric and/or adult cohort. More specifically, for the low-resolution dataset, only LAx, L-PRx, and UL-PRx could be computed to derive CPPopt. Of 34 studies meeting inclusion criteria, 41.45% of reported CPPopt-derived summary measures showed significant association with the Glasgow outcome score (GOS).

The multiverse analysis revealed a near-uniform p-value distribution for CPPopt-derived predictors in relation to GOS at 6 months, with only 7.73\% with a p-value below 0.05. Significant results arose mainly from 3 out of 10 statistical tests, with varying influence of other methodological choices such as exclusion criteria, autoregulation metric, or summarization strategy of CPPopt. These findings suggest that the strength of associations of CPPopt-derived variables with outcome appear to be highly sensitive to the choice of parameters and tests used, and highlight the need for methodological standardization.

Chapter 6: Development of an active cerebrovascular autoregulation model using representation learning: a proof of concept study with experimental data

&

Chapter 7: Representation learning for cerebrovascular autoregulation: A model study with experimental data classification

Reliable real-time CA monitoring in intensive care units (ICU) remains challenging. Using a porcine cranial window dataset (n=20) with concurrent ABP and ICP recordings, a deep representation learning model was developed to learn an active CA representation from 300-second signal segments.

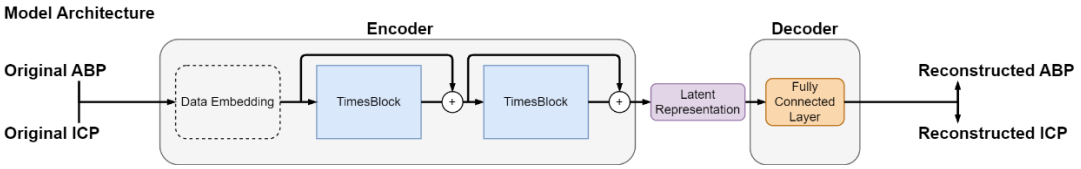

Figure 7. Model architecture. The architecture consisted of a sequentially connected encoder and decoder neural network which was input with 300 seconds of concurrent arterial blood pressure (ABP) and intracranial pressure (ICP) segments sampled at 20 Hertz (Hz). Model is optimized to reconstruct the original ABP and ICP input segments utilizing mean squared error (MSE) loss. The encoder network comprised of a data embedding layer and two residually connected TimesBlocks38, each selective for the top eight frequencies. Inception blocks consisted of 6 kernels with 0.1 dropout. Activation functions used were Gaussian Error Linear Units. Latent representation comprised of eight dimensions. A simple linear feedforward layer reconstructs the input data from the latent representation in the decoder network. The model comprised of 73 522 trainable parameters and was optimized on a single NVIDIA V100 GPU for 100 epochs with a batch size of 32 and the adaptive moment estimation (Adam) optimizer with a starting learning rate of 0.0001. Model training was halted when the validation loss did not decrease for 5 consecutive epochs.

The model achieved high reconstruction accuracy for active CA segments, while reconstruction errors increased steadily as CBF deviated from baseline, reflecting failing CA, with distinct frequency components driving errors during inactive states.

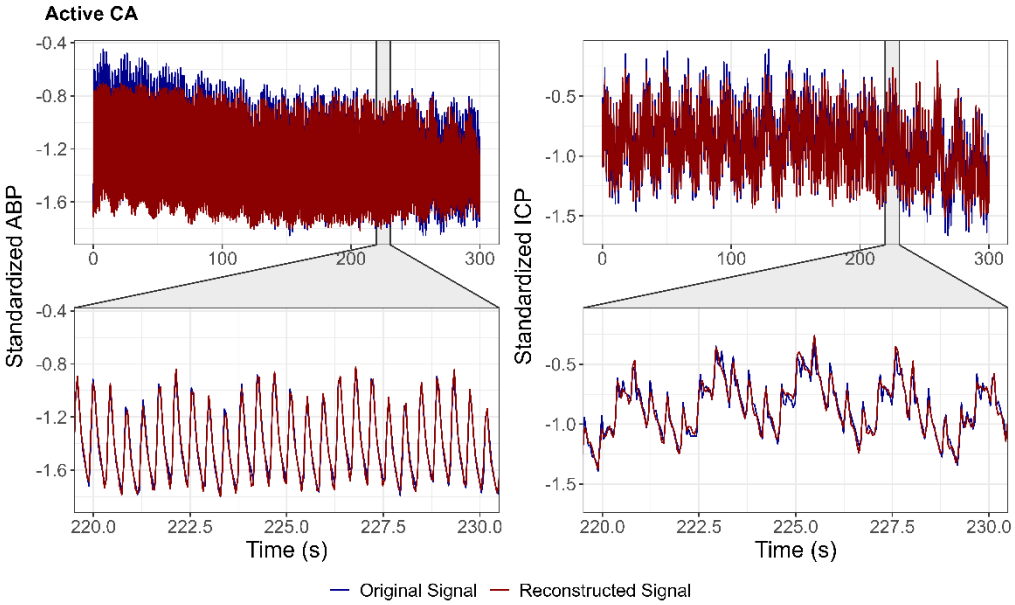

Figure 8. Original arterial blood pressure (ABP) and intracranial pressure (ICP) input traces (blue) and respective reconstructions (red) from the test data set active CA. The left and right figures per row concern concurrent standardized ABP and ICP traces, respectively. Each ABP and ICP trace is provided with a 10 second zoom window derived from the 300 seconds input - reconstruction segments.

Moreover, the observed interaction between increases in error for ABP and ICP and CA state suggests that the model can capture differential CA behavior with increasing versus decreasing CPP.

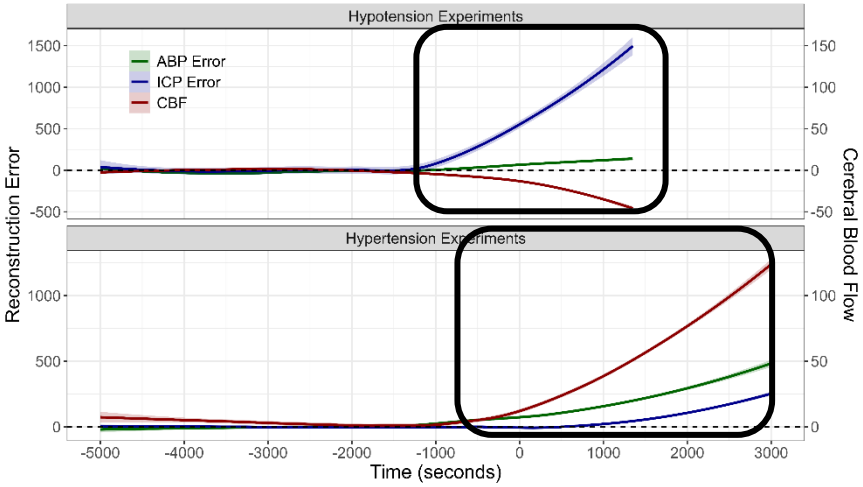

Figure 9. Change of arterial blood pressure (ABP) and intracranial pressure (ICP) reconstruction error and change of cerebral blood flow (CBF) from baseline over time visualized per piglet experimental condition. LiCA, HPaCA and HiCA reflect the Low inactive CA, hypertensive High Partially active CA, and High inactive CA, respectively. The onset of LiCA and HPaCA was centred as time point 0 to combine the data of all piglets within the experimental condition. Data points were smoothed using a locally weighted non parametric regression with an α of 0.75 and visualized with their 95% confidence intervals.

Classifiers built on reconstruction error or latent features substantially outperformed the conventional pressure-reactivity index (PRx) in classifying CA state, improving precision from 0.14 to 0.77 and recall from 0.62 to 0.87. This approach further suggested that relevant CA state information resides in both lower and higher frequency components of ABP and ICP, and that leveraging the full complexity of these signals can provide more accurate dynamic monitoring of CA than traditional correlation-based methods.

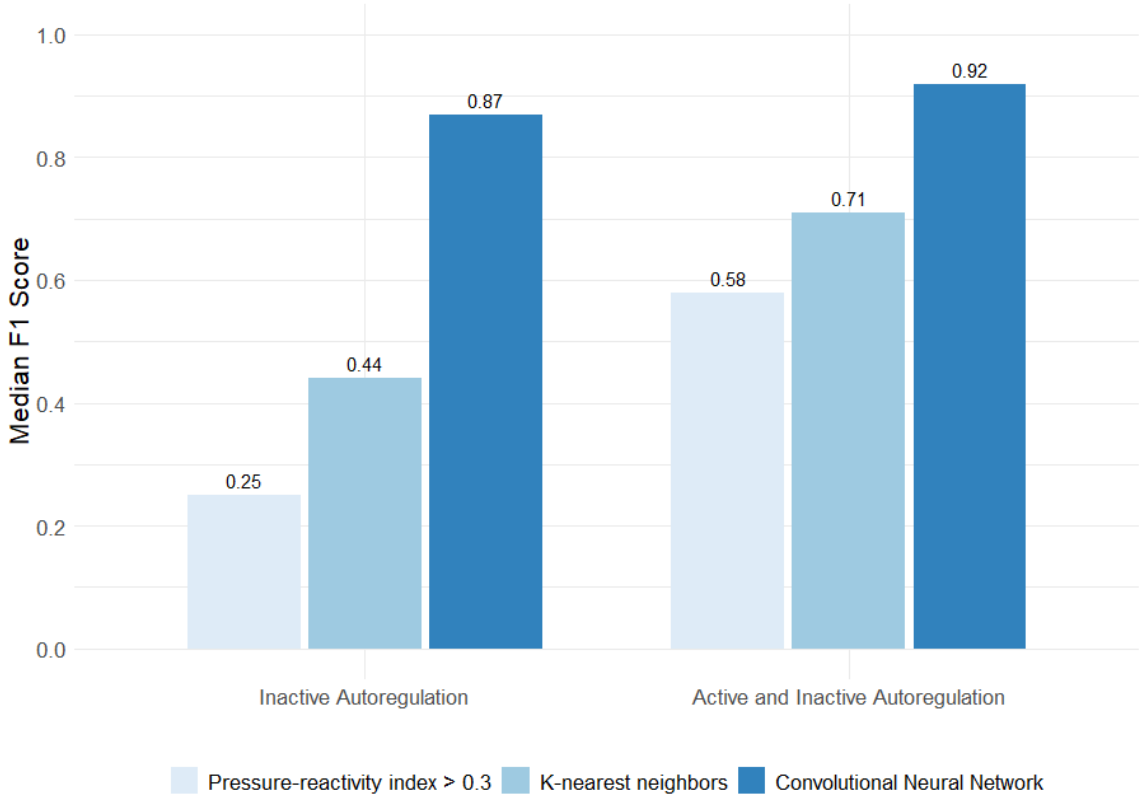

Figure 10. Classification performance of the pressure reactivity index versus representation learning – based features derived from the novel model. The F1 score is reflects how precise and complete the models could predict inactive autoregulation (left) or active and inactive autoregulation (right), with 0 being the lowest score and 1 the highest score. The baseline performance was shown for the pressure reactivity index (light blue). Reconstruction error – based features were used with a K-nearest neighbors classifier, while latent based features were used with a novel convolutional neural network classifier.

Chapter 8

Lassen's classic triphasic CA curve defined lower and upper limits of autoregulation (LLA and ULA), however, recent evidence suggests a more complex pressure-flow relationship.

Building on the concept of gradual CA failure derived from the quadriphasic model (cf. Figure 2), a novel nonlinear continuous static autoregulation metric (cSARm) was introduced.

This method generates a pentaphasic CA curve, capturing gradual transitions between active and impaired CA at both the hypotensive and hypertensive side of the CA curve and provides refined estimates of gradually failing CA between the respective autoregulatory limits. Using a porcine cranial window dataset (n=20) with pial arteriole diameter measurements, cSARm determined multiple LLA and ULA breakpoints (LLA2, LLA1, ULA1, ULA2) based on mathematical features of the first and second derivatives of a nonlinear fit, while mapping the progression from fully active to fully inactive CA onto a continuous 0 to 1 score.

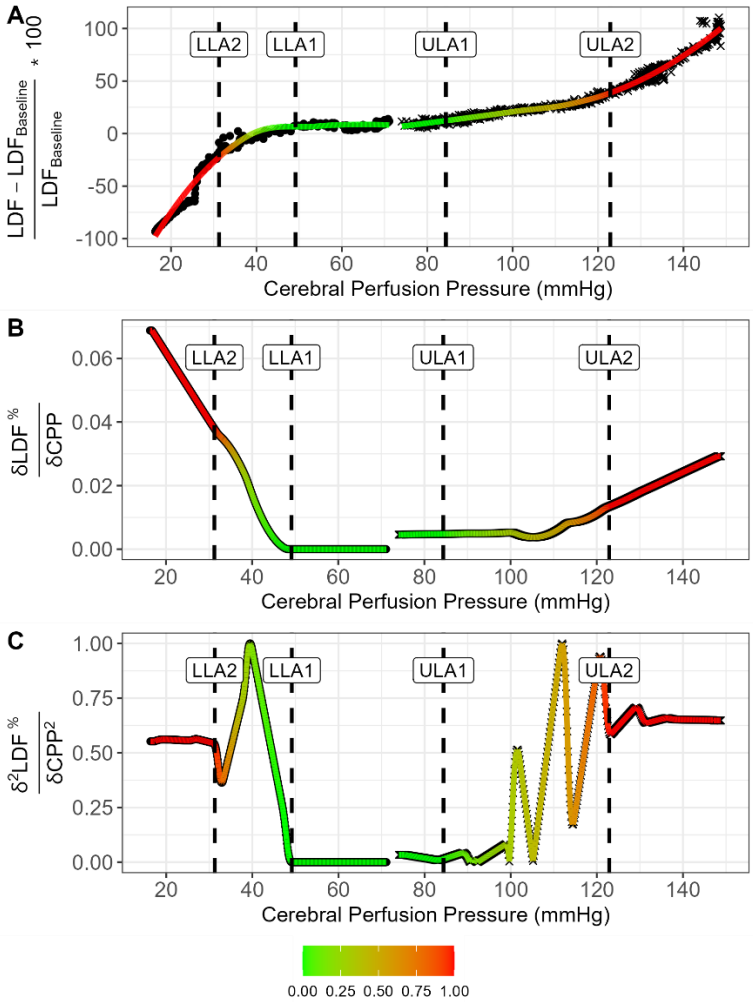

Figure 11. Example of a hypotensive -and hypertensive piglet and the continuous cerebrovascular autoregulation metric (cSARm). A) Data points of a hypotensive (dots) and hypertensive (crosses) piglet overlayed by their fitted locally weighted polynomial regression (LOESS) of degree 2 with a smoothing factor α of 0.75. B) The first derivatives computed on the LOESS regressions displayed in (A). C) Second derivatives computed on the LOESS regressions displayed in (A). All curves were colored using the novel cSARm algorithm.

Both quadriphasic and cSARm breakpoints were successfully computed, with cSARm generally estimating higher lower limits (LLA1) and slightly lower upper limits (ULA1) in absolute CPP, and provided a smoother transition from fully active to fully inactive CA. Namely, LLA1 reflected the onset of gradual CBF decline, while ULA1 captured early hypertensive responses in smaller arterioles, offering greater granularity than the quadriphasic model. Overall, cSARm extends the quadriphasic framework into a pentaphasic model that more continuously characterizes partially impaired CA zones under both hypotensive and hypertensive conditions.

Chapter 9: Predicting cerebral blood flow from standard intensive care unit variables

Adequate CBF is critical for brain health, yet continuous bedside monitoring in the ICU is limited by surrogate measurement techniques, since no reliable and robust CBF probe exists to date. Therefore the feasibility of using Temporal Fusion Transformer (TFT) models - trained on a porcine cranial window dataset (n = 60) - to predict continuous CBF from routinely collected ICU time series was investigated, offering a data-driven approach to real-time monitoring.

Raw signals were preprocessed and downsampled to 0.1 Hz, and four modeling conditions were tested to forecast 60-second CBF time series, with two conditions incorporating historical CBF data. Hyperparameter optimization using a Bayesian Tree-structured Parzen Estimator was performed, followed by full model training. TFT models without historical CBF data failed to capture meaningful temporal dependencies, whereas autoregressive models including past CBF substantially improved predictive performance, enabling accurate short-term forecasts.

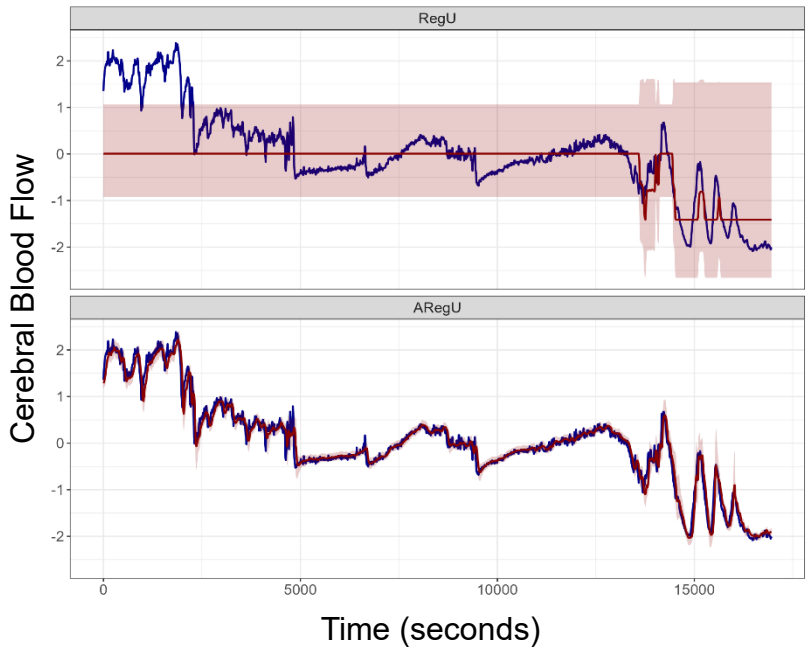

Figure 12. Real cerebral blood flow (blue) versus prediction (red) for a single exemplary subject. The top panel (RegU) did not use historical cerebral blood flow data to predict future cerebral blood flow, while the bottom panel (ARegU) did use historical cerebral flood flow data in addition to routinely collected ICU time series.

Attention analyses suggested that poor performance in non-autoregressive setups was driven by misallocated attention and standardization mismatches of the CBF signal. These findings indicate that autoregressive TFT models can support short-term CBF prediction in neurocritical care, particularly when brief calibration periods are available.

Chapter 10: General Discussion

To conclude, ICUs have developed rapidly in data-rich environments which can be exploited to improve outcome following severe TBI. Monitoring CA and ensuring an adequate CBF play a pivotal role in patient-specific care to help preserve brain health and prevent secondary injuries.

The results from chapter 6 – 9 may form the basis of a novel CA monitor depicted in Figure 13. Here, we outline how our developed models could aide clinicians and may result in a state-of-the-art monitor. The goal of the system in Figure 13 is to give doctors two clear messages: whether the brain's blood-flow control (cerebral autoregulation, CA) is working, and why it might be failing. Figure 13A shows the method developed in Chapters 6 and 7, where a model tracks blood pressure and brain pressure and flags dangerous patterns that signal a harmful cycle before

permanent brain damage occurs. Chapters 8 and 9 add key pieces by estimating how impaired CA is and predicting future blood-flow trends, which helps explain the cause and direction of worsening or recovery. Figure 13B represents the long-term goal: showing a personalized, real-time curve of a patient's brain blood flow and where they currently sit on it. Together, the chapters move from detecting problems, to understanding them, to eventually guiding doctors with a live, patient-specific picture of brain circulation.

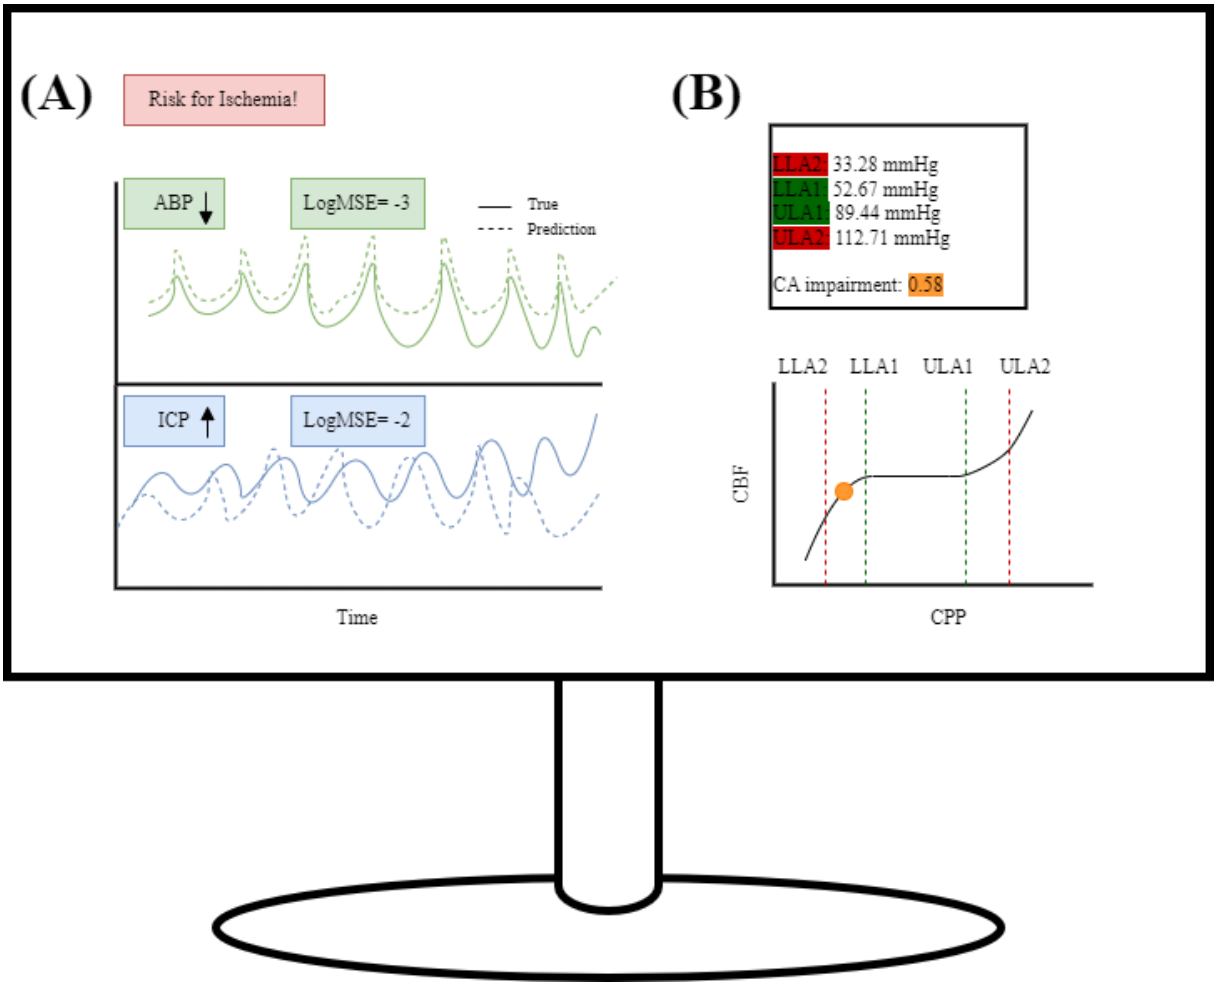

Figure 13. Envisaged cerebrovascular autoregulation monitor. (A) Shows the true arterial blood pressure (ABP) and intracranial pressure (ICP) and their reconstructions, trend progression (arrows), and reconstruction errors. (B) Shows the simulated CBF-CPP curve, annotated with the dynamically computed limits of autoregulation and the continuous static autoregulation metric (cSARm) score of CA impairment. Moreover, the patient's position is indicated on the CBF - CPP curve.

In summary, in this thesis significant contributions were made to TBI research: 1) advance the basic knowledge of CA (chapter 2), 2) evaluate existing popular methods in novel clinical datasets (chapter 3 – 5), 3) and develop novel methods to model CA, the static CA curve, and CBF, pioneering advanced modelling strategies (chapter 6 – 9).