

MORPHODYNAMIC EVOLUTION OF BRAIDED RIVERS: A MARKOV CHAIN APPROACH

THIS IS A TEMPORARY TITLE PAGE
It will be replaced for the final print by a version
provided by the service académique.

Thèse n. 1234 2025
présenté le 24 septembre 2025
à la Faculté de l'environnement naturel, architectural et con-
struit
Laboratoire d'hydraulique environnementale
Programme doctoral en mécanique
pour l'obtention du grade de Docteur ès Sciences
par

Clemente Mauricio GOTELLI ALVIAL

acceptée sur proposition du jury:

Prof F. GALLAIRE, président du jury
Prof C. ANCEY, directeur de thèse
Prof M. REDOLFI, rapporteur
Dr F. METTRA, rapporteur
Prof M. MUSA, rapporteur

Lausanne, EPFL, 2025

EPFL

*A Candei,
el sol que ilumina mis días.*

Acknowledgements

I would like to begin this section by expressing my deepest gratitude to the “grand jefecito,” Christophe. Thank you for believing in me from the very beginning, after only having seen my CV. Thank you for being there to encourage and support me whenever I doubted my ideas. I am especially grateful for the long—sometimes seemingly endless—conversations we had every time you passed by my office. Because of these chats, you often had to leave without talking to my colleagues, and then you blamed me, saying that I talk too much—when in fact, it was you who couldn’t stop talking. Through those conversations, I learned so much about history, politics, and philosophy. They enriched my mind immensely. From my experience as your student, I take with me countless lessons and a profound admiration for your capacity for reflection and intelligence. However, what I will cherish the most is the bond we built over the years. Your passion for knowledge, good food, and quality family time are also lessons I deeply value. Thank you for everything.

I am also very grateful to the members of the LHE. First, to Tomás—my compatriot—who gave me the opportunity to come here without knowing me beforehand. Thank you for your guidance at the beginning, for helping me navigate life in Switzerland and in the lab, and for your continued advice about academic life. My thanks also go to Barbara, the unshakable rock of the LHE, who showed me the essence of Switzerland and with whom I was able to practice and learn French. I could not forget Bob, the person who makes sure everything runs smoothly in the lab. Without your help, this thesis would not have been possible—thank you sincerely.

I cannot fail to mention my “communist comrades,” Yanan and Sofi, with whom I shared most of my PhD journey. Thanks to you, life in the lab was incredibly fun. I learned a lot from you—about your cultures, traditions, and perspectives. I’m truly happy to have spent more than two years working side by side. I wish you all the best for the future, and I hope that in a decade we’ll meet again in China, invited by Professor Chen ;) I would also like to thank the lab’s mascots, Pedro and Aline. You brought a different, lighter, and more cheerful “Ancey touch” to the lab. Thank you for the good times.

To my Politbüro comrades, Antoine and Tristaños: thank you from the bottom of my heart. You were the ones who welcomed us at the beginning and showed us the warm side of Switzerland, often hidden to those who come from far away. Through countless cafés, beers, political discussions, and lessons of GeoGuessr, I learned so much about your culture, your roots, and above all, your kind hearts. These memories are treasures I will carry with me forever.

I would also like to thank the rest of our Swiss friends, who have been a fundamental part of our last years in Switzerland. In particular, the HEC group—Camille, Jordan, Danika, and Matthias. You welcomed us with open arms and made us feel part of your circle from the very beginning. Through your warmth, generosity, and friendship, you showed us a side of Swiss life that we might never have discovered on our own. Each of you has contributed, in your own way, to making these years truly unforgettable. Thank you for all the laughter, the shared moments, and for making us feel that even far from our families, we were never alone.

Ahora en español quisiera agradecer a mi familia, tanto sanguínea como política, por su apoyo y cariño a la distancia. Las largas conversaciones con cada uno de ustedes, en las que hablábamos de todo, nos ayudaron a sobrellevar los momentos difíciles en que la nostalgia nos invadía. Del mismo modo, cuando el cansancio aparecía, la cercanía de las vacaciones en las que volveríamos a vernos nos daba el impulso necesario para seguir trabajando con fuerza. Sin ustedes, su compañía constante y su preocupación genuina, esta tesis tampoco habría sido posible.

Finalmente, quiero agradecer a la persona a quien dedico este trabajo de cinco años: mi Candei. Cuando te conocí, cambiaste mi vida para siempre. Con tu luz iluminas mis días y con tu cariño incondicional me das la energía que necesito para seguir adelante. Sin ti, sin tu apoyo, tu rigor y tu valentía, no habría llegado hasta aquí. Gracias por lanzarte conmigo a esta aventura hacia lo desconocido, por dejar atrás la comodidad de nuestro entorno y construir, juntos, un camino lleno de desafíos y aprendizajes. Mirando hacia atrás, puedo decir que estos cinco años pasaron volando, a pesar de todo lo vivido; mirando hacia el futuro, nos imagino compartiendo muchas más aventuras, siempre de la mano. Te amo, Candei.

Lausanne, October 24, 2025

Clemente Gotelli

Abstract

Braided rivers switch between quiet and active periods of bedload transport while their planform changes quickly. This makes both simple descriptive indices and heavy morphodynamic models hard to use in practice. This thesis offers a practical middle path: it treats channel change as a sequence of morphological states read from planform images and models how the river switches between them with a continuous-time Markov model. The aim is to turn images into probabilistic forecasts of sediment transport with stated uncertainty.

From flume imagery, binary water masks are compared with two complementary measures that capture edge movement and area overlap. We reduce these pairwise differences and cluster the images to obtain a small, readable set of recurring states, ranging from narrow and simple to wide and partitioned. The time the river stays in a given state is well described by an exponential law, which allows us to estimate transition rates and jump probabilities for the Markov model.

The resulting ensemble recovers means, variances, extremes, and main time scales, and it shows clear morphological control of transport, including a negative link with wetted width and with a braiding index. Splitting the variance indicates that differences between states explain a meaningful part of the instant variability, while the rest arises within states. The learned states and transitions remain stable across independent runs.

For image-only cases, two variants extend the method: one that preserves the long-term mean, and another that uses stream power from images to scale state-wise means. Weighting by how long the river stays in each state keeps the overall mean accurate when direct bedload data are missing. Overall, the framework provides a clear path from images to forecasts, explains intermittency as switching among states with different export capacity, and enables practical predictions with quantified uncertainty.

Key words: braided rivers; sediment transport; morphodynamics; continuous-time Markov chains; planform imagery.

Résumé

Les rivières tressées alternent entre des périodes calmes et actives de transport de fond, tandis que leur forme en plan change rapidement. Cela rend peu pratiques à la fois les indices simples et les modèles morphodynamiques lourds. Cette thèse propose une voie intermédiaire pratique : elle décrit l'évolution du chenal comme une suite d'états morphologiques lus dans des images en plan, et modélise les passages entre ces états avec un modèle de Markov en temps continu. L'objectif est de transformer des images en prévisions probabilistes du transport de sédiments, avec une incertitude explicitée.

À partir d'images de canal, des masques d'eau binaires sont comparés avec deux mesures complémentaires qui saisissent le déplacement des bords et le recouvrement des surfaces. Les différences par paires sont réduites en dimension puis regroupées pour obtenir un petit ensemble d'états récurrents, lisibles, allant de configurations étroites et simples à larges et partitionnées. Le temps passé dans chaque état suit bien une loi exponentielle, ce qui permet d'estimer les taux de transition et les probabilités de saut du modèle de Markov.

Sur cette base, des simulations aléatoires simples avec un bootstrap selon l'état génèrent des séries synthétiques de transport. Une étape de filtrage retient seulement les séries qui reproduisent la distribution observée et les rythmes temporels. L'ensemble obtenu retrouve les moyennes, variances, extrêmes et principales échelles de temps, et il met en évidence un contrôle morphologique net du transport, y compris un lien négatif avec la largeur mouillée et avec un indice de tressage. Une décomposition de la variance montre qu'une part notable de la variabilité instantanée vient des différences entre états, le reste provenant de la variabilité à l'intérieur des états. Les états appris et les transitions restent stables sur des séries indépendantes.

Pour les cas avec seulement des images, deux variantes étendent la méthode : l'une qui préserve la moyenne à long terme, et une autre qui utilise la puissance du courant déduite des images pour mettre à l'échelle les moyennes par état. Un pondérage par le temps de séjour dans chaque état maintient une moyenne globale fidèle quand on ne dispose pas de mesures directes du charriage. Au final, le cadre propose une voie claire des images vers la prévision, explique l'intermittence comme des passages entre états à capacité d'export différente, et permet des prédictions pratiques avec incertitude quantifiée.

Mots clefs : rivières en tresses ; transport de fond ; morphodynamique ; chaînes de Markov en temps continu ; images en plan.

Resumen

Los ríos trenzados alternan períodos tranquilos y activos de transporte de fondo mientras su forma en planta cambia con rapidez. Esto hace poco prácticos tanto los índices simples como los modelos morfodinámicos pesados. Esta tesis propone una vía intermedia: trata el cambio del canal como una secuencia de estados morfológicos leídos en imágenes en planta y modela los pasos entre esos estados con un modelo de Markov en tiempo continuo. El objetivo es convertir imágenes en pronósticos probabilísticos del transporte de sedimentos, con la incertidumbre claramente indicada.

A partir de un conjunto de imágenes de un canal experimental, se construyen máscaras binarias de agua y se comparan con dos medidas complementarias que capturan el movimiento de los bordes y el solape de las áreas. Las diferencias por pares se reducen en dimensión y luego se agrupan para obtener un conjunto pequeño y legible de estados recurrentes, que van desde patrones estrechos y simples hasta anchos y particionados. El tiempo que el sistema permanece en cada estado se describe bien con una ley exponencial, lo que permite estimar tasas de transición y probabilidades de salto para el modelo de Markov.

Sobre esta base, simulaciones aleatorias sencillas con un bootstrap por estado generan series sintéticas de transporte. Un paso de filtrado conserva solo las series que reproducen la distribución observada y los ritmos temporales. El conjunto resultante recupera medias, varianzas, extremos y las principales escalas de tiempo, y muestra un control morfológico claro del transporte, incluido una correlación negativa con el ancho mojado y con un índice de trenzado. Una descomposición de la varianza indica que las diferencias entre estados explican una parte apreciable de la variabilidad instantánea, mientras que el resto surge dentro de los estados. Los estados aprendidos y las transiciones se mantienen estables en corridas independientes.

Para casos en los que solo hay imágenes, dos variantes amplían el método: una que preserva la media de largo plazo y otra que usa la potencia de la corriente estimada a partir de imágenes para escalar las medias por estado. Un ponderado por el tiempo de residencia en cada estado mantiene la media global cuando no hay datos directos de transporte. En conjunto, el marco ofrece un camino claro desde las imágenes hasta la previsión, explica la intermitencia como cambios entre estados con distinta capacidad de exporte y permite predicciones prácticas con incertidumbre cuantificada.

Palabras clave: ríos trenzados; transporte de fondo; morfodinámica; cadenas de Markov en tiempo continuo; imágenes en planta.

Contents

Acknowledgements	i
Abstract (English/Français/Español)	iii
List of figures	xiii
List of tables	xv
1 Introduction	1
1.1 What is a braided river?	2
1.2 Why morphology forecasting matters	3
1.3 Complexity and variability	6
1.4 Morphology as a long-standing and practical lens	7
1.5 An informational viewpoint on morphology	8
1.6 Braided rivers as a state–transition process	9
1.7 The challenge of forecasting bedload	12
1.8 Knowledge gap, aim, and research questions	15
1.9 Roadmap of the thesis	17
2 Materials and Methods	19
2.1 Overview of the Process	20
2.2 Flume Characteristics	21
2.3 Image Processing	25
2.4 Distance Matrix Construction	26
2.4.1 Modified Hausdorff Distance	27
2.4.2 Dice Index	28
2.4.3 Combined Matrix	29
2.5 Dimensionality Reduction	30
2.5.1 What is dimensionality reduction?	30
2.5.2 Principal Component Analysis (PCA)	31
2.5.3 Multidimensional Scaling (MDS)	31
2.5.4 t-Distributed Stochastic Neighbor Embedding (t-SNE)	32
2.5.5 Uniform Manifold Approximation and Projection (UMAP)	33
2.6 Density-Based Clustering	33

2.6.1	DBSCAN: Density-Based Spatial Clustering of Applications with Noise	33
2.6.2	HDBSCAN: Hierarchical Density-Based Spatial Clustering of Applications with Noise	34
2.6.3	Metrics for Cluster Quality Assessment	35
2.6.4	Metrics for Cluster Morphological Description	36
2.7	Probabilistic modeling with Markov chains and Monte Carlo	37
2.7.1	Markov Chains	37
2.7.2	Continuous-Time Dynamics: CTMC	39
2.7.3	Bootstrap Method	40
2.7.4	Monte Carlo Simulations	41
2.7.5	Bayesian Inference	42
2.7.6	Decision-tree surrogate as a conservative prefilter	42
2.7.7	Cross-experiment comparison of transition matrices	43
2.7.8	Temporal Diagnostics: Autocorrelation and Power Spectrum	43
2.7.9	Extreme-Value Validation: Return-Period Curves	44
2.7.10	Variance decomposition and the fraction explained by state switching	45
3	Braided Rivers Morphology as Markov Chains	47
3.1	Distance Matrix	48
3.2	Dimensionality Reduction	51
3.2.1	Principal Component Analysis (PCA)	51
3.2.2	Non-linear Dimensionality Reduction	52
3.3	Clustering	57
3.3.1	Characterization of Clusters	61
3.4	Markov Model Construction	70
3.4.1	Insights from the Markov Representation of the System	74
3.4.2	Transport changes associated with state transitions	76
3.4.3	Export efficiency, accessibility, and uncertainty across regimes	80
3.4.4	Cross-Validation of the Markov Model	83
3.5	Summary of Key Points	87
4	Markov-Based Monte Carlo and Bootstrap Analysis	89
4.1	From the Transition Matrix to Synthetic Sediment–Transport Series	90
4.1.1	Monte Carlo–Bootstrap Framework	90
4.1.2	Bayesian Screening and Decision–Tree Extraction	92
4.1.3	Validation of the Accepted Time Series	94
4.2	Variance Decomposition and the Role of Cluster Granularity	99
4.2.1	Fraction of Variance Explained	99
4.2.2	Sensitivity Analysis to Cluster Granularity	101
4.3	Hybrid Semi-Deterministic Model	104
4.4	Summary of Key Points	109
5	Conclusions and Future Work	111
5.1	Conclusions	111
5.2	Future work	112
Appendix A		117

CONTENTS

Appendix B	123
Appendix C	127
Appendix D	133
Appendix E	135
Bibliography	137
Curriculum Vitae	145

List of Figures

1.1	Upper reach of the Navisence River above Zinal, Canton of Valais, Switzerland	2
1.2	Channel-pattern spectrum	4
1.3	Time series of sediment transport rate	13
1.4	Normalized cumulative mean of sediment transport rate	13
1.5	Cumulative sediment mass over time	14
2.1	Markov model flowchart	20
2.2	Experimental flume	23
2.3	Experimental setup and sensors	25
2.4	Image processing workflow for masks and edges	26
2.5	MHD computation example	27
2.6	Illustration of the Dice Index calculation	29
2.7	Comparison of dimensionality reduction techniques	32
2.8	DBSCAN and HDBSCAN on a two-moons dataset	34
2.9	Graphical Markov representation	38
3.1	Distance Matrix (MHD)	48
3.2	Distance Matrix (DICE)	49
3.3	Most different morphologies	50
3.4	Cumulative explained variance of principal components	52
3.5	Example 2-D metric MDS embedding	54
3.6	Representative 2-D t-SNE embedding	54
3.7	Representative 2-D UMAP embedding	55
3.8	DBSCAN vs. HDBSCAN on the same UMAP embedding	58
3.9	3-D UMAP embedding colored by HDBSCAN clusters	60
3.10	Binarized centroids of the 16 clusters	62
3.11	Empirical sediment transport distributions by morphological state	66
3.12	Sediment transport vs. entropic braiding index	67
3.13	Sediment transport vs. mean wetted width	68
3.14	Correlation matrix of sediment transport and planform metrics	69
3.15	Sediment transport time series colored by cluster	69
3.16	Temporal distribution of morphological states	71
3.17	Dwell-time histograms of representative clusters	72
3.18	Transition-probability matrix P for the embedded jump chain	73

3.19 State persistence versus activity	75
3.20 Frequent bidirectional transitions between clusters	76
3.21 Change in mean sediment transport for ordered transitions between clusters	77
3.22 Relative importance of morphological clusters	80
3.23 Mean first-passage time mosaics	81
3.24 Entropy–rate contributions of morphological clusters	83
3.25 Validation of morphological projection	84
3.26 Histogram of sediment transport for all experiments	85
3.27 Transition matrices comparison across experiments	86
3.28 Autocorrelation function of sediment transport	86
3.29 Power spectral density of sediment transport signals	87
4.1 Schematic workflow	90
4.2 Comparison of experimental and synthetic Q_s time series	92
4.3 Minimal decision tree	93
4.4 Screening outcome in frequency–Frobenius space	94
4.5 Mean and variance of accepted series	95
4.6 Return period validation	97
4.7 Temporal dependence validation	98
4.8 Variance explained for $K = 16$	100
4.9 Effect of cluster number on ρ_B	102
4.10 Mean ρ_B as a function of K	103
4.11 Acceptance rate vs. cluster number	103
4.12 Performance of the hybrid model	105
4.13 Adapted Bertoldi relationship	106
4.14 Performance of the hybrid model with Bertoldi scaling	108

List of Tables

2.1	Experimental flume dimensions, characteristics, and runs.	23
3.1	Explored parameter values for non-linear dimensionality reduction methods.	53
3.2	Explored hyperparameter values for DBSCAN and HDBSCAN.	57
3.3	Recorded parameters and quality metrics for every clustering run.	57
3.4	Selected parameters for UMAP embedding and HDBSCAN clustering. . . .	59
3.5	Clustering performance metrics for the selected configuration.	59
3.6	Cluster-wise statistics: image count, wetted width (cm), wetted area (m ²), sinuosity, and braiding indices (scaled image counts, rescaled min=19.7 cm)	63
3.7	Cluster-wise mean and standard deviation of the sediment transport rate Q_s (g s ⁻¹), computed from the time series while each state was active. . . .	66
3.8	Estimated rate parameter λ for the exponential distribution fitted to the durations of continuous presence in each cluster. Values are expressed in inverse minutes.	72
3.9	Five transitions with the largest absolute change in mean transport. Group membership follows Figure 3.19. Red means anchor, blue means intermediate, green means transit.	77
3.10	Row wise transition probabilities between color groups	78
4.1	Ensemble means and 95% confidence intervals of ρ_B for different levels of cluster granularity.	101

1 | Introduction

Rivers are not merely conduits of water. They supply drinking water, energy and navigation; recycle nutrients; recharge groundwater; dilute pollutants; and sustain mosaics of riparian habitats with high biodiversity. For much of the twentieth century, hydraulic engineering prioritised flood control and hydropower, often overlooking these broader ecological and geomorphic functions. However, contemporary restoration and rehabilitation seek a more holistic balance among safety, economic needs and ecosystem integrity, motivating interventions such as removal of bank protection, side-channel reactivation and mitigation of hydropeaking effects (Bertoldi et al., 2009b). Delivering such measures cost-effectively demands a robust understanding of river dynamics across chemical, biological, ecological, hydrological and morphodynamic dimensions.

Despite sustained progress, core morphodynamic questions remain unresolved. Nowhere is this more evident than in the strongly mobile reaches where bedload dominates and channel change is rapid. There, sediment pulses, braid-plain mobility and the attainment (or absence) of transport equilibrium are difficult to measure and harder to predict, even under steady forcing (Ancey, 2020a, 2020b). This chapter frames that challenge and establishes the conceptual ground needed to discuss why braiding matters, what is known about its controls, and why prediction remains difficult.

1.1 What is a braided river?

A braided river, like the Navisence in Figure 1.1, is a multi-thread channel network in which individual threads repeatedly bifurcate and rejoin around unvegetated bars (Jagers, 2003). Although there is no singular definition, several concise descriptions are commonly accepted:

- “A system of channels, splitting and rejoining around islands” (Jagers, 2003);
- “An unstable network of multiple channels separated by ephemeral bars” (P. Ashmore, 2024);
- “A network of bars and interweaving channel threads” (Y. Wang et al., 2024);
- “Multiple channel threads separated by mobile, inter-channel bars” (Li & Limaye, 2025).



Figure 1.1: UAV photo of the upper reach of the Navisence River above Zinal, Canton of Valais, Switzerland. Credit: Bob de Graffenreid, EPFL/LHE.

Braided rivers occur across a wide range of climates and scales. They are especially frequent in proglacial corridors, arid and semi-arid basins, and mountain piedmonts where coarse sediment supply is high and riparian vegetation is scarce (P. Ashmore, 2013). Typically set in unconfined valleys, these systems combine (i) large fluxes of unconsolidated sediment, (ii) stream power sufficient to mobilize that load, and (iii) banks prone to erosion (Wohl, 2020). The channels reorganize rapidly, with lateral migration and bar turnover occurring on short time scales (E. W. Lane, 1955).

Although no single criterion uniquely identifies a braided planform, the literature converges on several recurring ingredients:

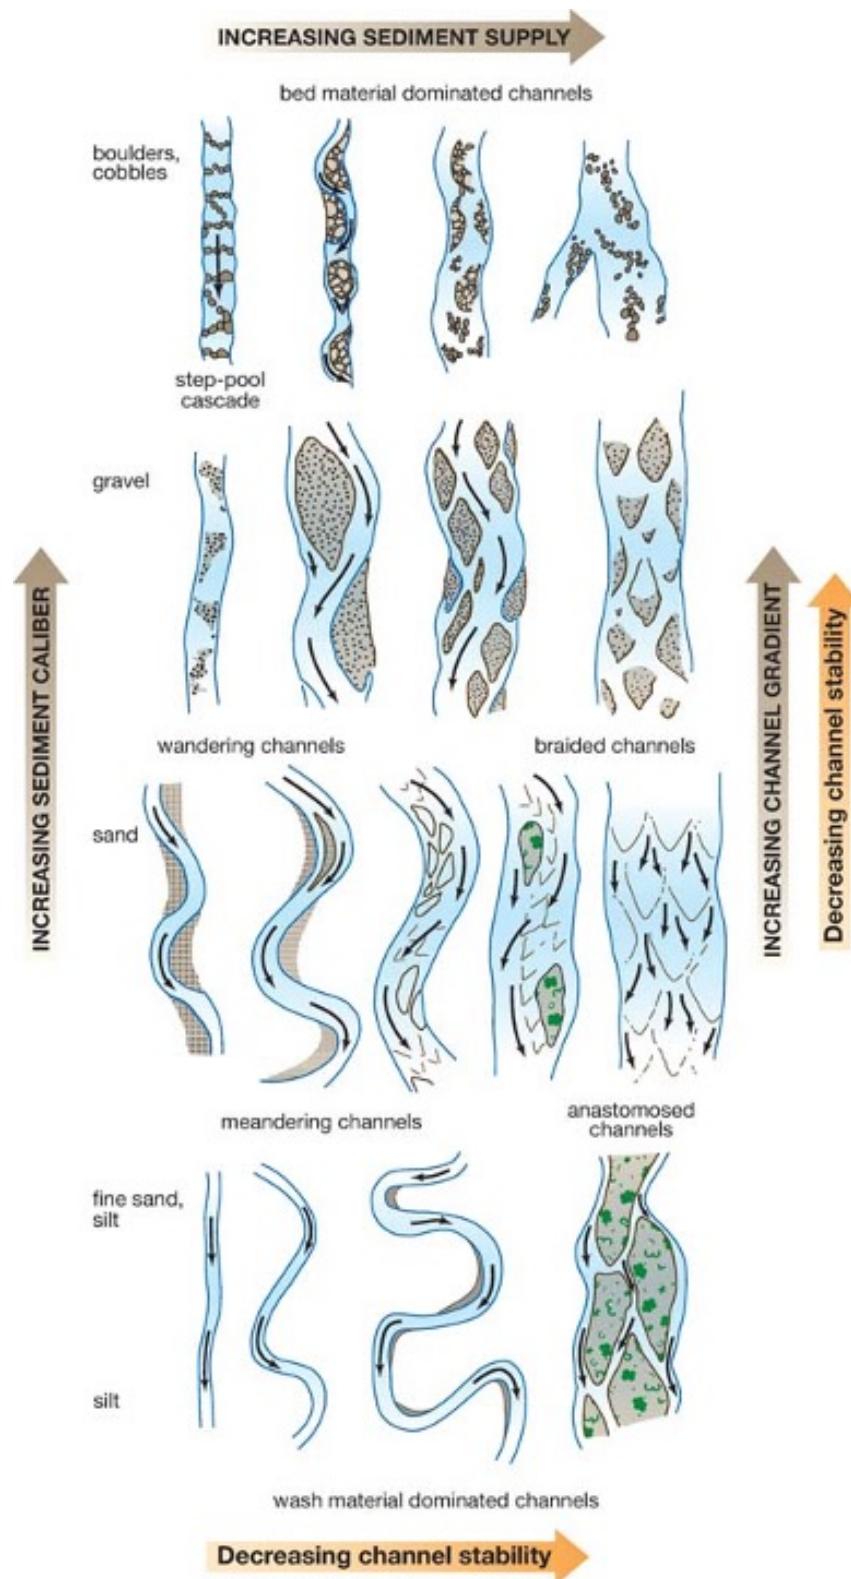
- ◊ **High stream power.** Steep valley gradients (roughly 1–10%) favor braiding, but unit stream power (discharge times slope) is a more robust discriminator: for a given grain size, braided rivers tend to exhibit larger values than meandering counterparts (P. Ashmore, 2013; Knighton, 2014).
- ◊ **Abundant bedload supply.** When sediment input exceeds transport capacity, deposition initiates mid-channel bars that deflect flow and promote new bifurcations (Leopold & Wolman, 1957; Williams et al., 2015).
- ◊ **Erodible, non-cohesive banks.** Easily eroded banks enable channel widening and bar proliferation; cohesive substrates or dense vegetation promote single-thread or anastomosing patterns (Eaton et al., 2010).
- ◊ **Rapid discharge fluctuations.** Large floods or strong melt-season hydrographs enhance bank erosion, reset bedforms, and clear stabilizing vegetation, maintaining braiding activity (Burkham, 1972; L. Wang et al., 2019).

In essence, an active braided river is a wide, shallow corridor of non-cohesive sediment with limited in-channel vegetation, driven by high rates of potential-energy dissipation. Several measurable attributes help characterize planform and short-term dynamics in practice: the braid index (number of active threads), junction density, bar size and turnover, and lateral-migration rates (Egozi & Ashmore, 2009; Limaye, 2020). These indicators vary with discharge regime and with how channels and bars are delineated from imagery or topography (Stecca et al., 2019). The continual creation, erosion, and re-entrainment of bars produce large, rapid swings in bedload transport and planform, consistent with documented nonlinear responses in geomorphic systems (Phillips, 2006).

Braiding also sits on a broader continuum of channel patterns governed by sediment caliber, supply, and slope. Figure 1.2, adapted from Church (2006), places braiding within this spectrum, between wandering and anastomosed patterns, emphasizing that modest shifts in external drivers can move a river among planform states without qualitative changes in material properties.

1.2 Why morphology forecasting matters

Braided rivers reshape valleys quickly, rerouting flow and redistributing sediment in ways that affect safety, infrastructure, navigation, and habitat. Changes in planform alter flood conveyance and bank stability, influence bridge and pipeline exposure, and restructure the mosaic of bars, side channels, and shallow ponds that support riparian and aquatic biodiversity (Bristow & Best, 1993; Church & Ferguson, 2015; Foufoula-Georgiou & Sapozhnikov, 2001; Jagers, 2003; Surian, 2015). Because these systems can reorganize on event to seasonal time scales, decisions about corridor protection, asset siting, and habitat maintenance benefit from anticipating where channels will split, migrate, and reconnect.



 Church M. 2006.
 Annu. Rev. Earth Planet. Sci. 34:325–54

Figure 1.2: Morphological continuum of alluvial channels as a function of sediment caliber, supply, and gradient (Church, 2006). Braided rivers occupy the high-supply, coarse-sediment, high-slope domain, bracketed by wandering and anastomosed patterns.

From a management and restoration standpoint, forecasting morphology ties directly to the transition from form-based interventions toward process-based restoration: projects are more effective when they restore regimes of flow, sediment, and wood and provide rivers with room to adjust, rather than imposing fixed forms (Beechie et al., 2010; Kondolf, 2006; Poff et al., 1997; Polvi et al., 2020; Wohl, 2020; Wohl et al., 2016). Large restoration programs and policies (e.g., Switzerland's framework for allocating space to rivers and planning rehabilitation) seek to re-establish connectivity and sediment continuity, yet outcomes depend on how morphodynamics and sediment supply interact locally (Bernhardt et al., 2005; Biron et al., 2014; Kurth & Schirmer, 2014; Woolsey et al., 2007). A probabilistic, state-transition description grounded in observable morphology delivers the type of information practitioners need, this is ranges rather than single values, dwell times, and lead times, so they can identify where dynamic widenings are likely to activate, how long states persist, and how transport intermittency may reorganize habitats and risks, thereby informing process-based design, monitoring, and adaptive management (Beechie et al., 2010; Polvi et al., 2020; Wohl, 2020).

Over the past century, many braided river corridors have simplified under human pressures such as channel incision, narrowing, embankment construction, and vegetation encroachment, reducing braiding intensity and altering ecosystem services (Bertoldi et al., 2014; Habersack et al., 2008; Stecca et al., 2019; Surian, 2015). Activities like gravel extraction, flow regulation, and channelization disrupt sediment budgets and disturbance regimes, thereby reducing the likelihood of bar formation and limiting thread mobility (Bertoldi et al., 2014; Redolfi, 2015). Forward-looking assessments can clarify whether common restoration actions, such as sediment augmentation, removal of bank protections, or side-channel reactivation, are likely to shift a reach toward desired geomorphic and ecological conditions, while also balancing cost and risk (Church & Ferguson, 2015; Ettema & Armstrong, 2019; Redolfi et al., 2020).

Forecasting is also motivated by measurement limits and process variability. Direct bedload monitoring rarely scales across a braidplain, and reported transport signals show strong intermittency, pulses, and threshold effects that complicate extrapolation (Ancey, 2020a; Antoniazza et al., 2019; Bertoldi et al., 2009a; Egozi & Ashmore, 2009). Morphodynamic adjustments depend on antecedent topography and junction dynamics, so exact trajectories are hard to predict; credible forecasts therefore represent uncertainty explicitly and emphasize likely ranges of change rather than single paths (S. N. Lane, 2006; Phillips, 2006; Singh et al., 2009; Stecca & Hicks, 2022; Warburton & Davies, 1994). In practice, morphology serves as both signal and proxy: repeat digital elevation models, orthophotos, and satellite imagery locate erosion and deposition, support inverse estimates of fluxes, and reveal patterns of lateral migration and bar turnover relevant for reach-scale planning (Redolfi et al., 2020; Sambrook Smith et al., 2006; Vericat et al., 2017).

Taken together, these considerations make morphology forecasting a practical requirement for managing braided rivers. It links observed planform dynamics to risks, costs, and ecological outcomes, and it provides a transparent basis for comparing scenarios under present constraints of regulation, sediment harvesting, and changing hydrologic regimes.

1.3 Complexity and variability

Braided rivers are strongly nonlinear, multiscale systems. Even under steady forcing they display intermittent pulses, sensitive dependence on initial conditions, and dynamics that produce divergent configurations from nearly identical starts, something that makes us think of deterministic chaos (Phillips, 2006; Sivakumar, 2004; Stecca & Hicks, 2022). Coherent structures (i.e., bars, chute cutoffs, confluences–bifurcations) interact across scales in ways that resemble turbulence cascades, yielding emergent organization within patterns that appear chaotic (Doeschl et al., 2006; Jagers, 2003; Stecca & Hicks, 2022).

Variability is not an artifact of external disturbances alone. Laboratory and field studies document autogenic fluctuations in bedload and planform at constant discharge, including phases of rapid change alternating with quiescent periods (P. Ashmore, 2024; Jagers, 2003; Warburton & Davies, 1994). Time series of transport and morphology show long-tailed distributions, bursts of transport, and weak short-term predictability, while stationarity, if there is any, often emerges only in an averaged, statistical sense (Ancey, 2020a; Stecca & Hicks, 2022; Vesipa et al., 2018). Scale invariance and dynamic scaling further testify that braiding entails a hierarchy of lengths and times, with partial self-similarity in planform and topography (Doeschl et al., 2006; Foufoula-Georgiou & Sapozhnikov, 2001; Surian, 2015).

In steep, bedload-dominated settings, variability tends to be particularly high and equilibrium notions become subtle. Lateral instability is persistent, thresholds are episodically exceeded, and averages converge slowly (Ancey, 2020a; Church & Ferguson, 2015; Singh et al., 2009). These properties difficult point forecasting since attempting single best-trajectory predictions is quickly undermined by nonlinear amplification and intermittency. More informative forecasts, on the other hand, target distributions of state change rather than exact paths (S. N. Lane, 2006; Stecca & Hicks, 2022).

Because these dynamics show themselves on bars, nodes, and channel threads, the planform provides a clear, scalable record of the underlying processes. Then, morphology as the primary lens offers both practicality and signal for inference at reach to basin scales (Doeschl et al., 2006; Sambrook Smith et al., 2006; Vericat et al., 2017).

1.4 Morphology as a long-standing and practical lens

For more than half a century, braided rivers have been summarized through planform morphology. Classic metrics (e.g., the braiding index (BI), channel-count measures, bar wavelength, and confluence–bifurcation spacing) capture style and degree of braiding and remain widely used in lab and field studies (Chalov & Alexeevsky, 2013; Doeschl et al., 2006; Egozi & Ashmore, 2008; Surian, 2015). Work on “active” versus “total” braiding intensity showed that braiding responds systematically to discharge and stream power, emphasizing the need to separate threads that convey flow from those that are inactive at a given stage (Egozi & Ashmore, 2009). At the same time, critiques have noted that BI based solely on low-flow channel counts can be sensitive to stage and delineation rules, urging care in how channels and bars are mapped from images or topography (Stecca et al., 2019). Building on these insights, recent contributions have introduced the effective braiding index (eBI), an entropy-informed modification of BI that discounts uneven thread activity and improves robustness; the BI/eBI ratio provides a compact descriptor of thread heterogeneity (Tejedor et al., 2022). Together, BI, eBI, and related intensity measures offer a practical language for comparing reaches, tracking adjustments, and relating structure to transport behavior.

Although the literature has often summarized the behavior of braided rivers through a general characterization based on one or a few planform indices, such as the previously mentioned and characteristic lengths (Chalov & Alexeevsky, 2013; Doeschl et al., 2006; Eaton et al., 2010; Egozi & Ashmore, 2008, 2009; Surian, 2015), this work adopts a different perspective: **we propose a multidimensional characterization that describes the system as a collection of observable morphological states and their transition probabilities.** This shift in emphasis makes it possible to explicitly capture the diversity, persistence, and alternation of configurations, rather than compressing the entire behavior into a single number or set of parameters.

Morphology also supports a more “physical” line of inference: the morphological method for estimating sediment fluxes by mass balance of erosion and deposition between surveys. This inverse approach, rooted in the Exner equation, accumulates volumetric changes over mapped areas and divides by the elapsed time to obtain reach-scale bed-material transport rates (Antoniazza et al., 2019; Vericat et al., 2017). When repeat digital elevation models (DEMs) are available, difference maps (DoDs) localize scour and fill, enabling spatially distributed estimates and event-to-seasonal budgets (Vericat et al., 2017). Recent applications extend this idea to two-dimensional mapping of fluxes across braidplains, leveraging survey intervals aligned with hydrologic events (Antoniazza et al., 2019). The approach complements point measurements by capturing intermittent pulses and pathway shifts that are otherwise hard to sample, while requiring careful treatment of vertical error, thresholding, and survey frequency (Antoniazza et al., 2019; Vericat et al., 2017).

The rise of consistent plan-view data has reinforced both strands. Aerial photography, fixed cameras, UAVs, and satellites provide repeatable coverage across large extents and long durations, turning morphology into an accessible measurement target (Jagers, 2003; Sambrook Smith et al., 2006; Tejedor et al., 2015b; Vericat et al., 2017). In practice, planform mapping now supports (i) index-based summaries (BI, eBI, intensity) to assess braiding de-

gree and its variability, and (ii) change-detection workflows to quantify erosion–deposition budgets and infer reach-scale fluxes at manageable cost (Antoniazza et al., 2019; Egozi & Ashmore, 2009; Tejedor et al., 2022; Vericat et al., 2017).

Treating morphology as a time-ordered signal is a good way of integrating these pieces. Sequences of maps can be transformed into time series of thread centerlines, nodes, and bar fields; and summarized by distributions of active width, events of rapid morphological change, and migration rates (Doeschl et al., 2006; Li & Limaye, 2025). New methods like network abstractions and information-based descriptors further quantify connectivity, redundancy, and pathway diversity in ways that connect structure to function (Booker & Eaton, 2021; Heckmann et al., 2015; Tejedor et al., 2015b). In laboratory experiments and natural rivers alike, such signals reveal dynamic scaling and near-stationary statistics at aggregate levels, even as individual configurations remain highly variable (Doeschl et al., 2006; Foufoula-Georgiou & Sapozhnikov, 2001; Surian, 2015). This combination of index-based characterization, mass-balance budgeting, and time-series analysis of mapped morphology, makes planform a practical lens for understanding braided rivers at the scales where decisions are made.

1.5 An informational viewpoint on morphology

When we see planform evolution as time series of signals, questions about how much of the near future is already encoded in today’s map arise. Information theory (Shannon, 1948) offers a concise language for that purpose where entropy measures diversity of channel occupation; mutual information estimate shared structure between parts of a network or successive times; and conditional entropy expresses the remaining freedom for reconfiguration (Tejedor et al., 2015a). In channel networks, these ideas connect naturally to loops, alternative paths, and shared flux, which can be related to redundancy and vulnerability under perturbations (Tejedor et al., 2015a, 2015b). Building on this foundation, the entropic braiding index (*eBI*) implements a Shannon-entropy measure to summarize cross-sectional channel diversity from images, providing a resolution-robust complement to classical braiding counts (Tejedor et al., 2022).

Information theory provides a way to think about predictability. Low conditional entropy and high mutual information mean that the current channel shape constrains what comes next, while high entropy indicates more freedom for change (Tejedor et al., 2015a, 2022). Observations support this: long planform records show that channel threads move coherently over years to decades, even though the system looks disordered on the surface. This suggests that channel geometry does contain usable information about future migration (Li & Limaye, 2025). At the same time, braided corridors show evidence of dynamic scaling and self-organization. Both laboratory and field studies reveal scale-invariant patterns and evolution toward near-critical states—hallmarks of systems where forecasts are better expressed in probabilistic rather than deterministic terms (Foufoula-Georgiou & Sapozhnikov, 2001; Sapozhnikov & Foufoula-Georgiou, 1997, 1999).

Recent work strengthens this informational view across scales. Using Taylor's power law and fixed-mass multifractal analysis, Rizzello et al. (2024) showed that core planform variables such as the number of wet channels, their mean width, and eBI display multiscaling and multifractality across aggregation lengths in a large braided reach. In practice, this means that entropy-based summaries and variance–mean relations change systematically with observation scale, yet remain computable from repeat imagery. This behavior supports the use of scale-aware informational metrics as both inputs and targets for forecasting.

Another application of information theory is a complementary line of work that uses spectral graph theory to formalize connectivity in multi-thread networks. Network representations yield metrics such as the number of alternative paths, resistance distance, link or flux sharing, and leakage, each reflecting how topology partitions and redistributes flux. Although developed largely for deltas, these tools have been extended to multidirectional estuarine systems, where analyses show how loops, redundancy, and tidal asymmetry influence steady-state flux allocation (Tejedor et al., 2015a, 2015b). These results demonstrate that network structure and directionality jointly govern information sharing across subnetworks. The same insight can be applied to braided corridors when networks are consistently extracted from imagery or topography. Hiatt et al. (2020) provide an operational roadmap for extracting networks, weighting links by geometry, and computing direction-aware connectivity metrics from remotely sensed data.

In sum, an informational viewpoint turns images into computable, scale-explicit summaries such as entropy, mutual information, and connectivity metrics. These descriptors (i) are grounded in observations available at corridor scales, (ii) capture how much structure today's planform already constrains, and (iii) highlight the limits of predictability that arise from scale-dependent variability, multiple possible routing options, and internally generated fluctuations. Together, these properties position information theory as a natural building block for probabilistic assessments of planform reconfigurations.

1.6 Braided rivers as a state–transition process

Building on the informational perspective outlined above, the central idea here is that the present planform constrains what can happen next, and that these constraints are statistical and scale dependent rather than deterministic. In practice, the descriptors used to summarize braiding, such as channel counts, active width, junction density, and related planform measures, vary through time even when external conditions remain steady (T. Hoey, 1992; T. B. Hoey & Sutherland, 1991; Redolfi et al., 2017; Warburton & Davies, 1994). Laboratory and field studies show that autogenic bar growth, decay, and reoccupation can sustain oscillations in planform metrics under constant or quasi-constant forcing. As a result, index time series move within ranges set by sediment caliber, supply, and slope rather than around a single value (Jagers, 2003; Redolfi et al., 2018; Stecca et al., 2019). This variability is not mere noise. Long records reveal alternation among a limited set of values or morphologies (Doeschl et al., 2006; Limaye, 2020). Recent work in bedload transport reinforces this probabilistic view, emphasizing that fluctuations are intrinsic and that improved prediction requires representing noise and pattern together (Li & Limaye, 2025).

Such behavior is consistent with the broader picture of nonlinear river dynamics. Geomorphic systems are nonlinear and can display signatures of chaos without being everywhere or always chaotic (Phillips, 2006; Sivakumar, 2004). In braided corridors, experiments and remote sensing have documented dynamic scaling and near-critical behavior, characteristic of systems that self-organize toward statistical equilibria with stable aggregate descriptors and scale-invariant fluctuations (Foufoula-Georgiou & Sapozhnikov, 2001; Rizzello et al., 2024; Sapozhnikov & Foufoula-Georgiou, 1999). From this point of view, **it is natural to see evolution as alternation among a finite catalogue of recurrent, observable morphological configurations, or states**. This framing reduces continuous, high-dimensional variability to a tractable set of recognizable patterns, while transitions among them capture the alternation seen in planform indices and provide a basis for probabilistic prediction.

These mentioned points motivate dividing planform imagery into discrete morphological states. The construction is empirical and data driven. States could be defined by thresholds or clustering in active width and channel counts, by bar stages and junction patterns, or by topology extracted from images or DEMs that represent threads and bifurcations as graphs with quantified connectivity and directionality (Hiatt et al., 2020; Tejedor et al., 2015a, 2015b, 2022). We believe morphology is a practical vehicle for this since it is directly observable at corridor scales with repeatability. UAV surveys, orthophotos, satellite composites, and repeat DEMs yield consistent time series over reaches and years (Sambrook Smith et al., 2006; Vericat et al., 2017). Moreover, morphology coevolves with flow and sediment pathways, so its patterns carry information about near-future reconfigurations, as suggested by coherent thread migration and curvature-linked adjustments documented over multiyear windows (Li & Limaye, 2025; Limaye, 2020; Rajbanshi et al., 2022). While concepts from information theory motivate the idea that the present pattern reduces uncertainty about the next, our implementation relies directly on morphology itself (i.e., on the geometry and arrangement of wetted threads) rather than on any single informational metric (i.e., indices like BI and eBI).

The idea of treating patterns as the modeling currency is not new, as it has been used recently in different areas. In deltaic and braided settings, multiple-point geostatistics (MPS) uses training images to encode higher-order spatial structure and to sample new, statistically consistent patterns. These ideas show how a few configurations can reproduce most observed variability, aligning naturally with a discrete-state view and enabling distance-based comparisons between realizations and observations (Scheidt et al., 2016a). Related developments by Mariéthoz, Renard, and collaborators demonstrate that realistic channel organization can emerge from stochastic recombination of patterns learned from images using algorithms such as SNESIM and Direct Sampling, with realizations drawn by Metropolis–Hastings or genetic search to honor multipoint constraints and conditioning data (Mariethoz & Caers, 2014; Mariethoz et al., 2010; Strebelle, 2006).

Combining time evolution with the previous methods can also be done by calibrating against image sequences. Hoffmann et al. (2019) proposed a Bayesian framework that (i) defines summary morphodynamic statistics, (ii) generates large ensembles of synthetic sequences by combining a Markov transition structure with randomized dwell times and within-state resampling of images, and (iii) accepts only those sequences whose statistics match the data. This method quantifies uncertainty in mode occupancies and switching and operationalizes the link between a small catalog of morphological modes

and stochastic alternation among them. The work from Hoffmann was a big inspiration for our work.

We believe a Markov state–transition model is a natural way to operationalize these insights. The construction is straightforward: select a state space from observable planform attributes (here, planform water distribution), choose a diagnostic time step suited to the data, and estimate the transition matrix $P = \{p_{ij}\}$ by counting how often one state follows another in time-ordered imagery. This yields multi-step forecasts $\mathbf{p}_{t+k} = \mathbf{p}_t P^k$, dwell-time and return-time diagnostics, and a stationary occupancy vector π satisfying $\pi P = \pi$ (Arkov et al., 1999; S. N. Lane, 2006). Our assumption is pragmatic rather than physical. Our goal is not to argue that river physics is memoryless, but to pick a resolution and a state definition at which unresolved mechanisms appear approximately Markovian; even if persistence or latent drivers matter, we could use higher-order or hidden-state variants to keep the state–transition logic (Arkov et al., 1999). Precedent exists in fluvial contexts, where alongstream variability in the number of active channels has been modeled as a Markov chain over cross-sectional states, reproducing salient statistics of braiding (Jagers, 2003); analogous random-walk and Markov formulations on channel networks connect topology to stationary flux allocation and vulnerability via Laplacian eigenstructure (Tejedor et al., 2015a, 2015b). In general terms, it has been proved that coarse-grained Markov descriptions are effective in those cases where regime occupancy and switching frequencies are statistically stable even when detailed paths are not (S. N. Lane, 2006; Phillips, 2006; Sivakumar, 2004).

This approach, even if not strict in conceptual terms, offers practical advantages. First, it separates what can change (the states) from how often these changes occur (the entries of P), making return times, absorption probabilities, and k -step likelihoods interpretable in morphological terms and potentially scalable from reaches to networks. Second, it is data efficient, since states come from imagery and transitions are simply counted, so new acquisitions update P and refine uncertainty without retuning complex physics-based models (Doeschl et al., 2006; Tejedor et al., 2015a). Third, it can accept nonstationarity since seasonal hydrology or sediment-supply shifts can be represented by allowing P to vary through time or by conditioning transitions on external covariates while keeping the same core structure (Stecca & Hicks, 2022).

The pertinence of a Markov state–transition model is reinforced by the coupling between morphology and sediment transport. Active width, junction asymmetry, and bar activity co-vary with discharge and stream power, and bedload signals display intermittency and pulses that mirror configuration changes rather than smooth trends (Ancey, 2020a; Bertoldi et al., 2009a; Egozi & Ashmore, 2009; Recking et al., 2024). At reach scale, the morphological method converts mapped erosion–deposition into flux estimates by mass balance between surveys, localizing scour and fill, and revealing when and where pulses originate (Habersack et al., 2008; Vericat et al., 2017). By characterizing each morphological state with transport summaries (means, variances, tail behavior) and using the Markov chain for occupancies, dwell times, and switching frequencies, we obtain a forecast mixture for transport. Morphology thus becomes a practical bridge from images to probabilistic predictions of transport magnitude and variability, consistent with evidence across laboratory and field studies and aligned with modern monitoring workflows (Ettema & Armstrong, 2019; Sambrook Smith et al., 2006).

1.7 The challenge of forecasting bedload

The state–transition perspective makes clear that sediment flux depends not only on hydraulics but also on which morphological configuration is active and how long that configuration persists. This distinction is important because pointwise bedload prediction remains highly uncertain even in controlled experiments and after decades of study. Coarse-bed transport is inherently intermittent, nonlinear, and influenced by multiple scales (Ancey, 2020a). Measurement is also challenging because high flows restrict direct sampling, surrogate sensors require site-specific calibration, and rating-curve extrapolations smooth out real variability. Reach-scale estimates therefore often show substantial and unavoidable scatter (Habersack et al., 2008; Recking et al., 2024; Vericat et al., 2017). Grain-scale processes such as hiding–exposure and incipient motion, together with bed-surface changes such as armoring or partial transport, alter entrainment thresholds. These dynamics generate long periods of low transport that are interrupted by sudden bursts (Ancey, 2020a; Marti & Bezzola, 2006; Meunier et al., 2006). Field and laboratory records both confirm that even under steady forcing, bedload can vary from near zero to several times the mean with only weak persistence over time. This pattern underscores the fundamental limits of predicting transport on a path-by-path basis (Ancey, 2020a; Warburton & Davies, 1994).

Our long flume experiment (≈ 1200 h) reinforces these observations. The instantaneous series in Figure 1.3 is highly intermittent, showing spikes well above the mean and long quiet periods. This agrees with previous reports that pulses and lulls alternate even when discharge remains nearly constant (Dhont & Ancey, 2018). The normalized running mean in Figure 1.4 approaches the long-term average only slowly and later drifts away, which supports analyses suggesting that very long windows of several hundred hours are required to estimate the mean within a modest error, while short records can deviate by orders of magnitude (Ancey, 2020a). The cumulative mass curve in Figure 1.5 departs from the linear trend for extended intervals, which reveals regime-like contributions where long periods of above-average accumulation alternate with quieter phases. Similar patterns have been associated with episodic bar migration and pool aggradation–degradation, although a considerable share of the transported volume can occur without clear bar translation. This indicates that several mechanisms may produce the same flux signature (Dhont & Ancey, 2018). Taken together, these results show that experiments lasting only ~ 100 –200 h risk misrepresenting what counts as “typical” transport conditions in braided or bar-forming channels (Ancey, 2020a; Doeschl et al., 2006).

These dynamics have practical consequences. First, replacing a fluctuating signal with a single mean is not innocuous because transport–shear relations are convex, so averaging inputs underestimates volumes (Jensen’s inequality). Taking natural variability into account usually increases predicted yields by ~ 10 –30% compared with deterministic runs that use mean inputs (Recking et al., 2024). Second, fluctuations depend on configuration because the same reach alternates among states with different hydraulics, thresholds, and connectivity. A single “bedload rate” is therefore less informative than a distribution of likely fluxes conditioned on the active configuration and its dwell time (Ancey, 2020a; Jagers, 2003; Recking et al., 2024). Third, the instruments and workflows that scale most effectively such as UAV imagery, orthophotos, satellite composites, and repeat

topography, deliver morphology reliably, while direct flux measurements remain scarce and uncertain. Morphology is therefore both a signal and a proxy for transport at corridor scales (Sambrook Smith et al., 2006; Vericat et al., 2017).

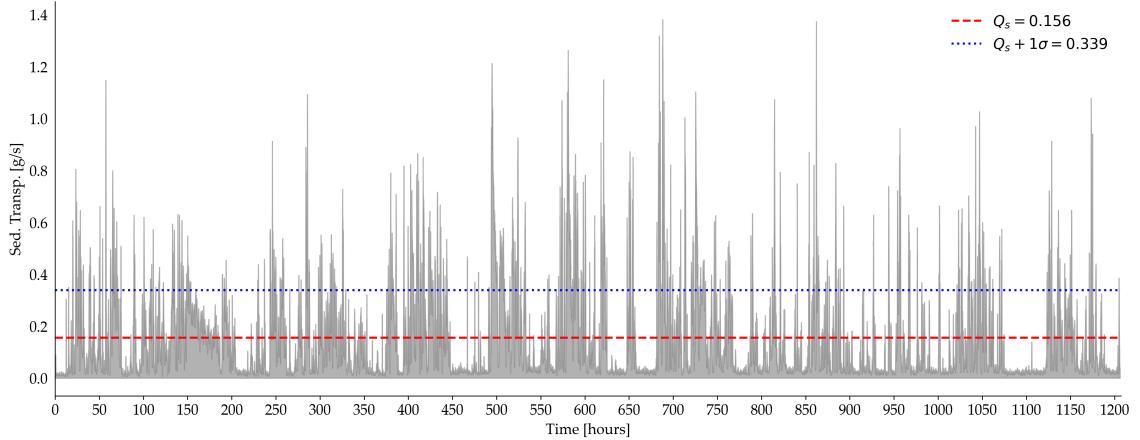


Figure 1.3: Time series of sediment transport rate measured in the experimental flume. The shaded area represents the instantaneous bedload transport [g/s] over time. The red dashed line indicates the mean transport rate $\mu = 0.156 \text{ g/s}$, while the upper and lower blue dotted lines represent one standard deviation above and below the mean ($\mu + 1\sigma = 0.283 \text{ g/s}$ and $\mu - 1\sigma = 0.029 \text{ g/s}$, respectively).

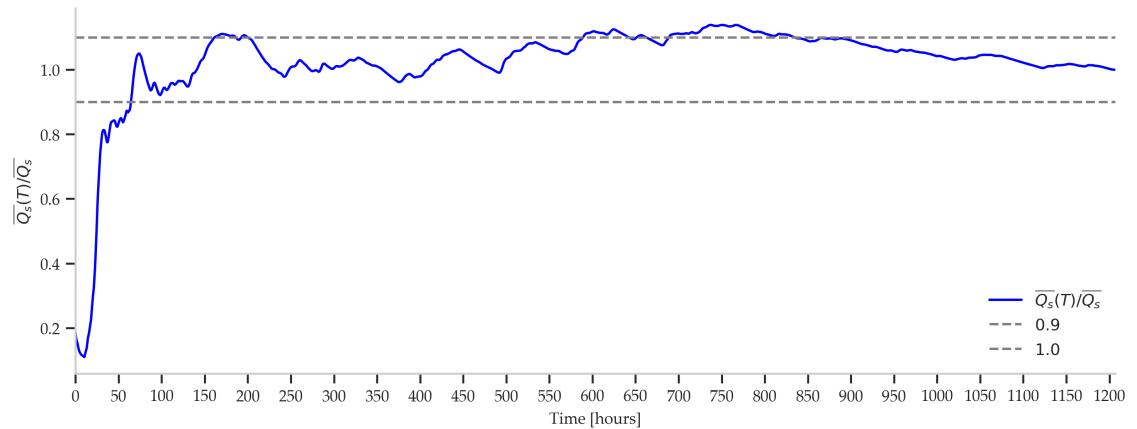


Figure 1.4: Cumulative mean of sediment transport rate $\overline{Q}_s(T)$ over time normalized by the long-term mean transport rate, $\overline{Q}_s(T)$. The blue curve shows how the cumulative mean initially converges toward the mean value ($\overline{Q}_s = 1.0$, dashed gray line), with a transient phase stabilizing near 0.9–1.0, followed by a gradual deviation at later times. This behavior reflects the interplay between local fluctuations and long-term sediment transport dynamics in the flume.

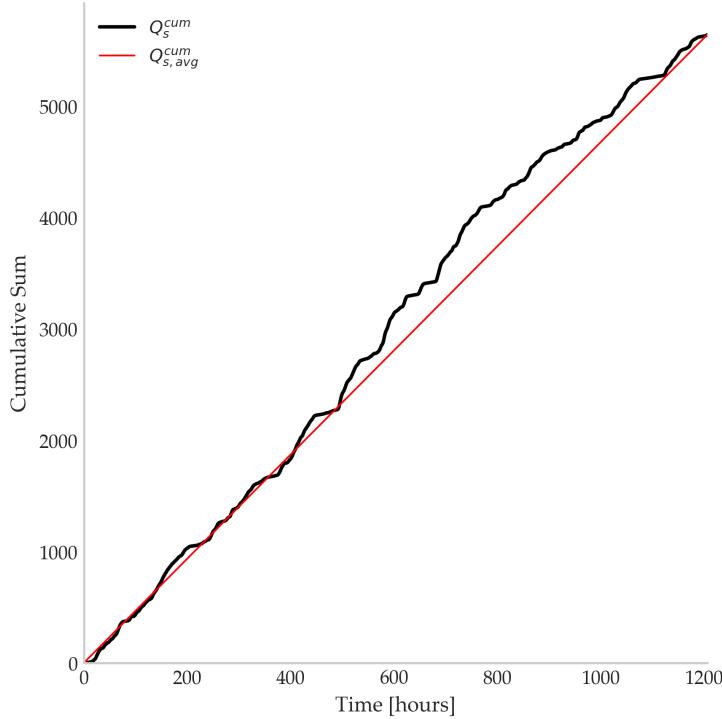


Figure 1.5: Cumulative sum of sediment mass over time. The black curve (Q_s^{cum}) shows the progressive accumulation of sediment transport during the experiment. The red line ($Q_{s,avg}^{cum}$) connects the first and last points, representing the net linear trend between initial and final states.

A state-based approach closes this gap. If images and DEM differencing identify the configuration that the corridor occupies, then Markovian state–transition models can forecast occupancies and switching frequencies; relating each state to transport summaries (median, spread, exceedance probabilities) turns morphology into probabilistic flux forecasts rather than single-value predictions (Arkov et al., 1999; S. N. Lane, 2006). This perspective aligns with recent recommendations to represent natural variability explicitly and to report volumes with confidence intervals instead of point estimates (Recking et al., 2024). It also integrates with modern monitoring like surrogate sensors (hydrophones, geophones) that capture high-frequency fluctuations, while plan-view mapping and DEMs supply the configuration context needed to interpret those fluctuations at reach scale (Recking et al., 2024; Vericat et al., 2017).

In applied terms, forecasting bedload is not about chasing the exact trajectory of $Q_s(t)$; it is about quantifying how morphological regimes allocate time and flux. Probabilistic forecasts tied to observable states are more actionable for flood protection, reservoir sediment budgeting, gravel-mining quotas, habitat maintenance, and post-event recovery than a single site-wide mean. They respect the irreducible intermittency of coarse-bed transport, leverage information contained in form, and acknowledge that management decisions must plan for ranges rather than point values (Ancey, 2020a; Habersack et al., 2008; Recking et al., 2024; Vericat et al., 2017).

1.8 Knowledge gap, aim, and research questions

Although planform descriptors such as BI/eBI, active width, junction density, connectivity analyses, and entropy-based metrics are well established, and although synoptic imagery and repeat topography are now routinely available, most applications still stop at description or bulk inference rather than moving toward forward prediction with quantified uncertainty (Booker & Eaton, 2021; Sambrook Smith et al., 2006; Tejedor et al., 2015a, 2015b, 2022; Vericat et al., 2017). At the other end of the spectrum, physics-based morphodynamic models struggle to propagate uncertainty and to remain reliable at braidplain scale under strong nonlinearity and changing configurations (Church & Ferguson, 2015; S. N. Lane, 2006). Between these two extremes there is still no operational framework that can translate widely available image data into transport-relevant morphological states, represent their alternation with a tractable stochastic dynamic, and use that representation to produce probabilistic transport forecasts with explicit uncertainty.

This missing middle ground defines the knowledge gap. Current research has shown that braided planforms are not purely descriptive but carry predictive signal when cast as multi-state systems (Ancey, 2020a; Jagers, 2003; Recking et al., 2024; Tejedor et al., 2015a; Vericat et al., 2017). However, existing work has not yet provided a systematic way to transform image-based descriptors such as wetted width or eBI into probabilistic transport predictions that account for intermittency, burst structure, and configuration dependence. **A framework that can achieve this would provide the needed bridge between descriptive metrics and physics-heavy models.**

Building on this gap, **the aim of this thesis is to construct a practical link from image-derived morphology to probabilistic forecasts of braided dynamics and sediment transport.** Concretely, our work seeks to (i) define observable morphological states from plan-view data, (ii) represent their alternation with a continuous-time Markov chain (CTMC) calibrated to data, and (iii) use that representation to generate transport time series under different data-availability scenarios, with quantified uncertainty and objective acceptance criteria.

Research questions

RQ1: Image → state. How can planform image sequences be converted into a compact set of discrete, transport-relevant morphological states while preserving the gradient from simple to highly braided configurations?

RQ2: State dynamics. Do the observed transitions among states obey approximately Markovian behavior with exponential sojourns, and do estimated rates organize consistently with planform complexity and transport capacity?

RQ3: Forecasting and uncertainty. To what extent can CTMC-driven Monte Carlo sequences, combined with state-conditioned bootstrap of Q_s , reproduce the empirical mean, variance, and bursts of transport under explicit accept/reject screening and simple image-only prefilters?

RQ4: Morphology–transport links. Which image descriptors (e.g., wetted width, eBI) most strongly control inter-state contrasts in transport level and intermittency, and how do state sequences organize multi-minute transport bursts?

RQ5: Transferability. How robust is the image-to-forecast workflow across data regimes (laboratory vs. image-only) and methodological choices (distance metrics, dimensionality reduction, clustering), and how can uncertainty be partitioned across these stages?

Objectives

- O1:** Build a scalable image analysis pipeline to extract binary water masks, compute complementary distances (e.g., Modified Hausdorff and DICE), and obtain discrete morphological states via reduction and clustering (UMAP/PCA + HDBSCAN), yielding an interpretable state space.
- O2:** Estimate a CTMC from the state sequence by fitting exponential dwell times and jump probabilities, thereby specifying the generator for analysis (occupancies, mean first passage) and forward simulation.
- O3:** Quantify inter- and intra-state transport behavior by associating empirical Q_s samples to each state, enabling state-conditioned bootstrap that preserves between-state contrasts and within-state scatter.
- O4:** Generate large ensembles of synthetic trajectories via Monte Carlo + bootstrap, enforce acceptance on transport moments (ABC-style), and derive conservative, image-only decision rules to pre-screen acceptable sequences.
- O5:** Diagnose morphology–transport couplings (e.g., strong negative associations of eBI and wetted width with the mean and variability of Q_s) and relate state cascades to observed multi-minute transport bursts.
- O6:** Quantify and report uncertainty throughout, and test sensitivity/robustness to methodological choices using the Chapter 4 framework.

Contributions

This thesis contributes: (i) an end-to-end image-to-state pipeline yielding an interpretable state space; (ii) a data-calibrated CTMC whose residence rates align with planform complexity and transport capacity; (iii) a Monte Carlo + state-conditioned bootstrap generator of transport series with explicit acceptance; and (iv) conservative, image-only screening rules that approximate the physical filter, enabling forecasting from imagery even when direct transport measurements are scarce. Together, these elements form a practical bridge from images to probabilistic forecasts of braided-river evolution, emphasizing interpretability, scalability, and quantified uncertainty, and they motivate semi-Markov extensions to better capture low-frequency variability identified in Chapter 4.

1.9 Roadmap of the thesis

This document is organized into five main chapters:

- ◊ **Chapter 1 (this chapter):** introduces the overall research context, highlights the importance of studying sediment transport and morphological dynamics in braided rivers, reviews the state of the art with emphasis on chaos theory, Shannon's information theory, and probabilistic/Markov models in hydrology and geomorphology, and states the objectives and hypotheses guiding the study.
- ◊ **Chapter 2:** describes the methodology, including experiments focused on building the probabilistic Markovian model, as well as the techniques used for data acquisition and analysis.
- ◊ **Chapter 3:** presents the results related to the construction and selection of the Markovian model, covering the distance matrix, dimensionality reduction methods, clustering alternatives, and the justification for the final model configuration. It also includes the characterization of the obtained cluster space.
- ◊ **Chapter 4:** addresses the statistical validation of the selected model, reporting Bayesian inference results, bootstrap analyses, decomposition of variability between transitions and within states, and sensitivity and robustness tests.
- ◊ **Chapter 5:** provides the general conclusions, summarizes the main contributions, and offers recommendations for future research on the morphological and sedimentological dynamics of braided rivers.

2 | Materials and Methods

This study investigates how sediment is transported and how channel forms evolve in braided rivers using a controlled laboratory setting. We designed a long-duration experiment that captures slow morphological adjustments while keeping boundary conditions constant. The configuration allows us to observe hydrodynamic and morphological processes with the resolution needed to evaluate mechanisms of sediment motion and pattern change.

The experiment targets the long-term evolution of a braided pattern under steady discharge in order to test a probabilistic description of system dynamics. We explore whether the sequence of planform configurations can be represented as a Markov process with states defined by shape and by time. The flume was run for more than one thousand hours so that a broad set of bed configurations could emerge without external forcing transients.

The remainder of the chapter documents the facility and operating conditions, the image acquisition and processing workflow, the construction of spatial similarity measures, the dimensionality reduction and clustering used to define states, and the probabilistic modeling and validation steps that connect state sequences to transport behavior.

2.1 Overview of the Process

As mentioned above, this experimental facility was conceived to test the hypothesis that braided river systems behave in a probabilistic manner and can be modeled as a Markov Chain. To achieve this, our main assumption is that the river can be conceptualized as a combination of distinct morphological states, each defined by its specific shape and spatial distribution over both space and time.

We conducted three runs under identical boundary conditions. The main, continuous run lasted approximately 1200 hours (hereafter *ExpL*) and provided the record used to learn the morphological state space and calibrate the continuous-time Markov chain. To test external validity, we performed two additional, independent realizations of about 200 hours each (*ExpA* and *ExpB*), processed with the same pipeline and reserved for out-of-sample evaluation of state recurrence, transition structure, and transport statistics. All runs were imaged every minute and accompanied by continuous outlet bedload monitoring, yielding a long, high-cadence dataset suitable for estimating dwell times, transition probabilities, and state-conditioned transport.

Figure 2.1 illustrates the comprehensive process undertaken to analyze and model the braided river system using a Markov Chain approach. The process begins with the acquisition of high-resolution images of the experimental flume every one minute, which are then processed into binary masks where only the water morphology is represented in white and all other elements in black. This transformation isolates the morphological features of the river at each discrete time step.

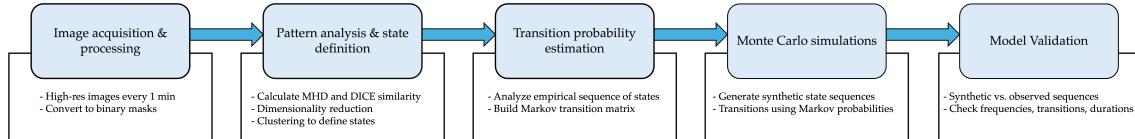


Figure 2.1: Flowchart of the process used to build the Markov model. The procedure starts with image acquisition and processing to extract morphological features. Pattern analysis and state definition are performed through similarity metrics, dimensionality reduction, and clustering. Transition probabilities between states are estimated from the empirical sequence, followed by Monte Carlo simulations to generate synthetic series. Finally, the model is validated by comparing synthetic and observed data across multiple metrics.

The next step involves defining the discrete states of the system through a clustering process. To achieve this, spatial similarity metrics such as the Modified Hausdorff Distance (MHD) and the DICE index are calculated to quantify the similarity between different morphological configurations. These metrics generate a distance matrix that captures the morphological similarity, allowing the differentiation of distinct river states. To manage the complexity inherent in high-dimensional morphological data, dimensionality reduction techniques are employed, simplifying the data while preserving essential structural information. Following dimensionality reduction, we apply density-based clustering algorithms to group similar morphological states into distinct clusters, each representing a unique state in the Markov model.

Once the states have been established, empirical transition probabilities between these states are calculated on the basis of the observed sequence of states throughout the experiment. This results in the construction of an experimental Markov Chain that effectively captures the dynamic behavior of the river system. Each state within the Markov Chain is further characterized by its morphology and sediment transport properties, with unique sediment transport distributions assigned to each morphological state, thereby linking morphological configurations to sediment dynamics.

The final stage of the process involves the application of Monte Carlo Simulations to generate synthetic sequences of states that emulate the experimental observations. Monte Carlo Simulations enable the exploration of a wide range of potential future scenarios by randomly sampling state transitions based on the established Markov transition probabilities. To ensure that these synthetic sequences accurately reflect the real system's behavior, each simulation is evaluated against the original experimental data using a series of distance-based metrics. These metrics assess the similarity of synthetic sequences to the observed data across various dimensions, including state frequencies, transition probabilities, duration distributions, and overall sequence similarity. Simulations that meet predefined thresholds across all metrics are considered acceptable, resulting in a big ensemble of synthetic series that faithfully represent the observed system dynamics.

This section will describe each component of the aforementioned process in detail. We will begin by describing the characteristics and instrumentation of the experimental flume, followed by an in-depth discussion of the data analysis methods employed, including image processing, spatial similarity metrics, dimensionality reduction, and density-based clustering. Subsequently, we will outline the probabilistic modeling approach through Markov Chains and Monte Carlo Simulations, and conclude with the tools and techniques used to study and validate the results of the Markov model. This structured approach ensures a comprehensive understanding of how braided river morphodynamics can be effectively captured and analyzed within a probabilistic framework.

2.2 Flume Characteristics

Note on the experimental facility selection.

This project began with a larger facility (2 m wide, 11 m long), on which I worked for more than two years developing the setup, instrumentation, and calibration. Early trials, however, revealed that the wide channel allowed water to move with greater freedom, creating an enormous range of possible configurations. As a result, experimental runs had to be extremely long in order to capture transitions. Because safety and budget constraints limited operation to a maximum of 12 hours per day, the calendar time needed to obtain long records effectively doubled. To overcome this, we shifted to a smaller flume that could operate continuously, accepting the extra effort required for rebuilding and recalibration. This adjustment preserved the scientific objectives while making data collection feasible. Results from the large-flume experiments are not presented here, but have been archived for potential future use beyond the scope of this thesis.

A natural question was whether the small flume (3.3 m × 1.0 m) was indeed “too small”. First, the smaller channel is easier to manage and can run without interruption, which is essential for collecting long image sequences and capturing transition statistics. Second, the narrower corridor, while still supporting braiding and bar dynamics, limits the number of admissible planform configurations compared to a wider channel. This restriction reduces the effective state space and shortens the experimental time required to observe recurring morphologies and estimate transition probabilities with precision. Small to moderate laboratory channels of similar dimensions (1.0–2.0 m wide, 2–12+ m long) have been used before in braided- and bar-forming studies (Bertoldi & Tubino, 2005; Davies & Lee, 1989; Métivier & Meunier, 2003; Paola, 2001; Pryor et al., 2011). Moreover, recent small scale experiments confirm that laboratory rivers capture robust transport–morphology coupling and self-organization, demonstrating that reduced geometries can reproduce the essential physics needed here (Abramian et al., 2020).

The main compromise of this change of facility involved instrumentation. The smaller flume could not accommodate as many monitoring tools as the original setup, and full bed reconstruction via structure from motion, which was part of the initial plan, was no longer possible. We therefore concentrated on two observables directly aligned with the probabilistic framework: (i) an extended, minute-resolution planform image series (binary water masks) for state definition and transition counts, and (ii) continuous outlet bedload measurements for transport diagnostics. This combination was able to provide what the Markov chain based analysis required: repeatable planform states, well-sampled dwell times and switches, and transport summaries conditioned on configuration.

As we mentioned before, the experimental flume used in this study (see Figure 2.2) measures 3.3 m in length and 1.0 m in width. Detailed dimensions and additional characteristics of the flume are summarized in Table 2.1. In total, three experiments were conducted under identical boundary conditions: a long continuous run of about 1200 hours (*ExpL*), used to learn the morphological state space and calibrate the Markov model, and two shorter independent realizations of about 200 hours each (*ExpA* and *ExpB*), used for out-of-sample validation of recurrence, transition structure, and sediment transport. All runs were imaged every minute and paired with continuous outlet bedload monitoring, ensuring a long, high-cadence dataset for both morphological and transport analyses.

The flume is uniformly filled with unimodal sand characterized by a median grain size (d_{50}) of 1 mm, ensuring consistent sediment transport dynamics throughout the experiments. The channel slope is maintained at 3%, providing a steady gravitational force to drive the water flow and sediment movement. A constant flow rate of 0.15 L/s is controlled using an automated valve system connected to a National Instruments USB-6001 data acquisition card and managed via a MATLAB script (Figure 2.3a-b), allowing precise regulation of the water supply across all runs.

Table 2.1: Experimental flume dimensions, characteristics, and runs.

Parameter	Value
Length [m]	3.3
Width [m]	1.0
Slope [%]	3
Sediment type	Unimodal, $d_{50} = 1$ mm
Flow rate [L/s]	0.15 (constant)
Sediment feeding method	Vibrator, funnel, and container with water
Sediment feed rate [g/s]	1.5 (constant)
Main run (<i>ExpL</i>)	~1200 h, continuous, calibration dataset
Validation runs (<i>ExpA</i> , <i>ExpB</i>)	~200 h each, independent realizations
Imaging	1 image/min
Transport monitoring	Load cell system

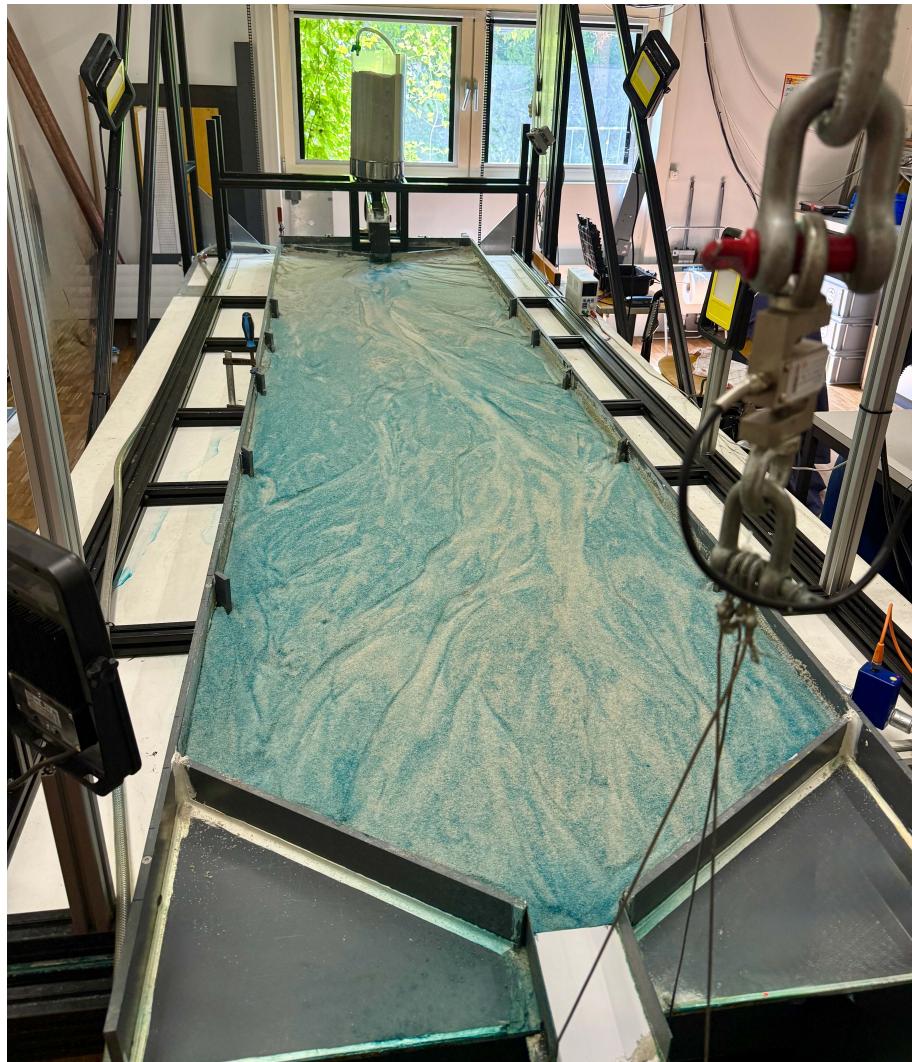


Figure 2.2: Photograph of the 3.3 m \times 1.0 m laboratory flume used in this study. Water is dyed blue to enhance planform visibility; the overhead frame supports the imaging and feeding systems.

Sediment feeding into the flume is achieved through a combination of a vibrator, funnel, and sediment container (see Figure 2.2). To prevent blockages and ensure a steady supply of sand, a small amount of water is added to improve the coupling between the vibrator and the container. This setup was calibrated during preliminary runs and adjusted to deliver a constant sediment input rate of 1.5 g/s, which keeps the inlet free of erosion and deposition. The vibratory mechanism dislodges sand particles, and the funnel directs sediment into the channel without interrupting the flow.

Sediment transport is measured with a strain gauge-based system. A collection basket captures sediment leaving the flume and is suspended from a 50 kg strain gauge with a precision of 0.03% of the maximum load (15 g). The strain gauge is interfaced with the USB-6001 and a computer running MATLAB. Measurements are acquired at 10000 Hz, then averaged over 60-second intervals and stored in text files. This averaging reduces noise and improves the reliability of transport estimates. As shown in Figure 2.3c–e, the strain gauge cell and its power/DAQ electronics provide continuous, precise monitoring of sediment weight.

Initially, a video-based sediment measurement system similar to Zimmermann et al. (2008) was considered. In that approach, high-resolution cameras and a backlit surface are used to track sediment particles continuously. In our setting this proved impractical since the particles of sand are too small; a backlit surface would have been too close to water, creating safety issues; front lighting produced shadows that degraded particle detection; and the data volumes required heavy nightly post-processing with limited options for re-analysis. For these reasons, the strain gauge-based system was adopted as a safer and simpler alternative.

Additionally, the flume is equipped with a Nikon D700 camera fitted with a wide-angle lens and fixed at a height of 2.5 m above the channel floor (Figure 2.3d). This configuration covers the entire channel in a single frame, capturing high-resolution images that delineate water boundaries. To enhance visibility, the water is dyed blue using food colorant.

In summary, the facility and instrumentation were configured to enable long, uninterrupted image sequences and continuous sediment measurements across three experimental runs. Together, *ExpL*, *ExpA*, and *ExpB* constitute the core dataset of this study, providing the basis for state definition, transition counting, and probabilistic modeling in the following sections.

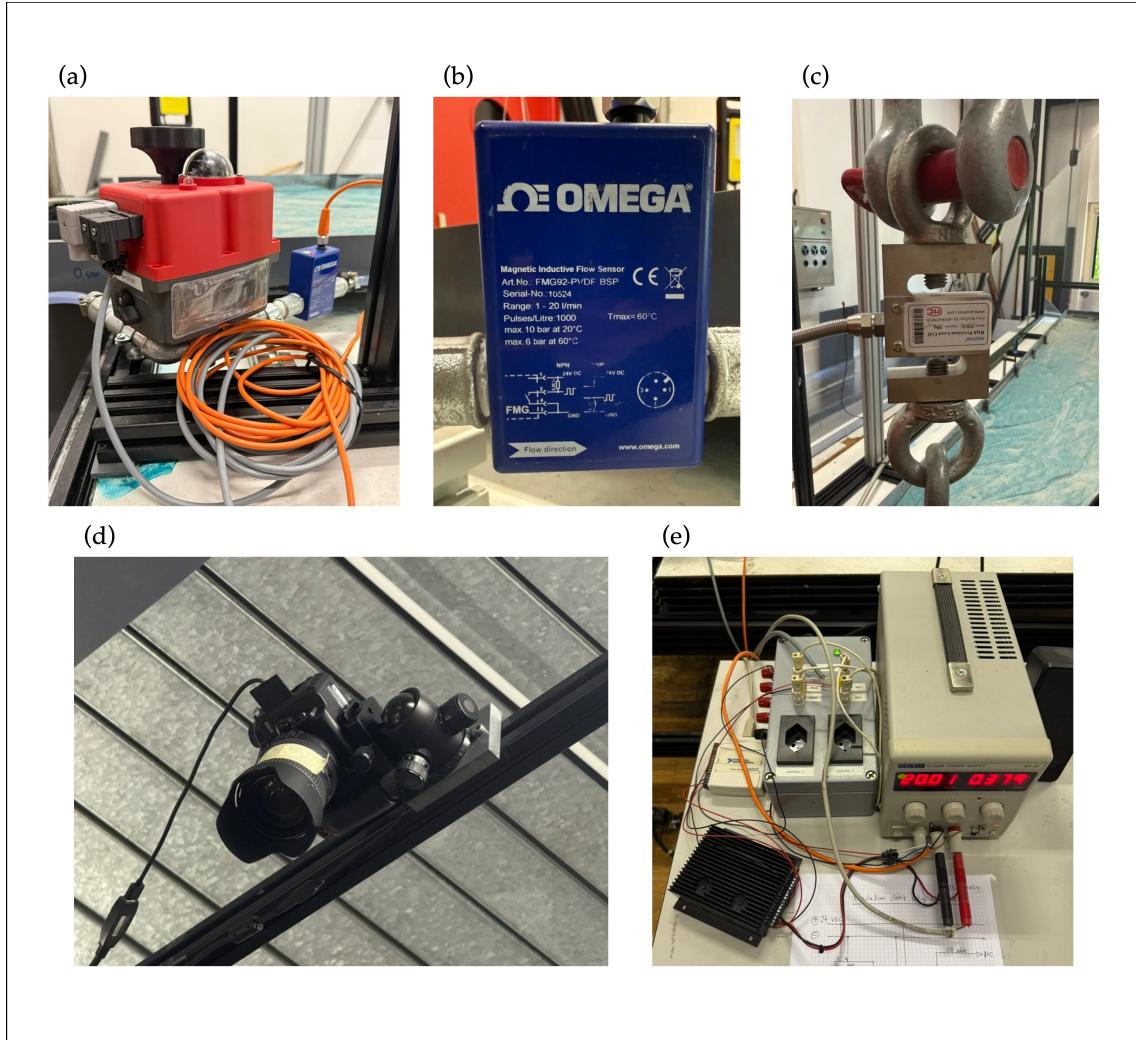


Figure 2.3: Components of the experimental setup: (a) flow-control valve; (b) magnetic inductive flow meter; (c) force sensor (strain gauge) measuring the mass of sediment exiting the channel; (d) fixed camera (Nikon D700) capturing an image of the channel every 1 min; (e) power supply and data-acquisition board.

2.3 Image Processing

The process of constructing masks from the original images was a critical step in the workflow. The same procedure was applied to all images, ensuring a reproducible process for the entire dataset. The steps were:

- (i) Crop.** We first cropped each image to the interior of the flume, retaining only the region where water flows over the sand and removing borders and hardware.
- (ii) HSV water mask.** We then generated a binary water mask using a filter in HSV color space tuned to the blue dye, which effectively separated water from dry bed material. This follows the approach of Scheidt et al. (2016a).
- (iii) Downscaling.** Masks were resized to reduce memory footprint and speed up subsequent calculations, which is essential given the size of the image archive.

(iv) Edge map for MHD. Finally, we computed a one-pixel edge map from the binary mask (i.e., the water boundaries). This edge representation is what we use in the Modified Hausdorff Distance (MHD) calculations between images, as it preserves planform structure while reducing sensitivity to small width fluctuations and greatly accelerating distance computations.

All steps were implemented in Python using open-source libraries; the code is provided in Appendix A so the procedure can be replicated or adapted.

The evolution of the images through each step is shown in Figure 2.4: original photo → cropped region → HSV mask → edge map used for the MHD analysis.

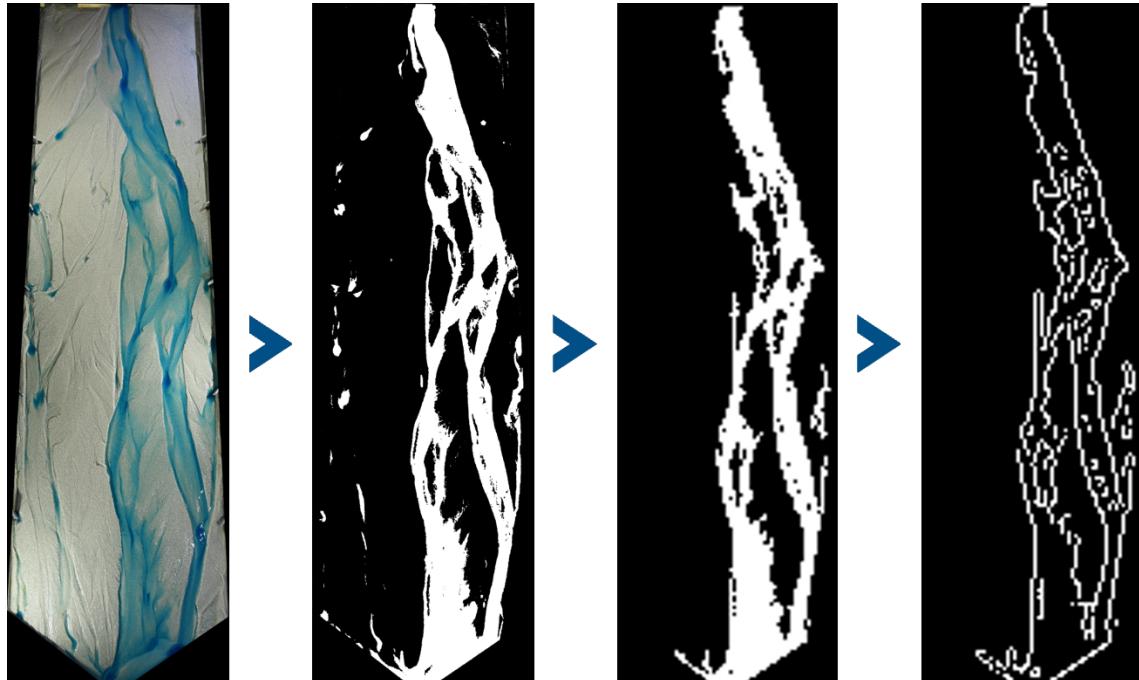


Figure 2.4: Image processing workflow. From left to right: original image, cropped flume interior, binary water mask obtained via HSV filtering, and the final edge map (one-pixel water boundaries) used for the Modified Hausdorff Distance (MHD) computations.

2.4 Distance Matrix Construction

After establishing the characteristics and instrumentation of the experimental flume, the next step is to analyze the captured morphological data to identify and group similar river configurations. Spatial Similarity Metrics are crucial for quantitatively determining how alike different morphological states are, which is essential for our objective of classifying and modeling distinct river shapes based on their spatial features. The following subsections will discuss the specific metrics used in this analysis, namely the Modified Hausdorff Distance and the DICE index, highlighting their roles in distinguishing and grouping similar morphological forms.

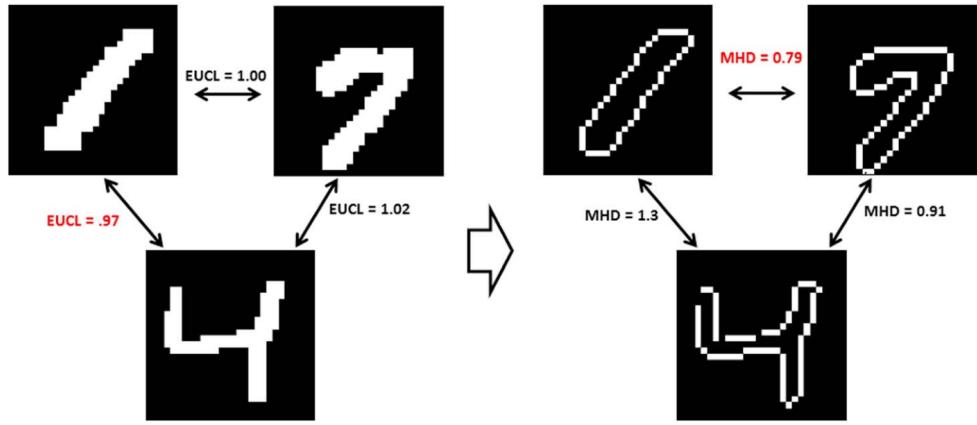


Figure 2.5: Comparison of similarity metrics using binary representations of digit-like shapes. The Euclidean distance (EUCL) fails to differentiate structural differences effectively, while the Modified Hausdorff Distance (MHD) captures edge-based structural variations more accurately. The figure highlights the advantages of MHD in capturing morphological differences between shapes. Image from Scheidt et al. (2016b).

2.4.1 Modified Hausdorff Distance

The Modified Hausdorff Distance (MHD) was first introduced by Dubuisson and Jain (1994) to address the limitations of the traditional Hausdorff Distance (HD). HD is a metric designed to quantify the similarity between two sets of points by evaluating the maximum distance between the closest points of the two sets. For two finite sets $A = \{a_1, a_2, \dots, a_N\}$ and $B = \{b_1, b_2, \dots, b_M\}$, the HD is mathematically defined as:

$$d_H(A, B) = \max \left\{ \max_{a \in A} \min_{b \in B} d(a, b), \max_{b \in B} \min_{a \in A} d(b, a) \right\}, \quad (2.1)$$

where $d(a, b)$ represents the Euclidean distance between two points a and b , calculated as:

$$d(a, b) = \sqrt{(a_x - b_x)^2 + (a_y - b_y)^2}. \quad (2.2)$$

Although effective in capturing global dissimilarities, the HD is highly sensitive to outliers because of its reliance on maximum distances. For example, a single point significantly displaced in one set can dominate the computation, leading to an overestimation of dissimilarity. In contrast, simpler metrics, like the mean pairwise Euclidean distance, fail to capture structural differences between point sets because they do not consider spatial arrangements. These limitations are particularly evident when analyzing complex shapes, as shown in Figure 2.5.

Figure 2.5 illustrates the comparison between the Euclidean distance and the MHD for binary masks representing digit-like shapes. The HD and Euclidean distance metrics can fail to differentiate structurally distinct shapes due to their sensitivity to outliers or the averaging effect of interior points. In contrast, the MHD better captures shape similarities by focusing on the distribution of the points along the edges.

To address these challenges, Dubuisson and Jain proposed the MHD, which replaces the maximum operation in the HD with an averaging approach, thereby reducing the influence of outliers. The directed MHD is defined as:

$$MHD(A, B) = \frac{1}{|A|} \sum_{a \in A} \min_{b \in B} d_H(a, b), \quad (2.3)$$

where $|A|$ is the number of points in set A . The symmetric MHD combines directed distances in both directions:

$$MHD_{sym}(A, B) = \frac{1}{2} (MHD(A, B) + MHD(B, A)). \quad (2.4)$$

As demonstrated before, the MHD better captures structural differences by focusing on edge points, which define the shape of the objects in the binary masks. Including interior points would unnecessarily increase the computational load without providing additional relevant information, as the similarity between shapes is determined primarily by their boundaries. This approach ensures that the computational effort is concentrated on the most informative regions, significantly improving efficiency.

In this study, the MHD is applied to compare binary masks generated from experimental flume images. These masks identify areas of water and dry regions, and the edges of these masks are used to compute the MHD, as they represent the morphology of the flow patterns. While focusing on edge points reduces the per-comparison cost, both the MHD and the Dice coefficient become computationally demanding at this scale because pairwise distances must be evaluated for the full minute-resolution record. To make the task feasible, we implemented from scratch a parallel, GPU-accelerated version of the MHD and validated its accuracy against the `skimage.metrics` implementation (van der Walt et al., 2014). Despite this acceleration, computing the full MHD matrix required approximately 2 months of continuous runtime on a dedicated workstation. Dice coefficients (see 2.4.2) were computed with vectorized routines (`skimage.metrics`), and although cheaper per comparison, they still contributed substantial wall-clock time due to the number of pairs. The code is provided in Appendix B.

2.4.2 Dice Index

The Dice Index, introduced by Lee R. Dice in 1945, is a well-established measure of similarity that evaluates the overlap between two sets. It is mathematically defined as:

$$\text{Dice}(A, C) = \frac{2|A \cap C|}{|A| + |C|}, \quad (2.5)$$

where A and C are the two sets being compared, $|A|$ and $|C|$ represent their respective sizes, and $|A \cap C|$ is the size of their intersection. The coefficient ranges from 0 (no overlap) to 1 (perfect match).

As illustrated in Figure 2.6, the Dice Index can be represented by a Venn diagram where the overlapping region B corresponds to the intersection between sets A and C . The calculation is visualized as twice the area of B divided by the total areas of A and C , highlighting its focus on shared regions.

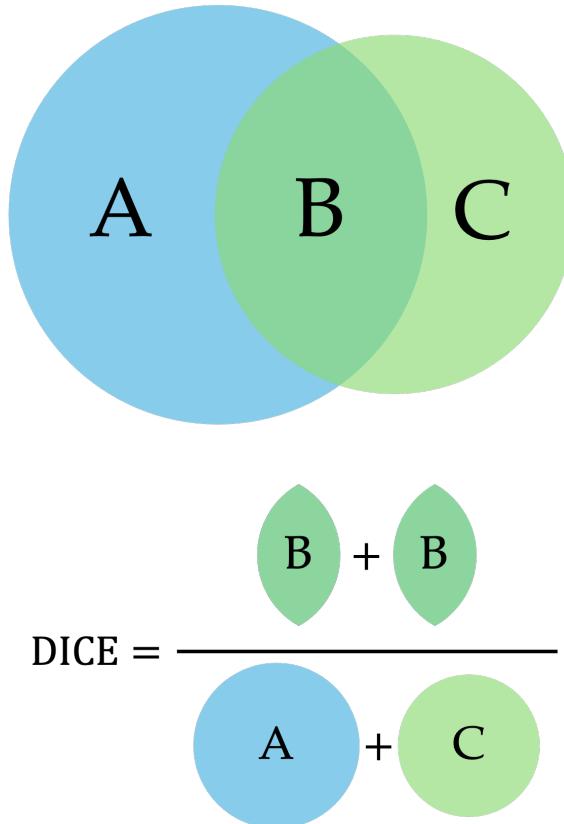


Figure 2.6: Illustration of the Dice Index calculation using a Venn diagram. The overlap region B represents the intersection between sets A and C . The Dice Index is calculated as twice the intersection area divided by the total area of both sets, emphasizing the shared elements between the sets.

In this study, the Dice Index is used alongside the Modified Hausdorff Distance (MHD) to provide a comprehensive comparison of binary masks. While the MHD captures edge-based structural differences, the Dice Index emphasizes volumetric overlaps (e.g., water-covered regions in the masks). This combination ensures that both structural and spatial relationships are represented, offering a robust framework for analyzing the distribution of water and dry areas in the experimental flume images. The Dice Index was computed using the implementation available in the Python library `scikit-image`, ensuring precision, reliability, and computational efficiency without the need for custom code.

2.4.3 Combined Matrix

To integrate the complementary information provided by the Modified Hausdorff Distance (MHD) and the Dice Index, we constructed a combined matrix by concatenating the two distance matrices rather than merging them through weighted averaging. This decision was motivated by the absence of prior knowledge regarding the relative importance of

edge-based (MHD) and area-based (Dice) similarities and by the need to avoid introducing arbitrary weighting factors.

Specifically we computed two separate distance matrices, one based on normalized MHD values and another based on Dice-derived dissimilarities. These matrices were concatenated along the feature axis to form a single composite matrix that jointly represents both structural and spatial aspects of morphological similarity. This composite matrix effectively doubles the feature space, allowing downstream analyses such as dimensionality reduction and clustering to learn the relative contributions of each metric without manual tuning.

By adopting this approach we ensured that both edge configurations and volumetric overlaps were equally available to the clustering algorithms, while preserving the flexibility to let data-driven methods such as principal component analysis or UMAP determine the most informative patterns. This concatenated matrix served as the input for the clustering pipeline described in Section 2.6.3, providing a robust foundation for defining the discrete morphological states used in the Markov modeling framework.

2.5 Dimensionality Reduction

2.5.1 What is dimensionality reduction?

Dimensionality reduction provides a practical way to analyze complex datasets such as the high-dimensional morphological information extracted from flume images. In fluid mechanics, Proper Orthogonal Decomposition (POD), which is essentially PCA applied to snapshot data, is widely used to extract energetic modes and build reduced-order models (Cammilleri et al., 2013; Taira et al., 2020). In the same spirit, here we project data onto low-dimensional spaces to enable efficient visualization, interpretation, and computation while preserving the structures most relevant to the task (Fukunaga, 1990; van der Maaten et al., 2009).

A key motivation to do this is the "curse of dimensionality". As the number of features grows, distance metrics lose discriminative power, noise can overwhelm signal, and algorithms become unstable or prohibitively expensive. This phenomenon, formalized by Bellman (1961), is directly relevant here because the combined MHD and Dice descriptors yield feature vectors with thousands of components, many of them low-variance or redundant, which complicates downstream analysis.

No single dimensionality-reduction method is universally best. Datasets differ in structure, noise, and how information is distributed, so methods make different trade-offs between preserving variance, neighborhood relations, or global geometry (van der Maaten et al., 2009). Some techniques emphasize global structure, while others prioritize local relationships or manifold assumptions.

In this study we compare four complementary approaches (i.e., PCA, MDS, t-SNE, and UMAP) to represent morphological patterns. Each highlights different aspects of the data. PCA provides a linear projection that retains maximal variance under orthogonal-

ity constraints; MDS seeks an embedding whose pairwise distances match the original dissimilarities, aiding interpretation of global spatial relationships; t-SNE concentrates on local neighborhood structure suited to visualization of nonlinear clusters; and UMAP targets both local and meso-scale structure with favorable runtime and scalability.

2.5.2 Principal Component Analysis (PCA)

Principal Component Analysis (PCA) is a linear dimensionality reduction technique that finds orthogonal directions of maximum variance called principal components. Projecting the data onto the first components reduces dimensionality while keeping most of the variability. In Figure 2.7 top left, PCA forms a broad cloud with partial mixing of colors, which reflects good capture of global variance but limited separation of classes. This behavior is typical for our morphological descriptors when the structure is only weakly linear.

Mathematically, given a zero-mean data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$, PCA solves the eigenvalue problem for the covariance matrix $\mathbf{C} = \frac{1}{n-1} \mathbf{X}^\top \mathbf{X}$

$$\mathbf{C} \mathbf{w}_k = \lambda_k \mathbf{w}_k, \quad \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p \geq 0,$$

where λ_k is the variance explained by the k -th principal component \mathbf{w}_k . The data are then projected as $\mathbf{Z} = \mathbf{X} \mathbf{W}_k$, with $\mathbf{W}_k = [\mathbf{w}_1, \dots, \mathbf{w}_k]$.

Before computing PCA, each feature of the composite MHD+Dice matrix is standardized to zero mean and unit variance

$$\tilde{x}_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j},$$

so that PCA ranks components by informative variance rather than by scale.

The cumulative explained variance

$$V(k) = \sum_{i=1}^k \frac{\lambda_i}{\sum_{j=1}^p \lambda_j}$$

is inspected and the smallest k with $V(k) \geq 0.99$ is retained. The resulting \mathbf{Z} serves as input for the following embeddings (MDS, UMAP and t-SNE) and for the clustering algorithms described later.

2.5.3 Multidimensional Scaling (MDS)

Multidimensional Scaling preserves pairwise distances when mapping data to a lower dimensional space. Unlike PCA, which emphasizes variance with a linear projection, MDS aims to reproduce dissimilarities. In Figure 2.7 top right, MDS arranges the samples in a roughly circular layout with moderate overlap. This representation helps read global relationships but separation of classes is still limited. MDS accepts many distance choices, which is useful for morphology, although its cost grows quickly with dataset size and results can depend on the selected metric.

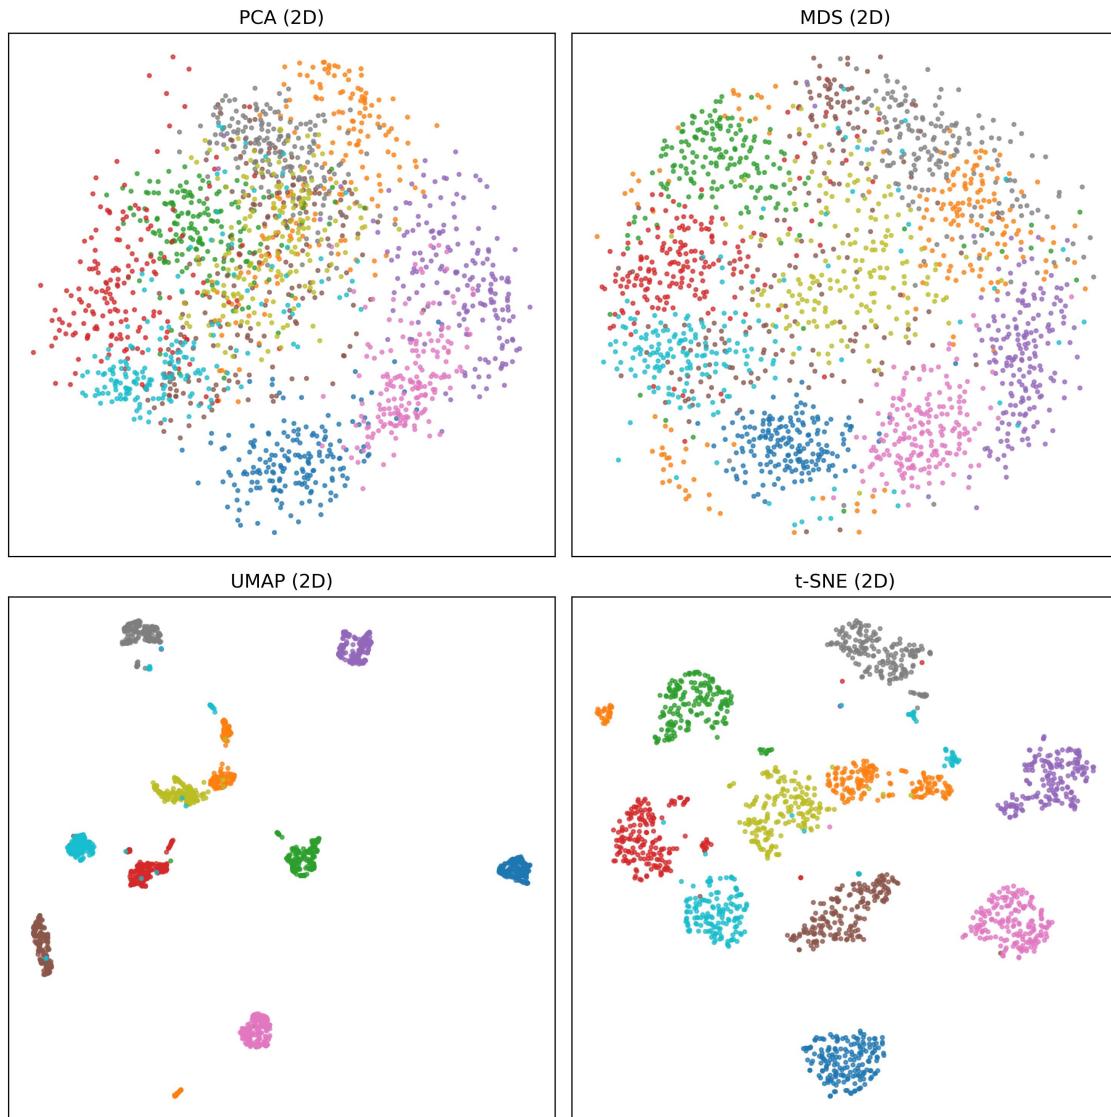


Figure 2.7: Comparison of dimensionality reduction methods on the Digits dataset. Panels show 2D projections using PCA top left, MDS top right, UMAP bottom left, and t-SNE bottom right. Colors denote digit classes. PCA and MDS give diffuse arrangements with gradual mixing. UMAP and t-SNE produce compact groups that are easier to cluster.

2.5.4 t-Distributed Stochastic Neighbor Embedding (t-SNE)

t-SNE models local neighborhoods and keeps them coherent after projection, which is ideal when the data form tight, non-linear clusters. In Figure 2.7 bottom right, t-SNE produces very clean group separation with sharp boundaries. This clarity is valuable for defining discrete morphological states and for clustering. The Student- t kernel alleviates the crowding problem and increases contrast between nearby groups, often revealing fine substructure that other methods merge. Results can change with perplexity and global distances are not metric faithful, yet within-cluster geometry and neighborhood purity are often best preserved with t-SNE on our descriptors.

2.5.5 Uniform Manifold Approximation and Projection (UMAP)

UMAP builds a graph of local neighborhoods and optimizes a low-dimensional embedding that balances local and broader structure. In Figure 2.7 bottom left, UMAP yields compact groups while maintaining meaningful spacing among them, which helps interpret relations between clusters and gradual transitions. UMAP is usually faster and scales well, though performance depends on hyperparameters such as the neighborhood size. Using both embeddings is helpful. t-SNE excels at isolating crisp clusters for state definition, whereas UMAP complements it by arranging those clusters with more interpretable inter-cluster geometry.

2.6 Density-Based Clustering

Clustering techniques are helpful for identifying natural groupings in data without requiring predefined categories. These methods could be particularly useful for analyzing complex datasets, such as the morphological patterns observed in our flume experiments. In this study, two density-based clustering algorithms were employed: DBSCAN (Density-Based Spatial Clustering of Applications with Noise) and HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise). Both techniques are good at detecting clusters of arbitrary shapes and handling noisy data, which makes them well suited for embeddings where point distributions are curved and densities vary.

Each of these algorithms operates under unique principles, which influence their behavior and performance in different contexts. Below, we provide a detailed description of DBSCAN and HDBSCAN, focusing on their mechanisms and on what the side-by-side result in Figure 2.8 reveals about their practical use in this work.

2.6.1 DBSCAN: Density-Based Spatial Clustering of Applications with Noise

Density-Based Spatial Clustering of Applications with Noise (DBSCAN) groups points according to local neighborhoods in the feature space (Ester et al., 1996). The algorithm partitions the dataset into dense regions separated by sparser areas, which is effective for detecting clusters of arbitrary shapes and sizes. DBSCAN defines two key parameters: ϵ and minPts. Epsilon specifies the maximum distance between two points for them to be considered neighbors, and minPts represents the minimum number of points required to form a dense region or cluster. A point is classified as a core point if its ϵ -neighborhood contains at least minPts points, which enables the iterative expansion of clusters by including other density-reachable points.

Figure 2.8 (left) shows DBSCAN applied to a two-dimensional synthetic set generated with `make_moons` ($n=500$, $\text{noise} = 0.2$). Parameters were fixed at $\epsilon = 0.1$ and $\text{minPts} = 5$, which are consistent with the code used to create the figure. Under these settings the algorithm identifies many small clusters that correspond to pockets of high local density along each moon and it assigns a large fraction of points to the noise label -1 . This outcome illustrates two typical behaviors. First, a single global ϵ can fragment curved manifolds when density varies along the structure. Second, the noise fraction can grow quickly when the chosen ϵ is smaller than the typical spacing in sparser parts of the data. In

our morphological embeddings, which exhibit gradients in point density across the state manifold, this sensitivity implies that DBSCAN is best used for rapid probing, parameter sweeps, and identification of conservative cluster cores rather than as the sole basis for defining final states.

When using DBSCAN, several challenges and considerations should be kept in mind. The selection of appropriate parameter values for ϵ and minPts is crucial, as it significantly affects the results. Determining useful values often requires experimentation and domain knowledge. DBSCAN is sensitive to variations in data density and noise levels, which may necessitate careful preprocessing of the dataset to ensure meaningful results (Schubert et al., 2017). Despite these challenges, DBSCAN remains a powerful tool for identifying spatial patterns and structures within datasets. Its simplicity and ability to detect nonconvex shapes make it valuable in exploratory phases, including the early analysis of our flume-derived state clouds.

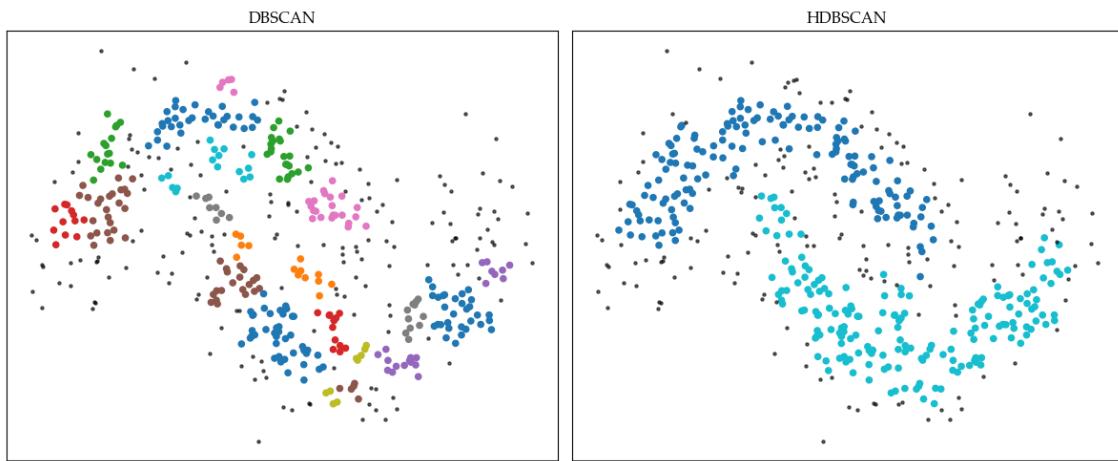


Figure 2.8: Comparison of DBSCAN (left) and HDBSCAN (right) on a synthetic two-moons dataset generated with `make_moons` ($n=500$, $\text{noise} = 0.2$). DBSCAN parameters: $\epsilon = 0.1$, $\text{minPts} = 5$. HDBSCAN parameter: $\text{min_cluster_size} = 10$. Colors indicate cluster labels and points with label -1 are treated as noise. DBSCAN fragments the curved manifolds into many small clusters and classifies many points as noise under a single global density threshold. HDBSCAN recovers the two dominant arcs as coherent clusters while leaving sparse outliers as noise, which illustrates its adaptation to varying density.

2.6.2 HDBSCAN: Hierarchical Density-Based Spatial Clustering of Applications with Noise

Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) extends DBSCAN to datasets where densities vary across clusters (Campello et al., 2013, 2015). Instead of fixing a single ϵ , HDBSCAN evaluates structure over a range of density levels. It does so by computing core distances to the k -nearest neighbors, transforming them into a mutual reachability graph, and building a hierarchy of cluster splits as the density threshold changes. Cluster selection is then based on stability, which measures how persistent a cluster is across that hierarchy. This process reduces the sensitivity to a single neighborhood scale and enables robust extraction of clusters that differ in internal density (McInnes & Healy, 2017).

The right panel of Figure 2.8 uses `min_cluster_size = 10` and shows that HDBSCAN recovers two large, nonconvex clusters that follow the moons while assigning only the sparsest points to the noise label -1 . This matches the intended use in our application. State clouds obtained from image-derived embeddings frequently present curved structures and density gradients, which are shaped by sampling frequency, dwell times, and the intrinsic variability of morphological states. A hierarchical evaluation of density better respects these gradients and avoids the over-fragmentation observed with a single global ϵ . In addition, HDBSCAN provides measurement of cluster membership strengths that help identify core exemplars and boundary points, which is useful when we later select representative medoids of each morphological state.

HDBSCAN’s ability to extract clusters of varying densities and its reduced dependence on user-defined neighborhood radii make it particularly suitable for high-dimensional datasets such as the morphological data analyzed in this study. By combining stability with density-based principles, HDBSCAN enhances the interpretability and robustness of clustering results. As with DBSCAN, the algorithm distinguishes noise points, which yields cleaner state definitions for downstream estimation of transition matrices and dwell-time distributions.

It is worth noting that trying both DBSCAN and HDBSCAN can provide complementary insights. DBSCAN’s simpler parameterization and efficiency make it suitable for datasets with more uniform densities or with well-separated cores. HDBSCAN excels in more complex scenarios where clusters have varying densities or hierarchical relationships. Since no single method works optimally for all datasets, applying both ensures a more comprehensive exploration of clustering structures. In this work we relied on DBSCAN to probe parameter ranges and to identify conservative cluster seeds, and we used HDBSCAN to obtain the final partitions that align with the continuous, curved manifolds observed in the embeddings and that support stable, morphology-based state definitions for the Markov model.

2.6.3 Metrics for Cluster Quality Assessment

To assess the quality of the clustering solutions obtained from the grid search, we computed three widely used internal validation indices: the Silhouette coefficient, the Davies–Bouldin index, and the Calinski–Harabasz index. Additionally, we calculated the fraction of points labelled as noise, which is particularly relevant for density-based clustering algorithms.

The **Silhouette coefficient** measures the cohesion and separation of the clusters. For each point, it compares the mean intra-cluster distance (how close the point is to other points in its cluster) to the mean nearest-cluster distance (how far it is from points in the closest neighboring cluster). Values range between -1 and 1 , where higher values indicate better defined and more separated clusters.

The **Davies–Bouldin index** quantifies the average similarity between each cluster and its most similar one, based on the ratio of within-cluster dispersion to between-cluster separation. Lower values reflect better clustering solutions, characterized by compact, well-separated clusters.

The **Calinski–Harabasz index**, also known as the variance ratio criterion, compares the between-cluster dispersion to the within-cluster dispersion. Higher values are preferred, indicating well-separated and dense clusters.

The **fraction of noise points** corresponds to the proportion of data points that were not assigned to any cluster and instead labelled as noise. This metric is crucial when using algorithms like DBSCAN or HDBSCAN, as it helps balance the trade-off between detecting meaningful clusters and overfiltering the data.

These metrics were computed for all combinations of hyper-parameters tested in the grid search. They provide complementary perspectives on cluster quality because the Silhouette coefficient emphasizes individual point assignments, while the Davies–Bouldin and Calinski–Harabasz indices assess overall cluster configuration. It is crucial to note that none of these metrics should be interpreted in isolation. For example, a high Silhouette score might trivially result from assigning each point to its own cluster, while low noise fractions could reflect overly coarse clustering.

As explained later, the final selection of clustering configuration involved a combined approach of visual inspection of cluster patterns (see Figure 3.8) and critical examination of the numerical metric matrix (Table 3.3). This balanced evaluation ensured that the chosen clustering solution was both quantitatively robust and qualitatively meaningful for subsequent analyses.

2.6.4 Metrics for Cluster Morphological Description

To quantitatively describe the morphological characteristics of each cluster, we computed five planform metrics capturing key aspects of channel geometry and complexity.

- (a) **Wetted width** — Mean cross-stream distance between the wetted banks, providing a first-order measure of channel scale.
- (b) **Wetted area** — Horizontal area occupied by water in the binary mask, reflecting the extent of active flow.
- (c) **Sinuosity** — Ratio of centerline length to straight-line valley length, quantifying the degree of channel meandering.
- (d) **Braiding Index (BI)** — Defined as the average number of independent wetted threads per cross-section, the BI is a classical and widely used metric to characterize braided rivers (Egozi & Ashmore, 2008). However, it is purely a count and does not account for differences in channel scale (e.g., discharge or width), making it sensitive to observational resolution. As image resolution increases, BI may increase artificially by detecting minor channels that contribute little to the total flow (Tejedor et al., 2022).

(e) **Entropy-based Braiding Index (eBI)** — To address BI limitations, we also use the entropic braiding index (eBI), introduced by Tejedor et al. (2022), which incorporates channel heterogeneity using Shannon entropy. Given N channels at a cross-section, with q_i denoting the discharge (or width) of channel i and $Q = \sum_{i=1}^N q_i$ the total discharge, the Shannon entropy H is computed as:

$$H = - \sum_{i=1}^N \frac{q_i}{Q} \log_2 \left(\frac{q_i}{Q} \right). \quad (2.6)$$

The eBI is then defined as:

$$eBI = 2^H, \quad (2.7)$$

which can be interpreted as the *effective number of channels*, that is, the equivalent number of equally sized channels conveying the same entropy. When all channels have equal discharge, $eBI = BI$; when a few channels dominate, $eBI < BI$, reflecting the reduced complexity. Additionally, the ratio BI/eBI quantifies the degree of channel disparity, helping to differentiate braided from anastomosed systems and assess cross-sectional stability under flow variability.

Together, these metrics provide a comprehensive morphological description of each cluster, capturing both spatial scale and structural complexity, while the combined use of BI and eBI ensures a robust, resolution-independent assessment of the multi-thread river dynamics.

2.7 Probabilistic modeling with Markov chains and Monte Carlo

As previously mentioned, our goal is to model the temporal evolution of our experiment as a Markov model. To achieve this, the Markov chain framework and Monte Carlo simulations are two fundamental tools used to model and analyze the evolution of morphological states in the experimental flume. These methods allow for a probabilistic representation of the system's dynamics and enable the generation of synthetic series that reflect the natural variability of sediment transport processes.

2.7.1 Markov Chains

A Markov Chain provides a compact probabilistic framework to model how the system transitions between discrete morphological states over time. In our context, we assume the flume can occupy one of n well-defined configurations $\{S_1, S_2, \dots, S_n\}$, each representing a characteristic spatial pattern. The key assumption is the Markov property, this is that the probability of transitioning to the next state depends only on the current state.

Formally, the process is described by a transition probability matrix \mathbf{P} , where each entry $P_{i,j}$ gives the probability of moving from state S_i to S_j in a single time step. By construction, the rows of \mathbf{P} sum to one:

$$\sum_{j=1}^n P_{i,j} = 1 \quad \forall i.$$

This formulation allows us to model the time evolution of the system as a sequence $\{X_t\}$, where X_t denotes the state at time t . The Markov property is then written as:

$$\Pr(X_{t+1} = S_j \mid X_t = S_i, X_{t-1}, \dots, X_0) = \Pr(X_{t+1} = S_j \mid X_t = S_i).$$

Figure 2.9 illustrates a Markov chain with $n = 5$ morphological states. Each node represents a state, and each directed edge indicates a permitted transition between states, labeled with its corresponding probability $P_{i,j}$. The absence of an arrow implies that a transition is not allowed (i.e., $P_{i,j} = 0$).

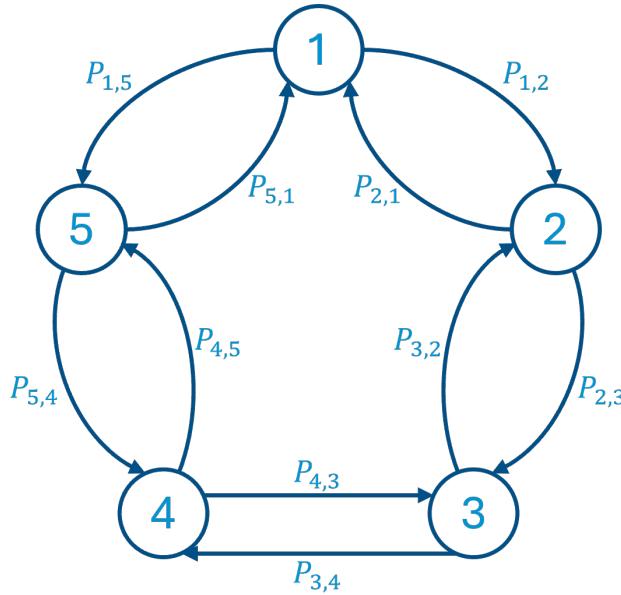


Figure 2.9: Five-state Markov Chain derived from flume data. Arrows indicate allowed transitions, and labels show the associated probabilities $P_{i,j}$.

For example, from state S_3 , the system can only move to states S_2 and S_4 , with respective probabilities $P_{3,2}$ and $P_{3,4}$. These allowed transitions define the structure of the system's evolution.

Once our state space has been discretized through clustering, the transition matrix \mathbf{P} is built by counting the number of observed transitions between states. Specifically, for each pair (i, j) , the raw count $C_{i,j}$ records how often the system moved from S_i to S_j . The final transition probabilities are then obtained by row-normalizing the count matrix:

$$P_{i,j} = \frac{C_{i,j}}{\sum_j C_{i,j}}.$$

This matrix summarizes the observed transition dynamics and serves as the core structure for generating synthetic sequences using Monte Carlo simulations in the next stage of our modeling framework.

2.7.2 Continuous-Time Dynamics: CTMC

The discrete-time Markov chain in Section 2.7.1 captures which state-to-state moves are possible and with what probabilities at the sampling grid. In the experiment, however, dwell times within each state are not fixed. Empirically they vary and are well approximated by an exponential distribution at the state level. To represent this, we model the evolution as a continuous-time Markov chain (CTMC) defined by a jump kernel without self-transitions and by state-dependent exit rates.

Let the system occupy one of n states $\{S_1, \dots, S_n\}$ and let $X(t)$ denote the state at time t . When the process is in state S_i it remains there for a random holding time τ_i that follows an exponential distribution with rate $\lambda_i > 0$. At the end of the holding time, the process jumps to a new state S_j with probability p_{ij} for $j \neq i$. The pair (λ_i, p_{ij}) fully specifies the CTMC.

Generator matrix The CTMC is described by the generator $\mathbf{Q} = [q_{ij}]$ with

$$q_{ij} = \lambda_i p_{ij} \quad (i \neq j), \quad q_{ii} = -\lambda_i, \quad \sum_{j \neq i} p_{ij} = 1,$$

so that each row of \mathbf{Q} sums to zero. The discrete-time transition matrix $\mathbf{P} = [p_{ij}]$ is the jump kernel of the embedded chain and has zero diagonal.

2.7.2.1 Estimation from observed trajectories

Let C_{ij} be the number of observed transitions from S_i to S_j with $j \neq i$. Let $N_i = \sum_{j \neq i} C_{ij}$ be the total number of exits from S_i . Let T_i be the total time spent in S_i across the record. The maximum-likelihood estimates are

$$\hat{\lambda}_i = \frac{N_i}{T_i}, \quad \hat{p}_{ij} = \frac{C_{ij}}{N_i} \quad (j \neq i), \quad \hat{q}_{ij} = \hat{\lambda}_i \hat{p}_{ij}, \quad \hat{q}_{ii} = -\hat{\lambda}_i.$$

Units follow the time base used to compute T_i . If some $N_i = 0$, we either merge S_i with its nearest state in the clustering space or apply a small ridge to the counts before normalizing.

2.7.2.2 Simulation of CTMC paths

To draw a synthetic path:

1. Set the initial state $X(0)$ from a chosen distribution (empirical occupancy or stationary distribution).
2. At state S_i , sample a holding time $\tau \sim \text{Exp}(\hat{\lambda}_i)$.
3. Draw the next state S_j with probability \hat{p}_{ij} for $j \neq i$, set $X(t + \tau) = S_j$, and iterate.

When a regular output grid is needed, we hold the last state constant between jump times and report the state at the grid instants.

2.7.2.3 Stationary distribution

If the chain is irreducible, the stationary distribution π solves

$$\pi^\top \mathbf{Q} = \mathbf{0}^\top, \quad \sum_{i=1}^n \pi_i = 1, \quad \pi_i \geq 0.$$

We compute π either from \mathbf{Q} or from the embedded chain with state-dependent residence times.

2.7.2.4 Mean first-passage time (MFPT)

Let m_{ij} be the expected time to hit state S_j for the first time starting from S_i with $i \neq j$. Collect these values in the vector $\mathbf{m}^{(j)} = (m_{1j}, \dots, m_{nj})^\top$. The MFPT satisfies the linear system

$$\sum_{k=1}^n q_{ik} m_{kj} = -1 \quad \text{for } i \neq j, \quad m_{jj} = 0.$$

Equivalently, removing row and column j from \mathbf{Q} gives $\mathbf{Q}_{-j,-j} \mathbf{m}_{-j}^{(j)} = -\mathbf{1}$. For a target set $A \subset \{1, \dots, n\}$, set $m_{aj} = 0$ for all $a \in A$ and solve on the complement.

2.7.2.5 Jump entropy rate

We summarize dynamical uncertainty with the jump entropy rate

$$H_{\text{rate}} = \sum_{i=1}^n \pi_i \lambda_i H(\mathbf{p}_i), \quad H(\mathbf{p}_i) = - \sum_{j \neq i} p_{ij} \log p_{ij},$$

expressed in nats per unit time when the natural logarithm is used. The factor λ_i weights how often departures occur from S_i , while $H(\mathbf{p}_i)$ measures the unpredictability of the destination conditional on leaving S_i . We estimate H_{rate} by using $\hat{\pi}$, $\hat{\lambda}_i$, and $\hat{\mathbf{p}}_i$.

2.7.3 Bootstrap Method

The bootstrap is a resampling technique used to estimate the uncertainty in a statistic by drawing repeated samples from the observed data. In our context, it is employed to account for intra-state variability in sediment transport while preserving the transition dynamics imposed by the Markov Chain.

Let $\{q_1, q_2, \dots, q_m\}$ be the set of observed sediment transport values associated with a given morphological state S_i . A single bootstrap replicate consists of sampling m values from this set with replacement and computing the desired statistic for that time step.

The procedure to generate a synthetic transport series for a given simulated state sequence is as follows:

1. For each time step t , identify the current morphological state $S_i = X_t$.
2. Draw a bootstrap sample from the empirical transport values corresponding to state S_i .
3. Assign the resulting value as the transport rate $Q_s(t)$ for that time step.

This process is repeated for every time step in the sequence, and the entire operation is applied across multiple bootstrap replicates. Specifically, for each state sequence, we generate $N_{\text{boot}} = 1000$ synthetic realizations of $Q_s(t)$, capturing the range of possible sediment transport trajectories consistent with the observed intra-state variability.

By combining these resampled values with the state transitions generated from the Markov model, the bootstrap method ensures that both inter-state dynamics and intra-state transport variability are reflected in the synthetic trajectories. This step provides the foundation for generating a comprehensive ensemble of sediment transport scenarios that will later be evaluated and filtered using Bayesian criteria.

2.7.4 Monte Carlo Simulations

Monte Carlo simulations are used to generate synthetic realizations of the system's temporal evolution by coupling a continuous-time Markov chain (CTMC) for state sequencing with bootstrap-generated transport rates. Each simulation produces a possible morphological and transport trajectory, enabling the exploration of variability and uncertainty in sediment dynamics.

Given the empirically estimated transition probability matrix \mathbf{P} from Section 2.7.1, we draw state sequences in continuous time by sampling exponential dwell times and then jumping according to \mathbf{P} . Specifically, each trajectory starts from a random initial state (unless otherwise specified). We draw one rate λ per trajectory from a prescribed range and, at the current state, sample a holding time $\tau \sim \text{Exp}(\lambda)$. To align with the 1-minute sampling used throughout the study, τ is rounded to the nearest minute and the state is held for that many steps before drawing the next state from the corresponding row of \mathbf{P} . This gives a 1 min step sequence $\{X_t\}$ that built by using dwell times from a distribution and reconfiguration through \mathbf{P} .

To assign transport values to each simulated sequence, we apply the bootstrap method described in Section 2.7.3. For every minute t , we draw $Q_s(t)$ with replacement from the empirical bag of sediment-transport observations associated with the active state X_t . This produces a synthetic time series that preserves both inter-state contrasts (differences in state means) and intra-state variability (within-state scatter).

By repeating this process across a large number of simulations, we construct an ensemble of synthetic trajectories. In this study, we generate $N_{\text{seq}} = 100000$ Markov state sequences, each spawning $N_{\text{boot}} = 1000$ transport replicates, for a total of 10^8 synthetic time series. Not all simulations are equally realistic; the next section describes an Approximate

Bayesian Computation (ABC) screening that evaluates each realization using moment tolerances and image-based similarity metrics, retaining only those that most closely resemble the original experimental data.

2.7.5 Bayesian Inference

Bayesian inference updates a prior distribution with new information contained in observed data. The updating rule is Bayes' theorem

$$P(H | D) = \frac{P(D | H) P(H)}{P(D)},$$

where H is a hypothesis about the system, D is the data, $P(H)$ is the prior belief, $P(D | H)$ is the likelihood of observing D if H is true, and $P(H | D)$ is the posterior belief after seeing the data.

In this work the hypothesis space is the set of all synthetic sediment-transport time series that can be produced by the Markov model.

1. **Prior.** A total of $N_{\text{seq}} = 100000$ state sequences are simulated from the empirical transition structure. Each sequence spawns $N_{\text{boot}} = 1000$ bootstrap replicas of the transport rate $Q_s(t)$, yielding 10^8 prior trajectories.
2. **Likelihood (ABC filter).** A trajectory is accepted if its mean and variance satisfy

$$|\mu^* - \mu_{\text{exp}}| \leq \varepsilon_\mu, \quad |\sigma^{2*} - \sigma_{\text{exp}}^2| \leq \varepsilon_\sigma,$$

with fixed tolerances ε_μ and ε_σ set from the experimental record. These tolerances are kept constant across all comparisons and sensitivity checks.

3. **Posterior.** The accepted subset is treated as the posterior ensemble and is used for all downstream analyses. The acceptance rate is reported alongside results and sensitivity analyses.

2.7.6 Decision-tree surrogate as a conservative prefilter

We use a shallow decision tree to screen synthetic state sequences using image-only information before any costly transport computation. The target of the classifier is whether a sequence would match the first two moments of the experimental transport record. For each candidate sequence A drawn from the prior (Section 2.7.5), we compute the expected transport mean $\hat{\mu}(A)$ and variance $\hat{\sigma}^2(A)$ implied by its state visits and transitions (using the calibrated state-wise summaries). We then compare these values with the experimental moments μ_{exp} and σ_{exp}^2 and assign a binary label,

$$y(A) = \mathbf{1}\left\{ \left| \hat{\mu}(A) - \mu_{\text{exp}} \right| \leq \varepsilon_\mu \text{ and } \left| \hat{\sigma}^2(A) - \sigma_{\text{exp}}^2 \right| \leq \varepsilon_\sigma \right\},$$

where ε_μ and ε_σ are the predefined tolerances (absolute or relative, as reported). Only sequences whose expected mean and variance are within tolerance are labeled $y = 1$. This produces an imbalanced dataset with rare positives.

Each sequence is described by six inexpensive, image-only descriptors computed against the experimental reference: state-frequency difference d_{freq} , Frobenius distance between transition matrices d_F , Jensen–Shannon divergence of pooled dwell times $d_{\text{JS}}^{\text{dwell}}$, diagonal distance d_{diag} for self-transitions, length-normalized Levenshtein distance d_{lev} between state strings, and average-row-entropy difference d_{entr} . We standardize all features with training-set mean and variance and use a small ridge in transition counts to avoid undefined entropies or divergences.

The classifier is a depth-3 decision tree with Gini impurity, class weights to counter imbalance, and a minimum number of samples per leaf to prevent overfitting. We split the data in a stratified way, tune hyperparameters by cross-validation to maximize precision subject to a recall floor, and optionally calibrate the tree’s probabilities on the validation set. An operating threshold τ on the predicted probability $\hat{p}(x)$ is chosen so that precision is at least 0.99; this enforces conservative behavior. In deployment, only sequences with $\hat{p}(x) \geq \tau$ pass to full transport assignment and ABC; the rest are discarded early.

Beyond filtering, the tree also reveals which descriptors, or combinations of them, best predict acceptable moments. We report feature importances and the learned split rules, which indicate the metrics that most strongly separate sequences that achieve mean and variance within tolerance. This analysis identifies the most informative descriptors for forecasting moment-level acceptance using only state-sequence information.

2.7.7 Cross-experiment comparison of transition matrices

We compare the transition structure across experiments by measuring the distance between their discrete-time transition matrices. Let $\mathbf{P}^{(A)} \in \mathbb{R}^{n \times n}$ and $\mathbf{P}^{(B)} \in \mathbb{R}^{n \times n}$ be the row-stochastic matrices estimated for experiments A and B (rows sum to one). The baseline discrepancy is the Frobenius norm

$$d_F(A, B) = \left\| \mathbf{P}^{(A)} - \mathbf{P}^{(B)} \right\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^n \left(P_{ij}^{(A)} - P_{ij}^{(B)} \right)^2}.$$

Estimation protocol

For each experiment we estimate \mathbf{P} from transition counts C_{ij} as in Section 2.7.1, add a small ridge to zero rows if needed, and row-normalize. We then align states (if required) and compute d_F and the diagnostics above. For multiple experiments we assemble a pairwise distance matrix and summarize with medians and interquartile ranges.

2.7.8 Temporal Diagnostics: Autocorrelation and Power Spectrum

We use two standard diagnostics to compare the temporal structure of experimental and synthetic time series: the autocorrelation function (ACF) in the time domain and the power spectral density (PSD) in the frequency domain. Both are computed on the de-meaned series and on a uniform time grid.

Autocorrelation function (ACF)

Let $\{Q_t\}_{t=1}^T$ be the sediment-transport series sampled every Δt . Denote $\bar{Q} = \frac{1}{T} \sum_{t=1}^T Q_t$. The sample ACF at lag k (with $k \in \{0, 1, \dots, L\}$) is

$$\rho(k) = \frac{\sum_{t=1}^{T-k} (Q_t - \bar{Q})(Q_{t+k} - \bar{Q})}{\sum_{t=1}^T (Q_t - \bar{Q})^2}.$$

We report $\rho(k)$ over a fixed lag range and use the integrated autocorrelation

$$\text{IACF}(L) = 1 + 2 \sum_{k=1}^L \rho(k),$$

as a scalar summary of persistence. For visual comparison we plot the empirical $\rho(k)$; for quantitative comparison we compute a weighted L^2 distance between curves over the chosen lag range.

Power spectral density (PSD)

The PSD quantifies how the variance of a stationary series is distributed over temporal frequencies. For an evenly sampled record with sampling interval Δt (1 min in our case), the one-sided PSD $S(f)$ integrates to the sample variance,

$$\text{Var}(x) = \int_0^{f_N} S(f) df,$$

where $f_N = 1/(2\Delta t)$ is the Nyquist frequency. We estimate $S(f)$ with Welch's method, which averages periodograms from overlapping, windowed segments to reduce variance. The series is demeaned and linearly detrended, split into segments of length L with 50% overlap, and weighted with a Hann window w . For each segment we compute the discrete Fourier transform, form the window-corrected periodogram, then average across segments. The estimate is converted to a one-sided spectrum by doubling power for positive frequencies (excluding $f = 0$ and f_N). The resulting units are (series units)²/Hz.

2.7.9 Extreme-Value Validation: Return-Period Curves

This subsection evaluates whether synthetic series reproduce the tail behavior of the experimental transport. We use return-period curves built from peaks-over-threshold events.

Let $\{Q_t\}_{t=1}^T$ be the transport series sampled every Δt . We detect extreme events using a peaks-over-threshold (POT) scheme with declustering:

1. Choose a high threshold u (e.g., the empirical q -quantile with $q \in [0.95, 0.995]$).
2. Identify all indices with $Q_t > u$. Consecutive exceedances within a run length τ_{run} are grouped into one cluster.
3. For each cluster, keep the local maximum Q_k^* and its time stamp. The set $\{Q_k^*\}_{k=1}^{N_u}$

are the POT event magnitudes.

For a grid of levels $x \geq u$, define $N_{>x}$ as the number of POT events with $Q_k^* > x$. With total observation time $T_{\text{obs}} = T \Delta t$, the empirical return period at level x is

$$\hat{T}(x) = \frac{T_{\text{obs}}}{N_{>x}},$$

with units of time. Equivalently, if working on the regular grid without declustering, the exceedance probability is $\hat{p}(x) = N_{>x}/T$ and $\hat{T}(x) = \Delta t/\hat{p}(x)$. For uncertainty bands, we treat $N_{>x}$ as Poisson with rate $\nu(x) = N_{>x}/T_{\text{obs}}$ and report exact Poisson confidence intervals, mapped to $T(x) = 1/\nu(x)$.

2.7.10 Variance decomposition and the fraction explained by state switching

We decompose the variability of the transport series Q_t into a within-state component and a between-state component using the law of total variance. Let $S_t \in \{1, \dots, n\}$ be the morphological state at time t . Define

$$\mu_i = \mathbb{E}[Q_t \mid S_t = i], \quad \sigma_i^2 = \text{Var}(Q_t \mid S_t = i), \quad \pi_i = \Pr(S_t = i).$$

The total variance satisfies

$$\text{Var}(Q_t) = \underbrace{\mathbb{E}[\text{Var}(Q_t \mid S_t)]}_W + \underbrace{\text{Var}(\mathbb{E}[Q_t \mid S_t])}_B = \sum_{i=1}^n \pi_i \sigma_i^2 + \sum_{i=1}^n \pi_i (\mu_i - \bar{\mu})^2,$$

where $\bar{\mu} = \sum_{i=1}^n \pi_i \mu_i$. We report the fraction of variance explained by state switching as

$$\rho_B = \frac{B}{W + B}.$$

Estimators from a single time series

Let the record be sampled on a regular grid with step Δt . Denote by $\mathcal{T}_i = \{t : S_t = i\}$ the index set of time steps spent in state i , with size $N_i = |\mathcal{T}_i|$ and occupancy weight $\hat{\pi}_i = N_i / \sum_k N_k$. The statewise sample moments are

$$\hat{\mu}_i = \frac{1}{N_i} \sum_{t \in \mathcal{T}_i} Q_t, \quad \hat{\sigma}_i^2 = \frac{1}{N_i - 1} \sum_{t \in \mathcal{T}_i} (Q_t - \hat{\mu}_i)^2 \quad (N_i \geq 2).$$

The mixture mean and the two components are then

$$\hat{\mu} = \sum_{i=1}^n \hat{\pi}_i \hat{\mu}_i, \quad \hat{W} = \sum_{i=1}^n \hat{\pi}_i \hat{\sigma}_i^2, \quad \hat{B} = \sum_{i=1}^n \hat{\pi}_i (\hat{\mu}_i - \hat{\mu})^2,$$

and the fraction explained is $\hat{\rho}_B = \hat{B}/(\hat{W} + \hat{B})$. For runs with missing values or irregular sampling, replace N_i by the total time spent in state i , T_i , and use $\hat{\pi}_i = T_i / \sum_k T_k$.

Sensitivity to the number of states

To assess robustness with respect to state granularity, we recompute $\hat{\rho}_B$ across a grid of cluster counts n . For each n , we estimate $\hat{\pi}_i, \hat{\mu}_i, \hat{\sigma}_i^2$ and summarize $\hat{\rho}_B(n)$. We report trends and plateaus rather than a single value when changes in n lead to material differences.

3 | Braided Rivers Morphology as Markov Chains

Building on the methodological overview in section 2.1, this chapter implements the Markov modeling of braided-river planform using the experimental record. Our aim is to define a discrete, data-driven set of morphological states and to quantify how the system moves between them. The workflow follows the first three green boxes in Figure 2.1: image acquisition & processing, pattern analysis & state definition, and transition probability estimation. The remaining steps, namely Monte Carlo simulations and model validation, are addressed in the next chapter.

The analysis begins by deriving a measure of dissimilarity between time-ordered planform configurations, using complementary image-based metrics to capture different aspects of morphological change. These distances are then projected into a lower-dimensional space using non-linear reprojection techniques, followed by unsupervised clustering to identify recurrent patterns. The resulting groups are taken as the states of the Markov model, and are compared in terms of their morphology and sediment transport signatures to aid interpretation.

Finally, the temporal sequence of states is used to estimate the transition structure, including the discrete-time transition probability matrix and measures of state persistence. Each subsection concludes with a brief summary of its main findings and their relevance to the objectives stated in the Introduction.

3.1 Distance Matrix

To quantify the morphological dissimilarity between every pair of images in the experiment, we computed two complementary metrics: the Modified Hausdorff Distance (MHD), described in Section 2.4.1, and the Dice coefficient-based distance, detailed in Section 2.4.2. Each metric captures a different aspect of morphological change. MHD emphasizes geometric displacement between contours, while the Dice Index reflects the degree of areal overlap between wetted areas.

Figures 3.1 and 3.2 show the resulting normalized distance matrices. Figure 3.1 shows that the color scale ranges from dark purple (low distance, i.e., high geometric similarity) to bright yellow (large displacement between contours). The dark main diagonal reflects zero distance between each image and itself. Narrow dark bands parallel to the diagonal identify intervals of morphological persistence, while bright zones indicate sudden geometric rearrangements.

Figure 3.2 shows the normalized values range from dark purple (low similarity, i.e., poor overlap) to bright yellow (high Dice similarity). The diagonal is bright by construction. Off-diagonal bright blocks highlight periods where the wetted area remains largely unchanged, even across distant time steps, suggesting morphological stagnation in terms of wet surface extent.

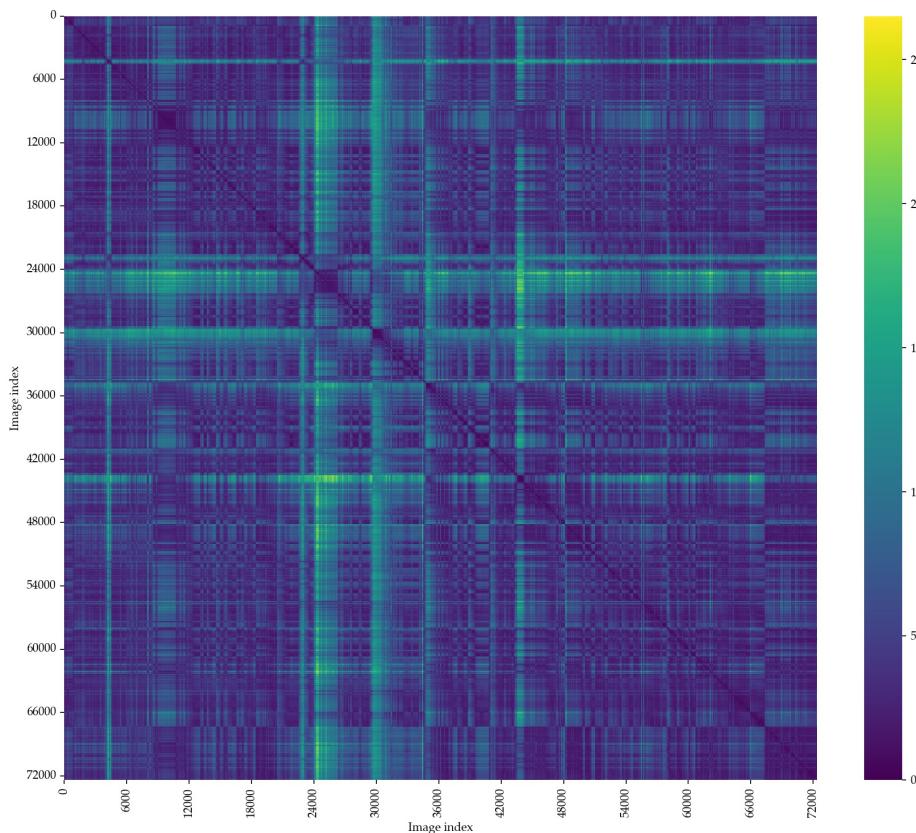


Figure 3.1: Normalized Modified Hausdorff Distance (MHD) matrix. Brighter colors represent larger geometric distances. Light diagonal bands indicate persistent morphology, whereas bright patches mark significant geometric change.

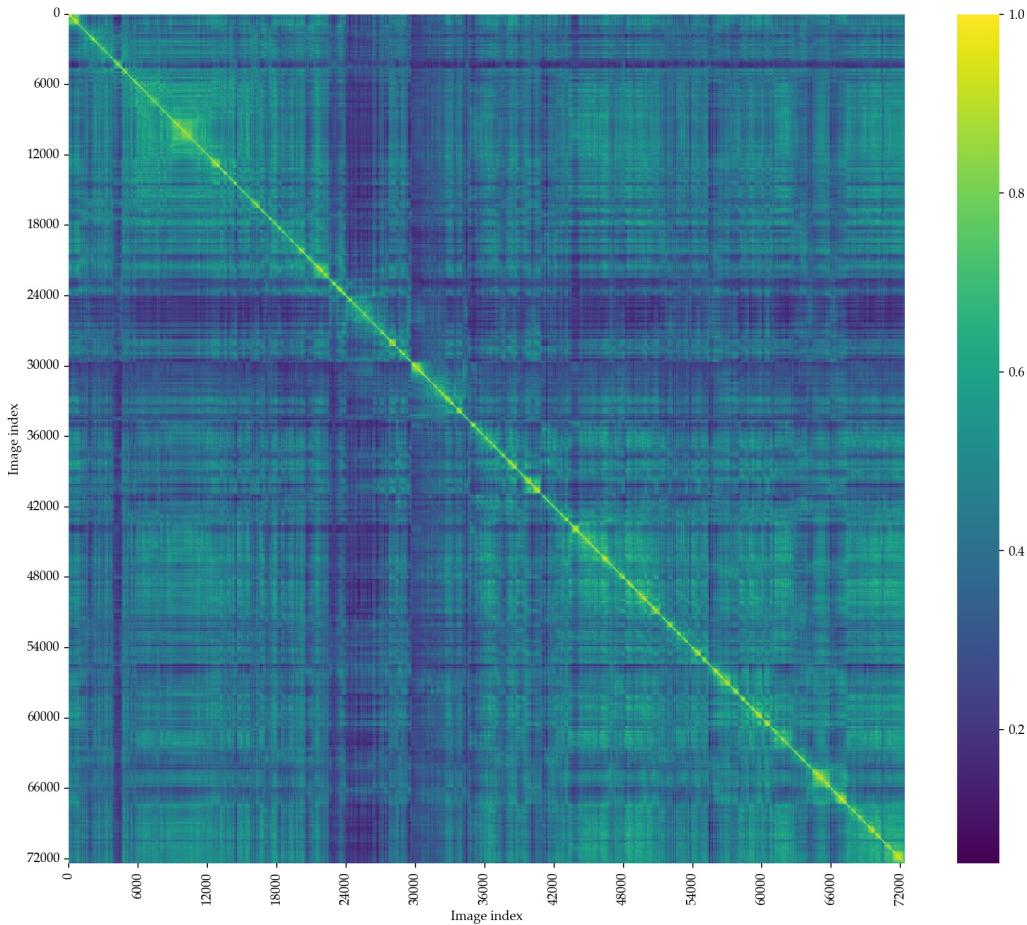


Figure 3.2: Normalized Dice-based distance matrix. Bright colors indicate high similarity in wetted areas. Bright off-diagonal blocks suggest consistent wet regions across distant time steps.

Although the two matrices appear broadly similar and both highlight intervals of persistence and transitions, they capture different morphological information. The Modified Hausdorff Distance (MHD) is particularly sensitive to contour displacement and geometric misalignment, whereas the Dice-based distance responds more directly to the spatial overlap of wetted regions (Karimi & Salcudean, 2020; Taha & Hanbury, 2015). This distinction means that evaluating with a single metric can leave blind spots: time steps that preserve overlap but shift the channel laterally yield low Dice penalties yet large MHD values, while local widening or narrowing can substantially affect Dice with only modest changes in MHD.

Concrete examples from our experiment illustrate this complementarity. A lateral migration of the channel increases MHD because contours are displaced, while also lowering Dice as the overlap between wetted regions decreases, even if total area changes little. Conversely, a localized expansion or contraction of the channel can substantially alter the Dice score without producing large geometric shifts of the contours, as long as the channel centerline remains fixed.

Because each metric emphasizes different aspects of change (i.e., shape deformation and structural rearrangement for MHD; areal continuity or fragmentation for Dice), their joint use is informative rather than redundant. In fact, in machine learning it is common to combine partially correlated “views” to improve robustness and stabilize downstream structure; multi-view learning shows that fusing complementary (and even mildly redundant) descriptors often enhances stability and generalization (Li et al., 2016; Xu et al., 2013). In the same spirit, keeping both distances reduces the risk of missing either boundary-dominated or overlap-dominated transitions, and leads to more reliable patterns in our dimensionality-reduction and clustering analyses.

Figure 3.3 provides a visual instance of this complementarity. The left panel depicts a slightly more laterally expanded, multi-thread morphology, whereas the right panel shows a narrower, single-thread channel. The wetted areas occupy opposite sides of the domain, yielding minimal overlap and clear contour misalignment. Consistently, both MHD and Dice flag this pair as the most dissimilar images, underscoring how their combination captures displacement-dominated and overlap-dominated contrasts in tandem.



Figure 3.3: Most dissimilar morphologies across the experiment according to both MHD and Dice. The left panel shows a slightly more laterally expanded, multi-thread pattern; the right panel shows a narrower, single-thread channel.

Taken together, the MHD and Dice distance matrices expose both periods of persistence and episodes of rapid reorganization, and they do so by emphasizing complementary aspects of change. We therefore combine these two signals into a single composite dissimilarity that balances contour shifts with wetted-area changes. This composite will serve as the input for the dimensionality-reduction stage and, later, for clustering into recurrent configurations that define the state space used by the Markov model.

3.2 Dimensionality Reduction

Now that we have both distance matrices, we assemble a single feature matrix by horizontally concatenating the normalized MHD and Dice matrices. Each column encodes the distance of all images to one reference image, so columns are the features. With n images, this yields an $n \times 2n$ matrix $X = [X_{\text{MHD}} \mid X_{\text{DICE}}]$. Before projection, we standardize features to place MHD- and Dice-derived columns on comparable scales.

Applying non-linear embeddings (UMAP, t-SNE) or even classical MDS directly to the $2n$ -feature space was computationally heavy and produced unstable layouts, a typical symptom of the “curse of dimensionality,” where distances lose contrast and redundant features accumulate (Köppen, 2000). To mitigate this, we use Principal Component Analysis (PCA) as a preprocessing step to compress variance into a smaller set of components, reducing noise and redundancy and yielding more stable, interpretable downstream embeddings.

3.2.1 Principal Component Analysis (PCA)

PCA was performed after column-wise standardisation of the composite matrix (zero mean, unit variance). The cumulative explained-variance curve (Figure 3.4) shows a steep rise that flattens quickly: the first 67 principal components already capture 99% of the total variance, a dramatic reduction from the original ~ 144000 features. Retaining only these 67 components thus preserves nearly all information while reducing dimensionality by more than four orders of magnitude, greatly accelerating the following stages of the process.

Because the mathematical foundations and numerical implementation of PCA were detailed earlier (Section 2.5.2), we limit the discussion here to the outcome:

- Number of retained components: $k = 67$ (99% cumulative variance).
- Transformation matrix $\mathbf{W} \in \mathbb{R}^{2n \times k}$, used to project the composite feature matrix into a k -dimensional space.
- Projected data matrix $\mathbf{X}_{\text{PCA}} \in \mathbb{R}^{n \times k}$, which serves as input for the non-linear projections analyzed in Section 3.2.2.

In summary, retaining $k = 67$ components preserves 99% of the variance while compressing the feature space from about 1.44×10^5 features to 67, which mitigates distance degeneracy, stabilizes following non-linear projections, and reduces computation. We therefore carry forward the PCA scores as standardized inputs for the next subsection, where we compare MDS, t-SNE, and UMAP across broad parameter sweeps to select the projection that best supports subsequent clustering into states.

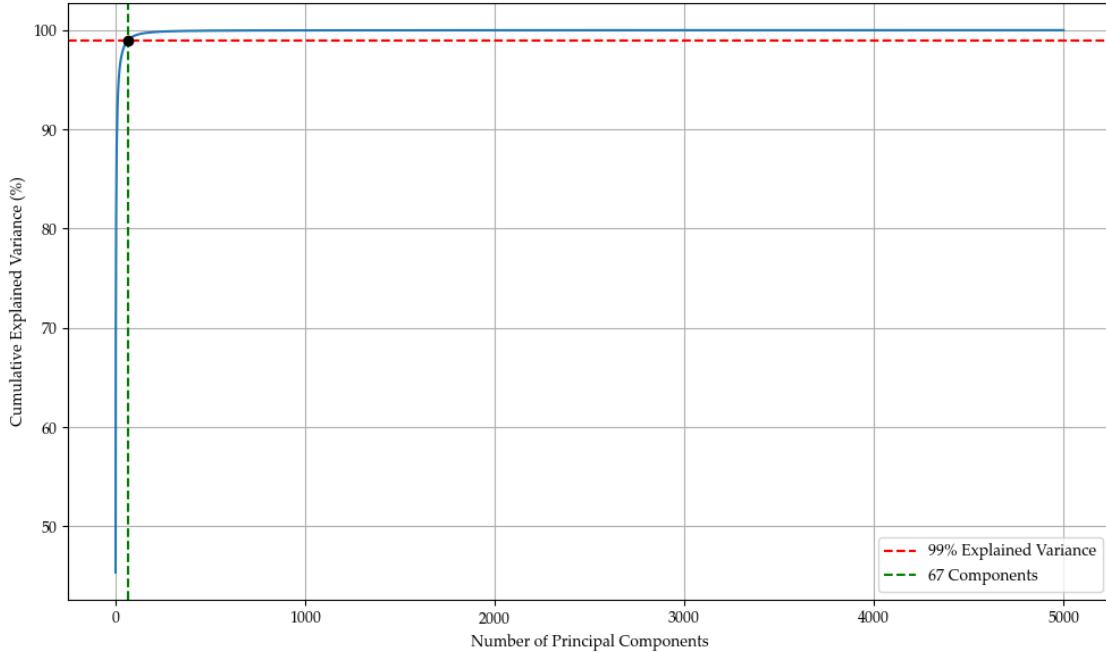


Figure 3.4: Cumulative explained variance of the principal components computed from the normalized composite distance matrix. The dashed line marks the 99 % threshold, reached at component 67.

3.2.2 Non-linear Dimensionality Reduction

While PCA efficiently reduces noise and dimensionality by retaining the most informative linear components, it is inherently limited to capturing linear relationships. In our dataset, however, morphological transitions are expected to involve complex, non-linear patterns that may not be adequately preserved through PCA alone. To uncover this latent structure and improve the separability of configurations, we complemented the linear reduction with three non-linear projection techniques.

We used the $k = 67$ PCA scores obtained in the previous subsection as input for:

- (i) Metric Multidimensional Scaling (MDS),
- (ii) t -Stochastic Neighbor Embedding (t -SNE), and
- (iii) Uniform Manifold Approximation and Projection (UMAP).

These three approaches were chosen deliberately rather than defaulting to a single “best” method. Although t-SNE and UMAP are often favored for non-linear embeddings, their performance is highly sensitive to parameter settings and dataset characteristics. Without direct comparison, it is difficult to determine *a priori* which method will best balance local neighborhood preservation, global geometry retention, and computational efficiency for a given problem. MDS, while older and less specialized for high-dimensional visualization, serves as a valuable baseline as it preserves global pairwise distances and offers a different bias–variance profile compared to manifold-learning methods.

Table 3.1: Explored parameter values for non-linear dimensionality reduction methods.

Method	Parameter	Values
MDS	n_components	2, 3
MDS	dissimilarity	precomputed
<i>t</i> -SNE	perplexity	30, 50, 100, 150, 200, 250, 300
<i>t</i> -SNE	learning rate	100, 200, 500, 1000, 2000
<i>t</i> -SNE	early exaggeration	12, 24, 36
<i>t</i> -SNE	n_components	2, 3
UMAP	n_neighbors	50, 100, 150, 250, 300, 500
UMAP	min_dist	0.0, 0.1, 0.2, 0.3, 0.5, 0.75
UMAP	n_components	2, 3

By including all three, we ensured a fair and systematic evaluation across complementary algorithmic principles: distance-preserving (MDS), stochastic neighbor-based (*t*-SNE), and manifold approximation (UMAP). This comparative approach reduces the risk of restricting our analysis to a single algorithm, increases the robustness of the final method choice, and provides a clearer understanding of how different projection biases affect clustering.

To evaluate their performance, we generated multiple embeddings in both two and three dimensions, systematically varying key parameters known to influence projection quality. Figures 3.5–3.7 show one representative 2D embedding for each method, and the full parameter grid is summarized in Table 3.1.

MDS effectively preserves global pairwise distances but is known to be sensitive to local minima and scales quadratically with sample size (Agarwal et al., 2007; Demaine et al., 2021). In our runs, the resulting projections often formed broad, curved manifolds with dense clusters interspersed with sparser regions (Figure 3.5). This structure suggests that while global relationships were largely maintained, local distortions in the input space may have introduced curvature and uneven density in the low-dimensional embedding, potentially affecting downstream density-based clustering.

t-SNE produced a scattered arrangement of small, elongated point clusters rather than well-separated compact “islands” (Figure 3.6). Many points formed short chain-like structures or arcs, with colors (normalized index) smoothly varying within clusters but transitioning sharply between others. While local neighborhoods were preserved, the irregular shapes and fragmented distribution highlight the method’s sensitivity to parameter tuning and its limited preservation of global geometry (van der Maaten & Hinton, 2008). Also, computational cost remained substantial at higher perplexities.

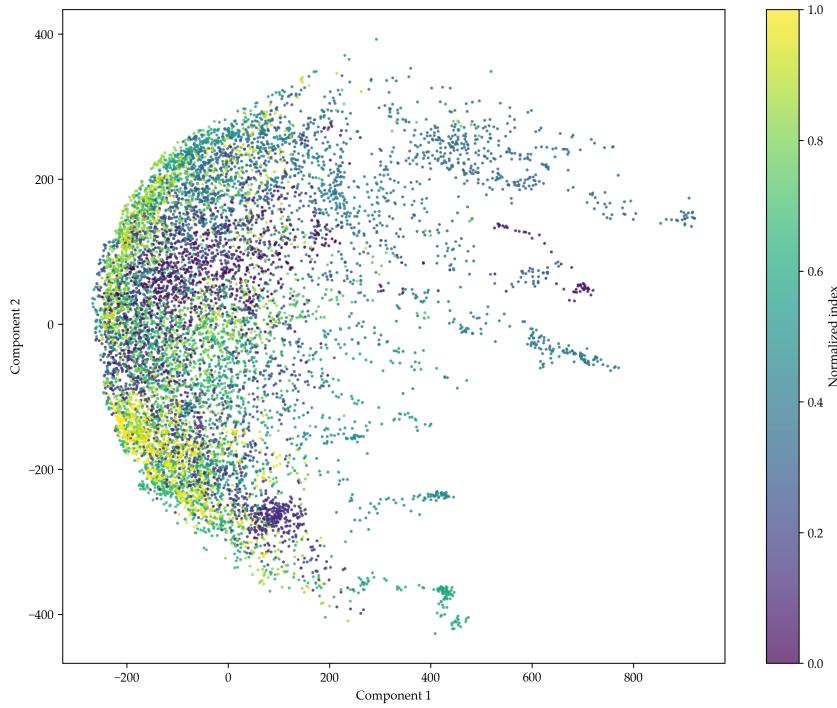


Figure 3.5: Example 2-D metric MDS embedding applied to our dataset. Each point represents one image; colour encodes the normalised time index (earliest = purple, latest = yellow). Note the elongated filaments and sparse central region, typical of the stress landscape for large data sets.

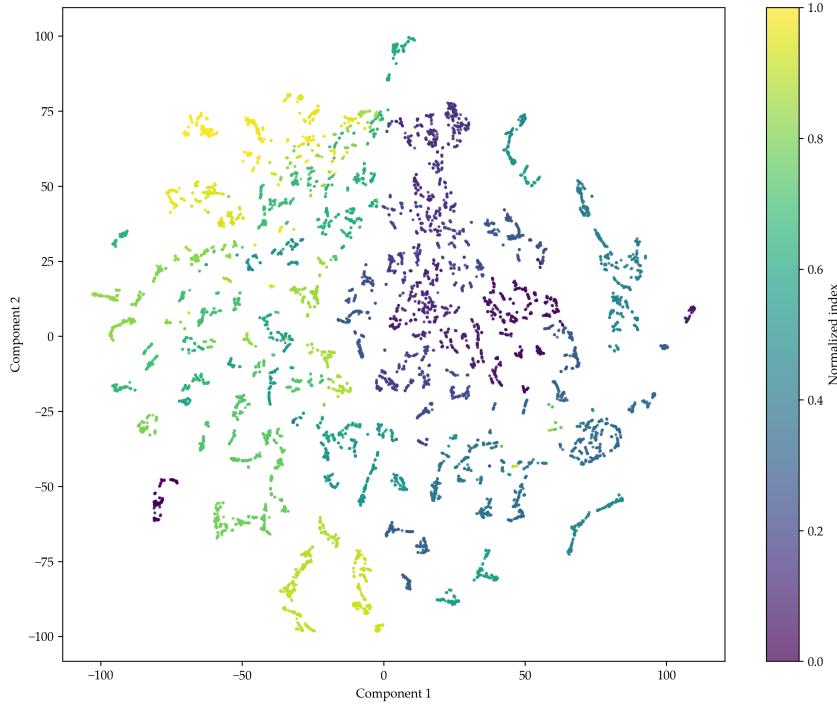


Figure 3.6: Representative 2-D t-SNE embedding applied to our dataset. Compact local islands are evident, but the global arrangement is sensitive to parameter changes, reducing geometric interpretability.

UMAP produced a sparse arrangement of small, often elongated clusters, with clear separation between many groups but also visible fragmentation into short chain-like structures (Figure 3.7). While points with similar normalized index values tended to cluster together, their shapes were irregular rather than compact. The `n_neighbors` parameter remained critical for balancing local continuity with global separation, and `min_dist` strongly influenced the observed elongation of clusters. Consistent with prior reports (McInnes et al., 2020), UMAP maintained high computational efficiency, scaling nearly linearly with dataset size and running substantially faster than *t*-SNE.

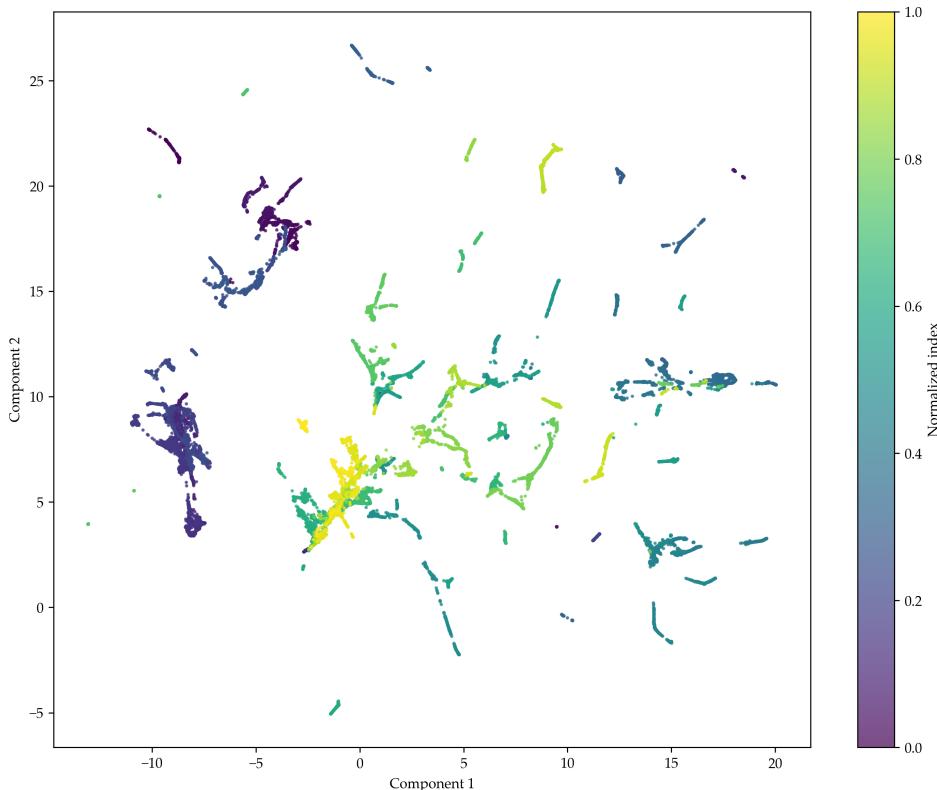


Figure 3.7: Representative 2-D UMAP embedding. Morphologically similar images cluster tightly, while dissimilar configurations are well separated—properties that favor density-based clustering and subsequent Markov-state definition.

Method selection

To identify the most suitable projection method, we generated a broad set of candidate embeddings by systematically varying the parameter combinations described above for MDS, *t*-SNE, and UMAP. Each embedding was assessed using two complementary approaches: (i) quantitative scoring with the Silhouette, Davies–Bouldin, and Calinski–Harabasz indices, and (ii) qualitative visual inspection. In addition, exploratory clusterings were performed with DBSCAN and HDBSCAN to evaluate how well each embedding supported meaningful and stable partitions. These same metrics will later be applied in the clustering analysis.

The evaluation yielded the following conclusions:

- **UMAP** consistently outperformed the alternatives, producing compact, well-separated clusters with structure that was coherent across both visual and quantitative assessments.
- **MDS** preserved global pairwise distances but frequently introduced local distortions, stretching clusters into elongated forms that hindered density-based clustering.
- **t-SNE** captured fine-grained local neighborhoods and yielded visually appealing groups, but the global arrangement was highly sensitive to parameter changes.

In addition, UMAP achieved these results at a fraction of t-SNE's computational cost, scaling efficiently with dataset size while maintaining a favorable balance between local detail and global structure.

This outcome underscores a well-known challenge in non-linear dimensionality reduction which is that optimal parameter settings are rarely known *a priori*, and performance can vary dramatically across configurations (Y. Wang et al., 2021). Reliable method selection therefore requires systematic exploration of the parameter space, which is computationally demanding but necessary to avoid bias toward a single algorithm. Moreover, clustering quality metrics, while useful, are insufficient on their own; they must be complemented with human inspection to detect distortions, fragmentation, or artifacts invisible to numerical scores (Machado et al., 2025). In our case, both quantitative metrics and visual evidence converged on the same conclusion: **UMAP offered the best compromise between interpretability, stability, and efficiency**, making it the best choice for subsequent clustering and state definition.

3.3 Clustering

Choosing between DBSCAN (Ester et al., 1996) and its hierarchical variant HDBSCAN (Campello et al., 2013) is not obvious *a priori* either. Both depend strongly on their hyperparameters and on the geometry of the input embedding. As in Section 3.2.2, we therefore adopted a comparative approach: for each UMAP embedding produced in the previous section we ran both algorithms over broad hyperparameter grids (Tables 3.1 and 3.2) and evaluated every run with the clustering metrics defined in Section 2.6.3.

Table 3.2: Explored hyperparameter values for DBSCAN and HDBSCAN.

Algorithm	Parameter	Values
DBSCAN	eps	0.10, 0.15, 0.20, 0.25, 0.30, 2.0
	min_samples	10, 25, 50, 75, 100, 150, 200
HDBSCAN	min_cluster_size	50, 100, 150, 250, 350, 400, 500, 800
	min_samples	15, 25, 50, 75, 100, 150, 250

Each clustering trial was logged with its parameter settings and scored with the quality indices listed in Table 3.3.

Table 3.3: Recorded parameters and quality metrics for every clustering run.

Column	Meaning / desirable trend
UMAP_n_neighbors	Local-global trade-off
UMAP_min_dist	Packing density; smaller \rightarrow tighter local groups
UMAP_n_components	2 or 3 dimensions in this study
DBSCAN_eps	Radius defining core points (ε)
DBSCAN_min_samples	Minimum points in an ε -ball
HDBSCAN_min_cluster_size	Minimum cluster size; larger \rightarrow broader clusters
HDBSCAN_min_samples	Density definition for core points
Num_Clusters	Clusters found (excluding noise)
Noise_Percentage	Points labelled -1; target < 20 %
Silhouette_Score	Higher \uparrow \Rightarrow tighter, well-separated clusters
Davies_Bouldin_Index	Lower \downarrow \Rightarrow better separation
Calinski_Harabasz_Index	Higher \uparrow \Rightarrow dense, isolated clusters

For clarity we illustrate one representative UMAP 2D embedding and vary clustering parameters across six panels in Figure 3.8. Noise points are hidden (assigned zero alpha) to emphasize how clusters change with the hyperparameters; shaded hulls outline each cluster.

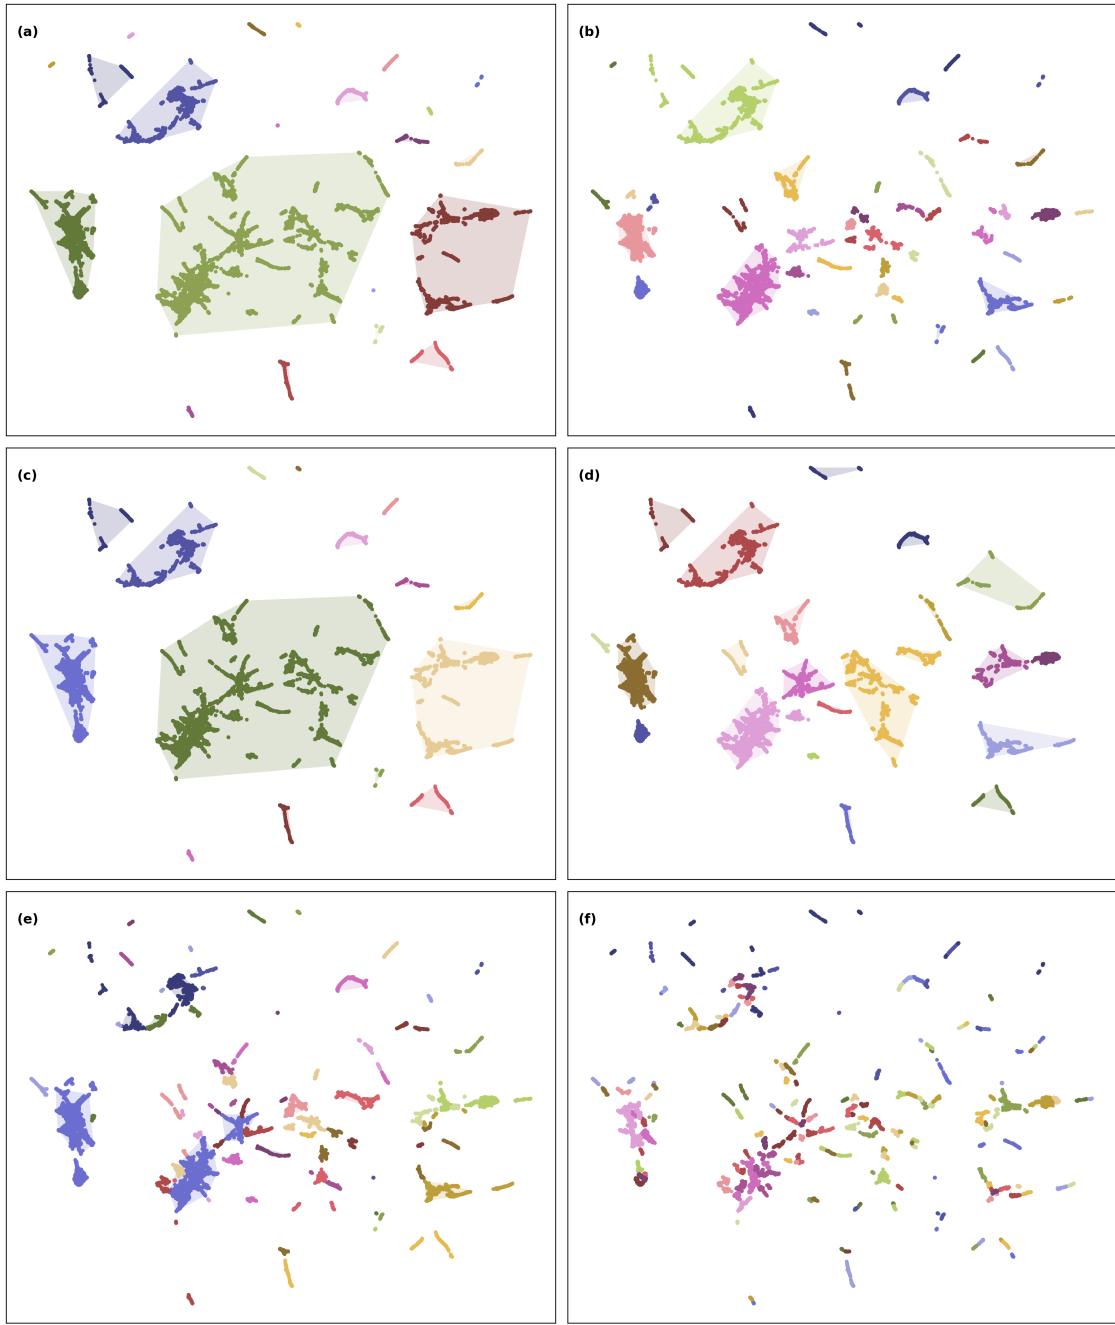


Figure 3.8: Effect of clustering hyperparameters on the same 2D UMAP embedding. Left column (a,c,e): DBSCAN with $(\epsilon, \text{min_samples}) = (2.0, 20)$, $(2.0, 200)$, and $(0.2, 20)$, respectively. Right column (b,d,f): HDBSCAN with $(\text{min_cluster_size}, \text{min_samples}) = (40, 150)$, $(400, 150)$, and $(40, 15)$, respectively. Points labelled as noise are not displayed for readability; colored convex hulls wrap the clusters. In DBSCAN, a large ϵ merges distant shapes (a); tightening the core definition with a very high min_samples removes sparse filaments as noise (c); a very small ϵ fragments structures into many small pieces (e). HDBSCAN adapts to local density: increasing min_cluster_size coarsens groups (d), while smaller min_cluster_size or min_samples accepts fine-grained but stable clusters without over-merging (b,f).

The comparison shows the core limitation of flat-density clustering for this dataset. With a single global ε , DBSCAN forces a trade-off in which generous radii over-merge elongated and compact shapes into heterogeneous aggregates (panel a). Strict core requirements remove most low-density branches as noise (panel c). Tight radii split coherent structures (panel e). In contrast, HDBSCAN builds a density hierarchy and keeps clusters that are stable across scales. Small structures are retained when they persist (panel b,f) and larger `min_cluster_size` values produce coarser but still morphologically coherent groups (panel d).

To move beyond visual assessment, we screened every configuration with Silhouette, Davies–Bouldin, Calinski–Harabasz, and the noise fraction. Good solutions must balance separation and coverage; excellent scores that arise from over-fragmentation (many tiny clusters) or from over-merging (few oversized clusters) were rejected. This joint visual–numerical procedure consistently favored HDBSCAN over DBSCAN for our UMAP embeddings, leading to the final configuration reported in Table 3.4.

Table 3.4: Selected parameters for UMAP embedding and HDBSCAN clustering.

Algorithm	Parameter	Value
UMAP	<code>n_neighbors</code>	300
UMAP	<code>min_dist</code>	0.0
UMAP	<code>n_components</code>	3
HDBSCAN	<code>min_cluster_size</code>	500
HDBSCAN	<code>min_samples</code>	250

Applied to the 3D UMAP projection, this setting yields 16 well-separated clusters while keeping the percentage of discarded points below 10% (Table 3.5); the arrangement is shown in Figure 3.9.

Table 3.5: Clustering performance metrics for the selected configuration.

Metric	Value
Number of clusters	16
Noise points	9.93%
Silhouette score	0.55
Davies–Bouldin index	0.52
Calinski–Harabasz score	20 479

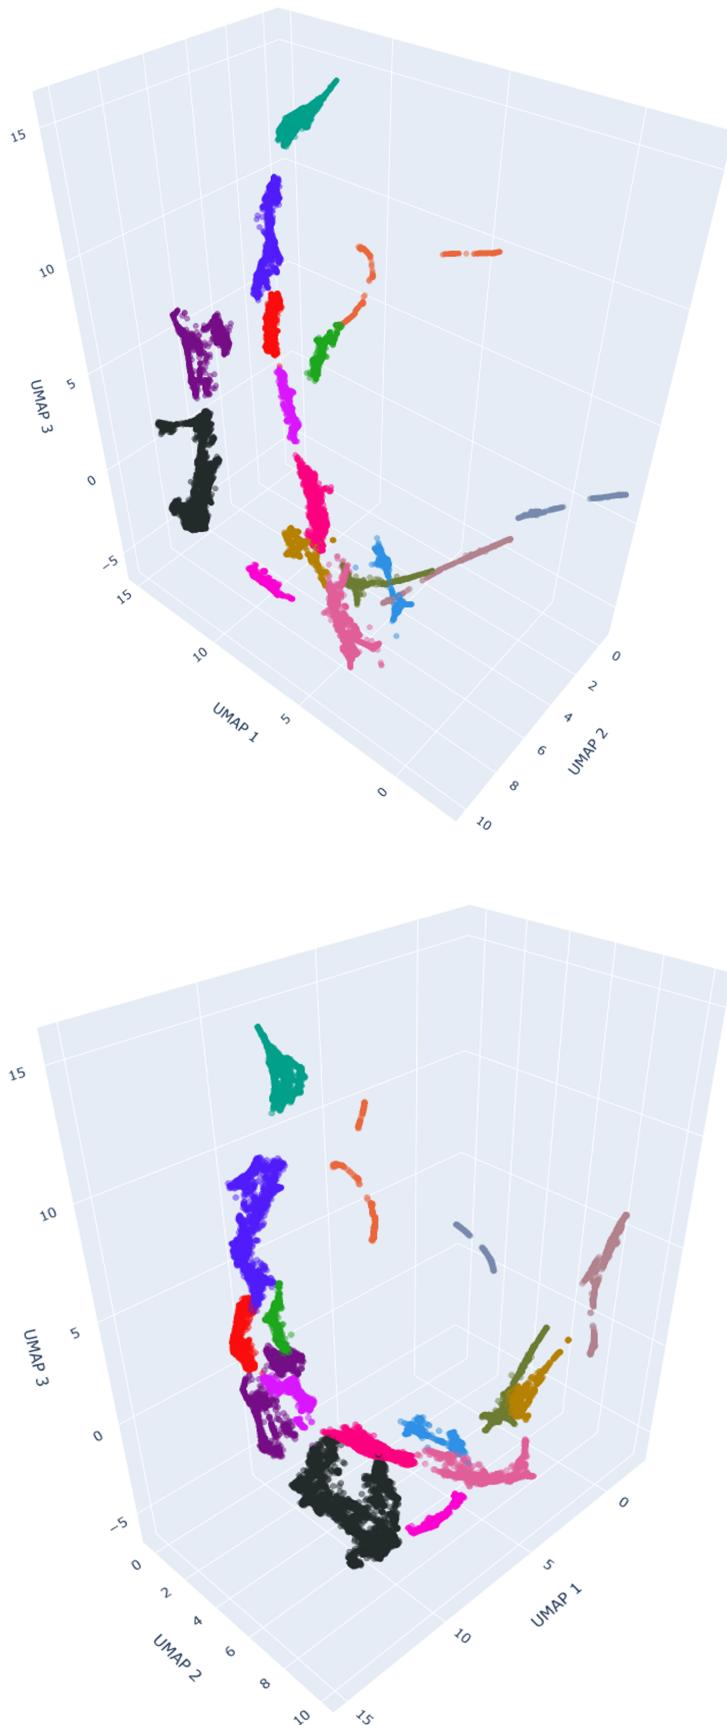


Figure 3.9: Three-dimensional UMAP embedding colored by the 16 clusters returned by the selected HDBSCAN configuration. Two viewpoints are shown to convey the spatial arrangement of clusters in 3-D space. An interactive version is available [online](#).

3.3.1 Characterization of Clusters

The goal of this section is to translate the purely data-driven clusters obtained with the UMAP + HDBSCAN workflow into physically interpretable river states. We proceed in two complementary steps. First, Section 3.3.1.1 documents the planform morphology captured by each cluster, combining a visual inspection of cluster medoids with five quantitative descriptors that are standard in braided rivers-geomorphological analyses. Second, Section 3.3.1.2 examines the sediment-transport signatures that were recorded synchronously with the imagery at one-minute resolution. Analyzing both morphodynamics and transport rates allows us not only to assess how distinct the clusters are in each domain but also to probe potential causal links between channel configuration and sediment flux.

3.3.1.1 Morphological Description

Figure 3.10 presents the binary medoid of every cluster, i.e., the image closest to the cluster centroid in the three-dimensional UMAP space of Figure 3.9. A rapid visual scan suggests a clear spectrum of planform configurations. Comparatively simple, single-thread states (Clusters 00 to 03, top row) transition into moderately sinuous and incipiently braided patterns (Clusters 04 to 08), and finally into fully developed multi-thread braids with large exposed bars and island complexes (Clusters 09 to 15, bottom rows). This qualitative progression is consistent with the classical channel-pattern spectrum, which ranges from single-thread to braided, as documented in fluvial geomorphology and associated with increasing lateral occupation and bar activity (Eaton et al., 2010; Leopold & Wolman, 1957; Surian, 2015).

Two additional, recurrent planform tendencies stand out. First, many clusters, especially at low to mid indices, show a persistent lateral bias in thread occupancy toward the left half of the flume. The thalweg hugs the left bank while emergent bars or partially wetted benches occupy the opposite margin. As braiding intensifies, the channel belt re-centers and discharge is redistributed more evenly across the corridor. This behavior is consistent with observations that thread motion can be coherent over long intervals and that channel belts widen and reorganize as braiding develops (Li & Limaye, 2025; Limaye, 2020).

Second, overall curvature increases alongside braiding intensity. Single-thread clusters display near-straight to gently curved thalwegs, whereas braided clusters exhibit pronounced alternating bends, frequent curvature reversals, and junction angles approaching orthogonality. These signatures are diagnostic of chute cut-offs and mid-channel bar driven avulsions, and are widely reported in experimental and field studies of multithread rivers (Micheli & Larsen, 2010; Nicholas et al., 2018).

Together with the widening of the active corridor and the richer texture of water–bar interfaces toward the higher-index clusters, these features depict a coherent increase in morphological diversity. The numeric labels (00 to 15) are therefore best read as identifiers of distinct states along a single- to multi-thread continuum rather than as an ordinal ranking of complexity.

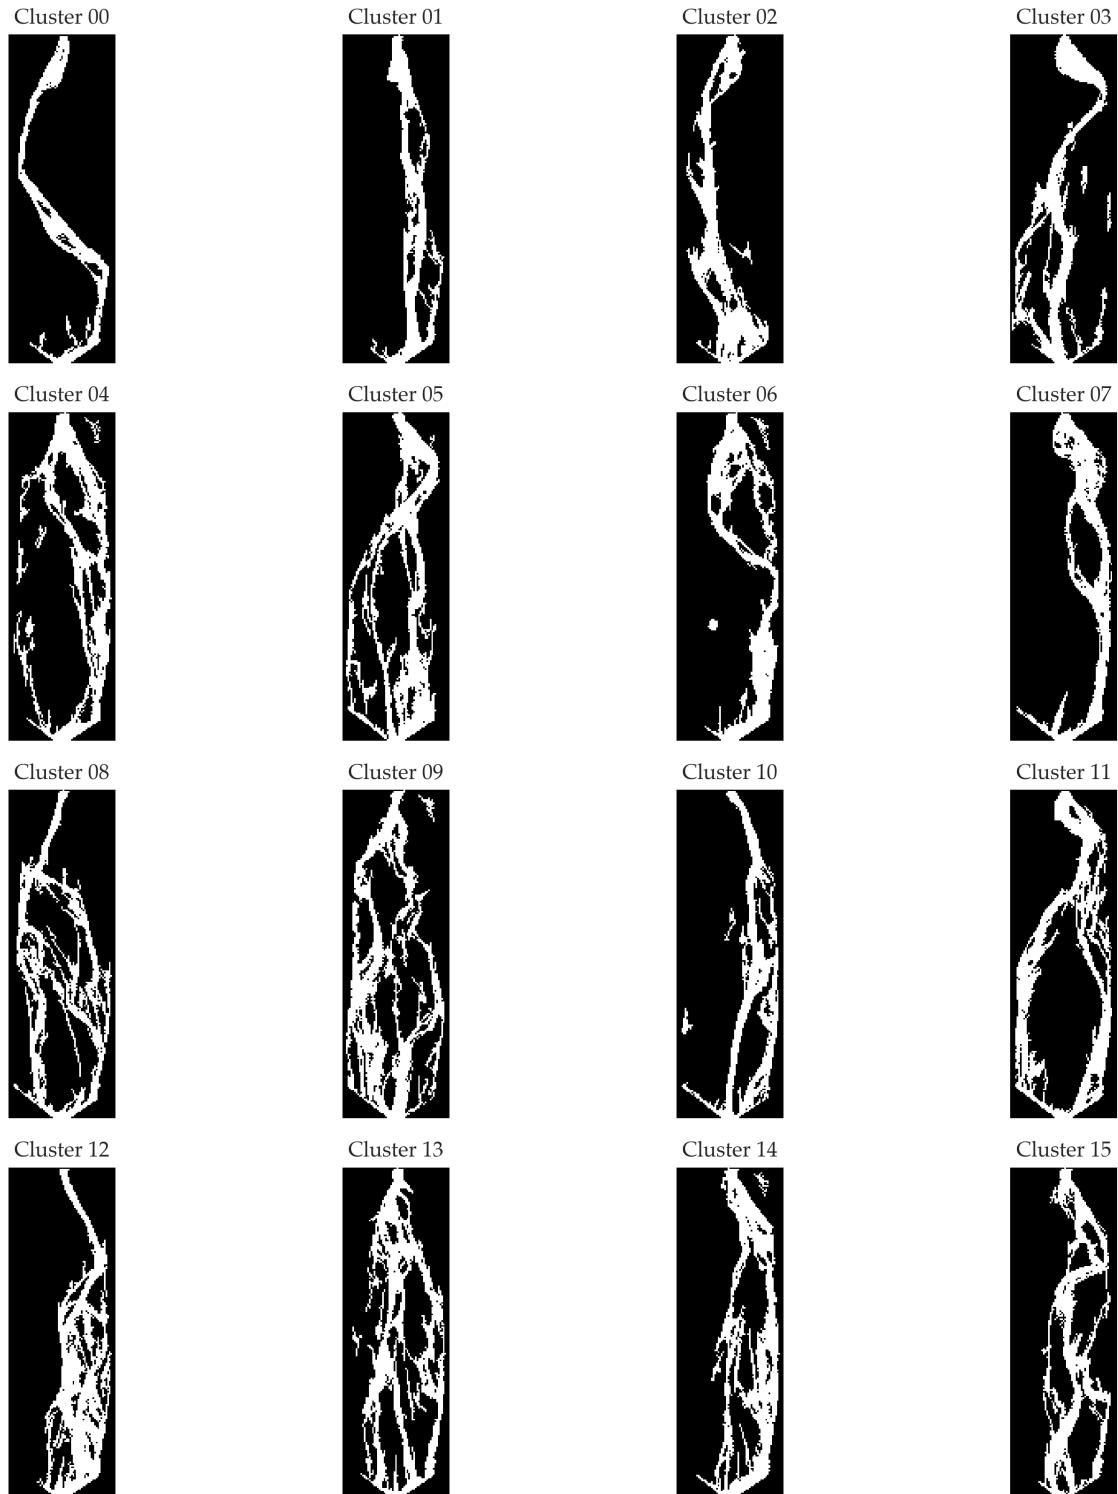


Figure 3.10: Binarized centroids of the 16 clusters (*Cluster 00–15*) obtained from the joint analysis of MHD and DICE distances. The 4×4 layout illustrates the morphological progression of channel states: the first clusters (00–03) exhibit relatively straight, single-threaded channels with limited lateral occupation; intermediate clusters (04–08) show increased sinuosity and the emergence of secondary threads; while the final clusters (09–15) display highly braided configurations, with multiple branches, exposed bars, and greater spatial complexity. These differences highlight the range of flow patterns and the gradual transition from simple to complex morphodynamics within the experimental river system.

Beyond this qualitative overview, each cluster was summarized by five planform metrics (see Section 2.6.4 for formal definitions) and by :

- (a) **Wetted width** — mean cross-stream distance between wetted banks;
- (b) **Wetted area** — total area of the water mask;
- (c) **Sinuosity** — centerline length divided by straight-line valley length;
- (d) **Braiding Index (BI)** — count of independent wetted threads;
- (e) **Entropy-based Braiding Index (eBI)** — Shannon entropy of thread occupancy, which down-weights ephemeral or minor channels.

The number of images per cluster is uneven (Table 3.6). Counts range from 1,954 (Cluster 00) to 13,729 (Cluster 11), with a median of $\approx 3,449$. Nearly half of all images ($\approx 47\%$) fall in Clusters 09–12, whereas six clusters contain fewer than 3,000 images each. This imbalance already hints at unequal state occupancies that we later quantify with the Markov model.

Table 3.6: Cluster-wise statistics: image count, wetted width (cm), wetted area (m^2), sinuosity, and braiding indices (scaled image counts, rescaled min=19.7 cm)

Cluster	Image Count	Wetted Width (cm)	Wetted Area (m^2)	Sinuosity (-)	BI (-)	eBI (-)	BI/eBI (-)
cluster 0	1 954	30.92	0.93	11.79	1.77	1.53	1.11
cluster 1	2 108	19.70	0.59	11.90	1.76	1.54	1.10
cluster 2	4 237	31.30	0.94	26.98	2.85	2.36	1.18
cluster 3	3 481	41.95	1.26	30.51	3.04	2.42	1.21
cluster 4	2 290	67.25	2.02	58.51	4.32	3.38	1.30
cluster 5	2 978	66.37	1.99	73.31	5.38	4.26	1.28
cluster 6	2 262	52.41	1.57	34.02	3.28	2.66	1.23
cluster 7	4 393	34.11	1.02	20.24	2.32	1.98	1.15
cluster 8	3 252	70.20	2.11	49.25	3.71	3.05	1.20
cluster 9	7 212	68.84	2.07	76.83	5.40	4.25	1.28
cluster 10	5 133	44.84	1.35	28.73	3.07	2.51	1.21
cluster 11	13 729	64.06	1.92	47.90	3.92	3.18	1.22
cluster 12	7 931	48.78	1.46	26.00	2.89	2.37	1.20
cluster 13	4 828	68.84	2.07	52.42	4.37	3.47	1.27
cluster 14	3 416	58.79	1.76	35.48	3.62	2.94	1.23
cluster 15	2 796	50.29	1.51	39.05	3.73	3.03	1.24

The statistics listed in Table 3.6 align with the medoids and quantify the gradient of planform complexity. Wetted width and area roughly double from low-index clusters (00–03) to strongly braided clusters (e.g., 09 and 13), indicating lateral expansion of active flow paths. Sinuosity follows the same trend, reaching its largest values in Clusters 09 and 05, while Clusters 00 and 01 remain close to straight.

Joint interpretation of BI and eBI. The entropic braiding index interprets the cross section as a probabilistic splitter. If channel i conveys a fraction p_i of the total discharge (we use width as proxy), the Shannon entropy $H = -\sum_i p_i \log_2 p_i$ measures pathway uncertainty and $eBI = 2^H$ is the effective number of channels. Hence eBI equals BI only when all threads have equal width, and $eBI \leq BI$ in every other case. Because very narrow threads contribute little to H , eBI is markedly more robust than BI to resolution and stage effects. The ratio BI/eBI quantifies channel-size heterogeneity. Values near one indicate uniform thread widths. Higher values indicate uneven partition with dominant and subordinate branches. These properties have been demonstrated on field transects and controlled numerical experiments, and the ratio helps distinguish braided from anastomosed patterns as well as interpret cross-section stability under discharge variability.

In our dataset the smallest values occur in Clusters 00–01 ($BI \approx 1.8$, $eBI \approx 1.5$), consistent with single-thread flow and low heterogeneity ($BI/eBI \approx 1.10$ – 1.11). The largest values occur in Clusters 09 and 05 ($BI = 5.40$ and 5.38 ; $eBI = 4.25$ and 4.26), followed by Clusters 13 and 04 ($BI = 4.37$ and 4.32 ; $eBI = 3.47$ and 3.38). These groups combine high channel count with elevated disparity ($BI/eBI \approx 1.27$ – 1.30). Most remaining clusters sit at intermediate complexity with moderate disparity ($BI/eBI \approx 1.20$ – 1.24). The range 1.10– 1.30 indicates that, even when multiple threads are active, discharge partition remains far from equal in the more braided states.

Two clusters may share similar BI yet differ in eBI and therefore in BI/eBI . This occurs when one state contains many minor chutes while another splits flow more evenly. The ratio is thus the concise summary of width diversity. Tejedor et al. (2022) showed that BI/eBI increases when small channels are activated at higher flows, that eBI is less sensitive than BI to added fine threads under higher image resolution, and that the ratio separates braided from anastomosed regimes. In our lab series this explains why strongly braided clusters (04, 05, 09, 13) pair high BI with the highest BI/eBI : many additional pathways exist, yet most carry a small share of the total wetted width.

Numerical experiments by Tejedor et al. show that cross-section stability depends on the ratio in a way that varies with the governing processes. Systems with vegetation or cohesive sediment tend to form stable sections with more even widths and higher eBI relative to BI . In sand-only braided runs, such as in our case, stability can occur with larger disparities and lower eBI relative to BI . Our facility uses non-cohesive sediment and no vegetation, so intervals dominated by high BI/eBI clusters likely represent configurations where a few dominant threads persist while minor paths appear and disappear between reconfigurations. This interpretation is consistent with the sustained high sinuosity and wetted width observed in the most braided groups.

In summary, our results of the triplet $\{BI, eBI, BI/eBI\}$ provide a coherent ranking of clusters from nearly single-thread to wide, bar-rich braids while remaining robust to mask resolution and short-term stage noise, in line with the properties documented by Tejedor et al. (2022)

3.3.1.2 Sediment Transport Description

We now examine how morphology-based states influence bedload transport. Table 3.7 shows the mean and standard deviation of the sediment-transport rate Q_s for each state, and Figure 3.11 presents the full distributions. Three main patterns appear in the violin plots. First, most states have a compact body below 0.2 g s^{-1} and a thin tail that reaches about $2\text{--}2.4 \text{ g s}^{-1}$. These long tails correspond to rare pulses that occur in many states and therefore do not distinguish one morphology from another. Second, the position of the median and the interquartile range (IQR) change systematically with planform: states with narrow, low-BI geometries have higher medians and wider IQRs, while strongly braided states concentrate near zero with narrow IQRs. Third, variability within a state is uneven. Some states show high central values with wide spread, while others remain consistently low and stable.

High-transport states. Clusters 00, 01, 07, and 02 exhibit the highest central tendency (means and medians), consistent with Table 3.7. In Figure 3.11 their violins are slender but centered at higher Q_s , with long upper needles. The thin bodies indicate that most observations cluster around a relatively stable transport level; the needles reflect occasional bursts to large values. These signatures match hydraulically efficient planforms with small wetted width and low BI/eBI (Table 3.6), typical of single-thread configurations that transmit momentum and sediment with minimal partition losses.

Low-transport states. Clusters 09, 05, 04, and 13 sit at the opposite end. Their violins are compressed near zero with tight IQRs, and the rare spikes appear only as tall, thin needles that do not shift the center of the distribution. These states are wide and highly braided (high BI and eBI), which favors storage and damped bedload export.

Between these extremes are intermediate states such as 03, 10, 11, 12, 14, and 15. They show middle values for both the median and the IQR, which matches their planforms of moderate width and braiding. Overall, Figure 3.11 and Table 3.7 reveal a clear trend. As wetted width, sinuosity, and BI/eBI increase, typical transport decreases and becomes more stable, while narrow single-thread states carry higher and more variable loads. This supports the use of morphological state as a condition for transport analysis in the Markov framework that follows.

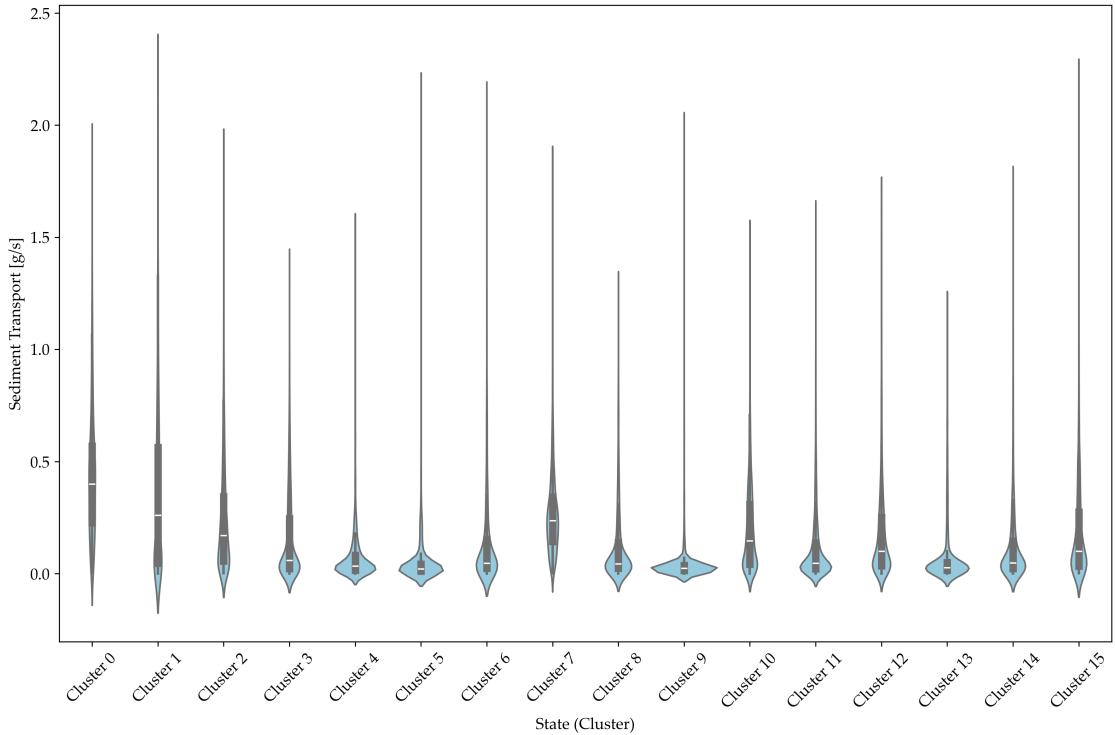


Figure 3.11: Violin plots of sediment transport rate Q_s across the 16 morphological states. The plot highlights inter-state contrasts in median transport and intra-state variability, providing a detailed view of how transport dynamics differ between states.

Table 3.7: Cluster-wise mean and standard deviation of the sediment transport rate Q_s (g s^{-1}), computed from the time series while each state was active.

Cluster	Mean $\langle Q_s \rangle$	Std. dev.
cluster 0	0.43	0.28
cluster 1	0.37	0.37
cluster 2	0.24	0.24
cluster 3	0.16	0.20
cluster 4	0.07	0.11
cluster 5	0.05	0.12
cluster 6	0.14	0.21
cluster 7	0.27	0.20
cluster 8	0.12	0.17
cluster 9	0.04	0.09
cluster 10	0.20	0.20
cluster 11	0.11	0.16
cluster 12	0.18	0.21
cluster 13	0.07	0.14
cluster 14	0.12	0.18
cluster 15	0.19	0.24
Global	0.155	0.208

Having established the state-level statistics, we now examine how morphology controls the mean bedload flux. Figures 3.12 and 3.13 plot the cluster average transport ($\langle Q_s \rangle$) against mean eBI and mean wetted width, with one point per state ($n = 16$). Both relations are strongly negative (Pearson $r = -0.920$, $R^2 = 0.847$ for eBI; $r = -0.959$, $R^2 = 0.919$ for wetted width) and the confidence bands remain narrow across the range. The trend is straightforward. States that are wider and more braided export less sediment on average. A unit increase in eBI corresponds to a drop in $\langle Q_s \rangle$ of about 10^{-1} g s^{-1} . An increase of roughly ten centimeters in mean wetted width corresponds to a decrease near 10^{-2} g s^{-1} . These effect sizes are approximate, but they convey the strength of the morphological control. In practical terms, higher eBI and larger width imply stronger partition of discharge, more wetted perimeter, and more bar-adjacent low-velocity zones, all of which promote temporary storage and reduce the outlet flux under fixed inflow.

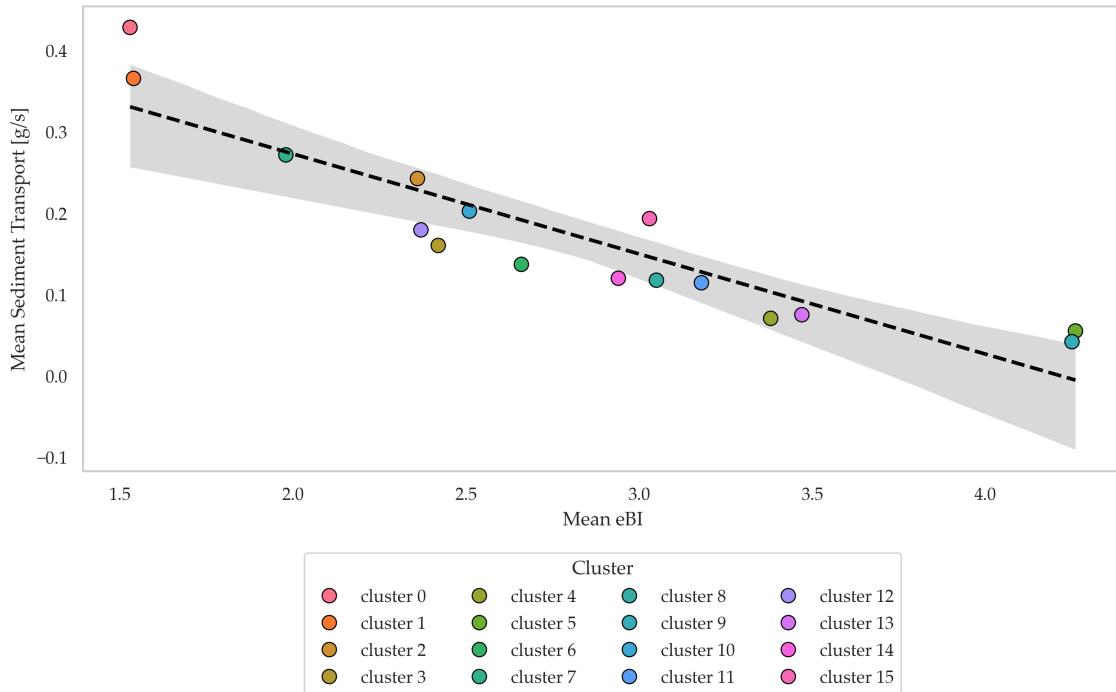


Figure 3.12: Cluster-average sediment transport ($\langle Q_s \rangle$) versus entropic braiding index (eBI). The dashed line is a linear fit; the shaded region is the 95% confidence interval. Higher eBI (more diverse, partitioned braiding) is associated with lower average transport.

Figure 3.14 summarizes co-variation among descriptors and transport statistics. Width, area, sinuosity, BI, eBI and the BI to eBI ratio are highly inter-correlated (coefficients near 0.9). Each of these geometry metrics is negatively correlated with both the mean and the standard deviation of Q_s with $|r| \gtrsim 0.8$. Hence, simple and narrow states tend to pass more sediment and do so more intermittently, while complex and wide states act as low-throughput and low-variability storage configurations. Mean and standard deviation of Q_s are themselves positively correlated ($r \approx 0.88$), so higher-flux states are also more volatile.

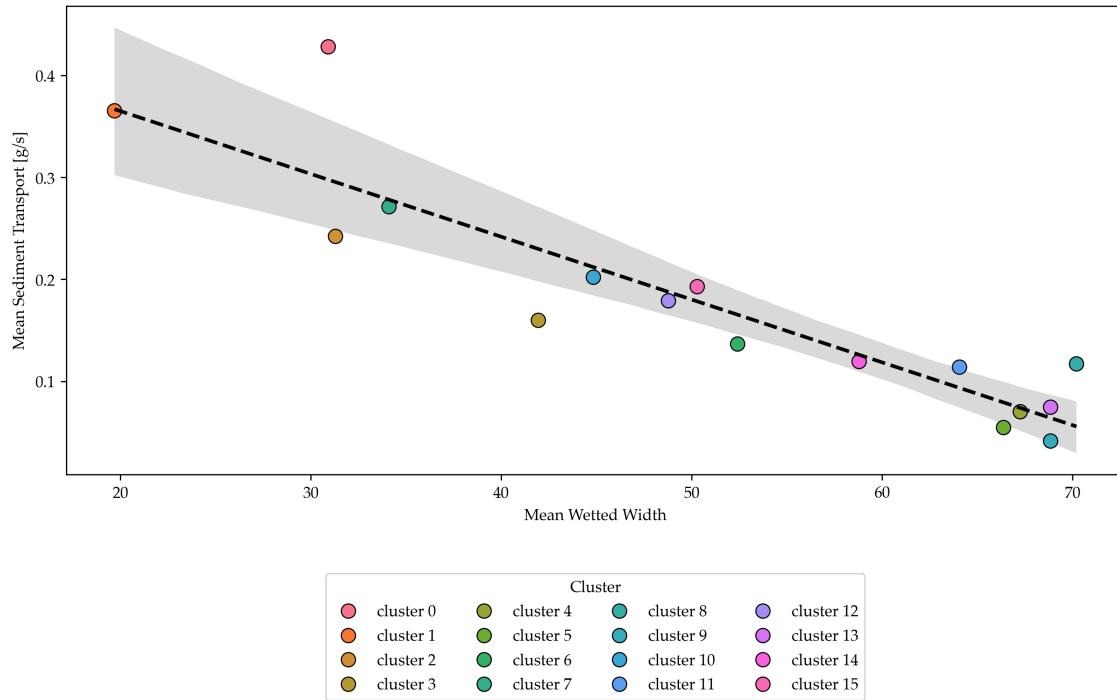


Figure 3.13: Cluster-average sediment transport $\langle Q_s \rangle$ versus mean wetted width. Colors denote states; the dashed line is a linear fit with 95% confidence band. Wider configurations tend to exhibit lower transport, suggesting that width acts as a proxy for partitioning and storage under the conditions of this experiment.

The colored time series in Figure 3.15 places these patterns in context. Large pulses line up with entries into narrow, low-eBI states such as 00, 01, 02 and later 10. Quiet periods occur during rapid shuttling among wide, high-eBI states such as 04, 05, 09 and 13. Peaks rarely occur as isolated spikes. They arrive in bursts that last tens to a few hundred minutes, during which the system dwells in a high-transport state and then cascades through one or two intermediates. State 07 often appears as a transient gateway in these sequences. This behavior is consistent with hysteresis reported in laboratory and field studies, where the geomorphic adjustments that sustain a pulse outlast the initial trigger (P. Ashmore, 1993; Benavides et al., 2022; Egozi & Ashmore, 2009).

Two practical inferences follow. First, either eBI or width can serve as a single proxy for partitioning and storage in this facility. Width is easy to measure and performs nearly as well as eBI. Second, because the geometry metrics are strongly collinear, a compact index or a first principal component would capture most of the predictive content with minimal redundancy, which is useful for following Markov state-based modeling.

Across hydrographs many studies report a positive association between active width and bedload rate. Those patterns reflect co-variation with discharge. Our experiment holds discharge fixed. Wider wetted width here occurs mainly in high-eBI states where discharge is split among several threads and bars store sediment. The negative relation between width and $\langle Q_s \rangle$ therefore reflects internal partitioning and storage under steady boundary conditions rather than contradicting discharge-controlled scaling (P. Ashmore, 1993; Benavides et al., 2022; Egozi & Ashmore, 2009).

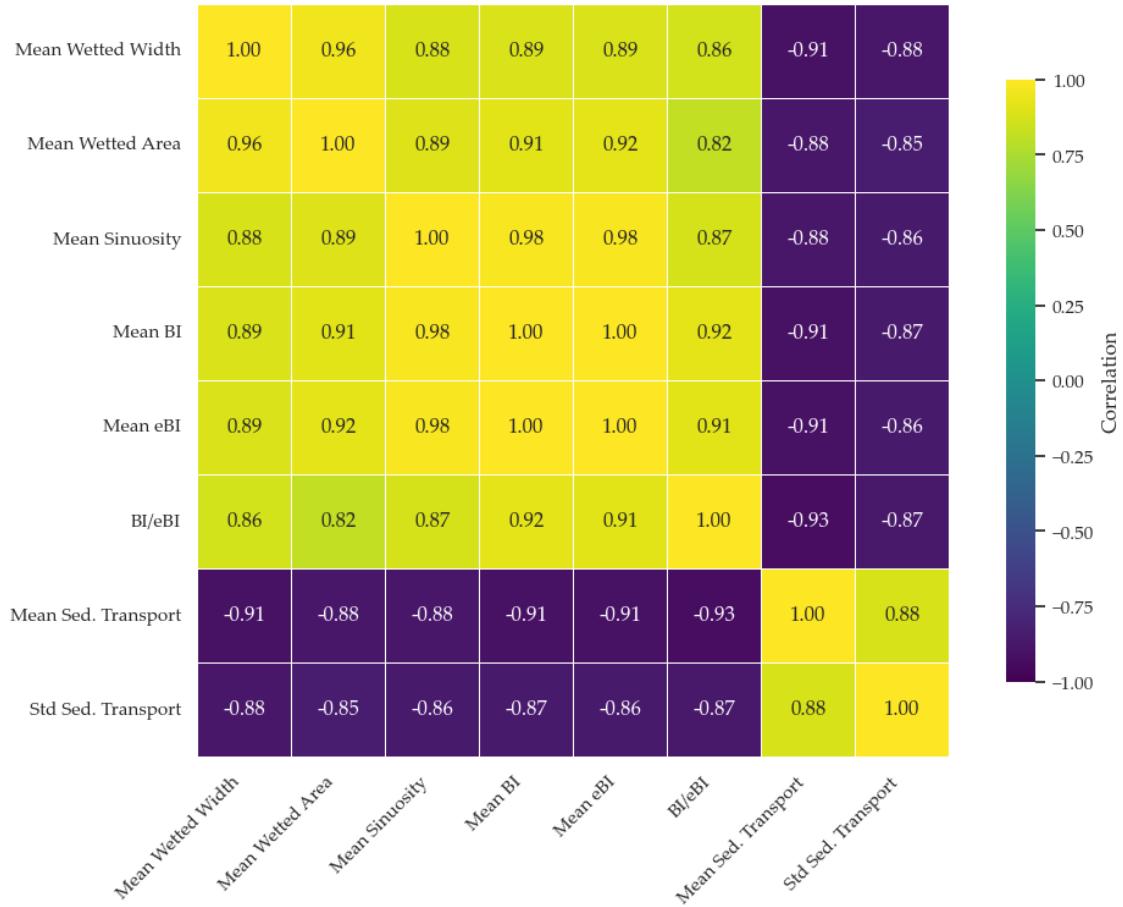


Figure 3.14: Correlation matrix between cluster-average sediment transport statistics ($\langle Q_s \rangle$, $\text{Std}[Q_s]$) and planform descriptors (wetted area, wetted width, sinuosity, BI, eBI, BI/eBI). Cool colors denote negative correlations; warm colors denote positive correlations.

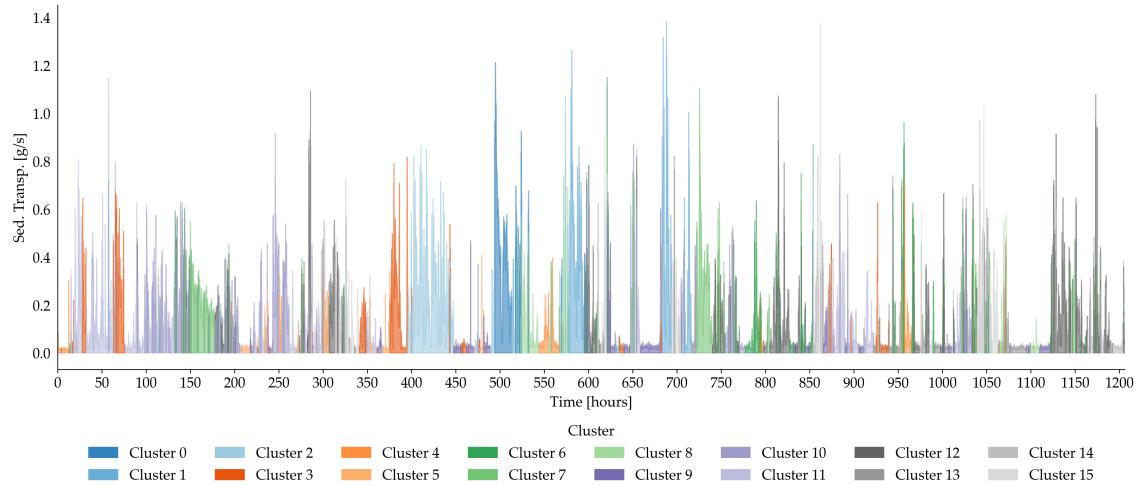


Figure 3.15: Time series of sediment transport rate Q_s during the flume experiment, colored by the active morphological state at each time step. Transitions from low-transport states (e.g., 04, 05, 09, 13) to high-transport states (e.g., 00, 01, 02, 10) align with major sediment pulses. Low-transport periods involve rapid switching among states, while high-transport phases often persist for 50–200 minutes and exhibit hysteretic sequences.

The evidence from bivariate fits, the correlation matrix and the colored time series all point in the same direction. State membership is a strong determinant of transport behavior. Accurately representing the timing and probability of transitions among states is thus essential for reproducing bursts and lulls in the outlet record, which motivates the Markov-state model introduced in Section 3.4.

3.4 Markov Model Construction

As stated at the beginning of this chapter, our aim is to capture the evolution of the braided-river experiment with a probabilistic description based on a Markov chain. In what follows we construct a continuous-time Markov chain (CTMC) from the image-derived state sequence. The model requires estimating two empirical components from data and relating them to the standard CTMC machinery. First, we estimate the state-specific residence (departure) rates $\lambda = \{\lambda_i\}$ from the observed durations of uninterrupted sojourns in each state i . Second, we estimate the embedded jump probabilities $P = \{p_{ij}\}_{i \neq j}$ by counting all observed transitions $i \rightarrow j$ and normalizing by the number of departures from i . These two ingredients jointly determine the infinitesimal generator Q of the CTMC: for $i \neq j$, the off-diagonal entries satisfy $q_{ij} = \lambda_i p_{ij}$, while the diagonal entries are $q_{ii} = -\lambda_i$. In this embedded-chain view, self-transitions are not represented in P ; state persistence is encoded entirely by the residence rate λ_i . Together, the pair (λ, P) fully specifies the model used throughout the remainder of the chapter.

The classification of each image into one of the sixteen clusters described in Section 3.3 allows us to construct a time series of the active morphological state at each sampling step. This labeled sequence is shown in Figure 3.16, where each minute of the experiment is assigned a corresponding state. Visual inspection of this timeline already suggests a strong heterogeneity in how long the system remains in each configuration as individual dwell times range from a few minutes to several hours. Representing such contrasts with a discrete-time chain defined at the 1-min sampling interval would force long segments to appear as hundreds of successive self-loops and short segments to rely on transition probabilities that are nearly zero, blurring the physical meaning of a model step and complicating comparison with the measured sediment-transport signal.

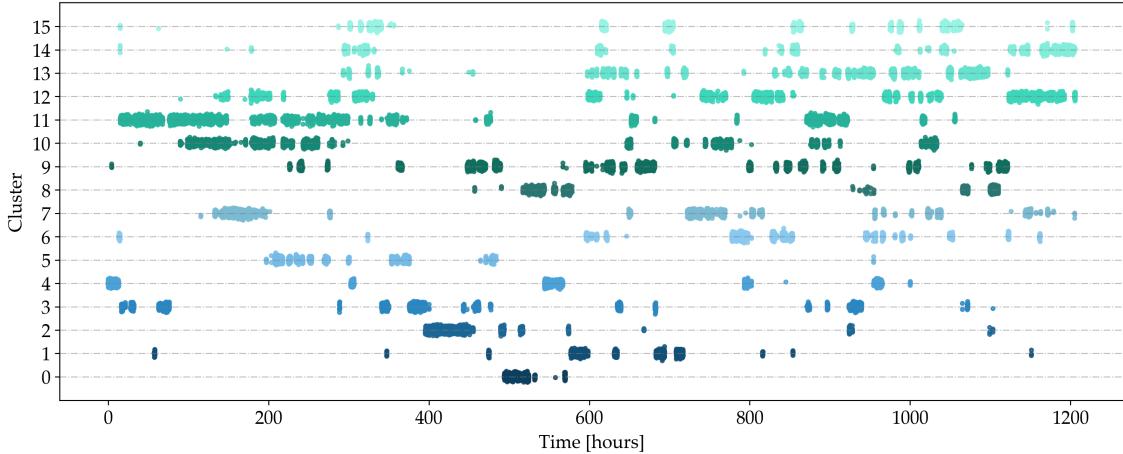


Figure 3.16: Temporal distribution of morphological cluster states throughout the experiment. Each dot marks the active state at a given minute. A small vertical jitter improves the visual separation of overlapping points and conveys the density of state occurrences over time.

A continuous-time formulation avoids this ambiguity. In a CTMC, the system remains in a given state for a random time that follows an exponential distribution (the only continuous distribution compatible with the Markov memoryless property), and then jumps to a new state according to the fixed probabilities P . This captures the wide variability in dwell durations without increasing the number of parameters and yields a clean separation between persistence (governed by λ) and reconfiguration (governed by P).

To construct the model from the observed state sequence, we proceed in two steps. For each state i , we collect all uninterrupted episodes and compute the sample mean $\langle \tau_i \rangle$; the maximum-likelihood estimate of the residence rate is then $\hat{\lambda}_i = \langle \tau_i \rangle^{-1}$. We next count all transitions $i \rightarrow j$ with $j \neq i$ and normalize by the number of departures from i to compute the empirical jump probabilities \hat{p}_{ij} . The resulting estimates are summarized as a set of representative dwell-time histograms with exponential overlays (Figure 3.17) and as the full transition-probability matrix (Figure 3.18); the numerical values of $\hat{\lambda}$ are reported in Table 3.8. In Figure 3.18 we omit the main diagonal by design, since self-transitions are not modeled in P ; state persistence is encoded by $\hat{\lambda}_i$.

Because the images are sampled every one minute, sojourn times are generated at that resolution. In practice, the vast majority of state episodes extend well beyond a single minute (often tens of minutes to hours), so discretization effects are minor at the time scales of interest. This is consistent with independent flume studies that resolve morphodynamic changes at short cadence (about 2 minutes) and document rapid adjustments superimposed on longer reorganizations in braided channels (Vesipa et al., 2018). High-frequency sampling is therefore essential to capture both short-lived excursions and sustained reorganizations, and our 1-minute frequency lies well within the window needed to observe such changes.

Table 3.8: Estimated rate parameter λ for the exponential distribution fitted to the durations of continuous presence in each cluster. Values are expressed in inverse minutes.

Cluster	λ	Cluster	λ	Cluster	λ
Cluster 00	0.004	Cluster 01	0.006	Cluster 02	0.006
Cluster 03	0.014	Cluster 04	0.005	Cluster 05	0.025
Cluster 06	0.025	Cluster 07	0.037	Cluster 08	0.013
Cluster 09	0.009	Cluster 10	0.049	Cluster 11	0.019
Cluster 12	0.030	Cluster 13	0.016	Cluster 14	0.027
Cluster 15	0.015				

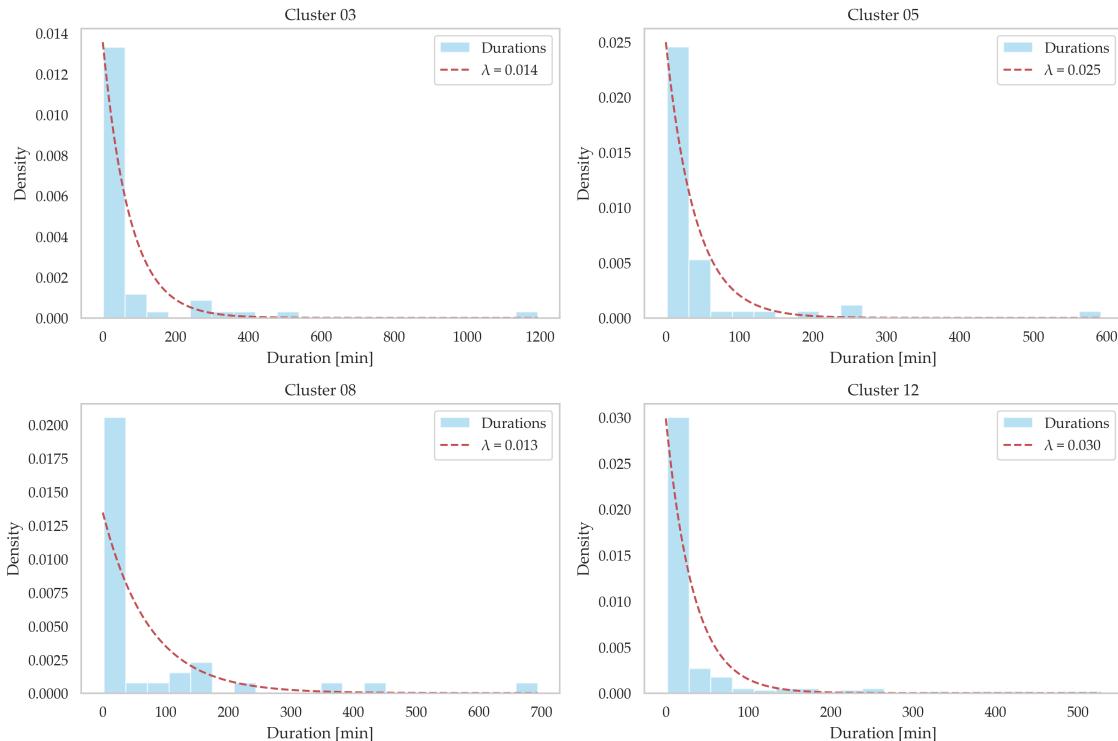


Figure 3.17: Empirical dwell–time histograms (blue bars) for four representative clusters, overlaid with the fitted exponential probability–density functions (red dashed lines) using the maximum–likelihood rates reported in Table 3.8. The panels illustrate the spread in residence behavior across the state space: Cluster 03 and Cluster 08 occupy an intermediate regime, whereas Cluster 05 lingers longer and Cluster 12 turns over more rapidly.

At this point a couple of implementation details need to be clarified.

- (i) Minutes labeled as “noise” by HDBSCAN in Section 3.3 were mapped to the nearest cluster in the UMAP space (nearest neighbor assignment), ensuring a continuous state series for inference.
- (ii) The experiment was designed and run long enough to reduce statistical uncertainty in $\hat{\lambda}$ and \hat{P} . Nevertheless, from a mathematical standpoint absolute proof of model adequacy is elusive, and finite samples, particularly for rare states, inevitably leave residual uncertainty in tail behavior and low probability transitions. When simulating from the fitted CTMC in Chapter 4 we initialize from a random state (a uniform draw over the sixteen states unless

otherwise specified), which keeps initial condition assumptions neutral.

The estimated rates are shown in Table 3.8, the full transition probability matrix is presented in Figure 3.18, and a set of representative dwell time histograms with their fitted exponential curves is provided in Figure 3.17. The complete set of dwell-time histograms for all clusters is provided in Appendix C.

Taken together, the residence time distribution visualized in Figure 3.17 and tabulated in Table 3.8, along with the jump probabilities in Figure 3.18, fully define the model that underpins the remainder of this chapter.

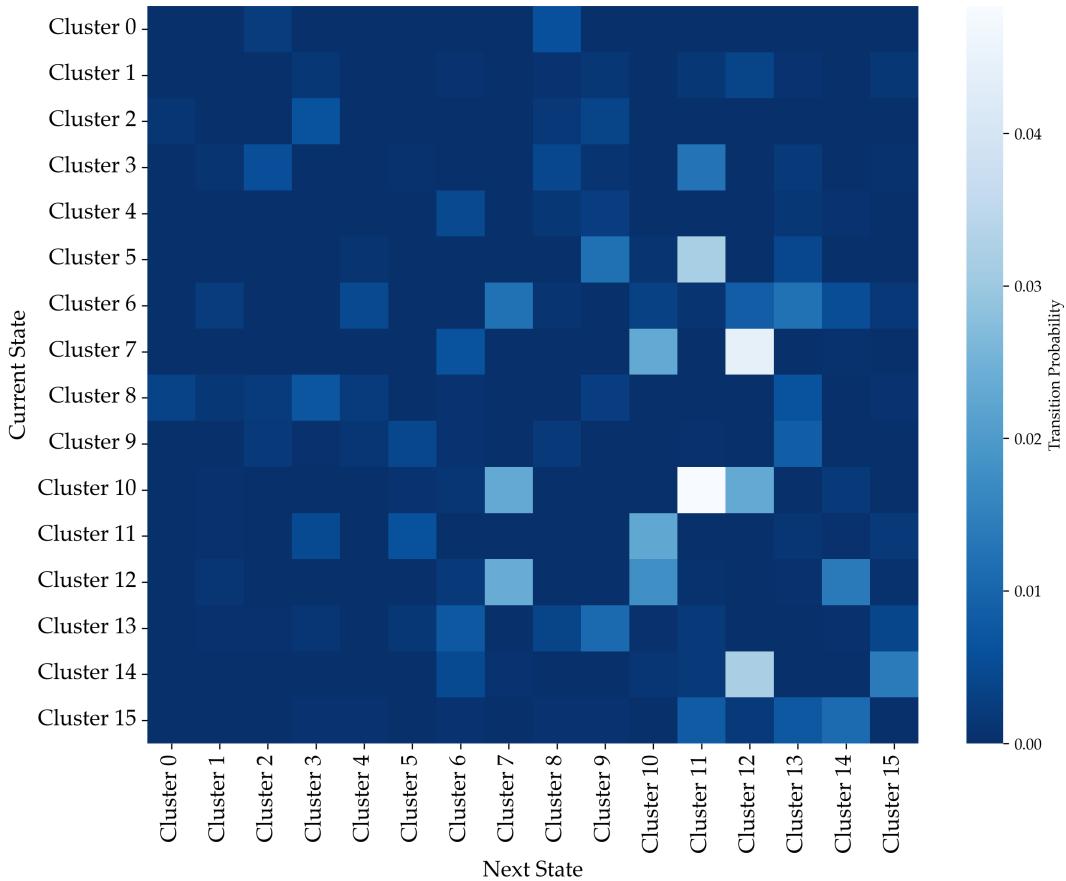


Figure 3.18: Transition-probability matrix P for the embedded jump chain (rows: current state; columns: next state). Entries on the main diagonal are omitted because state persistence is encoded by the residence rates $\hat{\lambda}_i$; off-diagonal intensities reflect the empirical probability of jumping from one morphological state to another upon departure.

The spectrum of $\hat{\lambda}_i$ shows a clear ordering that reflects both planform complexity and sediment transport. The smallest rates occur in wide, bar-dominated braids (Clusters 09, 13, and 05), with median residence times of two to three hours. These states have the widest channels, the highest braiding indices (BI and eBI), and the lowest mean sediment loads, indicating that once a wide multi-thread braid is established the system tends to store rather than export sediment. At the other extreme, Clusters 00, 01, and 02 have the highest rates and residence times shorter than ten minutes. They are the narrowest and least braided yet carry the largest mean sediment fluxes, consistent with frequent adjustments

under energetic, sediment-rich conditions such as short scour–fill cycles or rapid chute activation. Intermediate states (e.g., Clusters 03, 08, and 12 in Figure 3.17) fall between these extremes in both geometry and transport. Their fitted exponential curves align with the empirical histograms, supporting the exponential residence time assumption and showing that rates scale with morphological complexity and transport regime. Motivated by this pattern, we attempted to link $\hat{\lambda}_i$ to simple morphometric descriptors (i.e., wetted width, BI/eBI, junction density, active width), but with the parameter set used in this thesis the associations were weak and inconsistent, and did not yield a reliable predictive mapping of $\hat{\lambda}_i$ across states. Future research on this topic would be a good improvement for the model.

Overall, the residence rate ordering matches the geomorphic continuum in Table 3.6 and the two transport regimes in Table 3.7. Simple high load channels are short lived, while complex low load braids persist. This supports modeling state changes as a CTMC, where the estimated rates capture the asymmetric time scales of river morphodynamics. The CTMC is fully defined by the residence rates $\hat{\lambda}$ and jump probabilities \hat{P} , which specify the generator Q and allow both analytical summaries (such as state occupancies and mean first passage times) and forward simulation in Chapter 4. The dwell time histograms (Figure 3.17) and the transition matrix (Figure 3.18) provide evidence for this representation and link back to the morphology–transport contrasts described in Section 3.3.

3.4.1 Insights from the Markov Representation of the System

The Markov chain model provides a probabilistic view of how morphology, transport, and state changes interact in the braided system. We use it here to test how well the Markov assumption reproduces the observed dynamics. Cluster average sediment transport spans more than an order of magnitude, from $\langle Q_s \rangle \approx 0.04 \text{ g s}^{-1}$ in cluster 9 to $\approx 0.43 \text{ g s}^{-1}$ in cluster 0. Wetted width ranges from $\approx 19.7 \text{ cm}$ (cluster 1) to $\approx 70.2 \text{ cm}$ (cluster 8), while braiding indices increase from $\text{BI/eBI} \approx 1.76/1.54$ (cluster 1) to $\approx 5.40/4.25$ (cluster 9). These trends trace a continuum from narrow single thread channels to wide bar rich braids. Image counts are uneven, from 1954 in cluster 0 to 13729 in cluster 11, anticipating unequal occupancies that the Markov model quantifies.

Looking again at the sediment transport series colored by state (Figure 3.15), we see clear regimes linked to morphology. Low transport is often associated with rapid alternation among wide and highly braided clusters that favor storage, especially 5, 9, and 13 (see Table 3.6). In contrast, the largest transport pulses occur during excursions into simpler and straighter states such as 0, 1, 2, and 10. Intermediate clusters, including 7, 11, and 12, often appear along the hysteretic paths into and out of high transport episodes.

Figure 3.19 summarizes the trade off between persistence and activity. States with many transitions have short mean stays, while states with few transitions persist for much longer. This inverse envelope is consistent with a Poisson escape process where the departure rate λ_i controls both transition frequency and average dwell time ($\mathbb{E}[\tau_i] = 1/\lambda_i$). Bubble areas add the transport dimension. Within the green hull, larger markers indicate that highly active states tend to convey higher mean fluxes. Within the red hull, bubbles are smaller

except for cluster 0, which combines long residence with high transport.

Three regimes emerge:

- (i) **Red – anchor states (0, 1, 2, 4):** very few transitions (≤ 50) and mean dwell times well above 150 min, i.e., small escape rates. Only cluster 0 couples this persistence with very high mean transport (see Table 3.7); clusters 1 and 2 show moderate transport, while cluster 4 is long-lived but low-transport.
- (ii) **Blue – intermediate states (3, 5, 6, 8, 9, 13, 14, 15):** moderate activity (~ 80 –200 transitions) and intermediate persistence (~ 40 –110 min). These include the widest, most braided, storage-prone morphologies; bubble sizes are generally small to moderate.
- (iii) **Green – transit states (7, 10, 11, 12):** the shortest mean stays (< 60 min) and the highest activity (> 350 transitions). Bubbles are typically larger, highlighting the tendency of fast-reorganizing configurations to carry higher mean transport.

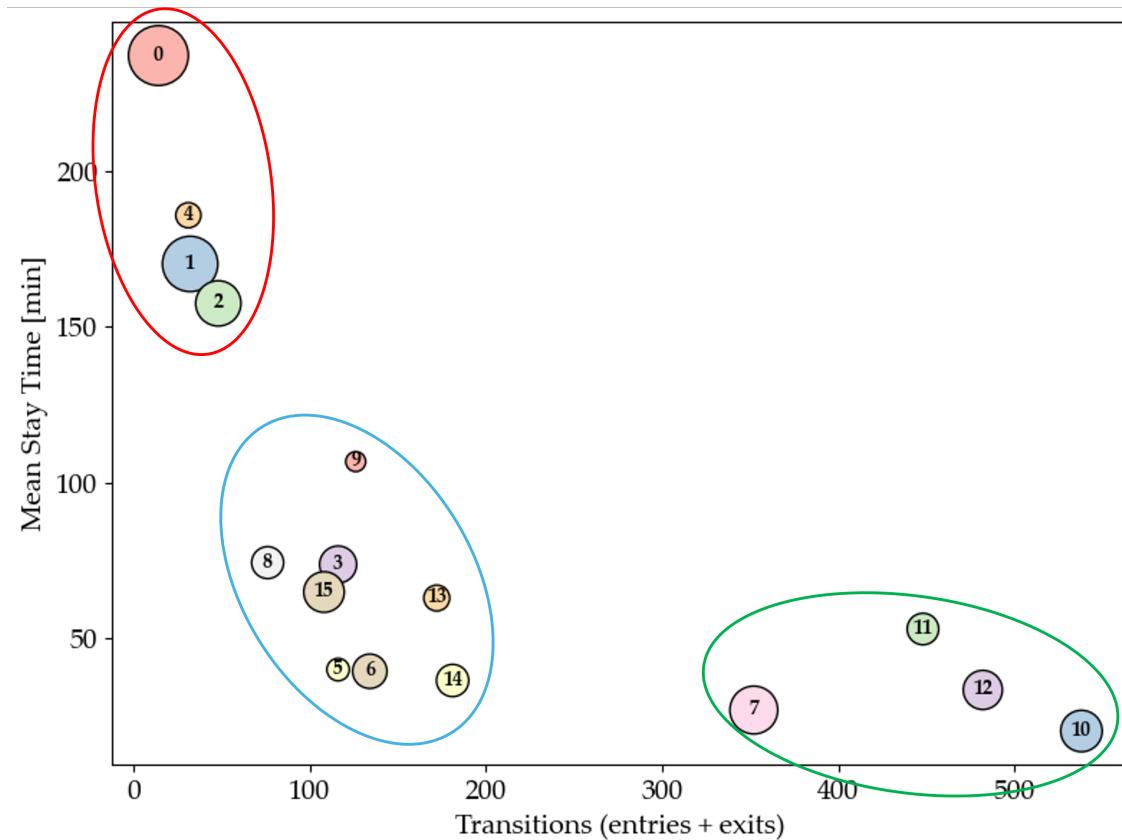


Figure 3.19: State persistence versus activity. Each point is a cluster (labelled by ID). The x-axis shows the number of transitions (entries+exits) and the y-axis the mean stay time in minutes ($\approx 1/\hat{\lambda}_i$). Marker area is proportional to the cluster's mean sediment transport (larger bubbles indicate higher flux). Colored ellipses group states with similar behavior: long residence and few transitions (red), intermediate regime (blue), and short residence with many transitions (green). The larger bubbles concentrated in the high-activity group reflect that states with frequent adjustments tend to convey higher mean transport.

The red group is interesting. Clusters 0, 1 and 2 combine very long dwell times with the high average transport (Table 3.7). Their medoids (Figure 3.10) are a narrow, weakly braided single thread with low BI/eBI, i.e., a hydraulically efficient pathway. Once such a confined thread establishes, scour–conveyance feedbacks deepen and maintain the thalweg, suppressing bar splitting and sustaining high transport. By contrast, cluster 4 is also persistent but low–transport, reflecting a stable yet partitioned, storage–prone morphology. Hence, longevity alone does not determine sediment yield; the type of persistent morphology (conveyance–efficient versus storage–dominated) is decisive.

3.4.2 Transport changes associated with state transitions

Looking back to the transition-probability matrix (Figure 3.18), we can identify dominant morphological pathways as certain pairs of states exchange more frequently than others, revealing recurrent morphodynamic cycles. This preferential behavior underscores the structured evolution of the system through a subset of favored transitions.

An alternative, more explicit view of jump frequency and directionality is given by the most frequent bidirectional pairs (Figure 3.20). Nearly symmetric exchanges emerge for some pairs (e.g., $10 \leftrightarrow 11$ and $7 \leftrightarrow 12$), which appear often despite short dwell times (green hull in Figure 3.19). Their prominence reflects high activity rather than persistence and helps explain how short-lived transit states can occupy a large fraction of the transition record. The full list of transitions is available in Appendix D.

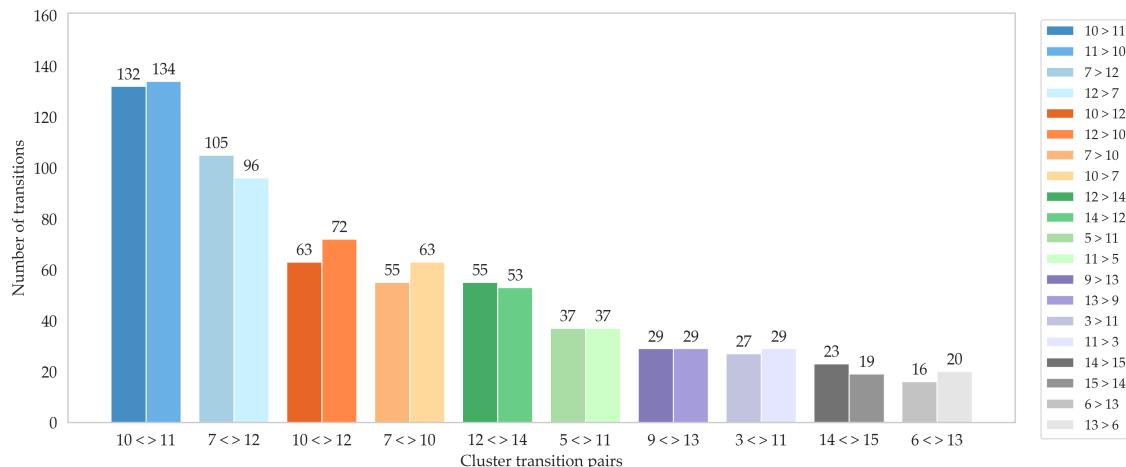


Figure 3.20: Most frequent bidirectional transitions between distinct morphodynamic clusters. Each cluster pair (e.g., $10 \leftrightarrow 11$) is plotted as two adjacent bars, one per direction ($10 \rightarrow 11$ and $11 \rightarrow 10$). Bar height indicates the number of observed transitions, and numbers above each bar give exact counts.

We now place these frequency patterns next to persistence and activity, and we evaluate their transport consequences with the ordered heatmap. The persistence-activity diagram in Figure 3.19 shows that green states are short lived and highly active, blue states are intermediate, and red states are long lived with few transitions. The heatmap in Figure 3.21 maps the change in mean sediment transport $\Delta Q_s = \langle Q_s \rangle_{\text{next}} - \langle Q_s \rangle_{\text{current}}$ for each ordered transition. Blue cells indicate that transport increases after the jump, and red cells indicate a decrease. Read together, Figures 3.19, 3.20, and 3.21 show that frequent exchanges

within the green set dominate counts while the largest transport swings are tied to regime crossings between red, blue, and green.

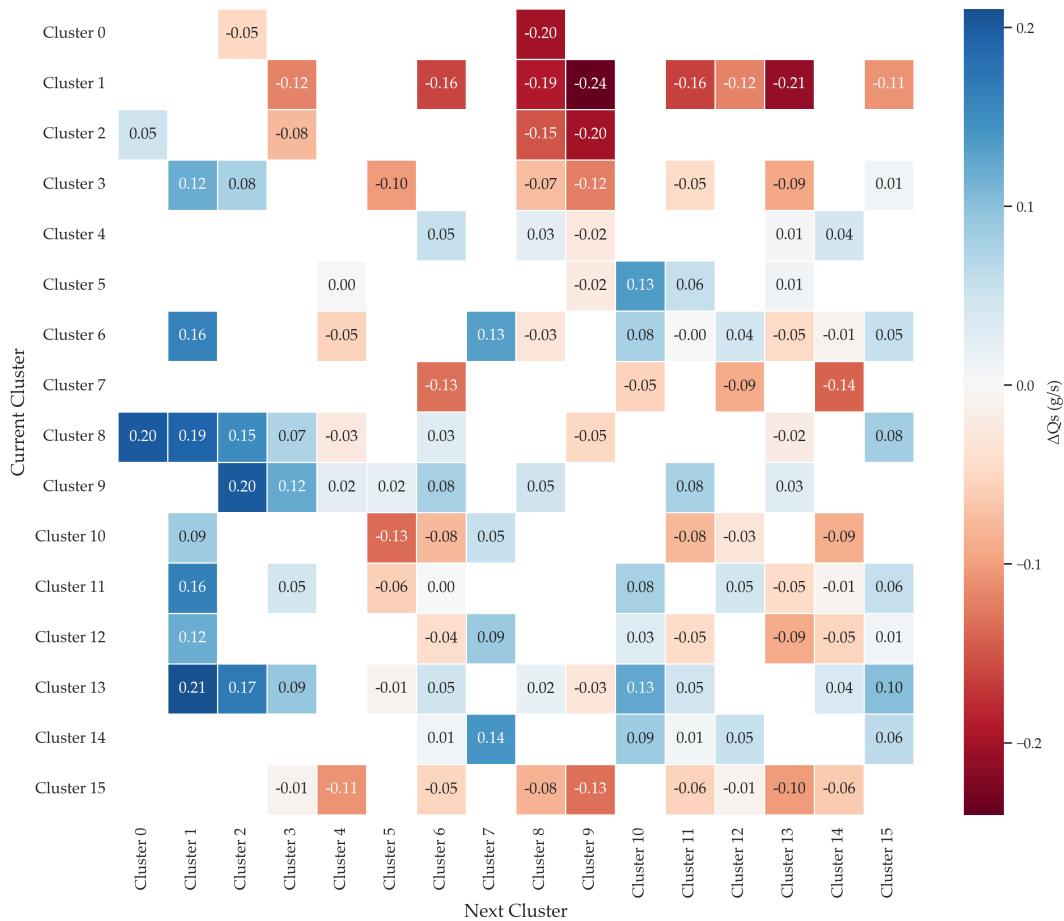


Figure 3.21: Change in mean sediment transport for ordered transitions between clusters. Each cell reports ΔQ_s in g s^{-1} for the jump from the row state (current) to the column state (next). Blue indicates an increase in transport and red a decrease.

Table 3.9: Five transitions with the largest absolute change in mean transport. Group membership follows Figure 3.19. Red means anchor, blue means intermediate, green means transit.

From	To	Group change	$\Delta Q_s [\text{g s}^{-1}]$	Rank by $ \Delta Q_s $
0	9	red \rightarrow blue	-0.24	1
0	13	red \rightarrow blue	-0.21	2
1	12	red \rightarrow green	-0.21	3
13	0	blue \rightarrow red	+0.21	4
8	0	blue \rightarrow red	+0.20	5

The five transitions with the largest absolute change in mean transport are listed in Table 3.9. They align with jumps that cross color groups identified in Figure 3.19. Below we interpret each pair using morphometrics reported in Table 3.6 (BI, eBI, mean wetted width) and the regimes in Figure 3.19.

- **0 → 9 (red to blue, $\Delta Q_s = -0.24$).** Cluster 0 is a narrow, weakly braided single thread with low BI and eBI and small wetted width in the long residence red group. Cluster 9 is a wide, bar rich braid with high BI and eBI and large wetted width in the blue group. The large negative ΔQ_s reflects a switch from a conveyance efficient pathway to a storage prone, partitioned morphology where transport drops.
- **0 → 13 (red to blue, $\Delta Q_s = -0.21$).** Same mechanism as above. Cluster 13 is also wide and braided with high BI and eBI. The jump out of the efficient single thread reduces mean transport sharply as flow spreads across multiple bars and threads.
- **1 → 12 (red to green, $\Delta Q_s = -0.21$).** Cluster 1 is among the narrowest and least braided states with very low BI and eBI and the smallest mean wetted width in the ensemble. Cluster 12 belongs to the high activity green group with short residence times and frequent reorganizations. The strong negative change indicates that leaving a narrow, efficient corridor for a transient configuration reduces mean transport.
- **13 → 0 (blue to red, $\Delta Q_s = +0.21$).** The mirror image of the 0 → 13 jump. Moving from a wide, storage prone braid with high BI and eBI and large wetted width into the confined single thread produces one of the largest transport increases in the dataset.
- **8 → 0 (blue to red, $\Delta Q_s = +0.20$).** Cluster 8 has the largest mean wetted width and high BI and eBI among the blue states. Transitioning into the narrow, low BI cluster 0 consolidates flow and raises transport markedly. The near tie with 9 → 0 reinforces that blue to red jumps are transport enhancing regime shifts.

Color level transition structure

We aggregated transitions by the three regime groups identified in Figure 3.19. The resulting directional counts are red to red 4, red to blue 52, red to green 7, blue to red 50, blue to blue 284, blue to green 180, green to red 8, green to blue 179, green to green 723. Row normalizing gives the probabilities in Table 3.10.

Table 3.10: Row wise transition probabilities between color groups

From group	To red	To blue	To green
Red	0.06	0.83	0.11
Blue	0.10	0.55	0.35
Green	0.01	0.20	0.79

From this data, three points emerge:

- ◊ **Probability of changing color versus staying within color.** Red almost always changes color when it moves, with a change probability near 0.94 and a same color probability near 0.06. Blue splits its moves between staying in blue and changing color, with a change probability near 0.45 and a same color probability near 0.55. Green is highly persistent at the color level, with a same color probability near 0.79 and a change probability near 0.21.
- ◊ **Where each group tends to go when it changes color.** When red changes color it goes to blue about 88 percent of the time and to green about 12 percent of the time. When blue changes color it goes to green about 78 percent of the time and to red about 22 percent of the time. When green changes color it goes to blue about 96 percent of the time and to red about 4 percent of the time. Direct green to red and red to green jumps are rare, which supports a two step pathway through blue for regime flips.
- ◊ **Volume of traffic by color pair.** The most common color preserved jump is green to green with 723 events. The most common cross color pathway is blue and green with 359 events combined, followed by blue and blue with 284 events and blue and red with 102 events. Red is the least connected basin by count. This pattern reinforces the view that blue acts as a bridge between the persistent green basin and the compact red set, consistent with the transport changes seen in the ordered heatmap and with the role of the green states as short lived but internally active configurations.

With this information, we can state that large transport changes align with regime crossings. Jumps from red to blue replace a confined, hydraulically efficient pathway by a partitioned, storage dominated braid and yield the largest transport decreases. Jumps from blue to red consolidate discharge into a single thread and produce the largest increases. Transitions involving the green group often encode brief, reorganizing episodes and tend to reduce transport when they originate in red states. These patterns are consistent with the persistence-activity structure in Figure 3.19 and with the morphometric ordering in Table 3.6. The barplot in Figure 3.20 confirms that many of the most common exchanges occur among green states or between green and blue, which drives the high transition counts without producing the most extreme ΔQ_s values. The heatmap and the color matrix together show that rare but decisive visits to red and the blue to red or red to blue jumps control the largest changes in transport.

Looking ahead, it would be useful to recover the three regime groups automatically rather than defining them *a posteriori*. A hierarchical or hidden semi-Markov model could learn a coarse regime layer above the 16 states. Other options include using nonparametric mixtures such as a Dirichlet process to "let the data determine the number of regimes". An automated procedure would yield reproducible boundaries, attach uncertainty to group membership, and facilitate cross-experiment comparisons. Certainly, this is a lot of job to be done, and more long experiments to perform.

3.4.3 Export efficiency, accessibility, and uncertainty across regimes

3.4.3.1 Export efficiency versus temporal occupancy

Another perspective that complements the persistence–activity and transition analysis is the relative contribution of each state to cumulative sediment export compared to its temporal occupancy (Figure 3.22). Read together with Figure 3.19, the barplot shows a clear color pattern. Green states (7, 10, 11, 12) punch above their weight in export despite short mean stays. Blue states with wide, braided planforms (5, 8, 9, 13, 14, 3, 6, 15) tend to occupy time without exporting much, which is consistent with storage-prone morphology. Red states (0, 1, 2, 4) are few and long-lived at the state level, and they exhibit high export efficiency per unit time, especially clusters 1 and 0.

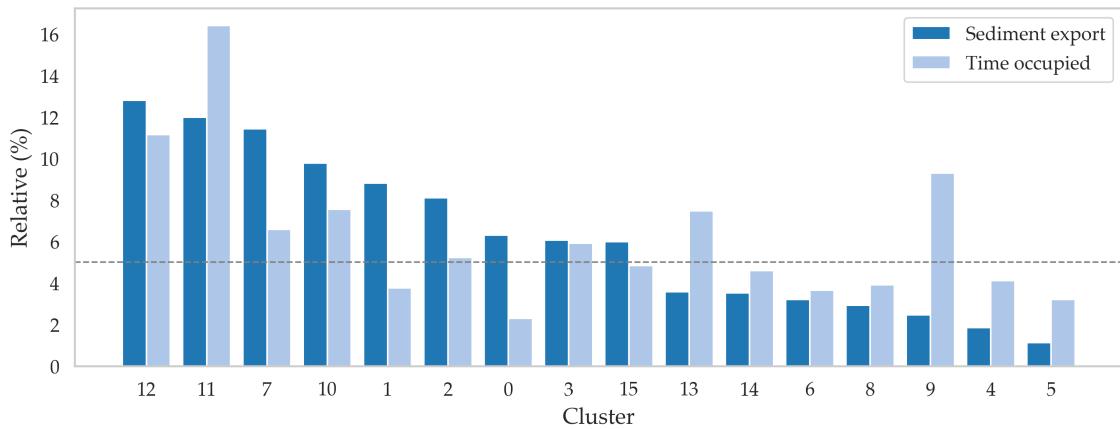


Figure 3.22: Relative importance of morphological clusters. Dark bars give the percentage of total sediment load exported while each state was active, and light bars show the percentage of experiment time spent in the same states. Values are sorted by sediment contribution, which highlights clusters that punch above (left of the dotted line) or below (right of the dotted line) their temporal weight.

Two examples anchor these contrasts. On the export-efficient side, cluster 10 contributes about one tenth of the total exported mass while occupying well under one tenth of the time, and cluster 1 contributes close to nine percent of export while present only a few percent of the time. On the reservoir side, clusters 9 and 13 spend a sizable share of the record yet contribute only a small share to export. These tendencies follow the morphometrics in Table 3.6. Clusters 9 and 13 are wide with high BI and eBI and thus partition discharge and favor storage. Clusters 1 and 10 are narrower and less braided, so they concentrate flow and export efficiently. The barplot therefore explains why the most active group in the transition record (green in Figure 3.19) can dominate export even with short individual dwell times.

3.4.3.2 Accessibility via mean first-passage times

Mean first-passage time mosaics provide a dynamic complement and quantify how quickly the system can access either the export-dominant basin or the peak-load set (Figure 3.23). Panel (a) reports the expected time to hit clusters 11 or 12. Panel (b) does the same for the set of peak-load states 0, 1, 2, 7, and 10. Results align with the color structure established

earlier. Access to the export-dominant basin is fastest from within the green set itself, with MFPT equal to 0 by definition for origins 11 and 12 and small values from green neighbors such as 7 and 10 (63 and 43 min, respectively). Access from blue storage states is slower. For example, starting in cluster 9 the expected time to reach the export-dominant basin is about 513 min, and from cluster 8 it is about 559 min. Access from red states is slowest. From cluster 0 the MFPT to the export-dominant basin is about 815 min and from cluster 2 it is about 624 min.

The peak-load set is much more reachable from many origins. From green and red states the MFPT is 0 for origins already in the set (7, 10, 0, 1, 2). From blue origins the times are moderate. For example, from cluster 8 the MFPT to the peak-load set is about 319 min and from cluster 9 it is about 409 min. From blue cluster 13 the MFPT is about 360 min. These numbers indicate that peak-load conditions are accessible on shorter horizons than the export-dominant basin, which requires longer reorganization.

Two route examples illustrate these patterns. Starting from a braided blue origin such as cluster 9, the system is expected to reach the export-dominant basin in about 513 min, whereas it reaches the peak-load set in about 409 min, typically via short visits to green transit states 7 or 10. Starting from a red single-thread origin such as cluster 1, the system is already in the peak-load set by definition and thus has MFPT equal to 0 to that set, yet it requires on the order of 389 min to reach the export-dominant basin. These read-offs translate the transition structure into concrete lead times and connect them to the color groups.

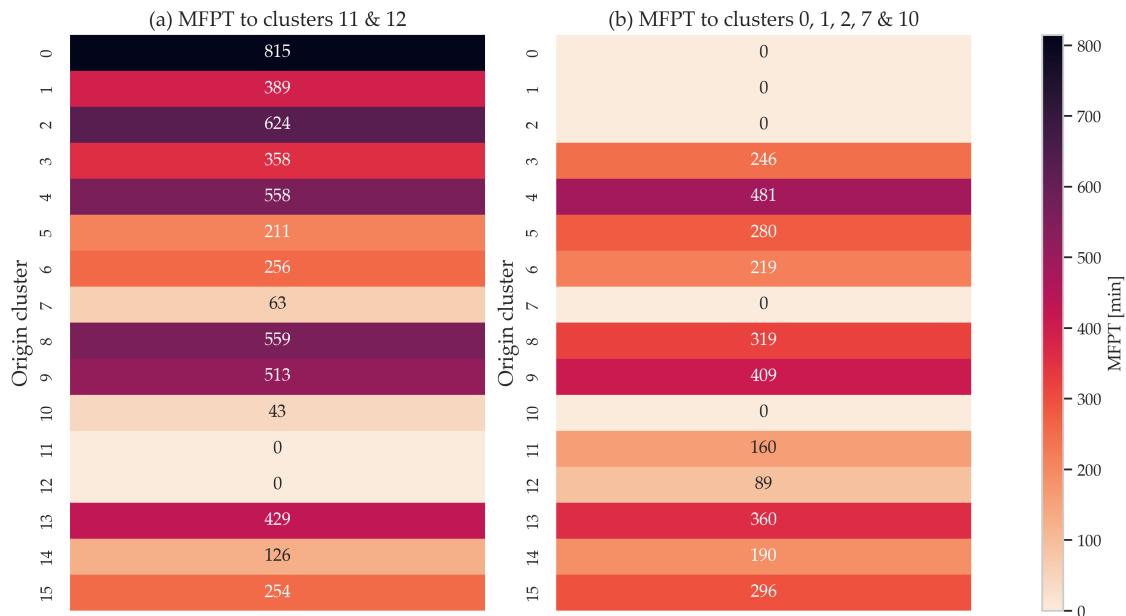


Figure 3.23: Mean first-passage time mosaics. Left shows expected time in minutes from each origin state to the export-dominant basin formed by clusters 11 and 12. Right shows the corresponding MFPT to the peak-load set formed by clusters 0, 1, 2, 7, and 10. Shorter times identify more direct pathways toward either long-term export or flash-transport conditions.

In summary, the three figures form a consistent story that matches the heatmap and color-matrix analysis in the previous subsection. Green states dominate export through many short visits and also act as gates to peak-load conditions. Blue states with high BI and eBI hold the system for long periods and delay access to the export-dominant basin. Red single-thread states are efficient exporters when active, but they are relatively isolated in color space and require long reorganization to enter the high-occupancy export basin.

3.4.3.3 Entropy rate and sources of dynamical uncertainty

The entropy rate measures how much dynamical uncertainty each state contributes to the sequence of morphologies. In a continuous-time Markov model it depends on two factors: the escape rate λ_i and the Shannon entropy of the corresponding transition probabilities. In simple words, entropy rate \propto “leave often” \times “many possible destinations.” Hence, short-lived, highly mobile states contribute most to overall uncertainty, unless their exits are concentrated toward a few partners that lower unpredictability.

The bar chart in Figure 3.24 confirms these expectations and connects cleanly with the color structure in Figure 3.19. The largest contributions come from clusters **6**, **8**, **1**, and **13**. Three of these belong to the blue group and combine moderate persistence with relatively diffuse transition rows, so $\lambda_i H(P_i)$ is large. Cluster **1** is a red state with high export efficiency and a nontrivial fan of destinations, which lifts its entropy contribution despite being in the long-residence family. At the other end, cluster **0** contributes the least. This is consistent with its long dwell times and with a transition pattern that is neither frequent nor highly dispersed.

Green transit states (7, 10, 11, 12) show sizable but not maximal bars. Their escape rates are large, yet a portion of their traffic is concentrated in symmetric pairs, most clearly the $10 \leftrightarrow 11$ and $7 \leftrightarrow 12$ exchanges highlighted by the bidirectional counts. That concentration lowers the Shannon entropy of the corresponding transition rows and prevents these states from dominating the entropy budget. This mechanism explains why the green group drives activity and often mediates regime crossings, while the blue group supplies much of the uncertainty injected per unit time.

Taken together with the transition heatmap and the persistence–activity map, the entropy–rate partition reinforces a coherent picture. Rapid–switch configurations inject most of the unpredictability when their exits are distributed across many targets. Quasi–deterministic swapping between preferred partners tempers uncertainty even under high activity. Long–lived, stable morphologies act as entropy sinks. At the color level this implies a division of roles. Blue states are the main entropy sources because they are numerous and their exits are comparatively diffuse. Green states are the kinetic hub that circulates the system quickly but not always unpredictably. Red states stabilize trajectories; within red, cluster **1** is the notable exception that contributes a meaningful share of entropy due to its richer set of exits.

Viewed together with the MFPT mosaics and the export–time barplot, the entropy analysis clarifies functional roles. Blue states act as main sources of dynamical uncertainty, delaying access to the export–dominant basin. Green states provide the fastest conduits to peak loads but generate less uncertainty due to paired exchanges. Red states, especially cluster 0, reduce uncertainty through long sojourns, while cluster 1 stands out as a faster yet less predictable exit from the red set. This triad of roles aligns with color–level transition probabilities and with the transport shifts observed at regime crossings.

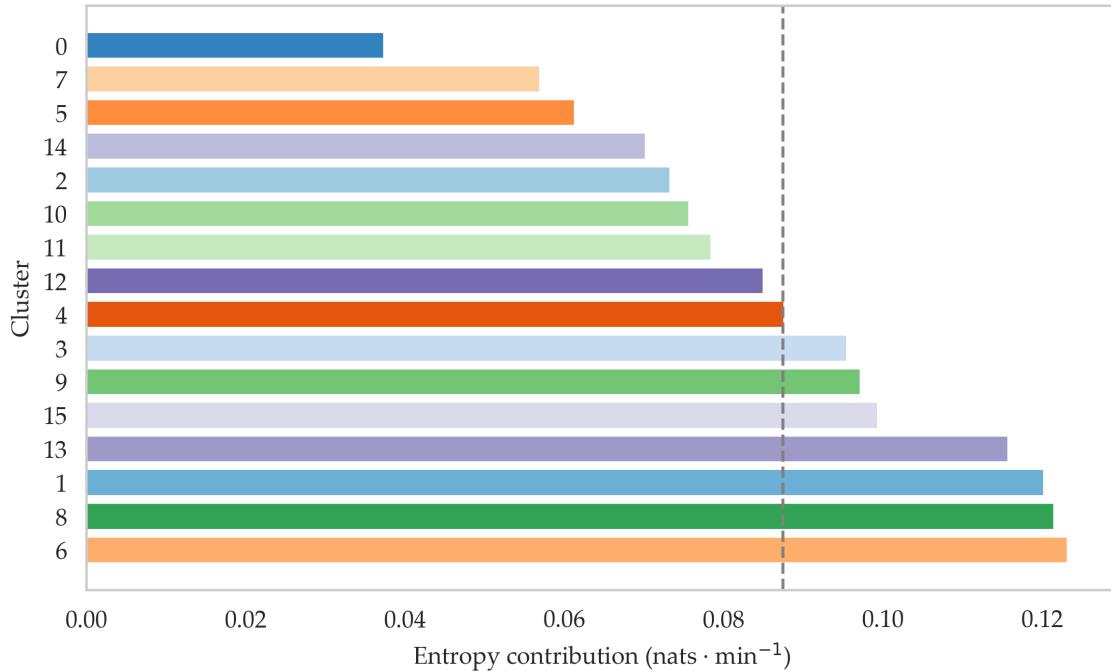


Figure 3.24: Entropy–rate contributions of each morphological cluster, in nats per minute. Bars quantify $\pi_i \lambda_i H(P_i)$ up to a constant of base; higher values correspond to states that the chain leaves often and that have many plausible destinations. The dashed line marks the ensemble mean contribution.

3.4.4 Cross-Validation of the Markov Model

To test the external validity and robustness of the Markov model derived from the 1200 h experiment (ExpL), we compare it against two independent runs of similar setup and boundary conditions, each about 200 h (ExpA and ExpB). A rigorous proof of Markovianity in geomorphic systems is generally not possible (Pawlowski & McCord, 2009). Our goal here is then empirical, to verify that the same morphological states reappear in new realizations and that the same dynamical features, including transition structure, short term memory, and multi scale variability, are expressed in the shorter records.

We organize the validation around two complementary questions. First, morphological coherence asks whether the states defined in ExpL reappear in ExpA and ExpB with comparable geometry and transport signatures. Second, dynamic fidelity asks whether the Markovian structure learned from ExpL reproduces the temporal organization observed in the shorter runs.

3.4.4.1 Morphological Coherence Across Experiments

We projected the binarized masks from ExpA and ExpB onto the frozen low dimensional space learned in ExpL (PCA then UMAP), using the same preprocessing and normalization. Each image was assigned to the nearest ExpL cluster centroid in that 3D space with a Euclidean metric. Two checks summarize the outcome. First, cluster coverage shows that all sixteen ExpL clusters received assignments from ExpA and ExpB, which indicates that the catalog of states from the long run is not unique to that realization. Second, the outlier rate was below 4.2 %, meaning that fewer than this fraction of projected images lay beyond the 95th percentile intra cluster distance observed in ExpL. This points to strong compatibility of the embedding and a low frequency of ambiguous assignments.

Figure 3.25 illustrates these results. Panel (a) shows that the distribution of distances from projected images to their assigned centroid closely follows that of ExpL, with only a thin tail beyond the dashed outlier threshold. Panel (b) plots distances from each projected point to all centroids. Clear minima (red and blue markers) appear across most cluster IDs, showing well defined nearest neighbor matches rather than flat valleys. Taken together, these results support that the morphological state space learned from the long run generalizes to independent realizations and that the clusters correspond to physically recurrent configurations rather than artifacts of a single dataset.

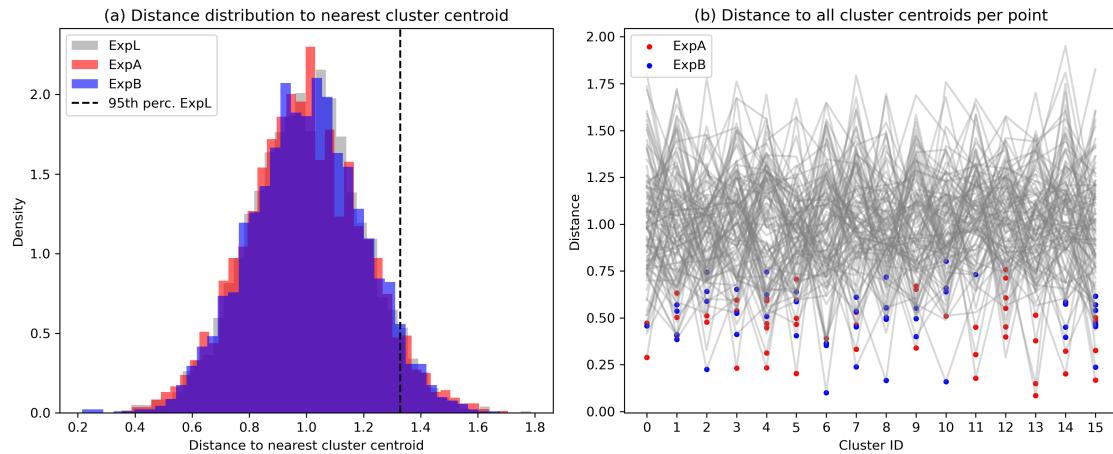


Figure 3.25: Validation of morphological clustering for ExpA and ExpB. **(a)** Distribution of distances from projected points to their nearest cluster centroid (ExpA and ExpB), compared to ExpL (shaded). The dashed line marks the 95th percentile of ExpL distances. **(b)** Pairwise distances between projected points and all centroids; red/blue markers indicate the minimum per point. The spread of minima across most cluster IDs supports broad coverage and unambiguous assignment.

3.4.4.2 Dynamic Fidelity of the Markov Process

Having established state space coherence, we now test whether ExpA and ExpB display the same dynamical structure. As a basic comparability check, Figure 3.26 shows that instantaneous sediment flux distributions overlap almost entirely across the three runs. This confirms that the dynamic comparisons that follow are not affected by differences in transport regimes.

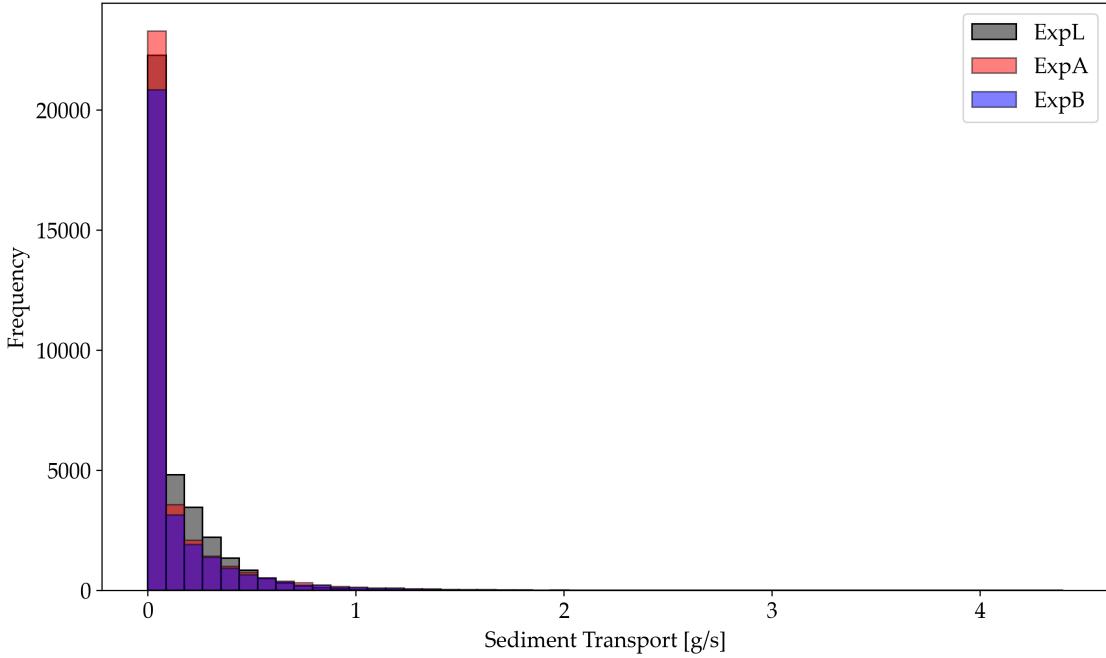


Figure 3.26: Global sediment transport distributions for ExpL, ExpA, and ExpB. The histograms overlap almost entirely, with a shared mode near zero and heavy-tailed behavior at high transport. The close agreement confirms that the shorter experiments sample the full range of flux regimes observed in ExpL.

We examine event to event, short term, and multi scale behavior through three diagnostics *(i)* the transition probability structure, *(ii)* the autocorrelation function (ACF) of sediment flux, and *(iii)* the power spectral density (PSD) of the same signal.

1. Transition probability structure Figure 3.27 compares the 16×16 matrices for ExpL, ExpA, and ExpB. High probability entries cluster near the diagonal and along a few off diagonal corridors (for example 10 to 12 and 4 to 7), while low probability regions coincide across runs. More than 80 % of ExpL entries with $P > 0.15$ reappear in both short runs. Frobenius norm differences relative to ExpL ($\|P_{\text{ExpA}} - P_{\text{ExpL}}\|_F = 0.27$, $\|P_{\text{ExpB}} - P_{\text{ExpL}}\|_F = 0.24$) remain within the variability expected from a much shorter experiment.

2. Short-term temporal correlation. The ACFs of the flux series (Figure 3.28) decay rapidly from $\rho(0) = 1$ to $\rho(\tau) \approx 0.3$ within 50–70 min and then taper slowly, crossing zero between about 1000 and 1400 min. ExpA decorrelates slightly faster at lags beyond 200 min, yet the overall shapes remain closely aligned, indicating a consistent memory time scale. This behavior matches the expectations of a CTMC fitted to ExpL, with short lag decay controlled by the leading eigenvalues and long tails governed by slower modes.

3. Multi-scale variability. Figure 3.29 compares the PSDs. Across 10^{-5} – 10^{-3} Hz the spectra collapse onto a common power-law. Minor divergence at the highest frequencies reflects the shorter duration of ExpA/ExpB, while agreement below $\sim 3 \times 10^{-4}$ Hz is good. Since the PSD of a first-order Markov process is set by its eigenspectrum, this concordance supports that the eigenstructure inferred from ExpL represents the system's variability.

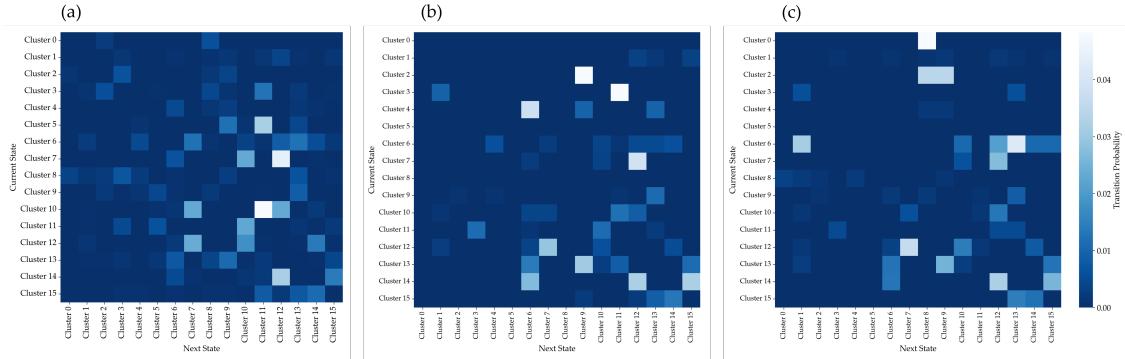


Figure 3.27: Transition matrices for the 16-state model across the three experiments: **(a)** ExpL, **(b)** ExpA, **(c)** ExpB. Light cells denote higher probabilities. The recurrence of diagonal bands and key off-diagonal corridors across runs indicates robust connectivity patterns.

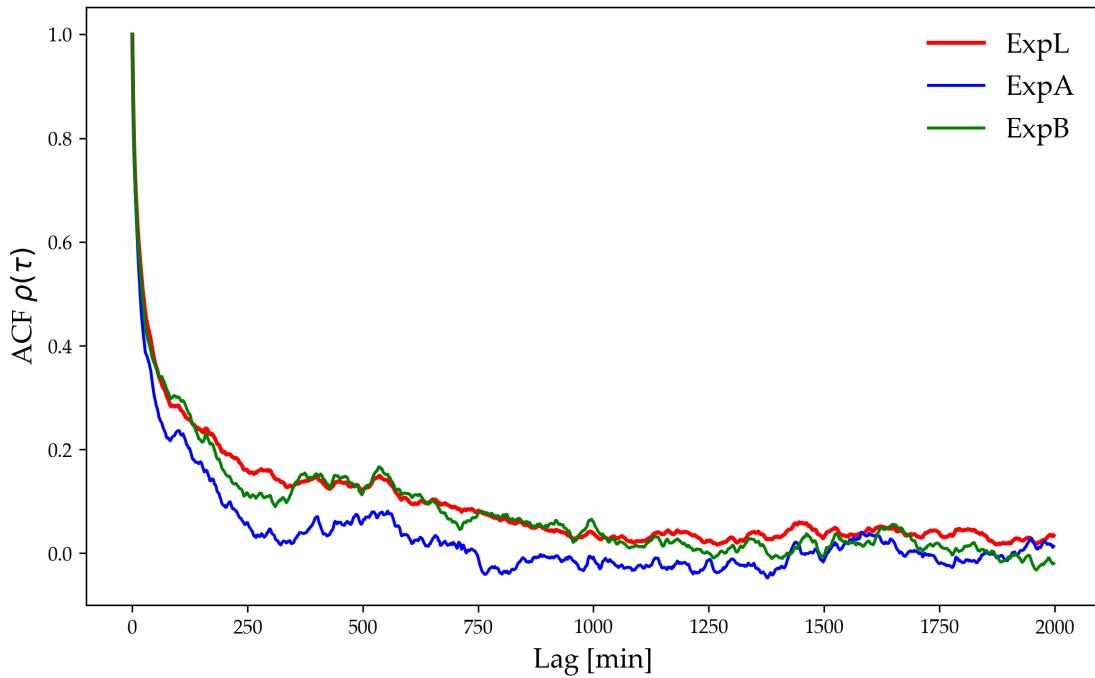


Figure 3.28: Autocorrelation functions of sediment transport for ExpL (red), ExpA (blue), and ExpB (green). A rapid decay within 50–70 min is followed by a slow tail crossing zero after $\sim 1,000$ – $1,400$ min. The similarity across runs suggests a common memory time-scale.

In sum, the overlap of transport distributions, the correspondence of transition probabilities, the shared ACF decay, and the very similar spectral scaling indicate that the Markov model from the 1200 h record captures both the range of morphological states and the key dynamical features seen in the independent 200 h runs. Under comparable boundary conditions this validation supports using the fitted model for interpretation and stochastic forecasting at experiment time scales.

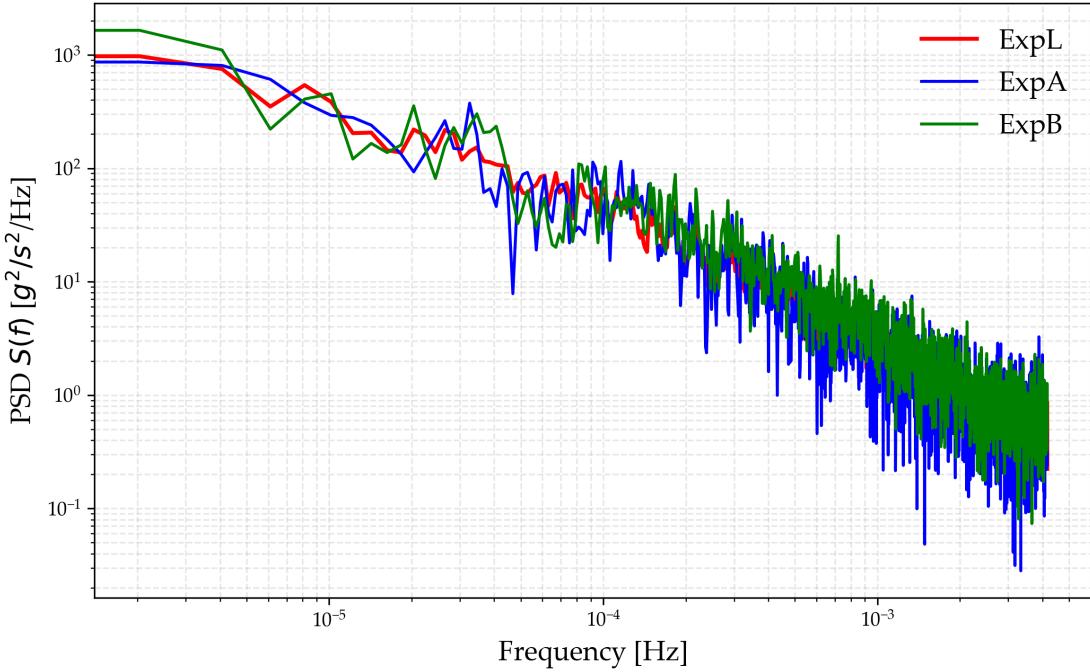


Figure 3.29: Power spectral density of sediment transport for ExpL (red), ExpA (blue), and ExpB (green). All exhibit a common power-law decay ($\sim f^{-1.6}$) over 10^{-5} – 10^{-3} Hz, indicating scale-invariant fluctuations across minutes to daily scales. Agreement is strongest at low frequencies, where record length is less limiting.

3.5 Summary of Key Points

- ◊ **A multi-state representation is feasible and useful.** The joint use of edge-based (MHD) and area-based (Dice) distances, low-dimensional embeddings, and density-based clustering yields morphologically coherent states that span a clear continuum from narrow single-thread planforms to wide, bar-rich braids.
- ◊ **Morphometric trends are systematic across states.** Wetted width, wetted area, sinuosity, BI, and eBI all increase toward multi-thread planforms. eBI acts as a stable partitioning that is more robust than BI to resolution and stage noise.
- ◊ **Transport is state-conditioned and anticorrelated with partitioning.** Cluster-average bedload rate $\langle Q_s \rangle$ is highest in simple, weakly braided states and lowest in wide, partitioned states. Across states, $\langle Q_s \rangle$ is strongly negatively correlated with eBI and with wetted width, consistent with the idea that partitioning promotes storage.
- ◊ **High-transport pulses are linked to regime changes.** The ordered heatmap of ΔQ_s shows the largest decreases for red to blue transitions and the largest increases for blue to red transitions, while green states act as a kinetic bridge. This explains why large pulses follow entry into conveyance-efficient states and why load decreases when the system spreads into storage-prone braids.
- ◊ **Persistence-activity structure is well represented by a CTMC.** Dwell-time histograms are well approximated by exponentials. Estimated exit rates $\hat{\lambda}_i$ order states

consistently with morphology. Long-residence anchors (red) have small $\hat{\lambda}_i$, transit states (green) have large $\hat{\lambda}_i$, and blue states fall in between.

- ◊ **Color-level transitions clarify regime connectivity.** Row-normalized counts show that red almost always exits to blue, green mostly exits to blue, and direct red-to-green moves are rare. Blue therefore acts as the bridge between a compact red set and a persistent green basin, matching the frequency barplot and the ΔQ_s heatmap.
- ◊ **Accessibility can be quantified with MFPTs.** Mean first-passage time mosaics identify shortest pathways and lead times. Peak-load states (0, 1, 2, 7, 10) are more reachable on shorter horizons than the export-dominant basin (11, 12). Access to 11–12 is fastest from green, slower from blue storage states, and slowest from red anchors.
- ◊ **Export efficiency versus time occupancy separates throughput from persistence.** Green states (7, 10, 11, 12) contribute disproportionately to cumulative export relative to their time share, blue storage states (5, 8, 9, 13, 14) occupy time but export little, and red single-thread states are efficient exporters when active.
- ◊ **Uncertainty sources are identified by the entropy rate.** Entropy contributions reflect the product of exit rate and row-wise dispersion. Blue states are major entropy sources because their exits are comparatively diffuse, green states inject activity but paired exchanges temper unpredictability, and red states stabilize trajectories.
- ◊ **Direct regression from morphometrics to $\hat{\lambda}_i$ is weak with the present descriptors.** Attempts to connect exit rates from simple morphometrics (wetted width, BI/eBI, junction density, active width) produced weak and inconsistent associations across states, suggesting that richer features or hierarchical models are needed for predictive mapping of $\hat{\lambda}_i$.
- ◊ **Results are stable across independent runs.** Global transport histograms overlap across experiments, high-pass filtered series share the same decay scale, and power spectral densities collapse to a similar slope over a big range of frequencies, supporting reproducibility of the patterns.

Representing braided rivers as a finite set of discrete morphological states, each with a measurable residence rate and a transition kernel, is both feasible and advantageous. It explains intermittency as time spent mixing among states with different conveyance efficiency, provides a compact (λ, P) summary that supports analytical diagnostics and forward simulation, and remains stable across independent runs. This multi-state framing is more informative than a single typical descriptor because it links planform, transport, and dynamics in a probabilistic model that will be used for forecasting in Chapter 4.

4 | Markov–Based Monte Carlo and Bootstrap Analysis

We begin with the key results from Chapter 3. The transition matrix \mathbf{P} and the exit rates $\hat{\lambda}$ define the continuous time Markov chain used here. Chapter 3 also provided a catalog of states with medoids and morphometric profiles, together with transport distributions $f_i(Q_s)$ for each state. Dwell times followed exponential patterns, which supports using $\hat{\lambda}_i$ and \hat{P} as a compact summary for simulation and diagnostics. These are the elements we now build on.

In this chapter we shift from state construction to transport surrogates and model performance tests, guided by the central question of whether the state–based Markov description carries sufficient information to reproduce transport statistics, capture temporal structure, and provide operational forecasts when direct flux measurements are unavailable. We first simulate multiple state sequences with the CTMC using both $\hat{\lambda}$ and \hat{P} , and then assign transport by bootstrap resampling from $f_i(Q_s)$, producing synthetic Q_s series that retain both inter–state differences and random variability. A Bayesian screening step filters the series to match experimental mean and variance, while image–based similarity metrics offer an additional check. Building on this, we train a decision tree that relies exclusively on image metrics to identify acceptable sequences. Finally, we validate the retained series through moments, extremes, autocorrelation, and spectra, quantify the variance attributable to state switching, test sensitivity to the number of clusters, and propose a lean hybrid model for image–only applications.

Section 4.1 explains the Monte Carlo, bootstrap, and Bayesian screening, and compares synthetic and experimental series. Section 4.2 studies the variance decomposition and how results change with the number of clusters. Section 4.3 presents a hybrid model that keeps state means but simplifies random variation, and shows how scaling with discharge or images can supply mean surrogates when discharge data are missing.

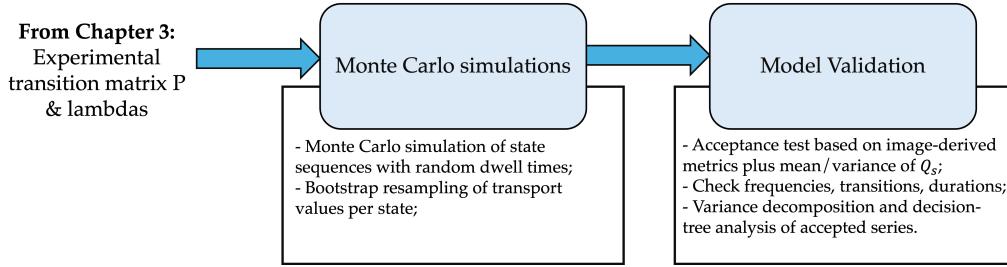


Figure 4.1: Schematic overview of the Monte Carlo and bootstrap pipeline used in Chapter 4. From left to right we show the experimental transition matrix \hat{P} and exit rates $\hat{\lambda}$ from Chapter 3, Monte Carlo simulations of state sequences with random dwell times, and model validation.

4.1 From the Transition Matrix to Synthetic Sediment–Transport Series

4.1.1 Monte Carlo–Bootstrap Framework

Our goal is to translate the information in the experimental transition matrix \mathbf{P} into a large ensemble of synthetic realizations. The procedure has two layers, a Monte Carlo simulation of the state sequence followed by a bootstrap resampling of sediment transport rates, but the bootstrap is applied only to state sequences that first pass a moment based screen. We generate $N_{\text{traj}} = 100000$ state trajectories of length T minutes, which serve as a broad prior over paths implied by \mathbf{P} and a global dwell rate. We chose T to be equal to the length of ExpL ($\approx 1200 \text{ h} = 72000 \text{ min}$).

1. State sequence simulation (Monte Carlo). We start from a random initial state and draw a dwell time τ from an exponential distribution $\mathcal{E}(\lambda)$. The rate λ is sampled once for the entire trajectory from the range $[0.001, 0.05] \text{ s}^{-1}$ and then converted to min^{-1} to match the one minute sampling. In principle each state could have its own λ , which would be richer but far more expensive, so we use a single λ for all states as a practical approximation that preserves the observed residence time range while enabling large ensembles. The system remains in the current state for τ minutes (rounded to the nearest minute) and then jumps to a new state sampled from Row i of \mathbf{P} , where i is the current state. The steps repeat until the trajectory reaches length T . This procedure is summarized in Algorithm 1.

Algorithm 1: State-sequence simulation (Monte Carlo)

- 1 Initialize random initial state i ;
- 2 Initialize time $t \leftarrow 0$;
- 3 Sample one $\lambda \sim \mathcal{U}[0.001, 0.05] \text{ s}^{-1}$ and convert to min^{-1} ;
- 4 **while** $t < T$ **do**
- 5 Sample dwell time $\tau \sim \mathcal{E}(\lambda)$ in minutes;
- 6 Hold state i for $\text{round}(\tau)$ one minute steps;
- 7 Update $t \leftarrow t + \text{round}(\tau)$;
- 8 Sample next state j from Row i of \mathbf{P} ;
- 9 Update current state $i \leftarrow j$;

Moment screen on state-only replicas. Let $m_s = \mathbb{E}[Q_s \mid S = s]$ and $v_s = \text{Var}(Q_s \mid S = s)$ denote the empirical mean and variance of transport in state s estimated from the experiment. For a simulated state path $\{S_t^*\}_{t=1}^T$ we compute the implied series moments without drawing transport values,

$$\mu^* = \frac{1}{T} \sum_{t=1}^T m_{S_t^*}, \quad \sigma^{2*} = \frac{1}{T} \sum_{t=1}^T v_{S_t^*} + \text{Var}_t(m_{S_t^*}),$$

which combine within-state scatter with between-state contrasts. A replica is accepted when

$$|\mu^* - \mu| \leq 0.002 \text{ g s}^{-1} \quad \text{and} \quad |\sigma^{2*} - \sigma^2| \leq 0.001 (\text{g s}^{-1})^2.$$

Out of the 100000 state trajectories approximately 0.12% satisfy these tolerances. Only these accepted paths proceed to the bootstrap layer. The same acceptance fraction determines the number of synthetic transport series used in subsequent analyses.

2. Bootstrap of sediment transport. Each state s is linked to an empirical bag $\mathcal{B}_s = \{Q_{s,1}, \dots, Q_{s,n_s}\}$. For every minute t in an accepted replica we draw $Q_s^*(t) \sim \mathcal{B}_{S_t^*}$ with replacement. Because this sampling is stratified by state, the synthetic series inherits both inter-state differences in mean transport and intra-state variability observed in the data. Algorithm 2 summarizes the procedure.

Algorithm 2: Bootstrap of sediment transport on accepted paths

- 1 **for** each minute t in an accepted state sequence **do**
- 2 Identify simulated state S_t^* ;
- 3 Define empirical bag $\mathcal{B}_{S_t^*}$;
- 4 Sample with replacement $Q_s^*(t) \sim \mathcal{B}_{S_t^*}$;
- 5 The resulting $Q_s^*(t)$ preserves inter-state contrasts and intra-state variability;

Figure 4.2 shows the experimental Q_s in blue together with a representative synthetic series Q_s^* in orange generated from an accepted state sequence. The synthetic series is stochastic at the minute scale yet reproduces range, intermittency, and burst structure, indicating that combining Monte Carlo state paths with state conditioned bootstrap captures the main statistical features of the system. Section 4.1.2 formalizes the acceptance rule and the surrogate screening based on image derived metrics.

The Monte Carlo layer generates state paths under \mathbf{P} with a single trajectory-level dwell rate consistent with the one minute sampling. The moment screen reduces this broad prior to a small and well calibrated set, and the bootstrap adds within-state variability while preserving differences between states. The workflow is reproducible, scalable, and efficient, as most paths are discarded before resampling. Its modular design ensures that future updates to \mathbf{P} , the state catalog, or the dwell time model can be incorporated without altering the screening or resampling steps. This completes the construction of the synthetic ensemble, which is used next for validation in Section 4.1.3 and later for variance decomposition in Section 4.2.

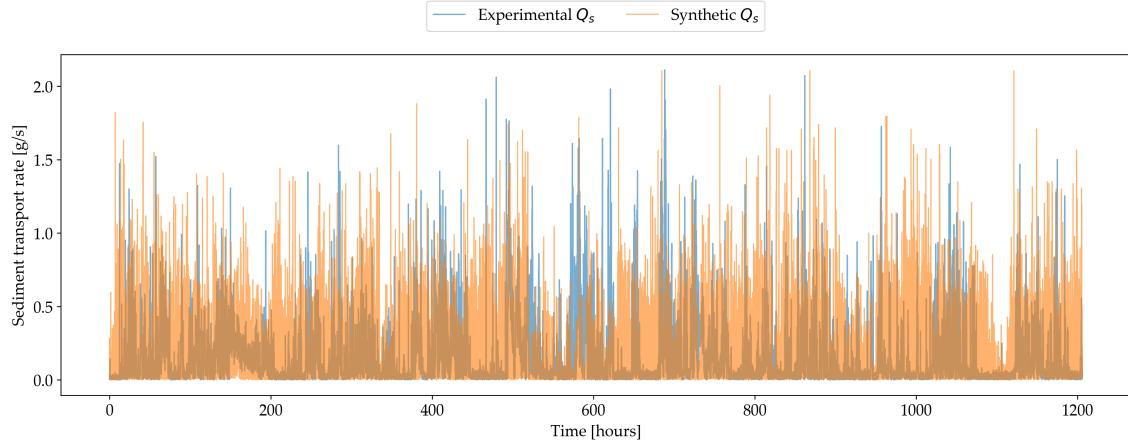


Figure 4.2: Overlay of the experimental sediment transport series (Q_s , blue) and a representative synthetic series generated by the Monte Carlo and bootstrap model (Q_s^* , orange). Both are plotted in hours to visualize variability and event structure

4.1.2 Bayesian Screening and Decision-Tree Extraction

Screening state paths before bootstrapping is valuable chiefly because it reveals image based predictors that flag acceptable paths even when Q_s is not measured, which enables fast preselection and practical use in imagery only settings. The same pipeline brings several additional gains. Bootstrap delivers uncertainty quantification by providing confidence bands for statistics and curves such as means, variances, ACFs, PSDs, MFPTs, and export totals. State conditioned resampling preserves heavy tails and heteroscedasticity without parametric assumptions and, with a block option, short range temporal dependence. The accepted replicas form a posterior predictive ensemble that supports validation against many diagnostics, not only the first two moments. Resampling is robust to outliers because it dilutes the impact of isolated extremes in Q_s while respecting their empirical frequency. Scenario testing is straightforward since one can modify \mathbf{P} , adjust dwell rates, or alter the state catalog and then rerun the screen with minimal code. Finally the screened set yields simple operational rules because feature importances identify image based thresholds that operators can apply when Q_s is unavailable. As an extra benefit, the screen avoids computation on sequences that would miss the experimental mean or variance, and it upgrades the prior of simulated paths to a posterior set that is physically comparable with the experiment.

The 100000 Monte Carlo state trajectories in Section 4.1.1 form a prior over paths implied by \mathbf{P} and a global dwell rate. Each path is first labeled using state only information. We compute the implied transport moments from the state means and variances,

$$\mu^* = \frac{1}{T} \sum_{t=1}^T m_{S_t^*}, \quad \sigma^{2*} = \frac{1}{T} \sum_{t=1}^T v_{S_t^*} + \text{Var}_t(m_{S_t^*}),$$

and accept a path when

$$|\mu^* - \mu| \leq 0.002 \text{ g s}^{-1} \quad \text{and} \quad |\sigma^{2*} - \sigma^2| \leq 0.001 (\text{g s}^{-1})^2.$$

About 0.12% of the 100000 paths pass this screen. Only those accepted paths are bootstrapped, which avoids resources on sequences that would miss the target moments.

We then ask whether simple image based descriptors can anticipate acceptance without any reference to Q_s . For every simulated path, accepted or rejected, we compute six metrics of the state sequence, namely state frequency difference, Frobenius distance between transition matrices, Jensen–Shannon distance of state durations, entropy difference, diagonal distance, and Levenshtein distance. A shallow decision tree trained on these features yields a conservative rule that closely reproduces the moment screen while remaining easy to apply.

Figure 4.3 shows the learned splits. The root threshold is $\text{freq_diff} \leq 0.015$. The left branch refines with $\text{levenshtein} \leq 0.974$ and a secondary $\text{freq_diff} \leq 0.012$. The right branch splits on $\text{frobenius_diff} \leq 0.95$, then on $\text{dur_JS} \leq 0.88$, and finally on $\text{frobenius_diff} \leq 0.952$. Accepted paths concentrate in the region with small state frequency difference and small Frobenius distance, which matches the visual envelope in Figure 4.4. For plotting we use round numbers $\text{freq_diff} = 0.016$ and $\text{frobenius_diff} = 0.956$ as dashed guides, consistent with the tree thresholds.

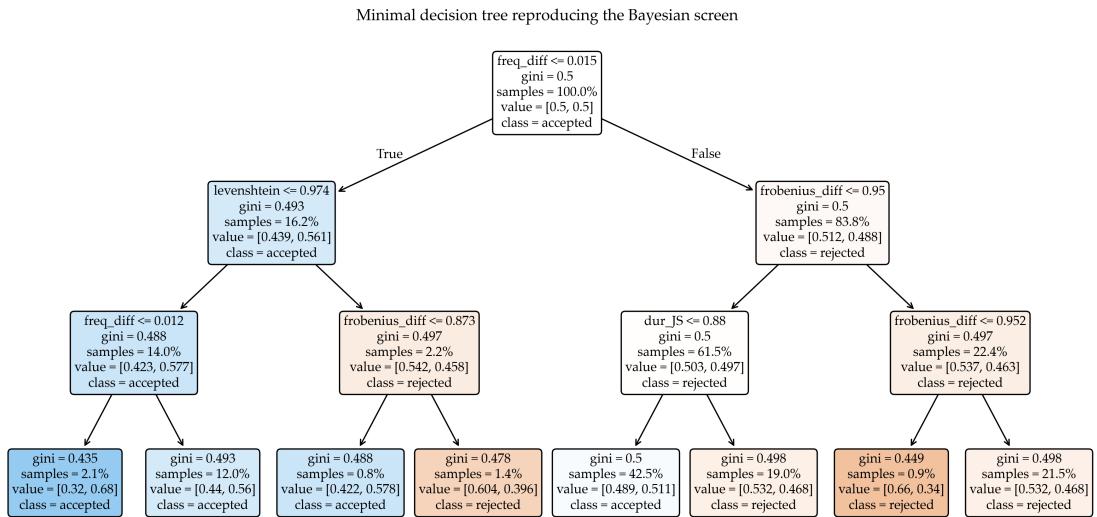


Figure 4.3: Minimal decision tree inferred from the six image based metrics. The root split is $\text{freq_diff} \leq 0.015$. Subsequent splits involve $\text{levenshtein} \leq 0.974$, $\text{freq_diff} \leq 0.012$, $\text{frobenius_diff} \leq 0.95$, $\text{dur_JS} \leq 0.88$, and $\text{frobenius_diff} \leq 0.952$. Boxes report Gini impurity, sample share, and class label

The resulting rule identifies the most informative parameters and provides a fast prefilter that yields synthetic series closely aligned with the experimental record. A simple operational version uses the triplet $\{\text{freq_diff} \leq 0.016, \text{dur_JS} \leq 0.88, \text{frobenius_diff} \leq 0.956\}$. Any path meeting these thresholds can be bootstrapped with high confidence that its mean and variance will match the laboratory series, while paths outside the envelope can be discarded early.

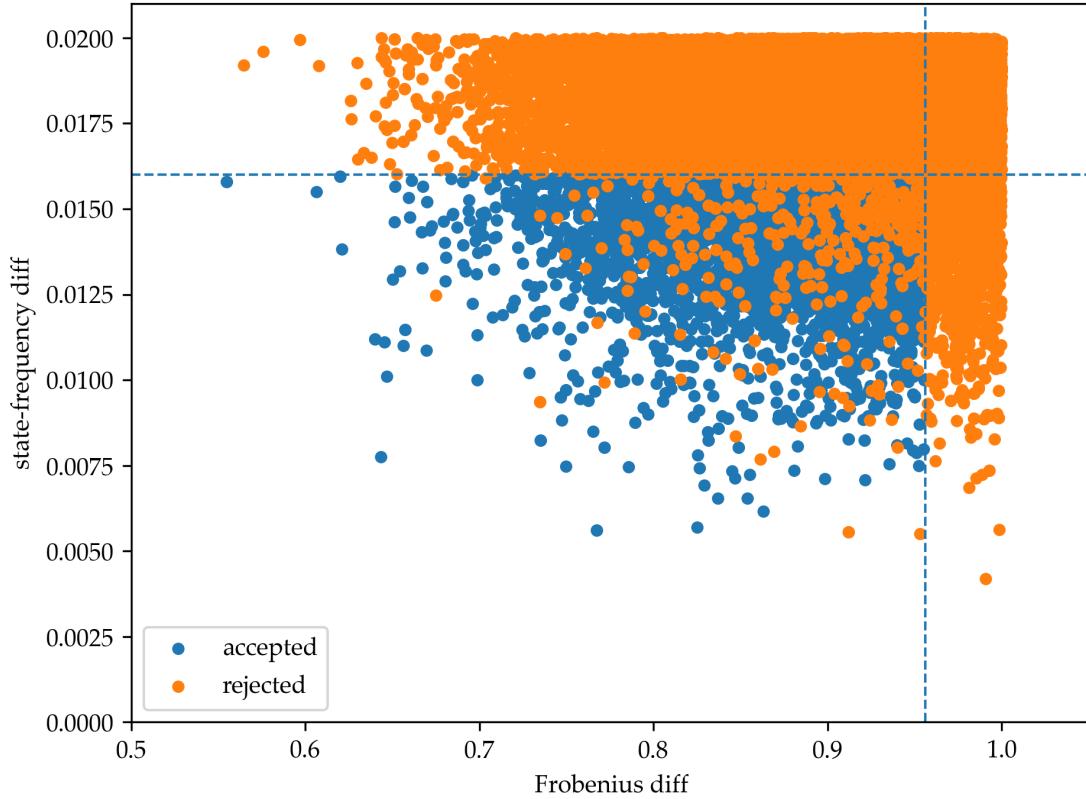


Figure 4.4: Trajectories projected on the {state frequency difference, Frobenius distance} plane. Dashed lines at 0.016 and 0.956 mark operational cutoffs. The lower left quadrant hosts nearly all paths that pass the moment screen, which confirms that small frequency mismatch and small Frobenius distance are the dominant predictors of acceptance

The two stage screen reduces the prior of 100000 simulated paths to a smaller posterior set that matches the experimental mean and variance by design. The decision tree then recovers most of these accepted paths using only a few image based thresholds. This provides a conservative and precise prefilter that saves computation, allows operation when Q_s is not available, and yields an ensemble suitable for variance decomposition and validation. The thresholds are data driven and easy to recalibrate for new runs or scenarios, and the feature ranking offers a compact checklist for operators who need fast triage without transport data.

4.1.3 Validation of the Accepted Time Series

Having obtained a high quality subset of synthetic trajectories, we now test whether the ensemble does more than get the average right. We ask whether it captures the full statistical fingerprint of the experiment, from bulk behavior to rare extremes and temporal structure. We use three complementary diagnostics (*i*) the marginal distributions of ensemble means and variances, (*ii*) the return period curve of extreme events, and (*iii*) the temporal dependence structure assessed with the autocorrelation function (ACF) and the power spectral density (PSD).

4.1.3.1 Global moments

The first requirement for a trustworthy ensemble is that each replica reproduces the bulk statistics of the experimental transport signal, namely its mean μ and variance σ^2 . After the two stage Bayesian screening in Section 4.1.2, 12436 realizations survived. Figure 4.5 shows the marginal distributions of the resulting μ^* and σ^{2*} .

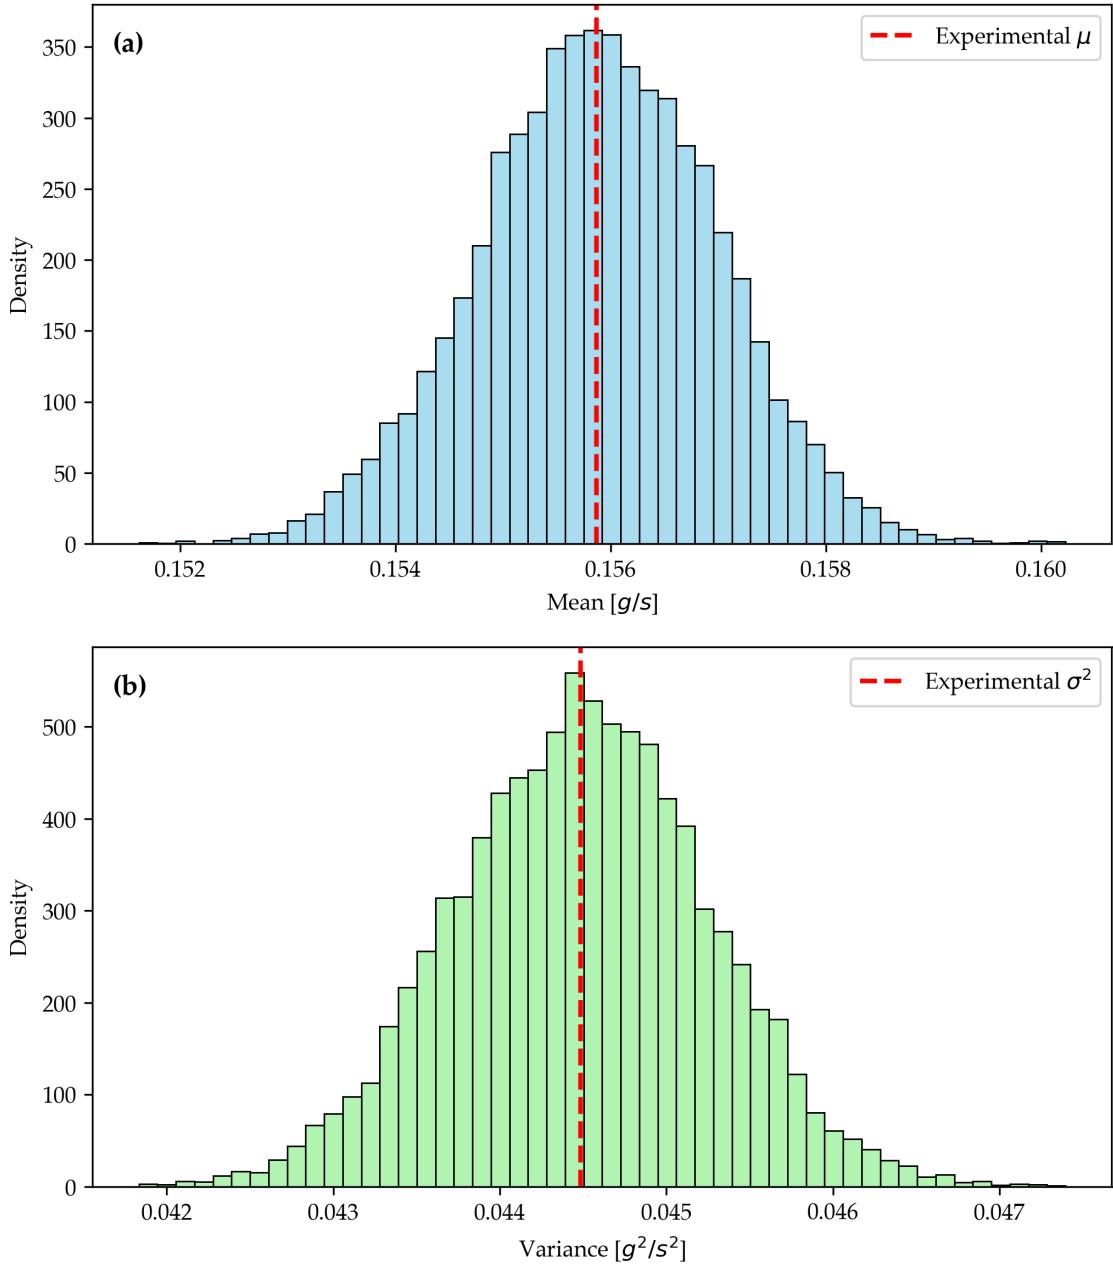


Figure 4.5: Normalized histograms of (a) the mean transport rate μ^* (blue) and (b) the variance σ^{2*} (green) for the 12436 synthetic series that passed the screening. The experimental benchmark is the red dashed line. Both distributions are tightly centered around the target values: $|\mu^* - \mu|/\mu = 1.3\% \pm 0.8$ and $|\sigma^{2*} - \sigma^2|/\sigma^2 = 0.9\% \pm 1.1$ (mean \pm SD).

Panel (a) displays a sharply peaked, nearly Gaussian histogram of replica means. The mode coincides with the experimental benchmark (red dashed line) to within 0.1 %. The spread, summarized by $SD(\mu^*)$, is only about 0.8 % of the target value. All series fall inside the prescribed ± 2 % corridor and the cloud is symmetric, which indicates negligible bias with $|\mu^* - \mu|/\mu = 1.3\% \pm 0.8$.

Panel (b) repeats the exercise for variances. Despite the broader natural range of σ^2 , the distribution remains centered on the experimental value with comparable relative error, namely $|\sigma^{2*} - \sigma^2|/\sigma^2 = 0.9\% \pm 1.1$. The bell shaped outline shows that screening did not distort variability and only removed pathological outliers.

Two additional checks strengthen this result. First, moments computed directly from the bootstrapped $Q_s^*(t)$ series are virtually indistinguishable from the state implied moments used in the screen, which confirms that the resampling step does not introduce bias in the first two moments. Second, the concentration of the histograms relative to the broad prior indicates strong posterior contraction of the ensemble, yet without mode splitting or skew, so the accepted set remains diverse rather than overfit.

Because the histograms are normalized, the vertical axis reports probability density. Peak densities exceed 300 in panel (a) and 500 in panel (b), a direct reflection of the tight bandwidth around the targets ($\sigma_{\mu^*} \simeq 7 \times 10^{-4} \text{ g s}^{-1}$ and $\sigma_{\sigma^{2*}} \simeq 4 \times 10^{-4} \text{ g}^2 \text{s}^{-2}$). The ensemble is therefore unbiased, precise, and individually compliant, which provides a calibrated baseline for the next tests on extremes and temporal structure.

4.1.3.2 Extreme value behavior

Rare high magnitude transport pulses are of upmost importance when studying scour risk and morphodynamic change, so the synthetic records must match not only the mean rate but also the frequency of large exceedances. Figure 4.6 compares the empirical return level curve from the experiment (red) with the ensemble summary from the accepted replicas (blue median with gray p5 to p95 band). Across the full range resolved by the record, from about 0.1 to 1200 h, the experimental curve tracks the ensemble median and remains inside the uncertainty band. The spread widens only near the record length, which is consistent with finite sample effects. This agreement indicates that the generator reproduces the tail slope and scale of exceedances and provides credible uncertainty bands for exceedance based diagnostics used in risk assessment.

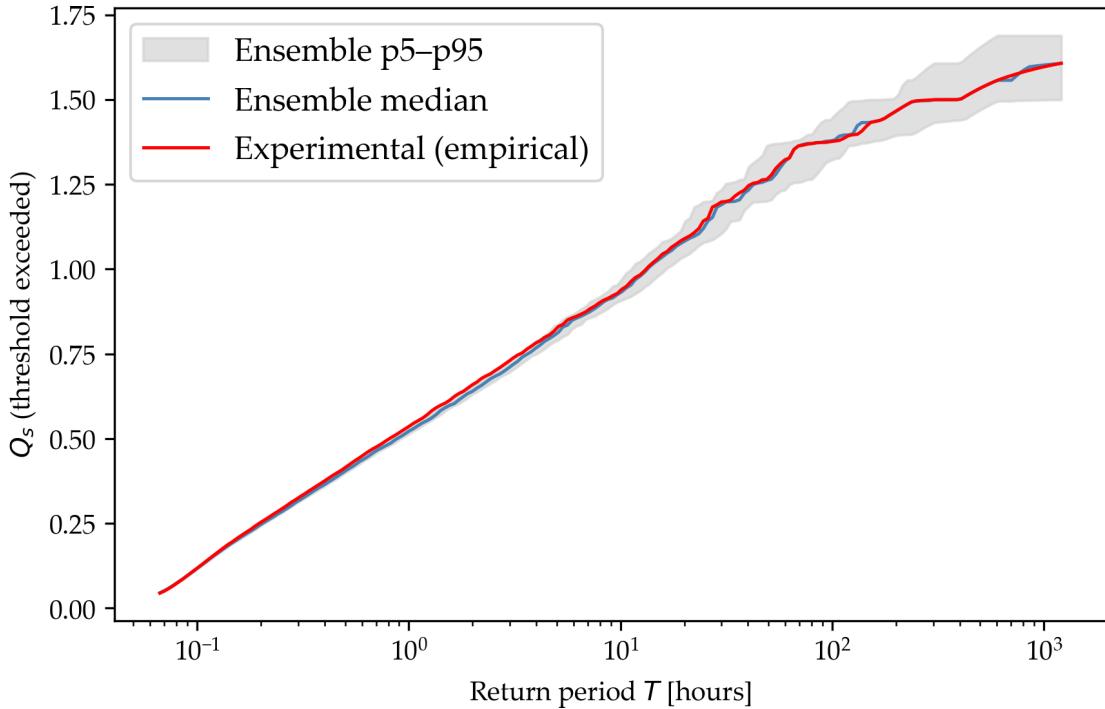


Figure 4.6: Return level of Q_s versus return period T on a logarithmic axis. The red curve shows the experimental empirical return levels based on peaks over a threshold. The blue curve is the ensemble median from the accepted replicas and the gray band is the p5 to p95 envelope. The experimental line stays within the ensemble band from about 0.1 to 1200 h and follows the median closely, with wider spread only near the record length.

4.1.3.3 Temporal dependence

We test whether the accepted ensemble reproduces the multi hour coherence of the experiment, not only pointwise statistics. Figure 4.7 contrasts two benchmarks. The left column uses a basic bootstrap that resamples $Q_s(t)$ from the full experimental series without regard to state. The right column uses the proposed Markov bootstrap that draws $Q_s(t)$ from the empirical bag of the currently visited cluster, so temporal structure can flow from dwell times and transitions.

In panel (a) the basic bootstrap collapses the ACF to near zero after the first few minutes. The experimental ACF in red remains positive for hundreds of minutes and decays slowly, with small oscillations at long lags. The gray band from the basic bootstrap sits tightly around zero and the blue median is essentially flat, which shows that shuffling values reproduces the marginal distribution of Q_s but erases persistence and burst clustering.

Panel (b) shows that the Markov bootstrap restores much of the missing memory. The blue median ACF begins at a positive value, decays gradually, and stays above zero for many hundreds of minutes. The experimental curve still lies above the ensemble for a wide range of lags, which indicates an underestimation of very long persistence, yet the shape and time scale are now comparable.

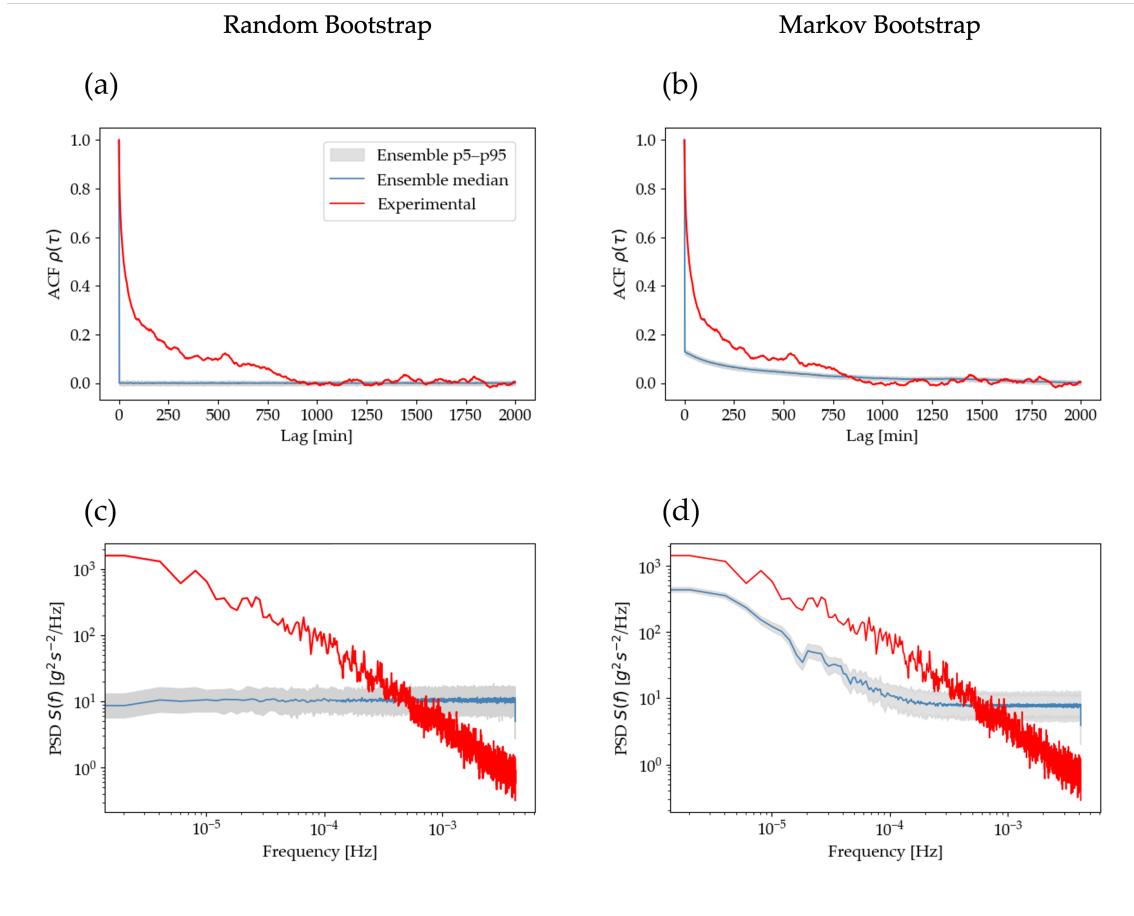


Figure 4.7: Temporal dependence diagnostics for the experimental record (red) and the synthetic ensembles (blue median, gray band 5 to 95%). Top row. Lag autocorrelation functions. Bottom row. Power spectral densities. Left panels use a basic bootstrap that resamples Q_s from the full series. Right panels use the state conditioned Markov bootstrap that draws Q_s from the bag of the visited cluster.

The PSDs in panels (c) and (d) tell the same story in the frequency domain. The basic bootstrap produces a nearly flat spectrum across the complete frequency range, a signature of white noise. This places too much energy at fast time scales and far too little at slow time scales, which explains the near zero ACF. By contrast, the Markov bootstrap generates a spectrum that bends down with frequency and approaches the experimental slope at low f . A gap in low frequency power remains, especially below about 10^{-4} Hz, but it is much smaller than for the basic bootstrap. At the highest frequencies the two spectra are close, showing that the generator reproduces pulse scale variability while improving the representation of multi hour modes.

Together the ACF and PSD comparisons show that state conditioning is essential for realistic temporal structure. A bootstrap that ignores state loses memory almost entirely, whereas the Markov version carries over the residence times and the transition kernel, which yields replicas with the correct decay shape, a comparable correlation horizon, and a spectrum with the right trend. The residual deficit at the very lowest frequencies highlights where additional structure, such as richer dwell time models or weak nonstationarity, could further close the gap.

The accepted ensemble reproduces the experimental mean and variance with small error, matches return levels across the observed range with credible uncertainty, and recovers the ACF decay and spectral trend from minutes to multi hour scales, although it underestimates the very slowest modes. These results show that the state conditioned Markov generator preserves the mechanisms that control typical loads and bursts and provides a calibrated basis for what follows, including variance decomposition driven by state switching, sensitivity to cluster granularity, and the construction of image only surrogates. The small low frequency gap suggests natural extensions through richer dwell time models or mild nonstationarity, but it does not affect the main conclusions supported by the ensemble tests.

4.2 Variance Decomposition and the Role of Cluster Granularity

The validated ensemble of time series from Section 4.1 provides a statistically rigorous setting to investigate the physical role of the morphological states. A central question is how much of the observed variability in sediment transport can be attributed to the system switching between states, and how much reflects intrinsic noise within states. Addressing this quantifies the explanatory power of our Markov model and clarifies how the choice of cluster granularity (K states) influences predictive accuracy and computational cost. We begin by measuring the fraction of variance explained for the baseline configuration ($K = 16$), and then assess sensitivity using coarser ($K = 10$) and finer ($K = 34$ and $K = 45$) partitions.

4.2.1 Fraction of Variance Explained

Separating variability caused by state changes (i.e., movement between morphological states) from fluctuations that occur within a state helps evaluate the informational value of the Markov representation. We use a two-layer description of the process and derive a variance decomposition that isolates these two sources, then evaluate it across the full ensemble of accepted trajectories.

Let Q_t be the sediment transport rate at discrete time t and $S_t \in \{0, \dots, K-1\}$ the morphological state given by the Markov chain. We write the process as

$$Q_t = \mu_{S_t} + \varepsilon_t, \quad \varepsilon_t | S_t = s \sim \mathcal{D}(0, \sigma_s^2), \quad (4.1)$$

where μ_{S_t} is the state-specific mean transport and ε_t represents within-state fluctuations. By the law of total variance,

$$\text{Var}(Q_t) = \underbrace{\mathbb{E}[\text{Var}(Q_t | S_t)]}_W + \underbrace{\text{Var}(\mathbb{E}[Q_t | S_t])}_B, \quad (4.2)$$

with

$$W = \sum_{s=0}^{K-1} \pi_s \sigma_s^2, \quad B = \sum_{s=0}^{K-1} \pi_s (\mu_s - \mu_g)^2, \quad \mu_g = \sum_{s=0}^{K-1} \pi_s \mu_s,$$

where W represents the average variability that remains within states (the internal noise around each state-specific mean) and B captures the variability explained by differences between state means (how much transport levels shift from one state to another). In other words, W measures within-state fluctuations, while B quantifies the contribution of state contrasts to the total variance.

The fraction of variance explained by between-state contrasts is

$$\rho_B = \frac{B}{W + B}, \quad 0 \leq \rho_B \leq 1. \quad (4.3)$$

Values near one indicate that contrasts between states dominate, while values near zero indicate that variability is mainly internal to states.

For the 12436 accepted trajectories we proceed as follows. We first count state visits to obtain π_s , and estimate μ_s and σ_s^2 from the bootstrap samples stratified by state. Using Eqs. (4.2)–(4.3), we then compute W , B , and ρ_B for each trajectory, and store (W, B, ρ_B) to summarize their ensemble distributions.

For $K = 16$ clusters, the ensemble average is

$$\rho_B = 0.151 \pm 0.007 \quad (\text{mean} \pm \text{standard deviation}),$$

showing that about 15% of the instantaneous variance in sediment transport is due to the system switching among states with different means, while the remaining 85% originates from fluctuations inside states. Figure 4.8 shows a narrow distribution. The 5th–95th percentile band extends from about 0.134 to 0.167, while the interquartile range lies between 0.144 and 0.158, with a mean close to 0.151. No replicate falls below $\rho_B = 0.122$, indicating that even in the least favorable cases a substantial portion of variance is still explained by between-state contrasts.

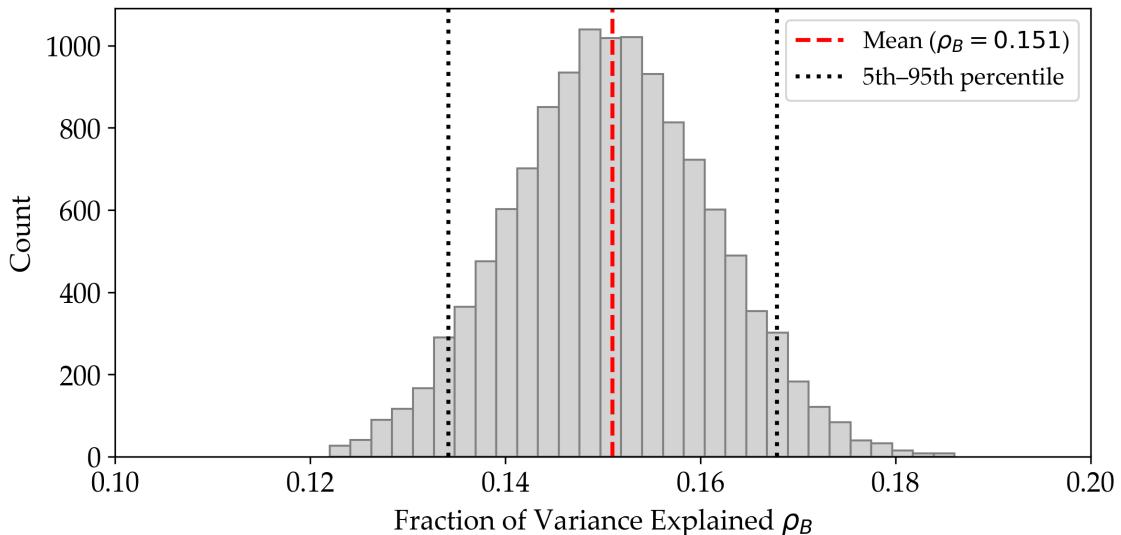


Figure 4.8: Distribution of ρ_B over 12 436 trajectories with $K = 16$ clusters. The dashed line marks the ensemble mean.

Taken together, these results indicate that the Markov partition explains a bounded yet meaningful share of instantaneous variability. For the baseline with $K = 16$ clusters, the between-state component accounts for about 15% of the variance (mean $\rho_B = 0.151$; IQR 0.144–0.158; 5th–95th percentile 0.134–0.167), with the remaining $\sim 85\%$ arising within states. The tight percentile band across trajectories highlights the robustness of this estimate. In practical terms, state switching provides predictive signal about transport levels and bursts, while most short-term variability occurs within states and would require additional noise models or richer dwell-time structure. We adopt ρ_B as a baseline reference and next analyze how it changes when the number of clusters is reduced or increased, balancing explanatory gain against model complexity.

4.2.2 Sensitivity Analysis to Cluster Granularity

The baseline estimate of $\rho_B = 0.151$ with $K = 16$ clusters raises a natural question. Does reducing the number of states remove important contrasts, and does increasing it add enough information to justify the added complexity? To address this, we rerun the full Monte Carlo, bootstrap, and Bayesian screening pipeline with the same tolerances and hyper-parameters for three alternative K : a coarse model with $K = 10$ clusters, a finer model with $K = 34$, and a very fine model with $K = 45$.

Figure 4.9 compares the resulting distributions of ρ_B , while Table 4.1 and Figure 4.10 track the ensemble means.

Table 4.1: Ensemble means and 95% confidence intervals of ρ_B for different levels of cluster granularity.

Number of Clusters (K)	10	16	34	45
Mean ρ_B	0.127	0.151	0.197	0.210
2.5th Percentile	0.120	0.144	0.189	0.202
97.5th Percentile	0.135	0.158	0.205	0.218

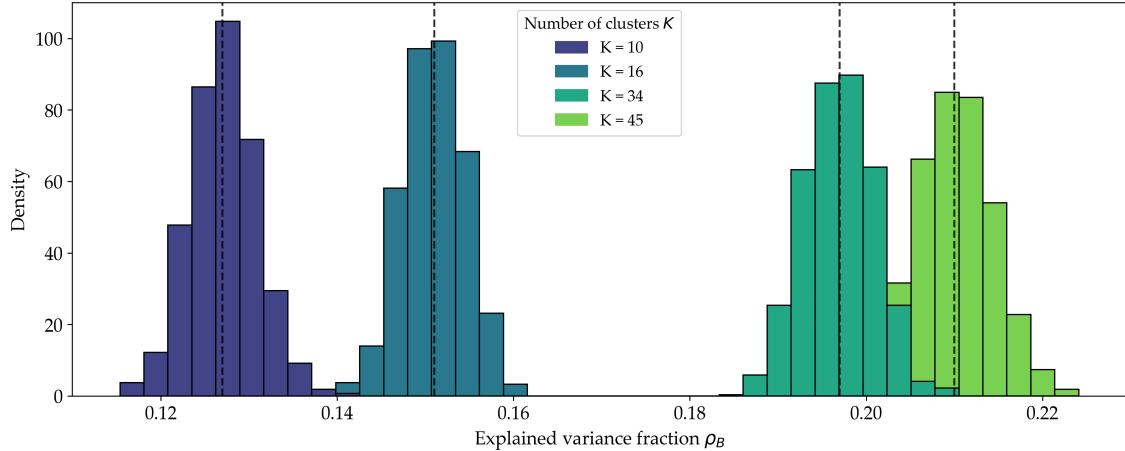


Figure 4.9: Distributions of the variance-explained fraction ρ_B for $K = 10, 16, 34, 45$. Dashed horizontal lines mark ensemble means: 0.127, 0.151, 0.197 and 0.210, respectively.

Key observations.

- ◊ **Coarse model ($K = 10$)**. The mean fraction of explained variance is 0.127 ± 0.0070 , about 16% below the baseline, so nearly 87% of the variability remains within clusters.
- ◊ **Baseline ($K = 16$)**. Provides a stable balance at 0.151 ± 0.0073 .
- ◊ **Fine model ($K = 34$)**. The mean increases to 0.197 ± 0.0080 , a gain of about 4–5 percentage points over the baseline, but this requires more than doubling the state space and reducing data per state, which increases estimation noise for μ_s and σ_s^2 .
- ◊ **Very fine model ($K = 45$)**. The mean rises further to 0.210 ± 0.0084 . The improvement beyond $K = 34$ is modest (about 2 percentage points), while data per state shrink further and parameter estimates become noisier.

The acceptance rate of simulated trajectories mirrors this trend (Fig. 4.11): $K = 10$ accepts 0.28 % of candidates, $K = 34$ only 0.02 %, and $K = 45$ falls below 0.01 %. Another point to note is that each additional percentage point of ρ_B beyond the baseline requires roughly an order of magnitude more CPU time.

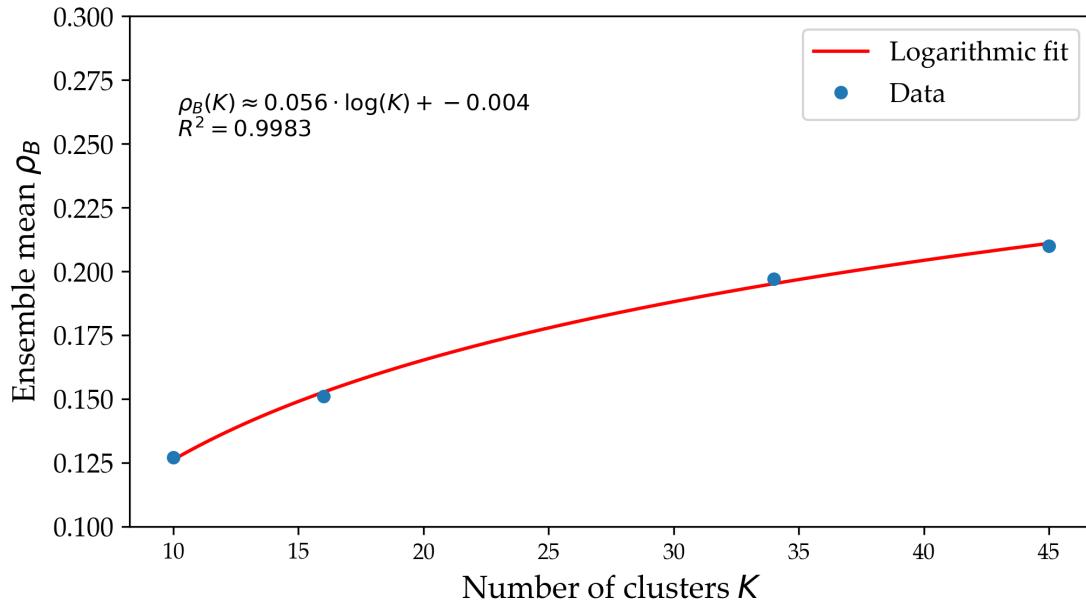


Figure 4.10: Ensemble mean of the explained-variance fraction ρ_B for $K = 10, 16, 34$ and 45 . The best fit (logarithmic, $R^2 = 0.9763$) suggests diminishing returns: increasing the number of clusters yields progressively smaller gains in explained variance.

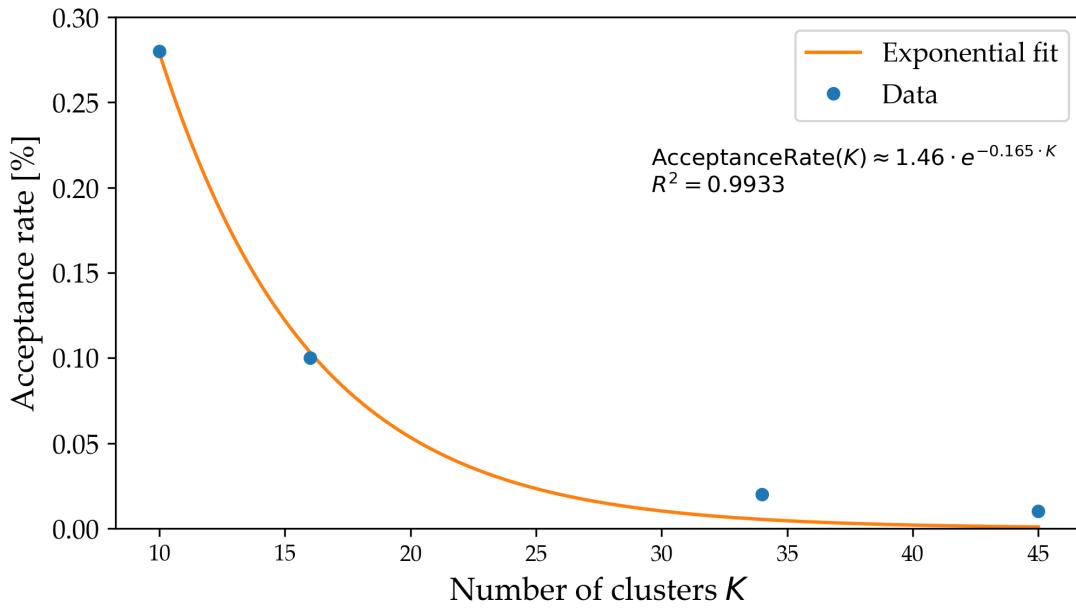


Figure 4.11: Acceptance ratio (accepted / attempted simulations) as a function of the number of clusters K . The observed decay closely follows an exponential trend, illustrating how model complexity drives a rapid drop in acceptance and emphasizing the computational trade-offs of high-dimensional sampling.

Taken together, the results indicate a non-linear relationship between cluster granularity and explanatory power (Fig. 4.10). Beyond $K = 16$, model complexity increases more rapidly than the gain in explanatory value. The configuration with sixteen clusters therefore offers a practical balance, capturing about 15% of the total variance while keeping computational cost and statistical stability within reasonable bounds.

The next section (Section 4.3) builds on this baseline configuration and examines how the Markov framework can be combined with a simple deterministic component to generate mean-preserving transport surrogates when only imagery is available.

4.3 Hybrid Semi-Deterministic Model

In many applications the key quantity of interest is the mean sediment transport rate, since it underpins long-term design and planning. Our experimental setting is unusually rich because it provides a full Q_s time series synchronized with imagery, which made it possible to build a comprehensive Markov+Bootstrap model that captures both mean and variance. Outside the laboratory, however, continuous transport measurements are rare. In practice, morphology snapshots (such as satellite images) and a global discharge estimate are often the only available information.

We first examine a mean-only hybrid in which the accepted Markov state sequences are retained, but the Bootstrap resampling is replaced by assigning, at each time step, the cluster mean transport for the active state. Figure 4.12 compares this approach (panel b) with the original Markov+Bootstrap ensemble (panel a).

Both histograms are centered near the experimental mean μ , which confirms that the Markov sequences correctly encode state occupancies. The difference lies in the spread. The Bootstrap model preserves within-state scatter, while the mean-only hybrid collapses that variability by construction. This makes the mean-only hybrid suitable when reproducing the long-term mean is the primary goal and only imagery is available. It favors parsimony over realism in short-term fluctuations.

We then turn to a more constrained scenario in which only water-distribution patterns (from imagery) and bulk flow are known. Following Bertoldi et al. (2009a), we relate transport to dimensionless stream power through

$$q_b^* = a (\omega_b^*)^b, \quad (4.4)$$

with dimensionless quantities

$$\omega_b^* = \frac{Q S}{b \sqrt{g \Delta d_{50}^3}}, \quad q_b^* = \frac{Q_s}{b \sqrt{g \Delta d_{50}^3}}, \quad (4.5)$$

where Q is discharge, S reach-averaged slope, b wetted width (image-derived), g gravity, $\Delta = (\rho_s - \rho_w)/\rho_w$ the submerged specific gravity, d_{50} median grain size, and Q_s the bedload mass flux (converted to volumetric using $\rho_s = 2650 \text{ kg m}^{-3}$).

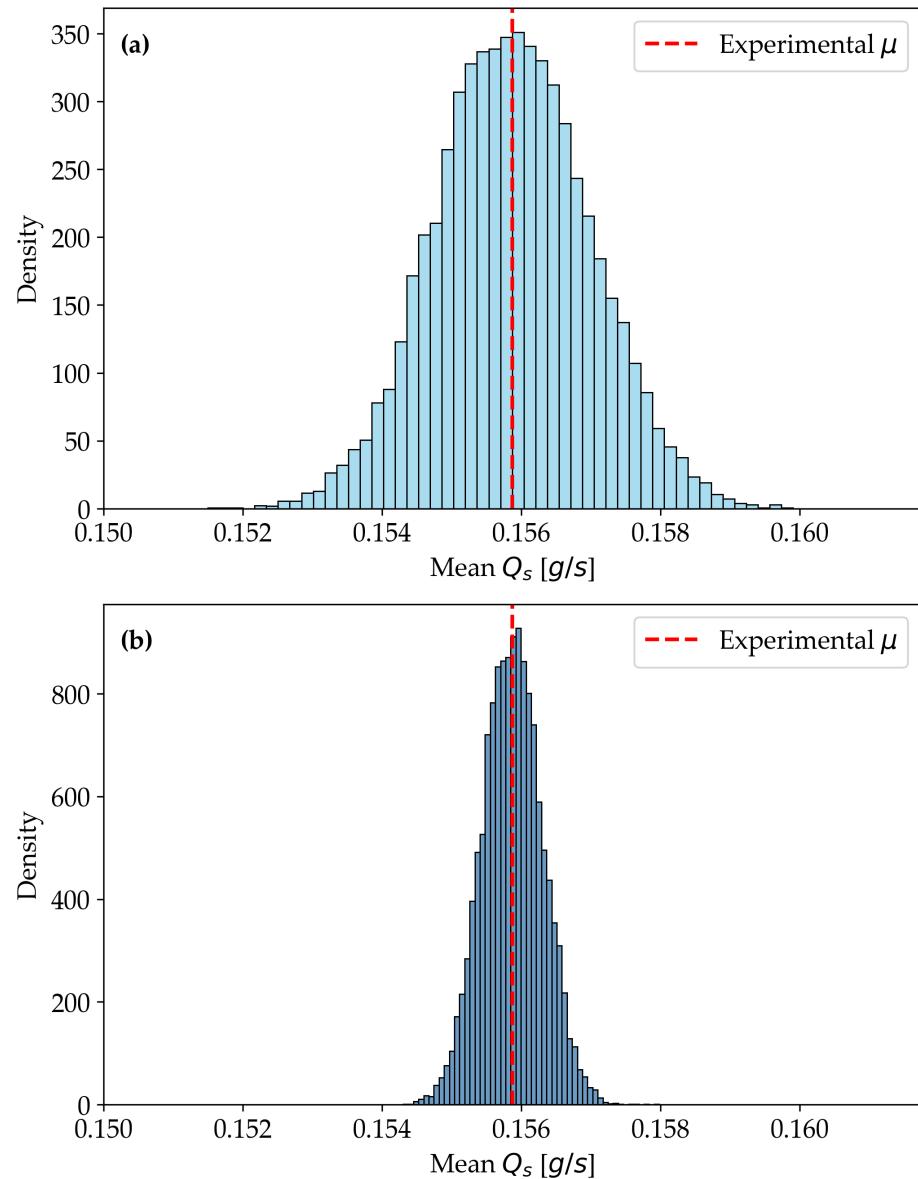


Figure 4.12: Comparison of predicted mean sediment transport rates under two approaches. (a) Mean-only hybrid: Markov state sequences with fixed clusterwise means (no bootstrap). (b) Full Markov+Bootstrap model. Both are centered near the experimental mean (dashed red line), but the bootstrap approach exhibits larger variance because it retains within-state variability.

Treating each morphological state as one observation, we compute statewise pairs (ω_b^*, q_b^*) from their mean attributes and fit (4.4):

$$q_b^* = 2.54 (\omega_b^*)^{2.72}. \quad (4.6)$$

As shown in Figure 4.13, the fit reproduces the transport–power trend but with a steeper exponent than the $n \approx 2.3$ reported by Bertoldi et al. (2009a), and with a downward shift that reflects the lower transport rates observed in the laboratory. This leads to a Bertoldi-scaled hybrid in which each state is assigned a predicted mean Q_s from (4.6) based on the image-derived ω_b^* , and the Markov state sequences are then used to generate synthetic means.

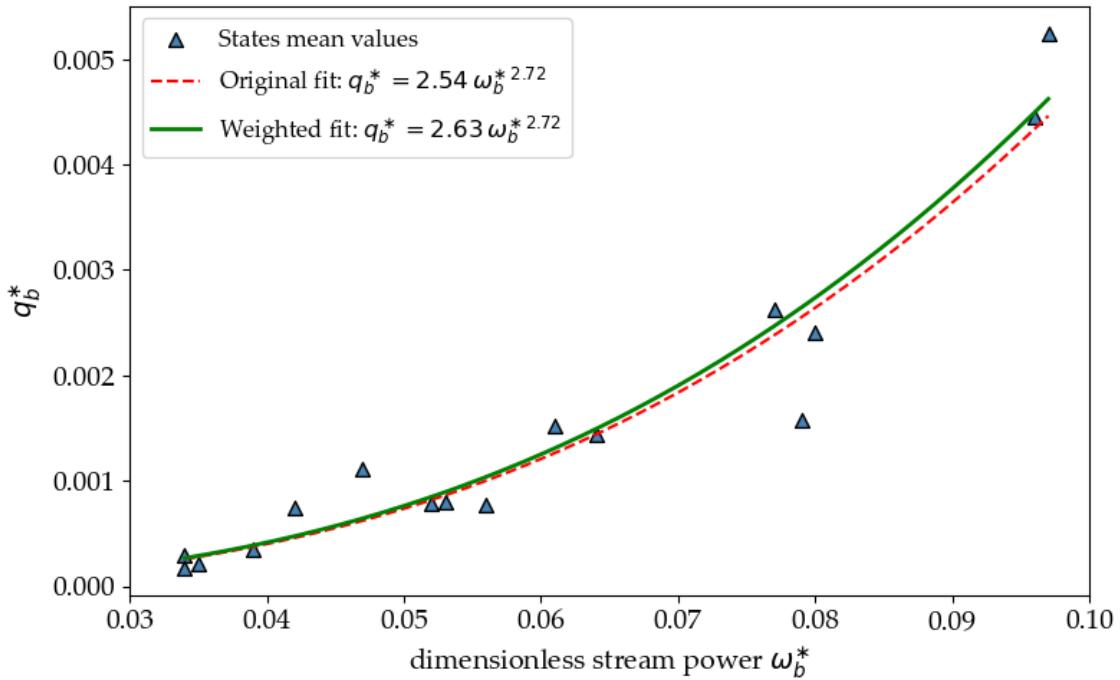


Figure 4.13: Clusterwise mean values of dimensionless stream power ω_b^* and transport q_b^* (triangles), with least-squares fit (red dashed). The fitted law $q_b^* = 2.54 (\omega_b^*)^{2.72}$ follows the form in Bertoldi et al. (2009a) but with a steeper exponent and overall lower transport—consistent with flume conditions.

If one performs a calibration of the Bertoldi-type law in the standard way, all clusters are given equal importance in the regression, so that the scale factor a is estimated as

$$a = \frac{\sum_s (\omega_{b,s}^*)^b q_{b,s}^{*,\text{exp}}}{\sum_s (\omega_{b,s}^*)^{2b}} = 2.54.$$

This approach implicitly balances frequent low-transport states against rare but transport-intense ones, even though their contribution to the long-term mean is very different. Because low- ω_b^* states are more numerous, they dominate the regression and bias the predicted mean downward. This effect is apparent in the histogram of synthetic means obtained from this calibration (Figure 4.14a), which underestimates the experimental reference. To correct for this, we modified the regression so that each state contributes in proportion to its temporal relevance, measured by its residence fraction π_s and mean

dwell time τ_s . The weighted estimator of the scale factor is then

$$\hat{a}_w = \frac{\sum_s w_s (\omega_{b,s}^*)^b q_{b,s}^{*,\text{exp}}}{\sum_s w_s (\omega_{b,s}^*)^{2b}} = 2.63, \quad w_s \propto \pi_s \tau_s,$$

with the weights normalized to sum to one. Predicted values are $q_{b,s}^{*,\text{pred}} = \hat{a}_w (\omega_{b,s}^*)^b$, and their weighted mean $\mu_w^{\text{pred}} = \sum_s w_s q_{b,s}^{*,\text{pred}}$ closely reproduces the experimental mean (Figure 4.14b). From a physical perspective, this modification acknowledges that sediment flux is governed not only by the intensity of transport in each morphology, but also by the duration of time the system spends in it. Emphasizing states that are both frequent and persistent restores the balance between common low-transport morphologies and rare high-transport ones, producing a mean-preserving surrogate from imagery alone.

This hybrid pathway is promising in data-limited contexts. With imagery (to estimate b and state identity), bulk hydraulics (Q , S), and grain size, it is possible to recover a mean-preserving transport surrogate. Future refinements could estimate state persistence directly from morphology or forcing, so that residence weights would not need to rely on a prior instrumented record. If such persistence predictions become feasible, the hybrid approach could provide mean-preserving, imagery-based estimates of Q_s even in the absence of bedload measurements, while the full Markov+Bootstrap method would remain the preferred option when short-term variability is important.

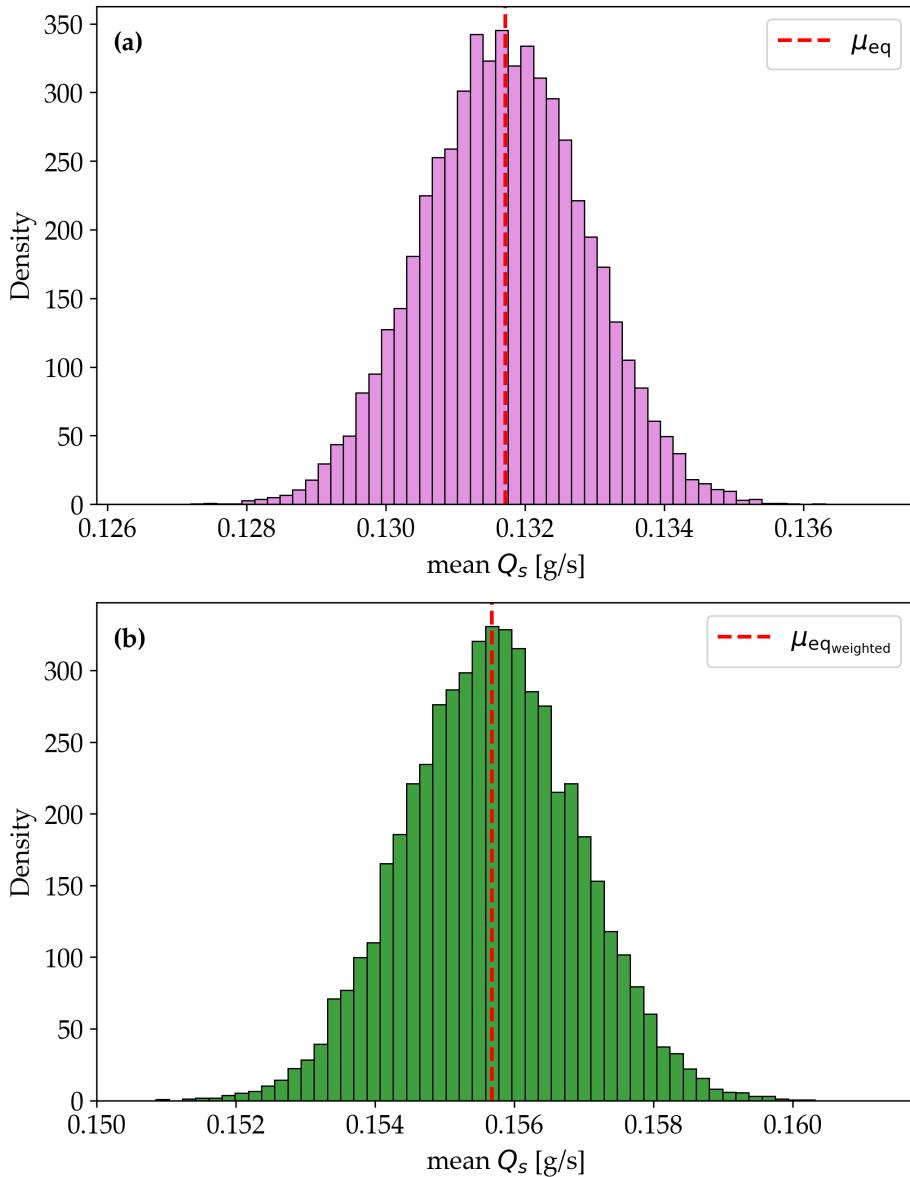


Figure 4.14: (a) Mean transport rates from the Bertoldi-scaled hybrid when all states are given equal weight. The predicted distribution underestimates the experimental mean (red dashed line). (b) Results when each state's predicted mean is weighted by its empirical residence fraction. The agreement improves markedly, highlighting the importance of accounting for state persistence.

4.4 Summary of Key Points

Monte Carlo + bootstrap surrogate recovers realistic transport dynamics Starting from the Chapter 3 transition matrix P , we simulate state sequences with exponential dwell times using one trajectory level rate λ and map them to transport with state conditioned bootstrap samples. The synthetic series reproduce the observed range, intermittency, and burst structure while preserving both inter-state contrasts and intra-state scatter.

Stringent Bayesian screening and an interpretable image only filter From 100000 simulated state paths, 12436 pass absolute tolerances on mean and variance and match laboratory moments inside a $\pm 2\%$ corridor (mean bias $\approx 1.3\% \pm 0.8$, variance error $\approx 0.9\% \pm 1.1$). A depth 3 decision tree based only on image derived metrics yields a conservative acceptance rule (`freq_diff` ≤ 0.016 , `dur_JS` ≤ 0.88 , `frobenius_diff` ≤ 0.956) with precision ≈ 0.99 and recall ≈ 0.29 , providing a fast prefilter when Q_s is unavailable.

Beyond moments the ensemble reproduces extremes and temporal structure The accepted replicas follow the experimental return level curve across the full record range (about 6 to 1200 h) and provide credible bands for exceedance based diagnostics. In the time domain and frequency domain the state conditioned generator recovers the ACF decay shape and the PSD trend from minutes to multi hour scales and clearly improves over a basic, state agnostic bootstrap. A small shortfall in very low frequency variance remains, pointing to gains from richer dwell time models or weak nonstationarity.

Variance decomposition quantifies the information carried by state switching For the baseline partition $K = 16$, the fraction of instantaneous variance explained by between state contrasts is $\rho_B = 0.151 \pm 0.007$ (IQR 0.144 to 0.158, minimum 0.122). State changes therefore account for a measurable yet bounded share of variability, with the remainder arising within states.

Granularity improves ρ_B with diminishing returns and higher cost Coarsening to $K = 10$ lowers ρ_B to 0.127, while refining to $K = 34$ and $K = 45$ raises it to 0.197 and 0.210. Acceptance rates decline steeply as K grows and the mean ρ_B follows a sublinear or logarithmic trend, making $K = 16$ a practical balance between interpretability, acceptance, and computational cost.

Practical hybrids for imagery only scenarios A mean only variant that keeps accepted Markov state sequences but replaces bootstrap draws with clusterwise means preserves long term transport means while sacrificing short term variance. A Bertoldi scaled hybrid links statewise transport to dimensionless stream power,

$$q_b^* = 2.54 (\omega_b^*)^{2.72},$$

steeper than the ≈ 2.3 reported by Bertoldi et al., enabling mean preserving surrogates from imagery and bulk flow alone.

Chapter 4 supports a view of braided rivers as multi state systems. Discrete morphological state sequences combined with state conditioned transport statistics generate statistically faithful surrogates, explain a quantifiable share of variance, and enable deployable mean preserving models when only images are available. The remaining low frequency mismatch is targeted and motivates semi Markov extensions focused on state specific dwell times, while the screening and decision tree provide an operational path for fast, image driven preselection and scenario testing.

5 | Conclusions and Future Work

5.1 Conclusions

This thesis builds a practical bridge from plan view imagery to probabilistic forecasts of braided river morphology and sediment transport. The central idea is to represent planform evolution as a sequence of discrete states and to model their alternation with a continuous time Markov chain. A lean Monte Carlo and bootstrap layer then converts state dynamics into transport surrogates that pass objective resemblance tests. The result is a workflow that sits between descriptive metrics and physics heavy models and shows that planform carries actionable predictive signal once cast as a multi state process.

Chapter 3 established a robust image to state mapping. Complementary distances based on MHD and Dice, followed by dimensionality reduction and density clustering, produced an interpretable morpho space with a finite catalog of recurrent states that spans the continuum from single thread to bar rich configurations. Medoids and morphometric profiles clarified that mean transport is highest in simple confined states and lowest in wide partitioned states. Across clusters the mean of Q_s is strongly and negatively associated with wetted width and with extended braiding index. Dwell times are well approximated by exponentials, which yields state specific exit rates $\hat{\lambda}$ and a transition matrix \hat{P} that together define a CTMC consistent with the observed persistence and activity.

Chapter 4 turned that structure into a working surrogate. We simulated 100,000 state paths under \hat{P} using one trajectory level rate λ and screened each path using moments implied by state means and variances. About 0.12% of the paths passed tight absolute tolerances on mean and variance, which produced 12,436 accepted sequences. State conditioned bootstrap on this set generated transport series that reproduce range, intermittency, and burst structure while preserving inter state contrasts and intra state scatter. An image only decision tree provided a conservative prefilter with simple thresholds on state frequency mismatch, duration distribution, and Frobenius distance between transition structures. Operational cutoffs around $\text{freq_diff} \leq 0.016$, $\text{dur_JS} \leq 0.88$, and $\text{frobenius_diff} \leq 0.956$ achieved high precision and moderate recall, which is an acceptable trade off for safe preselection when Q_s is unavailable.

The accepted ensemble passed tests beyond the first two moments. It tracked the experimental return level curve across the full record window from about 6 to 1200 h and provided credible bands for exceedance based diagnostics used in risk assessment. In the time and frequency domains the state conditioned generator recovered the ACF decay shape and the PSD trend from minutes to multi hour scales and clearly improved over a basic state agnostic bootstrap. A small shortfall remained at very low frequencies, which points to where additional structure could further reduce the gap without altering the core workflow.

Two results clarify what information the state sequence carries. First, transport is measurably state conditioned. Pulses concentrate in morphologically simple phases and lulls follow transitions into storage prone braids, in line with the negative relations between mean Q_s and both wetted width and extended braiding index. Second, variance decomposition showed that state switching explains a bounded yet meaningful share of instantaneous variability. For $K = 16$ the between state fraction was $\rho_B \approx 0.15$ with narrow dispersion. Finer granularity raised ρ_B with diminishing returns and steep computational costs because acceptance rates fell sharply.

The framework is ready for use in imagery only settings. A mean only hybrid that replaces bootstrap draws with cluster means preserves long term averages. A Bertoldi scaled option links statewise transport to dimensionless stream power and yields mean preserving surrogates from imagery and bulk hydraulics. Together with the full Markov and bootstrap pipeline these tools allow scenario testing and operations with transparent assumptions and traceable uncertainty.

In sum braided rivers can be usefully treated as multi state systems. A calibrated state space, a CTMC generator, and a lean Monte Carlo and bootstrap layer produce transport surrogates that are statistically faithful, transferable across realizations, and suitable for analysis and decision support at experiment time scales. The remaining low frequency mismatch is limited and well identified and it does not detract from the core result that imagery driven state dynamics already encode the levels, variability, intermittency, and extremes seen in the laboratory record.

5.2 Future work

We can think of many possibilities for future work. Below we outline some of them.

1. **Estimating the dwell rate λ from imagery only** Link the trajectory level rate to image derived covariates so that λ can be inferred without transport data. Start with parsimonious regressions that map wetted width, eBI, junction density, curvature proxies, active width to $\log \lambda$. Use generalized additive models, monotone splines, or hierarchical pooling across runs to stabilize estimates when some states are rare. A survival framing can model the hazard of leaving a state as $h_i(t \mid \mathbf{x}) = \lambda_i(\mathbf{x})$ with \mathbf{x} computed from masks only. Compare Cox style log linear effects, tree based survival, and shallow contrastive embeddings that predict λ from low dimensional image representations. Validate with cross run transfer, posterior predictive checks for dwell histograms, and sensitivity to image resolution.

2. **Forecasting when a configuration is out of distribution** Develop novelty detection and open set forecasting for planforms that do not match any learned state, for example after an engineered intervention. Use distance to centroids in the embedding, predictive entropy of the state classifier, and one class methods to flag novelty. Define an explicit unknown macro state with its own dwell prior and a transition stencil learned from the nearest valid states. For forecasting, propagate uncertainty with a mixture over the three closest macro states or with a physics guided bridge that nudges toward feasible neighbors while honoring engineered constraints. Allow online adaptation that spins up a new state if novelty persists and reaches support above a minimum occupancy.
3. **Synthetic morphology generators to feed the model** Build generators of plan view morphology and state sequences when observations are sparse. Cellular automata, neural cellular automata, and lattice gas models can reproduce channel splitting and bar dynamics under simple rules. Multiple point statistics and pattern based geostatistics can sample realistic planforms from training images. Diffusion models, adversarial generators, and autoregressive video models can be adapted with geomorphic constraints and conservation penalties so that generated masks respect area, width, connectivity, and mass balance. Use these generators to stress test the CTMC and to augment rare states.
4. **Three state coarse model and direct three class mapping from images** Explore a coarse representation with three macro states that reflect the clear grouping observed in the data. A hierarchical route merges the 16 states into 3 via community detection on the transition graph or by thresholding wetted width and eBI, then trains a supervised classifier to map images directly to the classes using engineered features or a shallow CNN. Impose temporal consistency with a three state hidden Markov model or a duration prior. Fit a three state CTMC and compare dynamic fidelity, extremes, and mean preservation against the 16 state baseline, including acceptance rates and computational cost.
5. **Effect of lateral confinement on state complexity** Vary corridor width and bank fixing to map how the number of states, their occupancies, and the transition rates respond to confinement. Build a phase diagram with axes defined by dimensionless stream power and a confinement index. Quantify thresholds where the system collapses to a small set of simple states and where it expands into bar rich repertoires. Use this diagram to anticipate how engineering works may alter the reachable state space.
6. **What triggers a transition between states** Attribute transitions to measurable precursors. Combine change point detection with covariates such as discharge, slope adjusted power, curvature, junction density, and bar migration. Train transition classifiers and use permutation importance or SHAP style attributions to identify proximal causes. Test for asymmetric gateways and for lagged signatures that foreshadow switches.

7. **Explaining the remaining variance** Decompose within state variability with hierarchical models. Represent transport inside each state with a stochastic component that carries short memory fluctuations and micro adjustments at scales below state resolution. Compare Markov plus within state noise against semi Markov alternatives to capture longer dwell variability and part of the missing low frequency energy. Report gains in ACF and PSD at long lags and in coverage of extreme bursts.
8. **Vegetation as a state modifier** Add vegetation cover and type as covariates that alter dwell times and jump probabilities. Build vegetation aware states or state annotations and test whether rooted patches lengthen persistence, redirect transitions toward storage prone configurations, and reduce burst frequency. Evaluate robustness across growth seasons and disturbance cycles.
9. **Role of floods** Quantify how floods move the system in state space. Condition kernels on hydrographs and track which states are visited during rising and falling limbs. Test whether the generator is time invariant or whether flood events cause a regime shift. Measure hysteresis in return paths and the residence penalties associated with large disturbances.
10. **Semi Markov and non exponential sojourns** Replace exponential dwell times with flexible families such as phase type, lognormal, or spline based hazards. Fit hidden semi Markov models where latent sub states explain long residence tails and conditioning on the next state is allowed. Implement duration aware simulation and compare ACF integrals and low frequency spectral energy against the exponential baseline.
11. **Markov chains with exogenous inputs** Let the generator depend on covariates like discharge, temperature, or sediment supply. A continuous time model with $\log q_{ij}(t) = \alpha_{ij} + \beta_{ij}^\top \mathbf{x}(t)$ can capture non stationarity directly. Use group sparsity to select which transitions respond to which inputs. Compare covariate driven kernels against piecewise stationary baselines and report gains in predictive skill and calibration.
12. **State definitions via representation learning** Learn embeddings from images with self supervised contrastive learning to improve cluster stability and transfer to new rivers. Compare learned states against expert defined planform classes and quantify which representation yields better transport surrogates. Use domain adaptation to port the embedding across sites with different scales and grain sizes.
13. **Graph based states and connectivity** Define states using metrics of the channel network graph such as node degree, path redundancy, and loop density. Test whether graph derived states better explain transport intermittency than purely areal or contour based clustering. Combine graph features with areal descriptors in a multi view clustering that preserves both topology and coverage.
14. **Coupling to physics based models** Use states to guide or emulate hydro morphodynamic solvers. Constrain bedload closures with state dependent parameters or build surrogates that map model outputs into state sequences. Inverse uses include calibrating friction or transport exponents by matching CTMC statistics. Design

twin experiments where physics based runs are nudged toward the observed state occupancies and transition corridors.

15. **Optimal sampling design** Determine the minimum image cadence and duration needed to recover the generator within a tolerance. Use information measures and synthetic experiments to set monitoring guidelines for cameras and satellites. Quantify identifiability of P and dwell rates as functions of cadence, duration, and signal to noise in segmentation.
16. **Experimental interventions** Use the state framework to assess effects of controlled changes such as local bank reinforcement, graded sediment pulses, or wood additions. Track how interventions shift dwell times and transition corridors and whether they create new states or push the system toward the boundary of the learned morpho space.

Appendix A – Image Preparation Code

The following Python script was developed to preprocess experimental images. It performs rotation, masking, cropping, and HSV-based filtering for further analysis.

Python Code

```
1 [Image Preparation Script]
2  \#
3  =====
4  \# LESO - Image preparation for the MPS process
5  \#
6  -----
7  \# Part 1: Cropping Images
8  \#
9  -----
10 \# Images are cropped using a mask to isolate the channel and
11 \# remove the background. The image is rotated to align the
12 \# channel, cropped to keep it centered, and masked.
13 \#
14 -----
15 import os
16 from PIL import Image, ImageChops
17 \# Define input/output folders and mask path
18 folder_path = r"C:\\FakePath\\Experiment\\RawImages"
19 output_folder_path = r"C:\\FakePath\\Experiment\\Cropped"
20 mask_path = r"C:\\FakePath\\Experiment\\Mask"
21 \#
22 \# Create output folder if it does not exist
23 if not os.path.exists(output_folder_path):
24     os.makedirs(output_folder_path)
25 \#
26 \# Load mask
```

```
27     mask_filename = "mask.png"
28     mask_path = os.path.join(mask_path, mask_filename)
29     mask = Image.open(mask_path).convert("L")
30     mask = mask.rotate(3.6, expand=True, fillcolor="black")
31
32     \# List all JPG files
33     jpg_files = [f for f in os.listdir(folder_path) if f.endswith
34                   (".jpg")]
35
36     \# Process each image
37     for jpg_file in jpg_files:
38         input_image_path = os.path.join(folder_path, jpg_file)
39         input_image = Image.open(input_image_path)
40
41         \# Rotate images
42         input_image = input_image.rotate(-90, expand=True, fillcolor=
43             "black")
44         rotated_image = input_image.rotate(3.6, expand=True,
45             fillcolor="black")
46
47         \# Apply mask
48         result_image = ImageChops.multiply(
49             rotated_image.convert("RGB"), mask.convert("RGB"))
50
51         \# Crop bounding box
52         bbox = result_image.getbbox()
53         result_image = result_image.crop(bbox)
54
55         \# Save
56         output_image_path = os.path.join(output_folder_path, "cropped_" + jpg_file)
57         result_image.save(output_image_path)
58
59         \# Close resources
60         input_image.close()
61         rotated_image.close()
62         result_image.close()
63
64         print("Image cropping completed.")
65
66         \#
67         \# Part 2: Binary Image Creation - Color Thresholding
68         \#
69         \# HSV filtering using blue ink as an indicator of the water
70         \# surface. MATLAB provides threshold values with its
71         \# Image Thresholding Toolbox.
72         \#
```

```
71
72     import cv2
73     import numpy as np
74
75     # Convert MATLAB-style normalized value into OpenCV scale
76     def convert_to_opencv_range(matlab_value, opencv_max):
77         return int(matlab_value * opencv_max)
78
79     # Process image to mask
80     def process_image(image_path):
81         img = cv2.imread(image_path)
82         hsv_img = cv2.cvtColor(img, cv2.COLOR_BGR2HSV)
83         mask = cv2.inRange(hsv_img, lower_hsv, upper_hsv)
84         return mask
85
86     # Example thresholds (values tuned experimentally)
87     channel1Min, channel1Max = 0.401, 0.662
88     channel2Min, channel2Max = 0.180, 1.000
89     channel3Min, channel3Max = 0.049, 1.000
90
91     # Convert MATLAB ranges to OpenCV ranges
92     channel1Min_opencv = convert_to_opencv_range(channel1Min,
93                                                 179)
93     channel1Max_opencv = convert_to_opencv_range(channel1Max,
94                                                 179)
94     channel2Min_opencv = convert_to_opencv_range(channel2Min,
95                                                 255)
95     channel2Max_opencv = convert_to_opencv_range(channel2Max,
96                                                 255)
96     channel3Min_opencv = convert_to_opencv_range(channel3Min,
97                                                 255)
97     channel3Max_opencv = convert_to_opencv_range(channel3Max,
98                                                 255)
98
99     # HSV filter bounds
100    lower_hsv = np.array([channel1Min_opencv, channel2Min_opencv,
101                          channel3Min_opencv])
101    upper_hsv = np.array([channel1Max_opencv, channel2Max_opencv,
102                          channel3Max_opencv])
102
103    # Input/output
104    folder_path = r"C:\\FakePath\\Experiment\\Cropped"
105    output_folder = r"C:\\FakePath\\Experiment\\Masked"
106    os.makedirs(output_folder, exist_ok=True)
107
108    # Process folder
109    for filename in os.listdir(folder_path):
110        if filename.endswith(".jpg"):
111            input_path = os.path.join(folder_path, filename)
112            output_path = os.path.join(output_folder, "masked_" +
filename)
```

```
113     mask = process_image(input_path)
114     cv2.imwrite(output_path, mask)
115
116     print("Masking completed.")
117
118     #
119     # Part 3: Resize, remove blobs, invert, and edge detection
120     #
121
122     def resize_remove_small_blobs_invert_and_edge_image(
123         input_path, output_path, target_size, size_threshold):
124         # Read grayscale
125         img = cv2.imread(input_path, cv2.IMREAD_GRAYSCALE)
126         height, width = img.shape[:2]
127         aspect_ratio = width / height
128
129         # Maintain aspect ratio
130         if aspect_ratio > 1:
131             new_width = target_size
132             new_height = int(target_size / aspect_ratio)
133         else:
134             new_width = int(target_size * aspect_ratio)
135             new_height = target_size
136
137         resized_img = cv2.resize(img, (new_width, new_height),
138             interpolation=cv2.INTER_NEAREST)
139
140         # Connected components
141         num_labels, labels_im, stats, _ = cv2.
142             connectedComponentsWithStats(resized_img)
143         output_img = np.zeros_like(resized_img)
144
145         # Keep blobs above threshold
146         for i in range(1, num_labels):
147             if stats[i, cv2.CC_STAT_AREA] >= size_threshold:
148                 output_img[labels_im == i] = 255
149
150         # Invert image
151         inverted_img = 255 - output_img
152
153         # Save result
154         cv2.imwrite(output_path, inverted_img)
155
156         # Process folder
157         input_folder = r"C:\\FakePath\\Experiment\\Masked"
158         output_folder = r"C:\\FakePath\\Experiment\\Processed"
159         os.makedirs(output_folder, exist_ok=True)
160
161         target_size = 200
```

```
159     filterSize = 30
160
161     for filename in os.listdir(input_folder):
162         if filename.endswith(".jpg"):
163             input_path = os.path.join(input_folder, filename)
164             output_path = os.path.join(output_folder, "resized_" +
165                 filename[:-4] + ".png")
166             resize_remove_small_blobs_invert_and_edge_image(input_path,
167                 output_path, target_size, filterSize)
168
169     print("Resizing and blob filtering completed.")
```


Appendix B – Parallelized MHD Code

This appendix presents the Python implementation used to compute the Modified Hausdorff Distance (MHD) between sets of points extracted from binary images. The code is designed to handle large datasets by processing the distance matrix in blocks, which are stored separately to reduce memory usage. The implementation leverages GPU acceleration through PyTorch to parallelize the calculations and significantly improve performance.

Python Code

```
1 [Hausdorff Distance Calculation Script]
2     import os
3     import time
4     from PIL import Image
5     import numpy as np
6     import torch
7
8     def load_images_to_points(directory):
9         """
10            Load all PNG/JPG images in a directory and extract points
11            from white pixels.
12            Each image is converted to grayscale, thresholded, and points
13            are extracted.
14        """
15        image_files = [f for f in os.listdir(directory) if f.endswith
16                      ('.png') or f.endswith('.jpg')]
17        sets_of_points = []
18
19        for image_file in image_files:
20            image_path = os.path.join(directory, image_file)
21            image = Image.open(image_path).convert('L') # Convert to
22            grayscale
23            image_np = np.array(image)
24
25            # Get coordinates of white pixels (threshold = 128)
26            points = np.column_stack(np.where(image_np > 128))
```

```

23
24     # Subsample: take only one of every two points
25     points = points[::2]
26
27     sets_of_points.append(points)
28
29     print("Points successfully extracted.")
30     return sets_of_points
31
32
33 def calculate_mhd_batch(tensor1, tensor2):
34     """
35     Compute the Modified Hausdorff Distance (MHD) between two
36     tensors of points.
37     """
38     def tensor_to_points(tensor):
39         return torch.nonzero(tensor.squeeze() > 0.1, as_tuple=False).
40             float()
41
42     A = tensor_to_points(tensor1).cuda()
43     B = tensor_to_points(tensor2).cuda()
44
45     if A.size(1) != B.size(1):
46         raise ValueError("Both point sets must have the same
47                         dimensionality.")
48
49     dist_matrix = torch.cdist(A.unsqueeze(0), B.unsqueeze(0)).
50         squeeze(0)
51     fhd = torch.mean(torch.min(dist_matrix, dim=1)[0])    # forward
52         distance
53     rhd = torch.mean(torch.min(dist_matrix, dim=0)[0])    # reverse
54         distance
55     mhd = torch.max(fhd, rhd)
56
57     return mhd.item()
58
59
60 def compute_and_save_distance_matrix_block(sets_of_points,
61     block_start, block_size, save_dir, device):
62     """
63     Compute a block of the distance matrix between sets of points
64
65     .
66     Saves the block as a NumPy file in the specified directory.
67     """
68     n = len(sets_of_points)
69     end = min(block_start + block_size, n)
70     block_height = end - block_start
71
72     tensors = [torch.tensor(points, dtype=torch.float32).to(
73         device) for points in sets_of_points]
74

```

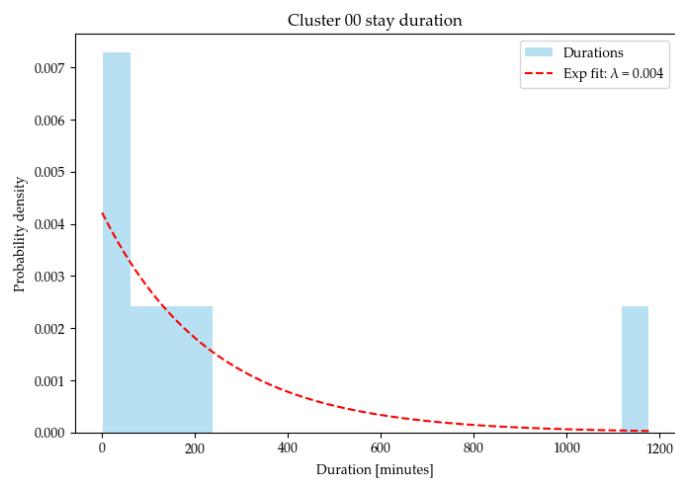
```

65     distance_matrix_block = np.full((block_height, n), np.nan)
66
67     for i in range(block_height):
68         tensor1 = tensors[block_start + i]
69         distances = []
70         for j in range(block_start + i, n):
71             tensor2 = tensors[j]
72             mhd = calculate_mhd_batch(tensor1, tensor2)
73             distances.append(mhd)
74             distance_matrix_block[i, block_start + i:] = distances
75
76     os.makedirs(save_dir, exist_ok=True)
77     block_filename = os.path.join(save_dir, f"distance_block_{block_start}.npy")
78     np.save(block_filename, distance_matrix_block)
79
80     print(f"Block {block_start} saved.")
81
82
83     def compute_distance_matrix_in_blocks(directory, block_size,
84                                             save_dir, device):
85         """
86         Compute the full distance matrix in blocks, saving each block
87             as a NumPy file.
88         """
89
90         sets_of_points = load_images_to_points(directory)
91         n = len(sets_of_points)
92
93         for block_start in range(0, n, block_size):
94             compute_and_save_distance_matrix_block(
95                 sets_of_points, block_start, block_size, save_dir, device
96             )
97             print(f"Block starting at {block_start} processed.")
98
99
100    if __name__ == "__main__":
101        # Example usage with fake paths for LaTeX
102        input_directory = r"C:\\FakePath\\Experiment\\BinaryMasks"
103        output_directory = r"C:\\FakePath\\Experiment\\
104            HausdorffResults"
105        block_size = 10
106        device = torch.device("cuda" if torch.cuda.is_available()
107                            else "cpu")
108
109        start_time = time.time()
110        compute_distance_matrix_in_blocks(input_directory, block_size
111                                         , output_directory, device)
112        end_time = time.time()
113
114        print(f"Total execution time: {end_time - start_time:.2f}
115             seconds")

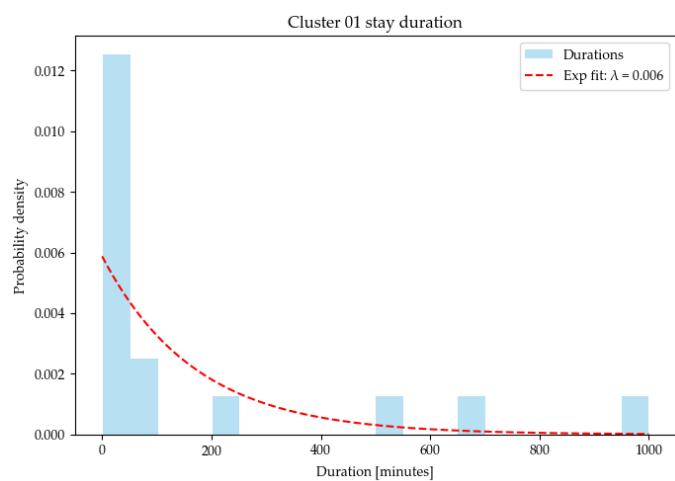
```


Appendix C – Dwell-Time Histograms

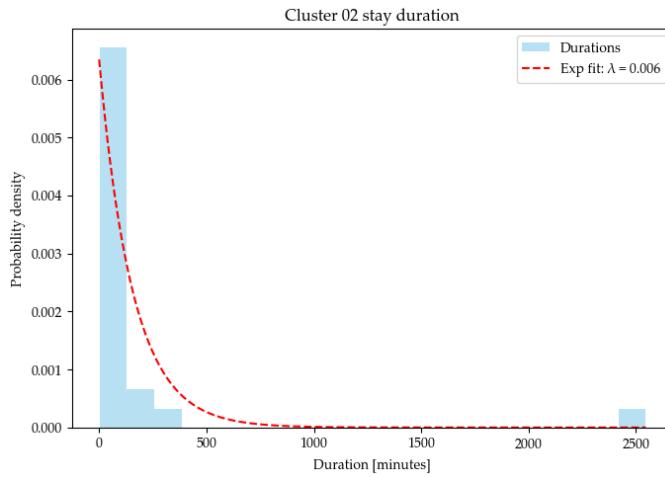
This appendix presents the complete set of dwell-time histograms for all 16 clusters.



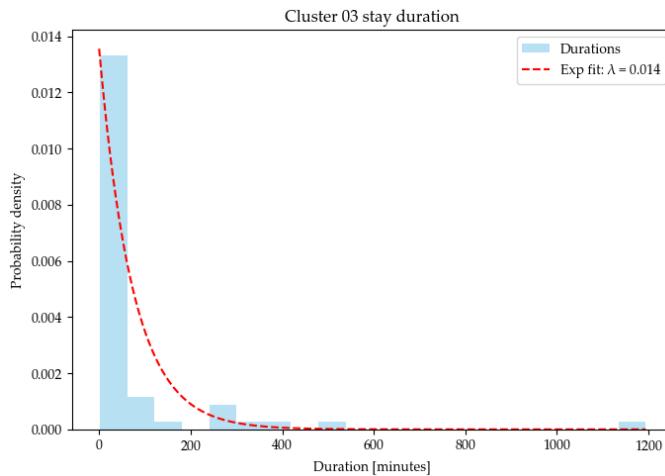
Dwell-time histogram for Cluster 00.



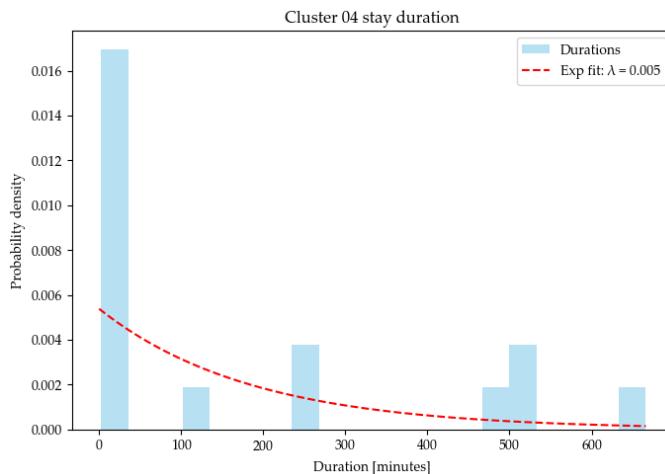
Dwell-time histogram for Cluster 01.



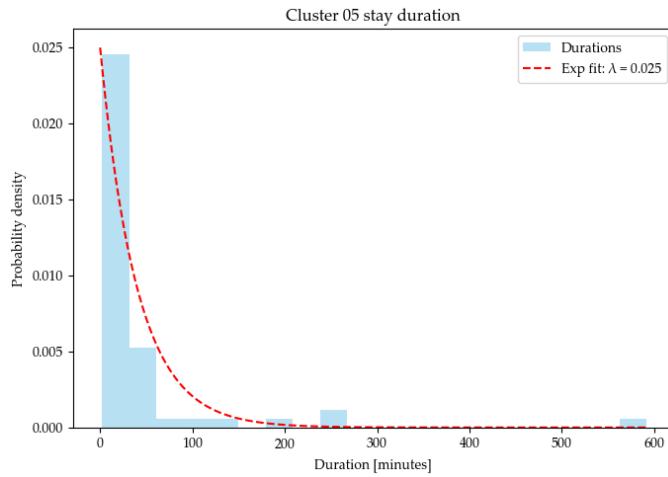
Dwell-time histogram for Cluster 02.



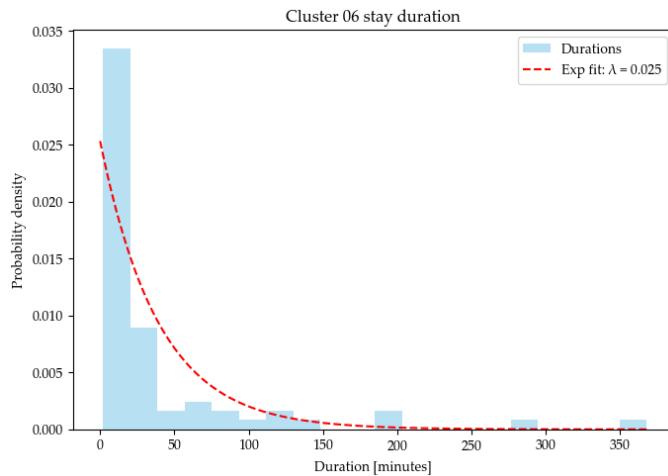
Dwell-time histogram for Cluster 03.



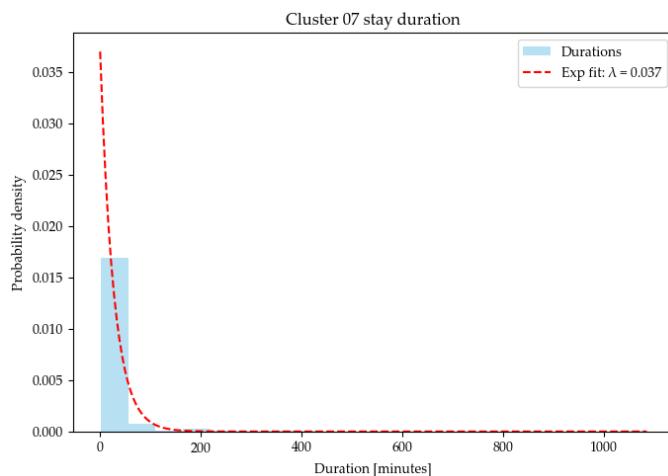
Dwell-time histogram for Cluster 04.



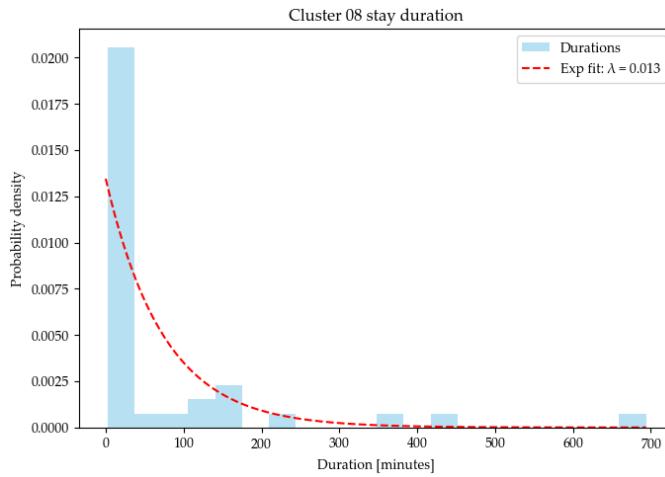
Dwell-time histogram for Cluster 05.



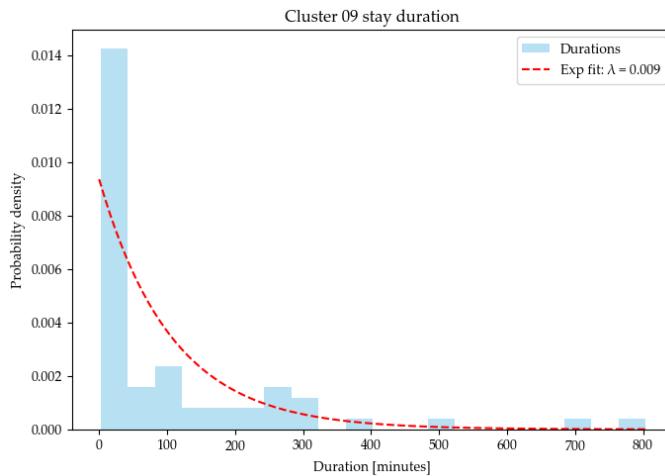
Dwell-time histogram for Cluster 06.



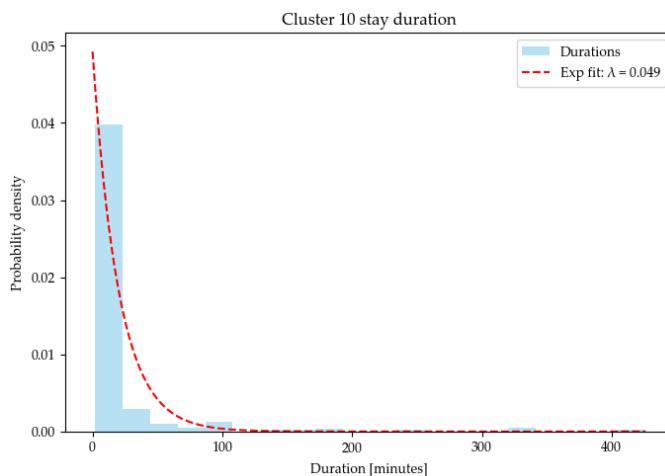
Dwell-time histogram for Cluster 07.



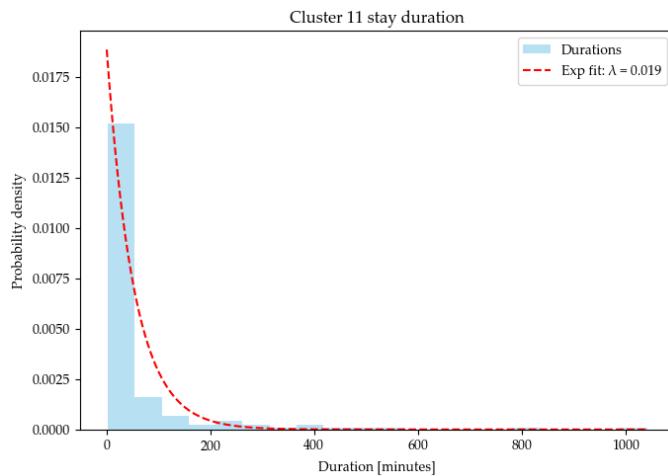
Dwell-time histogram for Cluster 08.



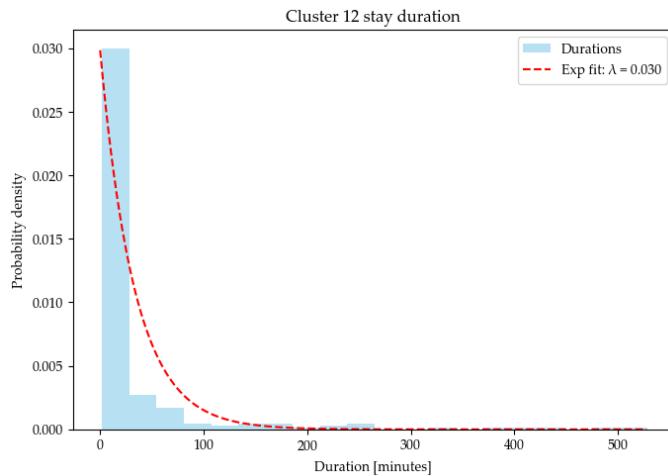
Dwell-time histogram for Cluster 09.



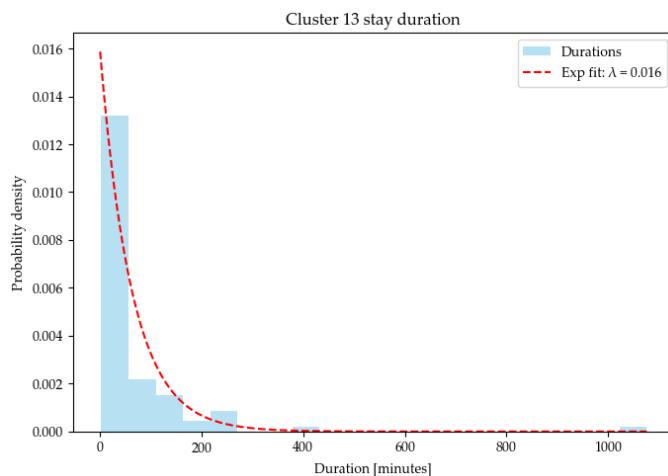
Dwell-time histogram for Cluster 10.



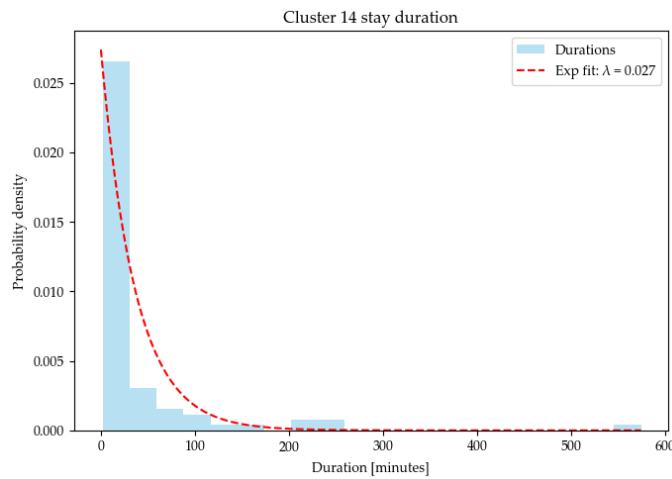
Dwell-time histogram for Cluster 11.



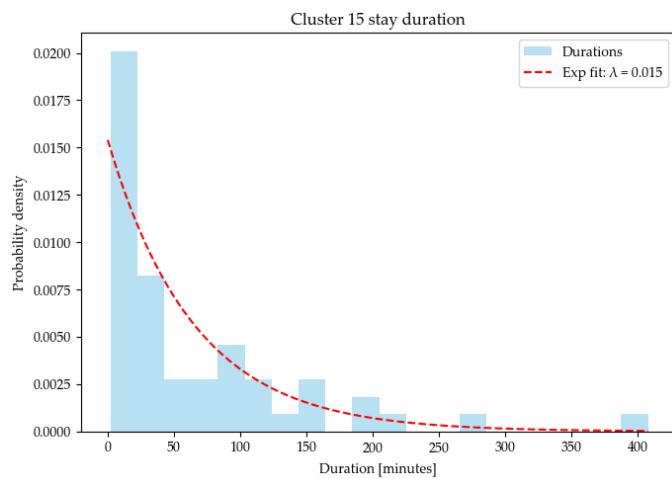
Dwell-time histogram for Cluster 12.



Dwell-time histogram for Cluster 13.



Dwell-time histogram for Cluster 14.



Dwell-time histogram for Cluster 15.

Appendix D – Top 64 bidirectional transitions with group labels

Top 64 bidirectional transitions with group labels

Rank	A	Group color	B	Group color	A→B	B→A	Total
1	10	green	11	green	132	134	266
2	7	green	12	green	105	96	201
3	10	green	12	green	63	72	135
4	7	green	10	green	55	63	118
5	12	green	14	blue	55	53	108
6	5	blue	11	green	37	37	74
7	9	blue	13	blue	29	29	58
8	3	blue	11	green	27	29	56
9	14	blue	15	blue	23	19	42
10	6	blue	13	blue	16	20	36
11	6	blue	7	green	16	15	31
12	5	blue	9	blue	14	14	28
13	11	green	15	blue	12	14	26
14	2	red	3	blue	12	12	24
15	13	blue	15	blue	11	13	24
16	4	red	10	green	12	11	23
17	8	red	12	green	10	12	22
18	8	red	13	blue	10	12	22
19	9	blue	14	blue	9	11	20
20	4	red	7	green	9	10	19
21	6	blue	9	blue	8	11	19
22	2	red	4	red	9	9	18
23	3	blue	10	green	7	11	18
24	9	blue	15	blue	8	10	18
25	7	green	13	blue	7	11	18
26	6	blue	15	blue	8	9	17
27	4	red	8	red	8	9	17
28	5	blue	13	blue	8	8	16

Continued on next page

Table 1 – continued from previous page

Rank	A	Group color	B	Group color	A→B	B→A	Total
29	3	blue	7	green	7	8	15
30	5	blue	15	blue	7	7	14
31	2	red	5	blue	7	7	14
32	2	red	7	green	6	8	14
33	8	red	14	blue	6	8	14
34	4	red	11	green	6	8	14
35	3	blue	9	blue	6	8	14
36	6	blue	12	green	6	7	13
37	3	blue	12	green	6	7	13
38	8	red	11	green	6	7	13
39	2	red	6	blue	6	7	13
40	2	red	9	blue	6	7	13
41	8	red	10	green	5	7	12
42	7	green	14	blue	5	7	12
43	4	red	6	blue	5	7	12
44	4	red	9	blue	5	7	12
45	2	red	10	green	5	7	12
46	8	red	15	blue	5	6	11
47	3	blue	14	blue	5	6	11
48	7	green	15	blue	5	6	11
49	2	red	11	green	5	6	11
50	3	blue	15	blue	5	6	11
51	6	blue	14	blue	5	5	10
52	6	blue	11	green	5	5	10
53	9	blue	11	green	5	5	10
54	9	blue	12	green	5	5	10
55	7	green	11	green	5	5	10
56	4	red	13	blue	5	5	10
57	4	red	15	blue	5	5	10
58	2	red	12	green	5	5	10
59	2	red	15	blue	5	5	10
60	3	blue	13	blue	5	5	10
61	5	blue	14	blue	5	5	10
62	5	blue	12	green	5	5	10
63	3	blue	5	blue	5	5	10
64	4	red	12	green	5	5	10

Appendix D – Weighted calibration of the Bertoldi-type law

The calibration of the Bertoldi-type law in the hybrid model requires special care when the goal is to reproduce long-term mean transport. The relationship is written as

$$q_b^* = a (\omega_b^*)^b,$$

where b is fixed by the regression slope in log–log space and a is a scale parameter. If a is estimated by ordinary least squares with equal weights, the regression takes the form

$$\hat{a} = \frac{\sum_s (\omega_{b,s}^*)^b q_{b,s}^{*,\text{exp}}}{\sum_s (\omega_{b,s}^*)^{2b}},$$

where the sum runs over states s . This approach gives each state the same importance regardless of its persistence. Since low- ω_b^* states are more numerous, they dominate the fit, and the synthetic means obtained from this calibration underestimate the experimental reference.

The relevant quantity is the long-term time average. If $S(t)$ denotes the state at time t and $q^*(S(t))$ its associated transport, then over a large horizon T

$$\mu_T(a) = \frac{1}{T} \int_0^T q^*(S(t); a) dt \rightarrow \sum_s \pi_s q_{b,s}^*(a),$$

where π_s is the empirical fraction of time spent in state s . In a Markovian setting with finite dwell times, not only the fraction of time but also the mean residence time τ_s matters for the variance of the sample mean over finite T . The variance of the block average can be approximated as

$$\text{Var}[\mu_T(a)] \approx \frac{C}{T} \sum_s \pi_s \tau_s \left(q_{b,s}^{*,\text{exp}} - a (\omega_{b,s}^*)^b \right)^2,$$

with C a constant independent of s and a . This expression shows that the contribution of each state to the error of the time average grows both with its time fraction π_s and with its persistence τ_s . Long and frequent episodes have more influence on finite-time averages than rare or short-lived ones.

Minimizing this variance with respect to a is equivalent to solving the weighted least squares problem

$$\min_a \sum_s w_s \left(q_{b,s}^{*,\text{exp}} - a (\omega_{b,s}^*)^b \right)^2, \quad w_s \propto \pi_s \tau_s.$$

The solution is

$$\hat{a}_w = \frac{\sum_s w_s (\omega_{b,s}^*)^b q_{b,s}^{*,\text{exp}}}{\sum_s w_s (\omega_{b,s}^*)^{2b}}.$$

Predicted values then follow as $q_{b,s}^{*,\text{pred}} = \hat{a}_w (\omega_{b,s}^*)^b$. The surrogate mean under the same measure is

$$\mu_w^{\text{pred}} = \sum_s w_s q_{b,s}^{*,\text{pred}}, \quad \mu_w^{\text{exp}} = \sum_s w_s q_{b,s}^{*,\text{exp}}.$$

The difference between the two is the weighted average of regression residuals,

$$\mu_w^{\text{pred}} - \mu_w^{\text{exp}} = \sum_s w_s \left(q_{b,s}^{*,\text{pred}} - q_{b,s}^{*,\text{exp}} \right).$$

By Cauchy–Schwarz inequality, this difference is bounded by the weighted root mean square error,

$$|\mu_w^{\text{pred}} - \mu_w^{\text{exp}}| \leq \sqrt{\sum_s w_s \left(q_{b,s}^{*,\text{pred}} - q_{b,s}^{*,\text{exp}} \right)^2} = \text{RMSE}_w.$$

Therefore, a good weighted fit (small RMSE_w) necessarily implies that the mean under w is close to the experimental mean under the same measure. Importantly, this guarantee does not require prescribing the experimental mean in advance, but emerges naturally from the choice of weights.

From a physical perspective, this procedure is justified because sediment transport depends not only on the instantaneous flux associated with a morphology but also on how long the system remains in that morphology. Frequent, long-lived low-transport states should influence the mean more strongly than rare, short-lived high-transport states. By using weights proportional to $\pi_s \tau_s$, the calibration emphasizes exactly those states that dominate the time dynamics, producing a surrogate law that is consistent with the physical process. This adjustment of the factor a thus transforms the Bertoldi relationship into a mean-preserving predictor without requiring prior knowledge of the global experimental mean, while retaining the form and exponent of the original law.

Bibliography

Abramian, A., Devauchelle, O., & Lajeunesse, E. (2020). Laboratory rivers adjust their shape to sediment transport. *Physical Review E*, 102(5), 053101.

Agarwal, S., Phillips, C. A., & Pachter, L. (2007). Multidimensional scaling for large genomic data sets. *BMC Bioinformatics*, 8(1), 1–11.

Ancey, C. (2020a). Bedload transport: a walk between randomness and determinism. part 1. the state of the art. *Journal of Hydraulic Research*, 58(1), 1–17.

Ancey, C. (2020b). Bedload transport: a walk between randomness and determinism. part 2. challenges and prospects. *Journal of Hydraulic Research*, 58(1), 18–33.

Antoniazza, G., Bakker, M., & Lane, S. N. (2019). Revisiting the morphological method in two-dimensions to quantify bed-material transport in braided rivers. *Earth Surface Processes and Landforms*, 44(11), 2251–2267.

Arkov, V., Kulikov, G., & Breikin, T. (1999). Fuzzy markov modeling in automatic control of complex dynamic systems.

Ashmore, P. (2013). 9.17 Morphology and Dynamics of Braided Rivers. In J. F. Shroder (Ed.). Academic Press.

Ashmore, P. (2024, August). What braiding reveals about river morphology, bedload and channel change. In *River Flow 2022* (1st Edition, pp. 9–18). CRC Press.

Ashmore, P. (1993). Anabranch confluence kinetics and sedimentation processes in gravel-braided streams. *Geological Society, London, Special Publications*, 75(1), 129–146.

Beechie, T. J., Sear, D. A., Olden, J. D., Pess, G. R., Buffington, J. M., Moir, H., Roni, P., & Pollock, M. M. (2010). Process-based principles for restoring river ecosystems. *BioScience*, 60(3), 209–222.

Bellman, R. (1961). *Adaptive control processes: a guided tour*. Princeton University Press.

Benavides, S. J., Deal, E., Rushlow, M., Venditti, J. G., Zhang, Q., Kamrin, K., & Perron, J. T. (2022). The impact of intermittency on bed load sediment transport. *Geophysical research letters*, 49(5).

Bernhardt, E. S., Palmer, M. A., Allan, J. D., Alexander, G., Barnas, K., Brooks, S., Carr, J., Clayton, S., Dahm, C., Follstad-Shah, J., et al. (2005). Synthesizing us river restoration efforts.

Bertoldi, W., Ashmore, P., & Tubino, M. (2009a). A method for estimating the mean bed load flux in braided rivers. *Geomorphology*, 103(3), 330–340.

Bertoldi, W., Ashmore, P., & Tubino, M. (2009b). A method for estimating the mean bed load flux in braided rivers. *Geomorphology*, 103(3), 330–340.

Bertoldi, W., Recking, A., & Surian, N. (2014, October 21). *Braiding rivers: state of the art and future challenges: Braided rivers workshop 2014;le domaine de sainte croix, france, 23–27 june 2014*. <https://eos.org/meeting-reports/braiding-rivers-state-art-future-challenges>

Bertoldi, W., & Tubino, M. (2005). Bed and bank evolution of bifurcating channels. *Water Resources Research*, 41(7).

Biron, P. M., Buffin-Bélanger, T., Larocque, M., Choné, G., Cloutier, C.-A., Ouellet, M.-A., Demers, S., Olsen, T., Desjarlais, C., & Eyquem, J. (2014). Freedom space for rivers: a sustainable management approach to enhance river resilience. *Environmental management*, 54(5), 1056–1073.

Booker, W. H., & Eaton, B. C. (2021). Morphodynamic styles: characterising the behaviour of gravel-bed rivers using a novel, quantitative index.

Bristow, C. S., & Best, J. L. (1993). Braided rivers: perspectives and problems.

Burkham, D. E. (1972). *Channel Changes of the Gila River in Safford Valley, Arizona, 1846-1970*. U.S. Government Printing Office.

Cammilleri, A., Gueniat, F., Carlier, J., Pastur, L., Memin, E., Lusseyran, F., & Artana, G. (2013). POD-spectral decomposition for fluid flow analysis and model reduction. *Theoretical and computational fluid dynamics*, 27(6), 787–815.

Campello, R. J. G. B., Moulavi, D., & Sander, J. (2013). Density-based clustering based on hierarchical density estimates. In *Advances in Knowledge Discovery and Data Mining* (pp. 160–172). Springer Berlin Heidelberg.

Campello, R. J. G. B., Moulavi, D., Zimek, A., & Sander, J. (2015). Hierarchical density estimates for data clustering, visualization, and outlier detection. *ACM transactions on knowledge discovery from data*, 10(1), 1–51.

Chalov, S. R., & Alexeevsky, N. I. (2013). Braided rivers: structure, types and hydrological effects. *Hydrology Research*, 46(2), 258–275.

Church, M., & Ferguson, R. I. (2015). Morphodynamics: rivers beyond steady state. *Water Resources Research*, 51(4), 1883–1897.

Church, M. (2006). Bed Material Transport and the Morphology of Alluvial River Channels. *Annual Review of Earth and Planetary Sciences*, 34(1), 325–354.

Davies, T. R., & Lee, A. L. (1989). Discussion of “influence of channel width on bed load transport capacity” by michael a. carson and george a. griffiths (december, 1987, vol. 113, no. 12). *Journal of Hydraulic Engineering*, 115(7), 1005–1007.

Demaine, E. D., Lee, N., Price, E., Schellekens, M., & Urschel, J. C. (2021). Multidimensional scaling and the nonconvex optimization landscape. *SIAM Journal on Optimization*, 31(4), 2840–2863.

Dhont, B., & Ancey, C. (2018). Are bedload transport pulses in gravel bed rivers created by bar migration or sediment waves? *Geophysical Research Letters*, 45(11), 5501–5508.

Doeschl, A. B., Ashmore, P. E., Davison, M., et al. (2006). Methods for assessing exploratory computational models of braided rivers. *Braided Rivers: Process, Deposits, Ecology and Management*, 36, 177–197.

Dubuisson, M.-P., & Jain, A. K. (1994). A modified Hausdorff distance for object matching. *Proceedings of 12th International Conference on Pattern Recognition*, 1, 566–568 vol.1.

Eaton, B., Millar, R. G., & Davidson, S. (2010). Channel patterns: braided, anabranching, and single-thread. *Geomorphology*, 120(3-4), 353–364.

Egozi, R., & Ashmore, P. (2008). Defining and measuring braiding intensity. *Earth Surface Processes and Landforms*, 33(14), 2121–2138.

Egozi, R., & Ashmore, P. (2009). Experimental analysis of braided channel pattern response to increased discharge. *Journal of Geophysical Research*, 114(F2).

Ester, M., Kriegel, H., Sander, J., & Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. *Knowledge Discovery and Data Mining*, 226–231.

Ettema, R., & Armstrong, D. L. (2019). Bedload and channel morphology along a braided, sand-bed channel: insights from a large flume. *Journal of Hydraulic Research*, 57(6), 822–835.

Foufoula-Georgiou, E., & Sapozhnikov, V. (2001). Scale Invariances in the Morphology and Evolution of Braided Rivers. *Mathematical Geology*, 33(3), 273–291.

Fukunaga, K. (1990). *Introduction to statistical pattern recognition* (2nd). Academic Press.

Habersack, H., Seitz, H., & Laronne, J. B. (2008). Spatio-temporal variability of bedload transport rate: analysis and 2D modelling approach. *Geodinamica Acta*, 21(1-2), 67–79.

Heckmann, T., Schwanghart, W., & Phillips, J. D. (2015). Graph theory - Recent developments of its application in geomorphology. *Geomorphology*, 243, 130–146.

Hiatt, M., Sonke, W., Addink, E. A., Dijk, W. M., Kreveld, M., Ophelders, T., Verbeek, K., Vlaming, J., Speckmann, B., & Kleinhans, M. G. (2020). Geometry and topology of estuary and braided river channel networks automatically extracted from topographic data. *Journal of Geophysical Research: Earth Surface*, 125(1).

Hoey, T. (1992). Temporal variations in bedload transport rates and sediment storage in gravel-bed rivers. *Progress in Physical Geography: Earth and Environment*, 16(3), 319–338.

Hoey, T. B., & Sutherland, A. J. (1991). Channel morphology and bedload pulses in braided rivers: a laboratory study. *Earth Surface Processes and Landforms*, 16(5), 447–462.

Hoffmann, J., Bufe, A., & Caers, J. (2019). Morphodynamic analysis and statistical synthesis of geomorphic data: application to a flume experiment. *Journal of Geophysical Research: Earth Surface*, 124(11), 2561–2578.

Jagers, B. (2003). *Modelling planform changes of braided rivers* [Doctoral dissertation].

Karimi, D., & Salcudean, S. E. (2020). Reducing the hausdorff distance in medical image segmentation with convolutional neural networks. *IEEE Transactions on Medical Imaging*, 39(2), 499–513.

Knighton, D. (2014, April). *Fluvial Forms and Processes: A New Perspective*. Routledge.

Kondolf, G. M. (2006). River restoration and meanders. *Ecology and Society*, 11(2).

Köppen, M. (2000). The curse of dimensionality. *5th online world conference on soft computing in industrial applications (WSC5)*, 1, 4–8.

Kurth, A.-M., & Schirmer, M. (2014). Thirty years of river restoration in switzerland: implemented measures and lessons learned. *Environmental earth sciences*, 72(6), 2065–2079.

Lane, E. W. (1955). Importance of fluvial morphology in hydraulic engineering. *Proceedings (American Society of Civil Engineers)*; v. 81, paper no. 745.

Lane, S. N. (2006). Approaching the system-scale understanding of braided river behaviour. *Braided rivers: process, deposits, ecology and management*, 36, 107–135.

Leopold, L. B., & Wolman, M. G. (1957). *River Channel Patterns: Braided, Meandering, and Straight*. U.S. Government Printing Office.

Li, Y., Yang, M., & Zhang, Z. (2016). A survey of multi-view representation learning. *arXiv preprint arXiv:1610.01206*.

Li, Y., & Limaye, A. B. (2025). Coherent motion of channel threads in the braided Brahmaputra-Jamuna River. *Journal of geophysical research. Earth surface*, 130(6), e2024JF008196.

Limaye, A. B. (2020). How do braided rivers grow channel belts? *Journal of Geophysical Research: Earth Surface*, 125(8).

Machado, A., Behrisch, M., & Telea, A. (2025). Necessary but not sufficient: limitations of projection quality metrics. *Computer Graphics Forum (EuroVis)*.

Mariethoz, G., & Caers, J. (2014). *Multiple-point geostatistics: stochastic modeling with training images*. John Wiley & Sons.

Mariethoz, G., Renard, P., & Straubhaar, J. (2010). The direct sampling method to perform multiple-point geostatistical simulations. *Water Resources Research*, 46(11).

Marti, C., & Bezzola, G. R. (2006). Bed load transport in braided gravel-bed rivers. *Braided rivers: process, deposits, ecology and management*, 36, 199–215.

McInnes, L., & Healy, J. (2017). Accelerated hierarchical density clustering. *arXiv [stat.ML]*.

McInnes, L., Healy, J., & Melville, J. (2020). Umap: uniform manifold approximation and projection for dimension reduction.

Métivier, F., & Meunier, P. (2003). Input and output mass flux correlations in an experimental braided stream. implications on the dynamics of bed load transport. *Journal of hydrology*, 271(1-4), 22–38.

Meunier, P., Métivier, F., Lajeunesse, E., Mériaux, A., & Faure, J. (2006). Flow pattern and sediment transport in a braided river: the “torrent de st pierre” (french alps). *Journal of Hydrology*, 330(3-4), 496–505.

Micheli, E., & Larsen, E. (2010). A multi-scale approach to understanding the role of riparian vegetation in channel migration on the sacramento river, california, usa. *Earth Surface Processes and Landforms*, 35(9), 1087–1101.

Nicholas, A. P., Ashworth, P. J., Sambrook Smith, G. H., & Sandbach, S. D. (2018). Avulsions in a simulated large, lowland braided river. *Earth Surface Processes and Landforms*, 43(11), 2283–2298.

Paola, C. (2001). Modelling stream braiding over a range of scales.

Pawlowski, C. W., & McCord, C. (2009). A Markov model for assessing ecological stability properties. *Ecological modelling*, 220(2), 86–95.

Phillips, J. D. (2006). Deterministic chaos and historical geomorphology: A review and look forward. *Geomorphology*, 76(1), 109–121.

Poff, N. L., Allan, J. D., Bain, M. B., Karr, J. R., Prestegard, K. L., Richter, B. D., Sparks, R. E., & Stromberg, J. C. (1997). The natural flow regime. *BioScience*, 47(11), 769–784.

Polvi, L. E., Lind, L., Persson, H., Miranda-Melo, A., Pilotto, F., Su, X., & Nilsson, C. (2020). Facets and scales in river restoration: nestedness and interdependence of hydrological, geomorphic, ecological, and biogeochemical processes. *Journal of Environmental Management*, 265, 110288.

Pryor, B. S., Lisle, T., Montoya, D. S., & Hilton, S. (2011). Transport and storage of bed material in a gravel-bed channel during episodes of aggradation and degradation: a field and flume study. *Earth Surface Processes and Landforms*, 36(15), 2028–2041.

Rajbanshi, J., Das, S., & Patel, P. P. (2022). Planform changes and alterations of longitudinal connectivity caused by the 2019 flood event on the braided Brahmaputra River in Assam, India. *Geomorphology*, 403, 108174.

Recking, A., Johannot, A., Horita, K., Nasr, M., Vázquez-Tarrío, D., Zanker, S., Fontaine, F., & Melun, G. (2024). An attempt to take into account natural variability in 1D bed-load prediction. *Journal of geophysical research. Earth surface*, 129(7), e2023JF007601.

Redolfi, M., Bertoldi, W., Tubino, M., & Welber, M. (2018). Bed load variability and morphology of gravel bed rivers subject to unsteady flow: a laboratory investigation. *Water Resources Research*, 54(2), 842–862.

Redolfi, M. (2015). *Sediment transport and morphology of braided rivers: steady and unsteady regime* [PhD Thesis]. Queen Mary University of London.

Redolfi, M., Guidorizzi, L., Tubino, M., & Bertoldi, W. (2017). Capturing the spatiotemporal variability of bedload transport: a time-lapse imagery technique. *Earth Surface Processes and Landforms*, 42(7), 1140–1147.

Redolfi, M., Welber, M., Carlin, M., Tubino, M., & Bertoldi, W. (2020). Morphometric properties of alternate bars and water discharge: a laboratory investigation. *Earth Surface Dynamics*, 8(3), 789–808.

Rizzello, S., Scaraggi, M., Nelson, A. D., Primavera, L., Napoli, G., Stecca, G., Vitolo, R., & De Bartolo, S. (2024). Multiscaling behavior of braided channel networks: An alternative formulation of Taylor's law variate transformations. *Physical review. E*, 109(3-1), 034306.

Sambrook Smith, G. H., Best, J. L., Petts, G. E., & Bristow, C. S. (2006). Braided Rivers: process, deposits, ecology and management.

Sapozhnikov, V. B., & Foufoula-Georgiou, E. (1997). Experimental evidence of dynamic scaling and indications of self-organized criticality in braided rivers. *Water Resources Research*, 33(8), 1983–1991.

Sapozhnikov, V. B., & Foufoula-Georgiou, E. (1999). Horizontal and vertical self-organization of braided rivers toward a critical state. *Water Resources Research*, 35(3), 843–851.

Scheidt, C., Fernandes, A. M., Paola, C., & Caers, J. (2016a). Quantifying natural delta variability using a multiple-point geostatistics prior uncertainty model. *Journal of Geophysical Research: Earth Surface*, 121(10), 1800–1818.

Scheidt, C., Fernandes, A. M., Paola, C., & Caers, J. (2016b). Quantifying natural delta variability using a multiple-point geostatistics prior uncertainty model. *Journal of Geophysical Research: Earth Surface*, 121(10), 1800–1818.

Schubert, E., Sander, J., Ester, M., Kriegel, H. P., & Xu, X. (2017). DBSCAN Revisited, Revisited: Why and how you should (still) use DBSCAN. *ACM transactions on database systems*, 42(3), 1–21.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal*, 27(3), 379–423.

Singh, A., Fienberg, K., Jerolmack, D. J., Marr, J., & Foufoula-Georgiou, E. (2009). Experimental evidence for statistical scaling and intermittency in sediment transport rates. *Journal of Geophysical Research*, 114(F1).

Sivakumar, B. (2004). Chaos theory in geophysics: past, present and future. *Chaos, Solitons & Fractals*, 19(2), 441–462.

Stecca, G., & Hicks, D. M. (2022). Numerical Simulations of Confined Braided River Morphodynamics: Display of Deterministic Chaos and Characterization Through Turbulence Theory. *Journal of Geophysical Research: Earth Surface*, 127(3).

Stecca, G., Zolezzi, G., Hicks, D. M., & Surian, N. (2019). Reduced braiding of rivers in human-modified landscapes: Converging trajectories and diversity of causes. *Earth-Science Reviews*, 188, 291–311.

Strebelle, S. (2006). Sequential simulation for modeling geological structures from training images.

Surian, N. (2015). Fluvial Processes in Braided Rivers. In P. Rowinski & A. Radecki-Pawlak (Eds.), *Rivers – physical, fluvial and environmental processes* (pp. 403–425). Springer International Publishing.

Taha, A. A., & Hanbury, A. (2015). Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. *BMC Medical Imaging*, 15, 29.

Taira, K., Hemati, M. S., Brunton, S. L., Sun, Y., Duraisamy, K., Bagheri, S., Dawson, S. T. M., & Yeh, C.-A. (2020). Modal analysis of fluid flows: Applications and outlook. *AIAA journal*, 58(3), 998–1022.

Tejedor, A., Longjas, A., Zaliapin, I., & Foufoula-Georgiou, E. (2015a). Delta channel networks: 2. Metrics of topologic and dynamic complexity for delta comparison,

physical inference, and vulnerability assessment. *Water Resources Research*, 51(6), 4019–4045.

Tejedor, A., Longjas, A., Zaliapin, I., & Foufoula-Georgiou, E. (2015b). Delta channel networks: 1. a graph-theoretic approach for studying connectivity and steady state transport on deltaic surfaces. *Water Resources Research*, 51(6), 3998–4018.

Tejedor, A., Schwenk, J., Kleinhans, M., Limaye, A. B., Vulis, L., Carling, P., Kantz, H., & Foufoula-Georgiou, E. (2022). The entropic braiding index (*eBI*): A robust metric to account for the diversity of channel scales in multi-thread rivers. *Geophysical research letters*, 49(16).

van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86), 2579–2605.

van der Maaten, L., Postma, E., & van den Herik, J. (2009). Dimensionality reduction: a comparative review. *Journal of Machine Learning Research*, 10, 66–71.

van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart, E., & Yu, T. (2014). Scikit-image: image processing in python. *PeerJ*, 2, e453.

Vericat, D., Wheaton, J. M., & Brasington, J. (2017). Revisiting the Morphological Approach. In *Gravel-Bed Rivers* (pp. 121–158). John Wiley & Sons, Ltd.

Vesipa, R., Camporeale, C., & Ridolfi, L. (2018). Hydraulics of braided river dynamics. insights from flume experiments (A. Paquier & N. Rivière, Eds.). *E3S Web of Conferences*, 40, 02020.

Wang, L., Cuthbertson, A., Pender, G., & Zhong, D. (2019). Bed Load Sediment Transport and Morphological Evolution in a Degrading Uniform Sediment Channel Under Unsteady Flow Hydrographs. *Water Resources Research*, 55(7), 5431–5452.

Wang, Y., Huang, H., Rudin, C., & Shaposhnik, Y. (2021). Understanding how dimension reduction tools work: an empirical approach to deciphering t-sne, umap, trimap, and pacmap for data visualization. *arXiv preprint arXiv:2012.04456*.

Wang, Y., Limaye, A. B., & Chadwick, A. J. (2024). Topography-based particle image velocimetry of braided channel initiation. *Water resources research*, 60(4), e2023WR035229.

Warburton, J., & Davies, T. (1994). Variability of bedload transport and channel morphology in a braided river hydraulic model. *Earth Surface Processes and Landforms*, 19(5), 403–421.

Williams, R. D., Rennie, C. D., Brasington, J., Hicks, D. M., & Vericat, D. (2015). Linking the spatial distribution of bed load transport to morphological change during high-flow events in a shallow braided river. *Journal of Geophysical Research: Earth Surface*, 120(3), 604–622.

Wohl, E. (2020). *Rivers in the Landscape*. John Wiley & Sons.

Wohl, E., Bierman, P., & Montgomery, D. (2016). Earth's dynamic surface: A perspective on the past 50 years in geomorphology.

Woolsey, S., Capelli, F., Gonser, T., Hoehn, E., Hostmann, M., Junker, B., Paetzold, A., Roulier, C., Schweizer, S., Tiegs, S. D., et al. (2007). A strategy to assess river restoration success. *Freshwater Biology*, 52(4), 752–769.

Xu, C., Tao, D., & Xu, C. (2013). A survey on multi-view learning. *arXiv preprint arXiv:1304.5634*.

Zimmermann, A. E., Church, M., & Hassan, M. A. (2008). Video-based gravel transport measurements with a flume mounted light table. *Earth Surface Processes and Landforms*, 33(14), 2285–2296.

Clemente Gotelli

HYDRAULIC ENGINEER, PHD CANDIDATE

clemente.gotelli@outlook.com

Rue de Bourg 15, 1003 Lausanne

+41 78 205 9383

linkedin.com/in/cmgotelli

Strengths

I am a hydraulic engineer completing a PhD in Hydraulic Engineering at EPFL, with a solid background in numerical modeling, fluid mechanics, and hydrology. I have experience working on projects of various scales in Chile and internationally, collaborating with companies of different sizes. I am looking for a role in hydraulics and hydrology that will allow me to continue developing my skills and contribute to innovative solutions in water resources management.

Education

PhD in Mechanics, Environmental Hydraulics Laboratory, EPFL November 2020 - Present

Stochastic modeling of mountain rivers evolution with a focus in the connection between morphology and sediment transport in a context of climate change. Assistant for EPFL courses: *Hydrological risks and land use development, Hydraulics: floods, dam breaks, wave; Fluid mechanics*.

MS. in Hydraulic Engineering, UC Chile 2016 – 2018

Fluid mechanics and numerical modelling specialization. Master thesis: Turbulence study of Marine Energy Turbines.

Bachelor's in Civil Engineering, UC Chile 2011 – 2016

Major: Civil Engineering | Minor: Hydraulic Engineering. Relevant courses: Advance hydraulics, hydraulic structures, advance hydrological modeling, environmental fluid mechanics, computational hydraulics, underground water, urban hydrology.

Experience

CDM Smith: Project Engineer | Santiago, Chile 2019 – 2020

Flooding Simulations and Design of fluvial protections

Conception and Design of Piping systems for desalination plants for Mining Companies in the Andes.

Transformation of calculation spreadsheets into Python with Graphic User Interface.

Experience working with international teams within the company for different projects around the world.

Project Engineer: UC Hydraulic Constructions Laboratory | Santiago, Chile 2018 – 2019

Construction and analysis of Hydraulic Scale Models for different dam projects in Chile.

3D Numerical Modelling of hydraulic structures to study different scenarios.

2D Numerical Modelling for fluvial protection design.

IT Skills

HEC-RAS, Iber, SWMM, Epanet, QGis/ArcGIS,

OpenFOAM, Ansys Fluent, Office 365 Suite

Python, Matlab/Octave, Fortran, VBA

Languages

Spanish – Native

English – Full Professional Proficiency (C1)

French – Professional Working Proficiency (B1/B2)

Projects

PhD thesis: Understanding the dynamics of braided river.

I performed experiments to understand the time evolution of braided rivers both in terms of morphology and sediment transport.

Personal situation

Citizenship: Italian & Chilean

Work Permit: *Permis B UE/AELE*

Marital Status: Married